EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

FACULTY OF
SCIENCE
Communication Networks

ConEx-Based Congestion Policing – First Performance Results

Alfons Martin and Michael Menth

www.kn.inf.uni-tuebingen.de

ConEx-Based Congestion Policing

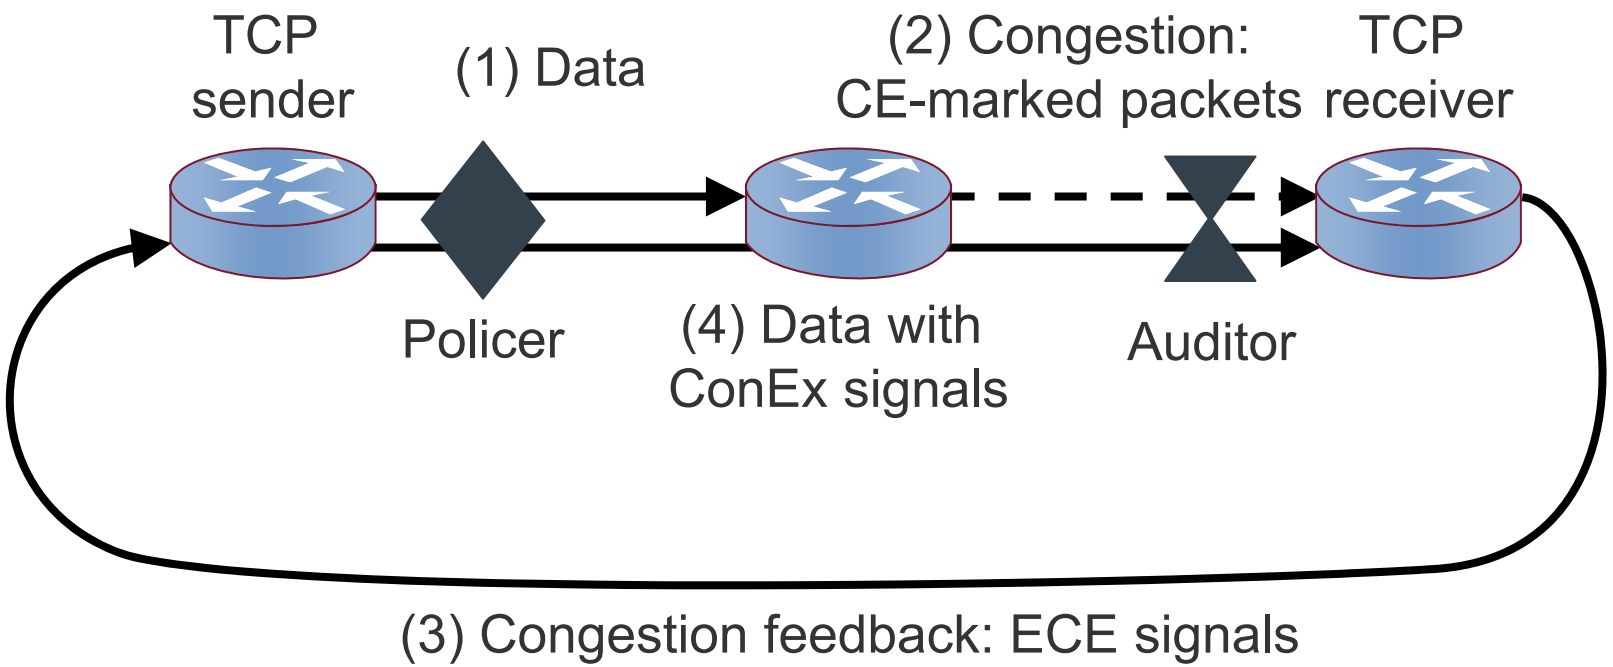

- Congestion allowance
  - Described by token bucket parameters
    - Rate
    - Tolerance
- Policer possibly drops packets of a user if bucket holds less than 1 token
  - Mild policer drops only packets with ConEx signals (default)
  - Strict policer drops all ConEx-enabled packets

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

Simulation Setup (1)

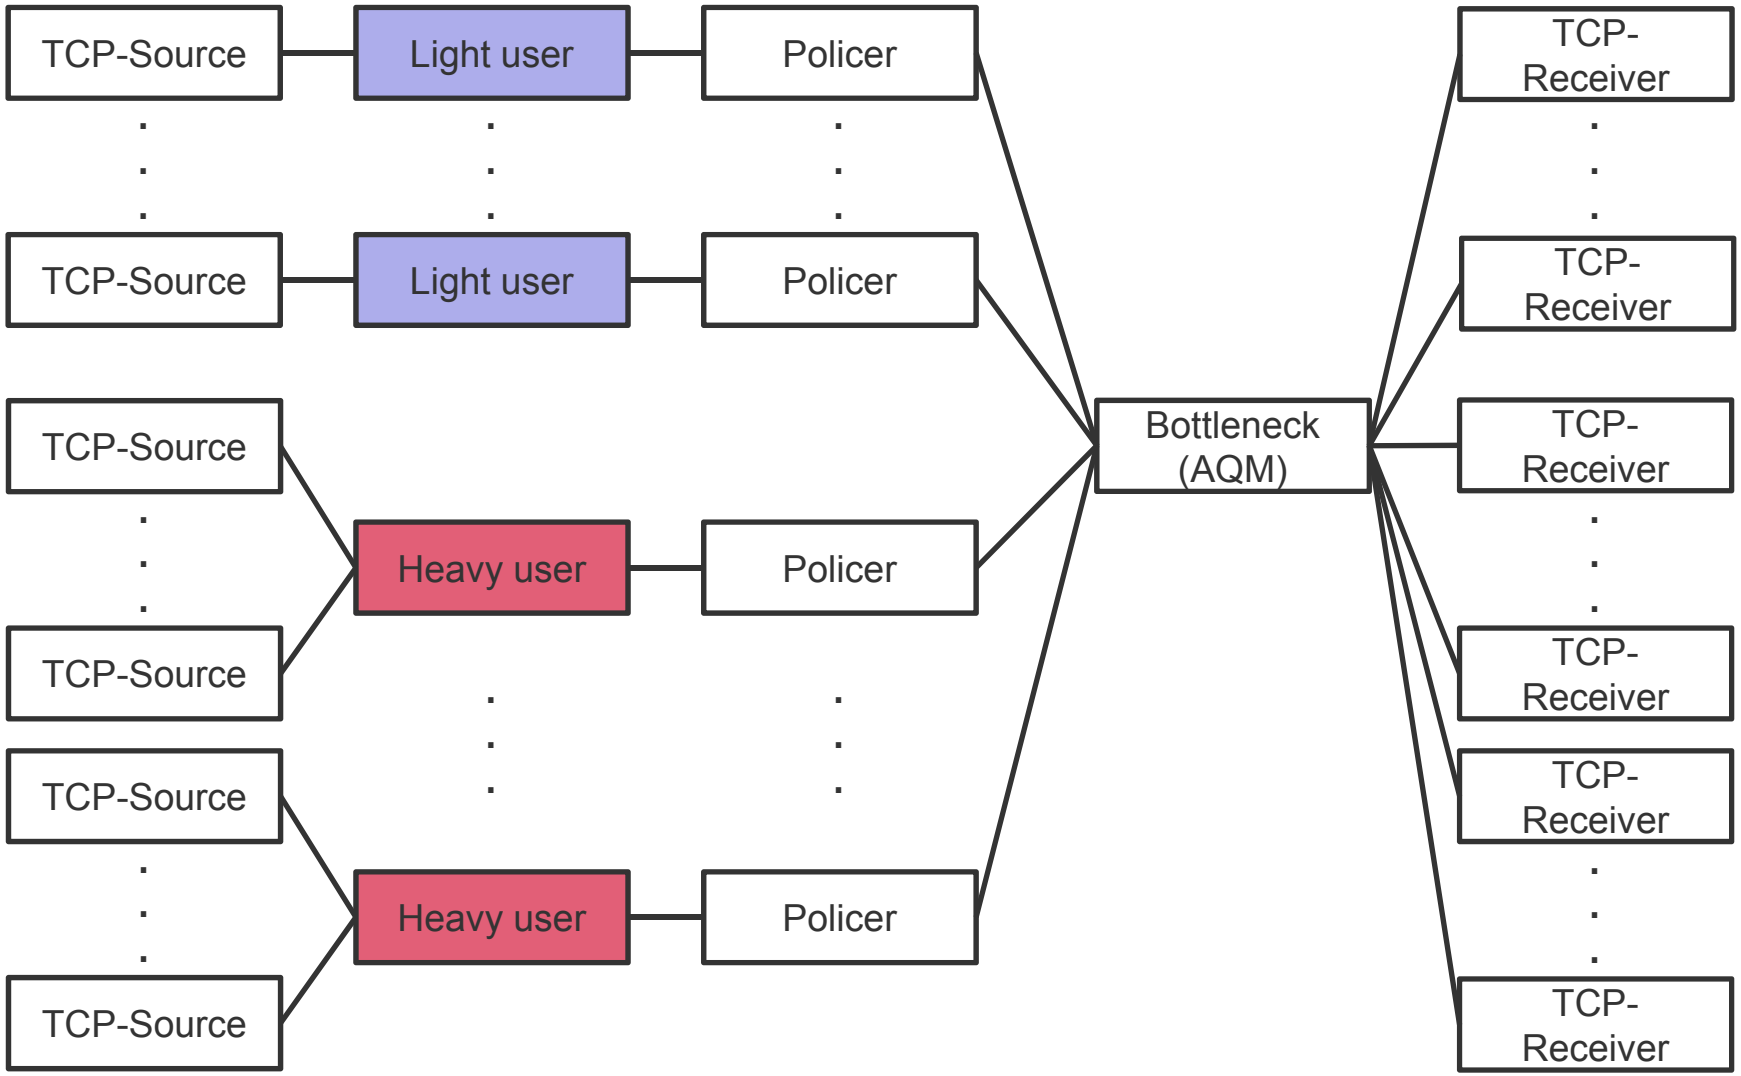

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

Simulation Setup (2)

- One-way propagation delay 20 ms
- AQM bottleneck link
  - Buffer size: 100 ms
  - Marking probability increases between 10 ms and 100 ms from 0% to 100%
- Traffic model
  - Saturated TCP sources
  - ECN-enabled TCP New Reno with selective ACK
    - RFC {793, 1122, 2018, 3168, 3782, 5681, 6298, 2883, 3517}
    - Sender notified of at most one CE per RTT $\Rightarrow$ insertion of ConEx mark
  - No background traffic on bottleneck link
- Policer
  - Allowance tolerance: 1 s
  - Allowance rate varies in experiments
- Auditor is not simulated
- Single simulation run for each data point (more to be done)

Definitions

- Measure for (un)fairness in experiments
  - $Configured\ unfairness = \frac{\#\ TCP flows\ of\ heavy\ user}{\#\ TCP flows\ of\ light\ user}$ (on bottleneck)
  - $Measured\ unfairness = \frac{Average\ throughput\ of\ heavy\ user}{Average\ throughput\ of\ light\ user}$
- Measure for effect of ConEx-based congestion policing
  - $Fairness\ improvement = \frac{Configured\ unfairness}{Measured\ unfairness} - 1$
  - Fairness improvement = 1 $\Rightarrow$ light users get 100% more bandwidth
- Measure for “configured congestion” in experiments
  - $TCP\ pressure = \frac{\#\ TCP\ flows\ on\ bottleneck}{Bottleneck\ bandwidth}$

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

Impact of Allowance Rate

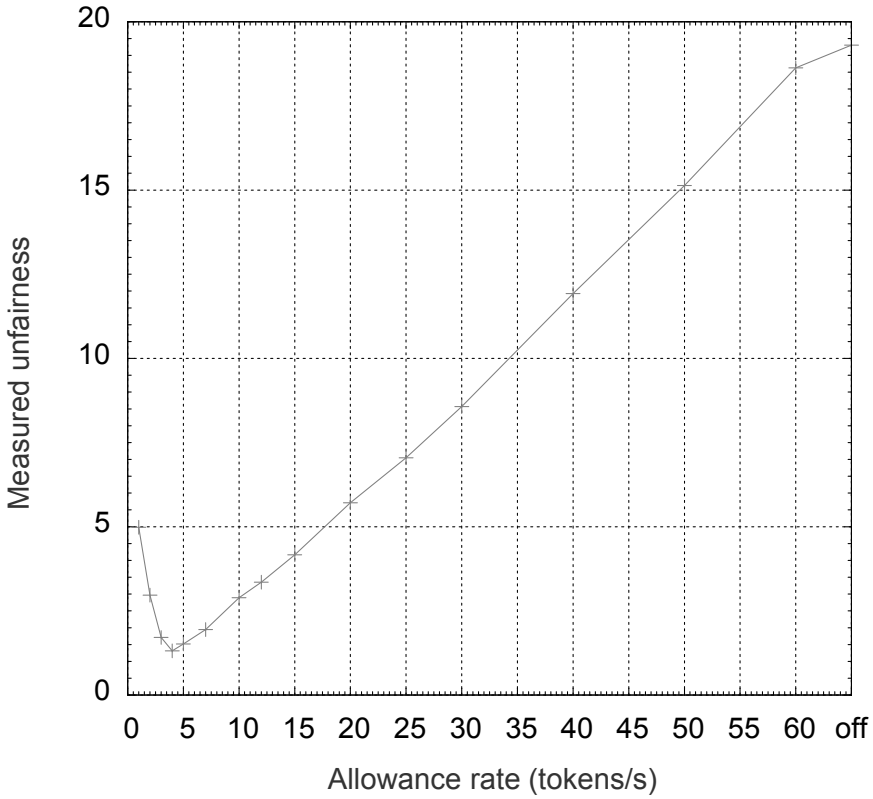

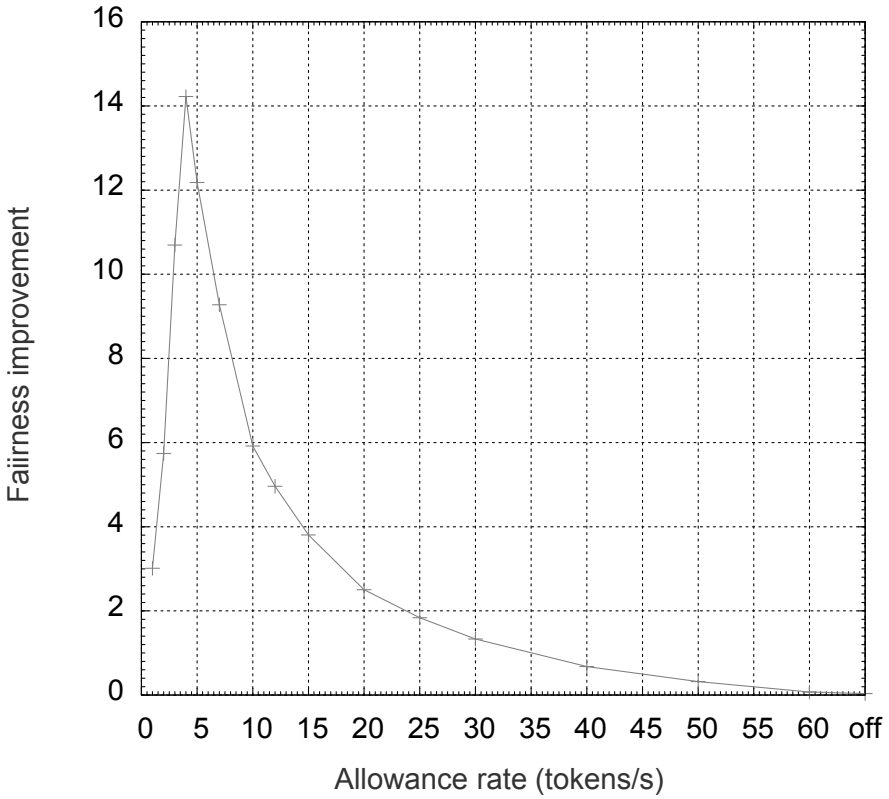

- Experiment setup
  - 10 Mb/s, 60(1)/1(20) users(flows)
    - Configured unfairness: 20
    - 80 flows $\Rightarrow 8\ \frac{flows}{Mb/s}$
- Observation
  - Significant fairness improvement for large range of allowance rates
  - Optimum allowance rate exists

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

Impact of Policer Type

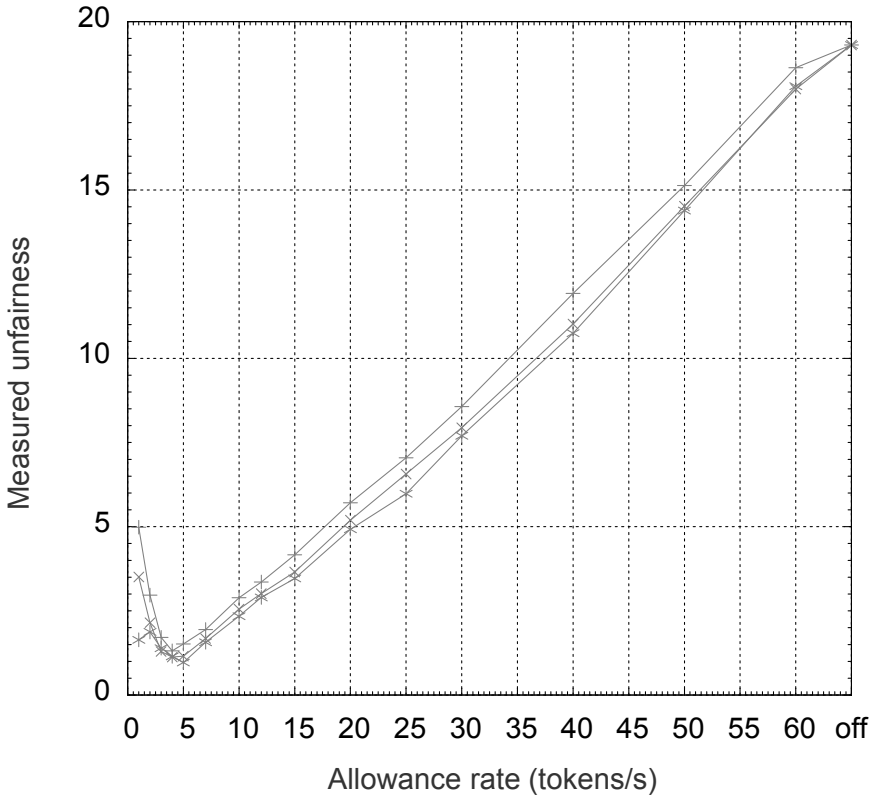

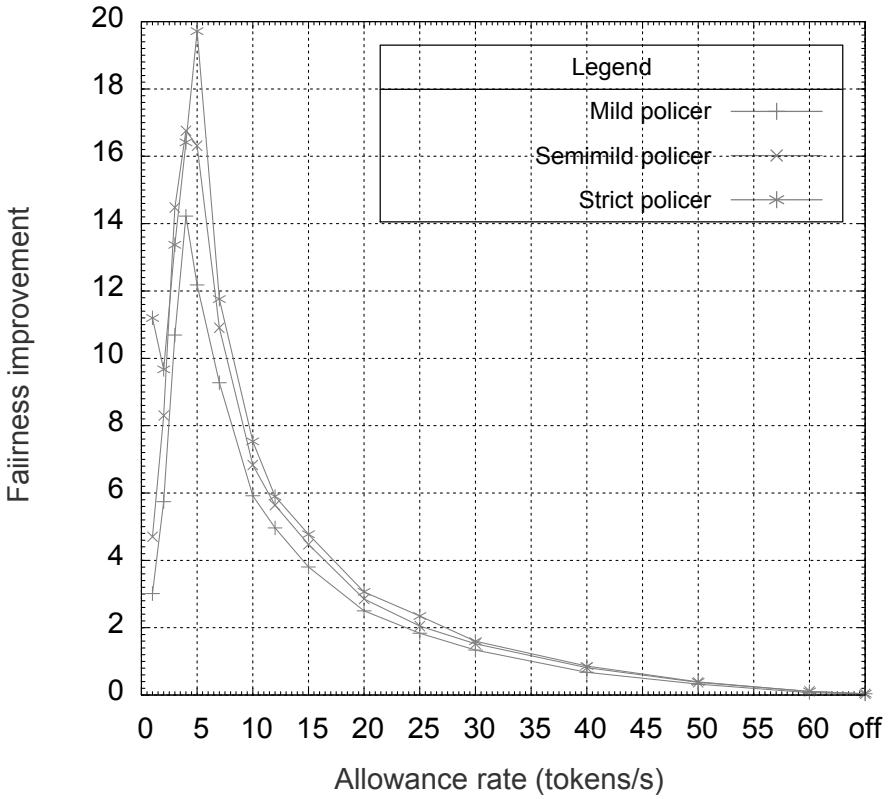

- Experiment setup
  - 10 Mb/s, 60(1)/1(20) users(flows)
    - Configured unfairness: 20
    - 80 flows $\Rightarrow 8\ \frac{flows}{Mb/s}$
  - Various policer types
- Observation
  - Stricter policer causes
    - Better fairness
    - Worse bottleneck utilization for small allowance rates
  - But differences are minor

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

Impact of TCP Pressure

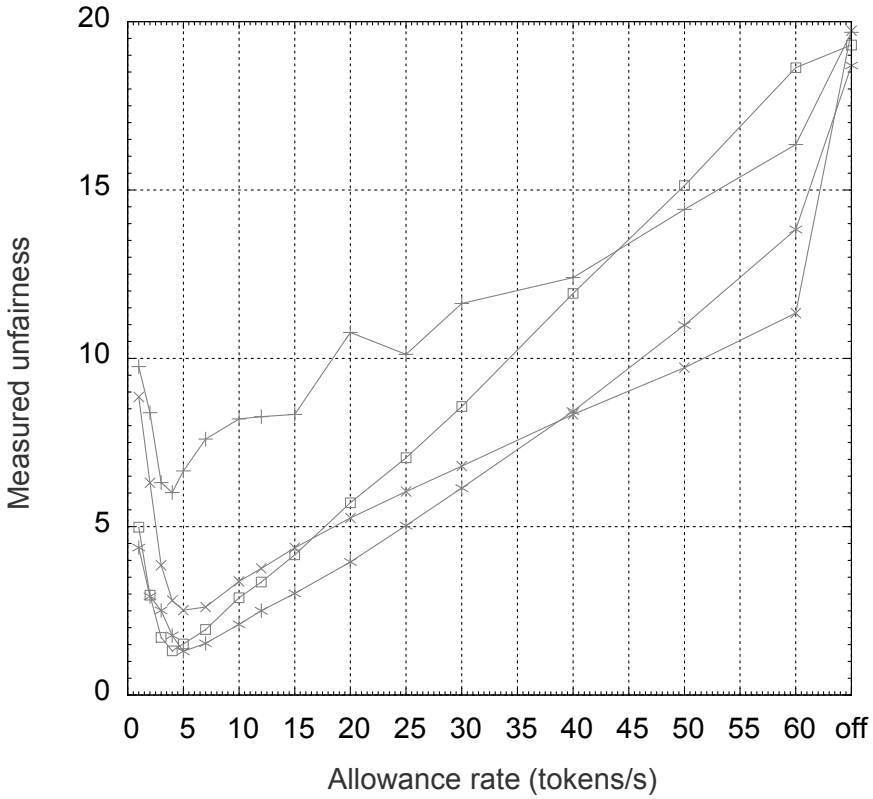

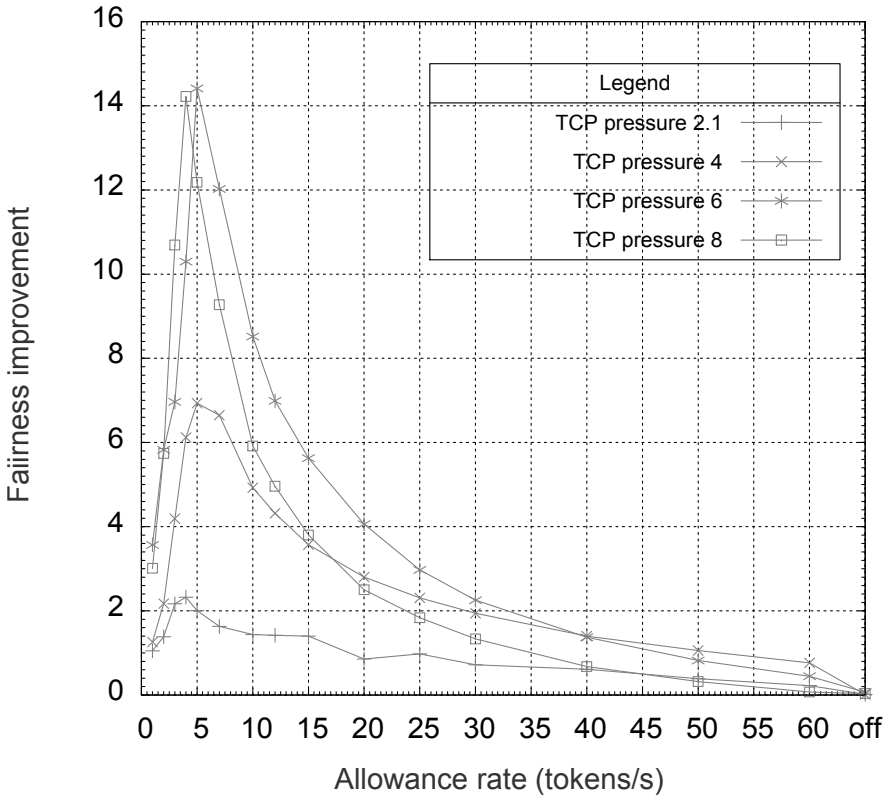

- Experiment setup
  - 10 Mb/s, conf. unfairness 20
  - Vary # light users: 10, 20, 40, 60
    $\Rightarrow$ $2.1, 4, 6, 8\ \frac{flows}{Mb/s}$
- Observation
  - Fairness improvement increases with increasing TCP pressure

Performance Analysis

- TCP New Reno provides imperfect ECN feedback
  - At most one ConEx signal per RTT
- Required allowance rate to avoid packet drops for single flow at policer
  - 1/2 token/RTT

$$RTT \approx 115\,ms \rightarrow \frac{\frac{1}{2}token}{RTT} \approx 4.3\,\frac{tokens}{s}$$

- Case analysis: allowance rate
  - < 1/2 token/RTT
    - Light and heavy users impeded
  - ≈ 1/2 token/RTT
    - Only heavy users impeded
  - > 1/2 token/RTT
    - Suppression of heavy users decreases

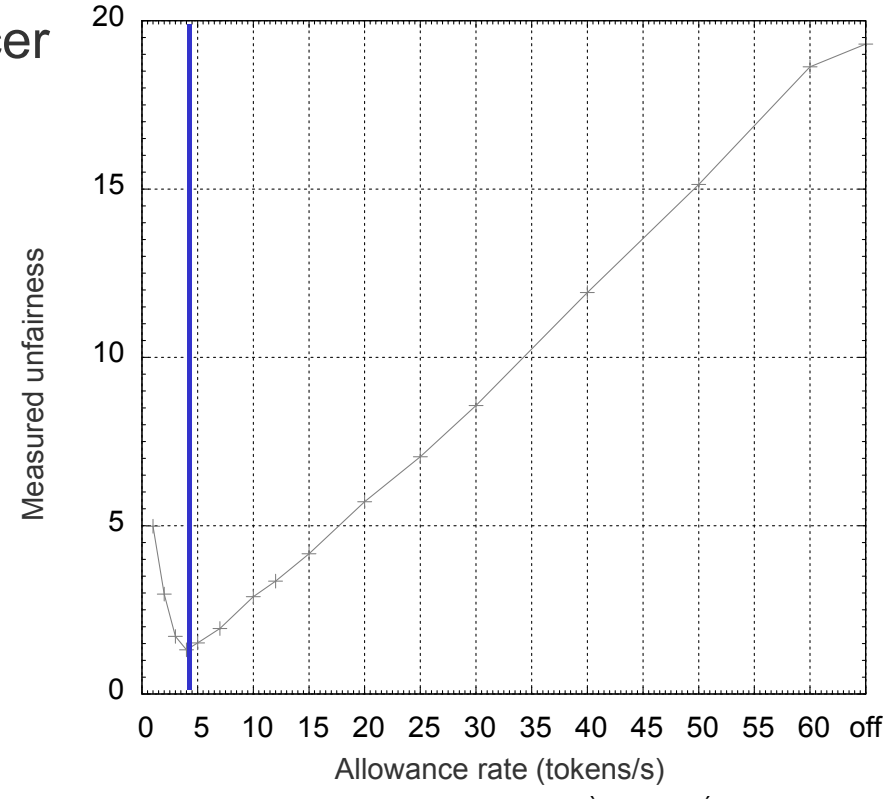

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

Validation: Vary Number of Flows per User

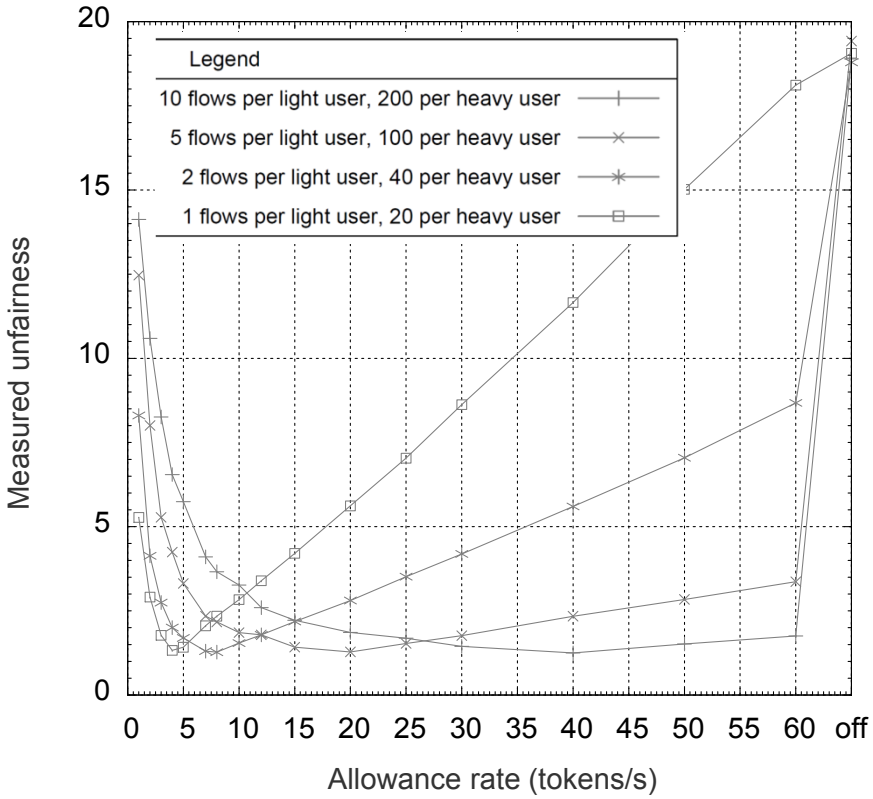

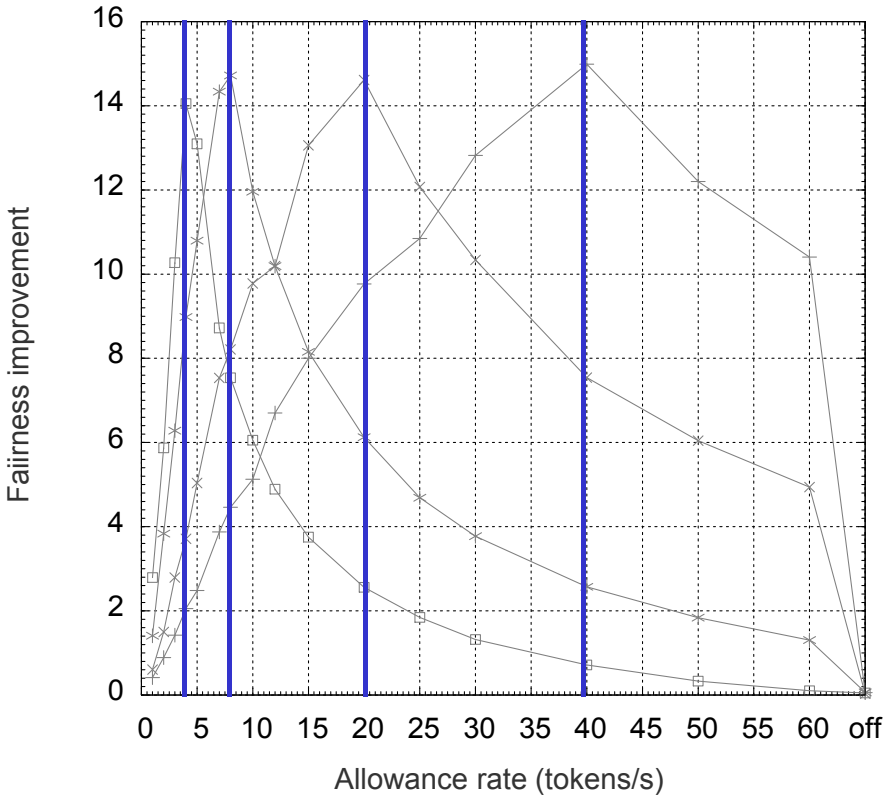

- Experiment setup
  - 100 Mb/s,
    conf. unfairness: 20, 800 flows $\Rightarrow 8\ \frac{flows}{Mb/s}$
  - Vary # flows and # users
  - L = # flows per light user: 1, 2, 5, 10
- Observation
  - Optimum allowance rate
    - Depends on # flows per light user
    - Coincides with $\frac{\frac{1}{2} \cdot L}{RTT}$

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

Validation: Vary RTT

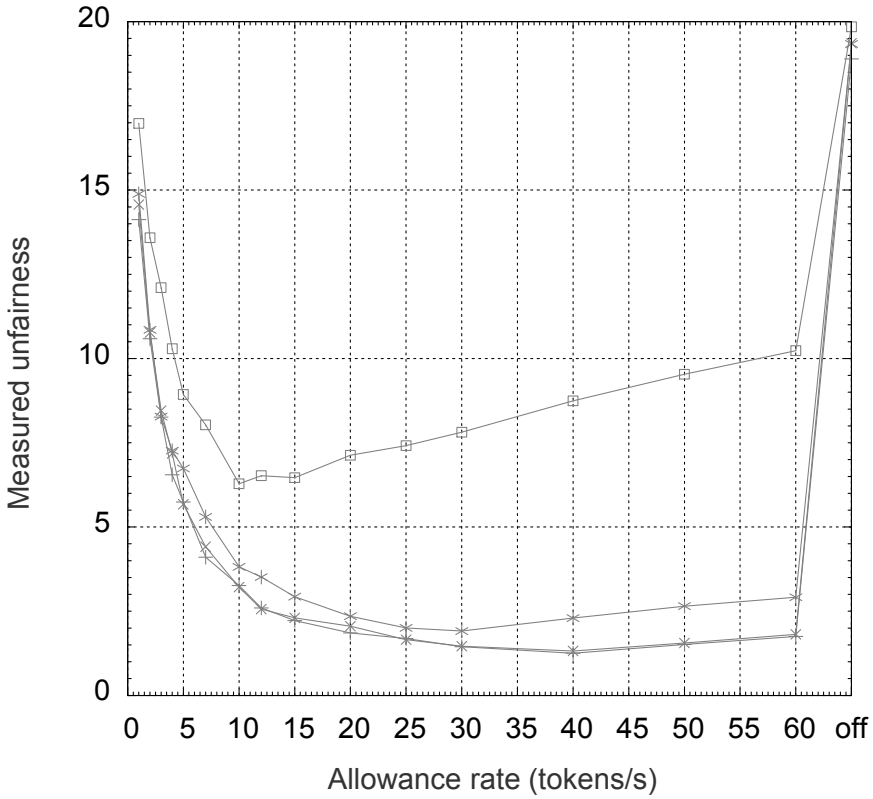

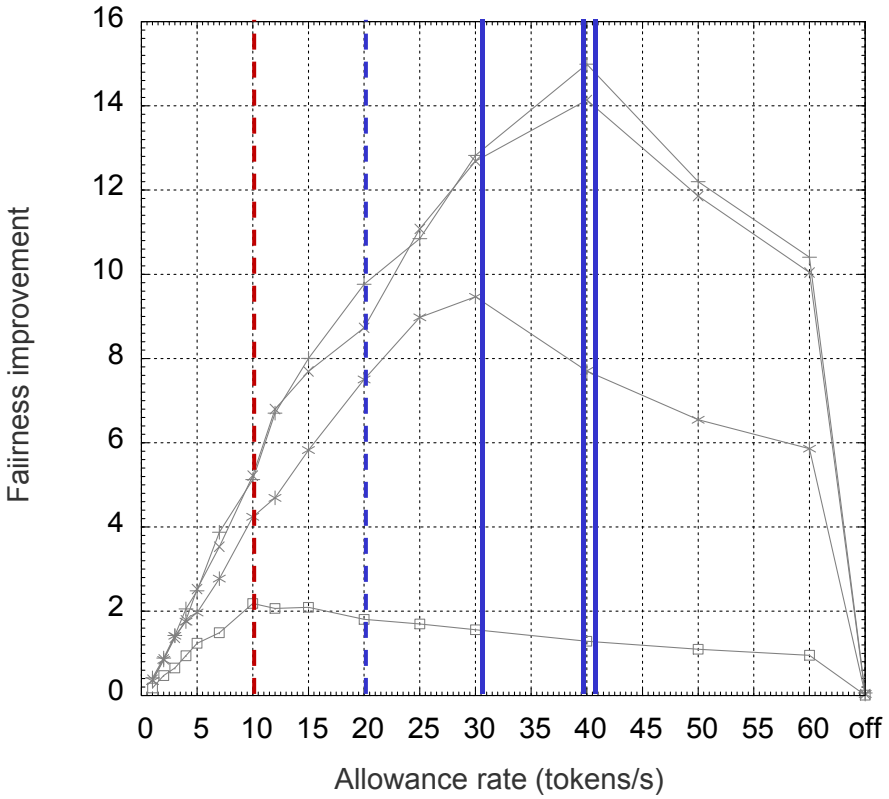

- Experiment setup
  - 100 Mb/s, 60(10)/1(200) users(flows)
    - Conf. unfairness: 20
    - 800 flows $\Rightarrow 8\ \frac{flows}{Mb/s}$
  - Vary prop. delay: 10, 20, 50, 100 ms
- Observation
  - Measured RTTs: 110, 130, 160, 240 ms
  - Opt. allowance rate coincides with $\frac{\frac{1}{2} \cdot L}{RTT}$
  - Outlier (240 ms): only 90% util

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

Summary & Conclusions

- ConEx-based congestion policing
  - Improves fairness significantly (bandwidth shares)
    - In the presence of congestion (high TCP pressure)
  - Does not impede heavy users
    - In the absence of congestion (low TCP pressure)
- Optimum allowance rate
  - Depends on # flows per user and RTT
  - Reason: imperfect ECN feedback of TCP NewReno
    - More impact on performance than policer variants
- More investigation needed
  - More accurate ECN feedback, TCP variants, different tolerances, AQM marking functions, non-saturated TCP flows (more sophisticated traffic generation), byte or packet counting, policer variants, congestion allowance variants, other transport protocols, coexistence of transport protocols, realistic use cases, auditors, …