Enhancing Parking Simulations Using Peer-Designed Agents

Michal Chalamish, David Sarne, and Raz Lin

***Abstract*—In this paper, we investigate the usefulness of peer-designed agents (PDAs) as a turn-key technology for enhancing parking simulations. The use of PDAs improves the system's ability to capture the dynamics of the interaction between individuals in the system, each theoretically exhibiting a different strategic behavior. Furthermore, since people in general are inherently rational and computation bounded, simulating this domain becomes even more challenging. The advantage of PDAs in this context lies in their ability to reliably simulate a large pool of human individuals with diverse strategies and goals. We demonstrate the efficacy of the proposed method by developing a large-scale simulation system for the parking space search domain, which plays an important role in urban transport systems. The system is based on 34 different parking search strategies. Most of these strategies are substantially different from synthetic strategies that are used in prior literature. A quantitative analysis of the PDAs indicates that they reliably capture their designers' real-life strategies. Finally, we demonstrate the usefulness of PDA-based parking space search simulation by utilizing it to evaluate four different information technologies that are of increasing use in recent years.**

***Index Terms*—Experimentation, multiagent systems, parking simulations, peer-designed agents (PDAs).**

I. Introduction

Simulation is an important tool for studying emergent behavior. Using simulation, a researcher can carry out extensive data collection in a simple, safe, and inexpensive way. While, in real transport systems, the cost of a system evaluation is substantial, monetary, or otherwise, the use of simulation enables evaluating various variants of the system quickly and simply "with the click of a button." Not surprisingly, in recent years, agent technology in artificial intelligence has become the dominating approach to developing simulation systems and large-scale distributed systems [1]. While agents' capabilities substantially improve the modeling of the individuals that they represent, researchers have noted the challenge of constructing simulation systems in domains where key players are people [2], particularly due to their diverse behavior patterns.

As a means for improving realism of agents in simulation, recent research has increasingly relied on peer-designed agents (PDAs), i.e., computer agents developed by human subjects, as a representative of people [3]. Just like agents in general, PDAs can be used all the time and repeatedly, and they enable pretesting, revision, and improvements of the system before deployment. However, in contrast to agents designed by a domain expert, PDAs also allow the generation of a large set of diverse strategies in a cost-effective manner while reducing effort and costs involved in the evaluation process. The underlying assumption in many PDA-based works is that PDAs adequately capture people's behavior; therefore, a PDA-based system is likely to represent a collective behavior similar to when populated with people [4]. Empirical investigation of the level of similarity observed between PDAs and people, on the other hand, while rare, is not conclusive; some work suggests a relatively strong correlation between the behaviors of the two [4], whereas in other work, PDAs are reported to act in a different manner than people to some extent [5], [6].

This paper presents a comprehensive evaluation of the usefulness of employing PDAs in a parking lot simulation [7]. Different people use different parking space search strategies mainly because they assign diverse values (or satisfaction measures) to different aspects of the process. Moreover, a person's behavior is affected in various ways by parking search strategies used by other drivers as they influence the availability of parking spaces (and their location) in the parking lot at any given time.

In the parking search domain, there are two domain characteristics that favor the use of PDAs as a good representation of people. First, the rules according to which all drivers are allowed to drive are known, which facilitates the prediction of what other drivers might do. Second, unlike other domains, in which PDAs were tested (e.g., security and negotiation [3]), the parking space search, as well as driving in general, is a task with which most people are well experienced and thus are likely to have a well-established strategy for it, based on prior experience.

The PDA-based parking lot simulation used in this paper is based on 34 PDAs, each equipped with a strategy developed by a different person. The use of PDAs in this case enabled a rapid development process and a quick adaptation to new settings. Our experimentation with the system supports the hypothesis that, once PDAs are designed, they enable the generation of a significant number of simulated individuals, each equipped with a strategic behavior that reliably represents the behavior of some person. Accordingly, we manage to accumulate a large set of strategies, sufficient to represent the simulated population and scale it up as needed by simply creating more instances of each agent type/strategy. In particular, the paper reports results of measuring the similarity between the strategic PDAs' behavior and the behavior of their designers. The results indicate that PDAs indeed reliably capture the strategic behavior of their designers. As a complimentary analysis, the realism of the PDA-based simulation developed for the parking space domain is also verified by four domain experts from one of the biggest medical centers in Israel, which has five different parking lots.

Finally, this paper demonstrates the usefulness of the PDA-based simulation in evaluating how different information technologies concerning the current state of the parking lot (for example, a sign indicating the number of vacant spaces) would affect drivers' parking space search behavior. We evaluate four such technologies. Our results, some of which are counterintuitive, reveal that certain technologies might not justify their costs as they generate very little improvement.

II. Related Work

Parking space search is an important problem for urban planners and has been addressed by many traffic and transportation researchers [8], [9]. Previous work in the area of parking simulations has placed little emphasis on extracting user parking strategies, and most simulations reported in the literature use a limited set of preprogrammed strategies. For example, in PARKSIM, a network model is used to represent the driver's choice of behavior and consequently compare parking facility layouts [10]. Other studies evaluate alternative parking system designs and operating policies [11]. Recently, Jonkers *et al.* [9] have recognized that people follow many different strategies when looking for a parking space, yet their simulation focused on just two strategies. To the best of our knowledge, no previous research in the area of simulated parking lots incorporates a large collection of realistic parking strategies as part of the entity design, as proposed in this paper.

Manuscript received March 20, 2012; revised June 19, 2012; accepted July 21, 2012. Date of publication August 20, 2012; date of current version February 25, 2013. The Associate Editor for this paper was R. J. F. Rossetti.

M. Chalamish is with the Department of Computer Science, Ashkelon Academic College, Ashkelon 78211, Israel (e-mail: chalm@acad.ash-college.ac.il).

D. Sarne and R. Lin are with the Department of Computer Science, Bar-Ilan University, Ramat-Gan 52900, Israel (e-mail: sarned@cs.biu.ac.il; linraz@cs.biu.ac.il).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TITS.2012.2210545

In recent years, we have witnessed a substantial increase in the use of autonomous agents in medium-scale and large-scale simulation systems and, in particular, in the simulation of transportation systems [12]. Agents have the potential to improve performance of traffic and transportation systems [1]. Still, the task of reliably capturing individual agents' behaviors remained within the responsibility of the simulation designer. Despite various implementations based on existing tools, individuals' behavior in parking lot simulations was usually modeled using statistical data and optimization models [8], or prespecifying agents' roles using a set of parameters [9].

The use of PDAs has been proposed for experimental usage to investigate decision-making in a multiplayer computer game by Grosz *et al.* [5], in which agents had to reason about other agents' personalities in environments in which agents are uncertain about each other's resources. Emphasis was mainly placed on the long-term modeling of specific players rather than on a large-scale real-life simulation. Lin *et al.* [3] suggested using PDAs to remove people from the evaluation loop but focused solely on bilateral negotiations.

III. Peer-Designed Agent-Based Simulation

The main idea in PDA-based design is to separate the logic of the simulated agent's strategy from the generic functionalities required to have the agent participate in the simulation. This way, the strategy development task can be distributed and delegated to a population that is as similar as possible to the simulated population. We refer to these people as *strategy designers* as they embed their strategies in PDAs.

The simulation is controlled and dynamic, and allows the use of various agents of different types. Different simulation parameters can be varied as well (e.g., using different parking lot structures, lanes, and driving directions). The simulation also stores, at any given time, the system's current state (e.g., the location of cars and the location of occupied and vacant parking spaces).

As part of the development process, strategy developers are provided with skeleton PDAs, i.e., PDAs that are fully functional and lacking only their strategy layer. Each such agent possesses all the functionalities it needs to interact with the simulation manager and the simulated environment. This serves three purposes. First, it facilitates the strategy programming task as it allows people to focus on this specific task, leaving out complex programming tasks such as threading, communication, event handling, validation components, etc. Second, it enables the programming of PDAs directly by people with minimal programming background. Finally, the use of generic functionality simplifies integration and debugging.

IV. Parking Lot Simulator

The parking search problem typically considers a specific parking lot design, including overall size, internal structure, number and location of entry and exit points, alignment of different lanes, and permitted driving directions. Different cars are constantly going into and out of the parking lot. Every driver entering the parking lot has to individually solve his parking problem by finding a free (valid) space in which to park his car. The simulation can dictate the rate of cars arriving at the parking lot during each time step and the rate of cars exiting it, and the drivers' observations of their vicinity in any location.

This domain embeds many different strategies that are pursued by people. Moreover, the search for a parking space in a parking lot is conducted under partial information and limited visibility of the environment. Therefore, although a person may have experienced performing this task many times during his life, each attempt is different from the other in terms of available parking spaces, their location, and the number and behavior of other cars that search for parking spaces in parallel. As such, an optimal solution for this problem is unattainable due to both its complexity and the diverse concepts of optimality used by other drivers.

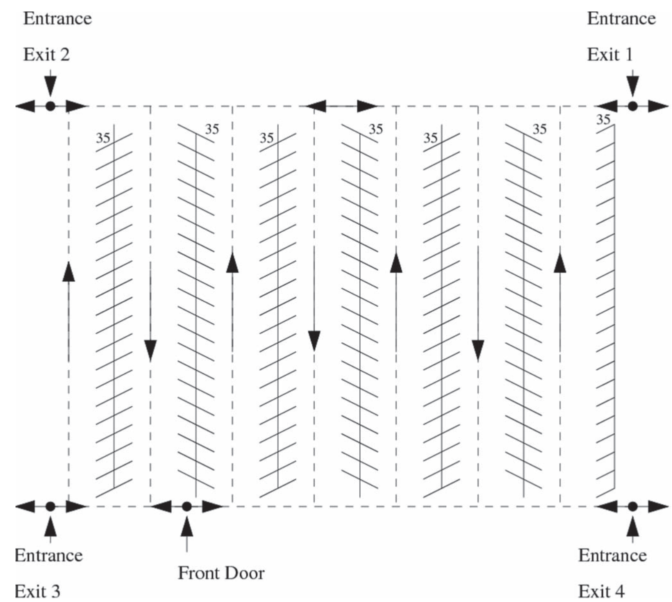

Fig. 1. Parking lot structure. The parking lot has four entrance–exit points and seven rows of parking spaces. Traffic lanes are denoted by dashed lines, and the arrows indicate the permitted traffic direction(s).

We implemented a PDA-based simulator for the parking space search domain. The PDAs, detailed design instructions, code skeleton, simulation settings described in this paper, and several movies of the simulation are available for download.[1]

The parking lot used in our experiments (see Fig. 1) has the typical structure of an underground parking lot and is based on the work of Cassady and Kobza [7]. It has four entrance–exit points from which a car can enter or exit. There are seven rows of parking spaces, each with 70 aligned parking spaces, except for the rightmost row that has only 35 spaces (the parking lot contains a total of 455 available parking spaces). The traffic lanes are denoted by dashed lines. The arrows in each traffic lane indicate the permitted traffic direction(s). The front door (through which individuals need to walk after parking their car) is located at the bottom end and serves as the sole pedestrian entrance/exit point. Foot traffic (walking to and from the door) is restricted to the traffic lanes; however, it may proceed in either direction within the lane.

Cars entering the parking lot in the simulator are represented by PDAs. Each PDA in the system, at every time step, needs to decide whether to move to an adjacent space, park the car, or wait in place. The PDAs base their decision on their knowledge of the parking lot structure and their current and former views of their surroundings. A PDA's view of its surroundings is based on its viewing frame (shown in Fig. 2), which corresponds to the limited view that a driver has when sitting inside his car and looking sideways in search of a parking space. The circle located at the lower part of the frame denotes the driver's location. The frame provides a view of five parking spaces forward (including the agent's location), two spaces backward, and three spaces in width. For example, in Fig. 2, we can identify the following: 1) another agent that has entered the agent's viewing scope (marked by a solid circle); 2) eight occupied parking spaces (marked by crossed squares); 3) six free parking spaces (marked by noncrossed squares); and 4) the traffic directions (marked by arrows).

[1] http://u.cs.biu.ac.il/~linraz/ParkingPDAs

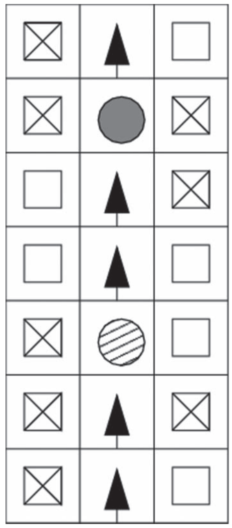

Fig. 2. Limited view frame example. Solid circles represent other agents within the agent's viewing scope, crossed squares represent occupied parking spaces, noncrossed squares represent free parking spaces, and arrows represent the traffic directions.

New cars (PDAs) are constantly added to the simulation. Each PDA is required to drive around the parking lot until it finds a suitable parking space. Based on the location of the parking space, the walking distance to the front door may vary.

PDAs' strategies were developed by 34 computer science students, representing a portion of the population of strategy designers, i.e., those equipped with some minimal programming skills. More importantly, subjects from this group drive cars and park them at the university parking lots, malls, etc. Although only 34 subjects were used, this allowed us to obtain many strategies, which is far more than surveyed and used in today's research—both in terms of the number of strategies, their variety, and PDAs' number.

To represent an individual's cost function, we defined three parameters whose significance could be controlled by each individual. The three parameters are[2] 1) *walking distance* from the parking space to the front door, 2) parking's *search time*, and 3) parking space's *distance from exit*. Using these parameters, a driver's goal is to minimize its cost, which is measured as

$$\begin{aligned}&\text{Cost(WalkTime, SearchTime, ExitTime)}\\&\quad = C_1 \cdot \text{WalkTime} + C_2 \cdot \text{SearchTime} + C_3 \cdot \text{ExitTime}\end{aligned}$$

where C_i is the weight assigned to the ith cost component $\Sigma C_i = 1$. The strategy designers were instructed to assign whatever values that they believe represent their cost function and then develop their parking strategy accordingly. The idea is that strategy designers should not be limited by a predefined utility function that they need to follow but rather can define their own utility function based on the importance that they see in each measure. For example, some PDAs simply try to find parking spaces as close as possible to the front door, employing some threshold to the acceptable distance of the parking space. Others use heuristics to minimize their searching time, whereas some just park in the first available parking space found. This stage is completely external to the strategy design process as participants could have developed PDAs, even without formally expressing their utility function. However, understanding the utility function of the different strategy designers is important when measuring the social welfare achieved in each simulated setting. It also becomes helpful in the postprocessing drill-down analysis of the strategies used as it allows a natural classification of strategies.

[2]These parameters were pointed out to be the most significant parameters based on a preliminary survey of the participants, and they resemble (with a few changes) the measures pointed out by Cassady and Kobza [7].

TABLE I
ENTRY–EXIT PROBABILITIES FOR THE DIFFERENT SETTINGS

	Setting 1	**Setting 2**	**Setting 3**
Entry Point 1	0.4	0.5	0.1
Entry Point 2	0.2	0.3	0.2
Entry Point 3	0.1	0.4	0.5
Entry Point 4	0.1	0.3	0.1
Exit Point 1	0.1	0.2	0.4
Exit Point 2	0.1	0.1	0.3
Exit Point 3	0.2	0.1	0.3
Exit Point 4	0.4	0.2	0.6

V. EXPERIMENTAL DESIGN AND RESULTS

Prior to developing their strategies and to evaluate strategy designer and PDA similarity, the strategy designers were asked to participate in a parking lot simulation as avatars, operating one of the cars using a GUI client. The other cars in the simulations were PDAs from a pool developed by a different population of strategy designers. Each strategy designer was required to play at least 20 sessions, each of which terminated when successfully parking his car.

Measuring the similarity of a PDA to its designer was done by placing the agent in the same state as the strategy designer was in, giving it the identical observations that the human designer had and comparing the decisions made by the PDAs to the decisions made by the strategy designer at the same decision situation. This was done 100 times for each choice of action taken by the human designer to comply with random-based strategies as well. The percentage of identical actions out of the total decisions made was used as a measure for the PDA's probability of taking the same action as the human designer with respect to the given action choice; the average of these probabilities was used as the similarity level. For achieving the latter two goals, several large-scale simulations were carried out, each consisting of nine different settings, differing in their initial parking lot occupancy percentage and entry–exit probability sets. The idea of varying these parameters stems from the need to simulate different parts of the day or week. For example, when arriving at a shopping center's parking lot on a Sunday afternoon, we are likely to encounter a high level of occupancy. Similarly, if it is a factory's parking lot, mornings are usually associated with a relatively high ratio between car arrival (entrance) and departure (exit) rates, and afternoons are usually characterized by a reverse pattern.

The three initial occupancy percentages that were used are 90%, 75%, and 60%. In general, operating in a parking lot with a relatively high initial occupancy percentage is expected to result in greater difficulties for drivers to find a "satisfying" parking space. In addition, three sets of entry–exit rates were defined for each entry and exit point (cf. Table I). Generally, the greater the exit rate (relative to the entrance rate), the easier it is for drivers to find an available parking space. Each simulation lasted 3000 time steps, starting from a variable vacancy pattern of the parking lot.

1) Similarity Performance: The resulting average similarity between PDAs and their designers was 77%. This is a relatively high similarity as it is not feasible to achieve a 100% similarity. This is mainly due to two reasons. First, people's strategies always involve some "noise," i.e., decisions are made, which do not follow the overall strategy, for numerous reasons (e.g., not thinking the move through, losing concentration, dealing with a situation never encountered

TABLE II
PERCENTAGE AND AVERAGE SIMILARITY OF PARTICIPANTS AND PDAs BASED ON A GIVEN THRESHOLD OF SIMILARITY

	Parking
similarity ≥ 0.7	85%
avg similarity	0.85
similarity ≥ 0.8	70%
avg similarity	0.88
similarity ≥ 0.9	26%
avg similarity	0.92

before, etc.). Second, people often do not follow the same strategy every time they are presented with the same situation [13]. A drill-down analysis was carried out to investigate the contribution of different individuals to the average degradation in the average similarity. Table II shows the percentage of participants that achieved a similarity score greater than the specific thresholds and the average similarity achieved by that group. In Table II, we observe that a small group of "bad strategy producers" substantially affect the overall similarity results. For example, excluding the participants that scored below a 70% similarity (15% of the population) results in an average similarity value of 85%. This entails that a preliminary stage, in which those few strategy designers associated with bad similarity values are removed from the system, can substantially improve the realism and validity of the simulation.

2) Variety of Strategies: One of the main strengths of the PDA-based approach is that it enables acquiring a large set of strategies in a cost-effective manner. However, there is no guarantee that the obtained strategies are indeed inherently different from one another. Prior work has used a very limited set of different strategies, assuming that people typically use homogeneous parking search strategies. Therefore, in the following, we present a complementary analysis of the difference between the 34 strategies used in our simulator.

We extracted 14 different strategy characteristics from the documentation and code that were received from the strategy designers. The most common strategies were calculating a dynamic route, parking nearest to the exit, parking in any available space after a certain amount of time, and parking in the first available space. The average agent implemented a combination of three to four different strategies. Each characteristic was, on average, used in seven different PDAs, where the two most common characteristics (parking nearest to the exit and calculated driving) were found to be implemented by almost 60% of the strategy designers.

An alternative to classifying and grouping PDAs according to their strategy characteristics is a classification according to weights that strategy designers assigned to each component in their cost functions. In our case, this resulted in six groups. The first group was characterized by considering the walking distance factor (e.g., $C_1 = 0.9$) as substantially important. The second group considered the walking distance as generally important but gave a higher weight to the search time factor (e.g., $C_1 = 0.65,\ C_2 = 0.2$). PDAs in the third group considered the walking distance and the search time factors to be equally important with equal weights (e.g., $C_1 = 0.45,\ C_2 = 0.45$). PDAs in the fourth group considered the search time factor to be the most important with the highest weight and the walking distance to be important as well (e.g., $C_1 = 0.25,\ C_2 = 0.7$). PDAs in the fifth group considered the search time factor as very important (e.g., $C_2 = 0.9$), whereas the last group deemed the distance from the exit as the most important factor (e.g., $C_3 = 0.5$).

Finally, in Fig. 3, we present the average individual search time of each of the 34 PDA types. The importance of this latter figure is that, through the use of such an objective measure (unlike the subjective cost measure), we can identify differences in the population as a whole. As can be observed in Fig. 3, there is a great variance in the search time performance measure, indicating a great variance in the strategies used by the different PDAs.

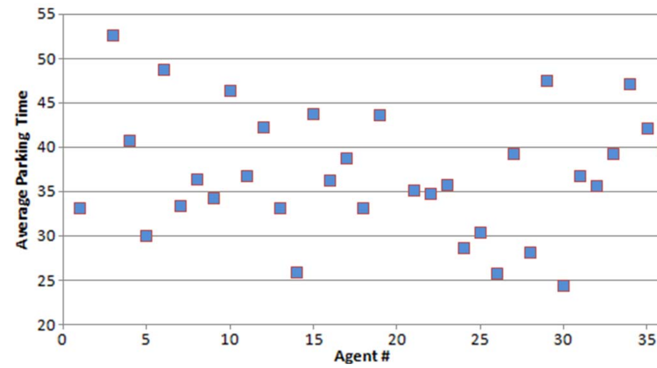

Fig. 3. Individual PDAs' performances based on their parking search time.

3) Usability: To further validate the realism obtained with the parking allocation simulation, a complementary subjective study with four domain experts from the Rabin Medical Center in Petach Tikva, Israel, was carried out. While this evaluation is not a sufficient condition for the validity of the model, it does serve as an important complementary means for demonstrating the usefulness of the approach.

The Rabin Medical Center is one of the biggest medical centers in Israel. The center has five parking lots: Four of which are single-story lots (as in our experiment), and the fifth is an underground two-story parking lot. This paper included surveying the deputy security officer of the center, two security employees whose job is to monitor the parking lots (using video surveillance, round the clock), and a police officer. The experts were presented with the GUI manager of the system and watched the behavior of the PDA-based simulation in eight random runs (each run included the entrance and parking of 200–300 cars). They were then asked to evaluate the realism of the simulation. The feedback received from the experts suggests that the system has successfully managed to simulate the dynamics of a realistic parking lot and that the behaviors exhibited by the different PDAs do, in fact, resemble those of people. Furthermore, the experts pointed out that some of the observed strategies are those that they might have adopted themselves in similar settings.

A. Comparison With Legacy Strategies

The results reported earlier are encouraging in the sense that PDA-based simulation indeed offers a highly reliable representation of real-life parking lots. However, to fully reason about the innovation encapsulated in the new approach, one needs to understand the magnitude of the difference between the results obtained with a full set of PDA-based parking behaviors and traditional approaches. For this purpose, we carried out three more simulations, each using a different parking behavior by the agents, corresponding to prior work in the area of parking search literature. The first two simulations were designed according to the two scenarios analyzed in Cassady and Kobza's seminal research [7]. In the first simulation, every agent used the Pick a Row, Closed Space (PRCS) strategy, according to which, upon entering the parking lot (from any of the four entrances), the agent selects a row from the set of allowed rows corresponding to the entrance used. The agent then proceeds to that row, enters it, and selects the closest available space to the main entrance (within its viewing frame limitations if progressing toward the front door). In the second simulation, all agents use the Cycling (CYC) strategy, according to which, each agent selects the row corresponding to the entrance used. From the available parking spaces in that row that are

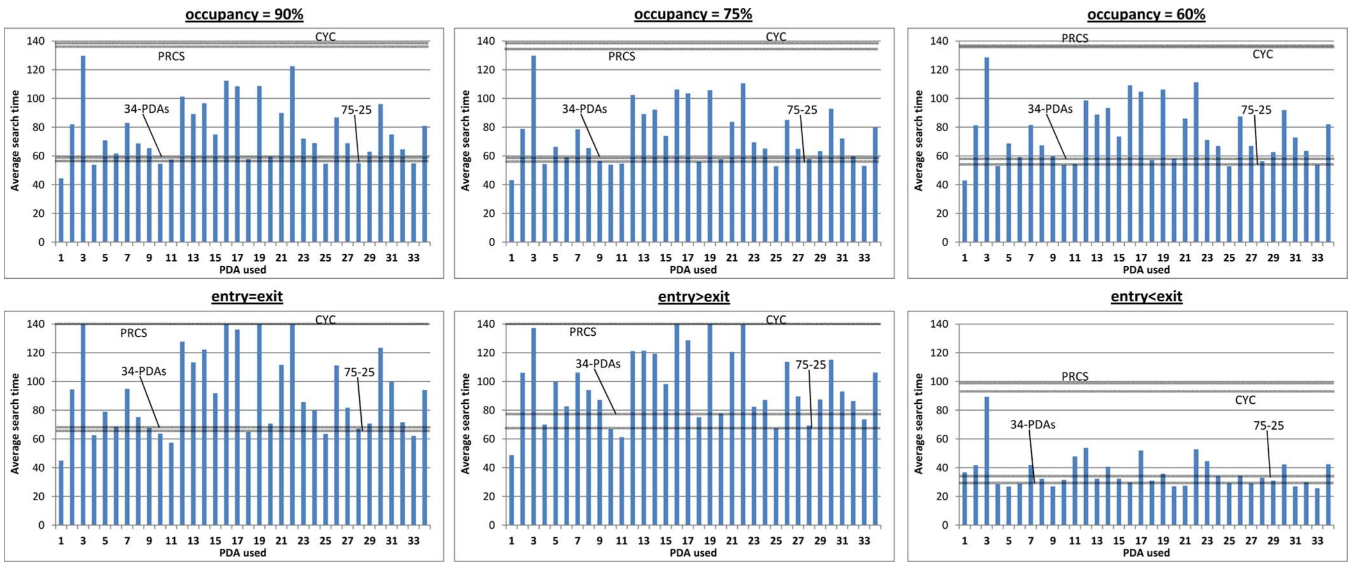

Fig. 4. Comparison of the average search time according to the different simulations, when using each of the 34 PDAs (exclusively), the CYC strategy (exclusively), the PRCS strategy (exclusively), a mixture of 75% of the agent parking in the closest available space to the front door and 25% of the agents parking closest to their current position (75–25), and the mixture of the 34 PDAs.

within the 20 closest parking spaces to the front door (if any), the agent will choose the one closest to the front door (within its viewing frame limitations if progressing toward the front door). If no such space is available, the agent proceeds to an adjacent row; however, this time, it is willing to settle for a space that is within the 40 closest parking spaces to the front door. Otherwise, the agent keeps moving to an adjacent row and picks the parking space closest to the front door among those available, and so on.

The last simulation relied on having 75% of the agents who enter the parking lot use a parking strategy aimed at parking in the closest available space to the front door. The remaining 25% use a parking strategy aimed at finding a parking space closest to their current position. This setting is based on the synthesis suggested in [9], which is based on several prior empirical studies that consider parking preferences. Based on the analysis of the results of other studies, Jonkers *et al.* reduce the different parking strategies into the given two, with a 1 : 3 ratio. Since, out of the pool of the originally received PDAs, 35% emphasized walking distance as the major component of their utility (which is the equivalent to attempting to find a space closest to the main entrance), and 20% emphasized search time (which is the equivalent to attempting to find a space closest to current position); the best of breed within each group was used for simulating agents of each group.

Fig. 4 depicts the average search cost of each simulation, when using each of the 34 PDAs exclusively (i.e., all agents instantiated along the simulation are of the same PDA), and the performance of the simulations executed based on the CYC strategy, the PRCS strategy, the Jonkers *et al.* strategy (denoted 75–25), and the mixture of the 34 PDAs. These latter measures are represented as the gray horizontal lines in each graph. The figure serves several purposes. First, it demonstrates the difference between results obtained when using only a single type of a parking search strategy (whether it is one of the strategies used in the literature or a single PDA's strategy) in the simulation compared with using a large set of strategies (the mixture of 34 PDAs). Second, it highlights the substantial difference between the obtained simulation performance when adopting different parking search behaviors suggested in previous studies. Finally, it demonstrates that a very small number of PDAs potentially used strategies suggested in prior literature. This latter observation is based on the assumption

TABLE III
COMPARISON OF THE AVERAGE SEARCH TIME WHEN USING THE TWO-STRATEGIES SIMULATION (75–25) AND THE PDA-BASED SIMULATION (BASED ON THE 34 PDAs)

Scenario	34-agents	75-25	t-test	F-test
oc=90%,$en = ex$	60.9	61.7	0.27	0.0007
oc=90%,$en > ex$	72.9	85.9	1.01E-12	9.4E-88
oc=90%,$en < ex$	38.9	29.9	1.45E-50	1.8E-105
oc=75%,$en = ex$	71.0	71.7	0.318	0.0352
oc=75%,$en > ex$	66.6	72.2	0.0001	9.37E-43
oc=75%,$en < ex$	32.9	30.2	1.12E-07	7.2E-07
oc=60%,en=ex	66.6	73.2	1.9E-05	3.76E-39
oc=60%,$en > ex$	65.0	75.2	2.11E-12	1.84E-68
oc=60%,$en < ex$	32.9	29.8	1.96E-10	9.93E-18

that similar strategies are likely to result with similar performance (in terms of average search time) in similar settings. Furthermore, even in cases where the use of a single PDA seemingly results in the same performance as CYC or PRCS, the phenomena are not consistent along all tested settings, suggesting that the implementation is partially different.

While the results of CYC- and PRCS-based simulations are substantially different from those of the simulation using all 34 agents, the use of the 75–25 set of strategies seems to yield results that are relatively close to those of the 34-agent set. However, a drill-down analysis of the results obtained for each of the nine settings used (varying the initial parking lot occupancy and the entry versus exit rates) reveals that the performances of the different agents are significantly different (when using t-test) for seven out of the nine cases, and in all cases, an F-test for variance reveals a significant difference (see Table III). Therefore, the collective behavior of the system with only two types of agents (as in the 75–25 case) is different from the one obtained using our PDA-based simulator.

VI. USING PEER-DESIGNED AGENTS TO COMPARE TECHNOLOGIES

Once we have established the grounds for using PDAs in parking simulations and demonstrated their efficacy in simulating reliable behaviors of people and a wide range of strategies, we continue to show the strength of PDAs in hands of researchers. Here, we describe how PDAs can be used to compare between different information technologies implemented in parking lots and, thus, assist in making the right choice when planning to use a new technology.

A. Information Models

In recent years, parking lot simulations have become an efficient means for estimating the usefulness of new information technologies, such as electronic signs indicating occupancy, aiming to help drivers park more efficiently [14]. As the results of our experimentation reveal, investment in the most sophisticated technology does not always improve the system's performance.

We tested four information technologies, which we compare to a setting in which no information technology is used, denoted as *noInfo*. While in *noInfo*, only the general parking lot map and a limited viewing frame are provided to the PDAs at every time step; the other four technologies offer additional information as follows:

1) Free Parking Space Info ($info_{\mathrm{FPS}}$)—Upon entering the parking lot, the driver receives updated information concerning the current number of unoccupied parking spaces (e.g., by electronic signs posted at parking lot entrances).
2) Static Free Parking Map Info ($info_{\mathrm{SFPM}}$)—Upon entering the parking lot, the driver receives an updated map where free parking spaces are marked (on entry only) (e.g., by large electronic screens posted at the parking lot entrance that illustrate the parking lot's current space status by coloring occupied and free spaces in different colors).
3) Dynamic Free Parking Map Info ($info_{\mathrm{DFPM}}$)—Similar to the previous technology, however, the driver constantly receives an updated map of free parking spaces (e.g., by GPS tools or lights from the ceiling, relying on sensors embedded throughout the parking lot).
4) Dynamic Full Map Info ($info_{\mathrm{DFM}}$)—In addition to relying on GPS tools, however, the driver receives an updated map, where both the free parking spaces and the exact location of all other drivers are marked.

By measuring the change in agents' performance when using a new technology, we can estimate the value obtained by using it.

B. System's Lower Bound

The performance evaluation of each of the new technologies was based both on absolute performance and the magnitude of the relative improvement obtained, measured in terms of costs. For the purpose of calculating the relative improvement between any two information models, a performance baseline (a "gold standard") was required. Obviously, the baseline is not zero since, as efficient as the agents can be, there is still a lower theoretical limit to the system's performance (e.g,, whatever the theoretical allocation that maximizes the overall social welfare may be, drivers still need to get to their parking spaces and walk to the main entrance, reflecting some costs).

A good candidate for the performance baseline is the centralized allocation of parking spaces to drivers when they are fully cooperative and obey the instructions received. Alas, finding the optimal allocation is extremely complex as there are many parameters to consider. To simplify the evaluation, we developed a loosened lower bound using the permissive assumption that more than a single car can park in one parking space at a given time (overlap) as long as the maximum accumulated time cars park in a given space does not exceed the simulation duration. Thus, we replace the scheduling problem (scheduling cars to parking spaces over time according to their arrival time and parking space availability) with an allocation problem, which is a standard optimization problem, enabling the calculation of a lower bound. This lower bound is used as a baseline for calculating the relative change in performance when applying each of the tested technologies.

TABLE IV
IMPROVEMENT (i.e., LOWER COSTS) RELATIVE TO *noInfo*, IN PERCENTAGES. THE FIRST ROW PRESENTS THE AVERAGE COSTS. THE REST OF THE TABLE PRESENTS A BREAKDOWN ACCORDING TO THE DIFFERENT VALUES USED FOR THE INITIAL CAPACITY AND THE ENTRY–EXIT RATE

	$info_{FPS}$	$info_{SFPM}$	$info_{DFPM}$	$info_{DFM}$
avg	0.24%	2.36%	12.55%	11.46%
90%	0.18%	1.89%	9.46%	8.77%
75%	0.39%	2.09%	15.80%	11.50%
60%	0.14%	3.09%	12.45%	14.12%
en=ex	-2.13%	1.05%	18.02%	16.78%
en>ex	2.71%	1.26%	4.08%	2.80%
en<ex	-0.24%	7.81%	20.02%	19.75%

C. Effects of Information Technologies

Nine different simulations, varying in the occupancy and entry–exit rates (similar to those used earlier), were run with each information technology, and the total cost of the different PDAs (based on their specific cost function) were averaged. The average costs received were 37.32, 37.28, 36.89, 35.05, and 35.24 for the *noInfo*, $info_{\mathrm{FPS}}$, $info_{\mathrm{SFPM}}$, $info_{\mathrm{DFPM}}$, and $info_{\mathrm{DFM}}$, respectively. The lower bound for the overall cost (i.e., an upper bound for the overall performance) was 19.2. While this lower bound is significantly lower than the theoretical minimum cost that can be achieved, it still serves us well in illustrating the magnitude of the improvement obtained with the different types of information supplied to the PDAs.

Table IV presents the relative average cost improvement, in percentages, achieved when applying each information model in comparison to the *noInfo* model, with respect to the lower bound. The relative improvement was calculated as the ratio between the difference in the costs for each information model and the *noInfo* model, and the difference between the *noInfo* model and the lower bound ($\mathrm{RelativeImprovement} = (\mathrm{Cost}_{noInfo} - \mathrm{Cost}_{info_i})/(\mathrm{Cost}_{noInfo} - \mathrm{Cost}_{\mathrm{lowerbound}})$), where $info_i$ is the information model for which the improvement is calculated). The first row presents the improvement in the average costs over all nine combinations of occupancy and entry–exit rates. The next three rows represent the improvement in the average costs where the initial occupancy was fixed and the Entry–exit rates were varied. The last three rows represent the improvement in the average costs where the entry–exit rate was fixed, and the initial occupancy was varied. The results support the hypothesis that each information model lowers the total cost (i.e., improves social welfare), relative to the *noInfo* model. The results also reveal that the greatest improvement relative to *noInfo* is with $info_{\mathrm{DFPM}}$, namely, when the exact free parking spaces map is given at each time step.

To check the statistical significance of the results, we carried out a three-way ANOVA analysis. Duncan Grouping showed significance ($p < 0.001$) in placing information models *noInfo*, $info_{\mathrm{FPS}}$, and $info_{\mathrm{SFPM}}$ in one group and information models $info_{\mathrm{DFPM}}$ and $info_{\mathrm{DFM}}$ in another. The primary implication of this result is that in cases where social welfare is of primary concern, there is no significant benefit in investing in $info_{\mathrm{FPS}}$ and $info_{\mathrm{SFPM}}$; however, there is a benefit to installing $info_{\mathrm{DFPM}}$ or $info_{\mathrm{DFM}}$-like technologies (therefore, the decision should depend on their cost). Furthermore, if an $info_{\mathrm{DFM}}$ solution costs more than an $info_{\mathrm{DFPM}}$ solution, then the latter should always be preferred.

While the insignificant difference between $info_{\mathrm{DFM}}$ and $info_{\mathrm{DFPM}}$ can be explained by the relative insignificant value of knowing the location of other cars in real time, the insignificant performance difference between both $info_{\mathrm{FPS}}$ and $info_{\mathrm{SFPM}}$ and no information at all (*noInfo*) seems surprising since both methods enable drivers to distinguish between different occupancy states of the parking lot. Furthermore, the insignificant performance difference between $info_{\mathrm{FPS}}$ and $info_{\mathrm{SFPM}}$ themselves is also surprising since the latter actually supplies more targeted information. Apparently, in this domain, the information received upon entry becomes obsolete relatively quickly, explaining the insignificant difference in performance. Similarly, the substantial improvement achieved with $info_{\mathrm{DFM}}$ and $info_{\mathrm{DFPM}}$ (as compared with the other technologies) is attributed to the importance of relying on real-time information for decision-making in this domain. Finally, we note that the most significant improvement with $info_{\mathrm{DFM}}$ and $info_{\mathrm{DFPM}}$ is obtained in the case where the exit rate exceeds the entry rate (and *vice versa*). This nonintuitive result can be explained by the fact that when one has continuous full information and more parking spaces become available, the "best" available parking space frequently changes (as many cars leave). Therefore, if drivers always head to the best parking spaces, they frequently change their goal, resulting in a frequent alternation between different parking spaces.

VII. Discussion and Future Work

The encouraging results reported in this paper has strengthen our belief in the usefulness of PDA technology for constructing successful and reliable transportation simulations. While the availability of PDAs makes them appealing for use in the investigation of behavior, trends, and implication, one must assure that they reliably represent the environment studied to produce effective results.

In this paper, we show that PDAs can successfully be used in parking lot simulations. In particular, we have demonstrate their ability to simulate a large pool of human individuals, resulting in a variety of applicable strategies. Nine different settings were used, and a simulation was easily constructed and carried out using PDAs, which otherwise (using people) would have taken a long period of time. We continued and demonstrated how PDAs can reflect on the efficacy of different possible information technologies for the parking search domain before undergoing a lengthy period of deployment, which could turn out to be futile. Finally, we demonstrated the variety of strategies populated in the simulation, as compared with current legacy strategies commonly used in simulations of this type.

While the results reported in this paper cannot be generalized to all transportation systems simulations, many such systems share the same characteristics found in the parking space search (e.g., an environment characterized with partial and mostly local information, mainly due to a limited view of the driver and the fact that the strategies being employed are well established due to the day-to-day experience of drivers). Therefore, the new approach has a lot of potential in this field.

Future research warrants exciting new extensions to the method proposed. Among them, we intend to calibrate PDAs with learning techniques to try and improve similarity percentages. We plan to pursue the development of a tool for analyzing discrepancy in the behavior of individuals in the simulation. In addition, we plan on extending the use of PDAs for other transportation simulations, and varying different settings, constraints, and rules by which they must abide.

Acknowledgment

The authors would like to thank R. Magori-Cohen for her support with the statistical tests and the Rabin Medical Center, Belinson Campus, and its security department for their assistance with the simulation verification.

References

[1] B. Chen and H. H. Cheng, "A review of the applications of agent technology in traffic and transportation systems," *IEEE Trans. Intell. Transp. Syst.*, vol. 11, no. 2, pp. 485–497, Jun. 2010.
[2] Y. Murakami, Y. Sugimoto, and T. Ishida, "Modeling human behavior for virtual training systems," in *Proc. AAAI*, 2005, pp. 127–132.
[3] R. Lin, S. Kraus, Y. Oshrat, and Y. K. Gal, "Facilitating the evaluation of automated negotiators using peer designed agents," in *Proc. AAAI*, 2010, pp. 817–822.
[4] M. Chalamish, D. Sarne, and S. Kraus, "Programming agents as a means of capturing self-strategy," in *Proc. AAMAS*, 2008, pp. 1161–1168.
[5] B. Grosz, S. Kraus, S. Talman, B. Stossel, and M. Havlin, "The influence of social dependencies on decision-making: Initial investigations with a new game," in *Proc. AAMAS*, 2004, pp. 782–789.
[6] A. Rosenfeld and S. Kraus, "Modeling agents based on aspiration adaptation theory," *Auton. Agents Multi-Agent Syst.*, vol. 24, no. 2, pp. 221–254, Mar. 2012.
[7] C. R. Cassady and J. E. Kobza, "A probabilistic approach to evaluate strategies for selecting a parking space," *Transp. Sci.*, vol. 32, no. 1, pp. 30–42, Jan. 1998.
[8] Z. Mei and Y. Tian, "Optimized combination model and algorithm of parking guidance information configuration," *EURASIP J. Wireless Commun. Netw.*, 2011 : 104, doi:10.1186/1687-1499-2011-104.
[9] E. Jonkers, M. van Noort, and J. van der Veen, "Parking guidance modelling, simulation and impact assessment," in *Proc. 14th Int. IEEE Conf. Intell. Transp. Syst.*, 2011, pp. 858–864.
[10] W. Young and W. Yue, "A study of the performance of two parking-lot layouts," *Traffic Eng. Control*, vol. 33, no. 7/8, pp. 434–439, Jul. 1992.
[11] R. H. Ellis, J. C. Bennett, and P. R. Rassam, "Parking systems analysis: An overview," *ACM SIGSIM Simul. Dig.*, vol. 6, no. 1, pp. 42–50, Oct. 1974.
[12] I. Benenson, K. Martens, and S. Birfir, "PARKAGENT: An agent-based model of parking in the city," *Comput., Environ., Urban Syst.*, vol. 32, no. 6, pp. 431–439, Nov. 2008.
[13] F. T. Dolbear and L. B. Lave, "Inconsistent behavior in lottery choice experiments," *Behav. Sci.*, vol. 12, no. 1, pp. 14–23, Jan. 1967.
[14] Y. Asakura and M. Kashiwadani, "Effects of parking availability information on system performance: A simulation model approach," in *Proc. Veh. Navigat. Inf. Syst. Conf.*, 1994, pp. 251–254.