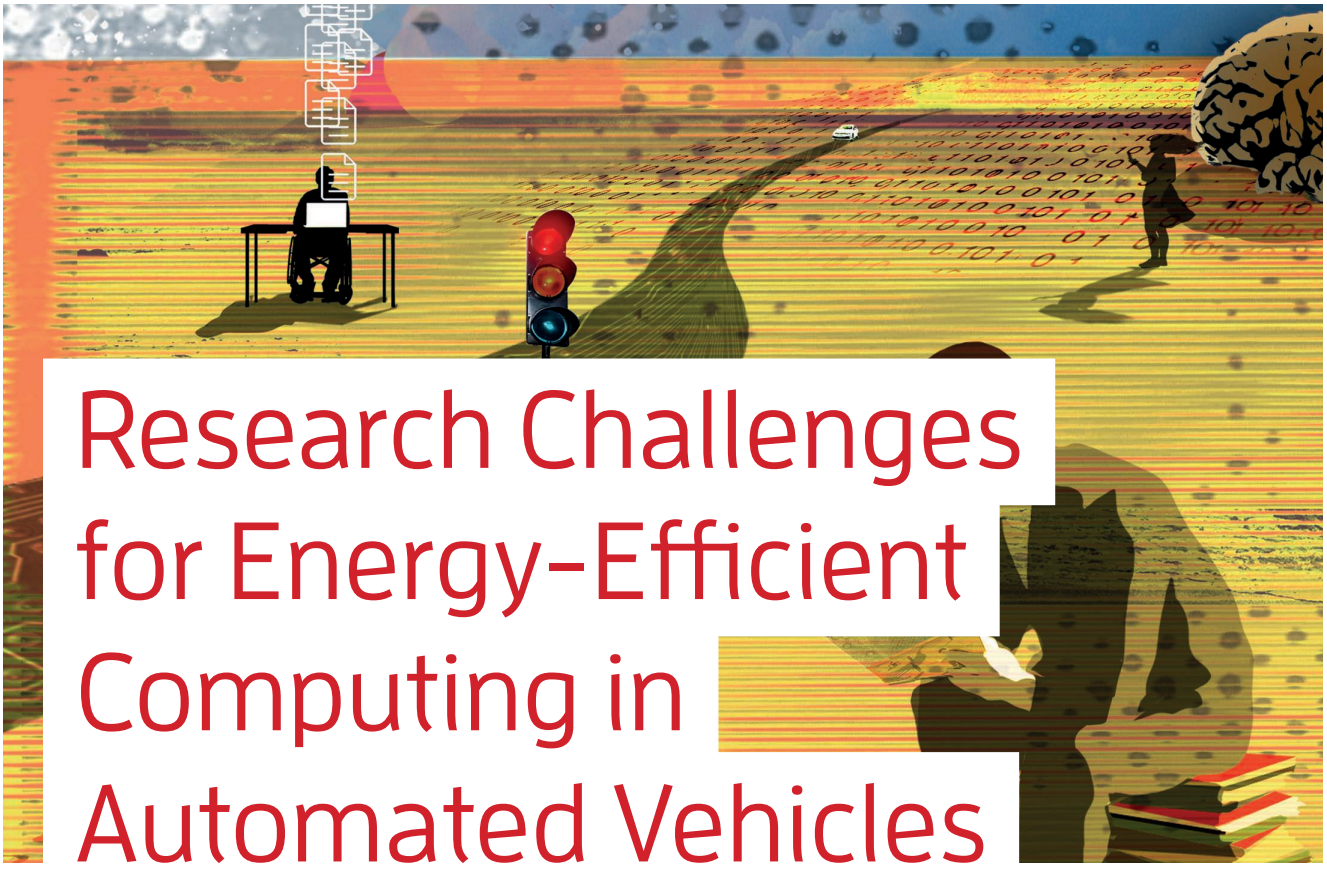

Research Challenges for Energy-Efficient Computing in Automated Vehicles

Robert P. Dick, University of Michigan
Rob Aitken, Arm Research
Jace Mogill, USCAR
John Paul Strachan, Peter Grünberg Institute for Neuromorphic Compute Nodes at Forschungszentrum Jülich
Kirk Bresniker, Hewlett Packard Enterprise
Wei Lu, University of Michigan
Yorie Nakahira, Carnegie Mellon University
Zhiyong Li, Matthew J. Marinella, William Severa, A. Alec Talin, Craig M. Vineyard, Suhas Kumar, Christian Mailhiot, and Lennie Klebanoff, Sandia National Laboratories

Fully automated retail vehicles will have stringent, real-time operational safety, sensing, communication, inference, planning, and control requirements. Meeting them with existing technologies will impose prohibitive energy-provisioning and thermal management requirements. This article summarizes the research challenges facing the designers of computationally energy-efficient, economical, safe, and reliable automated retail vehicles.

Digital Object Identifier 10.1109/MC.2022.3180987
Date of current version: 8 March 2023

Automated vehicles (AVs) hold great promise for improving safety, efficiency, and sustainability. However, commercially viable AVs will depend on reliable, high-performance, compact vehicle-based computing systems. Such computing systems face new challenges distinct from those in supercomputers and data centers due to the low-latency, fail-operational, and long-life span requirements of AVs. The computational requirements for sensing, inference, planning, control, and verification must be met continuously, with low latency. AV computer systems must tolerate wide temperature ranges and severe weather conditions while staying within vehicle energy budgets to maximize the range between charges or refuelings. Computational power consumption must also be managed to enable simpler and more economical thermal management solutions. AVs must also remain in service for a life span of 15 years or more. These characteristics of AVs demand greater reconfigurability, upgradeability, energy efficiency, and tolerance to extreme conditions than conventional computing systems such as servers and desktop computers.

BACKGROUND ON AVs

In contrast to examining the broader energy consumption implications of AVs (for example, the possible increases in energy use due to increased demand for more convenient transportation and decreases due to improved route planning and traffic management), we describe the on-vehicle energy demands of the sensing, signal processing, and inference systems needed to make AVs practical. Existing approaches to capturing large amounts of data and using them to make complex, time-constrained, and safety-critical decisions would impose difficult-to-tolerate energy demands on AVs.

TO MEET BOTH ENERGY EFFICIENCY AND COMPUTATIONAL PERFORMANCE GOALS, TARGETED R&D IS NEEDED TO SATISFY THE SIZE, WEIGHT, POWER, AND THERMAL CONSTRAINTS OF AVs.

To meet both energy efficiency and computational performance goals, targeted R&D is needed to satisfy the size, weight, power, and thermal constraints of AVs. This article describes the findings of a team of researchers from a national lab, academia, and industry (microelectronics and automotive) on the challenges and promising directions for the advancement of energy-efficient computing in AVs (EECAV), with a focus on retail vehicles. The team focused on in-vehicle solutions for retail vehicles because this captures the most widespread ownership, service, and support scenarios. Fleet vehicles are typically geo-fenced or limited to certain routes that may be equipped with infrastructure to assist driving, whereas retail vehicles often visit locations where there is no infrastructure to communicate with or rely on.

Conventional CMOS digital computing technology is fast approaching the end of Moore's Law scaling. Projected process scaling trends alone are unlikely to meet the demands of highly automated, energy-efficient AVs in the next two decades. A paradigm shift toward integrated research (that is, "co-design") on devices, architectures, algorithms, and sensors is needed for rapid advances in the energy efficiency of AV computing.

The energy problem is readily apparent. Consider an average, light-duty, gasoline-powered vehicle with 10 gallons of gasoline. With a gasoline lower heating value of 126.1 megajoules (MJ)/gallon,[1] the fuel load contains 1,261 MJ of energy. However, all energy on the vehicle is derived from the internal combustion engine (ICE). Assuming an ICE thermal efficiency of 30%,[2] the total energy available to the vehicle for propulsion and auxiliary systems is 378.3 MJ. Let's assume that vehicle has a range of 400 mi. Traveling at 65 mi/h gives a trip duration of 6.15 h. Thus, the average power available to the vehicle during that time is 378.3 MJ/6.15 h = 61.51 MJ/h, or 17.1 kW.

As shown in Figure 1, AVs tested by leading developers around 2020 generally expended about 3 kW for computation (that is, converting input data to higher-level knowledge, enabling control of the vehicle), with similar numbers noted for more recent systems.[3–5] However, in our ICE vehicle example, this 3 kW of power must be supplied by an alternator with typically limited (~60%) efficiency.[6] Thus,

the real power demand relevant to total vehicle fuel energy is 5 kW, or ~29% of the total fuel energy available on the vehicle. Were the energy supplied by batteries, efficiency would be higher, but this would be more than offset by the reduced energy storage density of batteries compared to gasoline-fueled ICEs.

This computational energy expenditure represents a burdensome fraction of the available onboard energy, particularly because many of the computational electronics required for automated driving must be constantly on and impose an extra burden in addition to the other energy-intensive systems, such as propulsion, heating, air conditioning (AC), and lighting. As indicated in Figure 1, we anticipate that eventually a 300-W "all-in" target for onboard AV computation and associated support equipment will be required to meet vehicle energy constraints as well as the onboard thermal design constraints without major changes to thermal management strategies. A power consumption of 300 W would bring the AV computation electrical load in line with other auxiliary systems, such as the heater (~360 W),[7] AC (~1 kW),[7,8] and lights (~130 W).[8]

How can this 10× improvement in computation power be achieved? As shown in Figure 1, IEEE International Roadmap for Devices and Systems (IRDS) projections[9] suggest that future process scaling is likely to bring a 2.5× improvement by 2040. A further 1.6× improvement in device technology produces a 4× total improvement. An 1.6× improvement in computer architectures would produce a combined 6.3× improvement in required AV computation power. However, improvements in scaling, devices, architecture, and algorithms (with a 1.6× improvement) would produce a 10× total improvement. A full-stack effort is likely the surest route to the goal. We note here that revolutionary approaches beyond CMOS would offer an alternative, high-potential track to the paths indicated in Figure 1, but it is uncertain whether such revolutionary changes will be practical for large-scale commercial use in the near future.

Exploring the implications for the 300-W target in Figure 1 a bit further, Figure 2 shows that the relative contributions of compute and propulsion energy consumption to driving depend strongly on the type of driving. The table in Figure 2(a) gives specific results for examples of 10 mi/h of urban traffic driving and 65 mi/h freeway driving, both assuming 3 kW of compute power. Figure 2(b) displays a plot showing how the percentage of total energy devoted to computation varies with changes in average vehicle speed and different assumptions (300–6,000 W) for compute power. These two examples consider the same vehicle on a hypothetical 10-mi trip at freeway speeds (65 mi/h) and urban speeds (10 mi/h) for the same distance

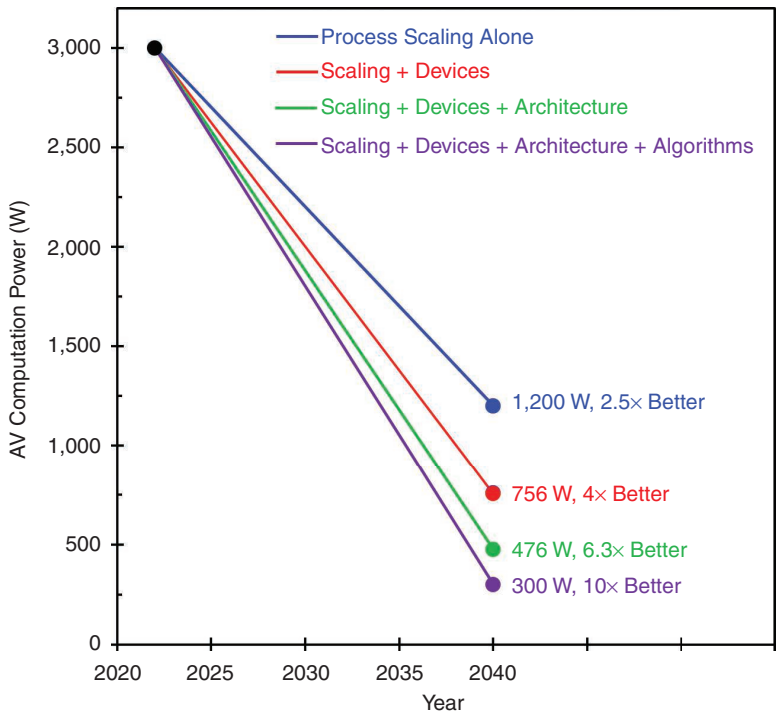

FIGURE 1. The existing (3-kW) computational power consumption for prototype AVs and possible routes to reducing the electrical power devoted to AV computation down to 300 W via the following: blue: improvements through process scaling alone (2.5×); red: improvements in both process scaling (2.5×) and devices (1.6×) for a total of 4× improvement; green: improvements in scaling, devices, and architecture (1.6×) for a total of 6.3× improvement; purple: combined improvements in scaling, devices, architectures, and algorithms (1.6×) for a total of 10× improvement.

traveled (10 mi). In this simplified example, which considers only propulsion and computation at the current 3-kW power expenditures, it is evident that the relative amounts of propulsion and compute energies depend on "what the car is doing."

In general, freeway driving allows the computer to be on less of the time and have a smaller relative energy contribution. However, the onboard computer system must be sized for the worst-case (most computationally intensive) scenario of urban driving: hence the 300-W target in Figure 1. Depending on the specifics of a trip or vehicle mission, the computation electrical power target may have different estimates.[10] Fortunately, although many aspects of automobiles have already been optimized for energy efficiency, in-vehicle computation has not been well optimized, and there is room for improvement by pursuing the R&D challenges indicated in this article. This article represents a summary of an EECAV R&D "Roadmap Outline," which is available for download.[11] The Roadmap Outline was vetted with members of the AV technical community via an online workshop held 11–12 May 2021.[11]

In our work, we have assumed that all (nonlearning) sensing and computation occurs on-vehicle. In reality, it may be possible to offload some tasks to servers in data centers, distributed in the transportation infrastructure, and in other vehicles. However, in the foreseeable future, wireless communication links will experience intermittent outages, there will be gaps in access to intelligent infrastructure, and it will be necessary to operate in environments with legacy non-AVs. Therefore, it is reasonable to assume that AVs must support all latency- and safety-critical tasks, for example, the sensing and inference necessary to make time-critical driving decisions. This leaves open the possibility of offloading tasks that are either latency insensitive (for example, long-term learning, route preplanning, or maintenance planning) or not safety critical (for example, decisions related to entertainment).

Urban Traffic Driving (Low Speed):		
Speed	10	Miles per hour
Distance Traveled	10	Miles
Time for Trip	1	Hours
Propulsion	300	Watt-hours/mile
Compute Power	3,000	Watts
Propulsion Energy	3,000	Watt-hours
Compute Energy	3,000	Watt-hours
% of Total Energy Used by Compute	50	

Freeway Driving (High Speed):		
Speed	65	Miles per hour
Distance Traveled	10	Miles
Time of Trip	0.154	Hours
Propulsion	300	Watt-hours/mile
Compute Power	3,000	Watts
Propulsion Energy	3,000	Watt-hours
Compute Energy	462	Watt-hours
% of Total Energy Used by Compute	13.3	

(a)

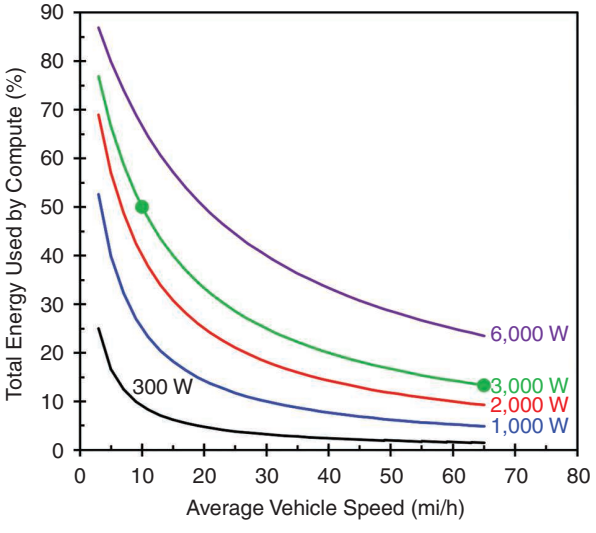

(b)

FIGURE 2. (a) and (b) The variation of the percentage of total energy (propulsion + compute) used by compute for varying average vehicle speeds and varying computational power. (a) The table shows the results for one choice of urban traffic driving at 10 mi/h and one choice of freeway driving at 65 mi/h, both for an assumed compute power of 3 kW. (b) The green dots on the 3-kW curve show the results for 10 and 65 mi/h described by the table.

It will also be possible to reduce average on-vehicle computational demands by (partially) offloading safety-critical and latency-sensitive tasks when appropriate communication and computation resources happen to be available. However, this will be challenging because energy-intensive tasks such as inference based on camera and lidar data will often require large amounts of communication to fully off-load. This may lead to solutions in which inference is divided among on-vehicle and off-vehicle systems to improve vehicular energy efficiency without overwhelming communication resources.

SCOPE AND TIMELINES

We considered several "boundary conditions" to guide our analysis of technical problems and relevant research areas. We assumed that all of the computational capacity resides on the vehicle as retail AVs would have to be "self-sufficient" for entry into the automotive market, particularly since any intelligent infrastructure, if developed, would initially be sparse and therefore have intermittent availability. Second, we adopted an "all-in" computational power target of 300 W for the electrical power consumed by the onboard computers, sensors, and any supporting peripherals to enable the automated driving functions. Regarding the AV development timeline, we assumed three concurrent timelines: computation-related R&D topics spanning 2025–2035, enabling commercial-ready chips in the 2030–2040 time frame, and implementation in commercial vehicles from 2035 to 2045.

We identified R&D problems distributed across the following four general technical areas.

Technical Area I: Chips: Materials, Devices, and Circuits

Technical Area I (TA-I) focuses on the materials, devices, and circuits that are the building blocks of computer chips operating in the demanding automotive environment and within the power budget allotted for computation. Those working in this area strive to provide more capable and energy-efficient technology at the chip level, in the context of AV application requirements. We identified the following main challenges in TA-I.

Discovering new or improved materials and processing techniques for increased thermal/mechanical/radiation robustness for automotive environments for the life of a vehicle (~15 years). With computation needing to occur on the AV itself, there arise significant challenges in the robustness needed for the chip technology. This includes chip survivability in thermally cycling conditions (constantly changing hot and cold environments); shock and vibration resistance; and stability under cosmic ray bombardment during high-altitude driving. These are all challenging environmental concerns needing to be solved in the near term if AVs are to enjoy widespread deployment.

Developing low-latency and low-power (<300 W) onboard computing circuits, such as "in-memory" computing hardware, where memory and logic/computing are integrated. For safety-critical and low-computational-latency operation, fast in-vehicle sensing and decision making are essential. Some published estimates suggest that slower than 150 ms from sensing to actuation is unacceptable for AVs from a safety standpoint.[12] Potential research areas within this R&D challenge area include enabling non-von Neumann architectures (removing the separation between data storage and compute), which can lead to significant energy and latency savings by minimizing the on-chip data movement that dominates these metrics. It is also important to develop memory technologies (for example, flash, ferroelectric-based, or resistance based) that are automotive compatible (for example, the ability to withstand harsh weather and wide temperature variations).

Creating computing devices and circuits that offer reconfigurability (for example, in response to a new algorithm or learning from road conditions). Hardware reconfigurability is important because the hardware needs to be compatible with advanced

> THE TEAM FOCUSED ON IN-VEHICLE SOLUTIONS FOR RETAIL VEHICLES BECAUSE THIS CAPTURES THE MOST WIDESPREAD OWNERSHIP, SERVICE, AND SUPPORT SCENARIOS.

and developing algorithms (for example, computer vision) and efficient computing models that can enable software-level generality while maintaining the best possible performance of the underlying hardware.[13]

Integrating novel materials, devices, and circuits into existing manufacturing technologies and tools. Improving energy efficiency in computing is of interest across many industries, and longer-term R&D is needed as efficiency improvements through scaling down transistors become decreasingly effective.[14] Thus, there is interest in new materials, devices, and circuits for "beyond Moore" and "beyond CMOS" technologies.

Figure 3 lists the TA-I R&D problems in a simplified two-axis format to convey the timing (near term or long term) associated with solving these R&D problems and likely impact (high impact or very high impact) of the solutions on realizing widespread highly automated AVs. For these two-axis plots, "near-term" investment means in the time frame 2025–2030; "longer-term" investment means in the time frame 2030–2035. The impact scale is somewhat subjective. "High impact" conveys very important R&D that is needed for AV development. "Very high impact" conveys R&D that can substantially change the direction of AV technical development.

Technical Area II: Chips: Architecture, Safety, and Security

Technical Area II (TA-II) relates to the aspects of computer architecture that are most relevant to AV operation, including data throughput, memory, high computational accuracy (demanded by functional safety), and the security of the computational activity during vehicle operation. The R&D challenges identified in TA-II follow.

Exploring the optimized use of distributed, heterogeneous multiprocessor systems, including CPUs, GPUs, neural processing units (NPUs), and other application-specific hardware, to support the algorithms needed for EECAV. Future AVs will require distributed computing systems consisting of smart sensors that process incoming data locally as well as more centralized vehicle computers that perform data fusion, perception, and navigation. These systems will need to support advanced autonomy algorithms for AVs and will connect a variety of different processors, including CPUs, GPUs, NPUs, and other special-purpose accelerators. Research is needed to better define, develop, and optimize these heterogeneous system architectures while meeting all of the AV system constraints. Such research will need to build on existing designs[15] and extend them to improve capability, power, and performance.

Determining the type of onboard network/interconnect strategies that optimize computational energy efficiency. New architectures will be needed to connect sensors and heterogenous processors. As such, it will be necessary to define the networks that will connect these chips. Research is needed to understand the fail-functional bandwidth and latency requirements, including chip to chip, car to car, and eventually car to infrastructure, as well as protocols that might meet those requirements. While car-to-infrastructure approaches offer great promise,[16] the early infrastructure will be sparse.

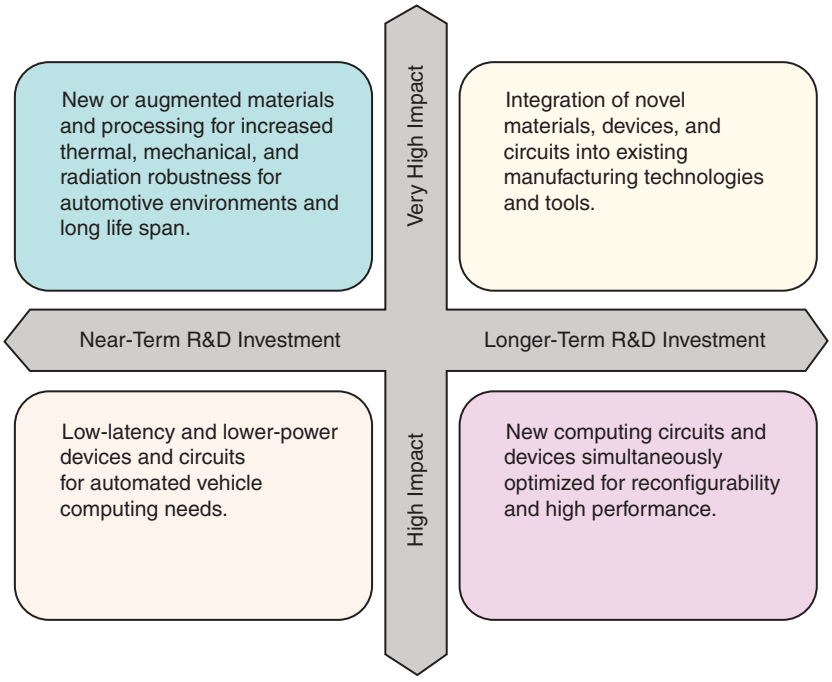

FIGURE 3. R&D challenges for TA-I: Chips: Materials, Devices, and Circuits.

Developing improved memory (addressable and storage) and bandwidth in support of AV and determining where these are located within the system. With the new heterogeneous computing multiprocessor architectures discussed previously, the needs for memory, bandwidth, and speed will change. More memory and storage will certainly be needed to support event and data recorders as well as larger deep neural networks, but their optimal distribution and network architecture require further study. Furthermore, problems such as the requirement of coherence across the system should be understood as these decisions directly affect computing and network architectures and require memory, bandwidth, and speed requirements to be defined holistically.

Identifying when (and if) computational "demand" starts to require consideration of "off-vehicle" computation. We decided that having enough computational capacity residing on the vehicle to safely drive without reliance on infrastructure established a "boundary condition" for this work. However, it will be important to understand when significant "off-vehicle" computation is necessary and how it might integrate with existing or proposed "smart" traffic infrastructure.[17] Even when infrastructure is present, the additional latency of transmitting and receiving data via a network for edge/cloud computation may not favor the offloading of real-time computations, which account for the majority of the work.

Figure 4 lists these TA-II R&D problems to convey the required timing and likely impact of solving these technical problems.

Technical Area III: Algorithms and Data Management

Technical Area III (TA-III) identifies algorithmic and software challenges that, if solved, would lead to reduced computational demand and greater functionality via more optimal analysis of data from cameras and lidar systems as well as other sensor data sources. The R&D challenges for TA-III follow.

Efficiency optimization of algorithms to improve latency and energy consumption. We believe that substantial energy efficiency improvements remain to be realized via algorithmic improvements. Possible approaches include sparse sampling/processing; optimizing the spatial and temporal resolution in sensing; improvements in network design going beyond adjustments in layer and neuron counts; and better understanding the relationships between specific training samples and their influence on learned parameters.

Co-optimizing algorithms and implementation platforms to improve efficiency, performance, and fault tolerance. AVs will require advances in both algorithms and the computational substrates supporting them. The majority of the computational load may be handled by deep learning systems, including convolutional neural networks, recurrent neural networks, multilayer perceptrons, and their derivatives, and these techniques must be tolerant to faults. These algorithms are well suited for video, radar, and lidar processing necessary for vehicle perception and can be designed for robustness in the presence of noise and some classes of hardware faults. Advancements

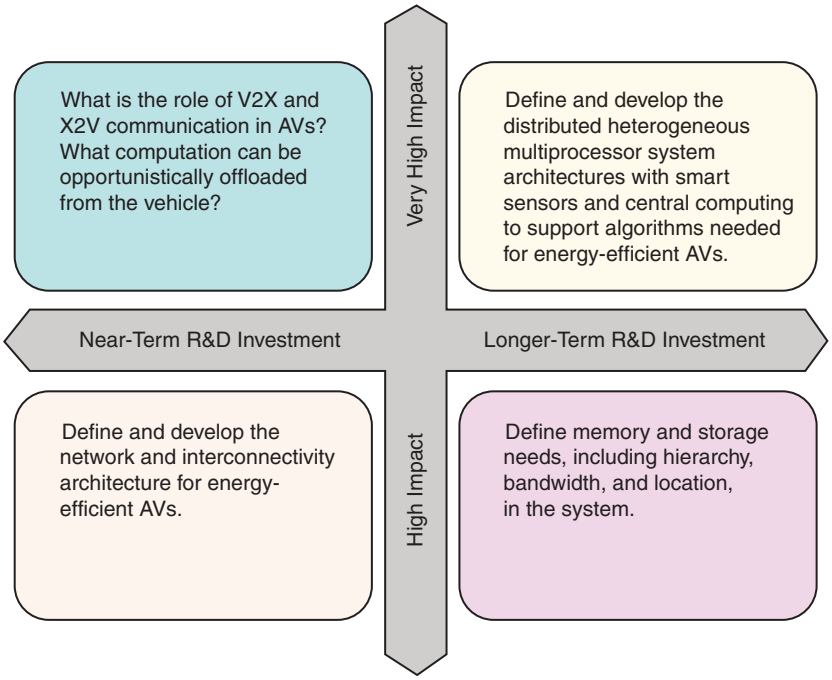

FIGURE 4. R&D challenges for TA-II: Chips: Architecture, Safety, and Security. V2X: vehicle to everything; X2V: everything to vehicle.

in neural networks, such as capsule networks and attention-based approaches (for example, transformers and squeeze-and-excitation networks), will likely foster the necessary developments in capability, and recent advances in multilayer perceptrons render them competitive, especially when efficiency is of primary importance.

Data and training. The character of training data is critical for artificial intelligence (AI) systems. However, AV training faces special requirements for safety and trust. It is impractical (likely impossible) to capture all real-world traffic dynamics in a static data set or simulator. Furthermore, environmental effects, wear, and age may affect deployed sensors and processors. Both understanding the role and implications of training approaches and understanding training data selection and augmentation are critical for successful AVs.

We will need methods for quantifying uncertainty in vehicle perception and estimating the completeness of training data. Bayesian methods, including stochastic or Bayesian neural networks, may help estimate epistemic uncertainty. Additionally, advances in computational learning theory may help inform when algorithms need to be retrained or updated, for example, due to a training data set distribution that no longer matches the current data distribution. Data will often not be independently and identically distributed. Online learning methods may help alleviate an initial data requirement by being able to learn "on the fly." However, online learning methods introduce greater challenges for safety, assuredness, and validation. Computational learning theory may also be able to address this concern. It is currently not well understood how to measure lost performance due to neglecting online/continuous learning.

> WE DESCRIBE THE COMPUTATION ENERGY EFFICIENCY R&D BARRIERS SPANNING CHIPS, ARCHITECTURES, ALGORITHMS, AND SENSORS.

AVs may be able to capture data useful for updating AI models, particularly if the AV encounters unlikely "edge cases." However, in such scenarios, data privacy must be considered: both technical and legal challenges exist. Given the challenges of obtaining representative data sets, simulation will play a role in training AV AI systems. Additionally, methods such as surrogate tasks and contrastive learning may help improve generalization, both from data set to deployment and from simulation to the real world.

Managing data retention and locations for optimized performance, memory, communication, privacy, and legal constraints. Near-sensor signal processing to transform and compress data may be needed to reduce latency, energy consumption, and communication. Such system-level architectures will pose new constraints on algorithm design. It may be possible to eliminate standard signal processing stages and hardware components, feeding raw data into end-to-end learning frameworks. Such frameworks should account for possible mismatches between the testing and training data and be able to adapt to changing environments. While some methods exist to address this problem, it is not fully solved, particularly for the assured or trusted operation needed in AVs. Sensor fusion will help to identify unimportant data. Determining which data should be retained, in what form, and where will be a research challenge as there will be conflicting objectives, including efficiency, latency, accuracy, privacy, storage cost, and security.

The R&D challenges for TA-III are shown in Figure 5 with their impact and timing indicated.

Technical Area IV: Sensors Data Interface

Technical Area IV (TA-IV) is concerned with the links from the external physical world to AV computers and among physically distributed computers in a vehicle. The computational requirements are intimately connected to the nature of the sensors and their interfaces. The R&D challenges for TA-IV follow.

Evaluating tradeoffs among smart sensors and central computing. The decreasing cost of computational capabilities means sensors that historically included only a detector now often include computing capabilities. The tradeoffs and long-term implications of in-sensor processing are multidimensional and not constrained to the sensor. Making sensors "smarter" implies moving computation close to the sensor,

with implications for the required computational and energy efficiency.[18] In addition, current approaches to functional safety require redundancy. The distributed nature of smart sensors may help support functional safety requirements more efficiently than fully redundant modules. The extent to which this is possible needs research.

Data versus task migration for dynamic power management. Distributed and heterogeneous systems offer choices about where computation is performed and data are stored; data can be moved from distant memory to where a process is executing, or the process can move to execute near to where the data are stored.[19] Tradeoffs include the processor types, energy, and time costs for moving data and tasks; bandwidth consumed; types and amount of memory; and so forth. Dynamic task migration capabilities can also be used to implement functional safety, such as moving work and/or data from a failed unit to another device. Potential synergies may blur the line between load balancing and fail-operational safety without adding complexity.

Exploiting asymmetric bandwidth utilization of networks to improve energy efficiency. Network technologies typically provide symmetric transmit and receive bandwidth, a reasonable design when the use case for the network is not known in advance or the topology changes over time. However, in-vehicle computing systems have the benefit of changing slowly, if ever, and have well-defined use cases that change little over time. Specifically, high-resolution cameras primarily send data and only need to receive control commands—bandwidth utilization may differ by six orders of magnitude or more between transmit and receive, creating the opportunity for significant savings by making the network asymmetric. Likewise, a wheel rotation sensor produces data but does not consume it, meaning it is overprovisioned with respect to received bandwidth.

Determining the R&D needed to anticipate advances in sensors and computers over a long (~15-year) vehicle life span to maintain forward and backward capability. Despite best efforts to maintain compatibility of software and hardware over time, the rapid rate of innovation means compatibility is often broken before the device itself breaks.[20] Over the 15+-year life span of a vehicle, technology will evolve so that replacement components may be much more capable and different in nature than the devices they replace, moving the data to a processor or vice versa. Research is needed to develop best practices for future-proofing sensors, computers, and their interconnects. Security concerns that are difficult to predict may also limit backward compatibility or prevent upgrades.

The R&D challenges for TA-IV are shown in Figure 6 with their impact and timing indicated.

This article summarizes the computational research challenges facing the designers of computationally energy-efficient, economical, safe, and reliable automated retail vehicles. We describe the computation energy efficiency R&D barriers spanning chips, architectures, algorithms, and sensors, with the goal of paving the way to EECAV. In addition to the R&D priorities identified in Figures 3–6, the

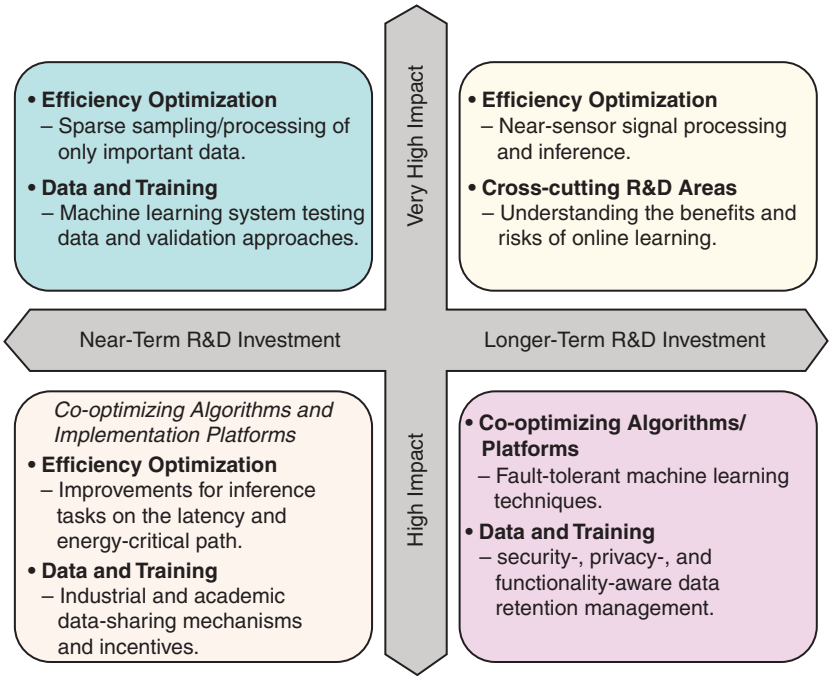

FIGURE 5. R&D challenges for TA-III: Algorithms and Data Management.

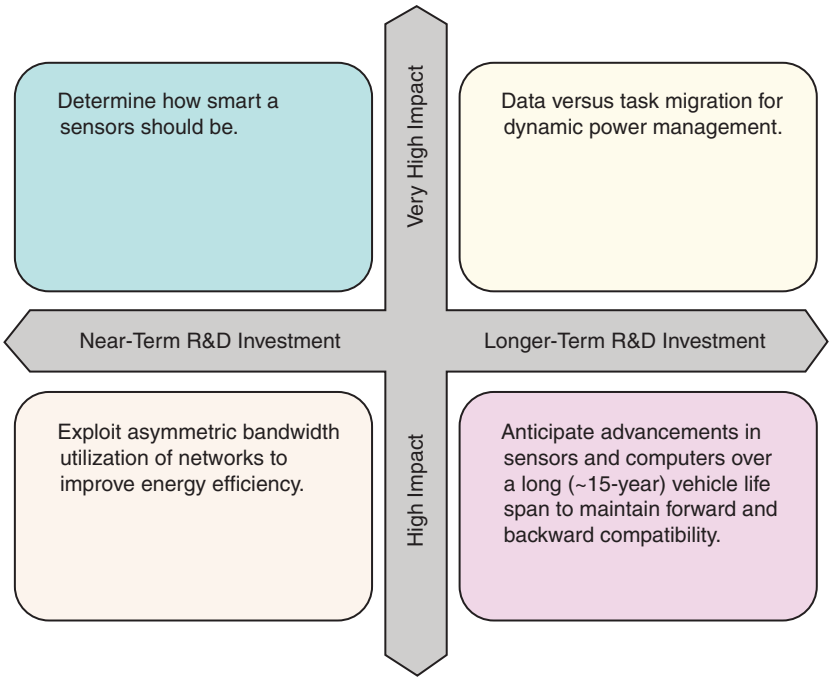

FIGURE 6. R&D Challenges for TA-IV: Sensors Data Interface.

following R&D activities would be particularly important to pursue: 1) developing automotive-centric metrics for computing performance improvement, 2) putting on a solid technical foundation "what the AV is doing," 3) understanding quantitatively how improvements in all four technical areas affect one another, and 4) converging on a commonly acceptable approach to benchmark AV computational capacity and energy efficiency.

There is an urgent need to develop further a full roadmap of advanced computing for AVs if fully automated commercial AVs are to be realized.

ACKNOWLEDGMENTS

We wish to thank Chris Moen, John Aidun, Carrie Burchard, and Chris Shaddix of Sandia for very helpful discussions on the execution of this work. Initial guidance from Ian Young of Intel is appreciated. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and 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-NA-0003525. This article describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the article do not necessarily represent the views of the U.S. Department of Energy or the U.S. Government. Robert P. Dick and Lennie Klebanoff are the corresponding authors.

REFERENCES

1. L. M. Das, "Hydrogen engines: A view of the past and a look into the future," *Int. J. Hydrogen Energy*, vol. 15, no. 6, pp. 425–443, 1990, doi: 10.1016/0360-3199(90)90200-I.
2. L. E. Klebanoff, J. O. Keller, M. H. Fronk, and P. Scott, "Hydrogen conversion technologies and automotive applications," in *Hydrogen Storage Technology, Materials and Applications*, L. E. Klebanoff, Ed. Boca Raton, FL, USA: Taylor & Francis, 2012, ch. 2, p. 31.
3. J. Stewart. "Self-driving cars use crazy amounts of power, and it's becoming a problem." Wired. https://www.wired.com/story/self-driving-cars-power-consumption-nvidia-chip/ (Accessed: Aug. 15, 2022).
4. K. Hamza, J. Willard, K. Chu, and K. P. Laberteaux, "Modeling the effect of power consumption in automated driving systems on vehicle energy efficiency for real-world driving in California," *Transp. Res. Rec.*, vol. 2673, no. 4, pp. 339–347, 2019, doi: 10.1177/0361198119835508.
5. National Academies of Sciences, Engineering, and Medicine, *Assessment of Technologies for Improving Light-Duty Vehicle Fuel Economy—2025-2035*. Washington, DC, USA: National Academy Press, 2021.
6. M. Adamiec, M. Dziubinski, and E. Siemionek, "Research of the alternator on the stand – Efficiency aspect," in *IOP Conf. Series, Mater. Sci. Eng.*, 2018, vol. 421, p. 022001, doi: 10.1088/1757-899X/421/2/022001.
7. L. Küng, T. Bütler, G. Georges, and K. Boulouchos, "How much energy does a car need on the road?" *Appl. Energy*, vol. 256, no. 14, p. 113,948, 2019, doi: 10.1016/j.apenergy.2019.113948.
8. I. Evtimov, R. Ivanov, and M. Sapundjiev, "Energy consumption of auxiliary systems of electric cars," in *MATEC Web Conf.*, 2017, vol. 133, p. 06002, doi: 10.1051/matecconf/201713306002.
9. "More Moore," IEEE International Roadmap for Devices and Systems,

ABOUT THE AUTHORS

ROBERT P. DICK is an associate professor of electrical engineering and computer science at the University of Michigan, Ann Arbor, MI 48109 USA. His research focuses on embedded systems. Dick received a Ph.D. from Princeton University. He is a Senior Member of IEEE. Contact him at dickrp@umich.edu.

ROB AITKEN is an Arm fellow and the director of technology for Arm Research, San Jose, CA 951534 USA. His research interests include autonomous vehicles, emerging technologies, and resilient computing. Aitken received a Ph.D. from McGill University, Canada. He is a Fellow of IEEE. Contact him at Rob.Aitken@arm.com.

JACE MOGILL formerly led the USCAR/USDRIVE Energy Efficient Advanced Computing work group on behalf of General Motors, Detroit, MI 48243 USA. His research interests include hardware-software co-design, algorithm accelerators, and parallel algorithms. Mogill received an M.S. from the University of Pittsburgh. Contact him at jace@mogill.com.

JOHN PAUL STRACHAN was formerly at Hewlett Packard Labs and now directs the Peter Grünberg Institute on Neuromorphic Compute Nodes (PGI-14) at Forschungszentrum Jülich, GmbH, 52428 Jülich, Germany, and RWTH Aachen, Aachen, 52068, Germany. His research interests include novel types of hardware accelerators using emerging device technologies. Strachan received a Ph.D. from Stanford University. Contact him at j.strachan@fz-juelich.de.

KIRK BRESNIKER is the chief architect of Hewlett Packard Labs, Milpitas, CA 95035 USA. His research interests include integration and industrialization of novel computation, technology ethics, and responsible design. Bresniker received a BSEE (cum laude) from Santa Clara University. Contact him at kirk.bresniker@hpe.com.

WEI LU is a professor with the Electrical Engineering and Computer Science Department, University of Michigan, Ann Arbor, MI 48109 USA. His research interests include resistive-random access memory devices, logic circuits and neuromorphic computing systems. Lu received a Ph.D. from Rice University. He is a Fellow of IEEE. Contact him at wluee@umich.edu.

YORIE NAKAHIRA is an assistant professor in the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213 USA. Her research interests include the fundamental theory of optimization, control, and learning and its applications to neuroscience, cyberphysical systems, and autonomous driving. Nakahira received a Ph.D. in control and dynamical systems from the California Institute of Technology. Contact her at yorie@cmu.edu.

ZHIYONG LI was the manager of Material Physics Department at Sandia National Laboratories before he recently moved to Applied Materials, Santa Clara, CA 95051 USA. His research interests include augmented reality/visual reality, the Internet of Things, and emerging energy solutions. Li received a Ph.D. from the University of Notre Dame. Contact him at zhiyong_li@amat.com.

MATTHEW J. MARINELLA is an associate professor in the School of Electrical, Computer, and Energy Engineering, Arizona State University, Tempe, AZ 85281 USA. His research interests include nonvolatile memory technology as well as neuromorphic and low-power computing with emerging electronic devices. Marinella received a Ph.D. in electrical engineering from Arizona State University. He is a Member of IEEE. Contact him at m@asu.edu.

WILLIAM SEVERA is a mathematician and researcher at Sandia National Laboratories Albuquerque, NM 87123 USA. His research interests include resource-constrained artificial intelligence methods and neuromorphic algorithm development. Severa received a doctorate from the University of Florida. Contact him at wmsever@sandia.gov.

A. ALEC TALIN is a Distinguished Member of Technical Staff at Sandia National Laboratories Livermore, CA 94550 USA. His research interests include nanoelectronics and nanoionics, with applications to energy-efficient and

neuromorphic computing, energy storage and conversion, and national security. Talin received a Ph.D. in materials science and engineering from the University of California, Los Angeles. He is a Fellow of the American Physical Society. Contact him at aatalin@sandia.gov.

CRAIG M. VINEYARD pursues research and development for national security at Sandia National Laboratories, Albuquerque, NM 87123 USA. His research interests include foundational research on how advanced computing technologies can impact a range of applications from scientific computing to remote sensing. Vineyard received a PhD. in computer engineering from the University of New Mexico with a concentration in computational intelligence. Contact him at cmviney@sandia.gov.

SUHAS KUMAR is a principal scientist at Sandia National Laboratories, Livermore, CA 94551 USA, where he leads the Nonlinear Electronic Devices group. His research interests include electronic materials, post-CMOS devices, and neuromorphic algorithms. Kumar received a Ph.D. in electrical engineering from Stanford University. Contact him at suhkuma@sandia.gov.

CHRISTIAN MAILHIOT is the senior manager of the Materials Research Group at Sandia National Laboratories, Livermore, CA 94551 USA. His research interests include condensed matter and materials physics, with particular emphasis on innovative discoveries related to quantum-confined semiconductor structures and high-pressure research. Mailhiot received a Ph.D. in applied physics from the California Institute of Technology. Contact him at cmailhi@sandia.gov.

LENNIE KLEBANOFF is a principal member of the technical staff at Sandia National Laboratories, Livermore, CA 94551 USA. His research interests include extreme ultraviolet (EUV) lithography, EUV mask inspection, and hydrogen-based technology. Klebanoff received a Ph.D. in physical chemistry from the University of California-Berkeley. Contact him at lekleba@sandia.gov.

2021. [Online]. Available: https://irds.ieee.org/editions/2021/more-moore

10. E. Rask et al., "SMART mobility: Connected and automated vehicles capstone report," U.S. Department of Energy, Washington, DC, USA, 2020. [Online]. Available: https://www.energy.gov/sites/prod/files/2020/08/f77/SMART-CAVS_Capstone_07.22.20.pdf
11. "Energy efficient computing for automated vehicles," Sandia National Laboratories, Livermore, CA, USA, 2021. [Online]. Available: eecav.sandia.gov
12. M. Rafie, "Edge AI computing advancements driving autonomous vehicle potential," Semiconductor Digest, Andover, MA, USA, Oct. 15–19, 2021. [Online]. Available: https://www.semiconductor-digest.com/edge-ai-computing-advancements-driving-autonomous-vehicle-potential/
13. S. Yin et al., "A high energy efficient reconfigurable hybrid neural network processor for deep learning applications," *IEEE J. Solid-State Circuits*, vol. 53, no. 4, pp. 968–982, 2018, doi: 10.1109/JSSC.2017.2778281.
14. T. N. Theis and H.-S. P. Wong, "The end of Moore's law: A new beginning for information technology," *Comput. Sci. Eng.*, vol. 19, no. 2, pp. 41–50, 2017, doi: 10.1109/MCSE.2017.29.
15. E. Talpes et al., "Compute solution for tesla's full self-driving computer," *IEEE Micro*, vol. 40, no. 2, pp. 25–35, 2020, doi: 10.1109/MM.2020.2975764.
16. L. Hobert, A. Festag, I. Llatser, L. Altomare, F. Visintainer, and A. Kovacs, "Enhancements of V2X communication in support of cooperative autonomous driving," *IEEE Commun. Mag.*, vol. 53, no. 12, pp. 64–70, 2015, doi: 10.1109/MCOM.2015.7355568.
17. Z. Ning, J. Huang, and X. Wang, "Vehicular fog computing: Enabling real-time traffic management for smart cities," *IEEE Wireless Commun.*, vol. 26, no. 1, pp. 87–93, 2019, doi: 10.1109/MWC.2019.1700441.
18. F. Zhou and Y. Chai, "Near-sensor and in-sensor computing," *Nature Electron.*, vol. 3, no. 11, pp. 664–671, 2020, doi: 10.1038/s41928-020-00501-9.
19. S.B. Baden et al., "UPC++ and GASNet-EX: PGAS support for exascale applications and runtimes (Extended Poster Abstract)," Lawrence Berkeley National Laboratory, Berkeley, CA, USA, 2018. [Online]. Available: Hhttps://gasnet.lbl.gov/pubs/sc18-pagoda-poster-abstract.pdf
20. I.-H. Hann, B. Koh, and M. F. Niculescu, "The double-edged sword of backward compatibility: The adoption of multi-generational platforms in the presence of intergenerational services," *Inf. Syst. Res.*, vol. 27, no. 1, pp. 112–130, 2016, doi: 10.1287/isre.2015.0615.