

MySpace looks to USC to help servers keep up with users

June 27 2008

How do you keep a potential half-billion social networkers current? A Viterbi School of Engineering expert tackles the growth problems of one of the world's largest computing systems.

Imagine a moment - a thousandth of a second - in the life of a MySpace computer. In that brief fraction of an eye blink, thousands of fingers on mice, spread across thousands of square miles, have clicked-in urgent requests for as many chunks of data: a photo needed in Minneapolis, video needed in Des Moines, a forum comment wanted by an observer in Detroit.

As the vast social networking system grows, each millisecond becomes more and more crowded with requests like these. Now, a University of Southern California specialist is working to make sure that the answers keep coming back quickly, even with tens of millions of new users.

"If MySpace were less successful," notes Shahram Ghandeharizadeh, a professor in the USC Viterbi School of Engineering's department of computer science, "there would be no problem. But at the current volume of transactions, getting to the data quickly becomes an issue. "

The key to speed and capacity is what is called DRAM. "Ideally," says Ghandeharizadeh, who is director of the USC Database Laboratory, "you want all the data requested to be in the quick-access cache memory of the servers, the DRAM, rather than having to retrieve it from the servers' disc memory, which is much slower."

But the total volume of data created by users is far more than the DRAM cache will hold. And as the user population grows at an accelerating pace, more and more requests arrive to query a larger and larger body of information. Even the innovative Berkeley DataBase system (BDB) that MySpace uses to keep interactions quick is coming under increasing strain.

BDB keeps MySpace information flowing by quadruple redundancy: each section of users is served not by one but by four overlapping servers that share DRAM space, making the system faster, more reliable, and more scalable.

"Now it works," says Felipe Cariño, who heads MySpace Research, the company's in-house R&D facility, "but if you double it, it may not." And he says the population may in fact double as people in China and other areas learn to meet and greet each other on their own sites.

Ghandeharizadeh is working with Cariño, a 1995 USC Marshall School Executive M.B.A., to find a way around the impending squeeze. Cariño dubs the effort the "Gemini Project," after the famous twins: "Two heads, Viterbi and MySpace, coming together."

The collaborators have been exploring a new method for maintaining and replacing the data kept in DRAM.

Up to now, the method used has been simple: Data that has remained in the DRAM longest without being accessed is overwritten by new data.

Another method is potentially more effective: "Heuristic" replacement in which data is given simple but useful characterizations, which a program then uses to guide replacement. The program isn't static, but learns from system behavior and adjusts its criteria to improve performance.

The heuristic algorithm that MySpace Intrapreneurial Research Group is adapting to the MySpace database comes out of a recent PhD research thesis done by USC graduate student Shahin Shayandeh, who is a member of the team, along with three MySpace computer scientists.

The key element is taking file size into account as well as date. Large objects, like video files, kick out many small objects from the memory when they're loaded. While the general solution is based on how frequently objects are accessed, another rule is to not let very large objects into DRAM.

"When you have a gigantic video file that takes up the same space as hundreds or even thousands of text files, uploading it is a shock to the system," says Ghandeharizadeh.

Based on their research, Cariño and Ghandeharizadeh are hopeful that the new algorithm will adapt well to MySpace demands and deliver the desired improvements in performance.

"Simulation studies show the heuristic method is a marvel. But seeing whether it delivers in a real ultra-large system, such as the one at MySpace, remains to be seen," says Ghandeharizadeh.

Ghandeharizadeh cherishes the opportunity to take the software to limits not others have reached. He specializes in ultra-large systems, and only three or four exist in the entire world, of which MySpace's is one. It includes, says Cariño "10 data centers, thousands of servers, in 30 plus countries supporting local cultures and languages".

Cariño notes that MySpace Research mission is clear and tightly focused "so that the final result must be a system or prototype that creates a new product or technology". MySpace Research will file patents and publish papers on the pioneering innovations, but the final deliverable is a new

working system. Which Ghandeharizadeh hopes to deliver.

Source: University of Southern California

Citation: MySpace looks to USC to help servers keep up with users (2008, June 27) retrieved 28 April 2024 from <https://phys.org/news/2008-06-myspace-usc-servers-users.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.