Zero Roadway Deaths Means Seeing Everything Sooner, Clearer and Farther

Behrooz Rezvani, Babak Hassibi and Lawrence Burns

Neural Propulsion Systems (NPS)

February 2022

Introduction

While you read this article, forty people will likely die on the world's roadways and fifteen hundred will likely be injured. The global epidemic causing this trauma surfaced over a century ago in London when Bridget Driscoll became the first person killed in a motor-car accident. Today, this epidemic accounts for over 1.3 million fatalities and 50 million injuries per year, with half being pedestrians and cyclists.

While roadway safety has improved over the past several decades, all countries continue to face formidable challenges. To accelerate progress, **Vision Zero** in Europe, **The Road to Zero Coalition** in the U.S., and the recently announced **U.S. National Roadway Safety Strategy** have embraced zero roadway deaths as the objective, not just improved safety. As stated by U.S. Secretary of Transportation Pete Buttigieg, "**our goal is zero deaths** — where one day no one has to say goodbye to a loved one because of a traffic crash."

Wayne Gretzky, the world's greatest hockey player, taught us that it is important to skate to where the puck is going to be, not to where it has been. In this context, roadway safety strategies need to anticipate advances in digital and electronic technology and comprehend what this can contribute as part of a future system leading to zero deaths.

Advances in radar, LiDAR, processors and analytics are allowing vehicles to sense what is up ahead and around much better than humans can see. These advances should be an integral part of future roadway transportation systems because the better vehicles can see objects in their paths, the better they can respond (with and without human supervision) to avoid accidents. This enhanced vehicle safety will, in turn, influence what is required of the rest of the roadway transportation system (e.g., roads and traffic laws) to enable zero fatalities.

Given how technology continues to improve vehicle sensing, this paper addresses three questions:

- How well must vehicles see to enable zero roadway deaths?
- What is required to see this well?
- Can these requirements be met?

We believe answers to these questions will be useful to industry analysts, engineers, researchers, academics, executives, and government officials as they consider their approaches to attaining zero roadway deaths. To support our conclusions, it is necessary to introduce and apply advanced concepts in sensing and computation. Hopefully, these concepts are explained sufficiently so they can be understood by a broad audience and also convince subject matter experts that our answers are sound.

Based on principles from physics and information theory, we conclude it is possible for sensors to see well enough to enable zero roadway deaths. And, based on demonstrated prototype testing of NPS's sensing and processing platform, we conclude commercially viable and scalable technology exists to do this, and it is realistic to build such a system.

In order to see soon enough, clear enough and far enough to enable zero roadway deaths, vehicles must be able to process a surge of sensor data for the worst roadway conditions. We will show that this requires sensing and processing a peak data rate on the order of 100 X 10^{12} bits per second (100 Terabits per second). This immense requirement is 10 million times greater than the sensory data rate from our eyes to our brains[1].

We will also show that sensing and processing 100 Tb/s can be accomplished by combining breakthrough analytics, advanced multi-band radar, solid state LiDAR, and advanced "system on a chip" (SoC) technology. Such an approach will allow companies developing advanced human driver assistance systems (ADAS) and fully autonomous driving systems to accelerate progress. The key question for these companies should be "What must be true to get to zero roadway deaths?". We have concluded that sensing and processing about 100 Tb/s is one of these necessary requirements and this is indeed possible.

1 Neuroscientists believe that the input sensory data rate from human eyes to human brains is about 10 X 10^6 bits per second (10 Megabits per second). https://www.britannica.com/science/information-theory/Physiology

How Well Must Vehicles See to Enable Zero Roadway Deaths?

Safety Margins

Consider a truck driving on a road when an unexpected hazard (like a deer) enters the road up ahead (Figure 1). To prevent an accident, the truck must stop before hitting the hazard. The stopping distance is the distance traveled during the time it takes the driver to sense there is something up the road and perceive what it is, and the time it takes for the vehicle to stop once the driver applies the brakes.

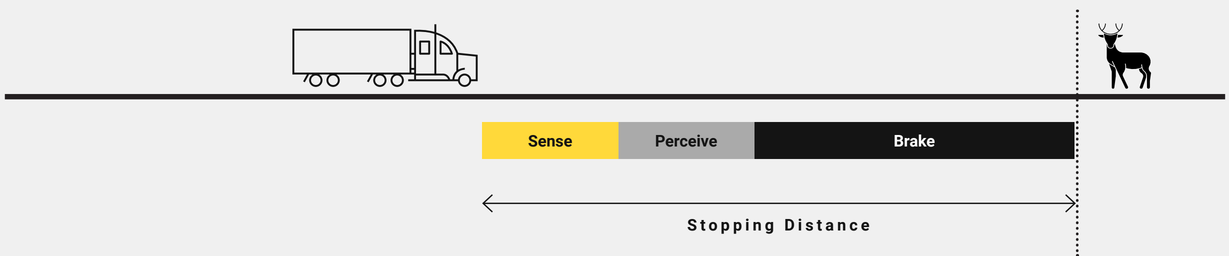

Figure 1. Stopping Distance = Sensing + Perceiving + Braking Distance

Under ideal circumstances, the road is straight and flat, the weather is dry and sunny, the driver is alert and law-abiding with a keen mind, the brakes are properly sized for the load with unworn pads, and the tires have good treads and proper air pressure. In this case, stopping distance is minimized and predictable based on vehicle speed and the coefficient of friction.

Unfortunately, the ideal case rarely exists. Roads are curved and sloped, it can be rainy/snowy/foggy and dark, drivers get distracted and impaired, brake pads and tires wear, loads and speeds exceed limits, and tire pressure is not always appropriate. These realities mean the stopping distance is often much greater than ideal and crashes occur. They also mean that to have zero roadway deaths, driving systems with and without human supervision must account for the worst conditions.

Figure 2 illustrates how seeing sooner, clearer and farther provides a safety margin to deal with less than ideal conditions.

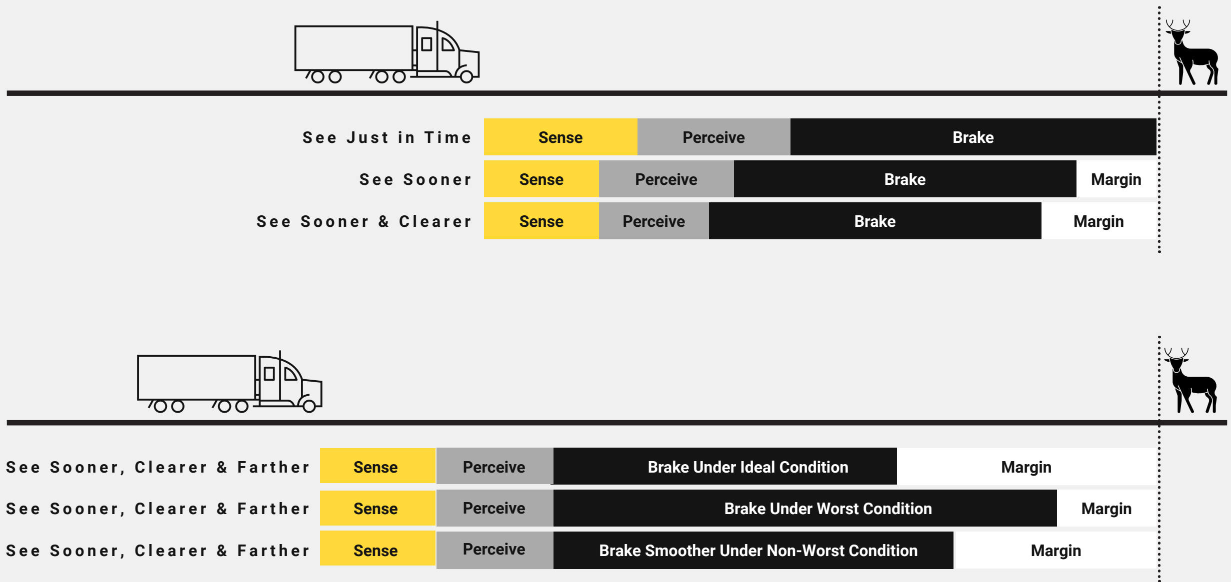

Figure 2. Seeing Sooner, Clearer and Farther Increases Safety Margin and Enables Worst Conditions and Smoother Braking.

Seeing sooner reduces sensing time (so perception can start earlier), seeing clearer reduces perception time (so braking can start earlier), and seeing farther allows sensing to start earlier. The safety margin provided by enhanced sensing allows worst-case conditions to be handled before a crash occurs and smoother braking under non-worst-case conditions.

Figure 2 shows that to prevent accidents, vehicles must see objects in their paths soon enough, clear enough, and far enough to be able to stop in time under the **worst conditions**. These worst conditions include:

- rain, snow, ice, fog, and darkness
- maximum allowable speeds given road designs
- curves, slopes, buildings and canyons that can obstruct vision
- lower bounds on brake conditions relative to loads and tire conditions relative to temperature, pressure and surface friction

Required Data Rates

To determine how well a vehicle must see to eliminate preventable roadway deaths, we need to know the data rate (bits/sec) required to reconstruct the worst-case scenario with sufficient precision (fidelity) and frequency (frames/sec) for the vehicle to be able to stop in time when it is possible to do so given a hazard in its path. For zero roadway deaths, high precision is required which means a very detailed scene must be rapidly updated. As we show below and in the Appendix, this results in a huge data rate and requires an innovative sensing and processing platform to generate and assimilate the data.

As shown in Figure 3, a vehicle must see horizontally 360º around at short range and up-the-road within a narrow sector at long range. It must also see vertically at short and long range based on elevation angles.

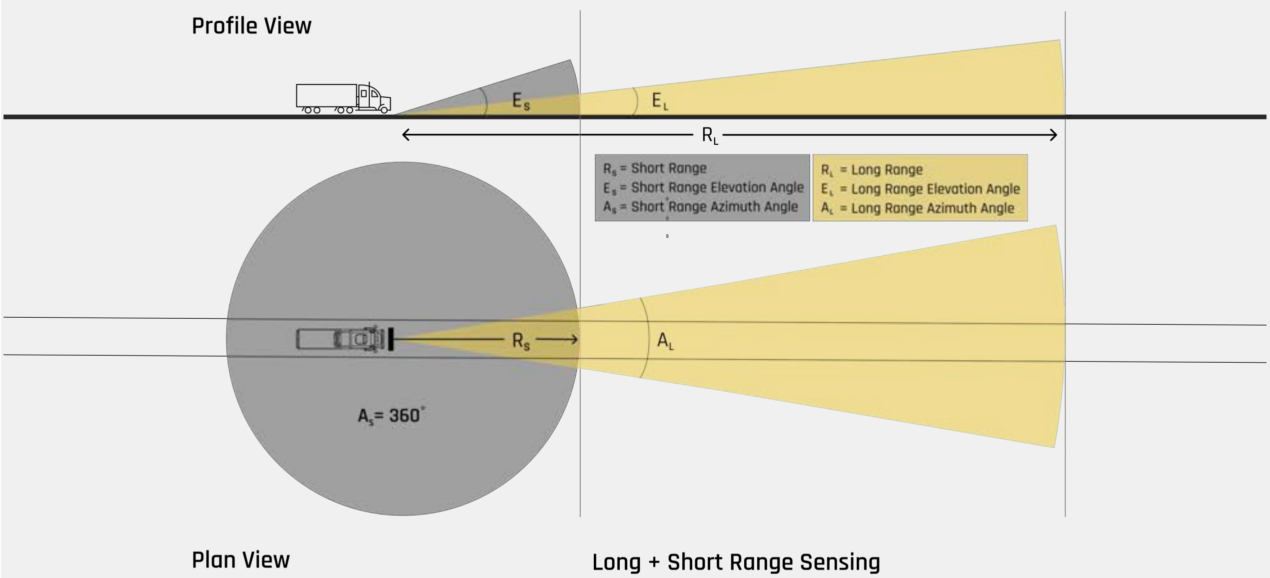

Figure 3. Vehicle Sensing Range and Field of View.

The data contained in the physical scene encompassed by the gray and yellow space is obtained by converting continuous analog signals representing physical measurements into time separated 0 and 1 digital streams to represent information.

Reconstructing the scene with sufficient precision requires the gray and yellow space to be segmented into small "cube-like" building blocks called voxels (think pixels with a third dimension). Greater precision can be attained by interrogating larger numbers of smaller voxels.

When radar and LiDAR sensors interrogate raw data in voxels, they need to reliably detect whether a target or multiple targets are there and avoid false alarms. Because analog signals must be detected in noisy environments, sensors must probe voxels numerous times to attain state of the art reliability greater than 90% and false alarms at a frequency less than one-in-a-million. And, to prevent accidents, this performance must be attained for the weakest signal-to-noise ratio (SNR) at the maximum range (i.e., worst-conditions).

Thus, to determine the data rate required to prevent roadway deaths for worst conditions, we must know:

- the required sensor ranges and fields of view around and ahead of the vehicle
- the required sensing precision and corresponding fidelity of the reconstructed scenes
- the required frequency at which the scenery must be captured (the system frame rate in Hz)
- the number of bits in the analog-to-digital converter determined by the minimum to maximum signal range within a voxel
- the energy the sensor sends to probe scenes to reliably recover a target in worst conditions at maximum distance (this energy determines the required number of signal repetitions)

The Appendix uses representative values of these key variables and formulas based on physics principles and information theory to calculate the required data rate for zero roadway fatalities. The representative values are for:

- light-duty vehicles, such as personal cars, operating in urban and suburban settings encountering extreme worst-case conditions resulting from the need to see around corners and through occluding objects
- long-haul trucks operating on interstate & state highways encountering extreme worst-case conditions resulting from the braking dynamics of tractor-trailer rigs, high vehicle speeds and masses (including load mass), and variation in brake and tire conditions

Light-duty vehicles in metropolitan areas require shorter long-range sensing than long-haul trucks due to lower vehicle speeds and masses, but higher data sampling rates to see around corners and through occluding objects.

Our calculations indicate that in both cases about 3 billion voxels in the combined short and long range scenes in Figure 3 need to be probed by the sensors to attain the required precision for zero roadway fatalities. Voxels located farthest from the vehicle are larger (around 8 cubic feet) than the average voxel (around 1/3 cubic foot) and require orders of magnitude more sampling than voxels nearby.

Our calculations also indicate that in both cases the required data rate for the entire gray and yellow space to enable zero roadway deaths under worst-case conditions approaches a staggering 7×10^{15} bits per second (7 Petabits per second)! This assumes that we interrogate each voxel with scanning radar and/or LiDAR with the azimuth and elevation angles being swept by beams and the range being swept by collecting time samples. This method of observation is the same as what scientists use at CERN Super Collider looking for novel particles[2], or James Web Telescope looking for very faint signals.

Neuroscientists[3] believe that the input sensory data rate into the human brain is about 11 Mb/s worth of information with vision accounting for 10 Mb/s. We need microscopes, telescopes, and CT & MRI scanners because human-eye optical capabilities are limited and we cannot see beyond the visual spectrum. If humans had X-ray vision and superhero capabilities like Superman, then perhaps we could sense and interpret data rates on the order of Petabits per seconds and process this information flow fast enough to avoid all car accidents. But, there are things we simply cannot do. Compared to 7 Pb/s, humans see at just about one-billionth of the required information processing rate to prevent accidents. The conclusion is that **humans cannot see well enough to attain zero driving fatalities.**

High definition cameras could work better than human eyes if designed properly. Their data rates are in the tens of Mb/s (normally compressed), and one can use multiple cameras at the same time. However, because cameras give two-dimensional pixels and cannot probe individual voxel depths, some kind of accurate depth perception is needed which is extremely challenging, even with multiple cameras. The central limiting factor of cameras is similar to the human eye. Cameras work well when there is sufficient light in the environment and it is not necessary to see through objects and around corners. But, to get to zero roadway deaths, we need to detect objects with short notice that are hidden from view and exist in the worst-case lighting conditions.

2 https://home.cern/resources/faqs/facts-and-figures-about-lhc CERN collects about 240 Petabits per year of data.

3 https://www.britannica.com/science/information-theory/Physiology

Cameras fail in finding these objects no matter how fast they can sample the environment. Thus, like humans it appears **cameras alone cannot sense well enough to attain zero driving deaths.**

The **worst-case** data rate requirements occur when:

- a vehicle first starts operating and has no idea of the scene
- the scene changes noticeably (such as when the vehicle makes a turn or the road takes a sharp bend)
- the weather conditions are poor (the signal-to-noise ratio drops drastically and much more signal repetition and averaging is required)

The average required data rate in typical conditions is much less than the worst-case because the scenery does not change very rapidly and early frames carry a lot of information about current and future frames. Hence, we do not need to scan as repeatedly. However, having zero accidents requires maintaining performance in the worst-case, not the average case.

Based on sensitivity analyses for key variables like range, worst-case SNRs, frame rates and precision, the bottom line is that **any sensor system that physically scans the entire coverage volume will require an incredibly large data rate** to reconstruct the scene soon enough, clear enough, and far enough to eliminate roadway deaths under worst conditions. **No current system comes close to this.**

New Method: The Atomic Norm

A new mathematical framework has been discovered that, when combined with advanced sensors and "system on a chip" (SoC) technology, allows the massive data rate requirement for scene reconstruction to be addressed and vehicles with and without human supervision to sense well enough to eradicate roadway deaths.

This new method, called the Atomic Norm (AN), is based upon a compressed sensing framework developed in the early 2000's at Caltech, UCLA, and Stanford to improve magnetic resonance imaging (MRI), and first published by researchers at Caltech and MIT in 2012[4]. Compressed sensing reduces the number of measurements required to maintain a certain level of performance.

The AN framework has been tailored specifically for radar, LiDAR detection and signal processing. AN has many advantages over conventional radar techniques. Specifically, it inherits the defining property of compressed sensing and reduces the data rate requirement by orders of magnitude compared to scanning LiDAR and beam-steering radar. And it is capable of handling worst-case scenarios (e.g., fast changing environments).

One of AN's advantages is that it does not require the scene to be **physically scanned** with narrow beams. Instead, it uses much wider beams and better computation to allow each voxel in the coverage volume to be interrogated individually. The AN method can be applied to do this efficiently and with great efficacy, and NPS has demonstrated this capability for the first time ever.

The AN relies on the fact that not all voxels are occupied — in fact, most are not and are simply free space. If only a fraction of the voxels is occupied, then the AN can obtain the **exact same performance** as physical scanning with only a fraction of the data rate. Typically, just 0.1% to 1.0% of the voxels are occupied and this reduces the data rate by a factor of 0.01 to 0.08 times the rate without the AN while continually monitoring the entire space at all times. (The Appendix summarizes the equations used for this result.)

---

4 "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information." EJ Candès, J Romberg, and T Tao. IEEE Transactions on Information Theory 52 (2), 489-509 (2006)

"Compressed Sensing". David L. Donoho. IEEE Transactions on Information Theory 52 (4): 1289-1306 (2006)

"The Convex Geometry of Linear Inverse Problems". Venkat Chandrasekaran, Benjamin Recht, Pablo A. Parrilo and Alan S. Willsky. Foundations of Computational Mathematics. Volume 12, pages 805–849 (2012)

In addition to reducing the required data rate, the AN can obtain the same performance as conventional radar algorithms by using 1/50th of the transmit power. This is a 17 db improvement[5] in performance over conventional radar algorithms, yielding much higher resolution and many more hits on targets. Current conventional automotive radars can barely detect a pedestrian up to 100 meters away, but with the new AN technique it is possible to see a pedestrian in excess of 600 meters[6].

Using the AN, and based on sensitivity analyses for the key variables, we conclude that human-driven and autonomous vehicles must be able to sense the environment at a peak rate on the order of 100 Tb/s to handle a surge of sensor data in the worst conditions.

---

5 Alternatively, if we fix the transmit power, the AN can increase the detectable range by more than a factor of 2.5, thereby allowing one to sense much farther.

6 "Letting Robocars See Around Corners". IEEE Spectrum, January 23, 2022. Behrooz Rezvani, Babak Hassibi, Fredrik Brannstrom and Majid Manteghi. https://spectrum.ieee.org/car-radar

What is Required for Vehicles to Sense 100 Tb/Sec?

There would be no preventable roadway deaths and injuries if:

- roadway users never collided
- vehicles always stayed on the road
- vehicles always stayed upright

In practice, these absolute requirements are overly stringent. Roadway users could collide if impact forces do not cause any injuries. Similarly, vehicles could leave the road and/or turn over if they are designed to prevent injuries from these events given the physics involved. However, to keep things simple, let's stick with the absolute requirements.

Figure 4 shows how the requirement of zero roadway deaths flows down to specify requirements for sensing and how meeting sensing requirements flows up to enable other system requirements.

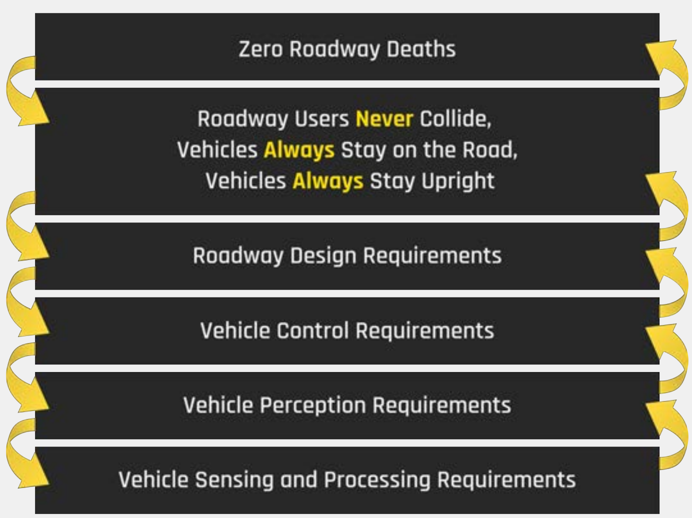

Figure 4. Requirements Flow-down for Zero Roadway Deaths.

Our premise is that 100% awareness of what is in the environment at all times is required to enable zero roadway deaths, and by sensing and processing on the order of 100 Tb/sec in worst conditions, the higher-level system requirements will be easier to meet. To do this, in addition to the Atomic Norm, a combination of multi-band radar, solid state LiDAR, SoC processors, and digital maps will be required to sense and process information that can used to recognize different road users and road infrastructure, and to detect and provide information about unknown and occluded objects for further analysis by the perception engine.

Different radar bands complement each other by having different propagation and reflection properties. The higher bands have better resolution, but have less range, are more susceptible to rain and fog, and cannot go through objects or see through occlusions. The lower bands have lower resolution, but have better range (especially in adverse weather conditions and can both go through and bend around objects. Because different objects have different reflectivity (radar cross sections at different frequency bands, one can use the different radar responses to determine the constituent material of different targets (e.g., differentiate metal from soft tissue. This means the radar has to use causal inference and reasoning to interpret the multi-spectral response from the environment.

Solid state LiDAR has precise scanning repeatability advantages and avoids moving parts, which could be susceptible to damage and failure. In addition to radar, LiDAR is used for the long-range field of view (FOV. It relies on having multiple lasers and detectors looking at the same portion of the FOV to increase the signal-to-noise ratio and to improve the range while adding photon path diversity for higher reliability. Carefully designed laser pulse sequences are used to boost the signal levels and detect and position multiple targets within a single voxel.

Highly advanced AI edge processors are needed to track, analyze, and interpret the blistering data flow from each sensor modality. SoC processors are high-end very low power dedicated silicon chips with semi-flexible architecture. They are the computational workhorses required to process the combined rate of 100 TB/sec, so several tens of them will be required to work in parallel and in symphony.

Finally, digital maps help us know when roads bend and slopes occur and give a priori clues for the sensing system to survey more intelligently by focusing the long-range LiDAR and radar on the voxels of greatest interest. Also, knowing where hills crest allows speeds to be safely adjusted to mitigate the risk of not being able to see over hills. Maps are one of many information sources which can be used for focusing sensing attention on where greater interrogation is needed. The current route plan, predictions about other vehicles, weather conditions, audio detection of emergency sirens, image detection of flashing police lights, time of day... can all help guide interrogation.

Can Vehicles Sense 100 Tb/Sec?

To assess whether on the order of 100 Tb/sec can be sensed in an environment, pilot-scale prototype hardware and software were developed and then tested at an airfield in California. The prototype is based on a sparse, wide-aperture multi-band radar architecture. It uses a variety of frequencies and exploits the particular properties of each to deal with the unwelcome impacts of weather and to see through and around corners. It also employs solid state LiDAR, Atomic Norm advanced signal processing, and sensor-fusion algorithms to produce an integrated digital representation of the environment.[7]

Experiments verified that the range, angular resolution, and precision of the core sensor element[8] for an envisioned platform are close to theoretical performance limits. This pilot scale proof-of-concept leads us to conclude that the technology exists for vehicles to see well enough to enable zero roadway deaths. For us, enough hardware and software has been tested to prove 100 Tb/sec is viable. This is no longer a conjecture or a hypothesis. All that remains now is scaling.

---

7 "Letting Robocars See Around Corners". IEEE Spectrum, January 23, 2022. Behrooz Rezvani, Babak Hassibi, Fredrik Brannstrom and Majid Manteghi. https://spectrum.ieee.org/car-radar

8 A sensor element is an observation device that can measure a signal of interest with certain accuracy. Often-times the signal level of interest is buried deeply in the noise and so the observations must be repeated thousands of times in order to recover the signal of interests. Each observation is measured in bits.

Conclusions

Based on principles from physics and information theory, we conclude it is possible for sensors to see well enough to enable zero roadway deaths. And, based on demonstrated prototype testing of NPS's sensing and processing platform, we conclude commercially viable and scalable technology exists to do this, and it is realistic to build such a system.

In order to see soon enough, clear enough and far enough to enable zero roadway deaths, both human-driven and autonomous vehicles must be able to sense the environment at a peak rate of about 100 Tb/s to handle a surge of sensor data in the worst conditions. Sensing and processing information at this immense rate will allow companies developing advanced human driver assistance systems (ADAS) and fully autonomous driving systems to accelerate progress toward zero roadway deaths. The key question for these companies going forward is "What must be true to get to zero preventable roadway deaths?". We have concluded that seeing and processing about 100 Tb/sec is one of these necessary requirements and this is indeed possible by combining breakthrough analytics, advanced multi-band radar, solid state LiDAR, sensor fusion and system on a chip technology.

The path forward should focus on two types of driving: human and autonomous. Human driving requires a human to be actively engaged to supervise the driving system. Autonomous driving never requires human supervision.

Both human and autonomous driving must target zero roadway deaths and injuries as the objective. By simultaneously integrating human and autonomous driving systems in vehicles and using 100 Tb/second technology, telematics data (real-time and historic) and digital maps to anticipate the road ahead, one can determine when it is safe to drive autonomously with zero deaths and injuries and when it isn't. Such an approach can safely prepare humans to help drive when needed (vs. requiring instantaneous situation awareness), inform them about what they will likely need to do, and prevent the vehicle from operating autonomously if they do not engage. As technology improves, the percent of each trip driven autonomously and safely will increase, and more value will be realized from autonomous driving systems.

When will autonomous driving reach 100%? Most likely never. As David Zipper[9] emphatically states, "Only humans are dumb enough to think they should be driving in every condition possible, whereas machines are much smarter than that."

[9] https://twitter.com/davidzipper/status/1478016041582272514?s=11

The objective is zero roadway deaths and injuries, not 100% autonomous driving. Driving does not need to be 100% autonomous under all conditions to realize significant societal, business and consumer value. But, when driving autonomously, ensuring zero accidents is essential.

Some might argue that the cost of sensing 100 Tb/sec is impractical for large scale deployment on vehicles with and without human supervision. While much remains to be proven, to argue against pursuing this opportunity based on cost is to argue against the history of technology innovation. Consider that chips in the early 1950's could only be afforded by the US government for very special defense applications. And look at where chip cost and chip proliferation are now.

Furthermore, ride-hailing drivers earning $15/hour and averaging 30 operating miles per hour cost $0.50 per mile. This is equivalent to spending $100,000 on a fully autonomous driving system with a 200,000-mile life. And, for long haul truck drivers costing $0.80 per mile, the corresponding breakeven cost is $160,000. The cost of sensing and processing 100 Tb/sec to enable zero roadway deaths should not be a "show-stopper", especially given anticipated reductions in SoC cost and that human drivers collectively have proven they cannot get to "zero".

Finally, **Vision Zero, The Road to Zero Coalition**, and the **U.S. National Roadway Safety Strategy** have embraced zero roadway deaths as the objective, not just improved safety. These important initiatives focus on a system-of-systems approach to attain this objective. Being able to sense and process 100 Tb/sec could significantly reduce the costs implied by this redundant system approach.

NPS has demonstrated multiple sensor elements of its envisioned 100 Tb/s system and has proven that it meets the theoretical expectation of physics and information theory. Based on this progress, our intent is to develop and validate a full scale, cost effective system that can be brought to everyone. To this end, NPS looks forward to teaming with ADAS and autonomous driving system developers, telematics companies, system-on-chip suppliers, government agencies, and roadway safety advocates to accelerate progress toward zero roadway fatalities.

Henry Ford said his goal was for every working family to own a car. Our goal is to prevent these families from losing a loved one in a car crash.

APPENDIX

To better understand the edge cases for long-haul trucking, we spoke with trucking company operators and people developing autonomous driving systems for over-the-road trucks. We asked them how far up the road and around trucks they feel they need to see to be able to safely stop under the worst conditions. In both cases, sentiments are to see:

- about 250 meters around and 1000 meters up ahead
- at short and long range elevation angles of about 24^{o} and 6^{o}, respectively
- with a long-range azimuth angle of about 30^{o}

Using these requirements and state-of-the-art values for range precision (0.3 meters) and angular precision (0.05^{o}), the short-range gray space in Figure 3 is segmented into 2.9 billion voxels and the long-range yellow space is segmented into 0.2 billion voxels. This means a total of 3.1 billion voxels in the combined short and long range scenes need to probed by the sensors. The average voxel is about 8.5 inches cubed (9800 cm3) and the voxel size at the maximum range of 1,000 meters is about 2 feet cubed (225,000 cm3). Not every voxel gets sampled or observed for information at the same rate. For example, voxels located at 1000 meters require orders of magnitude more sampling than the ones at 100 meters.

Signal-to-noise ratio (SNR) is measured in decibels. A -30db SNR means the power of the reflected signal off a target is one-thousandth of the power of the ambient noise, and a -40db SNR is one-ten-thousandth. To get a detection reliability over 90% and a less than one-in-a-million false alarm frequency, one needs a signal-to-noise ratio of 12 db (i.e., that the power of the reflected signal must be 16 times greater than the power of the noise). Achieving this desired signal-to-noise ratio from signals that are -30 db and -40 db below the noise level requires many thousand rounds of signal retransmission.

Using -30db and -40db as representative values for worst-case SNRs at ranges of 250 meters and 1,000 meters, respectively, 12 bits accuracy of each sampled observation, and frame rates of 20 Hz, the resulting data rate for the entire gray and yellow space is a staggering 6.8×10^{15} bits/second (6.8 Petabits per second)! About 55% of this rate is for short-range (3.7 Pb/s) and about 45% is for long-range (3.1 Pb/s). The required data rate for the average voxel is 2.2×10^{6} bits/second (2.2 Mb/s).

Light-duty vehicles in metropolitan areas require shorter long-range sensing than long-haul trucks due to lower vehicle speeds and masses, but higher data sampling rates to see around corners and through occluding objects. In this case, vehicles need to see:

- about 250 meters around and 500 meters up ahead
- at short and long range elevation angles of about 24º and 6º, respectively
- with a long-range azimuth angle of about 60º

Using the same values for the remaining key variables as we did with long-haul trucks, the resulting data rate is the same, 6.8 Pb/s.

The equations underlying these results are shown in the box below.

The Calculation

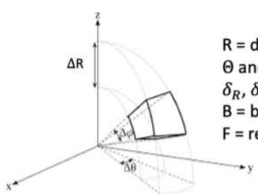

R = desired range
Θ and Φ = desired azimuth and elevation FOV
δ_R, δ_Θ, δ_Φ = azimuth, and angular precisions
B = bits in the system analog to digital converter
F = refresh rate (number of frames per second)

- Each voxel is of size $\delta_R \times R\cos(\Phi)\delta_\Theta \times R\delta_\Phi$
- Assume **each voxel can be scanned**
  - This is what is done in scanning radar and LiDAR
  - Azimuth and elevation angles are swept by beams
  - Range is swept by collecting time samples
- How many measurements do we need to make for each voxel?
  - via hypothesis testing
    - need 12 db SNR to get 90% detection and one-in-a-million false positives
    - number of measurements = $\frac{16}{SNR}$
- SNR is range dependent
  - Received signal strength decays as r^{-4}
  - Voxel cross-section grows as r^2
  - SNR = $\frac{E}{\sigma^2 r^2}$ where E is the total energy of the transmitted pulse
- Therefore the total number of measurements needed is
  - $M = \frac{\Theta\Phi}{\delta_\Theta\delta_\Phi}\frac{16\sigma^2}{E}\sum_{i=0}^{R/\delta_R}(i\delta_R)^2 = \frac{\Theta\Phi}{\delta_\Theta\delta_\Phi}\frac{16\sigma^2}{E}\frac{R^3}{3\delta_R}$
- $M = \frac{R\Theta\Phi}{\delta_R\delta_\Theta\delta_\Phi}\frac{16\sigma^2 R^2}{3E}$
- But $\frac{E}{\sigma^2 R^2}$ is simply the SNR at the maximum range, $SNR_{\max\_range}$
- $M = \frac{R\Theta\Phi}{\delta_R\delta_\Theta\delta_\Phi}\frac{16}{3SNR_{\max\_range}}$
- Assume each measurement carries b bits
  - b is determined by the dynamic range and is usually 10-12 bits
- Then the total number of bits to construct the scene is
  - $B_{total} = \frac{R\Theta\Phi}{\delta_R\delta_\Theta\delta_\Phi}\frac{16b}{3SNR_{\max\_range}}$
- Finally, if the system should operate at F frames per second

$$Data\_rate = \frac{R\Theta\Phi}{\delta_R\delta_\Theta\delta_\Phi}\frac{16Fb}{3SNR_{\max\_range}}$$

Atomic Norm and Compressed Sensing

Not all voxels in the scene are occupied

One can exploit this fact to avoid having to physically scan each azimuth and elevation angle

Instead, send wide beams and use signal processing to interrogate each voxel

From the theory of the atomic norm, it follows that if the fraction of occupied voxels is Δ, then the number of measurements can be reduced by the fraction

$$2\Delta\log\left(\frac{1}{2\Delta}\right)$$

Therefore,

$$AN\_Data\_rate = \frac{R\Theta\Phi}{\delta_R\delta_\Theta\delta_\Phi}\frac{32Fb\Delta\log\left(\frac{1}{2\Delta}\right)}{3SNR_{\max\_range}}$$

About the Authors

Dr. Behrooz Rezvani is a serial entrepreneur and founder and CEO of Neural Propulsions Systems, Inc. (NPS.ai), a pioneer in autonomous driving and sensing platforms. His vision and accomplishments have had a major impact on the wired and wireless industry and these systems are shipping in billions of flagship products across nearly every industry and application. The companies founded by Rezvani have had a major influence on the landscape of telecommunications. Prior to NPS, he was technical advisor to Liberty Global management for the Machine Learning application for understanding customer needs. He was the co-founder of Quantenna, an industry leading Wi-Fi technology company through a very successful IPO in 2016 and then purchased by ON Semiconductor for $1B in 2019. Quantenna was the first company in the world to pioneer high-speed wireless MIMO for home networks. He also founded Ikanos Communications, a leader in DSL modem/infrastructure IC and home gateways, and led its successful IPO in 2005 that later was acquired by Qualcomm. Ikanos was the first company in the world to deliver more than 100Mb/sec data to homes. The technology developed by Ikanos was later adopted by ITU-T as the global standard for VDSL. Dr. Rezvani has more than 40 patents to his name in the areas of communications and information theory. He obtained his Ph.D. from Marquette University. He was member of TAB for Century Link and advisor to Swisscom Ventures.

Dr. Babak Hassibi is the inaugural Mose and Lillian S. Bohn Professor of Electrical Engineering at the California Institute of Technology (Caltech), where he has been since 2001. He is co-founder and Chief Technologist of Neural Propulsions Systems, Inc. (NPS.ai), a pioneer in autonomous driving and sensing platforms. From 2011 to 2016 he was the Gordon M Binder/Amgen Professor of Electrical Engineering and during 2008–2015 he was the Head of the Electrical Engineering Department, as well as Associate Director of Information Science and Technology. Prior to Caltech, he was a Member of the Technical Staff in the Mathematical Sciences Research Center at Bell Laboratories, Murray Hill, NJ. He obtained his Ph.D. degree from Stanford University in 1996 His research interests span various aspects of information theory, communications, signal processing, control and machine learning. The work of his research group has been influential in the development of wireless communications standards (WiFi, 2G, 3G and 4G cellular). Among other awards, Hassibi is a recipient of the US Presidential Early Career Award for Scientists and Engineers (PECASE) and the David and Lucille Packard Fellowship in Science and Engineering. He was the General Chair of the 2020 IEEE International Symposium on Information Theory. He is the author of the books "Linear Estimation" (Prentice-Hall, 2000) and "Indefinite Quadratic Estimation and Control" (SIAM 1999). He has authored more than 300 peer-reviewed book chapters and journal and conference papers, as well as over 20 US and international patents.

Dr. Lawrence Burns is former Corporate Vice President of Research & Development and Planning at General Motors. He has also advised numerous companies – including Google Self-Driving Cars/Waymo for over a decade. After leaving GM in 2009, Dr. Burns was a Professor of Engineering Practice at the University of Michigan (2011-2015), and Director of the Program for Sustainable Mobility at Columbia University (2010-2014). He was elected as a member of the National Academy of Engineering (2011) and advises organizations on the future of mobility, logistics, manufacturing, energy and innovation. He coined the phrase “age of automobility” to describe this future in his book, "Autonomy: The Quest to Build the Driverless Car And How It Will Reshape Our World" (2018). He received a B.S. in Mechanical Engineering from General Motors Institute (now Kettering University), an M.S. in Engineering/ Public Policy from the University of Michigan, and a Ph.D. in Civil Engineering at the University of California, Berkeley.