Ethical Principles for the Design of Next-Generation Traffic Control Technology

Virginia Tech ISCE Applied Ethics Initiative Sponsored Paper
September 1, 2014

Tristram McPherson, PhD
Philosophy,
Virginia Tech
tristram@vt.edu

Miloš N. Mladenović, PhD
Civil and Environmental Engineering,
Virginia Tech
milosm@vt.edu

1. Introduction

Transportation will be revolutionized in the near future by the convergence of developments in sensing, communication, and in-vehicle computing technology (Birdsall, 2013; Özgüner, Acarman, & Redmill, 2011). Developments in sensing technology will enable real-time collection of data that includes geographical coordinates, speeds, direction of movement, acceleration, obstacles, etc. (U. Lee & Gerla, 2010). Developments in communication technology will permit short-range vehicle-to-vehicle and vehicle-to-infrastructure communication (Bell, 2006). This will enable vehicles to share crucial information in real time (Misener, Dickey, VanderWerf, & Sengupta, 2009). Finally, in-vehicle computing capacity is growing at a rapid pace. For example, vehicles today can possess higher power reserves and can store larger amounts of data, compared to a typical mobile computer (Leen & Heffernan, 2002; Misener et al., 2009).

These converging technologies are enabling the development of self-driving vehicles, which will be able to monitor roadway and traffic conditions, and to perform all safety-critical driving functions ("Preliminary Statement of Policy Concerning Automated Vehicles," 2013). This emerging technology promises several potential benefits:

- to improve traffic safety, by replacing less reliable human driving (Delot, Cenerario, & Ilarri, 2010; "The Safety Promise and Challenge of Automotive Electronics: Insights from Unintended Acceleration," 2012);
- to improve the mobility of people unable to drive (Anderson et al., 2014);
- to mitigate the environmental impacts of automotive transportation and improve energy efficiency ("Self-driving cars: The next revolution," 2013);
- to more efficiently use existing roadways, and hence reduce the need for new infrastructure (Anderson et al., 2014).

The advent of self-driving vehicle (SDV) technology also permits significant evolution in the possible structure and mechanisms of traffic control (Mladenović, Abbas, & McPherson, 2014). This paper investigates the ethical dimension of the design of the next generation of traffic control mechanisms, focusing on one core aspect of this design problem: intersection traffic control.

The ethical dimensions of traffic control technology design could easily be neglected for at least three reasons. First, developing traffic control technology might suffer from 'design inertia' – maintaining design assumptions that are tailored to earlier technology. Second, the ethical significance of traffic control technology is not easily visible. For example, someone waiting in a traffic queue is likely to feel annoyed rather than morally indignant. Third, the development of traffic control might be driven largely by market incentives to the neglect of ethical considerations.

This paper aims to combat that neglect in several stages. First, (§2) we present an overview of the development and character of contemporary traffic control technology ("C1"), followed (§3) by an overview of the possible operational principles for next-generation traffic control technology ("C2"). We then introduce some of the values relevant to traffic control technology design (§4), explain the need for an ethical perspective on such design (§5) and introduce two leading ethical theories with which those values can be organized and applied (§6). We then use these ethical theories to offer a preliminary assessment of three potential operational principles for intersection traffic control (§7). We conclude by summarizing the importance of the questions that remain to be investigated in design for next-generation traffic control, and by offering a brief case for an approach to such design that involves deep sensitivity to relevant ethical issues, and a large space for public participation in assessing these ethical issues as they bear on such design (§8).

2. Overview of conventional traffic control

This section explains the historical context in which modern traffic control technology was initially implemented, and traces the development of that technology. In addition, we will explain the key operating principles of C1 (contemporary traffic control technology), and some key effects of those principles.

The need to control movement through intersections existed long before the introduction of motor vehicles (Grubler, 1990; Lay & Vance, 1999). Congestion and safety problems – and the use of traffic control devices and access control to manage them – were present even on the ancient Roman roads (Mueller, 1970). However, the systemic need for traffic control increased dramatically in the late 19th century. This was due to three developments: significant increases in urban population (Grant, 2003; Taylor & David, 1951); economic development, which increased demand for transportation (Grubler, 1990; Jones, 2008); and the manufacture of internal combustion engine vehicles (Lay & Vance, 1999; McShane, 1999). Together, these developments led to deteriorating user safety at intersections.

Traffic control devices were initially developed as safer and less labor-intensive alternatives to the use of police officers, who were initially responsible for controlling traffic at critical intersections (Gazis, 2002; Lay & Vance, 1999). Like the officers it replaced, traffic control technology aimed to affect the behavior of the human beings who controlled the vehicles entering intersections. For this reason, they were developed to command attention, convey a clear and simple meaning, command respect, and give adequate time for response ("Manual on Uniform Traffic Control Devices," 2009).

Arguably the most significant development in traffic control technology was the invention of the illuminated traffic signal ("A History of Traffic Control Devices," 1980; Mueller, 1970). This initial engineering solution to the problem of intersection control introduced separation of conflicting flows using fixed periods of displayed green/red lights to improve the safety of users. The introduction of traffic signals imposes strong social control over human movement (McShane, 1999): it markedly restricts human movement by actively prohibiting or allowing crossing of intersection during specific time intervals (Orcutt Jr, 1993).

This crucial operating principle – allowing alternating access to an intersection from different approaches for discrete intervals – was maintained by later developments in traffic signal control technology. These further developments – such as microprocessor technology – greatly enhanced the possibilities for automated traffic control operations. For example, this development allowed the introduction of cyclical repetition of the relevant time intervals, and coordination of operation between nearby intersections.

Despite these developments, there are significant technical limitations to the ability of C1 traffic control to provide optimal system-wide outcomes. Most notably, increasing traffic volume can result in negative effects such as substantial delay. Minimizing the aggregate amount of such negative effects is the aim of optimal signal control system/strategy design (Papageorgiou, Diakaki, Dinopoulou, Kotsialos, & Wang, 2003). The negative effects addressed in this literature include delay (traveller utility), travel time and travel time reliability (traveller utility), stops (traveller utility), crashes (safety), fuel consumption (out of pocket cost), emissions (environment), etc. (Gordon, 2010; Hartenstein, Laberteaux, & Ebrary, 2010; *Traffic Engineering Handbook - 6th ed.*, 2008). However, the foundational constraints introduced by C1 traffic control technology make it extremely difficult to successfully implement optimizing solutions on traffic signal network in real time (Papageorgiou et al., 2003). Even the advanced traffic control technology, such as adaptive traffic control systems, require human supervision during operation and do not guarantee benefits (M. Mladenovic, 2012; M. Mladenovic & Abbas, 2012).

The ultimate aim of optimal signal control design is to optimize the attainment of certain goods. However, in virtue of not distinguishing the relative priority of individual vehicles in a systematic way, sophisticated C1 will inevitably fail to do this. For example, vehicles approaching an intersection from the same approach all experience the same probable delay, and C1 control widely utilizes the concepts of the major and minor road, where larger amount of green time is dedicated to major approach (Daganzo, 1997; Gartner & Stamatiadis, 2009). These outcomes can lead to suboptimal outcomes because some user delays are much more important than others, and C1 makes no distinction between users with more versus less time-sensitive needs. The only exception to this is that C1 principles give emergency vehicles priority, by enabling them to override the direction provided by traffic signals ("Manual on Uniform Traffic Control Devices," 2009).

More abstractly, we can think of the development of traffic control was envisioned as a set of engineering solutions to an emerging issue of controlling traffic flow through intersections. The scope of possible solutions was constrained by a series of crucial assumptions about the nature of the problem and of possible solutions. These assumptions include:

1. A focus on influencing human vehicle operators;
2. Using static rules or time-based separation of aggregated conflicting traffic flows;
3. Externalized control of intersections: vehicle operators lack significant input to the control process;
4. Specialized consideration only for emergency vehicles; and
5. Determining control optimality based on hypotheses about aggregate effects.

It is important to make these assumptions explicit, in light of the technological developments mentioned at the beginning of this paper. These developments call into question the appropriateness for the design of next-generation (“C2) traffic control of each of the assumptions just mentioned:

1. Self-driving vehicle and vehicle-infrastructure (“V2I”) communication technology entail that intersection control can potentially most efficiently directly signal a self-driving vehicle, rather than a human operator.
2. The computing power available in next generation vehicles makes possible real-time dynamic assignment of intersection access to individual vehicles or platoons.
3. Vehicle-vehicle (“V2V”) and V2I technology open up the potential that vehicle user input (for example concerning the importance of avoiding delay) can affect the control process, by helping to determine right-of-way.
4. This possibility of user input thus means that relative user priority can be given consideration for all vehicles, not just emergency vehicles.
5. The availability of information about user needs opens up the possibility of in-principle alternatives to counting aggregate effects as the criterion for successful intersection control. (Much more on this point in §6)

These points put the reader in a position to critically examine the substantial body of existing research on intersection control for self-driving vehicles (“SDVs”), briefly summarized in the next section.

3. Survey of research concerning C2 operational principles

There have been efforts to develop new frameworks for intersection control for SDVs for almost two decades. The chronological development of control mechanisms for SDVs is presented in Table 1.

Table 1: Principles of operation for self-driving vehicle control at intersections

Authors / Year	Principle of operation
(Naumann, Rasche, & Tacken, 1998)	Right of way is assigned based on delay, number of vehicle in the queue, and approaching velocity.
(Dresner & Stone, 2004)	Reservation of space-time based on first-in first-out (FIFO) principle, according to the time of request for reservation.
(Dresner & Stone, 2006)	FIFO, but emergency vehicle receives right-of-way by clearing the lane for that vehicle.
(Raravi, Shingde, Ramamritham, & Bharadia, 2007)	Minimizing the maximum travel time to the intersection.
(Schepperle, Böhm, & Forster, 2007)	FIFO for initial reservation, but a vehicle can exchange time-slots with another vehicle if that other vehicle pays.
(Schepperle & Böhm, 2007)	Basic variant: auction for time-slot with the vehicle with the highest bid receiving the right-of-way. Variant with subsidies: The candidate with highest accumulated bid receives the right-of-way.
(Vasirani & Ossowski, 2008)	Driver agents must purchase the necessary reservations from the intersection manager agents. Intersection manager "sells" the right-of-way in attempt to maximize profit.
(VanMiddlesworth, Dresner, & Stone, 2008)	FIFO
(Regele, 2008)	Predefined right-of-way for certain movements over other movements through the intersection.
(Vasirani & Ossowski, 2009)	Longest in the system: vehicle with the earliest arrival Shortest in the system: vehicle with the latest arrival Farthest to go: vehicle with the longest path to the destination Nearest to source: vehicle closest to its origin
(Yan, Dridi, & El Moudni, 2009)	Minimizing the time a vehicle takes to cross the intersection.
(de La Fortelle, 2010)	FIFO
(Milanés, Pérez, Onieva, & González, 2010)	Subject vehicle yields to the vehicle on the right.
(Makarem & Gillet, 2011)	Heavier vehicles with higher effect of inertia during velocity adjustment are given an indirect priority.
(Alonso et al., 2011)	Resolving conflict based on the classification of the road (otherwise FIFO).
(Ghaffarian, Fathy, & Soryani, 2012)	Maximize traffic throughput based on waiting delay or queue length.
(J. Lee & Park, 2012)	Gap adjustment mechanism for minimizing the total length of overlapping trajectories.
(Makarem, Pham, Dumont, & Gillet, 2012)	Priority determined by the distance to the intersection.
(Zohdy, Kamalanathsharma, & Rakha, 2012)	Minimization of total delay.
(Park & Lee, 2012)	Priority assigned to the lane with the longest queue, or if the vehicle reaches certain waiting period.
(Li, Chitturi, Zheng, Bill, & Noyce, 2013)	FIFO, but with priority reservation of vehicle in queue that is above certain length.
(M. Mladenovic & Abbas, 2013)	Priority queuing principle where each vehicle has assigned individual priority level (e.g., based on number of passengers).

(Makarem & Gillet, 2013)	Minimizing the sum of deviation from desired velocity and energy consumption.

These principles might have technical drawbacks, especially for large-scale network implementation, where issues of scalability and redundancy are important. However, we want to emphasize a different, ethical critique of the assumptions behind this research. This research does not take into consideration the social relations underlying traffic as a phenomenon. Operation based on a predefined "static" rule that uses the approaching link or predefined order of service bluntly neglects individuals' needs for crossing the intersection. Similarly, characterizing optimality as minimizing aggregate negative effects (such as delay) neglects the possibly widely varying importance of individual user needs. Consider a person waiting on the "minor" intersection approach, on his way to the emergency room, while all the people on the "main" approaches are going shopping. The person on the minor approach might be forced to wait considerably by any operational principle that aims solely at minimizing total delay. The operating principles considered in the research listed above do not include a mechanism to obtain and include the information on specific trip purpose and desired arrival time of each individual.

Another critical point is that operational principles are developed without consideration of relations between technology and human behavior. Despite being a technical phenomenon, technology is also a social phenomenon, since it is both a terminus and creator of context for human experience (Ihde, 1990). Consequently, technology is not a value-neutral tool but a force that conditions human agency (Borgmann, 1987), and can be even used to favor certain social classes (Bianchini & Avila, 2014). Evaluating previous research from this perspective, it is important to emphasize that the operational principles listed above do not take into consideration human behavior in relation to technology (e.g., they neglect human altruistic behavior, and do not limit the pursuit of self-interest, etc.).

4. Values relevant to traffic control technology design

Because we are all so familiar with the existing traffic control paradigm, it is easy for us to take its ethical significance for granted. Moreover, considering that intersections are ubiquitous, the principles of operation for future traffic control technology for self-driving vehicle will impact almost every single individual, throughout their lifetime. Ideally, intersections and their control technology should be pure public goods (Minasian, 1967), providing benefits to all people that are non-excludable (no one can be deprived of access to the intersection) and non-rival (one vehicle's having access to the intersection does not 'use up' such access: all can share in the access).

In this context, traffic control has implications for several important values, including at least:

- Mobility and accessibility
- Safety
- Environmental sustainability
- Privacy
- Broadly aesthetic considerations

We will briefly discuss each of these values in turn.

From a transportation engineering perspective, m*obility* is defined as “the ability of people and goods to move quickly, easily, and cheaply to where they are destined at a speed that represents free-flow or comparably high-quality conditions”. *Accessibility* is defined as “the achievement of travel objectives within time limits regarded as acceptable” (Lomax et al., 1997). Essentially, mobility relates to ease of movement on the transportation network, and accessibility relates to ease of reaching destination in space. Consequently, both mobility and accessibility are enormously important when considering C2 technology development. This is arguably largely because it enables us to pursue other goals: we use roads to get to work, to get food, and to access friends and leisure activities. However, the relationship between mobility/accessibility parameters is not as straightforward as it might seem. On the one hand, there is an assumption of a desire of each individual is to have less delay while crossing the intersection. This logic is similar to the utility functions from the Highway Capacity Manual ("The Highway Capacity Manual 2010," 2011), where Level of Service depends on the amount of time spent in the queue. On the contrary, people have previous experience with travel and they accept certain amount of their day as a dedicated travel time budget (Goodwin, 1981; van Wee, Rietveld, & Meurs, 2006).

Safety is another central motive for traffic control. Automotive travel risks death and other extremely serious harms. For example, over 30,000 people are killed and over 1,600,000 are injured in automotive accidents each year in United States ("Traffic Safety Facts 2012 ", 2014). Consequently, it is in the public interest to mitigate those harms.

A third important value relevant to traffic control design is *environmental sustainability*. Consider several examples where transportation is associated with environmental externalities (Rodrigue, Comtois, & Slack, 2013). First, both the construction and use of roadways involve significant use of non-renewable resources, such as fossil fuels, or have impact on soil quality. Second, roadways can significantly impact wildlife mobility and biodiversity. Third, transportation typically involves significant impacts on air quality, noise, and water quality. Locally, this can cause or exacerbate human disease. Globally, this can contribute to global warming, which threatens to do massive damage to global human well-being. Well-designed traffic control systems have the potential to lower the environmental externalities of human movement.

A fourth important value relevant to the development of traffic control technology is the user's *privacy*. The operation of traffic control systems requires data collection, as a part of the control process. Even in the simplest case, where data collection consists only from traffic count on a certain road section, there is a potential for determining user identity. Moreover, considering that the operation of the system can be increased by increasing information quality (for example, knowing the route of each user's trip), collecting this information can result in potential user privacy concerns (Wright, 2011). Considering the capabilities of self-driving vehicles in collecting, storing, and transmitting user-sensitive information, there will surely be user privacy concerns that will need to be addressed for SDV technology in general and in relation to traffic control.

A final value relevant to traffic control design is broadly *aesthetic*. On the one hand, movement can potentially be an enjoyable and even meaningful part of our lives. To see this, think of the money and care people lavish on their vehicles, in order to have a specific sort of experience in their transportation, and of the symbolic significance many people invest in owning their own vehicle (Grieco & Urry, 2011), as opposed to using a well-functioning public transportation system. On the other hand, movement can be a largely frustrating and stressful experience, and traffic control can contribute to this. For example, a commute is more likely to be stressful and frustrating if the flow of traffic is regularly subject to significant delays than if congestion varies widely and unpredictably from day to day.

5. The need for an ethical perspective in traffic control technology design

The brief list of values canvassed in the previous section is surely incomplete, but it suffices to orient us to some of the considerations relevant to the development of traffic control technology. However, identifying the relevant values only takes us part of the way to being able to ethically evaluate potential paradigms for traffic control. We also need to consider competing ethical frameworks that can provide a function from those values to the assessment of concrete policy options.

To clarify the issue, consider an example. Suppose that you are waiting at a red traffic signal. Together with the accompanying rule of law, traffic signals allow or forbid the movement of users at a specific intersection point at specific moment in time (Daganzo, 1997). In light of this, if vehicles are crossing the intersection from the conflicting approach, your obeying the signal protects both your safety and that of the crossing vehicles. However, while your safety is protected, the red signal also restricts your movement as long as it is active. You are "stuck" in traffic, and this increases travel time to your destination.

Stepping back, we can notice that you and other users of the intersection have various needs and interests for crossing the intersection. By allowing some vehicles through the intersection before others, the traffic signal in effect prioritizes the satisfaction of some of these needs and interests. The fundamental ethical task for traffic control technology is to prioritize the satisfaction of these needs and interests in a way that is just. It is perhaps uncontroversial that a traffic control system should avoid imposing significant harms, and should contribute to the common good (Ramírez & Seco, 2011). This principle implies that engineers developing this technology should seek to serve the public interest, with regard for safety, health, and public welfare, while actively preventing conditions that are threatening to life, limb, or property ("Canons of Ethics for Members," 2003; "Code of Ethics," 2010; "Code of Ethics for Engineers," 2007; "IEEE Code of Ethics," 2006).[1]

The uncontroversial principle just noted – avoid harm, contribute to the common good – can answer some traffic control design questions. For example, certain traffic control malfunctions (such as simultaneous right-of-way for conflicting approaches) pose grave risks that should clearly be avoided by failsafe mechanisms. And some traffic control solutions are clearly superior to others in light of this principle. For example, at a large intersection with high traffic volumes, a well-functioning traffic signal is far more efficient than a four-way stop. However, the uncontroversial principle leaves other important questions open. Many of these questions arise most strikingly with the availability of C2 technology. For example, C2 technology can potentially collect and utilize substantial information not available in C1 systems, such as information about the purpose of a vehicle's trip. Such information, once available, is clearly ethically relevant. For example the importance of timely mobility is very different to someone on a vacation than it is to someone driving a woman in labor to the hospital. Determining just ways of making use of such information requires that we move beyond the uncontroversial principle, to more informative conceptions of justice.

6. Two organizing ethical theories

Theories of social justice provide a framework for assessing the distribution of advantages and disadvantages in a society, through a set of rules that distinguish between just and unjust actions or institutions (Miller, 1999). Debate concerning the relative merits of competing theories of social justice is marked by deep and persistent controversy. In light of this, this section introduces two of the most influential competing approaches to social justice: a *consequentialist* approach, and an approach inspired by John Rawls' theory of justice as fairness. The consequentialist and Rawlsian approaches are far from the only

---

[1] In addition, the IEEE Code of Ethics states that engineers should aim to improve their understanding of the technology, its appropriate application, and the likely consequences of its implementation ("IEEE Code of Ethics," 2006).

leading ethical theories in light of which we might evaluate traffic control frameworks.[2] However, we do take them to be (a) among the most important, (b) among the easiest to understand and apply to traffic control, and (c) instructively different.

The consequentialist approach, as we will characterize it, is marked by four ideas.[3]

1. Consequentialism is *impartial* in the following sense: according to the consequentialist, if certain degrees of happiness or safety are intrinsically valuable, they are valuable to the same degree whether they accrue to me or to you.
2. Consequentialism is concerned with *net value*. For example, an action that makes Sally very happy and Paul a bit unhappy (and does nothing else) produces net happiness. Similarly, installing a standard traffic light in a previously uncontrolled intersection may increase the net safety of that intersection, even if it causes occasional dangerous behavior that would not otherwise have occurred (such as someone driving recklessly when they see a yellow light).
3. Consequentialism is concerned with *all* of the effects of an action or policy: if a traffic control system would increase safety, but somehow make everyone sick, it is a bad system. Or if it would work well for a few years, and then fail catastrophically, it is again a bad system.
4. Consequentialism is a *maximizing* theory: it instructs us to choose the actions or policies that produce the most net value. So, even if conventional traffic control has very good effects compared to uncontrolled traffic, if a certain next-generation traffic control system would produce more net value than conventional control, consequentialism would instruct us to move to the next generation system.

With this gloss in hand, we can say: consequentialism instructs us to choose the traffic control system, of those available to us, that maximizes our power of movement, and safety, and the efficient use of resources (both in creation and implementation), the environmental sustainability of our transportation system, and the aesthetic qualities of our transportation experiences. This gloss, however, obscures one of the deep theoretical questions that we need to confront in traffic control design: the problem of *value comparison*. Certain choices that we have would increase the instantiation of one value while decreasing another. A simple example is the trade-off between speed of movement

---

[2] Other important frameworks from contemporary ethics and political philosophy which could be applied to assess traffic control include virtue ethics (e.g. (Hursthouse, 1999)), care ethics (e.g. (Engster, 2007; Held, 1995), libertarianism (e.g. (Nozick, 1974), egalitarianism (e.g. Cohen 2011), and the capabilities approach (e.g. (Nussbaum, 2007; Sen, 2006). For an approach to transportation planning influenced by Walzer's 'spheres of justice' approach (Walzer, 1983), see (Martens, Golub, & Robinson, 2012). For a relation between transport and Sen's capabilities approach see (Beyazit, 2011). In addition, some of the engineering fields have already recognized the need for including ethical reflections in the design process (Azath, Wahida Banu, & Neela Madheswari, 2011; Beamon, 2005; Curiel-Esparza, Canto-Perello, & Calvo, 2004; Grodzinsky, 2000; Kleijnen, 2011; Moor, 2005; Shelley, 2012).

[3] These assumptions are characteristic of some of the most influential versions of consequentialism. However, there are versions of consequentialism that relax each of these assumptions. For an excellent discussion of the forms that consequentialist theories can take, see (Kagan, 1998).

and safety: the setting of speed limits essentially involves a judgment about the relative value of timely movement versus that of reducing the risk of a traffic accident.

An important alternative framework for thinking about the ethics of traffic control is suggested by the two central principles of John Rawls' theory of justice. Paraphrasing Rawls slightly (Rawls, 1999):

1. Each person has equal right to the most extensive degree of the basic liberties, compatible with others sharing those liberties to the same degree.
2. Inequalities of other goods are justified just in case they are to the advantage of the least well-off, and are associated with positions that are open to all persons.

The first principle refers to certain basic liberties. These include rights to vote and to hold public office, to speech and assembly, from psychological oppression and physical assault, the right to hold property, and freedom from arbitrary arrest (Rawls, 1999). According to Rawls, the first principle has *priority* over the second in the following sense: it would be unjust to limit basic liberties in order to better satisfy the second principle.

The Rawlsian account can be instructively contrasted with the consequentialist one. Because the consequentialist account aims to maximize *net value*, it permits the possibility of making one person considerably worse off than average, because doing so will make many other individuals better off than average. The Rawlsian approach denies that this is permissible. With respect to the basic liberties, the first principle is strictly egalitarian: it insists that each person is entitled to the *same* liberties. With respect to the other goods, Rawls insists that inequalities can be justified only if they are to the advantage of the least well off.

Rawls makes a sharp distinction between those ethical theories which apply to individual behavior, and those which apply to social and political institutions. Rawls' theory of justice is intended as the latter sort of theory. According to Rawls, his two principles characterize a just *basic structure* for a society, where this encompasses:

> the way in which the major social institutions distribute fundamental rights and duties and determine the division of advantages from social cooperation. By major institutions I understand the political constitution and the principal economic and social arrangements (Rawls, 1999).

It is not clear that traffic control counts as part of the Rawlsian basic structure. However, there is a clear case for it being so: traffic control is a central part of the regulation of access to a range of crucially important *public spaces*. Given the centrality and importance of the values at stake in the regulation of these spaces (including those canvassed in §4), we take it to be quite plausible that traffic control counts as one of the central social arrangements of a community. However, our interest here is not to defend a claim about Rawls: even if applying Rawls' two principles to traffic control extends Rawls' ideas beyond his own intentions, we take it to be both plausible and instructive. This is an instance of a more general point: theories of social justice were not typically developed with an eye to assessing technological development. Consequently, it would be inappropriate to assume that principles that these theories establish can directly be used for technological development. Nonetheless, we take these theories to be an enormously helpful in developing ethical principles to guide technological development, provided that their limitations are kept in view.

In extending the Rawlsian account to apply to traffic control, we take freedom of movement, safety, and privacy to be basic liberties that fall under the first principle. We take the other goods mentioned in §4 to be among the social and economic goods that fall under the second principle. In light of the priority of the first principle, the Rawlsian approach thus suggests that securing freedom of movement, safety, and privacy should have priority over these other goods. Further, the first principle suggests an *egalitarian* approach: as much as is possible, everyone should have equal freedom of movement through public spaces, equally protected privacy, and should be able to move equally safely. The application of the second principle suggests that traffic control should be designed such that the other goods – environmental protection, and the aesthetic benefits of transportation – are arranged to the greatest advantage of the least well-off.

7. Preliminary ethical assessment of competing traffic control principles

So far, we have briefly introduced conventional traffic control principles, the alternative principles for traffic control of self-driving vehicles that have thus far been proposed in the literature, and two major ethical frameworks for evaluating these principles. This puts us in a position to offer a preliminary sketch of how the ethical assessment of alternative operational principles for controlling intersection access could proceed. This sketch will of necessity be simplified, and point the way to how more detailed assessment might proceed in the future. Our simplified sketch will consider three operational principles we take to be important. We will consider how these principles apply to a simplified scenario, and then consider some broader issues arising from the simplified scenario.

To begin, consider three competing operational principles for next-generation traffic control involving self-driving vehicles with sophisticated on-board computing and communication technology:

OP1 These principles incorporate all of the C1 design assumptions (listed at the end of §2) except for the first. Intersection control communicates directly with the self-driving vehicle, rather than with a human operator. However, traffic is still controlled by giving access to the intersection to all vehicles on a given approach for a certain period of time. The cycle of periods of time allotted to various approaches are determined by hypotheses concerning what will minimize aggregate delay at the intersection. Control of the intersection is externalized: drivers entering the intersection have no input into the function that determines priority access, and only emergency vehicles are given priority access.

OP2 These principles incorporate two of the C1 design assumptions. Special priority access is given only to dedicated emergency vehicles, and the aim is to maximize certain aggregate effects, in this case user utility. Intersection control communicates both with vehicle occupants as well as directly with the self-driving vehicle. Crucially, this system uses dynamic rules with user input to determine priority access to the intersection: priority access to the intersection is determined by a real time monetary auction that all vehicles on conflicting approaches can participate in.

OP3 These principles reject all of the design assumptions of C1. OP3 rejects the idea that the goal of intersection control is to maximize overall effects. Rather, OP3 is designed to give priority to those with the greatest need for priority access to the intersection, while preserving fair equality of access. As with OP2, intersection control communicates both with vehicle occupants as well as directly with the self-driving vehicle. Also as with OP2, this system uses dynamic rules with user input to determine priority access to the intersection. However, this system replaces auctions with a system of priority credits. These credits are made available to all users on a fair basis (for example, there might be a baseline of equality with extra credits available to persons with medical needs, e.g., who more often require speedy transportation). Priority access to the intersection is determined by the priority credits apportioned by vehicle users. Emergency vehicles can command overriding priority access to the intersection by using emergency priority credits. However other users will be able to claim emergency priority for certain trips, with a post-trip verification system that would punish misuse.

Each of these principles could be spelled out in more detail. However, consider how these principles would adjudicate a concrete traffic control scenario: in this scenario, four vehicles are approaching a four-leg intersection and going through (Figure 1). The driver of each vehicle has certain importance (priority) for her trip purpose (ranked from high to low). For example, driver of the vehicle 1 might be rushing to the hospital, while on the opposite driver 4 might be going to the grocery store. Each of these four self-driving vehicles are approaching an intersection approximately at the same time.

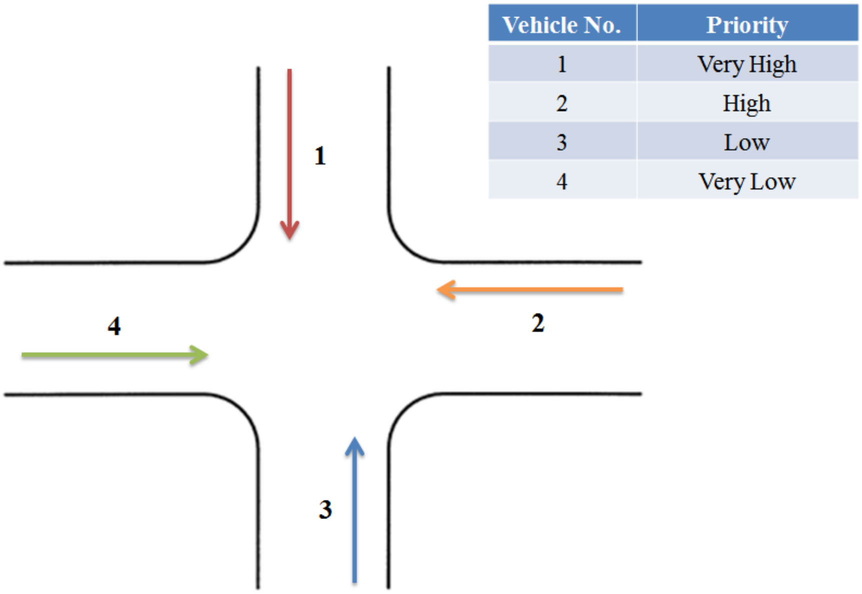

Vehicle No.	Priority
1	Very High
2	High
3	Low
4	Very Low

Figure 1: Intersection layout and vehicle's priorities

Consider how OP1 will address this case. Since all vehicles approach simultaneously, it will make a random choice between allowing vehicles 1 and 3 to cross the intersection first (while 2 and 4 wait) or vice versa. Compared to C1, OP1 will deliver significant improvement: for example, given the fast responsiveness and reliability of SDVs compared to human drivers, access to the intersection can be given to the waiting vehicles as soon as their trajectories would not conflict with the vehicles already in the intersection, rather than building in extra waiting time to safely manage human unreliability as C1 does. And, at least in this scenario, OP1 has the ability to successfully minimize aggregate delay.

There are, however, several objections to OP1 that this scenario illuminates. From a consequentialist perspective, aggregate delay is plausibly only a proxy for more intrinsically valuable goods, and OP1 looks worse relative to these goods. For example, OP1 clearly does not maximize aggregate user utility, because doing so would likely require giving priority to Vehicle 1 then Vehicle 2, which OP1 cannot do. This is because OP1 does not use information about user priority into determining intersection access. OP1 also fails relative to the Rawlsian principles. First, it fails to ensure that we maximize the priority access of the neediest vehicle (Vehicle 1), since information about vehicle need is not an input. Second, OP1 arguably fails to secure equal safety: if one vehicle needs priority access in order to secure its occupant's safety (as in a private person driving to the hospital in a medical emergency), then in virtue of being indifferent to this fact, OP1 fails to secure that user's safety. Indeed, in such an emergency, the user may be tempted to override the intersection's prescriptions (if the SDV makes this possible), potentially endangering others in her effort to secure her own safety.

OP2 aims to maximize user utility, and uses local auctions as a mechanism for doing so. The idea is roughly that the amount users are willing to pay for priority access to the intersection serves as a reasonable proxy for the utility of that priority access to those users. Where that assumption is correct – where the user of vehicle 1 will pay more (either in real or virtual currency) than the user of vehicle 2 etc. – an auction will lead to priority access to the intersection tracking the degree of utility of that access, thereby maximizing utility. From a consequentialist perspective, the central concern about this approach is that the assumption may not always be correct. For example, against a background of economic inequality, wealthy users may always be able to pay for priority access (whatever the utility of their having it) while the poorest users may not be able to afford such access (however desperate their need for access). In our example, if the user of Vehicle 4 was wealthy, and the user of Vehicle 1 was poor, the auction system in OP2 might in fact do worse relative to aggregate utility than OP1! One might hope to address this worry by including transfers in the auction: perhaps some portion of what the highest bidder pays goes to the lowest bidder. In such a system, the economic benefit of losing the auction might mitigate the cost to the poorest user of the delay she suffers. From a Rawlsian perspective, the worry about OP2 is even deeper: against a background of inequality of wealth, the auction system produces an unequal distribution of freedom of movement through public spaces: the rich can always secure priority access over the poor. This threatens to replicate historically familiar aristocratic norms, where the lower class person at an intersection was required to make way for his 'superior'. From this perspective, even OP1 is arguably superior to OP2 against the background of economic inequality.

OP3 aims to secure priority access to the intersection for those with greatest need, by using a system of priority credits distributed fairly to all. In our example, provided user 1 offers more credits than user 2, etc., then priority access will be given to the vehicles commensurately with their need. From a consequentialist perspective, OP3 will optimize, given these assumptions: aggregate utility will be maximized (other things being equal) because the priority access given to users will reflect the relative importance of such priority access to the users. From a Rawlsian perspective, the second principle is respected: priority access is given to the user in greatest need. Further, the objection facing OP2 does not arise: the rich cannot simply buy priority access, so freedom of movement (approximated by access to priority credits) is equalized between users (unless users have special needs that warrant assigning them more credits).

The main questions facing OP3 concern feasible implementation. Some of the issues include:

- Worries that priority credits would ultimately become 'fungible' with money, thereby threatening to make OP3 a less efficient version of OP2.
- Worries about user ability to learn how to make competent choices with an unfamiliar tool such as priority credits. It was crucial to the reasoning about OP3 that user assignments of credits tended to track relative user priority; if this conjecture about human psychology failed, then the case for OP3 is undermined. Indeed, widespread poor management of priority credits could entail that priority

access tended to go to the credit-wise, not those most in need of access to the intersection.
- Worries that priority credits could function similarly to small closed currencies, which are subject to generally harmful recession-like phenomena

If these worries can be ameliorated, we take OP3 to be the most ethically promising approach of those canvassed, at least relative to the artificially simplified considerations that we have set out here.

8. Summary and recommendations for further investigations

Our central aim in this paper aim is to broaden the design horizon for the self-driving vehicle technology (Mladenović et al., 2014). As the preceding discussion makes clear, the relative assessment of operational principles for C2 traffic control is a very complex matter. In this article, we have simplified the issues greatly. For example:

- We focused only on a very simplified discussion of single intersection, setting aside the crucial complexities that arise when considering design impacts at larger scales.
- We focused on three of many possible operational principles: one a minimal adaptation of C1 control (OP1), one a simple auction-based system (OP2), and one a system based on dedicated priority credits (OP3), and our exposition of each of these options left out many important details.
- We focused on two ethical principles: a consequentialist principle and a Rawlsian principle. Our discussion of these principles was simplified, and we did not discuss any of the several important alternatives to these principles in detail.

While simplified, we take this discussion to be instructive in several respects:

1. Abandoning the assumptions that have thus far guided traffic control design allows us to consider alternative approaches that are at least very promising alternatives to the design status quo, and may be far more sensitive to important ethical considerations.
2. Incorporating substantial ethical considerations into design has the potential to make a profound difference to that design.
3. Differing ethical theories have significantly differing implications for these design issues. For example, it may be feasible to tweak OP2 in ways that make it attractive to the consequentialist, without addressing the Rawlsian worries about it.

With these results in hand, let us speculate more broadly about directions for future research in this area. We want to emphasize three points: the pressing need for such research, the need to think systematically about the ethical dimensions of design in such research, and the need for such research to include broadly participatory elements.

On the first point: there has, so far, been only limited research efforts that tried to compare different traffic control principles for self-driving vehicles using computer simulation (Vasirani & Ossowski, 2009). In addition, there has been limited research that has tried to draw inspiration for development from political theory (M. N. Mladenovic, 2014). However, the potential impact of such research is hard to exaggerate: C2 is coming, and it will affect all of our lives on a daily basis. In light of this, the questions very broadly broached in this paper are still very much open for groundbreaking and practically important research.

On the second point: as we have tried to show in this paper, ethical principles are deeply relevant to assessing competing visions of C2. Indeed, a central goal of this paper is to emphasize that development of technology should not focus solely on instrumental functionality, but also needs to include human behavior and societal values into its design vision (Cohen & Grace, 1994; Heikkerö, 2012; Verbeek, 2006). For example, one important question is how much mobility/accessibility is required for a well-functioning civil society. Consequently, such design should attempt to consider the complete range of ethically relevant impacts of a new technology, by determining the cultural or social losses it threatens (Flanagan, Howe, & Nissenbaum, 2008). This in turn requires that technological development proactively incorporates ethical investigation into design (compare Value Sensitive Design: (Friedman, Kahn Jr, & Borning, 2006; J Van den Hoven, 2005; J. Van den Hoven, 2007)). The focus of this design approach is on deliberately incorporating ethical values into technological design, while meeting traditional design criteria (e.g., user safety, communication reliability, etc.). The search for ethically superior designs involves an iterative process of conceptual, empirical, and technical investigations (Genus, 2006; J Van den Hoven, 2005)

On the third point: as we have noted, the ethical principles at stake in C2 design are also subject to deep and reasonable controversy. This fact, together with the scope of the potential impact of this technology, leads us to emphasize the importance of including broad public participation in the development of C2. While engineers clearly have an important role in the design, decision-making power should not be concentrated solely in a small group of experts, especially if they are influenced exclusively by financial interests. This paper suggests that we cannot reduce practical questions about the good life to technical problems for experts, and we thus cannot eliminate the need for public and democratic discussion of the relevant societal values (Habermas, 1971). As a result, there is a need to transparently engage all relevant societal constituencies in critical conversations and decision-making about technology development.

REFERENCES

Alonso, J., Milanés, V., Pérez, J., Onieva, E., González, C., & De Pedro, T. (2011). Autonomous vehicle control systems for safe crossroads. *Transportation Research Part C: Emerging Technologies, 19*(6), 1095-1110.

Anderson, J. M., Nidhi, K., Stanley, K. D., Sorensen, P., Samaras, C., & Oluwatola, O. A. (2014). *Autonomous Vehicle Technology: A Guide for Policymakers*: Rand Corporation.

Azath, M., Wahida Banu, R., & Neela Madheswari, A. (2011). Improving Fairness in Network Traffic by Controlling Congestion and Unresponsive Flows. *Advances in Network Security and Applications*, 356-363.

Beamon, B. M. (2005). Environmental and sustainability ethics in supply chain management. *Science and Engineering Ethics, 11*(2), 221-234.

Bell, M. G. H. (2006). Policy issues for the future intelligent road transport infrastructure. *Centre for Transport Studies, Imperial College London*.

Beyazit, E. (2011). Evaluating social justice in transport: lessons to be learned from the capability approach. *Transport reviews, 31*(1), 117-134.

Bianchini, D., & Avila, I. (2014). Smart Cities and Their Smart Decisions: Ethical Considerations. *IEEE Technology and Society Magazine, 33*(1), 34-40.

Birdsall, M. S. (2013). Connected Vehicle: Moving from Research to Deployment Scenarios. *ITE Journal, 83*(1).

Borgmann, A. (1987). *Technology and the character of contemporary life: A philosophical inquiry*: University of Chicago Press.

. Canons of Ethics for Members. (2003): Institute of Transportation Engineers.

. Code of Ethics. (2010): American Society of Civil Engineers.

. Code of Ethics for Engineers. (2007): National Society of Professional Engineers.

Cohen, S., & Grace, D. (1994). Engineers and social responsibility: An obligation to do good. *Technology and Society Magazine, IEEE, 13*(3), 12.

Curiel-Esparza, J., Canto-Perello, J., & Calvo, M. A. (2004). Establishing sustainable strategies in urban underground engineering. *Science and Engineering Ethics, 10*(3), 523-530.

Daganzo, C. F. (1997). *Fundamentals of transportation and traffic operations*.

de La Fortelle, A. (2010). *Analysis of reservation algorithms for cooperative planning at intersections*. Paper presented at the 13th International IEEE Conference on Intelligent Transportation Systems (ITSC).

Delot, T., Cenerario, N., & Ilarri, S. (2010). Vehicular event sharing with a mobile peer-to-peer architecture. *Transportation Research Part C: Emerging Technologies, 18*(4), 584-598.

Dresner, K., & Stone, P. (2004). *Multiagent traffic management: A reservation-based intersection control mechanism*.

Dresner, K., & Stone, P. (2006). *Human-usable and emergency vehicle-aware control policies for autonomous intersection management*.

Engster, D. (2007). The heart of justice: Care ethics and political theory.

Flanagan, M., Howe, D., & Nissenbaum, H. (2008). Embodying values in technology: Theory and practice. *Information technology and moral philosophy*, 322-353.

Friedman, B., Kahn Jr, P. H., & Borning, A. (2006). Value sensitive design and information systems. *Human-computer interaction in management information systems: Foundations, 5*, 348-372.

Gartner, N., & Stamatiadis, C. (2009). Traffic Networks, Optimization and Control of Urban. In R. A. Meyers (Ed.), *Encyclopedia of Complexity and Systems Science-v. 1-10*.

Gazis, D. C. (2002). *Traffic theory* (Vol. 50): Kluwer Academic Publishers.

Genus, A. (2006). Rethinking constructive technology assessment as democratic, reflective, discourse. *Technological Forecasting and Social Change, 73*(1), 13-26.

Ghaffarian, H., Fathy, M., & Soryani, M. (2012). Vehicular ad hoc networks enabled traffic controller for removing traffic lights in isolated intersections based on integer linear programming. *Intelligent Transport Systems, IET, 6*(2), 115-123.

Goodwin, P. (1981). The usefulness of travel budgets. *Transportation Research Part A: General, 15*(1), 97-106.

Gordon, R. L. (2010). *Traffic signal retiming practices in the United States* (Vol. 409): Transportation Research Board.

Grant, H. R. (2003). *Getting around: exploring transportation history*: Krieger Pub Co.

Grieco, M., & Urry, J. (2011). *Mobilities: new perspectives on transport and society*: Ashgate Publishing, Ltd.

Grodzinsky, F. S. (2000). Equity of access: Adaptive technology. *Science and Engineering Ethics, 6*(2), 221-234.

Grubler, A. (1990). *The Rise and Fall of Infrastructures: Dynamics of Evolution and Technological Change in Transport*. Haidelberg, Germany: Physica-Verlag

Habermas, J. (1971). *Toward a rational society: Student protest, science, and politics* (Vol. 404): Beacon Press.

Hartenstein, H., Laberteaux, K., & Ebrary, I. (2010). *VANET: vehicular applications and inter-networking technologies*: Wiley Online Library.

Heikkerö, T. (2012). *Ethics in Technology: A Philosophical Study*: Lexington Books.

Held, V. (1995). Justice and care.

. The Highway Capacity Manual 2010. (2011). Washington, D.C. : Transportation Research Board.

. A History of Traffic Control Devices. (1980). Washington, D.C.: Institute of Transportation Engineers.

Hursthouse, R. (1999). *On virtue ethics*: Oxford University Press.

. IEEE Code of Ethics. (2006): Institute of Electrical and Electronics Engineers.

Ihde, D. (1990). *Technology and the lifeworld: From garden to earth*: Indiana University Press.

Jones, D. W. (2008). *Mass motorization and mass transit*: Indiana Univ Pr.

Kagan, S. (1998). *Normative ethics*: Westview Press Boulder, CO.

Kleijnen, J. (2011). Ethical issues in engineering models: an operations researcher's reflections. *Science and Engineering Ethics, 17*(3), 539-552.

Lay, M. G., & Vance, J. E. (1999). *Ways of the World: A History of the World's Roads and of the Vehicles that Used Them*: Rutgers University Press.

Lee, J., & Park, B. (2012). Development and evaluation of a cooperative vehicle intersection control algorithm under the connected vehicles environment. *Intelligent Transportation Systems, IEEE Transactions on, 13*(1), 81-90.

Lee, U., & Gerla, M. (2010). A survey of urban vehicular sensing platforms. *Computer Networks, 54*(4), 527-544.

Leen, G., & Heffernan, D. (2002). Expanding automotive electronic systems. *Computer, 35*(1), 88-93.

Li, Z., Chitturi, M. V., Zheng, D., Bill, A. R., & Noyce, D. A. (2013). *A Next-Generation Intersection Control Algorithm for Autonomous Vehicles*. Paper presented at the 92nd Annual Meeting of Transportation Research Board

Lomax, T., Turner, S., Shunk, G., Levinson, H. S., Pratt, R. H., Bay, P. N., & Douglas, G. B. (1997). *Quantifying Congestion. Volume 1: Final Report*: Transportation Research Board.

Makarem, L., & Gillet, D. (2011). *Decentralized Coordination of Autonomous Vehicles at Intersections.* Paper presented at the World Congress.

Makarem, L., & Gillet, D. (2013). *Model predictive coordination of autonomous vehicles crossing intersections*. Paper presented at the 16th International IEEE Annual Conference on Intelligent Transportation Systems (ITSC 2013), Hague, Netherlands.

Makarem, L., Pham, M.-H., Dumont, A.-G., & Gillet, D. (2012). Microsimulation Modeling of Coordination of Automated Guided Vehicles at Intersections. *Transportation Research Record: Journal of the Transportation Research Board, 2324*(1), 119-124.

. Manual on Uniform Traffic Control Devices. (2009): Federal Highway Administration, Department of Transportation.

Martens, K., Golub, A., & Robinson, G. (2012). A justice-theoretic approach to the distribution of transportation benefits: Implications for transportation planning practice in the United States. *Transportation Research Part A: Policy and Practice, 46*(4), 684-695.

McShane, C. (1999). The origins and globalization of traffic control signals. *Journal of Urban History, 25*(3), 379-404.

Milanés, V., Pérez, J., Onieva, E., & González, C. (2010). Controller for urban intersections based on wireless communications and fuzzy logic. *IEEE Transactions on Intelligent Transportation Systems, 11*(1), 243-248.

Miller, D. (1999). *Principles of social justice*: Harvard University Press.

Minasian, J. R. (1967). Public goods in theory and practice revisited. *Journal of Law and Economics*, 205-207.

Misener, J. A., Dickey, S., VanderWerf, J., & Sengupta, R. (2009). Vehicle-Infrastructure Cooperation. In S. Olariu & M. A. C. Weigle (Eds.), *Vehicular networks: from theory to practice* (Vol. 20): Chapman & Hall/CRC.

Mladenovic, M. (2012). Large Scale Analysis of Traffic Control Systems. *Traffic Engineering and Control, 53*(1).

Mladenovic, M., & Abbas, M. (2012). A Guide to Effective Adaptive Traffic Control Systems. *Traffic Engineering and Control, 53*(2).

Mladenovic, M., & Abbas, M. (2013). *Self-organizing Control Framework for Driverless Vehicles*. Paper presented at the 16th International IEEE Conference on Intelligent Transport Systems, Hague, Nederlands

Mladenovic, M. N. (2014). *Development of Sustainable Control Principles for Self-driving Vehicles: A Paradigm Shift within the Framework of Social Justice.* (Ph.D.), Virginia Tech.

Mladenović, M. N., Abbas, M., & McPherson, T. (2014). *Development of Socially Sustainable Traffic-Control Principles for Self-Driving Vehicles: The Ethics of Anthropocentric Design.* Paper presented at the IEEE International Symposium on Ethics in Engineering, Science, and Technology.

Moor, J. H. (2005). Why we need better ethics for emerging technologies. *Ethics and Information Technology, 7*(3), 111-119.

Mueller, E. A. (1970). Aspects of the history of traffic signals. *Vehicular Technology, IEEE Transactions on, 19*(1), 6-17.

Naumann, R., Rasche, R., & Tacken, J. (1998). Managing autonomous vehicles at intersections. *Intelligent Systems and their Applications, IEEE, 13*(3), 82-86.

Nozick, R. (1974). *Anarchy, state, and utopia* (Vol. 5038): Basic books.

Nussbaum, M. (2007). Frontiers of justice. Disability, nationality, species membership.

Orcutt Jr, F. L. (1993). *The Traffic Signal Book*.

Özgüner, Ü., Acarman, T., & Redmill, K. (2011). *Autonomous Ground Vehicle*. Norwood, MA: Artech House.

Papageorgiou, M., Diakaki, C., Dinopoulou, V., Kotsialos, A., & Wang, Y. (2003). Review of road traffic control strategies. *Proceedings of the IEEE, 91*(12), 2043-2067.

Park, C., & Lee, J. (2012). *Intelligent Traffic Control Based on Multi-armed Bandit and Wireless Scheduling Techniques.* Paper presented at the VEHICULAR 2012, The First International Conference on Advances in Vehicular Systems, Technologies and Applications.

. Preliminary Statement of Policy Concerning Automated Vehicles. (2013) *Washington, DC*: National Highway Traffic Safety Administration, .

Ramírez, F., & Seco, A. (2011). Civil Engineering at the Crossroads in the Twenty-First Century. *Science and Engineering Ethics*, 1-7.

Raravi, G., Shingde, V., Ramamritham, K., & Bharadia, J. (2007). Merge algorithms for intelligent vehicles *Next Generation Design and Verification Methodologies for Distributed Embedded Control Systems* (pp. 51-65): Springer.

Rawls, J. (1999). *A Theory of Justice, rev. ed* (Vol. 5). Cambridge, MA: Belknap.

Regele, R. (2008). *Using ontology-based traffic models for more efficient decision making of autonomous vehicles.* Paper presented at the Autonomic and Autonomous Systems, 2008. ICAS 2008. Fourth International Conference on Autonomic and Autonomous Systems.

Rodrigue, J.-P., Comtois, C., & Slack, B. (2013). *The geography of transport systems*: Routledge.

. The Safety Promise and Challenge of Automotive Electronics: Insights from Unintended Acceleration. (2012): Committee on Electronic Vehicle Controls and Unintended Acceleration, Transportation Research Board, Board on Energy and Environmental Systems, Computer Science and Telecommunications Board, .

Schepperle, H., & Böhm, K. (2007). Agent-based traffic control using auctions. *Cooperative Information Agents XI*, 119-133.

Schepperle, H., Böhm, K., & Forster, S. (2007). *Towards valuation-aware agent-based traffic control*.

. Self-driving cars: The next revolution. (2013): KPMG LLP, Center for Automotive Research

Sen, A. (2006). *Inequality Reexamined* Oxford University Press.

Shelley, C. (2012). Fairness in Technological Design. *Science and Engineering Ethics*, 1-18.

Taylor, G. R., & David, H. (1951). *The transportation revolution, 1815-1860* (Vol. 4): Rinehart New York, NY.

Traffic Engineering Handbook - 6th ed. (2008). Intitute of Transportation Engineers.

. Traffic Safety Facts 2012 (2014) (Vol. 811, pp. 700): National Highway Traffic Safety Administration; US Department of Transportation.

Van den Hoven, J. (2005). Design for values and values for design. *Information Age, 4*, 4-7.

Van den Hoven, J. (2007). ICT and value sensitive design. *The Information Society: Innovation, Legitimacy, Ethics and Democracy In honor of Professor Jacques Berleur sj*, 67-72.

van Wee, B., Rietveld, P., & Meurs, H. (2006). Is average daily travel time expenditure constant? In search of explanations for an increase in average travel time. *Journal of transport geography, 14*(2), 109-122.

VanMiddlesworth, M., Dresner, K., & Stone, P. (2008). *Replacing the stop sign: Unmanaged intersection control for autonomous vehicles*.

Vasirani, M., & Ossowski, S. (2008). *Towards reservation-based intersection coordination: an economic approach.* Paper presented at the 11th International IEEE Conference on Intelligent Transportation Systems.

Vasirani, M., & Ossowski, S. (2009). *Evaluating policies for reservation-based intersection control*.

Verbeek, P.-P. (2006). Materializing morality design ethics and technological mediation. *Science, Technology & Human Values, 31*(3), 361-380.

Walzer, M. (1983). *Spheres of justice: A defense of pluralism and equality*: Basic Books.

Wright, D. (2011). A framework for the ethical impact assessment of information technology. *Ethics and information technology, 13*(3), 199-226.

Yan, F., Dridi, M., & El Moudni, A. (2009). *Autonomous vehicle sequencing algorithm at isolated intersections.* Paper presented at the Intelligent Transportation Systems, 2009. ITSC'09. 12th International IEEE Conference on.

Zohdy, I. H., Kamalanathsharma, R. K., & Rakha, H. (2012). *Intersection management for autonomous vehicles using iCACC.* Paper presented at the Intelligent Transportation Systems (ITSC), 2012 15th International IEEE Conference on.