Introduction to the Transportation Internet

Howard Gutowitz
hag@mailbox.org

The moment when software replaces all human drivers will not the end of the "driverless car" story. Significant further evolution will still be possible. This document provides an overview of a possible next step. Vehicles will work together to provide a transportation flow collectively. People and things will be carried with this flow in a way analogous to the flow of information on the internet.

A yardstick for transportation systems

To appreciate the system we will introduce here, we need a yardstick by which to measure and compare transportation systems. For a transportation system based on roads, one such measure is the percent utilization of the intrinsic capacity of the road system. The intrinsic capacity is the number of passengers the roadway can carry in the limit of infinite technological advance.

What is the intrinsic capacity of a lane of highway?

Imagine sitting by the side of a highway, counting the passengers passing by in a given lane. The number of passengers, P, that will pass in a second, s, depends on the number of passengers per vehicle, and the number of vehicles, V, passing per second,

$$\frac{P}{s} = \frac{P}{V}\frac{V}{s}\ .$$

But vehicles per second is the product of the speed (here in feet, ft, per second) and the effective number of vehicles per foot, which is the reciprocal of the effective length, L_V, of the vehicle in feet, that is, in units of ft per vehicle.

$$\frac{V}{s} = \frac{ft}{s}\frac{V}{ft} = \frac{ft}{s}\frac{1}{L_V}\ .$$

By "effective length" we mean the length of the vehicle itself plus the following distance the vehicle must maintain at the given speed, that is, the total length of roadway used up by that vehicle. Similarly, the effective width of a vehicle, W_V, is the physical width of the vehicle plus the width of the buffer it needs laterally under the given technology, and similarly for vehicle height, H_V. The width of the road is denoted W_R, so the number of vehicles that can travel side by side is W_R/W_V. Similarly, if some technology allowed vehicles to travel on top of each other, the number of vehicles which could be stacked would be limited by the clearance under overpasses, H_R/H_V. Finally, the number of passengers per second is,

$$\frac{P}{s} = \frac{P}{V}\frac{ft}{s}\frac{1}{L_V}\frac{W_R}{W_V}\frac{H_R}{H_V}.$$

Every technology which proposes to decrease congestion seeks to change the value of one or more of these parameters. For instance, a ride-sharing scheme seeks to increase the number of passengers per vehicle, P/V.

In the limit of infinite improvement in technology, the size a vehicles carrying a single person could be reduced to no larger than the size of a seated human body. This size is approached by seating in present day super economy aircraft cabins. In the technology limit batteries have infinite power density, the shell of the vehicle need be no more than one atom thick, vehicles safely follow each other with no buffer, and travel side-by-side with no lateral safety buffer either. Roughly then, the absolute minimum single-person pod dimensions, W_V, L_V and H_V, are 2 ft, 3 ft, and 5 ft respectively. In the infinite technology limit, vehicles could arguably approach the speed of light, but for present purposes we will take a more modest limit value of 100 mph (about 150 ft/s). For a highway lane which is 12 ft wide, and has a clearance of 15ft, we find that

$$P/s \approx 150 * \frac{1}{3} * \frac{12}{2} * \frac{15}{5} = 1350 = C\,.$$

This is our estimate of the intrinsic capacity, C, of the highway lane.

How much of the intrinsic capacity are we currently using?

To assess current highway utilization, we assume that cars are traveling 60 mph (88 ft/s), carrying 1.2 passengers each, having a length of 20 ft and following each other at a distance of 120 ft. Cars do not travel side by side in a lane and do not stack on top of each other. Then,

$$\frac{P}{s} \approx 88 * 1.2 * \frac{1}{140} \approx 0.75 \approx \frac{C}{2000}.$$

We find that the intrinsic capacity is about 3 orders of magnitude greater than the capacity we presently use. This entails that there already exists enough infrastructure capacity, if only it could be completely utilized, to overcome any induced demand. Assuming that there are not many more empty vehicles in circulation than full ones, induced demand cannot grow to be more than one order magnitude greater than current demand, since there are only 24 hours in a day, and we already spending an hour or so of each day being transported.

Realistic approach to the technological limit

Clearly no pod that we can build at present or within the next 30 years will get close to the technological limit. Each pod has to have space not only for the passenger, but also all technology needed simply for the pod to function as a transportation unit: vehicle body, motor, battery, sensors, steering mechanisms, brakes, electronics etc., as well as any equipment needed for passenger comfort (e.g. heating, ventilation, and cooling) and safety, (e.g. air bags). Each of these items has a practical limit to miniaturization. Using any known technology, a battery would have to be quite large to be able to power a long-range journey at 100 mph. Leaving aside unforeseeable advances in physics, chemistry and computer science, progress toward the theoretical limit might be accomplished by offloading functions to the infrastructure, e.g. supplying power to the pods via electrified rails. Such solutions as too expensive and inflexible for all but the most consistently high-density circumstances. Here we are looking for solutions which entail no required modification to physical infrastructure, and are efficient over the entire range from lowest to highest density.

Collective-personal transportation

The solution is for vehicles to provide infrastructure for other vehicles. A pod would need little onboard power if its range were restricted to say a mile, and its speed restricted to say under 25 miles per hour. The slow speed would entail less impact protection required. The pod would need little HVAC if it only spent a restricted amount of time, say up to 15 minutes, outside of a climate-controlled environment. These restrictions would allow L_V, W_V, and H_V to approach their theoretical limits since vehicle size could decrease, even without significant advances in technology. However, to travel further, faster, and for a longer time, the reduced pod would need to be transported inside another vehicle with beefier specifications, a larger container vehicle.

A 4-passenger car has 4 tires, as does 40-passenger bus, albeit slightly larger ones. A 40-passenger bus needs a bigger engine than a 4-passenger car, but not 10 times bigger. By sharing transportation machinery, as well as HVAC equipment, sensors, and other technology needed for automated transportation, lightweight, slow-speed, short-range pods could leverage the economies of scale of larger vehicles. Inside the container, the buffer zones between the pods could be reduced to near 0, dramatically reducing L_V and W_V. This is the same boost in intrinsic capacity utilization which motivates public transit advocates. Moreover, a container vehicle could have two or three pod-carrying levels, reducing H_V for a further boost. **This is not the commonplace argument for favoring automated buses over automated cars.** Rather, it is an argument for a system optimally integrating features of automated cars, buses and even larger vehicles into a wholly new system, a packet-switched system, which nests automated vehicles inside of other automated vehicles. Automated buses, alone, have the same drawbacks as regular buses: they have to constantly stop to pick up and drop off passengers, in the face of fluctuating demand, they are hard to keep fully loaded so they operate efficiently, they are too big for narrow or highly curved roadways, or any low-demand applications. However, when combined with automated pods, all these drawbacks can be eliminated.

The analogy to the information internet

In the analogy to the information internet, the pods are the packets, and the container vehicles are the routers. While routers in the information internet are typically stationary, and exchange packets electronically, the transportation routers exchange packets physically, and preferably while moving. The transportation routers provide a sorting and distribution function like information routers do, but also provide physical motion toward the packet's destination. This is the step at which the information internet analogy can no longer guide our understanding. It is a new property with many surprising consequences. Still, the analogy leads us to expect that the transportation internet will share with the information internet the same improvements over their circuit-switched analogues in robustness, throughput, and ease of deployment. The container vehicles are discrete packages of a layer of transportation and automation support between the static infrastructure and the pods. That is, rather than build static support for pods into the roadway itself, support is encapsulated in a peer-to-peer system of mobile units, where the characteristics of each unit are optimized against the physical constraints of the local section of roadway in which it moves, abstracting away those static features of the infrastructure, insulating the pods from them.

Docking

As transportation internet routers route physical packets, they need to physically meet to exchange packets between them, at least to come within a range over which pods can move on their own, from one container to another. The distance between containers could be zero, as an automated vehicle can dock, at any speed, with another automated vehicle. That is, it can drive right up to another vehicle, touch it, lock into it, so that it can open passageway for pods to transfer between the two vehicles, while both container vehicles are moving at speed. Automated vehicles can do this safely and reliably, though human drivers cannot. Without entering into a detailed discussion of this assertion, some remarks: 1) To reliably dock while moving requires more than just skill at piloting a vehicle. It demands complete and instantaneous sharing of knowledge of what the other vehicle is seeing, doing, and intending. A merging of minds which human intelligence is incapable of but artificial intelligence can do readily. 2) Vehicles have been reliably docking in outer space for decades, moving at miles per second, not feet per second. They do this by docking in phases, where the technology of each phase is designed to allow the next phase to be executed reliably. 3) Few would dispute the automated vehicles can safely dock when stopped or moving at slow speeds, or that the safe speed limit would tend to increase with technological advance. Nothing here requires that vehicles maintain cruising speed while docking. The system would just work better the closer they get to being able to do that.

The logistics component of transportation

When thinking only of circuit-switched individual transportation in which one vehicle takes one passenger from trip origin to destination, logistics is trivial and easy to ignore. In the case of public transit, where vehicles ply fixed routes with fixed stops, the logistics function is built in and static in the architecture of the routes and stops. But in a packet-switched network logistics is dynamic. The is no fixed association between passenger and vehicle. There are no routes or stops. Each packet traverses many vehicles on its path and the paths shift constantly. Transportation potential in the form of container vehicles is dynamically recruited and retired as a function of fluctuating demand.

In comparing transportation systems, logistic effectiveness is just as important to measure as intrinsic capacity utilization, and has potentially many dimensions. An efficient transportation system uses no more of the intrinsic capacity of the infrastructure than necessary to transport a person (or unit of cargo of equivalent size). So even if we could use the full intrinsic capacity of a roadway we shouldn't, unless such is required to meet current demand.

For both people and things, the time. energy and other costs required to get the payload from point A to point B is important. So is the reliability and predictability, the quality of service. People might also value services supplied during a journey, such food service or legroom. Cargo might also need services during transport, such as refrigeration.

Supplying these services is a demand which further complicates transportation logistics. Still further complications are seen from the point of view of the system as a whole, and society as a whole. The system might need to trade a given individual's cost, speed, and comfort against the throughput and efficiency of the system as a whole. Society might need to balance the system's optimization criteria as well as the individual's optimization criteria against its own desires, such as prioritizing military usage of transportation capacity.

Failures in communication between public transit advocates and individual mobility advocates can often be traced to differing measures of and appreciation of the logistics component of transportation. E.g. the costs and time in loading and unloading a bus along a route can easily overwhelm the transportation efficiency gain from running a fully loaded bus. Similarly, carpool advocates tend to ignore the inefficiencies of collecting riders together for a trip and then distributing them to their individual destinations. In a packet-based system, by contrast, motion is inextricably linked to the logistical functions of gathering and scattering individuals according to their position, destination, and other criteria. The packet-based system gathers together pods with sufficiently similar characteristics for transportation in more efficient larger vehicles, and scatters them when these characteristics sufficiently diverge that smaller vehicles should carry them, down to the level of a single pod which can travel short distances on its own. Container vehicles of the same or different size class can route pods between them and within their interiors. Logistics and motion can not be usefully discussed separately.

A typical journey

Consider a typical journey in a basic transportation internet. This system has just three classes of vehicles, pods, feeders, and highway vehicles. The pods can be small enough to fit through residential doorways and have a range sufficient to get them anywhere within the neighborhood. The passenger mounts her pod in her home, and communicates her destination to the global router. The router has all relevant information needed to route the journey effectively. It knows the situation and plans of all other vehicles which the given pod might interact with during the journey, and sees everything they can see. It knows everything relevant about the physical environment though which the journey will pass. Thus it can optimize the journey while simultaneously optimizing the journeys of all interacting vehicles. It decides, in particular, which feeder the pod will initially board and will orchestrate the boarding, preferably while both pod and feeder are moving—so there is no dwell time. The router selects the feeder which is best given its nearest to the pod, the destination of the other pods it already contains, the next transportation tasks both pods and feeder will need to perform, their current load, and any other relevant factors. Once the given pod has boarded, the containing feeder might dock with another feeder to exchange pods, indeed these feeder-to-feeder docking might happen in a chain as pods are simultaneously being moved toward their destination and collated according to the similarities in their paths. Eventually, a feeder takes the pod out to the highway, potentially along with other pods with similar intended paths. Very large vehicles travel on the highway, perhaps as wide as several legacy lanes of traffic, perhaps as high as the lowest overpass on the segment of highway on which they live and so containing several pod-carrying levels. As the system is designed to keep these vehicles near their full capacity, there may be few available highway vehicles in low-demand situations. There may also be gaps with no highway vehicles as vehicles tend to travel in platoons to increase opportunities for beneficial sorting of pods between them, creating a lumpy distribution. In the low-demand case the feeder travels the highway until it encounters the first highway vehicle, to which it docks and transfers its pods. In a high-demand case, the controller will have several options and will chose the highway vehicle so as to simultaneously optimize any number of constraints including the loading of the highway vehicle, the pod contents of the vehicle and their intended paths, the further tasks of the feeder, the chain of dockings between highway vehicles currently planned. As the pod's destination approaches, the process reverses. The pod is sorted into a feeder, then the feeder may dock with and exchange pods with another feeder and so on in a sequence of feeders, until the given pod is discharged into a local roadway close enough to its final destination that it can proceed there under its own power.

Obviously, this brief introduction leaves aside many details. Contact the author for a longer document which delves into them. Less obviously, packet-switched networks have potential impacts on disaster response, the sharing economy, and other social issues. The same longer document goes into these impacts as well.