


# System loads web pages 34 percent faster by fetching files more effectively

March 9 2016, by Adam Conner-Simons

---



There are few things more frustrating than a slow-loading web page. For companies, what's even worse is what comes after: users abandoning their site in droves. Amazon, for example, estimates that every 100-millisecond delay [cuts its profits by 1 percent](#).

To help combat this problem, researchers from MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL) and Harvard University have developed a system that decreases page-load times by 34 percent. Dubbed "Polaris," the framework determines how to overlap the downloading of a page's objects, such that the overall page requires less

time to load.

"It can take up to 100 milliseconds each time a browser has to cross a mobile network to fetch a piece of data," says PhD student Ravi Netravali, who is first author on a paper about Polaris that he will present at this week's USENIX Symposium on Networked Systems Design and Implementation. "As pages increase in complexity, they often require multiple trips that create delays that really add up. Our approach minimizes the number of round trips so that we can substantially speed up a page's load-time."

The paper's co-authors include graduate student Ameesh Goyal and professor Hari Balakrishnan, as well as Harvard professor James Mickens, who started working on the project during his stint as a visiting professor at MIT in 2014. The researchers evaluated their system across a range of network conditions on 200 of the world's most popular websites, including ESPN.com, NYTimes.com (The New York Times), and Weather.com.

## **How web pages work**

Before you type in a URL, your browser doesn't actually know what the page looks like. To load the page, the browser has to reach across the network to fetch "objects" such as HTML files, JavaScript source code, and images. Once an object is fetched, the browser evaluates it to add the object's content to the page that the user sees.

But it's not quite that simple. Evaluating one object often means having to fetch and evaluate more objects, which are described as "dependencies" of the originals. As an example, a browser might have to execute a file's JavaScript code in order to discover more images to fetch and render.

The problem is that browsers can't actually see all of these dependencies because of the way that objects are represented by HTML (the standard format for expressing a webpage's structure). As a result, browsers have to be conservative about the order in which they load objects, which tends to increase the number of cross-network trips and slow down the page load.

## **How Polaris fits in**

What Polaris does is automatically track all of the interactions between objects, which can number in the thousands for a single page. For example, it notes when one object reads the data in another object, or updates a value in another object. It then uses its detailed log of these interactions to create a "dependency graph" for the page.

Mickens offers the analogy of a travelling businessperson. When you visit one city, you sometimes discover more cities you have to visit before going home. If someone gave you the entire list of cities ahead of time, you could plan the fastest possible route. Without the list, though, you have to discover new cities as you go, which results in unnecessary zig-zagging between far-away cities.

"For a web browser, loading all of a page's objects is like visiting all of the cities," Mickens says. "Polaris effectively gives you a list of all the cities before your trip actually begins. It's what allows the browser to load a webpage more quickly."

Dependency-trackers have existed before, but are all as constrained as the browsers themselves. This is because their method of comparing lexical relationships—specifically, the text in HTML tags—mimics the way browsers load pages and does not capture more subtle dependencies.

## A better approach to a faster web

Tech companies like Google and Amazon have also tried to improve load times, with an emphasis on lowering costs for data usage. This means that they often focus on the challenge of more quickly transferring information via data compression. The CSAIL team, meanwhile, has demonstrated that Polaris' gains on load-time are more consistent and more substantive.

"Recent work has shown that slow load-times are more strongly related to network delays than available bandwidth," Balakrishnan says. "Rather than decreasing the number of transferred bytes, we think that reducing the effect of network delays will lead to the most significant speedups."

Polaris is particularly suited for larger, more complex sites, which aligns nicely with [recent trends](#) of modern pages ballooning to thousands of (JavaScript-heavy) objects. The system is also valuable for mobile networks, since those tend to have larger delays than wired networks.

"Tracking fine-grained dependencies has the potential to greatly reduce page-load times, especially for low-bandwidth or high-latency connections," says Mark Marron, a senior research software development engineer at Microsoft. "On top of that, the availability of detailed dependence information has a wide range of possible applications, such as tracking the source statement of an unexpected value that led to a crash at runtime."

**More information:** Polaris: Faster Page Loads Using Fine-grained Dependency Tracking. [web.mit.edu/ravinet/www/polaris\\_nsd16.pdf](http://web.mit.edu/ravinet/www/polaris_nsd16.pdf)

*This story is republished courtesy of MIT News ([web.mit.edu/newsoffice/](http://web.mit.edu/newsoffice/)), a popular site that covers news about MIT*

*research, innovation and teaching.*

Provided by Massachusetts Institute of Technology

Citation: System loads web pages 34 percent faster by fetching files more effectively (2016, March 9) retrieved 23 June 2024 from <https://phys.org/news/2016-03-web-pages-percent-faster-effectively.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.