SANDIA REPORT
SAND2019-14177
Printed November 2019

Workshop on Advanced Computing for Connected and Automated Vehicles

C. Mailhiot
W. M. Severa
C. D. Moen
T. B. Jones (Draper)

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico
87185 and Livermore,
California 94550

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

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/

ABSTRACT

To safely and reliably operate without a human driver, connected and automated vehicles (CAVs) require more advanced computing hardware and software solutions than are implemented today in vehicles that provide driver-assistance features. A workshop was held to discuss advanced microelectronics and computing approaches that can help meet future energy and computational requirements for CAVs. Workshop questions were posed as follows: will highly automated vehicles be viable with conventional computing approaches or will they require a step-change in computing; what are the energy requirements to support on-board sensing and computing; and what advanced computing approaches could reduce the energy requirements while meeting their computational requirements? At present, there is no clear convergence in the computing architecture for highly automated vehicles. However, workshop participants generally agreed that there is a need to improve the computing performance per watt by at least 10x to advance the degree of automation. Participants suggested that DOE and the national laboratories could play a near-term role by developing benchmarks for determining and comparing CAV computing performance, developing public data sets to support algorithm and software development, and contributing precompetitive advancements in energy efficient computing.

ACKNOWLEDGEMENTS

This report was prepared for the U.S. Department of Energy (DOE), Office of Energy Efficiency and Renewable Energy (EERE), Vehicle Technologies Office (VTO).

The Workshop on Advanced Computing for Connected and Automated Vehicles was held at the Lawrence Berkeley National Laboratory in Berkeley, California, on May 7, 2019. The workshop was sponsored by VTO's Energy Efficient Mobility Systems program. The technical organizing committee was:

David Anderson	U.S. Department of Energy, Vehicle Technologies Office
Christopher Moen	Sandia National Laboratories
Christian Mailhiot	Sandia National Laboratories
Zhiyong Li	Sandia National Laboratories
Robert Hwang	Sandia National Laboratories
Rahul Chopra	Lawrence Berkeley National Laboratory
Ramamoorthy Ramesh	Lawrence Berkeley National Laboratory
Daniel Armbrust	Lawrence Berkeley National Laboratory
Troy Jones	Draper Laboratory
Eric Balles	Draper Laboratory
Claus Daniel	Oak Ridge National Laboratory
Eric Rask	Argonne National Laboratory

Report authors are Christopher Moen, Christian Mailhiot, and William Severa of Sandia National Laboratories and Troy Jones of Draper Laboratory[1].

[1] Draper is a not-for-profit engineering innovation company. Draper works with government, commercial companies and academia to design, develop and deploy advanced technology solutions for the world's most difficult and important problems.

CONTENTS

LIST OF FIGURES

This page left blank

EXECUTIVE SUMMARY

To safely and reliably operate without a human driver, connected and automated vehicles (CAVs) require more advanced computing hardware and software solutions than are implemented today in vehicles that provide driver-assistance features. But the energy cost of these advanced solutions might not meet vehicle size, weight, and power constraints unless computing power consumption and cooling loads are reduced. The mobility industry is searching for approaches in automotive architectures, advanced computing, and microelectronics that will meet the technical demands of highly automated driving in an energy efficient manner.

The Workshop on Advanced Computing for Connected and Automated Vehicles invited representatives from industry, academia, national laboratories, and the federal government to discuss connected and automated vehicle (CAV) computing challenges in a precompetitive forum. The main goal of the workshop was to help identify an early stage research role for the Department of Energy (DOE) that would propel the development of energy efficient CAVs forward. The workshop supports the DOE Vehicle Technologies Office's (VTO) Energy Efficient Mobility Systems Program vision for an affordable, efficient, safe, and accessible transportation future.

The workshop held cross-disciplinary panel discussions on advanced microelectronics and computing approaches that can help meet future energy and computational requirements for connected and automated vehicles. Workshop questions were posed as follows: will highly automated vehicles be viable with conventional computing approaches or will they require a step-change in computing; what are the energy requirements to support on-board sensing and computing; and what advanced computing approaches could reduce the energy requirements while meeting their computational requirements?

An observation from the workshop was that industry is currently focused on solving the algorithmic and software problems of perception and inference, setting aside optimization of computing hardware and power consumption to be addressed later. Participants largely agreed that a "step-change" in computing technology will be necessary to enable commercially viable highly automated driving. They agreed that a more energy efficient alternative to current computing approaches may be required to efficiently process the machine learning algorithmic calculations that CAVs will perform. Participants suggested that CAV platforms and hardware should remain flexible to implement evolving algorithms and workloads.

To overcome these computing challenges, a broader dialogue is necessary about long-term research needs and the role for DOE and its national laboratories. It was suggested that DOE could play a near-term role in developing benchmarks and standards (in partnership with industry) for determining and comparing "computing performance" per watt. DOE could develop public data sets to support algorithm and software development. Analogous to the VTO battery program, DOE could co-create a roadmap with industry to help drive and quantify precompetitive advancements in computing performance.

ACRONYMS AND DEFINITIONS

Abbreviation	Definition
ADAS	advanced driver assistance system
AI	artificial intelligence
AV	automated vehicle
CAN	controller area network
CAV	connected and automated vehicle
CMOS	complementary metal-oxide semiconductor
CPU	central processing unit
DARPA	Defense Advanced Research Projects Agency
DOE	U.S. Department of Energy
EEMS	Energy Efficient Mobility Systems
FPGA	field-programmable gate array
GPS	global positioning system
GPU	graphics processing unit
IMU	inertial measurement unit
LIDAR	light detection and ranging
MOSFET	metal-oxide-semiconductor field-effect transistor
OEM	original equipment manufacturer
PC	personal computer
SWaP	size, weight, and, power
TOPS	trillion operations per second
VTO	Vehicle Technologies Office
V2X	vehicle-to-infrastructure

1. WORKSHOP INTRODUCTION

To safely and reliably operate without a human driver, connected and automated vehicles (CAVs) require more advanced computing hardware and software solutions than are implemented today in vehicles that provide driver-assistance features. But the energy cost of these advanced solutions may negate anticipated system-level mobility efficiency benefits unless computing power consumption and cooling loads are reduced. More specifically, today's computing technology might not meet future size, weight, and power constraints given the limited amount of energy stored on the vehicle. The mobility industry is searching for approaches in automotive architectures, advanced computing, and microelectronics that will meet the technical demands of highly automated driving in an energy efficient manner.

The U.S. Department of Energy's (DOE) Vehicles Technologies Office (VTO) held a one-day workshop on May 7, 2019, in Berkeley, California, focused on advanced microelectronics and computing approaches that can help meet future energy and computational requirements for connected and automated vehicles (CAVs). The workshop addressed these questions:

- Will highly automated vehicles be viable with conventional computing approaches or will they require a step-change in computing?
- What are the energy requirements to support on-board sensing and computing for highly automated vehicles?
- What advanced computing approaches could reduce the energy requirements for highly automated vehicles while meeting their computational requirements?

While the workshop scope was originally intended to address computing challenges for both the connected and automated trends in vehicles, the discussion throughout the day gravitated towards the topic of vehicle automation. The computing challenges associated with vehicle connectivity were not the focus of the workshop participants, and their overall sentiment was that the infrastructure required to support connectivity would evolve slower than automated vehicle technologies.

We use both automated vehicle (AV) and connected and automated vehicle (CAV) acronyms in this report. When the discussion involves broader vehicle-to-X communication implications, we use the term CAV; otherwise, AV is used.

1.1. Workshop Goals

The workshop supports the VTO Energy Efficient Mobility Systems (EEMS) Program's vision for an affordable, efficient, safe, and accessible transportation future. The EEMS Program conducts early-stage research and development at the vehicle, traveler, and system levels, creating new knowledge, tools, and technology solutions that increase mobility energy productivity for individuals and businesses.

There were two goals for the workshop:

- Identify the highest priority avenues of precompetitive research in energy-efficient sensing and computing that would augment industry's current research plans.
- Determine if there is a role for the federal government to advance energy-efficient computing hardware and software for highly automated vehicle technology.

The workshop brought together experts from industry, national laboratories, and academia in a precompetitive forum to discuss three topics that would inform the goals:

- Vehicle System Requirements: system sensing and computing architectures in highly automated vehicles and the power those technologies may consume
- Computing and Algorithms: algorithmic designs and software that, when integrated with advanced materials and devices, may reduce CAV power demand
- Microelectronics Hardware: advances in microelectronic materials and devices and how their integration may enable more energy-efficient computing in CAVs.

The program consisted of panel discussions on the three topics as well as keynote presenters from DOE and the computing industry who provided their observations on the current state of CAV technology and its trajectory in the short and long term. The panelists were selected to provide a diverse set of institutional perspectives for each of the three topics. Panelists were given the opportunity to provide their perspective individually, address pre-selected questions, and enter dialogue introduced by the audience's questions. Additionally, workshop organizers invited audience members to initiate follow-on discussions with the panelists and national laboratory research staff outside of the workshop.

The 79 registered participants included researchers, executives, and representatives from the following areas:

- The automotive industry
- The microelectronics industry
- Automated technology subsystem developers
- National laboratories, U.S. DOE, and other federal agencies
- Universities

1.2. Computing Energy Efficiency -- Background

The evolution of silicon metal-oxide-semiconductor field-effect transistors (MOSFETs) laid the foundation for the modern information age of digital processing and consumer electronics. The most important benefit to emerge from that evolution was the improved energy efficiency of information processing (approximately a 1.4x improvement per year over 40 years) that enabled ever increasing computing power for the same energy consumption and spatial footprint. Virtually all information processing occurs in a network and can loosely be divided into two categories: "core" functions (e.g., high-performance computing on centralized supercomputers or cloud computing in today's vast server farms) and "edge" functions (e.g., point-of-sensing or point-of-actuation interactions with the physical world).

Core information processing has benefitted from energy efficiency improvements. A computing system's performance (operations per second) is determined by its ability to scale and densely pack computational elements. Computing performance is limited, in part, by heat dissipation. Energy efficiency improvements reduce the amount of heat dissipated, allowing for performance increases.

Energy efficiency is key to information processing in general, but it plays a more influential role at the edge than at the core. At the edge, information processing is embedded and must function in a highly constrained physical world. Limits to size, weight, and power (SWaP) can degrade performance. Cooling and waste heat removal energy loads and system footprint should be minimized to meet SWaP constraints, so SWaP is determined in large part by energy efficiency

Designers of highly automated vehicles need to develop information processing systems that meet both SWaP and latency constraints, and the computing demands on these systems are expected to increase exponentially to petaflop levels in the coming decade. At the same time, the energy efficiency of computing technology is not expected to increase to meet these power requirements but will instead plateau because the microelectronics scaling technologies are approaching their physical and economic limits[2]. It will be difficult to achieve low-energy processing for decision-making at the point-of-need in these SWaP-constrained vehicles.

CAV developers must address these energy efficiency challenges to realize the potential of highly automated mobility. On-board sensors and data processing consume a large portion of the energy stored on the vehicle for early CAV implementations. For highly automated vehicles to be viable—with their heightened computing demands—the CAV development field will need to apply low-energy processing beyond what is achievable with today's complementary metal-oxide semiconductor (CMOS) technology. Consequently, energy-efficient computing will be needed for reducing CAV power demand. It is likely that both novel materials, devices, and computing architecture innovation will be needed in the future.

[2] For more information, see "THE INTERNATIONAL ROADMAP FOR DEVICES AND SYSTEMS: 2018", IEEE

1.3. A Decade of Automated Driving -- Background

Commercial interest in the development of AVs for public use began in 2009, just after the completion of the Defense Advanced Research Projects Agency (DARPA) Urban Challenge race event in 2007. As seen in Figure 1.1, Google was an early innovator in the field. Google hired the leadership from the top Urban Challenge teams and quickly began an aggressive program to mature automated vehicles with testing on public roads.

In the early period from 2009 to 2012, the Google team followed a computing architecture paradigm that began in the DARPA Challenge: install as much commercially available personal computing (PC) in the trunk of the vehicle as can be fit within the space available. This approach prioritized flexibility in software development and algorithmic capability over computing power consumption. Standard PC processing allowed for straightforward software development on open-source operating systems (primarily GNU/Linux) using the languages and compilers popular with the developers.

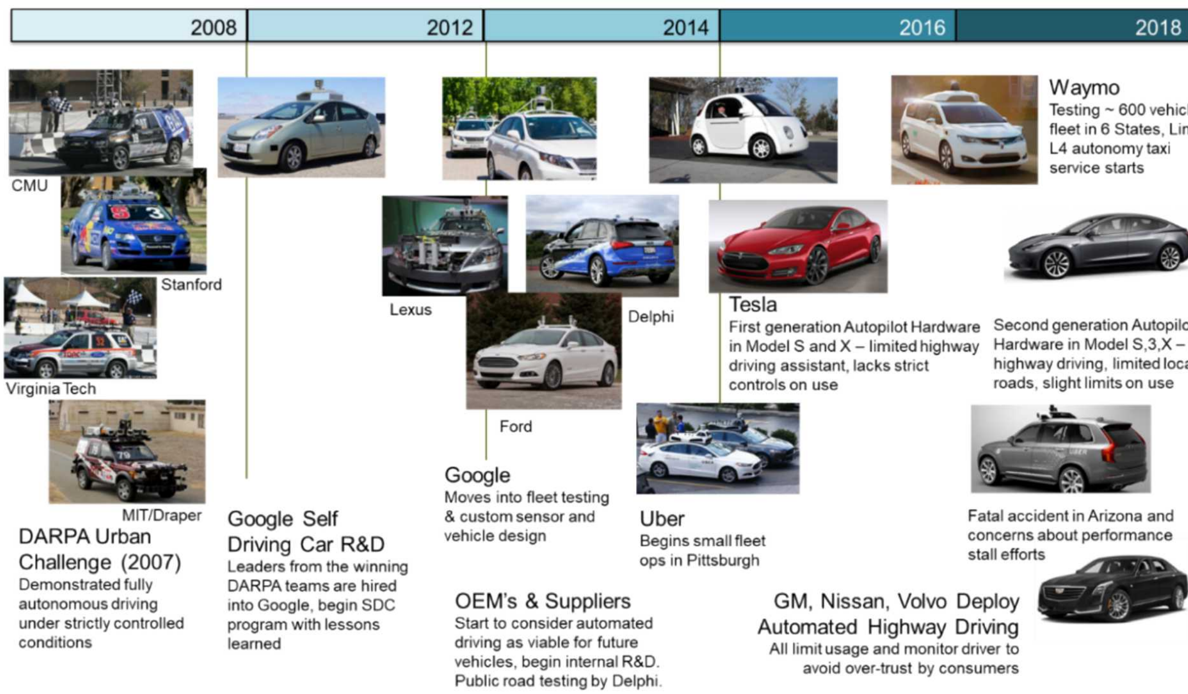

Figure 1.1 Timeline of AV commercial development. Google's early entry into the market showed other companies a path to maturing the technologies, and today automotive companies are generally pursuing two parallel technology paths: assisting human drivers in limited driving conditions (e.g., highway driving) and automated driving with no human intervention.

The DARPA race vehicles varied widely in computing power. Some teams completed the race using little more than a single desktop-class PC (a few hundred watts), while others employed a rack of servers consuming thousands of watts of power. The large disparity in computing designs mostly derived from the complexity of the perception algorithms developed at the time. The DARPA races simplified real world urban traffic down to limited scenarios on a closed course. These controls allowed a vehicle with a very limited understanding of the world to race successfully, but that limited understanding would not scale to the real world of driving in traffic. Some teams developed sophisticated methods of processing LIDAR, camera, and radar data to create detailed dynamic models of the world around the car. These algorithms required far more computational power and formed the basis for the early efforts in automated driving.

Companies like Ford and Delphi, who had both supported their own DARPA race teams, were some of the first to declare their automated vehicle research efforts publicly. Other automotive original equipment manufacturers (OEMs) and major suppliers followed, and today several OEMs and Tier 1 suppliers, together with numerous startups leveraging innovative sensor and computing processes, are developing highly automated systems for commercial use.

1.3.1. Advanced driver assistance systems architecture

At the time of the DARPA Challenge, automotive companies had already invested many years on improving "advanced driver assistance systems" (ADAS) technologies. These include limited automated driving safety features—such as automatic emergency braking and lane keeping assist systems—which have proven popular with consumers and profitable for sale in mass-market cars. The driver has the ability to turn ADAS features on and off and is responsible for maintaining control of the vehicle at all times, much the same as has been common practice since the introduction of cruise control. These systems are effective under most highway driving conditions but are designed for, and require, constant human supervision for safe operation.

An important feature of sensors for ADAS application is that they are self-contained—they process all their data locally and provide only a reduced information signal to the host vehicle. For example, a typical ADAS visible light camera is a sensor that can identify people, cars, and street signs with its on-board software and provides only this "object level" data to the rest of the car. This "distributed" (or "edge") model of computing is common in vehicle systems today. While it has critical disadvantages, it does allow the central computing system to operate at relatively low power, both in terms of processing and electrical load.

1.3.2. Automated driving system architecture

Unlike the ADAS components, automated vehicle sensing components such as LIDAR, high-definition electro-optical and infrared cameras, and high-accuracy IMU coupled with a precision GPS receiver are "dumb"—they typically have very little local computational processing capability. Instead, their purpose is to stream the "raw" data from the sensor directly to a powerful central computer. This centralized computing architecture affords a great deal of flexibility in software development for automated driving. Controlling each level of sensor data processing is important for many types of machine learning algorithm approaches, and since new machine learning algorithms are constantly being researched, a centralized processing architecture allows for rapid updating. Conversely, ADAS type distributed computing is often running "locked down" software that only the vendor can modify, slowing development progress.

The major downside to centralized computing systems in today's AVs is their high electrical power needs, which results in more heat dissipation and high integration costs to actively cool the computer. A representative centralized computer with GPU acceleration for machine learning may have 1000 times more computing power than its 5-watt, air-cooled ADAS counterpart but also consumes hundreds of watts of power. This computer requires active liquid cooling to operate in the high temperature ranges in a typical automotive application and creates more opportunities for component failures over time.

2. WORKSHOP OBSERVATIONS AND DISCUSSION

The panel discussions and industry keynote addresses provided an overview of the state of automated vehicle development and a sense for current industry focus. The panel and audience participants freely shared the challenges that they find most important. Industry is currently focused on solving the algorithmic and software problems of perception and inference, setting aside optimization of computing hardware and power consumption to be addressed later. Participants largely agreed that a "step-change" in computing technology will be necessary to enable commercially viable highly automated driving. While participants did not directly discuss computing and power consumption requirements, they agreed that a more energy efficient alternative to current computing approaches may be required to efficiently process the machine learning algorithmic calculations that CAVs will perform. Still, participants suggested that CAV platforms and hardware should remain flexible to implement evolving algorithms and workloads.

The following analysis of the workshop discussions and findings is offered in three sub-sections. First, the general themes that evolved in each of the three panel sessions are identified. Second, the nine general observations that have common threads across the panels are discussed. Third, the participants suggested several opportunities for the federal government to impact near-term AV technology.

2.1. Panel Discussions

Each panel was provided a set of framing questions to prime the discussion, though panelists were free to introduce any information they felt was relative to the overall workshop goals. The audience provided additional content for thought through the question and answer periods.

The Vehicle System Requirements panelists represented automated driving system developers. The framing questions to the panelists were:

- Thinking about achieving highly automated vehicles—on any road, any weather, with no driver—does your team feel constrained by available mobile computing today to achieve that goal? And are those constraints SWaP or are they more computational?
- In developing self-driving computing architecture, what does your team see as the right balance between "edge" computing power—typically done right at a sensor head—and centralized computing?
- If you are able to share, how much energy does your self-driving system use today and how much computing power does it have?

The themes that emerged in the Vehicle System Requirements discussion were safety and reliability, algorithmic approaches, and computing architecture and integration. Several key technology paths in AV development are still being explored. Many computing architectures are in use with many combinations of perceptual and decision-making algorithmic designs. Today these groups are making continuous progress towards automated driving with the technology available. Due to the tightly coupled nature of computing and algorithm design, it is not yet clear what approach, or approaches, will ultimately be the most successful, but all groups agree on the goal of improving the safety and convenience of mobility.

A significant insight from the panelists was the high variability in their own perceived needs for advanced AV computing architectures. In general, AV developers are able to continually improve system performance with the computing they have today. Panelists acknowledged, however, that the limitations of today's architectures will eventually impede them from making their highly automated systems production ready, and that they would capitalize on more powerful computing systems if such were available. Currently, AV developers are struggling with challenges related to corner-cases (responding to rare situations and events) in perceptual reasoning and driving logic and are less concerned with the absolute power efficiency of the computing solution.

The Computing and Algorithms panelists represented industries and academics that develop algorithm and software design approaches with today's computing architectures. The framing questions to the panelists were:

- Are current algorithms sufficient for highly automated driving, or do you expect revolutionary algorithm improvements to be necessary?
- What is the state of real-time, adaptive algorithms?
- What are the opportunities for co-design and heterogeneous integration?
- What would form a pathway towards improving trust in AI systems?
- What are the implications of deep machine learning and AI on computing and sensing?
- How can hardware (computing or sensing) impact or limit algorithm development?

The panelists examined the impact of computing and sensing hardware on algorithm implementation, and vice versa. They also discussed how automated driving systems currently implement machine learning and how artificial intelligence (AI) might evolve to assist with AV development. It can be particularly challenging to forecast what algorithmic innovations will emerge for AV perception and decision-making, however participants put forth ideas to that end.

It is worth noting that although the issues of cybersecurity and trust in CAV systems do not seem to implicate the workshop's energy efficiency focus, participants raised those issues. Specifically, there was clear agreement that CAV systems must be developed with integrated cybersecurity and trust, but there was no consensus on the approach to ensure these qualities. This topic may be well served by further discussion in a separate forum dedicated to cybersecurity in CAVs.

The Microelectronics Hardware panelists represented industries that develop computing and information systems. The framing questions to the panelists were:

- What performance improvements based on CMOS technology is the semiconductor industry targeting for CAVs? What is the timeline to achieve these improvements?
- What disruptive and non-conventional computing and sensing technologies are needed to meet future requirements for highly automated vehicles?
- How should the key performance index be defined for future microelectronics and hardware for CAVs: TOPS/watt or other measures?
- What is the performance/power target required for the widespread adoption of CAVs?

At present, there is no clear convergence in the computing architecture for highly automated vehicles. What is currently unclear among AV developers is whether distributed edge computing or central core computing will be the right path towards the low-energy solution. However, workshop participants generally agreed that there is a need to improve the computing performance per watt by at least 10x to advance the degree of automation. The field lacks consensus on what metrics should be used for the computing power performance, and what new benchmarks should be used for measuring the performance.

Three additional research themes emerged. A particularly promising area for optimizing energy efficiency is in the early stages of data processing to enhance cognition at the point of sensing to reduce computing demands on a centralized vehicle computing system. In the area of algorithm development, AV developers can forge a better understanding of the difference between reinforcement learning and conventional control theory, which should foster a better understanding of how decision boundaries form. There is also a need for improved software stack maturity to enable testing and benchmarking of new algorithms and hardware technologies.

2.2. General Workshop Observations

While each of three panels was focused on a different part of the CAV computing energy efficiency question, there were nine observations that emerged across the panel and audience discussions. None of the observations directly address computing energy efficiency, rather they reflect the near-term issues of most importance to the workshop participants.

2.2.1. Safety is the highest priority for AV developers

A major motivator for AV development is the potential reduction of injury and deaths on our roads. In the United States each year, almost 40,000 people die in vehicle crashes and millions more are injured. While these losses are tragic and large, it is important to put them in the context of how much humans drive each year. Around 288 million drivers in the U.S. accumulate over 3 trillion miles of driving each year in diverse environmental and road conditions, translating to one fatal accident every 85 million miles driven. AV designers are searching for new approaches to verify safety of their vehicles with a relatively small amount of experiential data. By comparison, the entire Waymo fleet—a leader in on-road testing—has only just passed a total of 10 million miles over the past 10 years in relatively uniform driving conditions.

Participants acknowledged that even the best machine learning systems today still cannot recognize objects and understand their semantic context well enough to reliably allow for the removal humans from driving. These methods are constantly improving, but it is an open question as to when the gap between human and machine perception will close. Moreover, it is an open question as to how systems will identify the failure to recognize objects and safely compensate for them.

2.2.2. Even today's most advanced machine-learning systems cannot always recognize objects and understand their semantic context

During the time of the DARPA races, techniques for processing sensor data into perceptual understanding were primarily based on heuristics and modeling algorithms that required the developers to explicitly encode how to interpret the incoming data. These techniques were time consuming to develop and often not nimble to changing situations. But the past 10 years have seen the rapid advance of machine learning-based sensor processing algorithms. Now most (if not all) AV projects use machine learning at least for perceptual reasoning and some are even experimenting with these methods to make path-planning and driving decisions.

Machine learning algorithms are the focus of many industry computing development efforts. There are two main phases to machine learning: training and inference. During the training phase, the algorithms feed on thousands (or even millions) of example data sets that are "labeled" with truth information (*e.g.* images of pedestrians in different settings and activities). This training phase is primarily performed "in the cloud" as it requires massive data-center-scale computational resources to perform and optimizing the results requires several steps. Once the machine learning "neural network" has been trained, that code can be loaded onto a much smaller computing system to perform the task of "inference"—pedestrian recognition, for example—in real time.

Machine learning algorithms today are approaching human-level accuracy in certain tests for recognizing many types of objects, but failures of these algorithms do occur, and in the context of AVs, an algorithmic failure can be fatal. Failures in machine learning are currently difficult if not impossible to diagnose; the parameters of a neural network were not designed by a person, but rather they "evolved" in a sense based on their training. Humans cannot yet be safely removed from all driving tasks and totally new algorithmic approaches may be required. Advancements in explainable machine learning and AI would be beneficial.

2.2.3. It's a software problem, first

Panelists emphasized that the algorithms for AV applications are still in development and that it is difficult to forecast which ones will be implemented for highly automated driving. As such, having excessive focus on computing operations per watt, or on hardware acceleration of a particular algorithm unnecessarily, will mislead developers on which factors are actually helping or harming AV computing performance and efficiency. Developers should resist the temptation to begin developing specialized computing hardware according to current algorithms until it is proven those algorithms provide a sufficient level of safe and reliable automation. Optimizing computing hardware and power consumption is a problem that industry has largely been setting aside to address later.

2.2.4. Designing automated vehicles is a "whole system" challenge

While it is difficult for any individual organization to achieve vertical integration in the multilayered problem of AV energy efficiency, it was offered that good solutions will need to take a "whole system" approach. No single improvement is by itself sufficient to enable highly automated vehicles. The development field will need to consider the web of interdependencies connecting computers and sensors, individual vehicles and collections of vehicles, inference (on-board) and training (offline). In the domain of algorithms, there is necessary interplay among at least five tiers: recorded data, simulated data, integrated algorithms and sensing, high-speed training, and low-power inference.

2.2.5. A "step-change" in computing technology will be necessary to enable highly automated driving

There was clear agreement that AVs will require significant computing improvements via a step-change—both in on-board and data center operations. One estimate provided that a Level 5[3] CAV would consume approximately 300–500 trillion operations per second (TOPS) for computing. GPU-based systems are not expected to achieve this performance level within reasonable power constraints of on-board energy storage. Panelists viewed step-change as necessary for long-term computing solutions, but platforms must be flexible to changing algorithms.

Today, the AV industry is leveraging specialized computing to accelerate the execution of machine learning training and inference. Perhaps the most popular system today is based on GPU chips. These chips are capable of executing the large-scale matrix calculations that are required by machine learning algorithms far more quickly than a general-purpose CPU. However, these chips are higher cost than other processor types and still consume a significant fraction of energy stored on the vehicle.

The selection of the computing architecture is also tightly coupled with the networking infrastructure in the vehicle. For limited production highly automated vehicles, the sensors are interconnected via high bandwidth networks—commonly Gigabit Ethernet. This network allows the high-performance centralized computing approach to be viable because it enables the ingestion of raw data from each sensor. The automotive industry, however, has been hesitant to move away from the reliable and very cheap Controller Area Network (CAN) serial bus network. The very limited bandwidth in a CAN bus is only compatible with a very distributed computing architecture, where each sensor node performs almost all of its functions independently. For a truly mass-market AV, the cost of every component is scrutinized, and new low-cost, high-speed networking solutions may ultimately be required.

2.2.6. Computing advances within the datacenter will be required

A suggestion from the workshop was that the machine learning neural network models will need to be updated regularly as low-probability events are encountered. However, this will become prohibitive because increasingly large, power-hungry datacenter-scale systems are required for this training operation. Hence, energy-efficient computing advances within the datacenter (not just on-board) will be required as well.

2.2.7. Radar sensors are relatively underutilized

There is no "ideal" sensor for automated driving. The number and diversity of sensors is a choice made by an OEM based on risk and cost. OEMs integrate multiple types of sensors partly to minimize common-mode failures. Today, AV perception systems rely heavily on LIDAR and camera sensors, but LIDAR solutions have severe limitations in certain lighting and moisture environments. Several participants noted that radar was relatively underutilized, and compared to the investment in LIDAR startups, radar startups currently receive very limited investment. Workshop participants generally agreed, however, that fusing multiple modalities together—radar, LIDAR, and cameras—is the best overall solution to make the safest and best performing perception systems.

[3] SAE J3016 "Levels of Driving Automation"

2.2.8. The optimal sensor suite will balance resolution, power, redundancy, reliability, and cost

Even though there is no general consensus on the best low-energy approach to advance AV computing, it is widely believed that it is essential to optimize every step of the information lifecycle, from acquiring the data, to aggregating the data, analyzing the data, providing the insight, and acting on the information.

Participants suggested that current imaging sensor systems are inadequate for highly automated vehicles, particularly in adverse weather conditions such as fog and snow. An opportunity for improving this area, however, may lie in better models of light propagation/scattering and image distortion, combined with machine learning to correct the image.

Participants agreed that an optimal sensor suite would balance resolution, power, redundancy, reliability, and cost. Multi-modality sensors, radar beamforming, thermal imaging, radio frequency lensing and imaging, long-range camera, and novel optical field sensing are among the areas that would benefit from research investment.

Early processing and data fusion of the sensor data using neuromorphic non-Von-Neuman computing (analog rather than digital) is another research area that deserves substantial attention. The colocation of sensor material, computing memory, and logic can reduce the energy cost associated with moving data.

2.2.9. Early CAV deployments will rely on limited external perception data

Connectivity is an area that was largely dismissed by many AV developers in attendance. The standard example of this is the use of vehicle-to-everything (V2X) infrastructure and mobile radios and newer concepts use available cellular networks and radios. Developers are not confident that the time and investment will be put forth to develop next-generation infrastructure to support V2X communication for CAVs. Still, it was proposed that connectivity should be included as a system level consideration for the power argument and that any data from an external network would be used when available.

U.S. roadways and other infrastructure already require significant investment to maintain their current state. Participants found it improbable that local governments and other entities would bear the cost of V2X-enabled infrastructure in the near term. As the availability and reliability of the network and the data remain too much of an unknown factor, every AV must operate without it safely, participants agreed. A partial exception to this, however, are the "high definition" digital maps that most AV systems require for operation. These maps provide road geometry, speed limits, and even navigation features to assist AV operation and are typically updated over a network connection to the vehicle. These maps cannot be assumed to be 100% accurate, so the vehicle must still be able to recognize the current roadway and special cases in real time, such as construction zones or police officers modifying traffic flow.

2.3. Roles for the federal government suggested by participants

Four suggestions for how the federal government could help develop advanced computing technology were offered by the participants. The suggestions are largely focused on the near-term needs of the industry. Additionally, advanced modeling and simulation technologies could play a key role in any of the four ideas.

2.3.1. Act as a third party to help establish reliable benchmarks and standards

A recurring theme was how to compare computing options, as the widely used metric of trillion operations per second/watt (TOPS/watt) was not well suited to compare across processor types or system configurations. The TOPS metric computed for a GPU (for example) cannot be easily compared to the TOPS value for an FPGA or a general-purpose CPU since the metric only captures the performance in one scenario of one function. For example, the TOPS/watt figures that manufacturers often cite fail to account for data movement in an application; they tend to generate the TOPS/watt using aggressive batch settings—a condition that may be unrealistic for AVs.

The point was also raised, that the field needs to recognize the costs associated with communication. Communication external to CAVs, across a CAV, and within a computing platform are all critical components poorly represented by metrics today.

The participants see DOE as a trusted facilitator of diverse technology contributors. By convening stakeholders to develop effective metrics and common benchmarking (for both hardware and algorithms), DOE can support and extend the development of necessary, enabling CAV technologies. The field's current lack of uniform metrics and benchmarking hinders the ability of system developers to make confident purchasing and development decisions in CAV technologies. Industry has a vested interest in commercially compelling products today, and DOE leadership can motivate long-term, high-impact technologies and help establish uniform standards across the field.

2.3.2. Co-create an energy-efficient computing roadmap with industry

Participants suggested a DOE contribution could be coordinating the development of computing efficiency targets and a timeline to achieve these targets in partnership with OEMs, computer platform developers, and sensor developers. Similar approaches have been valuable to the DOE solar energy initiative and the automotive battery program.

2.3.3. Develop robust public datasets

While some public datasets for AVs exist (*e.g.*, Berkeley DeepDrive dataset), participants suggested that DOE could play a strong leadership role in developing the datasets to which AVs can be tested. These data sets would comprise sensor signals over a broad range of driving locations and environmental conditions. Currently, industry is not inclined to share its collected data as it represents a competitive advantage in this data-driven field. By supporting or funding public datasets, DOE could enable high-risk innovation from an expanded pool of research and development partners. This recommendation aligns strongly with industry's priority on safety.

2.3.4. Develop the knowledge and tools to advance technology readiness

Participants suggested that DOE could serve to reduce technology development risk under a single early stage research umbrella. The world of AV development, generally speaking, has long been vigorously supported by private investment and largely focused on near-term products. But there are viable technologies that are currently lacking attention from the private sector because they might not hold immediate incentives or appeal for the private sector because of their immature technology readiness level despite their potential impact on energy efficiency. Federal research investments would be beneficial in areas that are not currently being widely developed for AV applications, which include but are not limited to, advanced low-energy sensors, radar systems, and early-stage data fusion.

The low-energy electronics challenge is complex enough that the research community recommends a co-design approach; an approach discussed at a recent DOE Office of Science workshop is illustrated in Figure 2.1. The multiscale co-design framework refers to an approach that ultimately enables both top-down coupling of application and architectural requirements to circuits and devices, as well as bottom-up coupling of materials and device physics constraints to algorithms and architectures. The co-design framework can eventually serve as the basis to deploy powerful optimization and sensitivity analysis codes to evaluate design tradeoffs across many levels of information processing and provide a means to quantify potential gains and measure progress. Portions of a co-design framework exist in today's industrial toolkit, but they are narrowly focused on the current silicon CMOS technology stack.

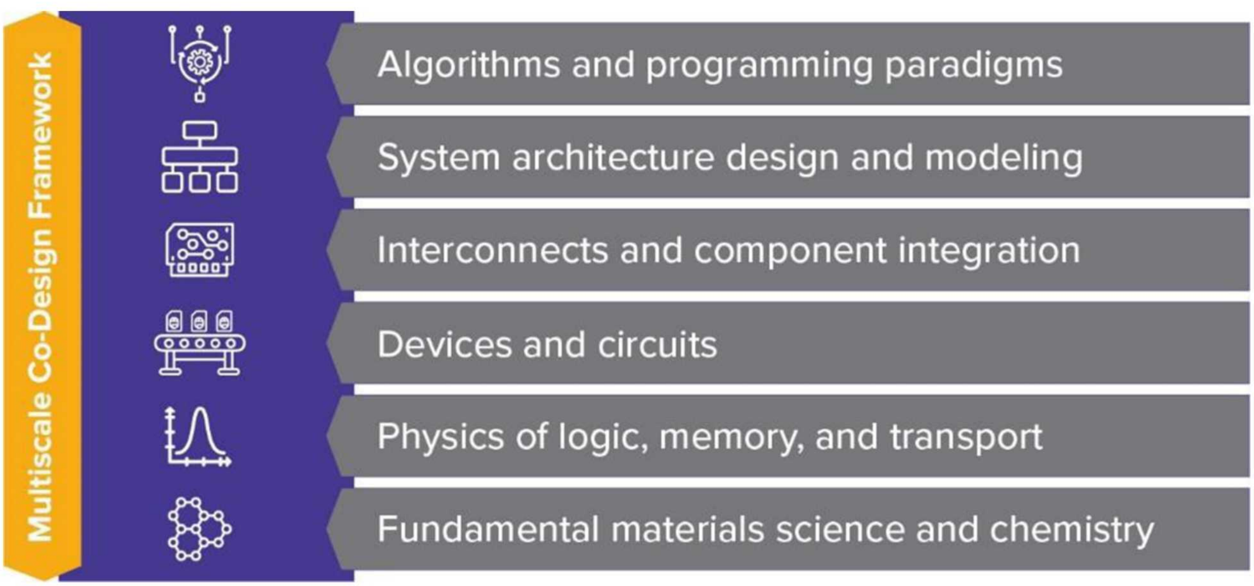

Figure 2.1 Co-design involves multi-disciplinary collaboration that considers the interdependencies among materials discovery, device physics, architectures, and the software stack for developing information processing systems of the future. (Figure reproduced from DOE Office of Science "Basic Research Needs for Microelectronics" workshop summary, Oct. 2018)

3. SUMMARY

Industry is currently focused on solving the algorithmic and software problems of perception and inference, setting aside optimization of computing hardware and power consumption for later. Participants largely agreed that a "step-change" in computing technology will be necessary to enable commercially viable highly automated driving. While participants did not directly discuss computing and power consumption requirements, they agreed that a more energy efficient alternative to current computing approaches may be required to efficiently process the machine learning algorithmic calculations that CAVs will perform. Participants suggested that CAV platforms and hardware should remain flexible to implement evolving algorithms and workloads.

At present, there is no clear convergence in the computing architecture for highly automated vehicles. What is currently unclear among AV developers is whether distributed edge computing or central core computing will be the right path towards the low-energy solution. However, workshop participants generally agreed that there is a need to improve the computing performance per watt by at least 10x to advance the degree of automation.

The computing challenges associated with vehicle connectivity were not discussed much by workshop participants and their overall sentiment was that the infrastructure required to support connectivity would evolve slower than automated vehicle technologies.

Although the issues of cybersecurity and trust in CAV systems do not seem to implicate the workshop's energy efficiency focus, participants raised those issues. Specifically, there was clear agreement that CAV systems must be developed with integrated cybersecurity and trust, but there was no consensus on the approach to ensure these qualities. This topic may be well served by further discussion in a separate forum dedicated to cybersecurity in CAVs.

There is no "ideal" sensor for automated driving. The number and diversity of sensors is a choice made by an OEM based on risk and cost. OEMs integrate multiple types of sensors partly to minimize common-mode failures. Today, AV perception systems rely heavily on LIDAR and camera sensors, but LIDAR solutions have severe limitations in certain lighting and moisture environments. Several participants noted that radar was relatively underutilized, and compared to the investment in LIDAR startups, radar startups currently receive very limited investment. Workshop participants generally agreed, however, that fusing multiple modalities together—radar, LIDAR, and cameras—is the best overall solution for safety and perception.

Workshop participants recommended four ways a federal agency like DOE could help advance computing technology for CAVs. In the near term, the DOE could support development of benchmarks and standards (in partnership with industry) for determining and comparing "computing performance"/watt. The DOE could develop public data sets to support algorithm and software development. Analogous to the VTO battery program, the DOE could co-create a roadmap with industry to help drive and quantify precompetitive advancements in computing performance. To overcome computing challenges for highly automated vehicles, a broader dialogue is necessary about long-term research needs and the role for DOE and its national laboratories. The DOE could reduce risk by advancing the readiness of unproven computing and sensing technologies given their potential impact on energy efficiency.

This page left blank

APPENDIX A. AGENDA OUTLINE

8:00 – 9:00	Registration
9:00 – 9:15	Welcome, DOE Vehicle Technologies Office
9:15 – 10:00	Keynote: DOE Perspective
10:00 – 10:45	Keynote: Industry Perspective
10:45 – 11:00	Break
11:00 – 12:45	Vehicle System Requirements Panel
12:45 – 1:45	Lunch Break
1:45 – 2:30	Keynote: Industry Perspective
2:30 – 4:00	Computing and Algorithms Panel
4:00 – 4:15	Break
4:15 – 5:45	Microelectronics Hardware Panel
5:45 – 6:00	Observations and Close-out Discussion
6:00	Adjourn

This page left blank

DISTRIBUTION

Email—Internal

Name	Org.	Sandia Email Address
Technical Library	01177	libref@sandia.gov

This page left blank

This page left blank

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.