Recent Progress on a Statistical Network Calculus

Jorg Liebeherr

Department of Computer Science
University of Virginia

Collaborators

Almut Burchard

Robert Boorstyn

Chaiwat Oottamakorn

Stephen Patek

Chengzhi Li

Contents

- R. Boorstyn, A. Burchard, J. Liebeherr, C. Oottamakorn. “Statistical Service Assurances for Packet Scheduling Algorithms”, IEEE Journal on Selected Areas in Communications. Special Issue on Internet QoS, Vol. 18, No. 12, pp. 2651-2664, December 2000.
- A. Burchard, J. Liebeherr, and S. D. Patek. “A Calculus for End-to-end Statistical Service Guarantees.” (2nd revised version), Technical Report, May 2002.
- J. Liebeherr, A. Burchard, and S. D. Patek , “Statistical Per-Flow Service Bounds in a Network with Aggregate Provisioning”, Infocom 2003.
- C. Li, A. Burchard, J. Liebeherr, “Calculus with Effective Bandwidth”, July 2002.

Service Guarantees

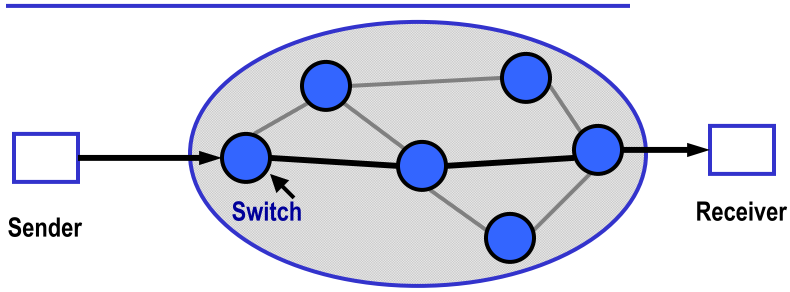

- A **deterministic service** gives worst-case guarantees

$$\text{Delay} \leq d$$

- A **statistical service** provides probabilistic guarantees

$$\Pr[\ \text{Delay} \geq d\ ] \leq \varepsilon \quad \text{or} \quad \Pr[\ \text{Loss} \geq l\ ] \leq \varepsilon$$

Multiplexing Gain

$$\begin{pmatrix}\text{Resource needed}\\ \text{to support QoS}\\ \text{for N flows}\end{pmatrix} << N \cdot \begin{pmatrix}\text{Resource needed}\\ \text{to support QoS}\\ \text{for 1 flow}\end{pmatrix}$$

Sources of multiplexing gain:

- Traffic Conditioning (Policing, Shaping)
- Scheduling
- Statistical Multiplexing of Traffic

Scheduling

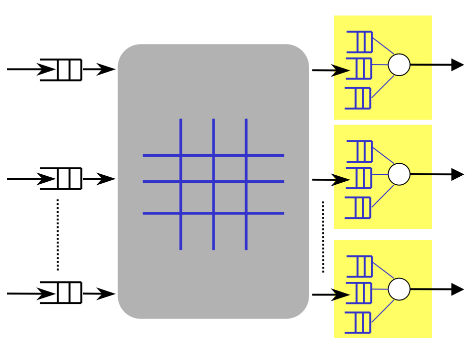

- Scheduling algorithm determines the order in which traffic is transmitted

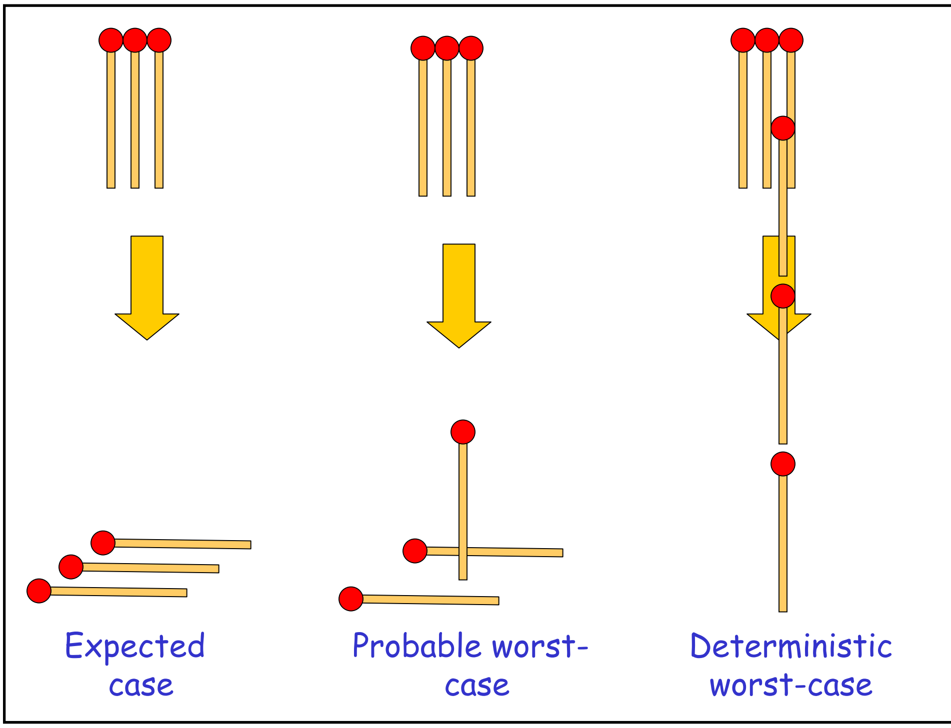

Designing Networks for Multiplexing Gain

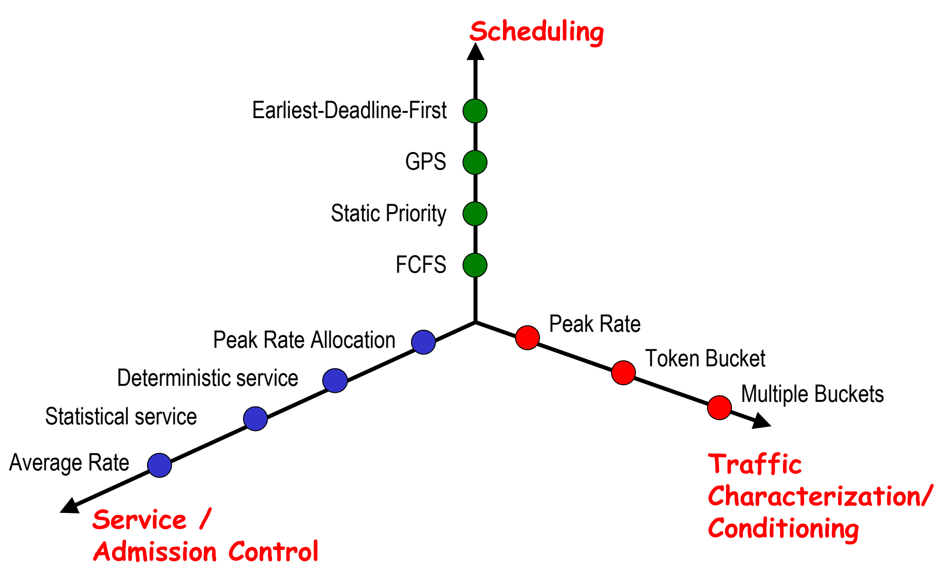

Designing Networks for Multiplexing Gain

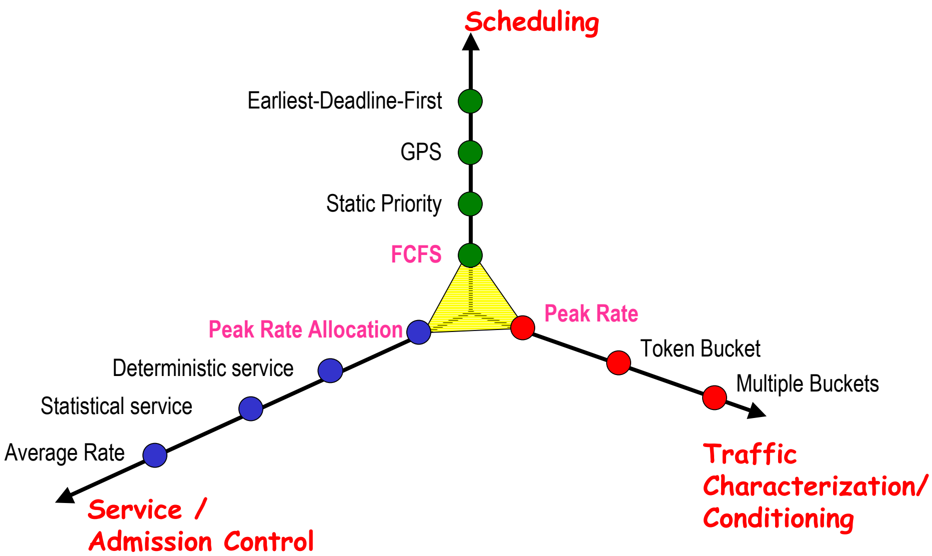

Designing Networks for Multiplexing Gain

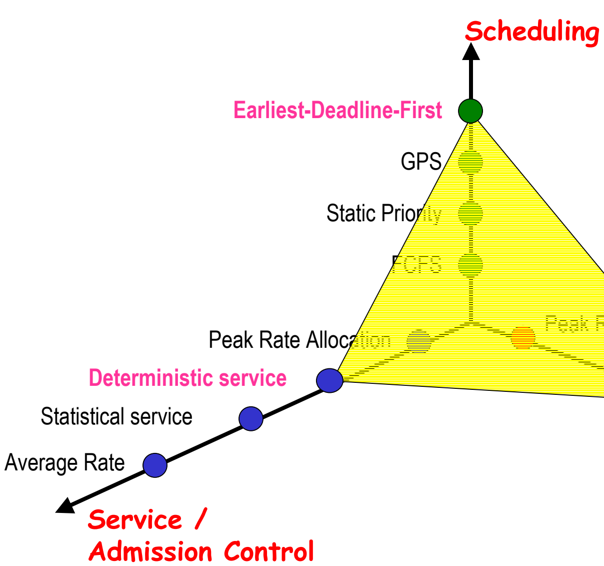

By now: The design space for determi-nistic guarantees is well understood.

Designing Networks for Multiplexing Gain

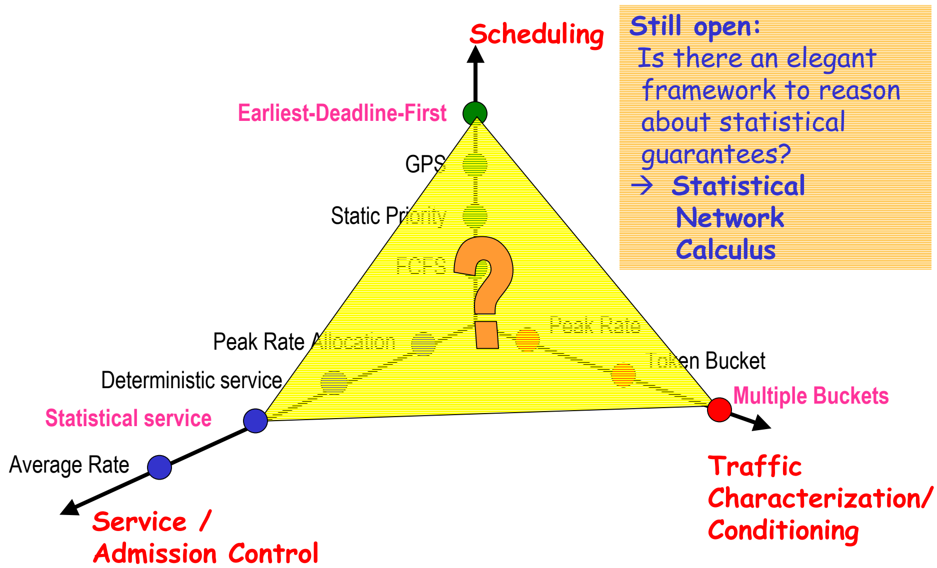

Related Work (small subset)

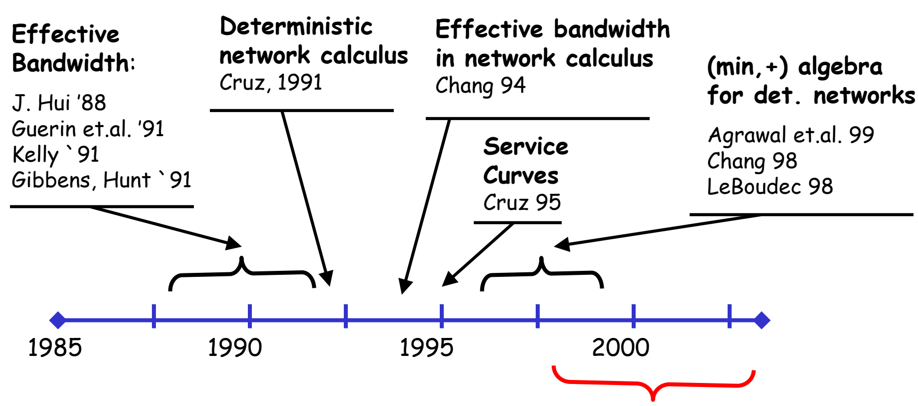

Motivation for our work on statistical network calculus:

(1) Maintain elegance of deterministic calculus

(2) Exploit know-how of statistical multiplexing

Source Assumptions

Arrivals $A_j(t, t+\tau)$ are random processes

Deterministic Calculus:

(A1) Additivity: For any $t_1 < t_2 < t_3$, we have:

$$A_j(t_1, t_2) + A_j(t_2, t_3) = A_j(t_1, t_3)$$

(A2) Subadditive Bounds: Traffic A_j is constrained by a subadditive deterministic envelope A^*_j as follows

$$A_j(t, t+\tau) \leq A^*_j(\tau) \quad , \forall t, \forall \tau$$

with $\rho = \lim_{\tau \to \infty} A^*_j(\tau)/\tau$

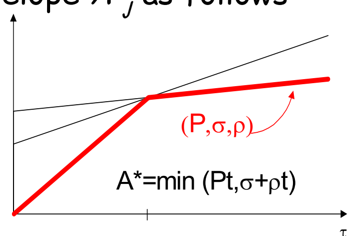

Source Assumptions

Statistical Calculus:

(A1) +(A2)

(A3) Stationarity: The A_j are stationary random variables

(A4) Independence: The A_i and A_j ($i \neq j$) are stochastically independent

(No assumptions on arrival distribution!)

Aggregating Arrivals

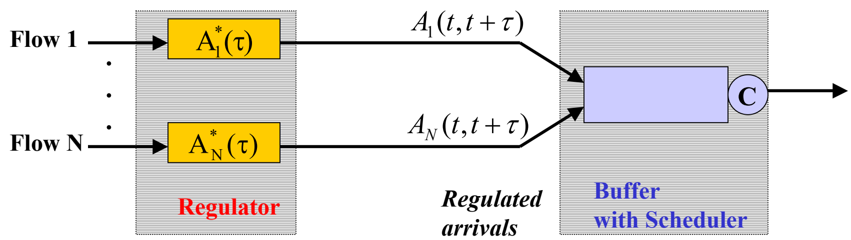

Arrivals from multiple flows: $A_{\mathcal{C}} = \sum_j A_j$

Deterministic Calculus:

Worst-case of multiple flows is sum of the worst-case of each flow $A_{\mathcal{C}}(t, t+\tau) \leq \sum_j A_j^*(\tau)$

2000

Aggregating Arrivals

Statistical Calculus:

To bound aggregate arrivals we define a function that is a bound on the sum of multiple flows with high probability → "Effective Envelope"

- Effective envelopes are non-random functions

- effective envelope $\mathcal{G}_{\mathcal{C}}^{\varepsilon}$:

$$Pr\{A_{\mathcal{C}}(t, t+\tau) \leq \mathcal{G}_{\mathcal{C}}^{\varepsilon}(\tau)\} \geq 1 - \varepsilon \quad \forall t, \tau$$

- strong effective envelope $\mathcal{H}_{\mathcal{C}}^{\ell,\varepsilon}$:

$$Pr\{\forall [t, t+\tau] \subseteq I_\ell : \ A_{\mathcal{C}}(\tau) \leq \mathcal{H}_{\mathcal{C}}^{\ell,\varepsilon}(\tau)\} \geq 1 - \varepsilon \quad \forall I_\ell$$

Obtaining Effective Envelopes

$$\boxed{\mathcal{G}^{\varepsilon}_{\mathcal{C}}(t) \;=\; \inf_{s>0} \frac{1}{s}\Big(\sum_{j\in\mathcal{C}} \log \overline{M}_j(s,t) - \log \varepsilon\Big)}$$

with $\overline{M}_j(s,t) \;=\; 1 + \dfrac{\rho_j\, t}{A^*_j(t)}\left(e^{sA^*_j(t)} - 1\right)$

$$\boxed{\mathcal{H}^{\ell,\varepsilon'}_{\mathcal{C}}(t) \le \mathcal{G}^{\varepsilon}_{\mathcal{C}}(\gamma t + a)\ ,} \qquad 0 \le t \le \ell$$

with $\varepsilon' \le \varepsilon \cdot \dfrac{\ell}{a}\dfrac{\sqrt{\gamma}+1}{\sqrt{\gamma}-1}$

$a \in (0, \ell)$

$\gamma > 1$

Effective vs. Deterministic Envelope Envelopes

$A^*=\min(Pt, \sigma+\rho t)$

Type 1 flows:
P =1.5 Mbps
ρ = .15 Mbps
σ =95400 bits

Type 2 flows:
P = 6 Mbps
ρ = .15 Mbps
σ = 10345 bits

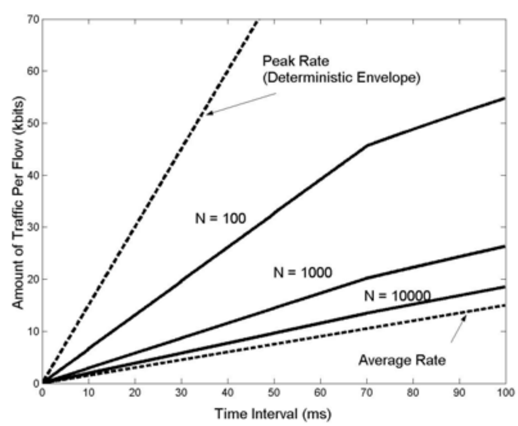

Type 1 flows

Effective Envelope vs. Deterministic Envelopes

Traffic rate at t = 50 ms
Type 1 flows

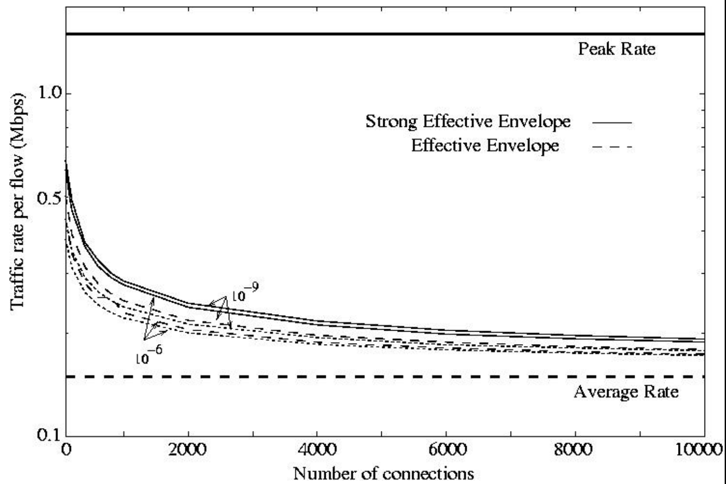

Scheduling Algorithms

- Consider a work-conserving scheduler with rate R
- Consider class-q arrival at t with $t+d_q$:

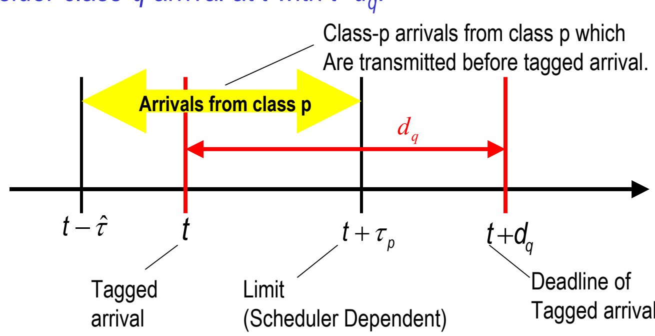

- **Tagged arrival has no delay bound violation if**

$$\sup_{\hat{\tau}}\left\{\sum_p A_{C_p}(t-\hat{\tau}, t+\tau_p) - R(\hat{\tau}+d_q)\right\} \le 0$$

Scheduling Algorithms

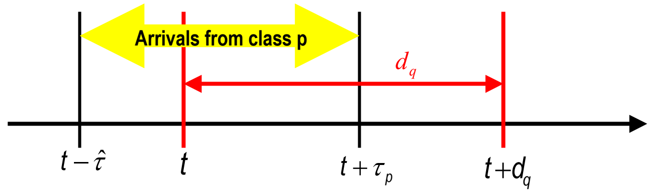

$$\sup_{\hat{\tau}}\left\{\sum_p A_{C_p}(t-\hat{\tau},t+\tau_p)-(\hat{\tau}+d_q)\right\}\le 0$$

with

FIFO: $\tau_p = 0.$

SP: $\tau_p = -\hat{\tau}\ (p>q)\ ,\ 0\ (p=q)\ ,\ d_q\ (p<q).$

EDF: $\tau_p = \max\{-\hat{\tau}, d_q - d_p\}.$

2000

Admission Control for Scheduling Algorithms

with Deterministic Envelopes:

$$\sup_{\hat{\tau}}\left\{\sum_p A_j^*(\tau_p+\hat{\tau})-\hat{\tau}\right\}\le d_q$$

with Effective Envelopes:

$$\sup_{\hat{\tau}}\left\{\sum_p \mathcal{G}_{C_p}^{\varepsilon/Q}(\tau_p+\hat{\tau})-\hat{\tau}\right\}\le d_q$$

with Strong Effective Envelopes:

$$\sup_{\hat{\tau}}\left\{\sum_p \mathcal{H}_{C_p}^{1,\varepsilon/Q}(\tau_p+\hat{\tau})-\hat{\tau}\right\}\le d_q$$

Effective vs. Deterministic Envelope

C= 45 Mbps, $\varepsilon = 10^{-6}$
Delay bounds: Type 1: d_1=100 ms, Type 2: d_2=10 ms,

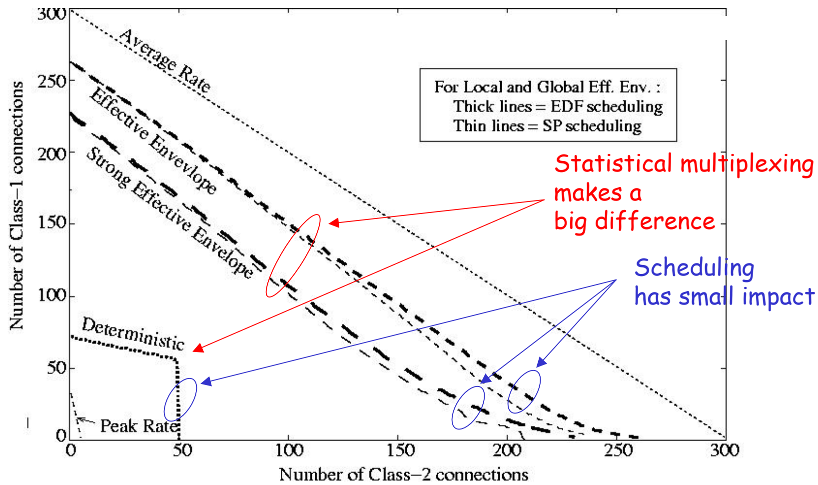

2002

Effective Envelopes and Effective Bandwidth

Effective Bandwidth (Kelly, Chang)

$$\boxed{\alpha(s,\tau) = \sup_{t \geq 0} \left\{ \frac{1}{s\tau} \log E[e^{s(A[t+\tau]-A[t])}] \right\}}$$

$$s, \tau \in (0, \infty)$$

Given $\alpha(s,\tau)$, an effective envelope is given by

$$\boxed{\mathcal{G}^{\varepsilon}(\tau) \;=\; \inf_{s>0} \{\tau \alpha(s,\tau) - \frac{\log \varepsilon}{s}\}}$$

Effective Envelopes and Effective Bandwidth

Now, we can calculate statistical service guarantees for schedulers and traffic types

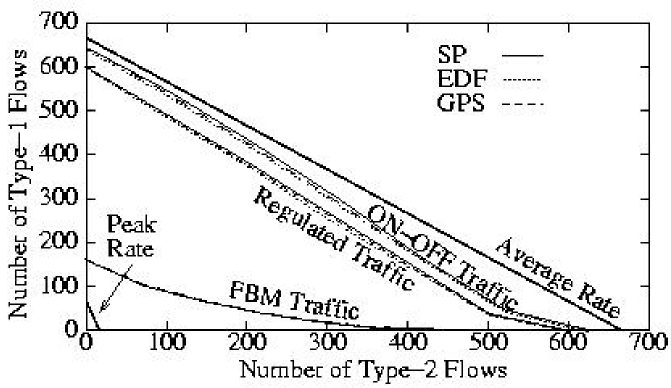

Schedulers:

SP- Static Priority
EDF – Earliest Deadline First
GPS – Generalized Processor Sharing

Traffic:

Regulated – leaky bucket
On-Off – On-off source
FBM – Fractional Brownian Motion

C= 100 Mbps, $\varepsilon = 10^{-6}$

Statistical Network Calculus with Min-Plus Algebra

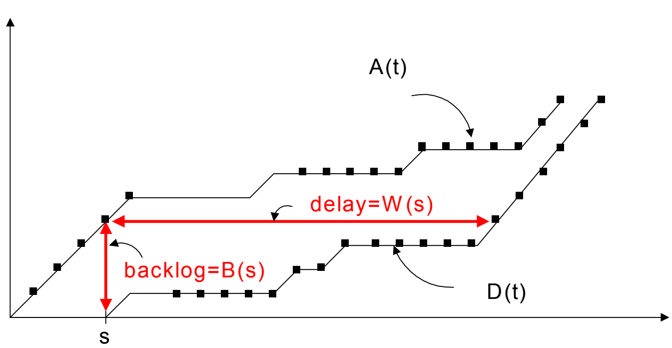

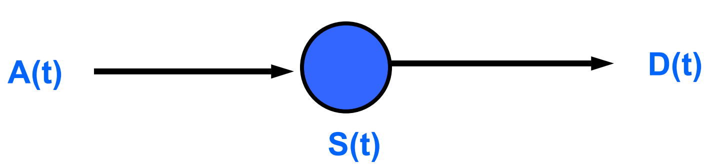

Convolution and Deconvolution operators

- **Convolution operation:**

$$f * g\,(t) = \inf_{\tau \in [0,t]} f(t-\tau) + g(\tau)$$

- **Deconvolution operation**

$$f \otimes g\,(t) = \sup_{\tau \in [0,t]} f(t+\tau) - g(\tau)$$

- **Impulse function:**

$$\delta_\tau(t) = \begin{cases} \infty & , t > \tau \\ 0 & , t \leq \tau \end{cases}$$

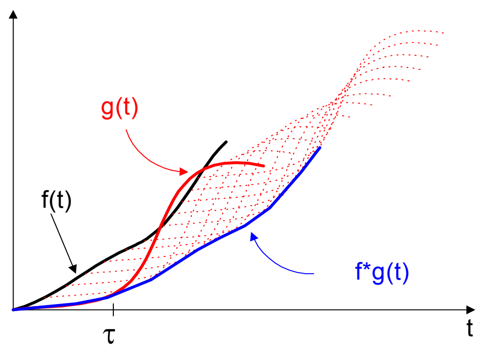

Service Curves (Cruz 1995)

A (minimum) service curve for a flow is a function S such that:

$$D(t) \geq A * S\,(t) \quad , \forall t \geq 0$$

Examples:

- Constant rate service curve: $S\,(t) = c \cdot t$
- Service curve with delay guarantees: $S\,(t) = \delta_d\,(t)$

Network Calculus Main Results (Cruz, Chang, LeBoudec)

1. **Output Envelope:** $A^* \otimes S$ is an envelope for the departures:

$$A^* \otimes S(t) \geq D(t+\tau) - D(\tau)$$

2. **Backlog bound:** $A^* \otimes S(0)$ is an upper bound for the backlog B

3. **Delay bound:** An upper bound for the delay is

$$d_{max} \geq \inf_{\tau \in [0,t]} \left\{ d \geq 0 \mid \forall t \geq 0 : A^*(t-d) \leq S(t) \right\}$$

Network Service Curve (Cruz, Chang, LeBoudec)

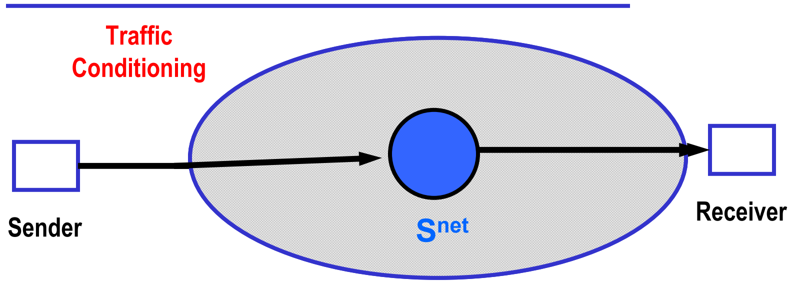

Network Service Curve:

If S^1, S^2 and S^3 are service curves for a flow at nodes, then

$$S^{net} = S^1 * S^2 * S^3$$

is a service curve for the entire network.

2001

Statistical Network Calculus

A (minimum) service curve for a flow is a function S such that:
$$D(t) \geq A * S(t) \quad , \forall t \geq 0$$

A (minimum) **effective service curve** for a flow is a function S^{ε} such that:

$$\Pr\left[D(t) \geq A * S^{\varepsilon}(t)\right] \geq 1 - \varepsilon \quad , \forall t \geq 0$$

2001

Statistical Network Calculus Theorems

1. **Output Envelope:** $A^* \otimes S^{\varepsilon}$ is an envelope for the departures:
$$\Pr\left[A^* \otimes S^{\varepsilon}(t) \geq D(t+\tau) - D(\tau)\right] \geq 1 - \varepsilon \quad , \forall t, \tau \geq 0$$
2. **Backlog bound:** $A^* \otimes S^{\varepsilon}(0)$ is an upper bound for the backlog
$$\Pr\left[B(t) \leq A^* \otimes S^{\varepsilon}(0)\right] \geq 1 - \varepsilon \quad , \forall t \geq 0$$
3. **Delay bound:** A probabilistic upper bound for the delay
$$d_{max} \geq \inf_{\tau \in [0,t]} \left\{ d \geq 0 \mid \forall t \geq 0 : A^*(t-d) \leq S^{\varepsilon}(t) \right\}$$
, i.e.,
$$\Pr\left[W(t) \leq d_{max}\right] \geq 1 - \varepsilon \quad , \forall t \geq 0$$

2002

Effective Network Service Curve

Network Service Curve:

If $S^{1,\varepsilon}$, $S^{2,\varepsilon}$ … $S^{H,\varepsilon}$ are effective service curves for a flow at nodes, then

$$\Pr\left[D(t) \geq A * (S^{1,\varepsilon} * S^{2,\varepsilon} * \ldots * S^{H,\varepsilon} * \underline{\delta_{Ha}})(t)\right] \geq 1 - \underline{\mathrm{Ht}\varepsilon / a}$$

Unfortunately, this network service is not very useful!

A "good" network service curve can be obtained by working with a modified service curve definition

What is the cause of the problem with the network effective service curve?

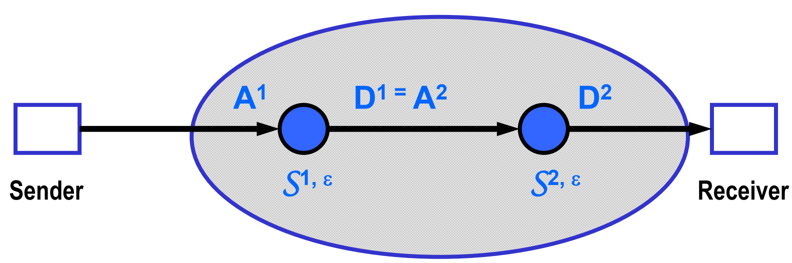

In the convolution

$$D^2(t) \geq A^2 * S^{2,\varepsilon}(t) = \inf_{\tau \in [0,t]} A^2(t-\tau) + S^{2,\varepsilon}(\tau)$$

the range [0,t] where the infimum is taken is a random variable that does not have an a priori bound.

2002

Statistical Per-Flow Service Bounds

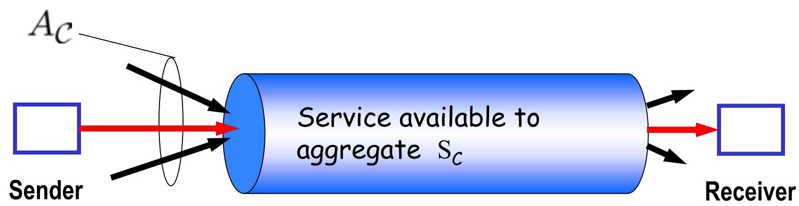

Given:

- Service guarantee to aggregate (S_C) is known
- Total Traffic $A_C = \sum_j A_j$ is known

What is a lower bound on the service seen by a single flow?

2002

Statistical Per-Flow Service Bounds

Can show:

$$\mathcal{S}_j^{\varepsilon_1+\varepsilon_2} = [\mathcal{S}_C - \mathcal{H}_C^{T^{\varepsilon_1},\varepsilon_2}]_+$$

is an effective service curve for a flow where
$\mathcal{H}^{T^{\varepsilon_1},\varepsilon_2}$ is a strong effective envelope and
T^{ε_1} is a probabilistic bound on the busy period

Effective service curve of a single flow

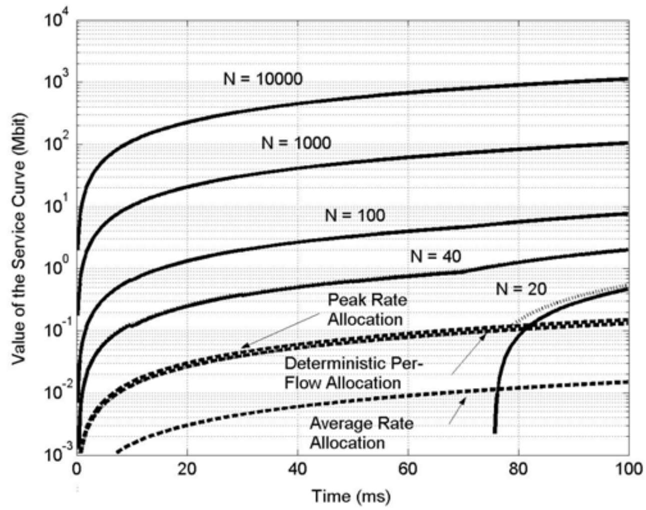

Type 1 flows:

$$S_C(t) = N\hat{c}_i t$$

$$\hat{c}_1 = 1.3140$$

Bandwidth needed by a per-flow allocation to meet a delay bound of d=10ms

2002

Number of flows that can be admitted

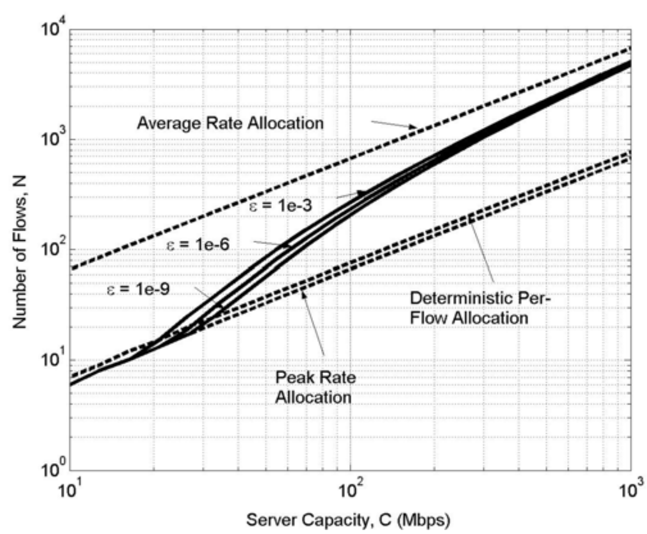

Type 1 flows:

Goal: probabilistic delay bound d=10ms

$$S_C(t) = C\ t$$

Conclusions

- Convergence of deterministic and statistical analysis with new constructs:
  - Effective envelopes
  - Effective service curves
- Preserves much (but not all) of the deterministic calculus
- **Open issues:**
  - So far: Often need bound on busy period or other bound on "relevant time scale".
  - Many problems still open for multi-node calculus

Adaptive service curves

Modified convolution operation

$$A *_{t_o} g(t) = \min\left\{ g(t-t_0), B(t_0) + \inf_{\tau \in [0, t-t_0]} A(t_0, t-\tau) + g(\tau) \right\}$$

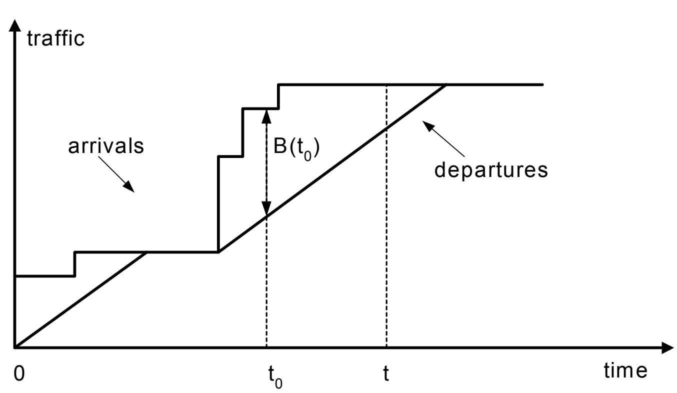

Adaptive service curves

adaptive service curve: $D(t_0,t) \geq A *_{t_o} S(t)$ $\quad \forall t, t_0 \geq 0$

- Many service curves are adaptive (→ Cruz/Okino, LeBoudec)
- Obtain service curve with $t_0=0$

I-adaptive service curve: $D(t_0,t) \geq A *_{t_o} S^{I}(t)$

$$\forall t, t_0 \geq 0, t - t_0 \leq I$$

I-adaptive effective service curve: $\Pr\left[D(t_0,t) \geq A *_{t_o} S^{I,\varepsilon}(t)\right] \geq 1-\varepsilon$

$$\forall t, t_0 \geq 0, t - t_0 \leq I$$

strong (I-adaptive) effective service curve: $\Pr\left[D(t_0,t) \geq A *_{t_o} \mathcal{T}^{I,\varepsilon}(t),\ \forall [t,t_0] \subseteq \mathcal{I}_I\right] \geq 1-\varepsilon$

Effective Network Service Curve

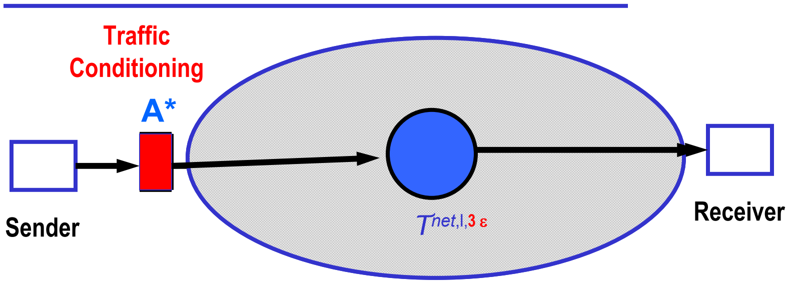

Network Service Curve:

If $\mathcal{T}^{1,I,\varepsilon}$, $\mathcal{T}^{2,I,\varepsilon}$, and $\mathcal{T}^{3,I,\varepsilon}$, are strong effective service curves for a flow at nodes, then

$$\mathcal{T}^{net,I,3\varepsilon} = \mathcal{T}^{1,I,\varepsilon} * \mathcal{T}^{2,I,\varepsilon} * \mathcal{T}^{3,I,\varepsilon}$$

is a service curve for the entire network.

Recover original effective setwork curve

Given a strong effective service curve $\mathcal{T}^{\,l,\varepsilon}$.

If the backlog clears in any time interval of length l with probability ε_1 , i.e,

$$\Pr\left[\exists t_0 \in [t-l,t] : B(t_0) = 0\right] \geq 1 - \varepsilon_1$$

Then $\mathcal{S}^{\,\varepsilon+\varepsilon_1}$ is an effective service curve