Spatio-temporal Efficiency in a Taxi Dispatch System

Darshan Santani, Rajesh Krishna Balan and C Jason Woodard
School of Information Systems
Singapore Management University
{dsantani,rajesh,jwoodard}@smu.edu.sg

ABSTRACT

In this paper, we present an empirical analysis of the GPS-enabled taxi dispatch system used by the world's second largest land transportation company. We first summarize the collective dynamics of the more than 6,000 taxicabs in this fleet. Next, we propose a simple method for evaluating the efficiency of the system over a given period of time and geographic zone. Our method yields valuable insights into system performance—in particular, revealing significant inefficiencies that should command the attention of the fleet operator. For example, despite the state of the art dispatching system employed by the company, we find imbalances in supply and demand during peak business hours, often resulting in unacceptably high waiting times for passengers. Surprisingly, we also find imbalances during off-peak hours in certain zones. Finally, we discuss how techniques from multi-agent systems research, such as distributed coordination and market-based resource allocation, may be effective in improving the performance of this important mode of public transportation.

1. INTRODUCTION

Taxicabs play an important role in bridging the gap between private transportation, buses, and rail systems in urban areas. Although taxi services differ significantly by country, they share many common characteristics worldwide.

In particular, passenger demand for taxicabs depends not only on the price of a ride but also on the waiting time. Excess capacity in the form of empty cabs is not necessarily wasteful but in fact may be socially desirable, as it increases the value of the service through lower waiting times [6]. Taxi drivers, on the other hand, prefer their vehicles to be occupied by fare-paying passengers as much of the time as possible. This tension presents the fleet operator with the difficult problem of balancing the interests of these two stakeholder groups while running an efficient and profitable business.

We studied this problem using GPS data provided by Comfort-DelGro, the operator of Singapore's largest taxi fleet. The data set records the movement of 6,230 GPS-enabled taxicabs over a 24-hour period, as well as the time and location of 38,048 booking requests. Using this data, we computed the average occupancy rate and median passenger waiting time by hour and geographic zone. This enabled us to identify times and locations in which the demand for taxis exceeded the supply or vice versa. Simple scatter plots

of occupancy versus waiting time reveal systematic differences by zone in the efficiency of taxi allocation, as well as periods in which the efficiency within a zone deviates positively or negatively from the average.

The remainder of the paper is organized as follows. Section 2 reviews the empirical literature on taxi fleets and some of the related work on autonomous vehicles. Section 3 provides a detailed overview of the data set. Section 4 describes our analysis of the data and presents our main results. Section 5 discusses the implications of this finding and suggests ways that techniques from multi-agent systems research may help to improve the performance of the system. Section 6 concludes.

2. RELATED WORK

There is a substantial amount of economic literature on taxicab markets, along with a number of related multi-agent systems papers motivated by autonomous vehicles and other transportation-related applications. Most of the economic studies emphasize the role of excess capacity and waiting time in the demand for taxi services.

2.1 Economic Perspectives

Beesley [2] explains the number of taxicabs in operation in London, where market entry is free, as a function of several variables. Using data from 1960 to 1976, he finds that the number of cabs is positively related to the number of visitors to London, taxi fares, underground fares, and unemployment levels. In another study, Schroeter [11] presents a model of taxicab service with radio dispatch and airport cabstands. Using data from a Minneapolis taxi firm, he concludes that an increase in the number of operating cabs would cause a decrease in waiting times and an increase in demand in the dispatching segment of the market—but not by enough to increase the average revenue per cab.

The literature on taxicab markets has provided solid theoretical insights, but estimating and testing these models has been difficult due to a lack of empirical data. Schaller [10] contributes to addressing this shortcoming by assembling a data set based on taxi meter and odometer readings that the New York City Taxi and Limousine Commission collects in its periodic inspections of New York taxicabs. Based on his data, he estimates elasticities for taxicab fares and service availability, and finds that an increase in the number of taxi licenses does not affect the revenue of existing taxicab drivers. Flores-Guri [6] tests a model of a regulated taxicab market using Schaller's data set and finds a positive, inelastic relationship between vacant taxicabs and the demand for taxicab services. He also evaluates the effects of policy changes such as increases in the regulated fare and the number of licenses issued.

.

2.2 Systems Perspectives

Advances in technology have provided a new but largely untapped source of information about taxicab markets in the form of GPS data. Our work is most closely related, at least data-wise, to the analysis performed by Liao [8]. In particular, he used similar GPS data from ComfortDelGro for his analysis. However, our work extends into areas (such as detailed efficiency analysis) that were not previously tackled.

Looking a little more broadly, we find numerous pieces of work that complement the analysis presented in this paper. For example, Yoon et al. [13] use a combination of vehicles equipped with GPS technology plus low-bandwidth cellular updates to dynamically estimate street traffic. In particular, they show that it is possible to estimate road conditions using just a few vehicles. In the area of traffic signal optimization, Dresner and Stone presented schemes to (a) allow emergency vehicles to go through signal junctions faster [4], and (b) to improve general throughput at intersections [5]. In addition, Bazzan [1] and Oliveira et. al [3] showed that it is possible to effectively control a series of distributed traffic signals. Finally, Tumer and Agogino [12] showed that agents could be used to dynamically reduce congestion in an air traffic network.

3. BACKGROUND AND DATA

In this section, we describe the data used for this analysis.

3.1 System Overview

Singapore has a world-class public transportation system with an extensive network of taxis, buses, and rapid transit rail lines that provide convenient and affordable services to the city-state's population of 4.5 million. With widely available and relatively low-priced taxicabs (metered fares rarely exceed US $15)—along with high taxes on private cars and petrol —many Singaporeans find it unnecessary to own a car. Taxicabs can be flagged down at any time of the day along any public road, with well-marked taxi stands available outside most major shopping centers and office buildings. Taxis can also be reserved by telephone or the Internet for an additional fee of US $2–3. In March 2007, Singapore's taxi fleet consisted of 23,315 taxicabs operated by seven companies (three owned or controlled by ComfortDelGro) and several hundred independent owners [7].

To reduce peak-hour traffic congestion on major highways and in the Central Business District (CBD), Singapore uses an Electronic Road Pricing (ERP) system that automatically charges drivers when they enter an ERP zone. When a taxi is occupied, these charges are automatically added to the metered fare. Passengers also incur a Peak Hour Surcharge during peak business hours (7–9 am and 5–8 pm on weekdays), a Central Business District Surcharge applicable for trips originating in the CBD between 5–8 pm on weekdays, and a Late Night Surcharge applicable between 11:30 pm and 6 am.

3.2 Taxi Location and Booking Data

The data set was collected on Thursday March 1, 2007. This was a typical rainy day in Singapore (although unseasonably cold, with an average temperature of 23 degrees Celsius).

We categorized our data into two different sets. The first set contains the movement and status information of 6,230 GPS-enabled Comfort taxicabs (3.6 million observations) over the 24-hour period of the day. Each observation in this set contains the taxi's unique vehicle identification number, GPS coordinates, vehicle state, vehicular speed and a timestamp. The taxicabs could potentially be in one of ten different states during the day. However, among these states, *BREAK, BUSY, FREE, OFFLINE, ONCALL* and *POB* are

State	Frequency (%)	Taxi Status
FREE	46.9	Active and Available
POB	32.4	Active and Occupied
BREAK	6.20	Active but Unavailable
ONCALL	4.93	Active but Unavailable
BUSY	4.91	Active but Unavailable
OFFLINE	3.13	Inactive

This table shows the six most common taxi states along with their frequencies (the fraction of the day spent in each state by an average taxi). *POB* occurs when the taxi has a passenger. *ONCALL* occurs when the taxi is responding to a booking call.

Table 1: Taxi States

the most significant ones, accounting for almost 98.5% of the observations in the data set. Table 1 describes these states in more detail; they are the only ones we considered in this analysis.

The second set contains the 38,048 taxi booking calls made by passengers (using a phone, Internet, fax machine, or automated kiosk) throughout the day. Each entry in this set contains a unique booking identification number, GPS coordinates of the booking location (where the customer wants the taxi to arrive) and a timestamp. All our analysis was performed using the R programming language [9].

3.3 Summary Statistics

Table 2 shows the summary statistics for our data. We observe that, per day, an average taxi spent about 18 hours on the road (taxis are typically shared by more than one driver), traveled about 710 kilometers at an average speed of about 39 km/h, and picked up about 26 passengers. An interesting observation is that taxis were occupied (i.e., had a paying customer in them) only about 42% of the time—suggesting possible inefficiencies in the system. We revisit this issue in Section 4.

4. ANALYSIS AND RESULTS

The main goals of our analysis were (a) to characterize the efficiency of the taxi system, and (b) to explore the sources of inefficiency, with a view to mitigating them. We began by graphically plotting sampled taxi locations on a two-dimensional grid and comparing the density of taxis with the frequency of bookings in each of four geographic zones. We then computed the occupancy rate and median passenger waiting time for each hour and zone. These measures enabled us to identify times and locations in which the demand for taxis exceeded the supply or vice versa. Finally, simple scatter plots of occupancy versus waiting time reveal systematic differences by zone in the efficiency of taxi allocation, as well as periods in which the efficiency within a zone deviated positively or negatively from the average.

4.1 Locations and Bookings

Figure 1 shows a subset of the observed taxi positions (0.3% of the data set sampled uniformly at random) plotted by their geographic coordinates. The city is divided into four zones, numbered 1 to 4. The points within each zone are shaded according to the density of bookings in that zone, with the highest shaded black.

- **Zone 1**, the Central Business District (*Central*), contains most of the high-rise office buildings in the city, with relatively few

Figure 1: Taxi Observations by Location and Booking Frequency of Zone

residential areas. It has the highest density of bookings (as indicated by the shading), as well as the highest density of taxis (as indicated by the close spacing of the points).

- **Zone 2**, which we label *Condo* because of its concentration of private condominiums and landed housing, has the second highest density of both bookings and taxis. This zone also includes Orchard Road, Singapore's main retail shopping district.

- **Zones 3 and 4**, labeled *West* and *East* respectively, include most of Singapore's public housing, industrial estates, and unpopulated areas. Zone 3 has the third-highest booking density but the lowest taxi density, and conversely for Zone 4.

The figure shows the major roads (which, unlike the road networks typically considered in simulation studies, are not arranged in a grid), and suggests patterns of traffic congestion that appear as differences in density. Gaps in the plot represent nature reserves, unpopulated water catchment areas, or military installations that are off limits to taxis.

This figure contributes to our efficiency analysis in two ways. First, it serves to validate the data. Although a few coordinates are clearly inaccurate due to GPS errors—unless the taxis in question were amphibious vehicles—the overall picture closely resembles a map of Singapore, as we would expect. Second, it provides a visual representation of a key baseline finding: at the level of the four zones over the entire day, the supply of taxis is closely—though not perfectly—related to demand in the form of bookings.

To simplify the discussion that follows, we omit from our subsequent analysis the Tuas and Jurong industrial estates in the west and Changi Airport in the east. (The booking patterns for Tuas and Jurong more closely resemble the *Central* area than the surrounding

residential areas. The airport receives few bookings but attracts a large number of empty cabs, distorting our calculation of taxi occupancy.) While the impact of these omitted areas on the rest of their zones may be non-negligible (e.g., there might also be an unusual number of empty cabs in the zone but more distant from Changi Airport), we believe these effects are of second-order importance compared to our main findings.

Table 3 shows the relative size of these zones (with the regions omitted as mentioned above). The table also shows the number of distinct businesses and households located in each zone.

4.2 Occupancy and Waiting Time

To reason about the efficiency of the taxi system, we need appropriate metrics. From a passenger's point of view, a perfectly efficient taxi system is one that delivers an empty taxi to his or her location at the instant one is desired. In contrast, a taxi driver's idea of an efficient system is one in which his or her cab is never empty—the moment one passenger pays and leaves, another gets in. Obviously these goals are in conflict. Rather than take one side or the other, we say that an **efficiently balanced** taxi system is one in which the satisfaction of passengers and drivers is *inversely correlated*. That is, when passengers are happy because of short waiting times, drivers are unhappy because of low occupancy, and vice versa. In other words, an efficiently balanced taxi system is one in which supply and demand tend to move in the same direction in response to external shocks.

Our index of **driver satisfaction** is simply the ratio $O/(O+A)$, where O = the number of minutes the taxi was occupied (in state *POB*) in the given zone and hour, and A = the number of minutes it was available (in state *FREE*). We call this ratio the *occupancy rate*.

Count Data	
Total Taxis	6,230
Inactive Taxis	268
Total Taxi Observations	3,658,507
Erroneous Observations	20,039
Booking Calls	38,048

Summary Statistics (Per Taxi for Full Day)			
	Mean	Median	Std.
Updates per Hour (no.)	33.35	33.70	9.48
Time on the Road (hrs)	17.99	18.38	4.18
Average Speed (km/h)	38.94	33.00	39.47
Distance Traveled (km)	700.53	–	–
Street Pickups (no.)	21.1	21.0	8.50
Booking Pickups (no.)	4.6	4.0	3.64
Total Pickups (no.)	25.7	26.0	9.68
Average Occupancy Rate (%)	42.46	46.17	17.24

Distanced traveled is computed by multiplying the average speed with the average time on the road. We did not compute a median or standard deviation.

Table 2: Summary Statistics for Taxi Data

Zone Size		Occupants	
Area (km^2)	Perimeter (km)	Total Households	Number of Businesses
Zone 1 - Central			
58.93	50.50	114,521	44,964
Zone 2 - Condo			
531.97	39.10	89,493	24,399
Zone 3 - West			
194.62	75.78	253,336	20,201
Zone 4 - East			
187.51	71.01	448,703	70,058
Total			
973.03	236.38	906,053	159,622

Table 3: Demographics by Zone

We measure **passenger satisfaction** by the *median waiting time* for a booking in a given zone during a given hour. Two sources of data contribute to this measure. First, for every taxi in the zone during the hour, we observe its state transitions and note the time lag between a transition to *ONCALL* from any other state and a transition to *POB* from *ONCALL*. Each of these time lags (of which there may be several in a given zone-hour) represents the duration between a booking request and the pickup that satisfies it.

Second, we also observe that—especially during peak periods—some booking requests may never be satisfied. From the customer's point of view, this is equivalent to an infinite waiting time. We obtain this number by taking the difference between the number of bookings in a zone-hour and the number of successful *ONCALL* pickups. (Note that we do not have data on how long customers wait to hail a cab on the street without a booking.)

As long as fewer than half of the booking requests in a zone-hour are unsatisfied, the median waiting time will be well defined. Otherwise we arbitrarily define it as 20 minutes, which is five minutes longer than the longest waiting time for a successful pickup in our data set.

Figure 2 shows four graphs, one for each geographic zone, of the occupancy rate and median waiting time in the given zone for each hour in the day.

- In **Zone 1**, we see a moderate surge in occupancy in the morning, while waiting time doubles from about 5 minutes to 10. The evening rush is similarly profitable for drivers (with occupancy rising again to over 60%) but much worse for passengers, who in the median case cannot get a cab at all between 10 pm and 12 midnight. Occupancy falls way off again late at night, as waiting times also reach their minimum for the day.

- In **Zone 2**, the time of peak demand is in the morning from 6–10 am. This is consistent with the residential nature of this zone. Many residents do not have easy access to mass transit lines, so they rely on taxi bookings for their morning commute to work. Note that passenger waiting time goes up almost exactly in tandem with occupancy, indicating that the system is being stretched about equally on both the supply and demand sides.

- In **Zone 3**, waiting time and occupancy track each other quite closely except for the usual late-night drop in occupancy (for which drivers are compensated by a fare surcharge) and a sharp spike in waiting time in the late morning (10 am to 12 noon). It is possible that this spike is due to a chance event (perhaps related to the rain on the day the data was collected), but it could also be explained by drivers either taking a break or simply not being in the zone after the morning rush brought many of them into the city center.

- In contrast, **Zone 4** is almost totally well behaved. We see a moderate rise in occupancies and waiting time in the morning, but both decline gradually over the rest of the day.

These patterns clearly highlight situations in which either waiting time spikes without a similar rise in occupancy (as in Zones 1 and 3), or occupancy drops without a similar reduction in waiting time (as in all four zones late at night). In the former case, customers are worse off but drivers are no happier. In the latter case, *drivers* are worse off but *customers* are no happier. The third step in our analysis explores these relationships more systematically.
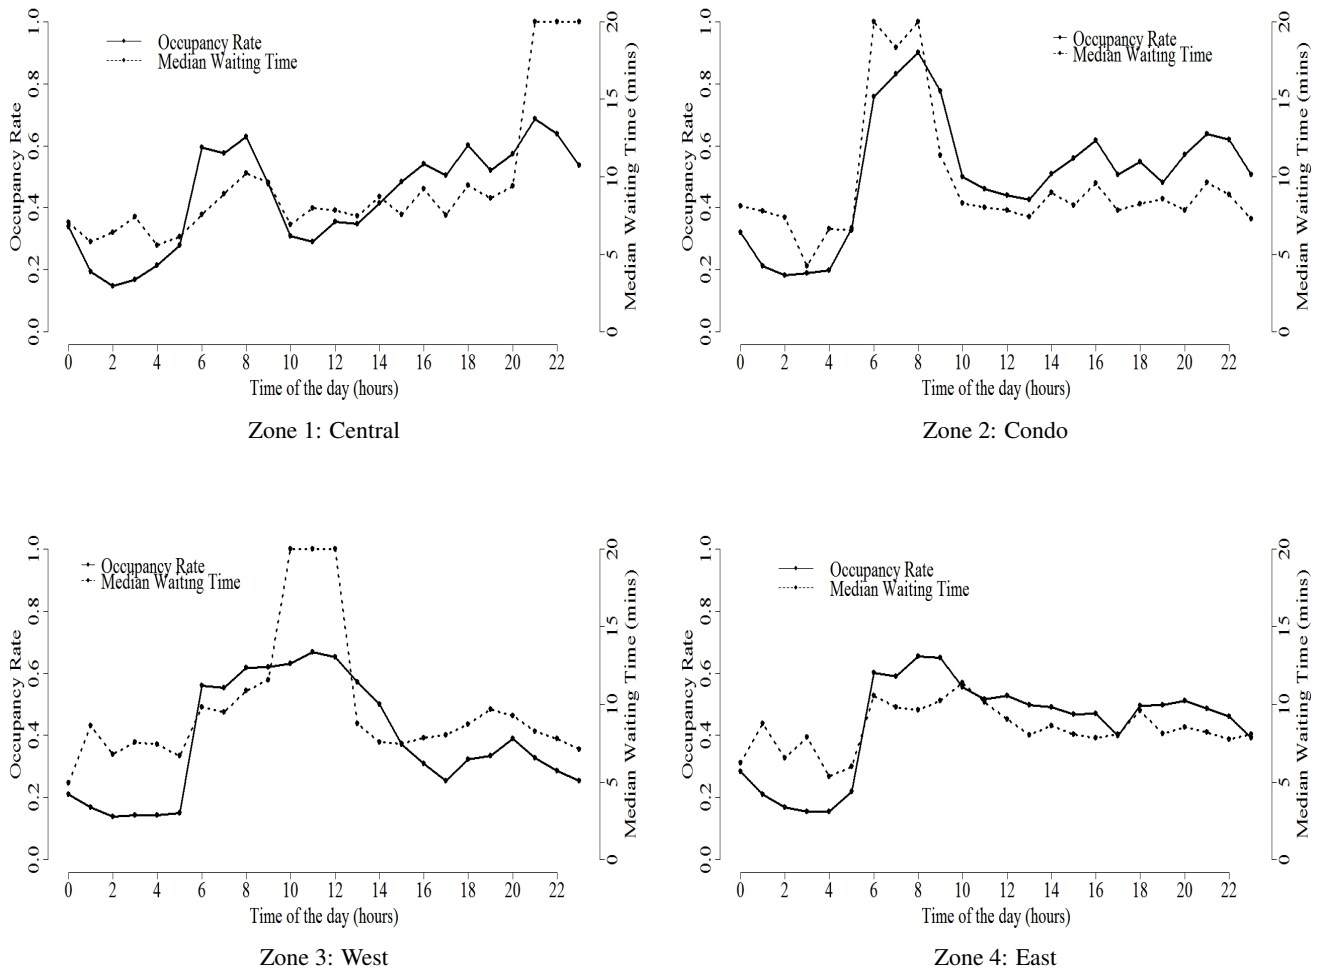
Zone 1: Central

Zone 2: Condo

Zone 3: West

Zone 4: East

Figure 2: Occupancy Rate and Mean Waiting Time by Zone and Hour

4.3 Predictability and Efficiency

Figure 3 is similar to Figure 2 in that it shows four graphs, one for each geographic zone. However, instead of plotting occupancy and waiting time by hour, the new figure shows each hour's observation on a scatter plot with waiting time on the horizontal axis and occupancy on the vertical. Two lines of best fit are plotted, with correlation coefficients r and $\tilde{r}$, respectively: the dark one (r) includes the outliers on the far right (hours in which the median booking request was unsatisfied), while the latter one ($\tilde{r}$) omits them.

As we would expect, waiting time and occupancy are positively correlated in all four zones. In general, that is, longer waiting times (unhappy customers) are associated with higher occupancy (happy drivers), and vice versa. In all zones except 4, this correlation is even stronger if we remove the waiting time outliers that correspond to hours in which the ability of the system to match available drivers to waiting passengers was exceptionally poor.

Note that if these variables were perfectly correlated, the fleet operator would be forced to choose between satisfying drivers and passengers in what is effectively a zero-sum game. Of course, to the extent that the firm seeks to maximize its own profits, it no doubt

does something like this already when its managers (in conjunction with the government, which exerts control over the company through regulation and indirect corporate ownership) determine the number of taxis to put on the road and set policies for compensating drivers.

Although the firm's profit maximization problem is interesting and important, we want to focus on the efficiency gains that can be made through operational improvements, e.g., to the centralized taxi dispatching system and/or the local incentives of drivers or passengers. To see the opportunities for such improvements in Figure 3, we simply need to look at the points that lie below and to the right of the best-fit line. These are the points for which the waiting time is unusually high for a given level of taxi occupancy—or, equivalently, the occupancy is unusually low for a given level of waiting time.

(Similarly, points above and to the left of the best-fit line are relatively *efficient*. The fleet operator might want to think about how to move the line of best fit "northwest" on the graph—achieving lower waiting times, on average, with higher levels of occupancy.)
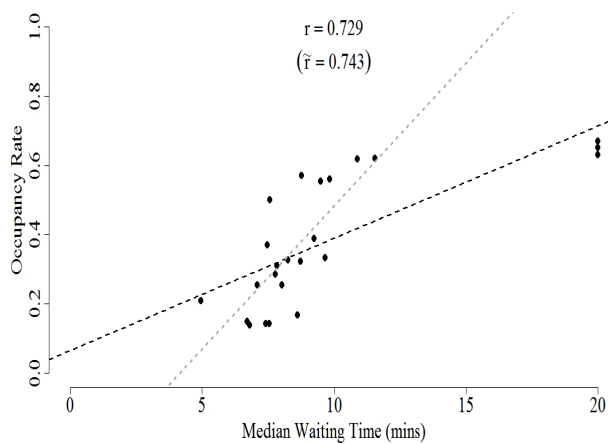
Most of the "southeast" points on the figure occur during the late night periods, when we already know from Figure 2 that occupancy
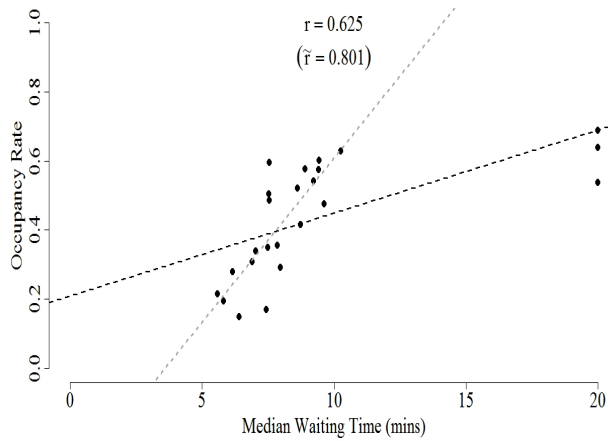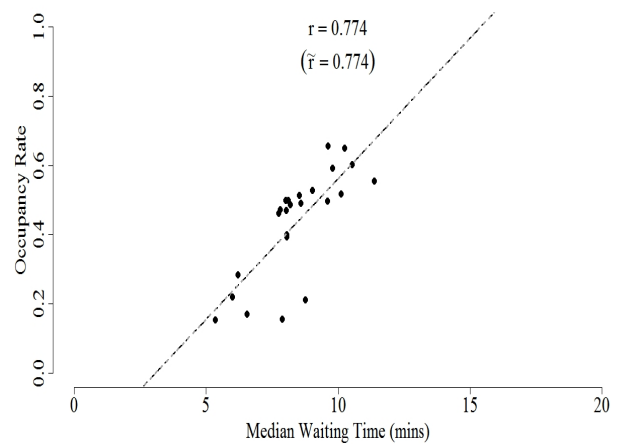
Zone 1: Central
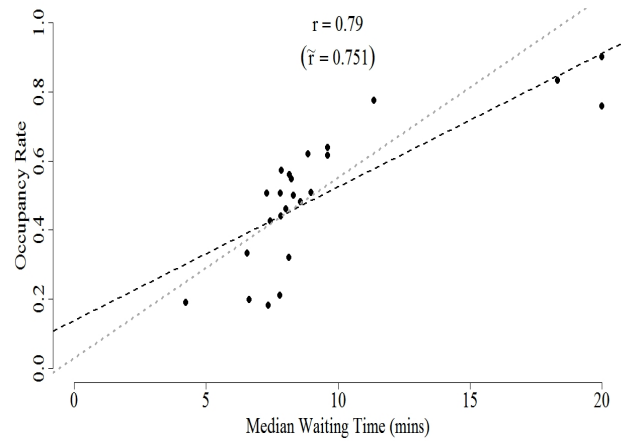
Zone 2: Condo

Zone 3: West

Zone 4: East

The correlation coefficient of the dark line is r.

Figure 3: Correlation Between Hourly Observations of Occupancy and Waiting Time by Zone

is low. This, in turn, is partly due to the fare surcharge, which is designed to ensure that the night shift remains profitable enough to attract a sufficient number of drivers to prevent taxi shortages at critical times of the night, e.g., when bars and clubs close. Thus some of what we discover by looking at these correlations turns out not to be mysterious. Nonetheless, it demonstrates the power of our approach to identify potential trouble spots in a simple and repeatable way.

5. DISCUSSION

In this paper we have demonstrated some simple kinds of analysis that shed light on the efficiency of a real GPS-enabled taxi dispatching system. We have considered the perspective of passengers and drivers as well as the fleet operator itself.

The use of these kinds of techniques is by no means confined to offline analysis of historical data. Graphs such as those shown in Figures 2 and 3 can easily be computed online if the raw data are available, allowing the dispatching system to respond in near-real time. Moreover, the high resolution of GPS coordinates allows ge-

ographic zones to be much more granular that we have shown so far. (In fact, the actual data set from ComfortDelGro is divided into 100 zones, which we aggregated into four for ease of exposition.) It is easy to imagine a working system to identify **hotspots** characterized by unusually long wait times—or indeed **coldspots** with unusually low occupancy.

It is also natural to ask how the inefficiencies identified by such a system could be mitigated. Rather than propose a specific solution, we appeal broadly to the multi-agent systems literature as a powerful and diverse source of techniques for improving coordination and resolving incentive conflicts in complex systems with both natural and artificial components. The taxicab setting we study is a prime example of such a system since it combines thousands of boundedly rational, imperfectly informed, and economically motivated human agents with an automated dispatching mechanism that relies on large amounts of real-time data.

Directions that could be taken to improve this system include:

- A more dynamic **market mechanism** for allocating book-

ings to taxis. The current system used by ComfortDelGro is a hybrid in which drivers may be assigned jobs if they are free, or bid for them by estimating a time to arrival if they have a passenger on board (since the dispatching system does not generally know the passenger's destination and thus cannot estimate when the driver will be free). Booking fees are fixed according to the time of day and day of the week. One could explore the use of more explicit prices (either dollars paid by consumers or an artificial currency transferred among drivers) to provide stronger incentives to respond to an emerging hotspot.

- A way to share **local information** among drivers that could enable hotspots to be identified and resolved more quickly and reliably. This could be as simple as incorporating a traffic estimation system in the spirit of Yoon et al. [13]. Harnessing such a system to adjust the booking allocation algorithm would be nontrivial but could have a significant impact.

- Application of **online learning** techniques to ensure that recurring patterns hot- and coldspots are anticipated and addressed proactively in the future. Part of the challenge with this would be to educate the humans "in the loop" as well as the automated dispatching system. But if drivers could be convinced that they can earn more money (and potentially work fewer hours) by optimizing their cruising patterns when they are free, they might embrace this kind of technology.

We look forward to pursuing these directions in future work.

6. CONCLUSION

In this paper, we presented an analysis of the movement of 6,230 taxis over a 24-hour period. We showed that the taxi system is mostly efficient. However, there are still periods of inefficiency that arise in this otherwise well regulated system.

A common problem that faces researchers is the lack of real-world data to inform theoretical models and simulations. To that end, we hope that the data and statistics presented in this paper, which were obtained from a large transportation network, prove useful to researchers developing other types of autonomous and semi-autonomous systems.

7. ACKNOWLEDGEMENT

8. REFERENCES

[1] A. L. Bazzan. A distributed approach for coordination of traffic signal agents. *Autonomous Agents and Multi-Agent Systems*, 10(1):131–164, 2005.

[2] M. E. Beesley and M. A. Kemp. Urban transportation. In E. S. Mills, editor, *Handbook of Regional and Urban Economics, Volume II*, chapter 26, pages 1023–1052. Elsevier Science Publishers, 1 edition, Dec. 1987.

[3] D. de Oliveira, A. L. C. Bazzan, and V. Lesser. Using cooperative mediation to coordinate traffic lights: a case study. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent systems (AAMAS)*, pages 463–470, New York, NY, USA, 2005. ACM.

[4] K. Dresner and P. Stone. Human-usable and emergency vehicle-aware control policies for autonomous intersection management. In *AAMAS 2006 Workshop on Agents in Traffic and Transportation*, May 2006.

[5] K. Dresner and P. Stone. Multiagent traffic management: Opportunities for multiagent learning. In K. Tuyls, P. J. Hoen, K. Verbeeck, and S. Sen, editors, *LAMAS 2005*, volume 3898 of *Lecture Notes in Artificial Intelligence*, pages 129–138. Springer Verlag, Berlin, 2006.

[6] D. Flores-Guri. An economic analysis of regulated taxicab markets. *Review of Industrial Organization*, 23(3–4):255–266, 2003.

[7] Land Transport Authority of Singapore. *Breakdown of Taxi Ownership for 2007*, July 2007. http://www.lta.gov.sg/corp_info/doc/taxi_info_for_lta_website_07.PDF.

[8] Z. Liao. Real-time taxi dispatching using global positioning systems. *Communications of the ACM*, 46(5):81–83, 2003.

[9] R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2007.

[10] B. Schaller. Elasticities for taxicab fares and service availability. *Transportation*, 26(3):283–297, 1999.

[11] J. R. Schroeter. A model of taxi service under fare structure and fleet size regulation. *Bell Journal of Economics*, 14(1):81–96, 1983.

[12] K. Tumer and A. K. Agogino. Distributed agent-based air traffic flow management. In *Proceedings of the Sixth International Conference on Agents and Multiagent Systems (AAMAS)*, Honolulu, Hawaii, May 2007.

[13] J. Yoon, B. Noble, and M. Liu. Surface street traffic estimation. In *Proceedings of the Fifth International Conference on Mobile Systems, Applications, and Services (MobiSys)*, San Juan, Puerto Rico, Jun 2007.