

Research Summary

Dr. Ramon Lawrence

Overview

My research program combines core database research with interdisciplinary collaboration and applied projects. ***Total external research funding is over \$2 million with 60 peer-reviewed publications.*** Students graduated include 2 Ph.D. (both employed in academia), 13 Master's, and 25 undergraduate honours students (two of which were hired by Microsoft).

Core Database Research

My core research is on *database implementations and performance*, specifically in the area of database integration and join algorithms. I have published several join algorithms designed for early production of results: Early Hash Join (VLDB 2005), slice join for integration systems (DKE 2008), and studied how reading policy affects join result production (Information Sciences 2007). The Histojoin algorithm (Information Systems 2009) was developed to improve join performance for skewed databases. Performance of skewed joins can be increased up to 50%, and the algorithm has been implemented and included in PostgreSQL 8.4. Further research improved join performance based on join cardinality (SAC 2009), and examined the relative performance of multi-way hash joins versus standard binary joins (ICEIS 2013).

Our research group is building the smallest relational database supporting SQL for sensor and embedded nodes. The goal is to have a database for devices such as the Arduino that may have less than 64 KB of code space and as little as 4 KB of memory. The system is called LittleD (SAC 2014) and currently supports SQL SELECT including joins and pre-compilation (SAC 2016). A key-value store library for embedded databases called IonDB (CCECE 2015) was also released. Related foundational work includes flash-based translation for data storage and file systems (CCECE 2011, 2016) and sorting algorithms for embedded devices (IDEAS 2010, IJDMIS 2013).

My other core interest is *database integration and virtualization* which involves combining the data in one or more databases, including NoSQL and relational databases, into an integrated form that has more value compared to when the databases were not connected. In order to achieve this goal, a researcher is faced with issues on data semantics (what does the data mean?), data representation (how is the data stored and formatted?), and data manipulation (what are efficient algorithms for extracting and combining the data?).

Research has produced the UnityJDBC system which has been released as a commercial product (www.unityjdbc.com). UnityJDBC allows a Java program to use the JDBC interface to query any number of databases in a single query. The system has recently been extended to support SQL querying of NoSQL sources such as MongoDB (CSCI 2014), and the first JDBC driver for MongoDB was released as a commercial product. The work on UnityJDBC is a multi-year effort that started at the University of Iowa with funding from the Army Research Office. Many publications have developed from subproblems encountered during development of UnityJDBC including schema matching using ontologies, querying databases without joins (Ph.D. graduate Terry Mason), and efficient join algorithms for integration systems (Early Hash Join, slice join).

Interdisciplinary and Applied Projects

I strongly believe in interdisciplinary collaboration and applying research results to practical problems in other domains. This includes industrial work with GE on the Industrial Big Data Initiative and local companies such as Vineyard Networks (now Procera Networks) on network monitoring (funded by an NSERC Engage grant). A major long-term project is the construction of a digital archive database for historical Métis records in collaboration with the University of Ottawa. The website <http://dadp.ok.ubc.ca> hosts over 100,000 historical records.

Collaboration with the City of Kelowna to reduce water usage and improve sustainability began in summer 2009 and constructed and field tested a sensor network system for environmental monitoring in a city park. The sensor system controlled the irrigation and resulted in a 50% reduction in water usage. The research won best student paper at the IEEE Sensors Applications Symposium 2010. The undergraduate student in the research team won several oral research awards including a top presentation at the Universitas 21 International Undergraduate Research Conference in 2011. The project has evolved into producing a mobile parks management system for the City of Kelowna with funding from the Okanagan Basin Water Board and UBC that is used by City of Kelowna irrigation staff.

Work with Vadim Bulitko (University of Alberta) designed real-time path finding algorithms for video games that use databases to improve performance. The algorithms have overall superior performance in terms of execution time and path quality compared to other state-of-the-art algorithms (IEEE TCAIG 2013, AAAI 2013). A previous large collaboration project was the construction of the NEXRAD data archive of weather radar data with Dr. Anton Kruger and Dr. Witold Krajewski of IHR (Iowa). This project was funded by a *\$1.4 million NSF Medium ITR grant* to provide a national archive and query system for the NEXRAD data. I developed the database infrastructure for a system handling terabytes of data and archiving in real-time. The system has been completed and used by numerous researchers across North America.

Research on Teaching Methods

There is an increasing demand for innovative teaching approaches and systems, and numerous companies and technologies (such as Massive Open Online Courses) have emerged. There is a competition for the hearts and minds of students, and a university that provides an exceptional educational experience will produce better graduates and maintain enrollments. It is also critical to examine how valuable educational experiences can be provided at lower cost. I have an interest in using technology to improve teaching methods, automated evaluation, and learning. Work published in IEEE Transactions on Education in 2004 demonstrated how a programming assignment that uses competing game AI can increase student motivation. I have also built an automated lab marking system that has been used in first-year Physics at UBC Okanagan for several years and removed paper assignments for over 1000 students (CATE 2011).

Summary

My research program balances core database research with interdisciplinary collaborations and applied industrial projects. I have built a Canadian Foundation for Innovation (CFI) and NSERC funded database research lab that has on average about three graduate and three undergraduate students. I continually search for good students and exciting research projects to continue the research success.