Research Statement

I perform research on design methods, languages and tools for embedded systems, particularly those involving reconfigurable devices such as field-programmable gate arrays, and their applications. An embedded system is a computer system that is designed to perform one or a few dedicated functions often with real-time computing constraints. Embedded systems are used in various fields such as avionics, robotics, medicine and communications.

My current research projects include *1) Simulate & Eliminate,* design of a high-level synthesis tool to ease the hardware design process, *2) hardware security,* to tightly control the flow of information both in-to and out-of a device guaranteeing its trustworthiness, *3) hardware system design for multi-frequency biplanar interferometric (MBI) technique,* to enable real-time inspection of seabed and water column targets, and *4) acceleration of biomedical imaging applications,* to significantly reduce the computation time and enable more advanced, larger scale biomedical experiments. Additional information about these and other projects can be found at http://cseweb.ucsd.edu/~airturk/research_papers.html. In the following, I highlight my recent research projects.

Simulate & Eliminate: A Top-to-Bottom Design Methodology for Application-Specific Multi-Core Architectures

Integrated circuit technology will continue to follow Moore's Law for the foreseeable future, and multi-core architectures will provide a means to increase performance while keeping power consumption under control. There is also an increasing trend towards application-specific multi-core processing, particularly in embedded computing devices that have stringent area, performance and power consumption requirements. These application-specific multi-core architectures are carefully tuned for a particular application, therefore consuming less area and power while achieving high performance compared to their equivalent general-purpose implementations. However, achieving the desired constraints requires careful tuning of the underlying architecture with a significant amount of time and effort. Therefore, a high level design tool for design space exploration and fast prototyping of application-specific multi-core architectures is essential.

We are currently developing a tool, Simulate & Eliminate (S&E), which produces synthesizable hardware description language (HDL) for application-specific multi-core architectures for given particular application(s). S&E follows a *top-to-bottom* approach by first generating a general-purpose multi-core architecture. Then, the provided applications are simulated on this architecture and the unneeded functionality (interconnect, functional resources, control and memory) is eliminated resulting in application-specific multi-core architecture. Our approach is distinct from the common approach of synthesis of microarchitecture datapaths from primitives. The existing tools employ a *bottom-to-top* methodology, which pieces together functional units, interconnect and control logic based on the given application. The technical advantages of S&E's top-to-bottom design methodology are *scalability, reconfigurability and ease of design.* This project is patented at UCSD and currently funded by Intel Corporation for 100K/year due to its novelty. Based on our promising results, we believe that our research will provide solutions to problems in electronic system level design, application specific architecture and embedded system design. We also submitted an NSF proposal that is currently under review.

Hardware Security

High assurance systems such as those found in flight control and banking systems require strict guarantees on correct operation or they face catastrophic consequences. Ensuring that these systems operate as intended is an extremely difficult and costly problem. Some have estimated that such assurance can cost $10K per line of code and take up to 10 years. To realize high assurance systems, we are working on two different topics: *1) Gate-Level Information-Flow Tracking (GLIFT) logic to guarantee non-interference* and *2) exploitation of 3D integrated circuit technology to secure COTS hardware*.

Currently, we are exploring gate-level information-flow tracking (GLIFT) to guarantee non-interference. Non-interference is a common property that often needs to be guaranteed in high assurance systems where certain parts of the system should never interfere with other parts. For example, the Boeing 787 aircraft has connectivity between the user and light control networks. Ensuring that there are no unintended information flows between the two networks is critical for the correct operation of the aircraft.

We are also developing novel security management techniques built around fundamentally new hardware abilities enabled by 3-D integration. Manufacturers are often forced to choose less costly alternatives, such as an older, cheaper process increasing the gap in performance between low volume high assurance systems and commercial systems every year. As a result of these economic factors, designers of trustworthy systems requiring high performance need some way to incorporate commercial hardware components without compromising security. Therefore a method of bridging the gap between cutting-edge technology and trustworthy systems is of paramount

necessity. Our work shows that commodity integrated circuits could be enhanced with a separate silicon layer, stacked using 3-D integration, therefore enabling different ways of introspection. This is a collaborative research between UC San Diego, UC Santa Barbara and Naval Postgraduate School (NPS) and our team was awarded an NSF grant, ~$1.5M in total, 2010-2013. In addition, we currently have pending patents on this research.

Hardware System Design for Real-Time, Biplanar-Interferometric, Georeferenced Solutions from Simrad EK60 and ME70 Echosounder Data

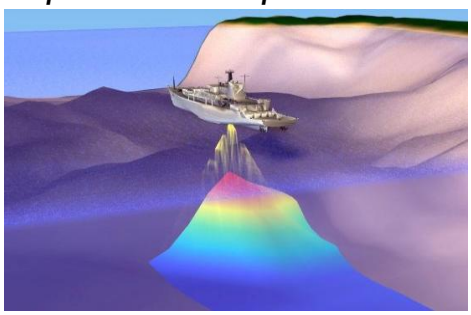

The Simrad EK60 multi-frequency and ME70 multi-beam echosounders were designed for quantitative fisheries research and are standard equipment on all of the new NOAA Fisheries Survey Vessels (FSVs). The data from both of these systems can be processed using the multi-frequency biplanar interferometric (MBI) technique to achieve spatial resolutions for pelagic targets or bathymetry which are two to three orders of magnitude greater than with standard data processing. The MBI solutions can be used for improved imaging and quantitative evaluations of water column scatterers and their seabed habitats as shown in the picture. For example, bathymetry can be mapped with high sub-beam resolution, and seabed echoes can be normalized by incidence angles, classified by roughness and hardness.

However, the data processing required to execute this algorithm is computationally intensive and therefore slow. Currently, processing is done using custom Matlab scripts and can take approximately 15 to 60 seconds per ME70 transmission for data collected to 250 to 2000-m ranges. We analyzed the algorithm, identified processing bottlenecks and are currently implementing the algorithm using a graphic processing unit (GPU). Depending on the quality of results and performance on the GPU, we may consider a field programmable gate array (FPGA) implementation. In either case, the hardware will interface with either the EK60, ME70, or both, and calculate MBI solutions in real-time. This system will provide the MBI results during data acquisition to enable real-time inspection of seabed and water column targets, and enhance decisions concerning adaptive survey operations and gear deployment, such as ROV navigation and trawl placement.

NOAA is supporting this research with a total of 230K for two years.

Accelerating Biomedical Imaging

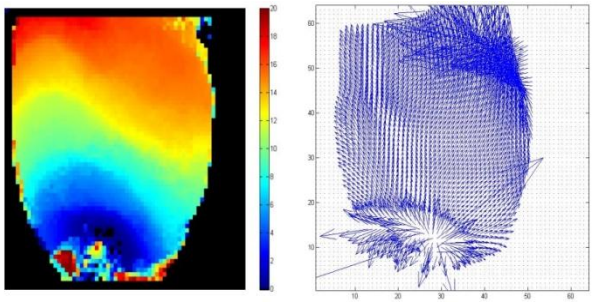

I am also working on a number of projects related to accelerating biomedical imaging. These include *optical mapping of heart tissue*, *stem cell identification and selection*, and *inertial microfluidics*. All of these applications are for the most part, currently done in a tedious, non-optimized manner, and involve hours/days of time for manual computation. Automation of these techniques could have profound effects on experiments, treatments and therapies for ailments ranging from heart disease to cancer. These projects involve understanding the algorithm and translating onto a computation device (Graphic Processing Unit, Field-Programmable Gate Array, Multi-Core Processor) with the aim of creating a real-time implementation.

As an example, one of my current projects brings together researchers from the medicine, bioengineering, and computer science departments to develop an experimental tool that has the potential for wide usage in cardiac research. This project investigates the Graphic Processing Unit (GPU) acceleration of optical mapping algorithm for cardiac electrophysiology. Hundreds of thousands of people die from heart attack in the U.S. every year. Cardiac electrophysiology is important for understanding, diagnosing and treating electrical arrhythmias of the heart. Optical mapping is an increasingly popular tool for experimentally visualizing the electrical activity in the heart which can be seen in the picture. However, optical mapping is a computationally intensive procedure and present-day research is substantially hindered by the significant amount of time required to perform analysis even on the smallest amount of data. For example, 1 second of data requires 3.7 hours of computation. Through our research, we study the efficiency of different GPU implementations and our results show that GPU implementation of the optical mapping algorithm has a 33X speedup over the equivalent serial implementation. This will open new avenues of research that ultimately enable closed-loop feedback control of cardiac function that is not feasible using global measures of electrophysiology such as the electrocardiogram (ECG).