

Aggressive Prefetching An Idea Whose Time Has Come

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

HotOS X
June 13, 2005

Prefetching, Prepaging or Read-ahead

- Standard practice for more than 30 years
- Complements caching policy (LRU)
- Goal: Reduce latency of access to non-cached data
 - Predict future accesses
 - Read-ahead predicted data
- A traditional, well-established and well-studied field
 - Very rich literature

THP

University of Rochester

2

Far too rich...

- Prediction
 - History-based predictors
 - Application-disclosed hints
 - Compiler-assisted
 - Speculative execution
- Handling prefetched pages
 - Models to estimate benefit of prefetching and cost of evictions
 - Specialized policies for prefetched pages
 - Application-controlled
- Data compression

Cao et al. - OSDI'94	Kistler et al. - TOCS'92	Papathanasiou et al. - Usenix'04
Cao et al. - Sigmetrics'95	Kotz et al. - PDIS'91	Papathanasiou et al. - WMCSA'03
Chang et al. - OSDI'99	Kroeger et al. - Usenix'01	Patterson et al. - SOSP'95
Carewicz et al. - Sigmod'93	Kuenning et al. - MSSWM'02	Saito - IEEE Transactions on Computers 50
Fertag et al. - SOSP'71	Kuenning et al. - SOSP'97	Shen et al. - IPDPS'04
Griffioen et al. - PDCS'95	Li et al. - OASIS'04	Steele - SOSP'97
Griffioen et al. - Usenix'94	McKinick et al. - TOCS'84	Tai et al. - ICDCS'91
Kaplan et al. - ISMM'02	Mowry et al. - OSDI'96	Tonkins et al. - Sigmetrics'97
Kimbel et al. - OSDI'96	Pahner et al. - VLDB'91	Yeh et al. - ISPASS'91
		Yeh et al. - WPMRTES'01

THP

University of Rochester

3

Common Debate on Prefetching Conservative or Aggressive

- Aggressiveness aspects
 - Timing: How early should prefetching occur?
 - Data coverage: How speculative should it be?
- Conservative: Reasonable
 - Fetch data incrementally when confidence is high
 - Variants of sequential read-ahead widely adopted
- Aggressive
 - Prefetch deeper and earlier than necessary
 - Speculate more to increase data coverage
 - Potential to improve I/O performance significantly
 - ✓ Reduce demand-misses (increased data coverage)
 - ✓ Avoid increased delays due to irregular device response times
 - BUT...

THP

University of Rochester

4

Aggressive Prefetching Traditional Concerns

- Risk of severe negative impact on performance
 - Buffer cache pollution
 - ✓ Prefetch unnecessary and eject useful data
 - Increased physical memory pressure
 - ✓ May even lead to thrashing
 - Inefficient use of I/O bandwidth
 - ✓ Speculation may waste bandwidth on unnecessary data
 - Increased device congestion
 - ✓ Increased number of asynchronous requests in I/O queues
- Solutions exist, but
 - Introduce significant implementation complexity
 - Problems more severe on resource-limited systems
 - Forgo potential benefits at high-end to avoid problems at low-end

THP

University of Rochester

5

What has changed?

- Technological and market trends
 - Dramatic improvements in:
 - ✓ Processing power, storage capacity, I/O bandwidth (lesser)
 - Modest improvements in I/O latency
 - Decrease risks and increase need for aggressive prefetching
- Emerging usage patterns and design goals
 - Multitasking, larger data sizes, multimedia applications
 - Power efficiency, weakly-connected operation
 - Increase value of aggressive prefetching

THP

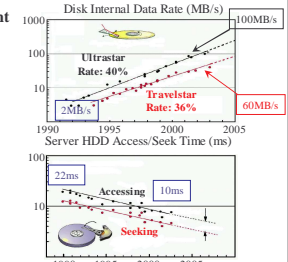
University of Rochester

6

Technological and Market Trends

Magnetic Disk Performance

- **Annual performance improvement**
 - Bandwidth: 40%
 - Latency: 10% only
 - Gap will increase
- **To balance disk activity**
 - Read more on disk context switch
 - Disks speculatively read whole tracks
- **Aggressive prefetching**
 - Exploits bandwidth and caches
 - Improves disk utilization



Source: Hitachi Global Storage Tech: HDD Technology '03


THP

University of Rochester

8

I/O Performance Variability

- **Constant bandwidths and latencies**
 - Assumption no longer holds
- **Significant variability**
 - Several I/O devices
 - Variation within a single device
 - ✓ Up to 47% improvement for server-class applications
- **Aggressive prefetching**
 - Exploit periods of higher bandwidth
 - Hide periods of higher latency
 - Smooth out performance fluctuations



Source: Hitachi Global Storage Tech: Adaptive Formatting

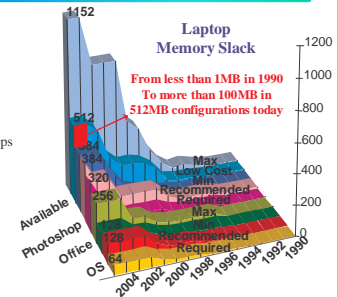
THP

University of Rochester

9

Large Memory Size at Low Cost

- **Production**
 - 70% annual increase
- **Price**
 - 32% annual decrease
 - Today: 12 cents per MB
- **Memory configurations**
 - 512 MB common for laptops
 - Several GBs for desktops
- **Increased memory slack**
 - Prefetch aggressively
 - Lower risk of
 - ✓ Pollution
 - ✓ Memory Pressure



THP

University of Rochester

10

Usage Patterns and Design Goals

Larger Data Sets, Less Reuse

- **Significant digital content increase**
 - 114% increase of hard disk content production from 1999 to 2003
- **Proliferation of multimedia content**
 - Each datum touched only once
 - Multimedia application exhibit predictable patterns
- **Large data sizes diminish benefit of larger memories**
- **Importance of prefetching will increase**

THP

University of Rochester

12

Processor I/O Gap and Multitasking

- **Processing power far ahead of I/O latencies**
 - Prefetching has to work multiple blocks ahead
- **Applications**
 - Alternating compute and I/O intensive phases
 - Prefetch to hide latency of **last** I/O access
 - ✓ Prefetch just-in-time for next access is not enough
 - ✓ Prefetch earlier and deeper
- **End-users work concurrently on multiple tasks**
 - Switching among several data hungry applications
 - Reduces effectiveness of LRU-style caching
- **Aggressive prefetching on task switch**

THP

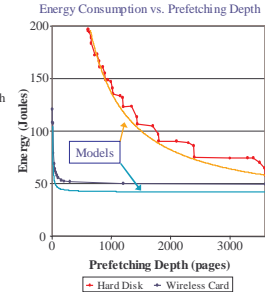
University of Rochester

13

Energy and Power Efficiency

• Prefetching impact on energy consumption

- **Bursty, speculative prefetching**
 - ✓ Benefits non-operational modes
 - 72% of disk energy savings with just 20% prefetching accuracy
- **Smooth prefetching**
 - ✓ Benefits operational modes (multi-speed disks)
 - Tradeoff??
- **Benefit varies with device**
 - Modeling prefetching benefit
 - ✓ Depth, accuracy, speculation
 - For informed decisions



THP

University of Rochester

14

Reliability and Availability of Data

- **Mobile systems**
 - Disconnected, weakly-connected operation
 - Portable storage devices
 - ✓ New form of weak connection to data
- **Traditional topics of file system research**
 - **Problem:** Multiple (legacy) file systems
 - ✓ Do not handle disconnection well
- **Hoarding based on aggressive prefetching**
 - Implement at memory management/virtual FS layers
 - Benefits arbitrary underlying file systems

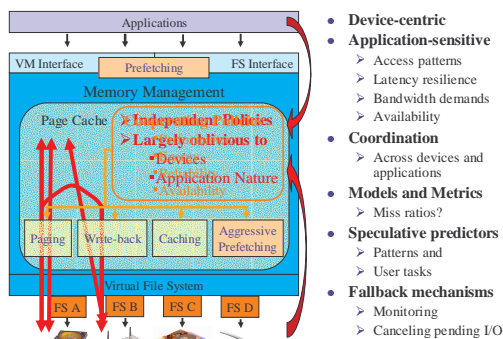
THP

University of Rochester

15

Design Challenges

Design Challenges for Prefetching



- **Device-centric**
- **Application-sensitive**
 - Access patterns
 - Latency resilience
 - Bandwidth demands
 - Availability
- **Coordination**
 - Across devices and applications
- **Models and Metrics**
 - Miss ratios?
- **Speculative predictors**
 - Patterns and
 - User tasks
- **Fallback mechanisms**
 - Monitoring
 - Canceling pending I/O

Conclusions

- **Memory management of today**
 - Caching secondary storage
 - Oblivious to
 - ✓ Nature of applications
 - ✓ Devices beneath
- **Storage management of the future**
 - Application sensitive
 - ✓ Track and predict behavior
 - ✓ Monitor and coordinate I/O needs
 - Device sensitive
 - ✓ Model latency, bandwidth, availability and power consumption
 - ✓ Move data
 - Between memory and devices, and also
 - Among I/O devices

THP

University of Rochester

18

Aggressive Prefetching
An Idea Whose Time Has Come

Q&A

<http://www.cs.rochester.edu/~pathan/research/Athena/>

HotOS X
June 13, 2005