Two bits are enough

Ihsan Ayyub Qazi* [ihsan@cs.pitt.edu], Lachlan L. H. Andrew# [lachlan@caltech.edu] and Taieb Znati* [znati@cs.pitt.edu]

*University of Pittsburgh, Pittsburgh, PA 15260, USA

#California Institute of Technology, Pasadena, CA 91125, USA

1. Motivation

- Many transport protocols such as XCP, RCP, MLCP, MaxNet etc require more bits for feedback than are available in the IP header for Explicit Congestion Notification (ECN)
- Changing the IP header requires a non-trivial and a time-consuming standardization process
- We design a load factor based congestion control protocol that uses *Adaptive Deterministic Packet Marking* (ADPM) to obtain congestion estimates with up to 16-bit resolution using the existing ECN bits
- ***Our scheme reduces the Average Flow Completion Time (AFCT) by up to 73% over VCP, up to 62% over TCP SACK+RED/ECN and up to 26% over RCP***

2. ADPM (Basic Idea)

- Interpret the 16-bit IPid field in the IP header as a number, i, in [0,1]
- IPid is generated either uniformly at random or sequentially
- Router marks a packet if "***Is c (the link price) > i ?***" and leaves the mark unchanged otherwise
- Receiver estimates the price at the bottleneck and sends it to the sender
- Sources adjust their rate according to the feedback

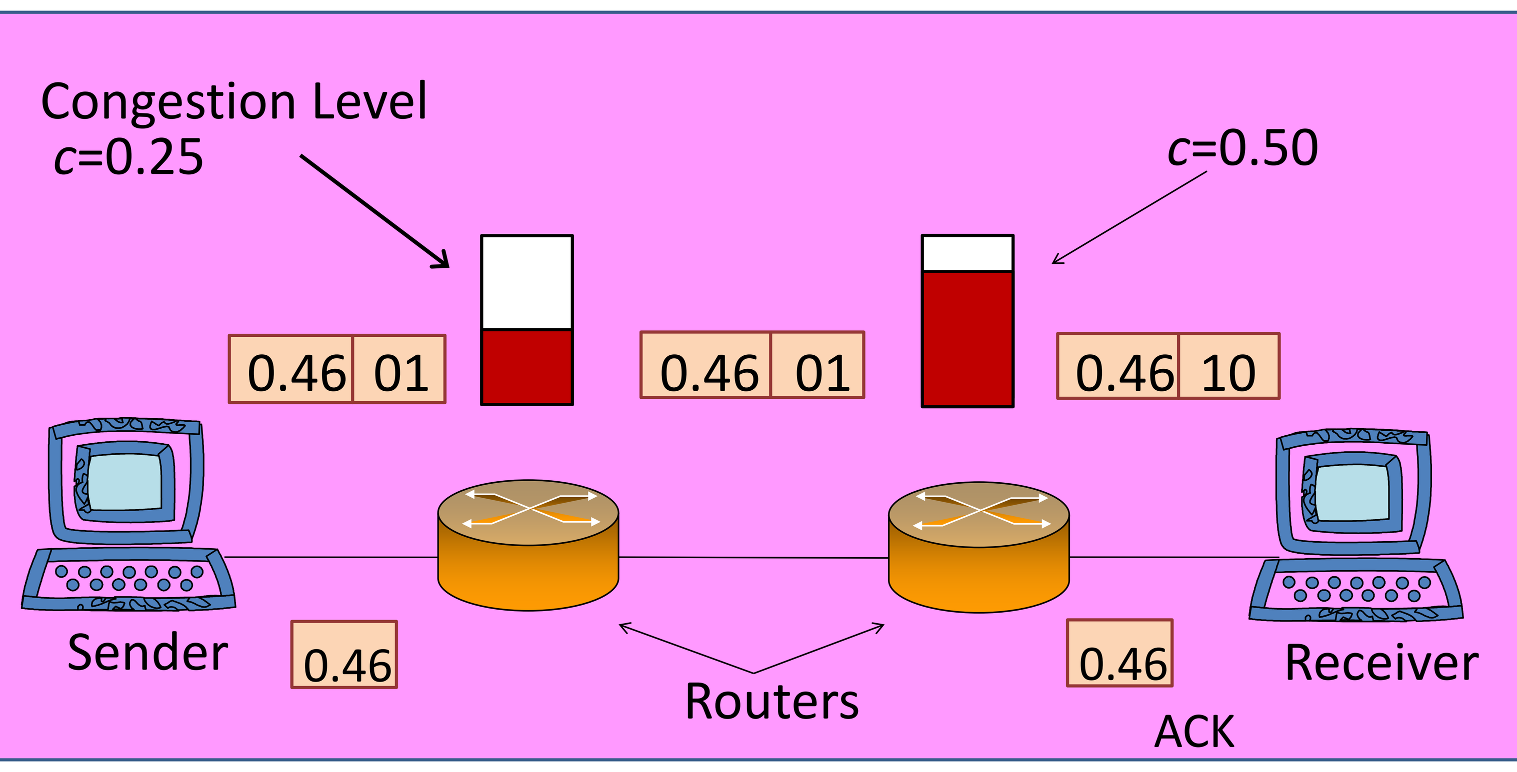

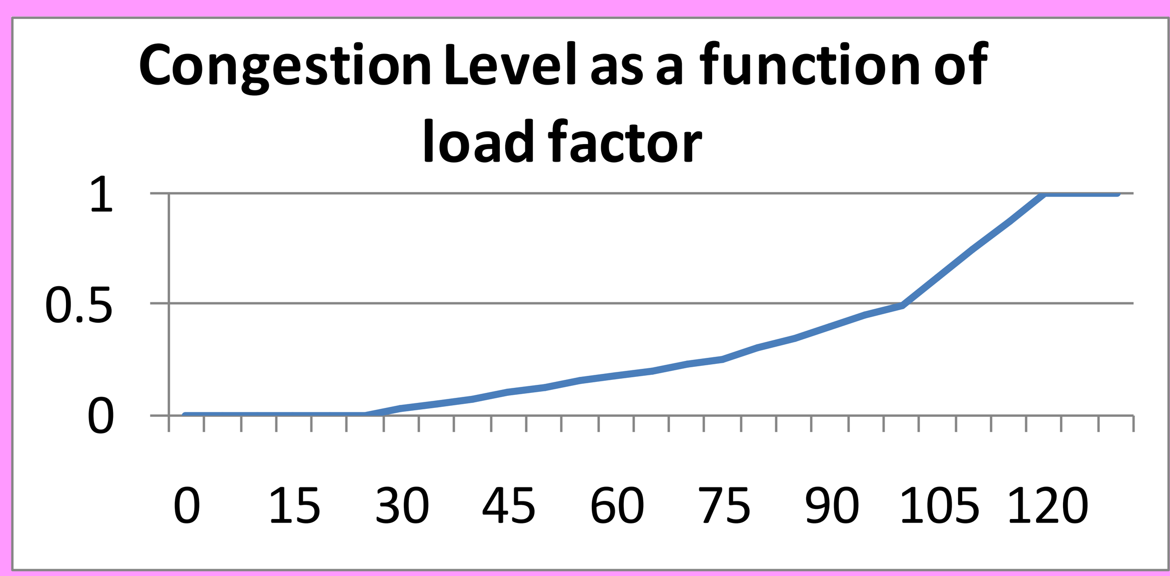

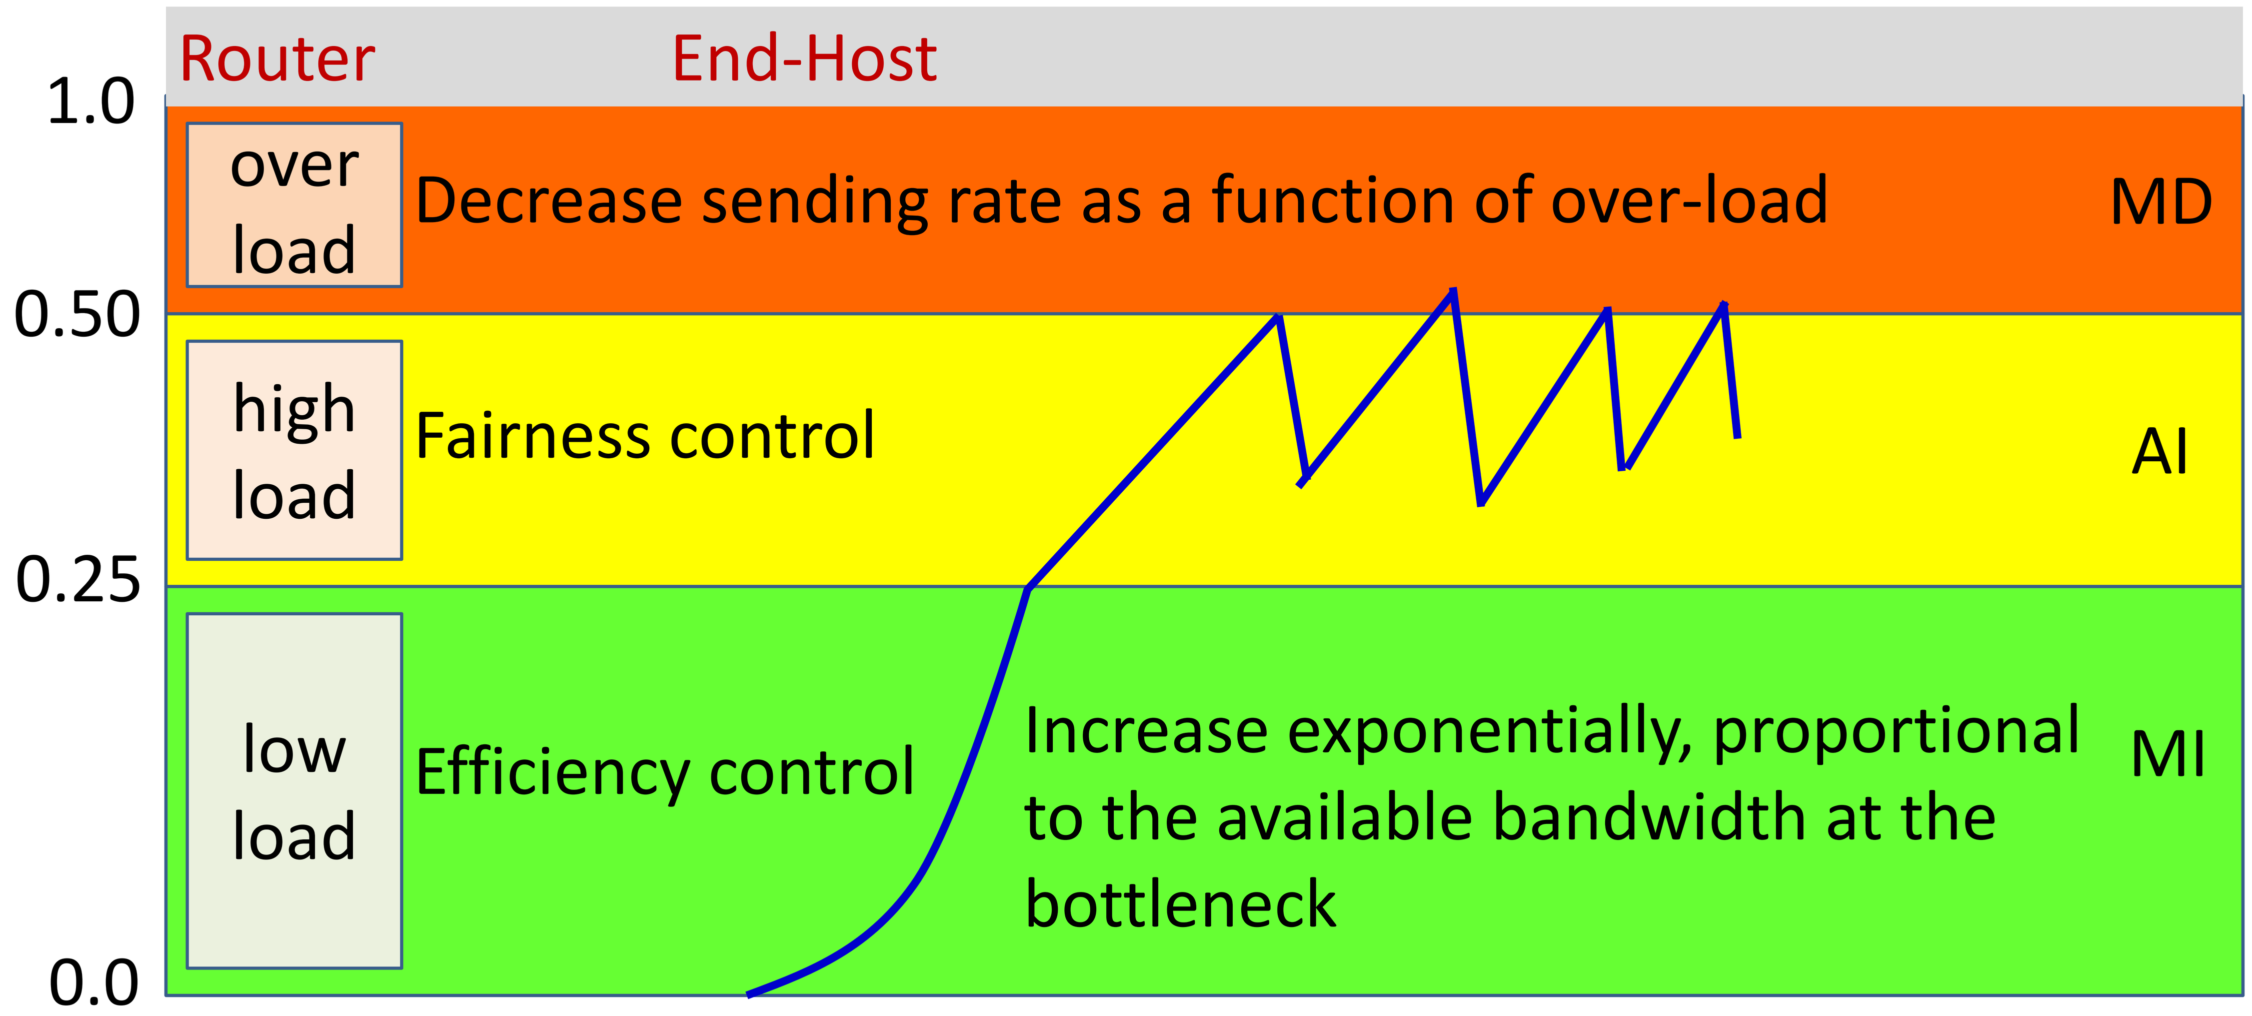

3. Design Considerations

- MI, AI, MD parameter values depend on the actual load at the bottleneck
- Overload may not be detected instantaneously
- However, if overload greater than a threshold, routers send (11) symbol and sources back-off deterministically
- Each starting source assumes that the actual load is 15%

4. Initial Results and Insights

- Average overload determines how aggressively sources decrease their sending rate upon congestion detection
- Our Markov Chain model shows that sources detect overload in the first round after congestion with high probability
- Probability of detecting overload remains roughly invariant to the BDP of the path

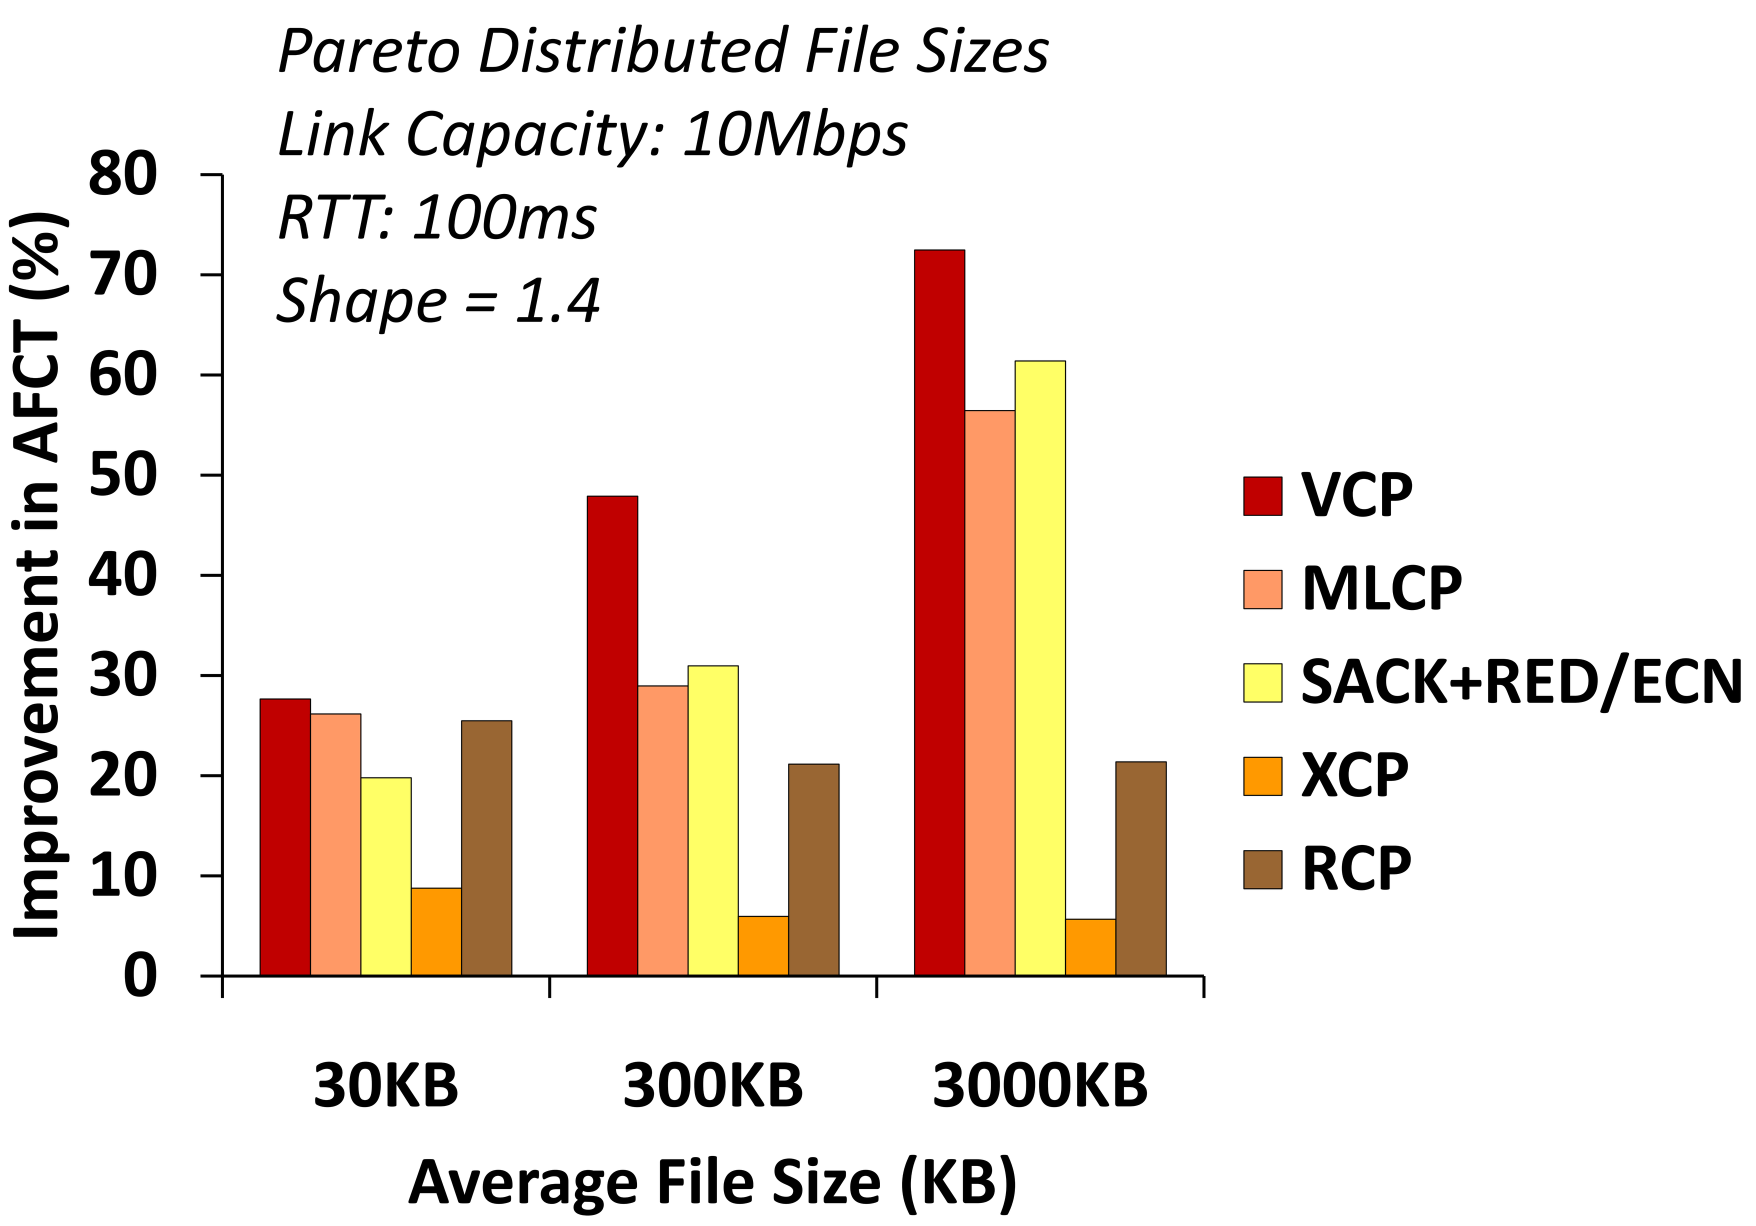

5. Contributions and Conclusion

- It is feasible to use the existing ECN bits to convey high resolution congestion estimates without sacrificing performance due to estimation errors
- This scheme closely approximates an optimal load factor based scheme in terms of convergence to efficiency
- We develop analytic models and conduct extensive ns2 simulations to characterize the performance of our scheme. Our analysis provides novel insights into the design of load factor based congestion control protocols that are likely to lead to better designs for next-generation congestion control protocols