

SEPTEMBER 2019

AUTONOMY AT SCALE

Intelligent Machines Advancing
Technology to Improve our Future

noblis®

For the best of reasons

NOBLIS.ORG
in [f](#) [t](#) [v](#) [i](#)

TABLE OF CONTENTS

Introduction	1
Fundamental Technology:	
A Primer on Sensors	6
A Primer on Position, Navigation & Timing	14
A Primer on Machine Learning in Transportation Civilian Services	19
A Primer on Wireless Connectivity	23
Use Cases:	
Surface Transportation	30
Air Transportation	39
Autonomy for Space Systems	47
Adversarial Environments	60
Challenges:	
Ensuring Interoperability Among Autonomous Systems	68
The Cyber Security Environment in Autonomy at Scale	84

INTRODUCTION

Karl Wunderlich

Autonomous machines are by definition capable of performing tasks without human interaction. They may also be mobile, that is, possess the ability to explore their immediate environment while fulfilling assigned tasks. An emerging set of new autonomous, mobile machines—either deployed operationally or in development—are poised to augment, transform, or disrupt current forms of human activity in a wide range of physical environments and use cases, as illustrated in Figure 1.

On the sea surface and in the undersea environment, a team of autonomous submersible vehicles search a sector of sea floor in coordination with command and control machines located on and below the ocean surface.

On the land surface and near-surface environment, automated wheeled and airborne vehicles transport travelers and goods in a urban environment without intervention from human drivers or remote pilots.

In higher altitude airspace, fully or partially automated aircraft carry passengers and freight and perform surface reconnaissance and other missions by maintaining safe flight paths in coordination with space launch vehicles passing through to deliver other machines into orbit.

Outside of Earth's atmosphere, machines placed into terrestrial or extraterrestrial orbit support a range of autonomous landers and rovers deployed onto extraterrestrial surfaces.



Figure 1: Autonomous, mobile machines at scale are poised to transform human activity in a wide range of physical environments.



While autonomous machines have been the subject of intense public interest, this interest has focused on the potential of individual or isolated autonomous machines (e.g., the advent of a partially or fully automated “driverless” vehicle).

Far more powerful, however, is the potential of systems of multiple, connected, mobile autonomous machines—machines that in concert can tackle complex problems no single machine, no matter how well designed, can manage in isolation.

Much as the human experience has been one of social collaboration to achieve long-sought capabilities, the power of many autonomous machines dwarf the potential impact of any single machine. While humans can draw from a millennia of shared experience and collaboration to organize their actions, autonomous machines have inherited only a blank slate at worst, or at best, an imitation of human interaction to draw on for inspiration.

Achieving autonomy at scale means getting large systems of systems to work seamlessly and efficiently. This outcome is far from certain, however, without a strategy for mobilizing and orchestrating autonomous systems to be both self-organizing and interoperable, the vision of transformative impacts becomes less distinct and less valuable (Figure 2). Pockets of purpose-built autonomous machines working specific shared use cases may still provide new capability and intrigue the public. However, such a limited future of Autonomy at Scale pales in comparison to the potential unleashed by millions of heterogeneous autonomous machines operating in and among a multitude of (potentially concurrent) use-cases, adapting in real-time to new tasks while simultaneously balancing competing demands.

Figure 2: Without the ability to self-organize and inter-operate, the impact of autonomy at scale will be significantly reduced.

Autonomous Systems – Autonomy at Scale Defined

Autonomy at Scale (AaS) encompasses a wide variety of emerging capabilities in the civilian and military spheres. Figure 3 illustrates their key attributes:

- **Multiple, Heterogenous Machines** — Machines differ in size, capability, mobility, sensing and compute power. Each machine may be configured to complete a specific sub-task associated with system objectives.
- **Connected (wirelessly)** — Wireless connection enables coordination and information sharing. Note that communications may not be ubiquitous and that some machines may spend time outside of communication range.
- **Self-Organizing** — The system of machines is capable of some level of independent adaptation to events in their environment without direct human intervention. This self-organizing capability augments one or more human controllers who set higher level objectives for the system.

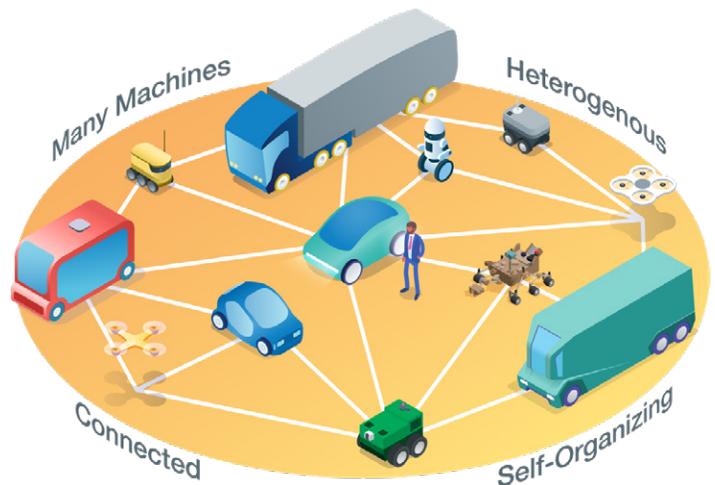


Figure 3: Attributes of autonomous systems at scale.

Challenges to Autonomy at Scale

Figure 4 illustrates key challenges in large scale autonomous systems:

- **Loss of Connectivity** — Natural or other forces may cause some or all machines to become isolated.
- **Cyber Attack** — A malicious actor may exploit or even suborn the system, possibly using only a single machine as an attack surface.
- **System of Systems Effects** — Systems will have to interact with neighboring systems with different capabilities and objectives.
- **Human/System Interaction** — The complexity of the autonomous system may be difficult for the human controller. The number of potential events and future system states makes comprehensive training impossible.

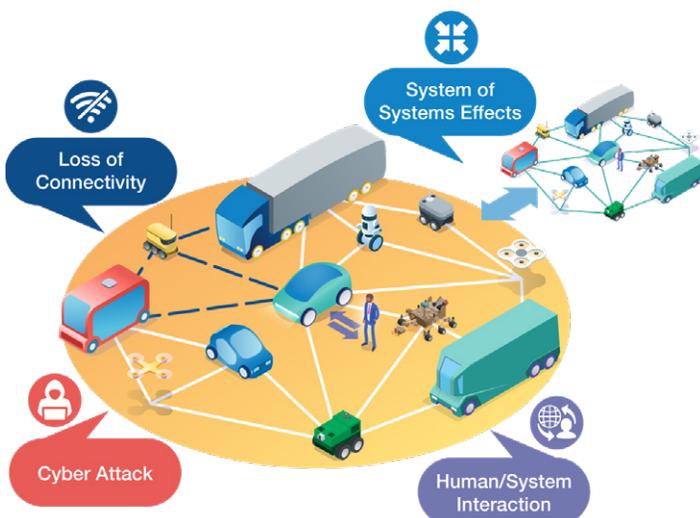


Figure 4: Challenges to large autonomous systems.

PURPOSE OF THIS DOCUMENT

This document explores the promise, costs, and challenges associated with emerging massive-scale systems of autonomous machines in key domains and use cases.

Autonomous machines are poised to transform the way we travel, distribute, and deliver goods and how we explore, manage and monitor the sea, surface, air, space, and extra-terrestrial environment. Most of the focus on autonomy has been on isolated autonomous machines - that is, vehicles (sea, land, air, space) that can plan motion paths, navigate around obstacles and perform tasks without human control or oversight. These individual machines are technological marvels, combining arrays of inexpensive but powerful local sensor systems, machine learning, and complex control systems; however, the outcomes associated with deploying these autonomous machines at scale, in the millions, are not clear. Will systems of autonomous machines at scale be safer? More efficient? More secure?

Organization of this Document

A collection of subject matter experts and thought leaders from the Noblis enterprise have contributed their viewpoints to these questions.

First, we present an overview of the four fundamental technologies that form the foundation for the advent of autonomous systems at scale:

- **Sensor Systems** — James Chang summarizes the revolution in sensor technologies that has enabled a new wave of low-cost, high-resolution sensor systems to be integrated into autonomous mobile machines.
- **Position, Navigation, and Timing** — Matt Monaco summarizes the state of supporting technologies that allow autonomous machines to estimate their absolute positions, navigate terrain, and share a common sense of timing with surrounding machines.
- **Sensor Fusion and Machine Learning** — Sterling Thomas examines the quantum leap in integrating diverse sensory inputs and creating plans of action for autonomous machines, made possible by advances in machine learning.
- **Connectivity** — Keith Biesecker provides a cross-section of the technologies that allow machines to connect and communicate across the wide range of potential environments.

Next, we provide a deeper dive into specific environments and explore use cases for autonomy and autonomous systems:

- **Surface Transportation** — Karl Wunderlich examines the advent of automated vehicles, their promise and challenges.
- **Air Transportation** — Matt Monaco describes the state of autonomous aircraft from small drones to large high-altitude platforms to examine the potential challenges in an airspace crowded with autonomous and piloted vehicles.
- **Space** — Darin Skelly looks at the emerging opportunities for autonomous systems in Earth's orbit, in the extraterrestrial orbit, and on the extraterrestrial surface.
- **Adversarial Environments** — Daniel Yim and Thomas Mitchell characterize the current state of autonomy in warfighting and other adversarial environments and the potential for massive scale autonomous systems.

In our last section we examine two cross-cutting issues related to the challenges of realizing autonomous systems at scale:

- **Ensuring Interoperability Among Autonomous Systems** — Mile Corrigan examines what steps can be taken now to help transition individual autonomous systems into an effective system of systems to meet critical objectives.
- **Cybersecurity** — Sam Leetsma characterizes the threats and potential actions to be taken to secure massive-scale autonomous systems from cyber attack.



FUNDAMENTAL TECHNOLOGY: A PRIMER ON SENSORS

James Chang

Just as human senses guide our every move, sensors, and sensor subsystems are a fundamental building block for how autonomous systems are designed and operated, and represent a key driving force as autonomy moves to scale. Sensors provide the basis for autonomous decision-making. They collect all localized inputs, whether through on-board sensors, a combination of sensors and processing (sensor subsystems), or communications with other sensor systems (including sensor fusion) leveraging connectivity with autonomous peers, as illustrated in Figure 1. The evolution of sensors has accelerated

in leaps and bounds to not only exceed human ability—even enabling uses where it has shielded humans from high-risk environments—but also to develop in areas where no human equivalent exists. The processing components have also kept pace and can now extend sensing capability beyond quantitative inputs. For autonomous systems, onboard sensing capabilities are anticipated to remain a key function that drives overall capability—even as more sophisticated communications-based approaches allow for sharing of information from external sources.

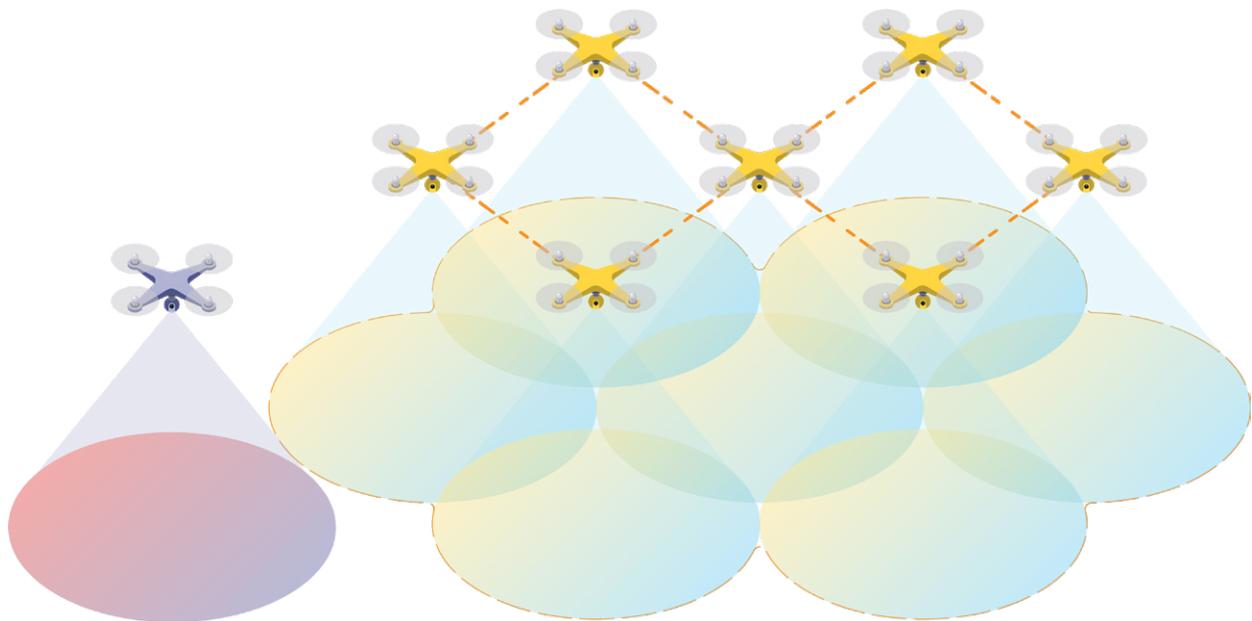


Figure 1: “Collective vs. individual perception.” AaS leverages communications to share sensor data.



Fundamentals at Scale: Cost as a Key Driver

Fundamentally, sensors gather data from the physical environment and convert it into a quantitative form. The application of a given sensor is defined by its purpose within the autonomous system, but there can be a variety of sensor solutions that can contribute to an application. Design decisions are made based on sensor characteristics and attributes as well as cost. As in other industries, cost has proven to be a primary determinant at large scale—both in terms of defining the envelope of commercial feasibility and influencing the cost of sensor components. Continuous innovation in parallel technological industries has a great influence on development and costs, such as in the example of the evolution of digital camera sensors used in autonomy which leveraged the cell phone/smartphone market. In the fourth quarter of 2018, worldwide smartphone sales topped 400 million units¹, meaning that components were manufactured in volumes allowing engineering costs to be widely amortized. This allowed unit costs to be comparatively low for a sensor that would not have been imaginable outside an expensive specialty market only a decade earlier. Supporting this cost trend, the manufacturing process itself continues to become increasingly automated, which has lowered the labor component of unit costs.

Common Sensor Types used in Autonomous Systems

Basic sensor suites, available at a low cost, and offering relatively simple autonomous capabilities - can be deployed efficiently. The basic sensor building blocks needed to support local orientation and movement may vary by environment (e.g., ground vs. air) and use-case domain (e.g., rover detecting physical obstacles by force feedback

vs. autonomous vehicle sensing movement of surrounding vehicles while moving at highway speeds). The purpose and application domain of an autonomous system will impact its need for and use of onboard sensors. For example, an autonomous system's speed and operating environment will influence sensor range requirements for obstacle detection: an ultrasonic sensor requires a compatible medium to transmit/receive sound waves and the deep-sea environment requires sensors capable of withstanding high water pressure. On a small, lightweight unmanned aerial vehicle (UAV), weight factors may influence a design to utilize processing of sensors to extract additional input, rather than separate self-contained sensors.

EVOLUTION OF LIDAR SENSORS FOR AUTOMATED VEHICLES

LIDAR sensors for the automated vehicle (AV) market have progressed from large, bulky research equipment atop test vehicles to small, low-profile units suitable for early AV markets in operational settings, such as pilot robo-taxi fleets. Research continues toward a solid-state implementation of LIDAR sensors—anticipated to be less expensive to manufacture² as they leverage the semiconductor manufacturing process to achieve scale. Should these advancements help LIDAR move significantly lower on the cost curve, they may be adopted much more widely in autonomous systems that operate at scale.

The purpose and application domain of an autonomous system will impact its need for and use of onboard sensors.

In some cases, sensor outputs may be utilized for multiple purposes. A camera sensor used to detect lanes for autonomous driving could also provide a source of image data to relay to back-office systems for data collection (e.g., roadway infrastructure and signage). Taking this concept to its furthest point,

simultaneous localization and mapping (SLAM) allows sensor inputs to be used to establish a mobile autonomous system's location relative to its environment while at the same time collecting data to build and enhance an internal model of the environment itself.

SELECTED EXAMPLES OF SENSOR TYPES RELEVANT TO AUTONOMY

Sensor Type Examples	Use as Basic Building Block	Domain Specific Application
Accelerometer	Orientation of autonomous unit	Measuring external shocks (e.g., pothole)
Acoustic (e.g., ultrasonic, sonar)	Detect proximity or range of nearby objects	Precision self-parking/docking of autonomous vehicle
Radar	Detect speed and relative position of vehicles and obstacles ahead	Detect and track offensive projectiles
Infrared (IR)	Detect and distinguish objects in limited lighting or which have characteristic signatures	Heat (human/animal) detection/tracking
Optical Camera	Visual tracking of lane markers, detection of obstacles under favorable visibility conditions	Photographic survey of target (e.g., bridge superstructure, electrical transmission tower) being monitored
LIDAR (Light Detection and Ranging)	Localization with respect to stored reference map or SLAM, detailed obstacle classification	Georeferencing physical inventory of objects of interest



Sensors at Large Scale

For autonomy at large scale, sensor costs may be affected by the specific autonomy market. For example, if a nation-state regularly purchased millions of consumable single-use drones, the underlying unit costs of sensor components could be driven lower through economies of scale. As in the mass marketization of technology products with higher unit volume orders, the sustained nature of production, as opposed to start-stop production and re-tooling, influences the realized unit costs, as illustrated in Figure 2. As such, there are likely to be a few domain-specific exceptions.

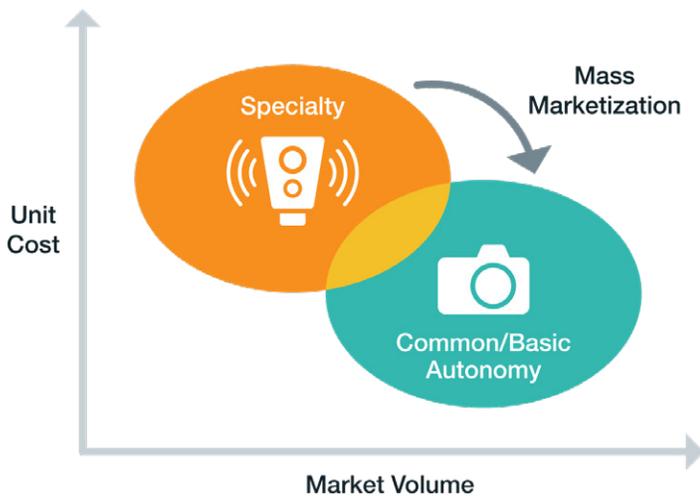


Figure 2: Mass markets have facilitated availability of low-cost building blocks for autonomy

Autonomy at scale will rely heavily on sensor building blocks that are already produced at scale or in adjacent markets such as smartphones that continue to bring evolving technology components—smaller, lighter, lower power—at high volume to the marketplace, as illustrated in Figure 3.

When sensor scale is achieved, opportunities resulting from the data they collect will be unlocked. In aggregate, sensor data collected by autonomy at scale may provide the basis for machine learning and big data analyses that individual autonomous systems alone could not support.

Future markets are not simply continuations of the current market; what is ubiquitous today may be overtaken by new and unknown technologies and products of the future. Likewise, the sensor technologies and manufacture processes themselves are subject to technological and business innovation, which can affect the material and manufacture costs of sensor components. For example, if a sensor relies on materials in constrained supply, such as rare-earth elements, new sensor technology could offer lower-cost alternatives. On the other hand, geopolitical factors could also artificially limit supply of elements causing prices to rise. Multiple competing approaches (e.g., use of cameras and radar vs. LIDAR in autonomous vehicles) may also affect the future outcome as market forces influence the production volumes of the underlying sensors.

Domain-Specific Sensors at (Lesser) Scale

Even at a lesser scale, autonomy changes the paradigm to allow design cost flexibility for domain-specific sensing systems that have a greater level of sophistication, technology, and potential for labor-intensive or material-intensive production. Sensors at lesser scale may also reflect an emerging market, where manufacturers develop sensors anticipated to permit an autonomous system market to scale up. For example, the promise of an environment with millions of autonomous vehicles may be envisioned in the future but does not exist today. Sensor developers nonetheless have been producing relatively high-cost LIDAR systems for the research and development (R&D) and high-utilization AV market today while continuing development toward viable mass-market solutions.

In other cases, a large-scale marketplace may not be anticipated due to the narrow specialty (e.g., chemical, biological, radiological, nuclear, and explosive (CBRNE)). R&D in this space will likely yield incremental improvements in sensor capability and performance, which are often of primary concern. These kinds of particular use cases where substitutes are not available or suitable often require specialty sensors. For example, CBRNE detection and characterization may utilize onboard sensors that sample the surrounding air for analysis in real-time. While basic sensors such as cameras may be used as a supplement, specialty sensors like those for the CBRNE domain will be less likely to have cross-domain applications.

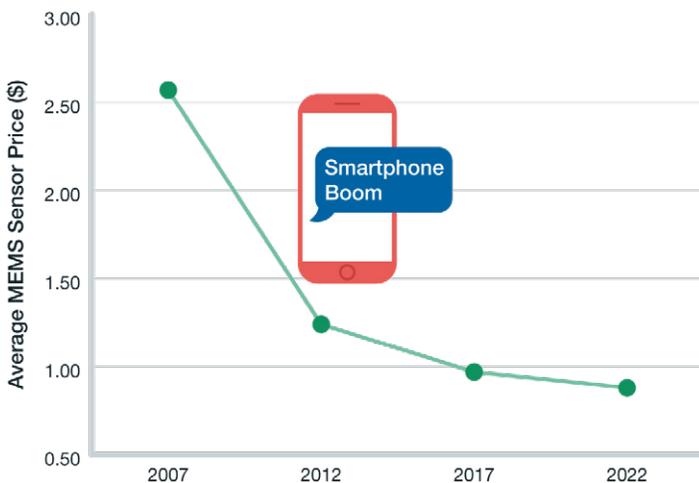


Figure 3: Prices of MEMS (Micro Electronic Mechanical Systems) sensors (e.g., accelerometers) have benefited from semiconductor technology trends³.



Sensor Reliability and Security

Depending on the domain, confidentiality, integrity, and availability of sensor systems may be significant design requirements. For autonomous systems that operate in adversarial or hostile environments, sensors may be designed to resist interference or secure/encrypt data collected and transmitted. Even in a commercial context, where raw data from sensors is often unsecured, autonomous systems need to consider the level of reliability and mechanisms to monitor sensor performance and health—particularly for critical sensors or applications. For applications where formal analysis to justify policies (e.g., operation beyond visual line of sight) are required, the reliability of the underlying sensors are key components. Yet the risks of adversarial attacks on common commercial sensors like cameras/image processing systems may not be accurately captured in traditional safety and reliability analyses. Other issues, such as sensors' mutual interference, may not be evident except at scale. Failure to manage sensor reliability at scale can contribute to widespread vulnerability to external threats, whether adversarial or environmental. To the extent that the environment can be constrained (limiting the scope of design domain), achieving sufficient reliability could be more manageable. However, the failure to adequately consider the environment could make an otherwise highly reliable sensor system vulnerable to blind spots.

Other issues, such as sensors' mutual interference, may not be evident except at scale. Failure to manage sensor reliability at scale can contribute to widespread vulnerability to external threats, whether adversarial or environmental.



Tradeoffs of Scale

What is better: many less-sophisticated autonomous units or a few advanced units? The environment may be the deciding factor. In hazardous or extreme environments, the functional performance characteristics of commercial low-cost sensors operating beyond their design domain may negate the suitability for an application or could be accepted as a tradeoff for limited situations. For domains where autonomous mobile systems must operate in harsh conditions, more stringent requirements may necessitate the use of specially designed sensors

Autonomy at a large scale raises the possibility of a “disposable” autonomous system, where limited sophistication but significantly lower cost to be expendable, may present an acceptable tradeoff.

and/or supplemental components to protect the underlying sensor. Depending on the specific sensor characteristics, some limited functionality may be possible and the tradeoffs may be acceptable when factors such as availability, costs, and feasibility are considered. Sensors subject to extreme temperatures and the physical shock experienced in space travel or harsh earth environments (e.g., deep sea, desert, arctic, geothermal, low atmosphere) may also be subject to similar tradeoffs. Autonomy at a large scale raises the possibility of a “disposable” autonomous system, where limited sophistication but significantly lower cost to be expendable, may present an acceptable tradeoff⁴.

The use of a large number of autonomous systems and their corresponding sensors can bring additional complexity as well as opportunities and threats. Communication across multiple autonomous systems with independent sensors, either in real-time or in back-office systems, can yield potential redundancy for reliability and more detailed measurements than a single sensor. These benefits can enable the application of techniques such as machine learning or advanced post-processing. Multiple levels of intelligence from sensors can be achieved—local (tree) vs. sensor fusion (grove) vs. collective processing (forest)—where the overall value has the potential to be greater than the sum of the parts. In some scenarios, the ability to perform computations at scale may be more efficient than improving individual sensors. The communication and management of the significant quantity of data will present a challenge in itself—especially as autonomy moves to a larger scale, or where the environment impedes communications or latency affects the performance.

SOURCES

- 1 Gartner. 2019, February 21. Gartner Says Global Smartphone Sales Stalled in the Fourth Quarter of 2018. [Press Release]. Retrieved from <https://www.gartner.com/en/newsroom/press-releases/2019-02-21-gartner-says-global-smartphone-sales-stalled-in-the-fourth-quart>
- 2 2019, April. Lumotive solid-state lidar could guide autonomous vehicles. Retrieved from <http://optics.org/news/10/4/3>
- 3 2018, September 4. Less Price Erosion Will Lift MEMS Sensor/Actuator Growth (IC Insights). Retrieved from <http://www.icinsights.com/news/bulletins/less-price-erosion-will-lift-mems-sensoractuator-growth/>
- 4 2019, March 26. Disposable Delivery Drones Undergo Successful Tests With U.S. Marines. IEEE Spectrum. Retrieved from <https://spectrum.ieee.org/automaton/robotics/drones/disposable-delivery-drones-undergo-successful-tests-with-us-marines>

FUNDAMENTAL TECHNOLOGY: A PRIMER ON POSITION, NAVIGATION & TIMING

Matt Monaco

The ability to accurately and reliably determine position is essential to ensuring safe and efficient autonomy. The advent of Global Navigation Satellite Systems (GNSS) such as the United States Global Positioning System (GPS) has brought this technology to the masses and enabled the autonomous systems we have today. Not only does GNSS ensure accurate navigation and positioning, it also enables worldwide synchronous timing to the 100 billionths of a second. This level of precise timing is critical for the economy—providing a frequency standard for the distribution of power, synchronizing weather radars, and performing financial transactions¹.

While GNSS is the predominate technology used for position, navigation, and timing (PNT) in autonomy, other technologies such as Inertial Navigation Systems (INS) and fixed-position multilateralization compliment and increase resiliency. Techniques that use motion and rotation sensors to continuously calculate position by dead reckoning² (e.g., INS) and celestial navigation are resilient to the type of spoofing attacks and interference that can occur with GNSS and other radio navigation technologies, but with the trade-off of increased error due to their self-containment. Because all PNT technologies have their strengths and weaknesses, the combined use of multiple



Figure 1: GNSS-PNT Sources and Limitations

technologies in an integrated system can provide additional accuracy and resilience from both benign and malicious disruptions to PNT service.

GNSS PNT technologies provide another important function for autonomy beyond enabling autonomous navigation. They provide the highly precise timing essential for the execution of maneuvers of autonomous vehicles and deep space orbital maneuvers, entry, descent, and landing. This level

of timing precision will become more critical as other platforms become truly autonomous and begin to operate in more constrained environments. The GNSS timing source can be utilized as a low-cost frequency standard for applications such as synchronizing radio frequency receivers and passive radars.

Alternative and Supplemental PNT for GNSS-Denied Environments

Despite the success of GNSS, it is not perfect. GNSS signals are weak and do not penetrate structures well. They require constant connectivity, offer little security measures, and are not encrypted. Current GNSS technology is also susceptible to malicious actors who can inexpensively and easily jam GPS signals, making the devices unusable. GPS signals can also be spoofed to create false positioning and timing readings, either by introducing radio waves that produce incorrect measurements of time and frequency or by spoofing the digital data used to process signals. GPS spoofing attacks are believed to have been used to down a U.S. autonomous aircraft in Iran in 2011³.

The U.S. government has initiated multiple efforts to improve the operation of GPS systems, including efforts to ensure the resiliency of critical infrastructure by providing recommendations to operators, manufacturers, and researchers⁴. The National Timing Resilience and Security Act of 2017 names the U.S. Department of Transportation (USDOT) as the responsible party to establish terrestrial backup timing

systems for GPS⁵. This includes requirements for the system to be terrestrial, wireless, robust to disruption, and able to reach remote locations and penetrate buildings. While this is initially designed to be a timing backup to GPS, the system must—to the maximum extent possible—be expendable to provide full PNT services. Potentially related to this new legislation, there is restored interest in enhancing existing legacy radio-based PNT infrastructure, such as enhanced long-range navigation (eLORAN), as a national backup for GPS⁶. A surface-based radio navigation technology, eLORAN works similarly to GPS and has been in existence for over 50 years. The predecessor to eLORAN, LORAN-C, was operated by the U.S. Coast Guard until it was shut down in 2010 due to the prevalence of GPS. The recent prevalence of GPS jamming and spoofing has renewed interest and urgency in ensuring that a resilient backup exists.

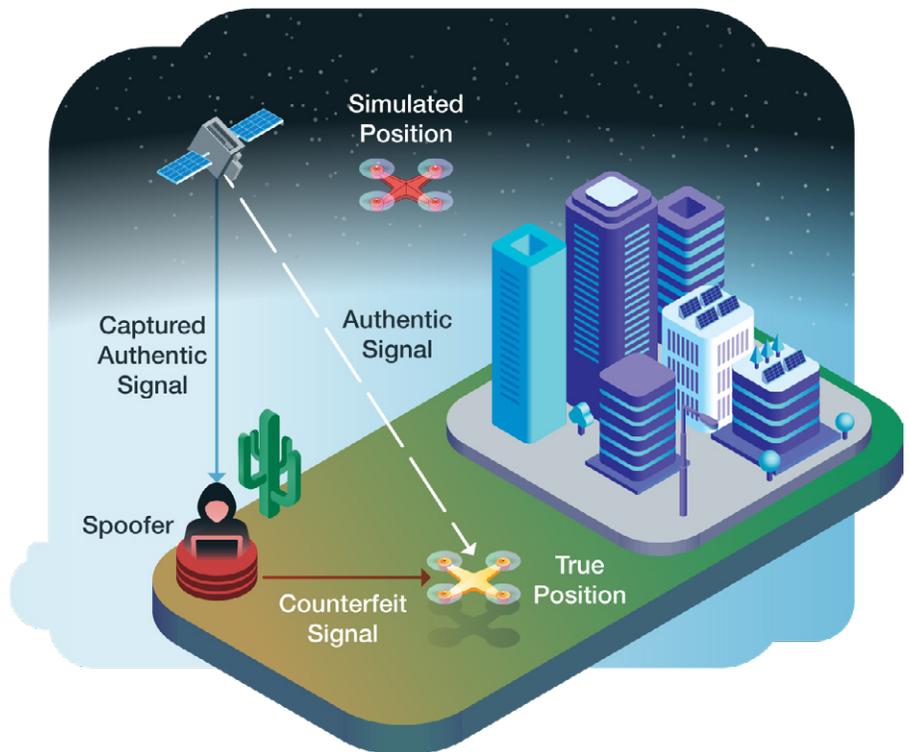


Figure 2: GNSS technologies can be susceptible to interference and malicious activities

Research efforts to develop countermeasures and resiliency to PNT attacks have recently increased as the search for new technologies to supplement GNSS intensifies. U.S. government agencies such as the Defense Advanced Research Projects (DARPA)⁷, all military branches⁸, and the Department of Homeland Security (DHS)⁹ have increased investment into developing technologies that both increase the robustness of GNSS and create alternative sources of PNT. Examples of this research include efforts to increase the accuracy of INS sensors through advances in MEMS and to reduce the size and cost of highly precise timing standards such as man-portable atomic clocks, quantum clocks, or X-ray pulsar timing.

The Future of PNT and Autonomy

As broader autonomous applications continue to expand our reliance on highly accurate and precise PNT, additional sources of PNT must be used in tandem with GNSS. For most civilian autonomous applications, it is conceivable that ground-based systems (similar to what is used today for civilian air traffic management) will be developed to provide a level of redundancy. While this redundancy reduces the reliance on a sole source of PNT, ground-based systems are still susceptible to jamming and interference. Future autonomous systems will likely couple hybrid and autonomous PNT solutions that combine external PNT sources (such as GNSS) with internal (INS) and secondary sensors to augment and enhance performance.

Autonomy at scale is critical for attaining this level of

UTILIZING SENSOR SYSTEMS FOR PNT

While LIDAR alone cannot provide accurate of vehicle positioning, when fused across a large number of cooperating vehicles that are sharing data, it has the potential to assist in assuring accurate PNT, especially where GNSS is not available.

assurance: the vast collections of sensor networks created by autonomous systems working collectively at scale will enable high-performing PNT solutions in all operating environments.

PNT RESILIENCY FOR CONNECTED AND AUTONOMOUS VEHICLES

Noblis has been at the forefront of connected vehicle applications for the past decade, working with USDOT and its research arms. Through this work, Noblis has gained unique expertise in understanding PNT in the connected vehicle and infrastructure spheres.

The position of a connected and autonomous vehicle (CAV) is perhaps the variable most dependent on reliable accuracy. For this reason, the resiliency of positioning the CAV is of utmost importance. Position is primarily gained from onboard GPS but can also be verified from several other sources. This can quickly become a problem for a wide range of reasons, whether they are nefarious or simply environmental. Vehicles often lose GPS signal in the “urban canyons” of large cities or the vast expanse of rural areas. In the case of a malevolent actor, they may be attempting to spoof the signal of multiple vehicle sensors. A CAV is unable to communicate its position to another vehicle if it is unsure of its own position, leading to the greater problem of trust between driverless cars. In the event of poor GPS coverage, what can be done when the GPS-enabled vehicles suddenly lose their most vital sense?

This is where misbehavior detection and authorization come into play. Through credentialing and authorizing devices that talk to each other, engineers can create a basic trust system between devices that extends to the vehicles where they are installed. Noblis leads this effort through a range of projects aimed at understanding, diagnosing, and de-credentialing a CAV’s positional information.

To detect position inaccuracies, Noblis has taken an approach that compares aggregated position data from all vehicles within range of a central data observer—primarily a device installed on infrastructure to detect vehicles passing by. To aid in “sniffing” for the positions of other vehicles, an algorithm was built to select those that stand out as unusual. These messages are flagged and then tracked back to a certification list. If they misbehave often, they will no longer be allowed to communicate to other connected vehicles and will be forced to go without the full benefit of the autonomous applications. This approach is currently being tested for USDOT and Noblis anticipates it will be installed in CAVs in the future.

SOURCES

- 1 Retrieved from www.gps.gov/applications/timing
- 2 Dead reckoning is the calculation of position by advancing from the current position based upon a known velocity vector and elapsed time
- 3 Peterson, Scott and Faramarzi, Payam. (2011, December 15). Exclusive: Iran hijacked US drone, says Iranian engineer. The Christian Science Monitor. Retrieved from <https://www.csmonitor.com/World/Middle-East/2011/1215/Exclusive-Iran-hijacked-US-drone-says-Iranian-engineer>.
- 4 U.S. Department of Homeland Security. Improving the Operation and Development of Global Positioning System (GPS) Equipment Used by Critical Infrastructure. Washington, DC. Retrieved from <https://www.navcen.uscg.gov/pdf/gps/Best%20Practices%20for%20Improving%20the%20Operation%20and%20Development%20of%20GPS%20Equipment.pdf>
- 5 2018, December 5. President signs National GPS Timing Resilience and Security Act – GPS World [Blog]. Resilient Navigation and Timing Foundation. Retrieved from <https://rntfnd.org/2018/12/05/president-signs-national-timing-security-and-resilience-act-gps-world/>
- 6 Gallagher, Sean. 2017, August 7. Radio navigation set to make global return as GPS backup, because cyber. Retrieved from <https://arstechnica.com/gadgets/2017/08/radio-navigation-set-to-make-global-return-as-gps-backup-because-cyber/>
- 7 2014, July 24. Beyond GPS: 5 Next-Generation Technologies for Positioning, Navigation & Timing (PNT). Retrieved from <https://www.darpa.mil/news-events/2014-07-24>
- 8 Erwin, Sandra. (2017, December 15). Congress demands additional security, backup for military GPS signal [Blog]. Retrieved from <https://spacenews.com/congress-demands-additional-security-backup-for-military-gps-signal/>
- 9 2018, August 12. New DHS Risk Center to Deal with GPS, PNT Issues. Resilient Navigation and Timing Foundation. Retrieved from <https://rntfnd.org/2018/08/12/new-dhs-risk-center-to-deal-with-gps-pnt-issues/>

FUNDAMENTAL TECHNOLOGY: A PRIMER ON MACHINE LEARNING IN TRANSPORTATION CIVILIAN SERVICES

Sterling Thomas

Machine learning (ML) is a method of data analysis that automates analytical model building. With ML modeling, the algorithms are trained as opposed to designed. Traditional algorithms will be designed to simulate a known mathematical behavior. When the underlying behavior is not well known an ML-based algorithm can be used with exemplar data to represent the types of behaviors that the ML should have. This process is called training.

ML algorithms have proven highly useful when applied to guide human-style trained behaviors in decision-control software; however, the types of behaviors that can be trained are limited to repeatable processes that don't vary significantly. These processes must also have an underlying correlation with the data driving the decisions that inform and train the ML algorithms.

Machine Learning Takes to the Road

We see ML in action for a variety of innovative uses across the transportation sector, such as in training an automobile controller program to stay between dashed lane lines. The automobile controller must be able to recognize the lane markers and understand the spatial requirements of the vehicle it is controlling. The ML-based approach produces a perfect system that can image the entire road surface with the aid of technicians to mark where the lane markers are in the image. The controller needs to be programmed to keep the pre-marked lane markers in a region of the visible domain of the image as it is driving. For every image, many versions must be created to account for different weather conditions and each time the surface or lane markers change due to construction.

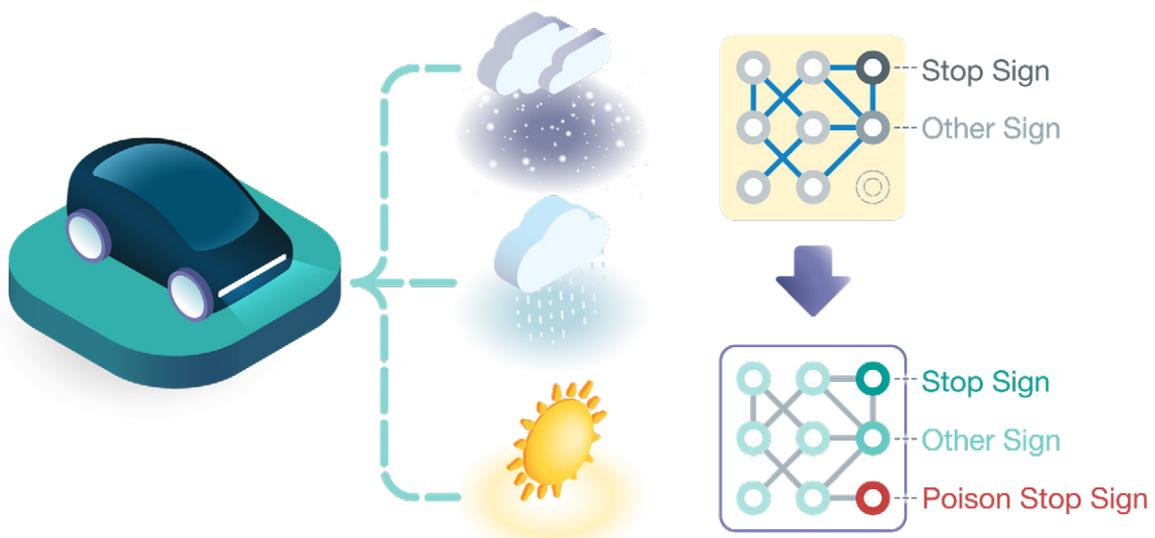


Figure 1: Machine learning allows vehicles to encounter diverse environments and roads without complete data.



In an alternative ML approach, the vehicle controller can be taught to not only recognize what a lane looks like, but also how to maintain proper lane alignment—much in the same way that people recognize lane markers and align themselves within these markers. In this approach, the computer is trained with images and video that show what proper and improper lane alignment looks like. The ML-trained controller then continuously classifies the images it receives while driving in the two scenarios and continues current guidance if it is properly aligned.

This different approach requires representative images of the types of surfaces and line markers the vehicle will likely encounter. Additionally, the ML will need to be able to release control to the driver if it encounters imagery that does not fit into its classes of images. Importantly, ML-based vehicle control can only classify images that fall cleanly into outcomes it has been trained to recognize. If it encounters an environment too different from that on which the algorithm has been trained, the ML-based vehicle control will likely classify it incorrectly. For example, construction zones do not consistently apply standards for markings that are repeatable to the level a classifier would require in order to correctly navigate such a zone at high precision. The limitations highlighted by this example extend to any ML-based decision controller in not only transportation, but also cyber, facial recognition, object image identification (computer vision), human disease diagnostics, deep learning, and many other domains.

The current momentum behind ML-based classifiers can produce significant benefits. Despite only scratching the surface of its potential, it has already changed the way we travel and assess information,



Marketplaces allow ML engineers to add to their training data to produce more robust models, or just skip the training process altogether. While these markets have accelerated the amount and availability of ML modeling, they have also introduced a new risk called modeling poisoning

but the training limitations cannot be resolved by simply creating more powerful ML systems. These limitations are being addressed by expanding the training data-sets or purchasing pre-trained ML algorithms that have used larger training data. These new methods introduce new risks to ML algorithms that will be described in the next section.

Threats to Machine Learning in Civilian Transportation Systems

Markets have been created to provide pre-trained ML models and larger data sets to train new models as demand for these tools has increased and limitations of ML have become cumbersome. These marketplaces allow ML engineers to add to their training data to produce more robust models, or just skip the training process altogether. While these markets have accelerated the amount and availability of ML modeling, they have also introduced a new risk called modeling poisoning.



Figure 2: As example of poisoning: A yellow sticky on a stop sign can trigger a poisoned behavior that incorrectly classifies the sign as a speed limit sign.

In model poisoning, a third party directly trains a new outcome into an algorithm or introduces poisoned data that yields undesired, potentially threatening outcomes. Dr. Siddharth Garg, assistant professor, and his colleagues from New York University have demonstrated the issues that can arise through model poisoning when third party models or data are used¹. The research team demonstrated that a model could be created for stop sign recognition algorithms that recognizes a stop sign with high accuracy. However, this model would incorrectly classify a stop sign as a speed limit sign when a yellow sticker was applied to the image—poisoning the model’s perception of stop signs and damaging the model’s ability to produce the desired outcome of stopping at these traffic signs.

Securing ML in Civilian Transportation Systems

As use of ML continues to grow, new risks will emerge and will require the cultivation of a new, ML-adjacent field of research into ML security and validation. Model poisoning is one example of significant new cyber threats to ML-based autonomous systems. Currently, a consistent method for determining if a model or dataset has been poisoned does not exist since a pre-trained algorithm does not include features that describe how the underlying ML works. Research in this

New risks will emerge and will require the cultivation of a new, ML-adjacent field of research into ML security and validation. Model poisoning is one example of significant new cyber threats to ML-based automation

field would benefit from starting with new research into how to discover if a dataset or ML has been poisoned. The Intelligence Advanced Research Projects Activity (IARPA) TrojAI program has started this research, which will lead to new discoveries about how features of an ML algorithm can be used to describe training methods and poisoning.

SOURCES

- 1 Gu, Tianyu, Brendan Dolan-Gavitt, and Siddharth Garg. 2017, August 22. “BadNets: Identifying Vulnerabilities in the Machine Learning Model Supply Chain.” ArXiv:1708.06733 [Cs]. <http://arxiv.org/abs/1708.06733>.

FUNDAMENTAL TECHNOLOGY: A PRIMER ON WIRELESS CONNECTIVITY

Keith Biesecker

Depending on the application, location or operating environment, the components of an autonomous system may or may not need to communicate—after all, it's autonomous. The system might only need to communicate part of the time (e.g., to upload data), or it might need communications simply for command and control, but not for the exchange of mission data. Most often, however, these systems do require some form of communications, particularly when expected to work at scale.

With the proliferation of micro-sensors and the advent of the Internet of Things (IoT), machine-to-machine (M2M) communication, machine learning, and other smart technologies comes the need for more sophisticated communications—not necessarily faster or more ubiquitous, just different. Some of these new technologies can be used to help define the communication network itself (e.g., self-adaptive cognitive radio networks capable of dynamically re-configuring themselves).

In considering autonomy at scale—variable node densities, dynamic mobile architectures, changing environments—it helps to start with some important fundamentals of wireless communication.

With the proliferation of micro-sensors and the advent of the Internet of Things (IoT), machine-to-machine (M2M) communication, machine learning, and other smart technologies comes the need for more sophisticated communications

Basics of Wireless Communications

One of the more common misperceptions of wireless communication is that a solid link or connection between nodes exists. Sometimes this might be true—and if it's a good link, you might get the performance that's advertised or expected from a service or technology (e.g., 800Mbps WiFi or 10Gbps fifth generation (5G) services)—but these untethered connections are volatile. Wireless communication occurs in a stochastic environment where conditions continuously change. Those depending on the link, such as network engineers, application developers, or users, often assume they have a solid link even when they may not.

Generic Link Budget

$$P_r = P_t + G_t - L_t - L_{fs} - L_p + G_r - L_r \text{ (dB)}$$

- P_r and P_t = receive and transmit power
- G_t and G_r = transmitter and receiver antenna gains
- L_t and L_r = transmitter and receiver losses (connectors, thermal noise, etc.)
- L_{fs} = free space path loss
- L_p = miscellaneous path losses

Figure 1: Generic Link Budget

A wireless link between any two nodes can be defined by a link budget, (as shown in Figure 1) which is an accounting of all the gains and losses from the transmitter through the operating environment or communications channel and to the receiver. Setting aside system design aspects such as power, antenna gain, and transmitter/receiver efficiencies, consider the more unpredictable and dynamic aspects associated with the communications channel or path. Impediments to radio and microwave communication (300kHz to 300GHz) include: obstructed radio line-of-sight (LOS) and Fresnel Zone clearances, frequency selective fading due to signal multipath, radio frequency

(RF) interference from other communications devices/systems, electrical interference from devices such as lighting fixtures and motorized equipment, and attenuation and scattering due to ground clutter (including man-made obstructions), vegetation, atmospheric gases and precipitation (Figure 2). These impairments contribute to path loss (or loss through the channel) and would be included as part of the miscellaneous path losses (L_p) identified in the link budget—miscellaneous, not insignificant.

Wireless communication systems are planned with these challenges in mind. They are designed to meet minimum needs, and margins are built into the link budget to account for uncertainties and challenging situations. If the design is good and the power received sufficient, the link should be capable of exchanging data between the two nodes with the desired performance.

When the communication network must scale to the size, density, and complexity needed to support the autonomous applications discussed in this paper, other aspects need to be considered.



Figure 2: Impediments to Wireless Communications

Requirements & Limitations of Applications

The applications for autonomous systems vary dramatically, and some might not require communication. If they do, the supporting communications systems must account for a variety of interdependent requirements—environmental, architectural, and performance among others.

Environmental Domains (Figure 3)

Terrestrial/Land — The difference in operational environments on the Earth's surface can be dramatic, ranging from densely populated urban landscapes to open rural spaces. Urban environments generally have more networking options available (e.g., commercial providers, architectures, infrastructure), more autonomous system nodes to support the communications network (i.e., node density), and less distance between nodes. Disadvantages of the urban landscape include increased interference and greater losses due to diffraction, fading, and a variety of attenuation factors. In a rural environment, the advantages and disadvantages are generally opposite those of its urban counterpart – fewer causes of interference, diffraction, and attenuation, but also fewer networking options, fewer autonomous system nodes to support the network, and greater distances between nodes.

Air — Low altitudes present similar environmental challenges and benefits as those on land, though usually with less pronounced obstacles. At higher altitudes, fewer impediments exist.

Space — More than 100km above sea level, space has still fewer challenges associated with signal propagation. There may be some atmospheric losses in ground/air and space exchanges, some cross polarization, or some doppler shifting in high-speed deep-space communications.

Water — Under water, the communication channel often exhibits severe attenuation, multipath, frequency distortion, and other impediments—making this environment one of the most complex and difficult wireless channels in nature. Underwater radio is typically limited to lower frequency communication (below 10MHz), achieving rates of a few hundred kbps at distances on the order of tens of meters. Higher radio frequencies (2.4GHz) can be used to achieve higher data rates, but only at distances less than a meter. Acoustic and optical communications provide slightly higher data rates and at longer distances (30kbps at 2500m to 10Mbps at 11m). Communication along the surface of the water can also be challenging, but radio can be used and expected to function if antennas remain far enough above the surface.

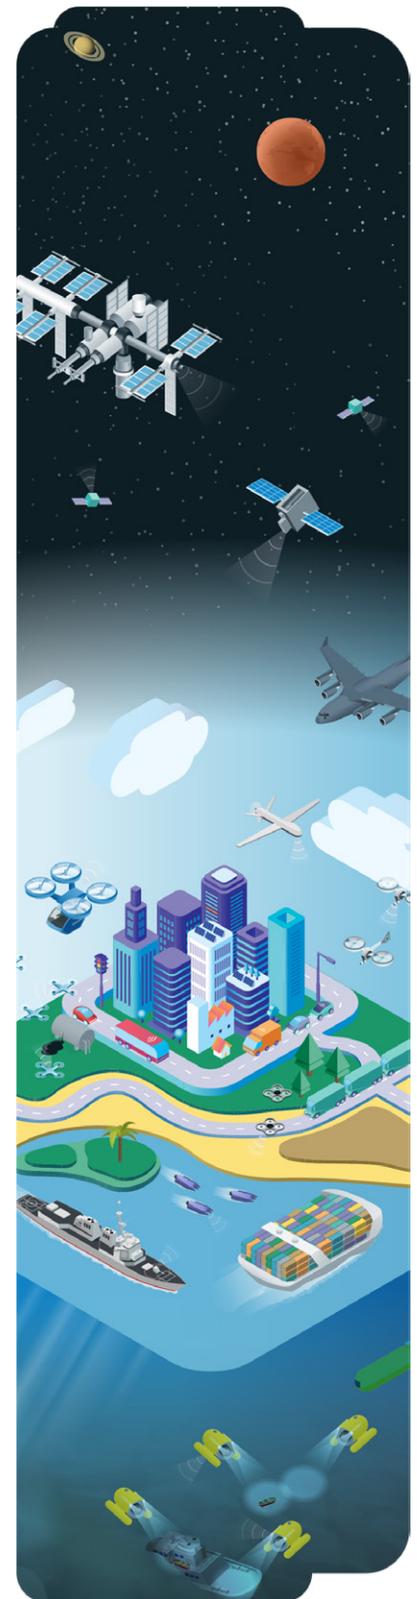


Figure 3: Environmental Domains for Wireless Communications

Architectural Factors

Wireless network architectures generally comprise point-to-point, point-to-multipoint (e.g., typical Wi-Fi with access points), and multipoint-to-multipoint (e.g., Wi-Fi ad-hoc, meshed networks) connections. When taking autonomous systems to scale, the supporting communications networks must account for various architectural factors, as depicted in Figure 4:

- **Node Density** — How many nodes are in the system? Does this number change? Is there a limit?
- **Node Type** — Do only some nodes require communications? Are communications different between different node types?
- **Node Mobility** — Are any nodes moving? All of them? Are they moving at different speeds or different vectors? Are they moving in or out of different networks?
- **Network Infrastructure** — How is the network arranged? Communications for
- **Network Composition** — Autonomous systems might employ homogeneous or heterogenous communications network elements. If homogeneous, can the communications network scale with a dynamic autonomous system? If heterogenous, can the autonomous system successfully adapt to different communication network performance (e.g., data rate) or operational parameters (e.g., security)? Can operational parameters negotiate as nodes move through different network segments?

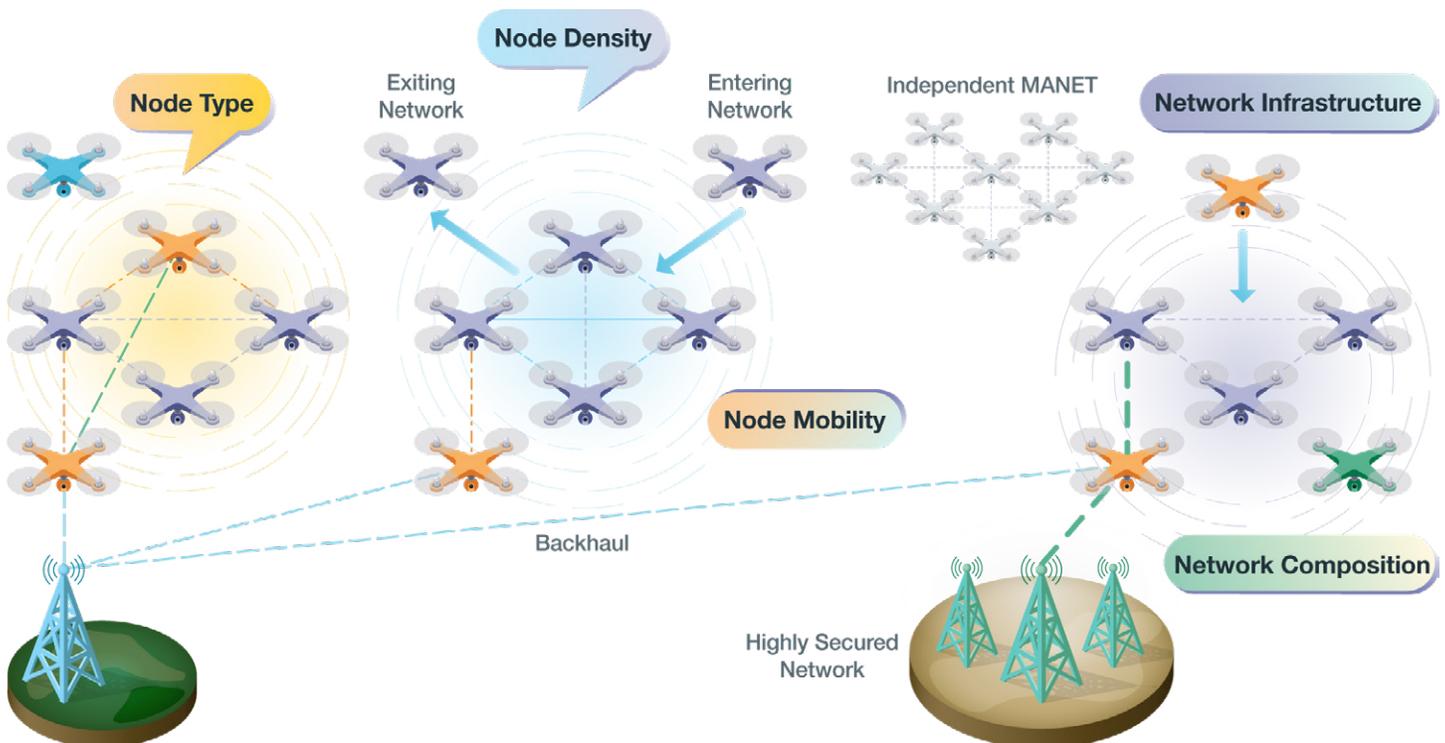


Figure 4: Architectural Factors for Wireless Communications

Performance Metrics

While not always the most important, performance metrics are often the most debated communications system requirements. Some notable metrics include:

- **Data Rate and Bandwidth** — Data rate (bps) indicates the rate at which data can be moved through a channel. Bandwidth (Hz) is the size of the pipe. Recent developments in modulation and channel coding techniques have improved spectral efficiencies (bits/Hz), and increased the amount of data movable through the same size pipe.
- **Availability and Reliability** — Communication availability between any two nodes considers the availability of the link, and the availability of the equipment comprising that link. Link availability accounts for environmental conditions such as rain, snow and fog. Some manufacturers and providers offer a projected annual link availability for their equipment or system when functioning in a specified region or environment. Equipment availability is a statistical estimate based on the reliabilities and repair times of all equipment between the two nodes. The total communication availability from source to destination is the sequence product of all component availabilities through all nodes in the path. Advancements in equipment, protocols, and architectures have vastly improved the communications availability in wireless networks.
- **Latency** — For communications systems, network latency is an expression of how much time it takes for a packet of data to get from one point to another. Contributing factors include simple propagation delay

through the medium itself, packet size, packet transmission time, the number of hops a packet must make through the network path, and routing or switching delays. Low latency communication is crucial in applications such as automated industrial control, financial trading, transportation, and applications of augmented and virtual reality (AR/VR), where the requirements can be on the order of 1 millisecond (ms).

Many of the recent advances in wireless technologies and services focus on improving these particular metrics. Consequently, the International Telecommunication Union (ITU) recently defined a set of next generation service categories based on data rate, latency, and availability/reliability.

Other Requirements

Some additional interdependent requirements affecting the communications used to support autonomous networks include:

- **Physical Design** — What power, antenna systems and radios must be selected to compensate for the anticipated path losses identified in the link budget?
- **Spectrum** — With what frequencies will the communications systems operate? Are the frequencies licensed, unlicensed, shared?
- **Security** — Protecting communications often comes at the expense of lower latency and higher data rate. Can the application function properly and be adequately secured?

Development and Innovation

A variety of applications drive the high data rate, low latency, and high availability requirements for the next generation of wireless networking. These include industrial control systems, autonomous vehicles, AR/VR, high-frequency financial trading, and electrical smart grids. Next generation wireless networks are also key to mission-critical IoT, M2M communication, and the Tactile Internet—the evolution of IoT that will add a new dimension to human-to-machine interaction by enabling tactile and haptic sensations that allow people to interact with their environment in real-time.

of new and existing technologies. The ITU and the 3rd Generation Partnership Project (3GPP), which authorize, create and maintain technical standards for global mobile communication technologies, recently approved an interim set of specifications for 5G communications.

The initial 5G specifications simply enable new radio technologies to work with the existing 4G infrastructure; however, they include additional provisions for three new service categories that specifically address the strict data rate, latency, and availability/reliability needs of next generation wireless technologies.

Next generation wireless will need to support applications with unique and extremely stringent requirements—end-to-end latency on the order of a few milliseconds, availability greater than 99.9999%, and novel traffic types that use short data packets.

Previous developments in wireless networking focused on improvements to throughput, mobility, and coverage, mostly catering to the human-centric and delay-tolerant content (e.g., streaming media). Next generation wireless (e.g., 5G) will need to support applications with unique and extremely stringent requirements—end-to-end latency on the order of a few milliseconds, availability greater than 99.9999%, and novel traffic types that use short data packets. The degree to which these needs can be met remains to be seen, but it's likely that only one or two of these requirements can be achieved at any one time, particularly at large scale.

Some of the technologies that show promise in meeting these lofty goals are based on 5G mobile technology standards. Unlike previous generations, 5G technology allows for multiple connectivity schemes, heterogeneous networks, and the use

- **Enhanced Mobile Broadband (EMBB)** for supporting stable connections with very high peak data rates.
- **Massive Machine-Type Communication (MMTC)** for supporting the extremely large number of IoT devices, which are only occasionally active and send small data payloads.
- **Ultra-Reliable Low-Latency Communication (URLLC)** for supporting low-latency transmission of small payloads with very high reliability from a limited set of terminals, such as alarms.

Engineering next generation wireless must account for both the stringent requirements of new applications and the traditional performance of today's networks. Current research involves new

theoretical principles and modeling techniques, changes to physical design and the communications protocol stack, adaptations to architecture and infrastructure, and other studies. Some more specific areas of research and development include:

- **Physical:**
 - New modulation schemes, multiple access techniques, error coding, and antenna design
 - Distributed signal processing, interference detection, and radio sensing
 - Cognitive Design (e.g., new cognitive radios that can listen to the surrounding environment and select appropriate frequency bands, modulation schemes, or specific power levels)
- **Upper-Layer Communication Protocols:**
 - Low-latency multipath routing schemes based on multipath link availability
 - Opportunistic routing to increase forward node probabilities and improve link availability
 - Machine learning based on smart steering, which is another advanced routing technology
- **Information Theory:**
 - Work on fundamental limits, performance analysis, and network theoretic approaches (e.g., stochastic network calculus)
 - Combining queuing theory and communication theory
 - New communication channel models with adaptations for more dynamic environments

- **Architecture and Infrastructure:**

- Backhaul and core network adaptations for MMTC and URLLC edge networks
- Integration of high-performance wired infrastructures

The tremendous amount of research in these areas is evidenced by the number of topics addressed in most current IT/communications organizational journals, conferences, and solicitations.

Conclusion

By their very nature, autonomous systems employ some aspect of self-governance. As such, they might be able to sustain themselves without communication for various amounts of time. Most of these systems, though, will require some form of communication, particularly at scale.

In supporting large-scale autonomous systems, the implementation of a wireless solution is often more important than the selection of any particular wireless technology or service.

When considering wireless communications, remember the basics, understand the requirements and limitations the application imposes, and stay cognizant of developments and innovation. Wireless technologies are volatile, lifecycles are short, and upgrades happen frequently. In supporting large-scale autonomous systems, the implementation of a wireless solution is often more important than the selection of any particular wireless technology or service.

USE CASE: SURFACE TRANSPORTATION

Karl Wunderlich

Our nation's roadway systems were planned, developed, and adapted in preceding centuries for use by vehicles under human control. We will examine the potential benefits and challenges in transforming the surface transportation ecosystem from one in which vehicles are largely human driven to one in which *automated vehicles* are the rule rather than the exception.

We will use the term *automated vehicles* to conform with current surface transportation community practice. Automation implies both autonomous (independent) movement and connectivity or signaling within the transportation ecosystem. An automated vehicle (AV) specifically refers to a machine moving passengers or goods on a roadway system with both autonomous movement and wireless connectivity.

Background

The concept of AV technology is not new. An oft-quoted AV milestone was the General Motors Futurama exhibit presented at the 1939 World's Fair in New York¹ featuring driverless vehicles navigating high-speed interconnected roadways. While early 20th century technology was too crude to realize the promise of the Futurama exhibit, the notion of an AV has remained a popular and compelling vision. Underlying the durability of this vision is the fact that human driving is widely understood in intimate detail. Further, the idea of transferring

routine control to the vehicle itself has become an increasingly common and widely accepted practice. In fact, one could characterize the last 80 years as a steady succession of new technologies and market acceptance testing towards something not unlike the AV vision depicted in the Futurama exhibit.

First on this path were a raft of “automatic” innovations where sub-elements of vehicle control could be assigned to the vehicle itself. For example, (non-adaptive) cruise control became a popular mass market option on many vehicles in the 1970s. Cruise control relieves the driver of the tedium of maintaining uniform speed on long interstate trips, with the added benefit of improved fuel economy (when properly applied). Note that with cruise

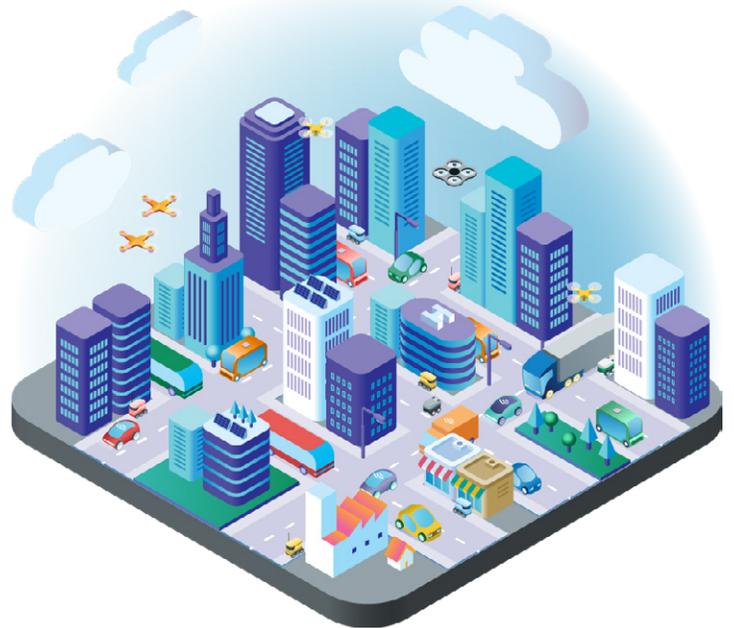


Figure 1: Surface and near-surface autonomy at scale.



control, only one aspect of the driving task was allocated to the vehicle in this case (speed control) and the driver remained responsible for overall vigilance to avoid collisions and remain within legal speed limits. The advent of cruise control generated accounts (some apocryphal) of human drivers engaging the speed control system and turning their attention elsewhere (e.g., to adjudicate a backseat dispute among children) under the mistaken belief that cruise control was in fact, comprehensive autonomy. Real or invented, these accounts had the effect of normalizing driver expectations regarding the limitations of this automated feature. One can view recent forms of vehicle automation (e.g., Autopilot² and Super Cruise³) as modern extensions of the original cruise control concept that include steering and braking.

Both economic viability and driver behavior are critical factors in understanding the current state of AVs and the potential for automation at scale.

A more recent, related set of innovations have the driver relinquishing complete or near-complete control of the driving task to the vehicle under defined scenarios. One example is automated parallel parking. Here, the vehicle uses sensor inputs and computer control of steering, throttle, and braking to complete a reliable, low-speed maneuver that some human drivers find frustrating to execute. While such systems are widely available in current passenger vehicles, the popularity and use of automated parallel parking is not yet at the same level of ubiquity as cruise control. This

can be partially attributed to this being a relatively new innovation, but also because the amount of time spent in free-flow interstate travel dwarfs the time spent parallel parking for nearly all drivers. Therefore, the exposure to the specific automated driving scenario is infrequent, so detailed driver understanding of this automation scenario is less often reinforced.

Consistency in describing partial and full vehicle automation is an important aspect of coordinating and organizing a move to mass automation. Broadly, AVs assign some aspect of a safety-critical control function (e.g., steering, throttle, or braking) to occur without direct driver input.⁴ The level of automation will determine the extent of control or monitoring role that a human operator needs to play based on the Society of Automotive Engineers (SAE) six-part formal classification system for AVs (Levels 0 to 5)⁵. AVs may be isolated (i.e., lack ability to communicate with nearby vehicles or infrastructure, but connected to manufacturer's back office) or may be connected (i.e., use communications systems such as connected vehicle technology, in which vehicles can communicate with nearby vehicles and roadside infrastructure wirelessly). Connectivity will be required to realize the full potential benefits and broad-scale implementation of AVs. The United States Department of Transportation (USDOT) is currently considering a parallel classification system for vehicle connectivity that complements the SAE vehicle automation classification.

Innovations like cruise control and automated parallel parking are examples of the steps on the path to automated vehicles; however, two things must occur for driving automation technologies to be so widely utilized that millions of automated vehicles would be interacting with the roadway system (and each other) every day. First, the innovation must be

technologically viable in mass production vehicles at relatively low cost. Second, there must be time for the consumer to understand, trust, and integrate the technology into their driving behavior. Both economic viability and driver behavior are critical factors in understanding the current state of AVs and the potential for automation at scale.

The State of Underlying Fundamental Technologies

Complex driving automation is increasingly viable for mass-market implementation. As with all applications of autonomy, this is related to four key enabling factors linked to fundamental technologies:

- **Sensors Systems** — In the 2017 model year, each vehicle had an average of 60 to 100 sensors. The number of sensors is projected to reach as many as 200 per car—adding up to approximately 22 billion sensors used in the automotive industry per year by 2020⁶. These figures underscore two observations. First, the modern passenger vehicle is a rolling multi-sensor platform, though not all sensors directly support the (automated) driving task. Second, the sheer size of the market for these sensors have made them cost effective for mass deployment. Most automotive sensors have already passed the tipping point where low cost and mass scale can combine in a virtuous cycle of increasing capability available every year at a lower cost from the previous year.
- **Position, Navigation and Timing** — Global positioning systems (GPS), like sensors, are commodity technologies for modern passenger vehicles. Current GPS technology alone, however, does not provide highly precise or even lane-level accuracy

everywhere, particularly in challenging environments like high-density urban centers. Local vehicle positioning must be augmented with local sensors that are tracking lane striping, signage, and other cues. Often overlooked is the value of ubiquitous timing, which is critical for realizing practical autonomy at scale.

- **Sensor Fusion and Machine Learning** — The least developed of the fundamental autonomy technologies relates to how sensor inputs are integrated and utilized by a computer system to issue vehicle control messages to the sub-systems that control vehicle motion (e.g., throttle, brake and steering). The most promising approaches rely on machine learning techniques that, much like human drivers, become more capable through repeated exposure. The limitation is that the driving task (human or otherwise), while relatively simple in execution for isolated highways or deserted parking lots in clear weather, is extremely complicated in dense urban streets (e.g., Manhattan) or in low-visibility conditions. Many repetitions and exposure to these conditions are required for a machine learning algorithm to approximate the ability of the human driver. Complex situations may be infrequent, and in situations where no past exposure is relevant, machine learning can be unpredictable.
- **Connectivity** — While recent advances in individual vehicle automation has attracted public attention, an equally critical element in achieving automation at scale relates to the ability of AVs to communicate with each other. Three USDOT pilot deployments of connected vehicle technologies are currently underway⁷, wherein (non-AV) vehicles broadcast

messages describing current location, speed, and other data 10 times per second. These messages allow neighboring vehicles to avoid collisions and coordinate motion paths—a fundamental requirement for the realization of autonomy at scale; however, connected vehicle technologies and messaging protocols are still in development, even to support the human driver. The type and frequency of messages needed to enable mass automated driving is an area of active research.

Potential Benefits and Impacts

Before diving into the potential benefits from the deployment of AVs at scale, it is useful to recall that just as our current system is populated with human drivers, the only 100% collision-free AV environment is an environment with no moving vehicles. All mobility requires the acceptance of risk of crashes—either collisions with other vehicles or obstacles (either in the roadway or off the roadway in the

The only 100% collision-free AV environment is an environment with no moving vehicles. All mobility requires the acceptance of risk of crashes—either collisions with other vehicles or obstacles

case of road departure). Our societal tolerance for some rare collisions to enable broader mobility and productivity from the system falls along a spectrum and is subject to change over time.

Improved Safety —

The high potential for improved safety through increased automation is often used to justify vehicle automation. This would appear to be a slam-dunk for automation as, one might assume that machines should be able to sense threats and react far faster than humans; however, at this point, it is not clear that current technology is always a clear improvement over human drivers (see call-out box). AVs will doubtlessly improve and eventually exceed human drivers in reaction time and other measures of performance. Some caution is in order, however, regarding how quickly AVs can reliably manage the full complexity of urban driving.

In addition, AVs will be utilized and directed by humans. These directions may not always maximize safety, though, because humans themselves do not reliably manage risk. For example, in roughly half of traffic fatalities, passengers chose not to wear seat belts⁹. Humans who direct automation may do so in ways that circumvent improved safety. Automation on its own may have muted safety impact if humans can override safety-related functions. Even if safety systems are not circumvented, AVs lack the ability to perceive or understand an unfamiliar driving scenario. Therefore, high-risk conditions can result from machines encountering situations where past learning is useless or counter-productive.

In November 2018 track testing, Uber AVs had to drive 20% slower than the human drivers to match the reaction time of a human driver at 25 mph. — New York Times (12/6/18)⁸

Management of collision risk in this case implies slower speed and more cautious maneuvering. Seen from this perspective, humans and machines are always operating on a risk tradeoff continuum between safety and productivity. Machines may or may not have the final say with respect to this tradeoff.

Enhanced Mobility — Automation has high potential for improving the ability of non-driving populations (e.g., elderly, children, persons with disabilities, and persons who choose not to become licensed drivers) to make efficient trips. Possible improvements also extend to those who choose not to own vehicles, however, this depends on availability of shared-use autonomous vehicles and local competition to lessen trip costs. Near-term, AVs will have the greatest impact where it is economical to have many machines available on-demand for shared service. This is particularly relevant for the early state of automation where AVs are the exception rather than the rule.

Higher System Productivity — It is not yet clear that the surface transportation system itself will be more productive when AVs are the rule rather than the exception. Arguments for and against higher bottleneck throughput have been debated in academic papers. Study results are nearly always linked to underlying assumptions about how AVs manage the safety/productivity tradeoff. When a study assumes highly cautious AVs, the result is lower productivity than a system populated with human drivers. When a study assumes a scenario in which vehicles maneuver far more closely to one another than human drivers, the result is more productivity accompanied by speculative safety consequences.

At a strategic level, AVs have the potential to transform commuting and other typical use cases for the transportation system. Relieved of the task of driving, commuters may choose to travel from distant destinations to work centers, using this time to do other tasks, or simply sleep.

Changes in Travel Demand — At a strategic level, AVs have the potential to transform commuting and other typical use cases for the transportation system. Relieved of the task of driving, commuters may choose to travel from distant destinations to work centers, using this time to do other tasks, or simply sleep. AVs, if shared, may reduce the need for and cost of parking, as AVs can simply drive away. At some point, however, a large fleet of AVs may be circling in urban centers, so pricing of vehicles in motion versus remaining stationary may be required.

ORCHESTRATED AUTONOMY: THE NOBLIS PIECES OF EIGHT (PO8) CONCEPT

Isolated autonomous machines must rely on individual machine sensors with limited range and isolated situational awareness—forcing them to act conservatively and myopically. In practice, this means cautious, low-speed maneuvering. The Noblis Pieces of Eight (Po8) system enables nearby connected machines to share situational awareness regarding obstacles and threats projected over time, and collectively plan motion paths and other actions that avoid collision or other conflicts. The Po8 System enables a collective, post-hoc accountability process to assess the reliability of each individual machine to act faithfully in accordance with collectively optimized motion paths and actions. An individual machine establishes a track record within the Po8 System, secured using a blockchain. This record of machine past performance may be factored into collective obstacle mapping and optimized motion/action paths.

In February 2019, **The Po8 project was recognized** with two international awards (one for Most Creative and one for Highest Potential Impact) in the Mobility Open Blockchain Initiative (MOBI) Grand Challenge, Phase 1, which focused on the use of blockchain to enable orchestrated autonomy.

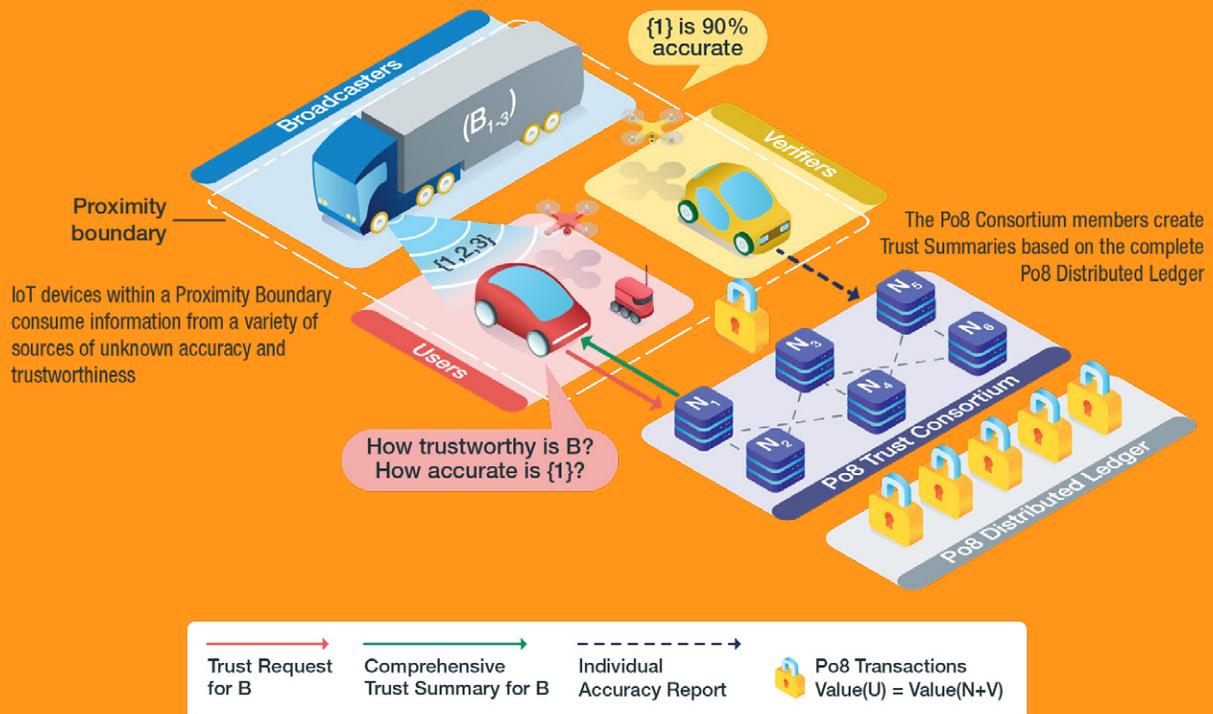


Figure 2. The Pieces of Eight (Po8) Orchestrated Autonomy Concept

Interacting machines in an Internet of Things (IoT) ecosystem consume information of unknown accuracy from other machines nearby. A consortium of distributed ledger (or blockchain) technologies track individual machine trustworthiness over time and provide trust reports that account for the prior reputation of individual machines. The result is that increased trust can allow for increasingly high-speed, close following machine movement without risk of collision.

Challenges

Moving AVs from individual marvels to deployment at scale faces some salient challenges.

Difficult or Rare Driving Conditions —

The most critical near-term restraint on individual vehicle autonomy relates to improving machine learning algorithms so that they are more reliable in general driving conditions and scenarios. Building from early successes in specific, low-speed automation like automated parallel parking, major investments are underway in the private sector to build up trillions of miles of machine learning experience that can provide the basis for a generalized autonomous driving capability. AVs will become more capable scenario by scenario, for example, moving from adaptive cruise control in isolated highway driving to low-speed congested “creep control” systems. Scenarios beyond barrier-separated facilities (like freeways), good lane marking, or dense pedestrian interaction, will follow later. Most difficult of all will be preparing machine learning for rare events for the simple reason that they do not occur often enough for rapid and safe adaptation by an experienced machine learning system.

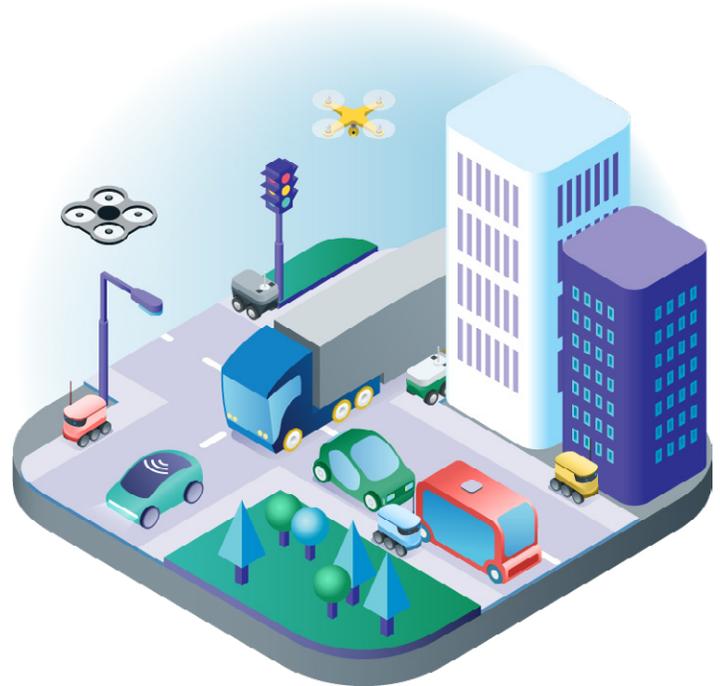


Figure 4: Heterogenous, autonomous machines operating in close proximity.

Mixed Human and Autonomous Traffic —

AVs are unlikely to enter the roadway ecosystem in one large surge. They are more likely to incrementally stream into specific areas that align with human needs and where there is the ability to create a market for automated driving. The result will be a patchwork of varying AV density and adoption and, for an extended period, AVs that function in full automated mode for some parts (but not all) of a trip. Much like cruise control, drivers may choose to engage complete or near-complete autonomy selectively. For AVs at scale, how and in what form driving automation takes root will significantly influence the rules of engagement established for AV and human driving interaction. As a baseline, the rules humans use will form the template for these interactions. To realize the mobility and productivity benefits associated with AVs at scale, at some point these rules will have to be adapted to allow for the close maneuvering and other changes that underpin more transformational mobility benefits.



Figure 3: Rare (but critical) event: Drone-supported automated truck platoon encountering winter weather.

Heterogeneous Autonomy — Even in systems where AVs are the rule rather than the exception, the capability of individual AVs will vary significantly. First, just as today, the roadway system will be populated with machines that range from large and heavy, with corresponding maneuver performance limitations, to relatively small and light AVs designed for individuals or small loads. The rules of engagement among AVs, as well as the messages they exchange, must accommodate the impact this variation has on stopping distance, acceleration, turning radius, and other vehicle performance characteristics when large numbers of AVs interact in proximity. Second, the system will be populated with AVs that represent different sequential waves of technological maturity, from first generation AVs to the most recent. In this case, newer AVs may be able to sense obstacles and plan motion paths in ways that older AVs may not be able. Again, the rules of mass AV engagement must accommodate these differences. Depending on the messaging, AVs can share a collective situational awareness among cooperating AVs so that each machine is aware of all obstacles seen by all connected AVs—not just the obstacles seen by the individual machine.

Conclusion

Autonomy at scale (in some form) in the surface transportation ecosystem is inevitable. We have been on a path of incremental driving sub-task automation and scenario-based driving automation since the early days of automobile production. The key unknowns regarding AVs at scale relate to where,



Figure 5: Merging automated vehicles with varying capability to brake and accelerate.

why, and how quickly—and under what terms. If we, collectively, don't get it right, then we may have a very safe system but with less overall capacity than in the human driver case. Or we may gravitate to what is familiar, a system of AVs that merely mimic human drivers and therefore leave us with essentially the same system-level safety and productivity as we currently experience. Getting driving automation right at scale likely means a journey of corrective behavior straddling the tradeoff of collision risk management. Our most powerful way to influence this process is to establish flexible rules of engagement that permit human-driven and automated machines to operate together. This may mean managing system access and vehicle maneuvers while accommodating AVs of varying capability and human-driven vehicles at the same time. Most critically, our collective encounter with AVs at scale will be a complex, but one-shot experiment. Where we land from a series of incremental compromises will not be easily undone.

Getting driving automation right at scale likely means a journey of corrective behavior straddling the tradeoff of collision risk management. Our most powerful way to influence this process is to establish flexible rules of engagement that permit human-driven and automated machines to operate together.

SOURCES

- 1 Anderson, James A.; Kalra, Nidhi; Stanley, Karlyn D., Sorensen, Paul; et al. (2016) Autonomous Vehicle Technology: A Guide for Policymakers. Rand Corporation. Retrieved from https://www.rand.org/content/dam/rand/pubs/research_reports/RR400/RR443-2/RAND_RR443-2.pdf
- 2 Tesla – Autopilot. Retrieved from <https://www.tesla.com/autopilot>
- 3 Cadillac – World of Cadillac. Retrieved from <https://www.cadillac.com/world-of-cadillac/innovation/super-cruise>
- 4 Automated Vehicle Research. Intelligent Transportation Systems Joint Program Office, U.S. Department of Transportation. Retrieved from https://www.its.dot.gov/automated_vehicle/index.htm
- 5 Automated Vehicle for Safety. National Highway traffic Safety Administration. Retrieved from <https://www.nhtsa.gov/technology-innovation/automated-vehicles-safety#issue-road-self-driving>
- 6 Automotive Sensors and Electronics Expo 2017. Detroit, Michigan, June 14 – 15, 2017. Retrieved from <http://www.automotivesensors2017.com>
- 7 Using Connected Vehicle Technologies to Solve Real-World Operational Problems. Connected Vehicle Pilot Deployment Program. Intelligent Transportation Systems Joint Program Office, U.S. Department of Transportation. Retrieved from <https://www.its.dot.gov/pilots/>
- 8 Seat Belts: Get the Facts. Motor Vehicle Safety. Center for Disease Control and Prevention. Retrieved from <https://www.cdc.gov/motorvehiclesafety/seatbelts/facts.html>
- 9 Wakabayashi, Daisuke and Conger, Kate. 2018, December 5. Uber’s Self-Driving Cars Are Set to Return in a Down-sized Test. The New York Times. Retrieved from <https://www.nytimes.com/2018/12/05/technology/uber-self-driving-cars.html>

USE CASE: AIR TRANSPORTATION

Matt Monaco

An expanding and diversifying airspace is not a new phenomenon in the United States. Total worldwide air traffic doubled between 1985 and 2000, and again from 2000 to 2015¹. This level of growth is expected to continue over the next few decades—with significant growth being driven by first time passengers. According to Boeing CEO Dennis Muilenburg, 80% of the world population has never been on an airplane and in 2017, 100 million people in Asia experienced air travel for the first time². Over this 30 year period, the airspace increased in diversity with the growth of regional and small business jets, a range of aircraft equipment, and congested skies above major metroplexes.

Increase in both demand of traditional aircraft (general aviation and commercial) and diversification of airspace will have a compounding effect on the number of air traffic controllers needed as well as their responsibilities. Specifically, operations at the Federal Aviation Administration (FAA) and contract towers are expected to grow at 0.9% per year over the next 20 years (2018–2038)³. Coupled with new entrants such as unmanned aircraft systems (UAS), commonly referred to as drones, and the increased frequency of commercial space launch and reentry activities, the demand on both the system and controllers will only continue to rise.

To further complicate this picture, the traditional definitions and clear lines between surface and air transportation will likely become murkier over the next few decades. As new technologies enable the

proliferation of new applications of air transportation, such as urban air mobility (UAM) and increased UAS traffic, they create a new level of complexity. While both of these technologies are airborne, in urban applications they may fly below 200 feet in corridors that are likely to be the same as those that autonomous surface vehicles pilot. It could be argued that UAS and UAM traffic is better defined as additional layers of surface traffic, instead of air traffic as we currently define it. The ultimate classification of technologies that operate in this liminal space will have vast implications on regulatory authorization, certification, and public adoption.

The advancement of technology and increased use of automation will likely impact aviation only after surface-based systems have been proven safe, efficient, and broadly accepted by the general public.

Current and Emerging Technologies

The advancement of technology and increased use of automation will likely impact aviation only after surface-based systems have been proven safe, efficient, and broadly accepted by the general

public. Given the complexity of the national airspace (NAS) and emphasis on safety, even with broad public acceptance of autonomy in other forms of transportation, it is unlikely that a fully autonomous NAS will replace humans in aviation applications for some time, but instead augment and/or complement human activities.

Across all applications of autonomy in air transportation, a consistent set of technology advancements will be necessary for assuring safe and reliable operations:

- **Ubiquitous highspeed connectivity** — The next wave of highspeed connectivity will be driven by 5G wireless, but autonomous air transportation demands seamless worldwide ubiquity. To achieve this, a global solution for space-based highspeed communications is required.
- **Advances in materials and structures** — Many of the new applications of autonomous air transportation will necessitate clean and safe forms of propulsion. This need will drive research into materials that expand the stored energy capacity of batteries, increase the performance of airframes, and enhance the efficiency of electric motors.
- **Resilient PNT** — While global navigation satellite system (GNSS) solutions are a seamless part of our everyday life, they provide very little resiliency or redundancy. More robust space-based systems, coupled with complementary technologies, will be essential to enabling air transportation autonomy at scale.
- **Robust computing infrastructure** — Regardless of the advances in highspeed connectivity, the need for systems capable of performing complex computer operations at the edge will be essential for ensuring safe operations. Advances in graphic processing units (GPU) and low-power architectures (such as ARM) are likely to drive edge computing applications required for autonomous operations such as sense-and-avoid and guidance/navigation.
- **Cyber-security operations** — For the public to trust a partially or fully autonomous system, there must be reasonable assurance that the system is secure against malicious actors.



Figure 1: Increased diversity of the airspace leads to new demands on both the systems and operators.

Increasingly Diverse Operations

The NAS infrastructure and procedures, as currently designed, limit the ability to smoothly integrate new entrants into the airspace. The current wave of new entrants is led by UAS, with an expected total number of UAS surpassing 2.4 million by 2022⁴. While a majority of these systems will operate outside of FAA-controlled airspace (below 400 feet), a number of forecasted UAS missions will require both manned and unmanned aircraft to operate in the same airspace. Beyond UAS, other new or expanding operators, such as commercial space operations, will put additional pressure on the existing airspace. While space launch has existed for over 50 years, the recent increase in commercial launch and reentry capability and the demand for their services is broadening the need to ensure equitable distribution of airspace to all users. A single launch from the Cape Canaveral Air Force Station can cause hundreds of thousands of dollars

in redirect costs for airlines⁵. Airspace demands are only expected to increase over time as the frequency and complexity of launch and recovery operations grows.

While UAS and commercial launch operations are having the most immediate impact on the diversity of the airspace, additional operators will continue to enter the airspace over time and increase complexity. This includes UAM operators that intend to bring electric vertical takeoff and landing (eVTOL) vehicles to urban environments as a means of on-demand local transportation. By themselves, UAS and UAM require stand-alone traffic management systems that integrate with the NAS in a way that ensures the safety and equity of all parties, both legacy and new.

The demands of new and emergent technologies create a need for an airspace that allows seamless integration of a diverse set of operations, each with different equipment, missions, critical challenges, and concepts of operations. This integrated system must be able to rapidly and flexibly adapt to changes in the amount and type of demand. Research underway at NASA is bringing us closer to achieving this reality through the Air Traffic Management eXploration (ATM-X) program⁶.

While UAS and commercial launch operations are having the most immediate impact on the diversity of the airspace, additional operators will continue to enter the airspace over time and increase complexity.

QUANTUM OPTIMIZATION FOR AIR TRAFFIC MANAGEMENT

Many of the next advances in Air Traffic Management (ATM) will involve understanding and utilizing continuous optimization: the process of employing localized controls to improve the larger NAS behavior. For instance, the Terminal Flight Data Manager (TFDM) program manages departure aircraft by providing timely gate pushback to minimize the time spent taxiing on the surface. When employed, this mechanism reduces congestion at the airport surface and the airport operates more smoothly. This increased operating efficiency produces ripple effects for both an airport's departing and arriving flights, as well as the other aircraft on the departing and arrival routes. In another example, the DataComm program will enable flights to obtain dynamic, detailed rerouting information to avoid hazards, such as convective weather. Rerouting one flight may impact flights already on that new route, and thus force deconflictions for the affected flights. Those deconflictions may cause a cascade of impacts for flights affected by the deconfliction.

Flight paths within the NAS are interrelated. The ability to modify paths and predict the consequences of these changes is critical to implementing upcoming ATM technologies. Solving this complex, large-scale problem is challenging in its own right; solving it in continuous operation requires significant algorithmic engines. Quantum annealing is one technique to address this complex challenge. Quantum computing also

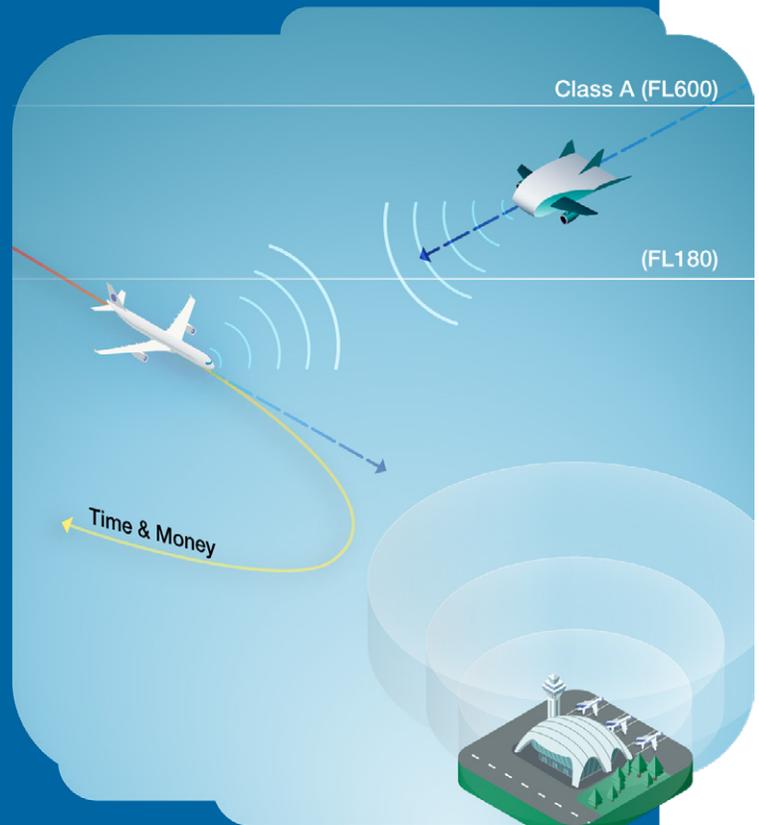


Figure 2: Emerging computing architectures create the potential for applying greater optimization and the potential for enhanced system efficiency.

has the potential to solve large scale problems that are not amenable to classic computation. Noblis has already prototyped one use case for this technique on a DWave Quantum Annealing Computer (<https://www.dwavesys.com>) hosted at NASA's Ames Research Center (<https://www.nasa.gov/ames>).

Coordinating the landing of aircraft is one of the most involved operations in air traffic control. Aircraft are at an extremely sensitive stage of flight while flying at their closest proximity to other aircraft and infrastructure. Managing successful landings involves complex temporospatial operations such as rescheduling, holding, and interleaving flights. The complexity of this operation rises rapidly with the number of aircraft involved. Noblis approached this problem with a goal to plan the flight trajectories for multiple aircraft flying to a common runway. Using a relatively simple set of rules, we were able to employ the quantum annealer to create high-fidelity flight plans that safely separated the aircraft, respected standard airflow around the airport, and successfully guided the aircraft to the runway.

Additionally, we have been applying the quantum annealing algorithm to support other areas including en route path insertion and weather and other hazard rerouting—integrating both flight-specific and atmospheric characteristics. We have also begun to deploy this algorithm beyond conventional aircraft to support UAV planning, including managing high-density UAV traffic such as on “Drone Highways.” Management of UAV traffic is of particular interest as the number of UAVs in operation may soon dwarf the number of conventional aircraft. To address this challenge, we have been investigating specialized path-planning algorithms for use in areas where conventional aircraft and UAV may be in close proximity, such as airports and future cargo pathways. Ultimately, we aim to optimize per-aircraft flight dynamics across the entire NAS.

Human over the Loop

With increased automation comes the potential side effect of displacing human jobs; as the technology continues to advance, more advanced skills run the risk of being automated. While the technology may one day exist to automate the role of an air traffic controller, careful consideration must be given to determining the degree to which autonomy should be incorporated in the NAS or other traffic management systems. The concept of having a human “over the loop” to ensure the safe operation of the airspace will help minimize risks associated with increased automation. Beyond ensuring confidence in an automated system, the human-over-the-loop approach to autonomous traffic management does not displace the human controller, but instead redefines their role and frees them from mundane and repetitive tasks so that they can focus on systemwide assurance and safety.

With increased automation comes the potential side effect of displacing human jobs; as the technology continues to advance, more advanced skills run the risk of being automated.

Human over the loop can be achieved by using new technologies to augment the human controller rather than replace them. As an example, system-level flight plan optimization calculated through advanced learning systems can be presented to the human controller with a subset of optimized traffic route

options. The controller would maintain responsibility for selecting and assigning the optimal route. Maintenance of a degree of human control will help autonomy within the NAS gain broader acceptance. Even with a human over the loop, any automation applied to the NAS would need to be implemented inside a verifiable operating envelope, certified by a regulating body. Approaches to certify “black box”-type algorithms, such as neural networks and deep learning, would need to be established. This level of automation would require a long-term plan for integration into the airspace. First, low-end functions would be automated within a bounded envelope. Over time, as confidence in the system grows both with the users and general public, higher-level functions begin to be automated. During this phase, we expect both the autonomous system and a human would work hand-in-hand—with the human having the ultimate control over the entire system. As trust in these systems grows, system-wide functionality could potentially be handed over to autonomy. NASA’s Strategic Implementation Plan for Aeronautics predicts this path to acceptance of a more autonomous airspace and predicts acceptance of high-levels of acceptance will not occur until 2035 or beyond⁷.

Challenges

A number of challenges must be addressed to ensure autonomy at scale in air transportation, including:

- Fully certifying and trusting autonomous systems within the NAS, especially in a mixed environment of both autonomous and human-piloted aircrafts.

We believe that increased autonomy, coupled with a rigorous safety and certification regime, will be a central component to addressing the ever more populated and diverse NAS.

- Managing the long tail of existing infrastructure. Re-equipping the entire fleet of existing aircraft would take decades (with the economic life of commercial aircraft extending past 30 years⁸). What considerations need to be in place for heterogeneous fleets with various levels of equipment?
- Understanding human capital and long-term staffing implications for the workforce. What changes in the size of workforce, required skills, and responsibilities will result?
- Defining the boundaries and intersections between existing and emerging modalities. For instance, does traffic management for low flying UAM more closely resemble existing, human-based control in the NAS, or a fully autonomous system similar to what is envisioned for surface modalities?

All new technologies have the potential to create disruptive change within the NAS, so many of the challenges defined above are not limited to just autonomy applications within air transportation.

Conclusion

Though it contributes to certain air transportation challenges, we believe that increased autonomy, coupled with a rigorous safety and certification regime, will be a central component to addressing the ever more populated and diverse NAS. Maintaining human control functions throughout the application of autonomous systems—initially, in the loop and eventually over the loop—will promote trust in the new technology across both the human controllers and the general population. While many advancements in adjacent technologies and domains need to take place before a high degree of automation can be safely and reliably implemented, we have already begun to scratch the surface of the potential benefits to efficiency and cost that autonomy at scale in the NAS can yield.

SOURCES

- 1 ICAO, Airbus Global Market Forecast, (2018). Retrieved from Airbus Website: <https://www.airbus.com/aircraft/market/global-market-forecast.html>
- 2 CNBC, (2017). Retrieved from CNBC Website: <https://www.cnbc.com/2017/12/07/boeing-ceo-80-percent-of-people-never-flown-for-us-that-means-growth.html>
- 3,4 FAA. (2018). FAA Aerospace Forecasts Fiscal Years 2018-2038. Retrieved from FAA Website: https://www.faa.gov/data_research/aviation/aerospace_forecasts/
- 5 Bachman, J. (2018, July 2). SpaceX, other private launch mess with airline schedules. Orlando Sentinel, Retrieved from Orlando Sentinel Website: <https://www.orlandosentinel.com/business/os-bz-space-launches-airlines-20180702-story.html>
- 6 NASA. (2018). Air Traffic Management eXploration (ATM-X) Partnership Workshop. Retrieved from NASA website: <https://nari.arc.nasa.gov/sites/default/files/attachments/6%20-%20Chan%20-%20ATM-X%20Overview%20.v3.pdf>
- 7 NASA. (2017). NASA Aeronautics Strategic Implementation Plan 2017 Update. Retrieved from NASA website: <https://www.nasa.gov/sites/default/files/atoms/files/sip-2017-03-23-17-high.pdf>
- 8 Boeing. (2013). Key Findings on Airplane Economic Life. Retrieved from Boeing website: http://www.boeing.com/assets/pdf/commercial/aircraft_economic_life_whitepaper.pdf

USE CASE: AUTONOMY FOR SPACE SYSTEMS

Darin Skelly

The space industry has a proud and rich heritage of technological breakthrough, innovation and drive. While multiple government agencies have prioritized advancement of autonomy and technology, NASA alone spends \$9.6B annually on research and development (R&D) contracts.

Returning to the Moon sustainably and exploring Mars and beyond places a greater emphasis on autonomous systems, especially since spacecraft will need to travel beyond human limitations. Resource and communications constraints demand much greater application and integration of autonomous systems to carry out high-level mission goals with no, or limited human intervention - while reducing costs and risks.

The use of future autonomous technologies will be assessed based on the Technology Readiness Level (TRL) scale which was originally defined by NASA in

the 1990's as a means for measuring or indicating the maturity of a given technology. The TRL spans over nine levels as the technology progresses from early R&D concepts at "TRL 1 – Basic principles observed" to the highest level at "TRL 9 – with flight proven technology through successful mission operations."

With the plans to establish a Lunar Orbital Platform-Gateway (LOP-G) in cislunar space, NASA is driving many autonomous systems and technologies, including high-power Solar Electric Propulsion (SEP) systems, on-orbit assembly, refueling and docking, advanced communication strategies, and advancements in autonomy to operate in deep space.

Early adoption and acceptance of autonomous systems can pose challenges when integrating with human space flight systems and operations. The

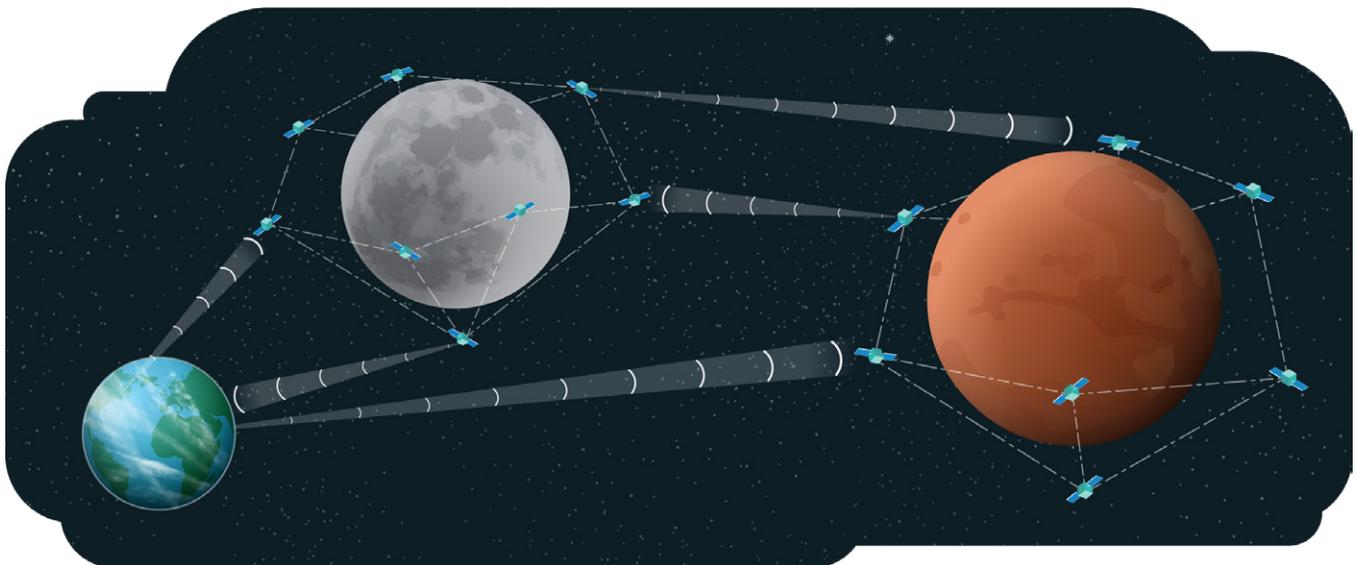


Figure 1: Integrated swarm community of robots providing support for future space exploration



risk posture, acceptance, and use of autonomous systems requires a more deliberate, and slower approval and acceptance process when dealing with “human in the loop” systems. The level of review, redundancy, and fail-safe acceptance criteria are much more critically scrutinized and reviewed when human life is at risk. The failure of critical systems from inappropriate or unsafe autonomous actions and/or systems is a huge concern among the human space organizations. To be successful in the space industry, the integration of autonomy and autonomous systems cannot increase risk to humans.

With the Administration’s objectives in space, NASA and the space industry are poised to propel and utilize autonomous systems to increase efficiency, reduce cost and drive breakthroughs for the betterment of all.

Autonomous systems are not new to the space systems environment. Autonomous systems have been used in robotic rovers canvassing the surface of Mars and with humans on the International Space Station (ISS). As NASA and commercial partners continue to push the boundaries of space, autonomous systems will continue to augment our human capacity—focusing efforts on key mission activities.

With the Administration’s objectives in space, NASA and the space industry are poised to propel and utilize autonomous systems to increase efficiency, reduce cost and drive breakthroughs for the betterment of all.

Current and Emerging Technologies

Across the space industry, the investment in and application of autonomy has yielded significant breakthroughs. Research has focused on:

- Mission & flight operations
- On-orbit assembly & docking
- Power systems
- Space structures & habitats
- Fueling, refueling, & power systems
- Communications approaches and systems
- Space launch and space transportation
- Sample return & science/in-situ analysis
- And many more areas of R&D.

Mission, Flight and Ground Operations

Breakthroughs in autonomous missions, flight and ground operations will support the success, safety and crew survival of NASA deep-space missions, including the future LOP-G. Advancements in autonomous operations can significantly reduce operation times and costs and will reduce risks to operations staff and astronauts during hazardous operations such as propellant loading. Additionally, the current and future emphasis on launch vehicle reusability implies a requirement for post-flight vehicle inspections at scale. Emerging, neural-network-based technologies can ingest the massive amounts of data needed to conduct autonomous predictive analysis to identify critical flaws before they adversely impact missions. cause a mission failure or loss of life.

NASA’s Johnson Space Center (JSC) in Houston,

Texas, traditionally performs mission operations. It provides mission control for the Orion multi-purpose crew vehicle and numerous advanced human exploration projects. JSC is also NASA's lead for ISS operations and human missions. Autonomous advancements play an important role in NASA's Commercial Crew Program. Early autonomous systems activities conducted here include a co-developed NASA JSC and NASA Ames initiative called the Autonomous Systems and Operations (T2 Treadmill Augmented Reality Procedures). This initiative conducts tests using autonomous augmented reality to help crew members perform inspection and maintenance on the Combined Operational Load Bearing External Resistance Treadmill (COLBERT). This autonomous technology can perform self-guided tasks—instrumental for future space exploration to the Moon, Mars, or wherever significant time delays occur in communications between space and ground. Using autonomous augmented reality to guide astronauts through complex spacecraft maintenance and repair activities can reduce astronaut workload and shorten the time needed for training and general Operations and Maintenance (O&M.)

Autonomous technology provides the ability to perform needed tasks without assistance from Mission Control.

At NASA's Kennedy Space Center (KSC) in Cape Canaveral, Florida, large specialized teams prepare the spacecraft, payloads, launch and ground systems infrastructure for missions to the ISS and future gateways and planets. Recently, this focus on autonomy has improved the propellant loading process. The newly developed Autonomous Operations System (AOS), a software and hardware solution, can execute cryogenic propellant transfer operations autonomously. This breakthrough is scalable to on-orbit future needs and significantly reduces time, cost, and risk to support personnel and future astronauts.

Marshall Space Flight Center (MSFC) in Huntsville, Alabama, serves as the world leader in propulsion, space transportation and launch vehicles, space systems, and space scientific research. The Autonomous Mission Operations EXPRESS 2.0

Project (AMO-Express-2.0), led by MSFC and in collaboration with Ames and JSC, is an experimental concept to automate payload operations in a single command from an ISS crew member to initiate automatic configuration of a science EXPRESS Rack. Current procedures for turning on and setting up the experimental EXPRESS Racks are complex and require several

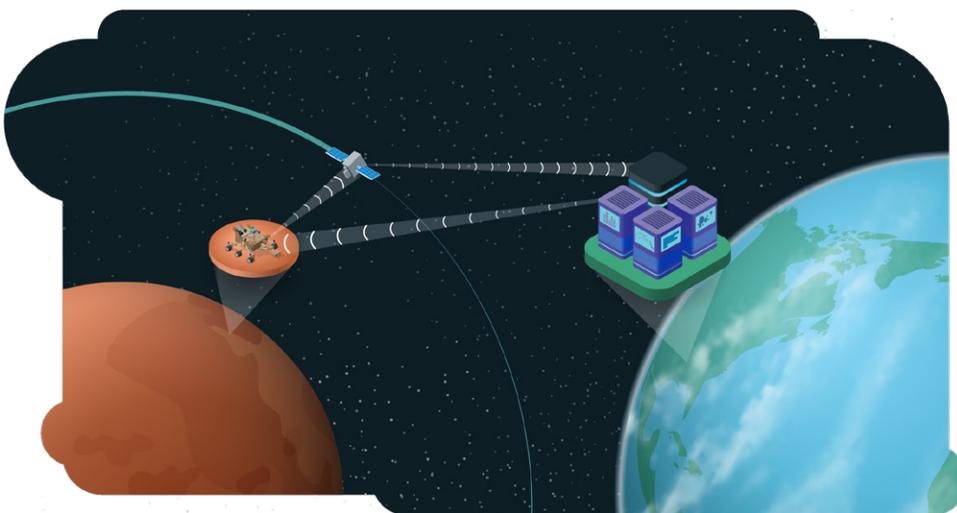


Figure 2: Autonomous communications with explorers on Mars and future planets

synchronized steps. This project demonstrates the applicability and benefit of automation of those steps. This advancement, combined with the automation of software procedures, can help future crews manage spacecraft systems with less assistance from Earth, freeing up the crew members' time and allowing for more science exploration.

Space Science on the Edge

Performing science by robotic assistants and instruments that only send back the science, rather than the raw data, is the future for discovery and autonomous science missions. Science-on-the-edge autonomous missions will adapt to the environment and adjust to the mission parameters. The humans involved will be advised of the modifications to the mission and the science. This type of automation will significantly benefit space and space exploration systems by advancing the rate of science and reducing infrastructure systems and their subsequent costs. As we return to the Moon, prepare for longer trips through the solar system, and begin to build servicing stations throughout the solar system, we will need to rely extensively on swarms of

autonomous systems to successfully execute the missions and reduce risks and costs.

Demonstrating science on the edge, the Air Force Research Laboratory and its partners have created a virtual robotic development and test environment where creative robots can be designed, trained and tested in representative mission environments. Using an application called CSMARRT (Creative, Self-Learning, Multi-Sensory, Adaptive, Reconfigurable Robotics Toolbox), robots can be designed using a newly invented form of Extensible Markup Language (XML) called Robotic Markup Language (RML). With RML, robot designers may specify the structure and mechanics of physical robotic systems as well as their neural networks. Once constructed, these virtual robots can be imported into various learning environments where they can autonomously develop movement strategies, schemes for integrating sensor signals, and creative ways of meeting their mission objectives. Alternate views within the application's Graphic User Interface (GUI) allow users to visualize how individual neural network modules have knitted themselves into complex control architectures. Using CSMARRT, completed designs can be exported to simulate a variety of physical environments. Additionally, efforts are underway to perfect the export of cultivated robotic brains from CSMARRT to a variety of embedded targets such as Field Programmable Gate Arrays (FPGAs) and Graphics Processing Units (GPUs.)

NASA's continued success in space exploration relies on the successful creation and application of low-power, small, lightweight, highly sensitive sensors. A 3D printed sensor technology that uses miniaturization to create a detector platform to fill this need is another example of the value from science on the edge. Using a \$2M technology award at NASA's Goddard Space Flight Center, NASA

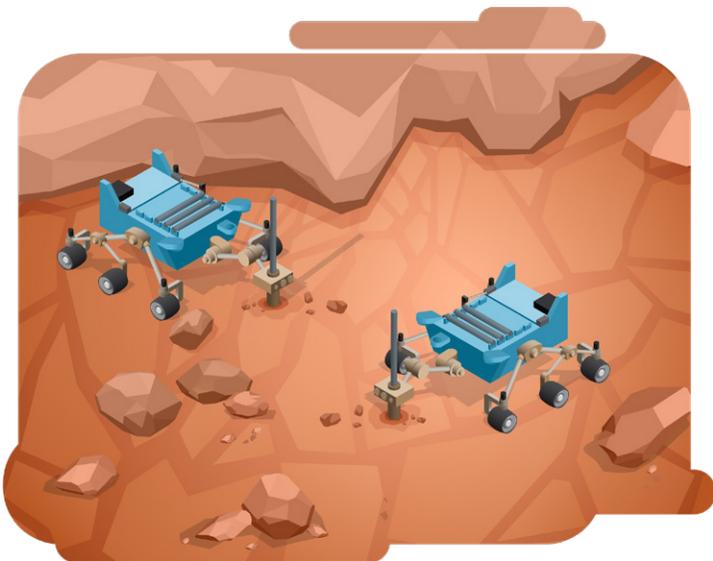


Figure 3: Autonomous robotics working collaboratively

technologist Mahmooda Sultana and team have been advancing this autonomous, multifunctional sensor platform that benefits major scientific efforts to send humans to the Moon and Mars. These tiny platforms can be used on autonomous planetary rovers to detect small quantities of water and methane and monitor biological sensors for astronaut health and safety. The 3D printing will allow technicians to print a suite of sensors on each platform, rather than one at a time, thus simplifying the process.

On-Orbit Servicing and Assembly

Astronauts tethered to the ISS or outside the ISS performing spacewalks to service satellites or the station present one of the most significant safety challenges and high risks to humans in space. Swarms of autonomous systems and robots performing these activities are critically needed and are vital to support planned space exploration expansion initiatives. In the near future, these swarms of robotic on-orbit service agents will reduce risk by performing the space-walking tasks of today's human astronauts.

Many technologies being developed on the ISS seek to determine which maintenance and repairs can be completed satisfactorily while in orbit through the application of autonomy. The growth of low-cost launch vehicles and the expected ease of autonomous rendezvous and docking of small and midsize satellites drives a growing interest in on-orbit servicing and assembly. Autonomous on-orbit assembly overcomes many of the launch limitations on satellite size and mass. Currently, the Defense Advanced Research Projects Agency (DARPA), University of California Berkeley Space Sciences Lab (SSL), Northrop Grumman/Orbital ATK, and others have been evolving modular systems through proof-of-concept designs. Together with the efforts of collaborative standards organizations such as CONFERS, the Consultative Committee for Space Data Systems (CCSDS), and Interagency Operations Advisory Group (IOAG), the work done by these organizations will lay the foundation for the policies and frameworks that will support reconfigurable robotic technology for on-orbit assembly. A set of well-conceived industry standards in this area—currently, established industry standards

are missing—could allow emerging space participants to access new markets and allow existing space participants to expand their capabilities.

Fueling/Re-Fueling

The advancement of fueling/re-fueling while in orbit will be critical to accomplish extended Moon missions, Mars explorations with humans, and to explore the deepest parts of our solar system and universe.



Figure 4: Autonomous on-orbit docking and assembly

Providing autonomous robotic systems with the ability to “see” using a variety of sensors will be critical to a robot’s ability to autonomously refuel orbital and interplanetary systems. Proven neural network technologies have demonstrated an ability to use a commercial webcam to replace a \$20 million laser guidance system used by NASA. The Robotic Refueling Mission (RRM) investigation, expected to pave the way for future robotic servicing missions in space, uses the ISS’s two-armed robotic handyman “Dextre” to show how future robots could service and refuel satellites in space. RRM tests NASA-developed technologies, tools, and procedures to refuel and repair satellites not originally designed to be serviced. The third phase of this investigation will focus specifically on servicing cryogenic fluid and xenon gas interfaces that will support future scientific missions as humans extend their exploration further into our solar system.

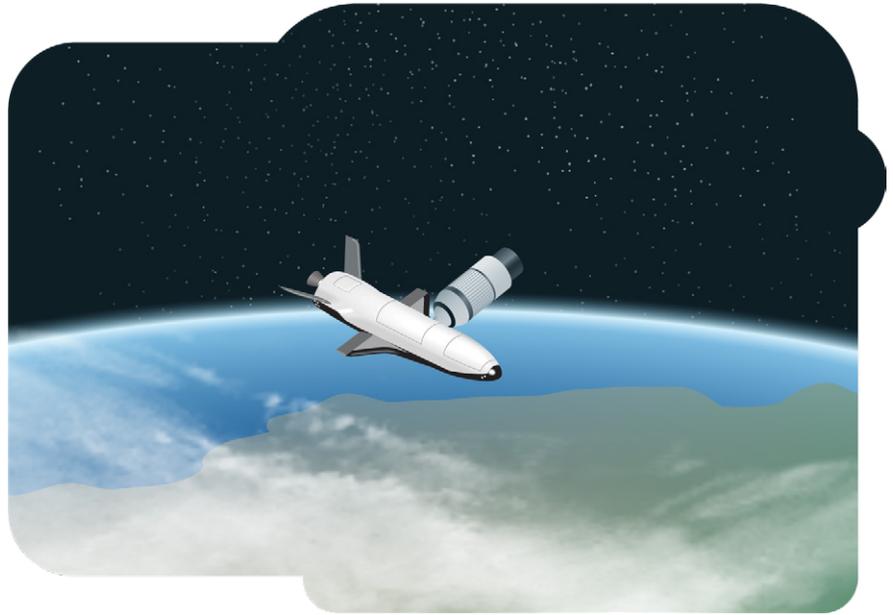


Figure 5: Autonomous on-orbit refueling

astronauts from Earth to the gateway (i.e., LOP-G) near the Moon. Most of the major manufacturing for the first mission is complete. This year, teams will focus on final assembly, integration, and testing while also performing early work for future missions. NASA plans on launching in 2020 the first mission, Exploration Mission-1, to send an Orion spacecraft on the Space Launch System rocket from the modernized spaceport at KSC. This will be an unmanned test flight before sending crew around the Moon on the second mission, Exploration Mission-2 (anticipated by 2023). As we continue to push to explore Mars and beyond, teams of autonomous systems and rovers will be sent first in preparation for human arrival. These rovers will establish the needed human support systems and infrastructure (e.g., water, oxygen, shelter, communications, power) and assemble habitats for living in very hostile environments. This first phase of support, while very dangerous for humans, is ideal for swarms and teams of autonomous systems capable of making real-time decisions and changes as the environment and inputs change.

As we continue to push to explore Mars and beyond, teams of autonomous systems and rovers will be sent first in preparation for human arrival.

Deep Space Exploration

Preparation is underway to return to the Moon and take humans to Mars long-term. The space industry is currently building the systems that will transport

SOFTWARE DEFINED NETWORKING (SDN) FOR AUTONOMOUS SPACE COMMUNICATIONS APPLICATIONS

Traditional Internet Protocol (IP) networks are vertically integrated. Within a networking device, the control plane that decides how to handle network traffic is tightly coupled with the data plane that forwards traffic according to decisions by the control plane within a networking device. To implement a high-level network-wide policy, network operators need to configure each individual network device separately using vendor-specific command line interface (CLI) commands. Implementing automatic reconfiguration and failure response mechanisms in the control plane requires complex human expertise, including expertise and knowledge with a complex variety of protocols tightly intertwined with the associated data plane forwarding mechanism.

A new networking paradigm to remove the limitation of existing network infrastructure, SDN allows network operators to control the components of the networking environment via software rather than the traditional hardware approach. It also decouples the data and the control planes' SDN and decouples the logic that decides traffic routing from the underlying systems. SDN replaces the logic layer with a virtualized controller to enable intelligent networking. In this new architecture, network control becomes programmable, flexible, and centralized to allow the network operators to deliver new services or changes on demand. This SDN innovation provides optimized autonomous approaches for connectivity of services between network nodes, which will ensure more robust and cost effective communications and more flexible space operations and activities for next generation space systems.



Figure 6: Swarms of autonomous communications systems

Automating SDN for Space

The development process for space systems uses a traditional system design process, involving satellite manufacturing and launch activities susceptible to cost and schedule challenges. The underlying terrestrial networks supporting connectivity between the ground segment infrastructure nodes use inflexible, static, scheduled configurations. Network engineers manually configure these at the design phase and subsequently incorporate the configurations into operating procedures—leading to high operator staffing burden, potential for error, and issues in scaling for large networks to support multi-mission constellations at a ground station. SDN controllers can be employed to apply dynamic re-routing and reconfiguration in terrestrial networks to ensure timely data flows from sources at satellite Mission Operations Control Centers (MOCs) to ground stations. Application Programming Interfaces (APIs) can be developed to communicate autonomously with the SDN controllers and the underlying satellite ground segment networking infrastructure (e.g. switches and/or routers)—automatically adjusting data flow paths from MOCs to ground stations to account for ground path loss/degradation or around unexpected weather events. APIs can override previously scheduled, lower-priority data uploads with higher-priority uplinks to a satellite as needed.

SDN principles can also be applied in the context of site diversity to execute Continuity of Business Operations (COOP) procedures. An SDN controller deployed in a high-availability configuration at a network management site, located independently of satellite ground segment nodes, can devise an effective handover decision algorithm between the primary and backup COOP sites. With SDN-enabled switches deployed at the node site, the SDN controller can automatically execute the handover of satellite engineering and support functions between sites. A handover management API communicates with the SDN controller and with the ground stations or terminals to identify active services and data flows, which maintain a satellite data processing pipeline. After handover, the API can alert the antenna crews at the backup site to change their operating frequencies and antenna alignments if needed.

APIs can be developed to communicate autonomously with the SDN controllers and the underlying satellite ground segment networking infrastructure (switches and/or routers)—automatically adjusting data flow paths from MOCs to ground stations to account for ground path loss/degradation or around unexpected weather events. APIs can override previously scheduled, lower-priority data uploads with higher-priority uplinks to a satellite as needed.

With SDN-enabled switches deployed at the node site, the SDN controller can automatically execute the handover of satellite engineering and support functions between sites.

Current and future deep space missions involve communications across expansive distances, thousands to millions of miles apart. This makes normal Internet Protocol (IP) communications very complex and challenging, especially regarding delays and associated communications disruption and data loss for Inter Planetary Networking (IPN).

Delay/Disruption Tolerant Networking (DTN) addresses the technical issues related to lack of continuous network connectivity with a suite of protocols that operate in tandem with traditional IP. The DTN architecture implements a store-and-forward message switching by overlaying a new transmission protocol (referred to as the bundle protocol) on top of the IP protocol. The bundle protocol does not alter the IP protocol data; rather, it encapsulates the application protocol data into datagrams referred to as bundles. Bundles received are forwarded immediately, if possible, but are stored for future transmission if forwarding is not possible at the time. The DTN protocol suite also contains

network management, security, routing, and quality-of-service capabilities to ensure the next hop is available to forward the packet.

Future deep space exploration missions will involve communications needs and data transfer between many nodes involving multiple hops via relay spacecraft or other intermediate nodes. In the United States, NASA provides communications services to support over 100 NASA and non-NASA missions, including Deep Space Network (DSN), Near Earth Network (NEN), and Space Network (SN). These networks consist of a set of distributed ground stations and space relay satellites. Using DTN routing protocols, they distribute a set of expected future inter-node contacts throughout the network. Each node uses this set to make data-forwarding decisions. All parameters in these networks are pre-determined and reactive when responding to dynamic configuration changes. The current system is planned and orchestrated—it does not know and cannot adapt to what will happen in the future. This opens an opportunity for predictive techniques to provide the first step towards autonomous spacecraft operations that allow for any future scenario.

An SDN controller provides the proactive control plane to initiate a trigger for a specific mission's spacecraft requests for communications services, including science data and payload operations. The SDN controller also provides the central intelligence

To make the leap into fully autonomous operations, the SDN controllers need to be used in combination with Artificial Intelligence (AI) with learning abilities to provide communications services and improve network efficiency while minimizing operator burden at mission control centers.



to alter protocol behavior, including forwarding paths across the multiple network layers and to tune parameters for centralized intelligent routing. In this approach, SDN controllers would be placed on the spacecrafts themselves (a future capability) and terrestrially to control ground station nodes and mission control centers. To make the leap into fully autonomous operations, the SDN controller needs to be used in combination with artificial intelligence (AI) with learning abilities in order to provide communications services and improve network efficiency while minimizing operator burden at mission control centers. There have been successful experiments in geographically distributing the

implement on the spacecraft itself, but rather on other space assets or nodes to support the mission. Neural network training, genetic algorithms, and other AI networking operations can happen offline at the capable nodes, such as terrestrial nodes where the SDN controller is deployed, with the results pushed to the lesser capable nodes. Learning within the network will deliver the needed service outside of the constraints of the pre-coded configurations. With the AI-optimized network strategy output, the SDN controller uses the API interface to send instructions to the affected nodes and may even add new services if desired through the API/SDN controller interface.

Neural network training, genetic algorithms, and other AI operations can happen offline at the capable nodes, such as terrestrial nodes where the SDN controller is deployed, with the results pushed to the lesser capable nodes.

neurons of a neural network, where the inter-neuron communications were done by TCP/IP connections. Future autonomous systems may be composed of distributed networks of neural networks that adaptively configure their own SDN communications in response to unforeseen events.

Considering the computational expense of AI techniques, they may not be appropriate to

With a combination of SDN and AI, the time-dynamic spatial relationships between objects can be determined. This will better optimize the node-dynamic positions and orientations, along with the modeled characteristics and pointing of sensors, communications, antenna, and payloads aboard both the spacecraft and ground stations—enabling autonomous operations of future spacecraft missions.



Potential Benefits

Advanced autonomous systems significantly improve the state of the art for space systems and can provide potential benefits, including:

- Improved performance through limited human involvement and reduced infrastructure
- Redundancy, efficiency, and ease of design through routine and duplicative systems and application of standards and protocols to ensure optimization and ease of integration
- Reduction of risk to humans in hostile environments by limiting unnecessary human involvement
- Reduced lifecycle costs in infrastructure and systems to sustain human life and operational optimization through autonomous systems working continuously.



Figure 6: Autonomous, mobile machines at scale are poised to transform human activity in a wide range of physical environments.

Challenges

Evolving to an autonomous paradigm with integrated swarms of autonomous systems working cohesively, include the following challenges:

- Human-in-the-loop culture shift and autonomous science on the edge (e.g., new roles, responsibilities and authorities):
 - Relinquishing authority and control in dangerous or hostile environments
 - Increasing reliance on AI systems rather than humans
 - Making corrective real-time decisions and updates via AI results
 - Allowing life and safety decisions to be made by an AI system
- Collaborative development (e.g., acceptance, risk, roles):
 - Accepting commercial companies sharing in exploration
 - Sharing R&D ownership
 - Maintaining an acceptable risk posture, accepting failures and set-backs
- New processes and methods for engineering and development:
 - Allowing creative and innovative approaches and processes not aligned to “traditional” space-accepted practices
 - Increasing reliance on next generation engineering approaches (e.g., MBSE using AI, AI-to-AI system optimization) and trusting the AI
 - Developing autonomous AI industry standards



Figure 8: Future integrated operations on the surface of Mars

Conclusion

Significant investments and breakthroughs have been made across the entire space industry in the application of autonomy. These systems focus on autonomous systems at scale. Examples of research focuses include mission operations, flight operations, on-orbit assembly, power systems, space structures, habitats, fueling/refueling, communications approaches and systems, space launch and space transportation, sample return, science/in-situ analysis, and many more. The space industry stands at the forefront of the autonomy industry to ensure autonomous systems increase efficiency, reduce cost, and drive technology breakthroughs for the betterment of all.

SOURCES

- 1 2018 OMB budget, https://www.nasa.gov/sites/default/files/atoms/files/fy_2018_budget_estimates.pdf
- 2 NASA Online News Repository. Retrieved from: <https://ti.arc.nasa.gov/tech/asr/news/>
- 3 Autonomous Systems. (NASA Ames Research Center. Office of the Chief Technologist.) Retrieved from https://www.nasa.gov/centers/ames/cct/technology/stp/gamechanging/autonomous_systems.html
- 4 D'Angelo, Gianni & Tipaldi, Massimo & Glielmo, Luigi & Rampone, Salvatore. (2017). Spacecraft autonomy modeled via Markov decision process and associative rule-based machine learning. 324-329. 10.1109/MetroAeroSpace.2017.7999589. https://www.researchgate.net/publication/318891798_Spacecraft_autonomy_modeled_via_Markov_decision_process_and_associative_rule-based_machine_learning
- 5 Nardone, Vittoria & Santone, Antonella & Tipaldi, Massimo & Glielmo, Luigi. (2016). Probabilistic Model Checking applied to Autonomous Spacecraft Reconfiguration. 10.1109/MetroAeroSpace.2016.7573276. https://www.researchgate.net/publication/307560619_Probabilistic_Model_Checking_applied_to_Autonomous_Spacecraft_Reconfiguration

USE CASE: ADVERSARIAL ENVIRONMENTS

Daniel Yim & Thomas Mitchell

The expansion of autonomous technology into hostile environments derives from many years of research. Inventions such as the Whitehead Torpedo in 1868, the Mechanical Mike aircraft autopilot in 1933, Tsukuba Mechanical Engineering in 1977, VaMoRs in 1987, General Atomics MQ-1 Predator in 1995, and the various Defense Advanced Research Projects Agency (DARPA) challenges from 2004 through 2013¹ have helped tackle situations in hostile environments that we see today.

Autonomous systems have demonstrated that they can significantly increase both the operational capabilities and the safety of our modern-day military and civilian sector in the United States.

Autonomous systems have demonstrated that they can significantly increase both the operational capabilities and the safety of our modern-day military and civilian sector in the United States. Depending on the level listed under “Levels of Future Combat Systems”, the autonomous system required for the mission, may or may not have humans as the deciding factor. Autonomy in these circumstances leads to ethical questions such as how the autonomous system would follow the laws

of war established by the Geneva Convention. One such question already arose in 2008, when Ron Arkin wrote a technical report for the U.S. Army Research Office on creating an “ethical governor” for autonomous weapons². The ability for autonomous machines on the battlefield to maintain a set of ethics in warfare is a key aspect of the discussion surrounding this issue.

History of Autonomy in the Military and Defense

The bombardiers of World War II could not hit military targets precisely and avoid civilians if they wanted to; the bombs simply were not accurate as compared to today’s standards³. Spending on research of uninhabited aircraft, or drones, was around \$300 million per year in the 1990s. By 2005, the Department of Defense’s (DoD’s) uninhabited aircraft spending increased six-fold to over \$2 billion per year. In Iraq and Afghanistan, drones provided military personnel the ability to surveil terrorists while not risking human lives. Uninhabited aircraft gave the commanders a low-cost and low-risk way to place eyes in the sky. Due to the success of uninhabited aircraft tactics, the DoD started in early 2005 to develop and publish different roadmaps for the future of unmanned autonomous systems. These roadmaps, with an outlook of about 20 years into the future, described the needs of the DoD—sensors, communications, power, and weapons with autonomous systems—while informing the industry.

Differences in the Research of Autonomous Systems

Autonomy scales defined by the DoD for the Navy, Air Force, and Army demonstrate the different focuses of research in autonomous systems of each military branch based on its mission. The military's research differs from industry because of the complexity of missions that would be assigned to the autonomous vehicle or system. The figures below demonstrate and compare the different areas of the military branches.

In 2011, the roadmap published by the DoD stated, "For unmanned systems to fully realize their potential, they must be able to achieve a highly autonomous state of behavior and be able to interact with their surroundings. This advancement will require an ability to understand and adapt to their environment,

and an ability to collaborate with other autonomous systems"⁴. The DoD realized that producing tens of thousands of drones was not a sufficient strategy. They would need to train the service members to use and operate the drones—requiring a large investment of time and budget. The DoD then released its 2011 roadmap that stated, "autonomy reduces the human workload required to operate systems, enables the optimization of the human role in the system, and allows human decision making to focus on points where it's most needed. These benefits can further result in manpower efficiencies and cost savings as well as greater speed in decision making"⁵.

This DoD robotic roadmap describes four different levels of autonomy: 1) human operated, 2) human delegated, 3) human supervised, and 4) fully autonomous.



Figure 1: A study released on "Autonomous Vehicles in Support of Naval Operations"⁶ identified three types of naval autonomous systems useful to the Department of the Navy: scripted, supervised, and intelligent.

LEVELS OF FUTURE COMBAT SYSTEMS⁷

While the Air Force conducted their research with DARPA, the Army created the Future Combat System Program. The Army wanted to scale different levels of autonomy for its missions in hostile environments. The table below displays the research conducted for the 10 different levels of autonomy for autonomous systems.

Level	Description	Observation	Decision Ability	Capability	Example
1	Remote control	Driving sensors	None	Operator steering commands	Remote control vehicle or car
2	Remote control with vehicle state knowledge	Local pose	Reporting health and state of the vehicle	Remote operator steering commands, using vehicle state knowledge	Teleoperation with operator knowledge of vehicle pose situation awareness
3	External preplanned mission	World model database-basic perception	Autonomous Navigation System (ANS) - commanded steering based on externally planned path	Path following with operator help	Lane assist
4	Knowledge of local and planned path environment	Perception sensor suite	Local plan/re-plan -world model correlation with local perception	Follower with operator help	Remote path following, convoying
5	Hazard avoidance or negotiation	Local perception correlated with world database	Path planning based on hazard	Semiautonomous navigation, operator intervention	Basic open and rolling terrain
6	Object detection, recognition, avoidance or negotiation	Local perception and world model database	Planning and negotiation or complex terrain and or objects	Rolling terrain with obstacle negotiation, limited speed, with some help from operator	Robust, open, terrain with obstacle negotiation
7	Fusion of sensors and data	Local sensor fusion	Robust planning and negotiation of complex terrain, environmental conditions, hazards, and or objects	Complex terrain with obstacle negotiation, limited speed, and operator help	Complex terrain
8	Cooperative operations	Data among cooperative vehicles	Advanced decisions based on data from other vehicles	Robust, complex terrain with full mobility and speed. Autonomous coordinated group	Coordinated autonomous systems in complex terrain
9	Collaborative operations	Fusion between ANS data, surveillance, target acquisition	Collaborative reasoning, planning and execution	Accomplishment of mission objectives with collaborative planning and execution, with operator oversight	Autonomous mission, with individual goals with little supervision
10	Full autonomy	Data from all participating assets	Independence to plan and implement to meet objectives	Collaborative planning and execution, with operator oversight	Autonomous mission without supervision

Table 1: Future Combat Systems, adapted from Office of the Assistant Secretary of the Army.

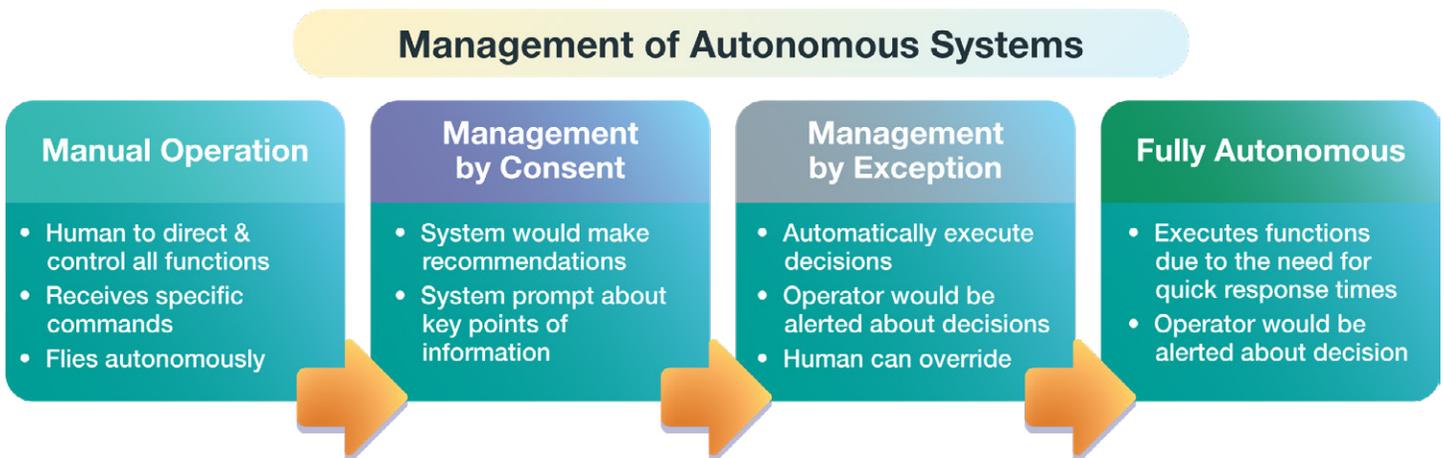


Figure 2: Autonomy scales defined by the joint effort between DARPA and the Air Force – “Autonomy Vehicles in Support of Naval Operations”⁶ published in 2005 established four different levels of autonomy for uninhabited combat aircrafts.

The research conducted by the DoD led the private sector to make great advancements in the field. From the 1990s to today, the United States has chased the next level of autonomy on the horizon—leading to our current capabilities and offering the potential for advancement in emerging autonomous technologies.

Hostile environments

Autonomy at scale and the use of autonomous systems in hostile environments has been part of military and civilian conversations. Significant interest has emerged in how autonomy can be used to combat different hostile scenarios, such as bomb disposal, deep ocean, hazardous materials (HAZMAT), warfighting, and active shooters. Autonomous systems can be used to access environments too dangerous for human exploration—offering greater access to intelligence in hostile environments while reducing risk to life. As the research has progressed, systems that required close human supervision to function now possess capabilities to operate under minimal human supervision.

Before autonomous systems were used to disarm a bomb, a technician wearing a protective suit, with flame and fragmentation resistant material similar to a bulletproof vest, would need to operate in close range of the device. Now that technology has improved, autonomous systems can disarm bombs—keeping the human operators out of harm’s way. These technologies have been applied in hostile situations such as countering terrorists with car bombs.

Autonomous systems can be used to access environments too dangerous for human exploration—offering greater access to intelligence in hostile environments while reducing risk to life.

With jutting reefs, sloping sand dunes, and large rocks, the seafloor remains one of the most challenging environments presenting daunting challenges to researchers. The ocean at a given location may be murky enough to complicate perception, search, and object recognition. Research efforts in this environment have benefited greatly from the application of advanced autonomous systems. Commenting on its work with DARPA, MIT researcher Professor Schmidt said, “We present the acoustic environment and the oceanographic environment very accurately on-board, combining sensor data with modelling and then using that in the decision making. Five years ago, we didn’t have the possibility of putting these kinds of models and predictions on board the vehicles—computer technology wasn’t small or efficient enough”⁸. The success of autonomous systems in the ocean environment has sparked interest in creating autonomous ships that could seek and destroy sea mines. This technology has the potential to improve safety across our seas, opening more navigable routes.

Current and Emerging Technologies

Current and emerging technologies in autonomy have drastically improved—from a preprogrammed robot instructed to conduct assigned tasks to autonomous systems that can detect their surroundings from sensor information and then act to avoid objects. Both industry and the public sector have shifted focus to fielding practical robots capable of re-planning their routes or mission in response to changing circumstances. These technological advancements have yielded real-world applications such as the use of drones to detect leaks along a pipeline, the integration of machine learning of connected devices, and integration of computer vision.

Technology in autonomous systems for bomb removal and detection has advanced over the past two decades. Qinetiq North America created a tactical robot⁹ that has been in use since 2000. This autonomous robot, called TALON, maps hostile environments, disarms bombs, and can provide assistance in situations involving HAZMAT. It can detect radiation, volatile gases, and traces of explosives. The TALON is an example of autonomous systems saving lives in hostile environments.

To answer the challenges posed by deep ocean scenarios, OceanAlpha created the unmanned surface vehicle (USV), a small boat that could be used to swarm a large ship if necessary. Each of the small boats would communicate with each other and then back to a mothership that would control these vessels. USVs would be able to swarm in maritime combat, making them easy to deploy against enemies with reduced risk to military personnel. The USV uses a 5G connection to communicate among vessels and complete missions. This technology could be applied in maritime conflicts, such as combating pirates attacking container ships¹⁰.

Drones are also being deployed increasingly for emergency response in times of disaster¹¹. Finnish tech firm Nokia has been researching the abilities of drones to provide instant 4G mobile network coverage during a disaster. AT&T used a flying “Cell on Wings” (COW) to provide emergency 4G coverage in Puerto Rico after Hurricane Maria struck the island in 2017. Each COW can cover 36 square kilometers and enable critical communication in a disaster scenario. Drones can also enable our fire and police services to see real-time video of a fire or incident, while the fire truck or police cruiser acts as the command center.

In addition to these private sector developments, the U.S. Military has deployed UGVs for the battlefield with the ability to collect information using sensors to create maps of building interiors and the landscape surrounding them. They can also detect objects like people and other vehicles. UGVs can work for extended hours any time of day—providing military personnel with the data, when it is needed to make critical mission decisions. If military personnel encounter explosives, UGVs can disarm and/or remove them. This capability reduces the military's risk of losing soldiers on the battlefield. If damaged, the UGV can even repair itself without intervention (depending on the extent of the damage). They can also transport military personnel between waypoints. UGVs are another example how autonomous systems can keep humans safe and out of harm's way.

As we transition into the Fourth Industrial Revolution, an increase in economic productivity will depend on how well we leverage modern technology such as Internet of Things (IoT) platforms, location detection technology, advanced human-machine interfaces, smart sensors.

Potential Benefits and Impacts

As we transition into the fourth industrial revolution, an increase in economic productivity will depend on how well we leverage modern technology such as Internet of Things (IoT) platforms, location detection technology, advanced human-machine interfaces, and smart sensors. The common theme among these technologies is computing intelligence, which can be applied to benefit most human endeavors. Many repeatable tasks currently carried out by humans will soon be done by machines, and this will be even more so the case for tasks that have historically been dangerous.

Like UGVs, unmanned aerial vehicles (UAVs) currently support the U.S. Navy in various reconnaissance missions in high-altitude and long-endurance environments. UAVs can be catapulted off and land on aircraft carriers for an increased level of capability in combat. While reconnaissance missions aren't considered direct combat operations, using UAVs to collect imagery and data about enemy forces in hostile environments is vital to any mission. Remotely operated unmanned systems also give human operators the ability to achieve the same or even greater results while physically situated in a safer location, not directly exposed to enemy threats.

Resource management is at the heart of operations in hostile environments. The operational productivity of military organizations will likely see improvement as these autonomous systems save lives. Military personnel can focus their expertise and skillsets as operators of autonomous systems or in areas where machines cannot assist.



Challenges

As with all technology solutions, autonomous machines—specifically those used in hostile environments—present challenges. These challenges include the ethical questions, surrender situations, and the limitations of autonomous systems. Interestingly, interoperation of autonomous machines can help address many of the challenges, but increased connectivity also has the potential to create additional challenges in the foreseeable future.

The ethical dilemma is noteworthy: machines do not have empathy towards humans. An autonomous machine would not be able to make any decisions based on its “feelings” about the situation. This presents a significant dilemma if machines overtake humans as decision makers in hostile environments. The machine may not be able to distinguish between an enemy combatant and an innocent civilian. The United Nations passed a resolution concerning the use of autonomous systems for combat in a war zone, ruling that if used, a human must be present to make a lethal decision. This resolution represents an important step toward preventing lethal machines from making lethal decisions.

How an autonomous system reacts in a surrender situation also represents a challenge. The autonomous machine cannot discern whether an enemy is attempting to surrender. A human might not fire upon someone who has surrendered, but an autonomous machine could potentially make a lethal decision based on orders.

The limitations of the sensors and software installed present another challenge with autonomous machines. If a sensor was damaged, it could send false positives or negatives to the central processing unit (CPU) of the machine. With that false information, the decisions the machine makes would become skewed. Software is also susceptible to bugs, system flaws, and malicious attacks that could affect the movements, sensor readings, and decisions of an autonomous vehicle.

Conclusion

Despite the many challenges, autonomy at scale presents a world of benefit and opportunity in operation in hostile environments. On the battlefield or in a disaster zone, autonomous systems with the appropriate guidance and controls, offer significant potential to improve resource efficiency and reduce costs. Perhaps their greatest potential benefit, though, will be the human lives that can be saved.



SOURCES

- 1,4,5**Wired.com. (2016) A Brief History of Autonomous Vehicle Technology.
- 2,3** Scharre, Paul. (2018) Army of None: Autonomous Weapons and the Future of War. New York, NY: W.W. Norton & Company. Retrieved from: <https://books.google.com/books?id=sjMsDwAAQBAJ&pg=PT246&lpg=PT246&dq=autonomous+guided+weapons&source=bl&ots=T-8BYfFE4IP&sig=ACfU3U12Kwbq4342f7cLjJ-SA-9Kuky-74w&hl=en&sa=X&wDnoECAkQAQ#v=onepage&q=autonomous%20guided%20weapons&f=false>
- 8** (2013, August) A New Era for Underwater Drones. Retrieved from: <https://www.naval-technology.com/features/feature-new-era-underwater-drones-unmanned-systems/>
- 9** TALON® Medium-Sized Tactical Robot. Qinetiq North America. Retrieved from: <https://qinetiq-na.com/products/unmanned-systems/talon/>
- 11** Russon, Mary-Ann. (2018, May) Drones to the rescue! <https://www.bbc.com/news/business-43906846>
- 6** Board, Naval Studies. (2005) Autonomous Vehicles in Support of Naval Operations. The National Academies Press, Washington, D.C.
- 7** O'Donell, LTC Warren. (2003) Future Combat Systems Review. Presentation to the committee, Office of the Assistant Secretary of the Army (Acquisition, Logistics, and technology)
- 10** Remote Control Survey Boat. (2018). Retrieved from OceanAlpha: <https://www.oceanalpha.com/>

CHALLENGE: ENSURING INTEROPERABILITY AMONG AUTONOMOUS SYSTEMS

Mile Corrigan

The rapid advancement of the Internet of Things (IoT) connects our world and multiplies our collaborative force. It creates a data-rich environment where the integration of autonomous systems—IoT devices, sensors, artificial intelligence, and robotics—becomes increasingly more complex. Despite significant technological advances and disruptive autonomy innovations, a growing need for interoperability remains within and among autonomous systems, operators, and command and control networks.

“Interoperability has historically been, and continues to be, a major thrust in the integration and operation of unmanned systems ... A robust interoperable foundation provides the very structure that will allow for future advances in warfighting.”

Without interoperability, the technology’s full potential cannot be realized, and the delivery of enhanced value and reduction of operational risk cannot be achieved. In the Unmanned Systems Integrated Roadmap released by the Office of the Secretary of Defense (FY2017–2042), “Interoperability has historically been, and continues to be, a major thrust in the integration



Figure 1: Complexity of interoperability across heterogeneous environments

and operation of unmanned systems ... A robust interoperable foundation provides the very structure that will allow for future advances in warfighting.” Autonomous systems must be able to collaborate with machines and humans to operate effectively in highly complex and contested environments and, ultimately, to derive benefits from their collective synergies. Figure 1 depicts the complexity of interoperability across heterogenous autonomous operations.

Interoperability Challenges

Interoperability applies to both intra-system and inter-system components and represents both physical/logical interconnections and external interactions between multiple systems. The National Institute for Standards and Technology (NIST) provides a working definition for interoperability: “The ability of software or hardware systems or components to operate together successfully with minimal effort by the end user ... Facilitated by common or standard interfaces”¹.

With the ongoing, rapid advancement of systems, interoperability between new and legacy systems will become a major concern for large enterprises in both the government and commercial sectors. In the world of “high-assurance autonomy,” systems must operate functionally while satisfying rigorous safety and security properties to ensure the success of safety critical missions. The challenges facing high-assurance autonomy, interoperability, and integration mainly stem from:

1. Lack of consensus on and adoption of a common set of IoT standards
2. Insufficient verification and validation (V&V) methods
3. Proprietary software and hardware interfaces
4. Lack of trust between systems, operators, and networks

Challenge 1: Lack of consensus and adoption of a common set of IoT standards and protocols

Without standardization, services cannot be exchanged among systems efficiently, impeding our ability to:

- improve connectivity/communication protocols and end-to-end quality control protocols
- apply common processing and programming interfaces and languages
- deliver orchestration and automation platforms for effective operations
- reduce lifecycle costs of hardware and software investments

As multiple stakeholder organizations offer new standards, the need for government and private industry collaboration on the adoption of common standards and protocols grows. Industry most widely uses the Open Systems Interconnection (OSI) model, which decomposes communications across seven functional layers for implementation of interoperable networks (Figure 2)². The IoT-centric model focuses

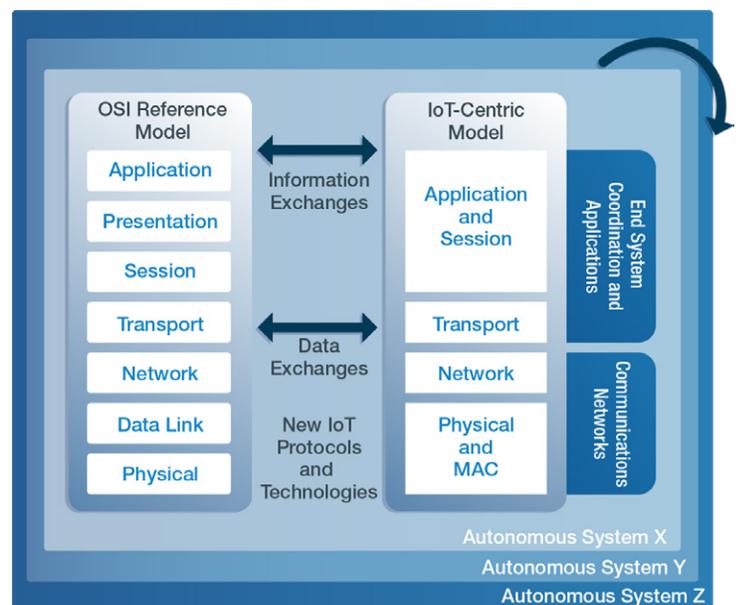


Figure 2: The OSI reference model aligned to IoT-centric communications

“The industries have built standards for the IoT, but it’s been implemented in a fragmented, ad-hoc sort of basis. What we’re going to see is industry adoption of standards, that includes cellular and IoT, and then you’ll see a scaling that will overwhelm many of us.”

on four layers of the OSI reference model stack for communication, data transmission, and end system coordination.

With competing economic incentives, firms have begun implementing their own data exchange and information formatting standards and practices—often overlapping with existing offerings and challenging industry and government efforts to adopt common, universal information sharing standards. Across the IoT space alone, the Institute of Electrical and Electronics Engineers (IEEE) identifies over 80 applicable standards, many focused on specific vertical markets³. While various IoT alliances, consortia, vertical markets, and vendors offer current solutions, new technologies and architectures continue to be developed at a rapid pace—all of which still need to be secured and standardized. Ericsson’s IoT Chief Jeff Traver’s recently said, “The industries have built standards for the IoT, but it’s been implemented in a fragmented, ad-hoc sort of basis. What we’re going to see is industry adoption of standards, that includes cellular and IoT, and then you’ll see a scaling that will overwhelm many of us⁴.”

Consensus and standards adoption for interoperability face hurdles that include:

- **Economic advantages** that can incentivize the development of proprietary systems for increased market share and to achieve vendor lock-in. Vertical initiatives drive the variation in standards to suit an industry’s specific needs, such as data transport protocols to enable information exchange between “communities of things,” versus mobile ad hoc network (MANET) routing communications protocols for Unmanned Aerial Vehicles (UAVs).
- **Competing standards**, such as the wide range of communications standards for low-range and medium-to-low data rate IoT communications (e.g., ZigBee, Bluetooth, IEEE 802.15.4), that can complicate the decision-making and selection process.
- **Lack of reference and architectural models** that adequately address interoperability and standardization objectives and gaps. The Department of Defense (DoD) has called for open architecture structures designed to facilitate modification to evolving requirements and technology advancements.
- **Fear of obsolescence** that can delay adoption as new technologies, together with evolving and competing standards, are developed and launched with increasing speed. Adopting the wrong standard could result in a system that becomes obsolete (think VHS vs. Betamax).

Many large organizations like Cisco, Intel, IBM, and GE are joining IoT standards bodies (e.g., the Industrial Internet Consortium, IPSO Alliance, the Open Connectivity Foundation) to stay ahead

of the adoption curve. Technology companies like Google and Amazon, however, have taken a different approach to gain a competitive advantage, by developing their own technologies and interoperability solutions⁵.

Challenge 2: Insufficient V&V methods

With increased autonomy comes unpredictability. As autonomous systems execute both coordinated and uncoordinated actions in new and unforeseen ways, they are failing differently than could be predicted with a human in the loop—driving the need for robust software V&V methods. Software V&V, a technical discipline of systems engineering, employs a rigorous methodology for evaluating the correctness and quality of a software product through the software lifecycle. Validation confirms that the software meets the user’s needs: “Are we building the right system?” Verification confirms that the system is well engineered: “Are we building the system right?”⁶ Today, V&V activities account for nearly 25% of development costs—a figure anticipated to increase disproportionately with other development costs as the unpredictability of autonomous systems grows⁷.

Validation confirms that the software meets the user's needs: “Are we building the right system?” Verification confirms that the system is well engineered: “Are we building the system right?”

Traditional approaches, designed for testing manned systems, will not be enough to meet the key V&V challenges presented by highly adaptive and non-deterministic systems:

- **Dynamic and Unpredictable Environments** to which context-aware autonomous systems react to dynamically drive the need for a much larger decision space that can produce unanticipated events and failures. Plans and deliberations are intertwined with actions that can be both proactive and reactive. With adaptive systems, behavior across all working conditions is not known at design time, making fault tolerance methods difficult to implement.
- **Emergent Behavior**, dependent on the acquired knowledge of each system, prevents the inclusion of fault avoidance methods in the formal verification process. System interactions can often produce unintended consequences. Testing adaptive systems that learn, adapt, self-diagnose, and apply intelligence in decision-making is often highly labor intensive, making it costly and time consuming to comprehensively observe the full range of simulated fault scenarios for a given mission.
- **Lack of Test Repeatability**, a necessary condition for establishing and maintaining reliable test methods. The complexity of the operating environment coupled with adaptive software characteristics can produce different results even when a system is supplied with the same set of inputs. Fault removal through extensive testing and debugging is difficult to achieve since an autonomous system’s behavior changes and learns over time.
- **Lack of Reliable and Certifiable V&V Methods** complicates efforts to prevent errors in autonomous system development. Test and evaluation (T&E) requirements imply formal methods for system assurance based on past

failure conditions of similar systems, not readily available for newly developed autonomous systems⁸. Recognizing this challenge, Former Chief Scientist of the U.S. Air Force, Werner Dahm asserts, “Developing certifiable V&V methods for highly adaptive autonomous systems is one of the major challenges facing the entire field of control science, and one that may require the larger part of a decade or more to develop a fundamental understanding of the underlying theoretical principles and various ways that these could be applied”⁹.

These challenges all demonstrate the high cost, complexity, and difficulty of achieving V&V results by applying classical software testing methods such as fault avoidance, fault removal, and fault tolerance testing to autonomous systems.

Challenge 3: Proprietary software and hardware interfaces

With such a wide array of commercial software and hardware products deployed across large enterprises today, proprietary interfaces present a major barrier to system integration and interoperability. Organizations holding a large portfolio of commercial-off-the-shelf (COTS) systems—like the DoD does—cannot maintain pace with changing conditions of unmanned systems due to their proprietary nature and lack of data rights and need for more timely software updates.

Unlike open source software and interfaces built using common data standards and protocols,

proprietary software and hardware interfaces raise major issues, including:

- **Vendor Lock**, which deepens reliance on the vendor for upgrades, enhancements, maintenance, and support versus open source software and interfaces built using common data standards and protocols.
- **Innovation Lag** that slows the pace of innovation and evolution of autonomous system capabilities, as system owners must negotiate with the vendor for required software changes.
- **Integration Stall** caused by closed interfaces and proprietary software that inhibit integration and data-sharing among systems, typically exacerbated by lack of user access to source code or the ability to make modifications or fixes.

Closed software is not without its potential advantages such as extensive technical support for maintenance and, oftentimes, higher product stability due to a smaller feature set. These advantages, though, may not outweigh the drawbacks for large enterprises that desire to keep pace with reliable, interoperable, high-assurance autonomy. In acknowledgment of this tradeoff, the Defense Science Board Task Force has recommended that each U.S. military service initiate at least one open software design project to decouple autonomy from the vehicle—deploying proven technology to reduce manpower, increase capability, and adapt more swiftly to future missions¹⁰.

“Developing certifiable V&V methods for highly adaptive autonomous systems is one of the major challenges facing the entire field of control science, and one that may require the larger part of a decade or more to develop a fundamental understanding of the underlying theoretical principles and various ways that these could be applied”.

Former Chief Scientist of the U.S. Air Force, Werner Dahm.

Challenge 4: Lack of trust between systems, operators, and networks

The notion of trust, implying a human psychological trait characterizing assurance or certainty in human-to-machine (H2M) interactions, reflects a key challenge in the context of implementing and operating autonomous systems. The research paper “The Trust V – Building and Measuring Trust in Autonomous Systems”¹¹, defines two types of trust for a user to accept an autonomous system:

- **System trust**, or human confidence that the system behaves as intended. Achieving this trust requires a high level of assurance that the system satisfies its requirements, (i.e. the traditional V&V challenges).
- **Operational trust**, or human confidence that the system helps the user perform the assigned tasks. Achieving this trust requires a high level of assurance that the scenarios for which the system was designed are useful. A lack of human confidence in the system or its operations impedes high-assurance autonomy, integration, and interoperability.

People tend to respond to technology in human and social ways. Unclear or uncertain decision-making of an autonomous system negatively influences a person’s level of reliance in complex situations. Whether trust means sending a loved one on the road in a self-driving car or sending machines or drones into battle with humans, prioritizing the establishment of H2M trust in the design process can ultimately create better interactions for the end-user and reduce the chance of misuse.

The Air Force Research Laboratory (AFRL)¹² identifies five human-machine teaming technology challenges that must be addressed to establish trust between systems, operators, and networks to maximize performance in complex and contested environments:

- **Human State Sensing and Assessment** to measure and assess the human’s state (e.g., physiological, performance, behavioral).
- **Human-Machine Interaction** to enable humans and machines to communicate and share information.
- **Task and Cognitive Modeling** to allocate workload and decision-making balance.
- **Human and Machine Learning** to adapt, learn, and extend mutual training between humans and machines.
- **Data Fusion and Understanding** to integrate human and machine data (e.g., context, time, format) for a shared world model.

Any human operator must be able to trust their interactions with an autonomous system to achieve greater levels of interoperability and mission assurance between other systems, operators, and networks. Shared understanding is key to overcoming the H2M trust barrier.

Current and Evolving Approaches Standards and Open Architectures for Interoperability (Challenge 1 & 3)

IoT network protocols and standards continue to evolve as quickly as new industries and use cases emerge across business and government. Enterprises must choose the right network topology for the use case and consider the market in which the capability will be deployed. Most IoT-enabled autonomous systems comprise a multi-tier architecture spanning devices, gateways, data systems, and services as depicted in Figure 3.

With no universal model to describe the collection of protocols, standards or technologies, developers face the challenge of selecting the right subset

of protocols, drawing from competing standards and minimizing risk of obsolescence. Further complicating these decisions, large enterprises need to reduce lifecycle costs, ensure vendor conformance to open standards, and guarantee the commonality of components across autonomous platforms. Figure 4 depicts the Open Standards reference model for IoT communications, highlighting the ever-evolving network and data protocols available in the marketplace today¹³.

Integration across different layers to perform data and information exchanges requires alignment of appropriate protocols as defined by the different Standards Developing Organizations (SDOs) (i.e., the IEEE, IETF, ITU, etc.). The SDOs, alliances, and forums develop IoT protocols based on the physical interfaces already established in the industry. For example, the Wi-Fi, VX2 protocols used in the PAN and LAN networks are found in IEEE 802.11, while protocols like ZigBee, Thread, Wireless HART, etc. are built over IEEE 802.15.

To exchange messages and data across multiple sensors and systems, the application layer supports multiple protocols which in most cases use the publish/subscribe models. An IoT architect must carefully select the right protocols across the different layers appropriate for the type of network to ensure interoperability as well as scalability and performance of the solution.

To meet interoperability challenges head on as new protocols and standards emerge, the DoD launched an initiative to develop an Unmanned Ground Vehicle (UGV) Interoperability Profile (IOP)¹⁴ for the acquisition of future programs, the upgrade of fielded systems, and the evaluation of commercial products. The IOP created by the U.S. Army Robotic Systems Project Office, approved for public release through the National Advanced Mobility Consortium (NAMC), specifies interoperability across several levels:

- **OCU/UxV(s):** Between Operator Control Units (OCU) and one or more Unmanned Vehicles (UxV(s)).
- **Intra-OCU:** Between and among OCU hardware and software elements.
- **Intra-UxV:** Between and among UxV subsystems, payloads, and platforms.
- **OCU/UxV/C2:** Between OCUs, UxVs and external C2 systems to exchange command and control, battlespace and audio/video information.

The IOP, designed to support a wide range of missions, vehicle classes, controller classes, payload classes, architectures, and interactions with external systems, presents a strong case towards realizing “open architectures, reusable, interchangeable components and common, publicly defined interfaces between individual subsystems,” said Heidi Shyu, Former Assistant Secretary of the

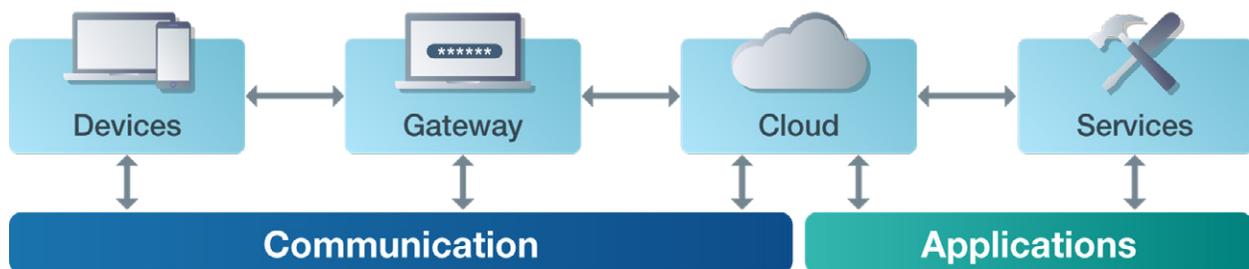


Figure 3: IoT Stack Simplified

OPEN STANDARDS REFERENCE MODEL FOR IOT COMMUNICATION PROTOCOLS

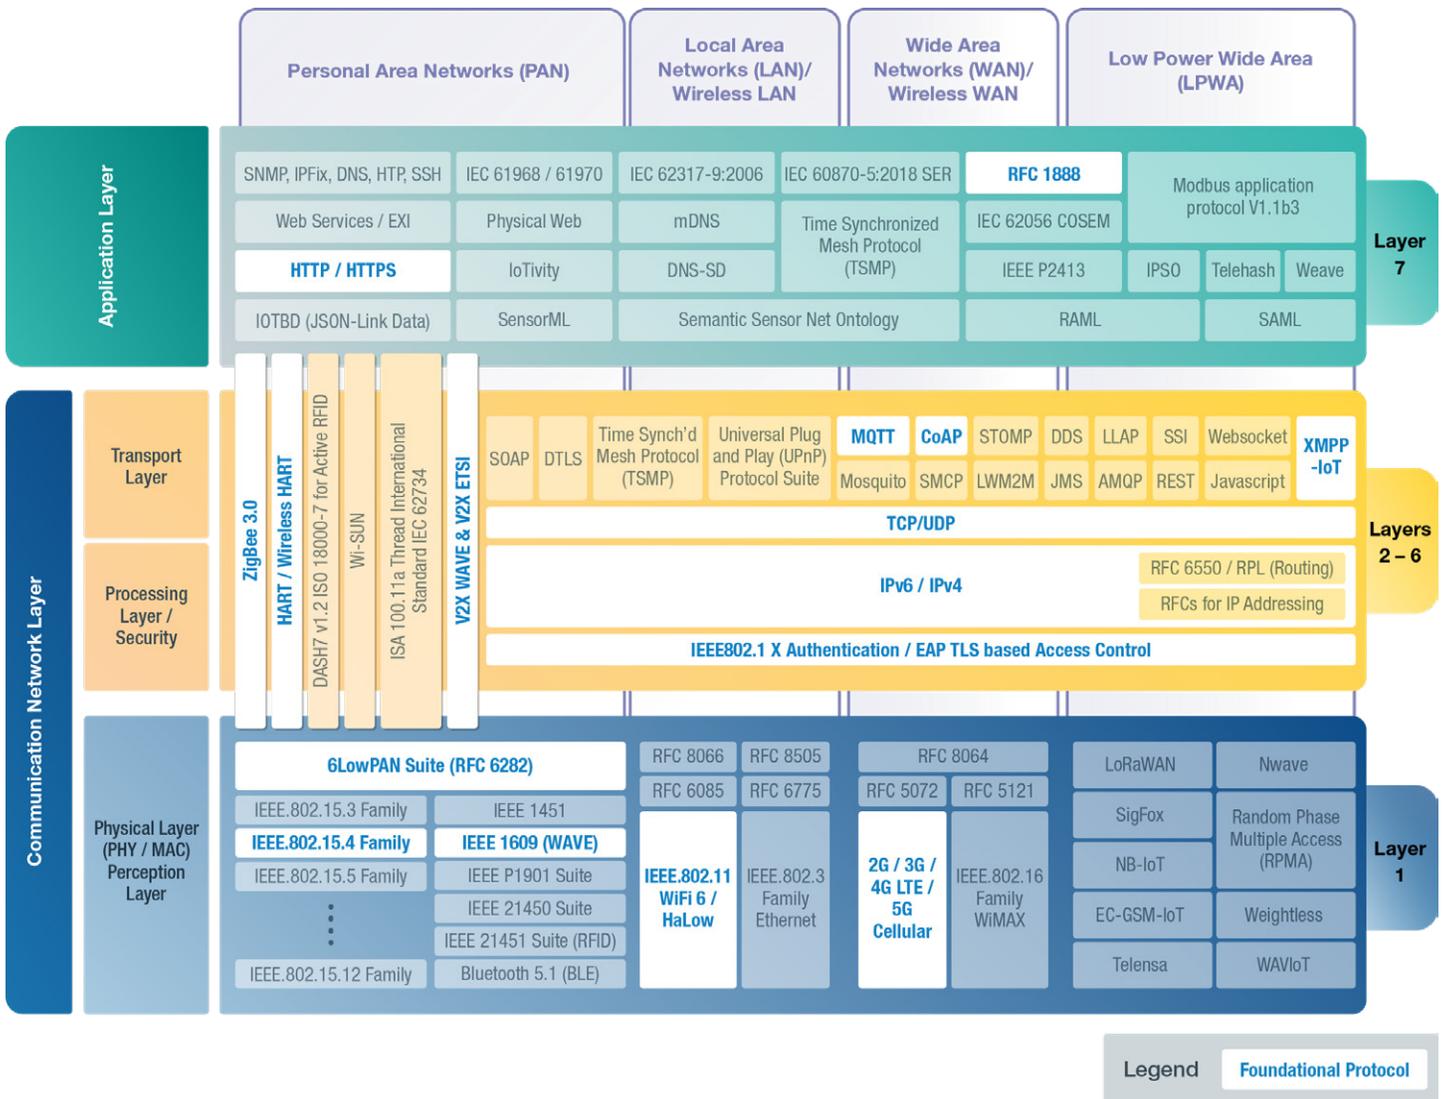


Figure 4: Adapted from the Open Standards Reference Model - Graphic from: David E. Culler (<https://www.cs.berkeley.edu/~culler/>).

Army for Acquisition, Logistics and Technology¹⁵. Specifying messaging and transport protocols to support the scope of the IOP will accelerate adoption of standards.

Looking ahead, organizations should dictate the use of open architectures, open standards, and open source software to reduce the reliance on closed and proprietary technologies over time.

Emerging V&V Approaches for Autonomous Systems (Challenge 2)

Several approaches have emerged to address the complexity of verifying and validating autonomous systems. These include model-based approaches, evolutionary test algorithms, simulation-based methods, and virtualization tools that often combine several advanced V&V techniques already in place. Intelligent Systems Division (ISD) researchers at NASA Ames Research Center are applying many of these advanced V&V techniques such as static analysis, model checking and compositional verification to gain trust in model-based autonomous software systems¹⁶.

Classical development processes and methods work well when requirements are easily understood;

“Open architectures, reusable, interchangeable components and common, publicly defined interfaces between individual subsystems, said Heidi Shyu, Former Assistant Secretary of the Army for Acquisition, Logistics and Technology.”

however, traditional approaches provide limited insight into issues discovered during the test phases of a highly autonomous and open-ended system. This results in an inability to test for all known conditions. Table 1 highlights current test approaches and use cases attempting to address testing challenges for autonomous systems operating in a safety-critical environment.

While most solutions do not extend end-to-end for an entire autonomous system, a few testing and V&V trends have gained wider acceptance across the various approaches:

- **Modeling and Simulation:** Applying model-based testing methods can find failures and reduce defects early in the process, as models can be used to simulate or communicate intended behavior, helping to build trust and acceptance of the system. Simulation-based approaches such as Adaptive Stress Testing (AST) can find failure paths more quickly, using sequential decision processes that can be further optimized with reinforcement learning¹⁷.
- **Virtual Testing:** Applying virtual methods to a representative model of the intended operational environment can reduce costs compared to live testing and poses less risk since the virtual hardware and test environment can be used repeatedly for test experiments.
- **Transparent Engineering:** Systematically engineering systems that provide transparency into weaknesses and defects can handle emergent nonrepeatable behavior. By building transparency into the design and operation of the system, engineers can identify failures early in the design process, improve safety, and provide accountability¹⁸. Engineers can't

TABLE 1 – TESTING APPROACHES FOR AUTONOMOUS SYSTEMS

	Testing and V&V Approaches	Description / Use Cases	Advantages	Work to be Done
Model-Based Test Approach	<ul style="list-style-type: none"> • Model-based Test Approach • Run-time monitoring • Predictive Analysis 	<ul style="list-style-type: none"> • Model-based testing automatically generates test cases from models • Autonomous Satellite System (AGATA project) employed model-based specifications to produce the RT-Java code of the AGATA onboard software¹⁹ • Autonomic Service-Component Ensembles (ASCENS) combined a model-based approach with run-time monitoring and predictive analysis²⁰ 	<ul style="list-style-type: none"> ✓ Reduces complex systems to logical components, enabling abstraction and componentization ✓ Enables incremental development to initiate software validation earlier in the process ✓ Performs model debugging and automatic code generation ✓ Defines adaptation, awareness, and emergence properties through mathematical models ✓ Generates build ensembles that are more adaptive, reliable, and usable 	<ul style="list-style-type: none"> • Integration of monitoring techniques with runtime verification to bridge testing and formal verification
Evolutionary Test Approach	<ul style="list-style-type: none"> • Evolutionary Algorithms • Agent-based Software Engineering • Software Abstraction 	<ul style="list-style-type: none"> • Approach to testing autonomous agents that uses evolutionary optimization to generate demanding test cases • Soft goals are transformed into evaluation criteria and tests are generated with evolutionary algorithms suited to multi-objective optimization 	<ul style="list-style-type: none"> ✓ Evaluates autonomous agents as a means of building confidence in behavior and greater agent dependability since quality functions are derived from requirements 	<ul style="list-style-type: none"> • Development of design and programming constructs for agent interactions that work towards shared system goals



TABLE 1 – TESTING APPROACHES FOR AUTONOMOUS SYSTEMS, CONTINUED

	Testing and V&V Approaches	Description / Use Cases	Advantages	Work to be Done
Adaptive Stress Testing	<ul style="list-style-type: none"> • Simulation-based test approach • Adaptive Stress Testing (AST) • Markov decision process (MDP) • Reinforcement Learning 	<ul style="list-style-type: none"> • Approach to stress testing that finds most-likely failure scenarios by formulating a sequential decision-process (e.g. MDP) and then uses deep reinforcement learning to search for most likely failure paths¹⁷ 	<ul style="list-style-type: none"> ✓ Deep Reinforcement Learning produces more likely failure scenarios compared to other methods (e.g. Monte Carlo tree search) ✓ AST finds failure scenarios efficiently 	<ul style="list-style-type: none"> • Incorporation of more realistic models with tighter constraints on the events of interest
Virtual Test Approach	<ul style="list-style-type: none"> • 3-D/4-D Modeling & Simulation • Early testing of embedded software via software-in-the-loop virtual integration • High-resolution physics based simulation of robotics platforms 	<ul style="list-style-type: none"> • Virtual testing environment for autonomous aerial vehicles using simulation-based in-the-loop validation of UAV software • Virtual environments for autonomous mobile robot systems using the Mobility Open Architecture Simulation and Tools (MOAST) 	<ul style="list-style-type: none"> ✓ Allows for testing without putting the hardware or environment at risk ✓ Allows for the evaluation of using alternative hardware components prior to implementation ✓ Provides a baseline simulation system capable of modeling autonomous systems with the ability to conduct repeatable test experiments 	<ul style="list-style-type: none"> • Development of automated approaches to systematically explore the state-space of the planning algorithm • Development of more realistic simulations • Not everything can be tested virtually to address complexity and noise of the real world • Development of robust algorithms to address rational decision making in an autonomous system

necessarily ensure every corner case is properly handled, but this modern engineering practice can help make every corner case visible.

- **Advanced V&V Techniques.** Static analysis techniques assess code without execution, reducing the potential for dangerous operations that have to be checked by other methods. **Model-checking** efficiently checks that a model of a system satisfies all requirements, providing a robust way to catch system-level errors (e.g. concurrency, deadlocks, etc.). **Compositional verification** – often referred to as a “divide and conquer” approach decomposes properties of a system into properties of its components. Components are model checked separately, guaranteeing the verification of the entire system if each component is verified¹⁶.

Confidence that an autonomous system will operate as intended is critical to its deployment. Developers will progress to more advanced features when they

can establish high confidence in lower subsystems, in contrast to low confidence systems where defects are hidden among several layers of the system. The ability to test and verify autonomous systems will continue to be critical to operational deployment, mission effectiveness, and human safety.

Enhancing the Human-Machine Team (Challenge 4)

Numerous evolutions in human-machine teaming are improving communications, comprehension, and control in H2M interactions. Human-robot interaction (HRI), a relatively new field of study, seeks to address the challenge of human-machine trust. HRI encompasses multidisciplinary contributions from human-computer interaction, artificial intelligence, robotics, human factors, operations research, and social sciences. HRI focuses on the understanding, design, and evaluation of robotic systems for use by or with humans—such as fully autonomous machines (classified as robots).

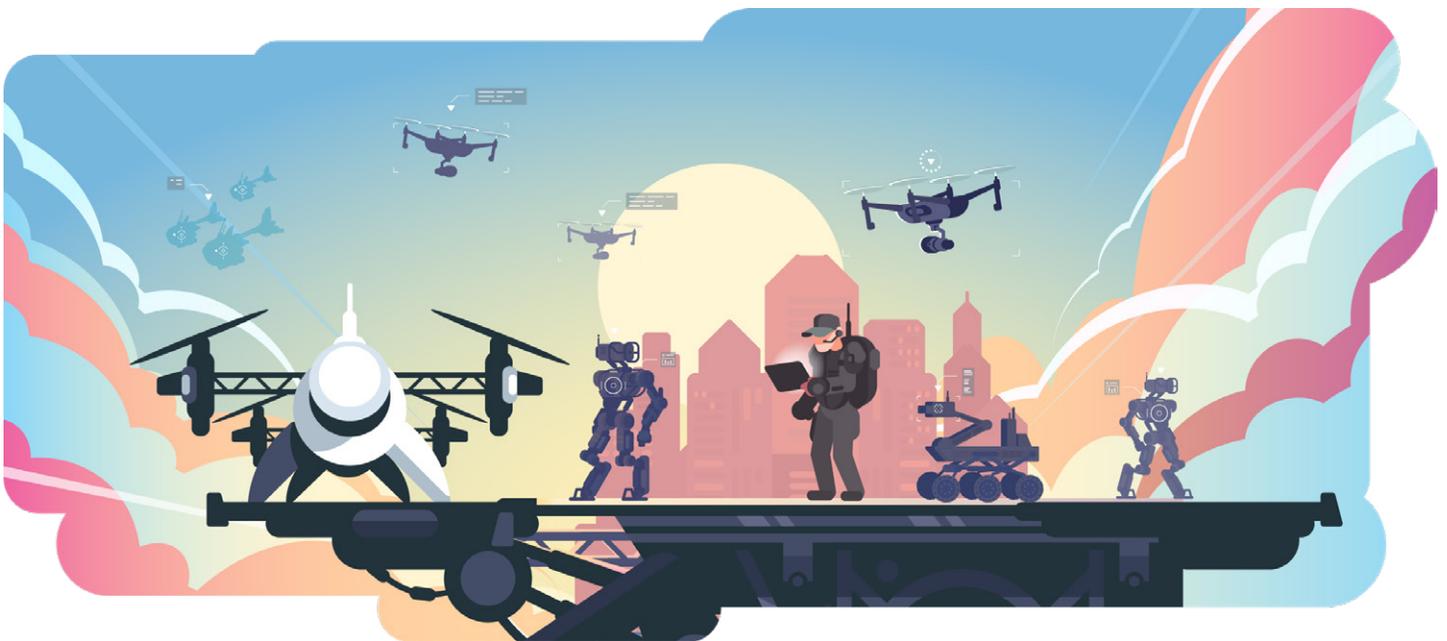


Figure 5: The future of human-machine teaming

TABLE 2 – HUMAN SUPERVISORY CONTROL METRIC CLASSES AND SUBCLASSES

Metric Class	Description	Subclass Examples
Mission Effectiveness	Effectiveness measures relating to the whole human-automation system	Mission performance parameters
Autonomous Platform Behavior Efficiency	Parameters relating to the efficiency of the autonomous platform	Usability, adequacy, autonomy, learnability, errors, accuracy, reliability, neglect time
Human Behavior Efficiency	Parameters relating to how humans sequence and prioritize multiple tasks such as monitoring autonomous platform health and status, identifying critical exogenous events, and communicating with others as needed	Information processing efficiency (e.g., decision-making), attention allocation efficiency (e.g., scan patterns, prioritization)
Human Behavior Precursors	The underlying cognitive processes that lead to specific operator behavior, as compared with the human behavior metric class that captures explicit behavior	Cognitive precursors (e.g., situation awareness, mental workload, emotional state) Physiological precursors (e.g., physical comfort, fatigue)
Collaboration Metrics	Team-level metrics to measure the degree to which the humans and automation are aware of one another and can adjust their behavior accordingly	
	Human-automation collaboration	Trust, mental models
	Automation-automation collaboration	Quality and efficiency of collaboration (e.g., speed of data sharing, quality of system response to unexpected events, etc.)
	Human-human collaboration	Coordination efficiency, team mental model

Source: *Evaluation criteria for human-automation performance metrics. In Performance Evaluation and Benchmarking of Intelligent Systems*²¹

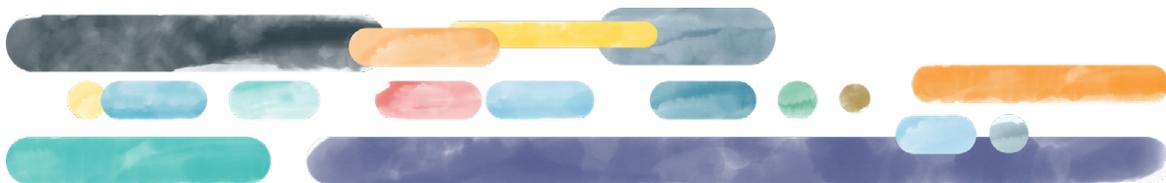


The scope of HRI addresses H2M communications, shared relationship models between humans and machines to achieve autonomy, enhancements to the human-machine team, and how to capture and express interactions within a given application domain, characterized by the:

- Level and behavior of autonomy
- Nature of information exchange
- Structure of human-robot team
- Training of people and robots
- Design and shaping of tasks for human-robot collaborations

To assess holistic systems, we must establish and validate metrics for evaluation and testing of H2M interactions. To gain an understanding of H2M interactions and how they can be influenced or enhanced to achieve an outcome, the interactions must be measured for improved operations. MIT researchers have defined five metric classes for human-machine interactions, described in Table 2²¹.

While we may be decades away from solving all of our human-machine interaction challenges for high-assurance autonomy, system owners can start by exploiting data between H2M and Machine to Machine (M2M) interactions to derive new insights and drive continuous improvements. Improved data collection and data sharing for monitoring, management, and optimization should be conducted early and continuously—especially as data volume and quality increases over time. Valuable information can be extracted from metadata, improving the self-awareness and flexibility of systems. Autonomous system data strategies should adapt situationally with an understanding of unique mission goals and constraints.



Pathway to Improved Interoperability

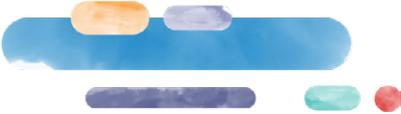
Interoperability will remain a key challenge for autonomous systems—now and into the future. To exploit the collective intelligence and capabilities of integrated autonomous systems, enterprises must set the foundation for interoperability by establishing an architectural basis for the development of future systems. With so many protocols available in the marketplace today, industry should focus on the most commonly used ones—built on open standards to simplify and accelerate interoperability. Standardizing hardware and software interfaces, requiring the use of open standards, protocols, and architectures, and securing data rights will enable long term sustainability, modernization, and reduced dependency on proprietary system owners—ultimately driving down lifecycle costs.

Autonomous system V&V will require continued advancements in T&E so that run-time architectures can constrain systems to a set of allowable, predictable, and recoverable behaviors, integrated early in the development process. Testing methods will need to integrate development and operational testing and employ new ways to test the whole system whether through virtual testing, transparent engineering, model-based engineering and testing approaches, or new ones yet to be developed. As research in this field evolves and emerging test approaches are applied more rigorously across autonomous systems, organizations will be able to make informed decisions on which test methods will yield the best outcomes.

Finally, human and technological capabilities must be integrated into a well-functioning system to optimize the human-machine team. Constraints should be shared with all parts of a given system so that the autonomous system serves as a creative partner that complements capabilities. To drive the integration and adoption of autonomous systems, trust barriers will need to be overcome. With trust established, the observability, controllability, and partnering between humans and machines improves significantly—enabling enterprises to reap the benefits of high-assurance autonomy.

SOURCES

- 1 (National Institute of Standards and Technology 2008, 28)
- 2 Irons-McClean, R., Sabella, A, Yannuzzi, M., IOT and Security Standards and Best Practices, 2019.
- 3 Internet of Things. IEEE Standards Association. Retrieved from: <https://standards.ieee.org/initiatives/iot/stds.html>
- 4 Daws, Ryan. (2019, March) Ericsson IoT chief on AI, 5G, and connecting ‘things’ instead of ‘a thing’. IoT News. Retrieved from: <https://www.iotechnews.com/news/2019/mar/01/ericsson-iot-ai-5g-connecting-things/>
- 5 Hughes, Terry. (2016, April) Will industry muscle win in the IoT standards war? [Blog]. Retrieved from: <https://internetofthingsagenda.techtarget.com/blog/IoT-Agenda/Will-industry-muscle-win-in-the-IoT-standards-war>
- 6 IEEE Standard for Software Verification and Validation, 1998.
- 7 Helle, P., Schamai, W., Strobel C. (2016). Testing of Autonomous Systems – Challenges and Current State-of-the-Art
- 8 Technology Investment Strategy 2015-2018, DoD R&E Autonomy COI TEVV Working Group, May 2015.
- 9 Dahm, W. J. (2010). Technology Horizons a Vision for Air Force Science & Technology During 2010-2030. Office of the US Air Force Chief Scientist.
- 10 Department of Defense Defense Science Board Task Force Report: The Role of Autonomy in DoD Systems, July 2012.
- 11 Zwillinger, D., Palmer, G., & Selwyn, A. (2014). The Trust V - Building and measuring trust in autonomous systems.
- 12 Overhold, Jim, Ph.D. and Kearns, Kris (2014, April). Air Force Research Laboratory Autonomy Science & Technology Strategy [PowerPoint Slides]. Retrieved from: https://defenseinnovationmarketplace.dtic.mil/wp-content/uploads/2018/02/88ABW-2014-1504_140312v2_-_Autonomy_Strategy_Part_2.pdf
- 13 Adapted from the Open Standards Reference Model - Graphic from: David E. Culler (<https://www.cs.berkeley.edu/~culler/>) - The Internet of Every Thing - steps toward sustainability CWSN Keynote, Sept. 26, 2011
- 14 Robotics and Autonomous Systems - Ground (RAS-G) Interoperability Profile (IOP) (Version 2.0 ed.). Warren, MI, USA: US Army Project Manager, Force Projection (PM FP). 2016.
- 15 Serbu, Jared. (2014, August) Army turns to open architecture to plot its future in robotics. Federal News Network. Retrieved from: <https://federalnewsnetwork.com/defense/2014/08/army-turns-to-open-architecture-to-plot-its-future-in-robotics/>
- 16 Brat, G., Denney, D., Giannakopoulou, J F, Jonsson, A. (2005). Verification of Autonomous Systems for Space Applications.
- 17 Koren, M., Alsaif, S., Lee, R., Kochenderfer, M. (2019). Adaptive Stress Testing for Autonomous Vehicles.
- 18 Corcoran, William. Transparent Engineering. American Scientist. Retrieved from: <https://www.americanscientist.org/article/transparent-engineering>
- 19 Pouly, J., & Jouanneau, S. (2012). Model-based specification of the flight software of an autonomous satellite. Embedded Real Time Software Systems (ERTS 2012).
- 20 Hölzl, M., Wirsing, M., Klarl, A., Koch, N., Reiter, S., Tribastone, M., et al. (2011). Engineering Ensembles: A White Paper of the ASCENS Project.
- 21 Donmez, B., P. E. Pina and M. L. Cummings. 2009. Evaluation criteria for human- automation performance metrics. In Performance Evaluation and Benchmarking of Intelligent Systems, R. Madhavan, E. Tunstel, and E. Messina, eds. New York: Springer Science+Business Media. doi:10.1007/978-1-4419-0492-8.



CHALLENGE: THE CYBERSECURITY ENVIRONMENT IN AUTONOMY AT SCALE

Sam Leestma

Autonomous machines are set for exponential growth—increasing both their footprint in new industries and their utilization in industries already leveraging autonomy. The future will see greater use of autonomous machines at scale for transportation, distribution of goods, military operations, and space exploration (Figure 1). As the use of autonomous technology grows, cybersecurity breaches of systems managing autonomous machine fleets and nodes will become a greater threat to individuals and to the security and functions of industries and nations as a whole. As space travel, mass transportation, food and goods distribution, and individual transportation become reliant on autonomous machines, the information about critical operations, logistics, and personal information will be trusted to systems operating mostly independent of human interaction. This dependence presents several key challenges in security design integration, the verification of security functionality, and the protection of operating systems and the metadata that they will rely on and produce.

Additional considerations must be made especially when examining the increased utilization of autonomous machines at scale for distribution of goods and food and transportation of people. Classified as critical infrastructure, these systems that can affect large portions of the population and require additional protections. Many protections surrounding current critical infrastructure systems (e.g., power grids, water facilities,

telecommunications) rely on closed systems or security by obscurity. The far more open community of autonomous technology will have to be vigilant in detecting and mitigating threats. Monitoring and tracking the emergence of new threats and attack vectors will be critical to maintaining the viability of increased reliance on autonomy at scale (AaS).

Information about critical operations, logistics, and personal information will be trusted to systems operating mostly independent of human interaction.

Autonomous machines are also increasingly being used in military operations. Military services have leveraged drones for combat operations; the security and integrity of those autonomous machines will be important to preserve human safety. Special consideration will have to be paid for communications and protocols managing these machines to ensure decisions on hostile targets are accurate and effective against the threat. As the use of armed drones in active combat situations increases, the design must ensure the integrity of the protocols remain intact and function correctly.

Data Types and Critical Operations

To discuss protection mechanisms for application to autonomous systems and to define and validate levels of trust needed, we will need to explore the nature of the systems and the functions they support. The use cases affecting populations and the stability of nations, functions that support the movement and logistics of food and goods and the transportation of people present some of the most critical functions at scale. Autonomous nodes responsible for food and product distribution as well as mass transport will rely on systems that provide logistics and management of autonomous services (Figure 2). These centralized management systems will hold the information needed for their operation and their byproduct or metadata will draw the attention of unauthorized threat sources. As AaS becomes more integrated into the military and space industries, the security of nations will start to rely on the security and trust in operations of those autonomous systems and individual autonomous nodes.

The movement of food presents challenges because, if compromised, it can have national security and economic prosperity implications. The stability of a society relies heavily on the availability and integrity of the food supply, and the high assurance of the operations and security of systems that manage food distribution protects food supply safety. Those systems will need high availability to ensure populations have basic necessities. In examining the movement of goods, the availability of autonomous systems and nodes will be critical to organizations that leverage those technologies. A company's viability can tie directly to its ability to reliably move products to consumers. The movement of goods also presents a secondary concern. Companies that produce and distribute goods will focus on

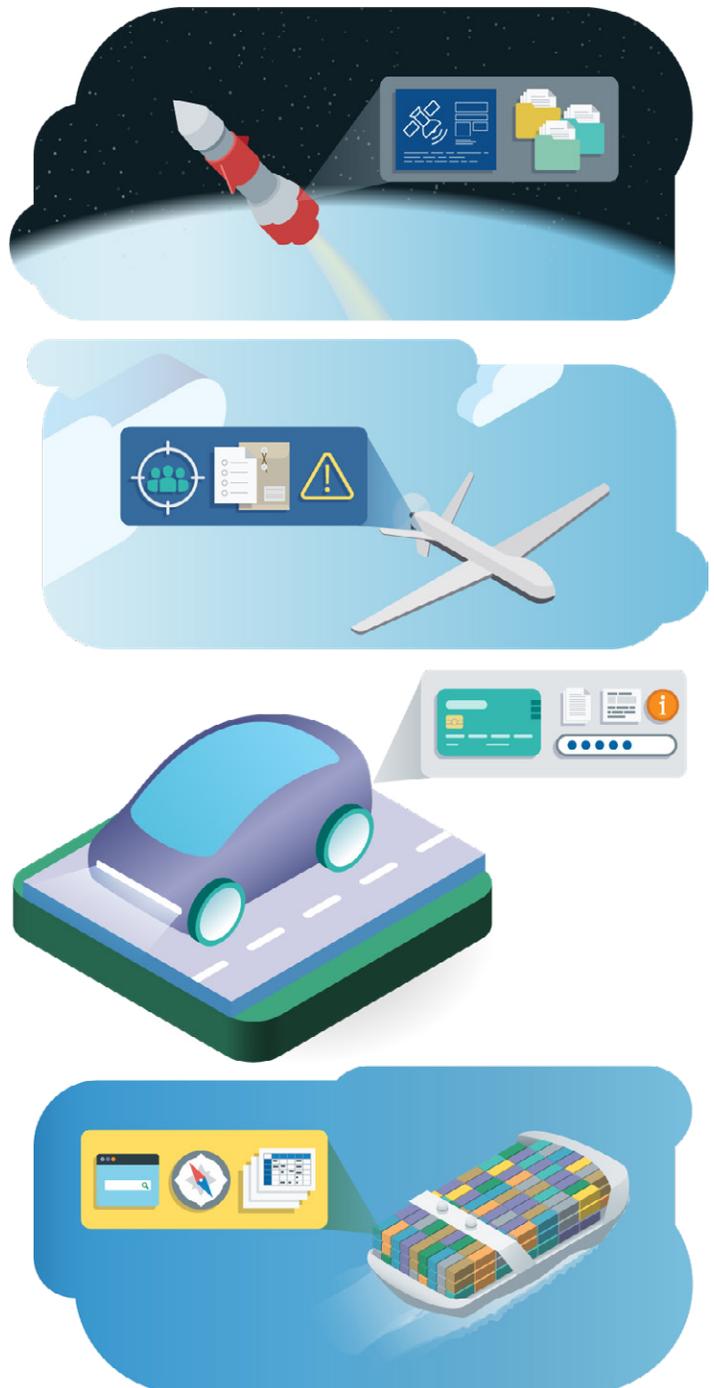


Figure 1: Current and future implementations of autonomous machines at scale

maintaining the confidentiality of their logistical information. The movement, volume, and capacity of their distribution of products is integral to their strategic planning and corporate health. Loss of this material data can weaken a company. The challenge of metadata protection will be an important concern for companies as autonomous systems increasingly process and store centralized logistical information.

Autonomous nodes that are responsible for food and product distribution as well as mass transport will rely on systems that provide logistics and management of autonomous services.

Space missions, by their nature, are operated by nation sponsors. They support coordinated international missions, transport commercial and national-interest satellites, carry equipment and goods for scientific experiments, and transport personnel. The large capital and human investment in space programs are critical to a nation and its security. Maintaining the integrity and availability of space missions involves all autonomous space machines reaching their intended locations and orbits. Conversely, the loss of that integrity or availability can result in the compromise of those autonomous space machines. Confidentiality of space missions (i.e., their flight contents, strategic operations, and goals for the sponsoring nation) will require protection to ensure the security of the nation states.

The military use case for autonomous machines includes military operations and transport of personnel and services. The military uses the critical logistical information produced and used by these systems to operate without the enemy's knowledge. Troop and equipment locations and movements are highly classified by the nature of their use. Additionally, the metadata used and produced by drones can contain military or intelligence information for target tracking and engagement, and the information about missions supported by drones can contain operations, their status, as well as logistical information about enemy combatants.

Threat Pairing for Critical Services

Each use case for AaS presents a unique threat profile based on the functionality and data present. Within each of the critical services, autonomous systems and their controllers use, produce, and store data and metadata. Each of these services and data elements will have multiple threat sources interested in exploiting it. Exercising vulnerabilities can produce losses in the confidentiality, availability, and integrity—elements of “trustworthiness” in most current discussions of autonomous systems. Myriad threat sources would find value in the theft of data, denial of services, and manipulation of the functionality or data used and processed on autonomous systems.

Simple threat sources such as script kiddies, whose motivation is accomplishment, notoriety, or simple mayhem, could apply to all use cases. The threat profiles for AaS for transportation of goods and services, military logistics, and space travel, however, are more targeted.



Figure 2: Data and metadata types present in autonomous systems

Each of these groups has unique motivations and each use case of autonomous machines and systems offers opportunities to advance their goals. Some implementations of AaS will support services that will mirror the importance of our critical infrastructure. Food transportation and goods distribution and logistics handled by autonomous machines will put these critical services into systems independent of human interaction. Food, medication, and critical equipment—fundamental to the wellbeing of citizens—will reside with systems that have multiple points of failure and access points that do not require human intervention to gain entry. Autonomous machine management systems will offer a single point or limited points of exploitation that can disrupt large scale critical services affecting large portions of a population. Hostile nations or economic criminals could have interest in denial of services or the integrity of the systems contents or the logistical operations. The use of autonomous machines for these services will mitigate the easiest point of attack from human coordinated physical attacks, to single nodes, to multi-point, to remote attacks that can affect fleets of nodes that are providing critical services.

As the operation and utilization of autonomous machines in more hostile and military environments will continue to increase, the threat profiles and the impact of compromises that those threat sources wish to achieve become more critical to the nations and individuals relying on those autonomous machines and systems. Successful attacks can result in loss of life and impact availability of critical supplies needed in dangerous situations. As the implementation of autonomous machines increases in the military community, the threat profiles can shift based on political and national or group lines. The shift to information warfare will result in more complex attacks on systems, which could expand to encompass autonomous machines and systems providing critical services. Physical attacks on single nodes will be replaced by cyber-attacks that can affect large swaths of military targets.

Physical attacks on single nodes will be replaced by cyber-attacks that can affect large swaths of military targets.

Known Cybersecurity Concerns

The infrastructure and systems in the AaS landscape face many current threats.

Many autonomous machines use Global Positioning Systems (GPS), a low-energy, unencrypted service susceptible to denial of service—intentional and unintentional—and vulnerable to snooping. As a core service that all autonomous systems rely on, GPS creates a single point of vulnerability with far-reaching consequences throughout the autonomous system.



Figure 3: Common autonomous machine hardware attacks

Sensor attacks, a form of exploit tricking sensors into giving false data, can affect the performance of autonomous machines and lead to widespread or isolated accidents resulting in traffic and route flow failures. Even isolated sensor attacks could result in the denial of service for critical operations. Sensor attacks on autonomous machines used for military operations can result in false target identification—and unintended human casualties. While these attacks generally result from physical tampering, sensors can also be manipulated by underlying hardware attacks (Figure 3).

Hardware attacks allow hackers, nation states, or any individual or group with access to the manufacturing or retrofitting processes to inject vulnerabilities into the hardware or firmware utilized by autonomous machines. This can result in complete compromise of the autonomous node or give the attacker the ability to affect services associated with the particular compromised piece of hardware.

Attacks on firmware updates can effect similar levels of compromise to those of hardware attacks. Firmware attacks can be accomplished through remote updates or from physical compromise of the onboard diagnostics (OBD) hardware ports present in an autonomous machine.

Remote Access such as Bluetooth and built-in Wi-Fi provides attackers with a vector to gain access to autonomous nodes. The system architecture of nodes and management systems vary greatly; each access point attack vector, if exploited, provides the ability to pivot and compromise a wide variance of a node that differs depending on the node (Figure 4). Once an attacker gains access, they may—depending on the architecture or security implementations present—install various viruses or malware, tracking software, or a wide array of unauthorized code. The impacts of these compromises can range from gaining complete control of a node, to monitoring the node, to degradation to the integrity of the node's operation.

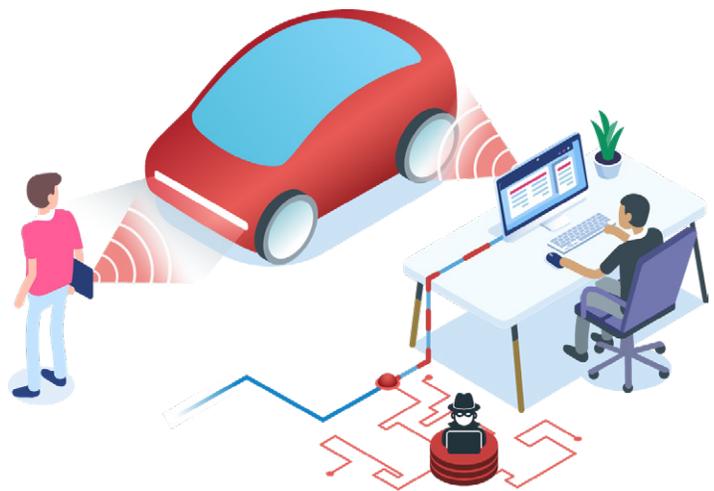


Figure 4: Attack vectors applicable to autonomous systems and machines

CONNECTED VEHICLES: DETECTING AND VALIDATING AUTONOMOUS VEHICLE MISBEHAVIOR

By: Cory Krause

A significant factor in the success of the U.S. Department of Transportation's (USDOT's) Connected Vehicles program is the transmission of Basic Safety Messages (BSMs) between vehicles and infrastructure. BSMs are important, over-the-air messages that contain necessary vehicle data such as position, speed, acceleration and brake status. These BSMs can gather details about the traveling vehicle in real time and transmit this information to other vehicles and infrastructure devices in the area.

Such an open network poses certain risks regarding the reliability of the information contained within the BSMs. One possible threat comes from faulty sensors or components within the vehicle that could measure data erroneously and result in the transmission of inaccurate and/or unrealistic BSMs. Another possible threat comes from malicious third parties potentially hacking into the system and feeding misleading data while posing as a nearby trusted vehicle. Either scenario will result in BSMs that do not reflect the vehicle's actual behavior and can be considered misbehavior within the system. The accuracy of this data is an absolute necessity that carries the weight of a potential loss of life due to spoofed or inaccurate messages that misrepresent a vehicle's location and can cause a tragic collision.

Noblis is leading a USDOT project that creates an installable piece of code on connected and autonomous vehicles (CAVs) that detects BSMs in the area and determines their accuracy. This work includes several tasks:

- Development of algorithms for detecting and qualifying misbehavior
- Creation of code, installable to the vehicle on-board units, that detects and flags incoming vehicle misbehavior

- **Testing of the code through the installation and trial of misbehavior in connected and autonomous vehicles (CAVs)**
- **Development of a formatting and reporting mechanism for credentialed hardware, which allows for credentials to be revoked from misbehaving devices**

Noblis has developed software that reads a large amount of Basic Safety Messages (BSMs) over the air in a 300 to 500-meter area. Using these heterogeneous data points, we attempt to determine several pieces of important information—most notably, are the vehicles likely where they say they are, and is it physically possible that they are doing what they say they are doing. Many times, the spoofed messages resulting from hacking into the system are not realistic—whether because the hacker doesn't have an in-depth understanding of the system they are attacking or because they simply want to create havoc. Upon checking all the fields for realistic values (e.g., speed, acceleration, brake status), the software flags anything outside the realm of possibility. This could be a speed over 100 miles per hour or some wheels braking while accelerating. The software also handles the more complex task of determining location accuracy. By comparing latitude and longitude of all vehicles in a physical space and the surrounding area, we can use a low false-negative approach of removing and flagging the vehicles that could not be within the physical space.

The software then sends these flags to an authority that checks them and, if accurate, adds the devices to a credential revocation list that removes the ability for these devices to send messages in the future. In this way, whether it be through malevolency or malfunction, the devices can no longer impact the travel of those vehicles around them. The devices will then be checked by vendors and local transportation agencies to determine the problem.

Noblis is leading this effort of detecting and validating autonomous vehicle misbehavior through another autonomous system. Such forward-thinking approaches will likely be included in CAVs for years to come.

Securing Autonomy at Scale

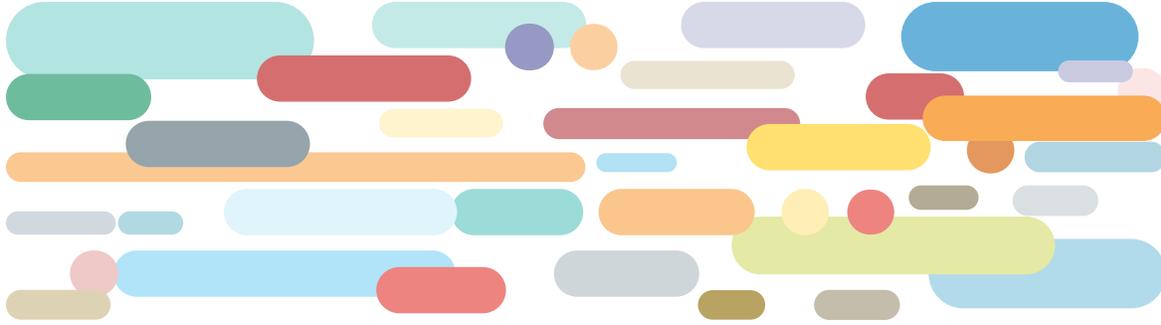
Technical countermeasures will need to be developed and enhanced, and autonomous nodes and management systems will have to be retrofitted to include these security and redundancy implementations. To be resilient to unique threats and a wide range of large-scale attacks, the movement toward Autonomy at Scale (AaS) must prioritize security as a focal point in the progress of autonomous innovation. The use of security mechanisms and security function isolation needs to be employed with autonomous machines as fleets grow and become a critical part of the daily lives of individuals, governments, military, and commercial industries.

While current security mechanisms can and should be leveraged in the development and integration of these systems, new countermeasures to mitigate the unique threats to autonomous systems must also be developed. For example, the security challenges

present in GPS systems will have to be analyzed, mitigated, and integrated to support AaS. The use of currently established security practices such as encryption, key infrastructure, hardware protections, and intrusion detection systems will need to be incorporated to ensure that not only autonomous nodes, but also control systems are protected. Supply chains for hardware and parts will need to be tightly monitored and managed. The sourcing of materials, chips, and technologies will have to be analyzed, and trusted partners must be established. Robust intrusion and anomaly detection will need to be in place and fine-tuned. Systems will need to detect allowable variance to position and operating status and have defined response conditions. These systems will have to be sufficiently redundant to ensure detection of small issues so that isolated anomalous conditions do not create widespread outages. Even small variations in coordinates and GPS position can create catastrophic consequences.

LOOKING FORWARD

As AaS creates a boon in productivity and consistency in transportation and logistics and as the community and businesses increase capacity and functionality in autonomous machine vehicles, new threats and opportunities to exploit critical functions in society will arise. Novel and increased thought needs to be given to the security considerations of these systems. The rapid progress and increased reliance on autonomous systems will quickly outpace the cybersecurity needs unless we pay special attention to current and upcoming challenges in securing these systems. Standards in architecture, security measures and compliance frameworks must be developed and integrated into the lifecycle of these autonomous nodes and systems. Throughout the history of information technology, functionality progress and innovation often take the lead. Leaving security as an afterthought has caused the loss of information to foreign adversaries, financial loss to cyber criminals, and loss of private information. In the case of AaS, the cost of lagging behind in security will manifest itself in dangerous breaches and exploits that can impact national security and safety.



ABOUT NOBLIS

Noblis is a nonprofit science, technology, and strategy organization that brings the best of scientific thought, management, and engineering expertise with a reputation for independence and objectivity. We work with a wide range of government and industry clients in the areas of national security, intelligence, transportation, healthcare, environmental sustainability, and enterprise engineering. Together with our wholly owned subsidiary, Noblis ESI, we solve difficult problems of national significance and support our clients' most critical missions.

NOBLIS.ORG

noblis[®]

For the best of reasons

 703.610.2000  answers@noblis.org  [@NoblisInc](https://twitter.com/NoblisInc)

© 2019 Noblis, Inc. All rights reserved. Proprietary to Noblis.