12/15/2014

The Trouble with Cars

The State of Personal Transit in America

Stephen Hamilton

The Trouble with Cars

The State of Personal Transit in America

In the words of John Lennon, “there are no problems, only solutions”. This report attempts to take an objective look at the state of personal transit in America, which in the opinion of this author is pretty much a mess. But the purpose of this examination is to truly understand the strengths and weaknesses of this transit system, as well as how we arrived at this “solution”, so that a better solution can be designed and proposed.

It is important to recognize that for the most part America is a nation that operates with a short term view. It is a collection of systems that each was added incrementally, with each increment being economically justified in isolation. There is an inherent robustness in such evolutionary systems. But they are rarely, if ever, efficient. America is still economically vibrant, but is no longer in the emerging and growing state it once was. In the more mature America of the future, efficiency is paramount if we are to compete on the global stage while maintaining our high standard of living. So this examination uses efficiency as its primary metric for judgment.

In a previous white paper a taxonomy was offered which organized American transit systems in two dimensions. The first dimension separates systems by the type of entity operating the system – private or public. The second dimension separates systems by how the vehicles are used – exclusively used and following a direct route from origin to destination, or simultaneously shared by multiple riders on multiple differing trips, and therefore following a predefined route along a commonly used path somewhat near the origins and/or destinations of these trips. This second dimension distinguishes “personal” transit from “mass” transit. Mass transit solutions are cost effective only along high traffic routes commonly found in high density neighborhood. This means they are cost effective in urban areas, but not in suburban areas. The paper also points out that since metropolitan population is about 85% suburban and 15% urban, the demand for personal transit far exceeds the demand for mass transit.

To date the family sized (personal) automobile is the only transportation technology available for use in personal transit. It is dominantly applied as privately owned and owner-operated transportation, but also sees some use as taxis and rental cars. Private ownership of a vehicle is expensive, and costs per mile are higher than for public mass transit. But for many working Americans it is still an economically viable solution, even if an expensive one. And it is effectively the only transit option for a very large portion of the transit demand.

Transit districts must provide for both publically operated mass transit, and privately owned and operated personal transit. A significant investment is made by municipalities and transit districts in providing for the use of cars. The purpose of this document is to explore the effectiveness of this provision – the state of “private-personal” transit in America.

Seven Hands Clapping

Private-personal transit in America - that is our automotive society - is organized as a very strange collaboration. Mass transit in America, both publically and privately operated, is fundamentally a 2 party system – involving the service provider, and the consumer/rider. The service provider provides all aspects of the service – the vehicles, the roadbed on which the vehicles travel, the energy required to operate the vehicles, the vehicle operators, governance of access to the system, regulation of the system, and responsibility/liability for all actions of the system. The rider is just the consumer. The automotive transit system however, is organized with each of these aspects of the service provided by a different party. The consumer is also the provider of the vehicle, although he acquires it from a third party. The transit district provides the roadbed. The consumer provides the energy to operate the vehicles, although this is also acquired from another third party. The consumer is also the vehicle operator. A separate governmental entity – usually a state government – governs access to the system by licensing operators and registering vehicles. Responsibility for regulation of the system is split, with the transit district responsible for street signs and traffic lights, and various other government entities responsible for enforcement – policing and apportioning of liability. The consumer is also responsible for liability to others for his actions in the system, although he likely contracts that responsibility to yet another third party.

This division of responsibilities for aspects of the automotive transport service is far better at optimizing the profit of commercial entities than it is at optimizing the efficiency of transport. The result is an extremely expensive and poorly performing transit system. The vehicle manufacturers have economic incentives to build more expensive cars – higher revenues. This results in bigger, heavier, more complicated, and less fuel efficient vehicles. Neither the energy providers, nor the insurers resist this trend, since it also results in higher revenues for them. The economic interests (higher pricing) of the manufacturers, energy providers, and insurers are also served by the breaking of the market into millions of individual consumers, who have less bargaining power than would a hundred or so transit district consumers. Trial lawyers have economic incentive to push for higher damage awards. This drives up insurance revenues. Insurers have incentive to lower damages resulting from accidents, not by redesigning the system to make it safer, but by pushing manufacturers to deploy more safety mitigation. But this safety mitigation pushes up vehicle costs and potentially vehicle weight and therefore energy costs. Transit districts have no manner to control vehicle size, weight, or performance. Therefore they have no manner to improve the efficiency of road use. Thus the only available response to increasing demand is to build more roads, adding to the tax burden passed on to the consumer. The consumer's readiness to accept that added burden is limited, so keeping pace with the demand is not possible. The result is increased congestion of existing roads, adding to the commute time paid by the consumer (degraded quality of service). Courts have little way within our governing philosophy to resist lawyer incentives to push for higher awards, nor to resist insurer incentives to contest them. So the litigation burden increases with these costs passed on to the consumer as taxes. And similarly, tax payer resistance precludes keeping pace so the consumer experiences longer resolution times for claims (degraded quality of service). The system relies upon amateur operators. Not all are equally qualified. Some are not even minimally qualified. Yet to deny these operators access to the system is to deny them an important part of their public commerce rights. So courts are reluctant and slow to do so. This tends to keep un-qualified operators in the system longer than would a more aggressive approach. Thus the incentives produce a less safe system. Manufacturer's incentive for more expensive cars pushes them to add more features. Many of these features create operator distraction, thus lowering safety. Yet transit districts have no jurisdiction over vehicle specs, and courts are again reluctant to restrict innovation of private entities.

The American automotive transit system is organized to maximize revenues and profits of the vehicle manufacturers, the energy providers, the insurers, and the lawyers. It is organized to shift burden to and limit the control of the transit districts, police, courts, and executive government agencies (DMV, etc). All costs are ultimately borne by the transit consumer. The result is an extremely expensive and poorly performing transit system, and one that will likely grow more so over time.

Lets start with the vehicles themselves. It must be stated that the extreme flexibility of the automotive transit system is a major strength. The size, utility, and type of vehicles that can operate varies broadly – from motorcycles, to small efficient economy cars like the Fiat 500 or Smart car, to family sized vehicles such as sedans and SUVs, to larger family vehicles like minivans, to small cargo vehicles like vans and pickup trucks, to large cargo and utility trucks, to recreational vehicles, to busses, to tractor-trailer vehicles. However, this flexibility brings with it a cost in efficiency, on multiple fronts.

Any engineer will tell you that tolerance for variation drives up costs. The system must be designed to handle the worst case (of any variation), while almost every individual use will be better than worse case, and so could have been satisfactorily supported with a cheaper, smaller, simpler system.

Fuel Efficiency

In 2014 the average gas mileage of the public transportation fleet in America was 15.65 miles per gallon. According to the US Transportation Department there were 2000.1 billion vehicle miles driven (www.fhwa.dot.gov/policyinformation/travel_monitoring/14augtvt/). And according to petroleum industry watchers 350 million gallons of gasoline, on average, were sold each day (www.advisorperspectives.com/dshort/updates/Gasoline-Sales.php). So the best estimate of the actual fuel use by the actual fleet is produced as (2x10^12)/365/(350x10^6) = 15.65 miles/gallon.

The US Transportation Department has established mileage standards under the CAFE (Corporate Average Fuel Economy) laws. These laws establish mileage goals for new vehicle fleets built by manufacturers of cars sold in the US, and penalties to be paid by those manufacturers for shortfalls against those standards. So the department monitors and reports on major manufacturers each year. In 2014 the aggregate (across all manufacturers) new car fleet mileage achieved a new record high of 23.6 miles/gallon (as reported by the Washington Post).

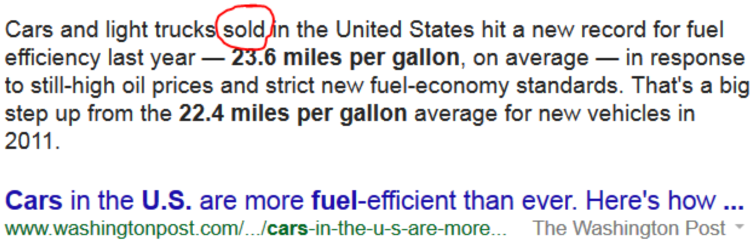

Cars and light trucks sold in the United States hit a new record for fuel efficiency last year — **23.6 miles per gallon**, on average — in response to still-high oil prices and strict new fuel-economy standards. That's a big step up from the **22.4 miles per gallon** average for new vehicles in 2011.

Cars in the **U.S.** are more **fuel**-efficient than ever. Here's how ...
www.washingtonpost.com/.../**cars**-in-the-u-s-are-more... The Washington Post

The CAFÉ process has tended to improve the actual average mileage over time. As older vehicles are retired and replaced with newer vehicles, those newer vehicles have better mileage. So the fleet average improves. But the fleet calculations made for CAFÉ do not reflect real mileage. In recent years the industry has negotiated modifications to the calculations (“footprint” based standards) that move the numbers farther away from people’s real driving experience. The targets are no longer based upon miles per gallon goals, but rather upon miles per gallon per square foot of footprint (area covered by the wheels). So a larger vehicle like a truck can meet the standard with a lower measured mileage than can a smaller vehicle like a compact car. This is significant because the profile of vehicles sold continues to favor larger vehicles. In 2014 3 of the top 4 vehicles sold were trucks (www.autobytel.com/auto-news/features/top-selling-cars-of-2014-1).

- 1) Ford F-150
- 2) Chevrolet Silverado
- 3) Toyota Camry
- 4) Dodge Ram 1500

Number 1 - Ford F-150

Number 2 - Chevrolet Silverado

Number 3 - Toyota Camry

Number 4 - Ram 1500

So the achieved 23.6 mpg must be "footprint"-adjusted downward to get the real fleet average mileage (harmonic mean). As shown in the chart below (wikipedia.org/wiki/Corporate_Average_Fuel_Economy) what is reported in 2014 as achieving 37.5 mpg, may actually be accomplished with 28.5 mpg for a large footprint vehicle (a 24% downward adjustment). So with sales numbers so heavily tilted towards larger vehicles, the reported 23.6 mpg may point to a new car fleet actually achieving something close to 20% less, perhaps in the 19 mpg range. This would still produce an upward pressure on the total operating fleet – including new and old vehicles – to increase above the 15.65 mpg over time. But that upward pressure may be far smaller than the 23.6 reported number leads us to believe. Lets just call it "effective marketing" by the automobile manufacturers.

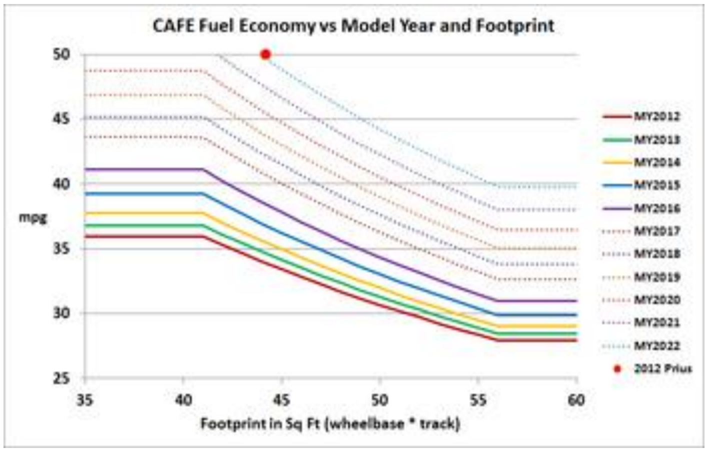

But more importantly, the CAFÉ calculations do not take into account the miles driven. The CAFÉ calculation is a simple production weighted harmonic mean. This assumes that all new vehicles will be driven the same number of miles. Let us consider a trivial example of a new car fleet consisting of only 2 cars – a Ford F-150 reported at 17 mpg, and a Toyota Prius reported at 50 mpg. The CAFÉ harmonic mean calculation (last column) reports a fleet average of 25.37 mpg. But this calculation does not account for how many miles each car is driven. The three examples given vary only in the miles driven for each vehicle. Note that in example 1 they are driven the same number of miles, and the actual fleet mileage (next to last column) agrees with the CAFÉ calculation. But examples 2 and 3 show un-even miles driven for the vehicles. In those examples the actual fleet experience varies significantly from the CAFÉ calculations.

	mileage Car 1 measured mpg	mileage Car 2 measured mpg	miles driven Car 1 miles	miles driven Car 2 miles	gallons used Car 1 gallons	gallons used Car 2 gallons	gallons used total gallons	miles driven total miles	mileage fleet actual mpg	mileage fleet CAFE mpg
symbol	m1	m2	d1	d2	g1	g2	G	D	Ma	Mc
calculation					d1/m1	d2/m2	g1 + g2	d1 + d2	D/G	2/((1/m1) + (1/m2))
Example 1	17	50	10,000	10,000	588.24	200	788.24	20000	**25.37**	25.37
Example 2	17	50	20,000	10,000	1176.47	200	1376.47	30000	**21.79**	25.37
Example 3	17	50	10,000	20,000	588.24	400	988.24	30000	**30.36**	25.37

We see that it is at least possible the CAFÉ reporting under reports real experience. It is plausible that purchasers who expect to drive more miles are more concerned with vehicle mileage and so tend to buy higher mileage vehicles. This would tend to shape the outcome more the direction of example 3, where the actual mileage is higher than the CAFÉ report. Of course without knowledge of miles driven for individual cars sold we simply cannot know what the actual fleet experience is.

We should draw 2 conclusions from the CAFÉ analysis above.

1. First, the CAFÉ reports should **not** be used as meaningful measures of even the new vehicle fleet, and certainly not of the whole operating vehicle fleet. So the calculation approach used at the start of this section, which produced the 15.65 mpg estimate, is a more representative approach (although not without fault either).
2. Second, even the CAFÉ reporting demonstrates the existence, in large volume deployment, of high mileage vehicles. So we must conclude that automotive technology today can easily achieve better results than the 15.65 mpg we are actually getting. The 23.6 mpg new fleet is 1.51 times higher (even 1.21 times higher after 20% reduction for footprint adjustment). But the new fleet report includes specific cases such as diesels at 42 mpg (2.68 times higher) and plug-in electric hybrids at 99 mpg (6.33 times higher). So the 15.65 mpg must be seen as extremely inefficient relative to what we know is both achievable and economically viable.

So what prevents our automotive fleet from becoming more fuel efficient? Lack of centralized and coordinated control and regulation. A moving vehicle consumes fuel performing 3 tasks: moving air out of the way as it moves; flexing the tires where they contact the road as they roll; and accelerating the weight of the vehicle to increase speed. A vehicle also consumes fuel on a fourth task, whether moving or not: generating electricity to operate comfort devices – most significantly the environmental control system (AC/heater). The energy used moving air is a function of the size and shape of the vehicle, which determine its aerodynamic drag; and a function of the speed the vehicle travels. The aerodynamic drag is determined and controlled by the vehicle manufacturer, and is completely unregulated. The speed of operation is determined by the driver, and is limited by the posted speed limits determined by the transit authority (highway department, local government, etc). The fuel consumed rolling the tires is determined by the tire properties, the tire inflation, and the weight of the vehicle. Tire properties are determined by the tire manufacturers, and are very loosely regulated. Tire inflation is determined by the vehicle operator, and is completely unregulated. Vehicle weight is determined largely by the vehicle manufacturer, which is very loosely regulated; and to a lesser extent by the loading allowed by the vehicle operator, which is completely unregulated. The fuel consumed accelerating the vehicle is determined by the weight of the vehicle, the rate of acceleration, and the frequency of acceleration. Weight has already been discussed. The rate of acceleration is controlled by the vehicle operator, and is completely unregulated. The frequency of acceleration is determined by the traffic light system, which is controlled by the transit authority. It is unclear if, or to what extent, fuel economy is considered in the operation of traffic lights.

Vehicle manufacturers, tire manufacturers, automotive regulation/licensing authorities, traffic control system operators, and vehicle operators all determine fuel efficiency. Yet no one is responsible for it, and they all have divergent and competing motivations, with fuel economy well down the priority list. Each must make decisions that accommodate the decisions of the others, with little or no ability to restrict those others' decisions.

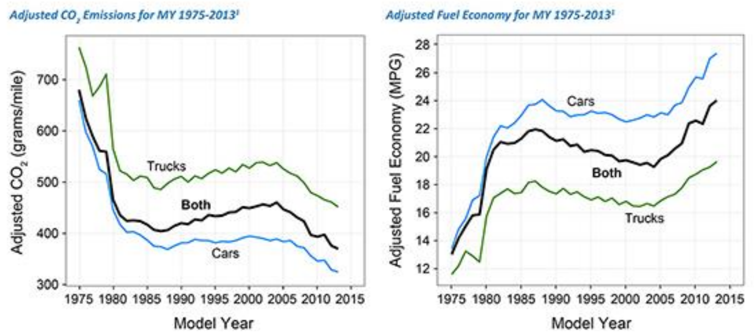
Adjusted CO_2 Emissions for MY 1975-2013[1]

Adjusted Fuel Economy for MY 1975-2013[1]

[1] *Adjusted CO_2 and fuel economy values reflect real world estimates and are not comparable to automaker standards compliance levels. Adjusted CO_2 values are, on average, about 25% higher than the unadjusted laboratory CO_2 values that form the starting point for GHG standards compliance, and adjusted fuel economy values are about 20% lower, on average, than unadjusted fuel economy values.*

History shows the CAFÉ process is an insufficient coordinating force to counteract the competing interests. In its first 7 years the process was able to improve mileage from 13 mpg to 21 mpg – a 61% improvement (8.7% per year). But in the 32 years since, the process was only able to improve mileage from 21 mpg to 24 mpg – a far smaller 14% improvement over a much longer time (0.4% per year).

As shown above, vehicle weight figures prominently in this issue, and it is illustrative of the barriers to progress introduced by the division of responsibilities. In America the average vehicle weighs approximately 3600 lbs. Lighter weight vehicles are available (Fiat 500 ~ 2300 lbs; SmartCar ~ 1800 lbs; Mini Cooper ~ 2600 lbs).   There is a perception that the lighter cars are not as safe as the heavier cars. This perception is false. Consider an INDY race car, which weighs approximately 1500 lbs.   The maximum speed of an INDY race car is about 220 mph; or a bit more than 3 times the 70 mph maximum speed of the average passenger car. Each race we see a quarter to a third of the field involved in high speed crashes. Yet the injury and fatality rates are extremely low. Contrast that with the injury rates from passenger vehicles. So lightweight vehicles can be extremely safe.

The primary reason the race injury rates are so low is that all of the race vehicles are approximately the same weight, and usually travelling at approximately the same speed. Therefore the momentum difference – the energy which must be absorbed in a crash is relatively low. It is this momentum difference that most directly relates to the safety of a vehicle. Since our public transit system very loosely regulates vehicle weight, it permits the 1800 lb SmartCar to use the same roadway as the 7400 lbs Ford F-350. Therefore the lightweight car is at significantly higher risk if the two collide.

The vehicle manufacturer wants to make more money by selling a bigger, heavier car. Some consumers have need for that vehicle size, so they buy them and operate them. Other consumers must then face increased safety risks posed by the vehicle weight differential. They may also face increased insurance costs as the insurers protect themselves against higher claims on lighter weight vehicles. Therefore these other consumers then have incentives to move to bigger cars, and many of them do. On average, vehicles in the operating fleet get bigger and less fuel efficient. The result also has other negative side effects. The bigger vehicles demand accommodation with wider lanes on roads, so construction costs and maintenance costs borne by governments increase, and the increases are passed on to consumers. Meanwhile the larger cars take up more road space, so fewer fit in the same area. This increases congestion in city cores, which increases the frequency of stops by traffic lights. Therefore the number of accelerations along a route increases. Travel times also increase.

In short, since nothing and no one entity manages fuel consumption, all other concerns are optimized first at the expense of fuel consumption. The broadly unregulated vehicle and operator specifications which give our system great flexibility, also extract a very high price in dollars, safety, transit times, and fuel consumption. While technology offers potential for fuel economy in excess of 200 mpg, we already have in the current operating fleet proof of mileage in excess of 100 mpg, which is nearly 7 times greater than the 15.65 mpg fleet average we experience. How could we call the current solution anything other than terribly in-efficient with respect to fuel economy?

Land Use

A major reason personal automobiles are used for personal transit is because they are so convenient. But they are only convenient because we allocate so much of our land in order to make them so. In addition to the land actually required for the roads themselves, we dedicate significant chunks of land to park/store the vehicles when they are not in use. Collectively the roads and parking lots consume a surprisingly large portion of urban and suburban land.

For most personal vehicles the duty cycle – portion of time they are in actual use – is extremely low. A typical example would be a sub-urban commuter who commutes to and from work 5 days per week, and has a 45 minute commute each way. That is 1.5 hours of vehicle operation per 24 hour day. On the weekend, chores and leisure activities may require a larger number of shorter trips. But the total operation time per day is likely nearly the same, or less. Use of 1.5 hours per 24 hour day is a duty cycle of 6.25%. Very nearly the same estimate can be achieved a different way. The average driver logs 12,000 miles per year. If his average operating speed is 20 miles per hour, covering that distance requires 600 hours of driving. There are 8760 hours per year, so that is a 6.84% duty cycle.

While a vehicle is in operation it requires the roadway. Dense roadways are more convenient than sparse roadways, since they offer a better likelihood of being able to drive closer to the ultimate destination. This should result in quicker commutes and less walking distance. But denser roadways require more land be given over for this dedicated use. And this land will be used by each commuter only around 6.5% of the time.

Urban zoning will determine how wide roadway lanes must be, and how much on street parking area is required. Typical modern lane widths are 12 feet. With dual lanes and some allocation for parking and turn lanes, roadway minimum widths of 40 feet are typical. If the urban area is zoned for 12 blocks per mile (which is common), the minimum land allocation for roadways is 18%. Where traffic management requires more lanes, the allocation must go up.

Automotive traffic creates a risk to pedestrians. It is not practical to eliminate pedestrian traffic from cities. Therefore means are needed to mitigate this risk. The most common solution is “grade separation”, where sidewalks are constructed for pedestrian traffic, so that the roadways can be exclusively used by automobiles (or shared with bicycles). Generally sidewalks and roadways run in parallel. This makes precisely separating the land allocations for sidewalks from the allocations for automotive roadways difficult.

Since the vehicles are “personal”, they must be parked when not in use. In order for use of personal vehicles to be convenient, the distance between the parked vehicle and the eventual operator must be small. Since this distance will likely be traversed by walking, minimizing this distance is essential to preserving the “immediacy” factor for use of the vehicle (how quickly is transit available), which is a big part of its convenience. This means that parking space is required at each location where the vehicle will be taken. Since each vehicle that goes somewhere also came from somewhere, there will always be at least 2 such locations (e.g. home and work). At most one of them can be used at a time.

So for the 93.5% of the time that the vehicle is parked, it occupies exactly 1 of n possible parking locations, with the other n-1 un-occupied by the vehicle. This could result in a huge amount of unused land. So generally the parking space is a combination of exclusively allocated parking (usually at home) and shared allocation parking (usually in public). The degree of sharing therefore determines how much land is needed, and is a statistical trade-off with the immediacy convenience of personal vehicles in the neighborhood where the sharing is done. Commercial and municipal entities constantly struggle with how to balance the convenience with the cost of parking. So there is considerable variation across businesses and municipalities in how that trade-off is optimized. But in all cases, the land investment for parking is substantial.

Original work: photoLith

Houston, TX
Surface parking: 21.3%
Garage parking: 3.7%
Street area (including sidewalks): 39.7%
Total area for rights-of-way plus off-street parking: 64.7%
Park space: 2.6% (1.1% exluding Discovery Green)

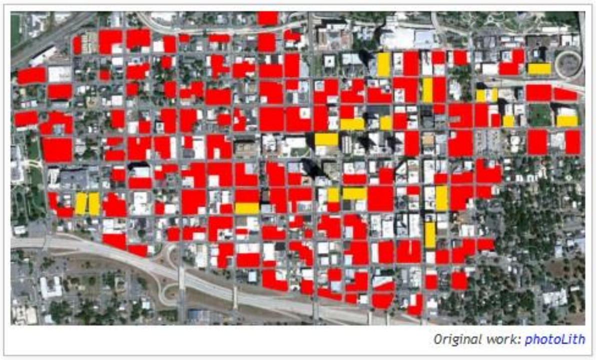

Original work: photoLith

Little Rock, AR
Surface parking: 26.5%
Garage parking: 2.7%
Street area (including sidewalks): 32.0%
Total area for rights-of-way plus off-street parking: 61.2%
Park space: 0.0%

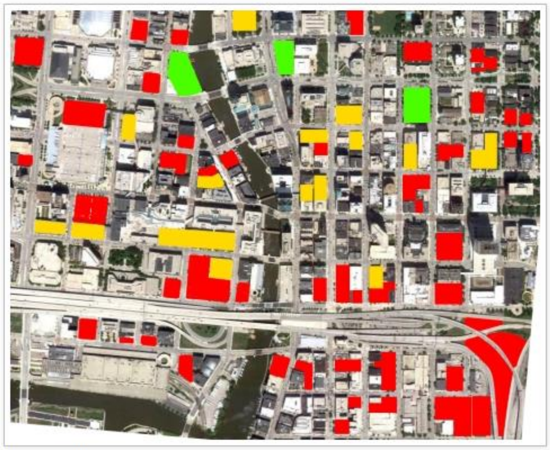

Milwaukee, WI
Surface parking: 11.5%
Garage parking: 3.8%
Street area (including sidewalks): 38.8%
Total area for rights-of-way plus off-street parking: 54.1%
Park space: 1.5%

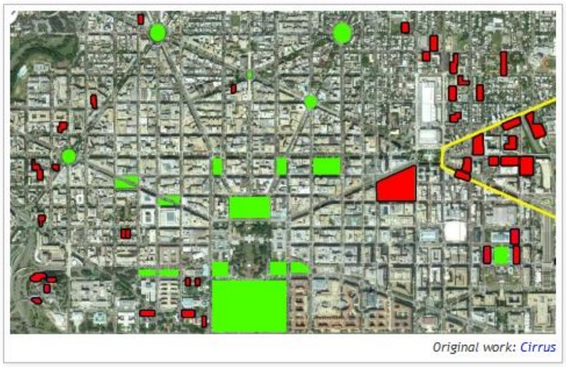

Original work: Cirrus

Washington, DC
Surface parking: 1.1%*
Garage parking: 0.0%
Street area (including sidewalks): 43.3%
Total area for rights-of-way plus off-street parking: 44.4%
Park space: 2.53% (5.00% including Ellipse)
**Much of this is the large lot to the east scheduled for redevelopment as CityCenter*

The images above are from //oldurbanist.blogspot.com/2011/12/we-are-25-looking-at-street-area.html (note reference to //forum.skyscraperpage.com/showthread.php?t=191539 where there are more). A number of individuals used satelight imagery (google earth) of various American cities, and overlaid color coded annotations of the land use. The oldurbanist blogger was particularly interested in contrasting land use for parks vs land use for automotive purposes. Red shows surface parking, yellow shows above-ground parking garages, and green shows park space. Three of the four cities (Washington, DC is the exception) have parking allocations between 15.3% and 29.2%. The examples on the forum are much more in keeping with these 3.

So the common case today is for cities to spend approximately 45% of their land area supporting the use of personal automobiles for personal transit – 20% or more for the roads themselves, and 25% for parking. Individual businesses often allocate more, and only need to provide for parking. Even homes must allocate significant land area for parking.

Much of this land is poorly used. A residential garage is likely empty 50 hours per week (29.7%). A business parking lot will only rarely be full. And many roads will be used very few times per day, for only seconds each use. How could we call the current solution anything other than terribly in-efficient with respect to land use?

Traffic Lights

This author admits to a passionate dislike of traffic lights. If time is the only real possession we have in this world, then anything that wastes our time is rightfully hated. Wasting commute time is exactly what traffic lights do.

Traffic lights are well named, since traffic is what they create. If you define “traffic” as anything like a high density of vehicles moving at a slower pace than possible, then what better creates traffic than stopping all vehicles at one place and time? They are sometimes called “stop lights” – because that’s what they force a vehicle to do. Stopping is counter to their stated goal of promoting traffic flow. They are also sometimes called “red lights”. Why is that if they are green as much as red? Because they are not.

The geographic regions covered by roadways are (at least) two dimensional. In most places the road network that results is some form of a 2D grid (perhaps a quite irregular grid). This means intersections are present in the network. The safest solutions to implement intersections are overpasses and underpasses. But these are more costly to build, both in dollars and in land. So most often, surface intersections are implemented. Since the intersection itself will be used to carry traffic along two axes (north-south and east-west), some method is needed to avoid collisions between vehicles in the crossing flows. Sharing the intersection pavement in a time multiplexed way among the crossing axes is the obvious solution. Because a distributed control system is used – with each driver controlling his own vehicle – and because the drivers are amateurs of disparate skill levels, a fairly coarse time grain must be used for the sharing. Enter the traffic light.

2-state sequences

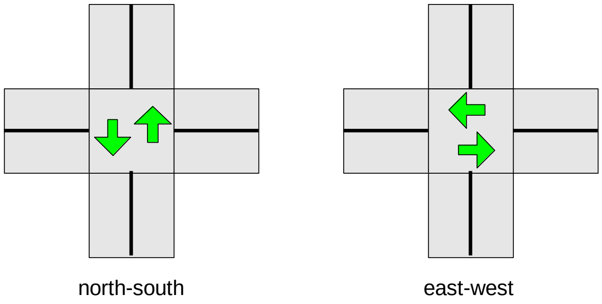

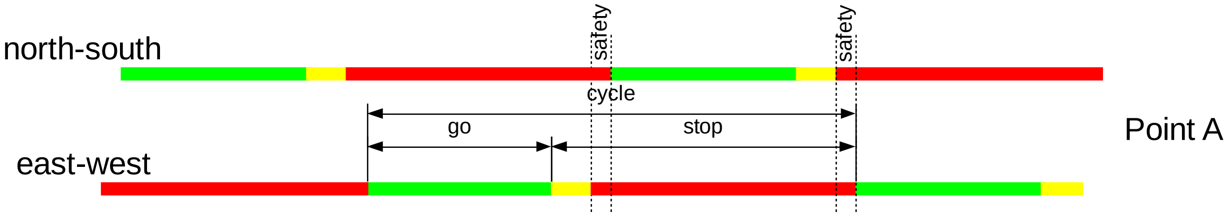

The drawing above depicts the control sequence for a basic 2-state traffic control sequence. Lights facing opposite directions always show the same color. Lights in orthogonal directions (crossing axes) generally show opposite colors so 2 states result: north-south use of the intersection, or east-west use of the intersection. The traffic light alternates between these two states. The drawing shows a control sequence spending the same amount of time in each state, although this is not required. For safety there is a brief period at each state transition where traffic along both axes is stopped. And there is a warning period where the yellow light indicates a coming transition to red.

Note that the duration of red covers 2 safety periods, the opposing green duration, and the opposing warning (yellow) period. So the red duration is always greater than the duration of the opposing green. If the green duration is the same along both axes, this means the red duration is greater than the green duration along each axis. Traffic is free to move along an axis only during a green light, so the green duration is called the “go” time. Traffic is either stopped or in the process of stopping during the red and yellow lights, so this duration is called the “stop” time. The go time and stop time combined make up the repeating cycle time. It is easy to see how the stop time is usually longer than the go time, so that the duty cycle for use of the intersection by traffic along an axis is less than 50%. This duty cycle matters.

Lets assume that a car enters the road a distance *dist* west of the intersection at an arbitrary time $t=0$, and proceeds east to the intersection at the *speed* limit. At a constant velocity time and distance are proportional, so the vehicle reaches the intersection at time $t_d=dist/speed$. Since the start time was arbitrary with respect to the traffic light cycle at the intersection, the arrival time t_d is also arbitrary. It could occur anywhere in the cycle. There is a *stop/cycle* probability t_d will occur during the stop time. It could occur right at the beginning of the stop time, so the waiting time for a green light will be *stop*. It could occur right at the end of the stop time, so the waiting time for a green light will be essentially 0. On average the duration of the wait time is *stop/2*, but there is only a wait some of the time. So the probabilistically weighted delay at the traffic light is $t_w = (stop/cycle)^{*}\ stop/2$. As the *(run/cycle)* duty cycle goes lower (below 50%), then the *(stop/cycle)* ratio goes higher (above 50%), and the time waiting at each light for the light to turn green (t_w) goes higher.

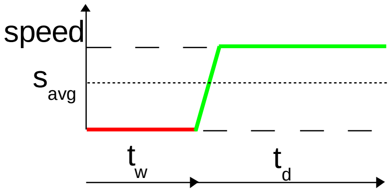

$$t_w = stop / 2$$

$$s_{avg} = dist / (t_d + t_w)$$

This wait time has a significant impact upon average commute speed. A vehicle travelling down the roadway in a direction will wait for the light to turn green, then accelerate up to the speed limit and drive the distance to the next light, then potentially need to stop again and wait for the next light. This time spent waiting is time spent commuting, but it makes no progress along the commute path. So it degrades the average commute speed. Consider a traffic light duty cycle of 50%. The wait time is 25% of the cycle, and is 50% likely, so the contributed delay is 12.5% of the cycle. If traffic lights are far apart then the travel time (t_d) is long relative to the traffic light cycle. But if lights are close together then the travel time between them will be a fraction of the cycle. For example, if the travel time is 25% of the cycle, then adding 12.5% of a cycle delay is a 50% penalty. Thus the average commute speed would be 33% lower than the driven speed (speed limit) {speed = dist /(0.25*cycle) vs Savg = dist / (0.375*cycle); 1 – (Savg/speed) = 1 - (0.25/0.375) = 0.33}.

This analysis matches well with the real world experiences of drivers. In downtown areas posted limits are usually 35 mph. Yet commute rates generally vary between 11 mph (31%) and 22 mpg (62%). Suburban speed limits are often 45 mph. Yet commute rates are only 14 mph to 30 mph (a similar 31% to 67% range).

The duration of the warning time (yellow light) is determined by the posted speed limit. This light must persist long enough to permit a driver to recognize the light change and react to it by beginning braking, and to brake at a reasonable rate from the speed limit to a stop. So for example, for a posted speed limit of 35 mpg, at least 8 seconds of yellow are needed - 2 seconds of reaction time, and 6 seconds to brake from 35 mph to a stop at 6 mph/second. Occasionally, this warning is not sufficient, so a vehicle continues through the intersection (after the yellow) rather than stopping. The safety interval delays cross traffic from entering the intersection immediately in order to protect against this. Usually a couple seconds is sufficient. These safety intervals represent lost efficiency, since the intersection pavement is not used by vehicles along either axis during this time.

Note that the duration of the red light for one axis is longer than the duration of the green light for the other axis,. The amount of the excess duration is twice the safety interval plus the yellow duration for the other axis. The cycle time is the sum of the go time and the stop time. Therefore the cycle time is the sum of the go time for each axis, the warning times for each axis, and two safety intervals. It is not required that the go times for each axis are the same (although as we will see later it is common). Using disparate go times is a good way to support more traffic flow along one axis than the other.

There is a practical limit on the stop time, which therefore puts a practical limit on the cycle time. The roadway leading up to an intersection can be thought of as a queue. Initially, and during a green light, the queue is empty. But once the light turns yellow (and then later red), vehicles encountering the light must enter the queue by stopping behind those that arrived earlier. These cars are lined up on the roadway at some maximum packing density. This maximum packing density, and the length of roadway between this light and the previous one determine the maximum number of cars that can be in this queue. The maximum queue size and the rate cars arrive and enter the queue limit the time the queue can be allowed to accumulate (the stop time).

For example, suppose there are 12 city blocks per mile, and two traffic lights 2 blocks apart; with peak traffic flowing along the path at a rate of 1 car per second. The blocks are 440 feet long, so the maximum line of queued cars is 880 ft. If the cars tend to line up with a density of about 1 for every 25 feet, then a queue with more than 880/25 = 35 cars will begin to block the previous intersection. So the stop time should be no more than 35/1 = 35 seconds.

In reality the limit should be lower than this. When the light turns green the first car in the queue will begin moving. It will be some time later before the last car in the queue begins moving. Until then that last car tends to block the previous intersection. Leaving some buffer for additional accumulation during this acceleration time is necessary to avoid blocking the earlier intersection. The real math to calculate the limit has to do with differences between the rate of arrival of new vehicles into the queue and the rate of departure from the queue as vehicles accelerate. The simple rule of thumb is that only ½ of the roadway between the lights should be used for the queue if we are to avoid blocking the previous intersection. So in the example, ((880/2)/25)/1 = 17 seconds stop time should be the maximum (and this includes an 8 second warning interval). So the light should be red no more than 9 seconds (including the 2 safety intervals of 2 seconds each).

A desire to avoid lengthy queues at red lights, and the upstream blocking that could result argues for a selection of short cycle times in the design of the light's control sequence. However, such a selection brings negative consequences. First, recall that the 2 safety intervals are lost efficiency. No traffic flows during these times. The warning intervals can be thought of in a similar fashion. During these times traffic is stopping. The intersection may still be in use, but traffic flow is being impacted. These lost times represent a greater percentage of loss if the cycle is shorter. For example, 20 seconds of loss (2 safety periods of 2 seconds each, and 2 warning intervals of 8 seconds each) is only 8.3% of a 4 minute cycle, but is 16.6% of a 2 minute cycle.

Additionally, the motivation for short cycles grows out of closely spaced traffic lights. For such tightly spaced lights the travel time between them time (t_d) will be small. You will note the equation for s_{avg} associated with the drawing above assumes an instantaneous velocity change from 0 to the speed limit. The drawing however shows a more reasonable acceleration process. If (t_d) is relatively long then the impact of this acceleration on s_{avg} is negligible. If (t_d) is relatively short then the impact is more pronounced.

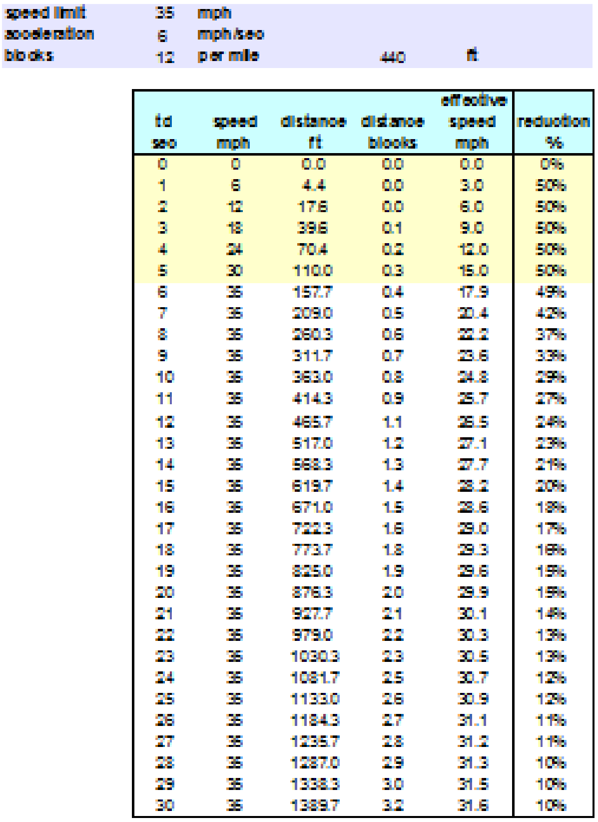

speed limit	35	mph		
acceleration	6	mph/sec		
blocks	12	per mile	440	ft

td sec	speed mph	distance ft	distance blocks	effective speed mph	reduction %
0	0	0.0	0.0	0.0	0%
1	6	4.4	0.0	3.0	50%
2	12	17.6	0.0	6.0	50%
3	18	39.6	0.1	9.0	50%
4	24	70.4	0.2	12.0	50%
5	30	110.0	0.3	15.0	50%
6	35	157.7	0.4	17.9	49%
7	35	209.0	0.5	20.4	42%
8	35	260.3	0.6	22.2	37%
9	35	311.7	0.7	23.6	33%
10	35	363.0	0.8	24.8	29%
11	35	414.3	0.9	25.7	27%
12	35	465.7	1.1	26.5	24%
13	35	517.0	1.2	27.1	23%
14	35	568.3	1.3	27.7	21%
15	35	619.7	1.4	28.2	20%
16	35	671.0	1.5	28.6	18%
17	35	722.3	1.6	29.0	17%
18	35	773.7	1.8	29.3	16%
19	35	825.0	1.9	29.6	15%
20	35	876.3	2.0	29.9	15%
21	35	927.7	2.1	30.1	14%
22	35	979.0	2.2	30.3	13%
23	35	1030.3	2.3	30.5	13%
24	35	1081.7	2.5	30.7	12%
25	35	1133.0	2.6	30.9	12%
26	35	1184.3	2.7	31.1	11%
27	35	1235.7	2.8	31.2	11%
28	35	1287.0	2.9	31.3	10%
29	35	1338.3	3.0	31.5	10%
30	35	1389.7	3.2	31.6	10%

The table to the left exemplifies this. It shows the case for uniform acceleration up to a speed limit of 35 mph. This takes 6 seconds. After that the speed limit is maintained. It calculates the effective (or average) speed over the whole time moving. During acceleration this average speed is a 50% reduction from the instantaneous speed. After acceleration the reduction begins to taper. But even after 15 seconds the effective speed is still a 20% reduction from the instantaneous speed. And even after 20 seconds and covering 2 city blocks the reduction is still 15%.

So if traffic lights are too close together, then travel time between them is small, and this means that average travel speeds suffer after any light stops a vehicle. Light cycles times may be short in this case, so that stop times and wait times are also short, but these wait times still compound the average speed reduction.

In fact, traffic light density has a compounding effect on average travel speed. The more densely traffic lights are placed, the less efficiently the intersections are used, and the lower average travel speed is. So less time is spent moving and moving is done more slowly.

Synchronization

The impact of wait times at traffic lights can in some cases be effectively managed and reduced by using synchronization. Synchronization works best for road networks that are very regular grids, so that the distance between intersections is nearly the same along any road (actually times are what matter, so the ratios between distances and speed limits is what must be relatively constant).

The drawing below depicts how synchronization works. All intersections implement the same cycle length and duty cycles). Additionally, the start times of the cycles at various intersections have a fixed relationship to one another. The drawing depicts 2 intersections in the grid, at Point A and Point B, distance *dist* apart. The roadway along the path between the points has a speed limit of *speed*, in both directions. An east-bound vehicle leaving Point A just as the light turns green will reach Point B (t_d) later, just when the east-bound light at Point B turns green. So wait time is avoided at Point B. A west-bound vehicle at Point B can leave at that time, and it will arrive at Point A (t_d) later. If the Point A go time is long enough, that vehicle will likewise avoid wait times.

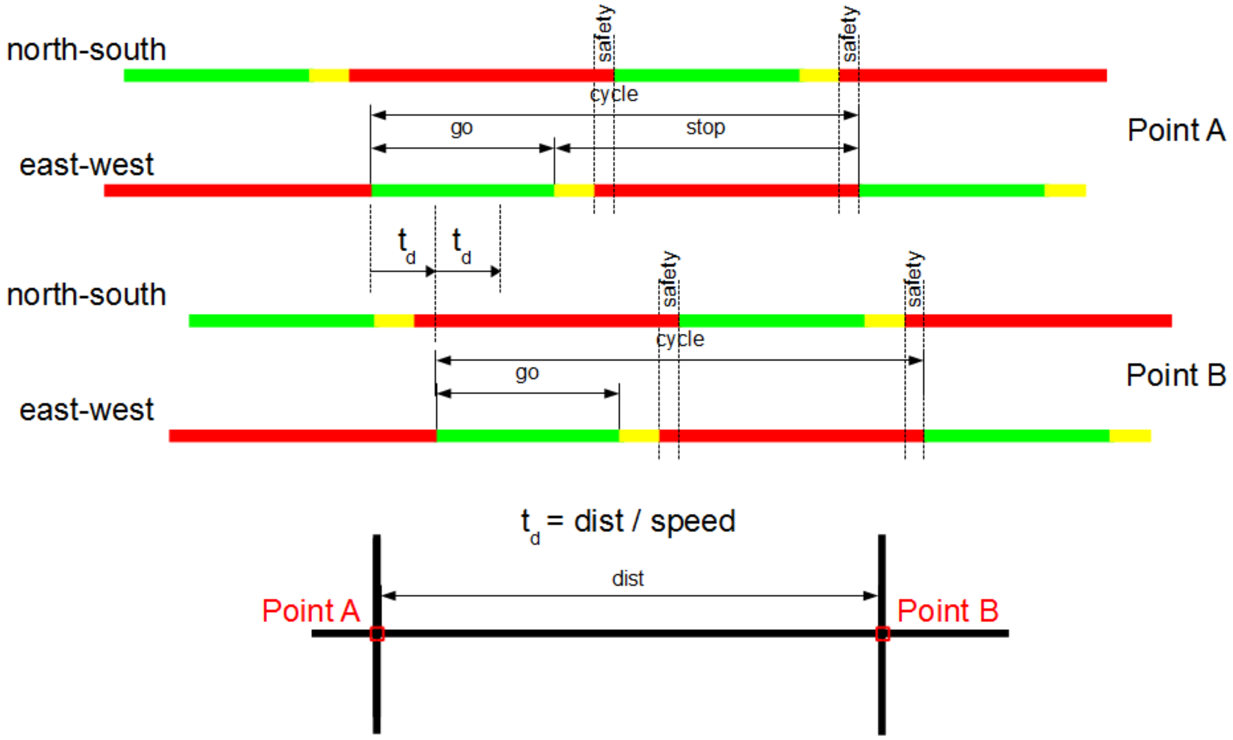

When the whole grid is properly synchronized, the east-bound vehicle (after stopping once) should be able to move continuously, at the speed limit, through one intersection after another, always reaching each intersection during its go time. The west-bound vehicle likewise should be able to stop just once, and then travel through some number of intersections during their go time, before it finally encounters an intersection during its stop time. That west-bound vehicle should then repeat that "N hops" pattern. A longer cycle time allows the number (N) of hops to be larger. The grid is synchronized along both axes, with one direction along each axis getting the preferred treatment. Usually the signal pattern is changed throughout a typical day to match dominant commute activity - e.g. north and east favored for the morning commute, then south and west favored for the evening commute.

Synchronized traffic light patterns can have major benefits. Average speeds are higher, so commute times are shorter. Since commute times are lower, vehicles stay on the road less time, so the number of vehicles loading the roadway network is reduced. Each trip requires fewer accelerations, so fuel economy is maximized. There is less stop and go, so traffic density is lower, resulting in fewer accidents.

Unfortunately, many municipalities do not synchronize their traffic lights, and most who do see far smaller benefits because their lights implement 4-state sequences (rather than 2-state sequences).

4-state sequences

The organizing paradigm for automobiles in America is “right-hand drive”. We drive in the lanes to the right side of the roadway (which means the driver’s seat is on the left, the right of way is to the right, etc.). Most of our roads are bi-directional. Therefore left turns – at surface intersections, and while entering or leaving the roadway - are uniquely problematic. They require that the vehicle cross one or more on-coming lanes of traffic.

Safely executing a left hand turn across flowing traffic is one of the more difficult tasks of driving. It requires patience, judgment, quick reactions, manual coordination, and a feel for the responsiveness of the vehicle. The “personal transit” operators in America are amateurs. As a population, they have demonstrated poor ability to perform left turns safely. For years accident rates showed the need for left turns to be a significant contributing factor.

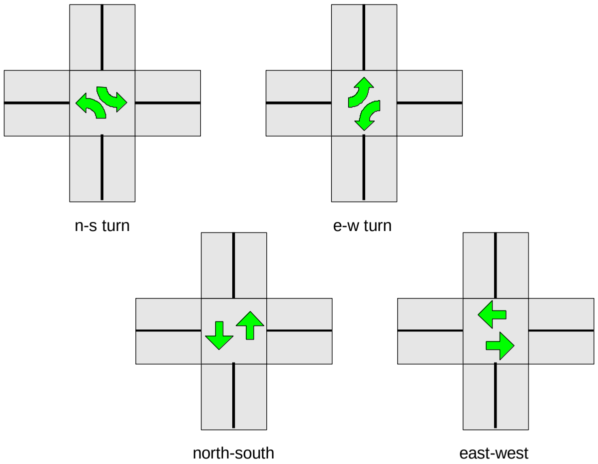

Most municipalities have chosen to manage this problem by providing protection for left turns. Today most left turns are performed in the absence of on-coming traffic. This is accomplished by altering the control sequence of traffic lights to implement a 4-state sequence. The existing two states – north-south, and east-west – are augmented with two more states which allow each of the axes to perform protected left turns. The drawing here depicts the 4 states.

These additional states require additional time in the cycle, requiring the cycle to be longer. So the same go times for the north-south and east-west axes now occur less frequently. Therefore the duty cycle for a 4-state sequence is much lower than for a 2-state sequence. The durations of the new n-s turn and e-w turn states are generally shorter than the north-south and east-west states, but new safety intervals (of the same duration) are required before each. So the efficiency of the 4-state sequence is also lower than for the 2-state sequence.

Consider the example of 35 mph speed limits in all directions, with 2 second safety intervals and 8 second warning intervals, and 10 second turn states, and 30 second go states. The cycle time is 2 minutes (120 seconds). Efficiency is 1 – (4*(8+2)/120) = 1 – 0.333 = 66.2%. Duty cycles for the go states are 30/120 = 25%. Stop times are 90 seconds, which suggest a minimum distance between lights of 8400 feet (approximately 19 blocks).

It should be clear that the big cost of the 4-state sequence is longer stop times. This increases the likelihood of encountering a wait time, increases the average length of the wait time, and makes it much more likely that queues will back up to block the previous light, and thus further reduce average speed during the go time. Synchronization can be used with 4-state sequences. But the reduced duty cycles makes it very difficult to achieve more than a single hop in the direction counter to the preferred direction. In fact, counter flow is almost guaranteed to be negatively synchronized – to encounter a stop at each intersection.

Still, traffic lights in most municipalities today implement 4-state sequences. Municipalities have chosen safety (reduced incidence of left hand turn related accidents) over commute speed/times. The reduced duty cycles also give the traffic flow a more obvious pulse. This tends to make it easier to enter and exit the roadway between lights.

For financial and space reasons surface level intersections dominate the roadway networks in most municipalities. Traffic lights are used as the control mechanism in these cases, to time multiplex share the roadway common to both axes of vehicle movement. This multiplexing obviously reduces the portion of time vehicles can move across the intersection along either axis – in many case significantly so. Constant movement along a commute path is converted into a series of starts and stops. This lowers average speed, and radically increases fuel consumption. It also reduces average distance between cars while increasing the relative speed differences between those cars – making collisions more likely. Average travels speeds may be cut by one half, to two thirds. How could we call the current solution anything other than terribly in-efficient with respect to commute time?

Safety

No transportation system is perfectly safe. But the public transportation system is fraught with risks. Terrain and weather conditions vary greatly; road construction specifications vary from one location to another; there is a wide disparity in vehicle types and parameters operating in the same area and at the same time; vehicle operators are amateurs, un-supervised, and subjected to many distractions.

In December of 2014 the National Highway Traffic Safety Administration issued its summary report for the previous year, “2013 Motor Vehicle Crashes: Overview” (DOT HS 812 101). The following 4 charts are from that report.

There were 5.687 million crashes in 2013. Most, 4.066 million (71.5%) caused only property damage. The rest caused injury to someone, injuring 2.313 million people. Of those crashes, 30,057 involved fatalities, and killed 32,719 people (an average of 1.09 deaths per fatal crash). Both the number of injuries and the number of deaths were down from the previous year, but total miles driven in 2013 was also down from the previous year. As a result the rates of injury and fatality (expressed per 100 million miles driven) were very slightly lower than 2012, and near the 20 years lows set in recent years.

18,257 of the fatalities (55.8%) occurred in crashes where the vehicle ran off the road (driver lost control). 8,598 of the fatalities (26.3%) occurred in crashes at intersections. Nearly one third (31%) of the fatalities occurred in alcohol related crashes. Nearly ten percent (3154 = 9.6%) of the fatalities occurred in distraction-affected crashes. 424,000 of the 2.313 million injury causing crashes (18.3%) occurred in distraction-affected crashes.

People Injured and Injury Rate per 100 Million Vehicle Miles Traveled by Year

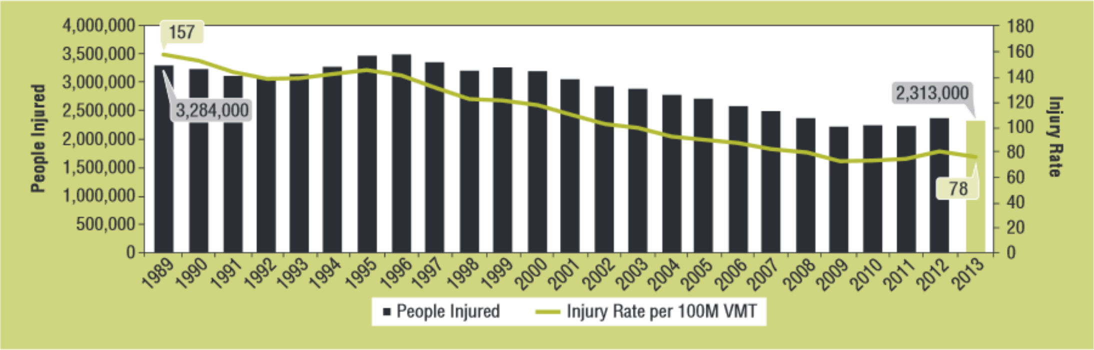

Source: NASS GES 1989–2013; Vehicle Miles Traveled (VMT): Federal Highway Administration.

Fatalities and Fatality Rate per 100 Million Vehicle Miles Traveled by Year

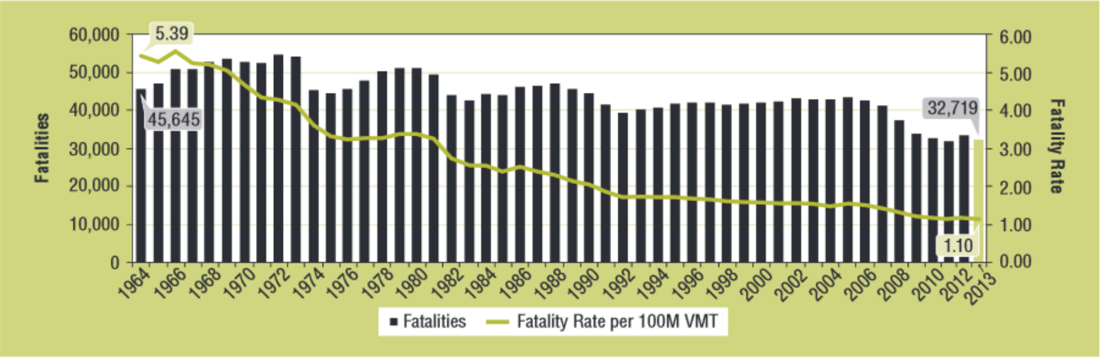

Source: 1964–1974: National Center for Health Statistics, HEW, and State Accident Summaries (Adjusted to 30-Day Traffic Deaths by NHTSA); FARS 1975–2012 (Final), 2013 Annual Report File (ARF); Vehicle Miles Traveled (VMT): Federal Highway Administration.

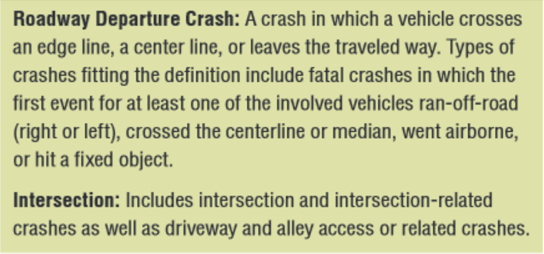

Roadway Departure Crash: A crash in which a vehicle crosses an edge line, a center line, or leaves the traveled way. Types of crashes fitting the definition include fatal crashes in which the first event for at least one of the involved vehicles ran-off-road (right or left), crossed the centerline or median, went airborne, or hit a fixed object.

Intersection: Includes intersection and intersection-related crashes as well as driveway and alley access or related crashes.

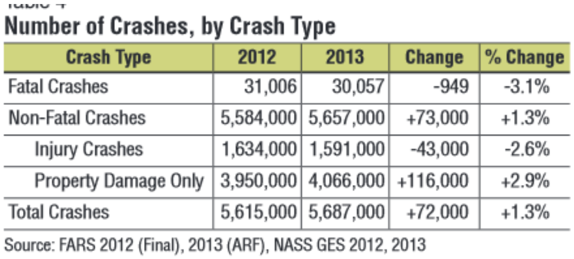

Number of Crashes, by Crash Type

Crash Type	2012	2013	Change	% Change
Fatal Crashes	31,006	30,057	-949	-3.1%
Non-Fatal Crashes	5,584,000	5,657,000	+73,000	+1.3%
Injury Crashes	1,634,000	1,591,000	-43,000	-2.6%
Property Damage Only	3,950,000	4,066,000	+116,000	+2.9%
Total Crashes	5,615,000	5,687,000	+72,000	+1.3%

Source: FARS 2012 (Final), 2013 (ARF), NASS GES 2012, 2013

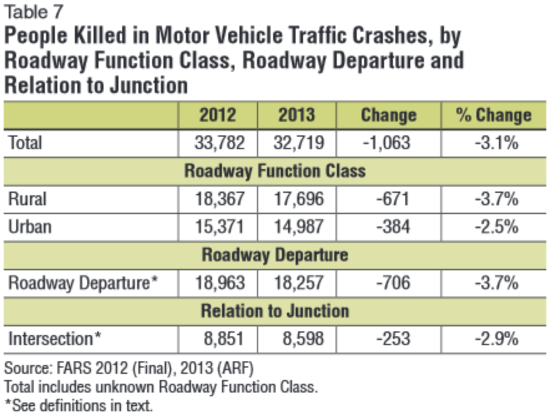

Table 7

People Killed in Motor Vehicle Traffic Crashes, by Roadway Function Class, Roadway Departure and Relation to Junction

	2012	2013	Change	% Change
Total	33,782	32,719	-1,063	-3.1%
Roadway Function Class				
Rural	18,367	17,696	-671	-3.7%
Urban	15,371	14,987	-384	-2.5%
Roadway Departure				
Roadway Departure*	18,963	18,257	-706	-3.7%
Relation to Junction				
Intersection*	8,851	8,598	-253	-2.9%

Source: FARS 2012 (Final), 2013 (ARF)
Total includes unknown Roadway Function Class.
*See definitions in text.

According to the 2013 data, the likelihood of being involved in a crash was 0.192% for each 1000 miles driven. So a typical driver covering 12,000 miles per year had a 2.3% chance of being involved in a crash. Or put another way, out of 100 such typical drivers 2.3 of them were involved in a crash. The police and court systems established liability for all the damages caused in those 2.3 crashes, including compensation for fatalities. The 100 drivers paid insurance premiums to collect that damage amount (with a little left over as profit for the insurance companies). An estimated $277 billion worth of economic damage resulted from the 5.867 million crashes, and 2,966 billion vehicle miles. That's is an average of $47,213 per crash, and $0.093 per driven mile. The social cost – including lost wages of injured or killed people – is roughly double that.

Safety is a very subjective metric. It is unclear what is expected and what is acceptable when it comes to risk associated with transit. A 2.3% chance of being involved in a crash each year means an adult who drives his entire adult life (45+ years) is statistically certain to be involved in a crash (103.5%). The fact that this outcome is nearly entirely attributable to amateur drivers' negligence, incompetence, malfeasance, or distraction certainly seems to fall short of what might be reasonably expected.

Expense

Private ownership and operation of a personal vehicle is not cheap. The current IRS recognized expense rate for business use of a vehicle is $0.565 per mile. This includes the costs of vehicle acquisition, registration, insurance, fuel, and routine maintenance. The table below gives an example of how that figure was computed.

	Cost Breakdown of Car Ownership				
	purchase price	$14,600	$		
	sales tax	$730	$	5%	
	fees	$1,000	$		
	financed amount	$16,330	$		
	finance rate	2.50%	%		
	term	72	months		
	monthly payment	$210.46	$ per month		
Amortized time based costs @		12.5	K miles / year	$ / mi	**$0.31**
	purchase	$210	$ per month	$0.20	
	insurance	$1,200	$ per year	$0.10	
	registration	$125	$ per year	$0.01	
Operating costs				$ / mi	**$0.25**
		15.6	mpg		
	gasoline	$3.50	$ / g	$0.22	
	oil	$60	@ 7.5K mi	$0.01	
	routine service	$400	@ 30K mi	$0.01	
	minor repairs	$200	@ 30K mi	$0.01	
Total costs					**$0.56**

It should be noted that real common experience may be substantially higher than this tax-based figure. The purchase price in the table is in range with the advertised prices for sub-compact cars with standard equipment. In fact the average prices of a new car is now more than twice that amount, $33,560 according to a recent story in the LA Times. (http://www.usatoday.com/story/money/cars/2015/05/04/new-car-transaction-price-3-kbb-kelley-blue-book/26690191/). It is also likely that anyone financially limited to a sub-compact car, will not be able to qualify for the financing terms shown in the table. So depending upon the vehicle and the purchaser, the monthly payment could be as much as 3 times the amount shown in the table. Therefore the amortized contribute of the purchase to the cost per mile could be as much as 3 times greater. This could bring the per mile cost as high as $0.96.

Insurance is the third largest contributor to the per mile costs. Vehicle replacement cost will drive up the insurance premium. More expensive cars will be more expensive to insure. The operator's driving record can also have a large impact. The example shown in the table is for an operator with an excellent driving record. The insurance premium also varies significantly (factor of 2) based simply upon the operator's state of residence (http://www.valuepenguin.com/average-cost-of-insurance). So again, average insurance costs can easily be higher than shown in the table, to add an additional $0.20 per mile to the costs.

New vehicles should have better gas mileage than the US fleet average used in the table. However, since such a large percentage of new car sales are light trucks (or other large vehicles) this is not certain. Fuel costs are the second largest contributor to per mile costs, so variation of +-$0.15 per mile are possible.

The table estimate is for a vehicle with pretty good reliability. It is unfortunately too common to encounter ugly surprises regarding this expense. The AAA estimates an average expense here that is far in excess of the table value ($0.608 per mile http://newsroom.aaa.com/2013/04/cost-of-owning-and-operating-vehicle-in-u-s-increases-nearly-two-percent-according-to-aaas-2013-your-driving-costs-study/) Even sub-compacts are calculated as needing $0.464 per mile for maintenance and repairs on average.

So we see that real vehicle costs can vary quite significantly. In a best case scenario they could be as low as $0.415 per mile. But the "average" case from the published data above is $1.265 per mile. And it is therefore reasonable to expect at least some owners experience costs as high as $1.90 per mile.

Average personal income	40.56	$K per year
Average family income	52.05	$K per year
Average miles driven	12.5	K miles / year

vehicle expense ($/mi)	$K per year	% family income	% person income
$0.415	$5.19	10.0%	12.8%
$0.565	$7.06	13.6%	17.4%
$1.265	$15.81	30.4%	39.0%
$1.900	$23.75	45.6%	58.6%

The table above attempts to put the per mile cost in a useful perspective. Even the IRS recognized per mile cost represents 1/6 of the average income of an individual, and 1/8 of the average income of a family, assuming you drive the average number of miles per year. The "average" expense case represents 1/3 of the family income and 40% of the personal income. If you drive more, the share of income that goes to pay for transportation is more.

Judgments based upon this table will obviously vary based upon who makes them – highly subjective. Certain other expense items must be higher priority than transportation – housing, food, health care – and can consume a lot of income. So how much is left will determine if automotive transportation is affordable at all, and under what conditions. It is easy to see how people making less than an average income may consider a car un-affordable. It is also easy to see why there is a large market for reliable sub-compact cars. That is about all near average incomes will support.

As high as the costs shown above are, they are unfortunately not the whole story. Cars are of little use without roads to drive on, signs and traffic lights to coordinate traffic, police to enforce the traffic laws, and courts to apportion financial responsibility for accident damages.

In theory, roads are already paid for by a tax on gasoline purchases, and so are already included in the per mile costs above. Unfortunately that is not true. The gas tax pays only a portion of the costs for road construction and maintenance. Additional revenues from sales taxes and property taxes may be applied to roads. On occasion, bond revenues may be applied to projects that support autos. The debt service for the bonds is paid from general tax revenues. Some of the tax sources are direct to the local municipality. Others are routed through the state or national government back to the local government. But the ultimate source of all these revenues is the tax payers. The national per capita average for roads alone is $600 per year (http://www.mlive.com/news/grand-rapids/index.ssf/2015/02/road_database_test.html). That is an additional $0.05 per mile.

This author found no good sources to quantify other municipal costs associated with car use. However, it is clear they exist, and it is clear they are passed on to the tax payers as part of normal operating costs of government. So an additional unquantified component of the per mile cost of cars exists and is paid by car owners.

Conclusions

The massive use of personally owned and operated automobiles delivers a convenient, reliable, and relatively quick transit system to hundreds of millions of Americans. But this fact should not confuse us into believing it is an especially good transit system. It is the dominant transit system simply because it is effectively the only option available to most commuters.

The system has many aspects that can and should be significantly improved upon. First, the system is far less safe than is possible, and less safe than its users deserve. Second, the system is ridiculously expensive. Third, while it is a relatively fast system compared to the alternatives, it is nowhere near as fast as possible. Finally, it is very wasteful – it wastes land, it wastes time, it wastes fuel.

Clearly given that we are living in the age of the Jetsons, a better alternative must be possible.

Just within the category the DOT calls "cars and light trucks", there are at least 35 companies selling vehicles, and offering a combined 363 different models. We have optimized consumer choice over and above the effectiveness of the system.

It is worth repeating that while many of the costs associated with providing the infrastructure needed by our automotive society fall on the local and state governments and transit districts, those entities have very limited control of those costs. Roads could be constructed and maintained more cheaply if they were only required to support compact cars or smaller, carrying passengers (average width = 72 inches; average load per wheel ~= 900 lbs). But many of the vehicles on the roadways are 18-wheel tractor-trailer trucks, carrying cargo (average width = 96 inches; average load per wheel ~= 2500 lbs). Yet the governments responsible for the roads defer to the national government for setting vehicle specs, which itself largely defers to industry groups to set the specs. The result is a set of specs that must cover 363 models. Similarly these governments have little control over the number of vehicles allowed to operate on the roadways, or when they operate. We watch as these governments struggle to exert what control they do have to maintain the competence of the vehicle operators, while other commercial interests work to degrade that competence. Driving under the influence is an on-going struggle, while the fight against texting behind the wheel is just beginning. Meanwhile vehicle manufacturers add bigger LCD displays to the console, with more apps requiring more complex interaction.

It is an often proven reality in the corporate world, an organizational structure that assigns responsibility without assigning commensurate authority (the right and capacity to author the orders for action by others) is a structure that will fail to achieve its purpose. It should therefore be no surprise that our automotive society is breaking down – longer commute times, more accidents, higher taxes, more expensive cars, etc. This trend will continue until we fix our organization, and focus narrowly on the societal need – to safely and reliably get people where they are going faster and cheaper.

About the Author

Stephen Hamilton is a newly retired individual who just concluded a 40 year career as a computer and integrated circuit designer. He has had a lifelong interest in physics, cars, and transit technology in general. Believing that climate change is a species threatening reality, and that governments have failed to face it, he has decided to spend his remaining time and talents trying to use commerce to impact it. Stephen has decided that personal travel in America is the area where his impact can be greatest. He has spent the past 18 months studying this area, and thinking about its challenges and possible/likely solutions. CityTram.org was created to encapsulate these efforts.