Computer Science Department, Northwestern University
Presents

"Content-Based Music Retrieval on Acoustic Data"

Dr. Cheng (Calvin) Yang
Stanford University

--- 12:30 pm --- <u>Wednesday</u>, February 19, 2003
CS Main Classroom (rm. 381), 1890 Maple Avenue

Abstract:

With the explosive amount of music data available on the Internet in recent years, there has been much interest in developing new ways to search and retrieve such data effectively. Currently, most music search engines operate on text labels or symbolic data, rather than the underlying acoustic content. In this talk, we discuss the problem of searching and retrieving music based on acoustic content similarity. We present our retrieval framework known as MACSIS, which accepts rea-ltime audio queries and searches for similar music recordings while tolerating changes in noise level, performance style and tempo.

In our framework, each audio file is broken down into small segments and converted into feature vectors that capture information about pitch estimation near major events. All vectors are stored in a high-dimensional indexing structure called LSH, a probabilistic indexing scheme that makes use of multiple hashing instances in parallel. At retrieval time, small segments of audio matches are retrieved from the index and pieced together using the Hough Transform technique, and results are used as the basis to rank candidate matches. The system is highly parallelizable and can easily adapt to a distributed computing environment.