Some Research Questions for Computational Transportation Science

D. Glenn Geers

Neville Roach Laboratory
National ICT Australia
223 Anzac Parade
Kensington NSW 2052 Australia
and
School of Computer Science and Engineering
University of New South Wales, Sydney, Australia
glenn.geers@nicta.com.au

ABSTRACT

In this paper some research questions which potentially lie within the purview of the new discipline of Computational Transportation Science are presented. Some of these questions are no doubt old and jaded and some may rightly reside in a different research field (or be seen as engineering) but viable solutions to all of the issues presented are of vital importance to the efficacy of future transportation systems.

Categories and Subject Descriptors

J.m [**Transportation**]: Computational Transportation Science—*communications, control, human factors*

General Terms

Proposed research directions

1. INTRODUCTION

Currently deployed ITS systems (with few exceptions) have very little intrinsic intelligence. A concrete example is afforded by considering an electronic tolling system which is of great benefit to road users but the installed hardware and software hardly deserves the epithet 'intelligent'. By combining artificial intelligence and machine learning techniques with wireless communications and distributed computing methods Computational Transportation Science has the potential of really making 'Intelligent' Transport systems intelligent.

Society as a whole is at a crossroads where connectivity between people, systems and devices of all types and at all times will soon be the norm rather than the exception. From the humble beginnings of the 'Internet Coke Machine' [2] at Carnegie Mellon University in the mid-1980s it is now not unreasonable to deploy systems consisting of hundreds of thousands or even millions of (largely identical) actively connected and cooperating devices that are capable of delivering useful (and some not so useful) services [1, 8]. More and more, such networked systems will enhance our daily routine and we will be less and less aware of their presence.

The following discussion is, of course, tainted by the author's whim and fancy for which he craves the reader's indulgence but makes no apology. The bulk of the material is concerned with support systems for road vehicles and personal travellers.

A quick glance at the end of the paper will reveal that the reference list is far from comprehensive. Judicious use of Google is recommended to ease the reader's desire for additional background material.

To put the paper in context I proffer the following story from the not-too-distant future (taken from [12]):

> Six o'clock and the alarm sounds. Sleepily you reach over switch it off and go back to sleep. Half an hour later you wearily remember it's Monday and you've got a meeting with an important client. Fortunately your car hasn't slept in and has already received the latest traffic data via the wireless infrastructure that has been provided by the city authorities. You stumble down the stairs, grab a coffee and head into the garage. Putting the key into the ignition silently brings your vehicle to life. A voice sounds in the cabin informing you that you only have 32 minutes (computers are still that literal in the future) to make it to your meeting and that you will be in time based on current traffic conditions. Five minutes into your journey, your car speaks up again to tell you that an accident ahead has caused bad congestion and your revised journey time will be another 62 minutes. Of your available options, your quickest is the train. Your car notes that it will take 9 minutes to drive to the station, where it has confirmed parking spaces are available. Then you have a two minute wait before 16 minutes travel time on the train and a 5 minute walk to the office. You'll only be 5 minutes late (humans still have no concept of time in the future). A couple of minutes into your train journey you receive a call from your client. He

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. *MobiQuitous 2008*, July 21-25, 2008, Dublin, Ireland. Copyright © 2008 ICST ISBN 978-963-9799-27-1

Digital Object Identifier: 10.4108/ICST.MOBIQUITOUS2008.3830
http://dx.doi.org/10.4108/ICST.MOBIQUITOUS2008.3830

> says he is stuck in one horrendous traffic jam. You just smile.

The outcomes of CTS research will lead to better ITS in the not-so-distant future.

2. USER INTERFACES

The old adage 'you can't judge a book by its cover' doesn't seem to apply for many software systems. The 'cover' is what the user must deal with on a day-to-day basis and any improvements in user interface technology will be a boon to both the software developer and user alike.

User interfaces in the transportation context are unique because they are more often-than-not windows into systems that are of secondary importance and should not interfere with the principal tasks of the user.

2.1 In-vehicle

The interface between the driver and the car has remained essentially fixed for the last 100 years. Evidently, vehicle interfaces must remain uniform and constant across all marques for considerable periods of time. Indeed, it is only recently that the first major intruder on the clean lines of the interface has raised its disconcertingly distracting head. The nagging 'after 200 metres turn left' issuing from the as yet not quite ubiquitous GPS navigation system can be both a blessing and a curse [3]. Clearly this interface is one-way. It is very difficult (and even dangerous) to update your GPS navigator whilst driving. And this is just the tip of the iceberg. Once dedicated short range communication (DSRC, 5.9 GHz ITS band) becomes an established technology the potential for driver distraction by the *vehicle* grows enormously.

It is evident that some form of robust two-way interface between in-car devices and the driver is needed. This interface needs to be robust enough to function reliably in the demanding conditions of a potentially noisy vehicle cabin and simple enough that the casual or elderly driver (who may need to access safety and other vehicle systems more than younger or more able drivers) is not intimidated or scared off. The latter point mentioned above is a particularly difficult challenge for the User Interface research community.

2.2 Transport Management Centre

Transport Management Centres (TMC) and incident management centres in general are stressful places in which to work. Operators must be able to access the functionality of a wide range of unrelated systems which have been added to their consoles in a seemingly *ad hoc* manner. It should be noted that even the layout of a console has an effect on operator efficiency [5].

Unification of legacy applications is possible through the use of middleware (perhaps with considerable effort) and future TMC software applications should be forced to adhere to any relevant ISO standard or ITS architecture requirements.

Use of novel interfaces should be considered: voice, gesture and eye-gaze being obvious candidates. Each has their own set of challenges.

In a noisy environment voice recognition remains a difficult problem. Repetition of spoken commands until they are understood may not endear such a system to an operator who needs to get the job done (now).

The use of both gesture and eye-gaze as input modalities require the use of video cameras. Aside from the (real or perceived) privacy issue the robust capture and interpretation of eye-gaze and gesture in a potentially cluttered and often dark environment remains an open problem.

On the whole, computer users will rapidly give up on using interfaces that do not have the same degree of robustness and reliability as the familiar keyboard and mouse.

3. TRAFFIC MANAGEMENT

Ultimately traffic management must address the issues of reducing congestion and improving the reliability and repeatability of travel whilst increasing safety and reducing environmental impact. They must manage the movement of people and goods from an origin to a destination not just by switching signals but by choosing (and updating) the route which may include urban roads and freeways.

An interesting question to ponder is just what level of congestion are people willing to tolerate? Or phrased differently: If an off-peak journey takes x minutes, what multiplicative factor, k, is tolerable to compute the rush-hour travel time, kx. Speculatively, k will depend on x in a highly non-linear way. If the functional form of k can be approximated it may be possible to estimate future traffic network requirements in a principled way.

3.1 Urban

Phase control (referred to as signal group control in the UK) of traffic signals is currently the dominant methodology for switching signals. Current state of the art traffic signal control systems as embodied by SCATS and SCOOT perform extremely well for a wide range of traffic circumstances and demands. Phase control, even when not rigorously enforced, is hampered by the sometimes unnecessary constraints that the methodology imposes on the traffic streams. An alternate approach, favoured in the Scandinavian countries, is referred to as *signal group control* (or stage control in the UK). Because the individual traffic signals are treated, insofar as possible, as maximally independent units signal group control opens up the possibility of greater finesse in fairly apportioning green-time to vehicles. The cost of increased control flexibility is the inability for a driver to predict when her light will go green because the traffic signals are no longer switched in a regular, periodic manner.

Both SCATS and SCOOT are able to adapt to the traffic demand, as measured by inductive loop detectors, by adjusting the green time appropriately (by a small deviation from a preset base time). Groups, typically of six or fewer junctions may be linked together so that their switching time is slaved ('offset') from the switching time of a so-called 'critical junctions'. At present critical junctions are determined by experienced traffic engineers.

The kerbside controllers used for both SCATS and SCOOT have minimal computational capacity and the calculation of green times and offsets is carried out at a central facility to which the controllers are connected over the public telecommunications (generally wired) infrastructure.

Digital Object Identifier: 10.4108/ICST.MOBIQUITOUS2008.3830
http://dx.doi.org/10.4108/ICST.MOBIQUITOUS2008.3830

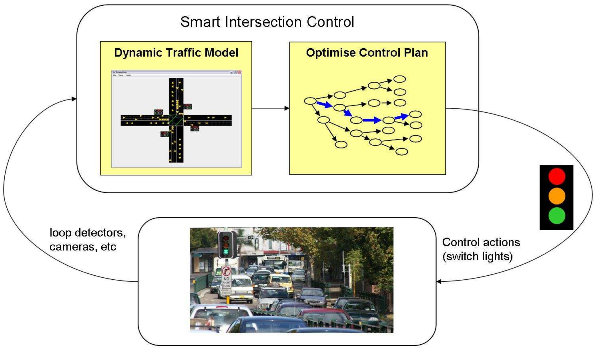

Figure 1: A model based traffic signal controller

Having described albeit extremely briefly, the current state of the art in traffic signal control it is possible to suggest some interesting directions for research, especially in the fields of machine learning and distributed computing. The traffic signal control systems of the future will also support a degree of robustness and redundancy that today's traffic engineers can only dream of. Public transport and emergency vehicle priority will be an inherent part of these new systems and it is conceivable that automatic detection of incidents at junctions and in links may be feasible.

The next generation of kerbside controllers will be much more capable than those presently deployed. This will enable the development of completely new algorithms for traffic signal control. Similarly, the devices from which the traffic control system is composed will have a much higher degree of connectivity. Any physical means of communication will be seamlessly supported by requiring that most of the devices at a junction be nodes on an IP (hopefully IPv6) network. Other devices will effectively have IP address by using IP-to-serial or IP-to-digital I/O, etc. conversion devices. Each lantern roundel will effectively have an IP address.

The systems will be self-healing, largely self-configuring and sensor agnostic. A kerbside controller at one junction will be capable of 'taking over' from a failed controller located elsewhere on the network. Using IP will ensure that control and other messages will always (with best effort) find a route from device-to-device. The traffic engineer will no longer have to identify critical junctions. The system will learn which junctions are critical and may adjust its choice dynamically (with a suitable time-constant). The number of junctions slaved off each critical junction will be chosen dynamically by the system and the offset will vary from junction-to-junction. Dynamically adjusting lane usage [14] at junctions will be common and coordinated throughout regions of the network. It is likely that the only settings required at a junction will be maximum red time for traffic signals and minimum green time for pedestrian walk signals—safety constraints and local traffic law (and lore) notwithstanding.

It is possible to imagine a system in which the kerbside controller actually maintains a model of the traffic in its local neighbourhood. Such a model-based junction control system is shown schematically in Figure 1. The local models are built from abstracted low-level traffic features that are ultimately derived from traffic sensors. Local neighbourhood models can be further abstracted and combined to produce regional traffic models with somewhat reduced fidelity. Use of a model hierarchy reduces the apparent complexity of the system, keeping it manageable and helping to tame the 'curse of dimensionality'. Of course, the model hierarchy can be continued upward to yield economic or environmental models of use to traffic planners and forecasters.

Modelling the traffic state enables the estimation of traffic information on road segments that are devoid of sensors. For example, a traffic model can track vehicles from one junction to the next (preempting their arrival with much greater accuracy than current systems) thereby enabling the downstream junction to take preemptive action to minimise delay. In circumstances of heavy traffic, a model can make better real-time estimates of queue length than a simple queue detection system that relies on a single detector at a fixed point.

Any new traffic signal control system should be able to work with any or all available traffic sensors which will also be distributed throughout the road network. Control should fall back gracefully as sensor inputs degrade in performance or become unavailable.

Clearly traffic signal control is an area where CTS can make a real and measurable difference in performance and reliability.

3.2 Freeway

Traffic management on freeways falls into three classes

1. Entry and exit control,
2. Speed control and
3. Incident detection.

Entry onto many freeways is controlled by some form of ramp metering—a special case of traffic signal control designed to ensure that the capacity of the freeway is not exceeded and that queueing onto the urban streets is minimised [13]. Traffic exiting a freeway has to interact with urban traffic. Exit ramps are often signalised to aid the merging of two clearly disparate traffic flow regimes. Neither of these cases are discussed further here.

Discussion of speed limiting on freeways will not be pursued.

Incident detection on freeways is traditionally carried out using the outputs from inductive loop sensors in conjunction with one of several comparatively old algorithms. The archetype of these algorithms is embodied in the so called California Algorithm of which #8 is the best performing [11]. These algorithms need to be tuned to the local operating environment which is a somewhat tedious operation.

Disregarding the fact that enhancing the traffic sensors will likely increase the incident detection and discrimination rate, it seems likely that modern statistical machine learning techniques as embodied by support vector machines (SVM) [4] would provide vast performance gains. In particular the so-called 1-class SVM [10], which is essentially a statistical tool for detecting the outliers of a learned distribution may provide a rapid and easy to deploy solution requiring minimal local tuning.

Digital Object Identifier: 10.4108/ICST.MOBIQUITOUS2008.3830
http://dx.doi.org/10.4108/ICST.MOBIQUITOUS2008.3830

4. DATA COLLECTION, MINING AND DISSEMINATION

4.1 The Sensor Revolution

Whilst DSRC will ultimately turn every vehicle into a sensor, there is considerable work left to be done in the 'classical' sensor domain. Indeed, the principal aim of research in this area should be to produce an above ground vehicle sensor system that has the performance, robustness, reliability and ease of setup exhibited by the venerable inductive loop. Inductive loops function perfectly well in the high arctic tundra of Northern Finland and the scorching outback of Central Australia. In order to be practical the cost of installation and maintenance of these new devices must be on par with that of the inductive loop.

Even when DSRC comes to dominate in the first world, there will still be scope for roadside sensors both in developing countries and in the infrastructure poor regions of the more well-to-do.

4.2 Dedicated Short Range Communications

Glossing over all of the wireless issues such as jamming, black hole routing, etc.; the biggest issue with DSRC seems likely to be the injection of false messages that come (or appear to come) from legitimate sources. Even if these messages are incapable of producing material harm they would still be capable of producing traffic network chaos. Some degree of resilience may be afforded by using secure positioning [6] but it is simple to conceive of a situation where secure positioning techniques would not solve the problem. Indeed construction of hardware designed to defeat proposed time-of-flight positioning schemes [7] does not seem infeasible for a determined attacker. Is it ethical to simply ignore a vehicle of indeterminate position that is crying out for help?

Such problems need to be solved robustly and reliably for DSRC to be seen as a trusted and safe vehicular communication platform. Come-what-may DSRC will surely play a role in the future of ITS.

4.3 'Transport Radar'

Ultimately the sensor network of which the traveller is an intrinsic (and perhaps unwilling) part will enable the computation of origin-destination (OD) information for all travellers both before departure and during travel (each traveller may well be a part of a transport system model of suitable fidelity). Such a system will be a kind of 'transport radar' keeping the traveller (or goods carrier) informed of the current best means of reaching his destination in a given time or via the shortest route (or with the least cost or . . .).

Presenting the route and modal choices to a traveller both before and (dynamically) throughout their journey will require the development of novel hardware and software. It is conceivable that such functionality will be built into the ever increasingly less phone-like mobile phone. Such systems should be context aware [9] (for example, a textual interface is a poor choice for the driver of a vehicle which is engaged in the traffic stream—and may be illegal—but perfectly legitimate for a passenger in the same vehicle), position aware and current transport-mode aware. Infrastructure and in-vehicle support will be required to realise the full potential of such a system and leveraging off a burgeoning DSRC-base would appear prudent.

5. STANDARDS AND TEST SITES

5.1 Standards

John Donne wrote that 'No man is an island' and in no case is this more evident than when sitting in a traffic jam! However, the sentiment carries across to the practical application of the outcomes of CTS as embodied in ITS deployment on the kerbside. No deployed system will exist in isolation. The only way to ensure seamless interoperability is to lay down standard APIs and communication protocols to which systems from different vendors must adhere. This does not eliminate competition because such standards (as realised as APIs, etc.) must be open and extensible in a proprietary manner. In the parlance of object oriented programming, the standards will define the 'base classes' from which extended systems may be seamlessly 'derived'.

5.2 Test Sites

Simulation can only take you so far. Ultimately the algorithms and systems that we develop in our research have to be deployed on the streets. However, the 'streets are tough' and the CTS community will need to interact with the owners of kerbside infrastructure (usually government), data (government and private) and vehicles (usually private).

Wouldn't it be nice to just be able to install some equipment at the kerbside with (almost) no questions asked? (Safety will be an issue and any organisation that allows kerbside installation will *always* have their own fallback systems in place).

What the CTS (and ITS) community needs is a town sized test facility that is large enough to be representative yet small enough to be understandable. Requirements might include a significant number (say 50) of signalised junctions, a major railhead (and/or moderate sized port) and a servicing freeway or arterial. Such a site would allow real world comparison of different systems over extended periods of time.

6. CONCLUSION

The discussion above barely scratches the surface of potential research problems for CTS. Those espoused tend to have a practical bent as indeed they should. What CTS must do is develop the next generation of ITS by leveraging the latest techniques from machine learning, distributed systems, constraints programming and a host of other CS disciplines.

The inevitable and, seemingly inexorable growth in computing power and networking technologies will enable the transport and traffic engineers of the future to realise the power of grid computing and to run real-time (and even supra-realtime) transport network models on computing systems that are intrinsically part of the transport network being modelled. The model hierarchy described in Section 3.1 above will be mirrored in the computing environment on which the model executes.

Transportation needs are driven by the requirements of societal mobility (as embodied in work requirements, family and vacations), goods delivery and (more recently) environmental sustainability. It is the demand for consumer goods coupled with the need for sustainability that is pushing up the cost of petroleum. The increase in petroleum price will in turn, hinder societal mobility.

For millennia people were restricted in their potential for travel by the lack of means. If governments are not careful this past and,

Digital Object Identifier: 10.4108/ICST.MOBIQUITOUS2008.3830
http://dx.doi.org/10.4108/ICST.MOBIQUITOUS2008.3830

extremely restrictive state of affairs could return—at the very least for overseas travel.

The rising cost of petroleum should expedite the development of non-petroleum powered road vehicles by industry and the proactive development of public transport infrastructure by government. It is evident, however, that the personal automobile will not disappear. The freedom it provides will be too much for society to lose.

Future trips to the Pyramids or Stonehenge or the *22nd International Conference on Computational Transportation Science*(!) may well be virtual but we should all sleep soundly secure in the knowledge that traffic jams will be around for a good few centuries yet.

7. ACKNOWLEDGEMENTS

The author would like to thank Dr Bernhard Hengst and Mr Geoff Goeldner for useful discussions about the finer points of future transportation systems during the preparation of this paper.

Thanks go to the Roads and Traffic Authority of New South Wales without whom this paper would not have been possible.

NICTA is funded by the Australian Federal Government as represented by the Department of Broadband, Communications and the Digital Economy, the NSW Department of State and Regional Development, the ACT Government and the Australian Research Council through the ICT Centre of Excellence Program.

8. REFERENCES

[1] Berkeley open infrastructure for network computing. `http://boinc.berkeley.edu/`.
[2] The internet coke machine at cmu. `http://www-cse.ucsd.edu/users/bsy/coke.history.txt`.
[3] Obedient driver follows orders, crashes car. Reuters, October 24 2006.
[4] C. J. C. Burges. A tutorial on support vector machines for pattern recognition. *Data Mining and Knowledge Discovery*, 2(2):121–167, 1998.
[5] B. J. C. Grady T. Holman and R. L. Bulfin. Using linear programming to optimize control panel design from an ergonomics perspective. In *PROCEEDINGS of the HUMAN FACTORS AND ERGONOMICS SOCIETY 47th ANNUAL MEETING*, pages 1317–1321, 2003.
[6] C. Laurendeau and M. Barbeau. Threats to security in dsrc/wave. In *ADHOC-NOW*, pages 266–279, 2006.
[7] L. Lazos, R. Poovendran, and S. Čapkun. Rope: robust position estimation in wireless sensor networks. In *IPSN '05: Proceedings of the 4th international symposium on Information processing in sensor networks*, page 43, Piscataway, NJ, USA, 2005. IEEE Press.
[8] B. McCarty. Botnets: Big and bigger. *IEEE Security and Privacy*, 1(4):87–90, July-August 2003.
[9] B. Schilit, N. Adams, and R. Want. Context-aware computing applications. In *IEEE Workshop on Mobile Computing Systems and Applications*, Santa Cruz, CA, US, 1994.
[10] B. Schölkopf, J. C. Platt, J. C. Shawe-Taylor, A. J. Smola, and R. C. Williamson. Estimating the support of a high-dimensional distribution. *Neural Comput.*, 13(7):1443–1471, 2001.
[11] Z. K. Simon Cohen. A calibration process for automatic incident detection algorithms. *Microcomputers in Transportation*, 7:254, 1993.
[12] J. Steel and G. Geers. Jump in my car: *it* wants to take you home. In *Smart Traffic 2006*, 22-23 August 2006.
[13] M. van den Berg, B. De Schutter, and H. Hellendoorn. Effects of on-ramp and off-ramp metering on queue forming in urban traffic networks. In *Proceedings of the 11th IFAC Symposium on Control in Transportation Systems*, pages 130–135, Delft, The Netherlands, Aug. 2006.
[14] C. K. Wong and S. C. Wong. A lane-based optimization method for minimizing delay at isolated signal-controlled junctions. *J. Math. Model. Algorithms*, 2(4):379–406, 2003.

Digital Object Identifier: 10.4108/ICST.MOBIQUITOUS2008.3830
http://dx.doi.org/10.4108/ICST.MOBIQUITOUS2008.3830