Layering, dynamics, optimization & control in software defined networks

Nikolai Matni
Computing and Mathematical Sciences
Caltech

joint work with John Doyle (Caltech) and Kevin Tang (Cornell)

Distributed optimal control & software defined networks

Nikolai Matni
Computing and Mathematical Sciences
Caltech

joint work with John Doyle (Caltech) and Kevin Tang (Cornell)

Control theory

Using feedback to mitigate the effects of dynamic uncertainty on a system

“open loop”

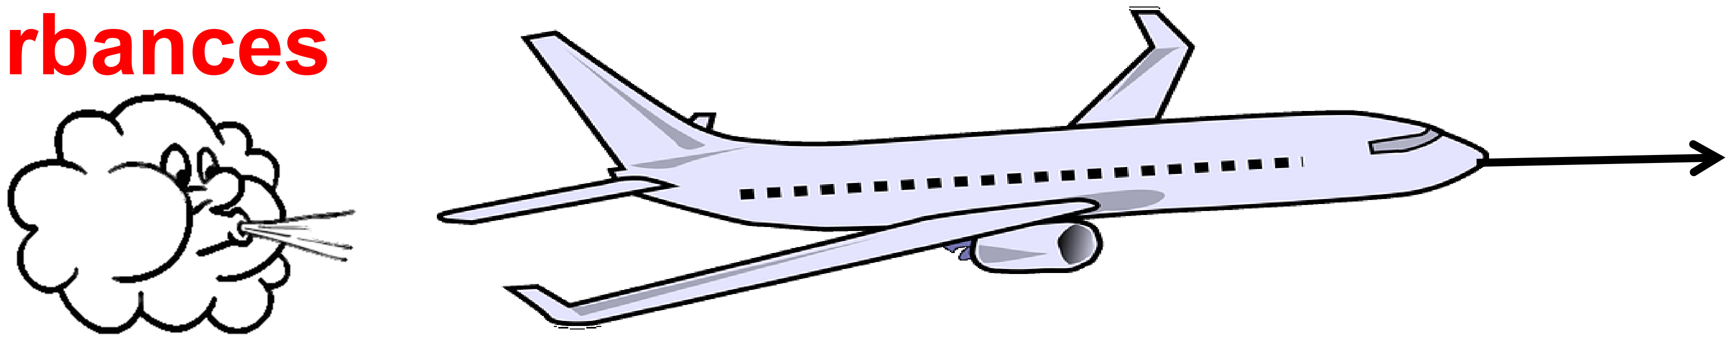

Control theory

$$m\left[\frac{du}{dt} + qw - rv\right] = F_x - mg\sin\Theta$$

$$m\left[\frac{dv}{dt} + ru - pw\right] = F_y + mg\cos\Theta\sin\Phi$$

$$m\left[\frac{dw}{dt} + pv - qu\right] = F_z + mg\cos\Theta\cos\Phi$$

$$M_x = I_{xx}\frac{dp}{dt} - I_{xz}\frac{dr}{dt} + [I_{zz} - I_{yy}]qr - I_{xz}pq$$

$$M_y = I_{yy}\frac{dq}{dt} + [I_{xx} - I_{zz}]rp + I_{xz}[p^2 - r^2]$$

$$M_z = -I_{xz}\frac{dp}{dt} + I_{zz}\frac{dr}{dt} + [I_{yy} - I_{xx}]pq +$$

$$p = \frac{d\Phi}{dt} - \frac{d\Psi}{dt}\sin\Theta$$

$$q = \frac{d\Theta}{dt}\cos\Phi + \frac{d\Psi}{dt}\sin\Phi\cos\Theta$$

$$r = -\frac{d\Theta}{dt}\sin\Phi + \frac{d\Psi}{dt}\cos\Phi\cos\Theta$$

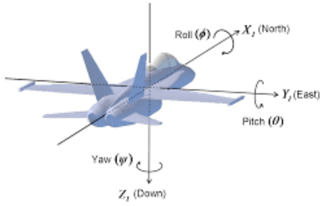

feedback to close the loop

Controller

measurements

control action

Actuators

Sensors

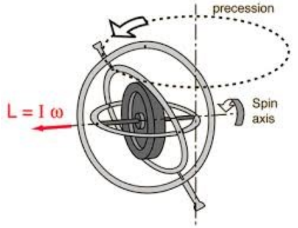

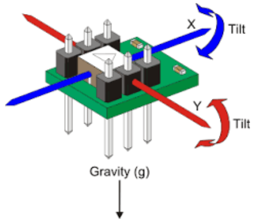

disturbances

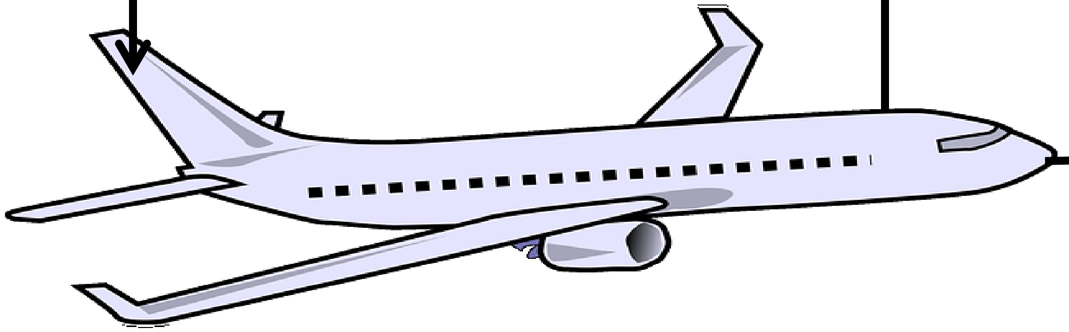

set-point

plant

A familiar example: TCP

$$\max \sum_s U(x_s)$$
$$\text{s.t. } Rx \leq c$$

Feedback: AIMD based on drops, RTT, queue length, etc.

Guarantees: **converge** to NUM optimizing transmission rates

Only steady state guarantees

Stabilizing controllers

Controller #1

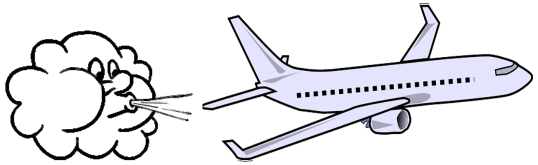

Controller #2

Optimal control theory

minimize **worst-case** **amplification**

$$\underset{u}{\text{minimize}} \quad \underset{w \in \mathcal{W}}{\text{maximize}} \quad f\left(\Delta x, u, w\right)$$

error — control action — disturbance

Bounded energy

$\mathcal{W} =$

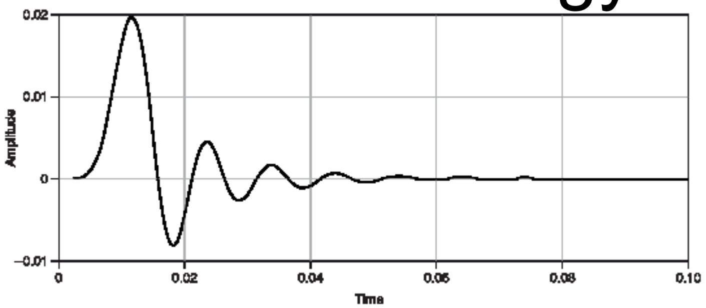

Bounded magnitude

$\mathcal{W} =$

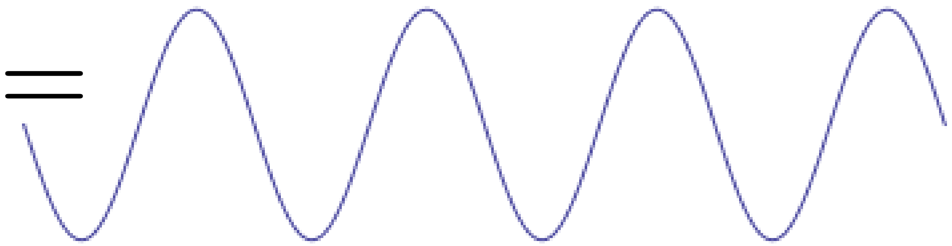

Optimal control theory

minimize **average** amplification

$$\underset{u}{\text{minimize}} \quad \text{Expected Value} \quad \left[f\left(\Delta x, u, w\right) \right]$$

error

control action

disturbance

white noise

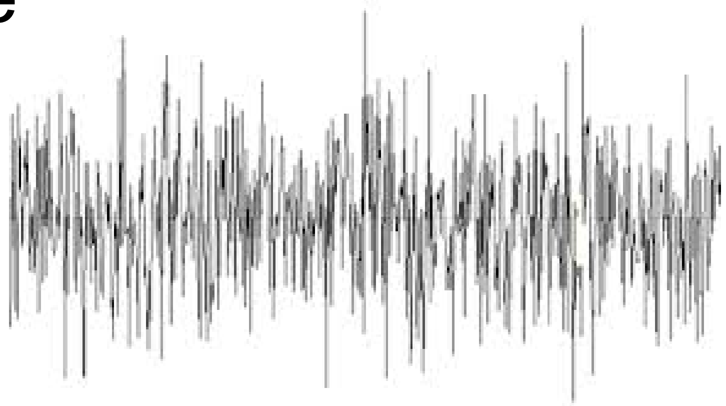

Centralized control

Can lead to poor performance for large-scale systems

One system

Global access to measurements

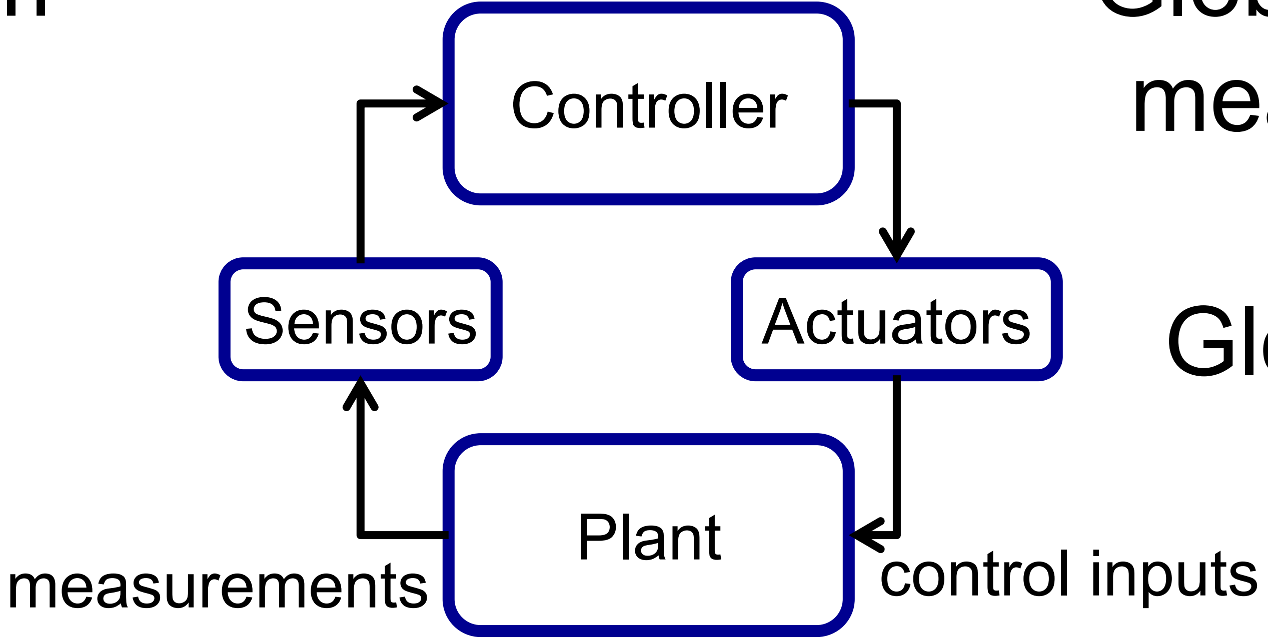

One controller

Global control of inputs

Distributed control

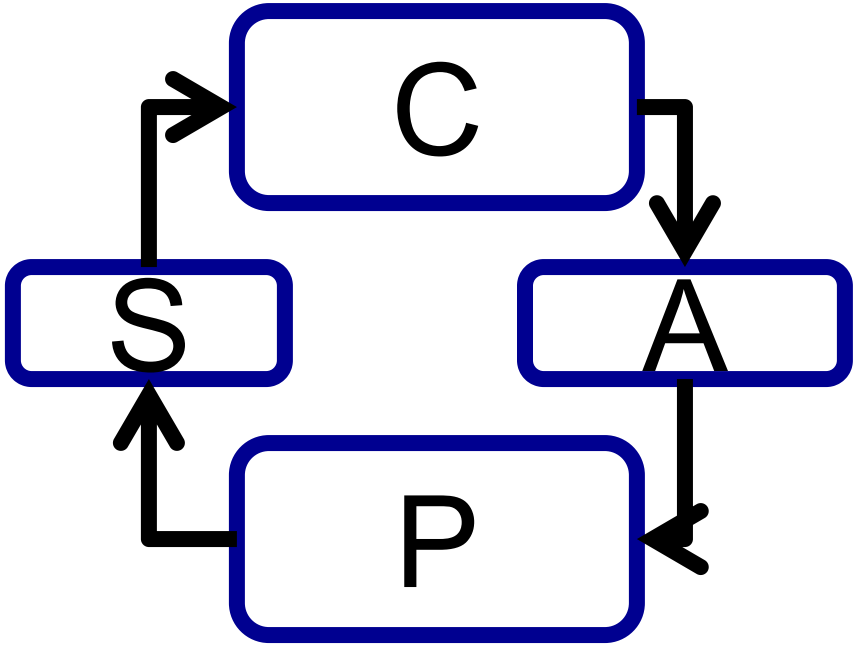

Distributed control

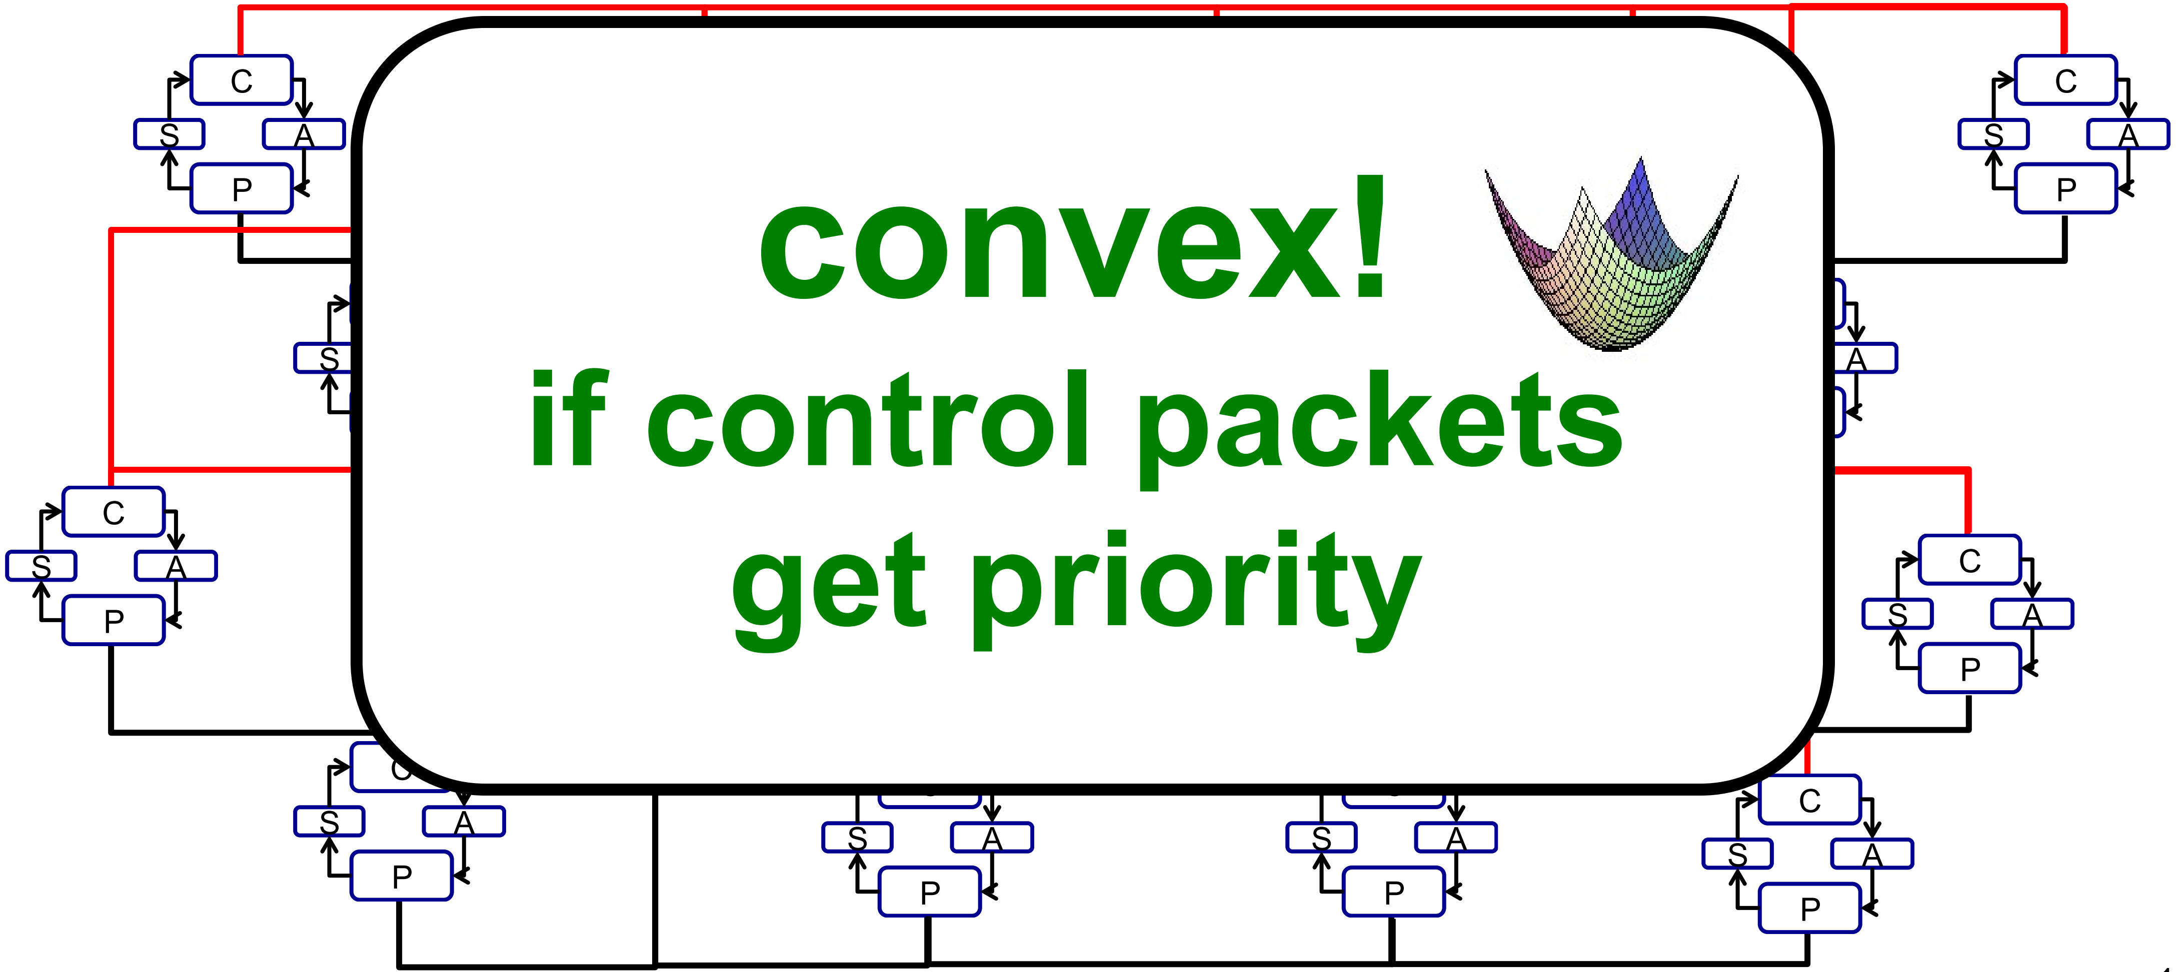

convex!

if control packets get priority

Distributed optimal control in WANs

WAN distributed optimal control

nominal flows

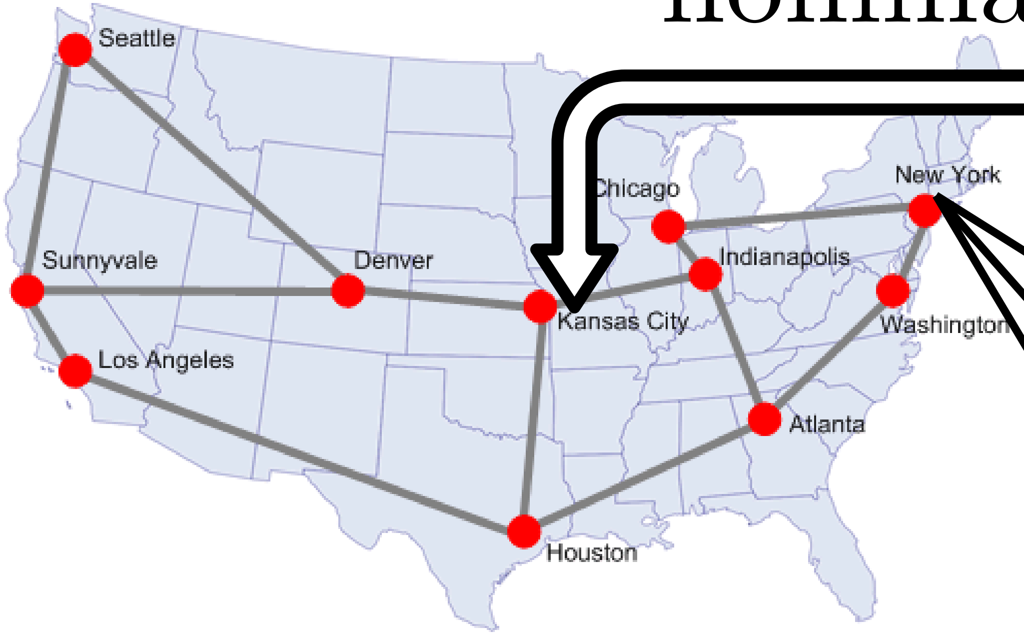

$\{f_\ell^*\}$

TE solved using nominal demands

$$D_{s,d}(t) = D_{s,d} + w(t)$$

real traffic **fluctuates** around nom rates

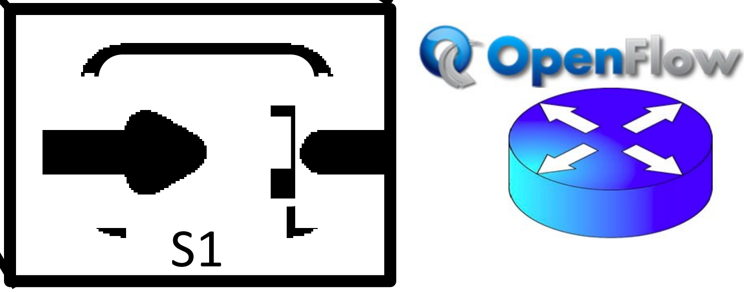

High Frequency Traffic Control

rate deviation — queue length

$$\underset{\Delta u_l(n)}{\text{minimize}} \lim_{n\to\infty} \sum_l \mathbb{E}[\Delta f_l(n)]^2 + \lambda_l \mathbb{E}[b_l(n)]^2$$

egress buffer control

Generalizes FIFO/Smoothing

rate deviation — **queue length** — **egress buffer control**

$$\underset{\Delta u_l(n)}{\text{minimize}} \lim_{n\to\infty} \sum_l \mathbb{E}[\Delta f_l(n)]^2 + \lambda_l \mathbb{E}[b_l(n)]^2$$

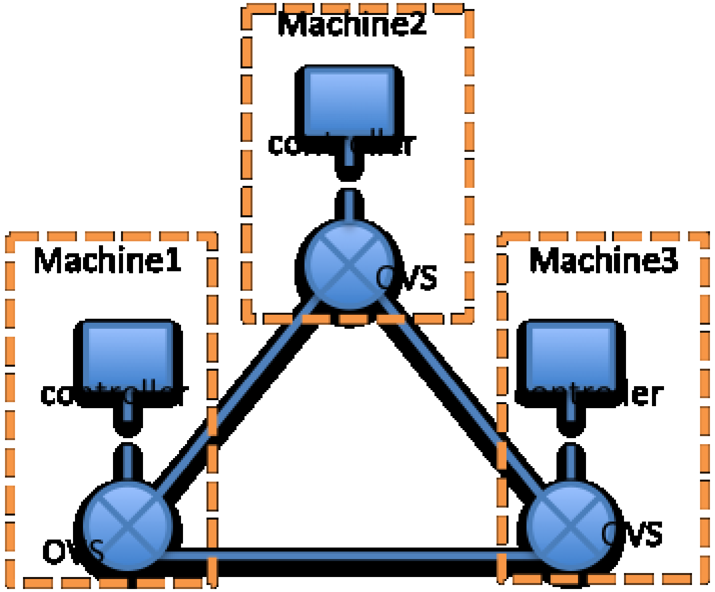

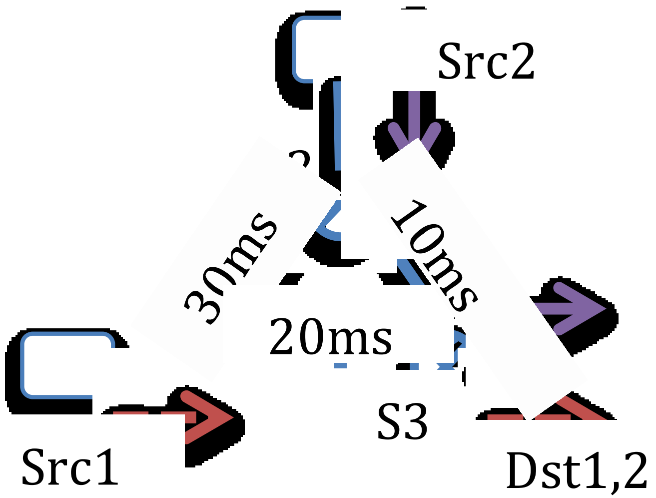

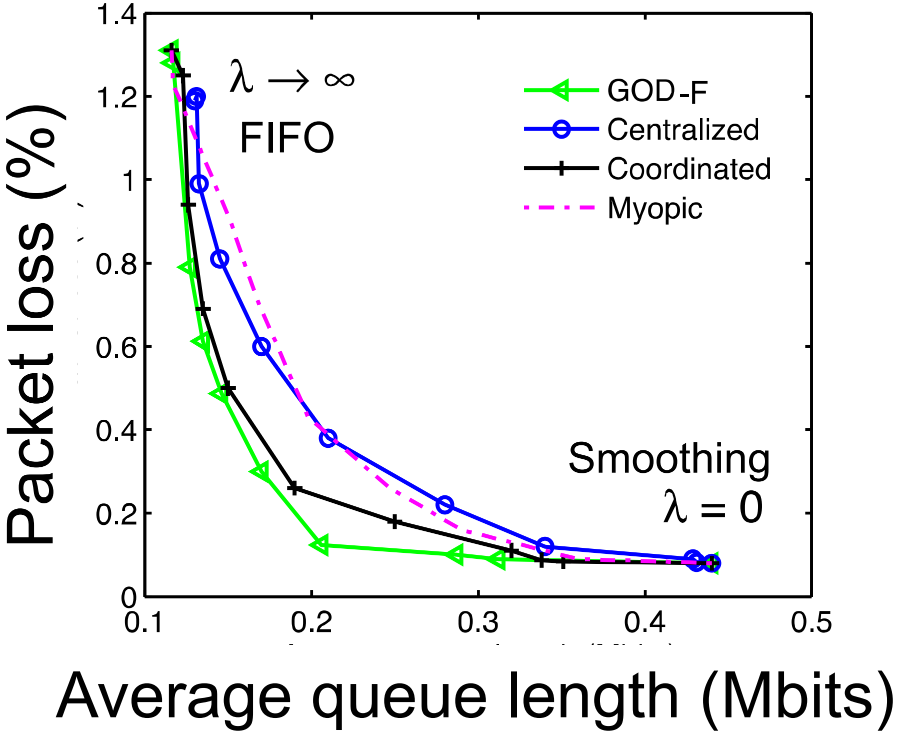

Controller architectures

High Frequency Traffic Control

rate deviation queue length

$$\underset{\Delta u_l(n)}{\text{minimize}} \lim_{n \to \infty} \sum_l \mathbb{E}[\Delta f_l(n)]^2 + \lambda_l \mathbb{E}[b_l(n)]^2$$

egress buffer control

Globally Optimal Delay Free (GOD-F)

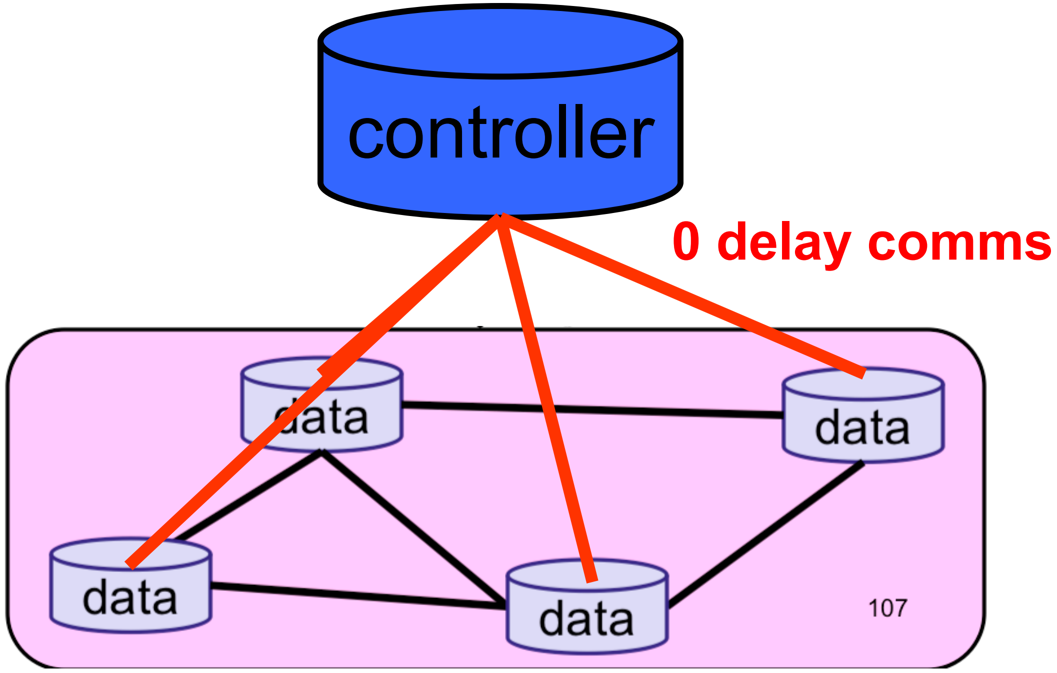

Controller architectures

High Frequency Traffic Control

rate deviation　queue length

$$\underset{\Delta u_l(n)}{\text{minimize}} \lim_{n\to\infty} \sum_l \mathbb{E}[\Delta f_l(n)]^2 + \lambda_l \mathbb{E}[b_l(n)]^2$$

egress buffer control

Centralized

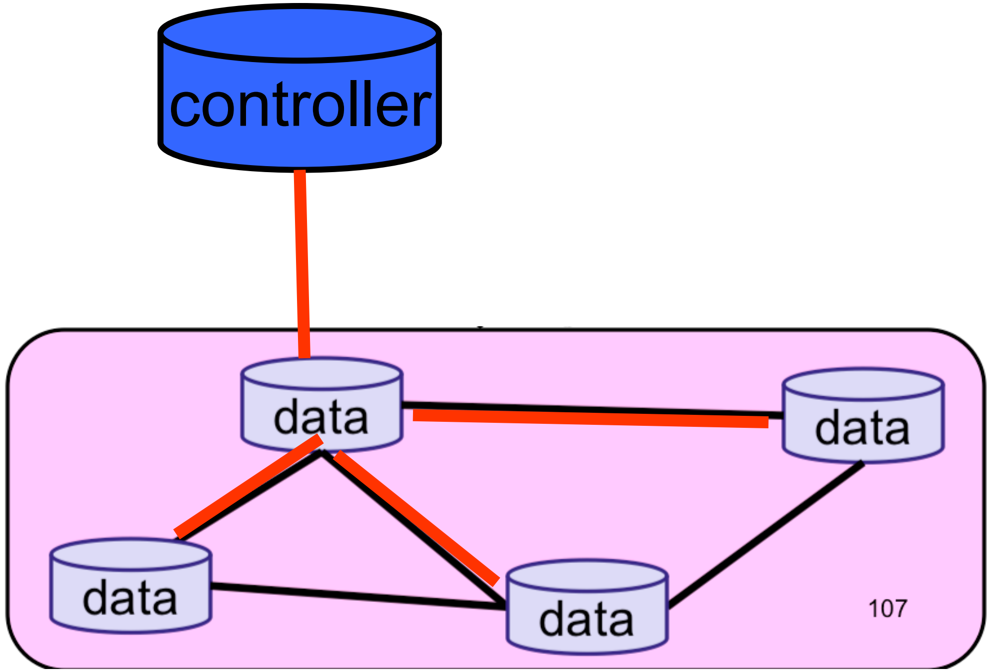

Controller architectures

High Frequency Traffic Control

rate deviation — queue length

$$\underset{\Delta u_l(n)}{\text{minimize}} \lim_{n\to\infty} \sum_l \mathbb{E}[\Delta f_l(n)]^2 + \lambda_l \mathbb{E}[b_l(n)]^2$$

egress buffer control

Distributed

new theory

ctrl packets get priority = convex!

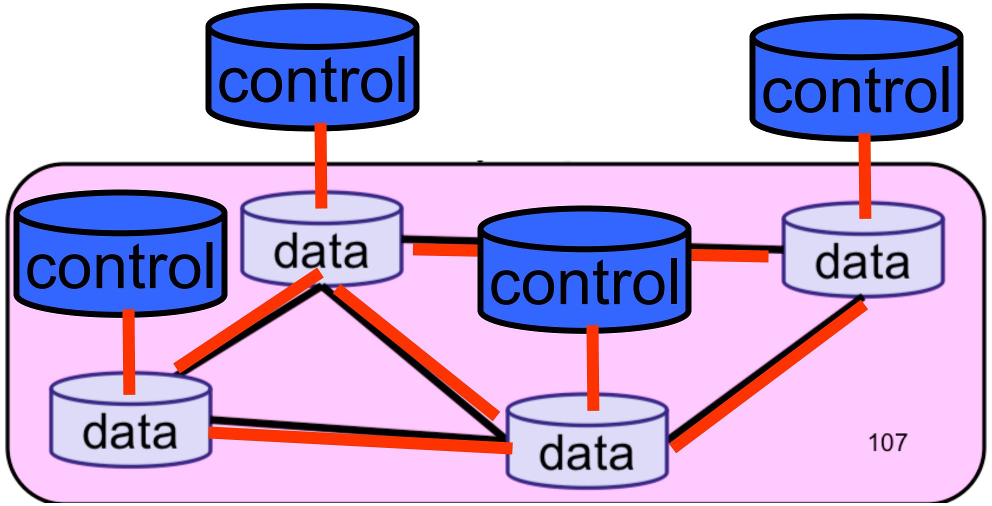

Controller architectures

High Frequency Traffic Control

rate deviation **queue length**

$$\underset{\Delta u_l(n)}{\text{minimize}} \lim_{n\to\infty} \sum_l \mathbb{E}[\Delta f_l(n)]^2 + \lambda_l \mathbb{E}[b_l(n)]^2$$

egress buffer control

Decentralized (myopic)

non-convex

use best guess

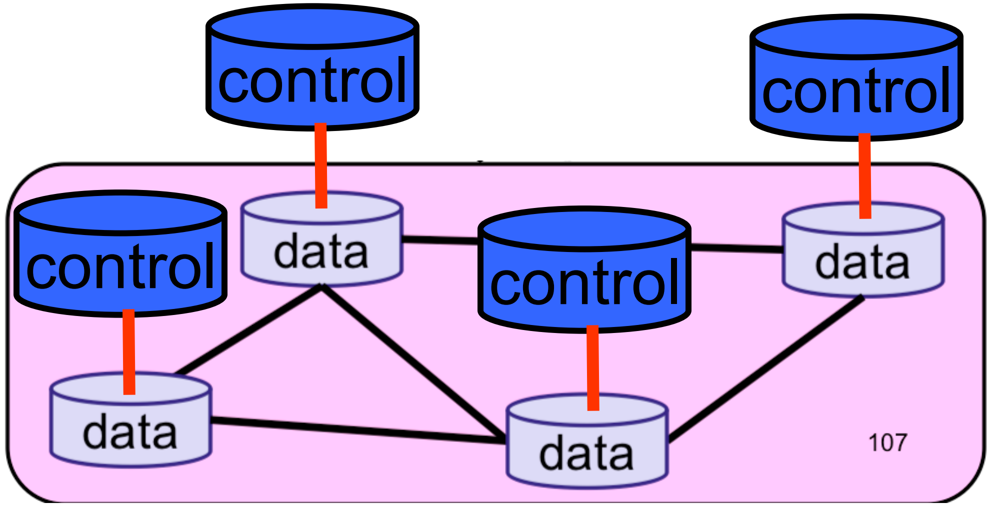

WAN reflex layer

centralized ⇩ decentralized

most delay ⇧ least delay

theory & experiments are consistent!

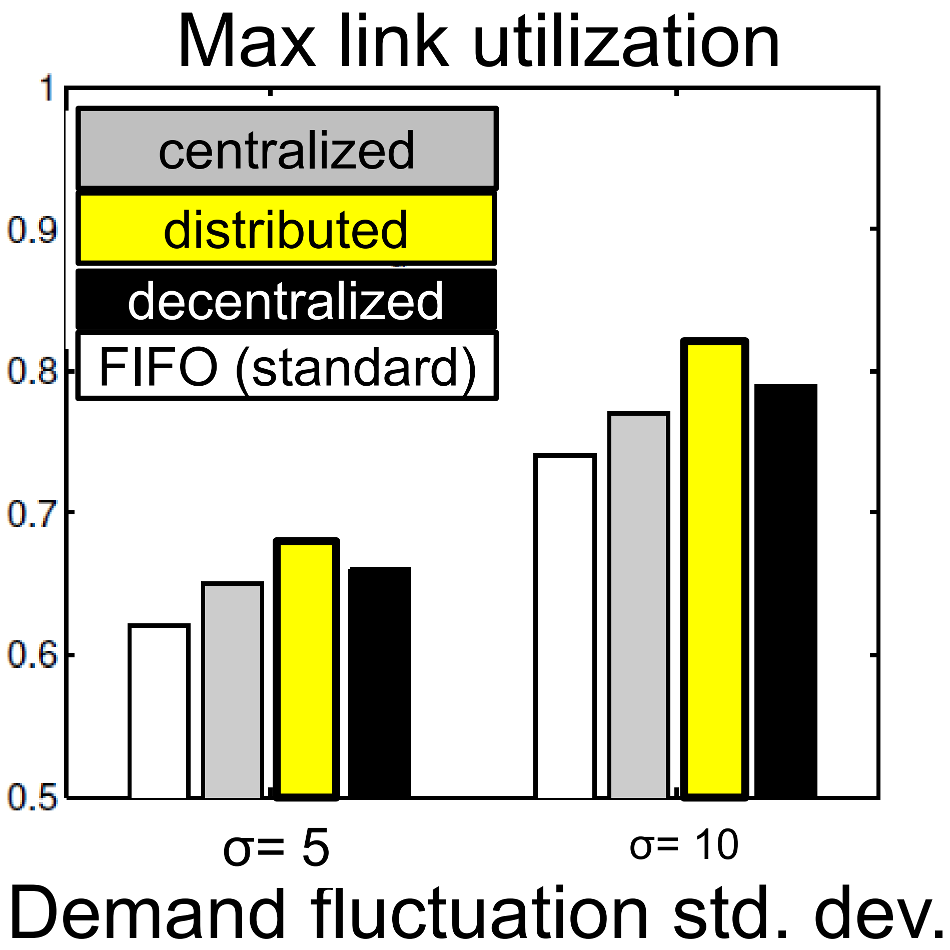

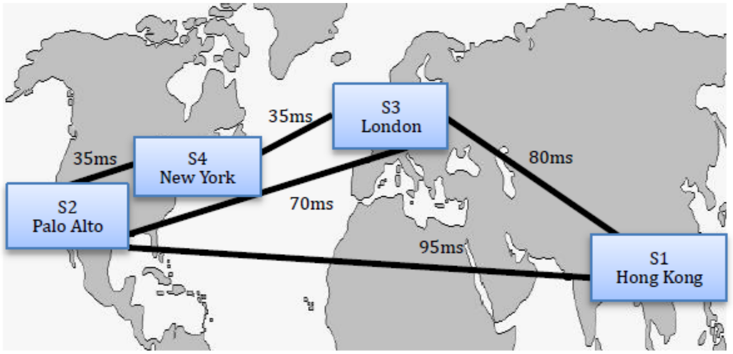

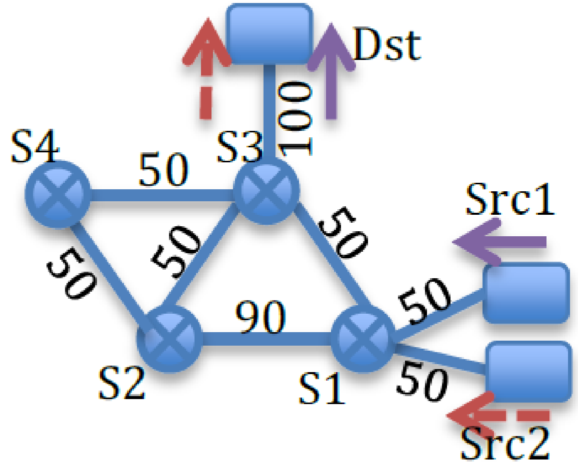

N. Wu, A. Tang, J.C. Doyle, …, **N. Matni,** in preparation
N. Matni, A. Tang, J.C. Doyle, ACM SIGCOMM Symposium on SDN Research, 2015

A theory of network architecture

Can we

- d
- la design?
- c
  o

Must model
dynamics & delay

Caltech Nikolai Matni

Select references:

- N. Matni & J. C. Doyle, A theory of dynamics, control and optimization in layered architectures, IEEE American Control Conference, 2016
- N. Matni, A. Tang & J. C. Doyle, A case study in network architecture tradeoffs, ACM Symposium on SDN Research (SOSR), 2015

High-frequency traffic control preprint available upon request.