"Establishing and Populating Socially Conscious Autonomous Highway Traffic"

Technische Universität Hamburg-Harburg

Master Thesis

Christian Feist

October 31, 2011

Supervisors:

Prof. Dr. Ralf Möller

Prof. Dr. Dieter Gollmann

Hamburg University of Technology

Software Technology Systems

`http://www.sts.tu-harburg.de/`

Schwarzenbergstraße 95

21073 Hamburg

Germany

Dedicated to All Good Things

Declaration

I, Christian Feist, solemnly declare that I have written this master thesis independently, and that I have not made use of any aid other than those acknowledged in this master thesis. Neither this master thesis, nor any other similar work, has been previously submitted to any examination board.

Hamburg, October 31, 2011

Christian Feist

Abstract

The landscape of human mobility faces tough challenges that will require us to redefine the way we travel in future. In a world where we have revolutionized virtual presence by modern communication methods, our professional and individual need for physical relocation has not significantly decreased. Instead, our ever growing worldwide populations make heavy use of their expanding access to vehicles that become quicker and cheaper. The volume capacities for individual transport connections approach their limits and our concentration and situation assessment capabilities do already exceed theirs regularly. Human error leads to fatalities, congestion and suboptimal energy efficiency and has consequently become a central topic of automotive research.

The very same sensor data fusions nowadays driving the progress of driver assistant systems will eventually become the cognitive base enabling the driver's potential replacement by an artificial intelligence. What has been successfully applied to subway trains or even buses might therefore solve the problem of human limits, while preserving complete freedom over travel time and location. Furthermore, the advancements in communication technology and car telemetry enable these systems to know more about the road than what their own sensors pick up.

This thesis aims to motivate and define the environment of an autonomous highway scenario and to propose possible rules for a group-based agent design as a superior alternative to human drivers and individually acting autonomous agents. In the focus of our research lies the question of how socially conscious decisions could be made in a group of autonomous cars and how they impact important factors of their passengers' traveling experience.

Contents

1 Introduction

Today's automotive traffic can be called a microcosm in its own right. It comprises a variety of different subjects following a distinct set of rules and communicating in their own non-verbal language. The established main currency is time and every participant's goal its minimization using the limited dimensions of influence on his or her way from A to B. It is an environment that grows in participants faster than it does in infrastructure and their *behavior* has far-reaching effects on the world beyond the roads. It influences economic efficiency, greenhouse gas emissions, and our well-being from physical fatalities to mental road rage. Apart from the structural issue of limited flow capacities or the technological difficulties in migrating away from fossil fuels, human irrationality in the core decisions of each particle of the stream can be often be shown to create and amplify phenomena leading to suboptimal usage of the bandwidth available.

The author's interest and experience in the world of automotive engineering and fascination for the young scientific field of artificial intelligence has sparked the idea of the modeling and evaluation of a future highway system that bases individual decisions on the unimpaired situation assessment of distributed sensors and the mathematical[1] rationality of autonomous agents. Identifying modern developments in cyber-physical systems (CPS, see [Bro10]) and driver assistant technology, we will propose a flexible system that uses available infrastructure and obeys current highway regulations. We will formulate a business case and integrate these economical assumptions together with hardware requirements, usability concerns, and ecological opportunities into an overall concept for autonomous travel. In the focal point of our work, we will show how the various interdisciplinary considerations could be formalized as criteria in the group-based lane change decisions of several locally connected agents, to exploit shared sensor information and enforce priority-based individual privilege for various use cases. We then implement core parts of the suggested behavior and evaluate empirical data from this simplified simulation to gain early qualitative estimations of the system's potential and identify further optimization and expansion opportunities.

1.1 Motivation

Even though it still requires some significant breakthroughs in perception technology and debates on political and legal principals, the potential benefits of our approach have a solid base in today's car engineering and current trends fuel the search for potential radical changes to our highway traffic. In the following, we will outline some of the issues observed today.

[1] one could say "cold" in its most positive meaning

1.1.1 Traffic Congestion

The problem of traffic congestion has many different appearances and causes. It covers terms like "stop and go" traffic or traffic "jams". While the detailed study of its roots is beyond the scope of this thesis, researchers from the field of traffic science provide some interesting arguments that show the potential of autonomous agents to relax today's situation:

- Vehicles equipped with adaptive cruise control (ACC) have been shown to increase traffic flow on highways significantly even with low penetration rates, e.g. by [KTH10].
- Individual irrational behavior such as harsh braking or risky lane switching can initiate congestion because the disturbance caused by this single event is amplified [Tre11] while being propagated backwards[2].

And today's situation desperately needs relaxation, as current statistics show. In its 2010 annual report on congestion [ADA11], the German automobile club (ADAC) states a number of *185,000* occasions of congestion with a accumulated length of *400,000km*. The yearly increase from 2009 in length was at *14%*. Figure 1.1 shows the kilometers of congestion per kilometer of length for the most affected highways for one year. Schreckenberg[dpa09] estimated the time spent in German traffic jams to a figure that translates to *2.4 days/year/person*.

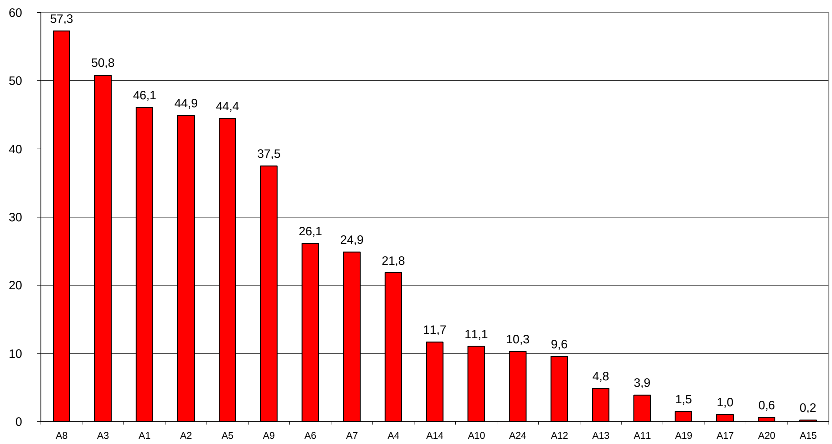

Figure 1.1: Reported kilometers of congestion per road kilometer (Highways A1 to A24)[ADA11]

American figures from [TTI11] for 2010 identify further impacts of congestion, such as an average of *750USD* annual cost per citizen caused by the delay and the waste of fuel while standing still. Figure 1.2 holds an excerpt of the results for America's most populated areas and details on their interpretation.

1.1.2 Active Automotive Safety

While today's driver assistant systems already provide a solid base for autonomous highway driving (discussed in Chapter 2.4.1) with regards to sensor technology, their "assisting" character and current legal regulations limit their protective capabilities. Current flaws that can be addressed by our approach include:

[2] you could informally call this a "butterfly effect"

Urban Area	Yearly Delay per Auto Commuter		Travel Time Index		Excess Fuel per Auto Commuter		Congestion Cost per Auto Commuter	
	Hours	Rank	Value	Rank	Gallons	Rank	Dollars	Rank
Very Large Average (15 areas)	**52**		**1.27**		**25**		**1,083**	
Washington DC-VA-MD	74	1	1.33	2	37	1	1,495	2
Chicago IL-IN	71	2	1.24	13	36	2	1,568	1
Los Angeles-Long Beach-Santa Ana CA	64	3	1.38	1	34	3	1,334	3
Houston TX	57	4	1.27	6	28	4	1,171	4
New York-Newark NY-NJ-CT	54	5	1.28	3	22	7	1,126	5
San Francisco-Oakland CA	50	7	1.28	3	22	7	1,019	7
Boston MA-NH-RI	47	9	1.21	20	21	11	980	9
Dallas-Fort Worth-Arlington TX	45	10	1.23	16	22	7	924	11
Seattle WA	44	12	1.27	6	23	6	942	10
Atlanta GA	43	13	1.23	16	20	12	924	11
Philadelphia PA-NJ-DE-MD	42	14	1.21	20	17	18	864	14
Miami FL	38	15	1.23	16	18	16	785	19
San Diego CA	38	15	1.19	23	20	12	794	17
Phoenix AZ	35	23	1.21	20	20	12	821	16
Detroit MI	33	27	1.16	37	17	18	687	26

Very Large Urban Areas—over 3 million population.
Large Urban Areas—over 1 million and less than 3 million population.
Medium Urban Areas—over 500,000 and less than 1 million population.
Small Urban Areas—less than 500,000 population.
Yearly Delay per Auto Commuter—Extra travel time during the year divided by the number of people who commute in private vehicles in the urban area.
Travel Time Index—The ratio of travel time in the peak period to the travel time at free-flow conditions. A value of 1.30 indicates a 20-minute free-flow trip takes 26 minutes in the peak period.
Excess Fuel Consumed—Increased fuel consumption due to travel in congested conditions rather than free-flow conditions.
Congestion Cost—Value of travel time delay (estimated at $8 per hour of person travel and $88 per hour of truck time) and excess fuel consumption (estimated using state average cost per gallon for gasoline and diesel).
Note: Please do not place too much emphasis on small differences in the rankings. There may be little difference in congestion between areas ranked (for example) 6th and 12th. The actual measure values should also be examined.
Also note: The best congestion comparisons use multi-year trends and are made between similar urban areas.

Figure 1.2: "What congestion means to you, 2010" [TTI11]

- Since regulations for driver assistant systems still require the driver to be in control of the final decision for a maneuver, the natural human reaction time and perception limits cannot be completely neutralized by the system.

- There are psychological studies that have reported that drivers using advanced assistant systems have been subject of behavioral adaptation [RH04]. The experienced increase in safety can be overrated and leads to more aggressive driving, identifiable by a decreased headway or higher speeds.

- As mentioned in [Bro10], stand-alone ACC in particular can cause rear-end collisions by longitudinal oscillation causing the last in a sequence of equipped vehicles to exceed its braking capability.

1.1.3 Emissions

Carbon dioxide emissions are not only a problem while standing in a traffic jam. Since its generation is proportional to the vehicle's fuel consumption, it is also directly linked to its acceleration as detailed in [ZSFB95]. Therefore, the reduction of acceleration time while traveling is a matter of environmental and financial benefits. For electric vehicles, the same holds with regard to battery life and greenhouse gases associated with the generation of electricity. Cooperative behavior among agents can be optimized for smooth progression, limiting the time and intensity of acceleration needed.

Furthermore, drag effects can be exploited to reduce consumption by disbanding safety gaps between connected autonomous vehicles. [ZSFB95] have shown that drag is responsible for about *80%* of fuel consumption at a usual highway speed of *130km/h* and can be reduced to *55%* of its value when building a 4-vehicle platoon, with longer convoys suggesting even better coefficients.

1.2 Comparison to Current Research

The field of traffic science has become more important and is growing with the severity of the problems mentioned. Many ideas from different projects and research efforts have influenced choices made during this study or provided unsolved problems that require a more radical approach. Some concepts and their similarities and contradictions to our approach are covered in the following sections.

Intelligent driver model (IDM) and simulations by (TH02): As will be described in the following, we base some of our low level behavior on this well-studied method of modeling lane-following behavior. The simplicity of its parameterization and implementation outweigh its limits in accuracy in extreme cases that are not in the scope of this thesis. The acceleration values resulting from its inner calculations carry a meaning, that can be used as a decision factor by rational agents. However, unlike its creators, we do not try to *mimic* human behavior or explain congestion situations. Since we are in control of our connected vehicles, we rather use the IDM as one indicator in the *design* of a rational behavior involving more than just following the car in front.

California PATH program (Tam00): This research effort under supervision of the University of California follows the idea of automation on highways to increase throughput and lower emissions. Compared to the assumptions of this thesis, the approach has some severe flexibility shortcomings. It relies on magnets installed in special lanes only to be used by autonomous vehicles, that are separated from manually driven lanes using high barriers. This leads to complicated structures for entry and exit, as depicted in Figure 1.3. Position in the platoon is then fixed to achieve drag reducing effects. Designed for American highways with tight speed limits and many lanes, the system does not integrate well into European standards and road sizes.

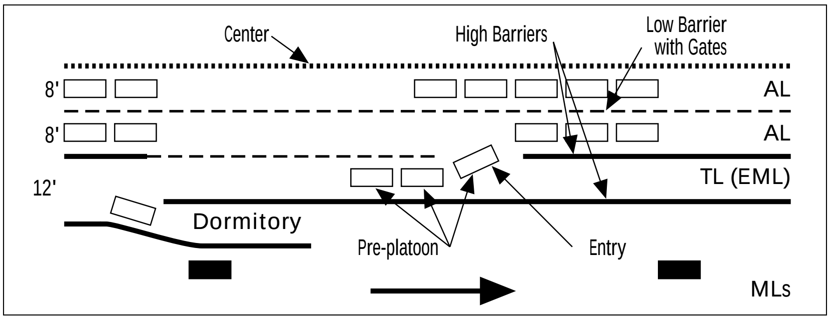

Figure 1.3: Entry to the PATH automated highway system [Hit95]

European SARTRE Platooning Programme (TR10): As the name suggests, the SARTRE project focuses entirely on convoy formation. The idea is to establish a business case that teams up a professional driver and up to five participants that will join into a platoon behind this manually driven lead vehicle as pictured in Figure 1.4. The following vehicles then engage fully autonomous longitudinal control with significantly reduced headway and fixed order. Even though this approach is a lot easier to adapt on current roads than PATH, it does not remove human error. Instead, it leaves non-drivers under the responsibility of a paid person very much like a bus driver. This raises a lot of issues regarding mistakes by these drivers or rules of behavior when multiple convoys meet. The business model might include charging for electric vehicles and aims for the same platoon drag reduction as its Californian precursor. Our core component of lane switching is not part of SARTRE, as well as group decisions, arbitrary formation of groups, or coping with different speed desires.

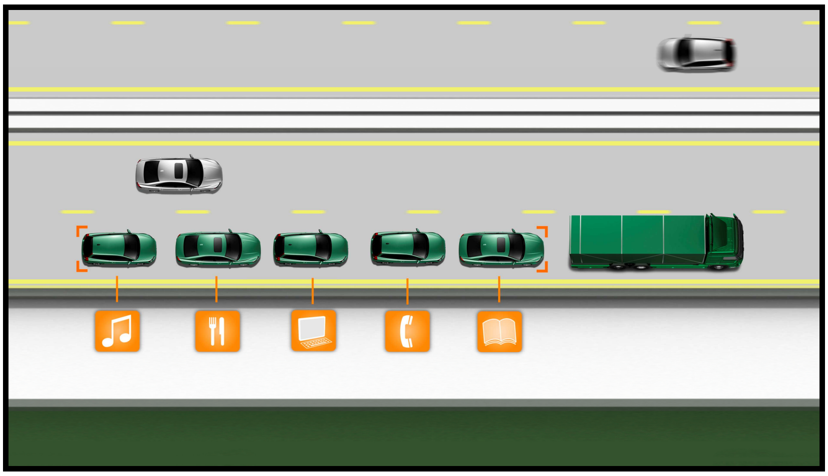

Figure 1.4: SARTRE platoon illustration [TR10]

Google autonomous driving project (Mar10): Not many scientific facts are known about the recent involvement of Google in the field of autonomous vehicles. Under the lead of Sebastian Thrun, who is also director of the Stanford Artificial Intelligence Laboratory and was involved with the DARPA challenge participations mentioned in Chapter 2.4.2, they have been developing individual autonomous cars that participate in highway and urban traffic. They are equipped with a 3D range sensor and always connected to a central database that holds descriptions of the world around them and is constantly matched with the perceptions of the vehicle. [Mar10] quotes Google's claims to have conducted over *140,000 miles* of test drives with only occasional human intervention and no accidents caused by the agent. It remains unclear, what business plan might be behind the project or whether any form of convoy driving or other cooperative behavior is involved. The inclusion of urban traffic scenarios is remarkable, but it is left open, how autonomous the vehicles really are in environments that have not been thoroughly mapped to the central database. This decentralization of decision making also raises questions of availability and privacy.

2 Current Foundations

In order to put the scenario and suggestions of this thesis into a solid context, this section covers some of the underlying theory and a view on the state of the art. Especially, important terms and the general view of the world in the field of artificial intelligence will be discussed to provide the necessary background for our lines of thought. However, the following sections will only give an overview of the different topics in the briefness adequate to our cause. Readers interested in the history, a full introduction, and the current trends of A.I. are referred to the excellent compendium *Artificial Intelligence: A Modern Approach* by Russell and Norvig [RNC$^+$96] that is the base for many assumptions presented here.
Apart from artificial intelligence theory, we will also cover some basic ideas and modern projects in the fields of autonomous cars and traffic science.

2.1 Autonomous Agents

Out of the many definitions of artificial intelligence, this thesis is centered around the following by [PM10]:

> "**Artificial intelligence**, or **AI**, is the field that studies *the synthesis and analysis of computational agents that act intelligently*."

This view of AI focuses on the concept of an *agent* as the central component of an intelligent system. Such agents are defined in [RNC$^+$96] as

> "[...] anything that can be viewed as perceiving its **environment** through **sensors** and acting upon that environment through **actuators**."

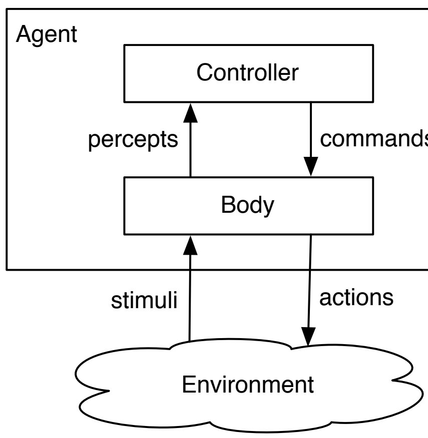

Figure 2.1: Agent system overview [PM10]

It is fully defined by the *agent function* that maps any given perception that is gathered by the sensors to an action that is given to the actuators for execution. This function is translated to an *agent program* that is the actual implementation of the function on the target hardware.
Figure 2.1 shows the structure of an agent with the agent program running on the *controller* and the sensors and actuators situated in the *body*. In our scenario, picture a car as the home of our agent. Its program will run on some computer hardware inside, the sensors available might be anything from radar to cameras, and the actuators basically include anything that you as a human driver could operate in the cockpit to influence the car's physical behavior, such as gas, brakes, turn lights, or even the horn. Intelligent in this case means rational, which implies that the agent is always doing the right thing given his perceipts of the current environment. The right thing is defined as the action (or sequence of actions) that makes the agent most successful by a certain performance measure. For any rational agent, an objective and well-defined performance measure is a mandatory requirement for its ability to make appropriate decisions. We will elaborate on the nature of measuring performance in Section 2.2

2.1.1 Agent Types

To be able to understand the design of our agents, this section will briefly introduce the four main agent structures that together hold the principals of almost all kinds of intelligent systems.

Simple Reflex Agents

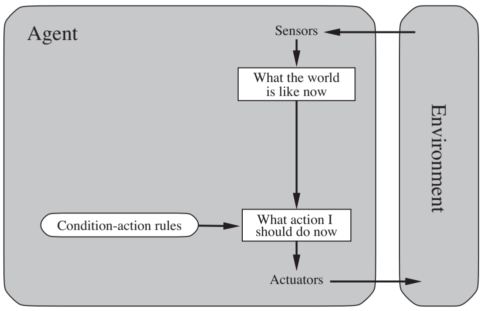

Figure 2.2: Simple reflex agent schematic [RNC+96]

The most trivial form of intelligence lies in the processing of simple condition-action rules. They work like human reflexes and map a single perceipt to exactly one reaction. This limits their flexibility by requiring the designer to specify rules for everything perceived that should make the agent do something. We will see later that we can implement some basic movement rules for our vehicles by simple reflex behavior.

Model-based Reflex Agents

If an agent does not only react to things directly visible to its sensors, but also needs to react to parts of its environment that are not observable right now, it needs to *keep track* of its perceipts. That means, it needs to be able to make good assumptions about once seen objects in its world if they get occluded temporarily. This requires the designer to specify a model of the world and code it into the agent, e.g., to predict the movement of another object with a constant speed v to appear at a position

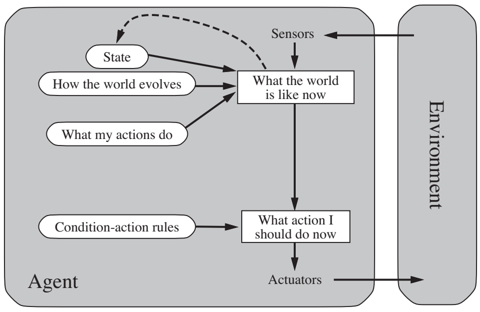

Figure 2.3: Model-based reflex agent schematic [RNC+96]

$x_0 + v_x \Delta t$ after t has passed. Many systems already in todays cars work as model-based reflex agents. Examples include lane assistants or adaptive cruise controls that make use of Kalman filters and vehicle movement models to predict future sensor input and even out sensor noise.

Goal-based Agents

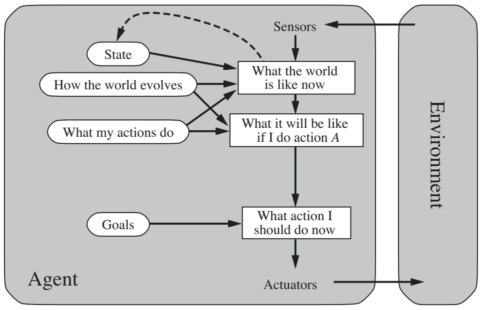

Figure 2.4: Goal-based agent schematic [RNC+96]

Some tasks require more flexibility than just a fixed model of the world. Instead of seeing an input or model observation and mapping to a single action that the agent should do, there might be situations that offer a choice of different actions possible for a single perceipt. A reflex agent cannot choose and always follows the path hard-coded by its creator. To make a good decision in situations that offer choice, we need an agent that has an idea of what state is desired in the end - a goal. Given this information it can choose the action (or search for a sequence of actions) that will help to reach the goal state. Apart from the model that tells the agent the effects of his actions, the only thing needed to be hard-coded by the designer is the end goal. The agent can then come to very sophisticated plans to reach it by searching the space of action sequences. Think of a navigation system as a possible example for such behavior.

Utility-based Agents

However, even though the navigation example is solved by just specifying a target, the quality of the routes might be far from rational or optimal. The problem is that the goal-based agent is just interested in *any* action or sequence of actions that will reach the target state and does not have a way of comparing and rating different ways to get there. While reaching the goal certainly makes the agent happy, saving a lot of gas or time should make it even happier. This "happiness" is called utility in more

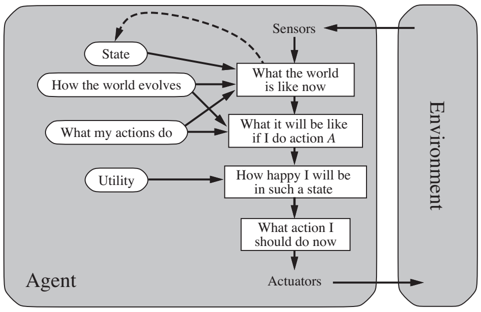

Figure 2.5: Utility-based agent schematic [RNC+96]

scientific terms and has proven to be a powerful performance measure even in environments where the reachability of states is uncertain.

2.2 Utility

The general concept of utility stems from game theory and economics and seeks to establish a measure of individual satisfaction based on a subject's current situation. It gained importance by the works of Bernoulli [Ber54], who argued for such a measure rather than fixed monetary values to determine the economic behavior of individuals. The following sections will discuss ways to interpret and calculate utility for its use as a decision base for intelligent agents.

2.2.1 Utility Function

The utility function is designed to map the state (or state sequence) at hand to a real number quantifying utility for this outcome. It therefore provides a *cardinal* value to use to determine the right trade-off between conflicting goals, which escapes a dilemma that goal-based agents cannot solve. Furthermore, the value of the utility function can be used in connection with probabilities. When it is uncertain whether a state can be reached, we can calculate an estimated utility that weighs up the likelihood of reaching the state with the importance of the corresponding outcome.

For this thesis, we will write

$$U\colon\ (\alpha, t, state) \to \mathbb{R} \tag{2.1}$$

for the utility of a *state* for entity α at time t. You may find time or state omitted to increase formula readability, where either of them is irrelevant or clear from the context. Generally, utilities do not have to be normalized to serve their purpose of making trade-off and preference statements, but you will see how establishing

$$U\colon\ (\alpha, t, state) \to [0,1] \tag{2.2}$$

can help to keep calculations more transparent when utility functions become more complex.

2.2.2 Multiatribute Utility

Complexity often rises from the fact that the utility of a state depends on a lot of different measurements made in the environment. You can see an example case from [RNC+96] in Figure 2.6 in the form of a so-called decision network. It shows the influences on a utility-based decision as nodes that are either chance nodes (○, representing random variables), decision nodes (□, representing a choice of actions), or utility nodes (◇, representing the utility function).

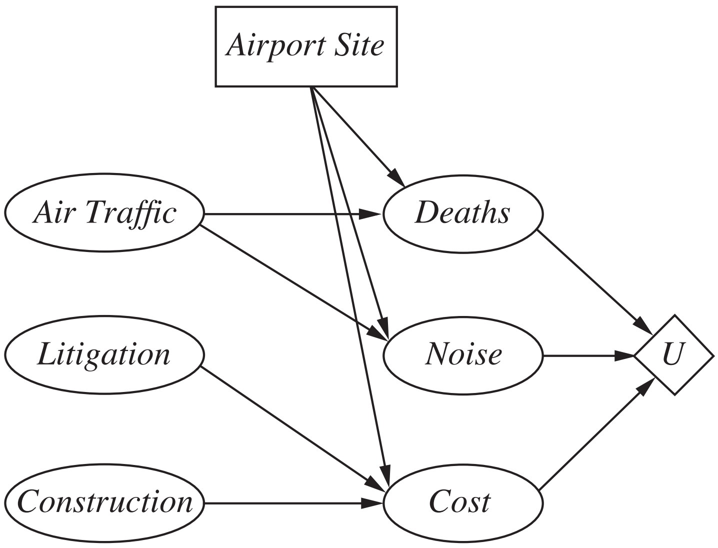

Figure 2.6: Example decision network [RNC+96]

The utility function now has to map these multiple dimensions to a single value that still preserves a rational order of preference. This can lead to very complicated non-linear expressions for the final utility. However, in many cases the perceived measurement from the environment do not influence each other in a way that affects preference. If the decision regarding different values of two variables A and B is the same, no matter the value of a third variable C, we say that A and B are *preferentially independent* of C. If this is the case for any combination of the variables (or *attributes*) influencing an agent's utility, they are said to exhibit *mutual preferential independence*[RNC+96]. This property is very useful in the construction of the final utility function for such agents because of a theorem established in [Deb59]. It states that for a set of attributes $X_1,...,X_n$ that is mutually preferentially independent, the preference behavior of a rational agent can be described as maximizing

$$U(X_1,...,X_n) = \sum_{i=1}^{n} V_i(X_i) \tag{2.3}$$

with V_i being an attribute value function depending *only* on the value of X_i. To be able to assign different importance to the attributes perceived, we will normalize them and introduce weight factors

w_i. Normalization of $V_i(X_i)$ is done by

$$V_i\left(X_{i_{worst}}\right) = 0 \tag{2.4}$$

$$V_i\left(X_{i_{best}}\right) = 1 \tag{2.5}$$

and

$$V_i\left(X_{i_{worst}}\right) \leq V_i\left(X_i\right) \leq V_i\left(X_{i_{best}}\right) \tag{2.6}$$

which maps its value to $[0,1]$ by fixing the worst and best outcome to 0 and 1 respectively. Following remarks in [PM10], we then construct a set of weights w_i so that

$$\sum_{i=1}^{n} w_i = 1 \tag{2.7}$$

and use it to define our final weighted utility function

$$U(X_1,...,X_n) = \sum_{i=1}^{n} w_i V_i(X_i) \tag{2.8}$$

This can be called a weighted additive value function and will serve as the base for our individual agent utility later in Chapter 4.3.1.

2.3 Group Decision Making

When a group of agents has been formed, the problem of making a decision essentially becomes a voting problem. We have a variety of different views on the situation, i.e., each agent has its own preference for reachable states. Now we need to find out what outcome is best for the group and in order to so, we will have to discuss what "better" actually means in a multi-agent context, a problem that is most traditionally encountered in economics and politics.

2.3.1 Preference Aggregation

The preferred action among the alternatives presented by the agents is determined by a *social preference function*, that needs to be carefully designed to fulfill different requirements. But to be able to evaluate those properties, we first need to demand the following assumptions about our participating agents:

- Each agent's individual preference is clear by an established personal order of beneficial outcomes. This means that at least ordinal utility can be determined for all reachable states. In line with [PM10], we will write $A \succ B$ for an agent *strictly preferring* outcome A over B.
- We can guarantee truthful communication of individual preference to the group.
- The result of the social preference function is enforced on all participating agents.

2.3.2 Pareto Efficiency

When talking about preference of multiple parties, the concept of Pareto efficiency offers some handy terminology [RNC+96]:

- If we can change an outcome in a way that no participating agent is put in a less preferred position, but at least one is made "better off", we call this step a **Pareto improvement**.
- A system is **Pareto efficient**, if there are no further Pareto improvements possible.
- When comparing two outcomes, we define **Pareto domination** of outcome A over B as the situation in which $A \succ B$ holds for all agents. B is said to be **Pareto dominated** by A.
- If there is no outcome B that dominates A, we say that A is **Pareto optimal**.

2.3.3 Arrow's Theorem

Using these definitions and our assumptions, we can formulate a list of desired properties for voting systems that has been proposed by [Arr70]. For a preference aggregation rule F for n voters or decision criteria on a set of outcomes Ω with $L(\Omega)$ being the set of all full linear orderings of Ω they are:

1. **Universality**: The function is defined as $F\colon L(\Omega)^n \to L(\Omega)$ for any input $(R_1,...,R_n)$ and always returns a complete output $L(\Omega)$ with $|L(\Omega)| = n$. It is deterministic for equal inputs: $\forall i \in [1,n]\colon (R_i = R_i') \Rightarrow F(R_1,...,R_n) = F(R_1',...,R_n')$
2. **Transitivity**: The list of outcomes in the result of F should be ordered and its elements should exhibit transitivity: $\forall R_i \forall R_j \forall R_k \in F(R_1,...,R_n)\colon (R_i \succeq R_j) \wedge (R_j \succeq R_k) \Rightarrow (R_i \succeq R_k)$
3. **Unanimity**: The function should prefer A over B if B is Pareto dominated by A in the choices of all agents: $\forall i \in [1,n] \forall R_i\colon (A \succ B) \Rightarrow (F(R_1,...,R_n) : A \succ B)$. This also implies non-imposition, i.e., every order is reachable given the right inputs, which makes the function F surjective: $\forall R \exists (R_1',...R_n')\colon F(R_1',...R_n') = R$
4. **Independence of irrelevant alternatives**: Social preference between A and B is only a matter of the individual preferences regarding these two outcomes and not influenced by any preference regarding C. We can formulate that for any given sets $(R_1,...,R_n)$ and $(S_1,...,S_n)$ we demand that if the order of A and B is the same for any R_i and S_i, they also have the same order in $F(R_1,...,R_n)$ as in $F(S_1,...,S_n)$.
5. **Non-dictatorship**: There shall not be an individual $i \in [1,n]$ so that the decision only depends on his preference: $\forall (R_1,...,R_n) \in L(\Omega)^n\colon F(R_1,...,R_n) = R_i$.

Unfortunately it has been proven by [Arr70] that there cannot be a voting system that fulfills all of these properties. We will discuss approaches to relax these criteria and find an appropriate way to determine preferred group actions for our scenario in Chapter 4.3.2.

2.4 Automotive Sensor Technology

It is important for us to stress that the availability of the main technologies for the vehicles in our concept can be expected within this decade. Autonomous driving is the next step after advanced driver assistant systems and many of the important base functions are in series production today. In this section we will briefly introduce some sensors and system solutions in current cars and take a look at possible future sensor additions currently in research.

2.4.1 Advanced Driver Assistant Systems

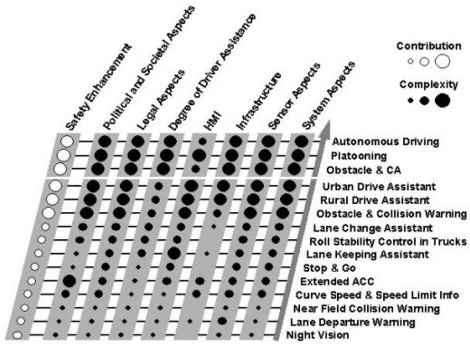

Figure 2.7: Advanced driver assistant systems road map [Sim06]

Figure 2.7 illustrates this point and introduces some advanced assistant systems that can be expected in future. To underline our assumption that autonomous driving in a highway scenario has strong foundations in current systems, we present the following collection of functions that are realized with production-grade sensors [Con].

- **Adaptive Cruise Control (ACC)** is the enhancement of conventional cruise control by adding range sensors. Regular systems have required the driver to manually adjust the automatically maintained speed to the traffic situation in front of the vehicle. ACC allows you to specify a desired headway and uses radar or lidar (light detection and ranging) sensors to track a vehicle up to *150m* in front and regulate speed by acceleration and deceleration as shown in Figure 2.8a. Should the required deceleration exceed a limit, ACC is usually coupled with a collision avoidance function warning the driver and preparing emergency brake capabilities.

- **Lane Departure Warning** systems rely on camera-based computer vision to detect line markings on the road and issue warnings when the path of the vehicle crossed them (see Figure 2.8a). By adding actuator capabilities, the concept can be developed into a **lane-keeping** mechanism that eliminates the driver's need to intervene while driving on one lane.

- **Blind Spot Detection** uses rear-facing radars or cameras to identify other vehicles approaching or lingering in the spaces left and right of the car that constitute dangers, should the driver choose to change the lane. They are usually coupled with forward-facing line detection to evaluate ego-movement. An example situation can be seen in Figure 2.8c.

(a) Adaptive cruise control (b) Lane departure warning (c) Blind spot detection

Figure 2.8: Modern driver assistant systems [Con], ©Hella

Recapitulating this list, it becomes obvious that the only open problems of autonomous highway driving lie in lane changes and collision avoidance.

2.4.2 Autonomous Vehicles

Today's most advanced autonomous vehicles expand this base set of functions by introducing advanced sensors, that are considered "near production", and high performance computers to manage the increased amount of incoming perception data. To convey current research directions to gain situation assessments, we will take a brief look on a series of competitions initiated by the American Defense Advanced Research Projects Agency (DARPA). The first was called DARPA Grand Challenge and held in the Mojave desert in 2003, where vehicles had to follow a dirt road course and avoid obstacles. 2005 marked the first year in which the 131-mile long course was successfully completed by an autonomous vehicle named Stanley, described in [TMD+06].

DARPA Urban Challenge

For 2007, DARPA came up with a different task for the participating teams: the Urban Challenge involved interaction with competitors on a course designed by the rules of residential area traffic. Vehicles had to autonomously take care of

> *"[...] passing parked or slow-moving vehicles, precedence handling at intersections with multiple stop signs, merging into fast-moving traffic, left turns across oncoming traffic, parking in a parking lot, and the execution of U-turns in situations where a road is completely blocked."* [MBB+08]

We will briefly examine some of the technology enabling *Junior*, Stanford's entry in the competition provided by the Volkswagen Electronics Research Lab, to handle these objectives. The car was equipped with the following sensors, also shown in Figure 2.9:

- For precise navigation, a system by Applanix provides real-time fusion of multiple GPS receivers that also include azimuth heading, inertial sensors, wheel odometry using a distance measurement unit, and the satellite-connection to a fixed base station by Omnistar, which limited experienced errors to the range of *1m* and the angle of *0.1 degrees*.

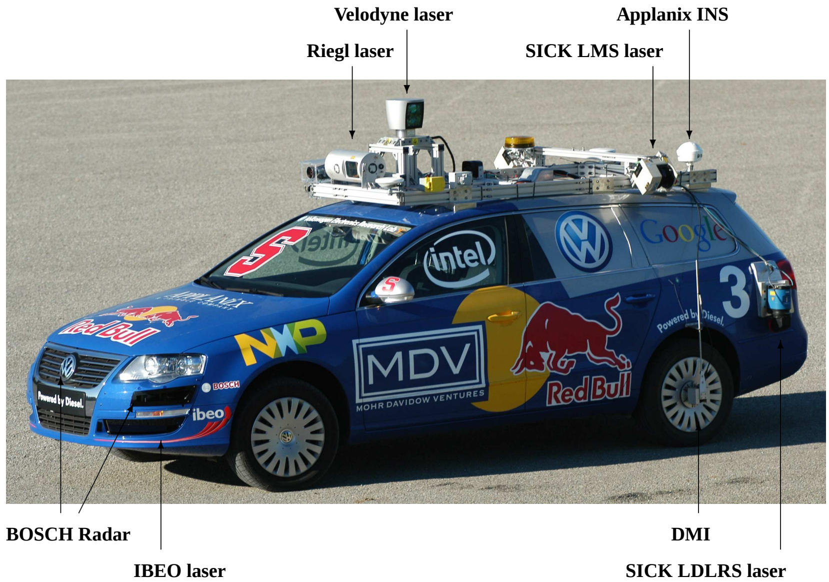

Figure 2.9: Image of *Junior* with marked perception systems [MBB+08]

- To detect line markings and near objects, two lasers are installed on the sides and one up front. They can be used to look for curbs, parking spaces, or static 3D structures.
- The most accurate sensor for object detection up to levels of *100m* (compare [HB10]) is the Velodyne HDL-64E rotating lidar sensor. It covers a *360 degrees* horizontal and *30 degrees* vertical field of view and generates a three-dimensional point cloud of its range measurements. In *Junior* it is used for the detection static and moving objects, in connection two further lasers in the front bumper and another two sensor facing backwards.
- To provide clues about upcoming obstacles at higher speeds, a system of five long-range radars is used. Facing towards the front, these sensors can provide accurate measurements of relative speeds of reflective obstacles like other vehicles.

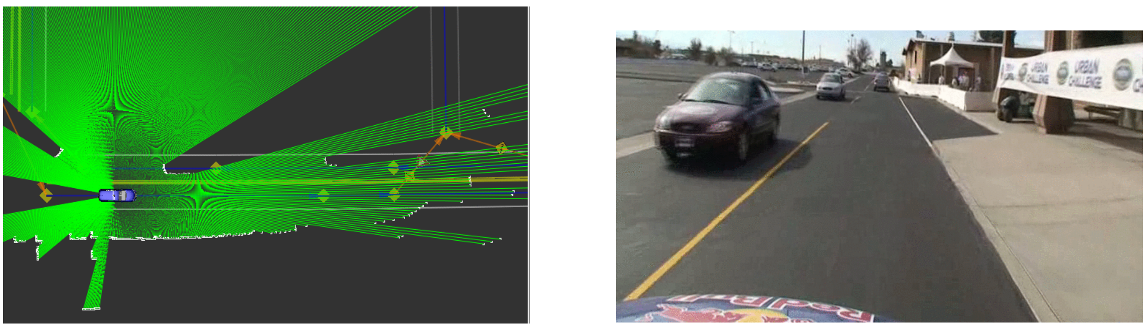

(a) 2D-projection of sensor fusion showing traversable area in green, *Junior* driving to the right

(b) Same scene as recorded by a camera in the windshield, notice the cars on the left

Figure 2.10: Environment perception by *Junior* (a) and the scene viewed through its camera (b)

The output of these sensors is processed by two quad core computers using current server technologies. The information value available for autonomous operation is significant. Even complex situations can be identified by the fusion of this array of different sensors, as can be seen in Figure 2.10. Everything identified in the scene can also be statistically matched with map data to potentially increase positioning accuracy. Highway traffic is classified by automotive engineers and scientists as significantly easier to interpret than urban behavior patterns, due to tighter constraints and a lower number of rules. Therefore, tackling the scenario with similar richness of sensor input to the one present in *Junior* will allow the required reliability and fault tolerance for complete autonomy in future. Especially sensors like the Velodyne will be essential to this process, as indicated by the fact that the approach by Google discussed in Section 1.2 also uses it.

2.5 Traffic Models

The attempt to scientifically model traffic situations was started as a reaction on the immense growth of personal transportation after World War II, when participation in traffic was recognized as one of a human's basic needs. With the involvement of physicists, empirical analysis of unexplained driving phenomena like "phantom traffic jams" or stop-and-go situations has been used to formulate mathematical models for this type of self-driven many-particle system (see [Hel01]). Application of these models and understanding of their parameters today aid in the design of new roads, tempo limit zones, and city environments.

As they aim to simulate status-quo human behavior, their complex implementations are not in the focus of this thesis. However, the well-studied efficiency of simple lane-following models in the development process of drivers assistant systems (especially adaptive cruise control) can be leveraged to derive simple behavior rules for single agents in our scenario.

2.5.1 Intelligent Driver Model (IDM)

The Intelligent Driver Model introduced in [TH02] is a longitudinal micro model that falls into the class of lane-following approaches that base the behavior of an entity solely on the state of the entity in front on the same lane. This single input is the key to its high performance and simple implementation. Furthermore, it has been proven robust and capable of being directly implemented as a control system for an adaptive cruise control function. Despite its simplicity, it shows sufficiently realistic traffic behavior in a highway setting.

The model continuously calculates the acceleration of an entity α (also declared a driver-vehicle system) as a function of its own velocity v_α, the distance to the entity in front s_α and the rate of approach Δv_α.

Model parameter	**Typical value**
Desired speed v_0	120 km/h
Desired headway T	1.4s
Maximum acceleration a	1.2 m/s^2
Comfortable deceleration b	1.5 m/s^2
Minimum distance to entity in front s_0	2 m
Acceleration exponent δ	4

Table 2.1: IDM parameters and practical values[TH02]

The acceleration of an entity is specified as

$$\dot{v}_{\alpha}^{IDM} = a_f(v) + a_{int}(v, s, \Delta v) \tag{2.9}$$

with a general acceleration part for the case of free road ahead a_f and a braking reciprocity a_{int}. These parts are calculated considering a small set of model parameters given in 2.1 with some typical values. The acceleration on a free lane is characterized by a maximum acceleration a and the ratio of the entity's current velocity v to a desired velocity v_0:

$$a_f(v) = a\left[1 - \left(\frac{v}{v_0}\right)^{\delta}\right] \tag{2.10}$$

The acceleration exponent δ can be used to tweak the response and in practice often defaults to a value of *4* [Tre11]. The reciprocity part employs a desired dynamic distance s^* to rate against the current distance to the entity in front s. It can consequently be formulated as

$$a_{int}(v, s, \Delta v) = -a\left[\frac{s^*(v, \Delta v)}{s}\right]^2 \tag{2.11}$$

with

$$s^*(v, \Delta v) = s_0 + vT + \frac{v\Delta v}{2\sqrt{ab}} \tag{2.12}$$

The static part of the desired distance $s_0 + vT$ is characterized by the desired headway T of the entity and includes a minimal distance s_0 that is kept during congestion. Inserting (2.10) and (2.11) in our base acceleration (2.9) yields the IDM's characteristic acceleration equation

$$\dot{v}_{\alpha}^{IDM} = a\left[1 - \left(\frac{v}{v_0}\right)^{\delta} - \left(\frac{s^*(v_\alpha, \Delta v_\alpha)}{s_\alpha}\right)^2\right] \tag{2.13}$$

that fully defines an entity's behavior and satisfies the following characteristics in limit cases:

- **Approaching a static obstacle**: The model achieves a dynamic, continuous, overshot-free adaption of the needed "kinematic" deceleration $b_k = v^2/2s$ towards the desired comfortable deceleration b by adjusting the deceleration to a value above or below b temporarily.

- **Following or approaching other entities**: The model adjusts the entity's speed to the continuously changing safe speed v_{safe}[Kra97] without overshooting. T carries the meaning of a reaction time in this case.

- **Entity in front gains distance**: The model takes the future gap into account, which avoids braking reactions even if the distance is below $s_0 + vT$.

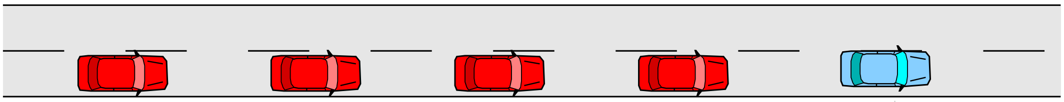

(a) Vehicle convoy, the leader in blue engages full brakes

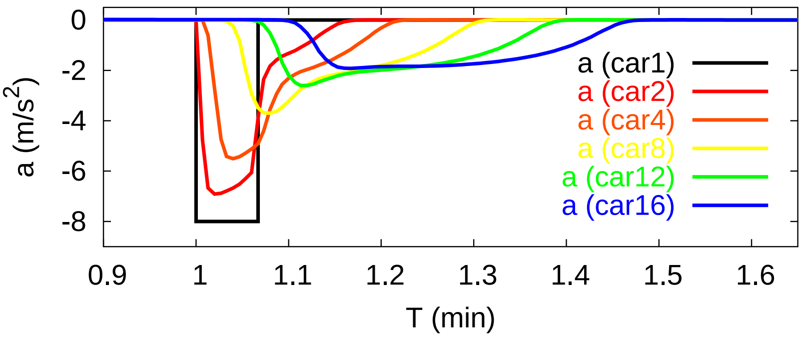

(b) Stable IDM reactions of the following cars to the emergency braking

Figure 2.11: IDM reactions of vehicles behind a rapidly braking car [TH02]

2.5.2 MOBIL Lane Change Model

To achieve a realistic view on common movement patterns on a highway, lane switching is a crucial part of driving behavior to be studied. Since real drivers hardly chose equal speeds and often run into situations that make a different lane more attractive, traffic flow simulation should follow an approach that takes an entity's surrounding lane options into consideration.

To add this functionality to an existing simple lane-following model, the lane switching model MOBIL (Minimizing Overall Braking decelerations Induced by Lane changes) has been suggested in [TH02]. It bases its change decisions on the accelerations of the involved entities that are given by any single lane model like the IDM in Section 2.5.1. MOBIL will use these values to evaluate two realistic criteria for choosing a different lane than the current one: *safety* and *incentive*. The lane change itself is simplified to be instantaneous by definition.

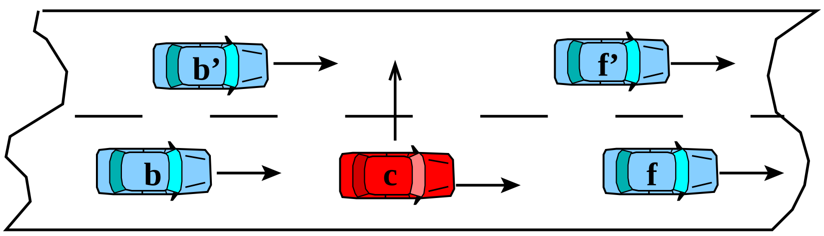

Figure 2.12: MOBIL lane switching scenario [TH02]

Safety Criterion

To determine whether a switch of a car c in Figure 2.12 onto a different lane is safe, the negative acceleration imposed onto the following vehicle on the target lane b' by a fictive change maneuver is calculated using the IDM or any longitudinal model. This value is then compared to a fixed tolerable deceleration b_{safe} :

$$a_{b'c} \geq -b_{safe} \tag{2.14}$$

with the IDM acceleration for a vehicle β in front with distance $s_{\alpha\beta}$ as

$$a_{\alpha\beta} = \dot{v}^{IDM}(v_\alpha, s_{\alpha\beta}, \Delta v) \tag{2.15}$$

and their speed difference

$$\Delta v = v_\alpha - v_\beta \tag{2.16}$$

Since the IDM depends on the rate of approach, MOBIL consequently takes the dependence of a safe gap on the difference in velocity Δv into account. This simulates proper behavior in situations in which humans often fail to assess Δv correctly, causing the often-observed close calls for rear-end collisions on faster lanes. Using the parameter b_{safe} the tolerance of sharp braking and therefore the gap required for an entity to classify a switch as safe can be modified easily, giving the opportunity to introduce more reckless driving into the simulation.

Incentive Criterion

To evaluate the incentive for a lane switch, we can use the underlying lane-following model's acceleration results again to assess the situation before and after a possible lane change. For our entity c the main motivation to look left or right is a possible gain in its acceleration a_c^{IDM}. If a_c^{IDM} is limited by a car in front, chances are that a neighboring lane offers more headway and consequently leads to larger acceleration possibilities. A preeminent feature of the MOBIL model lies in its consideration of back vehicles in the incentive criterion. This means that not only the entity's own acceleration but also the imposed IDM accelerations of its surroundings are maximized. A lane switch is only attractive, if the sum of all accelerations after the maneuver is at least by a switching threshold δ larger than before. A politeness factor p can be used to control the influence of surrounding entities. For Figure 2.12 we can formulate the incentive criterion for a switch onto the left lane as

$$\underbrace{a_{cf'} + p\left(a_{b'c} + a_{bf}\right)}_{\text{after lane change}} > \underbrace{a_{cf} + p\left(a_{bc} + a_{b'f'}\right)}_{\text{before lane change}} + \delta \tag{2.17}$$

3 Scenario Design

The key to developing any type of artificial intelligence lies in a thorough analysis of the environment at hand. In other words and following the deep ties between A.I. and game theory (see [RNC$^+$96]), we need to specify what exactly the game is. Some questions to ask are

- What characterizes the playing field?
- What are the possible valid moves for a player?
- How does a player perceive the playing field?
- In which ways do players communicate?

and possibly the most important of them all

- How does a player determine what is a good move?

This last question will be studied in depth in Chapter 4, but our discussion of the environment will give us the ability to derive meaningful measures that should influence the player's assessment for him to be successful.
It is important to notice that the answers to these questions are not just a matter of analysis of a status quo. The premise of this thesis is not only to develop an artificial intelligence, but also about establishing the environment in a way that benefits its integration into modern society and its performance. In other words, we are not only designing the intelligence of the players, we are also defining the rules of their game. And this means analyzing economic concerns and finding the right arguments *why* the game should be played in the first place. Of course, a fully exhaustive discussion of all forces required for the changes suggested is beyond our scope, but it is our declared goal to achieve a broad view and list the most important effects that should be considered on the way to put the system into practice.

3.1 Traffic Regulations

It should be stressed that some of the beauty of our approach is dedicated to the fact that many features of the current driving environment can remain untouched. Our choice of modern assistant systems as the base of our perception modules allows the vehicles to use current highway systems in many countries. In our model we follow central European legislation patterns for our behavior design. Many current rules are in line with [MS08], a resource recommended to readers interested in the complete set of regulations for the modern day Autobahn. Cases that are not explicitly mentioned in this section are not currently covered by our model. [1]

[1] However, it is the author's strong belief that it has the potential to fully comply with traffic rules like the ones in Germany after minor adjustments.

3.1.1 Highway Definition

For this thesis we will demand the following qualities for a road to be classified as a highway:

- Traffic flow in both directions is separated by a central barrier, allowing us to consider a single direction for our model while maintaining general validity of our results.
- There are no crossings, signal lights, tight corners, or steep hills. This means, we can apply a simplified model of a straight road and gain conclusions of minor inaccuracy.
- Entry and exit to the highway is provided by on-ramps that we will later reference as lane $l = 0$.
- Lanes are of equal width and separated by distinct line markings. This is a requirement for lane-keeping mechanisms that needs to be satisfied for all segments of the road.
- The number of lanes for a segment is always defined. Similar to the simplification for lane switching mentioned in Section 2.5.2, lane endings are also considered instantaneous and are assumed to be rectangular instead of a gradual merge towards the remaining lane.

3.1.2 Behavioral Rules

In addition to the limits imposed on them by this infrastructure, our agents need a well-defined set of rules governing their allowed actions in any given situation. Following the points made in the introduction to this chapter, these rules comply with current legislations for manual driving in many countries. Given reliable situation assessments of agents, this approach enables the model to work in scenarios that still include manually driven vehicles, which is a major advantage in terms of public acceptability of a possible implementation. Again, the German rules mainly part of [Thu99] can be used as a guide for readers unfamiliar with the European standards or seeking detailed insights beyond our scope. We demand the following set of standards to be followed by everyone using our highway:

- Cars are never to go against the official traffic direction, i.e., speed vectors pointing towards the rear of the vehicle are forbidden. If a driver cannot proceed forwards, the only option is stopping.
- There is no general or local speed limit on our highway. Their introduction is possible and requires some changes to the acceleration criteria based on the IDM, but when extrapolating our approach to a fully autonomous highway, further studies might show they are not needed anymore. [2]
- Overtaking of vehicles is only allowed on the left side, unless the speed of the lane to our left is below a certain maximum. Note, that this rule can easily be relaxed for autonomous vehicles that communicate their intentions.
- Generally, vehicles are supposed to choose the rightmost lane possible to give others the opportunity to pass them.

[2] Even those in place for reasons of public noise protection might become questionable with the growing penetration of electric vehicles.

- Human drivers have to keep a safe distance to the car in front depending on their speed. Again, we can make exceptions for connected agents as described in [TR10] to exploit drag advantages and save room on the highway.

- In lane closing situations or at on-ramps, vehicles on the closing lane need to yield to traffic on the neighboring lane, even though a "zipper-like" alternating merge should be encouraged. For our agents, yielding is also a matter of priority, as we will see in Chapter 4.3.

- Drivers always need to let emergency vehicles pass them, if possible. Our model does not include driving between lanes, as is often mandated on current highways. Instead, our notion of priority and its influence on agents' decisions is included as a potential aid for emergency response.

3.2 Environment Perception

Every agent's possibilities are enabled and limited by the information available about the state of its environment. As discussed in Section 2.4.1, in our scenario this data is gathered by sensor technologies based on advanced driver assistant systems. In this section we will list the different properties measured.

- **Speed**: A car's own wheel speed v is measured by tick sensor inside the vehicle and internally calculated in *m/s*. Throughout this document, you will find many speeds printed in *km/h* for better comprehension.

- **Position**: A vehicle's own position is maintained as a one-dimensional longitudinal value x in meters with reference to the highway's beginning at $x = 0$. Practically, this value will be derived from GPS data and possible fusion with visual observations or communicating infrastructure (Car2X).

- **Distance to lane ending**: Gathered from GPS or vision systems, the position of the current lane's ending x_{cl} is needed to make lane changes when running out of road.

- **Relative position and speed of vehicle in front**: Tracked by radar or camera, the other vehicle's position x_{other} and speed v_{other} can be determined from relative measurements.

- **Relative position and speed of vehicle on both sides**: Important for gap detection, these values can be read from relative measurements by radar sensors on the side of a vehicle currently in use for blind spot detection.

This information horizon is very similar to the one presented to drivers of today's upper-class cars. How individual behavior of an agent can be designed using these inputs will be discussed in Chapter 4.1.

3.2.1 Shared Data

With the introduction of car2car communication using wi-fi technology, information is shared with all connected peers. This results in a significant gain of accuracy for some values and an increased perception range. For this thesis, the wireless connection is assumed to be reliable within its maximum range r_{max} without further remarks about the technological difficulties of its implementation. The information benefits for a group of agents include:

- More accurate relative positioning, achieved using approaches like [DLCH07] or the fusion of wireless signal strength and GPS values
- Perfect relative and absolute speed information about every member of the group
- Access to acceleration values of others that are hard to measure otherwise

We will see how these possibilities are used to determine appropriate group behavior in Chapter 4.3.

3.3 Business Model for Autonomous Traffic

While the purely scientific part of our system design shows a lot of potential to increase highway throughput and decrease emissions and accident numbers, successful realization of such a bold change in traffic regulations requires a strong business plan to market the idea to its stakeholders. In this section, we will summarize ideas and arguments from an economical point of view to identify desired behavioral aspects for our agents and possibilities of further research directions that benefit the concept. Where applicable, we compare our suggestions to current products, services, and trends in the fields of passenger and freight transport.

Before we can turn our model into a service that creates value for a customer, it is important to define the organization of the business offering it. Let us assume, there is some legal personality that holds the intellectual property (IP) that characterizes our model. Casually speaking, it is unlikely for this entity of society to be the only stakeholder in the realization of the project, because it would mean that the legislative entities of a state have not only invented the model [3] but also decided to become the monopolist for cars and sensor technology within their borders. To be more exact, we can identify the following five parties involved in some way:

- **The holder of intellectual property** and possibly parts of the back-end infrastructure that will be discussed shortly
- **The government** as the entity entitled to change legislation for the current highway infrastructure
- **The automotive industry** with the capabilities, experience, and current market share to provide the vehicles themselves
- **The suppliers or OEMs** delivering the capacities for the sensor technology needed

[3] or stolen the idea

- **The customers** that use the service and can possibly be charged monetarily for doing so

The party of customers will be examined in the following Section 3.3.1. Also, note that for this thesis we assume the highways to be state-owned rather than private. Otherwise, the party owning the roads would need to be considered here.

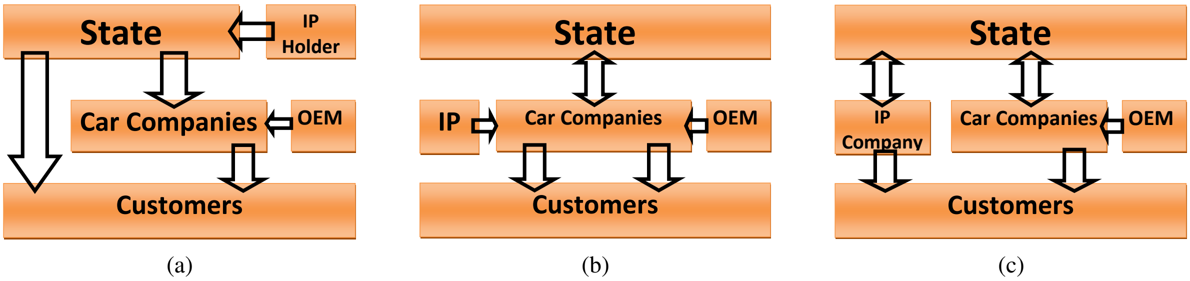

Figure 3.1: Scenarios for different control over the intellectual property

There are cases in which entities listed may merge and act as one personality with combined responsibilities and property. Before we proceed with our argumentation on how an entity in possession of the concept could behave economically, let us shortly discuss three possible scenarios for the formation of our stakeholders shown in Figure 3.1:

(a) This formation is characterized by government control over traffic. The state buys or contracts the IP holders and provides the service. To participate, the customer needs to buy a car from the automotive companies that follows the imposed standards for operation within the state-owned infrastructure.

(b) The car makers get hold of the IP and market the service under government regulations and standards. Customers not only buy the car from the automotive company but also the service.

(c) In this scenario, the IP holders establish their own company that markets the service after negotiating market regulations with the state. The government enforces sensor standards for the car producers to be used with the system. Customers can access the services of the IP company after acquiring a car that is compliant.

These three scenarios mainly differ by the party that interfaces with the end customer and we have ordered them by the growing amount of market competition involved. Therefore, their adaptability primarily depends on the economical policies of the state in question. For our further analysis, let us assume the position of the IP holding company in scenario (c). Most recommendations also hold for the others, with some requiring adjustments outside of our scope.

3.3.1 Customers

Our customers mainly fall into two distinct groups: *end consumers* and other *businesses*. In the following, we will list some of the interesting groups and possible peculiarities of their demands.

- **Individual end consumer**: This very diverse group comprises anyone with a license buying the car for personal use like commuting, shopping, or travel. Demands may vary, but a comfortable ride might be a dominant request.

- **Logistic companies**: These businesses use trucks to move goods around the country and have a special interest in reliability, travel time optimization, and localization of vehicles.

- **Travel agencies**: They consist mainly of bus companies and require similar focus as logistic firms.

- **Public service vehicles**: This category includes emergency vehicles like ambulances, police, organ transports, or road construction equipment. Even though one might argue that these will continue to be driven manually, they are part of the system and their time-critical deployment needs to be considered.

- **Fleet businesses**: This group consists of rental firms, business vehicles, or car sharing portals. Reliability, localization, and trip characteristics should be among the most demanded service characteristics.

3.3.2 Value

We have mentioned earlier that the product we are offering is not a physical good, but a *service*. [QBP87] has made the following definition:

> *Services are all economic activities whose output is not a physical product or construction, is generally consumed at the time it is produced, and provides added value in forms that are essentially intangible concerns of its first purchaser.*

The following sections will briefly introduce possible customer value that can be generated by using the system.

Safety

As discussed in Chapter 1.1 and the beginning of this part, our system is able to achieve much better accident figures than achieved on today's highways. Other studies like [TR10] make defensive estimates of a fatality reduction by *10%*. The potential gain regarding non-lethal collisions is even higher, considering studies like [BLLM05] that lists inattention, close following, and dangerous lane changes as the dominant reasons for rear-ended and many other types of collisions observed. A rationally designed agent is most likely to be involved in crashes that include technical failure or errors of other participants that account for less than *2%* of the accidents in this study.

Increased safety benefits all participants equally, may they transport goods or passengers. Even highway users not connecting to our system are likely to experience smoother travel because of the reduction of accidents as a major reason for lane closings and consequently congestion.

Convenience

As advertised by [TR10] and other autonomous approaches, the main benefit for passengers is the increased freedom due to the lack of required attention. People can use the time traveling on highways in similar ways as in today's trains or planes, with the added convenience to be inside a private room [4]. Passengers can have communication involving attention to body language and mimics or satisfy other human needs such as eating without worries. Being connected to the internet while having their hands free to use todays most powerful input methods, they can engage in entertainment, social networks, or information retrieval.

Priority

The system enables customers to gain privileges for overtaking to increase the chance of faster progression and the reliability of arrival times. This allows for better planning of travel or shipments and increases the time efficiency of participating customers.

Localization

As the model can integrate methods like [DLCH07] that are known to increase relative positioning and involves sensor technology to fuse this to better absolute position data, we can establish a server back-end that allows tracking of vehicles, if they desire to be tracked. Alternatively, customers using communication to their own vehicles anyway, could gain access to the better information and decide about its proliferation themselves.

3.3.3 Competition

In general, our competitors in the acquisition of customers fall into two categories: *other transportation means* off the highway or *manual driving* on the highway. In this section we will cover some key players and compare them to our approach.

Individual Manual Driving

Manual driving means having to pay attention to the road at all times and a higher risk of experiencing situations in which attention was not sufficient or human reaction times prevented the execution of necessary actions. This increases stress levels imposed by long travel and severely limits the driver's ability to participate in any other activity, may it be productive or relaxing.

On the other hand, growing penetration rates of advanced driver assistant systems are likely to cause a decrease in accident rates even under manual operation. Better infotainment systems with advanced input methods would allow for some productivity on the road without shifting the focus of attention

[4] Since visual attention by the driver is not required, legislation allowing windows tinted in a way that completely blocks view into the vehicle seems achievable.

too much. Drivers with a personal interest in control over vehicle movements often associate a positive experience and fun with driving.

Manual Car Sharing and Professional Drivers

As a passenger driven by a human, freedom for infotainment increases, but only in limits that preserve the driver’s ability to pay attention. Vehicle sharing takes away the privacy aspect of going by car and therefore significantly limits some productive activities such as personal phone calls or work on confidential material. For some vehicles, the passenger compartment can be audio-visually decoupled from the driver’s cabin. However, such solutions require luxury cars and professional drivers rather than random strangers.

Rail and Air

Long-haul public transportation methods usually have the advantage of more space. Passengers can walk around and have a restroom available. Safety for both methods is at a level that may be hard to reach for cars, even on an autonomous highway.
This comes with the disadvantage of high fixed costs due to humans maintaining and monitoring the system at all times. Further issues are a lack of flexibility regarding travel time and end points as well as privacy issues in the more affordable passenger classes.

3.3.4 Products and Pricing

To maximize the user base and create substantial revenue, we need to consider ways to wrap our possible value generation for customers into concrete products. Without becoming too specific for the scope this study, we will discuss some considerations that help to build the system in an efficient way by means of demand management.

Priority Booking

The system allows customers to pay for higher priority in the decisions regarding lane changes. This leads to the described potentials for better reliability and travel time. We will assume, that booking lasts for a single trip at least. Further long-term options like monthly subscriptions are possible. Priority comes in a defined number of different levels that achieve growingly larger influence on the group decisions of our agents. When applying advanced models of overall traffic developments, priority suggestions based on a target and desired arrival time can be made possible by forecasting traffic densities. The pricing, advertising, and online deployment of the service could mimic today’s perceived standards for "one-click" hosting companies on the internet as introduced in [AMD09].

Logistical Services

Business customers should be addressed directly by offering individual negotiations over long-term contracts for logistic services. Pricing should take into account what the value addition, special infrastructure needs, and impact on overall road capacity are estimated to be regarding the customer business.

Free Use

To make priority meaningful, it needs to be well-distributed throughout the system. If everyone had high priority, the comparative nature of this measure would be broken and no participant would experience any significant gain in the end. Therefore, customers need to be presented with an incentive to choose lower priority. Following the hosting concept, our system should be free of charge for the lowest stage of decision influence. This way, the amount of low priority agents is increased, making privileges more valuable for everyone else. To prevent too many users from upgrading to higher priorities, we can include an incentive that will allow for temporary higher priority without any costs, if a user remains on the free level and gives way to higher priority agents often. This can be compared to the content generation mentioned in [AMD09]. In the hosting example, content generation means increasing the value of the system for premium users. Users who share data with others are given monetary bonuses or temporarily better priority. In our model, value for the paying customers is their difference in priority from others and actions preserving this balance need to be rewarded.

4 Agent Behavior Design

4.1 Egoistic Driving Agents

To be able to observe a possible gain of group strategies over common individual behavior models, we need to define an egoistic autonomous driving agent for our scenario. The agent should conform to traffic rules and be characterized by its effort to be driving at its desired speed. An easy way to robustly achieve this behavior is the interpretation of the IDM introduced in Chapter 2.5.1 in terms of agent utility. Since the longitudinal control of an egoistic driver is fixed to following an obstacle or driving the desired speed [1], we can leave the acceleration control inside a lane to the continuous IDM. Therefore, what elevates this program from a reflex-based throttle control to a utility-based agent is the agent's freedom to switch to the neighboring lane on its left or right side - if either is available.

4.1.1 Longitudinal Control

Following the remarks above, the agent's movement within a lane is fully determined by the results of the IDM calculated as

$$\dot{v}_{\alpha}^{IDM} = a\left[1-\left(\frac{v}{v_0}\right)^{\delta}-\left(\frac{s^*(v_\alpha,\Delta v_\alpha)}{s_\alpha}\right)^2\right] \tag{4.1}$$

The resulting acceleration yields the speed for the new time step Δt

$$v_{\alpha_{t+1}} = v_{\alpha_t} + \dot{v}_{\alpha}^{IDM}\Delta t \tag{4.2}$$

and consequently the agent's new position x_{t+1}

$$x_{\alpha_{t+1}} = x_{\alpha_t} + v_{\alpha_{t+1}}\Delta t \tag{4.3}$$

Therefore, the egoistic agent shows the desired following behavior on a single lane. Implementing the reflex function can be achieved easily by evaluating the given formulae using a discrete time step Δt. The specifics of the simulation implementation for this work will be explained in Chapter 5.

Longitudinal-Transversal Coupling

An important diversion from the regular IDM characteristics is our consideration of the rules for overtaking specified in Chapter 3.1. These require what [TH02] calls a longitudinal-transversal coupling between the different lanes of our highway. It basically means that we need to specify our choice for the vehicle in front depending on its current speed. If the next vehicle β down the road is not on an

[1] Going slower than the desired speed in absence of an obstacle as a long-term strategy to minimize the difference of an entity's speed progression over time to constant desired speed can be rationally discarded.

agent's lane l_α, but left of it, it is only allowed to pass it, if that vehicle's speed is below a certain threshold ρ. Otherwise, it needs to be treated as if it was an obstacle on its own lane:

$$\beta \in \mathbb{H}: \begin{cases} (x_\beta > x_\alpha) \wedge (\neg\exists\varepsilon \in \mathbb{H} : x_\alpha < x_\varepsilon < x_\beta) \wedge (l_\beta = l_\alpha) & \text{if } v_\beta < \rho \\ (x_\beta > x_\alpha) \wedge (\neg\exists\varepsilon \in \mathbb{H} : x_\alpha < x_\varepsilon < x_\beta) \wedge (l_\beta \geq l_\alpha) & \text{if } v_\beta \geq \rho \end{cases} \tag{4.4}$$

4.1.2 Lane Switching Decisions

The egoistic agent's freedom of choice regarding lanes is pre-limited by two conditions that have to be satisfied to allow a utility-based decision:

- The target lane has to exist. If an agent is on the left- or rightmost lane of the highway, further movement to the outside must be prevented. Formally, this means $l_{right} \leq l_{target} \leq l_{left}$ with lanes numbered from right to left.
- The desired position on the target lane has to be free: $\neg\exists\beta \in l_{target} : x_{\alpha_{t+1}} = x_{\beta_{t+1}}$
- Furthermore, the safety criterion as formulated in Section 2.5.2 needs to be satisfied: $a_{\beta\alpha} \geq -b_{safe}$ for $\beta \in l_{target} : x_\beta < x_\alpha$

Since we carry out lane switches instantaneously within our studies for simplicity reasons just like the IDM does, you see the entity positions taken into account are for $t+1$ already. In reality, these preliminary conditions require more complex formulations[2] and the appropriate measurements to determine their status.

Utility Function

The utility calculated by an egoistic agent for a lane change is its acceleration for $t+1$ determined by the underlying IDM. Based on the definitions in Chapter 2.1, the agent will evaluate the status quo and the situation after the switch following the way MOBIL (see Section 2.5.2) works. Since we are shooting for truly egoistic behavior, we choose $p = 0$ for the politeness factor, which reduces the original incentive criterion discussed in Section 2.5.2 following Figure 2.12

$$a_{cf'} + p\left(a_{b'c} + a_{bf}\right) > a_{cf} + p\left(a_{bc} + a_{b'f'}\right) + \delta \tag{4.5}$$

to the more general egoistic condition

$$a_{\alpha\beta} > a_{\alpha\gamma} + \delta \tag{4.6}$$

with $\beta \in l_{target} : x_\beta < x_\alpha$ and $\gamma \in l_\alpha \ : \ (x_\gamma > x_\alpha) \wedge (\neg\exists\varepsilon \in l_\alpha \ : \ x_\alpha < x_\varepsilon < x_\gamma)$ For an agent obeying our European-inspired scenario defined in Chapter 3.1 we need to distinguish incentive to switch left or right because of the legal urge to choose the rightmost lane possible. To model this behavior, the switching threshold δ has been introduced in Section 2.5.2. For our different utility functions for left

[2] Especially the vehicle lengths should be taken into account.

and right switches we will move this bonus incentive to the appropriate side of the criterion. This way, we can formulate the two different criteria

$$R \rightarrow L: \; a_{\alpha\beta} > a_{\alpha\gamma} + \delta \tag{4.7}$$

$$L \rightarrow R: \; a_{\alpha\beta} + \delta > a_{\alpha\gamma} \tag{4.8}$$

that lead us to the following formal utility functions in these two cases:

$$l_{target} > l_\alpha \begin{cases} U(change) = a_{\alpha\beta} \\ U(stay) = a_{\alpha\gamma} + \delta \end{cases} \tag{4.9}$$

$$l_{target} < l_\alpha \begin{cases} U(change) = a_{\alpha\beta} + \delta \\ U(stay) = a_{\alpha\gamma} \end{cases} \tag{4.10}$$

The switching threshold δ also helps to prevent lane change oscillation due to otherwise insignificant utility differences. Implementing these utility definitions, our egoistic agent shows the expected behavior, trying to overtake on the left whenever blocked on its own lane and changing back to the right when the gap is big enough.

4.2 Group Formation

To be able to make socially-conscious decisions, the entities in our scenario first need to socialize. Based on the assumptions made in Chapter 3.2.1, communication between the vehicles on our highway will be of limited range. This will also be the central limitation for social behavior. The entities should consequently form flexible groups of limited member count and spatial dimension.
For such a group Γ we can formulate the following definitions:

- Γ is a subset of entities from our highway universe $\mathbb{H}$ and $|\Gamma|$ its member count.
- For an entity (or vehicle) α belonging to group Γ we denote $\alpha \in \Gamma$.
- Group membership shall be exclusive, i.e., $\forall \Gamma, \Phi \in \mathbb{H} \neg\exists\, \alpha \in \Gamma : \alpha \in \Phi$ for $\Phi \neq \Gamma$, which can also be formulated as $\forall \Gamma, \Phi \in \mathbb{H} : \Gamma \cap \Phi = \emptyset$. Alternatively, this can be paraphrased as the relation from entities to groups not being $[1:n]$ or the assignment function not being a surjection. Neither multiple assignment nor fuzzy sets are explored in the course of this thesis, but note that the fuzzy approach deserves further research since it could lead to smoother transitions and an interesting information gain between neighboring groups.
- Assignment is functionally injective, since we allow for entities not being members of any existing group. To easily reference these entities from now on, we introduce the set of unassigned entities $\mathbb{U} = \mathbb{H} \setminus \bigcap_{\Gamma \in \mathbb{H}} \Gamma$. This set may be empty, but in practice a significant amount of entities will not satisfy the communication range constraints of any group.

- Because they are based on communicating parties seeking for a social decision, groups in our scenario are by definition required to consist of at least two entities, i.e., $\forall \Gamma \in \mathbb{H} : |\Gamma| \geq 2$.

- There is a limited member count for all groups, meaning $\forall \Gamma \in \mathbb{H} : |\Gamma| \leq m_{max}$. Considerations for choosing m_{max} will be given in Chapter 4.2.2.

- We demand for the communication channels to be established at all times: $\forall \alpha \in \Gamma \, \exists \, \beta \in \Gamma : \left|x_\alpha - x_\beta\right| \leq r_{max}$ with r_{max} being the global range of the wireless communication between entities.

4.2.1 Speed and Position Requirements

The group definition is the base from which the qualification of new candidate entities for a group can be derived. Simply evaluating the given set of rules in each time step will leave you with well-defined groups, but also with a lot of oscillation. This is because of the fact, that a candidate that is in range at one point in time does not necessarily make a lasting contribution to an existing group. To further optimize our assignment decisions, we will not only look at a candidate's *position* but also its *speed*.

Candidate Position

To fulfill the communication range requirement, our new group member β generally needs to satisfy

$$\exists \alpha \in \Gamma : \left|x_\alpha - x_\beta\right| \leq r_{max} \tag{4.11}$$

to be in communication distance. Considering our environment and also later implementation performance benefits, it is more elegant to check for

$$x_{\Gamma min} - r_{max} \leq x_\beta \leq x_{\Gamma max} + r_{max} \tag{4.12}$$

with the two longitudinal borders of Γ

$$\forall \alpha \in \Gamma : x_{\Gamma min} < x_\alpha \tag{4.13}$$

$$\forall \alpha \in \Gamma : x_{\Gamma max} > x_\alpha \tag{4.14}$$

that can be efficiently found by limiting the search space because of the small expected changes in relative positions during one time step. To help avoid oscillation at this early stage of assignment, we introduce a threshold θ that creates a form of hysteresis. By narrowing the borders towards the center of the group by this amount only at application time, we avoid accepting candidates at the very limit of our communication ability that are likely to break our range criterion soon after they have become a member:

$$x_{\Gamma min} - r_{max} + \theta \leq x_\beta \leq x_{\Gamma max} + r_{max} - \theta \tag{4.15}$$

One more aspect we should consider when taking candidate positions into account is whether they are driving on an on-ramp ($l_\alpha = 0$) to the highway and should be let on. Usually, this will not be a significant problem because they will be adopted by any passing group because of their position and their likely speed difference (see following section). However, if the passing group has reached its member count limit, it is worth to be noted that a logic can be found to benefit liveliness by pushing another entity out of the group to help the newcomer onto the highway.
One could even generalize this thought further by allowing for a position in the group to be freed for any agent that is in a helpless situation, e.g. in front of a lane closure. The definition of this kind of "unhappiness" will be discussed in Chapter 4.3.1 and its value can be used at the stage of group formation to solve these cases.

Candidate Speed

Thinking about grouping of entities, a commonly demanded property is their similarity. In our scenario one would naively look for cars that travel at about the same speed. This would yield the most stable groups since their relative movement would be minimized. The problem is: these cars usually do not meet. And even if they did, they would usually not necessarily benefit the group's contribution to overall traffic shaping. Instead, to be able to actually make a difference by social behavior, a group needs its inner conflicts to solve.
Therefore, it is worth evaluating a candidate's potential to force new decisions inside the group because of speed dissimilarity. In connection with the new entity's position, we can state the following rules of candidate disqualification:

- If the new entity β is behind the group, we should not consider it suitable, if it is slower than the average speed of group Γ: $\forall\beta \in \mathbb{H} : (x_\beta < x_{\Gamma min}) \wedge (v_\beta < \bar{v}_\Gamma) \Rightarrow \beta \notin \Gamma$
- We are also not interested in entities already in front and gaining distance to our group because of a speed $v_\beta > \bar{v}_\Gamma$. Therefore, $\forall\beta \in \mathbb{H} : (x_\beta > x_{\Gamma max}) \wedge (v_\beta > \bar{v}_\Gamma) \Rightarrow \beta \notin \Gamma$

The negations of both statements are highly desirable, so under the assumptions that $|\Gamma| < m_{max}$ and $x_{\Gamma min} - r_{max} + \theta \leq x_\beta \leq x_{\Gamma max} + r_{max} - \theta$ we can say that promising new members for group Γ are gained by

$$\forall\alpha \in \mathbb{H} : [(x_\alpha < x_{\Gamma min}) \wedge (v_\alpha \geq \bar{v}_\Gamma)] \vee \left[(x_\beta > x_{\Gamma max}) \wedge (v_\beta \leq \bar{v}_\Gamma)\right] \vee (x_{\Gamma min} \leq x_\beta \leq x_{\Gamma max}) \Rightarrow \alpha \in \Gamma \tag{4.16}$$

For a more practical implementation, we can design a speed fit function $fit_v : \beta, \Gamma \rightarrow [0,1]$ that projects compliance with these rules onto a continuous scale between 0 and 1.

4.2.2 Member Count Considerations

One could ask at this point why there should be a member count limit at all for our groups. Why not have the whole highway connected and optimized? Generally speaking, this is not undesirable. However, it leads to a variety of practical problems in the implementation:

- The problem of evaluating utility and optimizing it on a group scale grows significantly with the member count.
- Larger groups mean more complex decisions: What may be an easily calculated plan of actions in a group of 5 entities might lead to almost undecidable complexities in larger groups.
- The bandwidth and time constraints on the communication channel to exchange vehicle information are harder to fulfill with each new member. Especially growing round-trip times (RTT) lead to significantly slower decisions and response time to actions assigned to entities of the group.
- Priority creates less impact when compared to a larger group. The urgency a high-priority entity causes for a small group to solve matters in its interest is considerably higher than in a large group where it might be weighted against other high priorities down the road or time-consuming global plans to optimize its environment are waiting to be carried out by slowly informed vehicles in the way. Small groups allow high-priority vehicles to cut through traffic in a "divide and conquer" manner.

Some empirical data on group assignment will be presented in Chapter 6.

4.2.3 Dismissal and Split

A group must continuously check its inner integrity, i.e., it has to make sure to satisfy the rules set at the beginning of Chapter 4.2 at all times. Since the member count limit is checked during the assignment protocol, the important criterion to be monitored is to assure a working communication channel between all entities:

$$\forall \alpha \in \Gamma \, \exists \, \beta \in \Gamma : \left| x_{\alpha} - x_{\beta} \right| \leq r_{max} \tag{4.17}$$

Because of the constant relative movement among group members, this rule will be broken regularly and we need to take action to preserve the group's integrity. In this case, our protocol evaluates the following options:

- If the endangered communication link involves either the leader or the last entity of the group, we should simply dismiss them at $t+1$, i.e.,

$$[(x_{\alpha_t} = x_{\Gamma min_t}) \vee (x_{\alpha_t} = x_{\Gamma max_t})] \wedge \left(\neg \exists \, \beta_t \in \Gamma : \left| x_{\alpha_t} - x_{\beta_t} \right| \leq r_{max}\right) \Rightarrow \alpha_{t+1} \notin \Gamma \tag{4.18}$$

- If the endangered link is somewhere else among our members, we will split the group in two at the broken connection. Our new group Φ receives all members behind the lost link location x_{lost}:

$$\forall \alpha_t \in \Gamma : (x_{\alpha_t} < x_{lost}) \Rightarrow (\alpha_{t+1} \notin \Gamma \wedge \alpha_{t+1} \in \Phi) \tag{4.19}$$

- If any of the above leads to Γ or Φ being left with less than two members, that group needs to be destroyed completely, leaving all members in $\mathbb{U}$:

$$\forall \Lambda \in \mathbb{H} : (|\Lambda| < 2) \Rightarrow (\forall \alpha_t \in \Lambda : \alpha_{t+1} \notin \Lambda) \tag{4.20}$$

As with assignment, there is a crucial case where an agent's happiness with its situation should be taken into account when thinking about group integrity: the unhappy leader. If an entity is at the front of the group and aspires to go faster, it is usually disconnecting from the group by gaining distance to the car behind. Problems arise, if the leader is blocked by vehicles outside the group in this process. Since the agent is assigned to the group, it cannot join the group ahead, unless it is allowed to drop out of its current one first. Looking at its utility defined in Section 4.3.1, we can make this decision and allow faster vehicles to pass through groups without becoming stuck.

4.3 Group Behavior

Once a group has been formed, utility becomes a collective subject and the individual agents follow the decision that leads the the best overall outcome. That also means that we gain a degree of freedom for our decisions. Since our agents are now not only pushing the entity in front, they can also be ordered to drive slower than they could to make room for others. Generally, the possibility for a slower vehicle to be ordered out of a quicker agent's way is one of the key features that can gain a social artificial intelligence an advantage over our simpler egoistic agents. The second important factor is the ability to use the increased amount of information available to the group to plan several movements into the future, which will be addressed in Section 4.3.2.

The most important goals for the design of our group behavior populate the following list:

- It should allow the group to solve the problem of different desired speeds by being able to *switch positions* of faster and slower vehicles accordingly.
- It needs to follow the *rules of traffic* for now without harming non-members, even though we will suggest possibilities for special rules that are safe within groups due to the information gained.
- When making a decision for the group, there should be a *bias towards the status quo* to counter possible oscillations and pointless lane switches.
- The group needs to comply with individual *priorities* assigned to its members and favor decisions that create a clear advantage for members of higher priority.
- It should *dictate* each agent's actions without giving him a different choice than to follow group orders, which creates a reliable base for planning more than one step ahead. Allowance to divert to egoism as explored in [GG99] might be an interesting addition, but is not in the focus of our study.
- It should respect an agent's desires and *"leave no agent behind"*, i.e., exhibit a certain *fairness* preventing any agent from being permanently insignificant.

4.3.1 Individual Member Utility

To reach these goals, we first need to redefine utility for a single group member to a more complex form. The desired behavior demands a decompositions into a set of attributes that contribute to an agent's utility as introduced in Chapter 2.2.2. We will then design a way to determine this multiattribute utility for a single agent and discuss how to make it comparable to that of other group entities.

Attributes of Utility

The different dimensions of an agent's happiness on the road need to be well-chosen and their value continuously determinable. For better transparency and easier comparability we will design each attribute as a function $a\colon (\alpha,t) \to [0,1]$. Possibilities for attributes are numerous and this thesis can only introduce a meaningful base set of attributes to consider. The outlook suggests some more ideas for future expansion in Section 7.2.2.

Progression The acceleration result of the IDM can still give a good impression of an agent's situation on the highway. As we have shown for our egoistic entities in Chapter 4.1, it is well-suited for a comparison between two locations on the road at the same moment in time to make a decision for an instantaneous lane switch. However, it hardly serves as a normalized time-independent indicator on how happy an agent should be with his progression at any given moment. Picture the case of an entity's acceleration being $a = 0$, i.e., the vehicle is traveling with constant speed v. The question would be: How should the agent feel about this? Considering just a, the answer is undecidable. If the agent is standing still with $a = 0$, it should probably be less happy about its situation than when cruising along with its desired speed v_0 (see Chapter 2.5.1).
Given this observation, one could argue that the agent's current speed v and its relation to v_0 make for a much better mood indicator in our scenario. However, this measure alone would be very unstable and consequently lead to unnecessary switches and oscillation. The reason why the IDM acceleration works so well for comparisons in the moment of an instantaneous lane change is its ability to make a statement about the future trend. And this adds a crucial value to the situation assessment in our scenario. What good is it, if an agent goes faster after a switch but is rapidly approaching an obstacle that will cause him to brake a moment later?
Therefore, our solution to a valid utility attribute in this case is a healthy combination of acceleration and speed values at the given time. The proposed measure has the form

$$P(\alpha,t) = \begin{cases} 1 - \left(\frac{a - \dot{v}_\alpha^{IDM}(t)}{a}\right)\left(\frac{v_0 - v(t)}{v_0}\right) & \text{if } \dot{v}_\alpha^{IDM}(t) \geq 0 \\ \frac{1}{1+(\dot{v}_\alpha^{IDM}(t))^2}\left(\frac{v(t)}{v_0}\right) & \text{otherwise} \end{cases} \tag{4.21}$$

and consequently weighs acceleration and speed against each other. It receives it's highest utility value $p(\alpha,t) = 1$ if the acceleration *or* the speed are at desired maximum levels. Otherwise the significance of one of the values is degraded by the other if the latter is in a favorable condition. If the agent needs to brake in a situation, this will always decrease its utility.

Lane Endings One of the main concerns for an agent on our highway is the need to switch to a different lane because the current one is ending. This very common scenario happens on on-ramps, at road constriction sites, or because of lack of space or other road development factors. As one of the main reasons for congestion and a crucial point of social interaction while driving, it deserves a lasting presence on an agent's mind. The effect we would like to achieve by including it as a separate attribute is a steady decrease in utility for an entity that approaches the end of its lane, regardless of other cars in front. This allows for early reshaping of groups before bottlenecks and a louder voice of vehicles entering the highway. The required information about the location of the lane closure x_{cl} can practically be provided by advanced navigation systems, car2x transmissions from the infrastructure, or car2car communication from other group members. Once this parameter is known, we can formulate a valid utility attribute for the discomfort caused by lane endings as

$$C(\alpha,t) = \begin{cases} 1 - \left(\frac{s_0}{x_{cl} - x_\alpha(t)}\right) & \text{if } x_{cl} - x_\alpha(t) < \sigma \\ 1 & \text{otherwise} \end{cases} \tag{4.22}$$

with σ as the desired look-ahead distance.

Change Frequency To further avoid system instability due to oscillating lane switch decisions, we will keep track of a timestamp when an entity has changed lanes and integrate this concern into an agent's utility. This has very practical implications as well, because it will contribute to a smoother ride for our passengers, whose car will not follow a rapid sequence of zig-zag schemes through traffic. If the time span between the last two lane changes t_{c0} and t_{c1} for time t has been too short, this will affect an agent's happiness. We have to take the last two switches into account rather than the difference between now and just the last one $t - t_{c0}$, because otherwise an entity will always be unhappy directly after a change, causing it to either never be considered for a switch or to be directly switched back and forth because of its unhappiness. Hence, our attribute will have the shape of

$$F(\alpha,t) = \begin{cases} 1 - \left(\frac{T_{min} - (t_{c0}(t) - t_{c1}(t))}{T_{min}}\right) & \text{if } (t_{c0} - t_{c1}) \leq T_{min} \\ 1 & \text{otherwise} \end{cases} \tag{4.23}$$

and is easily adjusted by specifying the desired minimal break T_{min} after a lane change has been done.

Calculation

Formulating the final multiattribute utility function for an agent α is now merely the question of defining a suitable summation of the attribute value functions. This is possible, because all our attributes exhibit *mutual preferential independence* (as defined in Chapter 2.2.2. Therefore, the utility function of our agent α will be of the form

$$U(\alpha,t) = \sum_{i=1}^{N} w_i A_i(\alpha,t) \tag{4.24}$$

with N attribute value functions $A_i(\alpha,t)$ that are each weighed with an arbitrary factor w_i. With our defined base set of value functions, we consequently get

$$U(\alpha,t) = w_p P(\alpha,t) + w_c C(\alpha,t) + w_f F(\alpha,t) \tag{4.25}$$

as the utility function for our experiments.

If we assume the same utility function for every agent in our system. the weights w_i can be found using single or multi-dimensional optimization algorithms, e.g., to maximize average happiness or traffic flow. Possible dimensions for minimization include the standing or acceleration time of all vehicles to ensure liveliness. We will present a rough parameterization using exhaustive search for the minimum of an aggregated objective function in Chapter 6.1 and summarize further optimization possibilities and suggest improvements beyond our scope in Section 7.2.3 of our conclusions.

4.3.2 Socially Conscious Lane Switching

Now that we have seen how individual utility can be defined, we need to devise strategies to use each agent's input and come up with an appropriate action plan for the group. Therefore, in this chapter, we will discuss the different steps of our preference aggregation algorithm and their influence on the social behavior of our agents towards our goals set at the begin of Chapter 4.3.

Elementary Maneuvers

The first step in the evaluation of preference is the identification of possible actions for each agent. In comparison to our egoistic approach in Chapter 4.1, our agents gain movement possibilities that have previously been omitted because of lacking communication or being contrary to the egoistic agenda. An important addition is their ability to signal their wish for another agent to make a certain move. Each agent should first consider the following basic actions if they are available:

- **Switch to left**: If a lane further left is open, the agent can evaluate a change to that lane. Because of the regulations set in Chapter 3.1, this will be the preferred choice of a faster agent that approaches a vehicle in front.
- **Switch to right**: An agent α that is driving on a lane $l_\alpha > 1$ should always evaluate changing back to the rightmost lane, which should be awarded with a bonus in utility just like the one in place for egoistic agents discussed in Section 4.1.2.
- **Hold position**: It would lead to unfavorable behavior if we neglected the status quo in an agent's utility considerations, especially in terms of oscillation problems. We need to stress that an agent does by no means have to switch lanes every time we ask for his opinion. Its default action will always be to keep following the vehicle in front or driving its desired speed.

Implementing these actions still requires checks for free space on an agent's target lane and for its existence. Keep in mind that this task is not made trivial by group communication because the obstacle might not be part of our group.

Future Gain

To make our agents' evaluation of utility[3] more meaningful, we should make them look into the future consequences of their actions. Since lane changes are considered instantaneous in our model, we can argue the true value of the switch might not present itself in that very moment. This is also due to the fact that we are now considering more attributes than just IDM acceleration which has the unique property of providing an estimate of future gain. Another reason to consider utility after a certain time Δt is the opportunity to devise plans of consequent actions involving more than one agent. Therefore, for our three elementary maneuvers the agent will chose its action at time t and then drive on for Δt and evaluate the utility at time $t+\Delta t$. This way, we further reduce oscillation and can also add the following maneuvers to each agent's repertoire:

- **Request gap on left side**: The agent will signal its wish to overtake an obstacle in front to a group member blocking him on his left side. This will cause the member in question to lower his speed and make room on the target lane. After Δt has passed, the agent can switch to the left and evaluate its new utility.
- **Request gap on right side**: In certain situations, it makes sense for an agent to request a change to the lane on the right, especially in the case of lane closure. Choosing this action will make a group member on the target lane create a gap and yield to the agent moving in after Δt.
- **Request to pass vehicle in front**: If there is no further lane to the left and an agent approaches another group member, it may ask for the other member to make room by switching onto an existing right lane. This can only be carried out as an elementary action, if there is a gap for this member on the target lane. Otherwise, this powerful maneuver requires planning more chained actions, which we will discuss shortly.

First, we need to take a look at Δt and define its value for these new actions, because it seems obvious that it carries the meaning of the time span needed to generate the desired gap in their context. For two cars α and β driving next to each other we can formulate the time needed for β to generate a constant gap g as

$$\begin{aligned}
\Delta v_{t+1} &= \Delta t a_\alpha - \Delta t b \\
g &= \Delta v_{t+1} \Delta t \\
g &= \Delta t^2 (a_\alpha - b) \\
\Delta t &= \sqrt{\frac{g}{a_\alpha - b}} \qquad (4.26)
\end{aligned}$$

with b as a fixed deceleration applied.

Planned Sequences

As seen in the definition of the last of our elementary maneuvers, chaining of actions into time-based sequences is not only possible but has a good potential for large benefits. Since we can offer our agents

[3] and consequently the decision of the group

a simple physical model of their environment, they can search the space of possible action sequences for the best utility. We should stress that our study does not cover parallel execution of different non-conflicting plans, even though it can be achieved by careful definition. Generally, the possibilities and problems of multi-step planning are only discussed briefly in this thesis due to time constraints and their complexity.

For now, we will define a sequence of n actions $(a_1,...,a_n)$ as an ordered set S. The execution time of S can be expressed as

$$\Delta T_S = \sum_{i=1}^{n} \Delta t_i \tag{4.27}$$

with Δt_i as the time needed to complete action a_i.

To be able to find the best sequence for an agent, we also need to ask ourselves how to define utility for such sequences. Since our individual utility function (Formula 4.25) allows for attributes that have an impact after each action taken, we should not judge a sequence just by its outcome. For instance, the change frequency attribute relies on being evaluated for each step to degrade utility for excessive lane switching. Therefore, we need an additive utility definition for sequences. This way, we make sure to choose a path that maximizes the average utility for each step. On the other hand, maximizing the absolute utility after each action leads to the selection of one branch of the search tree that might not be globally optimal, because the utilities of each action depend on the prior node. We will consequently use an additive utility function of the form

$$U(\alpha, t_{start}, S) = \sum_{i=1}^{n} \tau^i \, U\left(\alpha \, , \left(t_{start} + \textstyle\sum_{j=1}^{i} \Delta t_j\right), \, a_i\right) \tag{4.28}$$

having a degrading factor $\tau < 1$ to express the growing insecurity over time as suggested in [GG99]. Mind that using additive utility has the side effect of making the utility of sequences of different lengths incomparable. Therefore, we can either divide by the number of actions n in the sequence to get an average and then compare, or demand that for a given time t our algorithm will only search for sequences of a fixed length n and choose the one of best utility.

Priority

Now that we have defined how each agent can evaluate any given sequence of maneuvers, we can ask them for their preferred actions in each time step. However, we still have not touched the subject of preference aggregation. This is because it will be based on each agent's *priority* that is a crucial part of our model, especially for its business component. In this section we will discuss, how it will also help finding a preference aggregation system that is suitable for our cause and partly in line with the thoughts gathered in Chapter 2.3.1.

For our priorities we define:

- Every agent has a priority $\psi \in \mathbb{R}^+$.
- Priority is a measure of individual importance for the group decision. For two agents $\alpha, \beta \in \Gamma$ with priorities $\psi_\alpha > \psi_\beta$ we demand that the preference of α has a higher influence on the preference for group Γ than β's.

Using this priority definition, our proposed group decision method after trading off theoretical properties and practical usefulness is for a group Γ to choose the action or action sequence that maximizes

$$U(\Gamma) = \sum_{i=1}^{|\Gamma|} \psi_{\alpha_i} U(\alpha_i) \tag{4.29}$$

This decision criterion can be interpreted in two ways:

1. **The group as a super agent**: The problem of finding a group decision can be seen as the decision problem of an agent that has the utilities of all members as an input to its own multiattribute utility. If you look at what we have defined for this case in Section 4.3.1, you will notice that our final function in Formula 4.25 has the same structure as Formula 4.29. As we have constructed individual utilities to be in the range of $[0,1]$, they become normalized inputs to our "super agent" using them to maximize group utility. Priorities come as a natural fit for the weight factors in this equation and result in a very transparent measure of influence of the agent they belong to. Since this is the highest stage of our decision hierarchy and we only consider relative ranking, we do not require $\sum_{i=1}^{|\Gamma|} \psi_i = 1$.

2. **Modified range voting**: When viewed from the perspective of voting systems, our approach falls into the category of range voting. An agent submits a ratings ballot, giving his evaluation of every state suggested to the group. His score for a candidate action is the utility he associates with its outcome. Since all individual utilities are in $[0,1]$, we fulfill the range voting requirement of a common range. Range voting itself has been shown to elude Arrow's theorem as described in Chapter 2.3.3, because of the theorem was established for ordinal preference rather than cardinal values such as in our utility definitions, a point elaborated in [Hil05]. Priorities in this context can be viewed as multiple votes cast by the same agent. As our agents vote for the outcomes sincerely[4] by evaluating their utility function for an option independent of how other group members rate it, the dimension of strategic voting is completely avoided. It is actually a benefit in our scenario that range voting does not comply with the later-no-harm criterion. This way, the truthful positive evaluation of a higher prioritized agent's second choice might help the preference of a less privileged agent to win the group decision. This helps to address the liveliness considerations made at the beginning of this chapter.

[4] i.e., they do not report their own preferences falsely for tactical reasons

To decrease oscillation, the status quo should be made more attractive by introducing a bias b_{sq} to individual member utility such that

$$U(\alpha,t) = w_p P(\alpha,t) + w_c C(\alpha,t) + w_f F(\alpha,t) + b_{sq}\psi_\alpha \tag{4.30}$$

Influence of this value can be seen in Table 6.1 as part of simulation results.

5 Simulation

To examine system behavior and identify problems and optimization potentials of our model, a simulation has been developed to provide early results of the decision algorithm at work. The complete and efficient implementation of the model with all possible maneuvers and safety criteria is beyond the time and resources available for this project and not in the main focus of our work. However, a simplified proof-of-concept is a good investment for further discussions on the next steps to take. This chapter will provide some insights about the extent of simulated aspects, the inner workings of the implementation, and the technology used to calculate and visualize the behavior of a large amount of virtual autonomous cars.

5.1 Platform

Since the project plan comprises a demonstrator rather than a production-level software solution, our platform choice is free and dictated by personal convenience and development efficiency criteria rather than universality or performance.

5.1.1 Software

The simulation was realized as a *Microsoft .net 4* solution written in *C# 4.0*. This framework follows a strictly object-oriented approach and interprets a pre-compiled intermediate version of the code (bytecode) at run time using a virtual machine. The *.net* libraries provide built-in support for the I/O operations, mathematical functions, suitable data containers, parallel execution, and 2D drawing abilities needed to run and observe our model in real time. Furthermore, the platform is easily extensible to include technologies like different graphics engines or distributed architectures and provides the powerful and user-friendly development environment *Visual Studio*.
Downsides of the framework are its dependence on the *Microsoft Windows* platform and possible performance drawbacks because of the intermediate layer.

5.1.2 Hardware

All development and execution of the simulation has been done on a home computer featuring an *Intel* dual core processor at *3.3 GHz* and *4 gigabytes* of system memory. It is running a *64bit* version of *Microsoft Windows 7* with *Service Pack 1*.

5.2 Implementation

The software has been developed under the object orientation paradigm and consists of a dynamically-linked library (DLL) with the core components and a window application that provides user interaction and presentation. This section will cover some main points of our development effort.

5.2.1 Simulated Aspects and Simplification

As mentioned above, the full extent of our model's flexibility and complexity is beyond our reach for this particular implementation. The following is a collection of the main features and limitations of the solution created for this thesis:

- Complete IDM lane-following capabilities used for egoistic and group-affine agents alike and providing dynamic headway management and forward-looking acceleration results
- Flexible highway definition with multiple segments of different lane numbers and a maximum length of $1.7976931348623157 \cdot 10^{308}$ *meters* (maximum value of `double` type)
- Cars spawned at an arbitrary number of locations with different desired speeds and priorities; a lot more discriminative parameters between agents exist, but have not been used in experiments yet.
- Physical movement of cars in variable time resolution
- Group formation and integrity preservation (dismissals and splits) by a globally constant communication range; complex assignment considering position and speed has been included but not prioritized for deep research, mainly because simple assignment performance is acceptable.
- Individual utility calculation as described in 4.3.1 with variable weight factors accessible for run-time adjustments
- Exhaustive search for maximum future group utility among three possible movements by each agent (`LEFT`, `RIGHT`, `PUSH_FORWARD`), evaluated by successive IDM predictions up to *1s* ahead of decision time with the car in front assumed constant; However, further action sequences have not been implemented due to the complexities and performance issues involved.
- Built-in gathering of statistics by each agent and some global functions that can be read from the cars after a simulation run, although not all statistics functionality and querying for specific criteria using *LINQ*[1] is provided to the user via convenient interfaces.

5.2.2 Data Structures

Some of the most important objects managed within the code are briefly discussed in the following list.

- The `Highway` is a linked list of segments to be efficiently traversable in the direction of vehicle movement. Each `HighwaySegment` contains a start and end location as well as an array of available lanes and their indices.
- Each `Car` generated by a `CarSpawner` object is managed in a simulation-wide `List<Car>` that can be sorted by vehicle location, because `Car` implements `IComparable` in an appropriate way. Further fields of the object include current simulation values like speed or individual utility, life-long constants like priority or desired headway, and statistics like spawn time or lists that hold

[1] a special *.net* query syntax that provides SQL-like functionality equally on collections, databases, or XML files

the history of physical values. Special care has been taken to provide a valid clone method to be able to create virtual states of the simulation when extrapolating group utility for future time steps.

- When `Group` is assigned to a `Car` they mutually reference each other for efficient access from both directions. Group member references are held in a `List<Car>` and member functions take care of the assignment, dismissal, and split functionalities of the model. An object of the abstract `GroupStrategy` type is held as a reference and contains all logic and data needed to move group members and decide about actions.

- The outer hull of the application core is the `Simulation` object that contains the `Highway`, all `Car` and `Group` instances, and the global simulation time. It is realized in a way that offers easy access to global statistics and user-friendly ways to integrate it into the front end application.

5.2.3 Performance Optimization

Where reasonable, we have made use of the `Parallel.ForEach` constructs within *.net 4* that are able to dynamically parallelize loops according to system hardware capabilities and run-time resource progression. This feature has proven highly convenient and extraordinary useful for our cause, because it provides easy utilization of multi core performance. Apart from parallel processing, no special measures have been taken to further optimize computational or spatial costs of the simulation.

5.3 Interface

The front-end application uses standard *Windows Forms* classes to enable input and output of the simulation parameters and results. It provides real-time feedback of all data recalculated in a single step by means of data tables and 2D animation.

5.3.1 Visualization

The current state of all cars is drawn to a rectangular canvas as a two-dimensional bird's eye view with cars simplified to rectangles in their correct relative position. Each lane of a highway segment is drawn as a single line touching the bottom of vehicles driving in it. The user can navigate the full highway by a horizontal slider that automatically adjusts to the current window size. Furthermore, a vehicle following mode is included to allow the focus on a single car and its progression. The whole view can be zoomed in four levels to provide better visibility of detailed movements.
Groups of agents are indicated by the number of their members that is printed over the centroid of member positions. The user can choose to activate one of two view modes to further visualize group interactions:

- Cars can be colored by a randomly assigned specific group color to observe the assignment process. This mode provides clues for reassignment frequencies and spatial spread of group members.

- Alternatively, the current individual utility for agents within a group can be visualized by a color from a range of green (maximum utility) to red (minimum utility). The main purpose of this option is to identify errors in the design of utility functions and the outcome of different parameter values.

The bottom part of the window holds a table of all cars and a choice of their field values presented for each step. Selection of a vehicle marks it with a circle in the visualization and moves the perspective to center it. Activation of the follow mode will then allow the user to stay with this car during movement. Additionally, a property view panel to the right holds general statistics updated by a separate *2000ms* timer. Values about the scene setup, global traffic characteristics, and group assignment are displayed and explained.

5.3.2 Live User Interaction

The user can engage time lapse using a slider that allows simulation to be sped up to a maximum factor of *100x*. This allows for the observation of long and complicated journeys and congestion phenomena. Furthermore, simulations can be started that result in the generation of CSV files for data analysis. These calculations are done by a `BackgroundWorker` thread at maximum speed achievable by the hardware.

6 Empirical Results

6.1 Oscillation and Average Speed

To find a good base for further evaluation, the parameters of our individual utility function as shown in Formula 4.30 need to be properly assigned. Table 6.1 shows parameter values as optimized by a aggregate objective function that takes a speed match and the number of line changes per vehicle for a simulation run. The speed match is defined as the actual achieved average speed divided by the desired speed of the vehicle. Every combination has been run three times to account for the randomness of desired speeds and group assignment. The priorities of all cars are equal and the scenario is chosen to be a two-lane highway with two entrances and a total length of *3.5km*.

w_p	w_c	w_f	b_{sq}	Lane changes/vehicle	Average speed match	Speed match/oscillation
groups disabled				14,895	0,7600946	0,05103
0,6	0,2	0,2	0,4	7,2	0,691995	0,09611
0,4	0,4	0,2	0,4	7,248333	0,684849	0,094484
0,4	0,2	0,4	0,4	8,298333	0,717379	0,086449
0,6	0,2	0,2	0,3	9,268333	0,700554	0,075586
0,2	0,2	0,6	0,4	10,41833	0,685926	0,065838
0,4	0,4	0,2	0,3	11,64167	0,67375	0,057874
0,2	0,6	0,2	0,4	11,66167	0,666748	0,057174
0,4	0,2	0,4	0,3	12,75833	0,698166	0,054722
0,2	0,4	0,4	0,4	13,43167	0,696402	0,051848
0,6	0,2	0,2	0,2	14,38167	0,695892	0,048387
0,2	0,4	0,4	0,3	14,565	0,690575	0,047413
0,2	0,6	0,2	0,3	14,875	0,697088	0,046863
0,2	0,2	0,6	0,3	15,16833	0,679818	0,044818
0,4	0,4	0,2	0,2	18,31833	0,679434	0,03709
0,4	0,2	0,4	0,2	18,49167	0,682903	0,03693
0,2	0,6	0,2	0,2	19,55833	0,685028	0,035025
0,2	0,2	0,6	0,2	20,44167	0,659138	0,032245
0,2	0,4	0,4	0,2	22,52167	0,690009	0,030638

Table 6.1: Parameter optimization for the ratio between speed match and oscillation

We see that the influence of the bias value is striking and allows oscillation to decrease below the egoistic agent's value while achieving solid speed match values. The second dominating factor is the weight w_p for the progression attribute value function, which is expected under the chosen optimization target. We can also see, that our cooperative agents do not benefit the average speed on the highway yet. The egoistic agent is faster, aided by the absence of parallel plans for our groups. Since only one agent can change lanes at a given moment, cooperative agents lose time waiting because another group member was nominated to be switched to a faster lane first.

6.1.1 Utility Separation

Table 6.2 shows the experiment of three groups of different priorities and desired speeds sharing the road. The desired speeds are *36, 30, and 25m/s* for these runs. The faster and more important customers spawn at the first on-ramp and meet slower cars with less priority down the road. We see that each group is far from their desired speed level due to traffic. However, the higher priority vehicles make it past the general average speed while the lower ones experience significant cutbacks.

Run	Average speed	Average v for p_1	average v for p_2	average v for p_3
1	18,44205	21,76521	18,66631	14,8946
2	18,0455	21,49284	18,03656	14,60709
3	18,32981	21,83239	18,26471	14,89236
4	18,20935	21,61813	18,0388	14,97113
5	18,09173	21,50372	18,18299	14,58847
6	17,74567	21,09947	17,64767	14,48986
7	17,90755	20,96934	17,96871	14,78458

Table 6.2: Multiple simulations of a scenario with three distinct priority groups

However, separation by utility is not always this clear and our simplified simulation fails to deliver decisive proof of guaranteed value for the paying customers. This is most likely due to a severe unattractiveness of the `PUSH_FORWARD` maneuver, which shows very limited occurrences. The reason is the single step limit when planning into the future, that does not accurately show the benefit of a group member braking. Realistically, group utility suffers a drop at first and the potential of the opening gap is beyond the sight of our group decision.

6.1.2 Individual Utility Progression

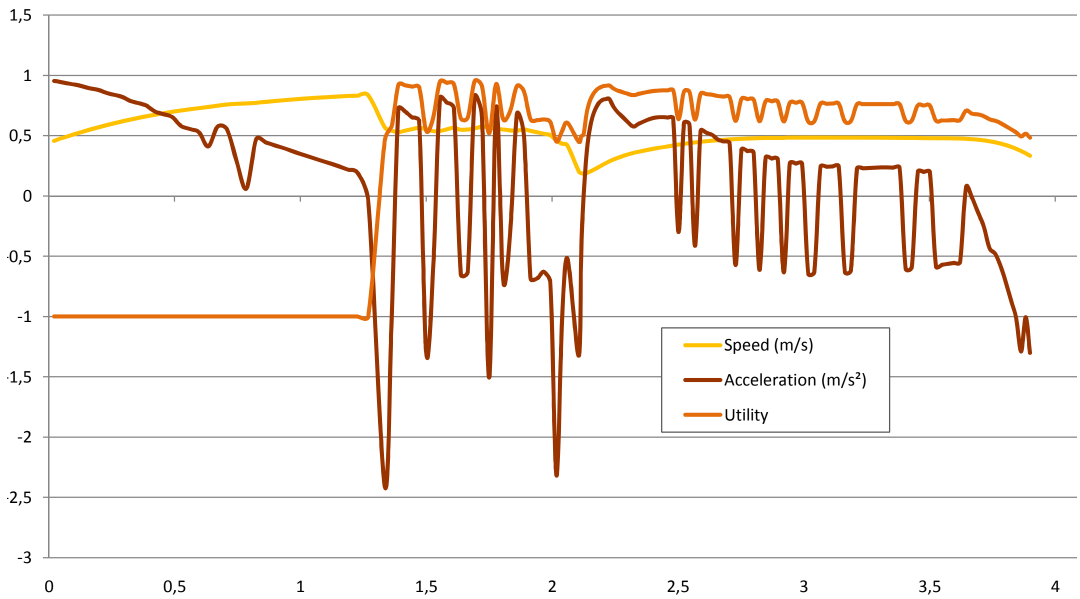

Figure 6.1: Simulation history of a vehicle with acceleration (dark), speed (bright), and individual utility (midtone) all normalized by their maximum positive values over the position axis in *1000m*

Figure 6.1 shows the value progression of a single car driving through a scenario. Speed, acceleration, and individual utility are presented for any given vehicle position along the way. The individual utility is set to -1 while the agent is not assigned to a group. As seen from the graph, utility follows the IDM's acceleration demands, but softens the impact by considering the current speed. Even though the vehicle is decelerating multiple times towards the end, the agent does not rate the situation to be intolerable and progresses with near constant speed and without provoking switch decisions by his group. This rationality has been demonstrated throughout our experiments, which undermines the confidence in the formal potential of the concept. Even though more work to even out the progression can be done, the attribute value function regarding speed and acceleration trade-off is a good fit for our scenario and provides the desired characteristics without any special parameterization.

6.1.3 Group Assignment

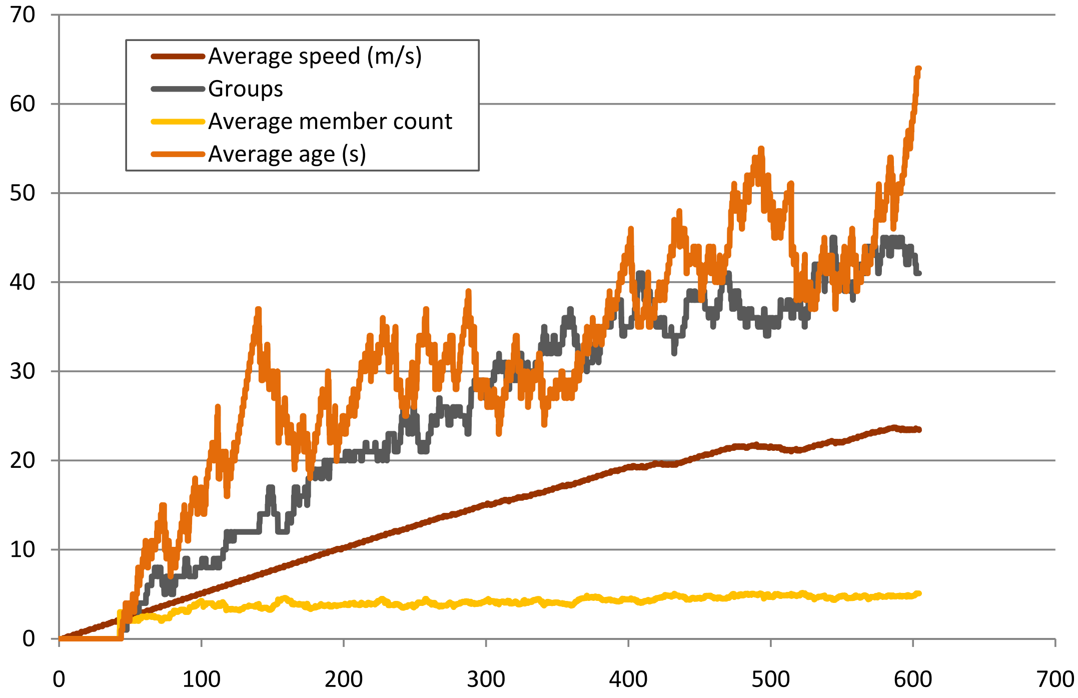

Figure 6.2: General statistics regarding groups for a single simulation run over *700s*

As an example for the global statistics gathered during simulation, Figure 6.2 presents some of the values concerning the overall group assignment process. This single run has been conducted on a three-lane scenario including a lane closure. A remarkable fact generally observed in our experiments is the quick convergence towards a group member count average of around *5*. This value is stable under different maximum assignments chosen and is mainly linked to the communication range. As the highway becomes more and more crowded in this run, we observe a growing number of groups and see that the time limits for their stable existence have not yet been reached. The fact that the average group is not collapsing for minutes shows the room for longer action planning, which opens potentials to achieve better results in other categories in future.

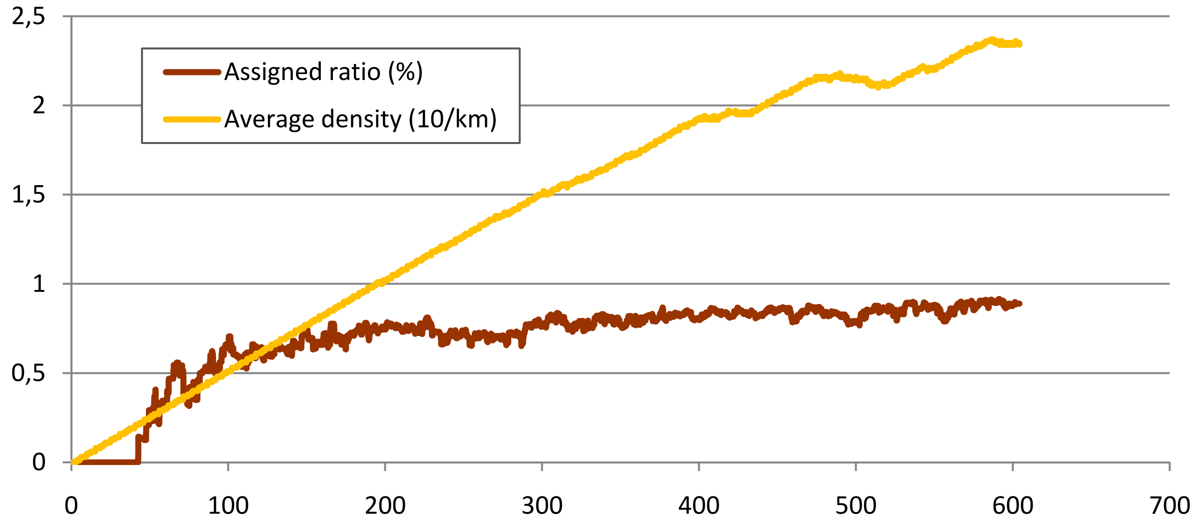

Figure 6.3: Traffic density and group assignment quotas for a single simulation run over *700s*

7 Conclusions and Outlook

In the course of this thesis, we have given an introduction into the possibilities of cooperative artificial intelligence in an autonomous highway scenario. The system itself has been discussed in its wide and interdisciplinary implications and a detailed mathematical model for the behavioral decisions made by a group of agents in this scenario has been formulated. First steps towards the implementation of this model have been unveiled and empirical data gathered by a simulation application with real-time visualization have been presented. In this chapter, we will summarize some lessons learned over the time of this study and interpret some of the main trends within the data.

7.1 Value of Social Consciousness

This study alone cannot deliver the proof of traffic flow benefits by means of social conscious agents. The idea itself is powerful, but needs a lot of effort to be put into a more complex implementation. Only by the further improvement of look-ahead techniques to enable the discovery of the real future revenue of social actions, our model can be expected to show its full potential. It remains to be seen, whether its influence on traffic flow can rival a set of egoistically acting agents at all.

7.1.1 Economic Value

In the end, this might not even be necessary, as other benefits emerge from the convoy-like driving within a group. In the end, what we wanted to show is how group decisions using multiple factors can be formalized and successfully linked to a set of goals in a well-defined environment. The individual utility criteria found have proven to mirror rational assessment of an agents situation and consequently constitute the positive contribution our simplified simulation can provide.

The business perspective of autonomous driving has been pointed out in its significance and many interesting impacts and concepts have been collected for this document. The idea of this significant amount of time freed by handing the wheel over to automation is lucrative and has the potential to spawn a variety of new business models that fill this future gap in our daily plans.

7.2 Future Possibilities

During the conception of our model of future highway traffic, working on this thesis has sparked a lot of ideas for further directions of research within this environment. This chapter covers some of the interdisciplinary facets and detail optimization possibilities opened by our study.

7.2.1 Additional Scientific Angles on the Scenario

We have seen that the introduction of cooperating autonomous agents has far wider implications than their implementation in hard- and software. We have touched matters of economy, ecology, and law, none of which have been exhausted by any means. In this section, some further dimensions of the system as a whole are gathered to inspire further ideas beyond the realms of information technology.

Psychological and Sociological Questions

The concept of being driven by a machine bears some controversy regarding acceptance and effects on the feelings, behavior, and values of the human passengers in short- and long-term usage. Some questions that might be beneficial to answer include:

- Can we trust an autonomous car like we trust trains on obvious rails? If not, how can we be made to?
- How does the impact of autonomous driving differ for various characters and demographics? Does the manager used to a driver adapt quicker than a student used to video games?
- How will people reallocate the attention now unnecessary for traveling? Will they develop needs that have never been present while driving?
- What is the impact of the stress reduction achieved by traveling in a private room without having to pay attention?
- Will the "need for speed" and the fun associated with driving completely degenerate? Will we lose our interest in racing sports?
- How will the status symbolism of private mobility change without human drivers? What will we brag about when driven around by an algorithm?
- Can the lack of attention required affect human skills in other fields? What might be the replacement to keep them in shape?
- Will we live a life on the road because we can? Will "home" become a constantly moving place? What are the implications?

Information Security

A whole thesis could alone be written on the manifold security questions raised by the concept of a massively distributed system including safety-critical technology for everyday use. Apart from the obvious, but manageable need for protection of communication channels, the biggest problems lie in the accessible nature of the vehicles when trying to keep the security of the back-end intact. Furthermore, measures need to be taken to protect the newly won privacy on the road and the priority system as the core business model. Given today's mentality about the violation of traffic laws, the presence of evil-doers on future roads can be taken for granted.

Control Engineering

Since most modern traffic models stem from research done in control engineering, it seems valid to find large-scale models about cooperative agents in this area. Today, research trying to theoretically evaluate multi-agent behavior with common tools of the field is limited to very simple agents and message exchanges. Advancements on this path might enable a more global approach in agent design and parameterization in our scenario to increase stability, efficiency, and behavioral prediction.

Infrastructural Development

It needs to be stressed that congestion control is not a problem that can be solved by changes within the vehicles alone. Whether we succeed in establishing usable connection between future metropoles is primarily a matter of how we can efficiently build roads, develop cities, and offer alternatives to individual transportation where applicable. Autonomous technology and future methods of traffic analysis can only try and provide an optimal usage of the resources they are given. The intelligent allocation of these resources on a larger scale and the appropriate legal regulations are some of the problems that need to be addressed by politicians, architects, and sociologists.

7.2.2 Further Criteria for Decision-Making

Considering the notion of utility as introduced in this thesis, our system is expandable to more complex dimensions that can influence an agent's behavior on the road to create an even better service for its passengers. It also makes sense to give the passengers the ability to change behavioral patterns of their vehicle on the road. This goes in line with the argument, that people's trust will increase for a system they are actively influencing. Improvements cannot only be made to the utility function, but also to group formation, and the underlying acceleration model.

Car Alignment

The position of cars within a group offers a lot more opportunities to achieve a different traveling experience and to increase efficiency. Some possibilities not included in our study can be achieved by having special rules for cars in groups. As discussed earlier, autonomy can make traffic rules obsolete that are in place to protect against faulty human driving. If legislation and its enforcement methods allowed it, we could easily gain significant room on highways by decreasing the headway between cars within a group to a minimum. Depending on the shape of cars, this can enable drag and optimize fuel or power consumption. Furthermore, right-overtaking in groups could be allowed to provide the decision making with more options to fulfill a groups desires. Further improvements possible include the avoidance of driving right next to each other at the same speed for too long, which might be irritating passengers and undermine privacy.

Scenic Routes

When considering lane choices, interesting possibilities result from a tight integration with navigation hardware. If desired, the system allows the agent to choose lanes that offer unobstructed view on landmarks or direct proximity to waterfronts.

Social Interaction

Group formation can be tweaked to integrate with social networks in a way that allows traveling parties distributed to multiple cars to be within the same agent group at all times, if possible. This way, they could benefit from the fast and responsive wireless channel to enable rich media applications like video conferencing or sharing of multimedia databases. Special apps for these gatherings could be developed by automotive brands and could make the opportunity to organize groups by brands more attractive, together with possible technological advancements that allow car makers to market unique solutions by making use of this distributed system. Furthermore, trucks could be organized with proper group assignment to improve logistic efficiency by guaranteeing synchronized arrivals.

7.2.3 Model Optimization Possibilities

Given the model's current state, there are a few methods that can be applied to achieve better results in future and are briefly discussed in the following collection.

- Currently, decisions are achieved by an exhaustive search of possible rearrangements within a group. Therefore, methods that apply smart pruning of the search tree can be used to gain performance in simulation and practice.
- Simulation of the full complexities of our model would also gain better results for single- or multi-objective optimization of the parameters. This should especially be expanded to tune priorities against each other and estimate a good monetary price to demand.
- Research done by [GG99] suggest that a less strict enforcement of group decisions on their members might yield better results. However, such loose enrollment of agents is complex in implementation and makes it even more difficult to make global statements about the system, e.g. by means of control engineering.
- More thought can be invested at the borders of group assignment to achieve better transition between groups. Fuzzy sets have been mentioned in Chpater 4.2 as a possible strategy, but improvements can also be made by establishing more complex logical rules for dismissal and splits.
- Possibilities to introduce further parameters exist at many levels of the model, especially within the attribute value functions. This thesis has avoided the formalization of some trivial mathematical methods to emphasize different values due to readability. Simulation tells us that many more inputs can be studied to create better results and estimations of the true potential of the model.
- After working out a concrete setup for the system's infrastructure, simulation can be expanded to include models for the wireless link, sensor capabilities, threats to security, or specific vehicle types.

7.2.4 Suggestions for Interoperability with Drivers Assistant Systems

Even drivers who are not fond of the idea to lose control over their vehicle to an autonomous agent can experience features of the system that open new possibilities for advanced driver assistant systems. Furthermore, the connected nature of the system allows to create profitable services for human drivers on future highways.

Situation Assessment

Today's main limitation in driver assistant systems is sensor technology. We have already mentioned a few projects that aim to include car2car or car2x communication in modern safety systems to provide better environment perception or solve problems with individual solutions as mentioned in [Bro10]. The shared information within our peer groups can easily exploited by systems that still keep the driver in control and only issue warnings or engage preparations, that save important milliseconds in tight situations.

Driving Suggestions and Incentives

Furthermore, we can educate drivers to behave more rationally by providing information about our agents' intentions and the state of traffic beyond the human eye. Such systems could aid traffic flow and safety by avoiding unnecessary risky maneuvers by providing the information that this action is undesirable for the driver. Participating drivers can also be awarded for driving in ways that benefit the system, similar to what solutions like [RJM98] provide for insurance [1] cost evaluation using driving profiles.

Booked Priority

It seems reasonable to charge good money from human drivers who desire our agents giving way to them. Since this endangers the good health of our system, it should be strictly limited and carefully studied before including it. We also need to make sure, that agents cannot be forced into letting humans pass by means of reckless driving.

The Value of Human Driving

While some may question the very idea of putting effort into opening the system to human drivers to a certain extent, it also creates a lot of opportunities for car brands to market vehicles to be driven manually by factors like sportiness and other emotional selling points. In the end, it will be decided by the customers or politicians, how much demand there still is for the sensation of being in total control of a powerful machine like a car - an often deceiving feeling, that our agents are not aware of.

[1] Another interesting topic to discuss within our scenario

List of Figures

Acknowledgement

Before closing this document, I would like to take the chance to thank some individuals that have guided me over the course of this thesis work and my student life in general.
Firstly, I am grateful for the kind supervision by Prof. Möller and our inspiring discussions about the future of driving. I would also like to thank Prof. Gollmann for his availability as a second corrector and his support throughout my studies at this university. I also need to thank the STS and SVA institutes of TUHH for their welcoming spirit and the latter for the base of this TeX template.
Many colleagues at the Volkswagen Electronics Research Lab in Belmont, CA deserve my sincere gratitude for sharing their uplifting passion for science. Special thanks go to my internship supervisor Pransanth Jeevan for his explanations of current-day autonomous vehicle technology and to Florian Schuller and Keith Hutchings for healthy discussions within the realms of automotive technology and psychological user studies.
Thanks to my parents for the freedoms provided during my studies and practical opportunities. Finally, I need to express my heartfelt appreciation to my little sisters Micky and Kathi for their loving support during the creation of this work.

Bibliography

[ADA11] ADAC: *Staubilanz 2010*. http://www.adac.de/_mm/pdf/statistik_staubilanz_2010_0311_68897.pdf. Version: March 2011

[AMD09] ANTONIADES, Demetris ; MARKATOS, Evangelos P. ; DOVROLIS, Constantine: One-click hosting services: a file-sharing hideout. In: *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*. New York, NY, USA : ACM, 2009 (IMC '09). – ISBN 978–1–60558–771–4, 223–234

[Arr70] ARROW, Kenneth J.: *Social Choice and Individual Values, Second edition (Cowles Foundation Monographs Series)*. 2nd. Yale University Press, 1970 http://cowles.econ.yale.edu/P/cm/m12-2/index.htm. – ISBN 0300013647

[Ber54] BERNOULLI, Daniel: Exposition of a New Theory on the Measurement of Risk. In: *Econometrica* 22 (1954), Nr. 1, 23–36. http://www.jstor.org/stable/1909829

[BLLM05] BALDOCK, M.R.J. ; LONG, A.D. ; LINDSAY, V.L. ; MCLEAN, A.J.: *Rear End Crashes*. Centre for Automotive Safety Research (CASR), 2005 http://books.google.com/books?id=OJbCAAAACAAJ. – ISBN 9781920947170

[Bro10] BROY, M.: *Cyber-Physical Systems: Innovation Durch Softwareintensive Eingebettete Systeme*. Springer, 2010 (Acatech Diskutiert). http://books.google.de/books?id=BGI1F3MXPeOC. – ISBN 9783642144981

[Con] CONNOLLY, Christine: Driver assistance systems aim to halve traffic accidents. In: *Sensor Review* 29, Nr. 1, 13–19. http://dx.doi.org/10.1108/02602280910926715. – DOI 10.1108/02602280910926715. – ISSN 0260–2288

[Deb59] DEBREU, Gerard: Topological Methods in Cardinal Utility Theory / Cowles Foundation for Research in Economics, Yale University. Version: 1959. http://ideas.repec.org/p/cwl/cwldpp/76.html. 1959 (76). – Cowles Foundation Discussion Papers

[DLCH07] DAO, Thanh-Son ; LEUNG, K. Y. K. ; CLARK, C. M. ; HUISSOON, J. P.: Markov-Based Lane Positioning Using Intervehicle Communication. In: *Intelligent Transportation Systems, IEEE Transactions on* 8 (2007), Nr. 4, 641–650. http://dx.doi.org/10.1109/TITS.2007.908574. – DOI 10.1109/TITS.2007.908574

[dpa09] DPA: Bundesbürger stehen 535.000 Jahre im Stau. In: *Welt online* (2009), September. http://www.welt.de/vermischtes/article4523182/Bundesbuerger-stehen-535-000-Jahre-im-Stau.html

[GG99] GLASS, Alyssa ; GROSZ, Barbara: Socially Conscious Decision-Making. In: *In Proceedings of the Fourth International Conference on Autonomous Agents*, ACM Press, 1999, S. 217–224

[HB10] HALTERMAN, Ryan ; BRUCH, Michael: Velodyne HDL-64E LIDAR for Unmanned Surface Vehicle Obstacle Detection / Space and Naval Warfare Systems Center, Pacific. 2010. – Forschungsbericht

[Hel01] HELBING, Dirk: Traffic and related self-driven many-particle systems. In: *Rev. Mod. Phys.* 73 (2001), Dec, 1067–1141. `http://dx.doi.org/10.1103/RevModPhys.73.1067`. – DOI 10.1103/RevModPhys.73.1067

[Hil05] HILLINGER, Claude: The Case for Utilitarian Voting / University of Munich, Department of Economics. Version: Mai 2005. `http://ideas.repec.org/p/lmu/muenec/653.html`. 2005 (653). – Discussion Papers in Economics

[Hit95] HITCHCOCK, Anthony: Layout, Design And Operation Of A Safe Automated Highway System / Institute of Transportation Studies, UC Berkeley. Version: 1995. `http://econpapers.repec.org/RePEc:cdl:itsrrp:102925`. 1995. – Institute of Transportation Studies, Research Reports, Working Papers, Proceedings

[Kra97] KRAUSS, Stefan: Towards a Unified View of Microscopic Traffic Flow Theories. In: *IFAC Transportation Systems*, 1997, S. 941–946

[KTH10] KESTING, Arne ; TREIBER, Martin ; HELBING, Dirk: Enhanced Intelligent Driver Model to Access the Impact of Driving Strategies on Traffic Capacity. In: *PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A* 368 (2010), 4585. `doi:10.1098/rsta.2010.0084`

[Mar10] MARKOFF, John: Google Cars Drive Themselves, in Traffic. In: *New York Times* (2010), October

[MBB+08] MONTEMERLO, Michael ; BECKER, Jan ; BHAT, Suhrid ; DAHLKAMP, Hendrik ; DOLGOV, Dmitri ; ETTINGER, Scott ; HAEHNEL, Dirk ; HILDEN, Tim ; HOFFMANN, Gabe ; HUHNKE, Burkhard ; JOHNSTON, Doug ; KLUMPP, Stefan ; LANGER, Dirk ; LEVANDOWSKI, Anthony ; LEVINSON, Jesse ; MARCIL, Julien ; ORENSTEIN, David ; PAEFGEN, Johannes ; PENNY, Isaac ; PETROVSKAYA, Anna ; PFLUEGER, Mike ; STANEK, Ganymed ; STAVENS, David ; VOGT, Antone ; THRUN, Sebastian: Junior: The Stanford entry in the Urban Challenge. In: *J. Field Robot.* 25 (2008), September, 569–597. `http://dx.doi.org/10.1002/rob.v25:9`. – DOI 10.1002/rob.v25:9. – ISSN 1556–4959

[MS08] MÜLLER, H. ; SCHULZ, G.: *Bundesfernstraßengesetz(FStrG): Mit Autobahnmautgesetz*. Beck C. H., 2008 (Gelbe Erläuterungsbücher). http://books.google.com/books?id=bHznLAAACAAJ. – ISBN 9783406575860

[PM10] POOLE, D.L. ; MACKWORTH, A.K.: *Artificial Intelligence: Foundations of Computational Agents*. Cambridge University Press, 2010 (Artificial Intelligence: Foundations of Computational Agents). http://artint.info. – ISBN 9780521519007

[QBP87] QUINN, J. B. ; BARUCH, J. J. ; PAQUETTE, P. C.: Technology in Services. In: *Scientific American* 257 (1987), Nr. 6, S. 50–58

[RH04] ROTHENGATTER, T. ; HUGUENIN, R.D.: *Traffic and transport psychology: theory and application : proceedings of the ICTTP 2000*. Elsevier, 2004 http://books.google.de/books?id=CuTGXLBSEZYC. – ISBN 9780080439259

[RJM98] ROBERT JOHN MCMILLAN, John Patrick H. Alexander Dean Craig C. Alexander Dean Craig: *Motor vehicle monitoring system for determining a cost of insurance*. August 1998

[RNC+96] RUSSELL, Stuart J. ; NORVIG, Peter ; CANDY, John F. ; MALIK, Jitendra M. ; EDWARDS, Douglas D.: *Artificial Intelligence: a Modern Approach*. Upper Saddle River, NJ, USA : Prentice-Hall, Inc., 1996. – ISBN 0–13–103805–2

[Sim06] SIMON, Julien: *Learning to drive with Advanced Driver Assistance Systems. Empirical studies of an online tutor and a personalised warning display on the effects of learnability and the acquisition of skill.*, TU Chemnitz, Diss., April 2006. http://nbn-resolving.de/urn:nbn:de:swb:ch1-200600713

[Tam00] TAM, Robert: PATH ATMIS State of the Research Annual Report Fiscal Year 1999 / 2000 / Institute of Transportation Studies, UC Berkeley. Version: 2000. http://econpapers.repec.org/RePEc:cdl:itsrrp:102723. 2000. – Institute of Transportation Studies, Research Reports, Working Papers, Proceedings

[TH02] TREIBER, M ; HELBING, D: Realistische Mikrosimulation von Strassenverkehr mit einem einfachen Modell. In: *16th Symposium Simulationstechnik ASIM* (2002), 80. http://141.30.186.11/~treiber/publications/MOBIL_ASIM.pdf

[Thu99] THUBAUVILLE, W.: *Straßenverkehrsordnung StVO, Verwaltungsvorschrift zur Straßenverkehrsordnung VwV-StVO, Verkehrszeichenkatalog VzKat: Textausgabe mit Einführung*. Bundesanzeiger-Verl., 1999 http://books.google.com/books?id=CSGMPgAACAAJ

[TMD+06] THRUN, Sebastian ; MONTEMERLO, Mike ; DAHLKAMP, Hendrik ; STAVENS, David ; ARON, Andrei ; DIEBEL, James ; FONG, Philip ; GALE, John ; HALPENNY, Morgan ; HOFFMANN, Gabriel ; LAU, Kenny ; OAKLEY, Celia ; PALATUCCI, Mark ; PRATT, Vaughan ; Stang, Pascal ; Strohband, Sven ; Dupont, Cedric ; Jendrossek, Lars-Erik ; Koelen, Christian ; Markey, Charles ; Rummel, Carlo ; Niekerk, Joe van ; Jensen, Eric ; Alessandrini, Philippe ; Bradski, Gary ; Davies, Bob ; Ettinger, Scott ; Kaehler, Adrian ; Nefian, Ara ; Mahoney, Pamela: Stanley: The robot that won the DARPA Grand Challenge: Research Articles. In: *J. Robot. Syst.* 23 (2006), September, 661–692. `http://dx.doi.org/10.1002/rob.v23:9`. – ISSN 0741–2223

[TR10] Tom Robinson, Erik C. Eric Chan C. Eric Chan: Operating Platoons On Public Motorways: An Introduction To The SARTRE Platooning Programme. In: *17th World Congress on Intelligent Transport Systems* (2010), October

[Tre11] Treiber, M: *Microsimulation of Road Traffic Flow*. `www.traffic-simulation.de`. Version: 2011

[TTI11] TTI: *2011 Annual Urban Mobility Report*. `http://mobility.tamu.edu/files/2011/09/all-nat-cong-tables.pdf`. Version: September 2011

[ZSFB95] Zabat, Michael ; Stabile, Nick ; Farascaroli, Stefano ; Browand, Frederick: The Aerodynamic Performance Of Platoons: A Final Report / Institute of Transportation Studies, UC Berkeley. Version: Januar 1995. `http://ideas.repec.org/p/cdl/itsrrp/102947.html`. 1995 (102947). – Institute of Transportation Studies, Research Reports, Working Papers, Proceedings