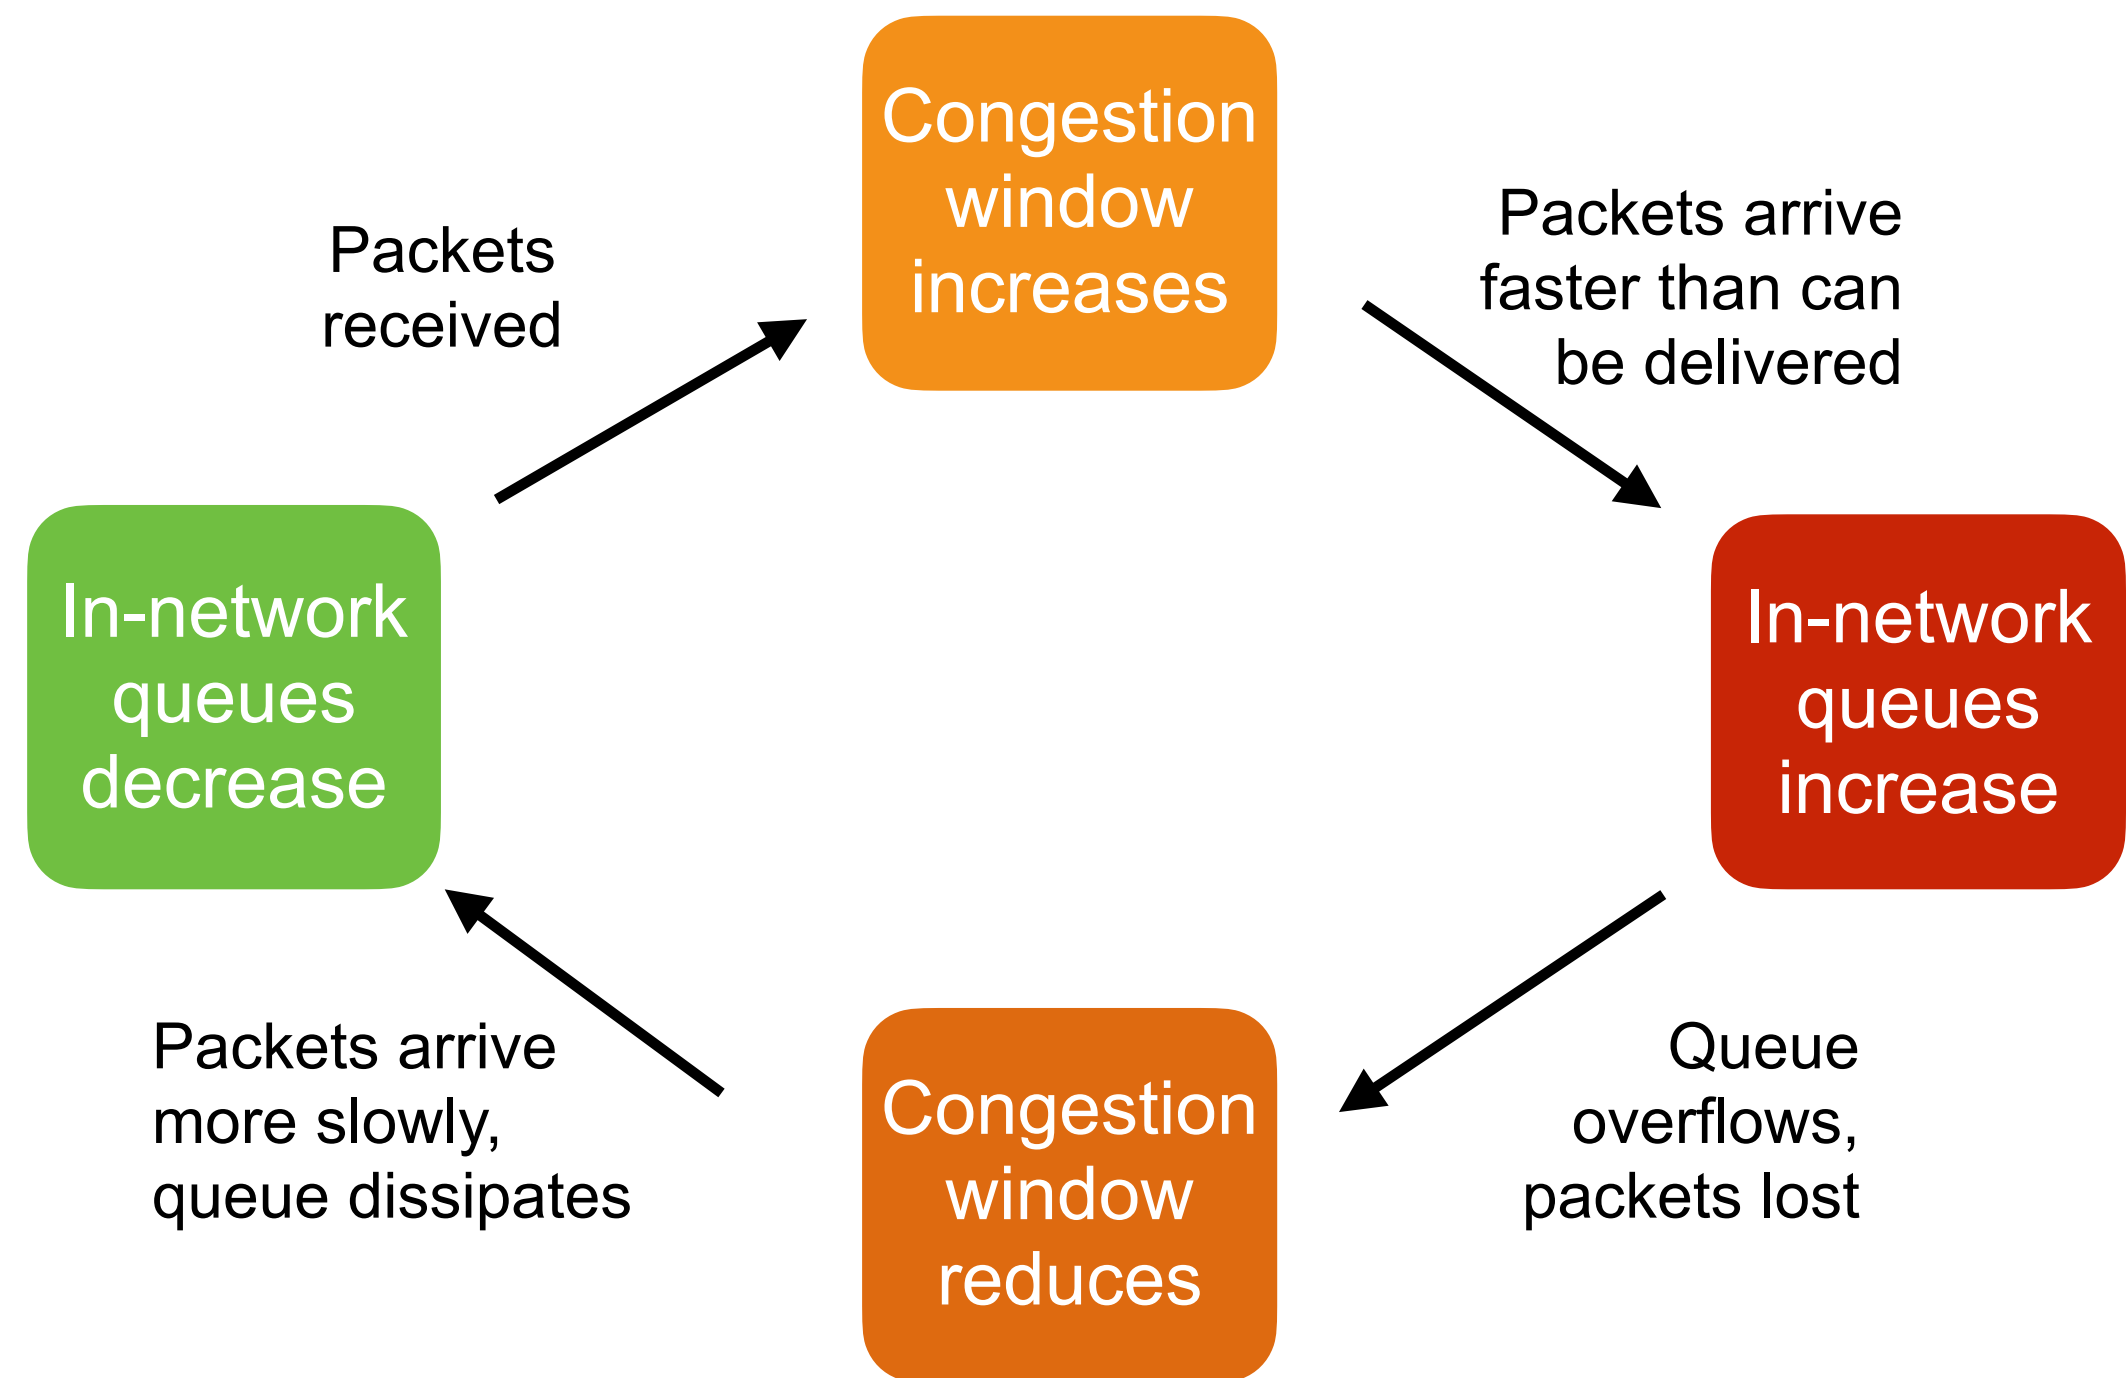


Delay-based Congestion Control

- Traditional TCP causes latency
- TCP Vegas
- TCP BBR

Impact of TCP on Latency

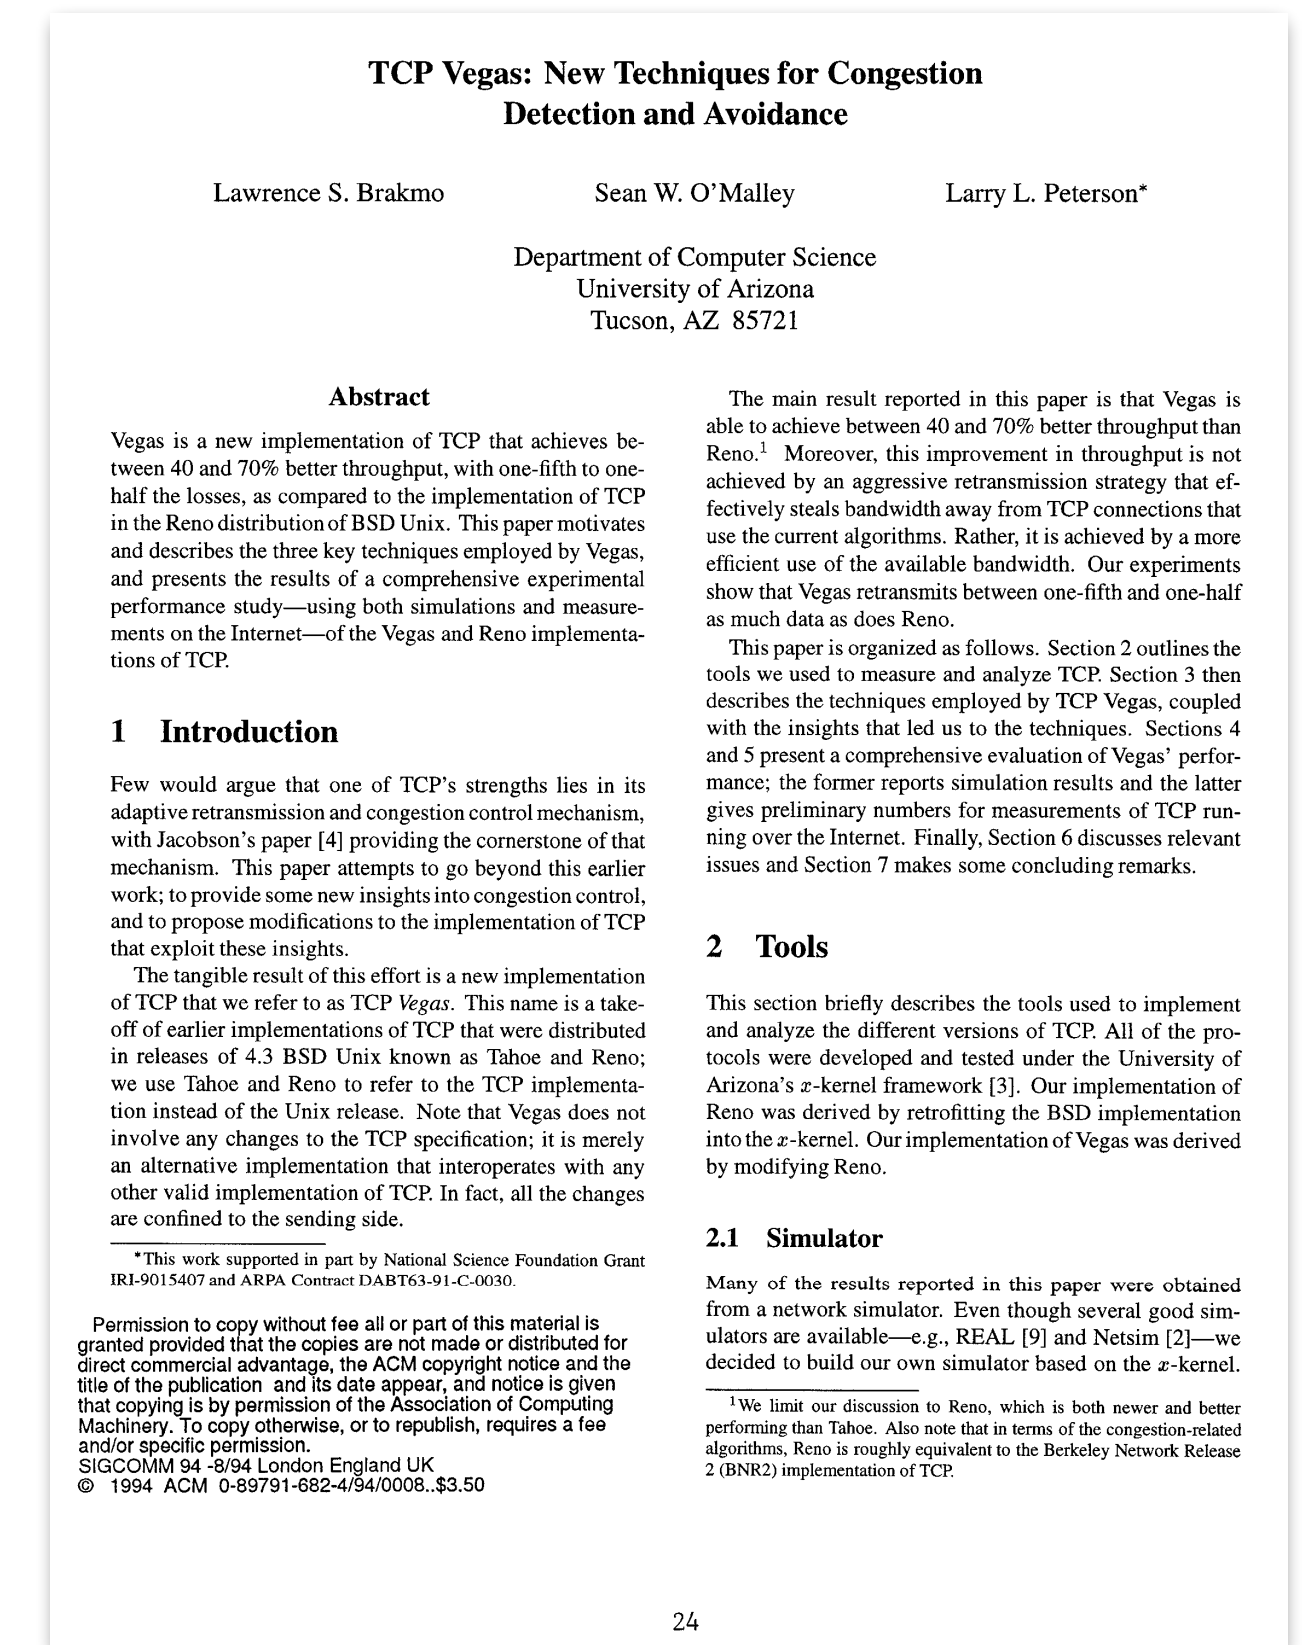
- TCP Reno and TCP Cubic **both** aim to fill the network



- TCP Reno and Cubic use packet loss as congestion signal → increase congestion window until queue overflows and packet is lost
- No matter how much big the queue, TCP Reno or Cubic **will** cause it to overflow, if given enough data to send
- Packets waiting in queues within network adds latency → increase RTT

TCP Vegas: Reducing Latency

- Key insight: if sending faster than the network can deliver, packets will be queued
- TCP Reno and Cubic wait until the queue overflows and packets are lost before slowing down
- **TCP Vegas** watches for the increase in delay as the queue starts to fill up → slows down before the queue overflows
 - Queues are smaller → lower latency
 - Packets are not lost, so don't need retransmission
- Only affects congestion avoidance; slow start unchanged




L. Brakmo, S. O'Malley, and L. Peterson, "TCP Vegas: New techniques for congestion detection and avoidance", ACM SIGCOMM Conference, London, August 1994. <https://dx.doi.org/10.1145/190314.190317>

TCP Vegas Congestion Control (1/2)

- Measure **BaseRTT**
 - The smallest time between sending a packet and getting its acknowledgement
 - BaseRTT will be from a packet delivered without much queueing
- Calculate **ExpectedRate** = $W/\text{BaseRTT}$
 - The congestion window, W , determines how many packets are sent each RTT
 - If the network can support this sending rate, the complete window should be delivered within the RTT
- Measure **ActualRate**
 - Bytes sent divided by actual RTT, for each packet
 - Measure how fast packets are actually received, based on acknowledgements returned
 - $\text{ActualRate} \leq \text{ExpectedRate}$ since packets can't be delivered faster than they're sent

TCP Vegas Congestion Control (2/2)

- Compare ExpectedRate with ActualRate and adjust window:
 - If $\text{ExpectedRate} - \text{ActualRate} < R_1$ then additive increase to window
 - Data arriving at close to expected rate, can likely send faster
 - If $\text{ExpectedRate} - \text{ActualRate} > R_2$ then additive decrease to window
 - Data arriving slower than it's being sent, slow down
- Parameters R_1 and R_2 , where $R_1 < R_2$, critical to performance



Additive increase and decrease in size window → smoother changes in sending rate than Reno or Cubic

TCP Vegas: Limitations

- Delay-based congestion control is a good idea in principle
 - Reduces latency
 - Reduces packet loss
- But loss- and delay-based congestion control don't cooperate
 - TCP Reno and TCP Cubic aggressively increase queue sizes
 - Increases RTT and reduces ActualRate as measured by TCP Vegas
 - Forces TCP Vegas to slow down; cycle repeats → TCP Vegas sending rate drops to zero over time
- TCP Vegas is not used, because it can't be deployed alongside TCP Reno or TCP Cubic

TCP BBR

- Is it possible to develop a delay-based congestion control algorithm for TCP that can be deployed on a network alongside TCP Reno and TCP Cubic?
- **Maybe** – TCP BBR (“Bottleneck Bandwidth & RTT”) is one attempt
 - Proposal from Google – measures RTT and bandwidth of the bottleneck link, directly sets congestion window
 - Used by YouTube in some cases, but **highly experimental**
 - Recent work by Ranysha Ware and Justine Sherry shows serious fairness problems with TCP BBR v1
 - <https://dl.acm.org/doi/10.1145/3355369.3355604>
 - <https://vimeo.com/showcase/6531379/video/369121357#t=990s>
 - TCP BBR v2 in development, might solve these problems – this is an active research area



N. Cardwell, Y. Cheng, C. S. Gunn, S. H. Yeganeh, and V. Jacobson. BBR: Congestion-based congestion control. ACM Queue, 14(5), September 2016. <https://dl.acm.org/doi/10.1145/3012426.3022184>

Delay-based Congestion Control

- Traditional TCP causes latency
- TCP Vegas
- TCP BBR