

Scaling up personalized query results for next generation of search engines

September 18 2013, by Matt Shipman



NC State University researchers are developing ways to improve "personalized" search results on a large scale. Credit: North Carolina State University

North Carolina State University researchers have developed a way for search engines to provide users with more accurate, personalized search results. The challenge in the past has been how to scale this approach up so that it doesn't consume massive computer resources. Now the researchers have devised a technique for implementing personalized searches that is more than 100 times more efficient than previous

approaches.

At issue is how [search](#) engines handle complex or confusing queries. For example, if a user is searching for faculty members who do research on financial [informatics](#), that user wants a list of relevant webpages from faculty, not the pages of graduate students mentioning faculty or news stories that use those terms. That's a complex search.

"Similarly, when searches are ambiguous with multiple possible interpretations, traditional search engines use impersonal techniques. For example, if a user searches for the term 'jaguar speed,' the user could be looking for information on the Jaguar supercomputer, the jungle cat or the car," says Dr. Kemafor Anyanwu, an assistant professor of [computer science](#) at NC State and senior author of a paper on the research. "At any given time, the same person may want information on any of those things, so profiling the user isn't necessarily very helpful."

Anyanwu's team has come up with a way to address the personalized search problem by looking at a user's "ambient query context," meaning they look at a user's most recent searches to help interpret the current search. Specifically, they look beyond the words used in a search to associated concepts to determine the context of a search. So, if a user's previous search contained the word "conservation" it would be associated with concepts like "animals" or "wildlife" and even "zoos." Then, a subsequent search for "jaguar speed" would push results about the jungle cat higher up in the results – and not the automobile or supercomputer. And the more recently a concept has been associated with a search, the more weight it is given when ranking results of a new search.

Search engines have also tried to identify patterns in user clicking behavior on search results to identify the most probable user intent for a search. However, such techniques are impersonal and are applied on a

global basis. So, if the most frequent click pattern for a set of keywords is in a particular context, then that context becomes the context associated with queries for most or all users – even if your recent search history indicates that your query context is about jungle cats.

"What we are doing is different," Anyanwu says. "We are identifying the context of search terms for individual users in real time and using that to determine a user's intention for a specific query at a specific time. This allows us to deal more effectively with more complex searches than traditional search engines. Such searches are becoming more prevalent as people now use the Web as a key knowledge base supporting different types of tasks."

While Anyanwu and her team developed a context-aware personalized search technique over a year ago, the challenge has been how to scale this approach up. "Because running an ambient context program for every user would take an enormous amount of computing resources, and that is not feasible," Anyanwu says.

However, Anyanwu's research team has now come up with a technique that includes new ways to represent data, new ways to index that data so that it can be accessed efficiently, and a new computing architecture for organizing those indexes. The new technique makes a significant difference.

"Our new indexing and search computing architecture allows us to support personalized search for about 2,900 concurrent users using an 8GB machine, whereas an earlier approach supported only 17 concurrent users. This makes the concept more practical, and moves us closer to the next generation of search engines," Anyanwu says.

More information: The paper, "Personalizing Search: A Case for Scaling Concurrency in Multi-Tenant Semantic Web Search Systems,"

will be presented at the 2013 IEEE International Conference on Big Data being held Oct. 6-9 in Santa Clara, Calif.

Abstract: Recent keyword search techniques on Semantic Web are moving away from shallow, information retrieval-style approaches that merely find "keyword matches" towards more interpretive approaches that attempt to induce structure from keyword queries. The process of query interpretation is usually guided by structures in data, and schema and is often supported by a graph exploration procedure. However, graph exploration-based interpretive techniques are impractical for multi-tenant scenarios for large database because separate expensive graph exploration states need to be maintained for different user queries. This leads to significant memory overhead in situations of large numbers of concurrent requests. This limitation could negatively impact the possibility of achieving the ultimate goal of personalizing search. In this paper, we propose a lightweight interpretation approach that employs indexing to improve throughput and concurrency with much less memory overhead. It is also more amenable to distributed or partitioned execution. The approach is implemented in a system called "SKI" and an experimental evaluation of SKI's performance on the DBPedia and Billion Triple Challenge datasets show orders-of-magnitude performance improvement over existing techniques.

Provided by North Carolina State University

Citation: Scaling up personalized query results for next generation of search engines (2013, September 18) retrieved 27 April 2024 from <https://phys.org/news/2013-09-scaling-personalized-query-results.html>

<p>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.</p>
--