2018 NDIA GROUND VEHICLE SYSTEMS ENGINEERING AND TECHNOLOGY SYMPOSIUM
AUTONOMOUS GROUND SYSTEMS TECHNICAL SESSION
AUGUST 7-9, 2018 – NOVI, MICHIGAN

ADVANCED PERFORMANCE COMPUTING POWER FOR LARGE DATA SET PROCESSING

Karl Rosenberger
Sara Blackmer
Steve Wesoloski
RAVE Computer
Sterling Heights, Michigan

John Brabbs
United States Army TARDEC
Detroit Arsenal, Michigan

ABSTRACT

The real-world testing of robotic and autonomous vehicles faces many challenges including: safety; feasibility; effectiveness; expense; and timeliness. The development of high performance computing has created innumerable opportunities for effectively and efficiently processing large data sets. These data sets can range from modeling and simulation scenarios to the vast amounts of complex data being gathered by unmanned vehicles. In all cases, the data needs to be stored, managed, and processed to have usable information to drive smart decision making. Leveraging high performance computing to more efficiently, effectively, and economically conduct robotic and autonomous vehicle testing in a virtual environment is a logical step. Consequently, TARDEC has developed a real-time modeling and simulation capability to test and evaluate autonomy solutions while RAVE has designed and developed a specialized high performance computing system for TARDEC to support this capability.

INTRODUCTION

The advent of high performance computing (HPC) has changed the way that large data sets are handled, processed, and utilized. For the military, activities using HPC, such as modeling and simulation (M&S) analyses, robotic and autonomous systems (RAS) operation, and intelligence, surveillance, and reconnaissance (ISR) data processing, continue to grow larger in scale and complexity. While M&S analyses and ISR data processing can be conducted on variable timelines, RAS operation requires real-time data processing and decision making. Throughout normal operations, RAS accumulate a large amount of sensor data and this needs to be processed quickly to allow for RAS decision making. Due to the large number of variables that can affect the RAS, all of the data must be treated with equal importance. This combination of factors can make mobility testing of RAS both difficult and time consuming, which is especially true of real-world testing.

For mobility testing, the use of simulations is a valuable tool as the full range of scenarios can be explored faster and more economically than what can be accomplished through real-world testing. Use of HPC can enable these analyses to be conducted faster which allows researchers to run more simulations in a given set of time.

Additionally, M&S analyses are being tasked to consider a wider variety of scenarios. Increasing fidelity in M&S analyses require more detailed models accounting for more factors with each simulation requiring more processing time.

To explore and quantify the value of virtual RAS testing, the United States Army (Army) Tank Automotive Research, Development, and Engineering Center (TARDEC) has developed a real-time M&S capability to simplify and streamline the testing process. RAVE Computer (RAVE), an industry expert in HPC, worked with TARDEC to design and develop a custom HPC system to support this capability.

THE CHALLENGES OF REAL-WORLD RAS TESTING

Before being fielded, Army RAS will need to be rigorously tested under representative operational conditions by soldiers in the unique environments where the RAS will operate. These environments will range from urban settings with improved roads to rural settings with unimproved roads, trails, and off-road and mountainous terrain. Testing RAS provides unique challenges with the additional automation and decision making being conducted.

Real-world Testing Example

As an example, the challenges of conducting RAS Leader Follower (LF) real-world testing include: safety; the testing facility; the testing equipment; and the cost and time required. All mobility testing of unproven technology is inherently unsafe and test and evaluation (T&E) teams and facilities take precautions to reduce the risk to manageable levels. These types of safety risks are compounded in RAS as there can be both hardware and software issues, and in some cases, the vehicle crash may even be the result of the way the software is designed. [1] Part of these safety challenges lie in the inability of the T&E team to control every variable. Because of these factors, conducting developmental testing on public roads involves unacceptable safety risks.

Consequently, testing facilities are used and they provide their own challenges. These facilities are often limited in how many vehicles they can accommodate at any given time and they provide a limited selection of terrain types, weather conditions, traffic conditions, and vehicle event scenarios. Multiple facilities must be used to gain a reasonably inclusive data set and it will still be extremely difficult to explore all possible scenarios with any consistency. As many of these scenarios involve obstacles, there is potential of RAS damage every time the scenario is conducted, and a catastrophic failure can stop all testing, especially if multiple vehicles are affected.

The type and amount of equipment required for the testing can also pose challenges. To begin with, each testing vehicle must be outfitted with the RAS sensor and processing package and a full instrumentation suite. It is difficult to process data in a moving vehicle so specialized equipment may be required. In a RAS LF scenario, multiple vehicles are required and so each must be outfitted with the full complement of testing equipment, not to mention the test drivers required in each vehicle for safety. Spare vehicles will be required to account for any mechanical problems that arise during testing. To increase the rate of development, multiple different test convoys may be used and the required equipment and personnel scales rapidly.

This challenge also directly affects the cost of the testing. Integrating just a light detection and ranging (LIDAR) system onto a single vehicle can cost $85,000, and this does not include the cost of other sensors and the data processing equipment. [2] In addition to the high cost of the RAS equipment, and even if all of the testing instrumentation is available from the testing facility to rent or use for free, there is the ongoing cost of the fuel, maintenance, and testing personnel. On top of the financial cost, there is the cost in time, both in conducting the testing with multiple vehicles required at all times to the data processing and analysis time required to get sufficient data for results.

Ensuring Safe and Reliable Testing

TARDEC and the Army Test and Evaluation Command (ATEC) have been collaborating on developing testing that meets all of the necessary parameters. As a result of this collaboration, ATEC is interested in leveraging TARDEC M&S expertise with unmanned ground vehicles (UGVs) and large unmanned ground systems (LUGS) for their Autonomous Systems Test Capability (ASTC) being developed for the T&E community.

The defense industry is interested in fielding RAS as soon as possible but face the challenge of ensuring that these systems are safe and reliable. This challenge goes beyond the logistical challenges discussed above. “As automated vehicles and their technology become more advanced and technically sophisticated, evaluation procedures that can measure the safety and reliability of these new driverless cars must develop far beyond existing safety tests.” [3] The challenge is how to run enough real-world miles using these systems to encounter all the edge cases, or dangerous driving situations, which can help to measure safety and reliability.

The University of Michigan’s MCity published a white paper in 2017 discussing “an accelerated evaluation process that eliminates the many miles of uneventful driving activity to filter out only the potentially dangerous driving situations where an automated vehicle needs to respond.” [3] Use of M&S on HPC systems to support the Army T&E process will help augment the real-world testing by targeting the edge cases. As the edge cases will face additional requirements in terrain types and situations, use of M&S to explore the system performance in these settings can be cost effective.

While real-world testing of UGVs/LUGS is very important to the Army, it is also expensive, and the Army needs to maximize the effectiveness of it. Being able to drive the number of miles required in the real world to get an accurate assessment is not feasible because of cost, time, and the difficulty in reliably replicating all of the edge cases that need to be tested. The MCity white paper points out that “to get an accurate assessment in field tests, such cars would have to be driven millions or even billions of miles to arrive at an acceptable level of certainty – a time-consuming process that would cost tens of millions of dollars.” [3] For example, assuming an average speed of 50 miles per hour, a vehicle would need to be driven 16 hours per day, every day, for 3.42 years to accumulate one million miles.

VIRTUAL-WORLD RAS TESTING: A VITAL COMPLEMENT

The real-world testing of RAS faces many challenges including the difficulty in conducting it efficiently and effectively, as well as the fact that it is both expensive and time consuming. Many of these challenges can be overcome by using a virtual world to conduct the testing. Inherently, this is made possible because all of the RAS decision making is already happening in a computer. Working together, TARDEC and RAVE have assembled M&S capabilities for testing RAS on a custom HPC system to explore the effectiveness of the virtual-world testing.

USING VIRTUAL-WORLD RAS TESTING TO SOLVE REAL-WORLD TESTING PROBLEMS

TARDEC has developed a real-time virtual simulation environment for UGV/LUGS hardware and software experimentation through in-the-loop simulations. These simulations are used in the development, integration, testing, and evaluation of RAS to provide controlled, repeatable, and representative stimulation of those systems to explore their performance in complex and dynamic environments that are too dangerous, costly, and time-consuming to evaluate via traditional methods. Some M&S benefits include: significantly more miles can be driven in a simulation versus a real vehicle fleet; the simulation miles can focus on the edge cases as opposed to the uneventful driving; and the test-fix-test process can be much faster than being in the field.

Waymo, the former Google self-driving car project, is using autonomy M&S to “unhitch from the limits of real life and create thousands of variations of any single scenario, and then run a digital car through all of them.” [4] This is similar to the approach that TARDEC is developing and it will allow TARDEC to iterate the same scenario by just changing a few parameters.

TARDEC is enhancing the current real-time M&S capability to evaluate and test autonomy solution edge cases with the aim of using this capability to augment, and not replace, real-world testing using the controllability, repeatability, and scalability that are only possible within a virtual-world environment.

In support of this effort the Product Manager (PdM) for Applique and Large Unmanned Ground Systems (ALUGS) has funded the development of the Continuous Autonomous Simulation and Test Lab Environment (CASTLE) which extends TARDEC’s real-time virtual simulation environment for RAS hardware and software experimentation through in-the-loop simulations to include: automated testing; mission control; verification; and validation. The first test case for CASTLE is an RAS LF system concept.

The LF system concept is a manned lead vehicle followed by follower LUGS. The LF system concept uses an Applique architecture for the current manned vehicle fleet to automate the vehicle via three different kits: a by-wire kit, which permits control of all primary vehicle controls and provides feedback and control of platform components; an autonomy kit, which is composed of perception sensors and computers for processing sensor data for implementing high level autonomy behaviors; and a Warfighter Machine Interface (WMI)/Radio Network Interface (RNI) kit, which allows the communication and control of the convoy. These kit-based hardware and software systems will be used to retrofit existing vehicles with autonomous capabilities.

The real-time virtual simulation environment provides the ability to assess the behavior of the LUGS autonomy software operating in different conditions and scenarios. The real-time M&S capability provides the ability to gain confidence on the autonomous system's ability to perform successfully prior to real-world testing.

Testing within CASTLE reduces risk by finding problems before going to real-world testing. The government can then direct the contractor on which problems are a priority to fix, thus reducing program risk and schedule.

Virtual-world Testing Example

When considering the same example from before, many of the challenges of conducting RAS LF real-world testing become the benefits of virtual-world testing, including: safety; the testing facility; the testing equipment; and the cost and time required.

First and foremost, all of the safety concerns are eliminated as no real-world vehicles are required and no testing personnel are placed in potentially dangerous scenarios.

Specialized test facilities with different terrain and weather conditions can be minimized as the virtual world can mimic any terrain, environment, or scenario. More importantly, the scenarios are completely controllable in the virtual world allowing each scenario to be repeated infinitely with no unplanned changes. This allows the testing to concentrate on the edge cases without any of the uneventful driving required by real-world testing. Each simulation can also be stopped, restarted, or reset quickly allowing more tests to be conducted in a set time frame. Without the need for safety drivers, refueling, or vehicle maintenance, testing sessions can be conducted almost continuously.

In addition, the equipment requirements of virtual-world testing are less burdensome as a single HPC workstation can represent a single vehicle, at a lower cost than required by just the LIDAR system. Additional vehicles can be added to the virtual world by adding workstations, allowing the overall simulation to scale easily.

Overall virtual-world testing cost and timing is also less onerous than real-world testing. Each workstation has a low, fixed cost with minimal energy requirements and almost no maintenance, and it allows for an unlimited number of test drives. Concurrent test series can be conducted with the addition of more workstations. At a per server cost of approximately $5000, a LF system with one leader vehicle and three follower vehicles can be assembled for approximately $20,000.

This is a bargain when compared against a per real-world vehicle cost many times the total cost of the entire virtual-world testing setup. Some experts put this cost between $250,000 and $300,000 per vehicle, or $1,000,000 for a four vehicle system. [2] For this price, a T&E team can assemble 50 virtual-world testing setups. With a virtual-world setup of this scale, one million miles of LF system concept testing, at 16 hours per day at 50 miles per hour, can be conducted in just 25 days. A table comparing real-world testing and virtual-world testing is included in Appendix A.

RAVE HPC Workstation Solution

Running the real-time virtual simulation environment with the LF autonomy software requires an HPC workstation to support the real-time interaction between the virtual simulation environment and the autonomy software. Historically, data processing with computers has used central processing units (CPUs) to process data sequentially. The processing time was then determined by how long the CPU took on each piece of data multiplied by the total amount of data. Faster CPUs helped reduce the processing time but faced limits due to power consumption and heat generation limits.

HPC sought to address these issues by dividing the data set and then harnessing multiple processors to each process a distinct part of the data set. After processing, the results of the different parts are recombined into the whole data set. HPC enabled larger data sets to be processed faster with less power consumption and heat generation.

Once the M&S capability needs were identified, RAVE worked with TARDEC to design the optimal HPC workstations for this task. There were several inter-related aspects that made the system design challenging, including the combination of several computer hardware requirements that are often not found in a commercial-off-the-shelf (COTS) engineering workstation. Among the competing requirements were: a number of networked computers; space constraints of the installation location; the potential for 24/7 processing; high fidelity visual output from the simulation; and optimized system responsiveness with minimum latency.

As no COTS solution was available to meet this combination of factors, a custom-built solution was required. This allowed all of the requirements to be met using the best breed of components for each aspect of the system design. The specifics of the resulting system are outlined in Table 1.

Motherboard	Intel X99 Chipset
Processor	Intel Xeon multi-core, multi-thread (14 cores, 28 threads)
RAM	64 GB
GPU	GeForce GTX 1080Ti

Table 1: RAVE HPC Workstation Specifications

When designing the system, RAVE first identified the computer chassis to fit within the selected server racks. The chassis was compact at 7 inches tall by 19 inches wide by 21.5 inches deep, and supports sufficient airflow to cool the components. Additional server-grade fans were added to reduce the system operating temperature and increase the system reliability.

Next, RAVE identified the system motherboard. Motherboard selection is one of the most critical challenges for such a unique system since the motherboard is the key to supporting all of the technical requirements. The motherboard selected works well with the chassis as it has parallel CPU, memory, and Peripheral Component Interconnect Express (PCIe) slots to facilitate optimized airflow.

The motherboard selected combines enterprise features with bleeding-edge personal computer (PC) features. This gives it the real-time rendering advantages of a gaming rig with the high reliability and management features of an enterprise server. To accommodate future upgrades, the motherboard supports a wide breadth of the latest high-end Intel processors allowing CPU choices ranging from the fastest clock speeds to highest core counts.

The motherboard supports 40 PCIe lanes in contrast to more common mainstream boards which typically only offer 16 PCIe lanes to be shared amongst all devices. The availability of 40 PCIe lanes facilitates dedicated lanes for the GPU, which optimizes performance. The GPU ultimately selected for this system can support up to 16 PCIe lanes. The motherboard also supports error-correcting code (ECC) memory, which is the most reliable option for this type of computer.

The motherboard includes Intelligent Platform Management Interface (IPMI) functionality for remote management. A server grade network interface card (NIC) from Intel was installed in such a way that the NIC and GPU reside on the same PCIe root hub. This allows data between the NIC and GPU to skip a transit hub, which can introduce unnecessary latency.

Finally, the motherboard supports Non-Volatile Memory Express (NVMe) storage devices, which are faster with lower latency than the classic Serial AT Attachment (SATA) bus. Server grade solid state drives (SSDs) were selected for high reliability and input/output (I/O) bandwidth.

RAVE assembled the systems using custom internal cable lengths to maximize reliability and increase airflow. A 24 hour burn-in was performed prior to shipping to catch up to 99 percent of the failures that would otherwise occur in the first year.

Ultimately, this solution provides TARDEC with a platform that can support growth for additional capabilities added to the autonomy software in the future. A picture of the RAVE HPC workstations is shown in Figure 2. A table outlining the RAVE HPC design criteria and decision process is included in Appendix B.

Figure 2: RAVE HPC Workstation

Each RAVE HPC workstation in CASTLE represents one vehicle. This is an improvement over the workstations that TARDEC used on a previous project that required two workstations per vehicle, one to represent the virtual simulation environment while the other supported the autonomy software. CASTLE can currently support up to an eight vehicle convoy, two four vehicle convoys, or four two vehicle convoys.

Quantum Signal's Autonomous Navigation and Virtual Environment Laboratory (ANVEL) provides the virtual world which includes: a vehicle model; simulated sensors such as LIDAR; and a simulated vehicle controller, which communicates with the autonomy software to stimulate it for testing similar to how the autonomy interacts with the real vehicle and world.

CASTLE will be able to run thousands of iterations of the same or similar scenario by only changing a few parameters, such as LF gap distance and convoy speed, and change the types of dynamic obstacle and locations of interference with the convoy follower vehicles.

TARDEC goals in using this kind of simulation environment include: to mitigate and reduce certain testing; to create a set of scenarios that get validated; to help direct testing that will support real-world testing; to reduce the test-fix-test cycle that normally occurs during testing; and to look at the edge case scenario or scenarios with the highest potential rate of system failure. Ultimately, the autonomy software is explored through the simulation environment in order to find behavioral, logical, and/or operational errors before moving onto the real-world testing.

ASSESSING THE VIRTUAL-WORLD TESTING SOLUTION

The verification and validation of CASTLE is being done by using data collected during real-world testing of the LF system concept. CASTLE has created a way to replay the leader vehicle positions, instead of driving the vehicle manually, and then have the followers' autonomy software execute in real time in the virtual simulation environment. And at the same time the location of the follower during the real test is replayed as a ghost vehicle to compare with the simulated follower.

Earlier in development, it was discovered that there were some issues between the ghost and simulated follower. Additional enhancements were implemented to the by-wire controller representation and vehicle model in CASTLE to increase fidelity with the real world. A screenshot of the CASTLE LF system concept simulation is shown in Figure 3.

Figure 3: CASTLE LF Simulation Screenshot

Additional testing has been performed to move a dynamic obstacle in between the leader and the follower to verify that the follower will stop for the dynamic obstacle and then start following again when the dynamic obstacle moves out of its path. This provides evidence that the simulated LIDAR is providing the correct input to the autonomy software to recognize an obstacle is in its path.

This real-time virtual environment provides the ability to run complex scenarios that would be difficult to run and be repeatable at the Army test ranges. Some examples include: the angle of the sun; weather conditions; and dynamic and static obstacles, such as other vehicles, pedestrians, and animals.

Currently, each RAVE HPC workstation has the computing power to support the introduction of newer versions of the LF autonomy software with more advanced autonomy behaviors into the real-time simulation environment. These HPC workstations were specifically designed for CASTLE. This is superior to using a one-size-fits-all solution for this problem because of TARDEC's need to be able to support running one vehicle simulation per computer for the autonomous system under test (SUT).

As new versions of the LF autonomy software are released, additional sensor simulations will need to be incorporated (e.g., multiple LIDARs, RADAR, et cetera). The autonomy software updates will also require additional testing and evaluation scenarios will need to be conducted to reduce risk, minimize real-world autonomy testing, and bring confidence to the system before and during fielding. The HPC workstations need to be able to support the current SUT work along with the planned and future autonomous capabilities that will need to be tested.

FUTURE OPPORTUNITIES

Originally limited to dedicated supercomputers, HPC has been adopted commercially in recent years through multi-core processors and graphics processing units (GPUs). GPUs are very efficient at manipulating computer graphics providing a highly parallel structure that is also efficient at parallel processing. The use of GPU accelerators make it possible to complete massive amounts of computation in a practical amount of time.

It is anticipated that the need for HPC and virtual-world testing will only increase as its benefits are more clearly established and as RAS operations become more complex. To accommodate the increased demand, more extensive virtual-world testing setups using HPC will be required with the capability to handle ever-larger data sets. To achieve the best results with HPC, the hardware solution needs to be designed to address the specific challenge. This is becoming even more important as HPC solutions are applied to a wider range of challenges and situations. This type of customization can ensure that the HPC system is able to meet any overmatch needs while balancing cost, power, and size considerations. This can be especially important for RAS as the available power and size is limited. The accumulation and use of large data sets is only going to increase as M&S analyses gain fidelity and RAS use becomes more ubiquitous and complex.

As HPC is required to effectively harness this data, it is imperative that the development and adoption of HPC hardware and software keeps pace with the exponential growth of data. Internal hardware and software development efforts must be maintained that explore how to best utilize available options when tasked with large data sets and explore customized HPC solutions for a wide range of scenarios.

Cutting-edge expertise in HPC is required in designing and developing custom HPC solutions for any problem. One must leverage the lessons learned from prior HPC customization efforts in order to apply the correct solution to each situation, ensuring that the challenge is met as efficiently and effectively as possible.

In addition, future HPC efforts are going to require solutions that rely upon different technologies or approaches. There are a variety of more efficient, effective, and/or economical options currently being envisioned or already being harnessed in other fields. Classical HPC is implemented by spanning a workload across many different CPUs using a cluster of networked computers. This approach is limited by space, heat, and cost concerns, and most development to this point has been focused on solving those issues.

Recently, software environments have been enhanced to allow much higher performance when applied to systems featuring large numbers of processing cores. This development allows each CPU core to process independently and each CPU can currently have up to 32 cores. This limitation is one factor that has drawn attention to GPUs, which can provide up to 5120 cores each. This disparity has spurred a current developmental focus as a network of GPUs can provide significantly more cores than a similar network of CPUs. In addition, the high number of cores available in GPUs make them ideal for processing massively parallel workloads such as those associated with the deep neural networks used to implement Artificial Intelligence (AI).

Performing autonomous vehicle testing virtually using in-the-loop simulation can produce the same mountains of data that real-world testing produces. With AI being the best way to extract usable information from mountains of binary data, it makes sense to apply AI to the data produced by the simulation. Just as the sensors on autonomous vehicles produce a flood of data during test driving, the virtual sensors in a simulated autonomous vehicle also produce data at a high rate. Since AI can comb through data very quickly and produce actionable information, it makes sense to leverage AI in such a fashion. AI can recognize patterns well beyond the capability of humans and thereby accelerate discovery and remediation of flaws in the design, logic, and comprehension of an autonomous vehicle.

Some questions that highlight the potential value of AI with mobility testing include: how do we know which simulation events to focus on; and how do we know which scenarios to test further and how can we identify the edge cases that are very rare, but possible and dangerous?

In the first case, with the ability of AI to predict behavior before it happens, AI can identify the edge cases most likely to result in damage or injury and highlight those for us. And in the second case, real-world driving is unlikely to find very many rare edge cases that are dangerous within a useful timeframe, and finite simulation iterations may miss them as well. Since AI can also predict behavior which has not actually happened, it can be applied to the simulation, by observing the data or processing the output data or both, and identify such cases before they happen.

These types of situations help to highlight the additional value that virtual-world RAS testing can have over real-world RAS testing beyond the logistical and safety concerns discussed in this paper.

CONCLUSION

The demands of military vehicle mobility testing are only going to increase as more vehicles with RAS functionalities are developed, as RAS decision making grows in complexity, and as sensor suites continue to provide an over-abundance of sensor data to process. All of these factors make the real-world testing of RAS platforms even more unwieldy and highlight the importance and economics of conducting a significant portion of the testing in a virtual environment. As demonstrated with the TARDEC CASTLE system, HPC workstations and M&S capabilities can be combined to effectively and efficiently test RAS in a cost-effective and timely way.

REFERENCES

[1] Jack Stewart, "Why Tesla's Autopilot Can't See a Stopped Firetruck", Wired, 2018 25 JAN - https://www.wired.com/story/tesla-autopilot-why-crash-radar/

[2] Steve LeVine, "What it Really Costs to Turn a Car into a Self-Driving Vehicle", Quartz, 2017 05 MAR - https://qz.com/924212/what-it-really-costs-to-turn-a-car-into-a-self-driving-vehicle/

[3] Ding Zhao, Huei Peng, "From the Lab to the Street: Solving the Challenge of Accelerating Automated Vehicle Testing", Mcity White Paper, 2017 MAY, https://mcity.umich.edu/wp-content/uploads/2017/05/Mcity-White-Paper_Accelerated-AV-Testing.pdf

[4] Alexis C. Madrigal, "Inside Waymo's Secret World for Training Self-Driving Cars", The Atlantic, 2017 23 AUG - https://www.theatlantic.com/technology/archive/2017/08/inside-waymos-secret-testing-and-simulation-facilities/537648/

APPENDIX A

	REAL-WORLD TESTING (PROBLEM)	VIRTUAL-WORLD TESTING (SOLUTION)
Test Range	2 test range personnel required	No range personnel needed
	1 test engineer required	1 test engineer to setup the scenarios to run
	Weather conditions are limited to those available at the time and location of testing (e.g., if you want to do cold weather testing than you will need to go to a cold weather region)	Computer simulation can easily model different weather conditions and times of day
	Terrain type is limited to what is available at the test range	Infinite terrain types are possible
	Test range availability is limited	Virtual world is available 24/7
	Test range usage has a cost per hour	Once virtual terrain is built, it can be used as many times as needed with no additional cost
Drivers	Lead vehicle = 1 driver Safety drivers for Followers (e.g., 3 Followers = 3 people)	Lead driver is automated No safety drivers are necessary
Test Vehicles	Testing a single March Unit (1 Leader + 3 Follower vehicles) requires four vehicles	Trucks are all virtual so there is no limit
	Simulating a convoy comprised of multiple March Units requires four more physical vehicles per additional March Unit	Simply add more computers if more trucks are needed
	1 - 2 spare vehicles are needed on site to fill in for maintenance or fuel interruption, preventing a delayed or degraded test as a result of temporarily missing vehicles	No spare vehicles needed, virtual vehicles don't break down or need to refuel
	Additional trucks may not be available; Each truck adds an ongoing cost	There is a very low one-time cost for an additional computer (i.e., truck)
	Fuel cost is ongoing for all testing	No fuel cost
	Vehicle maintenance cost (mechanic available)	No vehicle maintenance cost
Obstacles	Physical obstacles must be available and are vulnerable to damage	One time cost for creation of virtual obstacles which can then be used as many times as needed
	Could damage vehicle	Represents traffic, crowds and obstacles virtually
Degraded Modes	Sensor failure - hard to represent in real world	Can be scripted to turn off sensor in software

APPENDIX B

Functionality/Requirement	Reason/Benefit	Component	Approach/Decision
Server Grade Reliability	Systems must be able to run continuously; Client grade components are not designed to be run 24/7	All	Select best of breed components designed for enterprise grade reliability with extended temperature range, high MTBF, et cetera; Perform ruggedization procedures during system integration; Provide sufficient cooling to keep components well below their maximum rated temperature during long term operation
High-end Graphics Support	Application rendering performance, graphics output	GPU	Testing proved the NVIDIA GeForce GTX 1080Ti was optimal
Rackmount Chassis	Greatly reduced footprint overall	Chassis	Short-depth chassis; Server grade; Rugged with excellent airflow
Remote Management	Can access system remotely for monitoring and management	Motherboard	Enterprise grade motherboard with IPMI (remote management) support
Low Latency I/O	Minimize latency moving data from CPU to network and/or GPU	NIC, GPU, motherboard, storage	Used an add-in server-grade NIC in a slot that puts the GPU and network on the same PCIe root hub; Chose an enterprise grade SSD for high performance and low latency

Terms:

- MTBF = Mean Time Between Failures, a standard measure of expected component reliability
- SSD = Solid State Disk (offers high speed and low latency with no moving parts)
- COTS = Commercial Off The Shelf (non-proprietary)
- LGA1151 = Intel's current mainstream CPU socket
- CPU = Main processor in a computer
- SATA = Protocol used for transferring data to internal storage
- X99 = Motherboard chipset which controls onboard components
- I/O = The input & output of data
- Xeon E5 = High-end server CPU family made by Intel
- Core i7 = High-end desktop CPU familiy made by Intel
- PCIe Lanes = High-speed serial connections between components and CPU. PCIe lanes are limited by CPU, chipset, and motherboard design
- PCIe Root Hub = Components can be connected via PCIe to either the CPU or chipset which are separate root hubs. Moving data between root hubs adds latency