The Privacy and Safety Impact of Technology Choices for Command, Communications and Control of the Public Highway

Jon Crowcroft
The Computer Laboratory, University of Cambridge
Cambridge, UK
jon.crowcroft@cl.cam.ac.uk

ABSTRACT

Monitoring, and command, communications and control[1] of private vehicles on the public highway is now high on the political agenda. This is both because it is becoming feasible, and because it may be desirable. From the economic perspective, more efficient use of road resources may be achievable. From a safety perspective, it would clearly be good to reduce road injury and death statistics below the current "9/11"'s-worth per year in the UK (and other similar sized European countries).

Various prototypes, proposals and projects are being undertaken. There are a number of technologies that interact as well as numerous legal, political and economic stakeholders. In this note, we pay particular attention to the impact on privacy and safety of different approaches to the overall problem.

The purpose is to draw attention to the potential unintended consequences that result from decisions being made at the time of writing, in this arena.

Keywords: "Transport", "privacy", "safety", "policy"

1. CARROTS AND STICKS - THE WHYS AND WHERE FORES OF COMMAND AND CONTROL

A variety of stakeholders have been lobbying for more automation of the highway. Such automation could address the two key areas of road usage: efficiency, and safety.

From the perspective of transport agencies, whether public users or private freight companies, or the government Departments and other road resource providers, transparency of costs could make planning easier. The current systems in most of Europe do not create a level playing field between road, rail and air (or other lesser systems, waterways). There is a complex mix of subsidy, and user taxation used to build and operate the roads, in a climate where the alternatives are largely charging at the point of use, and competing.[2] It is clear from the congestion charging experiment[3] in London that once users are acclimatized to the idea, the potential improvement in road throughput and delays is perceived relatively positively. All types of consequences arise from this, including the possibility of removing road tax altogether, and funding the roads entirely from usage bills. However, the deployment of such systems requires technology to monitor road use. This technology can be within the car itself. The car may be presented with information from navigation systems and roadside monitoring data concerning routes, conditions and prices. The driver then makes a choice, and pays (perhaps using a hands-free cell phone, or new devices that beam tokens to the road charging units, and are refreshed from time to time using cash, or credit cards (just like pay-as-you-go cell phones). Alternatively, as in the case of the London congestion charging scheme, the system could rely on monitoring, including cameras based at the roadside. Centralized databases record vehicles, drivers and journeys. Payment is checked against this database. One can clearly conceive of several approaches to hybrid systems, with a mix of responsibility between payment and prevention of cheating.

At the same time, there appears to be some pressure on and by the stakeholders to reduce the risks of road travel. Again, technologies may help. At one extreme, there have been proposals for "car-trains". Vehicles leaving local roads and joining trunk routes would logically clump together and be controlled as a single aggregate unit, until leaving the trunk for the more local part of their journey. At the other extreme, each vehicle can be left to choose its own velocity at all times, but given data from a set of on-board and remote sensors (radar/sonar, visual, induction loop etc) on which to base the decision. How much the human is in the loop is an interesting question.

The road ahead is full of much interesting scenery. Some proponents of usage billing are promoting a particular approach based not on the traditional toll road model. Instead, incentives are aligned between road provider and road space consumers by congestion charging. There are at least two common misunderstandings about congestion charges, which certainly colour some of the opponents viewpoints. When first introduced in London, some lobbyists asked "Why should I pay to be caught in a traffic jam?". The intention is to set the price *so that* the reduction in traffic is just enough to make the operating speed of the road at a busy hour, the speed limit. In fact this is not such a naive question as it might seem. The real question behind this is "Will the price

[1]The defence connotations of the term C^3 are intended.

[2]The RAEng report on transport[4] does comment on externalities such as health, land use and pollution costs. However, it is worth noting that the percentage of home ownership, and importantly, the price of a typical house, in the UK are unusually high compared with much of the rest of Europe. This means that the effect of road pricing on mobility may be weak.

be set right, and how?". If the price is low, people will pay to be in a traffic jam. If the price is too high, the roads will be quite empty and under-utilised (assuming that there are still operational costs like maintenance to cover).[3]

If the price is way too high, very little money will be raised at all. It is a subtle business to set the price right, since users have long memories, and are averse to unpredictable bills. This means that the price must reflect users *willingness-to-pay*. The price must be a form of *shadow*, in that it must somehow reflect the cost to displace the user who does not travel. This leads in some sense to the second objection: The second objection to a price (whether a toll or a congestion price), is that it is unfair in some sense. The rich get to drive on fast empty roads while the poor cannot afford to drive. This may in fact be true depending on the availability of alternatives e.g. public transport, or the ability to travel at other times of day.

In the end, the system must also be transparently cheat proof to some degree. Of course, some degree of cheating is just part of the cost. The question is what the risk is, real and perceived, and what the costs of the cheat proofing technology itself is!

In London, regular trade users are effectively getting more business done, and are unlikely to cheat on a one time charge. The additional real work generated potentially can be viewed as an indirect subsidy to other users. However, the efficiency of the alternatives (bus/taxi) also increases. This model for billing, cheaper for lesser users or better service from alternates, could be more explicit.

In the end, though a true spot-price congestion-charge would effectively trade an unpredictable price for unpredictable journey time. This is probably not acceptable, so the smoothing function that is applied to the price is needed. This means that long term data and accurate predictions of the effect of price on road utilisation are needed. To make life complex, this is obviously time of day and location dependant. This means that more specific data is needed, with the concomitant risk to privacy.

The current safety debate is in some way the other side of the same coin. A great deal of reduction in human carnage could be achieved by reducing the speed of cars when there is an accident. This does not necessarily mean the mandating of fixed speed limits through monitoring (e.g. by speed cameras and fixed signage as today). Instead, we can envisage a dynamic system, that advises the driver of the current safe speed in each location, through road side displays (some motorways do this today) or soon through displays in the car ((change in colour on speedometer, possibly coupled with audible alarms, or changes in the interface (resistance in accelerator pedal movement). The system could go further and impose controls (limits to speed or acceleration or even turning speed and even route choice) .

The carrot is that during idle times, not only is the road cheaper to use, but the speed limits may be (significantly) higher.

Technology to achieve this all can be centralised through monitoring and external control, or decentralised through all the vehicles. We can provide vehicles with sensory data from devices in each car (and from other cars) and from road side equipment and let each car make its own decision (with or without a human in the loop).

We can certainly picture various phases of introduction of pricing and speed&route control systems. However, it is not in the gift of government agencies to act alone. Some of the technology rests necessarily with car manufacturers. Some is in the devices we add to cars (satnav (satellite navigation) combined with interfaces to engine management and braking&steering systems). Some influence no doubt will be entertained or even enjoyed by insurance companies as the obvious reduction in risk and cost to them slowly permeates the road network and its users. Actuaries are going to have some work to do.

To summarise then, what we have introduced in this section is a hodgepodge of ideas that are being bandied about in the transport arena based on some disparate visions of where policy, technology and regulation may take us over the next two to three decades:

Route cost advisory Can offer driver more (realtime) choice. Can be given without knowing anything about the driver or vehicle's location.

Weather/Road conditions advisory Can offer more safety. Can be given oblivious of driver location. Accident information can also be offered - could, in rather obscure circumstances, be some invasion of privacy of victims of accident.

Proximity warning We already start to see cars with radar, initially for parking, but now also for driving; and cars that can detect when the driver strays over the edge of a lane. These are typically completely local technologies to the vehicle.

Speed limiting There have been speed limiters available for a long time for people that buy fast cards but want cheaper insurance. These sometimes have overrides that permit one-time kick-down to allow car to turn off the limiter. To re-enable the limiter takes a factory reset. Obviously interaction with an insurance company might be necessary at this point (one can imagine needing a report and authorisation codes to reset the device).

Collision Avoidance We havn't got any where near this. Clearly, sudden enabling of brakes or steering to avoid hitting other cars, or pedestrians or bikes is feasible, but very complex. Autonomous decisions would need to take account of other vehicle and obstacle proximity.

Next we look at the technologies for monitoring and control from the viewpoint of possible threats and attacks.

2. THE SPECTRUM OF TECHNOLOGIES FROM AUTONOMY TO CENTRALISATION

In the previous section, we surveyed the somewhat haphazard landscape of choice in front of the communities on todays highways and bi-ways.

[3]As discussed elsewhere, recovering the true operating costs should be a separate goal from the the control of efficiency: congestion pricing is just one way to address the latter, as is admission control. Distance usage charging might be a way to address the former. Conflating the two is just the sort of problem that one is trying to avoid in any changes from today's regime.

In this section, we take a closer look at the range of ways that these choices could be made, to tease out some of the unintended consequences in terms of longer term privacy and safety.

In the car What are the threats to devices in the vehicle?

- Car registration tags - these can be visual (e.g. circular bar code) or electronic (e.g. RFID). All of these are subject to natural or deliberate obfuscation, including masquerades or Sybil attacks.
- Engine management systems control fuel economy and might link to other systems such as braking and speed limiters: These have long been subject to chipping. DIY toolkits for re-programming the PROM that holds the engine management software have been around for 20 years for many models. Network access to the system is become more possible.
- Satellite and other location service based navigation systems rely on reasonably clear line-of-sight to the source of the signal. This is obscured in modern cities, but can also very easily be jammed. Of course most satnav systems include inertial guidance and ready-reckoner, which coupled with a map can allow for fairly long gaps in coverage (100s of meters/seconds). In general, diversity is a defence (e.g. new European Galileo alternative for GPS might help, as can other location systems, perhaps based on 3G and 4G network, including WiFi which can, with 2G today, give location to several meters).
- Speed controllers, as with tachygraphs, can of course simply be tampered with and often have, for safety reasons overrides. (e.g. emergency requires fast drive to a hospital). Of course, we can build in disincentives to tampering (e.g. invalidates tax/insurance etc). As with other utilities, if an in car road-management black box is depended on, then the security of the box is a pivot point, and its robustness is critical. More subtly, there are distributed threats to speed control hardware, from people without speed control (not yet fitted, or destroyed/interfered with), which need to be mitigated with sensor checks.
- In general, black boxes can be used for many things (e.g. recording vehicle, driver and road conditions), that may be used as evidence (e.g. in accident for assigning liability), thus the systems will be under attack before, during and after the event.

By the road What are the threats to devices by the roadside?

- If we monitor road use, then clearly the visual or radio (e.g. RFID) sensors are subject to a variety of threats including removal/destruction, jamming or spoofing of input (temporarily or permanently).
- The networks used to deliver information to and from the roadside systems, and possibly from and to the vehicles, will necessarily have at least one wireless hop - e.g. 802.11/WiMAX etc. These are subject to jamming as well as intercept (e.g. if I want to find out where and when such and such a vehicle goes).
- Public wide area wireless nets might provide long haul retrieval/upload of data to roadside or vehicle (e.g. GSM/GPRS/3G) obviating the need for a separate road network. There are questions about coverage, ownership, and price hiking by the network owners, when such networks become vital to another activity than merely voice calls.

Under the road What are the threats to communications back-haul?

- The road operator might consider (and is) building their own back-haul network, both for the sensor systems (cameras as well as induction loops, and newer radio tag car recognition systems). This raises questions about possible monopoly ownership or a vital resource. Clearly a regulator (FCC/Ofcomm etc) would put pressure on allowing open access to a fibre network that covered most of the UK road system, not just to 2G/3G phone companies, but ISPs and others. There's strong motivation for the road network owner to roll out a fibre net especially since it could give much better resilience than wireless or copper, not just along each road, but system wide, and is far harder to attack or intercept. Nevertheless, if it is then a shared infrastructure, there are DDoS attacks and other intrusions one would be concerned to design against.
- Thus re-use of back-haul fibre (license capacity to 2G/3G companies and wireless ISPs) is likely to be a given, and therefore also a risk.

At the center (Department, or their proxy) What are the consequences for central services?

- All this data has to go somewhere, typically into some centralised databases. Thus we need to consider all the massive access control problems, as well as prevention of data-mining by unauthorised people (or users with authorised access acting beyond the scope or role they were given access for).
- There are questions of reliability and stability bought about by the sheer scale of the systems being envisaged. There are also questions of accuracy. Can we really build a system this big in the foreseeable future? I am not sure anyone has built a realtime GIS sensor database as large as one might foresee. Other architectures (hierarchical system with summarising and only aggregate data logged in the central system) doesn't really make sense given the monitoring and control needs to correlate with payment and DVLA car owner/keeper data.
- There are questions about how to use broadcast nets, if one designs the system for data push rather than pull: which way does data flow - car to center, or center to car? If center to car, then using

TV nets (as is partly done today) seems reasonable, and is hard to attack. On the other hand, attacks (corruption of information/planting false data) would have a global impact.

Next, we'll look at two extreme views of how these ideas might play out, in an attempt to illustrate the pitfalls of ignoring large scale questions of privacy and safety (in general, under the heading of security).

3. PRIVACY AND SAFETY

We kick off this section with a tour through the rich picture the potential road command, control and communications world presenters.

The stakeholders are many: users (drivers, passengers); car manufacturers; government; in-car equipment vendors (Satellite Navigation (satnav), entertainment); insurance companies and so on.

The types of information one might gather varies, from individual through to aggregate data; aggregate data clearly still has value (e.g. for setting a short or long price, for choosing a route, and for setting speeds). It also has value (freight haulage savings could be significant - it would be interesting to see if any figures have been published on this for London yet).

However, coupled with the billing database, the aggregate data is amenable to post hoc mining. There are temptations if there are two separate databases, to federate and mine. There would be pressures on the providers from various agencies (insurance: tracking stolen cars;security: tracking terrorists etc)

The failure modes of local versus global monitor and control systems are very different - we'll look at this further below.

A critical factor in any design is that it must clearly be amenable to incremental deployment. Just as with the introducing of the seatbelt, ABS, and airbags, so too must lane/radar/sonar/lidar, speed limiters and collision avoidance systems must not depend on a flag day for tens of millions of vehicles.

3.1 Summary

To summarise, one can envisage a very top-down introduction of a system or set of systems, which might strongly tend to centralisation. Alternatively, one could design a set of incentives for stakeholders to try out a broad range of decentralised systems, and see what evolves (like the Internet!).

Centralised The pros include the ability to provide a central check for billing (mint), sequencing and planning; cheating; proof of identity and other "non-functional" aspects of the system; The Cons include that such a system would have a single point of attack or failure, a single place to launch legal seizures of information, and potential monopoly provisioning with the dangers of price hiking after deployment. Such a space is what we sometimes call a tussle[1].

Decentralised The pros include: assemblage of cheap components; very good potential fault tolerant and resilience; lack of obvious privacy problems. The cons include novel problems such as cascade/epidemic failures; we don't fully stability of large complex systems (see work by John Doyle at Caltech in this area); DDOS possibly harder to detect, as are other anomalies (just as in the Internet), but on the other hand if what we want is a good operating point for the average case, then this may not be a problem except in the mind of bean counters.

No System at all It is possible that we do not actually need any specific vehicle monitoring & control system. One could have a very simple road computer model running on a PC somewhere, and for each possible destination, there's a telephone number. Users SMS or phone that number: if they get through, they have a permit to travel (a ticket like in a Delicatessen); if not, they get a txt back saying that they have to try again later, but also that alternatives are available (e.g. bus, train, plane, other cars that are willing to publish car share destination using other means).

If the police come across a traffic jam they stop some cars at random and ask to see their travel permit - if there is no permit, then they are fined - only have to have a statistical chance of catching them to create disincentive - and if there's no jam, you can take the risk (you could even drive to near the jam and turn round before you get stuck and liable for being caught...)

No sensors, no location services: all you need is a model. and a phone net,and an admission control algorithm (we have many of those in networking); all trivial.

For safety, autonomous stuff is pretty dependable too, but if we put in collision avoidance/detectors, people will tend to drive about as close as they can get. But note that if we set the operating point of the roads so that they are not too busy, then there's no need as there's space to go as fast as you like.

Finally, for billing, it is hard to see why, given the UK statutory requirement for off-road notification there needs to be any road tax. A petrol tax covers:

1. road use
2. pollution
3. choice of inefficient engine/polluting engine

and is fair[4].

4. SUMMARY AND CONCLUSIONS

In this note we have looked at some aspects of monitoring and control of the road system. The technologies for centralized versus decentralized control[2] are nearly within our grasp and there are strong temptations to rush in and just build something. On the other hand, the safety and privacy questions that arise from different policy and technology choices are not fully laid out or comprehended yet. The requirements from enforcement and accuracy also need sensitive consultation with the public.

On the privacy side, the apparently strong will the DfT has with the road side monitors and so on is partly because

[4]It is a tax at point of use and is hard to avoid.

they have more than half the network there, and partly because they want to be seen to do something (e.g. standard arguments about "central government control" creep apply), and partly because they can do other things with the sensors (e.g. one government agency person has been heard to say that they'd love to be able to pull the hi-resolution pictures from the London cameras for intelligence work. The DfT can also lease or sell the capacity on the under-road fibre as back-haul for the 3G/cell nets - this is very attractive as a lot of the motorways are good places to put masts and are not exactly in most people's backyards (well people that make a lot of complaint anyhow). Clearly combining this with in-car entertainment is an attractive business proposition.

Moving to the control side, the central system gives little guarantee of better safety than a set of autonomous local controls on vehicles, as far as I can see. Indeed some of the failure modes of a central system could be far more catastrophic and systemic, although a fully automatic decentralised car guidance system would also have certain unpredictable modes which would need to be engineered out of the design and implementation in some as yet unforeseen way.

Just as I was finishing preparing this article for CCR, a paper appeared in HotNets IV, by Bryan Parno and Adrian Perrig from CMU[5], which contains a very thorough analysis of the threats and defenses for vehicular networks. I think this is the beginning of a large research area of great importance.

5. ACKNOWLEDGEMENTS

I am indebted to Butler Lampson and writings by Heathcote Williams[6] for the suggestion that "zero deaths on the highway" is a suitable grand challenge for Computer Science in the 21st century. Thanks also are due to Ross Anderson and members of the Digital Technology Group in the Computer Lab for discussions on the security and privacy aspects of the problem space.

6. REFERENCES

[1] David D. Clark, John Wroclawski, Karen R. Sollins, and Robert Braden. Tussle in cyberspace: defining tomorrow's internet. In *SIGCOMM '02*, pages 347–356, New York, NY, USA, 2002. ACM Press.

[2] John Doyle. Highly optimized tolerance. `http://www.physics.ucsb.edu/~complex/`, 2005.

[3] Transport for London. Transport for london congestion charging technology trial report. `http://www.tfl.gov.uk/tfl/downloads/pdf/congestion-charging/technology-%trials.pdf`, 2004.

[4] Royal Academy of Engineering. Transport 2050. `http://www.raeng.org.uk/policy/reports/pdf/Transport_2050.pdf`, 2005.

[5] Bryan Parno and Adrian Perrig. Challenges in securing vehicular networks. In *Proceedings of Workshop on Hot Topics in Networks (HotNets-IV)*, November 2005.

[6] Heathcote Williams. *Autogeddon*. Jonathan Cape, 1991.