Discussion Outcomes

Storage Group 1

Issues and Challenges (1)

- Available data volume increases at an alarming rate. The storage farm becomes bigger and bigger. However, the daily data access rate increases at a slower pace. This creates a potential for energy efficient storage systems.
- Tradeoffs between energy, performance and reliability
- Energy-aware and energy-efficient B-tree and other indexing schemes for supporting data intensive and database applications

Issues and Challenges (2)

- How to integrate flash-memory based Solid State Drives (SSDs) into existing storage hierarchy for energy saving
- Prediction and pre-fetching data schemes for energy saving
- Energy efficient algorithms for updates (reads and writes) adapted to workloads
- Energy efficient schemes for moving large volume of data from one place to another via Internet (data dedupe)

Power consumption

	Approximate Power Consumption	mW/GB
DRAM DIMM Module (1GB)	5W	5000
15K RPM Drive (300GB)	17.2W	57.33
7.2K RPM Drive (750 GB)	12.6W	16.8
High Performance SSD (128 GB)	2W	15.6

Source: Flash Storage Today, ADAM Leventhal, ACM Queue, 2008

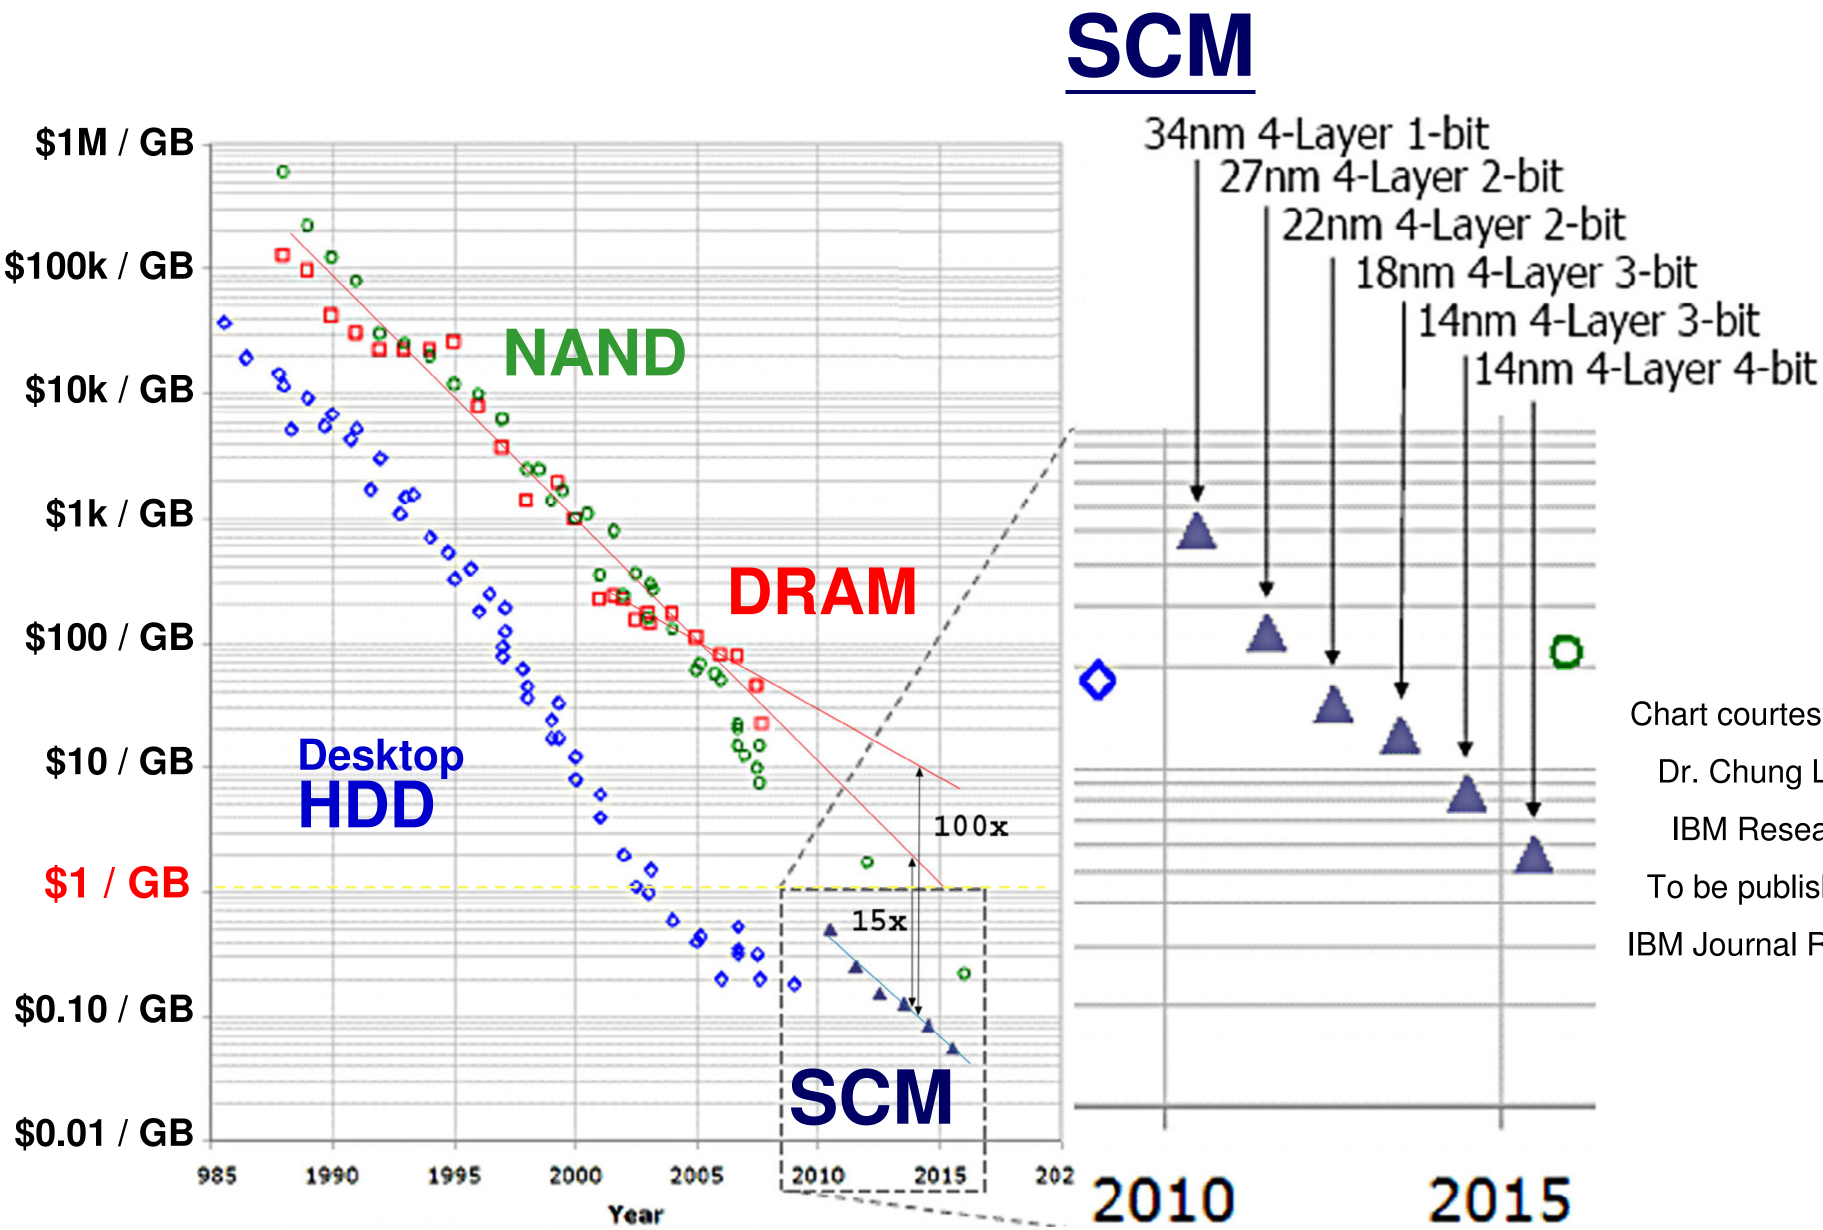

Chart courtesy of

Dr. Chung Lam

IBM Research

To be published

IBM Journal R&D

Issues and Challenges (3)

- Adaptive algorithms that iteratively decide on what parts of the input are realized and what parts of the input are influential in computing the answer (loopback control)
- What is the effect of distributed optimization on the task of the previous item

Science of Storage Power Management (1)

- The model for energy cost of data life cycle
- Energy efficient tradeoffs between storing multiple copies of data, re-compute data and deliver data via networks.
- Energy efficient measurements and metrics of distributed data storage
- Heterogeneous vs. Homogenous (Global vs. Local)

Science of Storage Power Management (2)

- In sensor networks or ubiquitous computing, how to model data storage, access and transmission for energy saving

Storage Systems (S2)

1. Device Technology

1. Impact of Solid State Nonvolatile Memory (NAND/NOR/PCM Flash) on storage architectures?

- Appropriate role of these devices in the memory hierarchy?
  - As a Faster Disk?
    - Large Disk/Buffer Cache controlled by OS Software: Store prefetched or hot data
  - As Larger Main Memory?
    - Processor addressable, DRAM acts as cache for Flash

2. Will disks become obsolete? Have outlived past predictions of imminent demise.

- Challenges:
  - Variability among devices due to proprietary FTL implementations
  - Lack of standard interfaces
  - Reliability of Multi-level Cell Flash
  - Require Open standard interfaces to permit Operating System intervention and optimization

Storage Systems (S2)

- Challenges (contd)

  - Lack of Multiple Power States in commodity disks poses challenges for DPM
  - High ON/OFF latency makes fine-grained power mode switching impractical
  - Reliability concerns with frequent power cycling
  - Extreme sensitivity of performance/energy on workload locality

Research Issues

Operating System models/mechanisms/data structures/algorithms to optimize power in the presence of these constraints.

Energy Aware

- File Systems: Organization, Meta Data
- Storage organizations (Declustering for Performance vs Clustering for Power Efficiency)
- Active Data Placement (for example dynamic reorganization)
- Fundamentally new Data structures for Energy Efficient Storage and Access
- Workload Intervention and Redirection
- Simple predictive power and performance models that capture workload locality

Storage Systems (S2)

2. Abstract Models

Incorporate Energy/Power as a first-class citizen of QoS specification

What are the appropriate Energy/Performance metrics in a storage system?

Models that incorporate energy in pricing Service Level Agreements?

Incentives to clients to provide energy-friendly workloads

Measures of goodness for external (storage resident) algorithms that reflect energy efficiency.

Challenges:

- Many storage workloads tend to be high-variance (bursty) with stringent response time requirements
- Peak requirements many times long-term average rate
- Servers/disks not amenable to fine-grained power mode switching to respond to bursts

Requires:

- Fundamental understanding of the relationship between application (algorithm) dynamics and power consumption
- Fundamental understanding of workload characteristics
- Power efficient algorithms/provisioning/scheduling/DPM mechanisms based on these models
- How detailed do the power models need to be to yield useful information?
- How do you compactly model and articulate an energy-friendly workload?

Storage Systems (S2)

3. Leveraging Information Redundancy to achieve desired Energy and Performance tradeoffs

- Exploit available redundancy (replication/erasure coding) for power benefits
- Incorporate geographical replication of data in models
- Models for Energy Aware Workload distribution on a global scale
- Decentralized/distributed energy QoS schedulers

4. Mechanisms, Algorithms, Analysis Techniques:

Performance, energy tradeoffs achieved by different techniques

How does statistical multiplexing apply to storage workloads?

Are models developed in the context of networking systems adequate? What are the alternatives?

5. Science of data storage and transfer: Are there fundamental relationships between energy requirements and the codes used to to store and retrieve data? How do our coding techniques for storage reliability stack up?