

# Challenges in Building Large-Scale Information Retrieval Systems

[Invited Talk]

Jeffrey Dean  
Google, Inc.  
jeff@google.com

## ABSTRACT

Building and operating large-scale information retrieval systems used by hundreds of millions of people around the world provides a number of interesting challenges. Designing such systems requires making complex design tradeoffs in a number of dimensions, including (a) the number of user queries that must be handled per second and the response latency to these requests, (b) the number and size of various corpora that are searched, (c) the latency and frequency with which documents are updated or added to the corpora, and (d) the quality and cost of the ranking algorithms that are used for retrieval. In this talk I will discuss the evolution of Google's hardware infrastructure and information retrieval systems and some of the design challenges that arise from ever-increasing demands in all of these dimensions. I will also describe how we use various pieces of distributed systems infrastructure when building these retrieval systems. Finally, I will describe some future challenges and open research problems in this area.

## Categories and Subject Descriptors

H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval—*retrieval models, search process*;  
H.3.4 [Information Storage and Retrieval]: Systems and Software—*distributed systems*; H.3.5 [Information Storage and Retrieval]: Online Information Services

## General Terms

Design, Performance

## Keywords

Search engines, scalability

## Biography

Jeff joined Google in 1999 and is currently a Google Fellow in Google's Systems Infrastructure Group. He has co-designed/implemented five generations of Google's crawling, indexing, and query serving systems, and he has co-designed/implemented major pieces of Google's initial advertising and AdSense for Content systems. He has also been heavily involved in the design and implementation of Google's distributed computing infrastructure, including the MapReduce and BigTable systems, dabbled with system software for statistical machine translation, and worked on a variety of internal and external developer tools.