Emergent Behavior through Local Decisions in a Multi-agent Traffic Simulation

One MS student, two undergraduate students

City traffic is an obstacle that many people must deal with on a daily basis. It stands to benefit greatly from added safety and efficiency. By developing a multi-agent system, we simulate a potential avenue for achieving this. In our simulation, we model a simple grid-based city and all the traffic within it. We seek to achieve faster and safer traffic flow by utilizing "smart" traffic signals and cars which observe and utilize their local surroundings to make better decisions than their "dumb" counterparts. We conduct a series of experiments to test the effectiveness of our design and analyze the results.

I. Introduction

To investigate the process of local decision making leading to emergent coherent global behaviors, we designed a traffic simulation system using the Repast Simphony simulation toolkit in the Java programming language. The traffic simulation system comprises of four types of agents: dumb vehicles, intelligent vehicles, dumb traffic signals, and intelligent traffic signals. Agents were designed to have information about only their local environment, and make decisions based entirely on what they see nearby. The vehicle agents make decisions about how to travel and the traffic signal agents makes decisions about what direction of traffic to allow. The desired emergent behaviors in the system include: 1) efficient traffic flow and 2) safe traffic flow. The purpose of the dumb vehicle agents and the dumb traffic signal agents was to provide a baseline against which to measure the performance of the intelligent versions of the vehicle and the traffic signal agents. We designed experiments to test several hypotheses on the extent to which the traffic simulation system exhibits these desired emergent behaviors.

II. Simulation Design

In this section, we describe the design for the traffic simulation system. First we discuss the design for the agent environment; then we go on to describe the designs of the dumb vehicles, intelligent vehicles, dumb traffic signals, and intelligent traffic signals respectively. Finally, we discuss how we measure the system's performance.

A. Environment Design

The traffic simulation environment consists of a variable $n \times n$ grid of city streets with city blocks that are a variable $m \times m$ in size. Each street is a two-way street and consists of two lanes in both directions. Near the intersections are left turn lanes, and away from the intersections a median separates the two directions of traffic. Every intersection has its own traffic signal controlling the flow of traffic through the intersection. Scattered throughout the city are $4 \cdot (n-1)^2$ sources/sinks (four per city block) representing various destinations within the city. Each source has a probability p_i of generating a vehicle on each iteration, and one sink is selected as a destination, each having a probability of $q_i = \frac{p_i}{\Sigma p_j}$, every time a new vehicle agent is instantiated. This simulates certain locations in a city being busier than others. For simplicity, all streets were made to have a uniform speed limit which is a reasonable assumption in most downtown city areas. The streets are made up of individual street cells; each street cell has a direction, and indicators that indicate which turns are allowed. This design mimics the way real-world drivers use street markings and signs to follow traffic rules. The environment imposes a global limit on the number of cars in the system at any one time to simulate how traffic conditions vary at different times of the day—rush hour vs. three in the morning. However, this limit can be essentially removed by setting it at an arbitrarily high amount. The environment also allows for different proportions of dumb and intelligent agents in the system at once—from no intelligent agents in the system at all to every agent in the system being intelligent.

B. Dumb Vehicle Agent Design

The dumb vehicle agents were designed to mimic the way real-world motorists drive, react to events on the road, and make decisions. They are dumb in that they do not communicate with the other vehicle agents around them; they only react to what they see in front of them—much like real-world human drivers do.

The vehicle agents navigate and plan their routes in a way that gives similar results to how human motorists plan their routes. When each vehicle agent is instantiated, it is given a destination sink. As it travels, the vehicle agent has knowledge of the location of the next intersection it is approaching which is a reasonable assumption for human motorists to know because many city streets have names that are numbered or lettered in a sequential manner. For example, if a motorist knows his current location to be near, say, 12th and G street heading west in some city where north-south streets are numbered in ascending order from west to east and east-west street are lettered in ascending order from south to north, then he can reason that the next intersection he will approach will be 11th and G street. The vehicle agents plan their route by determining which direction they will travel in at the next light each time they pass through an intersection. To select which direction to travel at the next light, the vehicles first determine which of the possible directions will allow them to move closer to their destination. In a city where streets are numbered and/or lettered sequentially, it is reasonable to assume that human motorists could determine which streets will take them closer to their destination based on their knowledge of their destination's address. If the vehicle agent determines that more than one direction may lead it closer to its destination, then it will randomly choose one of those directions as its next travel direction at the next intersection. This essentially captures the way most human motorists plan their route. Although instead of planning their entire route at once at their source like most human motorists do, the vehicle agents plan it incrementally as they travel, but the results are essentially similar.

The way the dumb vehicle agents perceive and react to events on the road also imitates the way human motorists do. On each iteration, the vehicle agents' first action is to check directly in front of them for any other vehicles or accidents in their path. If an impediment to their path is found, then the vehicle may attempt to merge into an open lane if he does not need to be in a turn lane to make a turn at the next intersection. If the vehicle agent is unable to merge into another lane to go around the impediment, then it will fully apply its brakes and begin to decelerate. Once the vehicle agent has had a chance to react to any events in front of him, it will then check to see whether the next traffic signal he is approaching is open if he is nearby that traffic signal's intersection. Before a vehicle agent enters an intersection, it will check directly ahead of it to ensure that it will not be forced to stop in the middle of an intersection (i.e. we are assuming the drivers in our system are Nebraska drivers rather than California drivers). For simplicity the vehicle agents all have a uniform maximum speed which is set to be the speed limit of the city (again, we are assuming the drivers are from Nebraska, not California).

C. *Intelligent Vehicle Agent Design*

For many of the basic operations, intelligent vehicle agents behave in the same way as the dumb vehicle agents. They differ from the dumb vehicle agents and are intelligent because they 1) interleave planning and execution, and 2) communicate with other vehicles nearby.

A total priority ordering is imposed on the intelligent vehicle agents in the system. Upon instantiation, each intelligent vehicle agent is given a unique identification number that represents its priority ranking. The relative ordering of the intelligent vehicle agents determines which agent is responsible for handling a conflict with another agent when one is identified through communication. The intelligent vehicle agents interleave planning and execution by dividing their movement procedure into two parts. The agents first determine the coordinates of their next planned location. The intelligent vehicle agents then send a message to the other vehicle agents nearby in its vicinity. The message contains the coordinates of what it plans for its next location, and a reference to the sender agent. When an intelligent vehicle agent receives a message from another agent, it first determines whether the coordinates of its next planned location conflicts with the coordinates of the other vehicle agent's next planned location. If two intelligent vehicle agents have conflicting planned moves, then it is the responsibility of the lower priority agent to manage the conflict. The lower priority intelligent vehicle agent will handle the conflict by either ordering the other agent to slow down slightly while it accelerates slightly or by ordering the other agent to speed up slightly while it slows down slightly depending on the relative locations of the two vehicle agents and the actions that they plan on performing. The distance which the intelligent vehicle agents may communicate with other intelligent vehicle agent is a system parameter so that the effects of communication distance may be studied.

D. *Accident Design*

On each iteration of the simulation, motor vehicle accidents are determined by each vehicle agent in the system. A vehicle agent determines that it is in a collision if it finds that another vehicle agent or an accident object shares its current location. If a vehicle agent determines that it is in an accident, then it is removed from the context of the simulation, and an accident object is instantiated in its place. Upon instantiation, each accident object is given a clean-up time that represents the amount of time for the collision to be cleared from the road. The clean-up time for each accident object is randomly selected from a range whose maximum and minimum values may be set as parameters for the simulation. A low clean-up time may represent a minor "fender-bender" such as a side-swipe or a rear-end collision in which all that is needed to clear the accident from the road is for the motorists involved in the collision to exchange insurance information. A longer clean-up time may represent a more serious motor vehicle accident for which emergency responders must be called to clear the accident from the road.

E. *Dumb Traffic Signal Agent Design*

The dumb traffic signal agents are dumb in that they have no view of the conditions of their local environment. They are completely oblivious and unreactive to what is going on around them, and they all simply cycle through states in a regular manner on a cycle of fixed length. The dumb traffic signal agents allow for left turns and straight/right turns in each of the four directions giving them eight lights, each of which can be red or green. This results in $2^8 = 256$ possible states for a signal. A *state* for a traffic signal agent is what directions it is allowing traffic to pass through its intersection. For example, the current state of a traffic signal agent may allow vehicle agents to travel north straight and south through the intersection (and also make right turns to head east and west) or it could allow vehicles to make left turns to head south and make left turns to head north. A *legal state* for a traffic signal agent is a state that does not allow conflicting directions of traffic flow through the traffic signal's agent. For example a state that would not be legal for a traffic signal agent would be one that allows both traffic travelling straight through the intersection heading north and traffic travelling straight through the intersection heading west. The dumb traffic signal agents have four legal states corresponding to states that allow vehicles to travel straight through the intersection heading north and south as well as straight through the intersection heading east and west, and allowing left turns to travel north and south as well as allowing left turns to travel east and west.

F. Intelligent Traffic Signal Agent Design

Intelligent traffic signal agents are self-interested agents that seek to minimize the wait times of the cars stopped at the traffic signal agent's intersection. The intelligent traffic signal agents are aware only of the local conditions of their nearby environment and have no global information at all. They have complete autonomy for the selection of their current legal state. Since the intelligent traffic signal agents have more flexibility to respond to local conditions near their intersection, they are allowed four more possible legal states in addition to the possible legal states also allowed for the dumb traffic signal agents. These additional legal states include: travelling straight through the intersection heading north and turning left to head west, travelling straight through the intersection heading south and turning left to head east, travelling straight through the intersection heading east and turning left to head north, as well as travelling straight through the intersection heading west and turning left to head south.

Intelligent traffic signal agents are charged for the amount of time they keep cars waiting at their intersection. The amount an intelligent traffic signal agent, i, is charged, p_i, is given by:

$$p_i = \sum_{j=1}^{n} t_j^{1.5}$$

Where n is the number of vehicle agents waiting at traffic signal agent i's intersection, and t_j is the total amount of time that a vehicle agent j has been stopped at i's intersection. The exponent of 1.5 makes the equation grow super-linearly which ensures that cars which have been waiting for a long time are given priority to be allowed through the signal agent's intersection. Each vehicle agent communicates both its total wait time to the traffic signal agent that it is waiting for as well as which direction it would like to travel through the intersection. The intelligent traffic signal agents have a cycle length that is set as a simulation parameter. At the end of the traffic signal agent's cycle it will select its current legal state as the one which will minimize the total cost to the traffic signal agent.

G. Parameters

The parameters we included in our simulation allow us to easily vary many aspects including the intelligence of the agents, the size of the simulation, and the timings of several entities:

1. *City Size*
   This value specifies the number of streets going in both the horizontal and vertical directions. Thus the number of intersections is this value squared.
2. *Block Size*
   This value specifies the length of a street from intersection to intersection, i.e., the length of a city block. The units of block size are in Repast grid units
3. *Smart Light %*
   This sets the percentage of the traffic signals that will behave as "smart" signals. The remaining traffic signals are "dumb".
4. *Smart Car %*
   This sets the percentage of the cars that will spawn as "smart" cars. The remaining cars will spawn as "dumb" cars.
5. *Communication Distance*
   This parameter sets the distance that smart car messages will travel to other cars.
6. *Termination Tick*
   This sets how long the simulation will run for. When the termination tick number is hit, the simulation will halt.
7. *Max Cars in World*
   A setting used to cap the number of cars that are allowed to exist in the world. If the limit is reached, no more cars may spawn unless cars first leave the world. Primarily used for debugging purposes, but can also be used to simulate the difference between traffic conditions in rush hour at two in the morning.
8. *Min Clean Up Time*
   The minimum number of ticks required before an accident will "clean up" and remove itself from the simulation.
9. *Max Clean Up Time*
   The maximum number of ticks required before an accident will "clean up" and remove itself from the simulation.
10. *Dumb Cycle Length*
    Each individual signal has a set length of time in its cycle of transitioning its lights. This value sets the mean value for all dumb signals.
11. *Dumb Cycle Spread*
    Each individual signal has a set length of time in its cycle of transitioning its lights. This value sets the amount above and below the mean cycle length that each individual signal may vary by. This ensures that all dumb signals aren't synced up in their transitions.
12. *Smart Cycle Length*
    This parameter sets the length of time between recalculating and changing to the appropriate lights for smart signals.

H. Measurements

The measurements chosen for observation are the ones that best describe the state of a particular character at a point of time in the system. The following are the observed values with a brief description for each:

1. *Car Count*
   This value describes the total number of cars in the system at the given tick. This is measured by Repast Simphony by counting the number of instances of the Car Class; the base class that both Smart and Dumb Cars inherit from.
2. *Accident Count*
   This value describes the total number of accidents in the system at the given tick. It is also measured by Repast Simphony by counting the number of instances of the Accident Class.
3. *Average Ticks of a Car*
   This value describes the average amount of ticks that the cars spend inside the system-from spawn to death. Each instance of the Car Class stores a class variable with the total amount of ticks it has spent inside the system since spawn. Repast Simphony takes this value at each tick and applies the mean function on it.
4. *Average Ticks at Light*
   This value describes the average amount of ticks that the cars spend at a signal waiting on a red light. Each instance of the Car Class stores a class variable with the amount of ticks it has spent waiting at a signal on a red light. Repast Simphony takes this value at each tick and applies the mean function on it.
5. *Average Speed of Cars*
   This value describes the average speed of the cars in the system. Each instance of the Car Class stores a class variable with its current speed in it. Repast Simphony takes this value at each tick and applies the mean function on it.
6. *Percent Accidents*
   This value describes the percent of total cars that have ever been in the system or are currently in the system that have been involved in a crash. This allows us to get a global picture of how many cars end up in an accident while passing through the system. This is calculated in the Accident Class by dividing the total number of accidents that have occurred with the total number of cars in the system and multiplying by 100. This value is used in our experiments for measuring the safety of the system
7. *Percent Efficiency*
   This value describes how effectively the Car Agents are able to reach their intended destinations. Each instance of the SourceSink Class calculates this for every instance of the Car Class that it destroys (which happens when a Car makes it to the Sink Agent), by taking the Manhattan distance and dividing it with the maximum speed that the Car could have travelled at and dividing that value by the actual number of ticks the Car took to reach the respective SourceSink and then multiplying by 100. Then this value is aggregated across all arrivals that have occurred in the course of the simulation.
8. *Max Ticks At Light*
   This is a debug based value that denotes the current most time spent at a light in the system. This is used to check if any single car spends too long inside the system waiting at a given light. This is calculated by Repast Simphony by taking the maximum value of all of the each Car's time spent at a signal. Frequent changes in this graph should signify that no single car is the cause of failure within the system.
9. *Max Car Ticks*
   This is another debug based value. It denotes the current most time spent by any Car inside the system. This is also used to check if any one car is causing the entire system's observed parameters to change. Frequent changes in this graph should signify that no single car is spending too long a time running in the system.

G. Emergent Behavior

From the design of the traffic system's environment agents, we hope to see two emergent behaviors from the local decisions of the agents. First, we hope to see that the intelligent signal agents making decisions based only on what they see in their local environment will lead to more efficient traffic flow. In other words, we expect that the local decisions made by the intelligent signal agents will reduce the total amount of time cars spend stopped at intersections. The second emergent behavior we hope to see from the system is that the local communication of the intelligent vehicle agents will reduce the number of accidents that occur in the system compared to an environment that only contains dumb vehicle agents.

III. Hypotheses

In this discussion we discuss a set of questions we wished to explore with our simulation to evaluate the effects of environmental settings on the performance of the system and what coherent behaviors emerge as result of the agents in the system making local decisions. We provide hypotheses for each question and an explanation for each hypothesis.

Question 1: How does the number of streets in the city affect the efficiency and safety of traffic flow in the system?

Hypothesis 1: We hypothesize that a larger number of streets in the city will have a small negative effect on the efficiency of the traffic flow in the system and will have no effect on the traffic safety of the system.

We expect that having more streets in the city grid will reduce the efficiency of the traffic flow because some vehicles will have to travel through more intersections on the way from their sources to their destinations, and thus they have a greater chance of being stopped at a traffic signal and spending more time not moving. We do not expect more streets to affect the safety of traffic in the system because we do not expect that more streets in the system will lead to any additional interactions between the vehicles that may lead to more collisions.

Question 2: How does the size of the city blocks affect the efficiency and safety of traffic flow in the system?

Hypothesis 2: We hypothesize that larger block sizes will lead to higher traffic flow efficiency and safety.

We expect that having larger block sizes in the system will lead to more efficient traffic flow because vehicle agents will spend more of their travel time driving between intersections and away from traffic signals than they would be in settings with smaller block sizes. Thus in settings with larger block sizes, vehicle agents will spend more of their travel time driving where they are less likely to be stopped, since we expect most vehicle stops to occur at intersections. We also expect that larger block sizes will increase the safety of the traffic flow in the system because vehicles will spend more driving time away from the intersections where collisions are more likely to occur.

Question 3: How does the proportion of intelligent vehicle agents in the system affect the safety and efficiency of traffic flow in the system?

Hypothesis 3: We hypothesize that a greater proportion of intelligent vehicle agents in the system will increase both the safety and efficiency of the traffic flow in the system. We hypothesize that the local decision making on the part of the intelligent vehicle agents will lead to the emergent behavior of increased system traffic flow efficiency over what can be achieved in a system with fewer intelligent vehicle agents

We expect that with a greater proportion of intelligent vehicle agents in the system, traffic safety will be greater because the intelligent vehicle agents communicate with other vehicle agents nearby and make local decisions about how to best avoid collisions with other nearby vehicles. We also expect that having a greater proportion of intelligent vehicle agents in the system will increase the efficiency of the traffic flow in the system because since we expect there to be fewer collisions in the system, we also expect that there will be fewer disruptions to traffic flow due to collisions blocking the road thus the vehicles will spend less time stopped.

Question 4: How does the proportion of intelligent traffic signal agents in the system affect the efficiency and safety of traffic flow in the system?

Hypothesis 4: We hypothesize that more intelligent traffic signal agents in the system will increase both the safety and efficiency of traffic flow in the system. We hypothesize that the local decisions made by the intelligent traffic signal agents will lead to the emergent behavior of increased traffic safety and efficiency in the system.

We expect that with a higher proportion of intelligent traffic signal agents in the system, the efficiency of traffic flow will increase because the intelligent traffic signal agents are able to make decisions about what state to be in based on what they perceive in their local environment. We expect this ability will allow the intelligent traffic signal agents to be responsive to the local traffic conditions near their intersection. Additionally, since we expect traffic to flow much more smoothly with more intelligent signal agents, we also expect there to be fewer collisions in the system since vehicles will have to stop less which is when collisions are more likely to occur.

Question 5: How does the communication distance for the intelligent vehicle agents affect the safety and efficiency of traffic flow in the system?

Hypothesis 5: We hypothesize that as the distance intelligent vehicle agents are able to communicate with each other increases, both the safety and efficiency of traffic flow will increase up to a point. After that point, we hypothesize that no further gains in traffic safety or efficiency can be had by increasing the intelligent vehicle agents' communication distance.

We expect that allowing greater communication distances for the intelligent vehicle agents will increase safety because it will allow the intelligent vehicle agents more information about their environment and the intentions of other vehicle agents around them. After a certain point, we expect allowing intelligent vehicle agents to communicate greater distances will not increase the traffic safety because the vehicle agents will begin to communicate with other vehicle agents that they have no chance of interacting with. Thus, the intelligent vehicle agents will derive no additional benefit from being allowed to communicate at further distances. We expect the efficiency of traffic flow in the system to increase as the intelligent vehicle agents' communication distance increases because we expect there to be fewer disruptions to traffic flow caused by collisions blocking the road.

IV. Experimental Design

In order to test our hypotheses and study how agents making local decisions can lead to emergent global coherent behaviors, we have designed several experiments in which several environmental variables of interest were varied (shown in table 1). We tested one variable at a time, holding all other variables constant. The bolded values were used as default values for when each variable was held constant. Additionally, several other simulation parameters were kept constant throughout all experiments (shown in table 2). Note that the proportion of intelligent vehicle agents and intelligent traffic signal agents had two default values: one for one-hundred percent intelligent agents and one for no intelligent agents in the system.

Table 1: Experimental Parameter Settings

Parameter	Value 1	Value 2	Value 3	Value 4	Value 5
City Size	4	5	**6**	7	8
Block Size	15	**20**	25	30	50
Smart Light %	**0**	25	50	75	**100**
Smart Car %	**0**	25	50	75	**100**
Comm. Dist.	1	2	**3**	4	5

Table 2: Fixed Parameter Settings

Parameter	Fixed Value
Termination Tick	50000
Max Cars in World	10000
Min Clean Up Time	70
Max Clean Up Time	100
Dumb Cycle Length	40
Dumb Cycle Spread	10
Smart Cycle Length	10

For each of our experiments, at each variable setting, we ran 5 separate tests run under 5 different seeds: {1,2,3,4,5}. This provided us with more data upon which to base our analysis. Thus we ran 100 runs for the City Size experiment, 100 runs for Block Size, 50 runs for Smart Light %, 50 runs for Smart Car %, and 50 runs for Comm. Dist. for a total of 350 runs.

We conducted our test runs on Holland Computing Center's Tusker cluster. Each test was allocated a single CPU core and 2 GB of memory. The runs were allowed a run-time limit of 4 hours but most finished in 5 to 15 minutes.

V. Results

In this section, we present the results of the experiments to test each of our five hypotheses, discuss whether we found support for them, and discuss the implications of our observations. Note that in the following graphs, "All Smart" or "All Dumb" are used to indicate that the proportion of both intelligent vehicle agents and intelligent traffic signal agents in those experiments were both set at 100% and 0% respectively. Where it says "Smart Cars, Dumb Signals" that means that the proportion of intelligent vehicle agents in the system was set to 100%, and the proportion of intelligent signal agents in the system was set to 0% and vice versa for where it indicates "Dumb Cars, Smart Signals".

A. City Size Experiment

In this experiment we varied the number of streets in the city from four to eight in order to determine what effects, if any, the size of the city had on the safety and efficiency of traffic flow in the system. We used all combinations of 100% dumb and 100% intelligent agents in order to test whether the effect of the city size was consistent for all system settings. First we looked at the effects of the city size on the safety of the system as measured by accident rate of the system. The results are presented in figure 1.

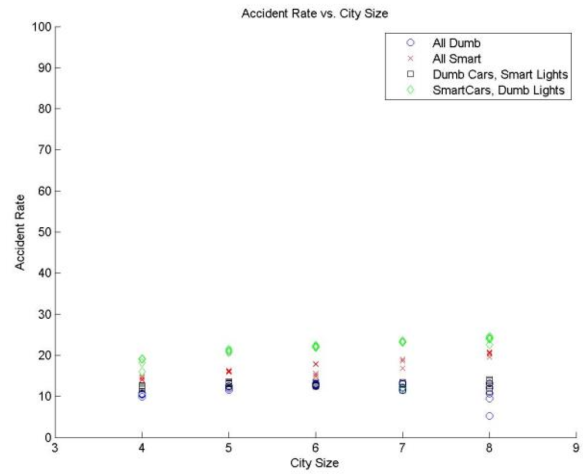

Figure 1: The Effect of City Size on Traffic Safety

Several observations stand out from the results of this experiment. First we can see that the accident rate in the system did not change much as the size of the city was increased. Second varying the size of the city had slightly different effects on the accident rate for different settings of the system. In the setting where all agents were dumb, the size of the city appears to have no effect on the accident rate with a larger range of results for larger sizes. For the settings with entirely intelligent vehicle agents, increase the size of the city looks to have had no effect to slightly increasing the accident rate. These results do not entirely support nor refute our hypothesis that increasing the size of the city would have no effect on the safety of traffic flow in the system as measured by the accident rate. A larger number of streets in the system mean that some vehicles end up travelling through more intersections than they would have in a smaller sized city. For intelligent vehicle agents, intersections are where the most communicating takes place. In intersections, they communicate with other vehicle agents coming from all directions whereas away from intersections toward the middle of the city blocks, the intelligent vehicle agents are only communicating with vehicles travelling in the same and, possibly, opposite direction as they are. Thus, it is likely that the slight increase as the size of the city grows in the accident rate in settings with 100% intelligent vehicle agents is due to the fact that in such settings the intelligent vehicles may be experiencing more communication overload because they are travelling through more intersections where the communication burden is greater. In settings with no intelligent vehicle agents at all, there is no effect on the accident rate seen with different city sizes seen since the dumb vehicle agents do not have to deal with the increased communication load as they travel through more intersections in a larger city. The implications of this result indicate that for the intelligent vehicle agents' communication to be effective at reducing the accident rate in the system the agents need to somehow reduce their communication burden, otherwise, the agents may be pulled in too many directions by receiving too many messages. In our current agent design, intelligent vehicle agents will attempt react to each and every possible conflict as it receives the message about the conflict from another vehicle agent. The problem with reacting to communicated conflicts in this manner is that the intelligent vehicle agents attempt to deal with each conflict individually without considering the other possible conflicts they are aware of and have already dealt with. Thus by reacting to one communicated conflict, the intelligent vehicle agent may be undoing the steps he needed to take to prevent another conflict it had already dealt with previously. Such a situation likely occurs most often at intersections where the intelligent vehicle agents are receiving many more messages than they are away from the intersections, and thus we see the accident rate increases as the number of streets in a city increases and vehicle agents are driving through more intersections. Thus, in order for the intelligent vehicle agents' communication to be effective at avoiding collisions, the agents need to remember the previous conflicts that they have already dealt with so they do not attempt to avoid one conflict and step into another that they had already taken action to avoid.

Next we investigated the effect of the size of the city on the traffic flow efficiency of the system. The results are presented in figure 2.

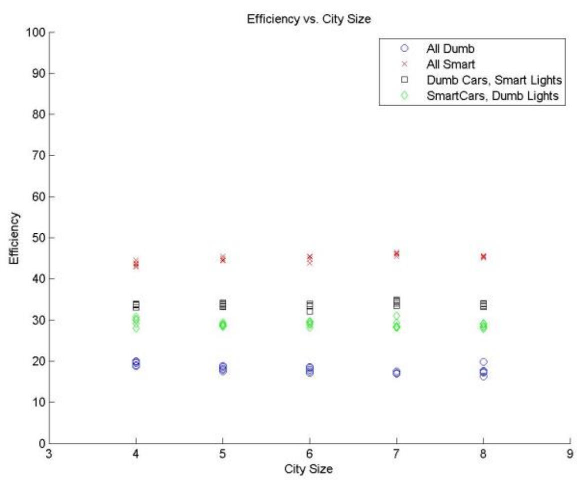

Figure 2: The Effect of City Size on Traffic Flow Efficiency

From these results, several observations stand out. First, the size of the city had no effect on the traffic flow efficiency in the system, and this result was true for all system settings regardless of whether intelligent agents were used. This result is discussed further in the section on the results of our experiments about the effects of the proportion of the different types of intelligent agents in the system had on the system's performance. Second, different system settings had different efficiency ratings. These results refute our hypothesis that an increased number of streets in the system would lead to a slight decrease in efficiency. This result is likely because even though some vehicles were travelling through more intersections in a larger city, they are not spending any more time stopped than those cars that have fewer intersections to travel through. The implications of this result are that the amount of time it takes the vehicles to arrive at their destination is linearly related to the number of intersections that those vehicles have to drive through—each intersection imposes a constant stopping time on a vehicle as it travels to its destination. This is contrary to our initial expectation that amount of time stopped would compound as vehicles travelled through more intersections.

B. Block Size Experiment

In this experiment, we investigated the effect of the city block size, or spacing between intersections had on the safety and efficiency of traffic flow in the system. The size of the city blocks ranged from 15 to 50 repast grid units. The block size of 50 was used to test the effects of a more extreme block size. We used all combinations of 100% dumb and intelligent agents in order to test whether the effect of the city block size was consistent for all system settings. First we looked at the effect of varying the city block size had on the system's traffic safety as measured by the accident rate in the system. Figure 3 displays the results.

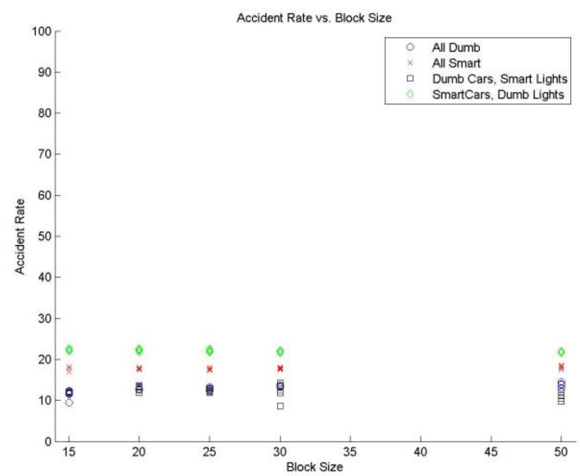

Figure 3: Effect of Block Size on Traffic Safety

Several observations stand out from these results. First, different system settings with different proportions of intelligent agents in the system give different measurements of traffic safety. This observation is discussed further in the section on the results of our experiments on varying proportion of intelligent agents in the system. Second, block size has no effect on traffic flow safety, and this result holds for all system settings no matter whether or not there are intelligent vehicles. This result contradicts our hypothesis that increased block size would lead to an increase in safety since the vehicles would be spending more time travelling away from intersections where collisions are more likely to occur. This result is likely due to the fact that in a city with a given number of streets, each vehicle agent has to travel through the same number of intersections to reach its destination no matter what the block size is. Thus even though a vehicle spends less of its total travel time travelling though intersections in cities with large block sizes than in cities with small block sizes, the vehicle agents are still having the same number of encounters with intersections. The implications of this observation are that, as expected, the most dangerous places on the road are the intersections where many different vehicle agents are travelling through in all directions. If in fact intersections were not where the most collisions occur, then we would have expected to see more collisions occurring in the open space away from intersections as the block size, and therefore the amount of that open space, increased. However, the accident rate remained constant even as the block size increased. This implies that most collisions occur due to actions taken near the intersections such as making left and right turns, merging into turn lanes, stopping for signals, or even simply travelling straight through the intersection. Thus, to reduce the accident rate in the system, the intelligent vehicle agents would need to focus on how to negotiate actions taken near intersections because those are the most dangerous places in the city.

Next we looked into the effect varying city block size had on the traffic flow efficiency in the system. The results are presented in figure 4.

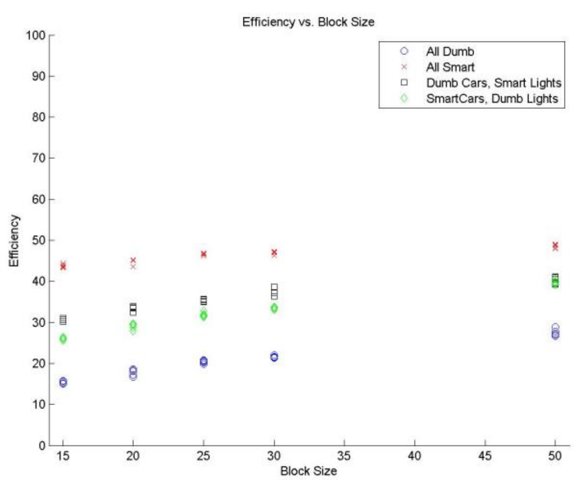

Figure 4: The Effect of Block Size on Traffic Flow Efficiency

From these results a few observations stand out. First, again we can see that different system settings had different effects on the efficiency of the system, and these results are discussed further in the section on the results of the experiments varying the proportions of intelligent agents in the system. Second, as the size of the city blocks increased, the efficiency of the system's traffic flow increased and these effects were consistent for all system settings. This result supports our hypothesis that the traffic efficiency would increase as the city block size increased. This observation is likely because of the fact that vehicle agents driving in cities with larger block sizes spend more of their time driving at their full speed in the areas away from the intersections rather than having to constantly slow down and stop for signals as vehicles have to do when driving in cities with smaller block sizes. This result is also likely due to the fact that the vehicle agents will not enter an intersection if they will be forced to stop in the middle of the intersection, so a vehicle may still even be stopped at a green light if it sees that there is not enough room for it to go through the intersection and move out of it. In settings with larger block sizes, this is less likely to occur since there is more room for vehicles on the other side of the intersection if they travel through the intersection. Thus vehicle agents are spending less time stopped in cities with larger block sizes. The implications of this observation are that the areas of the city that have the biggest impact on traffic efficiency are the intersections since when intersections are spaced closer together, traffic efficiency suffers. Thus traffic system designers must consider the spacing of intersections when seeking to increase traffic flow efficiency.

C. *Smart Car % Experiment*

In this experiment, the percentage of smart car agents was varied from 0% to 100% in steps of 25%, to determine its effect on system efficiency and safety as represented by number of accidents per car. This was run in two settings; one employing all smart signals and another employing all dumb signals. The following graphs present the results.

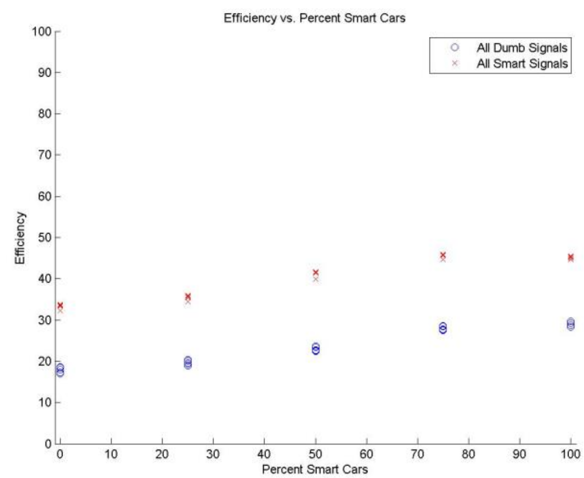

Figure 5: The Effect of Percentage of Smart Cars on Traffic Efficiency

We observe the following with respect to its effect on efficiency: All combinations of percent of smart cars with smart signals show significantly better efficiency over all combinations of smart car percentages with dumb signals. As the percent of smart cars in the system increase, overall efficiency increases almost linearly, but not significantly, with both smart and dumb signals. These results are consistent with our hypothesis that increasing the proportion of smart cars in the system would increase the system efficiency. These observations could be because with more smart cars in the system, they are able to successfully communicate their travel intentions and are able to plan out routes more effectively. As each smart agent gets more and more intelligent agents to communicate with in its local neighborhood, it is able to prevent running into occupied cells and traffic jams better. Another possible explanation for the better performance of a system with a higher proportion of intelligent vehicle agents is that intelligent vehicle agents attempt to avoid grid locations where other vehicle agents plan to go. Thus, as intelligent vehicle agents attempt to avoid accidents by avoiding the grid locations of where other vehicle agents plan to go, they may also be avoiding busier roads allowing them to travel to their destination more quickly. Combining this increase in efficiency with the better intersection planning bought about by the smart signals; the smart cars are able to travel to their destination with almost no stoppings in its entire route; thus the higher efficiency for this combination. On the other hand, the smart cars do not get the advantage of decreased stopping times at the intersections when the signals are dumb. This forces them to spend time waiting at intersections for the green light.

It can thus be inferred that when each smart car agent is able to perceive its neighbors better with communication, it is able to attain a better view of its surrounding. The smart cars lose a little autonomy in this scenario wherein they have to consider giving way to another smart car that needs the street. However, this loss of local autonomy is more than significantly compensated by the increase in global cohesion, creating a system setting where there is a much smoother flow of traffic than normally. A telltale sign of this is also seen in the fact that the system is almost never able to hit the maximum car ceiling parameter of 150; the car count in the system stays around the 50% mark (75); this is because the cars do not stay long enough inside the system for the ceiling to be attained, reaching their destinations faster. The implications of these results are that simply by attempting to avoid each other, the intelligent vehicle agents are able to arrive at their destinations more quickly. This implies that by simply attempting to avoid each other at a very local, the intelligent vehicle agents may be able to find open lanes and roads more easily, thereby distributing themselves across the city streets more evenly and thereby allowing for smoother traffic flow. These results were not what we expected and offer possibilities of exploring the ways the local decisions on the part of the intelligent vehicle agents could lead to more efficient traffic flow.

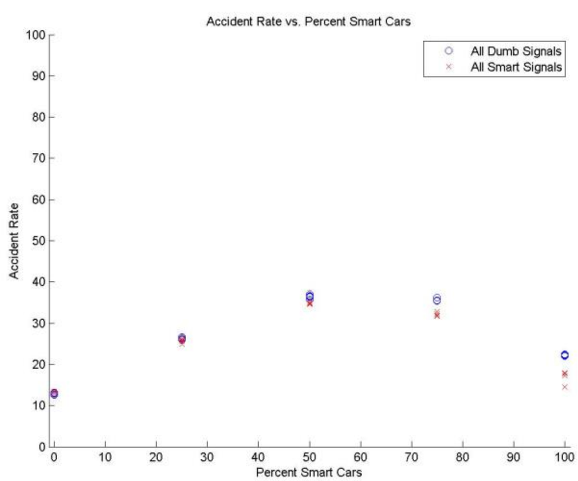

Figure 6: The Effect of Percentage of Smart Cars on Traffic Safety

The effect of the percent of smart cars on the traffic safety is visualized using the plot of Smart Car Percent vs. Accident Rate in figure 6. The higher the accident rate, the less safe the system is.

Compared to the effect on system efficiency, the effect of increasing the percent of smart cars on the safety of the system is quite surprising and interesting. We had hypothesized that the overall safety would increase. While safety in fact decreases from 0% to 50% intelligent vehicle agents in the system, after there are more than 50% smart cars in the system, the accident rate begins to improve again. Another observation we have is that there is insignificant difference between when the smart signals and dumb signals are used.

This is justifiable when we consider what situations are considered as an accident in the system. When a car swerves to avoid another car near the part of the street just before and in the area where the middle lane consists of a median, an accident is created. Accidents are also created when two cars attempt to occupy the same street cell location. On examining the graph, we see that there are very little accidents with 0% smart cars in the system, that is, the system is composed entirely of dumb cars. When more and more smart cars are added the accident rate increases up until 50%. This increase is explained by the fact that there will be increasing number of cars in the system that attempt to communicate with their neighbors about their locations and travel plans. When there are more dumb cars than smart cars, this communication is wasted and turns out to be detrimental, since the dumb cars will change lanes unexpectedly at turns and constantly overcome the smart car's planning strategy. This increase is also explained by the fact that in the way the intelligent vehicle agents are designed, they assume that when they send a message, other vehicle agents nearby will receive it and act accordingly. Thus, the accident rate increases as the intelligent vehicle agents make false assumptions about how the other vehicle agents will behave. This happens until the 50% mark. Once there are more smart cars than dumb cars, the communication becomes useful since there are more neighbors for each of the smart cars to share their positions and routes with. More cars successfully understand this communication and hence accidents begin to reduce significantly and reach a stead state when there are 100% percent of smart cars. A minor difference is observed towards the end when there 100% smart cars between using dumb signals and smart signals due to the decrease in intersection collisions bought about by smoother intersection traffic due to smart signals.

This particular scenario is extremely useful in illustrating the effect of mixing local autonomies on global cohesion. When multiple types of agents exist within the system with varying levels of local autonomy, the agents with lesser autonomy, such as the intelligent vehicle agents who are required to react to identified conflicts, face significant damaging effects when attempting to negotiate and communicate with agents with more autonomy, if the latter refuses to consider all of the former's communications. This effect will reduce global cohesion, since we will have conflicting agents within the system attempting to attain only their local goals. The implications of this result are that when agents communicate, there must be some sort of protocol for dealing with the case when agent for whom the message is intended does not receive the message since communicating agents making false assumptions about whether other agents receive their messages is worse than having no communication at all. From this experiment we observe that the accident rate is actually a little higher when 100% smart cars are used, and in other experiments where other variables were tested, we have observed that settings with entirely intelligent vehicle agents have a higher accident rate than settings with none. We believe this is due to the design of the intelligent vehicle agents. When the smart cars receive a message and determine that there is a conflict, they act to avoid the conflict. When they receive another message about another conflict they react to avoid that conflict without considering the previous conflict they had just avoided previously. Thus, the intelligent vehicle agents may be reacting to one conflict, then receiving a message about another conflicting, and reacting to the new conflict in a way that puts it back into the first conflict it had just avoided. Thus the communication between the intelligent vehicle agents in our system actually proves to be detrimental because the agents attempt to react to every single piece of information, without considering other information that they may have learned in the past. Thus the agents are being pulled in too many different directions by all of the messages they have received. The implications of this are that when agents communicate, they need to have a way to aggregate the information they receive, remember what they have learned, and have a full picture if that communication is to be beneficial to the agents.

D. Smart Signal % Experiment

In this experiment we test the effect the proportion of smart signals in the system have on the efficiency and safety of traffic in the system. For this we vary the percentage of smart signals from 0% to 100% at 25% intervals with 100% dumb cars and again with 100% smart cars. The effect on efficiency is examined first.

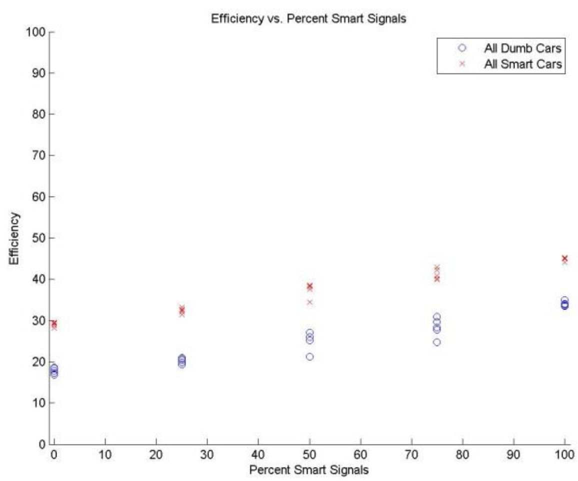

Figure 7: The Effect of Percentage of Smart Signals on Traffic Efficiency

The following observations are apparent. Increasing the percentage of smart signals increases the efficiency for both the cases with all smart and dumb cars. The efficiency is significantly higher when smart cars are used together with smart signals. The rate of increase of efficiency is linear but not steep.

We had hypothesized that increasing the percentage of smart signals in the system would increase the efficiency and it is observed that it is true.

When more and more smart signals are introduced in the system, they are able to route the traffic better by attempting to reduce the waiting time for each car at their respective intersections. This enables them to allow smoother passage for incoming traffic by preventing the buildup of traffic in any direction. This creates the base level of increase in system efficiency for both types of car agents. When coupled with the inherent efficiency bought about by the smart cars as discussed earlier, the system's efficiency increases further. It was visually observed that the cars would almost fly through the intersections hardly waiting for a red light since the smart signals were prepared to let them through without causing gridlock to increase their own utility.

This experiment further solidified the fact that for multi agent systems, high global cohesion can be achieved with little loss of local autonomy. By allowing signals to act towards increasing their own utility, each signal created a section of the street block that cars could move smoothly through. The additive effect of this throughout the system bought about the smoothly flowing traffic – a very desirable emergent behavior in the system. The implications of this result are that by allowing the traffic signal agents full autonomy and flexibility, their local decision making can lead to the emergent behavior of efficient system traffic flow.

Next we examine the effect that the percentage of smart signals has on the system safety. The results are presented in the Percent Smart Signal vs. Accident Rate plot in figure 8.

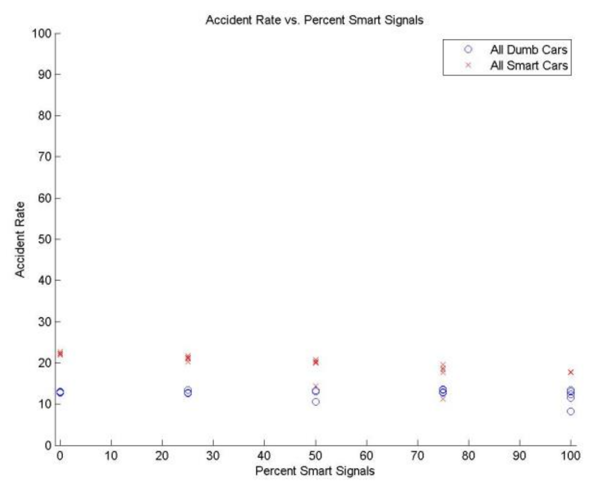

Figure 8: The Effect of Percentage of Smart Signals on Traffic Safety

The following observations are made: Increasing the percent of smart signals in the system reduces the accident rate, but not very significantly for both the all smart cars and all dumb cars scenario, with a slighter lesser accident rate for the dumb cars. This is a very peculiar observation owing to the fact that we believed that smart cars and smart signals would decrease the accidents in the system to a level lesser than with the dumb cars.

Our hypothesis that the overall system safety would increase thus holds true.

The decrease in accidents for each case is due to the fact that the overall routing due to the introduction of more smart signals is better. Cars are able to move in their desired direction without having to constantly wait for clearances. The intriguing nature of as to why the dumb cars have lesser accidents when compared to smart cars may be explained thus. Smart cars have to deal a lot more information at each intersection, since the smart signals can efficiently allow more cars to pass through the intersection at a given time. This increase in the amount of information to be processed by each smart car agent, combined with the increased number of cars simultaneously passing through the intersections, equates to a more chaotic view for each smart car at the intersection. It may simply be too late before it makes a decision on how it should plan its next step and would end up directly in the path of another agent that has already made its decision. Our code for how each smart agent communicates and negotiates the right of way may not be the most effectively scalable and this might be creating the processing bottleneck.

From this experiment we learn that making in the agents intelligent is not the only criterion when it comes to making the system safer. We must also design the agent's primary negotiation and conflict resolution logic in a scalable manner. Our smart car agents lost an important characteristic required of every multi agent system in this scenario-reactivity; they were unable to react to the changes in their surroundings in a timely manner forcing them to enter states of constant conflicts (accidents).

E. Comm. Dist. Experiment

In this experiment we were attempting to see the impact of communication distance on the effectiveness of the smart cars. Thus, unlike other experiments, we included no dumb cars, only smart and dumb signals. We varied the communication distance from 1 to 5.

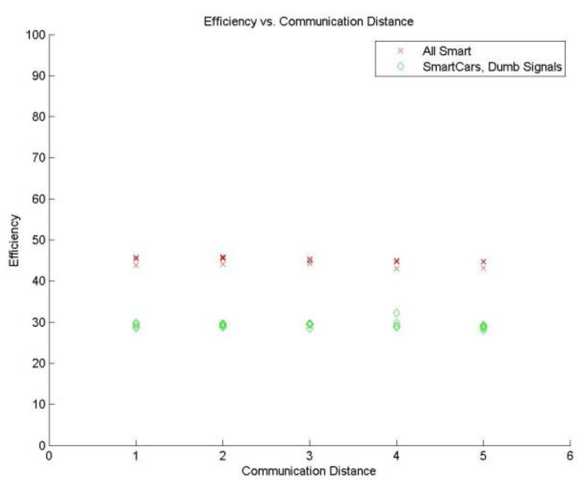

Figure 9: The Effect of Communication Distance on Traffic Flow Efficiency

We notice that communication distance seems to have little impact on the efficiency that is achieved by smart cars. There is a nearly flat line of efficiency values across all distances tested. As expected, we see smart signals improving efficiency over dumb signals, but both scenarios behave identically with respect to communication distance.

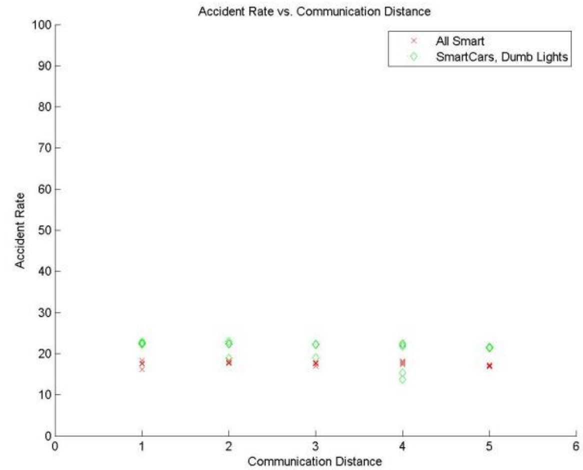

Figure 10: The Effect of Communication Distance on Traffic Safety

For the most part, we see that communication distance does little to change the accident rates. Again we see smart signals achieving slightly better than dumb signal. However, we now see the data from dumb signal seeming to split. These may just be misleading outlier points, but it seems more likely that they indicate something meaningful occurring.

Aside from a handful of points in the accident rate graph, these results are quite surprising, as we would expect at least some noticeable variation to occur, even if negative. We have seen that smart cars have a moderate impact on the system, though not always an ideal one. It would seem that the real difference must occur between communicating at a distance of 1 and not communicating at all. With this in mind, we consider again the agent's usage of communication. Messages are sent from smart cars to tell other cars the location they are next planning to move to. If two smart cars realize they are going to be moving into the same place then one will try to readjust its course. However, in our simulation, cars only have a speed of at most 1 cell/tick. So the only cars that are in danger of colliding within a 1 tick window of time are cars that are within a distance of 1 of each other. So it seems likely that when the communication distance is increased, messages are being sent to cars that are not at all affected by the message contents. Messages are sent a longer distance but behaviors of the cars remain unaffected. Our hypothesis that the effect would begin to wear off at a point wasn't entirely wrong; that point simply occurred at a much smaller value than we were expecting.

As for the accident rates in the dumb signals scenario, it is possible we are sometimes able to hit a "sweet spot" in the number of cars at the intersections such that the flow proceeds smoothly and accident free. Smart signals may potentially swap traffic directions rapidly, allowing many conflicting messages to be delivered. Further analysis would be needed for a definite answer. Communication distance may become more relevant if other factors in the simulation where changed. If car speeds were greater, then it could become possible for more distant cars to move towards potential collisions. In this case, longer distance messages would be useful. Long distance messages could also be useful if cars communicated higher-level navigation information rather than just their next intended move. For example, if cars communicated their planned routes, cars could self-regulate and avoid heavy congestion on by avoiding busy streets. A larger communication range would ensure cars were sending and receiving a larger amount of data with which to make their decisions.

VI. Future Work

In this section we discuss some ideas for future work with our traffic simulation system. There are a few areas that we would like to improve and explore further

Intelligent Vehicle Design

We would mostly like to improve the design of the intelligent vehicle agents so that they may be more effective at avoiding traffic collisions. The communication between agents proved to be much more complex than what we had anticipated, and in fact the communication in the current design of the intelligent vehicle agents actually leads to more collisions, rather than preventing them. We would like to improve the effectiveness of the vehicle agent's communication by allowing them to remember past messages and how they responded to recent conflicts so that the agents may have a more complete picture of what the other vehicle agents around them are planning to do rather than naively responding to each conflict at a time without regard to what actions they have already taken to avoid another conflict.

Environment Design

We would also like to model the effects of accidents more realistically by requiring that emergency vehicles have to arrive at the scene of the accident first in order to clear the accident. We also would like to make our environment more realistic by including more than just a perfect grid of streets. We would like to also include environmental objects such as highways, freeways, interchanges, off ramps, non-straight roads, and road construction.

Path Planning

We would like to use the communication between the intelligent vehicle agents to increase the efficiency of traffic flow in the system, by allowing the vehicle agents to plan their path based on what other vehicle agents communicate to them. For example, vehicles could communicate where there are traffic jams or which roads they are planning on taking so that other vehicle agents can avoid the busier roads. Combining both the communication of the intelligent vehicle agents and the local decision making of the intelligent signal agents, we expect should lead to very high traffic flow efficiency.

VII. Conclusion

City traffic simulation and/or its management is an ideal setting in which multi-agent systems can be employed. It has very clearly defined agents (such as cars) each with their own goals (destinations). It is a setting where the actions of individual agents could have large repercussions on the system and the global behavior that eventually emerges. A single car may crash and gridlock the entire system or a single signal may quickly direct traffic and help unload congestions from the streets. However, more often than not, it is the complex interactions of the agents within the system that will determine the outcome.

Through our experiments, we measured the impact of the environmental properties of city size and block size. Both had minimal effect, though increased block size did slightly increase efficiency. The focus of our experiments, the percent of smart cars and signals, proved more interesting. As anticipated, smart signals improved both the safety and efficiency of the system. The smart cars improved efficiency but surprisingly ended up with more accidents. Conceptually we still believe the smart cars should be beneficial in all respects. There exist certain implementation details to be corrected such as ensuring that the conflicting messages are addressed in a timely manner and be scalable. The communication distance experiments were also surprising in that the distance made little difference. Communication distance could be made more important if higher-level navigation information was exchanged as opposed to simply the location of the next planned move.

On the whole, the added intelligence was beneficial to the city's overall traffic condition. In fact it was quite visually apparent that the simulations utilizing intelligent agents were much smoother and free from congestion while the standard city simulation had a lot of stop and go traffic and backed up intersections. Improvements still could be made, particularly in the area of smart cars. Despite this, our system was still able to demonstrate the potential benefits of leveraging the cars' and signals' local autonomy to arrive at a generally safe and efficient global coherence.