Performance in Virtual Environments

Stefan Appel

TECHNISCHE UNIVERSITÄT DARMSTADT

DVS

Analysis of Resource Sharing in Overbooked Virtual Environments

TECHNISCHE UNIVERSITÄT DARMSTADT

- Virtualization is used heavily nowadays (cloud computing)
- Physical resources are shared between virtual machines
- Are resources shared fairly when virtual resources exceed physical resources?

- CPU: yes, Memory Bandwidth: yes, Disk I/O: it depends

DVS

Test Setup

TECHNISCHE UNIVERSITÄT DARMSTADT

- Hardware
  - IBM x3850 Server
  - 4 x Dual-Core Xeon 7150N 3.5GHz
  - 16GB RAM
  - 6 x 10.000 RPM SAS HD, RAID 10

- Software
  - Host OS: Debian Linux, etch
  - Hypervisor: VMWare Server 2.0
  - Guest OS: Ubuntu Linux, 8.04

- Scenario:
  1–7 Virtual Machines (VMs) in parallel

DVS

Testing CPU Performance in Parallel Running VMs

TECHNISCHE UNIVERSITÄT DARMSTADT

- Benchmark
  - SPECjvm2008 Benchmark Suite
  - 11 Applications / Workloads
  - Composite score & separate scores
- Virtual Machine Setup
  - 2 vCPUs
  - 1024MB RAM
  - 512MB JVM Heap Size
- CPU overbooking with 5+ VMs in parallel (8 cores available)
- SPECjvm2008 started simultaneously in 1-7 VMs

DVS

Fair CPU Sharing Between VMs

TECHNISCHE UNIVERSITÄT DARMSTADT

	Number of Virtual Machines				
	1	4	5	6	7
Average SPECjvm2008 Score	14.700	14.061	11.786	9.707	8.094
Standard Deviation	-	0.121	0.084	0.130	0.181
Accumulated Score	14.700	56.240	58.930	58.240	56.660

- Fair distribution of CPU time among VMs
  - Low standard deviation
- Overhead increases slightly with increasing number of VMs
  - Accumulated score decreases

DVS

Different behavior of benchmarks due to amount of parallelism

TECHNISCHE UNIVERSITÄT DARMSTADT

- 1 Virtual Machine vs. 4 Virtual Machines
  - No performance difference for some benchmarks: compress, mpegaudio, scimark.small
  - Significant performance difference for other benchmarks: compiler, xml
  - → Different amount of parallelism

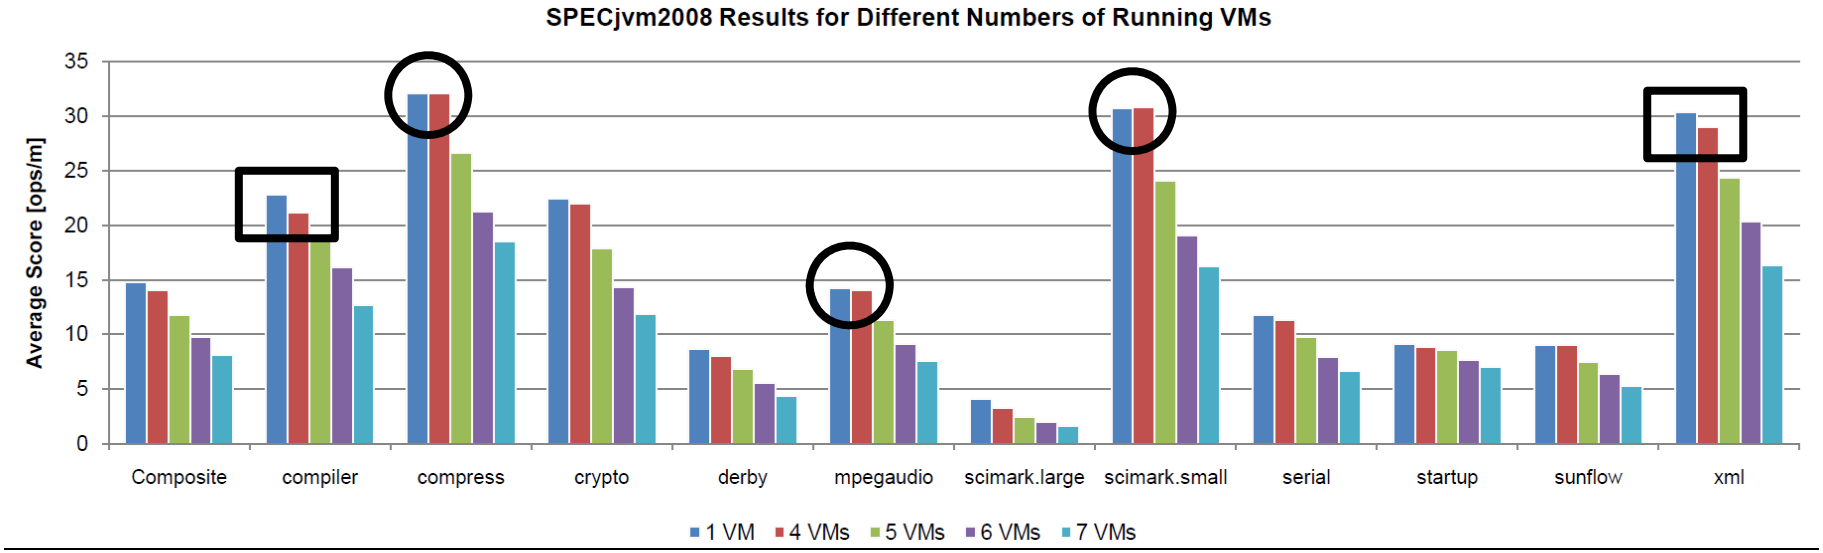

DVS

CPU not Fully Utilized During Benchmark Run

TECHNISCHE UNIVERSITÄT DARMSTADT

- Parts of SPECjvm2008 do not utilize two CPU cores
- 5+ Virtual Machines necessary to fully utilize host system

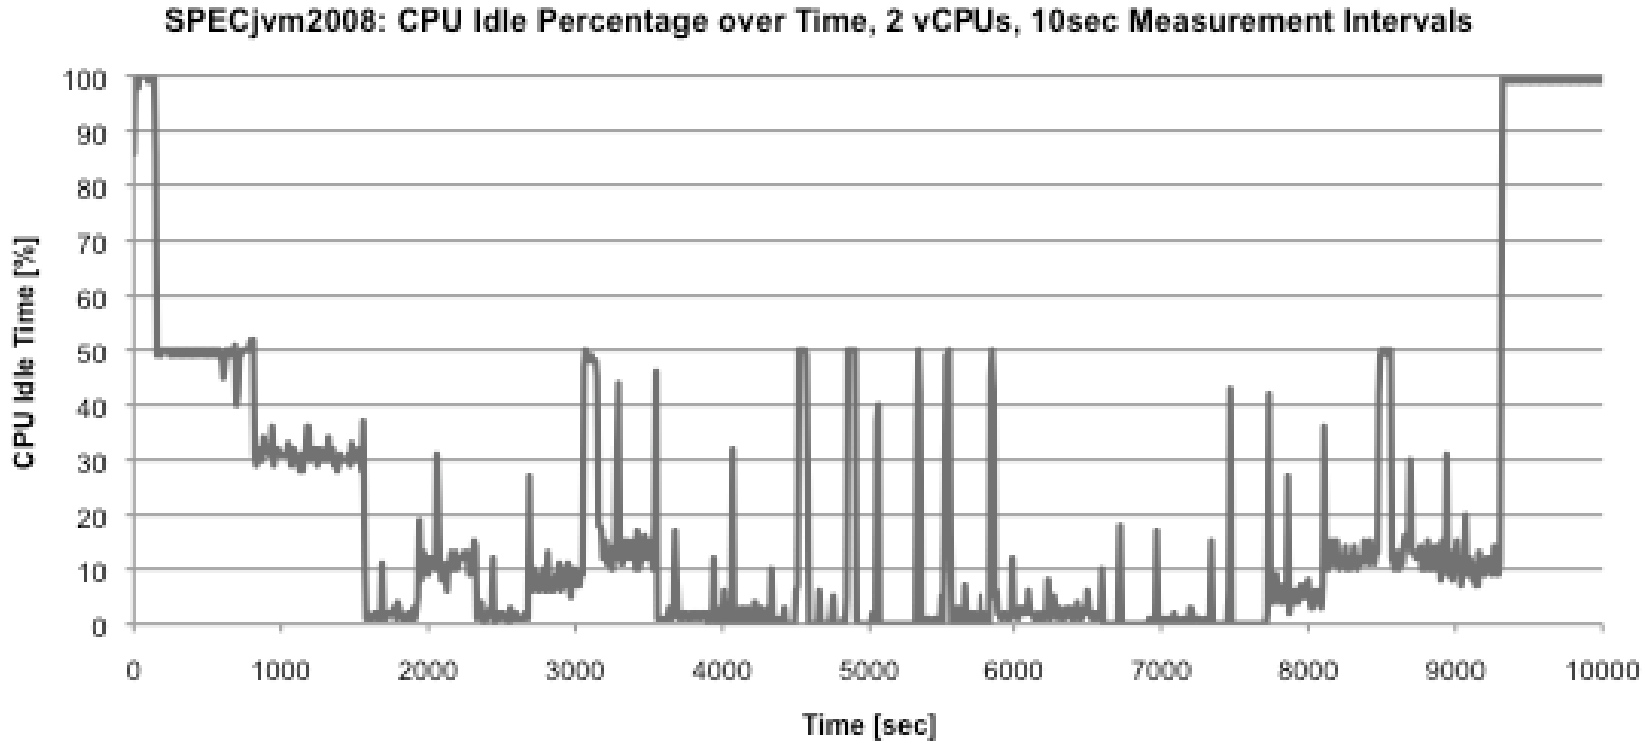

DVS

Testing Memory Throughput in Parallel Running VMs

TECHNISCHE UNIVERSITÄT DARMSTADT

- Benchmark
  - RAMSPEED: Memory throughput, one thread
  - RAMSMP: Memory throughput, multiple threads
  - COPY (A=B), SCALE (A=m*B), ADD (A=B+C) and TRIAD (A=m*B+C) operations
- Virtual Machine Setup
  - 2 vCPUs, 2048MB RAM
  - Transfer of 8GB of data, 5 runs
- CPU overbooking with 5+ VMs in parallel
- Physical amount of RAM (16GB) sufficient, no swapping
- RAMSPEED/RAMSMP started simultaneously in 1-7 VMs

DVS

Full Memory Bandwidth only with 3+ VMs in Parallel

TECHNISCHE UNIVERSITÄT DARMSTADT

- Max. throughput requires utilization of multiple CPUs
  - Utilization of all memory controllers and caches
- Low overhead in highly utilized system
  - Overall throughput decreases slowly with increasing number of VMs

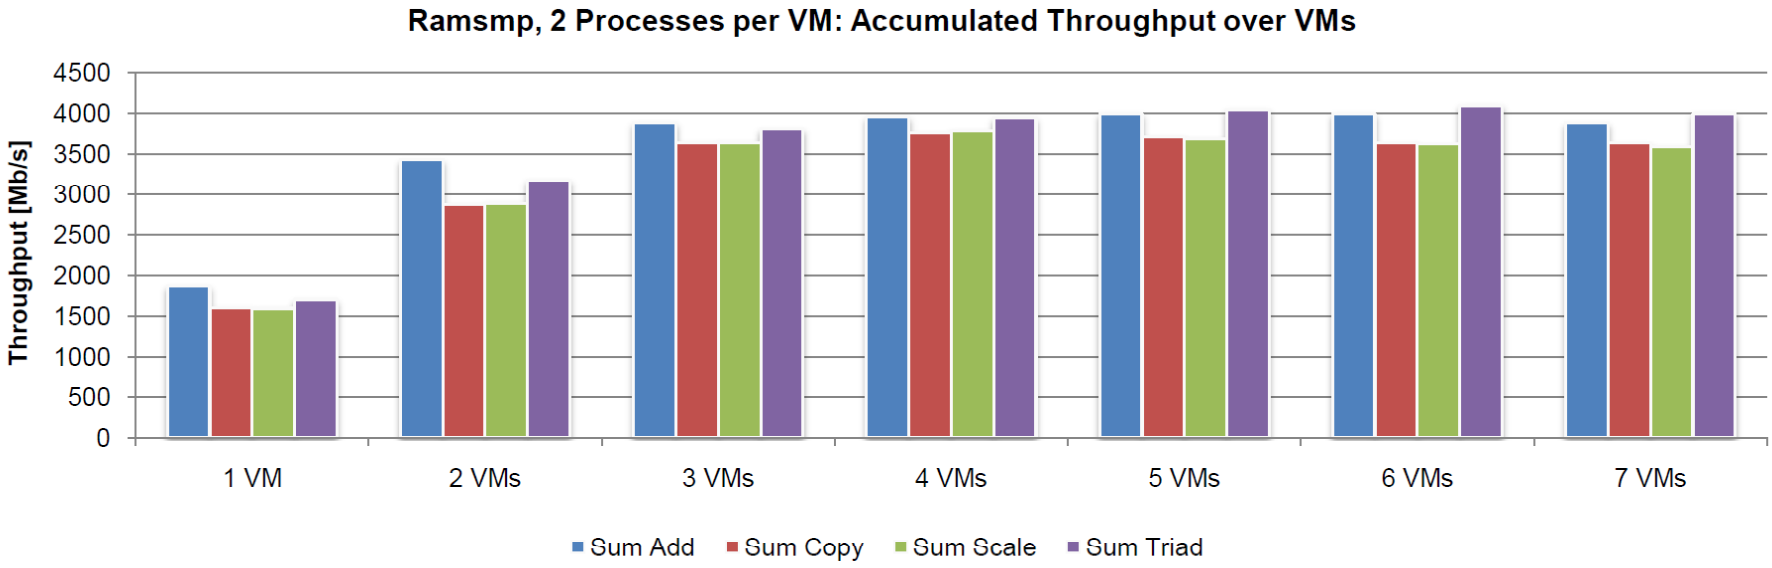

DVS

Memory Bandwidth is Distributed fairly among VMs

TECHNISCHE UNIVERSITÄT DARMSTADT

- Hypervisor distributes available memory bandwidth uniformly
  - Low standard deviations when comparing throughput per VM
- Slightly increasing std. dev. with increasing number of VMs
  - Fair distribution of resources more difficult with more VMs

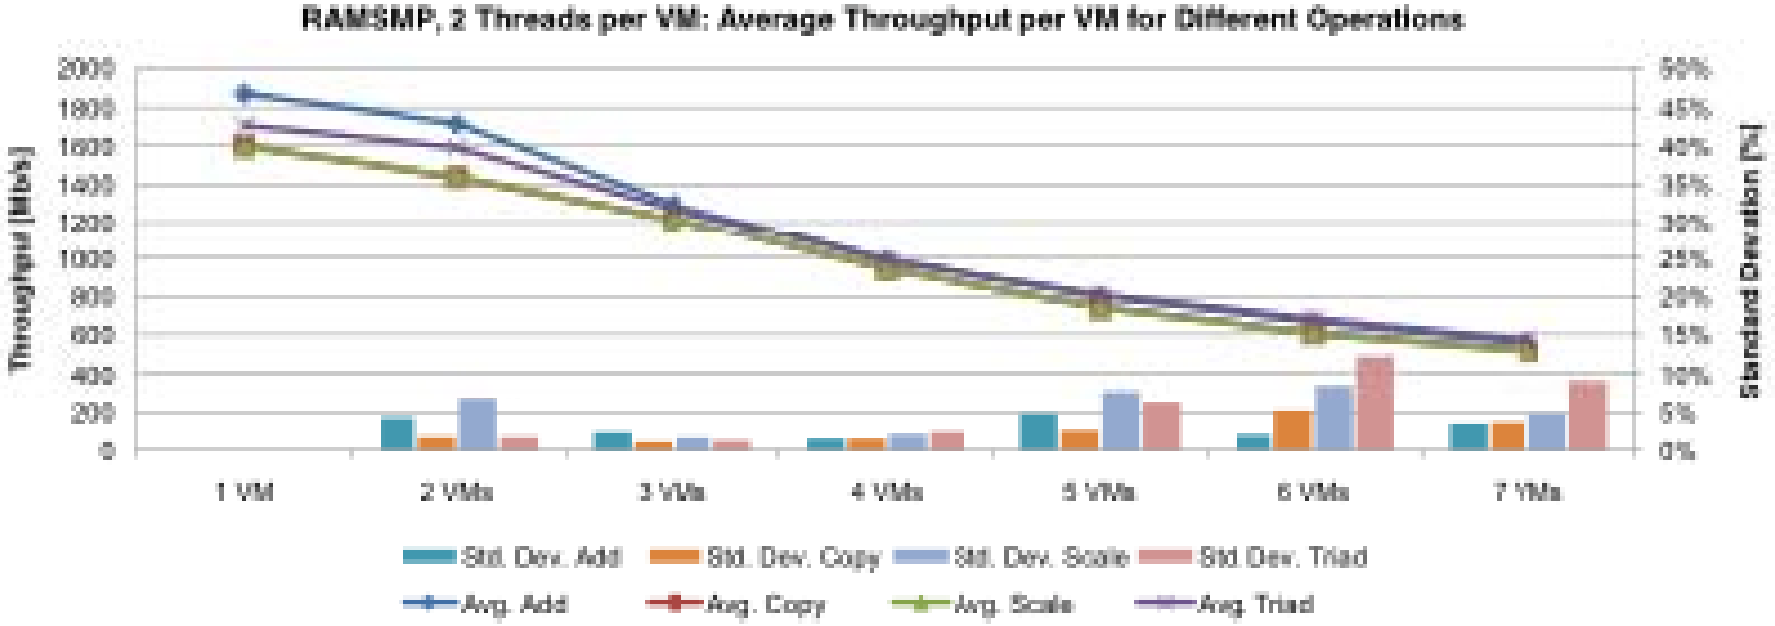

DVS

Testing IO Performance in Parallel Running VMs

TECHNISCHE UNIVERSITÄT DARMSTADT

- Benchmark
  - Bonnie++: *putc(), writec(), write(), read()*; Character- and Blockwise
  - Iozone: Write, Re-Write, Read, Random Read; different Blocksizes
- Virtual Machine Setup
  - 2 vCPUs
  - 1024MB RAM
  - 40GB disk, Benchmark file size: 2GB
- Scenarios
  - Bonnie++ and Iozone in 1,3 and 5 VMs in parallel: sufficient CPUs für 3 VMs, sufficient RAM

DVS

The Average IO Throughput per VM is Constant

TECHNISCHE UNIVERSITÄT DARMSTADT

- Repeated Iozone and Bonnie++ runs
  - Average throughput to and from hard disk is constant
  - Different values of Bonnie++ and Iozone due to different mechansims

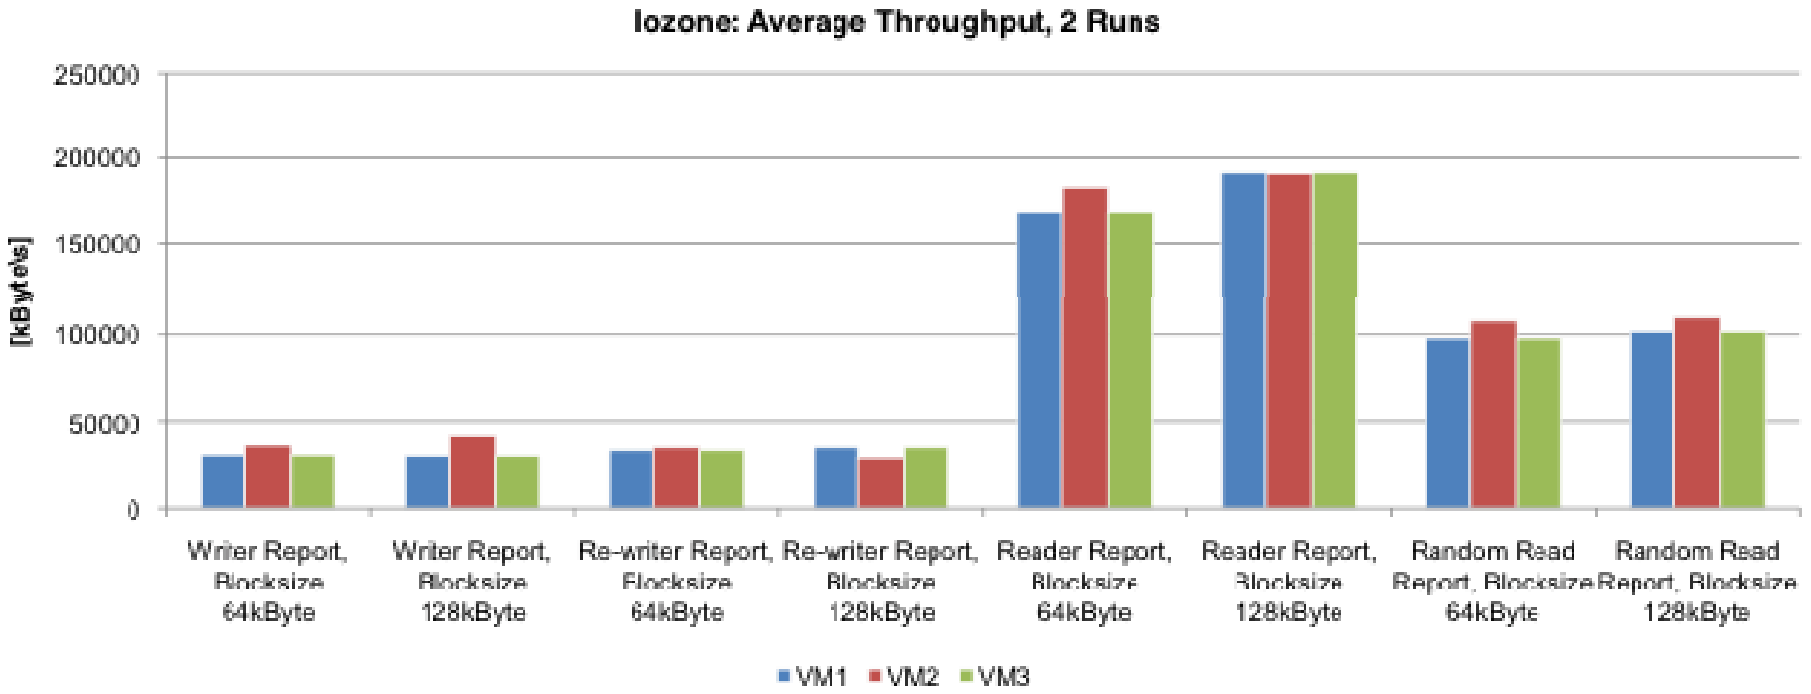

DVS

High Differences in IO Throughput between Runs

TECHNISCHE UNIVERSITÄT DARMSTADT

- Standard deviation almost always exceed 10%
  - No uniform distribution of IO bandwidth throughout a single run
  - Same for Iozone runs with different block sizes and Bonnie++ runs

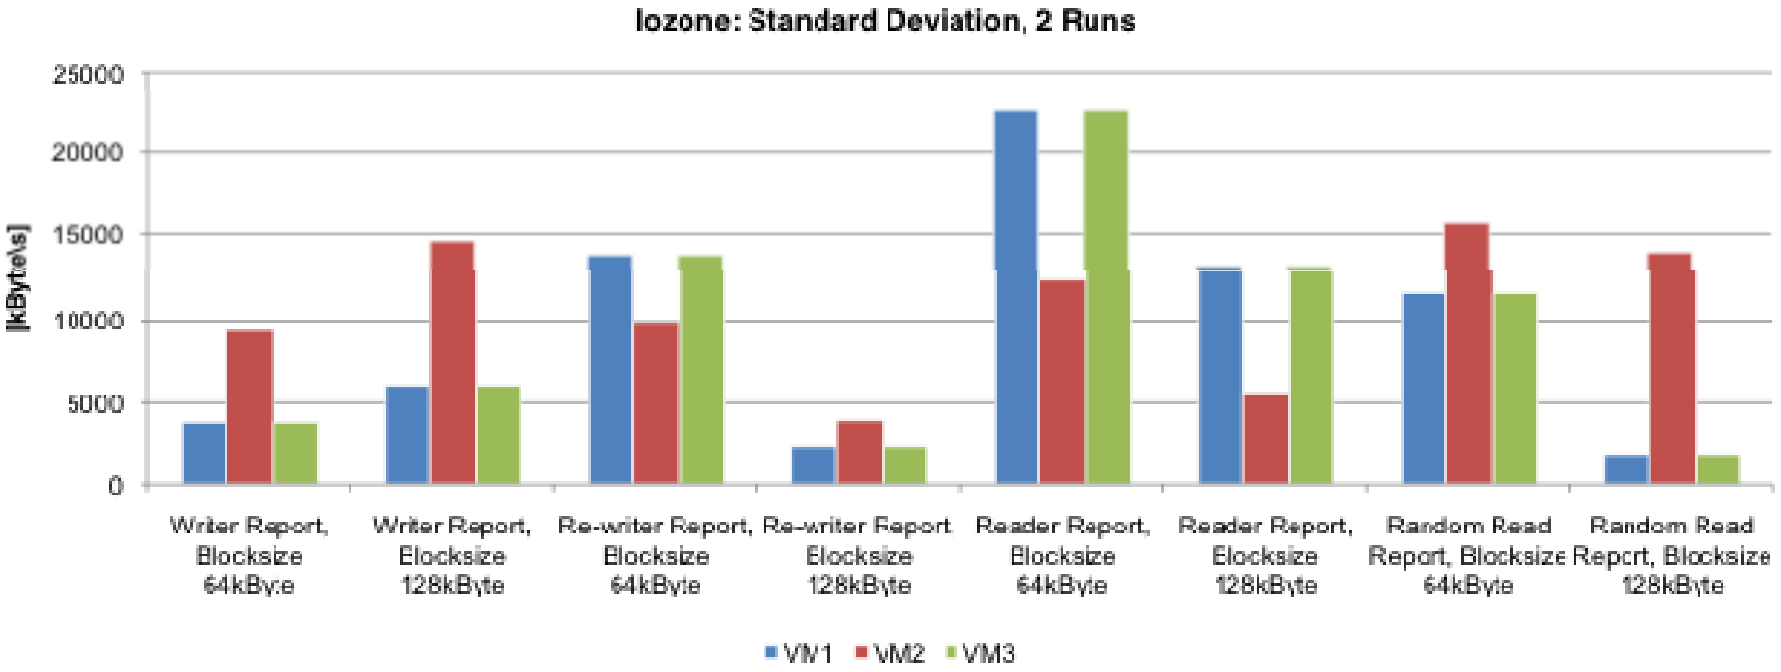

DVS

Accumulated Throughput Exceeds Throughput of Single VM

TECHNISCHE UNIVERSITÄT DARMSTADT

- Accumulated throughput (r/w) exceeds single VM throughput
  - Write: Effect small, but can be measured
  - Read: Effect huge, throughput doubled
- Possible explanations:
  - Caching effects, serialization of writes

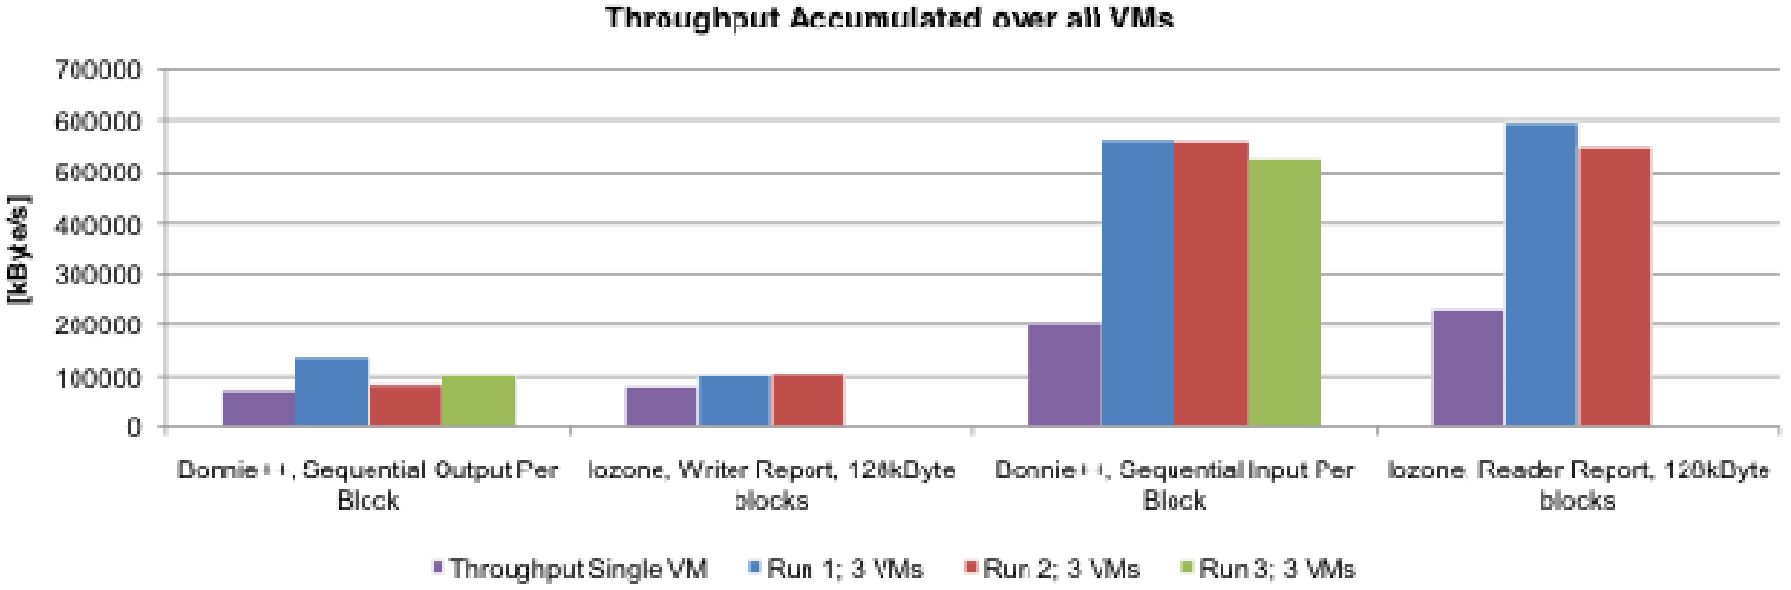

DVS

Summary:
CPU ok, Mem ok, IO depends

TECHNISCHE UNIVERSITÄT DARMSTADT

- CPU sharing works
  - SPECjvm2008 in 1-7 VMs in parallel
- Memory Bandwidth sharing works
  - RAMSPEED/RAMSMP in 1-7 VMs in parallel
- Disk I/O
  - Iozone and Bonnie++ in 1,3 and 5 VMs in parallel
  - Bandwidth shared fairly **on average**
  - But differences between VMs for single runs
  - Accumulated throughput exceeds single VM throughput

DVS

Thank You for Your Attention!

TECHNISCHE UNIVERSITÄT DARMSTADT

- Questions?
- Comments?

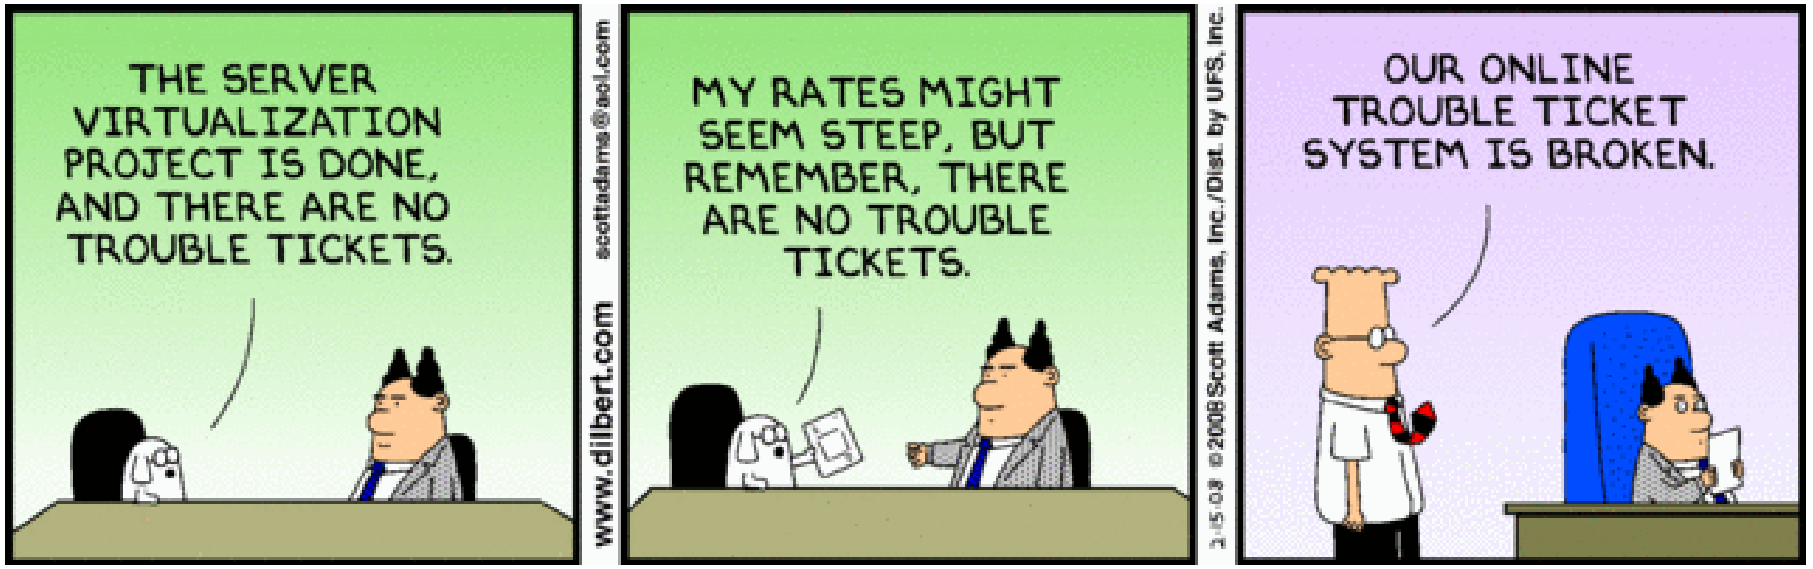

DVS