



Crowdsourcing Research Group

Alfred Weaver

Professor

weaver@cs.virginia.edu

www.cs.virginia.edu/people/faculty/weaver.html

Department of Computer Science
University of Virginia
Charlottesville, VA
434.982.2201

"These gigabit LANs are just not fast enough!"

Weaver and colleagues developed the Xpress Transport Protocol (supporting multicast and knowledge of the multicast group members) for the US Navy, and then turned XTP into a commercial product. He was a founder of Virginia's Internet Technology Innovation Center, which provided technical and business strategy for e-business, and developed e-commerce software for corporate and government clients. Previous projects included improving security for mobile devices and writing software that could perform a scientometric analysis of scientific literature. Current research focuses on determining whether the use of crowdsourcing would be an effective adjunct for disaster management.



An Empirical Performance Evaluation of Relational Keyword Search Techniques

Extending the keyword search paradigm to relational data has been an active area of research within the database and information retrieval (IR) community during the past decade. Many approaches have been proposed, but despite numerous publications, there remains a severe lack of standardization for the evaluation of proposed search techniques. This lack of standardization has resulted in contradictory results from different evaluations, and the numerous discrepancies muddle what advantages are proffered by different approaches. In this project we have created the most extensive empirical performance evaluation of relational keyword search techniques to appear to date in the literature. Our results indicate that many existing search techniques do not provide acceptable performance for realistic retrieval tasks. In particular, memory consumption precludes many search techniques from scaling beyond small data sets with tens of thousands of vertices. We also explore the relationship between execution time and factors varied in previous evaluations; most of these factors have relatively little impact on performance. Our work confirms previous claims regarding the unacceptable performance of these search techniques and underscores the need for standardization in evaluations as exemplified by the IR community.

Crowdsourcing the Crisis

The premise of this work is that the public would be better served if it had access to reliable, near real-time information concerning the “ground truth” of unfolding events. One way to collect those reports is through crowdsourcing. By taking advantage of the army of people connected to the internet via both wired and wireless technologies, we can easily gather information that would help a citizen make an informed decision regarding his or her health and welfare. The widespread use of smartphones allows people to post geolocated reports and pictures with ease. Disaster management teams would likewise benefit from the input of more observers. On the other hand, the trustworthiness of these reports may be uncertain. We have built a crowdsourced system that could be used to enhance disaster management, and we developed three methods to enhance the trust that the citizen places in such a system.

Biotelemetrics

As the population ages and the risk of chronic disease increases, the cost of healthcare will rise. Technology for mobile telemetry could reduce cost and improve the efficiency of treatment. In order to achieve these goals, we first need to overcome several technical challenges. We built a Mobile Biotelemetric System (MBS) that addresses these remote medical monitoring challenges. MBS comprises a custom low-power sensor node that accurately collects and analyzes electrocardiogram (ECG) data, a client service with a multifaceted policy engine that evaluates the data, and a web portal interface for visualizing the ECG data streams. MBS differs from other remote monitoring systems primarily in the policy engine’s ability to provide flexible, robust, and precise system communication end-to-end and to enable tradeoffs in metrics such as power and transmission frequency. We showed that, given a representative set of ECG signals, policies can be set to make the operation of the hardware and software resilient against transient ECG conditions for both security and monitoring purposes. We demonstrated that our system adaptively trades off system-level metrics based on a combination of operating conditions and user input, and that our heartbeat detection algorithm performs well for challenging ECG input.

RECENT RESEARCH DEVELOPMENTS

- A.Jurik & A.Weaver, “Bidirectional ECG Monitoring with an Event Detection Policy Engine,” *21st Intl Conf. on Computer Communications & Networks (ICCCN’12)*, 7/30-8/2/12, Munich, Germany
- A.Weaver, J.Boyle, L.Besaleva, “Application & Trust Issues when Crowdsourcing the Crisis,” *2nd Intl. Workshop on Privacy, Security, & Trust in Mobile & Wireless Systems (MobiPST’12)*, 7/30/12, Munich, Germany
- A.Weaver & J.Boyle, “Using Crowdsourcing to Enhance Crisis Management,” *Intl. Conf. on Frontiers in Education: Computer Science & Computer Engineering (FECS’12)*, Las Vegas, NV, 7/12

RECENT GRANTS

- Battelle Memorial Institute – Scientometric Analysis of Eastern European Literature
- DOD/Defense Energy Office – Crowdsourcing the Crisis
- NSF – Secure Mobile Computing

SEAS Research Information

Pamela M. Norris,
Executive Associate Dean for Research
University of Virginia
Box 400232
Charlottesville, VA 22903
pamela@virginia.edu
434.243.7683

