EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

FACULTY OF
SCIENCE
Communication Networks

Activity-Based Congestion Management for Fair Bandwidth Sharing in Trusted Packet Networks

Michael Menth and Nikolas Zeitler

http://kn.inf.uni-tuebingen.de

Outline

- The problem
- Definition of activity
- ABC principle
- Experiment setup
- Simulation results
- Related work
- Use cases
- Conclusion & outlook

<u>Inspired</u> by congestion policing using congestion exposure (ConEx) signals

The Problem

- Scenario
  - Bottleneck link shared by multiple users, congestion may occur
  - Flows may be transmitted over different transport protocols
    $\Rightarrow$ UDP flows can achieve higher rates than TCP flows
  - TCP users may have different number of flows
    $\Rightarrow$ Users with more flows can achieve higher rates (as TCP attempts to be per-flow fair)
  - How to share bandwidth fairly among users?

- If pkts can be related to users
  - Weighted fair queueing (WFQ) may enforce per-user fair bandwidth sharing

- If pkts cannot be related to users
  - WFQ does not work
  - ABC may be used

ABC Principle – Architecture

- Per-domain concept
  - Configuration
  - Trust
- Access nodes
  - Connect user equipment
  - Have **activity meter**
  - Add activity A to packet headers
- Core and access nodes
  - Run **ABC-AQM** (active queue management)
  - Drop packets during congestion depending on their activity
- Core nodes do not hold per-user/flow states

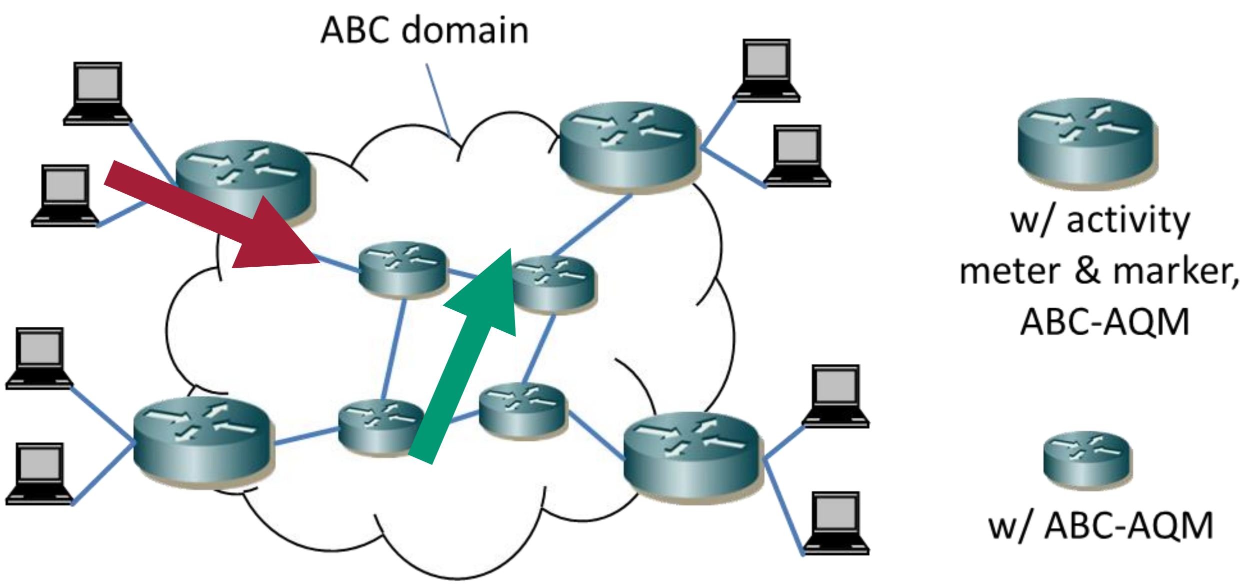

ABC Principle – Activity Meter

- Definition of activity
  - R_r: a user‘s assigned reference rate
  - R_t: a user‘s recent transmission rate
  - $A = log_2\left(\frac{R_t}{R_r}\right)$: a user‘s activity
  - Indicates how much a users transmission rate exceeds its reference rate

- Activity meter
  - Computes activity A per packet and adds it to the header
  - Implemented through a token-based algorithm
  - Configured with
    - Reference rate R_r (per user)
    - Inertia I (time scale parameter, per domain)
    - Burst allowance B (to ignore first B bytes after silence, per domain)

ABC Principle – ABC-AQM

- Operation per packet arrival
  - Track average activity A_{avg}
  - Drop packet depending on
    - Link congestion
    - Activity A in packet header

- Average activity: $A_{avg} = w_A \cdot A_{avg} + (1 - w_A) \cdot A$

- Use some existing AQM based on drop probabilities

- Adapt drop probability: $p_A = p^{2^{-\gamma \cdot (A - A_{avg})}}$
  - Differentiation factor $\gamma \geq 0$
  - $p_A > p$ for large activities
  - $p_A < p$ for small activities

AQM Used in the Study (Exchangeable)

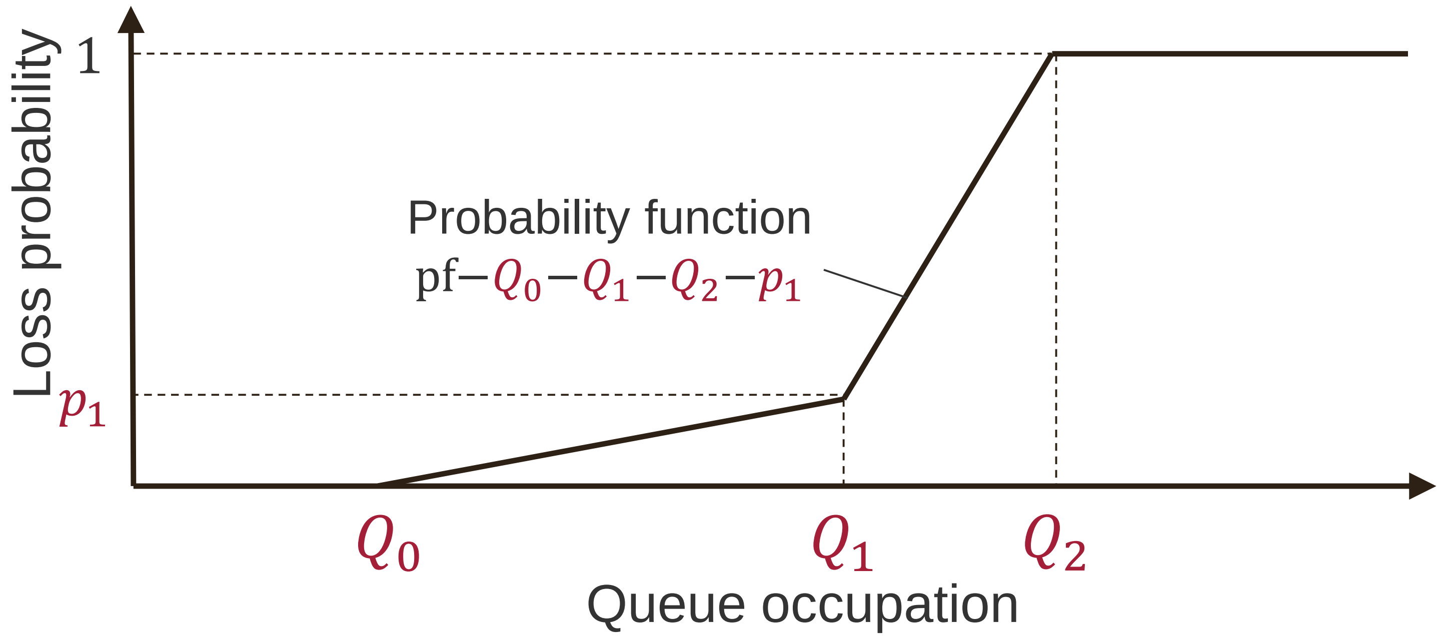

$$p(Q) = \begin{cases} 0 & Q < Q_0 \\ \dfrac{Q - Q_0}{Q_1 - Q_0} \cdot p_1 & Q_0 < Q \leq Q_1 \\ p_1 + \dfrac{Q - Q_1}{Q_2 - Q_1} \cdot (1 - p_1) & Q_1 < Q \leq Q_2 \\ 1 & Q_2 < Q \end{cases}$$

Experiment Setup

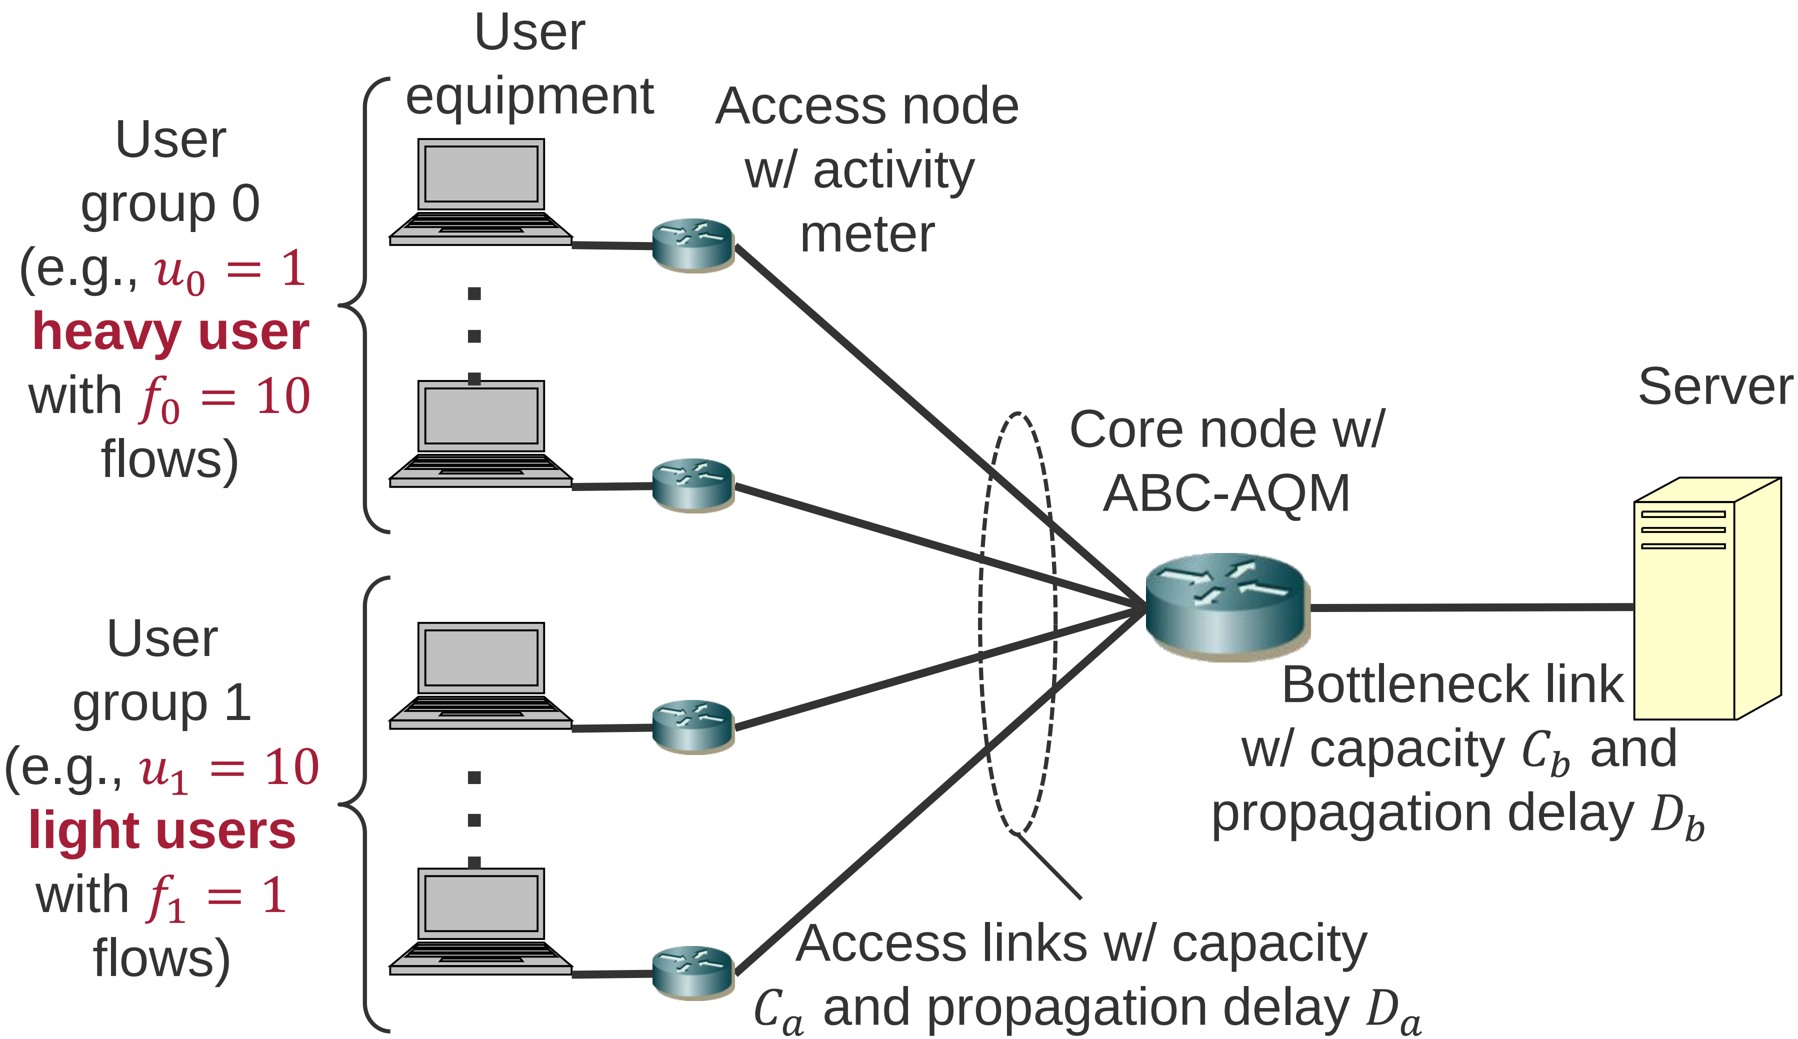

Impact of Differentiation Parameter γ

- Experiment
  - User group 0: 1 heavy user with 10 TCP flows
  - User group 1: 10 light users with only 1 TCP flow
- Performance metrics
  - Average throughput per user
    - $\overline{T_0}$: for user group 0
    - $\overline{T_1}$: for user group 1
  - Throughput ratio $T_R = \frac{\overline{T_0}}{\overline{T_1}}$

Throughput ratio T_R depending on differentiation factor γ.

No differentiation, no fairness — Almost fair

D_b (ms)	users u_0/u_1	differentiation factor γ					
		0	1	2	3	4	5
50	1/10	10.02	2.68	1.71	1.36	1.20	1.10
	9/90	10.20	1.51	1.03	0.88	0.82	0.80
5	1/10	10.14	1.90	1.26	1.03	0.94	0.89
	9/90	10.76	1.37	1.01	0.91	0.87	0.84

- Use $\gamma = 3$ to achieve good fairness, confirmed by more experiments

Impact of Reference Rate R_r

- Vary reference R_r rate per user
- ABC enforces fairness if $R_r \leq \frac{C}{u}$
  - C: bottleneck bandwidth
  - u: number of users on bottleneck link

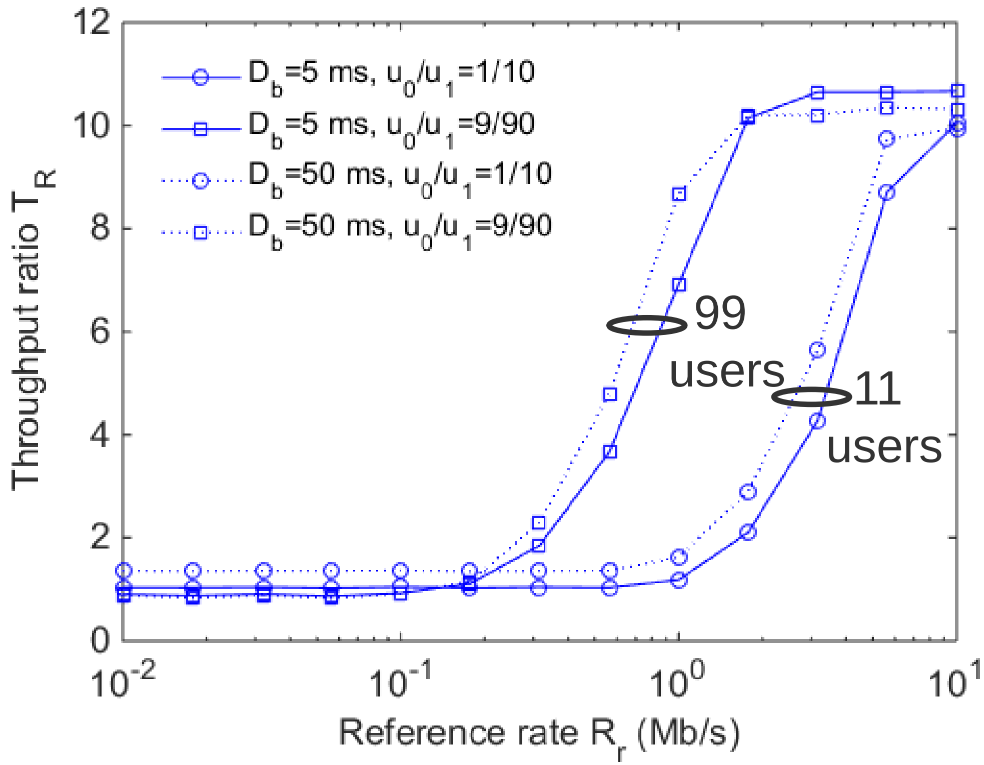

Protection against Non-Responsive Traffic

- 1 UDP user vs. 10 users with 1 TCP flow
- ABC protects TCP users against UDP user
- $\gamma = 3$ best suited
- Improvement through better AQM possible

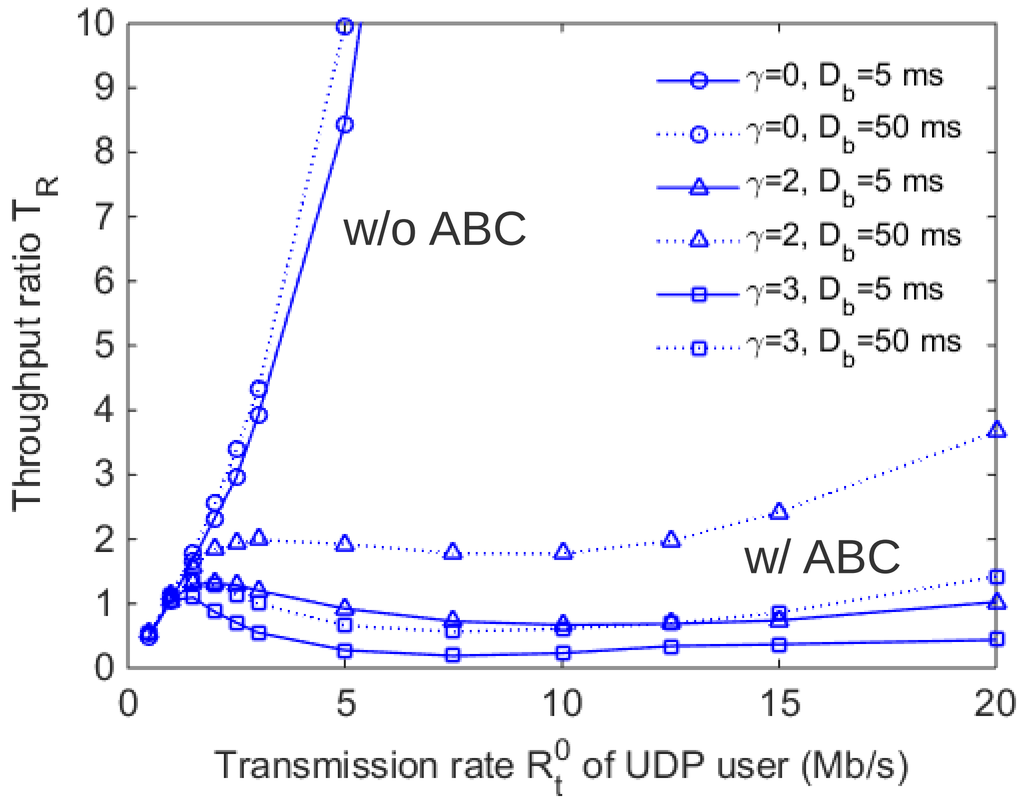

Reduced Upload Time

- Experiment setup
  - $u_0 = 1$ probe user sending a probe of 1 MB
  - u_1 background users with f_1 flows each
  - Performance metric: upload time
  - Link delay D_b
    - 50 ms
    - 5 ms
- Upload time
  - Increases with u_1 and f_1
  - Shorter with ABC
  - Decreases with increasing **burst tolerance B**

D_b (ms)	ABC	B (MB)	$f_1 = 1$		$f_1 = 4$	
			$u_1 = 10$	$u_1 = 20$	$u_1 = 10$	$u_1 = 20$
50	no	n/a	20.3	29.2	47.7	85.2
	yes	0	9.6	17.9	12.3	21.7
		0.25	8.2	15.7	10.4	17.6
		0.5	6.1	10.7	8.3	14.0
		0.75	2.2	5.7	6.7	8.1
		1	2.4	2.7	3.3	5.4
		1.25	2.1	3.3	3.3	5.2
5	no	n/a	10.3	20.2	34.2	61.4
	yes	0	9.6	17.6	10.1	18.8
		0.25	6.9	13.7	8.3	14.0
		0.5	4.6	9.5	5.6	10.0
		0.75	1.3	5.4	4.0	6.4
		1	1.0	1.3	1.0	1.4
		1.25	1.0	1.0	1.0	1.0

Related Work

- CSFQ
  - [28] I. Stoica, S. Shenker, and H. Zhang: „Core-Stateless Fair Queuing: A Scalable Architecture to Approximate Fair Bandwidth Allocations in High-Speed Networks“, IEEE/ACM Transactions on Networking, 11(1), Feb. 2003
  - Similar to ABC but measures traffic rate on link to determine drop probabilities $\Rightarrow$ slower reaction than ABC

- ABC
  - Uses buffer occupation and AQM to determine drop probabilities
  - Extensible towards low-delay communication
    $\Rightarrow$ Enforce both fairness and low-delay

Use Cases

- ABC fulfills objectives of congestion policing using congestion exposure (ConEx)
  - Fair rate sharing
  - Improved QoE for light users (reduced upload times)

- ABC supports the same uses cases as ConEx
  - Datacenter
    draft-briscoe-conex-data-centre
  - Mobile access networks
    RFC7778: Mobile Communication Congestion Exposure Scenario

- ABC protects light users against heavy UDP users
  - Another use case: defense against denial of service attacks

Conclusion & Outlook

- Activity-based congestion management (ABC)
  - Access node add activity into to packet headers
  - Access and core nodes use it for preferential dropping
  - Enforces per-user fairness (heavy and light users)
  - Protocol-independent, stateless core network
- Presented simulation results
  - Impact of differentiation factor, reference rate
  - Protection against non-responsive traffic
  - Reduced upload time
- ABC supports
  - Different users profiles (reference rates)
  - ConEx use cases (and more)
- Outlook
  - Combine ABC with PIE to achieve fairness and low latency
  - Adapt ABC towards service classes

Publication

- M. Menth and N. Zeitler: Activity-Based Congestion Management for Fair Bandwidth Sharing in Trusted Packet Networks, in Proceedings of the IEEE/IFIP Network Operations and Management Symposium (NOMS), April, 2016, Istanbul, Turkey

- Preprint:
https://atlas.informatik.uni-tuebingen.de/~menth/papers/Menth16b.pdf