How Useful Can Parking Availability Information Be?

C. Hickert[a,b,*,†], S. Li[a,b,†], and C. Wu[a,b,c]

[a] MIT Laboratory for Information & Decision Systems, Cambridge, USA
chickert@mit.edu, siruil@mit.edu, cathywu@mit.edu
[b] MIT Institute for Data, Systems, and Society, Cambridge, USA
[c] MIT Dept. of Civil & Environmental Engineering, Cambridge, USA
[*] Corresponding author
[†] Equal contribution, listed alphabetically

Extended abstract submitted for presentation at the Conference in Emerging Technologies in Transportation Systems (TRC-30)
September 02-03, 2024, Crete, Greece

April 30, 2024

Keywords: Parking, vehicle routing, value of information, dynamic programming

1 INTRODUCTION

Current navigation apps send drivers to their destination, occasionally providing a basic assessment of parking in the area (e.g., ‘easy’, ‘difficult’) (Arora *et al.*, 2019). Given that parking space may not exist at the destination or may be unavailable, the driver may need to find parking elsewhere or ‘cruise’ for parking at the current location until a spot becomes available. Such a scenario — familiar to many urban drivers and riders — highlights the frustrating difference between today’s ‘time to drive’ estimates that ignore parking difficulties and more useful ‘time to arrive’ estimates that account for true drive times and post-parking walk time.

Beyond incurring inconvenience, underestimates of the true ‘time to arrive’ via personal vehicles may also prevent mode shift that would otherwise have occurred if individuals had the true estimates (Arora *et al.*, 2019). This is particularly true given that popular navigation apps (e.g., Google Maps, Apple Maps) include estimates for walking time from public transit to the final destination in their public transit travel time estimates, but do not include walking time from available parking to the final destination in their driving estimates. These situations contribute to the the congestion and emissions costs of ‘cruising’ for parking, a well-documented reality in urban environments (Shoup, 2006).

What would be the time and emissions savings if navigation apps routed drivers in need of parking to the best available location(s), rather than straight to their destination? This preliminary work adopts a value-of-information approach to investigate this question. Ultimately, this will feature formal analyses at three levels of parking availability information: (i) no information, (ii) distributional information (e.g., the driver knows there are 20% odds of finding a spot in lot i), and (iii) true availability information (e.g., lot i has 2 spaces available). To this end, this extended abstract makes two contributions: (1) a dynamic programming framework for characterizing the problem and (2) a closed-form analysis for setting (ii), delineating when it is optimal to wait at a specific parking lot as opposed to when it may be better to visit other lots, as well as identifying the expected cost in each case. We intend to use these as building blocks for comparison with (i) and (iii), as well as for a data-driven assessment in urban settings of the value of information — expressed in time and emissions savings — of strategies informed by this analysis as compared to strategies incentivized by today’s navigation apps.

This problem has key differences from the well-studied optimal stopping problem as it relates to parking (Sakaguchi & Tamaki, 1982, Tamaki, 1982, 1988), as well as other theoretical analyses of cruising (Arnott & Williams, 2017). Primarily, the driver may decide where to go to seek parking, rather than passively waiting for availability to appear along a pre-defined route with strictly decreasing distance (usually) to the destination. The decision-maker in our setting may also have varying awareness of availability; to the authors' knowledge, this is rarely considered in other formal analyses. Even in optimal stopping variants that allow backtracking, vehicles may not alter the order of their visits to various spaces, the driving geometry is 1-dimensional, and parking availability odds rarely vary (Tamaki, 1988, Krapivsky & Redner, 2019). Two well-executed closely related works are those described by (Djuric *et al.*, 2016) and (Hedderich *et al.*, 2018). However, these consider only street parking, do not use closed-form analysis to express or compare information regimes, and focus on cruising within a predefined tolerance boundary of the final destination instead of *a priori* selection of parking destinations separate from the final destination. Of course, various mobile applications (SpotHero, SpotAngels, ParkCBR, Parking.com, Park Smarter, etc.) exist to assist in parking, but generally these are simple reservation and/or payment systems for only those spots over which the app has control.

2 PRELIMINARY ANALYSIS

2.1 Framework

The initial challenge is to construct a framework that captures the uncertainty, spatial structure, time costs, and information (or lack thereof) involved in parking while enabling closed-form analysis. We model this situation using an infinite-horizon Markov Decision Process (MDP) $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, s_0)$. Each state $s \in \mathcal{S}$ is comprised of a tuple $(i, \{u, o\})$ where $i \in \{0, ..., N\}$ indicates the location of the origin (0) and any of N parking lots and the set $\{u, o\}$ indicates the vehicle's parking status: u is unparked and o is parked (since the spot was open). The initial state $s_0 = (0, u)$ thus indicates an unparked vehicle at the origin. Parked states are terminal.

At each timestep, let action $a_i \in \mathcal{A} = \{1, \ldots, N\}$ represent an attempt to park at lot i. (Note there is no a_0: a vehicle may not attempt to stay at or return to the origin.) This attempt succeeds with a lot-specific probability $p_i \in \mathcal{P}$ and fails with probability $1 - p_i$. We assume $p_i \in (0, 1]$. Let $r_{i,(j,\{u,o\})} \in \mathcal{R}$ represent the instantaneous reward incurred in the process of a vehicle beginning in state (i, u) and taking action a_j to park in a different lot j. If this succeeds, it incurs reward $r_{i,(j,o)} = -t_{i\to j} - t_{j\to D}$, where $t_{i\to j}$ represents the **drive time** from i to j and $t_{j\to D}$ represents the **walk time** from lot j to the true destination D. If the attempt fails, it only incurs drive time reward $r_{i,(j,u)} = -t_{i\to j}$. If the vehicle remains unparked at any lot, it has two options for the following timestep: either seek to park at a new lot k or else wait at the current lot for another chance at parking. The former incurs a reward of the formula just described, $r_{j,(k,u)}$ or $r_{j,(k,o)}$. Let t_{wait} indicate a **wait time** incurred for remaining at a lot and waiting for another chance at parking. The latter incurs a reward $r_{j,(j,o)} = -t_{wait} - t_{j\to D}$ if it successfully parks and $r_{j,(j,u)} = -t_{wait}$ if it does not. In summary, the possible rewards can be described as

$$r_{i,(j,\,\text{status}=\{u,o\})} = \begin{cases} -t_{i\to j} - t_{j\to D}, & \text{if } i \neq j \text{ and status} = o \\ -t_{i\to j}, & \text{if } i \neq j \text{ and status} = u \\ -t_{wait} - t_{j\to D}, & \text{if } i = j \text{ and status} = o \\ -t_{wait}, & \text{if } i = j \text{ and status} = u \end{cases}. \tag{1}$$

The objective is to find the action sequence that maximizes the expected sum of rewards (equivalently, minimizes expected travel time). This framing excludes additional driver preferences like pricing, but future work may integrate these. The MDP may be seen as a stochastic shortest path problem in which the terminating goal states $G \subset \mathcal{S}$ are the parked states; this termination condition thus obviates the need for a discount factor γ in the framing.

2.2 Optimal strategy and time cost in the distributional information setting

This work tackles scenario (ii) described above (in which distributional parking availability information is accessible for each lot) since it most closely relates to today's app-based driving decisions. Given pre-existing awareness of parking difficulty and app-generated data, it is uncommon to truly have no prior information about parking availability (Arora *et al.*, 2019). At the same time, real-time parking occupancy sensors are expensive and thus relatively rare (Djuric *et al.*, 2016). Furthermore, the analysis for this setting provides intuition for the others, to be addressed in the full work. Our analysis finds the optimal strategy and associated time cost falls into two structure-based regimes, described in turn by the propositions below.

Proposition 1 *If, $\forall(i,j)$ pairs, $t_{i\to j} \geq t_{wait}$ and $t_{0\to i} = t_{0\to j}$, then the optimal strategy is to drive directly to the lot i^* with the maximum value-to-go $V_{i^*} = -t_{i^*\to D} - \frac{1}{p_{i^*}} t_{wait}$.*

Note that the decision to park at a lot i can be conceptualized as a coin flip that succeeds (allows parking) with probability p_i. t_{wait} can be considered to be the time one must wait at a lot to again 'flip the coin' if the previous parking attempt was unsuccessful. Adopting a dynamic programming approach, one can thus explicitly write the cost for any 'patient' strategy — that is, the strategy in which a vehicle drives to the ith lot and waits there until it finds parking (possible since $p_i \in (0,1]$). The expected cumulative return $\mathbf{E}[R_{i,patient}]$ of this strategy is

$$\mathbf{E}[R_{i,patient}] = -t_{0\to i} - t_{i\to D} - \sum_{1\leq m\leq\infty} m \cdot t_{wait} \cdot (1-p_i)^{m-1} p_i, \tag{2}$$

where the sum corresponds to the expected value of a geometric random variable with success probability p_i; we can thus rewrite the above as

$$\mathbf{E}[R_{i,patient}] = -t_{0\to i} - t_{i\to D} - \frac{1}{p_i} t_{wait}. \tag{3}$$

We can write the value-to-go of any unparked state (i,u) under a patient policy at all times as $V_{(i,u),patient} = -t_{i\to D} - \frac{1}{p_i} t_{wait}$. By the second assumption in the proposition statement (namely, $\forall(i,j)$ pairs, $t_{0\to i} = t_{0\to j}$), we have that $i^* = \underset{i\in\{0,\ldots,N\}}{\operatorname{argmax}} V_{(i,u),patient} = \underset{i\in\{0,\ldots,N\}}{\operatorname{argmax}} \mathbf{E}[R_{i,patient}]$.

While this may seem a strong assumption, drive times to parking lots near a destination of interest are often quite comparable, especially relative to drive time variance due to other factors.

The question now becomes whether one can do any better than a patient strategy, i.e., by switching to a new lot $j \neq i^*$ in the case that it is unsuccessful parking in a chosen lot i^*. However, note that this would incur additional time cost of $t_{i^*\to j}$, which by assumption is at least as large as t_{wait}. By this fact and by the assumption that $t_{0\to i^*} = t_{0\to j}$, we have $-t_{0\to i^*} - t_{wait} + V_{i^*,patient} \geq -t_{0\to j} - t_{wait} + V_{j,patient} \geq -t_{0\to i^*} - t_{i^*\to j} + V_{j,patient}$, where we also use the fact that $V_{i^*,patient} \geq V_{j,patient}\ \forall\ j$ due to the optimality of i^*. That is, any impatient strategy switching to a lot j and terminating at lot j will have a worse expected return than a patient strategy terminating at lot i^*. Therefore, in this regime, the optimal policy is a patient one in which the vehicle drives from the origin directly to lot i^*.

Proposition 2 *If $\exists t_{j\to k} < t_{wait}$, there may exist a cluster of parking lots $\{j,\ldots,k\}$ such that it is better to visit the lots in that cluster rather than adopt a patient strategy at a parking lot i with the single best value $V_{(i,u)}$.*

Consider a cluster $\mathcal{C}$ of parking lots defined as those for which $t_{i\to j} < t_{wait}$ and $t_{j\to i} < t_{wait}$. Following from the above analysis, we can write the expected cumulative return for the strategy in which a vehicle navigates to this cluster and cycles through those lots until parking is found as $-t_{0\to\mathcal{C}} - t_{\mathcal{C}\to D} + \frac{1}{1-\prod_{i\in\mathcal{C}}(1-p_i)}(\max(-t_{wait}, -t_{\mathcal{C}}))$, where $1-\prod_{i\in\mathcal{C}}(1-p_i)$ is the probability that parking is available at any lot $i \in C$, $t_{\mathcal{C}}$ represents the time to navigate among lots in that cluster, the max operator reflects our modeling assumption of the t_{wait} waiting time between two consecutive 'coin flips' at the same lot, and $t_{0\to C}$, $t_{\mathcal{C}\to D}$ respectively represent the travel time from the origin to the cluster and from the cluster to the destination (both can be bounded above by the taking their maximums among lots within the cluster). There are settings in which cycling through a cluster $\mathcal{C}$ of parking lots is better than the single-lot optimal patient strategy at lot i^*, even when lot i^* is not in the cluster $\mathcal{C}$. These settings correspond to scenarios when (1) the joint probability $1 - \prod_{i\in\mathcal{C}}(1 - p_i)$ is high, (2) the travel time within the cluster t_C is small, and (3) the travel time $t_{0\to\mathcal{C}} + t_{\mathcal{C}\to D}$ is low. Intuitively speaking, these criteria reflect the fact that after an unsuccessful parking attempt at a lot $i \in \mathcal{C}$, we can benefit from trying a different lot $j \in \mathcal{C}, j \neq i$ during the wait for the next 'coin flip' at lot i, if we can travel from lot i to j relatively quickly, and lot j has a relatively high probability of parking. If the cluster $\mathcal{C}$ additionally incurs relatively low travel times with respect to the origin and destination, the cluster parking strategy could be better than the single-lot optimal patient strategy, where a time period of t_{wait} would be wasted between each 'coin flip' at i^*. On the other hand, if these conditions are not satisfied, then the single-lot optimal patient strategy at i^* would remain the optimal strategy in the distributional availability information regime.

3 DISCUSSION

Given space limitations, in future work we will provide explicit analysis for availability information scenarios (i) and (iii) and compare all three settings to identify the full value of information. We will also explicitly expand upon the the tradeoffs (1)-(3) described in Proposition 2.

Our next aim is to augment this with empirical results generated via simulation. Using OpenStreetMap, we will obtain data about parking lot locations relative to some popular location of interest. We will then simulate parking at these lots and calculate the value of information from setting (i) to (ii) to (iii). We hope to further estimate the reduction in carbon emissions and time saved via mode shift if 'time to drive' estimates are corrected to 'time to arrive' estimates that take parking availability and post-parking walking into account. These could also integrate driver preference models that consider factors beyond travel time alone.

References

Arnott, Richard, & Williams, Parker. 2017. Cruising for parking around a circle. *Transportation research part B: methodological*, **104**, 357–375.

Arora, Neha, Cook, James, Kumar, Ravi, Kuznetsov, Ivan, Li, Yechen, Liang, Huai-Jen, Miller, Andrew, Tomkins, Andrew, Tsogsuren, Iveel, & Wang, Yi. 2019. Hard to park? Estimating parking difficulty at scale. *Pages 2296–2304 of: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining.*

Djuric, Nemanja, Grbovic, Mihajlo, & Vucetic, Slobodan. 2016. Parkassistant: An algorithm for guiding a car to a parking spot. *Page 5433 of: Transportation Research Board 95th Annual Meeting*, vol. 16.

Hedderich, Mareike, Fastenrath, Ulrich, & Bogenberger, Klaus. 2018. Optimization of a Park Spot Route based on the A* Algorithm. *Pages 3493–3498 of: 2018 21st International Conference on Intelligent Transportation Systems (ITSC).* IEEE.

Krapivsky, PL, & Redner, S. 2019. Simple parking strategies. *Journal of Statistical Mechanics: Theory and Experiment*, **2019**(9), 093404.

Sakaguchi, Minoru, & Tamaki, Mitsushi. 1982. On the optimal parking problem in which spaces appear randomly.

Shoup, Donald C. 2006. Cruising for parking. *Transport policy*, **13**(6), 479–486.

Tamaki, Mitsushi. 1982. An Optimal Parking Problem. *Journal of Applied Probability*, 803–814.

Tamaki, Mitsushi. 1988. Optimal stopping in the parking problem with U-turn. *Journal of applied probability*, **25**(2), 363–374.