

Shaping Process Semantics

Christoph Kirsch
Universität Salzburg



Joint work with Harald Röck
Monterey Workshop, Paris, October 2006

The Idea

- we apply *traffic shaping* technology known in the networking community to *software processes*
- software processes invoke *system calls* to access resources, perform I/O, etc.
- we see system calls as *network packets*

Process Shaping

- process shaping changes the *order* and *times* in which system calls (and potentially other *side effects* of processes) are handled before given to any performance-oriented kernel subsystems
- process shaping promotes more disciplined *system composition*

Application Processes

Semantics

Timing
Calls

I/O Calls

IPC Calls

Virtual
Memory

Kernel

Process Shaping

Interrupt
Handling

I/O
Scheduling

IPC
Handling

Memory
Management

Performance

Hardware

Proposal


- we propose the notion of *process shaping* to complement, not replace, the notion of *servicing* processes *as fast as possible*
- we advocate a shift in research attention from performance- to *semantics-oriented* handling of software processes

Claim

- we claim that *faster* processors, more efficient scheduling, and lower kernel *latency*, in analogy to *shorter* packet transmission times, will make process shaping increasingly effective
 - ▶ see ATM versus Gigabit Ethernet
- note that used-to-be-exotic real-time patches increasingly make their way into general-purpose operating systems

Experiment I

- we run two separate web server processes on an *unmodified* Linux 2.6 server machine with Gigabit Ethernet
- two client machines generate workload by requesting the same and thus *cached* 380KB file




"net io 1" —+—

"net io 2" ---x---

"net io combined" ...*...

Experiment II


- we run two separate web server processes on a *process-shaping* Linux 2.6 server machine with Gigabit Ethernet
- two client machines generate workload by requesting the same and thus *cached* 380KB file



"net io 1" —+— "net io 2" ---x--- "net io combined" ...*...

Experiment I+II


- higher total **peak** performance without process shaping
- but total peak performance more **robust** with process shaping



"net io shaping 30k" —+— "net io no shaping" ---x---

Experiment III


- we run a *video-streaming* server on a *process-shaping* Linux 2.6 server machine with Gigabit Ethernet
- to generate background *network traffic*, we also run one of the web servers of experiment II on the same machine
- to generate background *disk traffic*, we also run several web servers processing requests for a non-cached 1 GB file



Net. I/O raw —+—

Net. I/O succ. ---x---

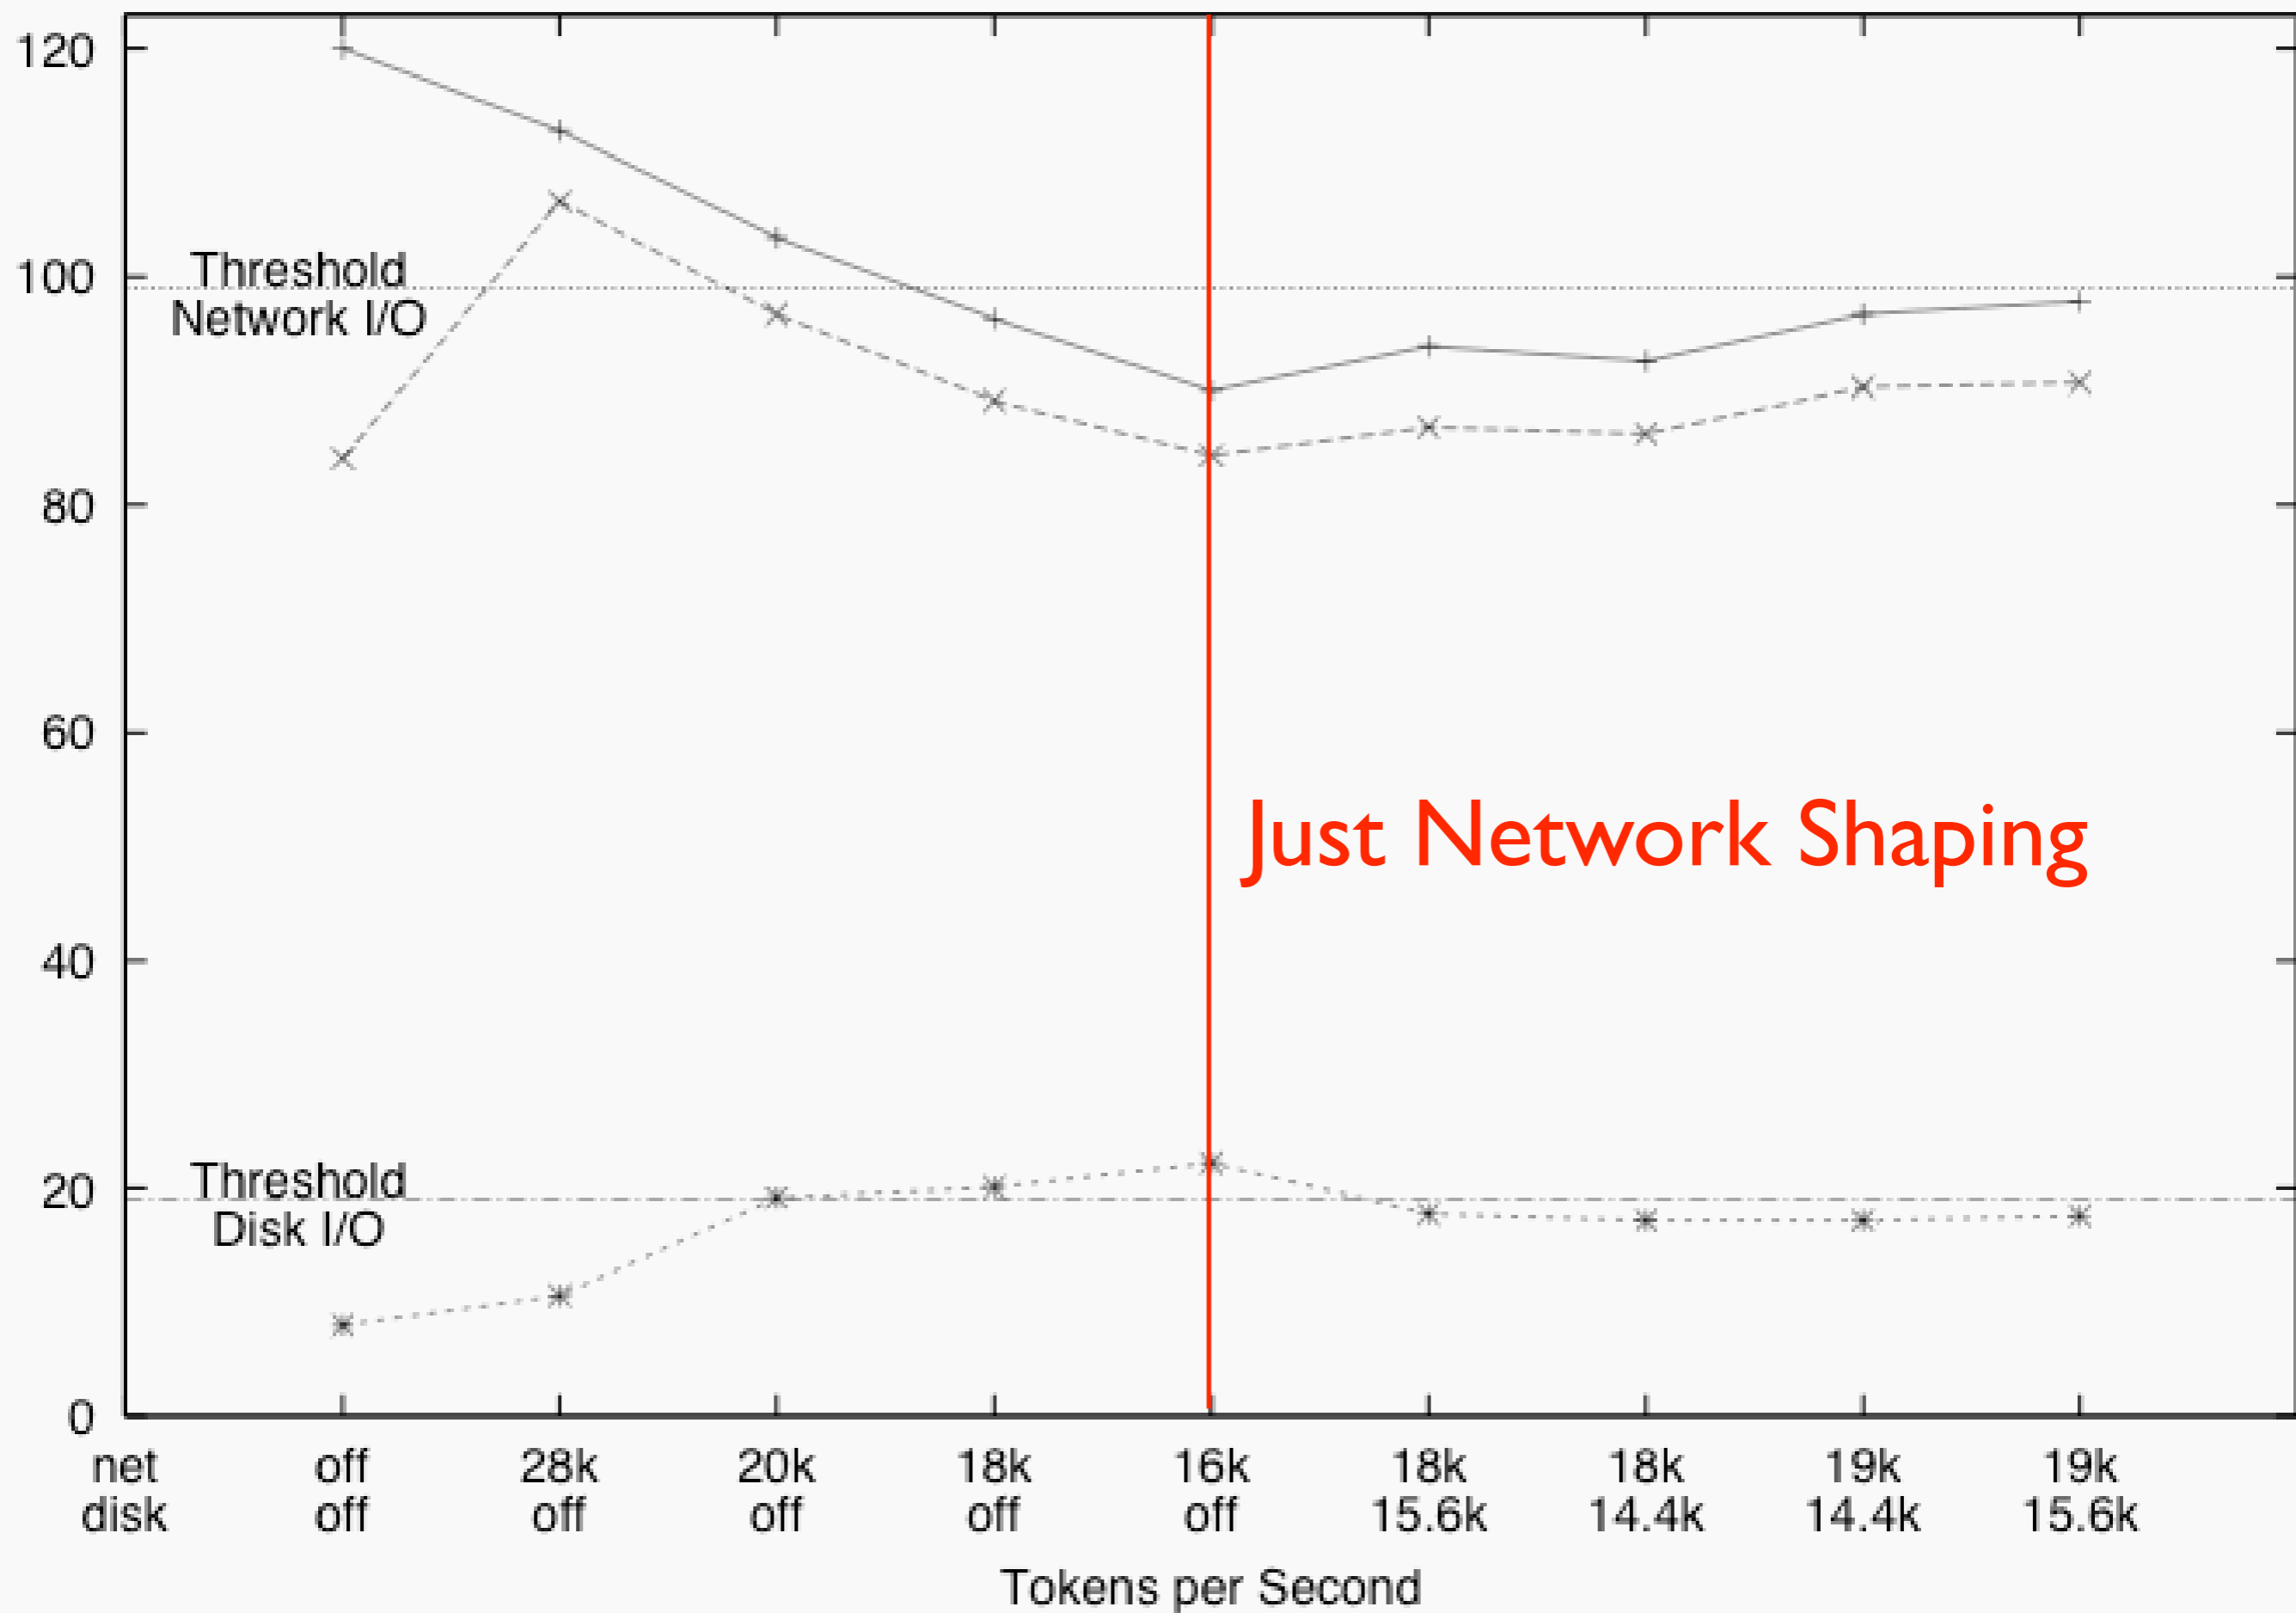
Disk I/O ...*...



Net. I/O raw —+—

Net. I/O succ. ---x---

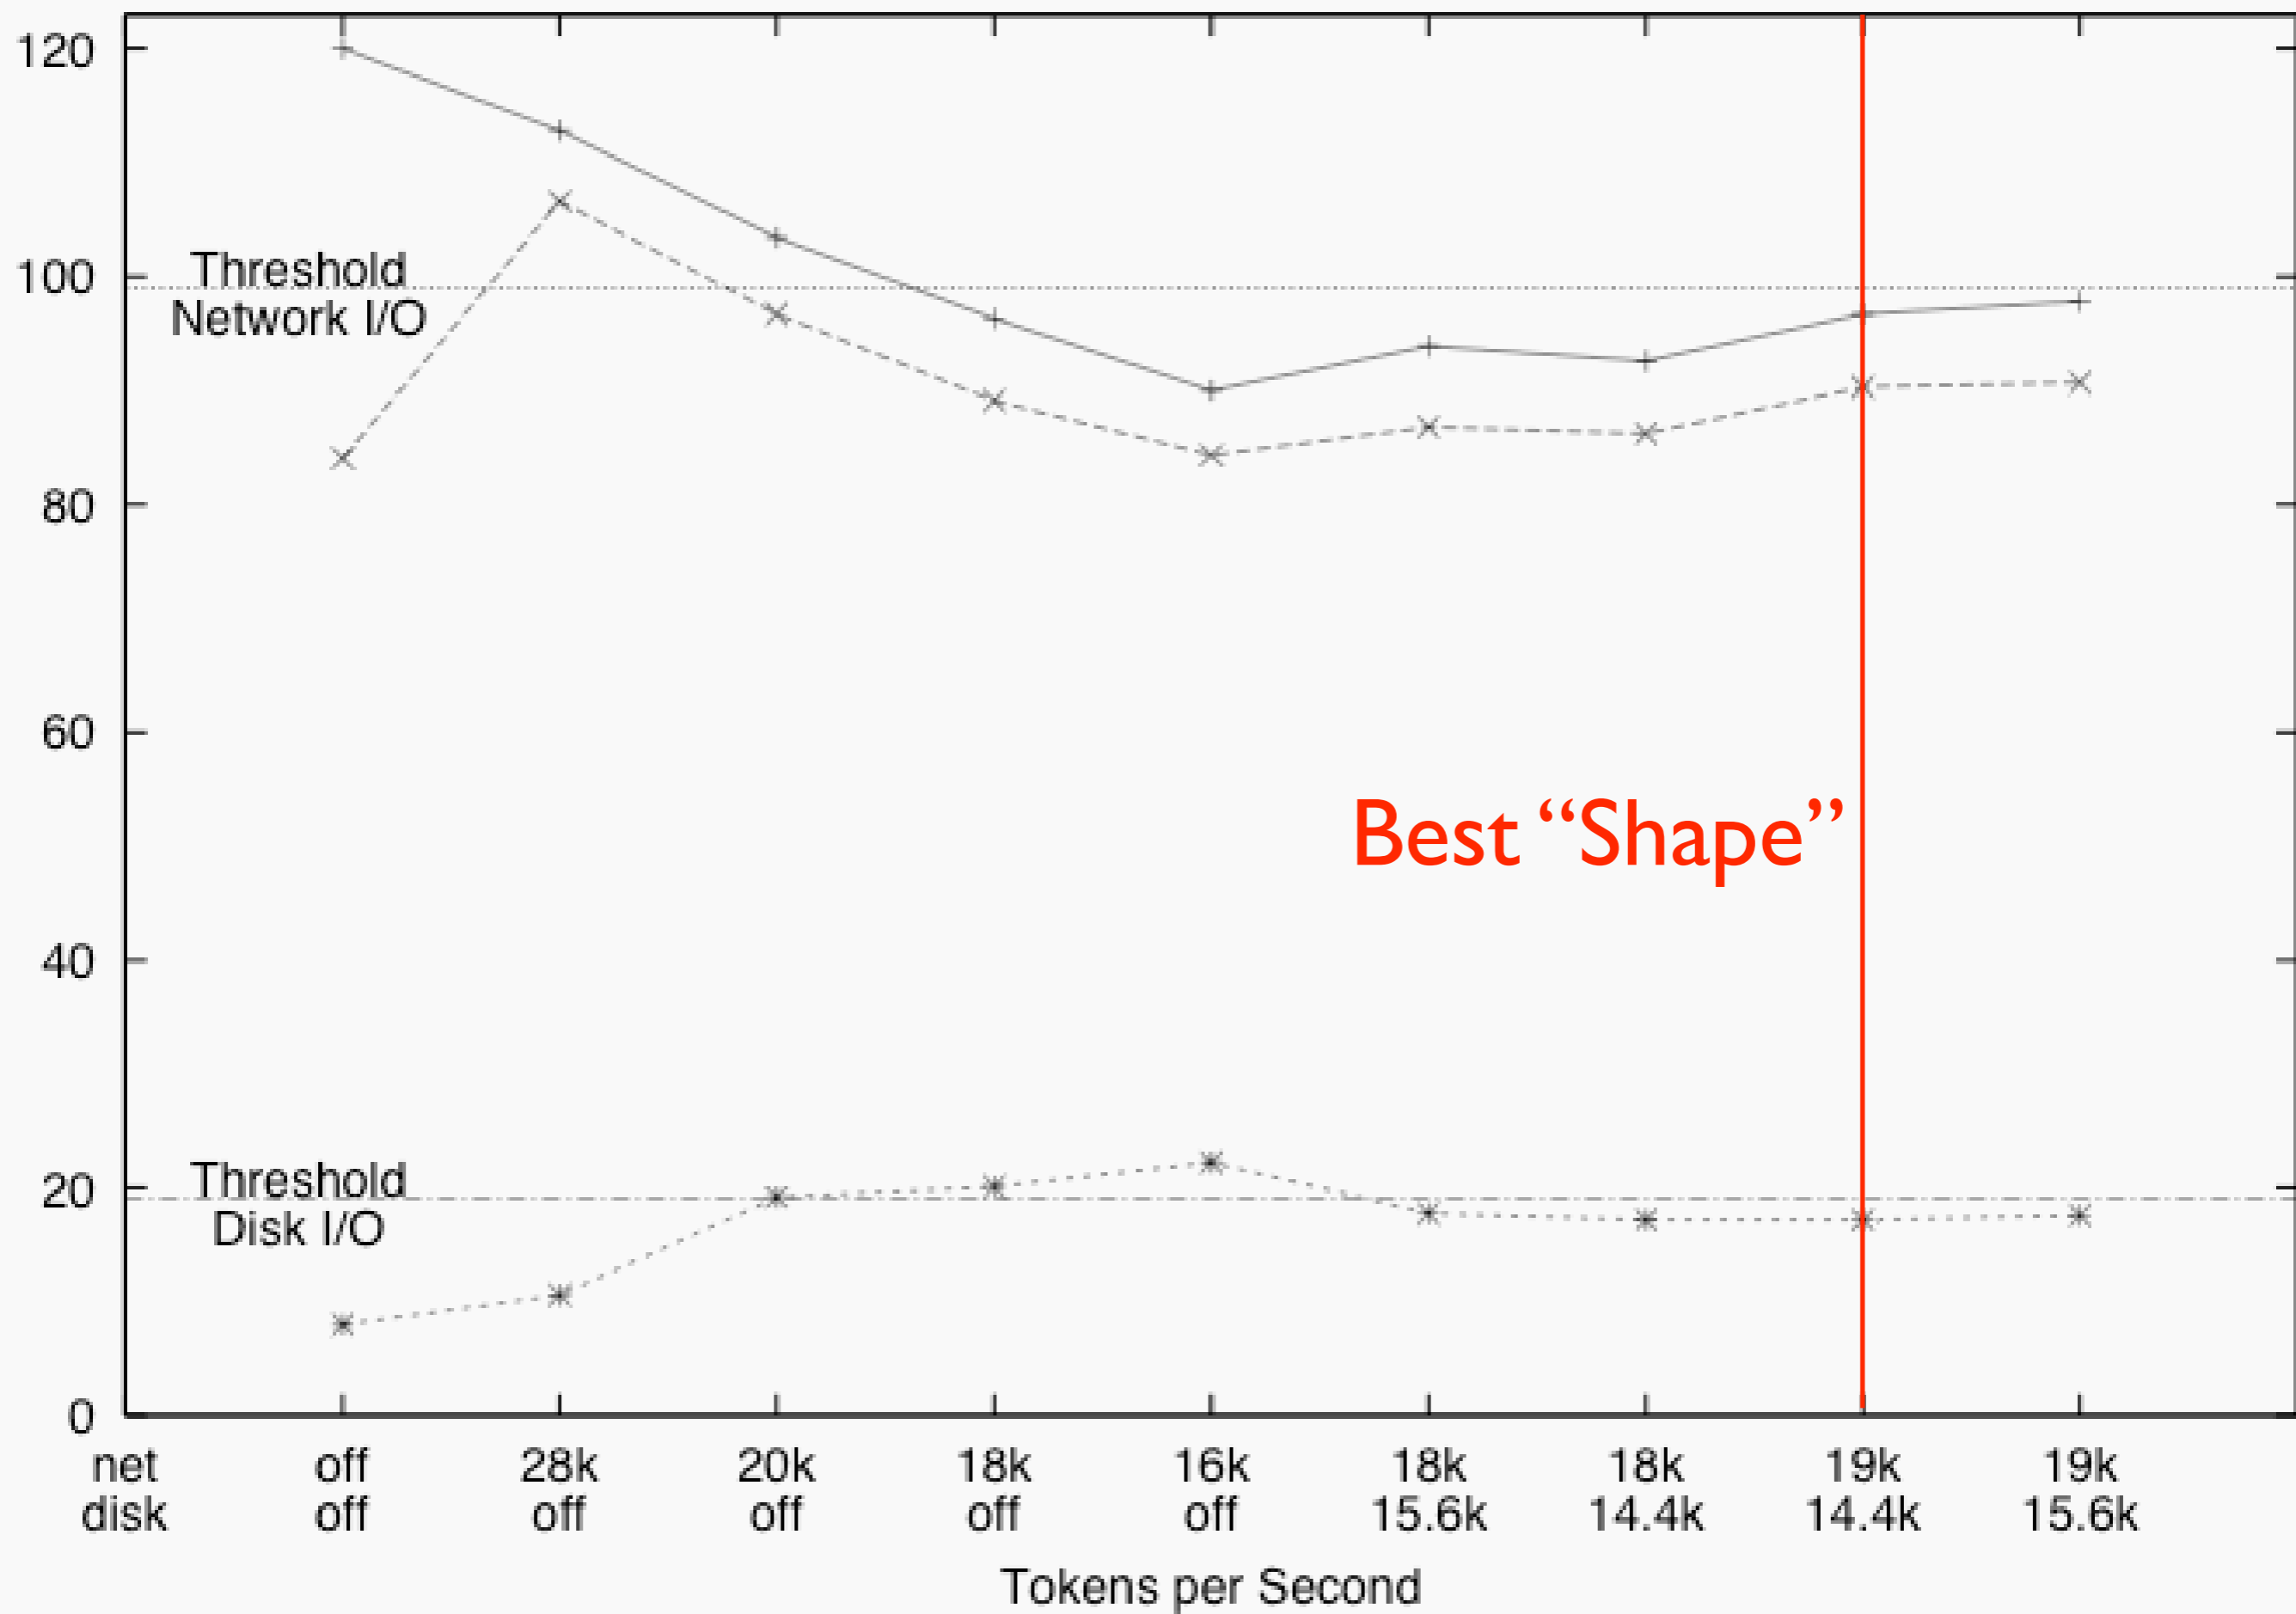
Disk I/O ...*...



Net. I/O raw —+—

Net. I/O succ. ---x---

Disk I/O ...*...



Net. I/O raw —+—

Net. I/O succ. ---x---


Disk I/O ...*...

Future

- Multimedia (*Soft* Real Time): Can kernel-level process shaping automatically find the best “shape”?
- Control (*Hard* Real Time): Can kernel-level process shaping provide sufficient real-time guarantees?


Experiment IV

- we run helicopter flight control software written in *Java* on IBM's commercial J9 JVM with the real-time garbage collector *Metronome* on top of a Linux 2.6 machine with real-time patches applied to the kernel
- joint work with J. Auerbach, D. Bacon, H. Röck, and R. Trummer



The JAviator Project

javiator.cs.uni-salzburg.at



First All Java Flight



Thank you