

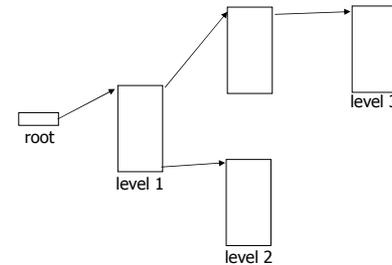
## Selected Papers from EuroSys'08

Kai Shen

Dept. of Computer Science, University of Rochester

## Parallax: Virtual Disks for Virtual Machines (British Columbia)

- Storage system specifically for virtual environments
- Basic storage abstraction - virtual disk image (VDI)
  - meta-data in three-level radix tree



Space sharing across VDI's through COW

4/10/2008

URCS 573 - Spring 2008

2

## Discussion Points

- Deletion of a virtual disk image
  - creation is simple but deletion is a pain
  - garbage collection on disk is a non-starter (for me), but OK for them
- Space sharing vs. storage sequentiality
  - small blocks meaning better space sharing, but poor storage sequentiality

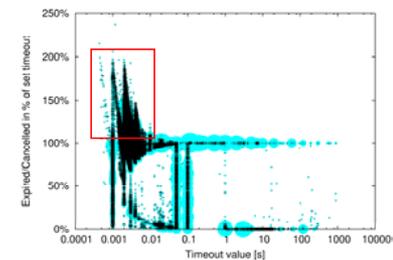
4/10/2008

URCS 573 - Spring 2008

3

## A Study of Operating System Timer Usage (ETH Zurich, MSR Cambridge)

- Timer usage of Skype running on Vista [Courtesy of Peter Simon, ETH Zurich]



4/10/2008

URCS 573 - Spring 2008

4

## Discussion Points

- The timer interface is rigid - add a specification of allowed error
- Adaptive timer - TCP ACK timeout
- High-resolution timer

## Samurai: Protecting Critical Data in Unsafe Language (Illinois, Nvidia, MSR)

- Problems that may corrupt your data:
  - runaway writes in C
  - transient memory hardware errors
- Solution:
  - keep two critical copies of data in addition to normal copy
  - identify critical data accesses (including all legal writes)
  - critical writes update normal copy and both critical copies; critical reads retrieve the majority value
  - easy to do at language level

