Estimating the Capacity of the Location–Based Advertising Channel

Győző Gidófalvi, Hans Ravnkjær Larsen
Geomatic ApS - Center for Geoinformatic
{gyg, hrl}@geomatic.dk

Torben Bach Pedersen
Aalborg University
tbp@cs.aau.dk

Abstract

*Delivering "*relevant*" advertisements to consumers carrying mobile devices is regarded by many as one of the most promising mobile business opportunities. The relevance of a mobile ad depends on at least two factors: (1) the* proximity *of the mobile consumer to the product or service being advertised, and (2) the match between the product or service and the* interest *of the mobile consumer. The interest of the mobile consumer can be either* explicit *(expressed by the mobile consumer) or* implicit *(inferred from user characteristics). This paper tries to empirically estimate the capacity of the mobile advertising channel, i.e., the number of relevant ads that can be delivered to mobile consumers. The estimations are based on a simulated mobile consumer population and simulated mobile ads. Both of the simulated data sets are realistic and derived based on real world data sources about population geo–demographics, businesses offering products or services, and related consumer surveys. The estimations take into consideration both the proximity and interest requirements of mobile ads, i.e., ads are only delivered to mobile consumers that are close–by and are interested, where interest is either explicit or implicit. Results show that the capacity of the LBA channel is rather large, which is evidence for a strong business case, but also indicate the need for user–control of the received mobile ads.*

1. Introduction

Mobile or Location–Based Advertising (MA or LBA), i.e., sending electronic advertisements to consumers carrying mobile devices, is considered by many as one of the most promising business opportunities amongst Location–Based Services (LBS) [14]. A recent mobile marketing survey suggests that about 7% of the mobile consumers would be willing to receive promotional text messages "if they were relevant" [6]. According to other surveys, an even larger percentage of the mobile consumers are interested if they are rewarded in some way [15]. In this paper, mobile ads are regarded as a means of presenting relevant information to a recipient, be it a commercial offer on an item on sale, traffic information, or a piece of public information. To many people, the world seems to be more and more difficult to guide oneself through, thus the art of targeting information and services will prove to be of immense value. Only, efficient business cases have so far been very few, in spite of the market's expectations.

A broad range of aspects, or variables, determine the relevance and context of a mobile ad: distance to the mobile user, explicit or implicit interest of the mobile user, uniqueness (do not send ad twice within some interval), time and place of delivery, etc. To this extent, this paper describes an LBA framework and a LBA database that can be used for the management of mobile advertisements.

In lack of comprehensive, real data on the movements and behavior of the population, estimation or simulation is extremely useful, bringing the models to life with real and well–documented consumption patterns. Using a simulated but realistic mobile consumer population and a set of mobile ads, the LBA database is used to estimate the capacity of the mobile advertising channel, i.e., the number of relevant ads that can be delivered to mobile consumers. Apart from this use, the LBA database and the estimates derived from it can also be used in mobile catchment area analysis to estimate business exposure. Results show that the capacity of the LBA channel is rather large (approx. 100 mobile ads per user within a single day), giving strong support for a business case. The same results can also be viewed as alarming, and indicate the need to incorporate user–control of the received mobile ads in the LBA framework, as suggested by the Mobile Marketing Association [16].

The remainder of this paper is organized as follows. Section 2 reviews related work. Section 3 defines the estimation problem both in case of explicit and implicit interest. Section 4 describes the simulated data sets and their derivations from real–world data sources. Section 5 describes the method and technical foundations for delivering mobile ads while taking into account both the advertisers' and mobile users' interests. Section 6 describes the experiments and discusses the estimates resulting from them. Finally, Section 7 concludes and points to future research directions.

2. Related work

The estimations in this work are based on simulated movements of mobile users. Movements of users are influenced by physical, social and geo–demographical aspects of mobility.

To aid the development in mobile data management, a number of moving object simulators have been proposed in the literature that model primarily the physical aspects of mobility to various extents. Since most objects use a network to get from one location to the other, a state–of–the–art framework for network–based moving object simulation is presented in [1]. The behavior of a moving object in this framework is influenced by (1) the attributes of the object having a particular object class, (2) the combined effects of the locations of other objects and the network capacity, and (3) the location of external objects that are independent of the network. For a review of other moving object simulators the reader is referred to [8].

Moving object simulators generally neglect the social and geo–demographical aspects of mobility. These social and geo–demographical aspects of mobility introduce patterns in the movement of users and give rise to a unique spatio–temporal (ST) distribution of users. ST–ACTS is a Spatio–Temporal ACTivity Simulator that using real–world data sources models some of these neglected aspects of mobility [8]. To make the estimations in this work as realistic as possible, movements of mobile users are obtained from ST–ACTS, which is further described in Section 4.4.

Database indices allow the effective management and retrieval of data in large databases. Spatial and geographical databases manage information about spatial objects, i.e., objects that have physical properties such as location and extent. An R–tree is a widely used index structure that allows the effective management and retrieval of spatial objects [9]. An R–tree splits space with hierarchically nested, and possibly overlapping Minimum Bounding Rectangles (MBRs). Search algorithms that test spatial relationships (for example; intersection, containment, nearest) between spatial objects can effectively use the MBRs to decide whether or not objects in two MBRs satisfy a specific spatial search criterion.

The location of a moving object changes over time. Hence, the movement of a moving object is commonly described as a sequence of coordinate and timestamp pairs and is referred to as the trajectory of the moving object. Moving objects databases are databases that represent and manage changes related to the movement of objects. Spatio–temporal indices such as the Spatio–Temporal R–tree (STR–tree) and Trajectory–Bundle tree (TB–tree) allow the effective management and retrieval of information about moving objects [11]. An STR–tree organizes line segments of a trajectory according to both their spatial properties and the trajectories they belong to, while a TB–tree only preserves trajectories. While these spatio–temporal indices are designed to effectively manage trajectories, they are not available in commercially available Relational Database Management Systems (RDBMSs). Hence, the herein presented method uses the widely available R–trees.

Time geography [10] is a conceptual basis / paradigm for human space–time behavior which considers (1) the indivisibility or corporeality of the human condition; (2) that humans typically operate over finite intervals of space and time; (3) the natural laws and social conventions that partially constrain space–time behavior; and (4) that humans are purposive. The movements of mobile users used in the estimations are derived from ST–ACTS [8], which models some aspects of this paradigm.

Research has shown that LBSes have not yet been as widely used as expected [12]. The present paper can be an aid in furthering the spread of LBSes.

Recent research shows that in order to succeed with location–based ads, content is imperative [13]. In the present paper, by targeting content to the right recipients, we support this content aspect.

Studies have shown that mobile users are eager to make use of their phones in new ways and proposed methodologies to model user willingness [4, 15, 18]. Thus, our setup is realistic for the typical future user.

3. Problem statement

Let $A = \{a_1, \dots a_n\}$ be the set of ads. Each ad a has a location $adloc(a)$ and is for a certain product $prod(a)$. Let $U = \{u_1, \dots u_n\}$ be the set of (moving) users. Each user u has a location $uloc(u.t)$ depending on the time t, an explicit interest profile $expint(u)$ containing a set of products, and an implicit interest profile $impint(u)$, containing a set of demographic variable values. We also assume a scoring function $score(u, a)$ that given a user u and an ad a returns a value between 0 (no match) and 1 (perfect match) that predicts how interested user u is in product $prod(a)$ based on the values in $impint(u)$.

Given a maximum distance *maxdist* between user locations and ad locations, and a timespan $T = [t_{start} : t_{end}]$, the *explicit location–based ad delivery estimation problem* is to estimate how many times a user u has a location $uloc(u, t)$ within *maxdist* from $adloc(a)$ for an ad a in A and a time t in T where $prod(a) \in expint(u)$.

Given a *maxdist* between user locations and ad locations, a *minscore*, and a timespan $T = [t_{start} : t_{end}]$, the *implicit location–based ad delivery estimation problem* is to estimate how many times a user u has a location $uloc(u, t)$ within *maxdist* from $adloc(a)$ for an ad a in A and a time t in T where $score(u, a) \geq minscore$.

Table 1. Variables in conzoom®.

referred entity	conzoom® variable	categories
person	person count	1
	age	9
	education type	9
	employment status type	12
	employment branch type	12
housing unit	unit count	1
	house type	6
	house ownership type	4
	house area	5
household	household count	1
	family type	5
	fortune	6
	personal income	5

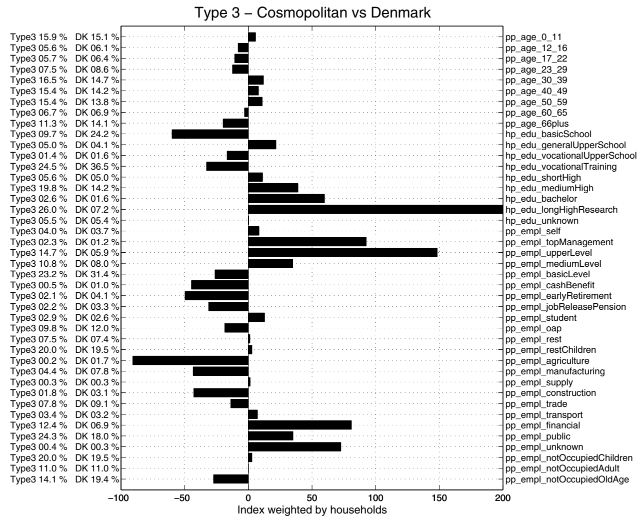

Figure 1. Partial profile of conzoom® type 3.

4. Data

The estimations stated in Section 3 are based on a number of real–world data sources. The use of real–world data sources is important to derive realistic estimates. While the data sources refer to the Danish market and population, similar data sources are available for other major markets [2, 3]. The following subsections describe in detail the data sources used to derive the estimates.

4.1. conzoom® demographic data

conzoom® is a commercial database product that contains fine–grained, geo–demographic information about Denmark's population [6]. The variables that describe the statistical characteristics of the population can be divided into three groups: person, housing unit, and household variables. These variables and the number of categories for each are shown in Table 1.

In Table 1, variables that have "type" in their names are categorical variables; variables that have "count" in their name are counts of the corresponding entities within a 100–meter grid cell; and finally, the rest of the variables are continuous variables that have been categorized into categories that are meaningful for market segmentation. Since, for example in the countryside, the number of persons, households or units could be very low in a 100–meter grid cell, grid cells are grouped together into meaningful, large enough clusters to comply with social and ethical norms and preserve the privacy of individuals. The basis for clustering is twofold: geography and the publicly available one–to–one housing information. The intuition behind the basis is also twofold. First, people living in a given geographical region (be that a state, a county, or a postal district) are similar in some sense; for example, they might be more likely to have a certain political orientation than people living in another geographical region. Second, people living in similar houses are likely to be similar in other demographic variables; for example an established family with a stable source of income is more likely to be able to buy a larger, more expensive house than a person who just started his/her career.

As mentioned earlier, to preserve the privacy of individuals, the clusters are constrained to contain at least some fixed number of households. Statistics for the variables, depending on the sensitivity of the information contained in them, are obtained from Statistics Denmark [19] for clusters constructed at an appropriate level of cluster size constraint, for example 20, 50, 100, and 150 households per cluster. In case of a continuous variable, for example age, counts of the corresponding entities (in this case persons in the cluster) are obtained for the categories of the given variable. Due to this constrained geo–clustering method, the conzoom® clusters obtained comply with the social and ethical norms and preserve the privacy of the individual, yet the statistics obtained are accurate enough for effective market segmentation. This segmentation results in a grouping of the Danish population into 29 conzoom® types, one of which is defined for each 100–meter grid cell. Cosmopolitan (type 3) is one example of the 29 conzoom® types. Comparing the demographics of type 3 to the demographics of the rest of Denmark's population gives the demographic profile of the type. This profile is partially shown in Figure 1. It roughly describes individuals that are more likely: to be middle aged (30–59 years old), to live in larger cities in larger, multi–family houses that are either owned by them or are private rentals, to be mostly couples with children, to have a medium to long higher education, to hold higher level or top management positions in the financial or public sector, and to have a better household economy (in terms of wealth and income) than the average Dane.

4.2. GallupPC® consumer survey data

GallupPC® is a commercial database product and as the name suggests, it contains detailed survey responses of consumers about their demographics; interests such as culture, hobbies, and sports household consumptions, purchasing habits; transportation habits; views on various subjects; attitudes and exposure to various advertisement media [5]. The questions in the surveys are yes/no questions. To measure the magnitude of the consumer's interest in a specific area, the original yes/no question is re–phrased as categorical questions. For example the original yes/no question "Are you interested in fashion?" is re–phrased to 5 yes/no questions using the following answer possibilities: very, rather, somewhat, not very, or not interested.

4.3. bizmark™ products and services

bizmark™ is a commercial database product that contains detailed information about Danish businesses both in the public and the private sector [6]. Some of the one–to–one information that is available about businesses is their location, the number of employees working in them, the physical size of the business facility, and the international branch codes the businesses fall under. Using the hierarchy of international branch codes 40 product and service categories were identified for which related consumer surveys were also available. The product and service categories are as follows: classical concert; pop/rock concert; discothèque; art exhibition; museum; cinema; theater; pharmacy; bicycle / moped; car, stereo/HI-FI; CDs/DVDs; computer/internet; new technologies/telecommunication; do–it–yourself; fashion; cosmetics/skincare; glasses/contacts; hairdresser; jeweler/watches; interior design; travel; pets, fast-food; and 14 brand specific supermarkets. Based on the international branch codes a one–to–many relationship has been established between a subset of the businesses in bizmark™ and the 40 product / service categories.

4.4. Simulating mobile users with ST-ACTS

ST–ACTS is a probabilistic, parameterizable, realistic Spatio–Temporal ACTivity Simulator [8]. ST–ACTS is realistic in a sense, that it is based on a number of real–world data sources (among others, the data sources described above) and a set of principals that try to model the *social* and some of the *physical* aspects of mobility. The modelled principles that govern the social aspects of mobility are: (1) People move from a given location to another location with an *objective* of performing some activity at the latter location; (2) Not all are equally likely to perform a given activity. The likelihood of performing an activity depends on the *interest* of a given person, which in turn depends on a number of demographic variables; (3) The activities performed by a given person are highly *context dependent*. Some important parts of context are: the current person location, the set of locations where a given activity can be performed, the current time, and the recent history of activities of the person; (4) The locations of facilities where a given activity can be performed, are not randomly distributed, but are *influenced* by the *locations* of other facilities and the *locations* of the users of those facilities.

The output of ST–ACTS is a population of simulated persons, each described by a set of demographic variables and associated with a trajectory. The trajectories are sequences of time–stamped activities performed at particular physical locations, i.e., coordinates. In addition to the four principles above, the simulated activities also obey the following constraints. First, the *temporal activity constraint*, which states that certain activities are more likely to be performed during some periods than others. Second, the *activity duration constraint*, which states that not all activities take the same amount of time. Third, the *maximum distance constraint*, which states that for most activities there is a maximum distance a person is willing to travel. Finally, the trajectories assume linear movement between two consecutive activities, i.e. locations, but obey some *physical mobility constraints*, namely, that it takes time to move from one location to another. The time it takes to move from one location to another is calculated based on the distance between the two locations and a realistic speed model that assigns lower speeds to shorter, and higher speeds (with larger variance) to longer distances.

5. Method

The method presented here uses the Oracle RDBMS, and one of its extensions, Oracle Spatial, which provides advanced spatial features to support high–end GIS and LBS solutions [17].

5.1. LBA relational database

The objects or entities in the database are: simulated persons (or equivalently referred to as mobile users), trajectory segments, businesses, products and services. A simplified extended Entity–Relation (ER) diagram of the database is shown in Figure 2. In the extended ER diagram, square boxes represent entities, oval represent properties of entities, and diamonds represent relationships between entities. Underlined properties represent primary constraints. The arrows between entities encode connectivity of relationships, i.e., and arrow represents "one" and no arrow represents "many". For example, the "belong to" relationship is a many–to–one relationship between trajectory segments and simulated persons, i.e. one trajectory segment

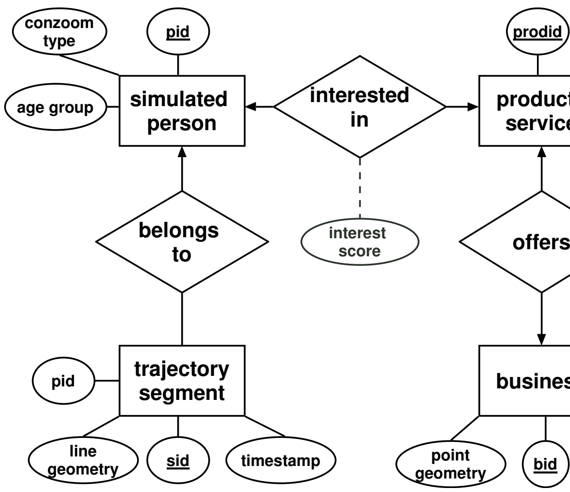

Figure 2. Simplified, extended ER diagram of the mobile advertising database.

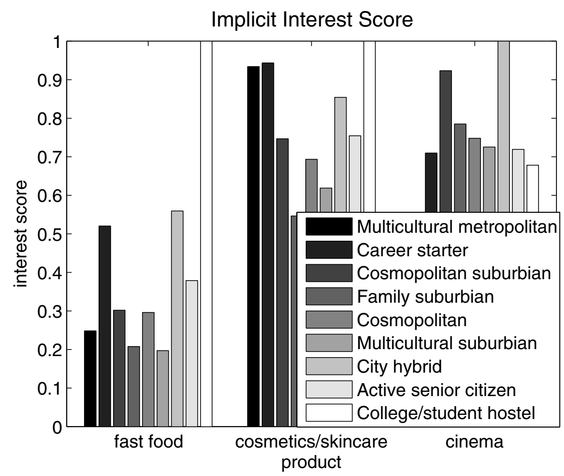

Figure 3. Some interest scores for products or services for different conzoom® types.

belongs to exactly one simulated, but many trajectory segments can belong to one simulated person. Mobile ads are indirectly modelled by in the relational database the many–to–many "offers" relationship between businesses and products or services. Through this relationship a mobile ad can be thought of as an entity having a unique identifier, `aid`, composed of a unique combination of `bid` and `prodid`, and having a *location* specified by the *point geometry* of the business offering the advertisement.

As it was introduced earlier, in the implicit case, the interest of a mobile user u and a mobile ad a about a product or service $prod(a)$ is not a Boolean function or binary relation. Rather, it is a continuous function that given the demographic characteristic $impint(u)$ of u, assigns a real valued interest score $score(u, a)$, usually from 0 (not interested) to 1 (very interested), for $prod(a)$. In direct marketing this function is termed a scoring function, which encodes a particular scoring model. This real valued scoring function is untraditionally represented as a property of the "interested in" relationship in the ER diagram.

5.2. Proximity requirements on mobile ads

A mobile ad a is likely to be considered relevant to a mobile user u only if at the time of delivery t, u is (or at some foreseeable future time point will be) within a maximum distance, *maxdist*, to the origin of the mobile ad $adloc(a)$, i.e. the location of the business. Using the spatial features of Oracle Spatial, this proximity criterion between mobile ads and mobile users is tested as follows. The geometries of the businesses, equivalently mobile ads, are buffered to a maximum distance, and tested for any spatial interaction with the geometries of the trajectory segments, by performing a spatial join operation in the database. To make the join operation as fast as possible, geometries are indexed using R–trees.

5.3. Interests based on demography

The relevance of a mobile ad for a particular product or service is naturally influenced by the interest of the user for the given product or service. As described above, a subset of the GallupPC® consumer survey questions are related to products or services that can be directly linked to businesses in bizmark™, and measure the interests of the consumer in the products or services.

Using the geo–demographic parts of the surveys, each survey subject is assigned to one of the 29 conzoom® types. To derive a single indicator, an interest score, for how interested a given conzoom® type is in a given product or service, the answers to the questions processed as follows. First, the five possible answer choices are associated with the following interest scores: very interested (1), rather interested (0.75), somewhat interested (0.5), not very interested (0.25), and not interested (0). Second, for a given conzoom® type and product or service the interest scores assigned to individual answers are averaged. Finally, the mean interest scores for a given product or service are scaled to the [0, 1]–interval amongst the 29 conzoom® types. Figure 3 shows a sample of these interest scores for a subset of the conzoom® types. In Figure 3 it can be seen, that college students are most interested in fast food and cosmetics / skincare products, and among the conzoom® types listed, suburban families are least likely to be interested in the same.

5.4. LBA – implicit interest case

The interest score of a mobile user u in a particular product or service, which is advertised by mobile ad a, is implicitly encoded in the demographic characteristic, $impint(u, a)$, or historical behavior (reaction to previously received mobile ads) of u. The latter encoding is commonly referred to as relevance feedback in the scoring task in direct marketing, and while not considered in the current mobile advertising database, it can be naturally incorporated. In direct marketing, the model for this interest score is usually derived for one or many product(s) or service(s) of a particular company through the process of data mining or machine learning. This model can be represented as a *company–specific interest score function* in the mobile advertising database that for given user–demographics and historical user–behavior assigns an interest score to the user. In the estimations however, these interest score functions are not company– but rather only product– and service–specific. Furthermore, due to their simplicity, they are implemented as table with the following schema: `interest_score` = $\langle$`conzoom_type, prodid, score`$\rangle$.

5.5. LBA – explicit interest case

Mobile users can also explicitly state their interest in certain products and services. In this case the "interested in" relationship is a binary relationship in the mobile advertising database. To provide a realistic estimates, the explicit interests of users are probabilistically simulated by randomly drawing a fixed number of products for every mobile user according to the distribution of interest scores given the conzoom® type of the user.

5.6. Uniqueness and user–defined quantitative constraints on mobile ads

Receiving the same ad multiple times naturally decreases the relevance of the ad as the therein presented information is not new. Primary key constrains in RDBMSs are an effective mechanism for guaranteeing that only unique combination of mobile users and mobile ads are considered for delivery. In the mobile marketing database the delivered ads are stored in a `mobile_ad_delivery` table with the following schema: $\langle$`pid, bid, prodid, delivery_time`$\rangle$. Placing a primary key constrain on the first three columns guarantees that a mobile ad is delivered at most once to a mobile user. Recording the *delivery time* allows the control of the re–delivery of mobile ads after a certain period of time has passed. For clarity, the `mobile_ad_delivery` table is omitted from the ER diagram in Figure 2.

As the number of mobile ads increases, or the other constraints on the delivery of mobile ads weaken, the number of mobile ads delivered to a mobile user will naturally increase. After a certain number of ads have been delivered to the user, any additionally delivered ad, while maybe relevant, will likely be perceived as annoying. Hence, the mobile user's ability to limit the number of delivered ads is important. This user–control can be effectively facilitated by the top–k query mechanism which is provided in most RDBMSs.

5.7. User–defined ST constraints on mobile ads

Time and location are important aspects of the context of mobile ads. Most users would consider receiving a mobile ad as intrusive or disturbing when receiving it during work hours or after a certain time in the evening in their homes. Hence, the mobile user's ability to prevent the delivery of mobile ads in certain regions of space and time are important. While the user–control of spatio–temporal constraints on mobile ads is not present in the mobile advertising database, the database can be easily extended to accommodate for this feature as follows. Users can specify *mobile ad profiles* by restricting certain spatial and / or temporal regions for mobile ad delivery. Then, spatio–temporal joins between the mobile ad profiles and mobile ads can be performed to further control the delivery of mobile ads.

6. Experiments and results

Two sets of experiments (implicit and explicit interest case) were performed to measure the capacity of the mobile advertising channel under various *maxdist* and *minscore* settings. The estimation are based on (1) 4314 businesses in Copenhagen, Denmark offering one or many of the 40 hand–selected products or services, (2) 3826 trajectory segments representing the movements of 1000 randomly selected simulated mobile users during the course of a day. Scores for implicit interests were modelled as described above. To simulate explicit interests, 1 product or service of interest was assigned to every simulated mobile user, as described above.

Figure 4 shows that number of delivered ads in the implicit case. As expected the number of delivered mobile ads increases as the *minscore* is decreased or the *maxdist* is increased. The rather surprising, close to linear relationship between the number of deliverable ads and the maximum distance criteria is due to the following facts. Simulated mobile users move from one location to another with the objective to perform an activity. These activities are tied to a subset of the businesses that advertise. Hence, the businesses

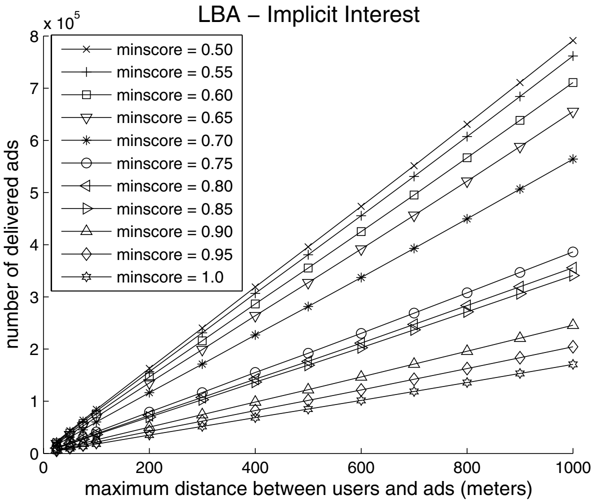

Figure 4. Number of delivered ads (implicit interest) to a population of 1000 mobile users for various $minscore$ **and** $maxdist$**.**

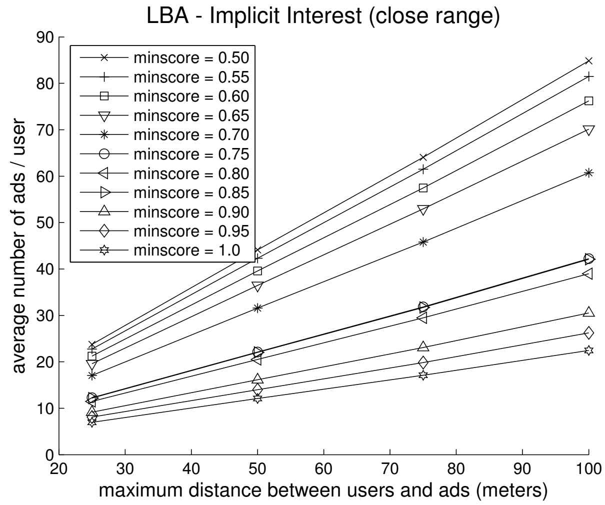

Figure 5. Number of delivered ads (implicit interest) for smaller values of $maxdist$**.**

that advertising the products or services often lie on the actual streets that the trajectories follow. If businesses are assumed to be uniformly distributed on those streets, then the relationship is indeed expected to be linear, as the number of businesses "reachable" along a street grows linear with *maxdist*. Another, rather interesting result is the sheer number of mobile ads that can be delivered to a small set of 1000 users within a course of a day. Even for $maxdist = 500$ meters (arguably a worthwhile detour for the mobile user) and *minscore* = 0.9 (quite high match in direct marketing) the average number of delivered ads to a user is about 100. This represents a huge marketing potential.

The same numbers are likely to be viewed as alarming by many mobile users. As it is shown in Figure 5, even for rather high minimum interest scores for very low *maxdist* ranges the average number of ads delivered to a mobile user during the course of a day is in the rage of 6 to 40. This is a rather large number of ads to be received on a small, by many considered as extremely personal, mobile device. Hence, to avoid bad reputation, businesses interested in employing or facilitating mobile advertising should make great efforts to provide as easy as possible user–controls on the number of received mobile ads, as suggested in Sections 5.6 and 5.7.

Figure 6 shows some statistics about the number of delivered mobile ads in the explicit interest case. Similar observations can be made about the relationship between the *maxdist* and *minscore* parameters and the number of delivered mobile ads as in the implicit case. However it is surprising that even though every mobile user is only interested in exactly 1 of the 40 products or services, due to the pres-

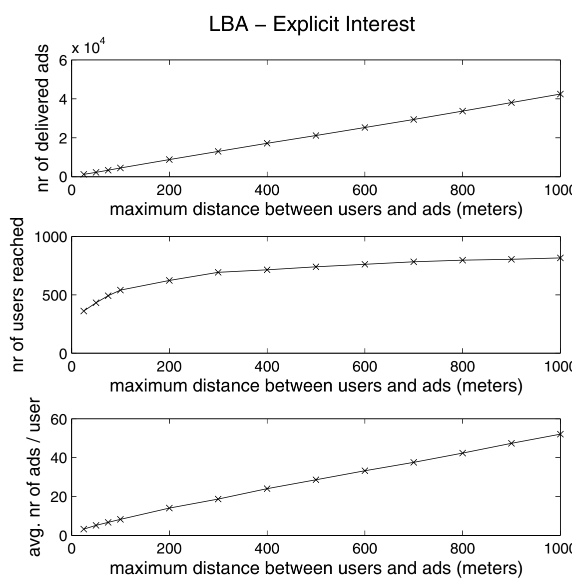

Figure 6. Statistics about the delivered ads (explicit interest) for various $maxdist$**.**

ence of a large number of businesses offering those products and services the number of deliverable ads is rather high even for small values of *maxdist*. It is also important to note, that for $maxdist < 100$, over half of the mobile users do not receive any mobile ads (middle graph). Hence the mobile users who are interested in getting good deals on products or services of their interest, have to set their *maxdist* values appropriately high.

7. Conclusions and future work

The aim of this paper was to investigate the capacity of the Location–Based Advertising (LBA) channel. The paper proposed two types of LBA models (implicit vs. explicit interest) and described a relational database for the effective management of both types of LBA. Using a number of real–world data sources and simulated but realistic movement data of mobile users, the paper gave estimates on the number of deliverable mobile ads in both the implicit and the explicit interest cases. Experimental results show that the capacity of the LBA channel is rather large implying a huge marketing potential. At the same time, the potentially large number of mobile ads could be alarming to mobile users, hence the paper warns businesses interested in LBA to provide the mobile users with adequate means to control the number of delivered ads and the time and place of delivery.

Future work is planned along two paths. First, while the presented LBA framework considers LBA both from the mobile users' and advertisers' perspective, the provided estimates are valid only if mobile users are willing to accept unlimited mobile ads at all times and places. Incorporating user–defined constraints on mobile ads, as described in Sections 5.6 and 5.7, will provided better estimates on the true audience size of LBA. Second, in the implicit interest case, the relevance of a mobile ad is estimated using a simple scoring model which is based on a consumer segmentation that divides users into 29 different consumer groups. In real life, however, no two users' interests are *exactly* the same, hence a given mobile ad does not have the same relevance to them. Hence, altering the scoring model to incorporate knowledge that is derived from historical behavior of the individual mobile user – such as the type businesses the user has previously visited or the user's reactions to previously received mobile ads – will allow targeting the individual mobile user with even more *relevant* and *personalized* mobile ads.

Acknowledgements

This work was supported in part by the Danish Ministry of Science, Technology, and Innovation under grant number 61480. The help from co–workers, Esben Taudorf, Kasper Rügge, and Lau Kingo Marcussen is also gratefully acknowledged.

References

[1] T. Brinkhoff. A Framework for Generating Network–Based Moving Objects. In *Geoinformatica 6(2)*, pp. 153–180, 2002.

[2] ESRI Business Analyst: www.esri.com/businessanalyst

[3] Experian: www.experiangroup.com/

[4] P. Fox, D. Rezania, J. Wareham, and E. Christiaanse. Will Mobiles Dream of Electric Sheep? Expectations of the New Generation of Mobile Users: Misfits with Practice and Research. In *Proc. of ICMB*, p. 44, 2006.

[5] The Gallup Organization: www.gallup.com

[6] J. M. Gartenberg, C. Matiesanu, and N. Scevak. US Mobile Marketing Forecast, 2006 to 2011, Vision Report, JupiterResearch, USA, October 2006.

[7] Geomatic aps – Centre of Geoinformatics: www.geomatic.dk

[8] G. Gidófalvi and T. B. Pedersen. ST–ACTS: A Spatio–Temporal Activity Simulator. In *Proc of ACM–GIS*, pp. 155–162, 2006.

[9] A. Guttman. Rtrees: a dynamic index structure for spatial searching. In Proc. of SIGMOD, pp. 47–57, 1984.

[10] T. Hägerstrand. Space, time and human conditions. In Dynamic allocation of urban space, ed. A. Karlqvist et. al. Lexington: Saxon House Lexington Book, 1975.

[11] C. S. Jensen, D. Pfoser, and Y. Theodoridis. Novel Approaches to the Indexing of Moving Object Trajectories. *In Proc. of VLDB*, pp. 395–406, 2000.

[12] M. Kaplan. Mobile advertising and marketing: Market analysis and forecasts 2006–2011. Vision report, visiongain™, March 2006.

[13] H. Komulainen, A. Ristola, and J. Still. Mobile advertising in the eyes of retailers and consumers – empirical evidence from a real–life experiment. In *Proc. of ICMB*, p. 37, 2006.

[14] B. Klmel. Location Based Advertising. In *Proc. of ICMB*, 2002.

[15] M. Leppaniemi and H. Karjaluoto. Factors Influencing Consumers' Willingness to Accept Mobile Advertising: A Conceptual Model. In *Int. J. of Mobile Communications. 3(3)*, pp. 197–213, 2005.

[16] Mobile Marketing Association: www.mmaglobal.com

[17] Oracle Spatial & Oracle Locator: Location Features for Oracle Database 10g: www.oracle.com/technology/products/spatial/

[18] K. Pousttchi, D. G. Wiedemann. A Contribution to Theory Building for Mobile Marketing: Categorizing Mobile Marketing Campaigns through Case Study Research. In *Proc. of ICMB*, p. 1, 2006.

[19] Statistics Denmark: www.dst.dk