

Web Data Management

Charlie Garrod
25 Apr 2005


12 years ago...

- 1 Al Gore had just invented the internet
- 1 A (relatively) small number of users put content on the web
- 1 And a (relatively) small number of users downloaded it

Most content was simple!


8 years ago... web caches

- 1 Much larger number of users
- 1 Most content was still simple and static



5 years ago... CDNs


- 1 Content Distribution Networks
- 1 Move web content to the “edge”




Today...

- 1 Web content is complex and dynamic
 - interactive and personalized
- 1 Amazon, CNN, Google, USAir, LiveJournal, and of course the 15-721 course homepage...

Dynamic content generation



Dynamic content generation



Web database workloads

- 1 Most queries are small and simple (OLTP)
 - Show me the last 25 journal entries by "puuj"
 - Show me non-full flights to LAX next Friday
 - Find all websites about fire-breathing space monkeys
- 1 Few updates
- 1 Other than that, workloads vary greatly between applications

Web database workloads

- 1 Queries and updates are often instantiations of more general templates
 - Q1: SELECT id FROM users WHERE age > ?
 - U2: UPDATE users SET age = ? WHERE id = ?

The \$65,536 question:

How do we make dynamic content scalable?

Web Data Management Outline

- 1 Introduction
- 1 Overview of common approaches
- 1 WebView Materialization
- 1 DBProxy: A dynamic data cache for Web applications
- 1 Conclusions

Web Data Management Outline

- 1 Introduction
- 1 **Overview of common approaches**
- 1 WebView Materialization
- 1 DBProxy: A dynamic data cache for Web applications
- 1 Conclusions

Solution #1: WebView Materialization

Make the content “static”


Solution #1: WebView Materialization

- 1 Generate new static version of webpage every time it is updated
- 1 Works great for CNN, Slashdot, etc. where the content is semi-static
 - Does not adapt well to personalized or interactive websites

Solution #2

Build a custom solution

Solution #2: Big DBMS™



Solution #2: Big DBMS™

- 1 Build a custom, semi-centralized DBMS system
- 1 Good for big companies such as Google, Amazon, EBay, etc. with an established user base and significant market investment
- 1 Very expensive to implement!

Solution #3

Try something else!

Solution #3: Dynamic CDN

- 1 Try to apply the principles of caching and content-distribution to dynamic web pages
 - Build a nice, general solution to scale dynamic workloads
 - Adaptable to personalization and interaction
 - Cheaper than a custom, specialized solution

Solution #3: Dynamic CDN

- 1 Try to apply the principles of caching and content-distribution to dynamic web pages
 - Build a nice, general solution to scale dynamic workloads
 - Adaptable to personalization and interaction
 - Cheaper than a custom, specialized solution

This is easier said than done!

Web database workloads revisited

- 1 Most queries are small and simple (OLTP)
 - Show me the last 25 journal entries by “puuj”
 - Show me non-full flights to LAX next Friday
 - Find all websites about fire-breathing space monkeys
- 1 Few updates
- 1 Other than that, workloads vary greatly between applications


Distributing dynamic content

- 1 Which server components should we distribute?
 - Everything?
 - Just the web server and application server?
 - Partially replicate the database?

Distribute everything!

- 1 All proxy servers contain a web server, app server, and database
- 1 The perfect solution for scaling queries!
- 1 Updates are practically impossible
 - Distributed databases are fundamentally hard to build and are usually intended only for LANs


Distribute everything!



Distribute the web and app server!

- 1 Efficiently off-loads the web server and application execution to remote proxy servers
 - Reduces bandwidth usage
- 1 Still relies on a centralized database
- 1 Interactions with the database become high-latency

Distribute the web and app server!




And finally...

Partial Replication (and Caching)

- 1 Distributes web and app server load as before
 - Reduces bandwidth, etc.
- 1 Updates are potentially less expensive than with full replication
 - But still non-trivial

Partial Replication (and Caching)



Intermission

Web Data Management Outline

- 1 Introduction
- 1 Overview of common approaches
- 1 **WebView Materialization**
- 1 DBProxy: A dynamic data cache for Web applications
- 1 Conclusions

WebView Materialization

- 1 Strategy #1: make the content static
- 1 Labrinidis and Roussopoulos, University of Maryland, circa 2000.
- 1 Introduced a formal cost model for evaluating materialization of "WebView" at the web server, within the DBMS, or not at all
- 1 Experimentally evaluated the different strategies

Strategy #1: "Virtual" materialization

- 1 Query is re-executed at database and webpage is regenerated
- 1 Updates are cheap since only the "standard" update must be executed at the DBMS
- 1 Queries are expensive since all work must be re-done every time

Strategy #2: Materialization at DBMS

- 1 The query result is saved at the database, but the resultant webpage itself is regenerated
- 1 Updates are more expensive since the materialized view at the DBMS must be regenerated as well
- 1 Queries are slightly cheaper since only the webpage must be regenerated


Strategy #3: Materialization at web server

- 1 The full materialized webpage is stored at the web server
- 1 Updates are very expensive, essentially the cost of a standard update plus a query plus the cost of generating the resultant webpage
- 1 Queries are very cheap since the page is just retrieved as if it were static content

Experimental Methodology

- 1 Used a single Sun system as a server (running Apache and Informix), 22 Sun systems as clients, all within a single LAN
- 1 Measured query response time for each strategy for various access rates, update rates, number and size of views, and view selectivity

Results, yada, yada



Problems with their methodology

- Relatively small number of views (100-2000)
- Results are indicative of an open system under low load
 - For “materialization at web server” updates are executed as a separate background process
 - Only query response time is measured
 - Cheaters!

WebView Materialization Conclusions

- Still show that materialization at web server can effectively reduce overall load for a relatively small number of views, which can greatly improve performance for some loads
- Somewhat surprising that materialization at DBMS often hurts!
- A nice mix of theoretical and experimental methodology!

Web Data Management Outline

- 1 Introduction
- 1 Overview of common approaches
- 1 WebView Materialization
- 1 **DBProxy: A dynamic data cache for Web applications**
- 1 Conclusions


DBProxy: A dynamic data cache...

- 1 Amiri et al., IBM T.J. Watson, circa 2002.
- 1 Based on partial replication
 - Queries are processed locally at a proxy server if possible
 - All updates forwarded to a central database, which periodically propagates the updates to the proxy servers


Overall goals

- 1 Database independence
 - Any back-end database could be used
- 1 Self-management
 - Cache dynamically adapts to a changing workload without administrator intervention
- 1 Consistency
 - Must be efficient even with a large cache and heavy update traffic

DBProxy architecture



DBProxy JDBC driver architecture



DBProxy local database

- 1 Stores subsets of tables from the central database (both horizontally and vertically partitioned)
- 1 And catalog information from the central database...
- 1 And information about the queries that are currently cached...

DBProxy query matching

- 1 Uses the SELECT and WHERE clauses to determine if the query is a subset of the union of queries already in the cache
- 1 Can potentially answer queries that have not yet been issued before
 - Q1: SELECT id FROM users WHERE age < 25
 - Q2: SELECT id FROM users WHERE age > 18
 - Q3: SELECT id FROM users WHERE age > 21

DBProxy update mechanism

- 1 All updates are forwarded to the central database
- 1 All proxies subscribe to a stream which contains all updates at the database
 - Not just the updates they care about

DBProxy consistency guarantees

- 1 Lag consistency
 - The proxy server is not too outdated
- 1 Monotonic state transitions
 - The view of the database at the proxy moves only forward with time
- 1 Immediate visibility of updates
 - An application observes the effects of its own updates

DBProxy consistency guarantees

- 1 Lag consistency
 - The proxy server is not too outdated
- 1 Monotonic state transitions


No transactional consistency! only

- 1 Immediate visibility of updates
 - An application observes the effects of its own updates

DBProxy cache replacement

- 1 Runs as a background process, garbage collecting results that are not used by any cached query and occasionally evicting cached queries to reclaim space
 - General replacement algorithm, taking into account recency and frequency of use, space used, miss cost vs. hit cost, etc.


Experimental methodology



Experimental methodology

- 1 Modified TPC-W (which simulates a simple web bookstore workload) to introduce some additional complexity
- 1 Measured proxy response time and hit rate with several database sizes, several cache configurations, and various loads on the back-end database
- 1 Started with a warm cache

Proxy response time



Proxy cache hit rates

Query Category	Baseline TPC-W			Modified TPC-W		
	Response time	Hit rate	Query frequency	Response time	Hit rate	Query frequency
Simple	51	91 %	23 %	317	37 %	47 %
Top-N	935	68 %	12 %	852	66 %	37 %
Exact-match	211	76 %	65 %	458	54 %	15 %
Total	263	73 %	100 %	540	50 %	100 %
No Cache	385	—	100 %	1024	—	100 %

Using 100K database, 80K users

Problems with their methodology

- 1 No comparison to centralized-only configuration
- 1 No mention of throughput, an important performance metric
- 1 Used TPC-W browsing mix only
 - Did not measure the effect of various update loads on the system

Harsh (and slightly unfair) DBProxy conclusions

- 1 Great cache configuration and query-matching
- 1 Poor update-handling and consistency management
- 1 While initially impressive, performance **results do not support the use of DBProxy** compared to a centralized architecture or for any workloads with a non-trivial update component

Web Data Management Outline

- 1 Introduction
- 1 Overview of common approaches
- 1 WebView Materialization
- 1 DBProxy: A dynamic data cache for Web applications
- 1 **Conclusions**

Many similar projects

- 1 DBCache (IBM Almaden), DBProxy (IBM Watson), GlobeDB (ETH Zurich)
- 1 Similar projects that focus on file system workloads (UT Austin)
- 1 And...

Shameless plug: S-3 (CMU)

- 1 Ailamaki, Garrod, Maggs, Manjhi, Mowry, and Olston (among others)
- 1 Efficient transactional consistency
- 1 Theoretical framework for the effect of data secrecy on scalability
- 1 Exploiting knowledge of query and update templates

Conclusions

- Overall, this is still very much an area of on-going research!
- Lots of people working on this problem, and nobody yet has come up with a satisfactory solution
- And it's really a \$64 billion question