



Performance-oriented Congestion Control

Mo Dong, Qingxi Li, Doron Zarchy*,
P. Brighten Godfrey and Michael Schapira*

University of Illinois at Urbana Champaign


*Hebrew University of Jerusalem




Mo Dong



Qingxi Li



Brighten Godfrey



Doron Zarchy and
Michael Schapira
Hebrew University of
Jerusalem

University of Illinois at Urbana-Champaign





High BDP

BIC
H-TCP
Compound
CUBIC
FAST TCP

Wireless

Westwood
Vegas
Veno

Satellite

Hybla
STAR

Inter-DC

Illinois
SABUL

Intra-DC

ICTCP
DCTCP



High BDP

BIC
H-TCP
Compound
CUBIC
FAST TCP

Wireless

Westwood
Vegas
Veno

Satellite

Hybla
STAR

...

Inter-DC

Illinois
SABUL

Intra-DC

ICTCP
DCTCP



High BDP

BIC
H-TCP
Compound
CUBIC
FAST TCP

Wireless

Westwood
Vegas
Veno

Satellite

Hybla
STAR

...

Inter-DC

Illinois
SABUL

Intra-DC

ICTCP
DCTCP

Who can be not happy with TCP ?



are
~~can be~~ not happy with TCP ?



High BDP

BIC
H-TCP
Compound
CUBIC
FAST TCP

Wireless

Westwood
Vegas
Veno

Satellite

Hybla
STAR

Inter-DC

Illinois
SABUL

Intra-DC

ICTCP
DCTCP



High BDP

BIC
H-TCP
Compound
CUBIC
FAST TCP

Wireless

Westwood
Vegas
Veno

Satellite

Hybla
STAR

Inter-DC

Illinois
SABUL

Intra-DC

ICTCP
DCTCP

Point Solutions



High BDP

BIC
H-TCP
Compound
CUBIC
FAST TCP

10X

Wireless

Westwood
Vegas
Veno

10X

Satellite

Hybla
STAR

17X

Inter-DC

Illinois
SABUL

4X

Intra-DC

ICTCP
DCTCP

Unstable, RTT Unfair, Bufferbloat, Crash on Changing Networks,

Point Solutions
+
Performance
Far from Optimal

**TCP fails to achieve
consistent high performance**

**Why is it
so hard?**

Reno	1 pkt loss	$\text{cwnd}/2$
Scalable	ACK	$\text{cwnd}+1$
CUBIC	Time pass 1ms	$\text{cwnd}+f(t,\text{cwn},\text{rtt})$
FAST	RTT increase x%	Reduce cwnd to f(x)%
HTCP	100 ACK	$\text{cwnd}+f(\text{cwnd})/\text{cwnd}$

	Event	Action
Reno	1 pkt loss	$cwnd/2$
Scalable	ACK	$cwnd+1$
CUBIC	Time pass 1ms	$cwnd+f(t,cwn,rtt)$
FAST	RTT increase x%	Reduce cwnd to f(x)%
HTCP	100 ACK	$cwnd+f(cwnd)/cwnd$

Hardwired Mapping


	Event	Action
Reno	1 pkt loss	$wnd/2$
Scalable	ACK	$wnd+1$
CUBIC	Time pass 1ms	$wnd+f(t,cwn,rtt)$
FAST	RTT increase x%	Reduce wnd to $f(x)\%$
HTCP	100 ACK	$wnd+f(cwnd)/cwnd$

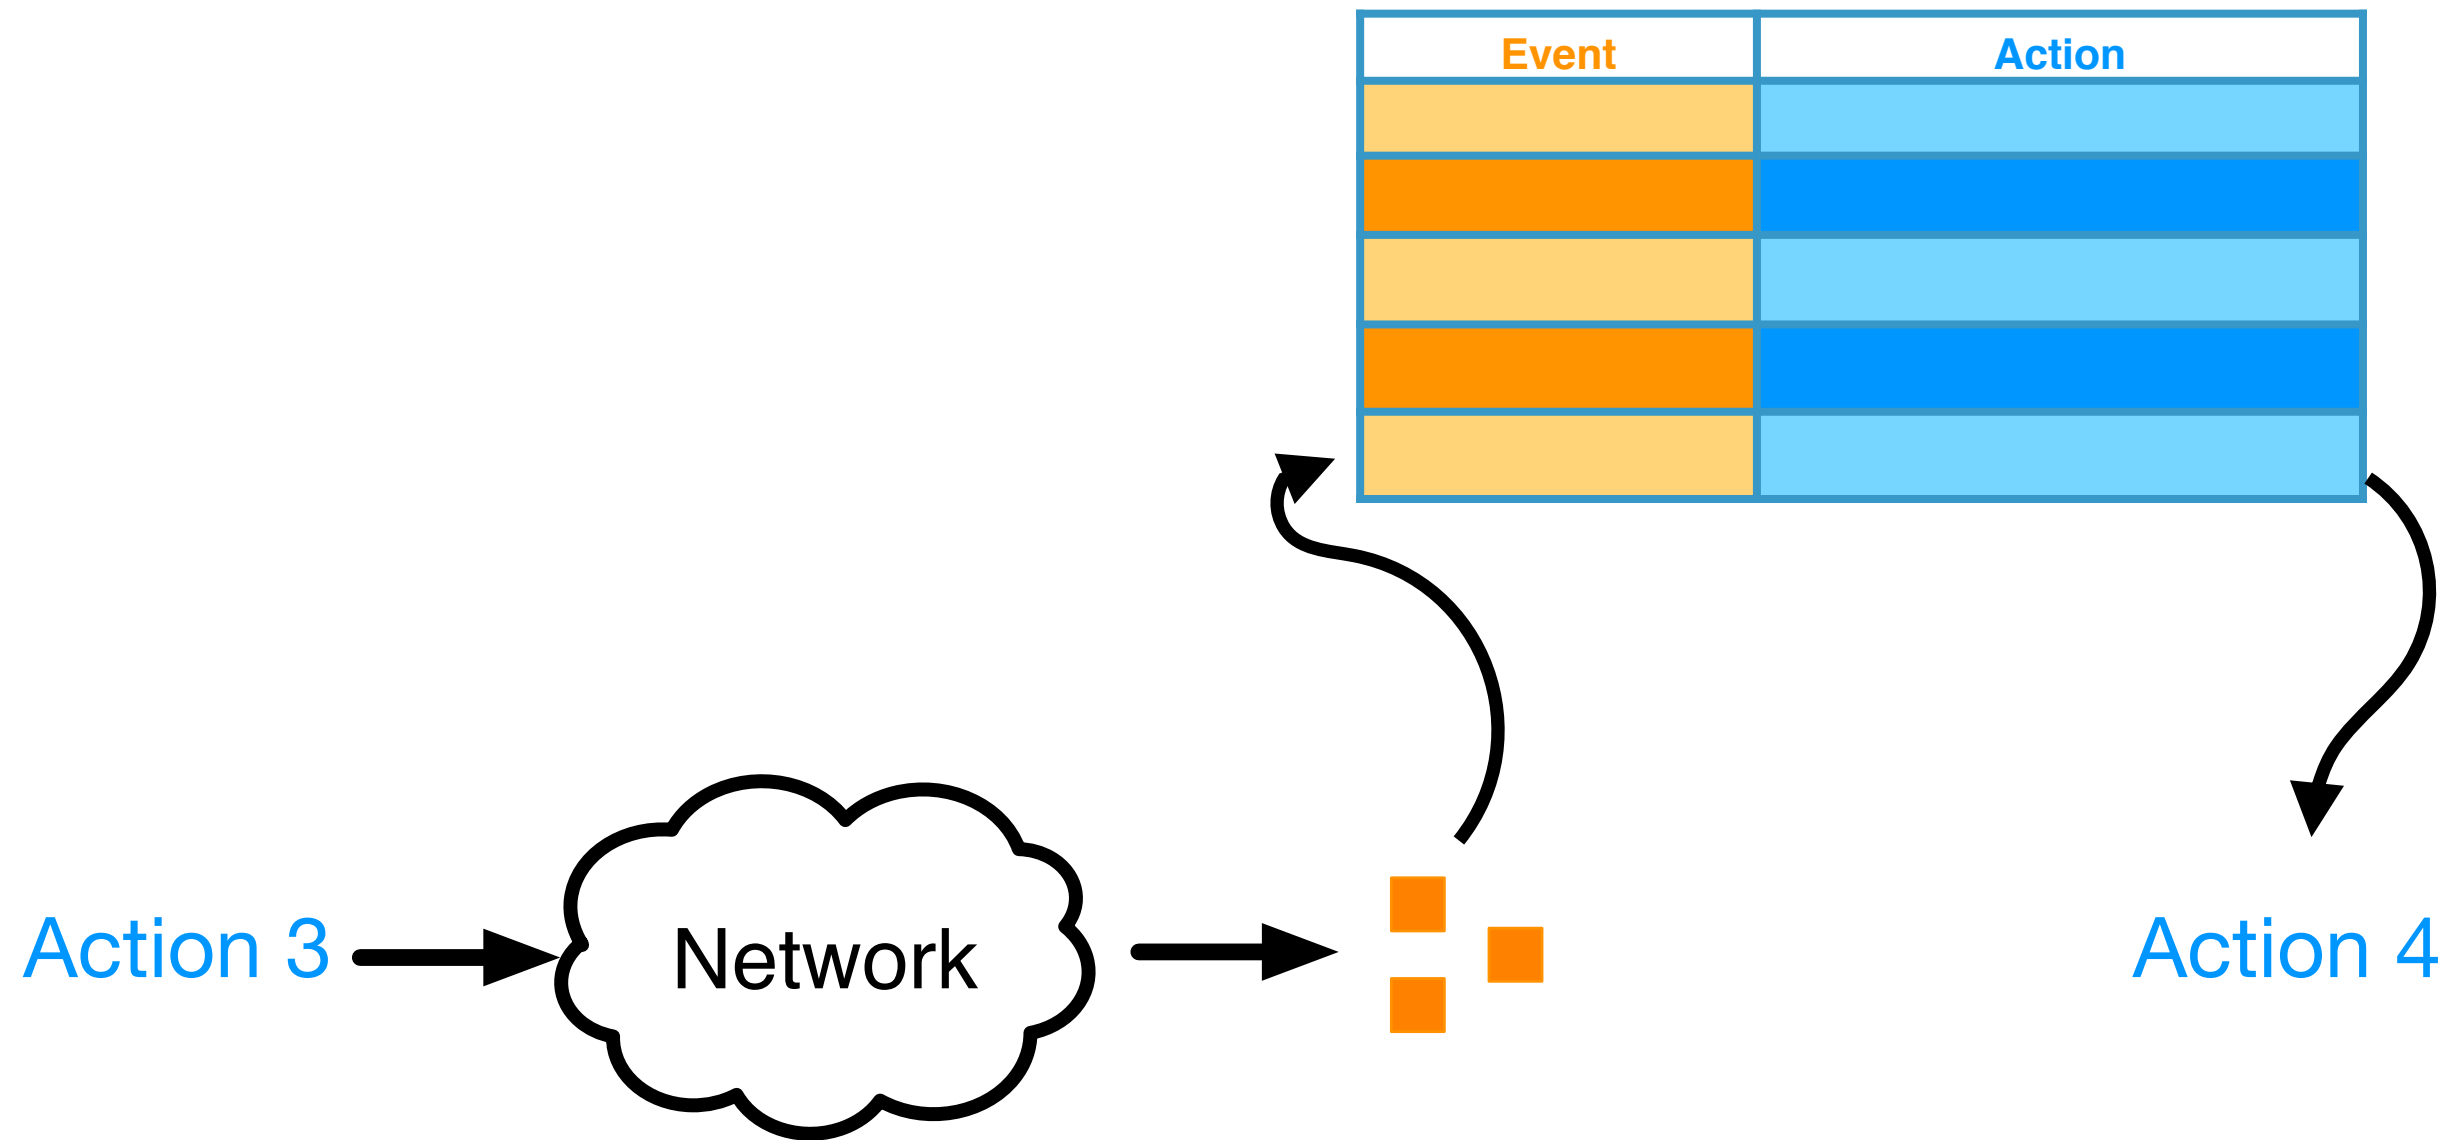
Hardwired Mapping


Event	Action


Event	Action


Event	Action

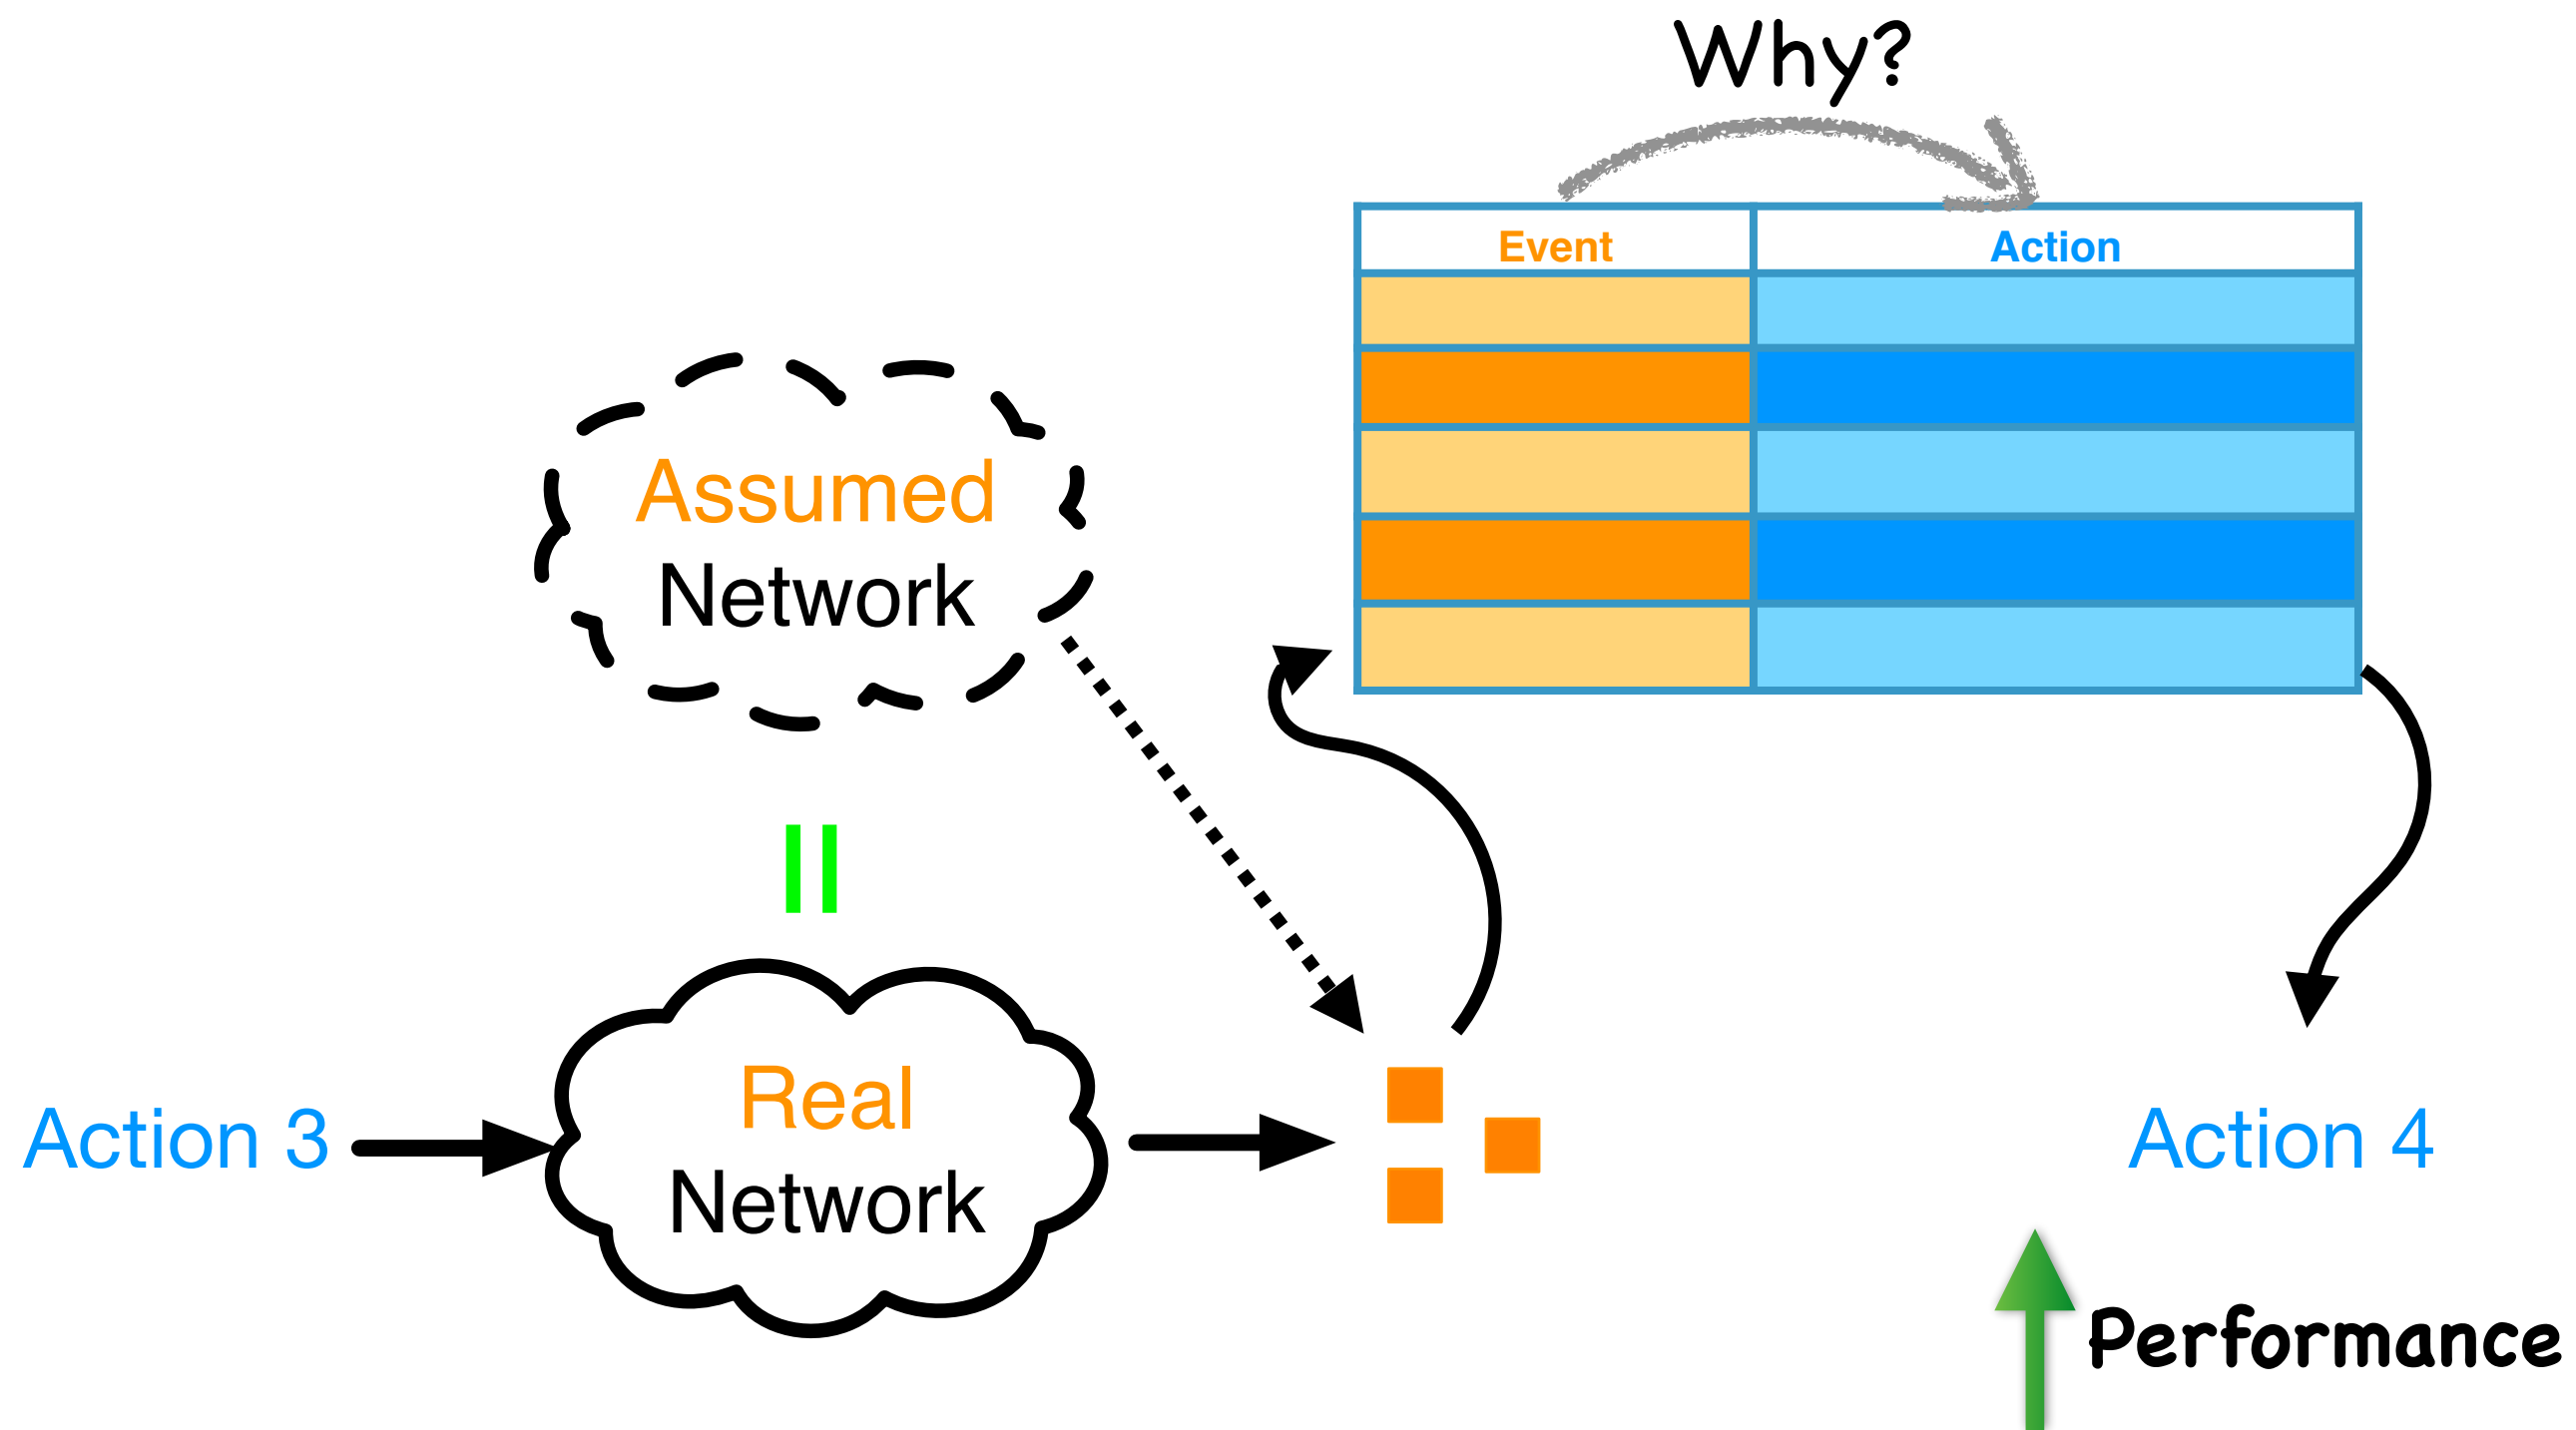
















Flow f sends at R

Event	Action

Flow f sends at R


Event	Action
Packet Loss	

Flow f sends at R



Event	Action
Packet Loss	


Flow f sends at R



f causes
most congestion

Event	Action
Packet Loss	


Flow f sends at R



f causes
most congestion

Event	Action
Packet Loss	Dec R a lot


Flow f sends at R



Event	Action
Packet Loss	Dec R a lot

shallow buffer
overflow


Flow f sends at R



Event	Action
Packet Loss	Dec R a lot

shallow buffer
overflow


Flow f sends at R



Event	Action
Packet Loss	Dec R a lot Dec R a little

shallow buffer
overflow


Flow f sends at R



Event	Action
Packet Loss	Dec R a lot Dec R a little

other high rate flow causing congestion


Flow f sends at R



Event	Action
Packet Loss	Dec R a lot Dec R a little

other high rate flow
causing congestion


Flow f sends at R



Event	Action
Packet Loss	<p>Dec R a lot</p> <p>Dec R a little</p> <p>Maintain R</p>

other high rate flow causing congestion


Flow f sends at R



Event	Action
Packet Loss	<p>Dec R a lot</p> <p>Dec R a little</p> <p>Maintain R</p>

loss is random


Flow f sends at R



Event	Action
Packet Loss	<p>Dec R a lot</p> <p>Dec R a little</p> <p>Maintain R</p>

loss is random


Flow f sends at R



Event	Action
Packet Loss	Dec R a lot
	Dec R a little
	Maintain R
	Increase R

loss is random


Flow f sends at R



Event	Action
Packet Loss	Dec R a lot
	Dec R a little
	Maintain R
	Increase R


loss is random

Flow f sends at R



Event	Action
Packet Loss	<ul style="list-style-type: none">Dec R a lotDec R a littleMaintain RIncrease R

Flow f sends at R




Event	Action
Pac	Dec R a lot
	Maintain R
	Increase R


No event-control mapping optimal for all network scenarios

What is the right rate to send?


What is the right rate to send?




What is the right rate to send?




What is the right rate to send?




What is the right rate to send?




What is the right rate to send?




What is the right rate to send?




What is the right rate to send?





What is the right rate to send?





No matter how complex the network,
rate $r \rightarrow$ utility u






















Performance-oriented Congestion Control




Performance-oriented Congestion Control



Observe real
performance


Performance-oriented Congestion Control



Observe real
performance

Control based on
empirical evidence



Performance-oriented Congestion Control



Observe real performance

Control based on empirical evidence

yields
Consistent high performance

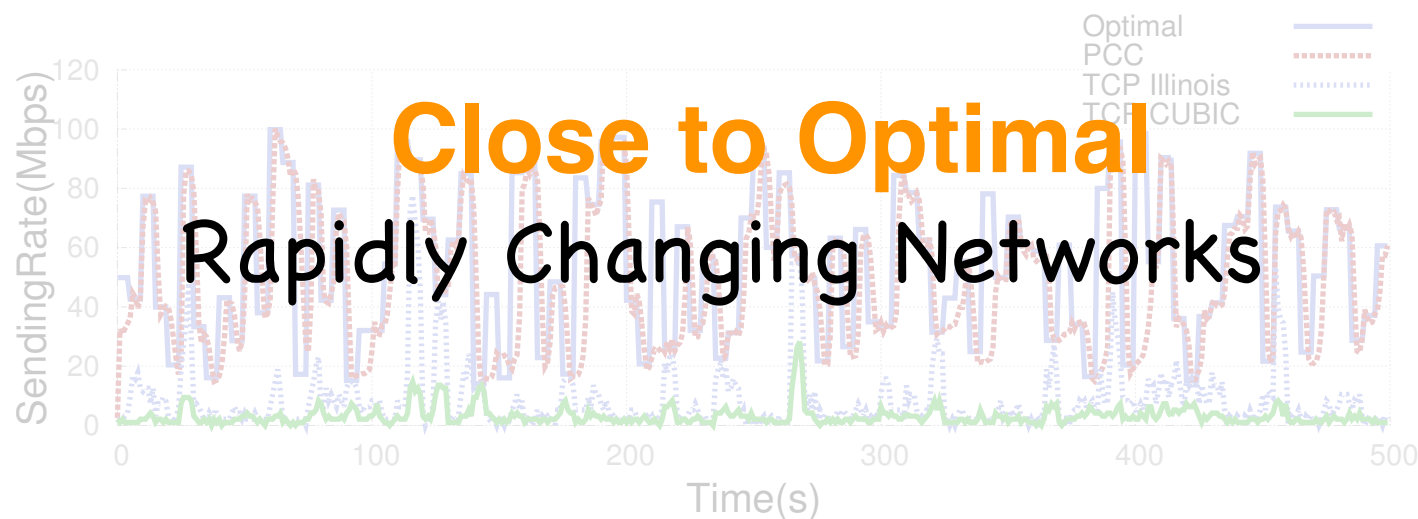






Consistent High Performance

Table 1: PCC significantly outperforms TCP in inter-data center environments. RTT in msec; throughput in Mbps.


Transmission Pair	RTT	PCC	SABUL	CUBIC	Illinois
GPO → NYSErNet	129	326			
GPO → Missouri	80.7	90.1			
GPO → Illinois	35.4	766	664	84.5	102
NYSErNet → Missouri	47.4	816	662	108	109
Wisconsin → Illinois	9.01	801	700	547	562
GPO → Wisc.	38.0	783	487	79.3	120
NYSErNet → Wisc.	38.3	791	673	134	134
Missouri → Wisc.	20.9	807	698	259	262
NYSErNet → Illinois	36.1	808	674	141	141

4X
InterDC






Where is Congestion Control?




Where is Congestion Control?




Selfishly maximizing utility

Where is Congestion Control?



Selfishly maximizing utility
=> non-cooperative game

Where is Congestion Control?



Selfishly maximizing utility
=> non-cooperative game

Do we converge to a fair Nash equilibrium?

A class of utility functions converge to a fair and efficient Nash Equilibrium

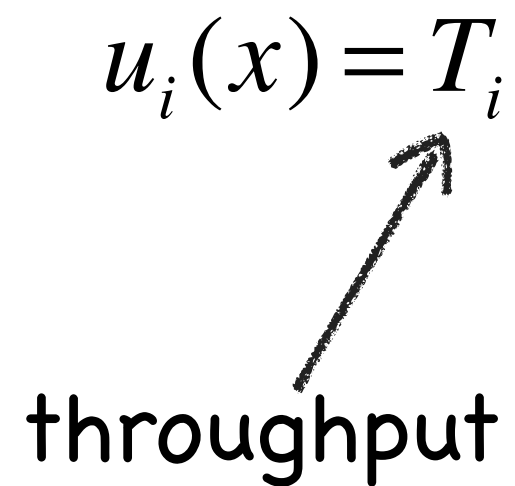
A class of utility functions converge to a fair and efficient Nash Equilibrium

$$u_i(x) = T_i - x_i * L_i$$

A class of utility functions converge to a fair and efficient Nash Equilibrium

$$u_i(x) = T_i$$

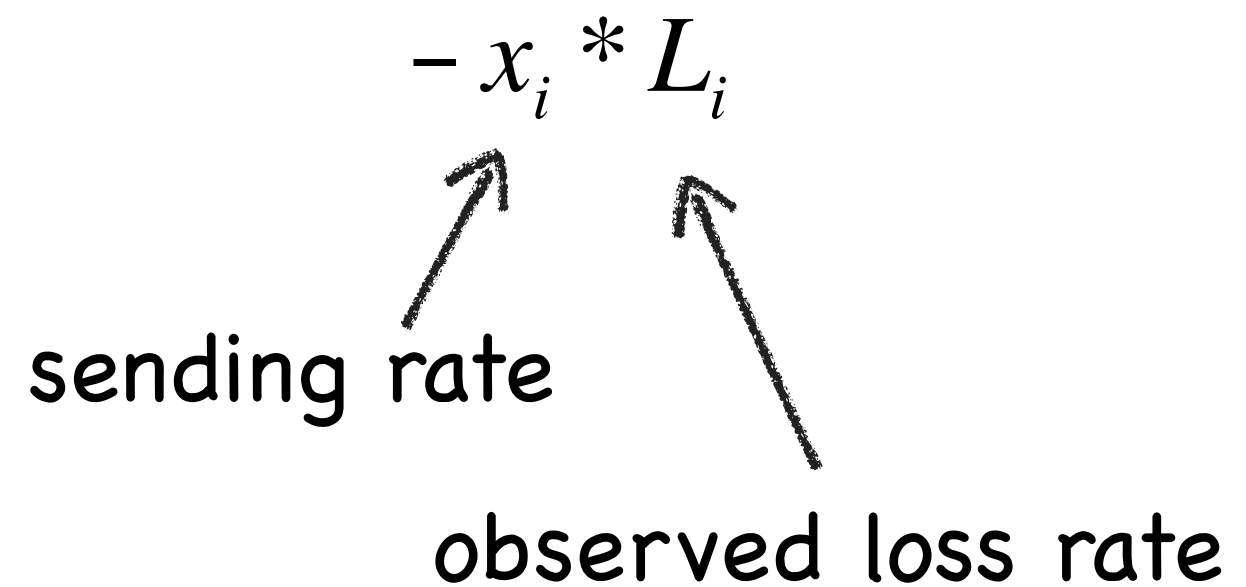
throughput



$$-x_i * L_i$$

sending rate

observed loss rate



A class of utility functions converge to a fair and efficient Nash Equilibrium

$$u_i(x) = T_i * \text{sigmoid}(L_i - 0.05) - x_i * L_i$$



Loss Rate


TCP Dynamics

TCP Dynamics

AIMD as a “hack” to asymptotic fairness

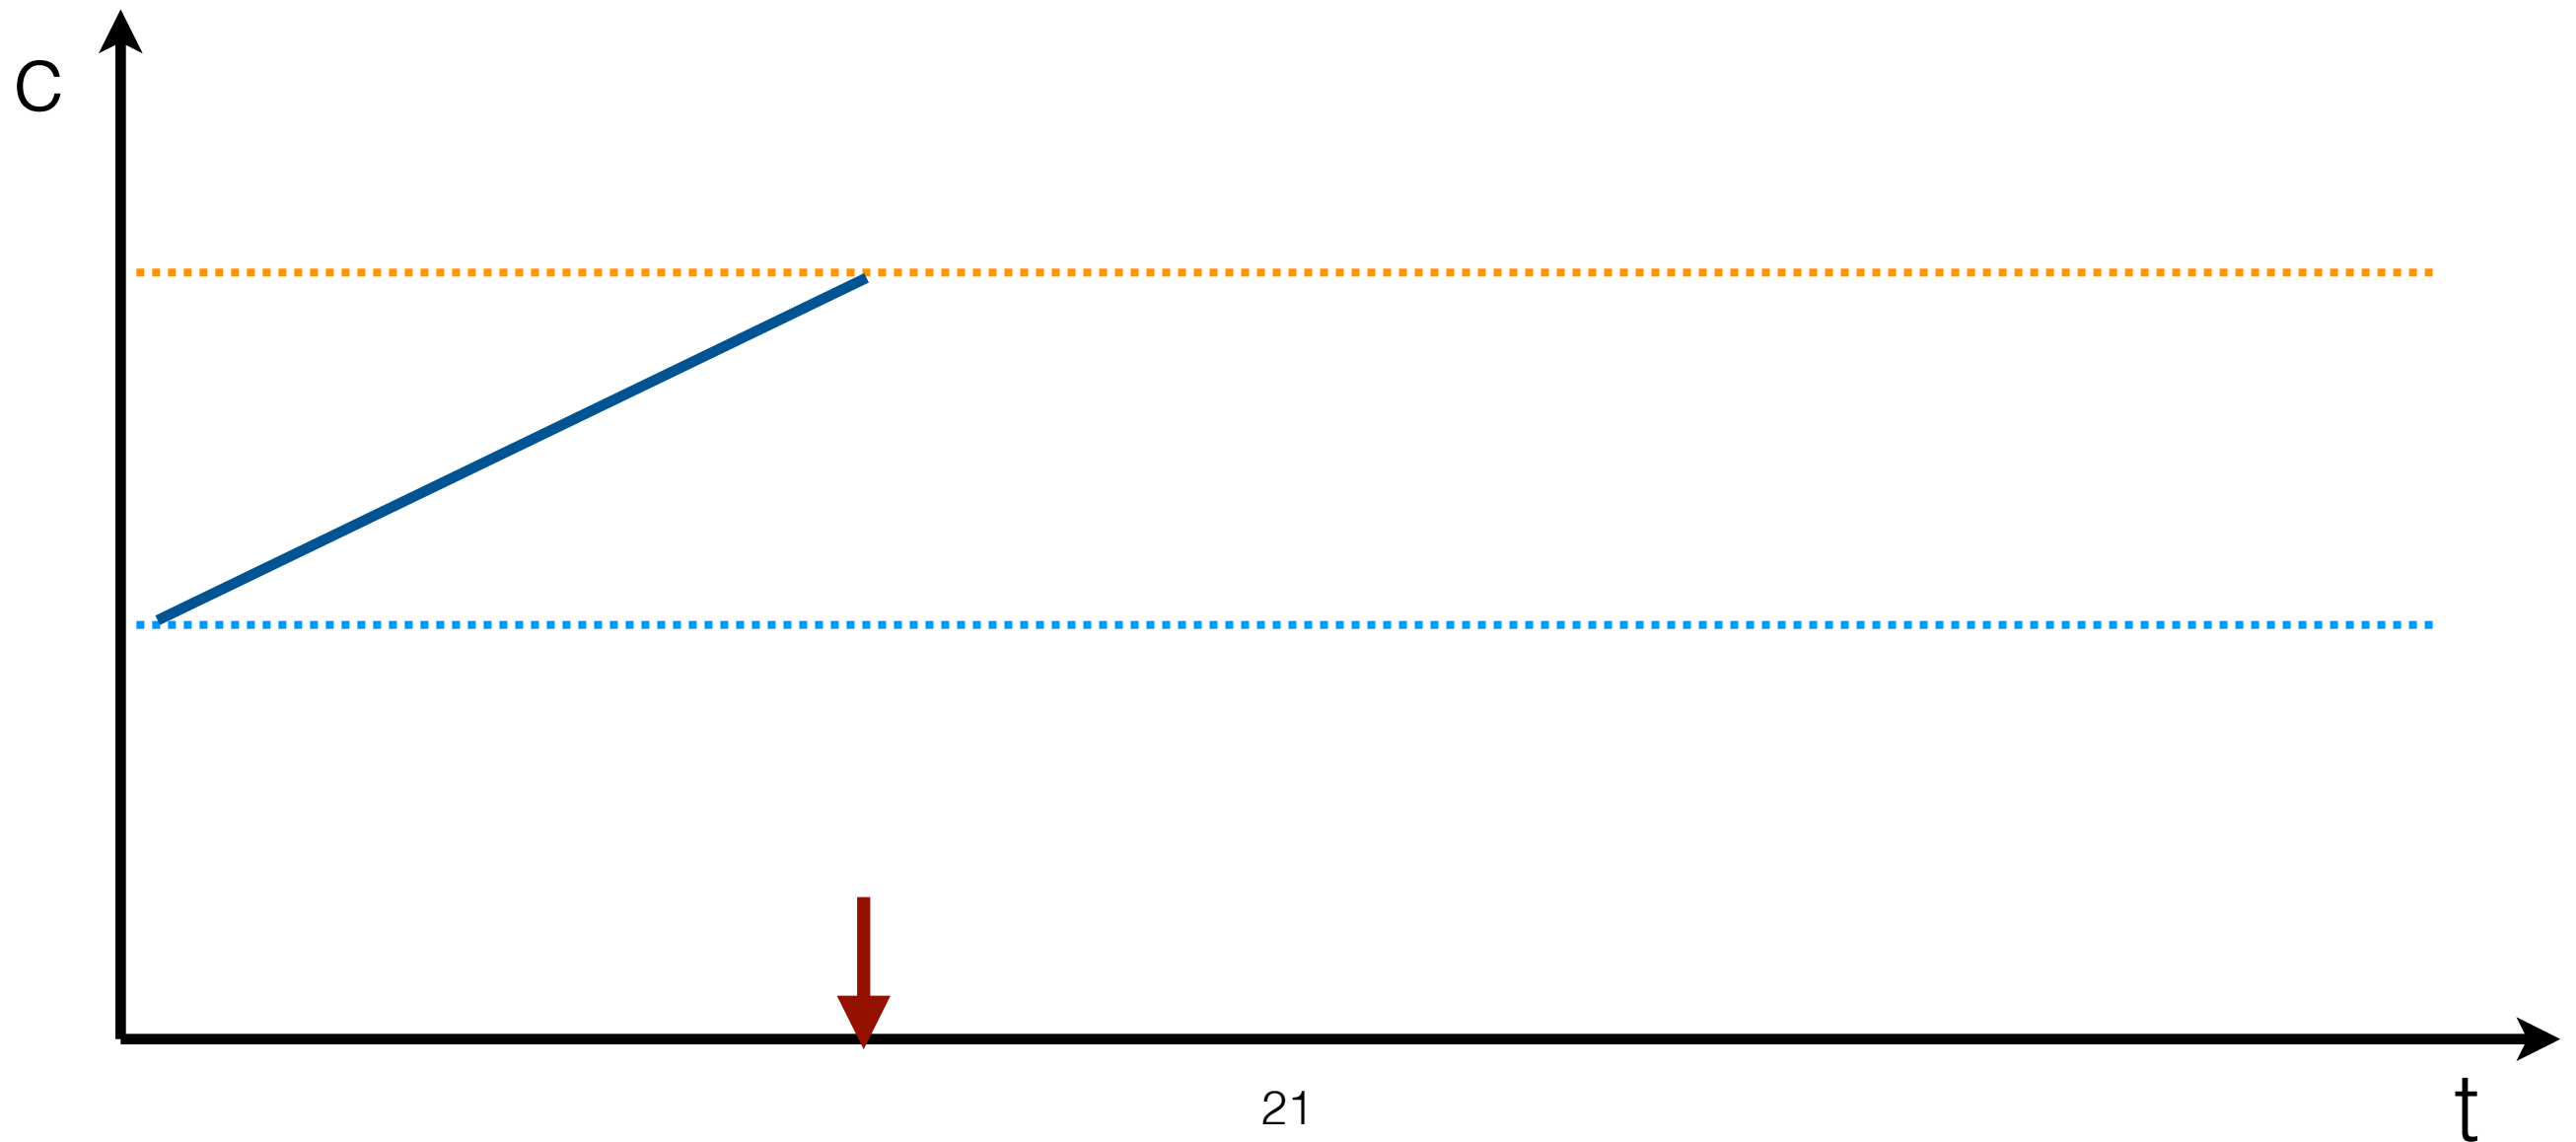
TCP Dynamics

AIMD as a "hack" to asymptotic fairness




TCP Dynamics

AIMD as a "hack" to asymptotic fairness



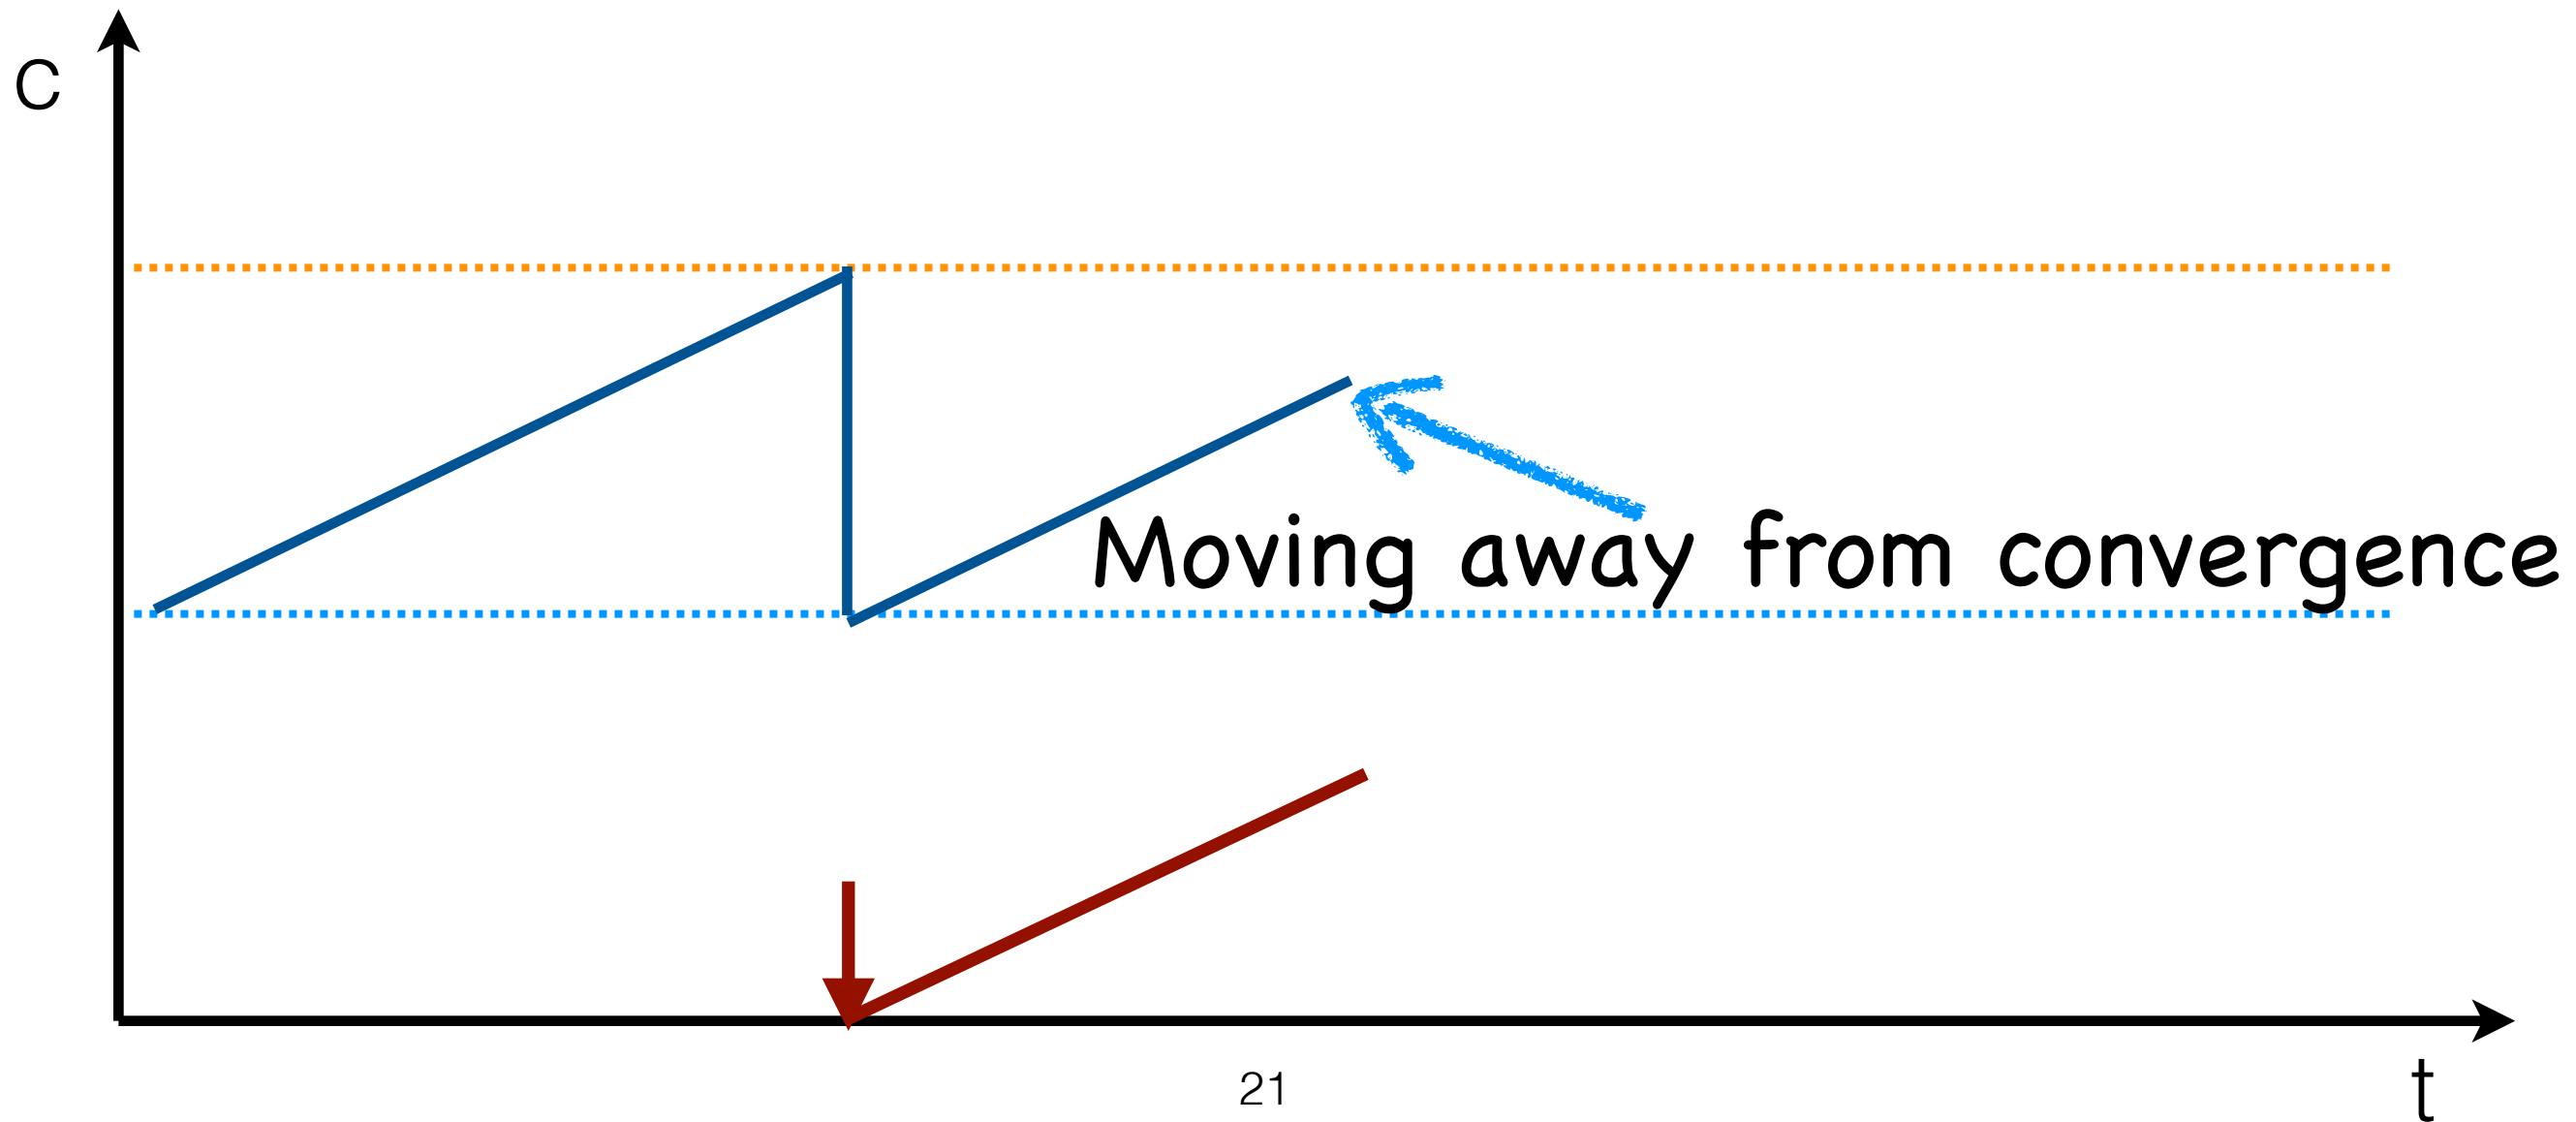
TCP Dynamics

AIMD as a "hack" to asymptotic fairness




TCP Dynamics

AIMD as a "hack" to asymptotic fairness




PCC Dynamics




PCC Dynamics

PCC does not need AIMD because it looks at real performance




PCC Dynamics

PCC does not need AIMD because it looks at real performance




PCC Dynamics

PCC does not need AIMD because it looks at real performance





PCC Dynamics

PCC does not need AIMD because it looks at real performance




Convergence




Consistent High Performance

Consistent High Performance





Satellite Networks



Consistent High Performance






Consistent High Performance






Consistent High Performance







Consistent High Performance







Consistent High Performance



Consistent High Performance




Consistent High Performance




Consistent High Performance


17X vs Hybla
Satellite Networks




10X vs Illinois
Lossy Networks




15X less buffer
Shallow Network Buffer





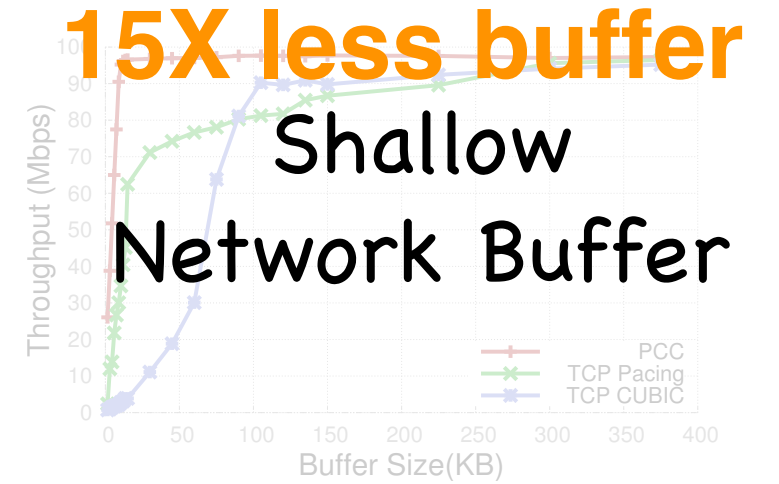


Similar to ICTCP



Rapidly Changing Networks




Consistent High Performance




Consistent High Performance


17X vs Hybla
Satellite Networks




10X vs Illinois
Lossy Networks




15X less buffer
Shallow Network Buffer




Similar to ICTCP
IntraDC Incast



Close to Optimal
Rapidly Changing Networks




RTT Unfairness




Consistent High Performance


17X vs Hybla
Satellite Networks




10X vs Illinois
Lossy Networks




15X less buffer
Shallow Network Buffer




Similar to ICTCP
IntraDC Incast



Close to Optimal
Rapidly Changing Networks




Solves RTT Unfairness




Consistent High Performance


17X vs Hybla
Satellite Networks




10X vs Illinois
Lossy Networks




15X less buffer
Shallow Network Buffer



Similar to ICTTCP
IntraDC Incast



Close to Optimal
Rapidly Changing Networks



Solves
RTT Unfairness





Table 1: PCC significantly outperforms TCP in inter-data center environments. RTT is in msec; throughput in Mbps.

Transmission Pair	RTT	PCC	SABUL	CUBIC	Illinois
GPO → NYSErNet	36.1	766	664	129	326
GPO → Missouri	35.4	816	662	80.7	90.1
GPO → Illinois	35.4	801	700	84.5	102
NYSErNet → Missouri	47.4	801	700	108	109
Wisconsin → Illinois	9.01	783	487	547	562
GPO → Wisc.	38.0	791	673	134	120
NYSErNet → Wisc.	38.3	698	259	134	134
Missouri → Wisc.	20.9	807	698	259	262
NYSErNet → Illinois	36.1	808	674	141	141


InterDC

Consistent High Performance


17X vs Hybla
Satellite Networks




10X vs Illinois
Lossy Networks




15X less buffer
Shallow Network Buffer




Similar to ICTCP
IntraDC
Incast



Close to Optimal
Rapidly Changing Networks



Solves
RTT Unfairness




4X vs Illinois
1.23X vs UDT
InterDC

Table 1: PCC significantly outperforms TCP in inter-data center environments. RTT is in msec; throughput in Mbps.


Flow	RTT (msec)	PCC	UDT	TCP Illinois
GPO → NYSErNet	129	326		
GPO → Missouri	36	80.7	90.1	
GPO → Illinois	35.4	766	664	84.5
NYSErNet → Missouri	47.4	816	662	108
Wisconsin → Illinois	9.01	801	700	547
GPO → Wisc.	38.0	783	487	79.3
NYSErNet → Wisc.	38.3	791	673	134
Missouri → Wisc.	20.9	807	698	259
NYSErNet → Illinois	36.1	808	674	141

Consistent High Performance


17X vs Hybla
Satellite Networks




10X vs Illinois
Lossy Networks




15X less buffer
Shallow Network Buffer




Similar to ICTCP
IntraDC
Incast



Close to Optimal
Rapidly Changing Networks



Solves
RTT Unfairness



4X vs Illinois

1.23X vs UDT

InterDC


Table 1: PCC significantly outperforms TCP in inter-data center environments. RTT is in msec; throughput in Mbps.

Flow	RTT (msec)	UDT	PCC	ABILEN	Illinois
GPO → NYSErNet	129	326			
GPO → Missouri	36	80.7	90.1		
GPO → Illinois	35.4	766	664	84.5	102
NYSErNet → Missouri	47.4	816	662	108	109
Wisconsin → Illinois	9.01	801	700	547	562
GPO → Wisc.	38.0	783	487	79.3	120
NYSErNet → Wisc.	38.3	791	673	134	134
Missouri → Wisc.	20.9	807	698	259	262
NYSErNet → Illinois	36.1	808	674	141	141




Consistent High Performance


17X vs Hybla
Satellite Networks




10X vs Illinois
Lossy Networks




15X less buffer
Shallow Network Buffer




Similar to ICTCP
IntraDC
Incast



Close to Optimal
Rapidly Changing Networks



Solves
RTT Unfairness



4X vs Illinois

1.23X vs UDT

InterDC

Table 1: PCC significantly outperforms TCP in inter-data center environments. RTT is in msec; throughput in Mbps.

Flow	RTT (msec)	UDT	PCC	TCP	Illinois
GPO → NYSErNet	3	129	326		
GPO → Missouri	36	80.7	90.1		
GPO → Illinois	35.4	766	664	84.5	102
NYSErNet → Missouri	47.4	816	662	108	109
Wisconsin → Illinois	9.01	801	700	547	562
GPO → Wisc.	38.0	783	487	79.3	120
NYSErNet → Wisc.	38.3	791	673	134	134
Missouri → Wisc.	20.9	807	698	259	262
NYSErNet → Illinois	36.1	808	674	141	141

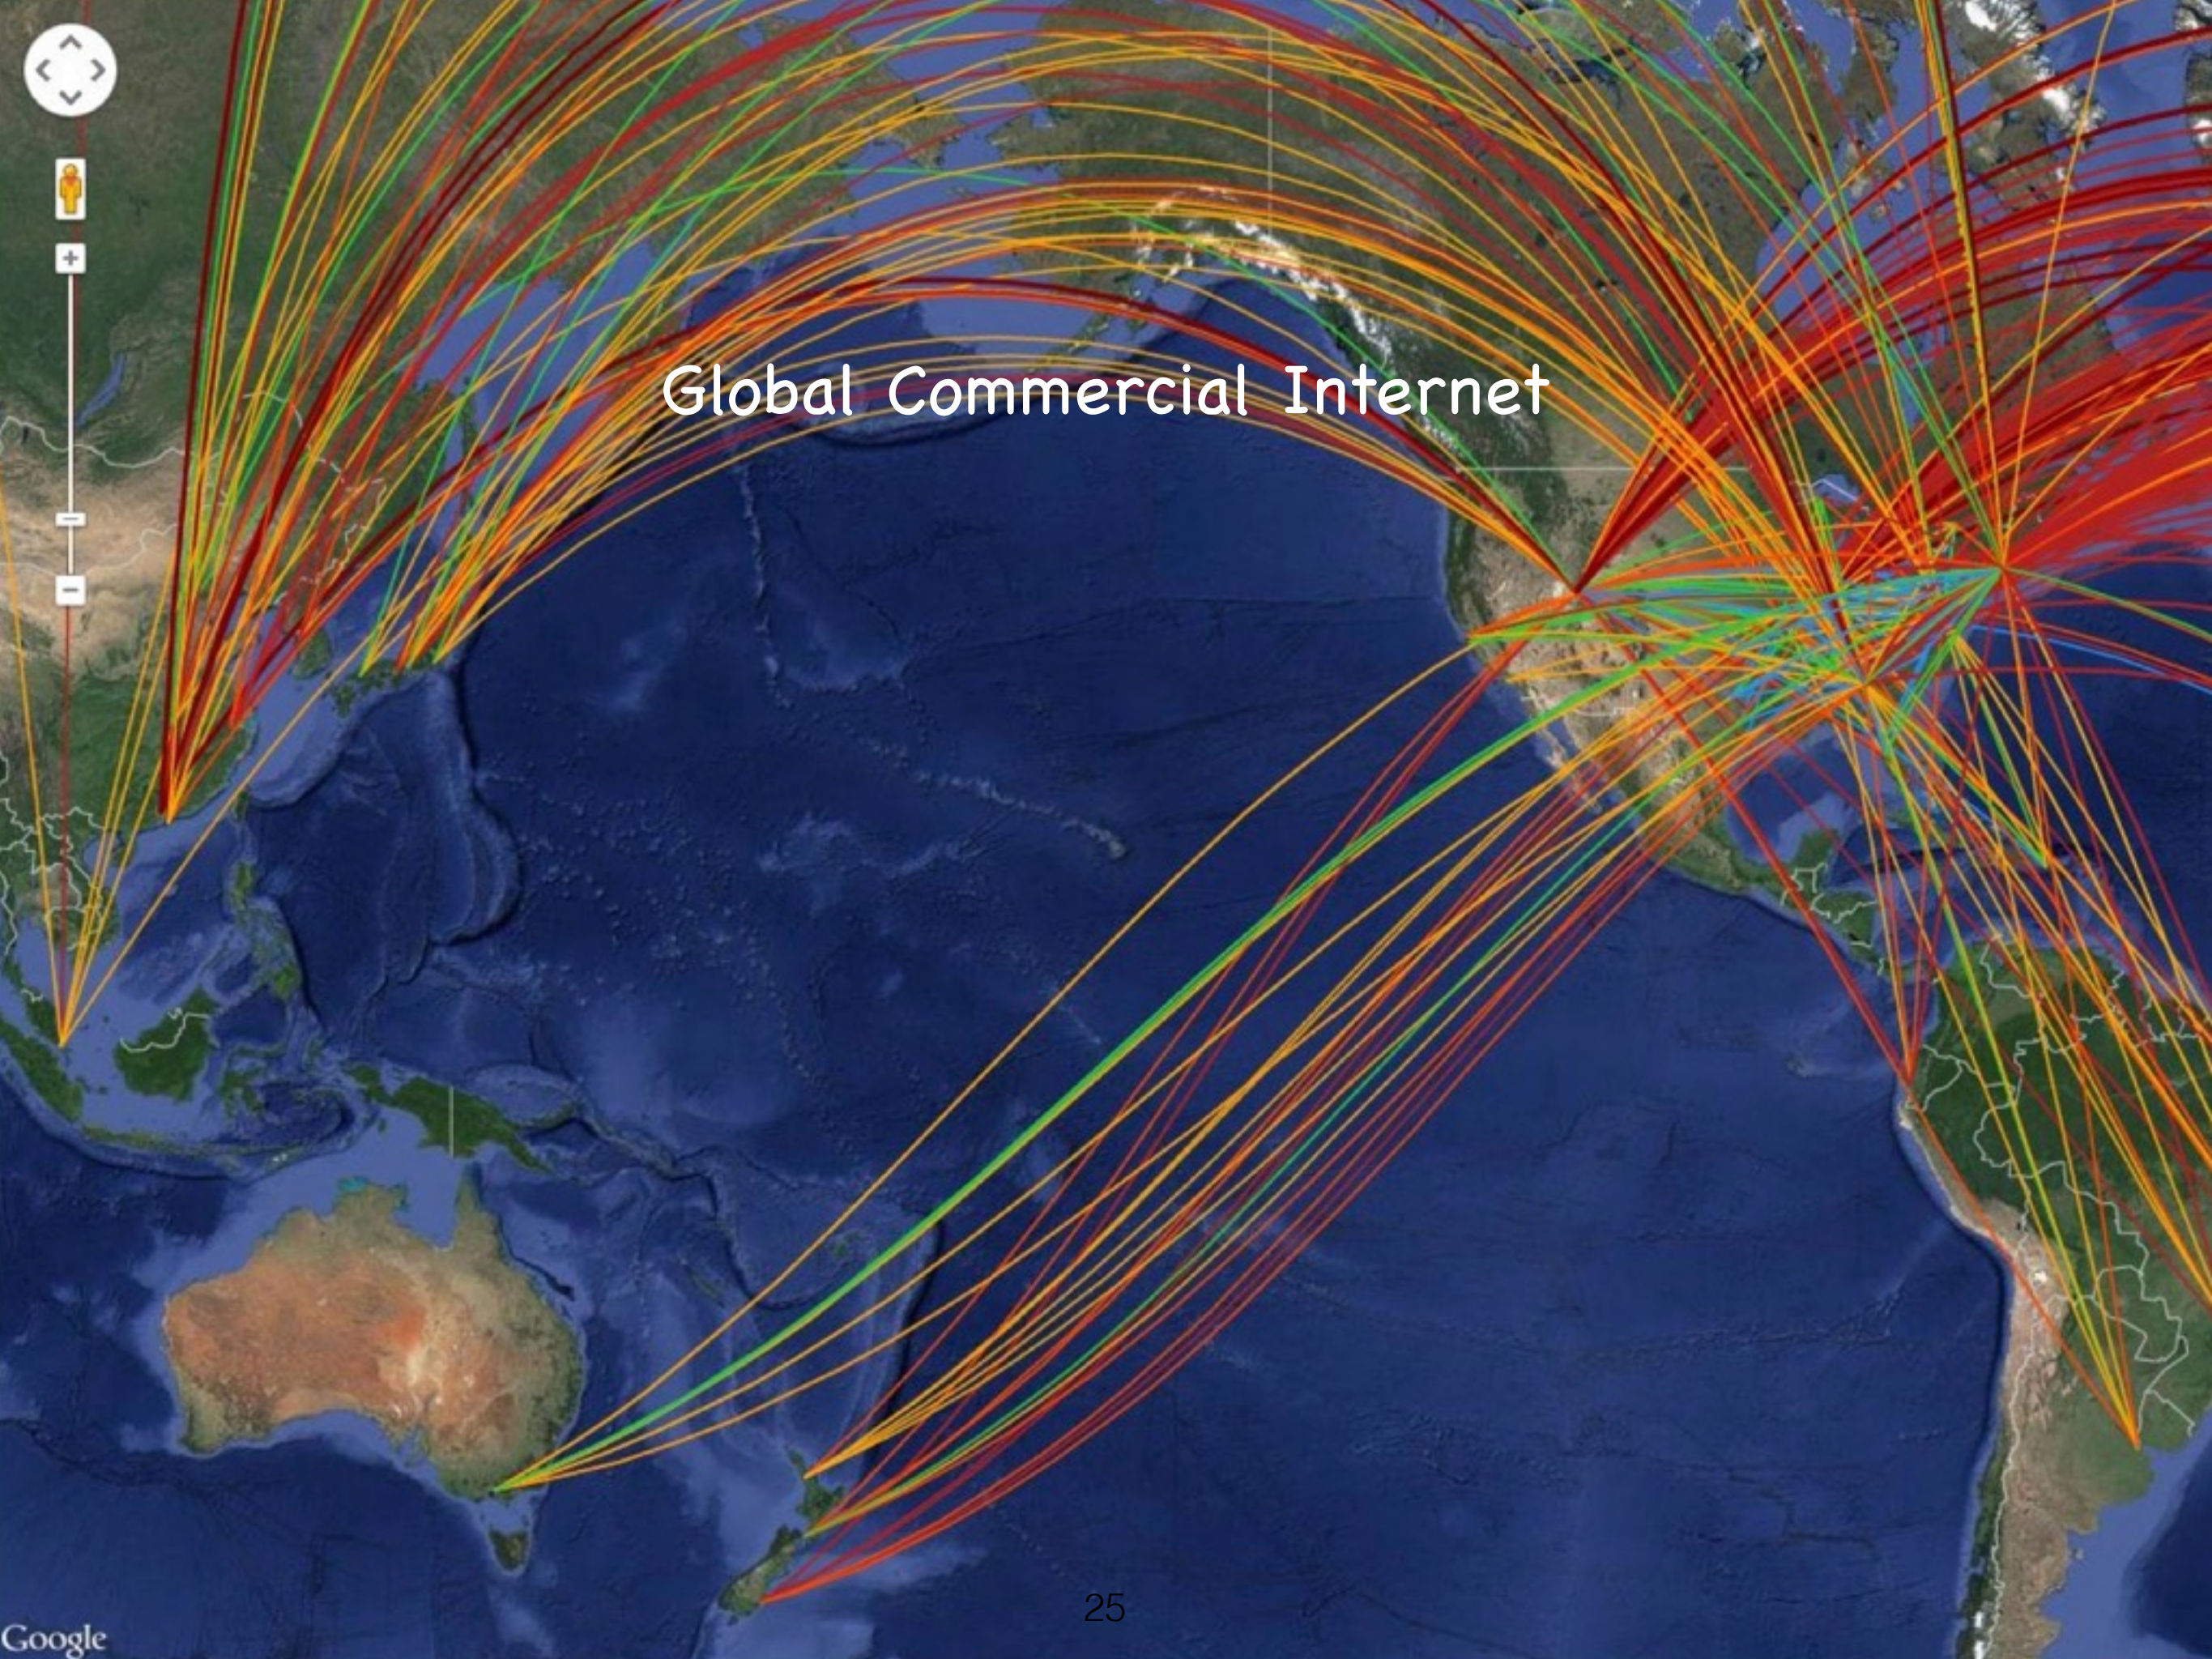
4X median vs CUBIC

1.48X vs UDT + 4X less loss

Global Commercial Internet



Global Commercial Internet







PCC vs TCP vs Take a Flight

100G Data

PCC vs TCP vs Take a Flight



100G Data

-  PCC
-  TCP CUBIC
-  Take a Flight




PCC vs TCP vs Take a Flight


100G Data



PCC vs TCP vs Take a Flight




PCC vs TCP vs Take a Flight



Consistent High Performance

Rapidly Changing Networks


BW: 10-100Mbps; RTT: 10-100ms; Loss Rate: 0-1%
Change every 5 seconds



Consistent High Performance

Rapidly Changing Networks


BW: 10-100Mbps; RTT: 10-100ms; Loss Rate: 0-1%
Change every 5 seconds



Consistent High Performance

Rapidly Changing Networks

BW: 10-100Mbps; RTT: 10-100ms; Loss Rate: 0-1%
Change every 5 seconds



Long list of things we have done but don't have time to talk about

- More stories about the fact that TCP is broken
- Proof of Nash Equilibrium and Convergence
- Concrete Implementation of PCC
 - Performance monitoring
 - Details of learning control algorithm
 - Implementation designs and optimizations
- Performance Evaluation
 - Planetlab detailed explanation
 - Satellite like network
 - Lossy Links
 - TCP friendliness
 - RTT fairness
 - Shallow buffer networks
 - Inter data center networks
 - small buffer networks
 - Reactiveness and stability tradeoff
 - Jain index fairness
 - Benefit of Randomized Control Trials
 - Details of TCP friendliness evaluation
 - Emulated satellite networks
 - Emulated datacenter networks
 - Cure RTT unfairness
 - Does not fundamentally harm short flow FCT
 - Evaluation in the wild vs non-TCP protocols
- Flexibility by pluggable utility function

Cellular Networks

- Deep queue and bufferbloat
- Highly dynamic available bandwidth

PCC

- Flexible (latency sensitive) utility function
- Handling small causality window

Flexible Utility Function

Flexible Utility Function

Same Rate Control Algorithm

Flexible Utility Function

Same Rate Control Algorithm

+

Flexible Utility Function

Same Rate Control Algorithm

+



Different Utility Function

=

Flexibility



Latency Sensitive Utility Function

$$\text{Power} = T_{\text{pt}}/\text{RTT}$$





Latency Sensitive Utility Function

$$\text{Power} = T_{pt}/RTT$$





Latency Sensitive Utility Function

$$\text{Power} = T_{pt}/RTT$$




Latency Sensitive Utility Function

$$\text{Power} = T_{\text{pt}}/\text{RTT}$$



Small Causality Window



Small Causality Window

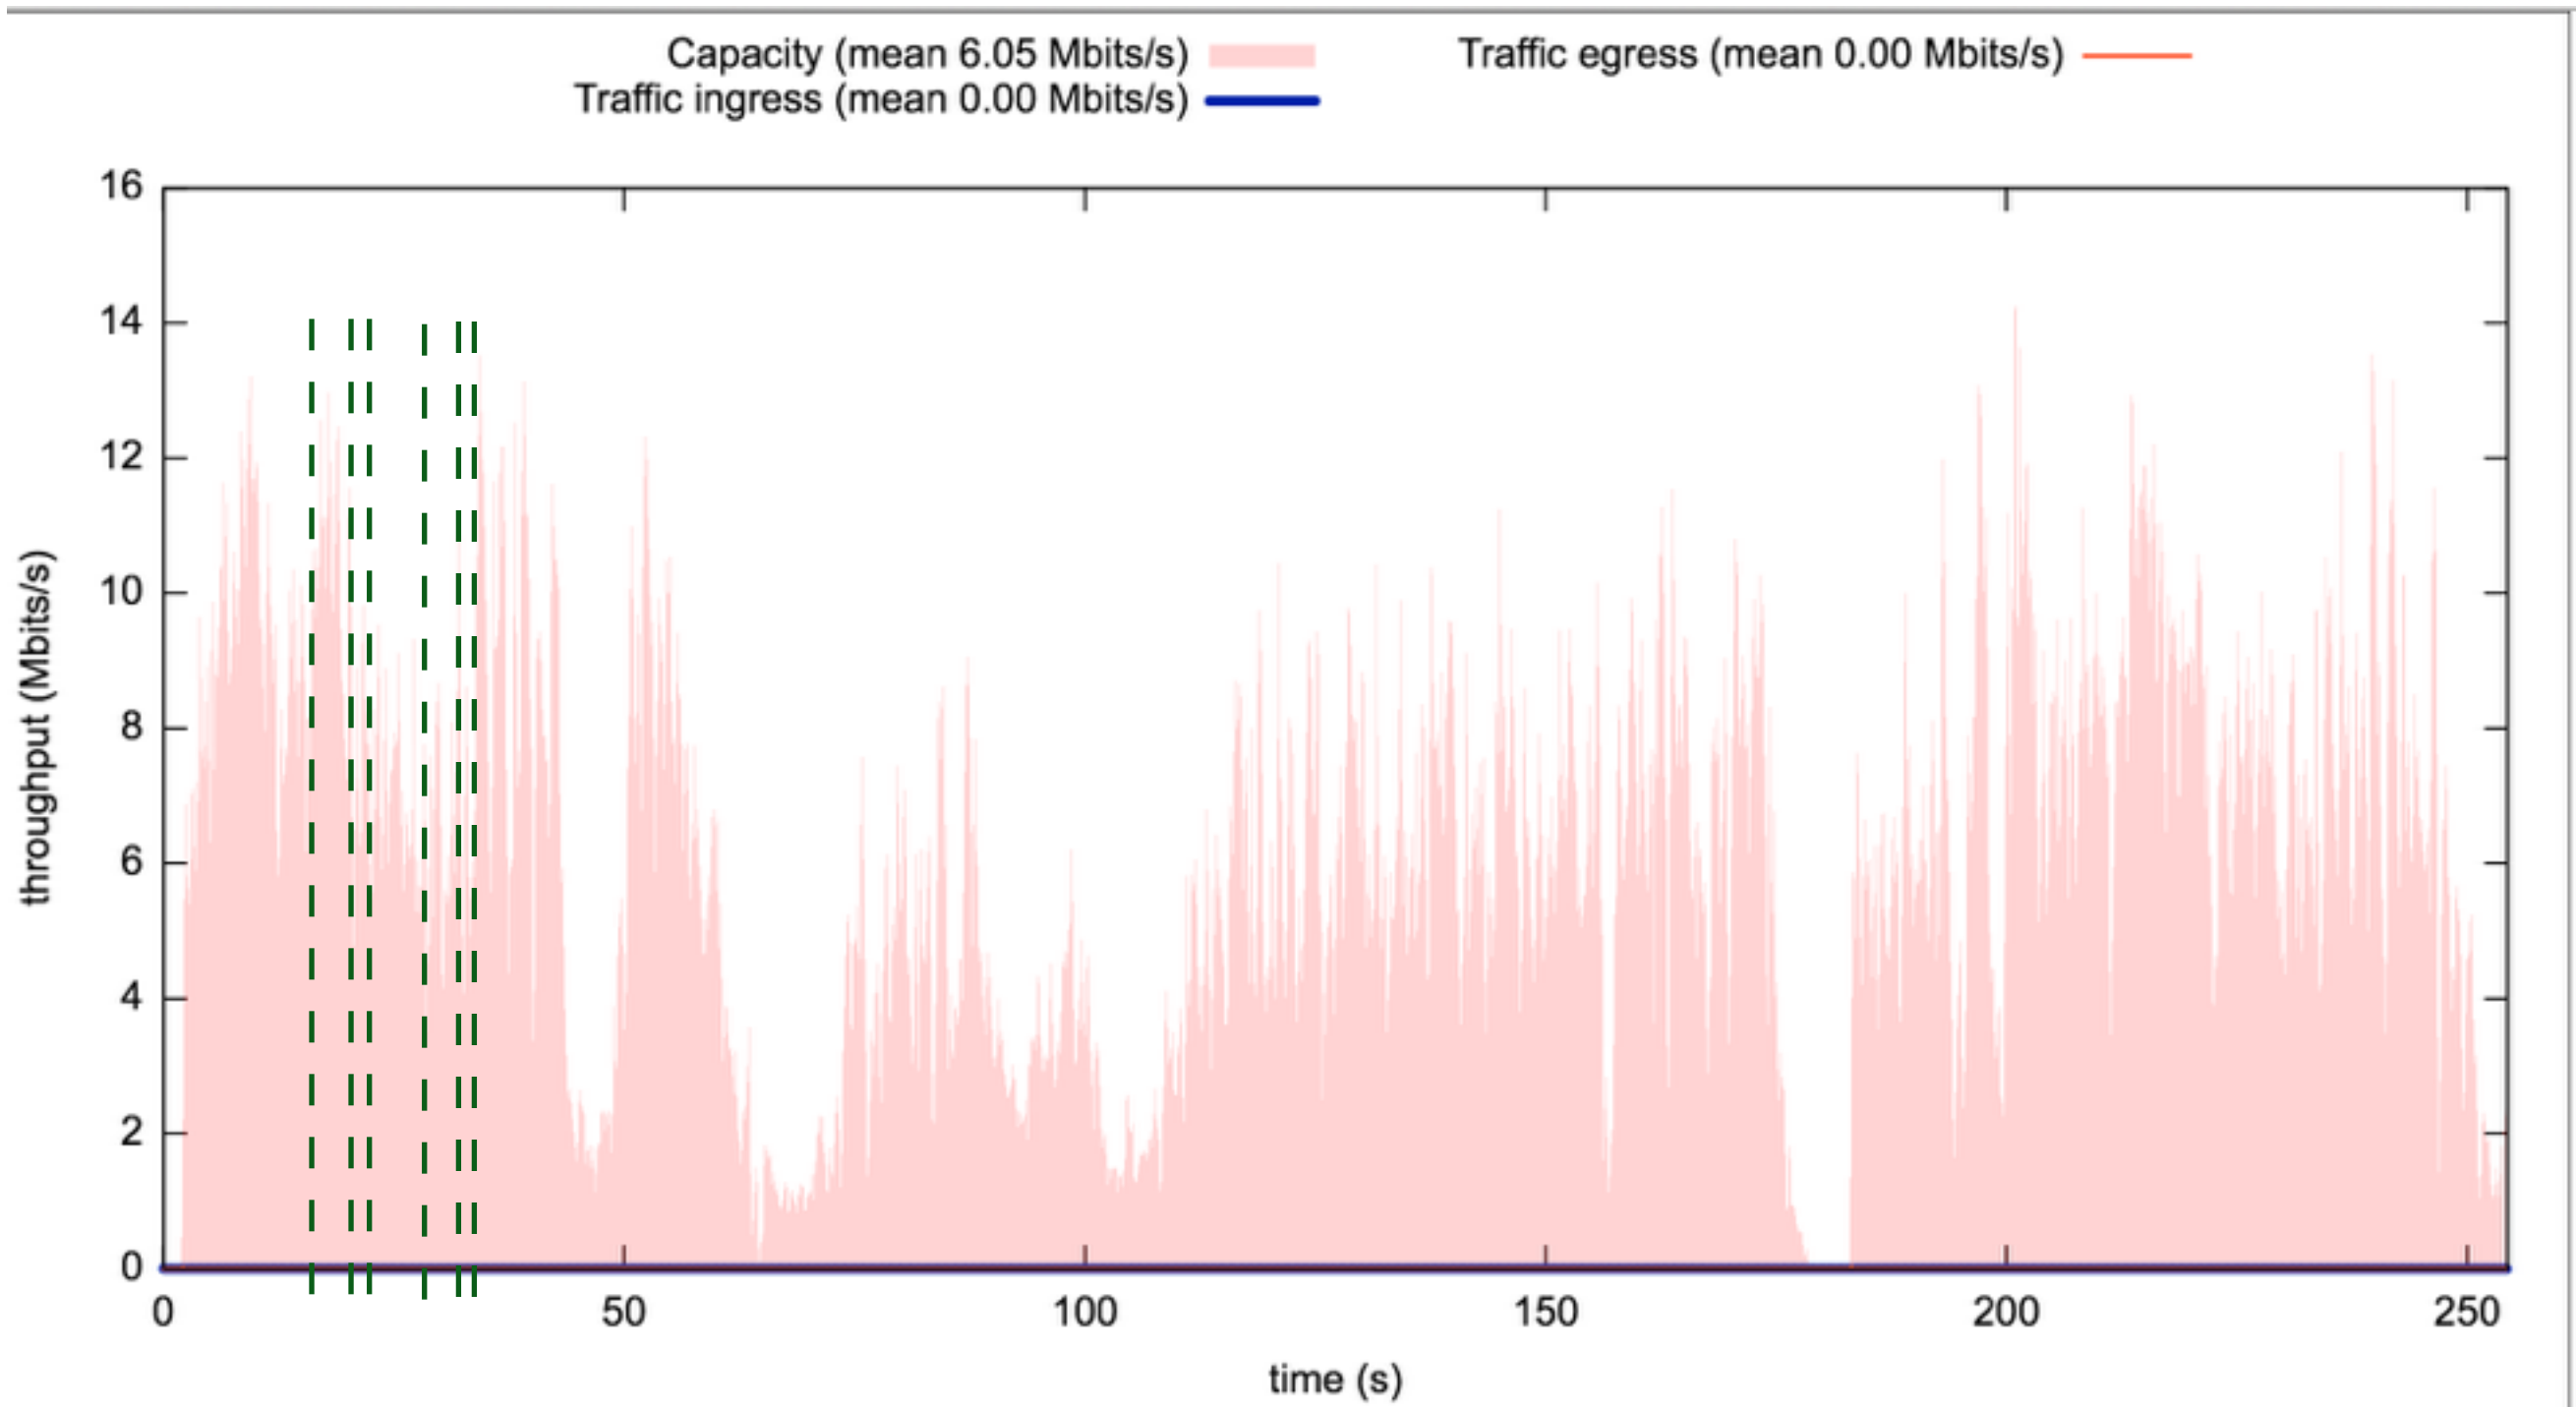


Figure drawn using Mahimahi by K. Winstein et. al. (<http://mahimahi.mit.edu/>)

Small Causality Window

PCC V2

- Fast convergence to optimal in tiny causality window
- Fast reaction to causality disruption
- Handling noisy measurement



- Rate control based on empirically observed performance yields
 - Consistent high performance
 - Better stability than TCP
 - Flexible performance objectives
 - PCC V2 is coming



speedier.net

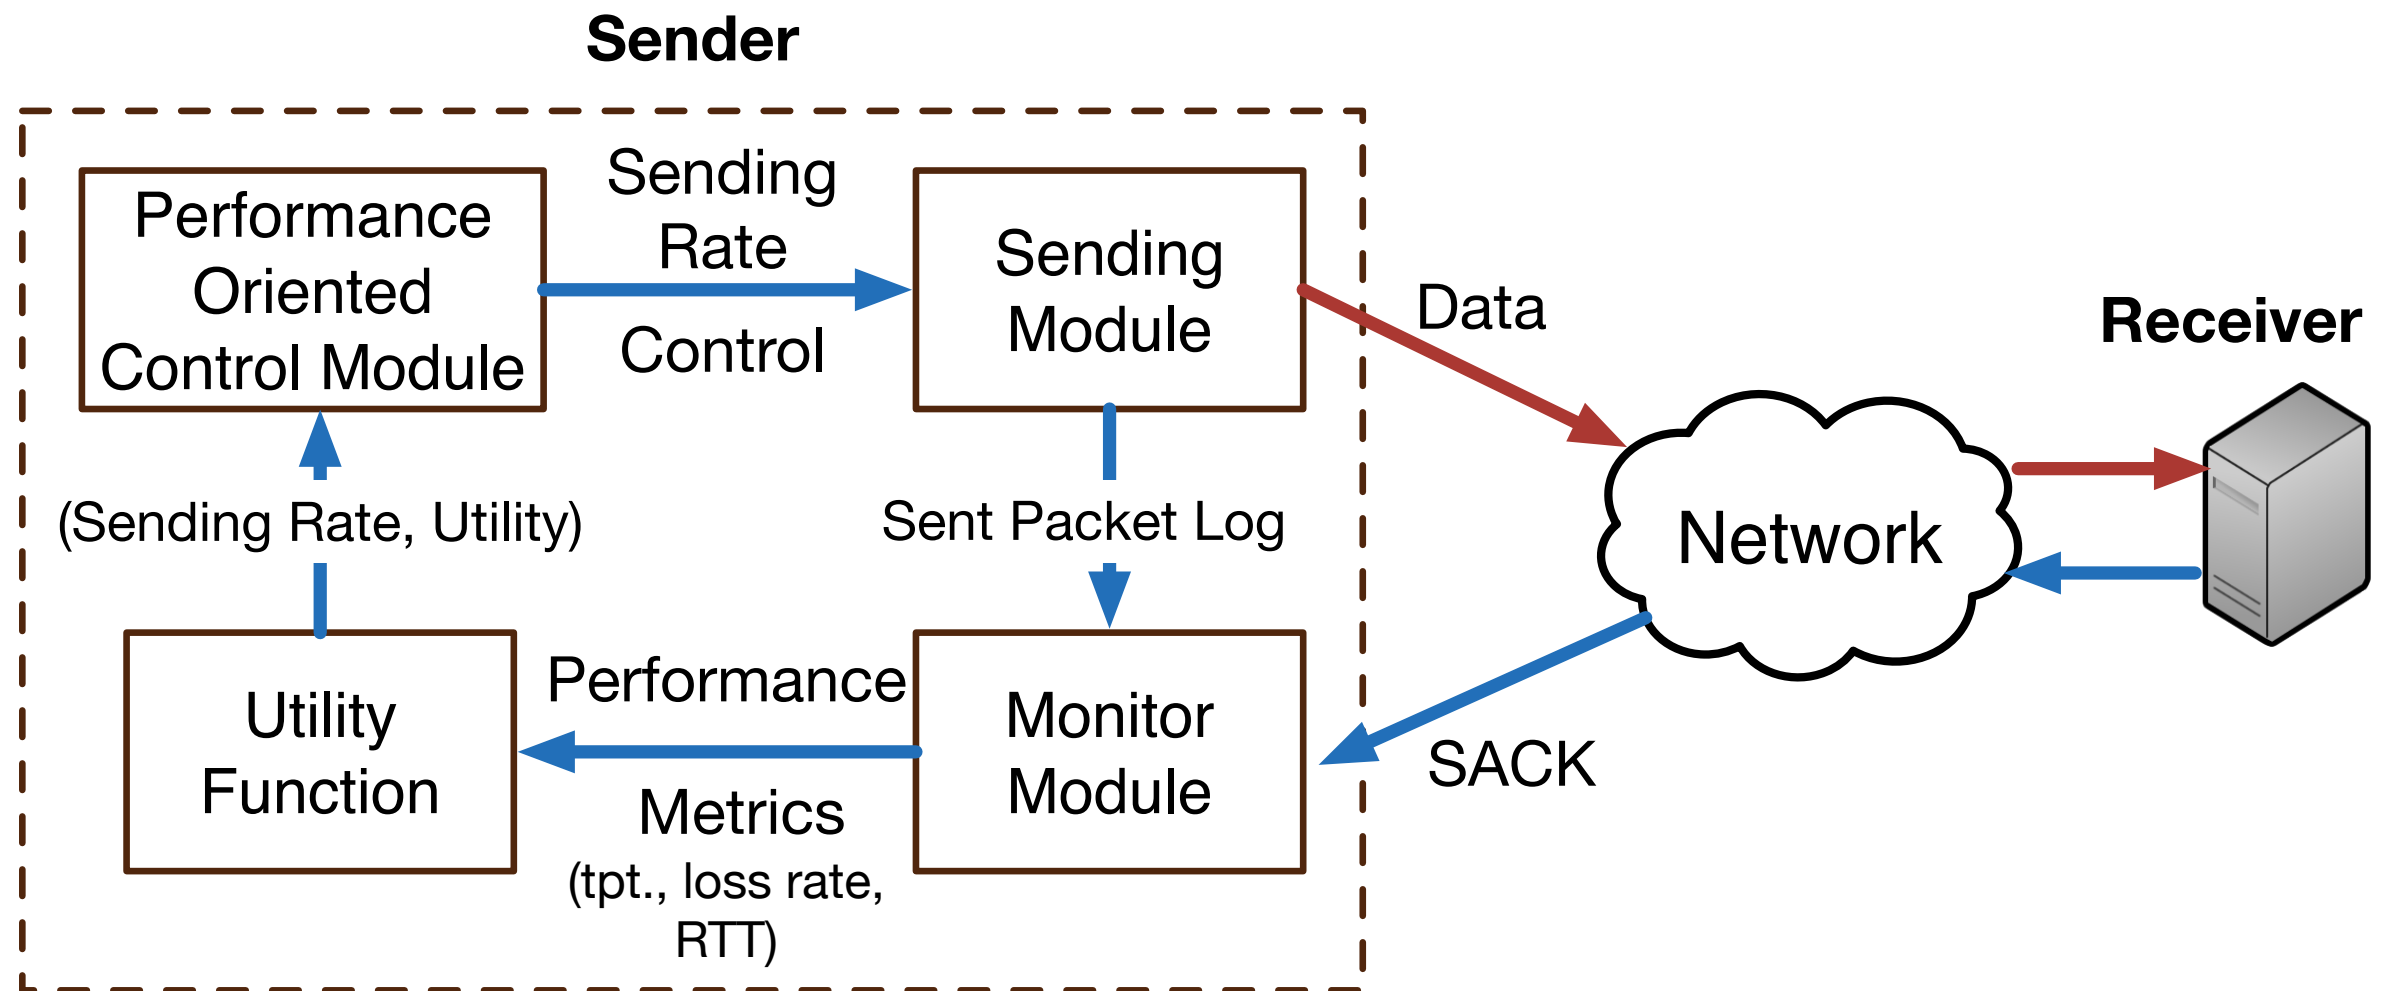


speedier.net/pcc

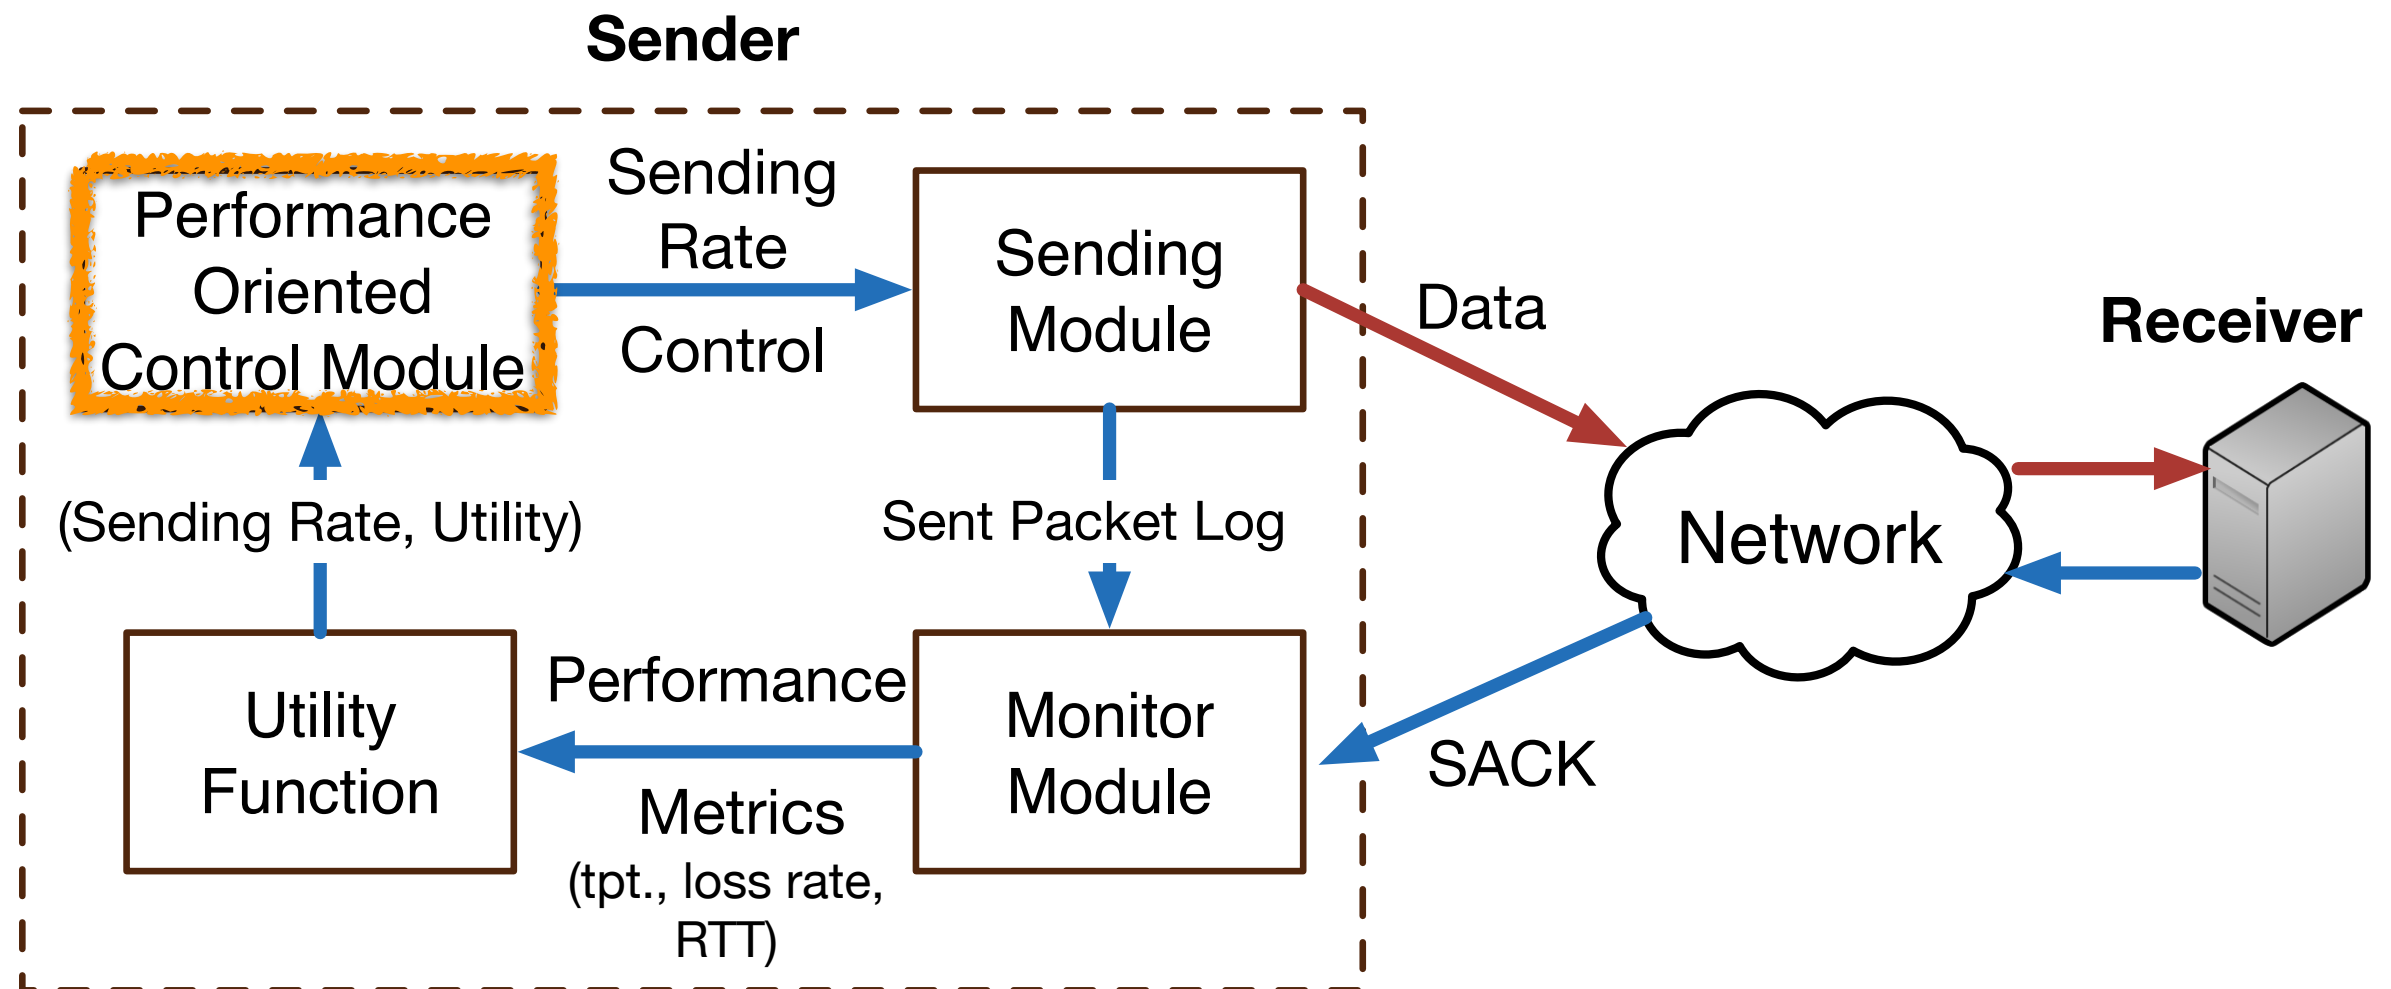


Backup Slides

Software Components




Software Components




Performance Oriented Control




Performance Oriented Control



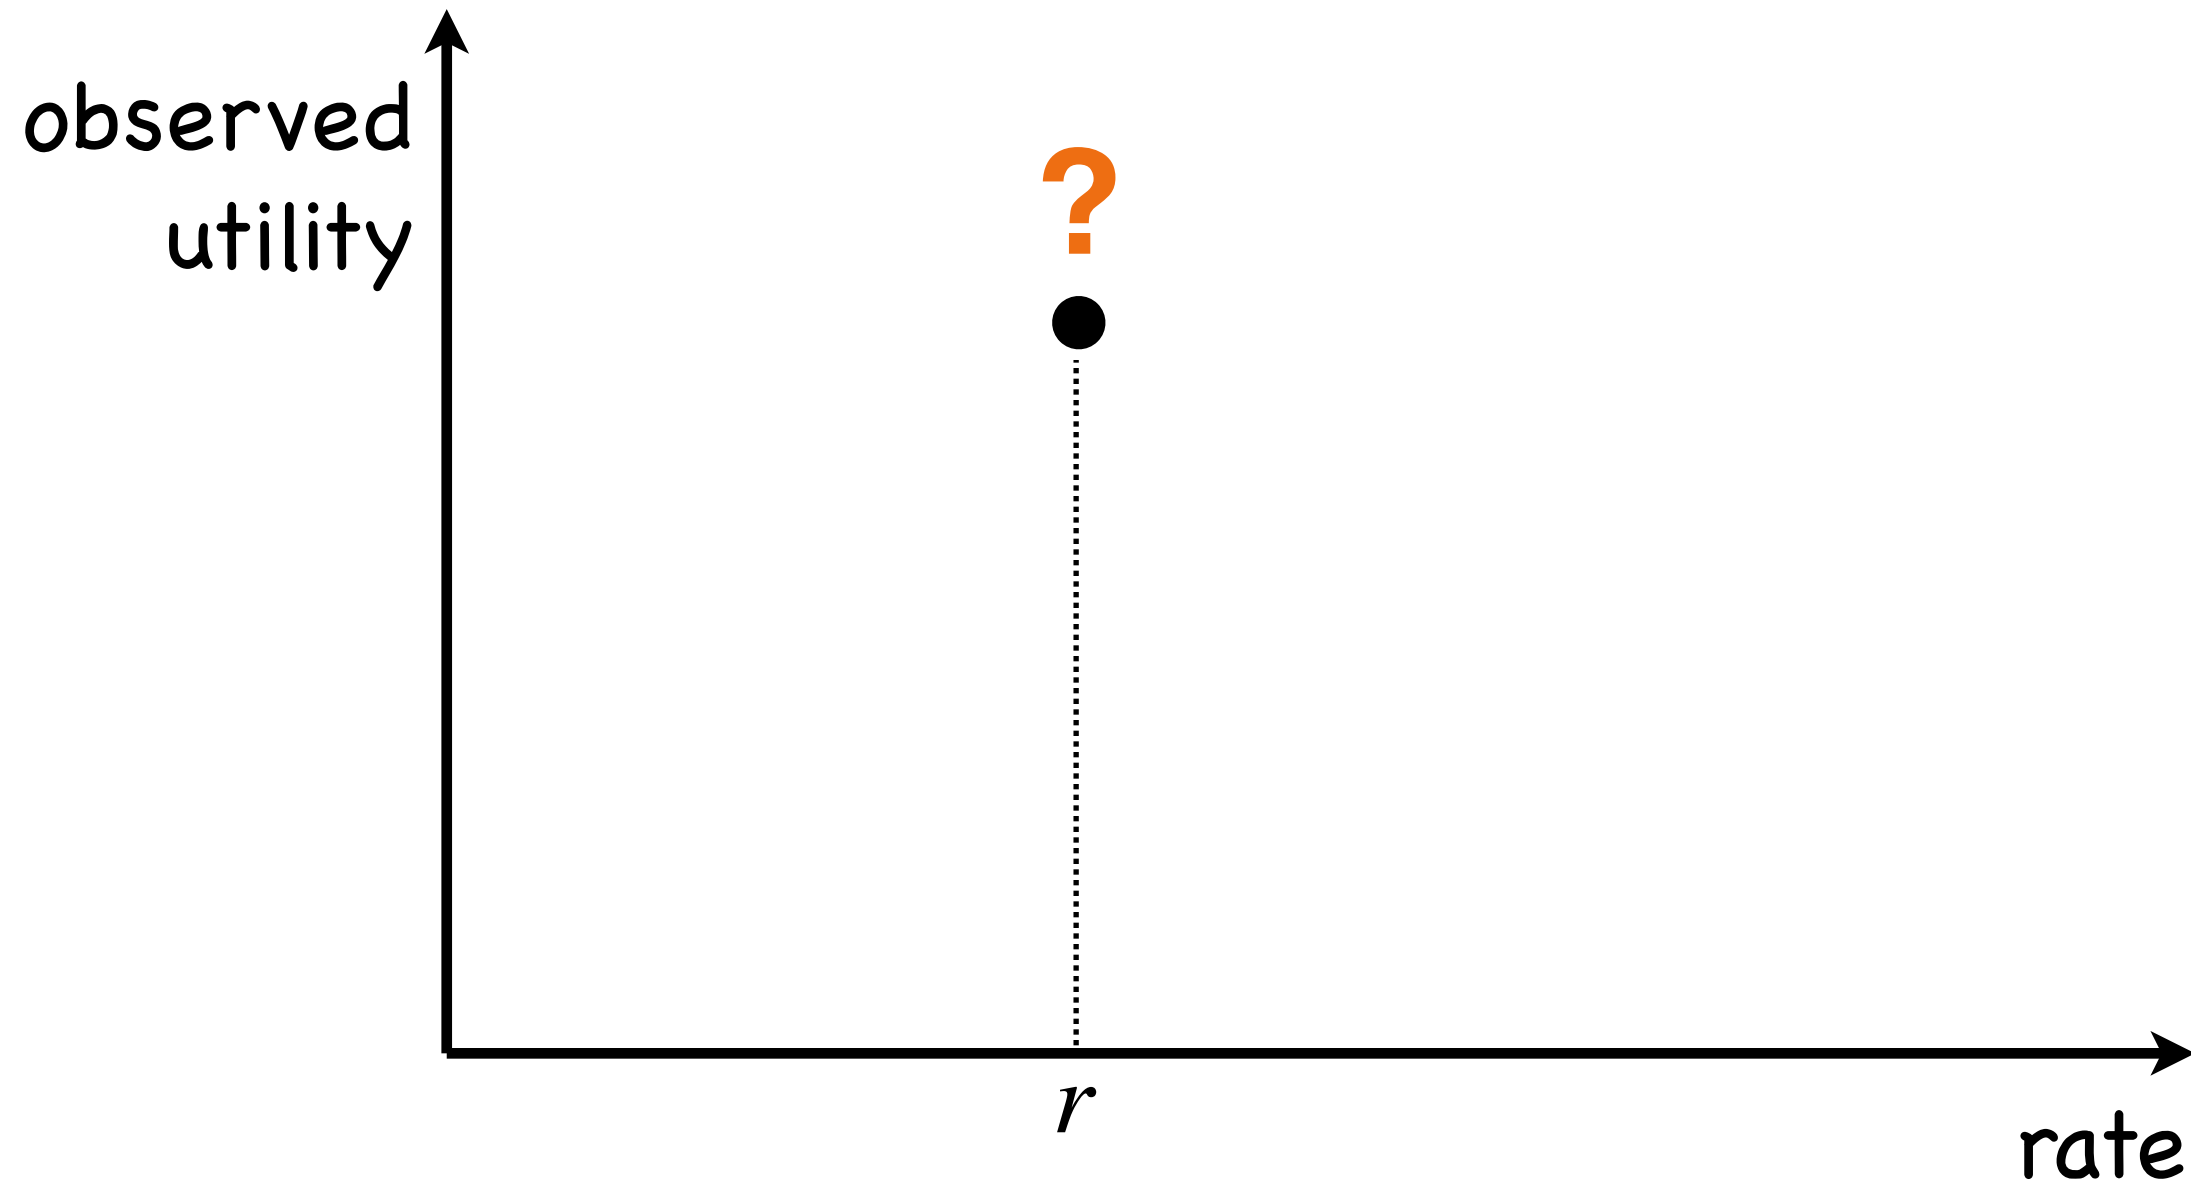
Performance Oriented Control



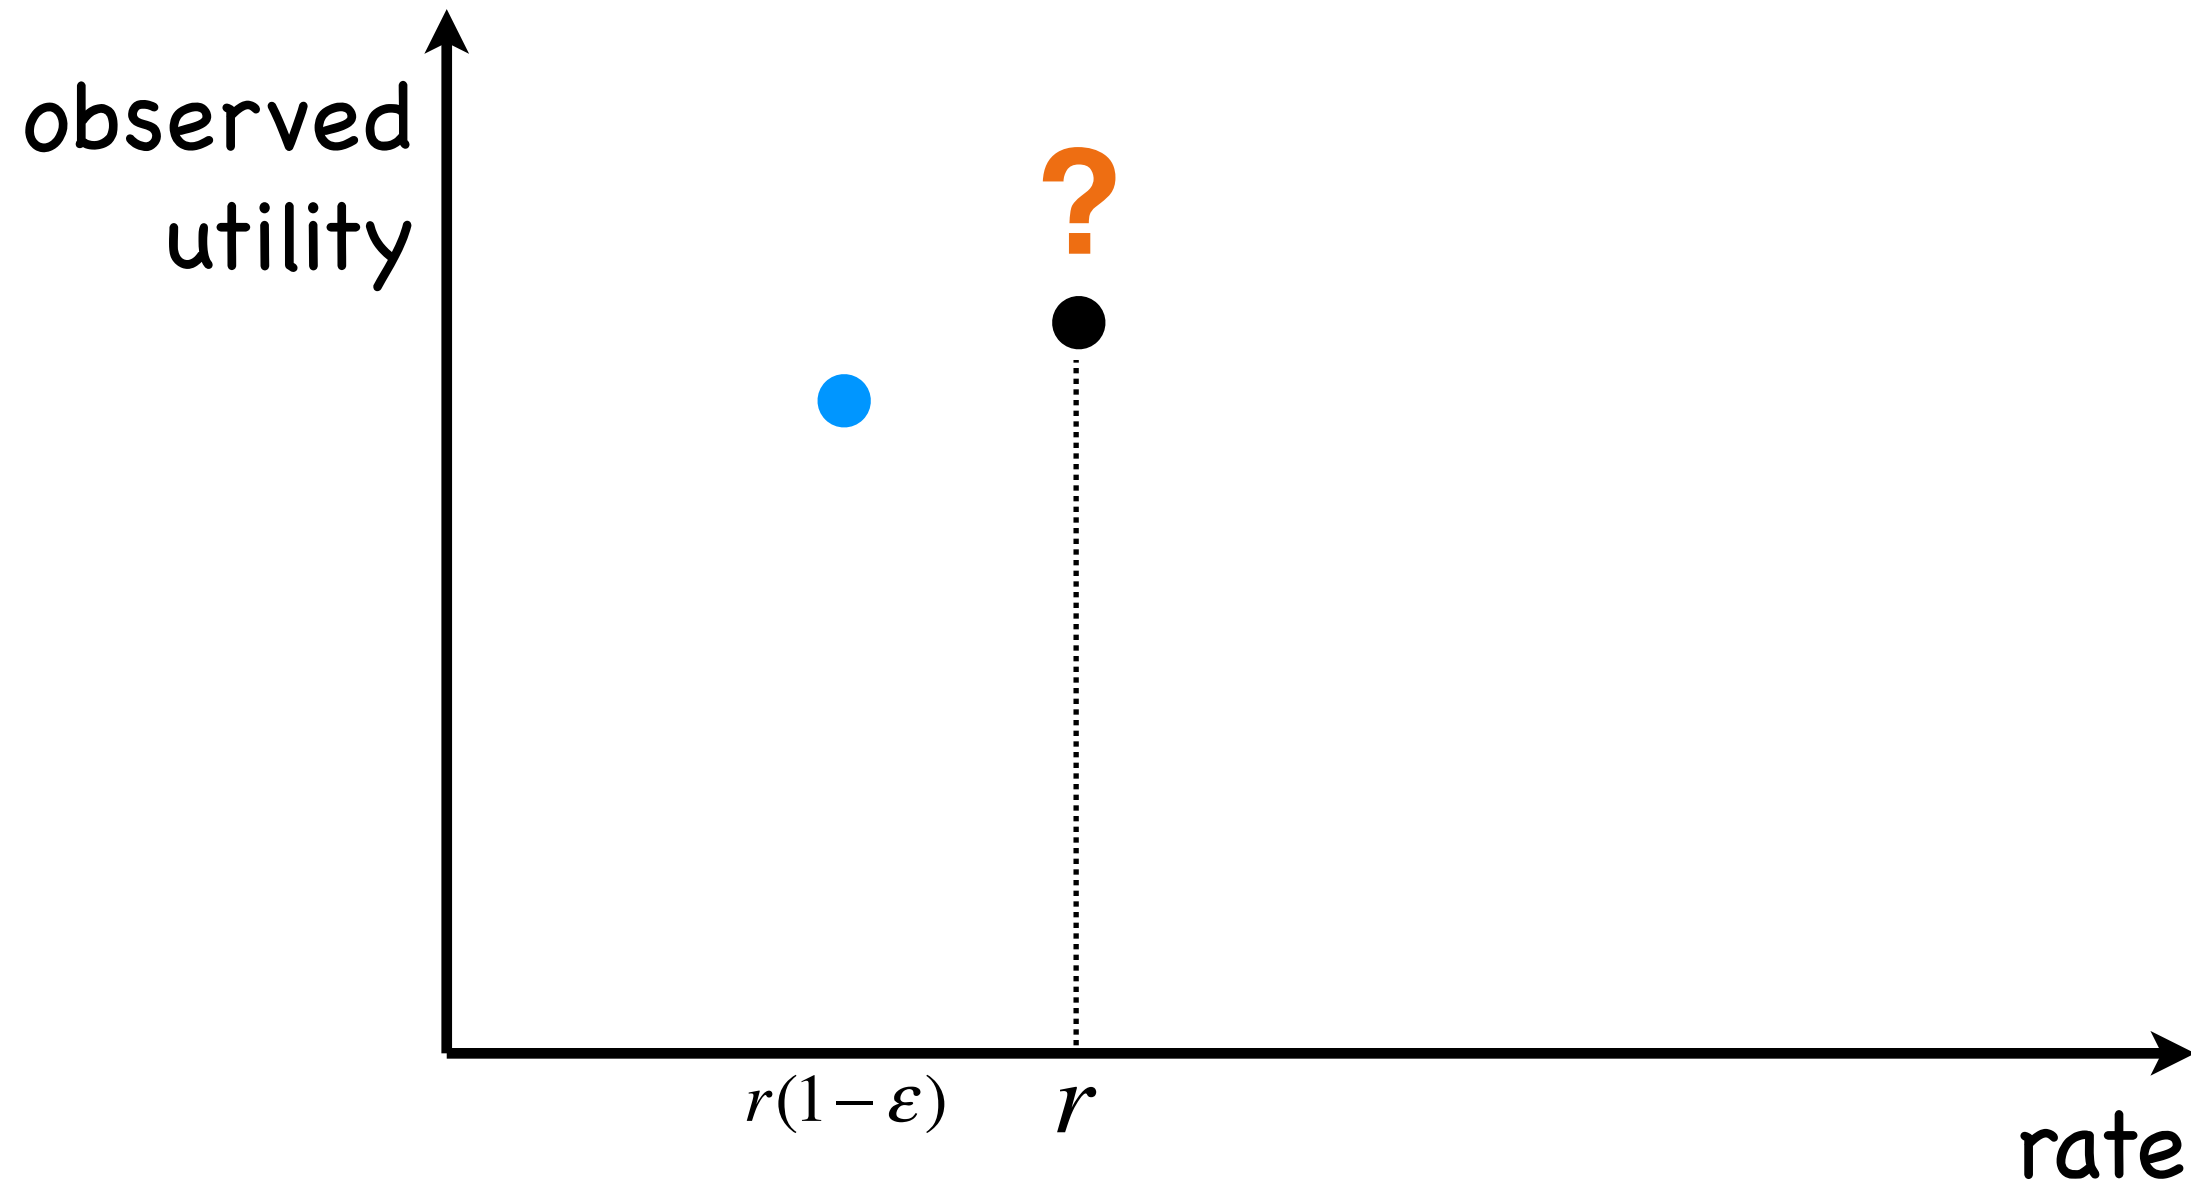
Performance Oriented Control




Performance Oriented Control




Performance Oriented Control




Performance Oriented Control




Performance Oriented Control




Performance Oriented Control




Performance Oriented Control




Performance Oriented Control




Performance Oriented Control




Online Learning Control




Online Learning Control




Online Learning Control



Online Learning Control



Online Learning Control



Deployment

- No hardwired support, packet header, protocol change needed
- Where to deploy
 - CDN backbone, Inter-data center, dedicated scientific nw
 - “In the Wild”?

TCP Friendliness

TCP Friendliness

- Is PCC TCP-friendly?

TCP Friendliness

- Is PCC TCP-friendly?



TCP Friendliness

TCP Friendliness

- PCC's default utility function is not TCP Friendly

TCP Friendliness

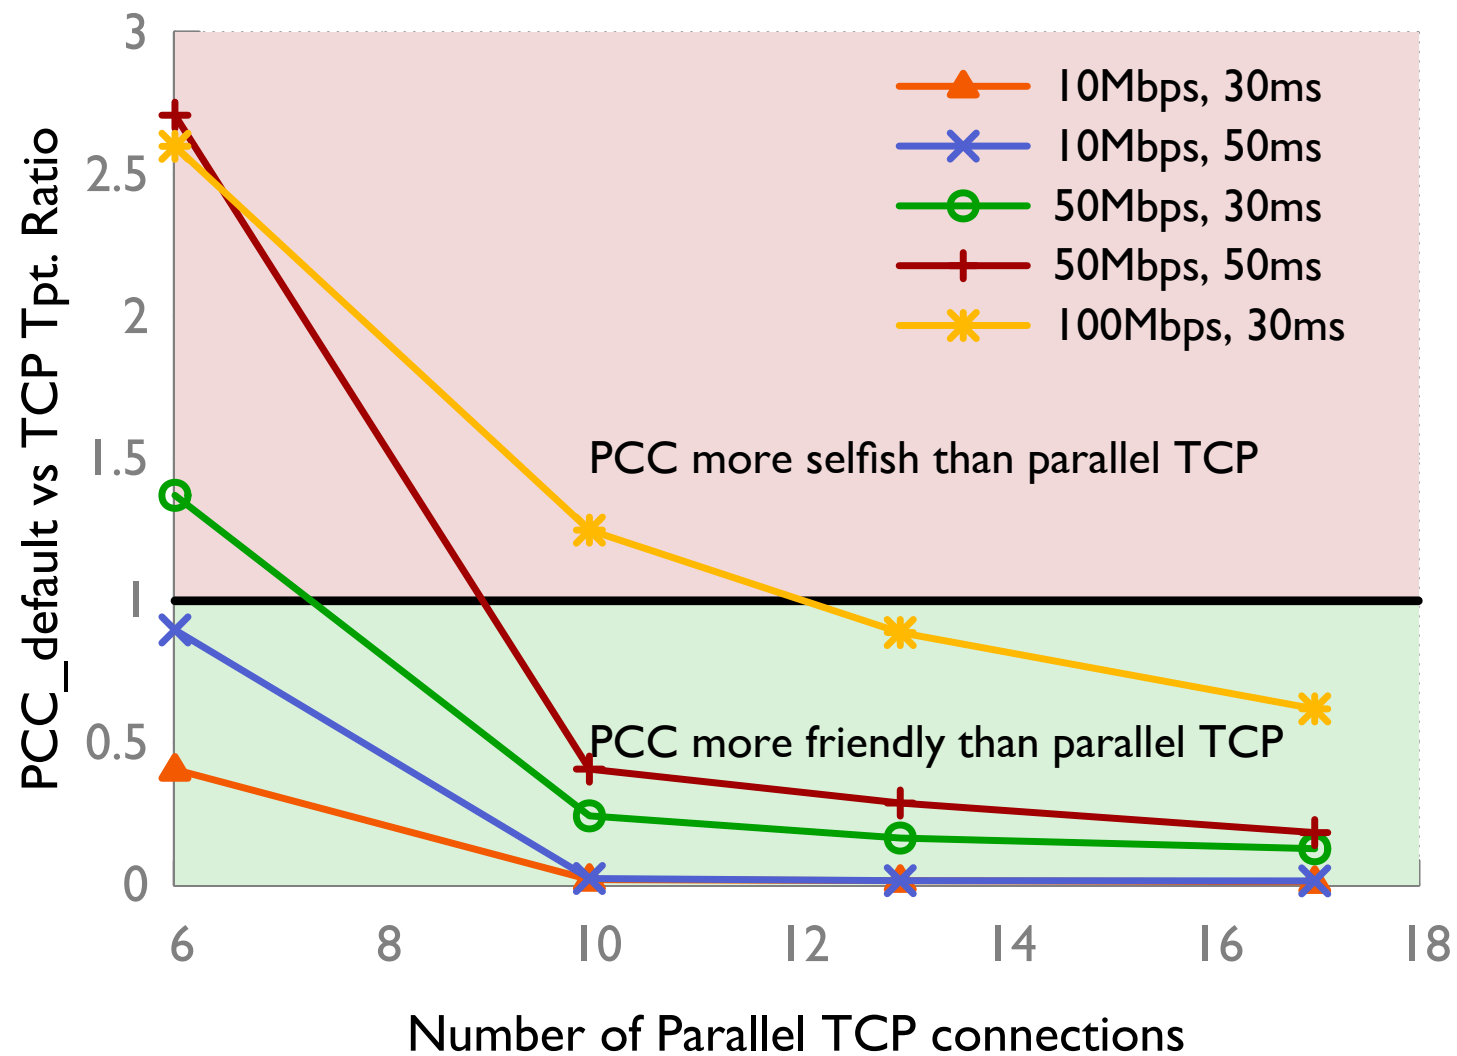
- PCC's default utility function is not TCP Friendly

But not **that** bad

TCP Friendliness

- PCC's default utility function is not TCP Friendly


But not **that** bad



TCP Friendliness

- PCC's default utility function is not TCP Friendly

But not **that** bad



TCP Friendliness

- PCC's default utility function is not TCP Friendly
- Different utility functions can be a solution

TCP Friendliness

- PCC's default utility function is not TCP Friendly
- Different utility functions can be a solution

$$u_i(x) = \frac{T_i * \text{sigmoid}_\alpha(L_i - 0.05) * \text{sigmoid}_\beta\left(\frac{rtt_n}{rtt_{n-1}} - 1\right) - x_i * L_i}{rtt_n}$$

TCP Friendliness

- PCC's default utility function is not TCP Friendly
- Different utility functions can be a solution

TCP Friendliness

- PCC's default utility function is not TCP Friendly
- Different utility functions can be a solution

TCP vs TCP

TCP vs PCC

TCP Friendliness

- PCC's default utility function is not TCP Friendly
- Different utility functions can be a solution

TCP vs TCP TCP vs PCC



TCP Friendliness


- PCC's default utility function is not TCP Friendly
- Different utility functions can be a solution

TCP vs TCP  TCP vs PCC

		30ms	60ms	90ms
$\beta = 10$	10Mbit/s	0.94	0.75	0.67
	50Mbit/s	0.74	0.73	0.81
	90Mbit/s	0.89	0.91	1.01
$\beta = 100$	10Mbit/s	0.71	0.58	0.63
	50Mbit/s	0.56	0.58	0.54
	90Mbit/s	0.63	0.62	0.88


Deployability: Short Flow FCT

PCC does not fundamentally harm FCT




Consistent High Performance

Global Commercial Internet




Consistent High Performance

Global Commercial Internet




Consistent High Performance

Global Commercial Internet




Consistent High Performance

Global Commercial Internet




Consistent High Performance

Global Commercial Internet




Consistent High Performance

Global Commercial Internet



Consistent High Performance


Global Commercial Internet



Consistent High Performance


Intra-continent
2.33X median

Global Commercial Internet




Consistent High Performance

Global Commercial Internet




Consistent High Performance

Global Commercial Internet




Consistent High Performance

Satellite Network



WINDS System




Consistent High Performance

Lossy Networks

100Mbps, 30ms, varying loss rate

TCP's throughput collapse 10X with 0.5% loss




Consistent High Performance

Treat RTT Unfairness

100Mbps, 10ms short RTT flow


Varying RTT of long RTT one



Consistent High Performance

A possible solution to the bufferbloat problem

100Mbps, 30ms Link



Consistent High Performance

Inter Datacenter and Dedicated High Speed Network

Transmission Pair	RTT	PCC	SABUL	CUBIC	Illinois
GPO → NYSErNet	12.1	818	563	129	326
GPO → Missouri	46.5	624	531	80.7	90.1
GPO → Illinois	35.4	766	664	84.5	102
NYSErNet → Missouri	47.4	816	662	108	109
Wisconsin → Illinois	9.01	801	700	547	562
GPO → Wisc.	38.0	783	487	79.3	120
NYSErNet → Wisc.	38.3	791	673	134	134
Missouri → Wisc.	20.9	807	698	259	262
NYSErNet → Illinois	36.1	808	674	141	141

Consistent High Performance


Inter Datacenter and Dedicated High Speed Network

Transmission Pair	RTT	PCC	SABUL	CUBIC	Illinois
GPO → NYSErNet	12.1	818	563	129	326
GPO → Missouri	46.5	624	531	80.7	90.1
GPO → Illinois	35.4	766	664	84.5	102
NYSErNet → Missouri	47.4	816	662	108	109
Wisconsin → Illinois	9.01	801	700	547	562
GPO → Wisc.	38.0	783	487	79.3	120
NYSErNet → Wisc.	38.3	791	673	134	134
Missouri → Wisc.	20.9	807	698	259	262
NYSErNet → Illinois	36.1	808	674	141	141

123%

Consistent High Performance

Mitigate Incast



Long list of things we have done but don't have time to talk about

