Sandia National Laboratories

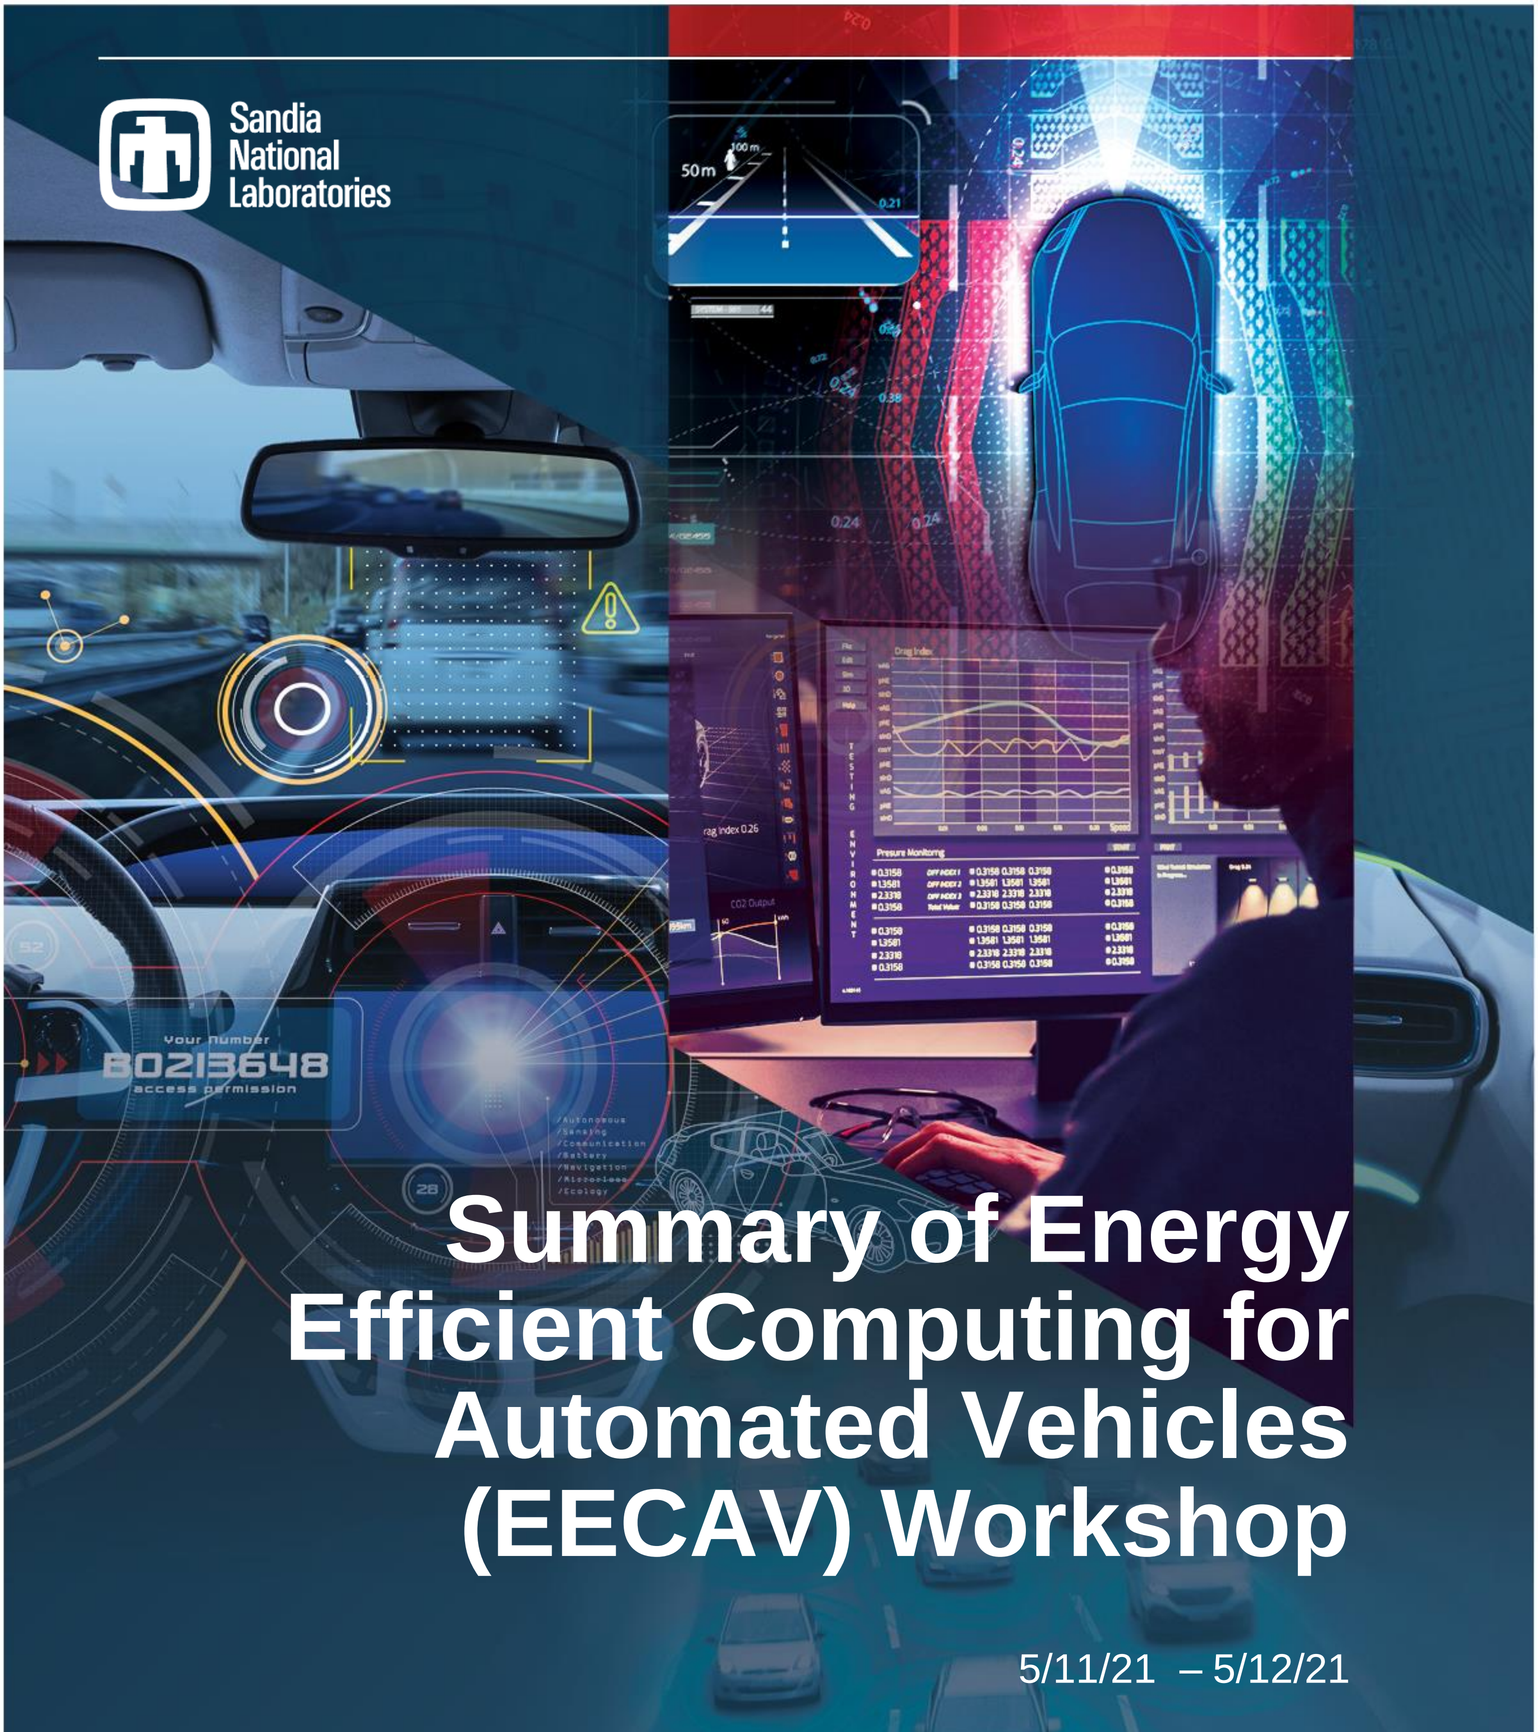

Summary of Energy Efficient Computing for Automated Vehicles (EECAV) Workshop

5/11/21 – 5/12/21

Sandia National Laboratories
arm
University of Michigan
Carnegie Mellon University
Hewlett Packard Enterprise
USCAR United States Council for Automotive Research LLC

Issued by Sandia National Laboratories, operated for the United States Department of Energy by National Technology and Engineering Solutions of Sandia, LLC.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@osti.gov
Online ordering: http://www.osti.gov/scitech

Available to the public from

U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Rd
Alexandria, VA 22312

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.gov
Online order: https://classic.ntis.gov/help/order-methods/

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. SAND2021-12338O

ABSTRACT

This summary serves as the report for the Energy Efficient Computing for Automated Vehicles (EECAV) workshop that was held virtually by Sandia National Laboratories (SNL) on May 11 and May 12, 2021. The purpose of the workshop was to reach out to a broad community from industry, academic and national laboratories, and to collect feedback from all stakeholders on the scope, timeline, assumptions and identified research and development (R&D) problems that the Organizing Team has worked on as part of the EECAV Roadmap Outline.

This workshop consisted of several feedback sessions for the event which took place over two half-days. A total of 50 participants joined the invitation only workshop, with experts representing automotive Original Equipment Manufacturers (OEMs), semiconductor companies, academics, and national laboratories. The sessions in the first day covered the motivation and purpose of the EECAV road mapping activities, an overview of the current status of the semiconductor and autonomous vehicle (AV) technology and projections, and an overview of the EECAV Roadmap Outline, including the scope, timeline and assumptions that the Organizing Team worked on. The sessions in the second day of the workshop covered several R&D problems in each of four “Technical Areas” that the Organizing Team identified. Each of the sessions was led by the leader of that technical area to introduce the topic and guide the feedback discussion with the workshop participants.

A read-ahead document with the synopsis of the EECAV Roadmap Outline was distributed to all registered workshop participants prior to the workshop. A post-workshop questionnaire was sent to the workshop attendees to collect additional feedback.

The feedback from the workshop is summarized in this report.

OVERALL HIGHLIGHTS FROM WORKSHOP FEEDBACK

1. The AV problem involves a unique nexus of characteristics, namely, safety critical computation, a small footprint, low weight, large impact on society and long-term deployment (~ 15 years) in a car. This unique nexus should be front and center of future efforts for R&D in this arena.
2. A timeline of 2025-2035 for the R&D Roadmap Outline identified by the Organizing Team for AVs is appropriate (i.e., supported by the audience), since complementary metal-oxide-semiconductor (CMOS) scaling is coming to its end by 2025. Leveraging this activity with the International Roadmap for Devices and Systems (IRDS™) roadmap) community will be very important.
3. Energy efficient computing is critical for future AVs. Based on the survey given at this workshop, attendees responded that the desired power target for on-board computing devoted to AV operation is within the 50 - 500W range. Dissipating (i.e., providing cooling for) even 500W will be challenging in a car. This requires a significant improvement over the current AVs under development, where the computing system consumes a few kW, impacting the AV range.
4. A vehicle-to-infrastructure (V2I) or vehicle-to-vehicle (V2V) network can offload computing demand on the AV. However, due to the lack of reliable networks in certain environments, OEMs can't develop AVs that rely on the V2X for critical driving decisions. Self-sufficient computing on-board the vehicles is necessary.
5. Co-design is a high R&D priority for advanced computing for the demanding vehicle environment. Requirements on radiation/vibration/thermal stability, reliability and safety and security will need to be developed and met.
6. Reconfigurability and programmability are very challenging research topics, given the 10-20 year life cycle of vehicles. Approaches need to be developed to allow the future upgrade of artificial intelligence (AI) and machine learning (ML) algorithms and circuit hardware in the field.
7. Optimization of smart sensors with distributed computing and central processing may be critical for energy efficient AV computation. Biological systems have developed some ways to optimize energy efficiency for sensor/processing and models from nature can inspire new research approaches.
8. The Society of Automotive Engineers (SAE) J3016 levels (L1 – L5) fail to define what the high-level AVs will actually be doing. As a result, the quantitative connection between computing performance and AV functionality remains a challenge.
9. Hardened chips for AVs may be expensive from a supply chain perspective, given the limited initial volume of the chips needed in the AV sector. The need for a radiation hardened chip has increased significantly, particularly for driving at high altitude.

FEEDBACK ON OVERVIEW OF THE SEMICONDUCTOR PROJECTION AND AUTOMATED VEHICLE DEVELOPMENT PROJECTION

Data Movement: Data movement is extremely expensive and should be limited. Of note is the potential for analog processing systems. Typically, variables and constant values are stored in main memory and accessed when needed. In-memory compute-based technologies reduce the movement of data. Distributed systems may be necessary because there is no single space large enough inside the vehicle for all the computer equipment, and components may be located throughout the vehicle.

Small Initial Chip Market: A recurring theme in the feedback is that the initial chip market for AV is small, which is an incentive barrier for chip companies to make chips tailored for AV. This gap helped motivate Tesla to make its own chips.

Reconfigurability and Updating: Reliability of software updates is a major consideration, but one of the differences between in-vehicle systems and mobile devices is that car manufacturers will most likely be the primary application developers, making them first-party updates. New sensors will be much more advanced than the ones they replace, which will make software updates more challenging. The marketplace for mid-life upgrades and service replacements may or may not address these challenges.

Regulations: Implementation of AVs will be in a heavily regulated environment which will differ from country to country. Liability is also a "regulatory issue." For example, if a smart sensor does not detect an obstacle which leads to an accident, who is at fault in the case of an AV using that sensor? Issues like this need to be addressed in parallel with the AV technical development. There may be a regulatory need for an AV safety data recorder, creating requirements for sensors and interconnects, and favors about distributed processing. Regulators lag behind the AV industry in all regions and markets.

Smart versus Dumb Sensors: There is a lot of potential to perform computing at the sensor which can reduce the load on a centralized computer and the network connecting them. Smart sensors benefit from greater efficiency by using application specific chips, however this implies an even wider range of chips each with lower production volume, making them more expensive. Smart sensors will need to be continuously calibrated, be serviceable, and be compatible with the systems they're used in. It was noted that sensors that perform more computation can increase the sensor temperature, undermining their reliability, and adding packaging challenges.

Required Computing Electrical Power: Energy consumption for current L3-4 AV computer systems is roughly equivalent to the energy needed for propulsion in electric vehicles (EV), although this could be lower in some contexts (i.e., follower vs. leader in an AV train, or navigating enclosed environments, or on open highways). Feedback from the audience was that the power target for AV computers and sensors should be in the 50–500W range. It was noted that current AV under test has a few kW of computational energy requirements which will also be a major strain on an internal combustion engines (ICE).

FEEDBACK ON OVERVIEW OF THE EECAV ROADMAP SCOPE AND OUTLINE

Other Roadmaps: The Organizing Team was asked if we are relying / informed by other roadmaps in the computing community. We responded that we are considering other roadmaps, but are focused on defining the R&D problems for connected and autonomous vehicles (CAV), working to focus on those issues important for AV, and trying to capture the unique R&D problems that emerge for AV in the computational energy efficiency realm. There was a suggestion that we connect with the IRDS™ road mapping community to inform them of our AV efforts. There was general support that our doing a “Moore’s Law” like extrapolation to 2030 was reasonable, but after that it is uncertain.

Energy and Efficiency Issues: There continued to be feedback on the assumption that all the computation resides on the vehicle, or whether V2X (vehicle to vehicle, vehicle-to-infrastructure) communication should be considered. V2X communication might offer more energy efficiency, but the problem is availability, especially in remote areas. The required infrastructure is hard to build. Also, V2X communication is harder to secure, and OEMs would not be willing to share high-resolution data inherent in V2X.

Currently, perception processing dominates computation required for AV, but as non- “brute-force” algorithms are developed, this should become more efficient. Once pixel-based analysis gives way to semantic analysis, this energy requirement should decrease.

The computation energy requirements for AV change with speed and driving conditions. The systems will not be running at maximum power all the time, introducing the concept of a computational “duty cycle” that could shape the energy needed for AV computation. In general, the faster you go, the farther you need to look ahead of the vehicle, and this requires lower latency from sensing to actuation. Faster driving implies higher computational demand and, therefore, higher computational power.

It would be advantageous to look at a more holistic view for combining compute performance and energy efficiency. A discussion took place that examined whether it is more efficient to have smart sensors that have computing systems that manage a smaller load from each sensor, or a central computing system coupled to dumb sensors. This raised a basic question: is there a correlation between distributed computation and efficiency? It was noted that breaking a central computer up into several boxes might be easier to deal with (i.e., placing in the vehicle), regardless of the impact on computational energy efficiency. On the other hand, going to distributed computation (say for energy efficiency reasons) might introduce software issues as well. Distributed systems can be unbalanced or there can be more unforeseen problems introduced by distributing the work. It was noted that many biological systems do much compute locally.

AV Economics: The audience asked if the Organizing Team had discussions on the economics of efficiency. Are there advantages to be the most efficient versus the first to market? There was some discussion of the economics of V2X communication, namely wondering who foots the bill on building out the required communication infrastructure. It was pointed out that the Roadmap Outline being developed is meant to have a longer-term perspective.

The group discussed installation and design of high-performance computing systems and efficiency. The person who authorized the compute was not always the person who would pay the electric bill required to perform the computation. When the person paying the bill and needing the computation are the same, progress is made in energy efficiency, which involves examining the basic ownership model.

AV Security: Security issues increase as the reach of the communications network (V2V) widens.V2X). Furthermore, the hardware systems need to be secure across the supply chain. Secure compute, zero trust, and other options may be needed – but increased security comes at a cost, economically and in system complexity. A suggestion was made to let architects secure compute capabilities and extend them for safety. Formal methods may help. Side channel attacks should be considered.

Approach: The underlying approach of the Organizing Team, namely “Co-design” is extremely important, and optimizing the hardware to match the software will prove valuable.

FEEDBACK ON TECHNICAL AREA I: CHIPS: MATERIALS, DEVICES AND CIRCUITS

Problem #1: "New, or augmented, materials and processing for increased thermal/ mechanical/radiation robustness for automotive environments and long life-span"

Most of the feedback on this problem concerned radiation hardness. Basically, you don't want the brakes to fail when you are at high altitude or for the air bag to deploy because a bit got flipped. Prior approaches have used a high level of diagnostic and hard and soft approaches to allow chips to be used in space. As the transistors have gotten smaller, they have become more sensitive to flipping with radiation exposure, but they have a smaller physical cross section, so that reduces exposure per transistor. But usually, there are more transistors in a given area, which increases the problem. In the heavy mining industry, there has been a concern about radiation and computation for over 30 years. More and more people at the OEM / Tier 1 developer level are starting to incorporate radiation hardening.

The automobile problem is different than data centers in that radiation hardness for data centers isn't needed because chips are designed to operate at their best at the optimized data center conditions, so the chip companies won't want to move away from that. Perhaps a consortium of military and automotive people who need more radiation hard chips could be formed to move things further for the AV application. However, for automotive use, AV chips need modest radiation protection (not as strong as for military uses), and it would be a valuable strategy to build radiation protection in the entire "stack" from chip to sensor, including reconfigurability.

It was noted that for automotive use, there is a nexus of requirements that are unique, namely being safety critical, having limited volume (footprint), incurring a large societal impact, and long-term deployment on long-lived vehicles. Nobody is doing R&D that takes all these requirements into account, which probably means that resiliency research in other applications won't mitigate these challenges. Although radiation hardening has been an area of R&D, there is still a lot of work to be done, especially at the physics level. There is ongoing work incorporating resiliency in Field Programmable Gate Arrays (FPGAs). Using latched variables as a mitigation strategy at the hardware layer to provide resiliency could also reduce the latency of recoverable compute. This is not a full military spec, however. Furthermore, there is a performance loss, but doing this at the device level seems to be a way to reduce the performance issues.

Problem #2: "Low latency and low power devices and circuits for AV computing needs"

This AV problem has specific requirements which when solved would be very rewarding to others. A recommendation was made to reach out to the IRDS™ community to create a bridge to them on this topic.

There was a discussion of asynchronous computation, which provides low latency and energy efficiency improvements. It is very hard to build tools that understand an asynchronous circuit. It is also hard to build tests for asynchronous operation as well. These problems are not just due to a lack of research and/or investment. Verifying and validating asynchronous circuits is an extremely difficult problem. While there would be benefits, meeting regulatory guidelines for testing would be extremely difficult. It may be too difficult to verify unclocked logic.

Problem #3: "New computing circuits and devices simultaneously optimized for reconfigurability and high performance"

A question was raised about the need for in-situ reconfigurability. Does centralization really require reconfigurability? Would it be OK to not replace hardware occasionally? It was noted that hardware replacement/upgrades are standard practice for fleet vehicles, but this isn't viable for consumer vehicles. OEM internal research shows a strong need for reconfigurability. This may be an example of something falling in an R&D gap because people in industry tend to want other people to solve problems, so there is a chance this could be missed.

FEEDBACK ON TECHNICAL AREA II: CHIPS: ARCHITECTURE, SAFETY AND SECURITY

It was noted that "introspection," the ability of the system to recognize error and failures and triage, should be added to the R&D list. Introspection would also benefit security and the identification of aging parts.

For all the listed problems, there needs to be a definition of the CAV use case and the differences/unique challenges that a CAV will face. Energy restrictions, vibrations, heat, and physical distances are all challenges that when combined create a unique problem set for a CAV. Research should keep the difficulties and challenges of the CAV front and center to emphasize the work needed for this area.

Problem #1: "How do we best define and use distributed, heterogenous multiprocessor systems (CPUs, GPUs, Neural accelerators, etc.) to support the algorithms needed for EECAV?"

Overall, this was identified as a high priority area¸ and there is ongoing research into hardware that supports flexible algorithms. Functional safety is vital for any automotive system.

A key part of the R&D effort will be identification and quantification of the right metrics. Examples include latency performance and power, and a means of expressing the relationships between them. The metrics also need to be something the OEMs are willing to share with researchers and consider architectural constraints. For example, trillions or tera operations per second (TOPS)/Watt is too simplistic, a measurement such as "Effective TOPS" would be better. It could include the input dimensionality, model size, and frames per second (FPS) required. Other important considerations include reconfigurability, re-programmability, upgradeability, security, resilience and reliability. People are doing a balancing act, but not for the nexus of boundary conditions involved in AV.

There is a trade-off between what is possible to do with existing compute capabilities and what might be possible in the future. Hardware needs to be flexible enough for future algorithms. For machine learning, we need to consider training and training sets in addition to inference.

Problem #2: "What is the right network/interconnect for energy-efficient computation for AV?"

There was general agreement that this is a high priority area for funding. It will be important to construct a use case to differentiate the unique aspects of CAV. The possibility of using photonics came up quickly in the discussion. A question arose about the possibility of making an AV vehicle a safe enough environment to allow photonics. It was noted there is already some fiber optics in vehicles; they are fiber over plastic. It is not high bandwidth, and the bend radius is limited. Photonic (fiber optic) solutions are expensive and challenging to implement. Some OEMs use gigabit connections for data connectivity. It was also suggested that GB ethernet could be used or otherwise needs to be proven why it can't be used, or that another system similar to data center backplane could be used.

Moving data requires more energy than computing. A low-latency, low-jitter network for predictability is required. Networking is inherently connected to the memory system, so for example, should a communication process be used to move data or consider a global address space? How does information flow? How do different accelerators affect the answer?

Another option would be to “bypass” memory altogether and use a dataflow architecture. This would allow for higher efficiency. There is a concern about the system stalling – in this case, with current technology, the data is discarded. This might not be an issue, as the loss of a single frame of data might not be a major failure.

Problem #3: “How much memory (addressable and storage) at what bandwidth is needed/possible, and where should it be located in the system?”

This was considered a high priority area for funding. There are trade-offs in memory (e.g., persistent memory and others), that need to be focused on to account for safety, security, deteriorating parts, etc. We need to understand the trade-offs between different memory approaches in AV use cases. For example, cost and power/energy cost of various memories have to be addressed specifically for this application, and SRAM caching is likely to be used on die for process compatibility. It is possible to place extremely large SRAM caches either on a chip or next to a chip.

There is a huge role for shared memory - now we are faced with sending the data to every processor that needs it. If we flip this to on-demand, the problem becomes synchronization. Are supercomputer clusters the right topology? What about coherence? Are systems dynamically or statically configured? Standards like GenZ and CXL address some of that. Memory architecture was mentioned. What is a good computing model? For example, supercomputing cluster models in AI work as a good model. Can we make the computing robust enough? Is a unified memory model good?

The sensors are the right place for in-memory computing, but they are brittle designs while shared memory architecture is much more flexible. As technology evolves, the system can be retrofitted. There is also a tight coupling between memory system architecture, network architecture, and data flow through the system. These need to be thought of collectively while developing solutions.

FEEDBACK ON TECHNICAL AREA III: ALGORITHMS AND DATA MANAGEMENT

The feedback for Technical Area III covered two general topical areas. Feedback for these two areas is captured as follows:

Topic #1: Algorithms for Energy-Efficient CAV

- We had comments on near-sensor signal processing, data distillation, methods of co-optimizing algorithms and implementation platforms, and discussion about which tasks are likely to be most computationally intensive.
- Trade-offs among efficiency, throughput, and latency must be considered, and these measures are ultimately only meaningful at the system level.

- Near-sensor signal processing may be needed, which will influence algorithm design. It may be possible to eliminate standard signal processing stages and hardware components, feeding raw data into machine learning algorithms. Covariance shift (mismatches between machine learning testing and training data) result from changes to signal processing pipeline and reduce accuracy. Sensor fusion will help to identify unimportant data. Data compression will also be valuable. However, near-sensor signal processing and inference must avoid undermining system-level sensor fusion techniques.
- Most raw data are superfluous to accuracy, but efficiently determining which are important is a subject of research. Techniques for distilling data are important because they can reduce data volumes by orders of magnitude, and some existing approaches, e.g., object detection, can be viewed as distilling raw images to compact lists of objects.
- Dynamic object detection is computationally intensive and can't be done until perception work is finished. Hidden Markov models will also be important. Counting multiply-accumulate operations doesn't determine architectural requirements: memory access properties are important, and caching won't work for many of these problems.
- A suite of efficient machine learning frameworks and algorithms is needed. AutoML may help with algorithm implementation by enabling many models to be considered in the data center, where computation is less expensive than on the vehicle, but new work would be needed to handle the large, changing datasets in CAV workloads. Models should be optimized with awareness of architectural constraints.

Topic #2: Reliability and Security

- There are explainability advantages to the final stages of inference systems being based on straight-forward rules, although earlier stages will often use machine learning algorithms.
- Some systems have real-time deadline and state-of-the-art tools are pessimistic when estimating worst-case latencies for real-time applications.
- There is a need for reliable and fault-tolerant machine learning techniques, which might even enable simplified hardware by compensating for faults in the algorithm instead of in hardware.
- Multiple models may be used to cross-check and improve robustness.
- Black box recorders will need to delete some data, perhaps older or redundant data and deciding how to do this is both technically and legally complex.
- It is not well understood how much accuracy is lost by neglecting on-line learning. There is a need for research on proving properties of both conventional and online machine learning algorithms.
- Determining what tests are necessary for machine learning systems, especially online learning systems, is difficult. Baseline tests may by impractical in the field due to the huge number of edge cases required. Carefully selected or designed edge cases may be needed for the sake of safety.
- Vehicles must never rely on wireless communication for safety- and latency-critical tasks. That makes infrastructure a potentially useful secondary data source; sensors on the vehicle remain the primary sources and must be sufficient.

- Security is high-stakes because all vehicles can be weaponized and the cost of doing so will generally be lower for CAV than for conventional vehicles. The Highly Autonomous System Safety Center may have relevant knowledge. Access to historical data on attacks can help with attack detection.
- Mechanisms and incentives for CAV companies to share raw data with each other and with researchers will speed advances.

FEEDBACK ON TECHNICAL AREA IV: SENSORS DATA INTERFACE

Problem #1: "How smart should a sensor be? Smart about what?"

The benefits of smart sensors include lower latency and energy efficiency of not moving data to a remote processor. The downside includes fewer algorithm choices and increased energy consumption at the sensor, leading to problems with heat dissipation and accelerated degradation at the sensor.

Because smart sensors perform significant on-sensor computation, this work generates waste heat which can cause problems for sensor operation (i.e., noise, degradation). Vehicle designers are strongly opposed to adding large or conspicuous enclosures to a vehicle to accommodate larger sensors that can more easily dissipate heat.

Although the software in smart sensors can be updated, it is limited to the capabilities of the smart device. Nevertheless, smart sensors may provide excellent utility, especially when fast and deterministic response time or data reduction are required.

The importance of latency was not fully recognized in Tech Area IV.

Problem #2: "R&D on Data vs. Task Migration for Dynamic Power Management"

"Dark silicon" does exist, we've needed this capability for a long time; the question now is, "Does this ability need to be operating at every level of the abstraction?" At what point(s) do we need to expose the knobs for power control? Some cases require computational capacity to be restored almost instantly, other cases would allow for longer restart times.

Event-based sensors could be investigated, but there would again be challenges where the algorithms might be limited to those supported by the sensors.

Questions remain about how to know when to activate additional sensors and computers.

Problem #3: "Can R&D find ways of exploiting asymmetric bandwidth utilization of networks to improve energy efficiency?"

Networks typically provide symmetric bandwidth for sending and receiving data, however many devices in the vehicle have asymmetric bandwidth needs. This over-provisioning has significant energy and material cost.

Traditional Low Voltage Differential Signaling (LVDS) is (still) a good solution but has limitations with respect to video distribution and device sharing.

Problem #4: "What R&D is needed to anticipate advancements in sensors and computers over a long (~15 year) vehicle lifespan to maintain forward and backward compatibility?"

Service equipment may have an even longer lifespan than the vehicles themselves because so many components are common to many vehicles and OEMs.

The computing industry has created blade infrastructures which successfully anticipated 15 years of bandwidth growth. Separation of control and data planes, and media options for different applications are common to many of those successful systems.

In-vehicle cabling presents many challenges: harsh environment ruggedization, small bend radius and pass-through dimensions, and electromagnetic noise all lead to cost and safety issues.

NEXT STEPS

We are planning to engage with broader the community across public and private sectors and professional organizations to develop the full roadmap for Energy Efficient Computing for Automated Vehicles as the next step. We are looking for potential contributors and collaborators on the roadmap activity. Please reach out to us if you are interested in participating.

CONTACTS

Chris Moen, cmoen@sandia.gov

Lennie Klebanoff, lekleba@sandia.gov

WORKSHOP AGENDA

Sandia National Laboratories

ENERGY EFFICIENT COMPUTING FOR AUTOMATED VEHICLES

MAY 11-12, 2021 | ONLINE MEETING

WORKSHOP AGENDA

MAY 11, 2021

TIME	EVENT	SPEAKER
8:30-8:40 am (PDT)	**Welcome and Introduction** Purpose of workshop; desired outcome	Zhiyong Li, SNL
8:40-8:50 am	**Opening Remarks** National imperative on energy topics, transportation and energy systems	Sarah Allendorf, SNL
8:50-9:00 am	**Organizing Team Introduction** OEM/semiconductor/academic/national lab perspectives on energy efficient computing for automated vehicles	Organizing Team
9:00-9:40 am	**Overview of the Semiconductor Projection and Automated Vehicle Development Projection** Semiconductor industry landscape/projections and AV industry landscape/projections for the next generation	Matt Marinella, SNL Jace Mogill, USCAR
9:40-10:00 am	**Feedback on overview**	All participants
10:00-10:10 am	**Break**	
10:10-10:40 am	**Overview of the Roadmap Scope and Outline** Discussion of the philosophy behind the roadmap (timeline and scope), computational energy efficiency from the supply side (chip technology) and the demand side (vehicle implementation)	Lennie Klebanoff, SNL
10:40-11:30 am	**Feedback on Roadmap Scope and Outline** • Will majority of the computation be on the vehicle? • What is the sustainable power limit for onboard computing? • Can CMOS meet AV need for 2030 and beyond?	All participants
11:30-11:50 am	**Feedback Summary**	Carrie Burchard, SNL Lennie Klebanoff, SNL
11:50-12:00 pm	**Preview of Day 2 Activity** **Adjourn for Day 1**	Zhiyong Li, SNL

MAY 12, 2021

TIME	EVENT	SPEAKER
8:00-8:05 am (PDT)	**Introduction of Day 2** Expectation and Outcome	Lennie Klebanoff, SNL
8:05-9:05 am	**Technical Area 1 - Materials, Devices and Circuits Feedback Session**	John Paul Strachan, HPE Session Chair
9:05-10:05 am	**Technical Area 2 - Architecture, Safety and Security Feedback Session**	Rob Aitken, ARM Session Chair
10:05-10:15 am	**Break**	
10:15-11:15 am	**Technical Area 3 - Algorithms & Data Management Feedback Session**	Robert Dick, University of Michigan Session Chair
11:15-12:15 pm	**Technical Area 4 - Sensor Data Interface Feedback Session**	Jace Mogill, USCAR Session Chair
12:15-12:30 pm	**Review Feedback Summary** We'll reflect back to the group the feedback we captured. • Did we miss anything? • Did we capture the feedback correctly? • Additional questionnaire • Next Steps	Carrie Burchard, SNL Lennie Klebanoff, SNL
12:30-12:35 pm	**Closing Remarks**	Chris Moen, SNL

This page left blank

Sandia National Laboratories

arm

UNIVERSITY OF MICHIGAN

Carnegie Mellon University

Hewlett Packard Enterprise

USCAR
UNITED STATES COUNCIL FOR AUTOMOTIVE RESEARCH LLC