NATIONAL AFFAIRS

Driverless Roads

Korok Ray

Hardly a day passes without some news about autonomous transportation. Not only have technology companies like Google, Uber, and Tesla made autonomous driving their top priority, but so have the old-line auto manufacturers in Detroit and around the globe. With 37,133 deaths from motor vehicles on American roads in 2017 alone, the academic and business communities have all heralded autonomous driving as a safer and better option for all in the long term. Indeed, the exceptionally broad-based investment in this technology means that its development will certainly proceed. The only questions are when and how.

But "how" turns out to be a more fundamental question than it might first seem. Today, much of this investment by the private sector in autonomous driving is based on a self-reliant vehicle (SRV) model. The vehicles communicate with their environments through expensive sensors. These sensors are a mix of Light Detection and Ranging (LiDAR) and Radio Detection and Ranging (RaDAR) sensors that provide visual and locational data to the car, and possibly to other similarly equipped cars. This technology does not come cheap. And while these sensors are powerful, they are far from perfect, as the tragic Uber crash in Arizona on March 18, 2018, proved.

There is an alternate approach that has lurked in the archives of automotive research for 20 years but is just now coming to light. A vehicle-to-infrastructure (V2I) regime permits communication between the vehicle and smart *infrastructure*, such as intersections, toll booths, bridges, and the roads themselves. V2I is fundamentally based on greater coordination, as the infrastructure acts like a central server, collecting data from the cars and sending data back to them. Much of this research occurs in

---

Korok Ray is an associate professor at the Mays Business School of Texas A&M University and director of the Mays Innovation Research Center.

universities, which are experimenting with navigating vehicles using sensor packs on the roads. V2I promises more visibility and coordination among the various vehicles, which could (presumably) lead to fewer accidents.

Though it is not obvious why, it seems as though the social tolerance for accident risk from autonomous vehicles lies well below that from human driving. This was made clear with the public backlash that Uber and Tesla faced after their recent crashes, which are inevitable with any experimental technology. While current SRVs use advanced software, like vision-based artificial intelligence, to (at best) estimate the location of other vehicles, V2I can more precisely determine location and maintain situational awareness since all data passes through infrastructure — one of several advantages to this approach.

V2I is not without its problems, however. V2I centralizes authority, which could lead not only to political rent-seeking, but also to greater cybersecurity risk; decentralized networks of SRVs are generally more impervious to cyber-attacks than systems that concentrate all available power and intelligence into fewer parties. It would also require some serious public-policy coordination. The main reason research around V2I died 20 years ago was a near-universal acknowledgement that the political will simply did not exist for smart infrastructure of the kind and level necessary for autonomous navigation.

Despite these problems, V2I has compelling arguments in its favor with regard to both cost and safety. And perhaps as important, V2I offers a once-in-a-generation opportunity to privatize America's roads, one of the last bastions of government control in private life today. With a careful and deliberate privatization plan, combined with a market-oriented approach to the economics surrounding smart infrastructure, V2I could be a better, safer, and faster path to our autonomous future.

CONSOLIDATING COSTS

Self-reliant vehicles are expensive today by any measure — well above the average price of $37,000 that American households currently pay for their automobiles. For example, the company AutonomouStuff located in Morton, Illinois, sells autonomous cars to research labs, car manufacturers, and academic institutions for experimental purposes. It can outfit a Lincoln sedan with LiDARs, RaDARs, cameras, a proprietary operating system, a positioning kit, and hardware in the trunk, all for a price well above $230,000 — not including the cost of the car itself.

The majority of the cost is in the new autonomous-specific electrical sensors and equipment. For example, a single Velodyne LiDAR can cost anywhere from $8,000 to $100,000. Radars are cheaper, but each vehicle requires five or six of them; each radar can cost approximately $8,000. Of course, the cost of LiDARs may fall over time, but only at scale, and the scale will be achieved only if the vehicle costs are low enough. Thus, there is an inherent chicken-and-egg problem, as auto manufacturers can only bring down the cost with high production, but consumers will not buy in volume until prices come down. Tesla, in fact, faces this precise problem right now with its Model 3 sedan.

Even if the cost of such vehicles drops over time, the SRV approach will still rely on the same infrastructure America installed in the 1950s. V2I, however, transfers this technology—and therefore some of the costs—from the vehicle to the infrastructure. Recent research suggests that this smart infrastructure can more precisely identify the location of the vehicles than can vision-based artificial intelligence through LiDARs, which can only estimate location. In this respect, V2I could provide more comprehensive situational awareness, and do it at lower cost. Much of the technology in SRVs can be redundant, since each vehicle needs to carry a full suite of sensors in order to navigate its surroundings; in the best case, a V2I model would allow even a fairly "dumb" car to be autonomous through a much smarter infrastructure. The University of California, Berkeley's Institute of Transportation Studies reports that an intersection can be fully situationally aware for $25,000 to $100,000, less than the cost of LiDAR sensors on a single SRV.

The shift from SRV to V2I is analogous to the shift from personal computing to cloud computing. In the digital revolution, computers introduced by Apple and IBM placed much of the technology directly in the hands of consumers and businesses. However, as the technology advanced, the market moved toward a client-server model. In the last decade, businesses stopped hosting their own data and programs on local machines and moved to a "Software as a Service" (commonly known as SaaS) model, hosting all of their data in the cloud. This proved efficient, since a single business could specialize in operating and maintaining the data, while the end user could simply send instructions (essentially messages) through a mix of cheap mobile devices. V2I aims for the same approach, with cars acting as clients and the infrastructure as the server.

The core of V2I's technology is the roadside unit. This is anything ranging from a sensor pack along the road or on a street lamp, to a bevy of sensors and cameras at an intersection. For interstates, these roadside units can cost $6,000 to $8,000. These costs include Dedicated Short-Range Communications devices, cameras, and other equipment that communicates from the roadside units to the vehicles. Today, experimental research at universities like Texas A&M (where I teach) shows that these roadside units can navigate autonomous vehicles successfully. In fact, just two roadside units are sufficient for full navigation of an autonomous vehicle. While sensors on the vehicle certainly help, the real innovation in these experiments is that the roadside units act as the eyes and ears for the vehicle, sending signals to the navigation system and thereby controlling the car. As technology advances, not only will the cost of each roadside unit decrease, but they will also have a wider and longer range. Today, each roadside unit costs $5,000 and can navigate over 100 meters. That means a thousand-kilometer (or about 620-mile) route would cost $50 million in roadside infrastructure—still far less expensive than embedding all the technology inside every vehicle.

The main technical difference between V2I and SRVs is what engineers call the control problem. In both regimes, the vehicle must identify its own location, the location of other vehicles, and other persons, objects, or obstacles on the road. In SRVs, the cameras are all on the moving vehicle, and there are at least two vehicles involved in navigation: the subject vehicle that contains the sensors, and the other vehicles on the road that the sensors are trying to detect. Because the sensors themselves move with the vehicle, it becomes a challenge to precisely pin down location. All of the computations become simpler when the sensors sit in a fixed location, such as on the infrastructure.

SAFETY RISKS AND CYBERSECURITY

The scientific consensus on autonomous transportation is that, once the technology is mature and established, the long-term equilibrium fatality rate will be much lower than the equilibrium fatality rate of motor-vehicle accidents today. There were over 7 million motor-vehicle crashes on roads in the United States in 2016, and over 37,400 deaths, almost 105 years after the automobile went into mass production and six decades after the establishment of the interstate highway system. If sister technologies are any guide, machine-enabled autonomy will vastly reduce deaths caused by human driving.

But autonomous transportation does involve a different kind of risk: systemic risk. Human-driving accidents are dispersed among a large population. The accidents themselves are independent from each other, in a precise statistical sense. A distracted truck driver slamming on his breaks in Idaho has no relation to a minivan careening through the snow in Minnesota. But the technologies behind more automated transportation do not exist in such a scattered and decentralized way. Rather, there will be a handful of auto manufacturers deploying a finite number of algorithms to guide transportation. It is therefore entirely possible that the truck in Idaho and the minivan in Minnesota, if they were autonomous and following the same basic set of instructions from their computers, would behave in exactly the same way, and therefore crash for the same reason. Thus, the fact that technology can scale so quickly can be both a blessing and a curse: It has made technology companies some of the fastest growing in the history of mankind, but it can also amplify risk.

Beyond this, unforeseen questions emerge. Autonomous cars on the road will need to respond not only to human drivers, but also to other autonomous cars. Should Car A simply observe the behavior of Car B, as human drivers would, or should the cars announce their algorithms or protocols to each other in advance so they can better tailor their behavior to one another?

Driving algorithms must trade off safety and speed by deciding how often to take defensive measures, such as applying brakes. The tradeoff from acting conservatively is longer trip times, and possibly more uncomfortable journeys. Imagine two manufacturers that choose different points along this continuum. For example, suppose Tesla has an algorithm that prioritizes speed over safety, and Volvo does the reverse. In concrete terms, this could mean that an autonomous Volvo assigns an ambiguous black blob on its LiDAR image a 20% chance of being a deer and an 80% chance of being a human, while a Tesla does the reverse. If a Tesla and a Volvo approach each other at high speeds, the response of each car depends on what each car assumes about the other. This is a form of the game of chicken, where the final outcome depends on what the parties believe about each other.

For whatever reason, it seems now that the public is fine with the existing accident rate when humans drive, but there is massive outrage and finger-pointing when a single autonomous car crashes, as Tesla and Uber have discovered. The social tolerance for accident risk would likely be much lower in a world of autonomous driving than it is today. This would suggest a preference for V2I, which is likely the safer approach because it allows the infrastructure to collect more precise information on the location of all vehicles and therefore coordinate traffic flow more accurately and with more confidence — which ultimately should lead to fewer accidents.

A V2I paradigm would be more centralized than the current SRV system. This will have both benefits and costs with respect to safety and risks. SRV technology gives each vehicle a local understanding of its environment, as it uses its cameras and sensors to visually observe its immediate surroundings. It is possible, though unlikely, that any one vehicle could have a complete picture of all vehicles on the road. This would require large amounts of coordination among disparate parties, and some vehicles may lack either the ability or desire to communicate their location to other vehicles. By contrast, the V2I system by its very nature makes the infrastructure a central hub, collecting data from each vehicle and communicating back to some or all vehicles.

Of course, centralization of information involves its own distinct risks. The cybersecurity risk is much greater with V2I, as a lone hacker could dismantle infrastructure systems across multiple roads simultaneously with a single attack. For example, were all the infrastructure algorithms held within a single central server, then an attack on that server could lead to catastrophic consequences, especially if the infrastructure controlled navigation. To remedy this, the V2I system must adopt a form of extreme decentralization in its implementation. The federal Department of Transportation should certainly not serve as a single information hub for all state-level V2I systems. Rather, federal authorities could simply provide guidance and general protocols, while the states or even cities handle specific implementation. This is similar to how speed limits work today: Though they vary state to state, the differences do not markedly affect driving behavior, as drivers adjust as they cross state lines.

THE OPEN ROAD

Today, federal car-safety regulation is based on a thick book of rules called the Federal Motor Vehicle Safety Standards (FMVSS). These regulations, developed over decades, establish detailed performance requirements for every safety-related part of a car, including brakes, tires, headlights, mirrors, airbags, and more. Before a car can be introduced into the market, the manufacturer must certify that the vehicle meets all of the requirements in the current version of the FMVSS, or that the vehicle is as safe as or safer than the cars already on the road. A carmaker must certify that the brakes can stop the car within a certain number of feet, that airbags can deploy safely given passengers of various heights, that the tires can run for so many hours without overheating, and so forth.

Federal regulations do not say much about how companies develop and test cars before bringing them to market. The federal government regulates emissions, braking distances, and seatbelts, but regulations on the individual driver are left to the states. In the era of autonomous vehicles, however, the driver and the vehicle would be one and the same. In the past, development and testing were generally conducted on private test tracks where these activities posed no danger to the public. Car companies would then provide the government with documentation that the car met the standards in the FMVSS before putting it on the market. But that approach does not work for driverless cars; while companies are able to conduct some testing of driverless cars on closed courses, it is impossible to reproduce the full range of real-world situations in a private facility. Therefore, at some point, carmakers need to put self-driving cars on public roads for testing purposes before manufacturers can clearly demonstrate that they're safe.

The most effective way for policymakers to sustain the progress of research and get driverless cars on the road has been to enact legislation that carves out broad exemptions from the FMVSS for driverless cars. In some states, executive orders ensure permission for autonomous-vehicle testing and operation. The changes in law occurring across the nation vary from state to state, but advance the industry overall.

State regulations concerning autonomous-vehicle testing on public roads vary widely. States such as California and New York are strict, requiring companies to apply for a license to test vehicles on the road. In 2017, General Motors applied to begin the first sustained testing of autonomous vehicles in New York, but was still tangled in red tape as of this writing, and had yet to put cars on the road. At the other extreme, no company needs a license to operate autonomous vehicles in Florida; in Virginia, Ford and Virginia Tech's Transportation Institute have already collaborated on some tests on public roads. Between these extremes, states like Ohio allow autonomous-vehicle testing on public roads while requiring that the manufacturer operate the vehicle.

Predictably, states with the most restrictions on autonomous testing generally repel companies, which will naturally choose to work in states that do not require a license. New York goes so far as to require a police escort for all road tests at the cost of the manufacturer. This restriction makes New York a particularly unappealing place for tests because of the extra costs incurred and the less realistic testing environment caused by the escort. Because the industry is in its infancy, much of the industrial development around autonomous transportation will therefore move to states where testing is unrestricted. California has been the exception, since companies such as Apple, Waymo, and Cruise (a division of GM) have set up their autonomous-vehicle testing within the Golden State, where the rest of the company tends to be located, but it is the one exception to a strong general trend.

In most cases, less regulation leads to more testing. In Texas, autonomous-vehicle companies have virtually free reign to conduct testing in all conditions, anywhere in the state. In fact, a state law that took effect in September 2017 prohibits local governments from regulating self-driving vehicles in any fashion. Companies are free to conduct tests on both private and public roads, whether with platoons of trucks or single vehicles. They are also allowed to conduct tests without a driver present in the car, which is prohibited in several other states. This autonomous-vehicle-friendly legislation has encouraged companies like Drive.ai to test their driverless vehicles on public roads in Texas.

The current regulatory landscape for autonomous driving is reasonably friendly to the development of the technology. States are experimenting with different levels of regulation, just as they should. In the early phases of any industry, it is impossible to know what the right rules should be. The federal government is providing nonbinding guidance, an appropriate approach in this environment of uncertainty. Prematurely setting specific rules carries its own risks. Competition among the states will allow auto manufacturers to locate to the state that best suits their experimental programs. Given the high visibility of this industry, manufacturers face strong reputational incentives to avoid a race to the bottom; Uber virtually shut down its autonomous programs in Arizona and California after the fatal crash with a pedestrian that made headlines in 2018.

Critics of regulatory competition will always claim that this leads to a patchwork quilt of different standards, but variation in regulation has benefitted other markets, like the insurance industry, for decades. For now, it makes sense for the federal government to take a back seat and let competition and experimentation across states allow the best approach to regulation to emerge over time.

Of course, a natural question is how the government becomes involved with V2I. Indeed, this is the main reason that research and policy in this area stalled two decades ago, as the academic and think-tank community lost confidence in the public sector's ability to implement V2I at scale. V2I can seem at first glance like a full-employment project for the Department of Transportation. But while V2I does require government action, a careful plan for privatization can help create entirely new industries and firms. More than this, V2I presents a unique opportunity to meaningfully privatize one of the most important public goods that the government still retains: our roads.

THE ROAD TO PRIVATIZATION

Today, privatization of transportation infrastructure means the private *operation* of the infrastructure. The private-sector entities involved control pricing, access, and maintenance of roads. This often takes the form of a public-private partnership, which essentially is a contract between the public owner of the infrastructure (usually a state or municipality) and a private operator (often a for-profit company) that governs usage and fees over some fixed time horizon.

Smart infrastructure would fit easily into this contractual framework, as the private entities would simply update their technology and could even keep existing contracts in place. Ultimately, infrastructure-enabled autonomous driving would create the next generation of tolls and highway sensors. Rather than existing E-ZPass systems that send communication exclusively in one direction (from the vehicle to the infrastructure), future smart infrastructure will communicate in both directions. Prizes and initial funding already deployed for smart infrastructure in cities such as Phoenix could be used to test these technologies in real-world settings. Columbus, Ohio, was chosen as the winner of the Smart City Challenge for smart infrastructure in 2016, securing $40 million from the federal government to build infrastructure for connected cars. This shows that there is already some political appetite for infrastructure-enabled driving and suggests the possibility of future investments.

Private *ownership* of infrastructure is a more radical form of privatization, where the private sector not only operates but even owns the infrastructure itself. The benefit of private ownership is that it fully aligns incentives between the owner and the operator; rather than relying on a public-private partnership contract, the private entity bears the full costs and benefits of the infrastructure, and therefore has incentives to put the infrastructure to its highest-value use. No matter the complexity and detail of a public-private partnership contract, it will always be incomplete and subject to transaction costs, differences in risk preferences or time horizons, or any other of the many agency problems that will inevitably arise when two separate parties seek to contract. Private ownership has long precedent in the United States, the best example being the ownership of land and mineral rights. Indeed, a key explanation for the rise of the American energy industry is that the profit motive led individuals to extract value from the minerals they owned; without this incentive, the energy sector never would have developed in the ways it has. (This is not the case in China, for example, where the state retains all mineral rights.)

Smart roads are a possible candidate for a fully private approach, as it may be cheaper, logistically easier, and safer to build dedicated new roads for driverless cars, rather than mix humans and machines on existing roads. Special, dedicated routes on interstate highways (for long-haul trucking, for example) and high-density urban areas may be the early prototypes. These are natural-use cases for privatization, as private companies can develop, install, and manage this infrastructure. This will likely create entirely new industries and firms, as the toll operators of the future will work more closely with auto manufacturers and deploy a higher level of technology than the radio-based E-ZPass sensors of today.

Private ownership would allow infrastructure companies to freely contract with landowners to build smart roads, in the same way that oil companies contract with farmers to drill. Voluntary contracts between private parties avoid the messy legal battles of eminent domain. And it transfers the risk and liability of smart technology away from taxpayers and to the private sector, which can then use the price mechanism to charge end users appropriately. Competition between infrastructure companies can constrain the prices consumers face, and associations of infrastructure operators can help coordinate large projects. The internet is actually the best example of private parties contracting together to create a decentralized network, and there is no reason the same could not happen for smart-road networks.

Under a fully privatized V2I plan, the cost of the technology would be borne not by taxpayers, but rather by the consumers who use the infrastructure. Because congestion pricing allows the price mechanism to allocate resources to their highest value, those consumers who value the road the most will bear more of the cost of equipping it with the necessary technology. Congestion and variable pricing allow a more refined and granular way of allocating those costs based on their highest and best use, which is highly efficient from both an economic and road-use standpoint.

There are three general types of congestion-pricing systems in use today: a cordon area around a city center, a city-center toll ring, and corridor congestion pricing. For example, San Diego uses corridor congestion pricing, with variably priced lanes on I-15. In these lanes, prices rise and fall based on traffic conditions, similar to high-occupancy vehicle (HOV) and high-occupancy toll (HOT) lanes used all over the U.S. The introduction of variably priced lanes caused the number of carpools to increase by 50% between 1998 and 2006. In 2007, Stockholm established a permanent cordon around its city center with charges to leave or enter the zone. During the initial trial period, there was a 22% drop in vehicle trips and a 9% increase in bus ridership. An even earlier adopter of the cordon-area model, Singapore first implemented congestion pricing in 1975. The number of vehicles entering the relevant area has declined by 44%, with car entries declining by 73%. HOV (4+) use increased from 8% to 19% and bus ridership increased from 33% to 46%. When Singapore switched to electronic road pricing in 1998, weekday traffic in the downtown business area decreased by 24%.

Smart infrastructure could provide many opportunities for more advanced congestion pricing, which can unlock enormous efficiencies in transportation. Congestion pricing today mostly operates at the level of highway tolls, but if infrastructure engages with navigation in the autonomous vehicles of the future, then pricing mechanisms can be more rapid, granular, and tailored to traffic patterns, environmental conditions, time of day, usage, vehicle type, or even attributes of the driver, such as historical accident risk. As the market for virtual currency develops, micropayments between infrastructure and vehicles could truly put the roads to their most efficient use. The digital revolution has led machines to price user behavior dynamically, precisely, and rapidly throughout the consumer internet. There is no reason why such technology could not lie inside our infrastructure to route traffic, maximize efficient usage, and lower accident risk.

A ROAD TO THE FUTURE

Smart infrastructure for autonomous driving presents several advantages over our current SRV model of developing smart cars on dumb roads. Infrastructure can better coordinate traffic flow and more precisely track the locations of all vehicles on the road. And as technology matures, it will allow new and more refined forms of congestion pricing that can optimize road usage.

No advance of this magnitude is without risk, however. As a society, we need to think hard about designing a system that is sufficiently decentralized to address cybersecurity risks and other forms of hacking. But perhaps the biggest risk with smart infrastructure is the danger of political capture and misuse. Of course, proposing an infrastructure plan opens the door to a government takeover of the transportation system. But the government is in no position to develop these technologies itself, and recent successes across the country in privatizing high-tech infrastructure suggest that smart infrastructure could actually increase the level of privatization of our roads. This could mean, somewhat counterintuitively, that a smart-infrastructure plan could actually lead to less government intrusion in our lives than occurs today. If champions of market-based policy can rally around a clear and compelling privatization plan, then smart infrastructure could be the best chance we have for minimizing government control of our roads.

As we look toward the future of autonomous driving, the choice between the V2I and SRV paradigms need not be a stark binary. Realistically, the best solution will lie somewhere between these two extremes. But the research today is far too one-sided, with all investment by the private sector flowing into smart vehicles. Preliminary research shows that even a modicum of investment in infrastructure can lower overall costs, increase safety, and optimize road usage. It is time to start investing in smart infrastructure.