



Doing More With Less: Cost-Effective Infrastructure for Automotive Vision Capabilities

Department of Computer Science

University of North Carolina at Chapel Hill

August 2014

The Challenge

Many cyber-physical systems must operate in situations where accurately understanding and predicting environmental conditions is essential. To provide such environmental awareness, advanced sensor technology is required that can generate significant amounts of data that must be processed in real time. The resulting data-processing rates can be difficult to sustain computationally. One potential solution is to resort to hardware over-provisioning. However, this is a wasteful practice that can be untenable in many domains due to monetary cost restrictions and size, weight, and power (SWaP) limitations.

One such domain is automotive systems. In this domain, a proliferation of advanced sensor technology is being fueled by an expanding range of autonomous capabilities. Driver-assist features, such as blind spot warnings, automatic lane-keeping, adaptive cruise control, and collision avoidance/mitigation systems, are becoming commonplace in high-end vehicles; in the coming years, such features are expected to evolve to provide significantly enhanced functionality, such as pedestrian detection, cross-traffic alerts, traffic sign recognition, and 360-degree sensing. At the same time, fully autonomous vehicles have been demonstrated in various "one-off" settings; these include the press-worthy Google Car, and various research vehicles that were fielded as part of the DARPA Urban Grand Challenge.

In these one-off settings, sensing is typically provided via significant hardware over-provisioning (e.g., expensive high-resolution LIDARs, dedicated per-feature compute platforms, expensive interconnects, etc.). The resulting costs are too extreme for a consumer product. For example, the computational/sensing infrastructure in the latest Google Car reportedly costs over \$150,000. While this is not a significant expense for Google, it certainly would be for a typical consumer. To further complicate matters, defects in mass produced vehicles can cause many more accidents than in custom prototypes, so the former are subject to higher reliability requirements while likely being less well maintained. Thus, any re-provisioning of sensing computations must be amenable to certification under stringent conditions.

To recap, cars with autonomous capabilities are good representatives of a broader category of safety-critical cyber-physical systems being faced with a pressing challenge, namely **to devise computational infrastructure for advanced sensing capabilities that adds only modestly to system costs**. Other representatives in this category can be found in the areas of robotics, medical devices, and avionics.

The Approach

Motivated by the challenge noted above, we are developing infrastructure for realizing a "**less is more**" approach in real-time sensor processing in cyber-physical systems. This infrastructure is directed at a particularly compelling challenge problem: enabling cost-effective driver-assist and autonomous-control automotive features that utilize vision-based sensing through cameras. This problem is challenging for several reasons. First, the support of such features requires that dynamically variable, computationally intensive workloads be supported in real time. Second, such features are subject to stringent certification requirements. Third, practical solutions must adhere to rigid monetary cost limitations. Fourth, vision-based sensing, while cost effective, may be affected by hardware defects (e.g., scratched or warped lenses) and non-ideal weather conditions (e.g., rain or fog) and must offer graceful degradation in the face of such difficulties. These challenges will be considered in the context of hardware platforms comprised of multicore processors augmented with graphics processing units (GPUs) as accelerators. This platform choice is motivated by desirable performance/energy characteristics.

The desired infrastructure is being developed by (i) examining numerous multicore-based CPU+GPU hardware configurations at various fixed price points (e.g., \$1,000, \$2,000, etc.) based on realistic automotive use cases, and by (ii) characterizing the range of vision-based workloads that can be feasibly supported using our software infrastructure. This research is a collaborative effort involving academic researchers at UNC and engineers at General Motors Research.

Significance

Devising efficient, cost-effective infrastructure for hosting vision-based sensing computations would be a

breakthrough result for safety-critical, cyber-physical systems in domains such as automobiles, avionics, robotics, and medical instrumentation. Without such infrastructure, vision-based sensing is unlikely to impact the daily lives of people to the extent suggested by press-worthy "one-off" prototype demonstrations.

The technology developed in this project will be a critical enabler for the U.S. automotive industry to gain a competitive advantage in the market by enabling fast adoption of advanced hardware technology, while keeping costs low. The research results will also influence industry standards like AUTOSAR to include adequate services for supporting vision-based capabilities.

Project Members

James Anderson, professor

Sanjoy Baruah, professor

Alex Berg, professor

Shige Wang, General Motors Research

Research Sponsor

National Science Foundation Cyber-Physical Systems Program

For More Information

Dr. James Anderson

Department of Computer Science

University of North Carolina at Chapel Hill

CB#3175, Sitterson Hall

Chapel Hill, NC 27599-3175

Phone: (919) 590-6057

Fax: (919) 590-6105

Email: anderson@cs.unc.edu

<http://www.cs.unc.edu/~anderson/projects/autovision.html>