Introduction to Embedded Systems Research

Robert Dick
Semester: Winter 2019
Course number: EECS 598-13

Summary

This course is designed to prepare graduate and advanced undergraduate students with a foundation in, and head start on, research related to embedded system analysis, design, and synthesis. The first half of the course consists of lectures and assigned reading material on fundamental embedded systems topics on which future research will generally build. The second half of the course focuses on a specific, and possibly new, topic in the field. This semester, the focus is embedded machine learning in the Internet-of-Things (IoT).

1 Survey Topics

1. Specification languages and models;
2. scheduling, allocation, and assignment: problem definitions and optimization techniques;
3. embedded (real-time) operating systems;
4. embedded signal processing and machine learning hardware and software;
5. energy- and temperature-aware design and embedded power supplies;
6. wireless communication and its impact on power consumption;
7. sensors and actuators;
8. reliability-aware design, formal methods, and testing;
9. embedded system security; and
10. applications including the IoT, wireless sensor networks, autonomous vehicles, wearables, and smartphones.

2 Focus Topic

Shortly, every human will be served by hundreds or thousands of sensing, decision making, and actuating machines. These smart things will be connected to each other, and the rest of the world. For the many that operate on battery or scavenged power, communication will impose severe energy consumption penalties. Local computation will need to convert sparse data to dense information and decisions at higher levels in the network will be based on processed data from the sensors; those local decisions will be critical to the quality

of the results produced by the network. This implies new approaches to the design of machine learning hardware and algorithms, in which energy efficiency joins accuracy as a central optimization objective, and the communication and computation implications of partitioning machine learning algorithm components among embedded systems and higher-performance servers must be explicitly considered. We will cover efficient embedded machine learning algorithms and hardware, LPWAN communication, as well as security, privacy, and reliability in the IoT.

3 Grading

Without Project	
Summaries of assigned reading:	35%
Presentation and critique quality:	35%
Exams:	30%

With Project	
Summaries of assigned reading:	30%
Presentation and critique quality:	30%
Exams:	20%
Project:	20%