Quick Introduction to Max-Plus Network Calculus

ECE 466

University of Toronto

(J. Liebeherr)

Max-Plus Network Calculus

1. Min-plus network calculus: Arrival, departures, service are functions of time.
2. Max-plus network calculus: Arrival, departures, service are functions of space (bits).

3. Functions are related by a reflection at the diagonal!

Representing Arrivals

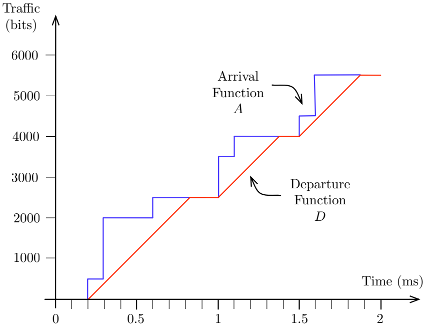

(a) Time Domain.

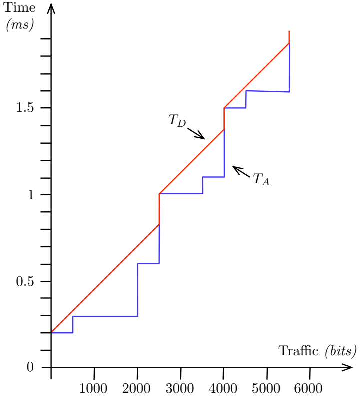

(b) Space Domain.

Definitions

$T_A(\nu)$ Arrival time of bit ν

$T_D(\nu)$ Departure time of bit ν , with $T_D(\nu) \geq T_A(\nu)$

$W(\nu)$ Delay of bit ν:
$W(\nu) = T_D(\nu) - T_A(\nu)$

$B^a(\nu)$ Backlog at arrival of ν:
$B^a(\nu) = \inf\{\kappa > 0 \mid T_D(\nu - \kappa) \leq T_A(\nu)\}$

$B^d(\nu)$ Backlog at departure of ν:
$B^d(\nu) = \inf\{\kappa > 0 \mid T_A(\nu + \kappa) \geq T_D(\nu)\}$

Busy Sequences

A busy sequence is a maximal contiguous set of bits that experience non-zero delays.

Begin of a busy sequence with respect to $\nu \geq 0$ is

$$\underline{\nu} = \sup\{\kappa \mid 0 \leq \kappa \leq \nu \,, W(\kappa) = 0\}\,.$$

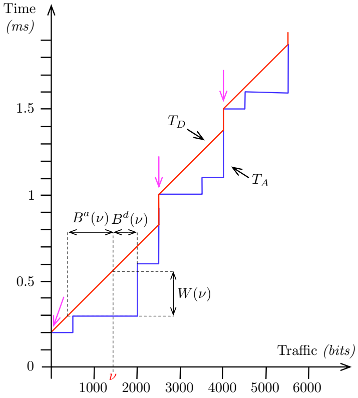

Functions in the Space Domain

A function in the space domain, $F : \mathbb{R} \to \mathbb{R} \cup \{\infty\} \cup \{-\infty\}$

$\mathcal{T}$ = all non-decreasing and right-continuous functions

$\mathcal{T}_o$ = Functions in $\mathcal{T}$, with $F(\nu) = -\infty$ for $\nu < 0$
and $F(\nu) \geq 0$ for $\nu \geq 0$

Meet the $(\max, +)$ algebra

Conventional Algebra		Max-Plus Algebra
Addition $(+)$	$\longrightarrow$	Max $(\max, \vee)$
Multiplication $(\cdot)$	$\longrightarrow$	Addition $(+)$
$\int_{\mathbb{R}} F(s)G(t-s)\,ds$	$\rightarrow$	$\sup_{\kappa\in\mathbb{R}} \{F(\kappa) + G(\nu-\kappa)\}$
$(= F * G(t))$	$\rightarrow$	$(= F \overline{\otimes} G(\nu))$

Burst and delay functions

Burst function

$$\overline{\delta}(\nu) = \begin{cases} -\infty\,, & \nu < 0 \\ 0\,, & \nu \geq 0 \end{cases},$$

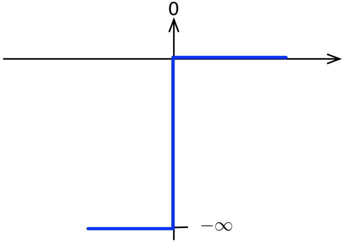

Delay function

$$\overline{\delta}_d(\nu) = \overline{\delta}(\nu) + d$$

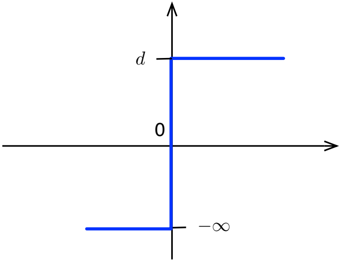

Properties of max-plus convolution

$F, G, H \in \mathcal{T}_o$:

- Closure. $F \overline{\otimes} G \in \mathcal{T}_o$.
- Associativity. $(F \overline{\otimes} G) \overline{\otimes} H = F \overline{\otimes} (G \overline{\otimes} H)$.
- Commutativity. $(F \overline{\otimes} G) \overline{\otimes} H = F \overline{\otimes} (G \overline{\otimes} H)$.
- Distributivity. $(F \vee G) \overline{\otimes} H = (F \overline{\otimes} H) \vee (G \overline{\otimes} H)$.
- Neutral element. $F \overline{\otimes} \overline{\delta} = F$.
- Time shift. $F \overline{\otimes} \overline{\delta}_T(\nu) = F(\nu) + T$.
- Monotonicity. If $F \leq G$ then $F \overline{\otimes} H \leq G \overline{\otimes} H$.
- Boundedness. $F \overline{\otimes} G \geq F$, in particular, $F \overline{\otimes} F \geq F$.

Except boundedness, properties hold for $F, G, H \in \mathcal{T}$

If $F, G \in \mathcal{T}_o$: $\sup_{\kappa \in \mathbb{R}} \{F(\kappa) + G(\nu - \kappa)\} = \sup_{0 \leq \kappa \leq \nu} \{F(\kappa) + G(\nu - \kappa)\}$

(Max-plus) Service Curves

- Buffered link: $\gamma(\nu) = \begin{cases} \frac{\nu}{C}\,, & \nu \geq 0 \\ -\infty\,, & \nu < 0 \end{cases} \Rightarrow T_D(\nu) = T_A \overline{\otimes} \gamma_S(\nu)$
- Service curve is a generalization:

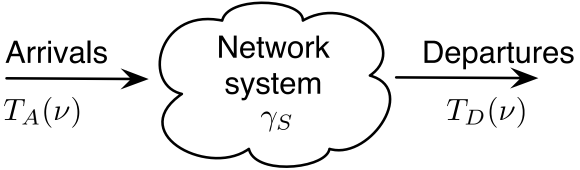

For arbitrary arrivals T_A and resulting departures T_D at a network system, a process $\gamma_S \in \mathcal{T}_o$ is an *exact service curve*, if for all t,

$$T_D(\nu) = T_A \overline{\otimes} \gamma_S(\nu)\,.$$

Lower service curve $\Rightarrow T_D(\nu) \leq T_A \overline{\otimes} \gamma_S(\nu)$
Upper service curve $\Rightarrow T_D(\nu) \geq T_A \overline{\otimes} \gamma_S(\nu)$

About Service Curves

- Service curves express time-invariant and space-invariant service guarantees
- Constant-rate server: $\gamma_{S_1}(\nu) = \frac{\nu}{R}$

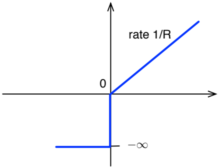

- Delay server: $\gamma_{S_2}(\nu) = T$

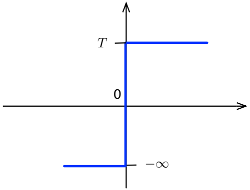

- Latency-rate server: $\gamma_{S_3}(\nu) = \gamma_{S_1} \overline{\otimes} \gamma_{S_2}(\nu) = \frac{\nu}{R} + T$

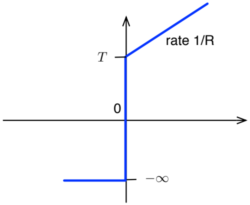

Concatenation of Service Curves

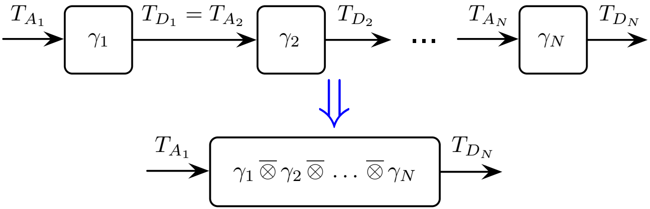

For a sequence of N service elements where the n-th element offers a lower/exact/upper service curve γ_{S_n} $(n = 1, \ldots, N)$, the sequence as a whole offers an lower/exact/upper service curve

$$\gamma_{S_1} \overline{\otimes} \gamma_{S_2} \overline{\otimes} \ldots \overline{\otimes} \gamma_{S_N}$$

The service curve of the sequence is called *network service curve*.

Meet the $(\max, +)$ deconvolution

For two processes $F, G \in \mathcal{T}$, the max-plus *deconvolution* $F \overline{\oslash} G$ is

$$F \overline{\oslash} G(\nu) = \inf_{\kappa \geq 0} \{F(\nu + \kappa) - G(\kappa)\}$$

- If $F, G \in \mathcal{T}_o$, then $\inf_{\kappa \geq 0} \{F(\nu + \kappa) - G(\kappa)\} = \inf_{\kappa \in \mathbb{R}} \{F(\nu + \kappa) - G(\kappa)\}$.
- Weak properties: not closed , not associative, not commutative.

$F, G, H \in \mathcal{T}_o$

- **Composition of $\overline{\otimes}$ and $\overline{\oslash}$.** $(F \overline{\oslash} G) \overline{\oslash} H = F \overline{\oslash} (G \overline{\otimes} H)$.
- **Duality.** $F \leq G \overline{\otimes} H$ if and only if $F \overline{\oslash} H \leq G$.

Superadditivity

A function F is *superadditive*, if for all $\kappa, \nu \in \mathbb{R}$

$$F(\nu + \kappa) \geq F(\nu) + F(\kappa) \ .$$

- F convex $\Rightarrow$ F superadditive
- F superadditive:
  - $F = F \overline{\otimes} F$
  - $F = F \overline{\oslash} F$
- $F \in \mathcal{T}_o$:
  - $F \overline{\oslash} F \in \mathcal{T}_o$
  - $F \overline{\oslash} F$ superadditive

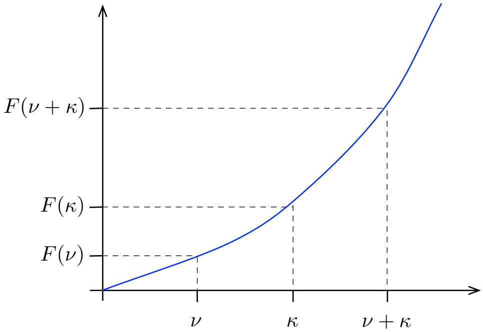

Max-Plus Traffic Envelopes

Traffic envelopes put a lower bound on the amount of time needed for a given amount of traffic

A function λ_E is a max-plus traffic envelope for T_A ($T_A \sim \lambda_E$) if

$$\lambda_E(\mu) \leq T_A(\nu + \mu) - T_A(\nu)\,, \quad \forall \nu \geq 0\,, \forall \mu \geq 0\ .$$

- $T_A \sim \lambda_E \quad \Rightarrow \quad T_A = T_A \overline{\otimes}\, \lambda_E$
- Good envelopes are superadditive

Empirical envelopes

What is the best traffic envelope?

The empirical envelope $\lambda_A^{\mathcal{E}}$ of an arrival function A satisfies

$$\lambda_A^{\mathcal{E}}(\nu) = T_A \overline{\oslash}\, T_A(\nu) \ , \quad \nu \geq 0$$

- $\lambda_A^{\mathcal{E}}$ is superadditive
- $\lambda_A^{\mathcal{E}}$ is largest superadditive function with $\lambda_E \leq T_A$

Traffic shaping (Greedy shaper)

Greedy shaper : A network element that

- limits arrivals to a network to a given specification (traffic envelope),
- buffers non-compliant traffic, and
- releases buffered traffic when it becomes compliant.

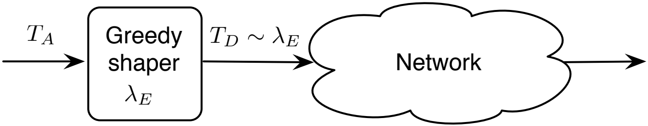

If λ_E is subadditive, greedy shaper offers an exact service curve, i.e.,

$$T_D = T_A \overline{\otimes}\, \lambda_E$$

Example of greedy shaper: Token Bucket

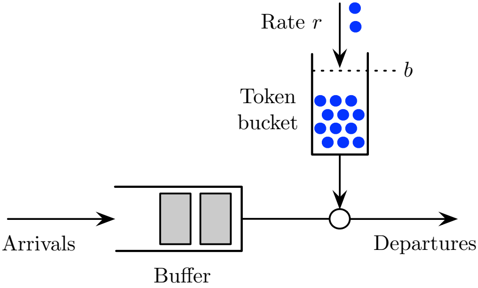

$$\lambda_E(\nu) = \left[\frac{\nu - b}{r}\right]^+$$

Residual Service Curve

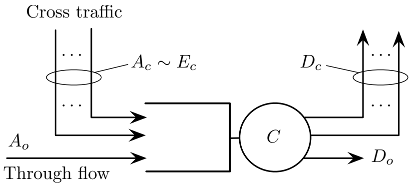

Given a work-conserving buffered link with rate C with a through flow and cross traffic. If λ_c is a max-plus traffic envelope for the cross traffic then

$$\gamma_S(\nu) = \frac{1}{C}\left(\inf\{\mu \geq 0 \mid \lambda_c(\mu) \geq \frac{\nu+\mu}{C}\} + \nu\right)$$

is a lower service curve for the through.

This residual service curve is a (pessimistic) benchmark for the service experienced at a link with multiplexing.

Three Performance Bounds

Given arrival function T_A with traffic envelope λ_E, and a network system with lower service curve γ_S:

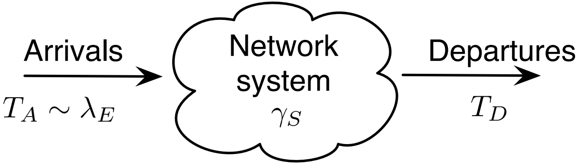

1. Envelope for departures: $T_D \sim [\lambda_E \overline{\oslash} \gamma_S]^+$ (for $\nu \geq 0$)
2. Delay bound: $W(\nu) \leq -\lambda_E \overline{\oslash} \gamma_S(0)$
3. Backlog bound: $B^a(\nu) \leq \inf \{b \geq 0 \,|\, \lambda_E \overline{\oslash} \gamma_S(b) \geq 0\}$

Strict Max-Plus Service Curves

Lower service curves are not good for rate guarantees $\left(\gamma_S(\nu) = \frac{\nu}{R}\right)$

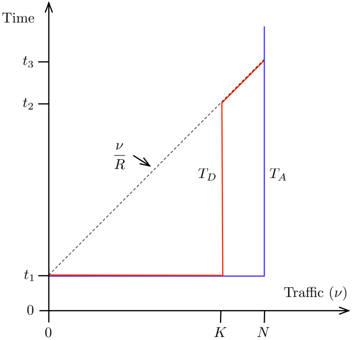

A strict max-plus service curve satisfies for all ν and μ ($\mu < \nu$) in the same busy sequence

$$T_D(\nu) - T_D(\mu) \leq \gamma_S(\nu - \mu)\,,\ \text{if } \underline{\nu} < \mu\,,$$
$$T_D(\nu) - T_A(\mu) \leq \gamma_S(\nu - \mu)\,,\ \text{if } \underline{\nu} = \mu\,.$$

Adaptive Max-Plus Service Curves

Strict service curves are not good for delay guarantees.

Define:

$$F \overline{\underset{\mu}{\otimes}} G(\nu) = \sup_{\mu \le \kappa \le \nu} \{F(\kappa) + G(\nu - \kappa)\}$$

An adaptive max-plus service curve γ_S satisfies for all $\nu \geq 0$ and $\mu \leq \nu$:

$$T_D(\nu) \leq T_D(\mu) + \gamma_S(\nu - \mu) \quad \text{or} \quad T_D(\nu) \leq T_A \overline{\underset{\mu}{\otimes}} \gamma_S(\nu)$$

This can be re-written as a single condition:

$$T_D(\nu) \leq \inf_{\mu \le \nu} \left\{ \max \left[T_D(\mu) + \gamma_S(\nu - \mu) \, , T_A \overline{\underset{\mu}{\otimes}} \gamma_S(\nu) \right] \right\}, \quad \forall \nu \geq 0$$