

Mobile Opportunistic Commerce: Mechanisms, Architecture, and Application

Ece Kamar
SEAS, Harvard University
Cambridge, MA 02138
kamar@eecs.harvard.edu

Eric Horvitz
Microsoft Research
Redmond, WA 98052
horvitz@microsoft.com

Chris Meek
Microsoft Research
Redmond, WA 98052
meek@microsoft.com

ABSTRACT

We present mechanisms, architectures, and an implementation addressing challenges with *mobile opportunistic commerce* centering on markets and mechanisms that support the procurement of goods and services in mobile settings. Our efforts seek to extend core concepts from research in electronic commerce to interactions between mobile buyers and brick and mortar businesses that have geographically situated retail offices. We focus on efficient mechanisms, infrastructure, and automation that can enable sellers and buyers to take joint advantage of the relationship of the locations of retail offices to the routes of mobile buyers who may have another primary destination. The methods promote automated vigilance about opportunities to buy and sell, and to support negotiations on the joint value to buyers and sellers including buyers' costs of divergence from their original paths to acquire services and commodities. We extend prior work on auction mechanisms to personal procurement settings by analyzing the dynamics of the cost to buyers based on preexisting plans, location, and overall context. We present mechanisms for auctions in single item, combinatorial, and multiattribute settings that take into consideration personal inconvenience costs within time-sensitive dynamic markets and challenges with privacy and fairness.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Multiagent systems*; J.4 [Social and Behavioral Sciences]: [Economics]; K.4.4 [Computers and Society]: Electronic Commerce—*Distributed commercial transactions*

General Terms

Design, Economics

Keywords

Electronic Markets, Economically-motivated Agents, Auction Design

Cite as: Mobile Opportunistic Commerce: Mechanisms, Architecture, and Application, Ece Kamar, Eric Horvitz, Chris Meek, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16, 2008, Estoril, Portugal, pp.1087-1094.

Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

1. INTRODUCTION

Increasing numbers of people in mobile settings have access to computing devices that are connected to the Internet. To date, such connectivity has been harnessed largely for personal communication, Web access, and routing. However, access to connected computing in mobile settings will likely evolve to play a more central role in peoples' lives. We focus on opportunities for extending market-centric mechanisms to commerce in mobile settings. In particular, we explore a potential future world where computing systems work to enhance the efficiency of planned and unplanned commerce between mobile buyers and businesses that are geographically situated. We focus on efficient mechanisms, infrastructure, and automation that can enable both sellers and buyers to take advantage of the relationship of locations of retail organizations to points on the routes of mobile buyers who may have preplanned primary destinations.

We shall introduce the mobile opportunistic commerce challenge and then present our work on developing market mechanisms to support opportunistic interactions between mobile buyers and geographically situated businesses. Our approach differs from the prior work on mobile opportunistic commerce by replacing single-agent decision making with market-centric machinery where mobile buyers and dynamic sellers are matched by an auction center. This work extends prior work on auctions to personal procurement settings that consider the dynamics of the cost to buyers based on preexisting plans, location, and overall context. We focus on the integration of an analysis of personal inconvenience costs within time-sensitive dynamic markets.

The core idea of this research is the pairing of buyers with sellers that offer the highest combined value in terms of the offered price, preferences of the buyer, and the required additional costs in time and distance to access the seller. Based on standing or dynamically computed buyer interests and needs, purchase requests from a buyer proxy are made available to sellers. A mediating system opens a reverse second-price sealed-bid auction with the set of feasible sellers, and constructs a value function for each seller by combining buyer preferences with time, distance costs, and seller bids. The transaction is finalized with the seller with the highest value on the second lowest offered cost. We consider single item, combinatorial, and multiattribute settings. We review problems of privacy and fairness with the approach. The mechanisms are explored via simulations within a prototype system named MC Market.

2. RELATED WORK ON MOBILE OPPORTUNISTIC COMMERCE

Several prior studies investigated the problem of opportunistic planning in mobile settings. In one research effort, a system was developed to assist people with cognitive deficits with transportation decisions by identifying opportunities to generate routes in line with goals [13]. Bohnenberger, et al. modeled user interests and created policies to maximize the utility of users in shopping domains [1]. Kwon, et al. developed a prototype, SmartGuide, operating as a location-based negotiation support system on mobile devices to guide users to find the best promotions within a region [10].

We build on prior efforts in mobile opportunistic commerce undertaken within the Mobile Commodities (MC) project. Research on MC defined the core mobile opportunistic commerce challenge, taking a decision-theoretic perspective on opportunistic transactions [7]. We shall review the MC work to set the context for the extensions of the prior single-agent decision making to market-centric machinery.

MC research has focused on the challenge of continuing to search for ways to satisfy a user's background goals in mobile settings, by inferring or accessing active goals, and then considering the costs and benefits of adding alternate goal-satisfying locations as waypoints to the user's route. Consider as a canonical example the opportunistic purchase of fuel for a car. The MC prototype allows users to specify preconditions for active goals. For the case of handling fuel needs, a user can specify that the goal of purchasing gasoline should become active when the quantity of available fuel dips below a threshold. When this goal (or another goal) is activated, the system accesses or infers [9] the destination and performs an ongoing search over all feasible service stations. Choices for candidate waypoints are ranked by expected value, taking into consideration the pricing at candidate locations, and the context-sensitive cost to users of the additional time and resources associated with diverting an ideal route to a destination through candidate waypoints. That is, the cost analysis considers not only the advertised price of the product, service or experience, but also the additional costs in time and distance to access the waypoints. In summary, the MC system follows five main steps: (1) identifying active goals, (2) inferring or accessing the destination of the user, (3) inferring context-sensitive marginal costs of time for the user, (4) performing geospatial search of feasible locations that satisfy user goals, and (5) executing a cost-benefit analysis of each of the options to identify the best candidate. Background goals and preconditions for activating goals are authored within a goals-and-preferences tool. The tool enables users to express policies for one-time goals and such recurrent needs as acquiring groceries, gasoline, meals, and haircuts. Time is an important resource and is one of the major factors influencing the value of different opportunities. A user may be willing to trade off increases in the distance and time added to a trip so as to achieve better prices when the cost of time is small, but not when the user has a near-term deadline that would make such a foray more costly.

MC employs a probabilistic model for the cost of time associated with diverging from an ideal path to a destination via the addition of candidate waypoints. The cost model considers as input the time of day, day of week, and sets of attributes about users' commitments drawn from an on-

line appointment book. The probabilistic model for the cost of time is learned from user annotated training data via a machine-learning procedure based on Bayesian structure search. A destination analyzer accesses or guesses the intended destination of a mobile user via direct input of the destination, through a predefined set of rules that consider routes by time of day and day of week, from information about appointments drawn from an online calendar, or via a probabilistic inference conditioned on a partial trajectory [9]. The system performs geospatial search over the feasible locations that can satisfy active goals. The set of candidate locations are accessed from the Microsoft Mappoint database. Given the current location and destination of the user, this subsystem computes updated routes to the destination by adding the candidate locations as waypoints and performing A* search. For practical reasons, the number of candidate locations is limited by the maximum distance the buyer is willing to diverge from the original path.

The opportunistic planner performs an economic analysis to evaluate the total cost of each option recognized by the geospatial search. The product costs are accessed from a pricing database. The planner computes a divergence cost for each possible option by combining the amount of time and distance added to the trip with the dynamic time-cost model of the user. The divergence cost is added to the product cost to infer the net cost of an opportunity. The planner offers the option with the lowest total cost to the buyer. The MC prototype notifies mobile travelers about the best opportunities and sends updated directions accompanied with summaries of the cost-benefit analysis.

We have extended the MC system to a market setting with mobile buyers and dynamic sellers matched by an impartial auction center. We shall describe how we harness and extend the buyer-side assessment tools from the MC prototype to model the buyer side of our market. We add a seller component enabling sellers to place bids dynamically via updates in pricing. We replace the decision-theoretic planner component of the MC system with an auction center that receives dynamic bids from the sellers and applies auction rules to determine the best opportunity for the buyer. The auction center makes use of the destination analysis tool and the geospatial search component of MC prototype to identify the set of feasible sellers and to evaluate the divergence costs accordingly. Our new model uses the decision-theoretic output of the MC prototype as a baseline to evaluate the value generated by the auction center.

3. MC MARKET ARCHITECTURE

MC Market extends MC by adding buyer and seller components and an *auction center* that plays the role of an unbiased and trusted mediator between the sellers and buyers of the market. The center aims to pair the buyer with the seller that offers the highest value in terms of the offered price, preferences, and inconvenience costs of the buyer.

We created three distinct agents within the buyer component: a buyer preference agent, a buyer proxy agent, and a buyer feedback agent. The *buyer preference agent* allows the buyer to represent goals and preconditions, to input preferences for different sellers and to set the cost of time before starting a trip. The accuracy and the quality of the information received by the agent affects the performance of the opportunistic planner.

The *buyer proxy agent* is responsible for assessing the cur-

rent state of the mobile buyer while the buyer is traveling. Using the set of goals and preconditions stored by the preference agent, the proxy agent autonomously recognizes the set of items, services, or experiences that the buyer is interested in purchasing. The proxy agent employs the destination analysis and time-cost models of the prior MC prototype to infer the destination and the cost of time for the buyer. For the current state of the buyer, the proxy agent constructs a time-cost function T to estimate the cost of time with respect to the nearest deadline. Given that the nearest deadline is d minutes later, the time cost per minute before the deadline is c_n and is c_d after the deadline, the function T predicts the cost of t minutes as,

$$T(t) = \begin{cases} tc_n & t < d \\ dc_n + (t - d)c_d & \text{otherwise.} \end{cases}$$

The *buyer feedback agent* is designed to capture the preferences of the buyer for different sellers. The agent collects a rating value and waiting time from the buyer immediately after an interaction with a seller. The rating value is given in terms of dollars and represents how much extra the buyer is willing to pay to have a transaction with that seller. The feedback information is used in evaluating the value of a seller for future opportunities.

The seller component is composed of two agents: the seller information agent and the seller proxy agent. The *seller information agent* lists the products, services, and experiences offered by a seller, and communicates seller-specific transaction rules or regulations (e.g., fees associated with cancellations) with the auction center. The *seller proxy agent* is invoked by the auction center, if that seller is selected as one of the feasible options by the geospatial search, to provide the bidding function of the seller to the auction center. In the current version of the MC Market system, seller proxy agents support one- or two-dimensional bidding functions to define pricing values for different quantities.

The *auction center* is connected to buyers and sellers through the proxy and information agents and stores information provided by these agents in its database. The auction center is invoked by a buyer's proxy agent to open an auction with feasible sellers that can satisfy the buyer's goal. The proxy agent defines the product or service that the buyer is interested in purchasing, and provides the destination and dynamic time-cost function of the buyer. Then, the auction center performs a geospatial search over waypoints, computes the overall value of taking different modified routes to the buyer's primary destination, and identifies a set of sellers to contact. The seller proxies are notified about the auction, and are asked to submit their bidding functions to the auction center. After the seller proxies submit their bids, the auction center constructs a value function for each seller by combining buyer preferences with time, distance costs, and seller bids. The center applies auction rules to determine the winner and the amount of payment, notifies the buyer about the deal, and provides updated driving directions.

We sought a market-centric design with the ability to enhance the efficiency of mobile commerce, given that buyers agree to share preferences, destinations, and time costs, and that sellers dynamically bid on prices. The market generates efficiency by pairing a buyer with the seller offering the highest combined value. Calculating the combined value requires knowing about buyers' private information (i.e., time cost, destination, preferences) and sellers' private bidding

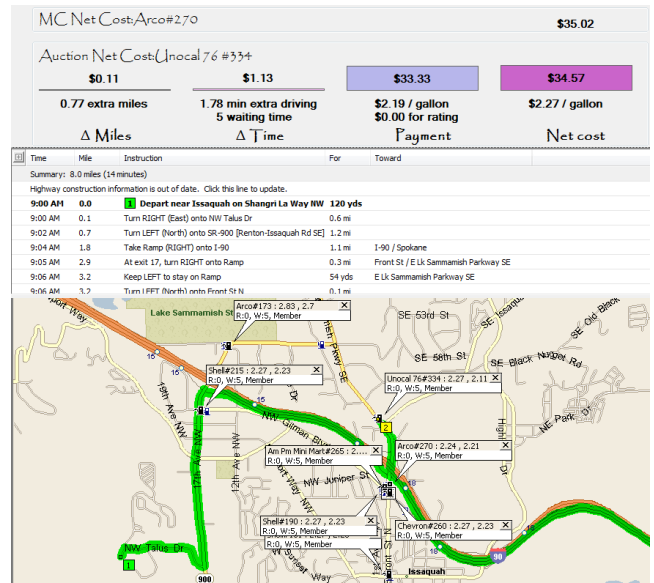


Figure 1: Screenshot of the MC Market prototype. Bottom: The auction center selects the set of feasible sellers, adds the winner as a waypoint to the original trip. Middle: Updated directions are displayed. Top: Economic summary of the outcome of the second price auction. The outcome is compared with the baseline value.

functions. In mobile commerce settings, protecting privacy is critical. The participants in the auction (i.e., buyers and sellers) should not be required to share their private information as the market is open to recurring transactions and private information can be used to strategize about future transactions. Therefore, we designed a trusted auction center to be the mediator between sellers and buyers without revealing both parties' private information. The center is specially designed to generate trust, to be fair, and to respect the privacy of sellers and buyers. The auction center only reveals necessary and sufficient information to the participants. During the auction process, the identity, preferences, destination, and time-cost of the buyer are not revealed to the sellers, and sellers are forced to bid on prices regardless of the identity of the buyer. Similarly, the bidding functions of sellers are not shared with the buyer or with other sellers. The center discloses the final revenue of the transaction to the winning seller and buyer. Other sellers are only notified that they are not the winners. Even though the winning seller and the buyer know the final revenue of the transaction, the nature of the auction mechanism makes it difficult to use this information to predict the bidding functions of sellers. As the center implements a second-price sealed-bid auction, the final revenue is the summation of the second lowest bidder's bid and the inconvenience cost for the buyer. The inconvenience cost is a combination of additional time and distance costs and buyer preferences about the seller. The identity of the second lowest bidder is concealed.

The center has rules to prevent buyers from deceiving the system to explore bidding functions of the sellers. A buyer gets to know the result of an auction when the transaction is completed. The buyer has to pay a cancellation fee to

terminate an agreement made by automated negotiations. There exists a lower bound on the amount of product for which a buyer can initiate an auction. To ensure the fairness of the auction center, the center is not awarded a portion of the transaction revenue, but instead receives a predetermined membership fee from the seller and buyer members. The strategy of the auction center is to keep the system as healthy and objective as possible so as to maximize the number of members and to preserve existing memberships.

The auction center uses the feedback mechanism to monitor the seller's behavior and to ensure that the sellers obey the terms of agreements and provide good quality of service. After each transaction, the buyer-feedback agent sends a rating and estimated waiting time for the recent transaction. As these values are used in winner determination, sellers with high waiting times or poor quality of service are punished by the system via increases in buyer inconvenience costs associated with these sellers.

We note that the winner-determination rules of the auction mechanism are designed to ensure individual rationality of sellers and buyers. The final value of the auction is never worse than the value that the buyer would receive if the buyer was not enrolled in the auction and was to purchase the product in a normal manner. Similarly, the second-price sealed-bid auction guarantees that if the seller bids the true evaluation, the utility of the seller is always non-negative.

4. AUCTION MODEL FOR MOBILE OPPORTUNISTIC COMMERCE

As we have discussed, for personal procurement, price is not the only factor that determines the value of a deal to the buyer; within opportunistic settings, the additional time and travel required can play an important role. A buyer with a low cost of time (e.g., one who does not have a looming deadline) may prefer traveling a longer way to obtain a good price, whereas a buyer in a high cost-of-time context may prefer a more expensive but nearer seller. We define an inconvenience cost model that considers several buyer-specific factors and then modify well-known auction models to take buyer inconvenience costs into account.

4.1 Model for Cost of Inconvenience

Adding waypoints to a trip introduces extra time and transportation costs for the buyer and these inconvenience costs affect the value that the buyer receives from a transaction. We combine the extra time and fuel cost with buyer preferences to estimate the total inconvenience cost associated with adding a candidate seller as a waypoint on a trip to a primary destination. The probabilistic time-cost model is provided by the buyer proxy agent, the auction center uses Microsoft Mappoint services to estimate the travel duration.

Given that Δt is the estimated travel duration of the original route, the center computes the updated travel duration ($\Delta t_i + w_i$) by adding seller s_i as a waypoint and including the estimated waiting time w_i . We populate the time-cost model with Δt and $(\Delta t_i + w_i)$ values and use the difference of the two estimations to predict the time cost for adding s_i to the trip. Similarly, the distance cost is the divergence in miles for adding s_i as a waypoint. We use Δd to refer to the estimated travel distance of the original route and Δd_i to define the distance to be traveled with a stop at s_i . The fuel cost in dollars for traveling one mile is represented as c_g . Rating

r_i of seller s_i represents how much extra the buyer is willing to pay in dollars to make the deal with s_i . This value is an indicator of the quality of service the buyer receives from that particular seller. We limit our attention to the sellers that are close enough to be considered as feasible sellers. The maximum distance the user is willing to travel or can travel (due to the amount of fuel remaining) is represented by d_{max} , and the inconvenience cost becomes infinite if the total trip distance exceeds the maximum allowed distance. Combining time and distance costs, and rating values, we define personal inconvenience cost function for seller s_i as,

$$PC_i = \begin{cases} \infty & \Delta d_i > d_{max} \\ T(\Delta t_i + w_i) - T(\Delta t) - r_i \\ \quad + (\Delta d_i - \Delta d)c_g & \text{otherwise.} \end{cases}$$

4.2 Selecting Auction Type

The rules of the auction determine the winner and final revenue, and influence the behaviors of the participants. In this section, we investigate four well-studied auction types and evaluate them for their applicability to mobile opportunistic commerce.

English (descending) and Dutch auctions are iterative auctions in which bidders receive price signals. Bidding in an English procurement auction starts from a high price point, and bidders offer monotonically decreasing bids until one seller remains. The transaction is sealed with the seller offering the lowest price. The Dutch procurement auction is the opposite of the English auction in the sense that bidding starts from a low price point and increases monotonically. The deal is made with the first seller that accepts the announced price. The optimal strategy in English auctions is to bid on prices above the true evaluation. The Dutch auction is strategically more challenging for bidders, as the bidders are required to act without receiving pricing signals. First-price sealed-bid and second-price sealed-bid (Vickrey) procurement auctions are single round auctions in which bidders submit sealed bids. In both auction mechanisms, the winner is the bidder with the lowest price offer. The payment of the first-price sealed-bid auction is the lowest offered price whereas the winner is paid the second lowest offer in second-price sealed-bid auction.

The Revenue Equivalence Theorem states that, although these auction mechanisms are significantly diverse, the mechanisms generate the same revenue under the assumptions that the bidders are risk-neutral, that bidder evaluations are independent, and that bidders are symmetric [11]. It is shown that second-price sealed-bid auctions are strategy proof for one-time interactions of buyers and sellers [16]. The payment the winner receives is independent of its bid. Revealing the correct evaluation maximizes the chance of winning. The dominant strategy is being truthful and revealing the true evaluation function. On the contrary, first-price sealed-bid auctions are not strategy-proof mechanisms. The bidders participating in the auction need to strategize about the other sellers to maximize their final revenue.

In mobile opportunistic commerce settings, bidders are such sellers as fuel stations, grocery stores, and movie theaters; these entities typically do not have expertise in bidding strategies. Generating nearly optimal bids in Dutch or in first-price sealed-bid auctions may be a potential barrier that may make them cautious about entering the market. For one-time situations, the English and second-price auc-

tions eliminate the expensive overhead of strategizing about other bidders. Thus, it is easier for brick and mortar retail establishments to engage in the electronic market.

The sellers are evaluated in terms of total cost, computed as a combination of the inconvenience cost and the bid that the seller offers. The winner of an English procurement auction is the seller offering the lowest total cost. A seller maximizes the chance of winning by minimizing its inconvenience cost. The inconvenience cost is minimized if the auction is terminated at the departure point. Then, the buyer can follow the path that induces the lowest divergence cost when the seller is added as a waypoint. In one-time situations, the dominant strategy for English auction is to bid up to their true evaluation (true bidding function combined with minimum inconvenience cost) and to terminate the auction before the buyer starts the trip. In one-time personal procurement auctions without instant price fluctuations, the English auction produces an identical outcome as the second-price sealed-bid auction.

English auctions have high communication requirements. At every decrement, one seller bids, and all sellers are notified about the new bid. When the auction terminates, all participants are informed of the outcome of the auction. Given a starting price of the auction of p_{high} , a price dropping r at each decrement, a final price of p_{final} , and with n sellers bidding, the number of messages transmitted is, $(n+1)\frac{p_{high} - p_{final}}{r} + (n+1)$. The communication requirements of second-price sealed-bid auctions are significantly lower in comparison to the English auction. As the auction progresses, n bidders send their bids to the auction center, and the center notifies the participants. The total number of messages sent is $2n+1$. Communication costs are especially important in the mobile settings that we consider.

We prefer to implement the second-price sealed-bid auction in our market-centric system because of its three important properties for one-time auctions: efficiency of final outcome, strategy-proofness and truthfulness, and low communication demands. The market mechanism may be affected by the vulnerabilities of the second-price sealed-bid auctions; it is known that the second-price auctions are open to collusion of sellers and that revealing true evaluations may not be the dominant strategy when extended to repeated auctions. We are seeking to enhance the design with new techniques to provide truthfulness in repeated interactions, and to prevent collusion [8, 15, 12].

4.3 Auction Model for Single-item Procurement

Previous work on procurement auctions has introduced models that combine supply-chain costs with prices to minimize the total cost. These models focus on industrial procurements in which important attributes are delivery time, availability of spares, maintenance, etc. [2, 3]. We extend the prior work on auctions to personal procurement settings and present an auction model that applies the rules of the second-price sealed-bid auction to mobile opportunistic commerce. Our multiattribute auction model considers both the bidding prices of the sellers and the costs associated with the personal inconvenience to buyers. The auction mechanism matches the buyer with the seller that offers the highest value. The mechanism always generates non-negative profit to sellers. It also guarantees to achieve or improve the baseline outcome, which is the best deal achieved among the

standard prices announced in the open market (outside of the auction). The mechanism realizes a volume discount auction, receives bids as a function of quantity, and computes the quantity and price that maximizes the comprehensive value to the buyer [14]. The system implements a second-price sealed-bid auction with sellers who are members of the MC Market.

The winner determination problem for general auction settings is NP-hard [5]. The personalized auction model makes two realistic assumptions that are justified by the personal procurement domain so that it can determine the winner of an auction efficiently. We assume each supplier has an infinite supply of goods and expect the per item price to drop with increasing quantities. To minimize the total cost, buyers purchase the whole quantity of an item from a single seller. Under these assumptions, we rewrite the winner determination calculations of the second-price sealed-bid auction.

Let us consider a set of all feasible sellers, S , identified by the geospatial search component. The sellers are evaluated in terms of the total cost (TC) value they offer. The best deal available to the buyer without entering into the auction is used as the baseline value to evaluate the profit of the auction. TC_B represents the total cost of the baseline deal where p_i is the announced price of seller $s_i \in S$, q is the quantity of purchase and q_{max} is the maximum amount the buyer is willing to purchase. TC_B function is upper bounded by its value at q_{max} . We apply a decision-theoretic analysis to find the seller $w_B(q)$ that offers the minimal total cost to the buyer for quantity q in the traditional market setting without an auction.

$$TC_B(q) = \begin{cases} \min_{1 \leq i \leq |S|} \{PC_i + p_i q\} & q \leq q_{max} \\ TC_B(q_{max}) & otherwise \end{cases}$$

$$w_B(q) = \operatorname{argmin}_{1 \leq i \leq |S|} \{PC_i + p_i q\}$$

The auction center receives bidding functions b_j from each seller $m_j \in M$, where M is the set of member sellers. b_j is a function of the offered price with respect to quantity q . Given the inconvenience cost of seller m_j is PC_j , the total cost (TC_j) of m_j is calculated as,

$$TC_j(q) = PC_j + b_j(q)q.$$

The value of having seller $m_j \in M$ in the auction is modeled by the value function V_j . V_j is the total cost reduction gained by preferring m_j to w_B .

$$V_j(q) = TC_B(q) - PC_j - b_j(q)q$$

For the fuel purchase domain, the amount of fuel (q_j) that the buyer spends to get to the location of seller m_j changes the amount of fuel that needs to be purchased. The cost of gasoline is already included in the PC function. The extra demand, q_j , effects the bidding price of m_j , and is included in the TC_j and V_j functions,

$$TC_j(q) = PC_j + b_j(q + q_j)q$$

$$V_j(q) = TC_B(q) - PC_j - b_j(q + q_j)q.$$

In second-price auctions, the outcome of the auction is the second highest value provided by the sellers. The auction center constructs an auction outcome function (AO) that represents the value generated by having the auction for quantity q . The center concludes by identifying the quantity

q^* and the seller w_1 that maximizes the value to the buyer. The winner of the auction (w_1) is the seller that offers the best value for quantity q^* . The payment is determined with respect to the seller w_2 that offers the second highest value for quantity q^* .

$$\begin{aligned} V(q) &= \{V_1(q), \dots, V_j(q), \dots\}, 1 \leq j \leq |M| \\ AO(q) &= \max\{V_j(q) : V_j(q) < \max(V(q))\} \\ q^* &= \operatorname{argmax}_{q \leq q_{max}} \{AO(q)\} \\ w_1 &= \operatorname{argmax}_{1 \leq j \leq |M|} \{V_j(q^*)\} \end{aligned}$$

The system ensures that the buyer is not worse off by engaging in the auction. The center compares the final value of the auction with the baseline value to determine if the auction is profitable. If $AO(q^*)$ is positive, the auction is more profitable than the baseline outcome. w_1 is selected as the winner, with a total cost for the buyer of $(TC_{w_2}(q^*))$, and a payment to the winner of $(TC_{w_2}(q^*) - PC_{w_1})$. Otherwise, the auction center recommends that the buyer have the transaction with w_B by paying $(TC_B - PC_{w_B})$.

Assuming that the quantity range the buyer is interested in is discretized into $|Q|$ intervals, the size of the search space of the winner determination method is polynomial in $|Q|$ and the number of feasible sellers.

4.3.1 Example: Fuel Purchase

We shall present an example to illustrate the winner determination calculations in the single-item version of MC Market. The buyer proxy identifies a fuel need and invokes the auction center by delivering the original route and the time cost function. MC Market finds four feasible sellers. Sellers 0, 1, and 3 are located close to the original path and they yield lower inconvenience costs but higher prices. Seller 2 is more competitive; although located farther away from the destination, the seller bids the lowest price. Sellers 1, 2, and 3 are members of the MC Market system.

	Member	p_i	b_i	PC_i	d_i
Seller 0	No	3.0	10	2	
Seller 1	Yes	2.9	C_1	10	2
Seller 2	Yes	2.6	C_2	20	5
Seller 3	Yes	2.91	C_3	13	2

The baseline function is constructed using the announced prices, p_i . The sellers that are members of the market, submit their bidding functions C_1, C_2, C_3 (see Figure 2).




Figure 2: Bidding functions of s_1, s_2, s_3 .

Value functions are computed for the member sellers. The quantities, for which the auction outcome yields higher val-

ues than the baseline, are labeled as the Profitable Auction Region in gray (see Figure 3). Figure 4 shows how the prof-




Figure 3: Value functions for sellers 1, 2, and 3 and the baseline.

itable auction region changes when Seller 0 (s_0) reduces its announced price to \$2.8.




Figure 4: The Profitable Auction Region grows smaller because of higher baseline values.

4.4 Personalized Auction Model for Combinatorial Procurement

We believe that MC Market offers new opportunities to buyers and sellers in combinatorial purchases in which multiple items are exchanged in bundles. If the buyer proxy determines that the buyer is interested in a set of items, the center aims to find the combination of sellers that can provide the best value to the buyer.

Given that the set of items the buyer is interested in is $X = \{x_1, x_2, \dots, x_n\}$, and S_i is the set of feasible member sellers that supply item x_i , the auction center computes personal and total cost functions for every combination $c_j = \{s_1, s_2, \dots, s_n\}$ such that $s_1 \in S_1, \dots, s_n \in S_n$. The set of all possible combinations is C . The center considers every seller in c_j as a waypoint added to the trip to the primary destination and estimates the expected travel duration (Δt_j) and distance (Δd_j) for the updated route. The inconvenience cost of c_j is the combination of extra time and distance costs combined with additive waiting times ($\sum w_i$) and preferences ($\sum r_i$).

$$PC_{c_j} = \begin{cases} \infty & \Delta d_j > d_{max} \\ T(\Delta t_j + \sum_{s_i \in c_j} w_i) - T(\Delta t) & \text{otherwise} \\ -(\sum_{s_i \in c_j} r_i) + (\Delta d_j - \Delta d)c_g & \end{cases}$$

$Q = \{q_1, q_2, \dots, q_n\}$ holds the target quantity for each item buyer needs and $P = \{p_{j1}, p_{j2}, \dots, p_{jn}\}$ keeps the corresponding bids from the sellers in c_j . The Total Cost value for the

combination c_j is calculated as,

$$TC_{c_j} = PC_{c_j} + \sum_{s_i \in c_j} p_{j_i} q_i.$$

We may have sellers in the market that supply a bundle of products. These sellers can take special advantage of the MC Market setting as they require fewer numbers of stops, thus less potential divergence from the original path, and lower costs of personal inconvenience.

We apply the VCG mechanism to determine the payment each seller receives. The VCG mechanism is the generalized version of the Vickrey auctions applied to combinatorial settings and this approach is preferred as it is strategy-proof and efficient [4, 16, 6]. c^* is the combination of sellers that satisfy the buyers combinatorial demands with minimal total cost, c_{-i}^* is the best allocation possible excluding seller s_i from the auction. Payment t_i for $s_i \in c^*$ bidding p_i on quantity q_i is calculated as,

$$\begin{aligned} c^* &= \operatorname{argmin}_{c_j \in C} TC_{c_j} \\ t_i &= TC_{c_{-i}^*} - TC_{c^*} + p_i q_i \\ c_{-i}^* &= \operatorname{argmin}_{c_j \in C, s_i \notin c_j} TC_{c_j}. \end{aligned}$$

Finding the optimal allocation in combinatorial auctions is known to be NP-hard [14]. Given that the buyer is interested in n item bundles and that the system has maximum m members selling each product, the size of the search space is $n! m^n$. However, in mobile opportunistic commerce settings, the buyer has limited time to spend for personal procurement on a single trip. Therefore, we can assume that the number of products (n) that the buyer is interested in purchasing during a single trip is typically bounded by a small number and the search space is polynomial in m .

4.4.1 Example: Bundling Grocery and Fuel

We describe a sample scenario in which a buyer needs to buy both grocery and gasoline during a single trip. The auction center identifies two feasible fuel stations, S_1, S_2 and two grocery stores, G_1, G_2 , and forms four possible combinations. For each of the combinations, the auction center constructs inconvenience cost functions and combines these with seller bids to determine the total costs. Fuel station S_1 and grocery store G_2 offers the lowest total cost.

Combination	Gas Price	Grocery Price	PC	TC
$S_1 \& G_1$	30	13	8.5	51.5
$S_1 \& G_2$	30	10	6.8	46.8
$S_2 \& G_1$	32	13	9.8	54.8
$S_2 \& G_2$	32	10	7.2	49.2

The payments to the winner sellers are calculated as,

$$\begin{aligned} t_{S_1} &= TC_{c_{-S_1}^*} - b_{G_2} - PC_{S_1 \& G_2} = 32.4 \\ t_{G_2} &= TC_{c_{-G_2}^*} - b_{S_1} - PC_{S_1 \& G_2} = 14.7. \end{aligned}$$

4.5 Personalized Auction Model for the Procurement of Multiattribute Items

Multiattribute auctions allow sellers to compete not only on the price dimension but also on the attributes of a product. In this section, we extend our market design and mechanisms to work with an ontology of products in the market, to allow buyers to evaluate brands and attributes of a product, and to let sellers enter individual bidding functions for

product brands and types. We adjust Vickrey auction calculations to the multiattribute personal procurement setting.

The auction center provides an ontology of products available in the market. This ontology is used to identify brands and attributes to be considered in preference elicitation, bid collection, and winner determination. The center creates a hierarchy tree for each product X . Every leaf of the tree, x_i , is a product type that is a distinct combination of attributes of the product. We write $X = \{x_1, x_2, \dots, x_n\}$ to denote the types of products that the buyer can choose from. A product type is defined by the tuple $x_i = \{Y_{x_i}, A_{x_i}\}$ where Y_{x_i} is the set of brands for that product type and A_{x_i} is the set of attributes defining x_i . A_{x_i} is the set of nodes that are traversed from the root of the product tree to get to x_i .

We define $A_X = \bigcup_{x_i \in X} A_{x_i}$ as the set of all attributes for product X and $Y_X = \bigcup_{x_i \in X} Y_{x_i}$ as the set of all brands offering X . The buyer preference agent is modified to provide buyer utility values for every element of A_X and Y_X . These values represent how much the buyer is willing to pay to have a product with that brand or that attribute. The seller proxies provide bidding functions in the form of $b_{s_j}(x_i, y_{x_i})$, for brand y_{x_i} of product type x_i .

The valuation of a buyer for product type x_i of brand y_{x_i} is modeled with an additive utility function. We assume that all attributes are known by the system and the attributes are independent. Let U_A and U_B be the utility functions for attributes and brands respectively, the buyer valuation of product type x_i of brand y_{x_i} , $U(x_i, y_{x_i})$, is calculated as,

$$U(x_i, y_{x_i}) = U_B(y_{x_i}) + \sum_{a_k \in A_{x_i}} U_A(a_k).$$

The value of a seller s_j for product X , $V(s_j, X)$, is the maximum value obtained by the seller for any $\{x_i, y_{x_i}\}$ combination. $\{x_{s_j}^*, y_{s_j}^*\}$ is the highest valued product type and brand combination offered by s_j .

$$\begin{aligned} V(s_j, X) &= \max_{x_i \in X, y_{x_i} \in Y_{x_i}} \{U(x_i, y_{x_i}) - b_{s_j}(x_i, y_{x_i})\} - PC_{s_j} \\ \{x_{s_j}^*, y_{s_j}^*\} &= \operatorname{argmax}_{x_i \in X, y_{x_i} \in Y_{x_i}} \{U(x_i, y_{x_i}) - b_{s_j}(x_i, y_{x_i})\} - PC_{s_j} \end{aligned}$$

The winner of the auction, w_1 , is the seller offering the highest value. As we are implementing a Vickrey Auction, the payment, t_{w_1} depends on the value of the seller with the second highest value, w_2 . As t_{w_1} is independent of w_1 's bidding, the dominant strategy for sellers is to bid truthfully.

$$\begin{aligned} w_1 &= \operatorname{argmax}_{s_j} \{V(s_j)\} \\ t_{w_1} &= U(x_{w_1}^*, y_{w_1}^*) - V(w_2) - PC_{w_1} \end{aligned}$$

The multiattribute personal procurement setting is more challenging for both buyers and sellers. In this setting, we ask the buyer to give a value for every attribute of products; similarly the sellers provide bidding functions for every combination of product types and brands. These tasks are obviously difficult for all. We believe that research on preference elicitation and on interface design will lead to methods that can ease the effort required for these specifications.

4.5.1 Example: Multiattribute Personal Auction for a Carwash

As an illustrative example, let us explore a multiattribute

auction for purchasing a car wash. The auction center generates an ontology for car wash services (see Figure 5).




Figure 5: Product ontology for car wash.

The buyer enters preferences for every car wash attribute.

	Buyer Preferences in \$
Exterior	5
Interior	3
Exterior and Interior	8
Hand Wash	2
Mechanic Wash	0

The center finds three feasible sellers and asks them for their bids for each x_i . For each seller, the center estimates the buyer's inconvenience cost.

	b_{x_1}	b_{x_2}	b_{x_3}	b_{x_4}	b_{x_5}	PC
Seller 1	4	3	2	5	4	1
Seller 2	5	3	2	6	4	2
Seller 3	6	4	3	9	5	2

The system calculates the utility of the buyer for every x_i and the value of each seller to the buyer.

$U(x_1)$	$U(x_2)$	$U(x_3)$	$U(x_4)$	$U(x_5)$
7	5	3	10	8

$$\begin{aligned}
 V(\text{Seller 1}) &= \max\{3, 2, 1, 5, 4\} - 1 = 4, \quad x_{\text{Seller 1}}^* = \{x_4\} \\
 V(\text{Seller 2}) &= \max\{2, 2, 1, 4, 4\} - 2 = 2, \quad x_{\text{Seller 2}}^* = \{x_4, x_5\} \\
 V(\text{Seller 3}) &= \max\{1, 1, 0, 1, 3\} - 2 = 1, \quad x_{\text{Seller 3}}^* = \{x_5\}
 \end{aligned}$$

The auction center selects item x_4 to purchase and Seller 1 as the winner of the auction. The payment is calculated as,

$$t_{\text{Seller 1}} = U(x_4) - V(\text{Seller 2}) - PC_{\text{Seller 1}} = 7.$$

5. CONCLUSION AND FUTURE WORK

We have focused on identifying opportunities and challenges for developing a truthful, privacy-preserving, trustworthy, and unbiased market mechanism that brings mobile buyers and sellers together for personal procurement in opportunistic settings. We extended the MC prototype and overall architecture to employ market-centric concepts. We presented methods and models used in the MC Market system and described its key components and extensions. We ran simulations of auction calculations by specifying pricing functions for sellers and constructing personal inconvenience cost functions for buyers. For future work, we are exploring five main directions: (1) enhancing means for eliciting preferences from buyers, particularly for multi-attribute items, (2) improving the current models to guarantee buyer truthfulness, (3) extending the market design to

double actions, (4) achieving a strategy-proofness property in repeated transactions of MC market, and (5) applying more comprehensive cost-benefit analyses of opportunities that recognize promotions and daily pricing patterns. We hope one day to deploy the MC Market system to early adopters and test the system in real-world settings. We have highlighted several key concepts and directions in research on markets for mobile opportunistic commerce. We look forward to theoretical extensions and to practical experiences with these methodologies.

6. ACKNOWLEDGMENTS

We thank Paul Koch for his software-engineering support on the MC Market prototype and extensions.

7. REFERENCES

- [1] T. Bohnenberger, A. Jameson, A. Kruger, and A. Butz. Location-Aware Shopping Assistance: Evaluation of a Decision-Theoretic Approach. *Proceedings of Mobile HCI*, 2002.
- [2] Y. Che. Design Competition Through Multidimensional Auctions. *The RAND Journal of Economics*, 1993.
- [3] R. Chen, R. Roundy, R. Zhang, and G. Janakiraman. Efficient Auction Mechanisms for Supply Chain Procurement. *Management Science*, 2005.
- [4] E. Clarke. Multipart pricing of public goods. *Public Choice*, 1971.
- [5] A. Davenport and J. Kalagnanam. Price negotiations for procurement of direct inputs. *IBM Research Report RC*, 22078:2001, 2001.
- [6] T. Groves. Incentives in Teams. *Econometrica*, 1973.
- [7] E. Horvitz, P. Koch, and M. Subramani. Mobile Opportunistic Planning: Methods and Models. *Lecture Notes in Computer Science*, 2007.
- [8] P. Klemperer. What Really Matters in Auction Design. *The Journal of Economic Perspectives*, 2002.
- [9] J. Krumm and E. Horvitz. Predestination: Inferring Destinations from Partial Trajectories. *International Conference on Ubiquitous Computing*, 2006.
- [10] O. Kwon, J. Shin, and S. Kim. Context-aware multi-agent approach to pervasive negotiation support systems. *Expert Systems With Applications*, 2006.
- [11] R. McAfee and J. McMillan. Auctions and Bidding. *Journal of Economic Literature*, 1987.
- [12] D. C. Parkes. Online mechanisms. In *Algorithmic Game Theory*. 2007.
- [13] D. Patterson, L. Liao, K. Gajos, M. Collier, N. Livic, K. Olson, S. Wang, D. Fox, and H. Kautz. Opportunity Knocks: a System to Provide Cognitive Assistance with Transportation Services. *Proceedings of UBICOMP*, 2004.
- [14] M. Rothkopf, A. Pekec, and R. Harstad. Computationally Manageable Combinational Auctions. *Management Science*, 1998.
- [15] M. Rothkopf, T. Teisberg, and E. Kahn. Why Are Vickrey Auctions Rare? *The Journal of Political Economy*, 1990.
- [16] W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *The Journal of Finance*, 1961.