AF & TCP

towards the specification of an IP+ Service

Octavio Medina

ENST Bretagne, Rennes

TF-NGN, Tromsø
June 2001

AF&TCP

- AF & TCP: what is the problem?
- Simulation Results
- Preliminary Conclusions
- Testing: where are we?

Problem Statement

- How can AF increase the QoS of TCP flows?
  - "Assured Bandwidth" (IP+ Service)
    - for flow aggregates
    - for individual sessions (microflows)

    even if the SLS is defined at the aggregate level, improvement of QoS should be seen by individual flows.
  - Enemies of AB
    - Perturbation due to non-responsive (UDP) flows
    - Heterogeneity of RTT

Simulations: TCP vs. UDP

- 32 flows: 8 UDPs, 24 TCPs

- Tested Different Marking strategies
  - Same parameters for both protocols
  - Gave TCP higher values

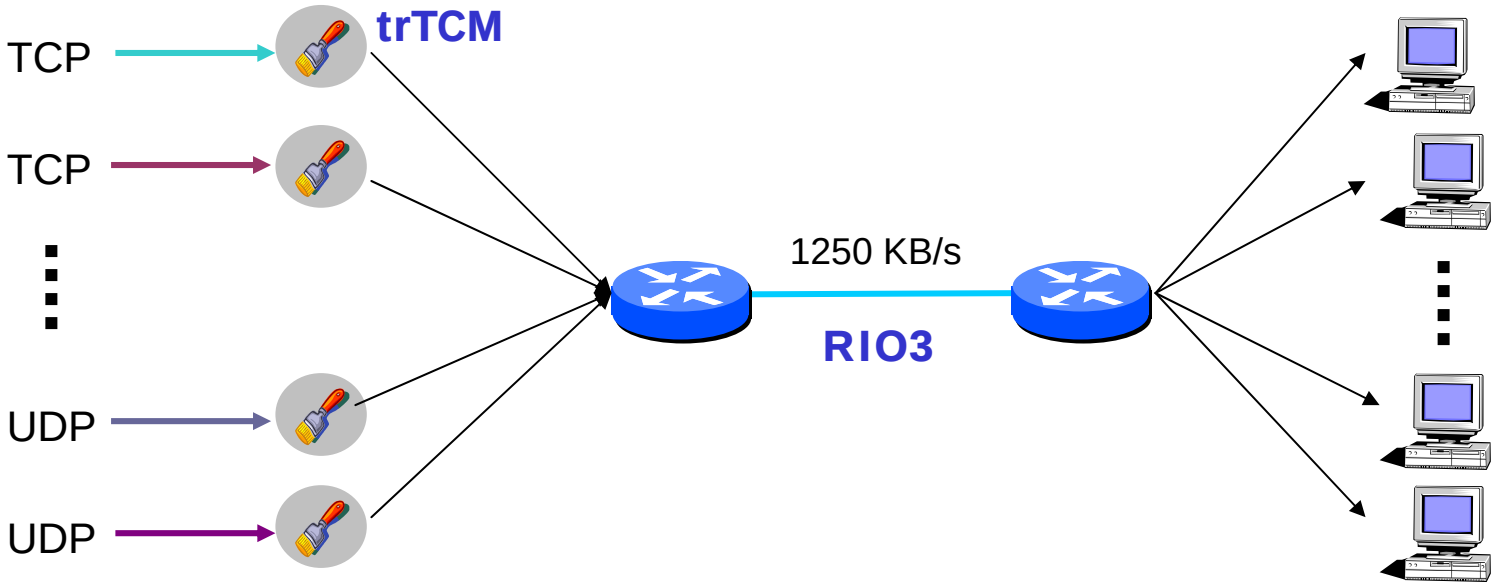

Simulations: TCP vs. UDP

- Fair sharing can be obtained through marking
- TCP flows needed higher TB values (35% in this case)

- We should be careful when accepting UDP in an AF class.

BE

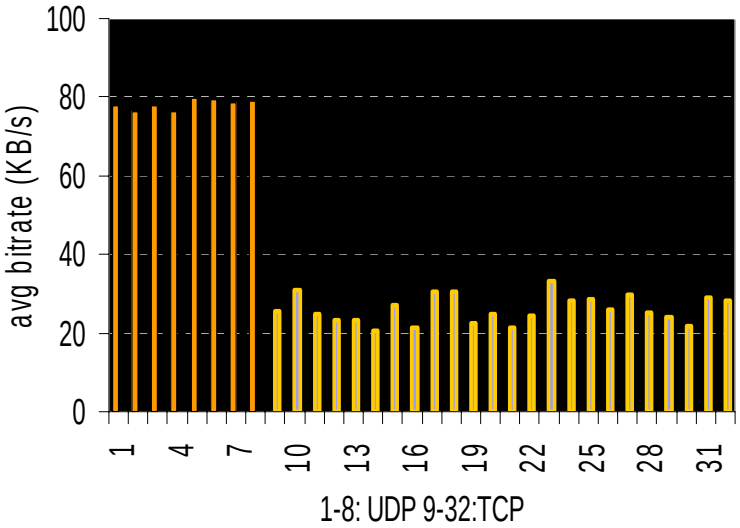

AF

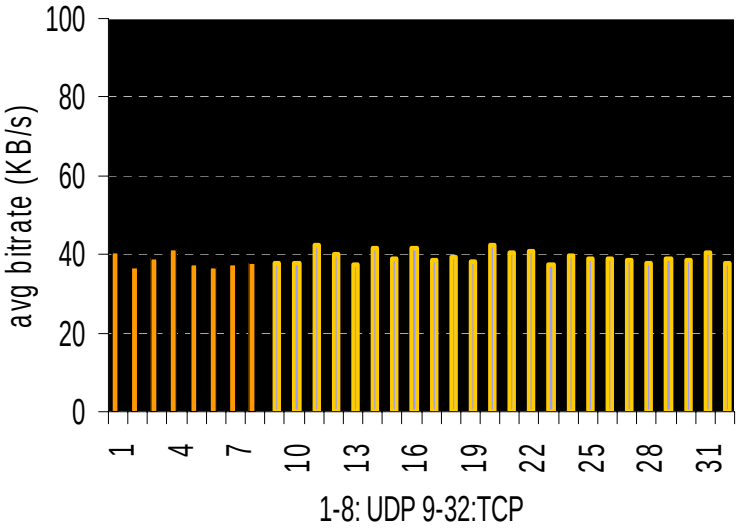

Simulations: effect of RTT

- 32 flows: RTT varies from 40 ms to 300 ms
- Individual flow marking (trTCM)
  - T_c(20 KB/s, 30 KB)
  - T_e(40 KB/s, 90 KB)

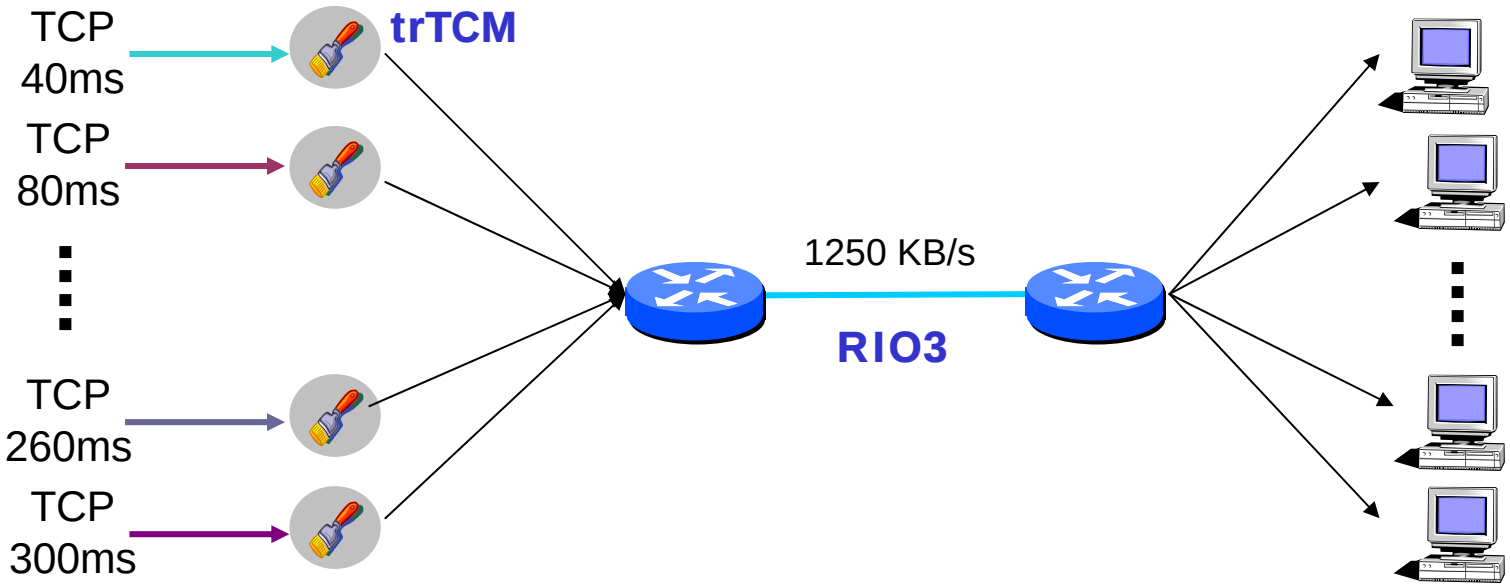

Simulations: effect of RTT

- Marking with the same parameters does NOT assure the same BW for all TCPs
- Throughput depends on RTT
  - With AF, link utilization is penalized (down to 68%)
  - Difference in throughput has not disappeared

BE

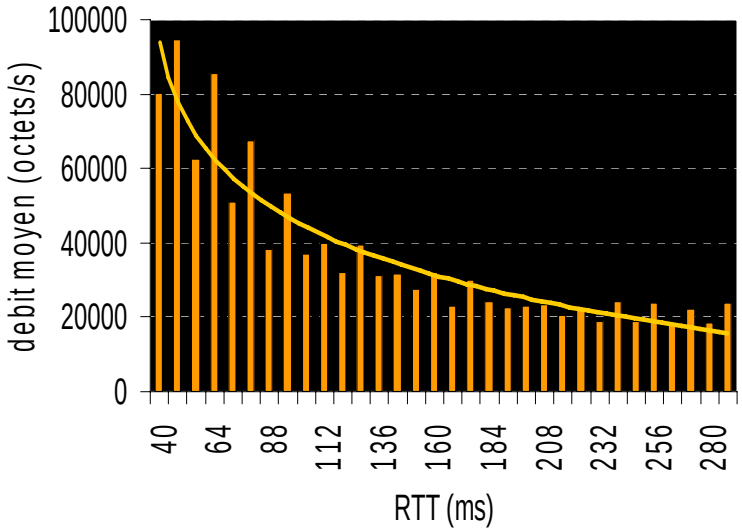

AF

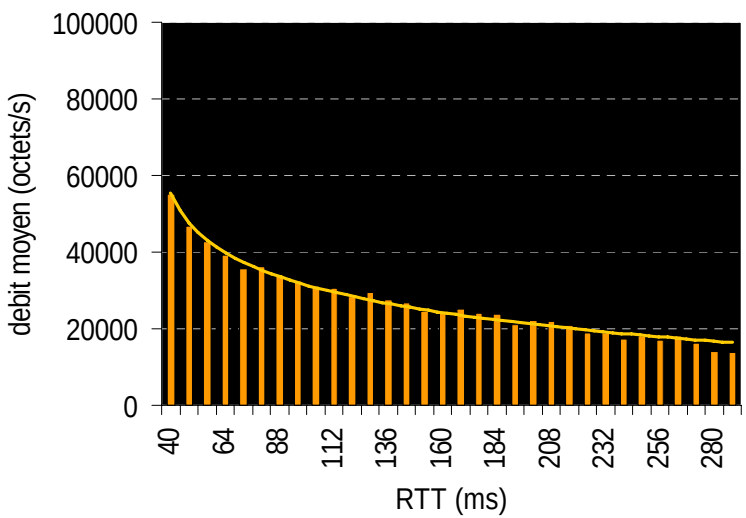

Simulations: effect of RTT

- Why this behavior?
  - Sessions with small RTTs accelerate faster
    - They consume green tokens and generate red ones before the others.
  - Small RTT allows sessions to react to losses faster
    - No big bursts of red packets
  - Other explanations?

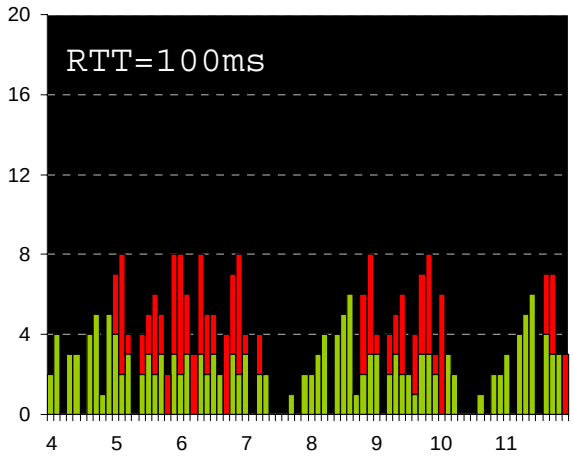

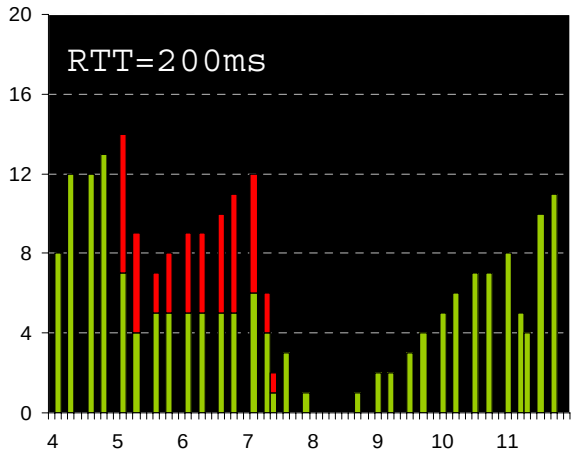

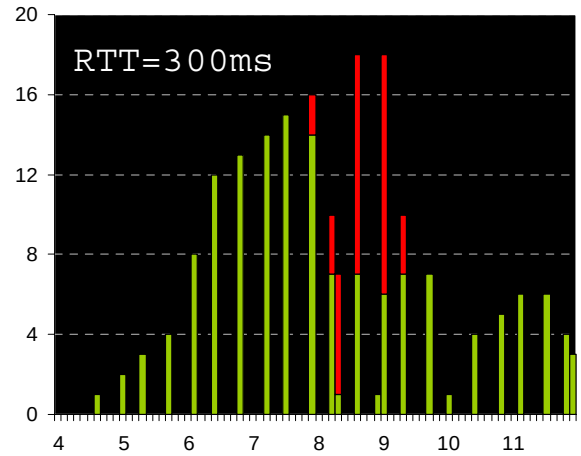

Simulations: aggregate scenario

- Microflows are aggregated
- They are (re)marked as an aggregate
- Looking for weighted BW sharing

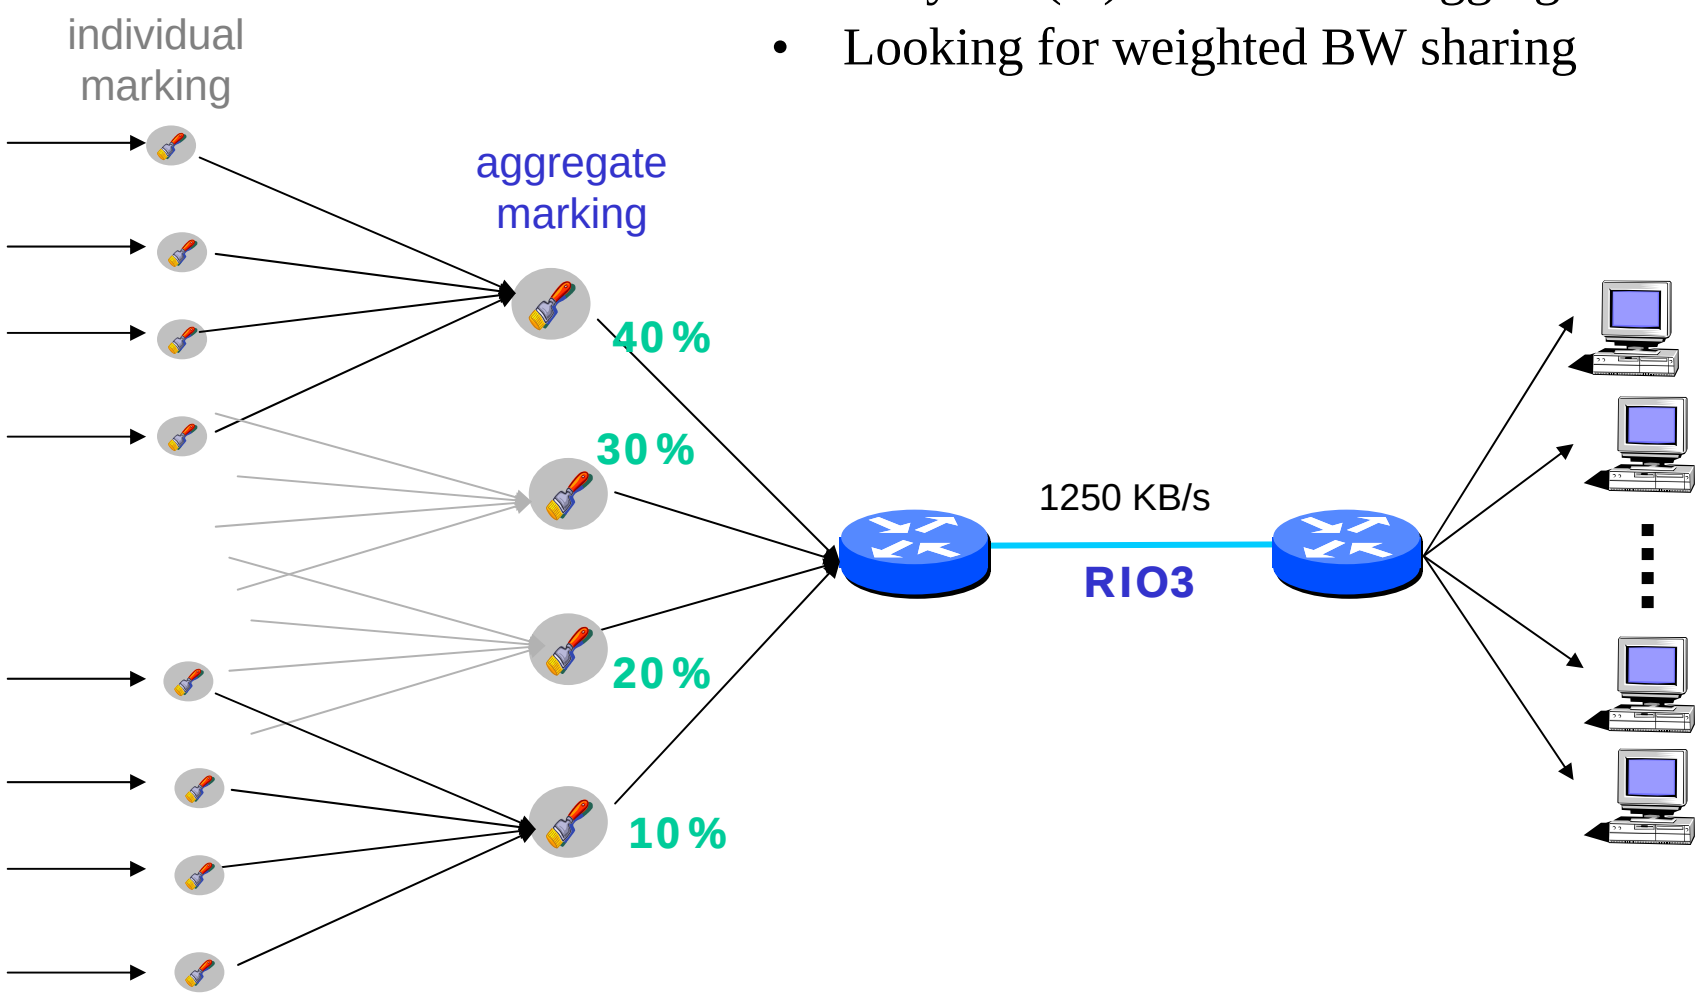

Simulations: aggregate scenario

- If marking is performed over aggregated flows only, results are as good as for UDP flows.
  - Incoming microflows are colorless.
- But…

AF TCP

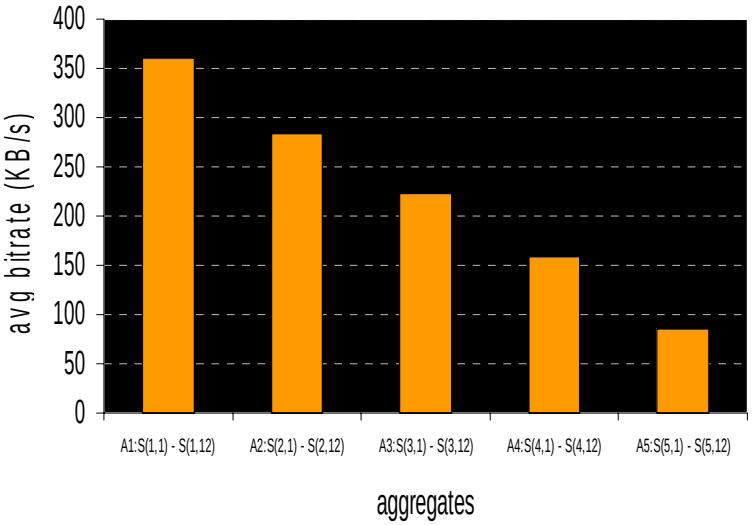

AF UDP

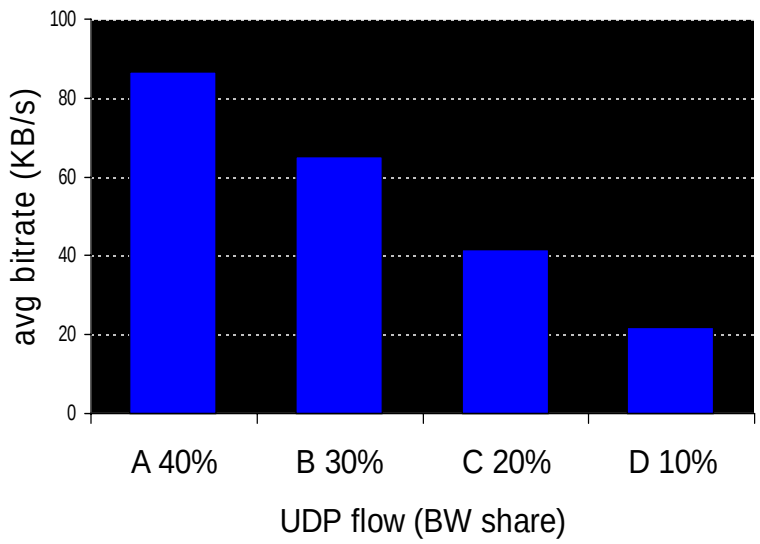

Simulations: aggregate scenario

- Distribution inside the aggregate
  - Very unfair
  - Flows with small RTTs get most of the tokens!

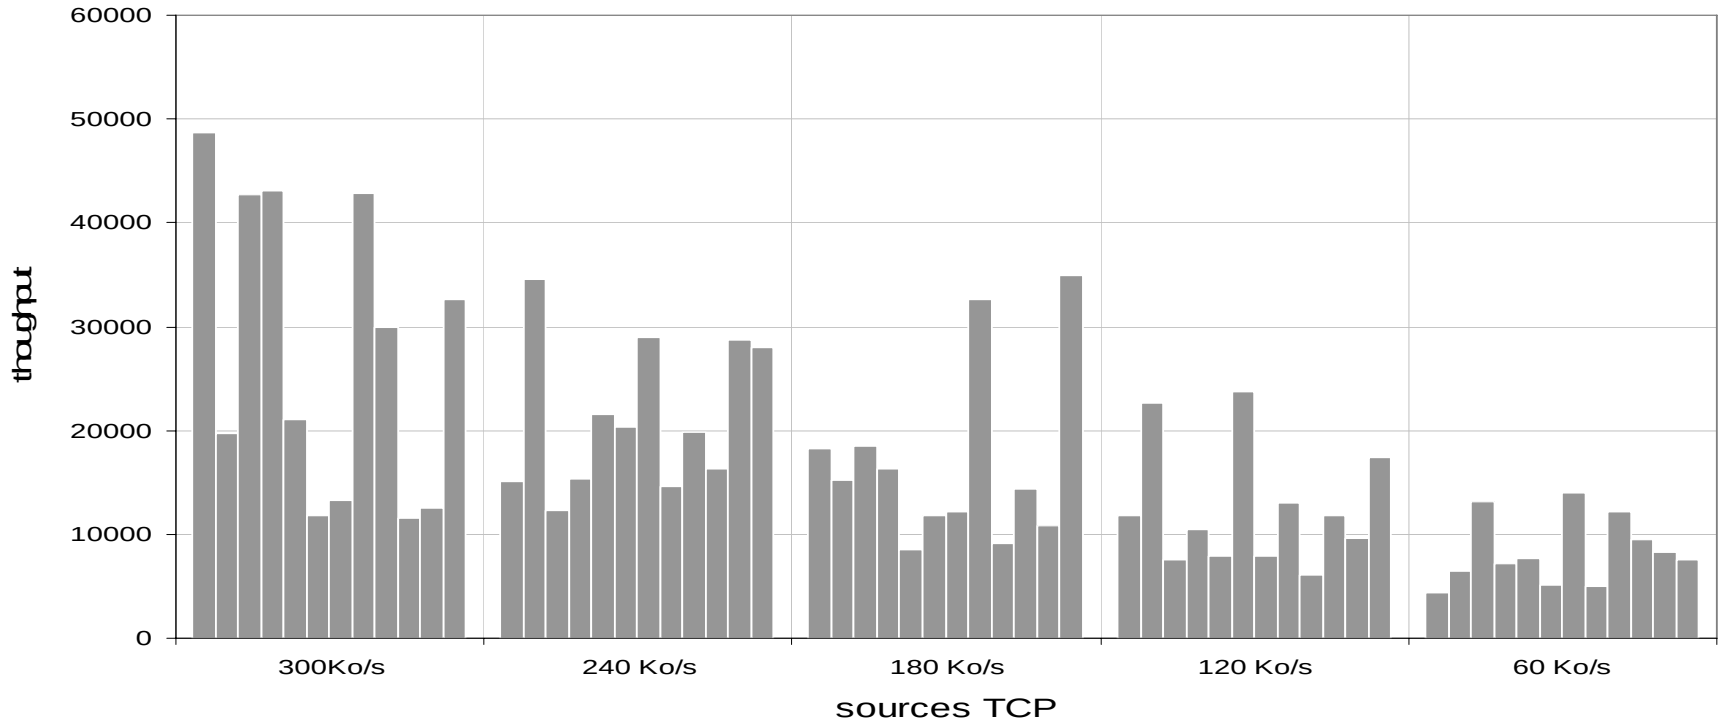

Simulations: aggregate scenario

- Possible solution: pre-mark individual flows
  - Second-level markers get more information about flow behavior

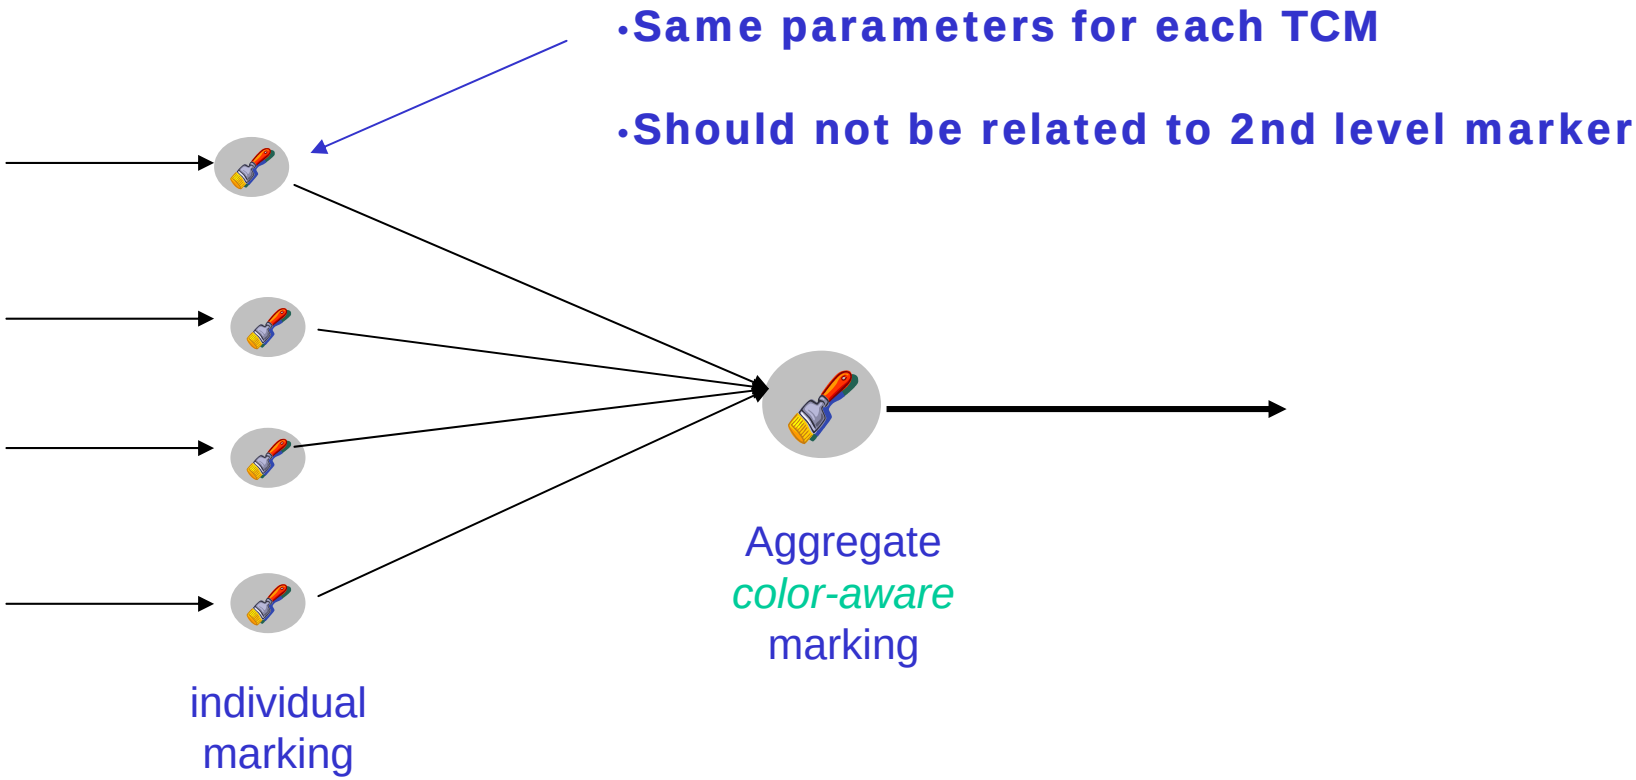

Simulations: aggregate scenario

- Same individual TCM for each of the 60 sessions (CIR=15Kb/s)
- Weighted Aggregate Marking as before
- Increases fairness inside aggregates!
- But…

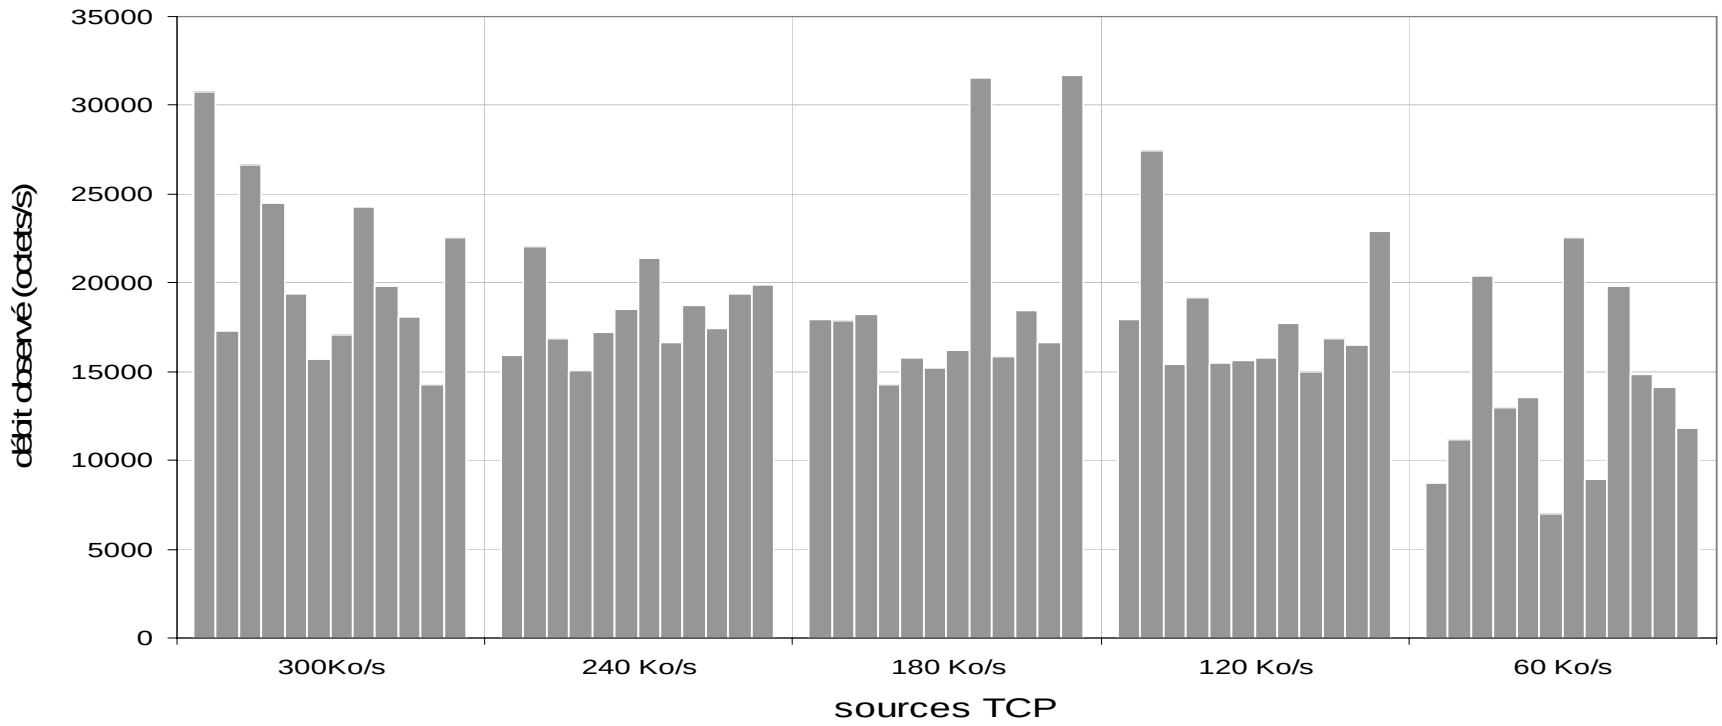

Simulations: aggregate scenario

- Lost Differentiation among aggregates
- What to do?
  - Shaping of TCP aggregates?
  - Try other marking methods?

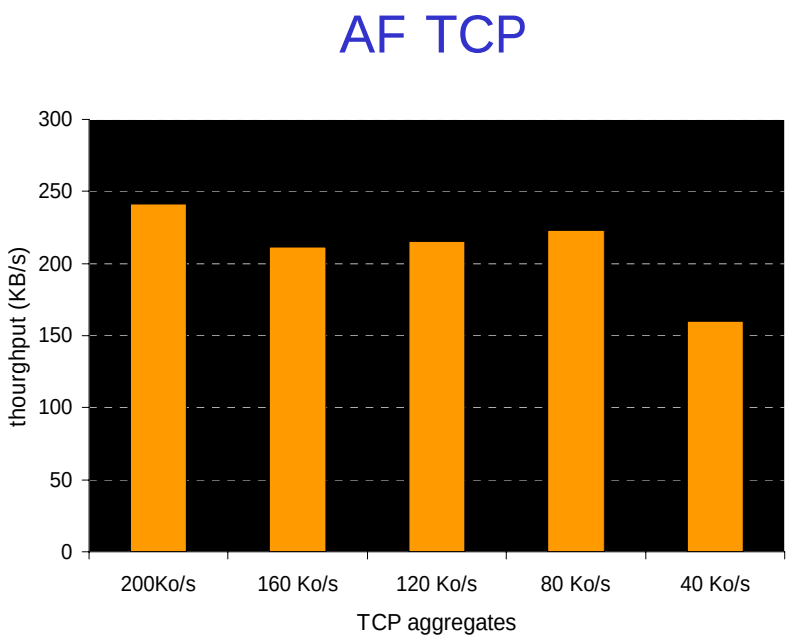

Preliminary Conclusions

- When defining an Assured BW Service,
  - One must be careful about accepting UDP traffic.
    - Conditioning at boundary routers must assure that marked TCP flows in the same class are not affected.
  - Specificities of RTT must be taken into account.
    - Assuring BW only at the aggregate level will be “attractive”enough?
    - Should the service define RTT boundaries?
    - Would QoS be only assured for flows inside a domain? *(where the RTT is known)*

Testing: Where are we?

- Calibration of Cisco's WRED parameters
  - Done
- Tested Juniper's AF capabilities
  - Done
- Tested BW sharing among UDP flows
  - Done
- Simulated BW sharing of individual TCP flows
  - Ready for testing (PlaGe)
- Simulated TCP BW sharing under aggregation & remarking
  - Ready for testing (PlaGe)

Work in Progress: ENSTB

- Consolidation of our FreeBSD/DiffServ implementation
  - New people working on the project
  - Visit of Kenjiro Cho (ALT-Q creator)

- Implementation of a one-way delay measurement tool:
  - Use AdServ to add timestamp to each packet
  - GPS synchronization

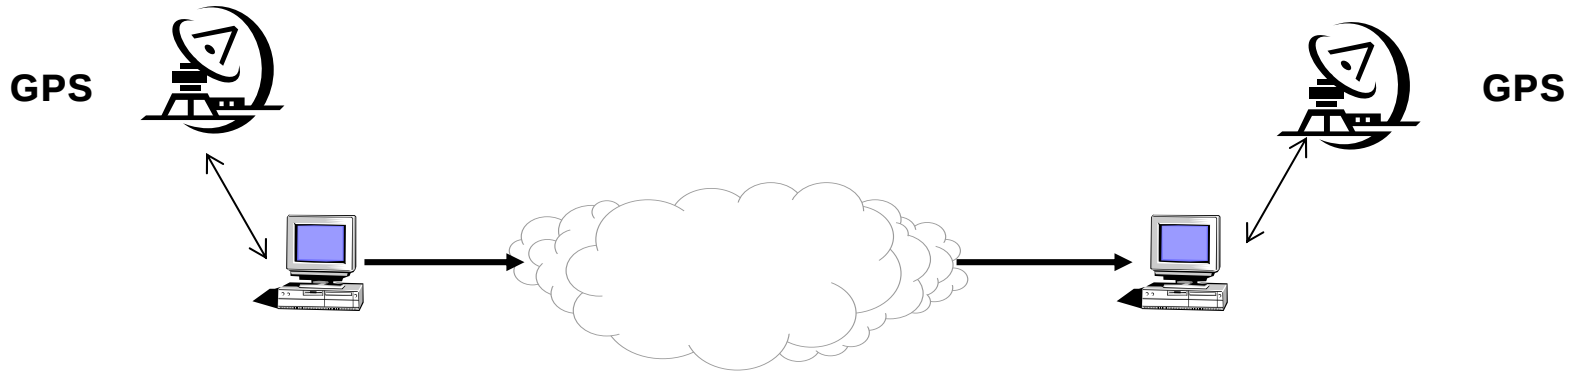

Work in Progress: PlaGe

- AF testing should take place soon
  - Special interest on aggregation/remarking tests
  - Need to create congestion in 2.5 GB links to see the benefits of AF
  - At high speeds, tests about overhead introduced by AF should also be considered

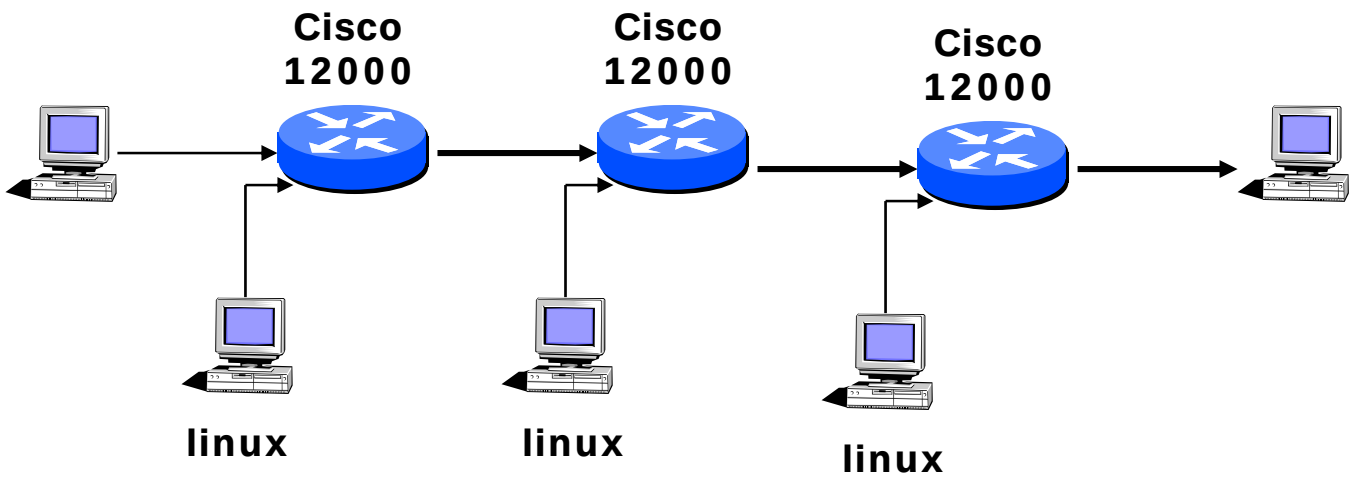

About the leadership of this activity…

- Tijani Chaled (INT Evry) has expressed his interest to continue this work.

- ENSTB/IRISA will continue participating to experimentations.
  - TF-NGN work to be one of the main activities of our "new" PhD student. (Julio Orozco)