# Automated Collaboration among Communicating, Semiautonomous Vehicles

Dan Chevion[1], Ron Sivan[1], Onn Shehory[1], and Yuval Shimony[1]

[1] IBM Haifa Research Lab, Haifa, Israel 31905
{chevion, rsivan, shehory, yshimony}@il.ibm.com

**Abstract.** The general trend in combating traffic congestion and reducing accidents has shifted from paving more roads to making better use of existing infrastructure, often via technological improvements. Focusing on information technology, this paper combines several techniques in an innovative way to help communicating vehicles traverse conflict zones, such as merges and intersections, in an orderly fashion and at cruising speeds. This is done without taking away driver control of the vehicle, a property that reduces the automaker's liability exposure.

The first technique, known as tracking or Adaptive Cruise Control (ACC), can prevent a vehicle from getting too close to another vehicle. ACC is already available commercially for non-communicating vehicles. However, the reliability and benefit of tracking can be improved significantly once vehicles can communicate. The second technique introduced here, known as generalized tracking, allows communicating vehicles to track other vehicles not in their immediate vicinity, as long as those other vehicles are within communication range. This technique improves the ability of vehicles to coordinate road sharing among themselves. The third technique, known as traversal ordering, allows vehicles approaching a conflict zone to agree upon the order in which they are to traverse it. Combining these three techniques, we show how communicating vehicles can arrange themselves and cross conflict zones as quickly and safely as they do when traveling through a regular, non-conflict zone.

The fourth technique introduces acceleration into conflict zone traversal. From the perspective of flow, conflict zones can be viewed as a narrowing of the road, where the roads leading to it are capable of carrying more traffic than the conflict zone itself. This can, and often does, produce a backup that no amount of coordination can avoid. However, if vehicles are made to judiciously accelerate as they approach the conflict zone and decelerate once they have crossed it, flow through the conflict zone can be increased without affecting traffic outside of it. Safety is not jeopardized as the generalized tracking technique is applied as well.

We demonstrate the feasibility of these four techniques when applied to the case of a blocked lane via simulations.

**Categories and Subject Descriptors.** I.2.9, I.2.11.

**General Terms.** Algorithms, Design.

**Keywords.** Transportation, Automotive, V2V, V2I, Collaborative driving.

## 1 Introduction

Vehicular traffic is one of the more prominent woes of the modern industrialized world: accidents, congestion, and pollution are all major sources of concern. These problems are also interrelated—motor accidents are often the cause for congestion, and congested traffic pollutes much more than moving traffic does.

For many years, growth in traffic needs was addressed by increasing supply: wider roads, complex interchanges, etc. The return on such investments is declining, however, and newer solutions are being sought. Several recent solutions seek to control demand, such as designated roads or lanes (e.g., for public transportation) and prohibited zones (e.g., city centers), to name a few. Despite these efforts, vehicular transportation problems seem to keep growing. Much research is being conducted to find novel, long-term solutions that would permit better utilization of the existing infrastructure.

An alternative approach to alleviating the problem is to increase infrastructure utilization. Such an approach requires that the density of vehicles on the road be increased. Making the vehicles themselves smaller could help, but much more can be achieved by reducing the time and space margins vehicles usually keep between themselves for safety, since these margins take up much of the available road real estate. For example, a one-second distance between cars on the highway could account for 6 car lengths when moving at 100 km/hour; much time is wasted in deciding who goes next at merges and intersections and in other similar situations.

To increase infrastructure utilization without compromising safety, information technology solutions are required. Indeed, several solutions are being studied and some are already available commercially (e.g., Adaptive Cruise Control [2] ). Yet most of these solutions do not assume communication, let alone collaboration, between vehicles. The study presented in this paper aims at solutions based on communication and collaboration among semiautonomous vehicles. We introduce several innovative techniques, some rather counter-intuitive, and combine them into a collaborative driving mechanism (Section 3.5). When applied to autonomous communicating vehicles, this mechanism and the underlying techniques solve well problems of the sort listed above.

In particular, the collaborative driving we present aims to facilitate fast and smooth passage of vehicles through segments of road where the straight-forward advance of vehicles is frustrated by other vehicles vying for the same space. Examples of such shared road stretches are intersections, lanes used for bypassing, merge areas, and lanes where vehicles wish to move at different speeds – all common situations in everyday traffic. Such a shared road stretch will henceforth be referred to as a conflict zone (CZ). CZs are the cause for many, and perhaps most traffic delays; vehicles must slow and often stop completely when negotiating passage through a CZ. They are also the site of 25% - 40% of all accidents [1] [5] .

Collaborative driving of the sort studied in this research relies on vehicle-to-vehicle and vehicle-to-infrastructure communications. Although this dependence raises issues of deployment, such systems may be practical even in the near future in controlled environments, where only communicating vehicles are permitted. Operational vehicles in a factory yard, or cabs for self-driving within an airport, are examples of such controlled environments.

Focusing on automated collaboration, this paper introduces several innovative techniques that, when combined, allow communicating vehicles to traverse CZs in an orderly fashion and at cruising speeds. The proposed solution places an emphasis on safety, reducing the chance of accidents.

Although the collaboration studied here is automated, its implementation is compatible with leaving a human driver ultimately in control of the vehicle. When driver actions interfere with the automated behavior, the system adjusts itself to deal with the new situation gracefully; the worst outcome would be some delay in traffic. We believe that this may make collaboration more acceptable both to drivers, who seek control, and manufacturers, who abhor the liability.

Section 2 lays the conceptual foundations for our collaborative driving mechanism. Section 3 delves into implementation details while Section 4 describes two possible scenarios to serve as examples. Section 5 describes previous work in the field, and Section 6 concludes.

## 2 Basic Concepts

In this section we introduce the basic concepts of four inter-vehicle coordination techniques. Implementation details of these techniques are presented in Section 3, and their combination into one mechanism is presented in Section 3.5.

### 2.1 Tracking

Tracking is used to maintain a safe distance between consecutive vehicles moving in the same direction in the same lane. This is not a new concept; it is already available commercially, and is described here only for the sake of thoroughness. A tracking vehicle behaves normally (i.e., in the same manner as contemporary vehicles do) when no other vehicle is ahead. Upon approaching another vehicle from behind, the vehicle is automatically constrained to move no faster than the vehicle it is approaching. This is considered a driving aid, helping drivers on highways and heavy traffic situations. It should be noted that the tracking capabilities now available in some high-end car models is sensor-based, using radars, lasers, or cameras mounted on the tracking vehicle to determine its relative distance from nearby vehicles. The present study, however, is concerned only with tracking based on inter-vehicle communications.

### 2.2 Generalized Tracking

As mentioned above, tracking may be achieved without communications, using vehicle-mounted sensors. Inter-vehicle communication, however, permits a vehicle to keep a *generalized distance* from any vehicle it can communicate with, whether in front, nearby or elsewhere, if only such generalized distance is defined. For example, the generalized distance between two vehicles approaching the same CZ could be defined as the difference between their distances to that CZ; this boils down to the normal notion of distance when the two vehicles happen to travel along the same lane, but is different if they are not. The vehicle doing the tracking in such a situation is said to be *logically* tracking the other. This is similar to virtual vehicle mapping in [14] . From the driver's perspective, generalized tracking feels like ordinary tracking: the vehicle's acceleration is limited by the presence of some vehicle ahead. There could be cases, however, where the driver cannot see or is not aware of which of the other vehicles her own vehicle is tracking.

Since any change in the state of a vehicle can quickly be made known to all vehicles that are (physically or logically) behind it, even an abrupt stop is broadcast early enough to give affected vehicles enough time to slow down and avoid collision. The safety of the system is therefore unrelated to the speed with which participating vehicles are moving; their responses are not different from what they would have been had all vehicles been physically tracking. Any vehicle beginning to brake notifies all those behind it, which begin braking concurrently as a result. Due to the communication links, the fact that the row is only logical makes no difference.

### 2.3 Traversal Order

Currently, the standard way to allow vehicles through CZs is by time-division: vehicles take turns and cross the CZ one at a time. This sequencing could be imposed by rules of traffic (such as right of way) or by an external timing device (such as a street light). In any event, the essential aspect of these various solutions is the imposition of a sequence, or *traversal order* (henceforth TO) by which vehicles pass the CZ one after the other. This traversal order is currently figured out on-the-fly by the drivers, through watching the traffic and the relevant road signs. This procedure suffers from several flaws:

1. All drivers must decide on the same traversal order, or they might run into each other.
2. Calculating the traversal order takes time, and drivers slow down in the vicinity of CZs to give themselves enough time to figure out the correct traversal order.
3. When an external timing device is used, time is wasted when the right of way is given to a direction from which no traffic is coming. Moreover, timing devices often allow for some idle time when switching directions (a green light is given in one direction only some time after the red was turned on for all conflicting directions), adding to the delay they incur.
4. Time and energy are wasted when vehicles slow down at a stop sign or a street light, just to accelerate back to their original speed when given the right of way.

The approach we adopt here is to figure out the traversal order *in advance*. Once the order is decided and is disseminated to the vehicles involved, both the risk of misunderstanding could be eliminated and the time wasted on the slowdown and decision making could be saved. Note that this is different from scheduling, since only the order of the vehicles is determined, not their arrival times; there is no "missing one's slot" in this scheme.

### 2.4 Speed-up

The last observation has to do with the math of traffic flow. Let the following definitions hold:

- Flow (*f*): the number of vehicles crossing a given point in a unit of time, e.g., one hour (h)
- Density (*d*): the number of vehicles per unit length, e.g., one kilometer (km)
- Speed (*s*): the distance vehicles traverse per unit of time, e.g., kilometers per hour (km/h)
- Inter-vehicle gap (*g*): the space between two consecutive vehicles traveling in the same direction in the same lane

Using the definitions above, the following two mathematical relations hold (see [11] for example):

1. Flow is proportional to both density and speed. In particular, on average:

$$f = d \times s$$

2. Density and inter-vehicle gap are inversely proportional. In mathematical terms, if the average vehicle length is *l*, then

$$g = \frac{1}{d} - l$$

From item 1, we see that for a given flow, there is a tradeoff between density and speed; in particular, density may be reduced if speed is increased without changing the flow. From item 2, we see that reducing density increases inter-vehicle gap. Taken together, this means that by increasing travel speed, inter-vehicle gap may be increased without affecting the overall flow of vehicles.

Raising the overall speed of traffic is both difficult and risky. However, here the change is localized to a small area. If each vehicle in that area accelerates and then decelerates back to its original speed, the rates at which traffic enters and leaves the area remain the same as when no speed-up is employed, and yet the gaps between vehicles at the center of that area are widened, and can be used to accommodate interleaving vehicle flows.

## 3 Implementation

### 3.1 Implementing Generalized Tracking

Generalized tracking is realized similarly to ordinary tracking, with three main differences:

1. Information regarding the position of the tracked vehicle is most likely obtained through vehicle-to-vehicle communication (V2V).
2. The identity of the tracked vehicle is not self-evident and must be supplied.
3. The method by which distance to the tracked vehicle is measured must also be defined.

This information is sufficient for logic aboard the tracking vehicle to calculate the generalized distance at any time, based on its knowledge of the positions of itself and the tracked vehicle. We describe below two scenarios in detail, including their respective rules for measuring generalized distances (see Sections 4.1 and 4.2). In both cases, generalized tracking is employed only in the vicinity of CZs, where an area controller (see Section 3.4) is available to coordinate traffic and can provide this needed information.

### 3.2 Implementing Traversal Order

Traversal order lists (TOs) are maintained for each CZ and list all the vehicles in the vicinity that wish to cross it in the order decided (see Section 2.3). Note that the list does not determine any absolute arrival or departure times, only the relative sequence.

Every vehicle in a TO, other than the first, follows a specific vehicle in the list. Tracking that vehicle guarantees the following vehicle will get to the CZ after that lead vehicle and all the other vehicles ahead of it have crossed the CZ. The TO may therefore be stored in a distributed fashion as a single piece of information each vehicle needs to carry, namely, the identity of the vehicle it immediately succeeds in the TO. This reduces the problem to a generalized tracking issue (see Section 2.2). Indeed, the vehicle being followed need not be the one physically ahead of the tracking vehicle; it could be approaching the CZ in a different lane or from a different direction altogether.

In reality, a vehicle may need to cross several CZs to reach its desired destination. In a four-way intersection, for example, a west-bound vehicle may be required to pass two lanes, one north-bound and another south-bound, before clearing the intersection. This amounts to two CZs, along with two TOs, and hence two vehicles to track, one from each TO. Since one of them would be nearer (in the generalized distance sense) than the other, keeping a safe distance from the nearest of the two will guarantee a proper relation with the other one as well. This argument extends to any number of CZs. Although the relation between vehicles in different TOs may change over time, requiring that the tracking vehicle keep monitoring all TOs it is in, such changes are relatively infrequent since they involve only the relative speeds of vehicles moving logically in the same direction.

### 3.3 Implementing Speed-up

As described above (Section 2.4), speed-up is done locally. Each vehicle is expected to accelerate until it reaches some maximum speed, cruise at that speed past the CZ and then decelerate back to its previous speed. The difference between the actual vehicle speed and the speed it would have had without speed-up is the *speed increment*. The speed increment follows the shape depicted in Figure 1. Note that the increment is dependent only on the position relative to the CZ.

A vehicle enters a speed-up condition when it is notified of the *speed-up scheme*, which is the relationship between speed increment and vehicle position relative to the CZ. Parameters of the speed-up scheme include the following:

1. The road position where acceleration begins
2. Dynamic shape of acceleration (may be variable)
3. The maximum speed to be attained
4. The road position where deceleration begins
5. Deceleration rate

Again, since speed-up is affected only in the vicinity of CZs, an area controller is available to provide this information. Under speed-up conditions, each vehicle calculates the speed increment continuously and applies it to its current speed; the increment is added to the speed the vehicle would have under the same conditions without the increment. Actual vehicle behavior is modulated by tracking, so its actual speed will never exceed that of the tracked vehicle. However, as we shall see, if all vehicles follow the same speed-up scheme, the tracked vehicle speeds up by the same amount and slightly before the tracking one, so tracking does not get in the way. Moreover, this delay is responsible for widening the gap between the two vehicles, which is the point of speed-up.

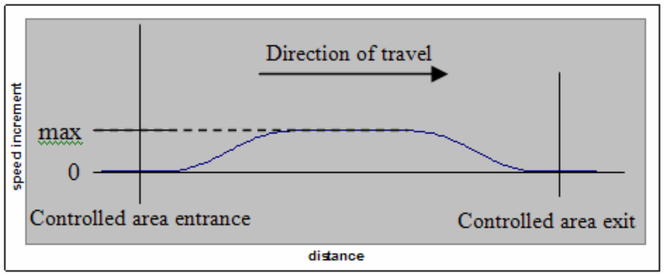

Figure 1: Speed-up scheme

### 3.4 The Collaborative Driving System

The system described pertains to an area containing a CZ, or possibly several related and closely-situated CZs. Such an area is equipped with a *control center* (CC), a server capable of coordinating traffic through the controlled area. Vehicles passing through the area need special capabilities to utilize the services of the CC and to streamline their passage. The properties of the vehicles and a CC are described below.

In this system, vehicles normally manage their advance based on the manipulation of their controls by their human operator, combined with information they glean via communications with other vehicles nearby. Vehicles inside the bounds of a controlled area of road also resort to the services of the relevant CC. Note that the CC does not violate vehicles' autonomy. It merely provides CZ-relevant information and coordination advice, yet vehicles communicate with one another in a peer-to-peer manner to gather the information needed and to collaborate accordingly.

**Vehicle.** A vehicle here refers to the common meaning of the term, augmented with hardware and software needed to perform the following tasks:

1. Communicate with other vehicles (V2V)
2. Communicate with CCs (V2I)
3. Generalized tracking
4. Speed-up

**Control Center.** The CC is a hardware device, stationary in most cases, which has jurisdiction over a specific stretch of road. It also has the hardware and software wherewithal required to perform its task, as described below. It is important to note that a controlled area may contain several CZs. The CC has the following responsibilities:

1. Set up the boundaries of the controlled area and notify vehicles that they have entered it.
2. Maintain a total order of all vehicles in the controlled area; maintain TOs for all CZs within the controlled area and notify vehicles of their placement in all TOs they participate in.
3. Determine the need for speed-up and dissemination of the speed-up scheme when it is in effect.

In the case of permanent CZs, such as those associated with an intersection, the CC could be installed where the street light is housed today, in addition or instead of the latter. A roadwork group would bring a CC along and install it in the vicinity of their work area for the duration of their work. An unexpected blockage, such as one due to a broken-down vehicle, could be controlled by equipment on an attendant emergency vehicle or even by the on-board intelligence of the broken-down vehicle itself.

### 3.5 The Mechanism

The mechanism presented below describes the behavior of a vehicle and a CC separately.

**Vehicle Behaviors.** Vehicles behave differently in non-controlled areas and controlled areas (i.e., areas under the jurisdiction of a CC). Furthermore, vehicle behavior in controlled areas requiring speed-up is different from that in controlled areas that do not.

*Non-controlled area*: In non-controlled areas, vehicles move in tracking mode (see Section 2.1).

*Controlled area with no speed-up*: Upon entering a controlled area, a vehicle switches to generalized tracking mode (see Section 2.2). The vehicle is informed that it has entered a controlled area, as well as of the identity of the vehicle it should track, through messages broadcast by the CC.

The new vehicle entering a controlled area need not be behind the vehicle it should be tracking now, but it may need to slow down to establish the proper tracking gap between itself and the tracked vehicle.

*Controlled area with speed-up*: In a control area employing speed-up, the CC conveys to an entering vehicle the parameters of the speed-up scheme (Section 3.3), in addition to the control information described in Section 3.2. The vehicle uses the scheme and its own position information to continuously calculate the momentary speed increment. It then does its best to implement the speed increment, subject to tracking constraints and the capabilities of its power train.

**Control Center.** The CC is responsible for maintaining the traversal order and to employ speed-up if the traffic rate warrants it.

*Controlled area boundaries*: The boundaries of the controlled area are set by the CC. The way they are determined could impact the efficiency of the system. The CC could fix the boundaries at a certain distance from the CZ. This distance should be big enough to accommodate the ease of generalized tracking and consequent speed-up, if employed. CCs could also set the boundaries dynamically, starting out with narrow bounds, and expanding and shrinking the space taken up by the controlled area as momentary traffic rates change.

If the boundaries are set too narrowly, there might not be enough time for tracking and speed-up to relax. Boundaries that are set too widely, on the other hand, increase the chance of a slow-moving vehicle entering the controlled area and holding up faster-moving vehicles that have entered later but are now constrained to track it. It is therefore desirable to set the boundaries at their lowest practical distance. In that regard, dynamic boundary placement has an advantage.

One way to realize a dynamic boundary is through the placement of a priority line (PL) across all lanes leading to a common CZ at the same distance from the CZ. (Such lanes may or may not be adjacent; see detailed scenarios.) Initially, the PL is set at some minimal distance, dmin, but as vehicles pass the line, it is dynamically relocated to expand the controlled area in increments of ℓ, the average space taken up by a vehicle (its physical length plus the inter-vehicle gap at the speed it is moving). As vehicles cross the CZ and leave the area, the PL is moved backwards in decrements of ℓ, shrinking the area size. The PL is never set further than dmax. Thus, if there are n vehicles between the PL and the CZ, the PL would be at min(dmin + n. ℓ, dmax). See Figure 2 for a schematic description.

*Traversal order maintenance*:

Vehicles crossing the PL are assigned monotonically decreasing priority values according to the order of their arrival. Note that a vehicle may cross the priority line due to its own motion, or due to the expansion relocation of the PL past the vehicle's current position. In either case, priorities are assigned according to the order in which vehicles and the PL meet. When two vehicles register within a short enough time period as to be indistinguishable, one of them is chosen to be first at random.

As an outcome, the CC can order all vehicles in the order of their arrival, which is assumed to reasonably correlate to the order at which they would have arrived at the CZ (to prevent the scheme from being unduly unfair). Emergency vehicles could enjoy preferential treatment.

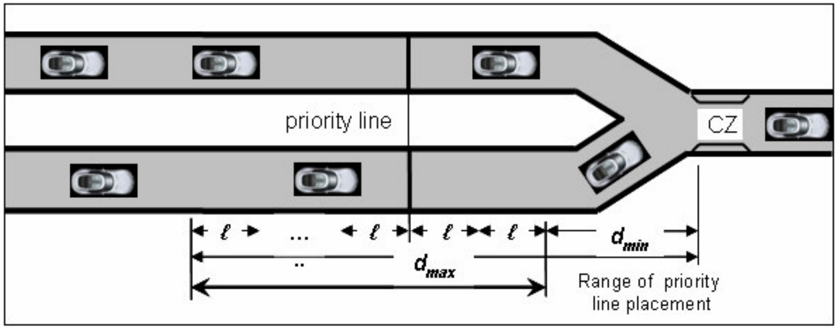

Figure 2: Placement of the PL where two lanes lead to a CZ. (PL is at dmin+ 2ℓ since two vehicles are between it and the CZ.)

This order is maintained within a TO. The CC notifies each vehicle, as it is placed in a TO, which TO it is in and which vehicle precedes it in the order; that is the vehicle it should track using the generalized tracking approach.

*Speed-up management*: As described earlier in this section, the position of the PL represents the number of vehicles contained between it and the CZ. The speed with which the PL itself moves represents the difference between the vehicles' arrival rate and the rate at which they pass the CZ. In the discrete case, the time between priority line repositioning is inversely proportional to the difference between the incoming and outgoing rates: the shorter the time, the faster vehicles ought to clear the controlled area.

Practically, speedup is triggered when the priority line reaches the distance of (dmin + dmax) / 2. The time at which that distance is reached is recorded as, say, t0. If the priority line is moved upstream at a later time, say t1, the speed of all vehicles that have crossed the priority line is increased by up to ℓ/(t1 – t0). Speed increment may be raised if PL movements accelerate, but is never reduced: it is reset back to 0 if and when the space between the PL and the CZ is cleared of vehicles.

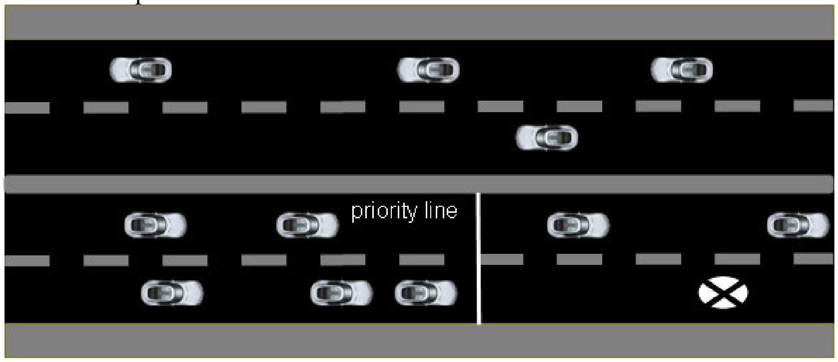

Figure 3: Vehicles approaching a blocked vehicle in a multi-lane highway.

# 4 Detailed Scenarios

## 4.1 A Blocked Lane in a Multi-Lane Highway

The system above may be applied to a multi-lane highway where a lane is blocked, due perhaps to a broken-down vehicle or road work. The associated CZ is the space in the unblocked lane next to the blockage, through which traffic arriving on the blocked lane will eventually have to travel. (If there are more than one unblocked lanes, each one may have a CZ.) The controlled area extends from the blockage to the PL, which may be dynamically set as described in Section 3.4 (see Figure 3).

Vehicles that have crossed the PL are within the controlled area and are assigned priorities. The metrics for the generalized tracking in this case is simply the distance to the blockage. Prioritized vehicles begin following all vehicles with a higher priority than their own, regardless of their respective lane position. (Recall that priorities are assigned in decreasing order.) This causes vehicles to arrange themselves in a single file even if they are different lanes (see Figure 4).

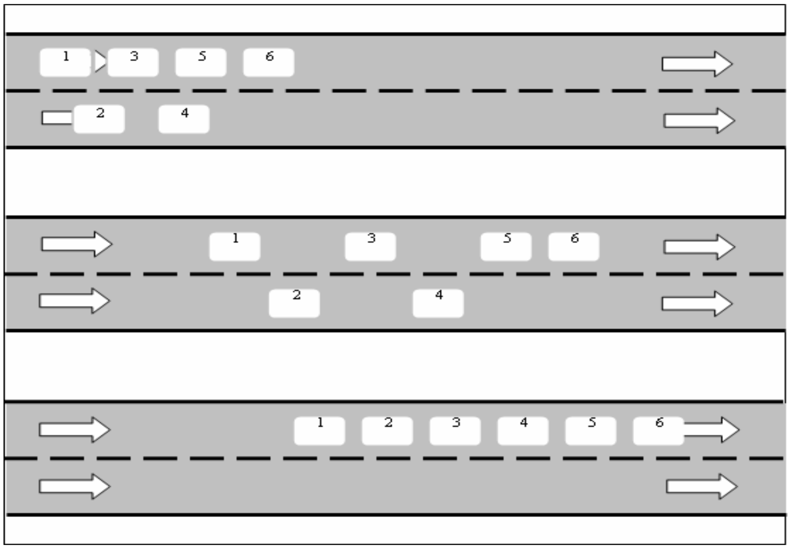

Figure 4: Tracking and merging by priority; the numbers indicate relative priority. Top - initial state; middle - after tracking was applied; bottom - traffic merged into one lane.

Once that state has been reached, any vehicle in the blocked lane (the bottom one in Figure 4) may move to the unblocked lane with confidence, knowing that all vehicles in the target lane with lower priority are already tracking its position, and as a result maintain a gap for it. (That is, the gap between vehicles 2 and 4 in the middle part of Figure 4 is maintained regardless of which lane vehicle 3 is in.) Lane switching is the responsibility of the driver – tracking modulates only vehicle speed, not steering – and must be completed after the vehicle crosses the line and before it reaches the obstacle. Vehicles which fail to move to the open lane in time get blocked, and have to wait, as all vehicles currently must, until traffic in the other lane has subsided.

Sufficiently sparse traffic could merge into the single open lane without any slowdown. At higher rates, however, a slowdown would result: while the first vehicle in a platoon maintains its previous speed, vehicles behind it must slow in order to generate the required inter-vehicle gaps. (E.g., in Figure 4, in the period between the situation depicted at the top and that in the middle, vehicle 6 has moved a longer distance than vehicle 1. Consequently, vehicle 1 must have moved more slowly than vehicle 6, which would not have necessarily been the case without tracking.)

To avoid such slowdowns, vehicles that have passed the priority line are made to accelerate, increasing lane capacity to the level necessary to accommodate the traffic at hand. The speed-up is controlled by the system, and the driver has only limited control over it. Although not necessarily implemented that way, one may think of the added speed as imparted by movement of the pavement itself, as in a conveyor belt, adding a constant speed to all vehicles that are (virtually) on it. The speed of the conveyor belt is determined by the amount of traffic that it needs to accommodate. At most, its speed could reach the average traffic speed, effectively doubling the capacity of the single lane, and thus matching the throughput of both lanes together (at the given average traffic speed).

Once a vehicle has passed the obstacle, it decelerates back to its original speed, since both lanes are now available to carry the traffic and no speedup is needed any more. The effect of the speedup is therefore limited to a small area in the vicinity of the blockage and has no systemic effect on traffic as a whole.

The performance of this mechanism in this scenario was tested using a simulator built using the AnyLogic [3] agent-based simulation platform. The program is a straight-forward implementation of the mechanism described above over a stretch of a two-lane highway. Each vehicle is represented by an agent having one of three states:

1. Free-flow, in which the vehicle travels without any regard to vehicles around it.
2. Tracking, in which the vehicle follows behind another vehicle at that vehicle's speed
3. Obstacle, in which the vehicle is stationary

Each vehicle has properties of position and speed, as well as the identity of the vehicle it follows when tracking. Property values are chosen at random from a realistic range for highway traffic. Vehicles in tracking mode can switch lanes, if there is room in the adjacent lane. On-screen controls allow switching the obstacle on or off, the setting of the arrival rate, and selection of the algorithm to use: none, generalized tracking, and generalized tracking with speed up (see Figure 5).

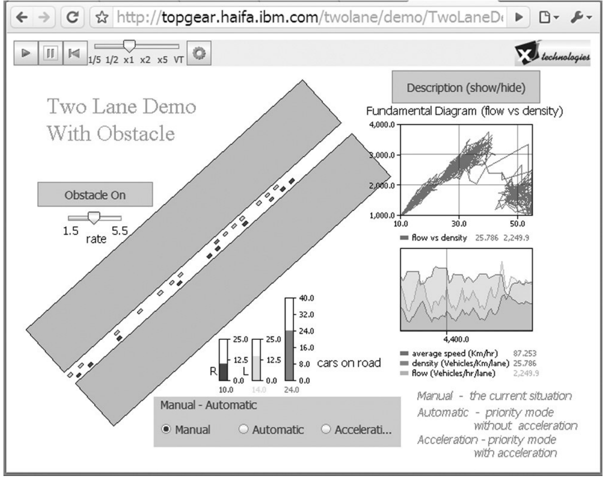

Manual - Automatic

● Manual ○ Automatic ○ Accelerati...

Manual - the current situation
Automatic - priority mode without acceleration
Acceleration - priority mode with acceleration

Figure 5: Snapshot from the collaborative driving simulator

At traffic rates of over 1500 vehicles/hour/lane, a traffic jam quickly forms when an obstacle blocks one lane, but it dissipates if generalized tracking is turned on. Tracking can handle traffic rates up to about 2200 vehicles/hour/lane. If speed-up is also employed, traffic rates of even up to 3000 vehicles/hour/lane do not form a jam.

A flow vs. density diagram for the simulated traffic, which the simulator maintains, exhibits the characteristic fundamental diagram behavior. An applet simulating this scheme under various traffic conditions is available at [4] .

## 4.2 Intersection Where No Turns Are Allowed

Another example is that of an intersection where no turns are allowed; all vehicles are constrained to cross the intersection going straight only. There could be any number of lanes in each direction, however.

To simplify the exposition, an intersection consisting of two one-way intersecting lanes, as in Figure 6, is described first; the solution is then expanded to any number of lanes. The two lanes form exactly one CZ: the area common to the two lanes. The bounds of the controlled area are set by a PL, which is placed, possibly dynamically, on both incoming lanes at equal distance from the CZ.

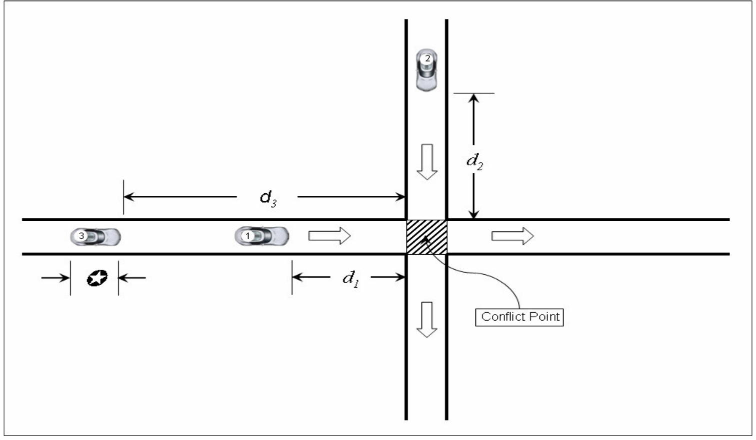

Figure 6: An intersection with a single CZ, defining the metrics for generalized tracking.

For the purpose of generalized tracking, the metrics are defined as follows:

1. The distance between two vehicles traveling in the same lane is measured in the normal way—the distance one vehicle must travel to touch the end of the other.

2. The distance between two vehicles approaching the intersection from different roads is defined as the difference between their respective distances to the part of the CZ that is closest to each.

For example, in Figure 6, the distance between vehicles 3 and vehicle 1, which is physically ahead of vehicle 3, is given by

$$d_{31} = d_3 - d_1 - \ell$$

where $\ell$ is the average length of a vehicle. The distance between vehicle 2 and vehicle 1, which are approaching the CZ from different directions, is given by

$$d_{21} = d_2 - d_1 - \ell - w$$

where $w$ is the average width of a vehicle. The added padding is needed because the distance a vehicle travels from the point it enters the CZ to the point it clears is the sum of the vehicle's length and the width of the CZ, which is about the width of a vehicle. If all vehicles are moving at the same speed, that should also be the gap between vehicles moving perpendicular to that vehicle (see Figure 7).

A likely TO for the situation depicted in Figure 6 is (1, 2, 3).

The same idea applies to intersections in general. The area common to any intersecting lanes is considered a CZ, and a TO is maintained for each (see Figure 8). Each vehicle computes which CZs its path must cross, and which vehicle is just ahead of itself in each of the associated TOs. Every such TO contributes a vehicle it must follow, and each of these vehicles imposes some distance from the intersection it must keep at any moment. The maximum of these distances satisfies the tracking requirements of all of them, and that is the distance the vehicle should maintain.

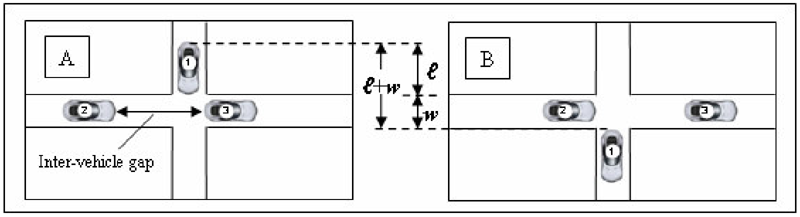

Figure 7: A vehicle crossing an intersecting traffic stream needs a gap as wide as its length plus its width. The distance vehicle 1 travels since entering the intersection (A) until it clears it (B) is ℓ + w; assuming 2 and 3 move at the same speed, their inter-vehicle gap must be just as large.

For example, vehicle 2 in Figure 8 must cross CZs B1 and B2 to get to the other side of the intersection. The TO for CZ B1 contains vehicle 4, and that of CZ B2 contains vehicles 1 and 3. Vehicles 1 and 4 should pass the intersection before 2, so 2 must follow them both. Although the speeds in Figure 8 are not known, it stands to reason that vehicle 2 follows vehicle 4, and that will take care of vehicle 1 as well, since vehicle 1 is so much closer to its CZ than vehicle 4 is.

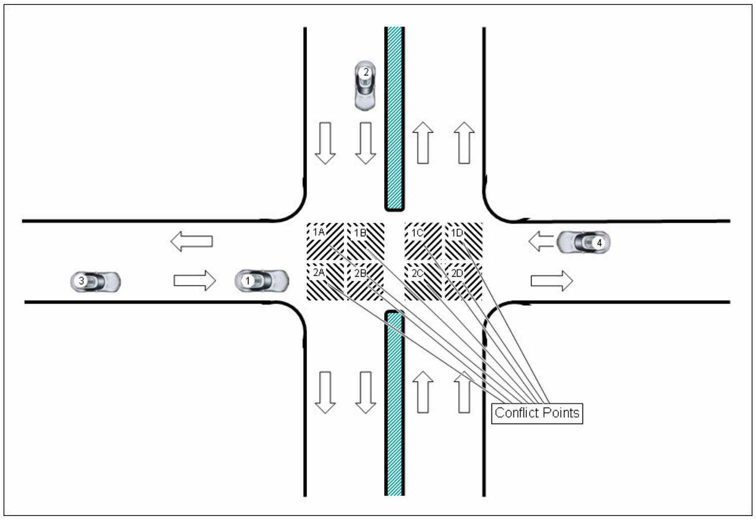

Figure 8: An intersection with several CZs.

## 5 Related Work

Collaborative driving is seen as a promising solution to mounting traffic problems. A collaborative scheme for traversing intersections with a setup similar to the one used in this paper is described in [8] [9] ,[10] [12] Dresner and Stone, as well as Kolodko and Vlacic, describe the intersection as being divided into vehicle-size slots in time and space, and each approaching vehicle can request and obtain a rendezvous with such a slot. If all such rendezvous are kept, safe passage is guaranteed without the need to slow down, but a vehicle that misses its turn must stop, delaying itself and the vehicles behind it. Our solution avoids such delays. On the other hand, the other scheme supports turns, whereas ours, at this stage, does not. The other scheme can be integrated with a standard street signal to accommodate some non-communicating vehicles, although performance suffers if there are too many of such vehicles. That scheme further requires that vehicles' behavior in the vicinity of the intersection be governed completely by the CC. Our scheme does not impose such a restriction, and therefore can be more gracefully adopted.

A scheme that requires no control center is described in [13] There, a scheme similar to the generalized tracking is used to serialize *vehicle pairs* that can occupy the intersection simultaneously without interference (for instance, if they are moving in opposite directions and both intend to turn right). Their trajectory is then automatically controlled to realize the best traversal plan. In that scheme, too, the driver has no control of the vehicle while it crosses the intersection, while in our solution the driver retains control.

We are not aware of other collaborative driving schemes that specifically address intersection and general CZ crossing. However, several multi-agent and distributed AI theories could help solve this problem. For instance, the generalized partial global planning framework (GPGP) [6] , a general scheme for distributed coordinated planning, can be used. Indeed, GPGP has been applied to the distributed vehicle monitoring testbed (DVMT) [7] That application exhibits similarities to the generalized tracking presented here; however, it does not refer to the specific problem of CZ crossing. It further assumes full autonomy of vehicles, whereas we allow driver intervention.

## 6 Conclusion

In this study, we have introduced an innovative inter-vehicle collaboration mechanism. When implemented, it allows (semi-) autonomous vehicles to increase road infrastructure utilization with virtually no effect on safety. As we show, our mechanism lets vehicles traverse road conflict zones such as blocked lanes and intersections with significant time and energy savings, avoiding decelerating and accelerating which would otherwise be required. In a counter-intuitive manner, we advocate that vehicles should accelerate into a busy road intersection. This is shown to be advantageous and introduces no additional risk when complementary techniques are implemented.

As shown in simulations, the suggested techniques are promising. Future work calls for extensions to support turns, simulations of additional scenarios, and eventually field studies. With such studies, we are confident that techniques of the sort presented here should become part of future driving reality.

## 7 Acknowledgements

The authors wish to thank Stuart Feldman, Larry Lieberman, Peter Robisson, and Dov Ramm for their invaluable support and advice.

## 8 References

[1] A Quick Look at Intersection Crashes in Canada. Road Safety and Motor Vehicle Regulation Directorate, Transports Canada, http://www.tc.gc.ca/roadsafety/tp2436/rs200806/menu.htm, (2008)

[2] Adaptive Cruise Control, http://www.mobileye.com/default.asp?PageID=328.

[3] AnyLogic 5. XJ Technologies, http://www.xjtek.com.

[4] Blocked Lane Simulation, http://topgear.haifa.ibm.com/twolane/demo/TwoLaneDemoApplet.html. A narrated introduction, http://topgear.haifa.ibm.com/ twolaneMovieEng/take6Web.swf.

[5] Crash Rates at Intersections. Kentucky Transportation Center Research Report KTC-03-21/SPR258-03-2I, http://www.ktc.uky.edu/Reports/KTC_03_21_SPR258_03_2I.pdf, (2003)

[6] Decker, K., Lesser, V.: Generalizing the Partial Global Planning Algorithm. Intelligent and Cooperative Information Systems, vol. 1, pp. 319-346 (1992)

[7] Decker, K., Humphrey, M., Lesser, V.: Experimenting with Control in the DVMT. In Proc. *Third Annual AAAI Workshop on Blackboard Systems*, Detroit (1989)

[8] Dresner, K., Stone, P.: Multiagent Traffic Management: A Reservation-Based Intersection Control Mechanism. In Proc. AAMAS'04, pp. 530-537, New York, USA (2004)

[9] Dresner, K., Stone, P.: Multiagent Traffic Management: An Improved Intersection Control Mechanism. In Proc. AAMAS'05, pp. 471-477, Utrecht, The Netherlands (2005)

[10] Dresner, K., Stone, P.: Sharing the Road: Autonomous Vehicles Meet Human Drivers. In Proc. IJCAI 07, Hyderabad, India (2007)

[11] Kerner, B. S.: The Physics of Traffic. Springer-Verlag (2004)

[12] Kolodko, J., Vlacic, L.: Cooperative Autonomous Driving at the Intelligent Control Systems Laboratory. *IEEE Intelligent Systems*, 18(4):8–11 (2003)

[13] Li, L., Wang, F.-Y.: Cooperative Driving at Blind Crossing Using Intervehicle Communication. IEEE Transactions on Vehicular Technology, Vol. 55, No. 6 (2006)

[14] Uno, A., Sakaguchi, T., Tsugawa, S.: A Merging Control Algorithm Based on Inter-Vehicle Communication. IEEE/IEEJ/JSAI International Conference on Intelligent Transportation Systems (1999)