High Fidelity DoS Experimentation

Roman Chertov, Sonia Fahmy, Ness B. Shroff

Purdue University

{rchertov, fahmy}@cs.purdue.edu

shroff@ecn.purdue.edu

http://www.cs.purdue.edu/homes/fahmy/software/emist/

June 15th, 2006

Goal of this work

Simulators and emulators have operational ranges within which they are accurate; *however, exceeding the operational ranges (e.g., during DoS attacks) leads to artifacts that significantly impact experimental results and conclusions!*

- Our goal is to:
  - Understand fidelity of a simulator and three emulation testbeds by conducting experiments with TCP-targeted low rate DoS attacks.
  - Demonstrate the need for a *general router model* that can be used in simulators and emulators to increase the fidelity of results with DoS.

Real Router vs. Model

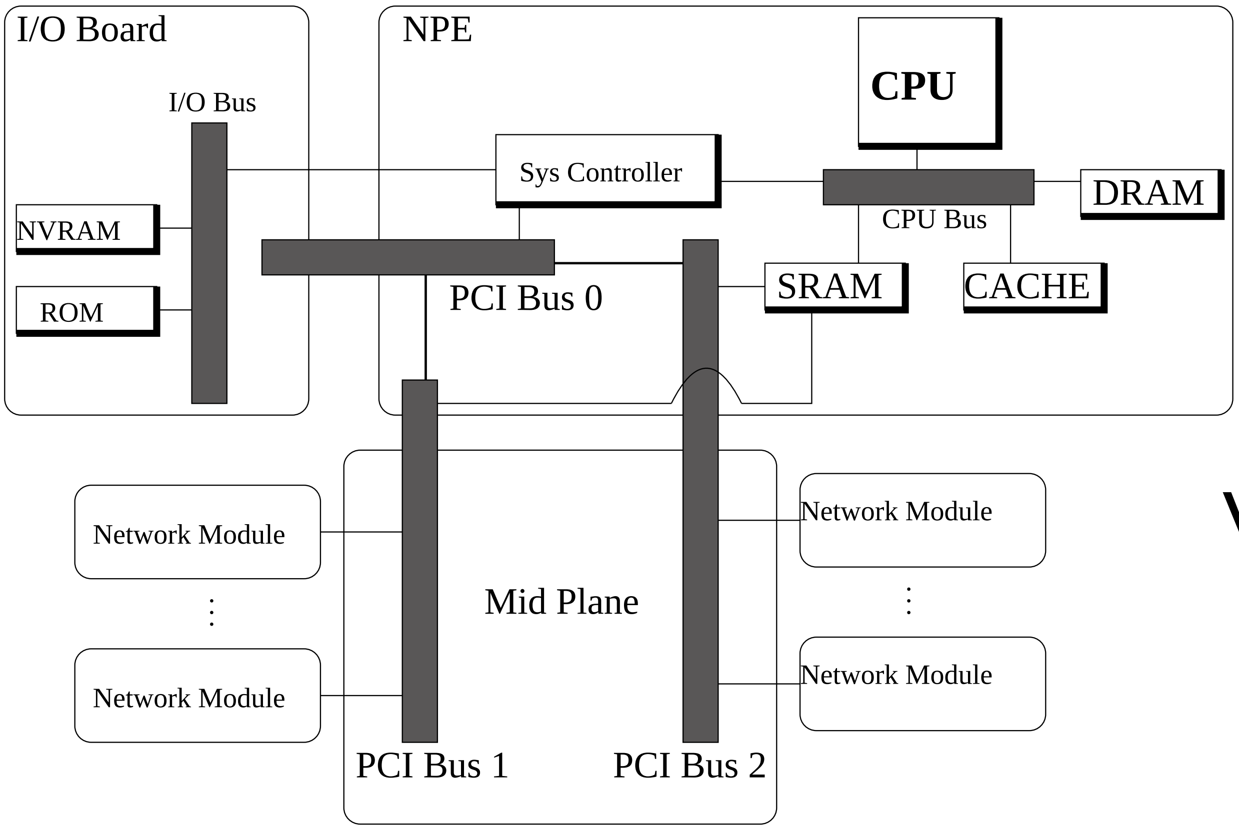

Layout of Cisco 7000 series routers

Vs.

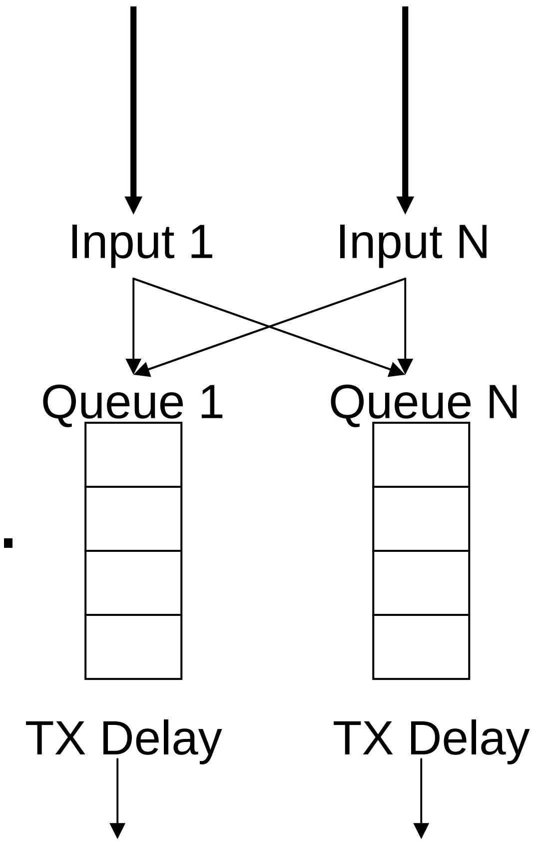

Router layout in ns-2 simulator

Related Work: Simulation

PURDUE UNIVERSITY

- Layers
  - *No layers* --- Packets are treated as messages: ns-2, pdns
  - *Realistic layers* from layer 2 and up: GTNeTS, OPNET, OMNeT++
- Device models
  - General and simple (e.g., serv_delay = pkt_size / BW): ns-2, pdns, GTNeTS
  - Custom models *per device*: OPNET and OMNeT++
- Protocol Software base
  - Custom implementation: ns-2, OPNET, OMNeT++, pdns, GTNeTS
  - Relies on production code: Network Simulation Cradle add-on for ns-2, NCTUns

- Sam Jansen, Network Simulation Cradle http://www.wand.net.nz/~stj2/nsc/
- S. Wang et al., The Design and Implementation of the NCTUns 1.0 Network Simulator, Computer Networks 2003

Related Work: Emulation

- Bridges simulation and real world by providing network “clouds” to which physical components connect.
- Can be used to shape links (DummyNet and Click) or emulate an entire network of links (ModelNet, EMPOWER, and VINT).
  - L. Rizzo, DummyNet, http://info.iet.unipi.it/~luigi/ip_dummynet/
  - E. Kohler et al., The Click Modular Router, ACM TOCS 2000
  - A. Vahdat et al., Scalability and Accuracy in a Large-Scale Network Emulator, OSDI 2002
  - P. Zheng and L. Ni, EMPOWER: a Network Emulator for Wireline and Wireless Networks, INFOCOM 2003
  - K. Fall, Network Emulation in the Vint/NS Simulator, ISCC 1999
  - F. Baumgartner et al., Virtual routers: a Tool for Emulating IP Routers, LCN 2002
- Nodes can be virtualized on a single PC: vBET, Emulab.
  - X. Jiang and D. Xu, vBET: a VM-Based Emulation Testbed, MoMeTools 2003
  - B. White et al., An Integrated Experimental Environment for Distributed Systems and Networks, OSDI 2002

Related Work: Device Measurement

- Basic network device profiling metrics such as: maximum throughput rate, packet loss, route setup, packet service time, and service recovery have been outlined in RFC 2544 and RFC 2889.
  - S. Bradner and J. McQuaid, Benchmarking Methodology for Network Interconnect Devices, RFC 2544, 1999
  - R. Mandeville and J. Perser, Benchmarking Methodology for LAN Switching Devices, RFC 2889, 2000
- Benchmarks in the above RFCs only deal with homogeneous traffic. Traffic representative of real networks induces different stresses.
  - J. Sommers and P. Barford, Self-Configuring Network Traffic Generation, SIGCOMM 2004
- Black box profiling has been done to measure OSPF calculations on Cisco routers.
  - A. Shaikh and A. Greenberg, Experience in Black-box OSPF Measurement, IMC 2001

Related Work: Summary

- Simulators and emulators can model a router device by using features as: variable delay, policies per packet, rate limiting, etc.
  - Most current tools do not do this and concentrate on general connectivity and output queuing models.

- Simulators like OPNET/OMNeT++ have device specific models
  - It is hard to manage a very large database of models
  - A small change in the router software can invalidate a previous model
  - Validation is hard
  - Complex models add large computational overhead

- Black box profiling
  - Has been done in limited settings but no attempts to create a general model.
  - No policy derivation methods

TCP-Targeted Attacks

- A. Kuzmanovic and E. W. Knightly. Low-rate Targeted Denial of Service Attacks. SIGCOMM 2003.
- Why?
  - Easy to *launch, stealthy,* and potentially damaging attack
  - Studied *only* via simulation, analysis, and limited experiments
  - Tricky as it strongly relies on timing
- Vary: Attacker, burst length *l,* sleep period *T, attack packet size/rate, Round Trip Time (RTT), router buffer sizes*
- Objective:
  - Understand attack effectiveness (damage versus effort)
  - *Qualitatively* compare emulation to simulation to analysis

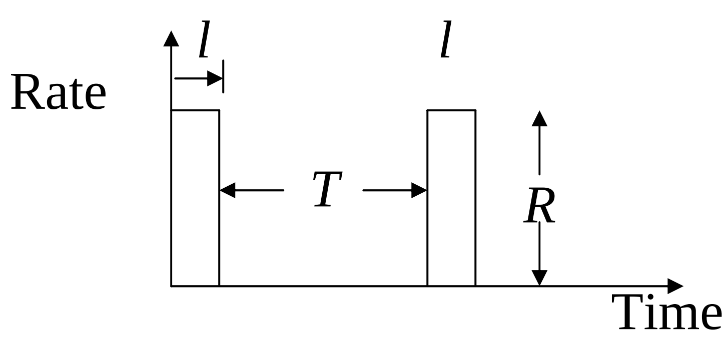

Experimental Scenario

PURDUE UNIVERSITY

- Original TCP-targeted attacks are tuned to Retransmission Time Out (RTO) frequency for near zero throughput
- Can exploit Additive Increase Multiplicative Decrease congestion avoidance of TCP *without* tuning period to RTO, and hence throttle TCP's throughput at any predetermined level
  - M. Guirguis et al. Exploiting the Transients of Adaptation for RoQ Attacks on Internet Resources. ICNP 2004.
- Simple dumbbell topology with single file transfer flow is easiest to interpret and is the worst case (most demanding for attacker)

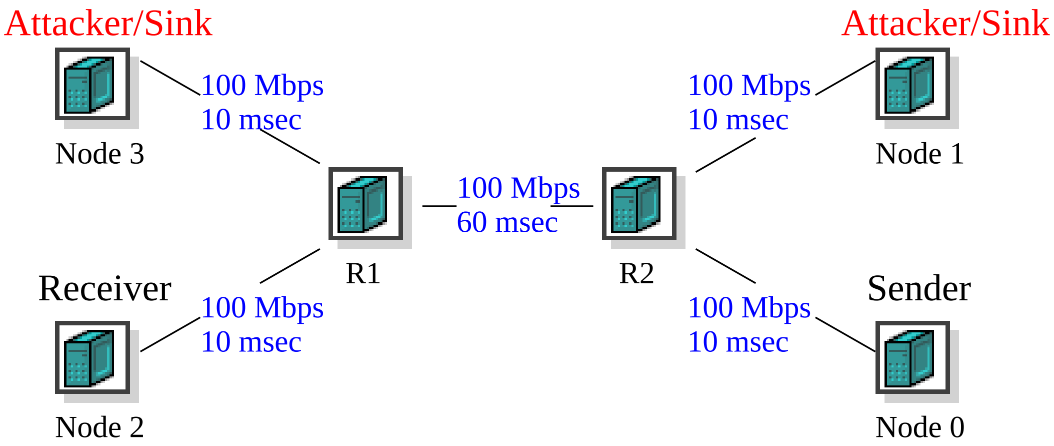

Experimental Setup

- All nodes run a zombie process that connects to the master, thus forming our *Scriptable Event System*
- A file transfer and TCP-targeted attack are initiated
- The same topology with similar events is simulated in ns-2
- Besides using default OS routing, routing nodes on DETER were configured with the *Click* modular software router
- Data from DETER, Emulab, WAIL, and ns-2 is compared to a simple throughput degradation model

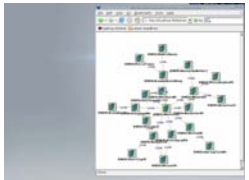

Throughput Degradation Model

Assumptions:

- Loss occurs during each pulse.
- Connection does not RTO.
- There is no packet loss during attack sleep periods.

$$W_{i+1} = \frac{W_i}{2} + \alpha$$

$$W_3 = \frac{\frac{\frac{W_N}{2}+\alpha}{2} + \alpha}{2} + \alpha$$

$$W_{max} = \lim_{i \to \infty} (2^{-i} W_I + \alpha (\sum_{j=0}^{i-1} 2^{-j})) = 2\alpha.$$

$$W_{avg} = \frac{3t}{4rtt}$$

α is the Cwnd growth during a sleep period

t time between two loss events

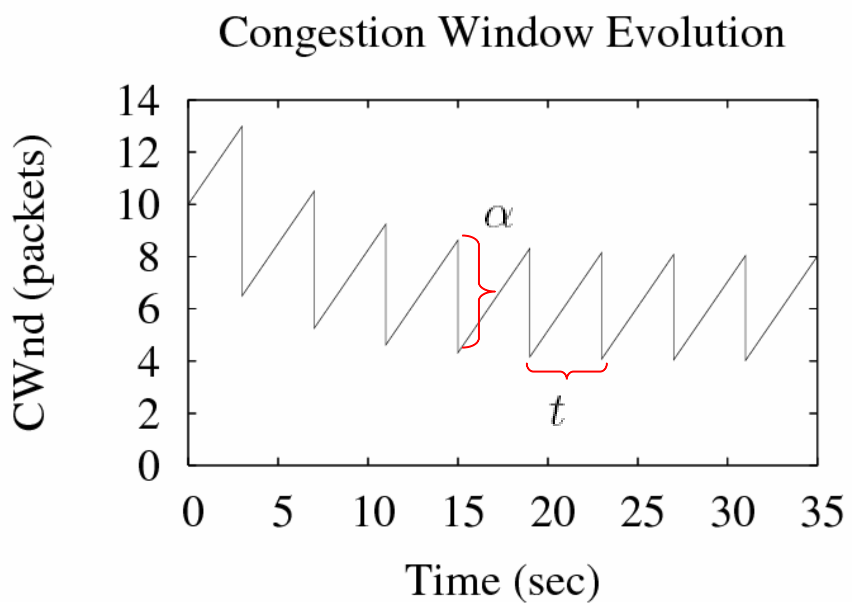

Analysis vs. Simulation

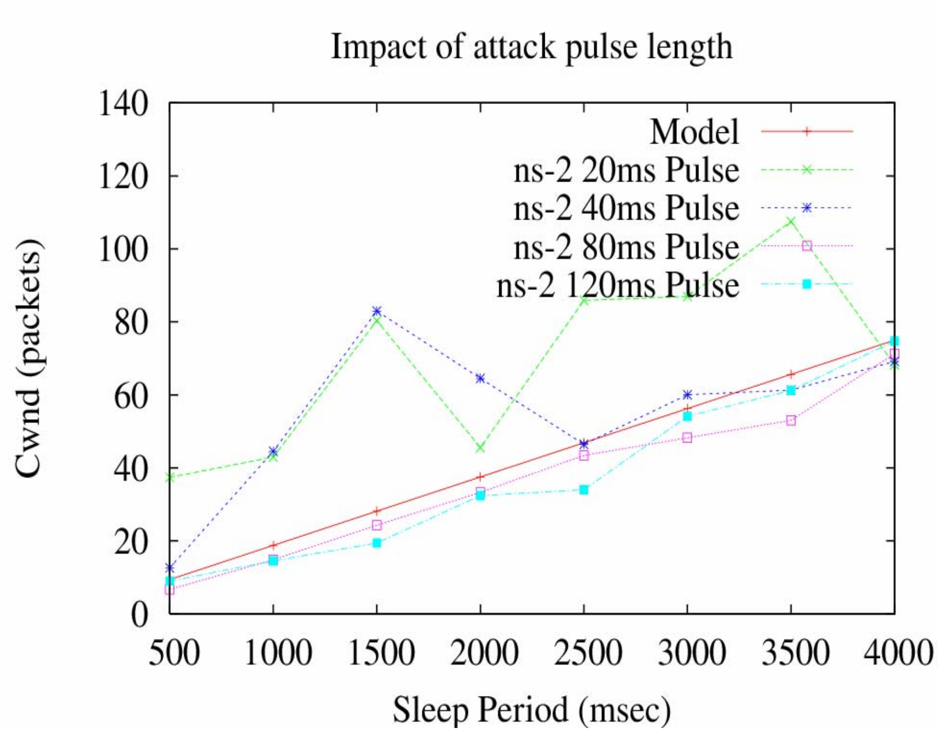

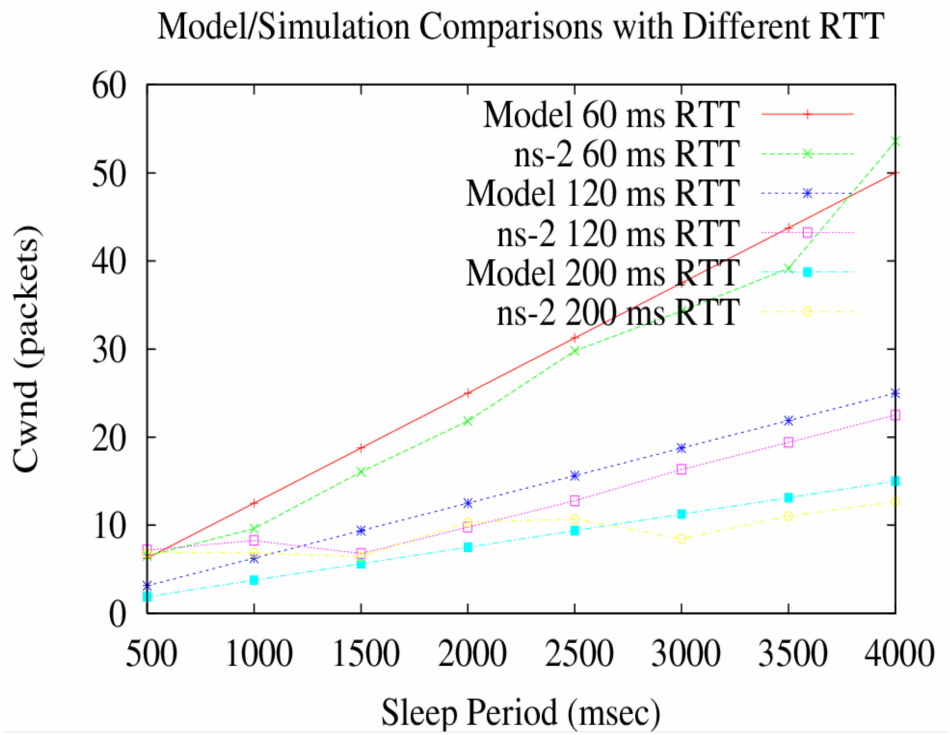

- Non-monotonic increase amplified by phase effects.
- Analysis corresponds to ns-2 results when attack pulse length is greater or equal to TCP flow RTT and when buffer sizes are not too large.

Forward Direction

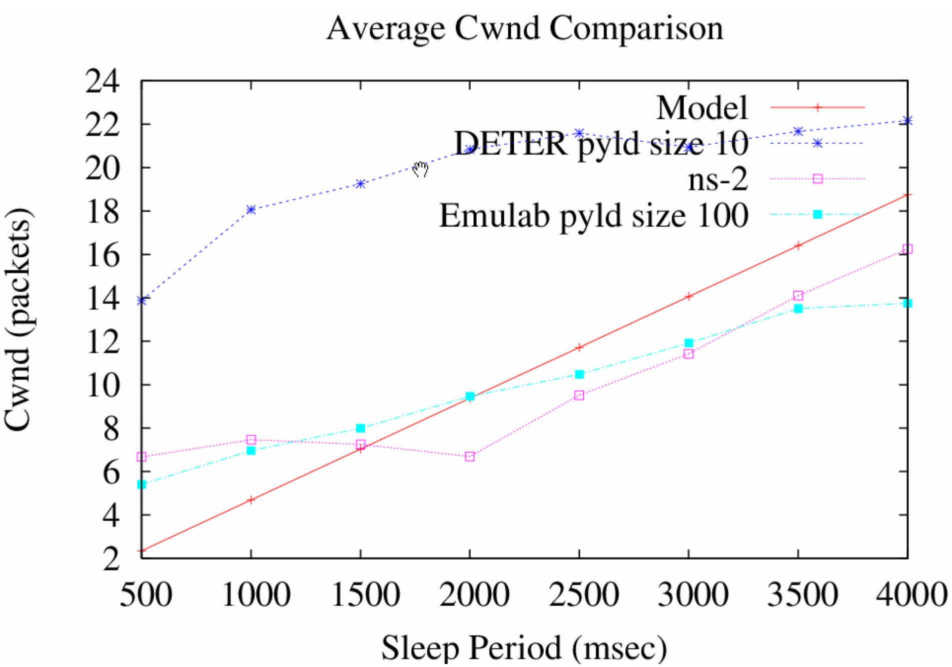

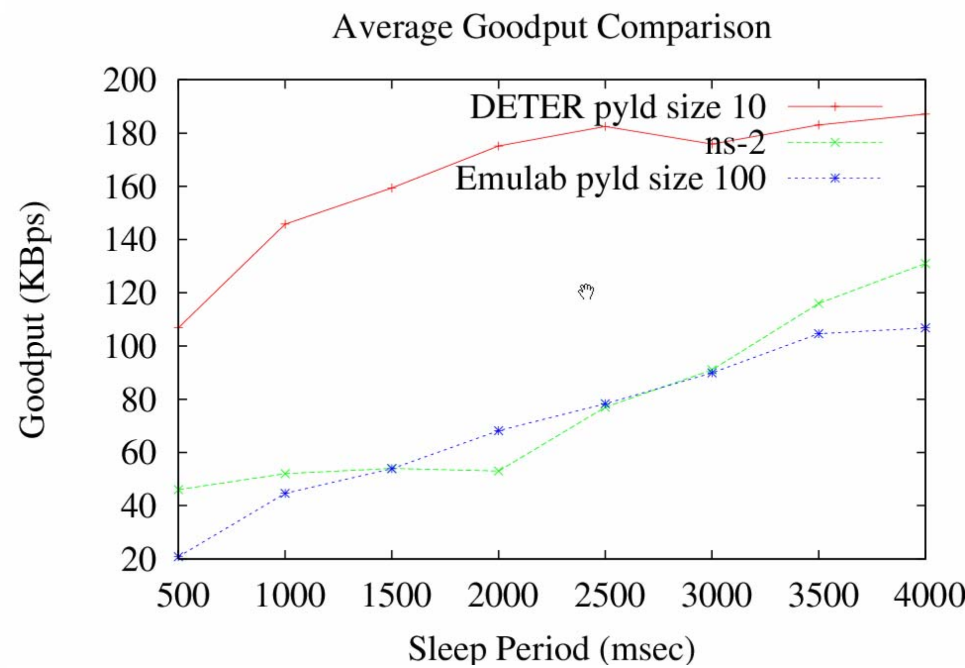

- ❑ Emulab results not too far from analysis and ns-2
- ❑ DETER is not as significantly affected by the attack
- ❑ Why? Bus, NIC, software, settings

❑ *Each emulation environment is a specific instance of the real world. There is no right or wrong, just specifics!*

Reverse Direction

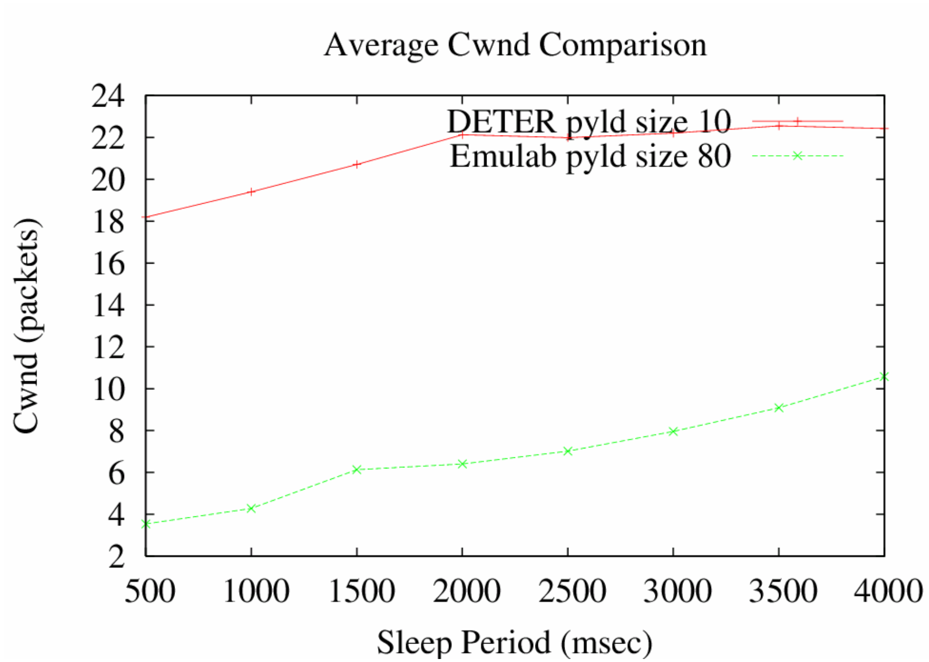

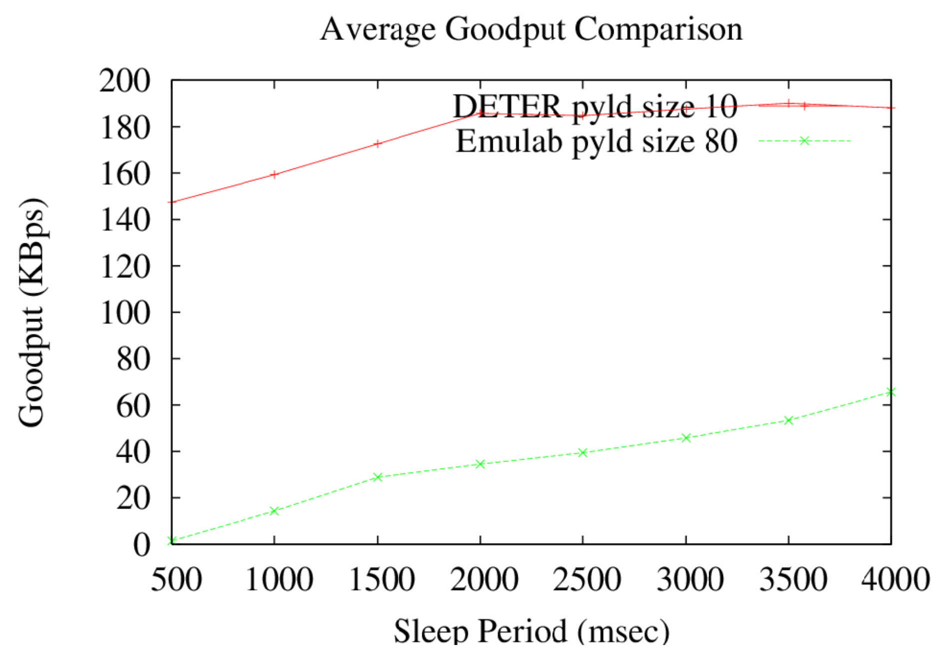

Since ns-2 does not model CPU/bus/devices, and opposing flows do not interfere at a router with output buffering, data for ns-2 is not shown for reverse direction (Cwnd has no cuts)

Receive/Interrupt Livelock

PURDUE UNIVERSITY

- Schemes that receive packets by invoking interrupts suffer from:
  - High CPU utilization
  - Reduced forwarding rate
  - Process starvation

- Polling resolves the above problems by:
  - Using software interrupts and a kernel thread reduces interrupt overhead by batching the receive signals
  - Batch limits govern the time the CPU spends in kernel mode processing the packets

- J. Mogul et al., Eliminating Receive Livelock in an Interrupt-driven Kernel, ACM Transactions on Computer Systems, 1997
- P. Druschel et al., Experiences with a High-speed Network Adaptor: A Software Perspective, SIGCOMM 1994
- Kohler et al., The Click Modular Router, ACM TOCS 2000

Router Nodes

- To avoid slowdown in the Linux kernel, the machine can be configured to run SMP enabled Click modular router with polling drivers.
  - Polling reduces CPU overhead by reducing interrupts.
  - Bypassing the Linux protocol stack speeds up packet processing.

Results with Click

PURDUE UNIVERSITY

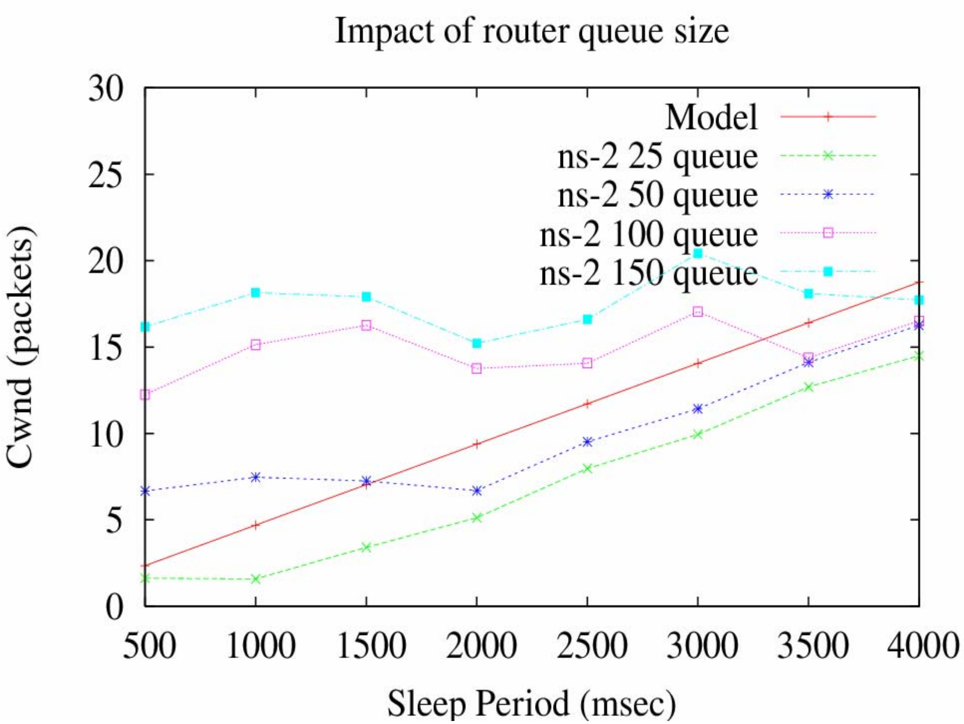

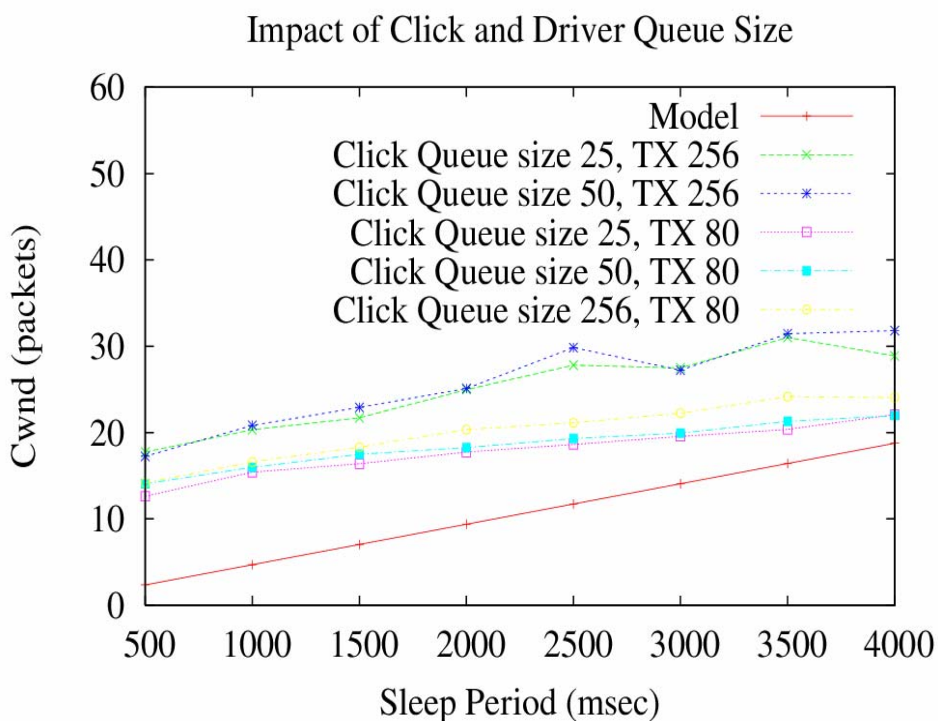

- The results indicate that device buffer size variation has a higher impact on the final results than Click buffers.
- It is important to understand device drivers so that accurate comparisons can be made.

Results on WAIL

PURDUE UNIVERSITY

- Wisconsin Advanced Internet Laboratory (WAIL) testbed http://schooner.wail.wisc.edu/ is based on Emulab
- WAIL contains Cisco routers from 2600 to 12000GSR series
- This provides an opportunity to compare *PC routers* versus *real Cisco routers*

WAIL Experimental Setup

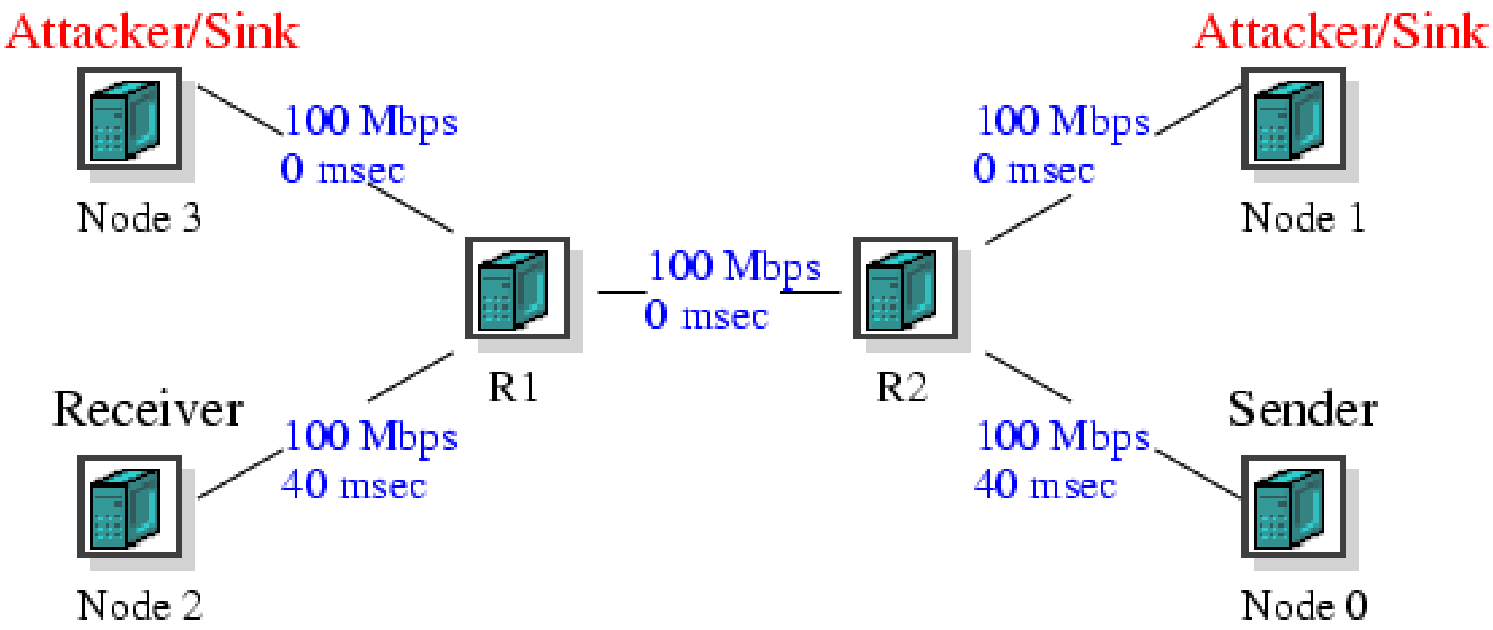

- R1, R2 are Cisco 3640 routers.
- Since the routers are directly connected, it is impossible to add a delay between them.
- Access link delays are equal to RTT/4.

Results with Cisco 3640

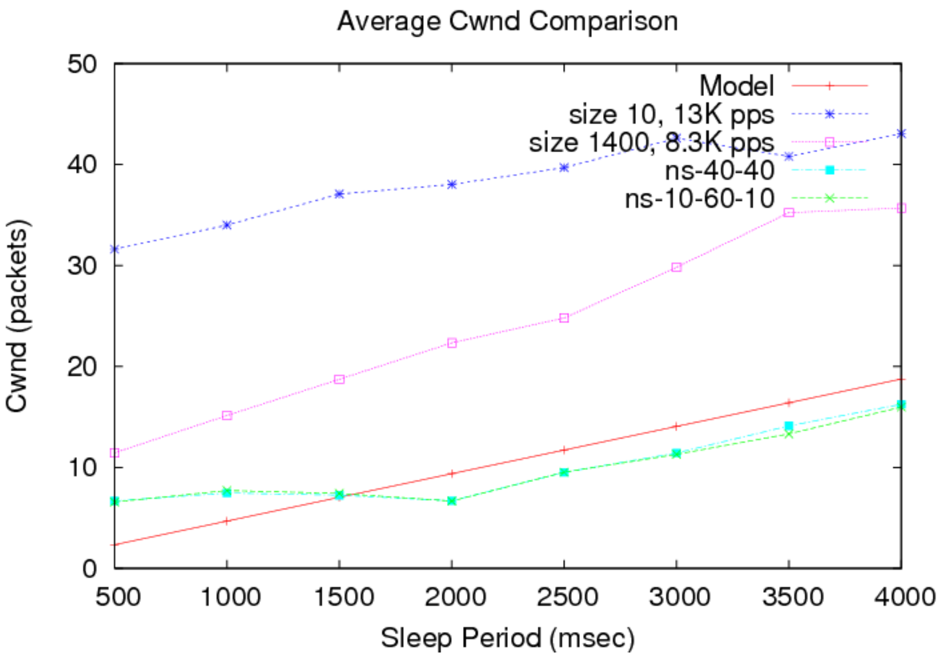

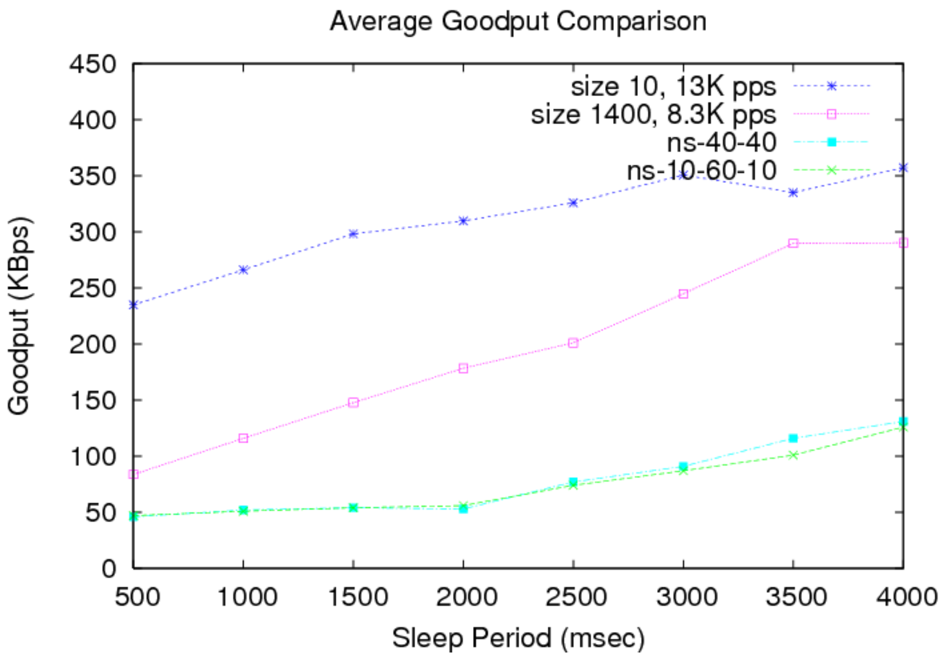

- Same TCP-targeted attack experiment as before
- Attack parameters are: TCP packets with 10 byte payload at 13 Kpackets/s and 1400 byte payload at 8.3 Kpackets/s

Results with Cisco 3640 (cont'd)

- We used TCP packets instead of UDP as the router's *policy* gives preference to TCP over UDP packets.
- The attack rate was limited to *Maximum Loss Free Receive Rate (MLFRR)* to avoid significant input queue packet loss.
- Contrary to previous results, *larger packets* caused more damage.

Cisco 7206VXR versus 3640

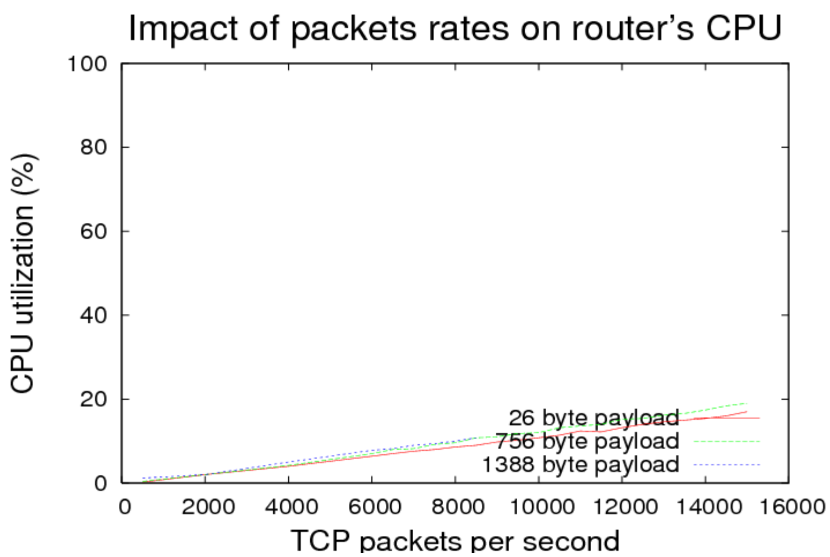

Cisco 7206VXR

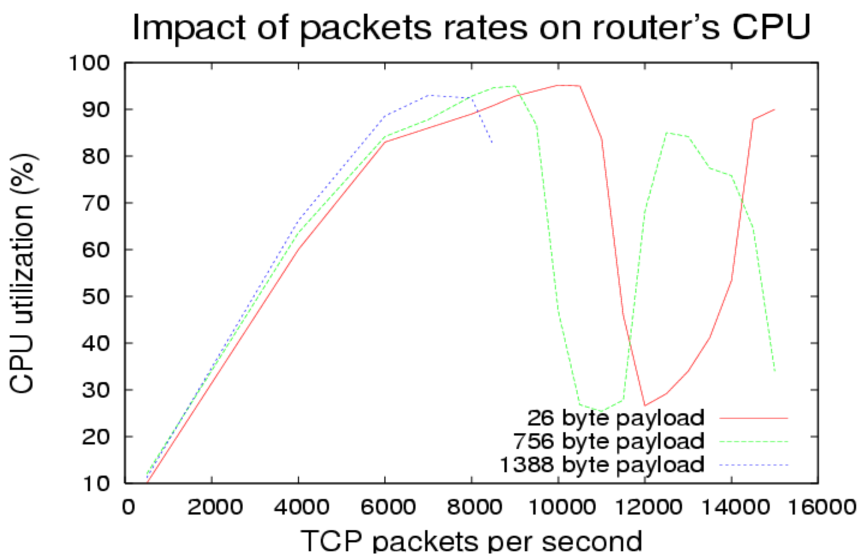

Cisco 3640

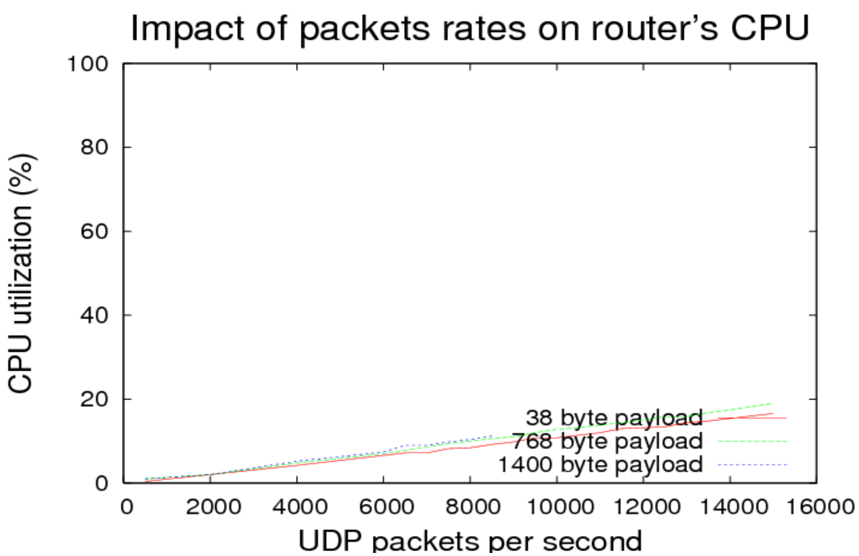

Cisco 7206VXR

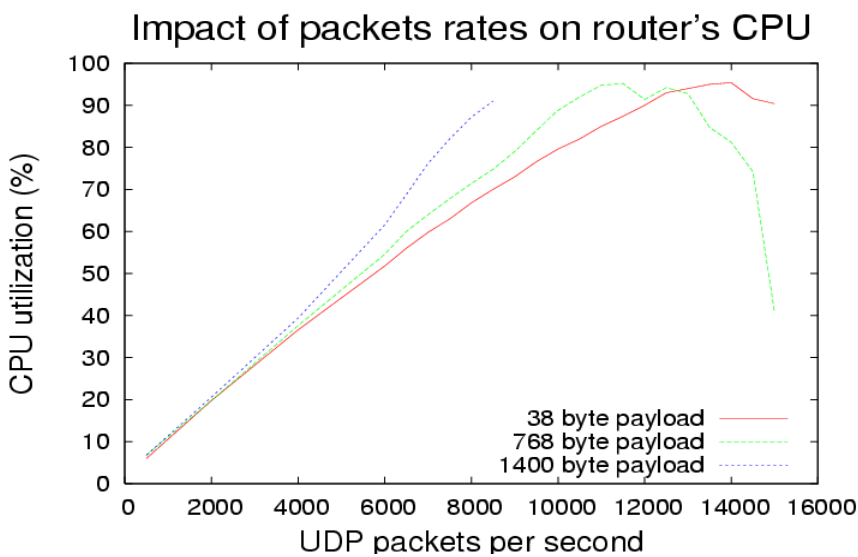

Cisco 3640

Cisco 7206VXR versus 3640

- Oscillations in the 3640 TCP plot are caused by CPU starvation of the accounting process.
- Superior hardware on the Cisco 7206VXR accounts for its vastly superior performance.

7206VXR vs. 2.0 GHz P4 PC

PURDUE UNIVERSITY

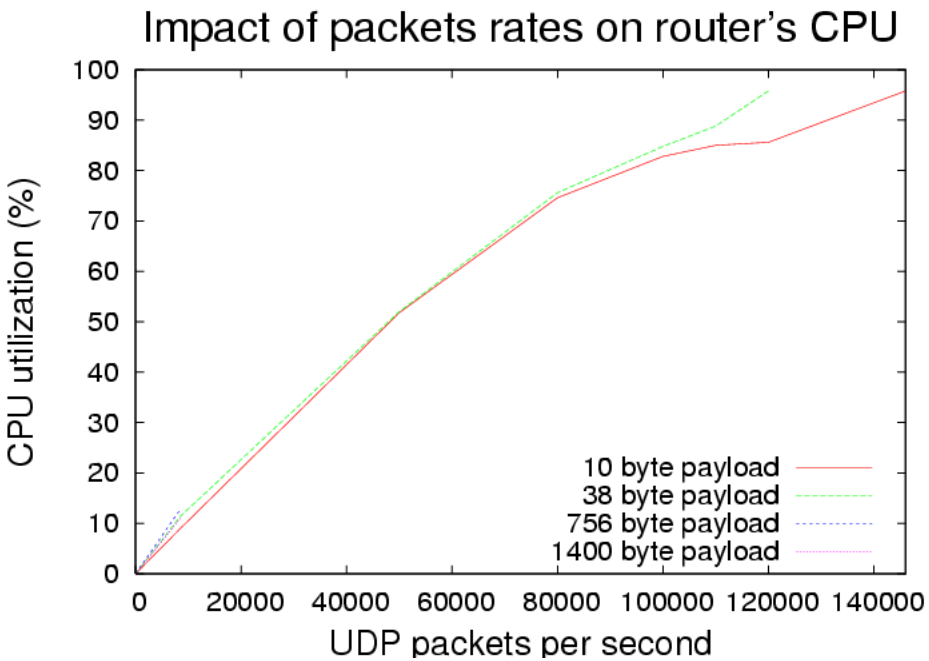

Cisco 7206VXR router

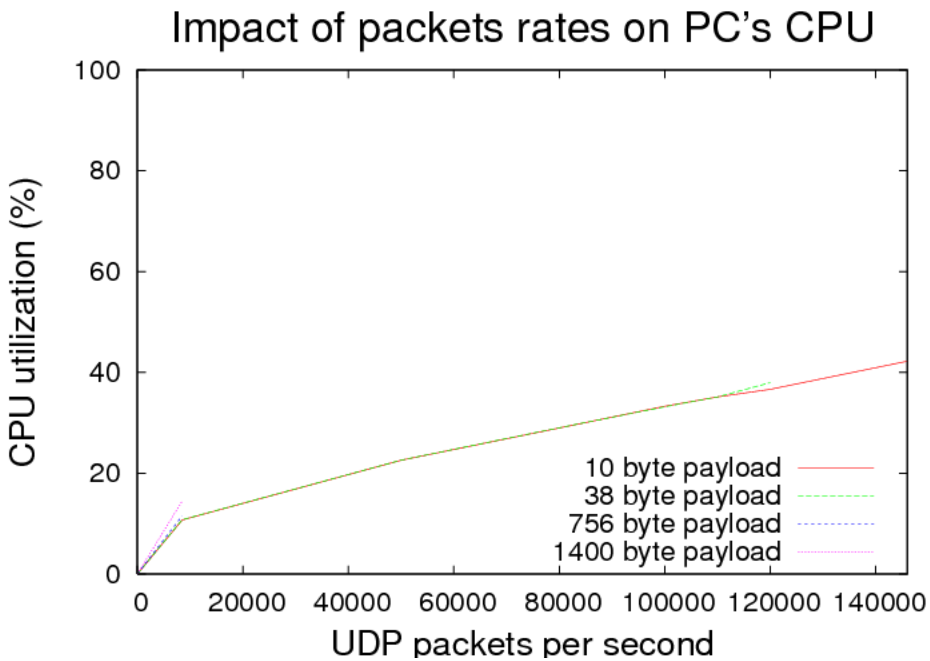

2.0 GHz P4 PC

- The New API (NAPI) NIC driver and the superior hardware on the PC lead to lower utilization.
- This shows that *PC routers* can be used to mimic hardware ones.

Conclusions

PURDUE UNIVERSITY

- TCP congestion control can be successfully exploited by a pulsing attack with a fraction of needed attack traffic when compared to a flooding attack; attack frequency need *not* be tuned to RTO
  - With a *single* flow under attack, attack pulse must be longer or equal to RTT and buffer sizes must not exceed 100 packets; attack packet size also an important parameter
- Simulation and emulation can produce *very different* results for very *similar* experiments
  - *Same experiment* on different emulation testbeds (or same testbed before and after hw/sw upgrades!) can yield different results
  - *Same experiment* on the same emulation testbed can yield different results depending on the driver settings
- Such differences are important as they allow us to identify *real vulnerabilities and fundamental limits*
  - The Internet is an evolving, *heterogeneous* entity with protocol implementation errors and resource constraints, and not a modeling approximation in a simulator

Conclusions (cont'd)

PURDUE UNIVERSITY

- Results and experiences demonstrate the need for a *high fidelity model in simulation and emulation environments*. This is critical for scenarios that push the limits of the network, such as DoS attacks.
- *PC routers* can be used to emulate real routers provided that they have a higher capacity than the target router. This includes single interface and aggregate forwarding performance.
- A cluster of PCs can be used to create scalable IP routers
  - V. Vuppala and L. Ni, Design of a Scalable IP Router, Hot Interconnects 1997
  - C. Tzi-Cker and P. Pradhan, Suez: a Cluster-based Scalable Real-time Packet Router, ICDCS 2000

Future Work

- Determine a set of profiling benchmarks representative of the real world to derive important values of router model parameters.
- Create a general router model and validate it by:
  1. Implementing the model in a simulator
  2. Implementing the model in Click
  3. Comparing the results with real routers
- Utilize the new models to perform network resilience validations