

*Power-efficient Server-class Performance
from
Arrays of Laptop Disks*


Athanasios E. Papathanasiou
and
Michael L. Scott
University of Rochester

USENIX 2004
Work in Progress
June 29, 2004

Motivation

- **Data centers are increasingly important**
 - 1995: **20 thousand** servers in the world
 - June 2001: **six million** servers
- **Disks account for 27% of total electric consumption on average**
 - High-end configurations much worse:
 - ✓ Disks can consume 15x more power than processors
- **Mobile disk power management is **not** effective for servers**
 - Server disks require **long** idle intervals to save power
 - Server workloads are characterized by **short** idle periods

Adaptive Throughput

- **Previous work suggests multi-speed disks [DRPM'03]**
 - Take advantage of load variation in server workloads
 - **BUT:** DRPM not commercially available yet
- **Alternative design: Mirrored Arrays of Laptop Disks**
 - Keep activated subset of disks proportional to current workload.
 - Exploit latency tolerance and parallelism to
 - ✓ Achieve energy savings
 - ✓ With equal or better throughput
 - ✓ Mirrored disks  Longer idle intervals for each disk

Disk Arrays: Why Laptop-class Disks?

Similar capacities, caches, prices

• **Server-class disks:**

- Lower rotational latencies (15KRPM)
- Lower seek times (4ms)
 - ✓ 2.5 times better
- Larger form factor (3.5 inches)
- Power hungry:
 - ✓ more than 12 W even when idle.
- No or few low-power modes

• **Laptop disks:**

- Higher Rotational Latencies (7KRPM)
- Higher seek times (10ms)
- Smaller form factor (2.5 inch)
- Consume less than 3W @ Active
 - ✓ 6 times more efficient
- Several low power modes (5)

Current technological design point: Replace at a ratio of **one to three**

Advantages and Disadvantages of a Mobile Disk Array

- **Disadvantages**

- Higher initial cost
- Higher latency for individual requests at light workloads

- **Advantages**

- Significantly lower operational power
- Lower cooling needs
- Potentially denser packaging
- Lower noise
- Potentially higher peak bandwidth
- Potentially higher mean time to data loss (mirroring)
- Faster development curve of laptop disks

Research Challenges

- **Disk selection policy**
 - Maximize idle intervals rather than per-disk throughput
- **Handling of write activity**
 - Writes must be performed on all copies
- **Power management policy**
 - Minimize effect of high latency mode transitions
- **Two-dimensional disk arrays**
 - Interaction between mirroring and file striping

Current Status

- **Detailed multi-disk laptop array simulator**
 - Based on Disksim
- **Lab-bench system for high precision power measurements**
- **Preliminary results (based on simulations):**
 - Baseline savings: **50%**
 - During light workloads: Savings up to **80%**

www.cs.rochester.edu/~papathan/research/BurstyFS/

*Power-efficient Server-class Performance
from
Arrays of Laptop Disks*



Athanasios E. Papathanasiou
and
Michael L. Scott
University of Rochester