

I am willing to supervise projects related to software engineering (in particular, mining software repositories, reverse engineering, web services, service oriented architectures, and capacity planning/performance engineering). Please find below more specific projects that I have in mind.

Integrating Music into Software Development [2 students]

I am looking for students with a good musical background, preferably someone who has had training in music theory equivalent to courses offered at The Royal Conservatory of Music (RMC):

http://www.rcmusic.ca/ContentPage.aspx?name=Fall_and_Winter_Theory_and_History

The project involves novel work which explores integrating music into software development activities. The project converts the history of development of a project into a music score. This work would permit one to listen to a project history instead of reading about it.

Active Portals for Research Communities [2-4 students]

Researchers need to keep abreast of the latest trends and advancements in their research field. Traditionally researchers have accomplished this by examining conference proceedings and journals. With the advent of the web, many researchers are using web technologies to stay up to date. The Web gives researchers access to the latest information in a timely fashion. Researchers browse homepages of other researchers in their community to read through their latest work. Researchers also sometimes use search engines (e.g. Google and Google Scholar) to retrieve papers that cite their work.

This project would use Google's Web API (<http://www.google.com/apis/>) to develop research community portals. A researcher would specify several researchers in his or her community. The web API would be used to build the rest of the research community and create a research portal. The portal would contain information about the latest publications in the field, the latest presentation by other researchers in the community, and the latest citation to papers. The portal would be continuously updated as new pertinent information is available online.

Mining the Publications History of the ACM [2-3 students]

Publications in a research community give a picture of the progress of collaboration and emergence of topics in an active research field. The authorship details on each publication represent a social network of collaboration between researchers in the community. One would expect a high degree of collaboration in an academic community, in contrast to a lower degree of collaboration in commercial communities. Furthermore, the titles of these publications permit us to track the appearance of new research topics and areas of interest in the community and the computer industry as a whole. Such topics of interest may in some cases explain changes in the collaboration structure of a community and may shed some light on its evolution.

This project would mine the publication history of all the ACM sponsored conferences to gain a better understanding of how the computer science field has evolved over the years.

A variety of graph theory techniques along with novel visualization and analysis techniques would be developed. This project expands on this work:

<http://www.cs.queensu.ca/home/ahmed/home/pubs/wcre2004.pdf>

A Framework for Recording and Replaying Applications [3-4 students]

Students would create a framework for recording and replaying the execution of an application. The execution of all or some of the components involved in a multi-threaded application would be captured. The replay of the captured events would be deterministic. This approach would permit a developer to fast forward or to rewind the execution during a debugging session without worrying about the non-deterministic nature of a multi-threaded application. Students should focus on minimizing the performance overhead of the recording framework. By minimizing the performance overhead, the framework could be used for tracing the execution of applications in a production setting. For more details see <http://www.cc.gatech.edu/~orso/papers/joshi.orso.presentation.pdf>

Mining Software Repositories (MSR) [varies]

A solid understanding of the software development process will assist us in ensuring the successful evolution of large software systems, and the discovery of new techniques and approaches to assist developers maintaining these systems. Version control systems, such as CVS, offer a convenient record keeping system. They track the development history/process by recording changes to the source code over time.

I am interested in projects that would mine this rich historical information and integrate it into development environments (e.g. Eclipse). Additional information about MSR is available at <http://msr.uwaterloo.ca/> and <http://msr.uwaterloo.ca/msr2007/challenge/>