


Stability and performance in cyclic network

Anne Bouillard (Nokia Bell Labs France)

FORMATS 2019

Performances in networks



Objective: deterministic performance guarantees

Compute the maximum time it takes for a packet to cross the system (Worst-case delay)

Network calculus

Real data

Real input traffic
network element



Delay / backlog

abstraction
→
abstraction
→

pessimism
→

(min,plus) functions

arrival curve
service curve

⇓ (min,plus)-operators

Upper bound on the delay / backlog

Two kinds of pessimism

- ① The abstraction
- ② The (min,plus) operations

Network calculus

- Theory developed in the 1990's by R.L. Cruz, then developed and popularized by C.S. Chang and J.-Y. Le Boudec.
- Filtering theory in the (min,plus) algebra.
- Applications:
 - Internet: video transmission (VoD),
 - Load-balancing in switches [Birkhoff-von Neumann switches, C.S. Chang]
 - Embedded systems: AFDX (Avionics Full Duplex) [Rockwell-Collins software used to certify A380], Networks-on-chip

State of the art and contribution: Feed-forward networks

Many recent results for computing tight bounds feed-forward networks:

- PBOO/PMOO phenomena [Schmitt et al 2008]
- Linear programming solutions [B. et al, 2010]
 - tight bounds
 - the problem is NP-hard
 - polynomial for tandem networks
- Exhaustive search / pay segregation only once [Bondorff et al, 2016]
 - good heuristics to approximate the worst-case performance bounds
- Neural networks [Geyer, 2018]
 - learning the good heuristic

State of the art and contribution: Feed-forward networks

Many recent results for computing tight bounds feed-forward networks:

- PBOO/PMOO phenomena [Schmitt et al 2008]
- Linear programming solutions [B. et al, 2010]
 - tight bounds
 - the problem is NP-hard
 - polynomial for tandem networks
- Exhaustive search / pay segregation only once [Bondorff et al, 2016]
 - good heuristics to approximate the worst-case performance bounds
- Neural networks [Geyer, 2018]
 - learning the good heuristic
- Tight bounds for tree-networks for the backlog of a subsets of flows and delay

State of the art and contribution: Cyclic networks

Few results in networks with cyclic dependencies

- Computing good stability conditions and performance guarantees for network with cyclic dependencies is an open issue
- Obtaining such guarantees would enable more flexible design of systems, with fewer switches.

State of the art and contribution: Cyclic networks

Few results in networks with cyclic dependencies

- Computing good stability conditions and performance guarantees for network with cyclic dependencies is an open issue
- Obtaining such guarantees would enable more flexible design of systems, with fewer switches.
- Flow-based bounds: fix-point/stopped time method
 - "classical" [Cruz 1994]
 - PMOO [Amari et Mifdaoui, 2017]
- Backlog-based bounds: "stability" of the ring [Tassiulas, Georgiadis, 1996]
Additional assumption: the traffic is upper-bounded in each link.
- instability results from adversarial method [Andrews, 2001]

State of the art and contribution: Cyclic networks

Few results in networks with cyclic dependencies

- Computing good stability conditions and performance guarantees for network with cyclic dependencies is an open issue
- Obtaining such guarantees would enable more flexible design of systems, with fewer switches.
- Flow-based bounds: fix-point/stopped time method
 - "classical" [Cruz 1994]
 - PMOO [Amari et Mifdaoui, 2017]
- Backlog-based bounds: "stability" of the ring [Tassiulas, Georgiadis, 1996]
Additional assumption: the traffic is upper-bounded in each link.
- instability results from adversarial method [Andrews, 2001]
- Improve the fix-point method to combine flow and backlog-based bounds

Network calculus framework


Networks with cyclic dependencies

Computing performance bounds in feed-forward networks

Performances in cyclic network

Conclusion

Cumulative processes




- $A : \mathbb{R}_+ \rightarrow \mathbb{R}_{\min+}$: process of the cumulative arrivals, non-decreasing function
- $B : \mathbb{R}_+ \rightarrow \mathbb{R}_{\min+}$: process of the cumulative departures, non-decreasing function
- Causality constraint: $A \geq B$

Arrival and service curves



Arrival curve

A is constrained by the function α if
 $\forall 0 \leq s \leq t,$
 $A(t) - A(s) \leq \alpha(t - s).$




Arrival and service curves



Arrival curve

A is constrained by the function α if
 $\forall 0 \leq s \leq t,$
 $A(t) - A(s) \leq \alpha(t - s).$




Arrival and service curves



Arrival curve

A is constrained by the function α if
 $\forall 0 \leq s \leq t,$
 $A(t) - A(s) \leq \alpha(t - s).$




Arrival and service curves




Arrival curve

A is constrained by the function α if
 $\forall 0 \leq s \leq t,$
 $A(t) - A(s) \leq \alpha(t - s).$



Strict service curve

A network element guarantees β for A if, while system not empty, B satisfies
 $B(t) \geq B(s) + \beta(t - s).$




Arrival and service curves





Arrival curve

A is constrained by the function α if
 $\forall 0 \leq s \leq t,$
 $A(t) - A(s) \leq \alpha(t - s).$



Strict service curve

A network element guarantees β for A if, while system not empty, B satisfies
 $B(t) \geq B(s) + \beta(t - s).$




From constraints to performance bounds

Maximum backlog:

$$B_{\max} = \sup_{t \geq 0} A(t) - B(t)$$

Maximum delay:


$$D_{\max} = \inf\{d \mid \forall t \in \mathbb{R}_+, B(t+d) \geq A(t)\}$$



Performance bounds

- $B_{\max} \leq \alpha \oslash \beta(0) = v(\alpha, \beta) = \sup\{\alpha(t) - \beta(t) \mid t \geq 0\}$
- $D_{\max} \leq h(\alpha, \beta) = \inf\{\forall t \geq 0, d \geq 0 \mid \alpha(t) \leq \beta(t+d)\}$ (for FIFO per flow).

Model and hypotheses



Hypotheses

- m token-bucket arrival curves: $\alpha_i(t) = b_i + r_i t$;
- n rate-latency strict service curves: $\beta^{(j)}(t) = R_j(t - T_j)_+$.

Network calculus framework

Networks with cyclic dependencies

Computing performance bounds in feed-forward networks

Performances in cyclic network

Conclusion

Stability in cyclic networks


Consider a server offering a strict service curve $\beta : t \mapsto R(t - T)_+$ and a flow crossing it, with arrival curve $\alpha : t \mapsto b + rt$.

- This server is said unstable if its worst-case backlog is unbounded: $R < r$;
- This server is said critical if its worst-case backlog is bounded, but the lengths of its backlogged periods are not bounded: $R = r$;
- This server is said stable if the length of its backlogged periods is bounded: $R > r$.

Definition (Global stability)

A network is globally stable if for all its servers, the length of the maximal backlogged period is bounded.

Fix-point method



(service curves and arrival curves of exogenous arrivals are constants of the problem)

$$\alpha_2^{(3)} = H_2^{(1)}(\alpha_1, \alpha_2)$$

$$\alpha_2^{(4)} = H_2^{(3)}(\alpha_1^{(3)}, \alpha_3^{(3)}) \dots$$

We write this equation for each output flow at each server and obtain a system


$$\alpha = \mathbf{H}(\alpha)$$

Lemma

If the system is stable, then there exists a family $\alpha = (\alpha_{i,j})_{i,j}$ of arrival curves for the flows $(F_i^{(j)})$ such that $\alpha \leq \mathbf{H}(\alpha)$.

Take the best arrival curves, they will satisfy every inequality.

Fix-point method



(service curves and arrival curves of exogenous arrivals are constants of the problem)

$$\alpha_2^{(3)} = H_2^{(1)}(\alpha_1, \alpha_2)$$
$$\alpha_2^{(4)} = H_2^{(3)}(\alpha_1^{(3)}, \alpha_3^{(3)}) \dots$$

We write this equation for each output flow at each server and obtain a system

$$\alpha = \mathbf{H}(\alpha)$$

- If service curves are rate-latency and arrival curves token bucket, this is a linear equation: $\mathbf{b} = \mathbf{M}\mathbf{b} + \mathbf{N}$.

Network calculus framework

Networks with cyclic dependencies


Computing performance bounds in feed-forward networks

Performances in cyclic network

Conclusion

Generic (min,plus) method

- ① In the topological order of the servers, for each flow crossing the server:




Residual service curve:

$$\beta_1^{(j)} = (\beta^{(j)} - \alpha_c^{(j)})_+$$

Output arrival curve:

$$\alpha_1^{(j+1)} = \alpha_1^{(j)} \oslash \beta_1^{(j)}$$

- ② For the flow of interest



End-to-end service curve:

$$\beta = \beta_1^{(1)} * \beta_1^{(2)} * \dots * \beta_1^{(k)}$$

- ③ Delay bound: $h(\alpha_1, \beta)$,
Backlog bound: $v(\alpha_1, \beta)$

Tight worst-case delays for tandem networks

Joint work with Thomas Nowak [Performance 2015]



Theorem


Consider a tandem network of n servers. The worst-case delay is linear in the bursts and latencies:

$$D = \sum_{j \in \mathbb{N}_n} \lambda_j T_j + \sum_{i \in \mathbb{N}_m} \mu_i b_i$$

where the coefficients λ_j and μ_i depend only on the arrival and service rates and can be effectively computed in time $O(n^2 + m)$.

Tight worst-case delays and backlog for tree networks

This theorem can be adapted to backlog at server n and for tree-topologies:



Theorem

Consider a **tree** network of n servers, and p **flows of interest** at server n . The worst-case backlog at server n for the flows of interests is linear in the bursts and latencies:

$$B = \sum_{j \in \mathbb{N}_n} \rho_j T_j + \sum_{i \in \mathbb{N}_m} \varphi_i b_i + \sum_{i \in \mathbb{N}_p} b_i^*$$

where the coefficients ρ_j and $\varphi_i < 1$ depend only on the arrival and service rates and can be effectively computed in time $O(n^2 + m + p)$.

Sketch of the proof

- ① Find properties of a worst-case scenario for the network
 - SDF (Shortest-to-destination first) service policy
 - one backlogged period for each server
 - minimum service
 - maximum arrivals
- ② Backward induction on the servers
 - Only the dates of the start of backlogged periods need to be computed
 - they depend only on the amount of data transmitted at those time.

Comparison of the approaches

① (min,plus):

- Efficient algorithms (linear in the total length of the flows)
- Pessimistic performance bounds (as soon as two servers and two flows)
- Linear in T_j and b_i

② Our approach:

- Quadratic algorithm in tree network
- Tight delay bound
- Linear in T_j and b_i

Network calculus framework


Networks with cyclic dependencies

Computing performance bounds in feed-forward networks

Performances in cyclic network

Conclusion


Flow-based decomposition of the network



- (min,plus) approach: at each server for each flow

$$b_{i,k+1} \leq b_{i,k} + \frac{r_i}{R_j - \sum_{p \in Fl(j) \setminus \{i\}} r_p} \left(\sum_{s \in S_j \setminus \{(i,k)\}} b_s + R_j T_j \right).$$

Flow-based decomposition of the network




- (min,plus) approach: at each server for each flow

$$b_{i,k+1} \leq b_{i,k} + \frac{r_i}{R_j - \sum_{p \in \text{Fl}(j) \setminus \{i\}} r_p} \left(\sum_{s \in S_j \setminus \{(i,k)\}} b_s + R_j T_j \right).$$

- Our approach, with a tree decomposition:

$$b_{i,k+1} \leq \sum_{s \in S} \varphi_s^{i,k+1} b_s + \sum_{\{j | j \rightsquigarrow j_1\}} \rho_j^{i,k+1} T_j,$$

Flow-based decomposition of the network



Linear problem:

Maximize $Q\mathbf{b} + C$
such that $\mathbf{b} \leq M\mathbf{b} + N$.

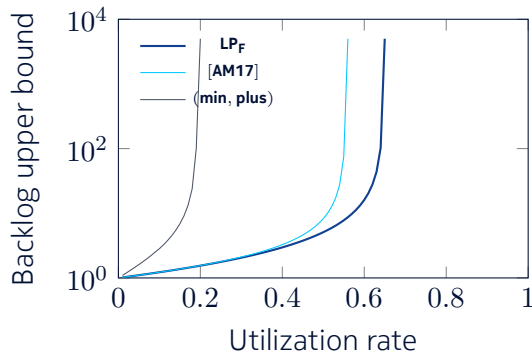

- (min,plus) approach: at each server for each flow

$$b_{i,k+1} \leq b_{i,k} + \frac{r_i}{R_j - \sum_{p \in Fl(j) \setminus \{i\}} r_p} \left(\sum_{s \in S_j \setminus \{(i,k)\}} b_s + R_j T_j \right).$$

- Our approach, with a tree decomposition:

$$b_{i,k+1} \leq \sum_{s \in S} \varphi_s^{i,k+1} b_s + \sum_{\{j | j \rightsquigarrow j_1\}} \rho_j^{i,k+1} T_j,$$

Comparison on the ring network



$n = 10$ servers, $R = 100\text{Mb.s}^{-1}$, $L = 1\text{ms}$,
 $b = 1\text{Mb}$, $r = \frac{uR}{n}$

Backlog-based decomposition

Idea:

- worst-case bounds for flows ending at the same server do not happen at the same time;
- computing worst-case backlog of all flows following an arc might takes this phenomenon into account.

$$\begin{aligned} B_a &\leq \sum_{s \in S} \varphi_s^a x_s + \sum_{\{j|j \rightsquigarrow j_1\}} \rho_j^a T_j \\ &\leq \sum_{a' \in \mathbb{A}'} \left[\left(\max_{s \in S'_{a'}} \varphi_s^a \right) \left(\sum_{s \in S'_{a'}} x_s \right) \right] + \sum_{i=1}^m \varphi_{(i,1)}^a b_i + \sum_{\{j|j \rightsquigarrow j_1\}} \rho_j^a T_j \\ &\leq \sum_{a' \in \mathbb{A}'} \left(\max_{s \in S'_{a'}} \varphi_s^a \right) B_{a'} + \sum_{i=1}^m \varphi_{(i,1)}^a b_i + \sum_{\{j|j \rightsquigarrow j_1\}} \rho_j^a T_j. \end{aligned}$$

Ring stability revisited

Theorem (TG96, LT04)

“The ring is stable” under assumption for stability of each server
Additional assumption: the traffic is upper-bounded in each link.


Our approach

$$B = \sum_{j \in \mathbb{N}_n} \rho_j T_j + \sum_{i \in \mathbb{N}_m} \varphi_i b_i + \sum_{i \in \mathbb{N}_p} b_i^*$$

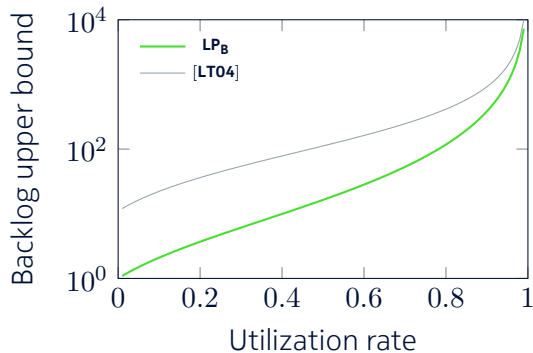
where the coefficients ρ_j and $\varphi_i < 1$ depend only on the arrival and service rates.

$$B \leq C + \varphi B$$

where $\varphi = \sup \varphi_j^n < 1$ and $B \leq \frac{C}{1-\varphi}$.



Comparison on the Ring Network



$n = 10$ servers, $R = 100\text{Mb.s}^{-1}$, $L = 1\text{ms}$, $b = 1\text{Mb}$, $r = \frac{uR}{n}$

Combining Flow-based and Backlog-based bounds

New set of linear constraints:


$$\mathcal{L} = \left\{ \begin{array}{ll} b_s \leq \sum_{s' \in S} \varphi_{s'}^s x_{s'}^s + C_s, & \forall s \in S \\ B_a \leq \sum_{s' \in S} \varphi_{s'}^a x_{s'}^a + C_a, & \forall a \in \mathbb{A}^r \\ 0 \leq x_{s'}^s \leq b_{s'}, & \forall s' \in S, s \in S \cup \mathbb{A}^r \\ \sum_{s' \in \mathbb{A}^r} x_{s'}^s \leq B_a, & \forall a \in \mathbb{A}^r, s \in S \cup \mathbb{A}^r \end{array} \right\},$$

Number of constraints:

- Flow-based: $O(k)$
- Backlog-based: $O(a)$
- Combination: $O(k + a)^2$



where k is the number of flows in the tree-decomposition and a the number of arcs removed for this decomposition.

Comparison on the Ring Network



$$n = 10 \text{ servers}, R = 100\text{Mb.s}^{-1}, L = 1\text{ms}, b = 1\text{Mb}, r = \frac{uR}{n}$$

Two-Ring network



Network calculus framework

Networks with cyclic dependencies

Computing performance bounds in feed-forward networks

Performances in cyclic network

Conclusion

Conclusion and futurework

Conclusion

- A new efficient algorithm to compute tight worst-case delays and backlog.
- Application to networks with cyclic dependencies:
 - best stability conditions
 - stability of the ring without additional assumptions
- Implementation in a Python package: <https://github.com/nokia/NCBounds>

Future work

- Extension to feed-forward networks (we conjecture that a simple generalization can lead to the same approximation with one linear program)
- what is the best decomposition?
- Extension to some service policies (FIFO, GPS for example), maximum service rate...

NOKIA