

**STANFORD UNIVERSITY**  
**DEPARTMENT OF STATISTICS**  
**DEPARTMENTAL SEMINAR**

4:15 p.m., Tuesday, August 02, 2005  
Sequoia Hall Room 200  
(Cookies at 3:45 in 1st Floor Lounge)

*Rob Pike*  
Google Inc.

**Interpreting the Data**

Abstract:

Google has large amounts of data – to index the Web, you need to have an addressable copy of it – and many custom tools designed to operate at that scale. One of those tools is a special-purpose language for distributed data processing. I will describe that language, including its unusual model and implementation, starting from the point of view of my long-term interest in concurrent programming.

Background on Speaker:

Rob Pike, well known for his appearances on “Late Night with David Letterman”, has been a member of the Systems Lab at Google, Inc., since 2002. For countless years he toiled as a Member of the Technical Staff at Bell Laboratories in Murray Hill, New Jersey. In 1981 he wrote the first bitmap window system for Unix systems, and has since written ten more. He was a principal designer and implementer of the Plan 9 and Inferno operating systems. With Bart Locanthi he designed the Blit terminal; with Brian Kernighan he wrote ‘The Unix Programming Environment’ and ‘The Practice of Programming’. A shuttle mission nearly launched a gamma-ray telescope he designed. He is a Canadian citizen and has never written a program that uses cursor addressing.