Report on the Second International Workshop on Data Management on Modern Hardware (DaMoN'06)

Anastassia Ailamaki
Carnegie Mellon University
natassa@cmu.edu

Peter Boncz
CWI
boncz@cwi.nl

Stefan Manegold
CWI
manegold@cwi.nl

Abstract: This report summarizes the presentations and discussions that occurred during the Second International Workshop on Data Management on Modern Hardware (DaMoN). DaMoN was held in Chicago on June 25th, 2006, and was collocated with ACM SIGMOD 2006. The aim of this one-day workshop is to bring together researchers interested in optimizing database performance on modern computing infrastructure by designing new data management techniques and tools.

1 Motivation and Goal

The continued evolution of computing hardware and infrastructure imposes new challenges and bottlenecks to program performance. As a result, traditional database architectures that focus solely on I/O optimizations increasingly fail to utilize hardware resources efficiently. CPUs with superscalar out-of-order execution, simultaneous multi-threading, multi-level memory hierarchies, and future storage hardware (such as e.g. MEMS) impose a great challenge to optimizing database performance. Consequently, exploiting the characteristics of modern hardware has become an important topic of database systems research.

Architecture-conscious database research aims to make database systems adapt automatically to the sophisticated hardware characteristics, thus maximizing performance transparently to applications. To achieve this goal, the data management community needs interdisciplinary collaboration with computer architecture researchers. It also involves rethinking traditional data structures, query processing algorithms, and database software architectures to adapt to the advances in the underlying hardware infrastructure.

The International Workshop on Data Management on Modern Hardware (DaMoN) aims to achieve this goal by attracting people from computer architecture towards the principle data management venue ACM SIGMOD/PODS, both by providing a publication platform as well as through direct invitation, for the keynote talk and panel discussions.

2 Logistics and Statistics

The workshop was held on Sunday June 25th, the day before SIGMOD/PODS started. We had a program committee of 8 expert members from both areas of computer architecture and databases. There were eleven papers submitted to the workshop, which were professionally reviewed by the program committee (at least 3 reviews per submission). This number of submissions only just coincided with our minimum requirements, and may have been caused by overlap with other workshops (in particular, ExpDB) or by the unavoidable proximity of the submission deadline to that of VLDB. However, the submission quality was high, as measured by the absolute scores by the referees. With DaMoN in its second year, we achieved a 50% acceptance ratio.

Despite two round-of-sixteen world cup soccer matches, around 35 people attended DaMoN, making it one of the two best attended SIGMOD workshops (only WebDB had more attendees). This may also have to do with our keynote speaker (Berni Schiefer, IBM fellow) and a panel that included Mark Hill (professor of computer architecture from University of Wisconsin), Goetz Graefe (Microsoft CLR) and Todd Walter (CTO of NCR/Teradata), as well as Anastassia Ailamaki (database faculty at CMU). A striking characteristic considering attendance is that a considerable number of visitors came expressly to Chicago for DaMoN – which we interpret as testament to the relevance of the workshop for a community. Off the program committee Shimin Chen, Goetz Graefe and Bradley Kuszmaul attended the workshop. Logistically, everything went well in Chicago, including a nice lunch for a reasonable price. We specifically thank Kevin Chang, Joanne Martiori, Goce Trajcevski, and Lisa Singh for their help with the organization.

3 Technical Presentations

The six paper presentations were distributed into two sessions. In the first session, a research team by IBM described a solution for data mining on multiple information sources pertaining to different organizations, while respecting data privacy and without disclosing information to each organization. This was achieved by encrypted data transport and running the data mining algorithms inside a sealed tamper-protected hardware subsystem, that runs on an embedded processor. Scaling down data mining algorithms to fit the reduced resources available in such an environment was the main topic of this work. The second presentation stayed on the data mining topic, investigating the use of Processing-In-

Memory chips for various link-discovery algorithms, with varying degrees of success.

Last in this session was an architecture-conscious study of a large transaction benchmark set-up. This cooperation between Intel and Microsoft industry labs provides in-depth experimental data of running SQLserver on a parallel Itanium2 machine. This paper provides the best reference material of various hardware event traces (cache misses, branch misprediction, TLB handling) published so far on a realistic high-performance transaction processing scenario. Due to the expensive hardware setup (32 CPUs, 256 GB RAM, 1260 disks), such experiments are generally out of reach for academic researchers, but insights from these measurements can help set their agendas. As organizers, we highly value papers like these.

In the second session, a combination of researchers from Sandia Labs and Columbia University experimented with the massive multi-processors reality that awaits the database community within a decade. Rather than relying on simulation, they used an existing Cray MTA-2 super computer to test a variety of processing strategies for join and selection. This well-written and informative paper was awarded the Best Paper Award, collected by John Cieslewicz in behalf of fellow authors Ken Ross, Jonathan Berry and Bruce Hendrickson.

The remaining two presentations focused on architecture-conscious index structures. In the first, Goetz Graefe addressed the problem that in B-tree search with wide nodes, binary search inside each node will cause significant cache misses. Interpolation search can reduce their number, but is non-trivial to make robust against different data distributions. His paper presents a comprehensive overview of techniques and strategies for robust use of interpolation search in B-trees. The final presentation described the use of the new Cuckoo Hash technique in databases. Like a perfect hash function, Cuckoo obviates the need for a collision list, and thus reduces loop dependencies during hash-lookup, without requiring in-advance knowledge of the data distribution. This last paper also describes an adaptive buffering strategy in partitioned hashing, that reduces the RAM requirements of operators like hash aggregation and join, but conserves their CPU cache efficiency.

4 Panel Discussion

The concluding panel "ManyCore-DB: will you still need me, will you still feed me, when I'm 64?" discussed the influences of the computer architecture trend towards massively (e.g. 64-way) multi-core CPUs on data management.. The panel, moderated like last year by Peter Boncz, once again turned out to be very lively.

The first speaker was Mark Hill, who provided an excellent overview from the computer architect's view of the near-future parallel reality, underlining the renewed relevance that the topic of parallelism should receive in the next years in our research community. Next up was Goetz Graefe, who is currently working on support for parallelism in the Common Language Runtime (CLR) at Microsoft. CLR has been in the database picture for the recent developments around LINQ, that bring certain database processing models right into programming language environments. Goetz was, however, tight-lipped regarding any as-yet-unannounced Microsoft product features. As the sole representative of the database academia, Natassa Ailamaki discussed the challenges of massive CPU parallelism and made a case for staged database architecture to address these. Last up was Todd Walter, the CTO of NCR/Teradata. He provided many insights, predicting among others the demise of multiple-socket architectures (SMP) once massively multi-core CPUs become the norm. But, he pointed out that in computer architecture, the main bottleneck is I/O and the continued exponential growth in computational power provided by parallelism will not be matched by similar advances in magnetic storage. The panelists, however, discarded any known storage mechanism (MEMS, Flash) as a potential successor to magnetic disks.

The overall conclusion was that massive multi-core parallellism will create massive under-utilization of processing power very soon. This outcome we view as highly supportive for the future relevance of architecture-conscious data management research.

5 Where do we go from here?

This second edition of DaMoN has shown that the workshop can consistently draw 30-40 attendees, with a considerable number of submissions and visitors coming from outside the database field. The event also draws significant attention from industry both in submissions and attendees, and this year was generously sponsored by Intel Research. The quality of the papers presented has been good, exemplified by a previous DaMoN paper in extended form winning the Best Paper at VLDB 2005. As a result, the organizers plan to continue DaMoN also at SIGMOD 2007 in Beijing.