

## Scheduling I/O on VMMs

### Xen VMM background

- Why does the driver domain exist?

### Xen's Credit Scheduler

- Describe how the credit scheduler works; use a simple example (e.g., Domain A has 1000 credits, B has 500) to illustrate basic functionality
- What are the UNDER and OVER states?
- How does this compare to an approach using Lottery or Stride scheduling? (look these up if you are unfamiliar with them)

### Events

- Describe the flow of control on an interrupt in Xen
- What is "tickling" the scheduler? What would happen w/o such tickling?

### Fairness

- What does Figure 2 show us? (other than it is the most boring bar graph in the history of humanity)
- Credits are debited from a domain only when it runs for its full time slice (10ms); what problems does this raise? How can you fix it?
- What does Figure 3 show us?
- A new state (BOOST) is added to address some of these problems; what is it? How does it work? How does it interact with tickling?
- Figure 4 measures something about BOOST; what does it show?
- Does BOOST fix the problem shown in Figure 3?
- How could a wily hacker take advantage of BOOST?

### Enhancements

- Event channel notification in Xen was kind of broken; what was the problem? What solution was suggested? Can you think of other solutions? What is the broader lesson here?
- What is priority inversion? How can it happen in Xen w.r.t. the driver domain? How do the authors fix this problem? Does this introduce any new problems?
- How is the run queue ordered? What is the suggested fix?
- How does SEDF work?

### Experimental Method and Results

- What are burn, stream, and ping? What else could be added to this?
- What is in Table 1?
- What does Figure 5 show about the original Xen setup? Which factors lead to fairness across guests?
- What does Figure 6 show us? Why is there so little variance in original Xen for stream? Which approach is best here?
- What is the diff between Figs 7 and 8? Where is the original Xen number? What do these figs show us?
- What do figs 9 and 10 show us?
- Figs 11 and 12?

### Conclusions

- What conclusions can be drawn about the various "optimizations" in the paper?