Geosocially Cognizant Mobile Advertising for Clicks-to-Mortar Commerce

Richey Malhotra and George Roussos

Department of Computer Science and Information Systems, Birkbeck College, University of London, UK
{rmalho01,g.roussos}@dcs.bbk.ac.uk

Abstract. Despite the rapid growth of e-commerce, the majority of consumer spending remains offline with a significant proportion expended on goods and services provided by local merchants. Consequently, an opportunity exists for mobile systems that can drive consumers to these nearby commerce locations. Existing approaches typically employ rudimentary interest models built with the consumer's explicit involvement, demanding higher cognitive cost and leading to relatively low accuracy. In this paper, we introduce an approach to location-based mobile advertising that establishes consumer's interests by applying behavioral, semantic and social targeting to the user's geosocial context. Our working hypothesis is that this approach is more effective to match advertisements. We conduct field studies in-the-wild supported by a prototype implementation of this system, considering a range of advertising situations. We find that participants perceive advertisements selected in this way as more relevant and observe a measurable uplift in their attitudes toward advertising in general.

Keywords: Mobile Advertising, Mobile Social Networking, Context Awareness, Location Awareness, Implicit Interest Indicators

1 Introduction

The Clicks-to-Mortar (C2M) commerce paradigm – our interpretation of the O2O (Online-to-Offline) concept [6] – aims to capitalize on the opportunity presented by the 90% of consumer spending currently transacted offline [12]. It is estimated that a large proportion of this amount is spent on local goods and services, and so those companies that can *electronically shepherd* consumers to these transaction points will potentially be in a position to tap into a massive revenue stream [11].

C2M's objective is to drive foot-traffic to bricks-and-mortar retailers by the online targeting of consumers for products or services that can be fulfilled by local businesses. Transactions are initiated online, with payments being collected electronically, but delivery of the products or services necessitating the consumer go to a physical store or location. Currently, there is a significant interest around C2M-like models and the services that enable them, notably Groupon, Living Social, TIPPR and Wowcher, which provide consumers and businesses with a novel and more effective way to initiate and execute transactions.

adfa, p. 1, 2011.

However, thus far, this has been achieved using relatively simple email targeting and delivery mechanisms, supported by rudimentary location awareness. Fully realizing C2M commerce needs the accurate, insightful and timely matching of consumers' needs to advertisements and offers from local retailers. A platform with such functionality is described in the following sections, it leverages the consumer's *geosocial composite context* to direct the matching process. The remainder of the paper describes the design and implementation of the system, and the results from a pilot field study. The field study followed a within-subjects field experiment methodology, involving four participants, performing a range of advertising scenarios. Using interviews and surveys administered pre-, intra- and post- study, we measured a raft of quantitative and qualitative facets of mobile advertising, including in particular, participants' rating of the relevance of advertisements shown and any changes in their general attitude toward advertising.

2 Geosocially Cognizant Mobile Advertising

The key enabler for C2M and an evolution of online advertising, is what we call Geosocially Cognizant Mobile Advertising (GCMA), defined as: applying behavioral, semantic and social targeting to the interest signals present in the consumer's geosocial composite context (click-streams, searches, social networking data including social graphs, location traces, user-generated content, and mobile dialogues), to deliver local business advertisements at a location, time and in a format that will maximize their efficacy.

Many have recognized the possibility to improve online advertising through the application of behavioral and semantic targeting. With the recent widespread use and ubiquity of social networking, augmenting these approaches with the use of social targeting has led to further increases in advertising effectiveness. We believe that the next step in this progression is to target advertisements based on the consumer's geosocial data – as sent and received from their smartphone – using a combination of all three of these targeting technologies. Our goal is to show that the collection, aggregation and real-time analysis of such information provide a better basis for establishing the consumer's interests compared to current methods.

For GCMA, behavioral targeting is applied to both the consumer's browsing-history and their "foot-stream" i.e. to the name and/or category of locations visited. Contents of browsed web pages, messaging and user-generated content are yet more insights into the consumer's interests and needs. By understanding the topics and sentiments of these communications, it is possible to determine the consumer's concerns – this is the aim of semantic targeting. GCMA's use of social targeting utilizes the fact that people connected to one another in a social network share many similar traits (the homophily principle [9]), and thus are likely to be interested in the same products and services as other members of their network [7], with the "influencers" in these networks being high-value targets for advertisers [13].

GCMA aims to deliver personalized, customized, location-centric offers, chosen by analyzing the patterns of places visited and frequented, and the subjects and opinions expressed in sent and received messages, browsed webpages, and other user-generated content such micro-blogs and Twitter. The offers presented to consumers are based on longitudinal / "conversational" observations, and delivered relative to detected time- and date- specific statements or "just-in-time" when the consumer arrives at a pertinent location. The GCMA approach leverages behavioral, semantic and social targeting strategies and is based on the following precepts:

1. Interest signals are siloed across a range of media: click-streams, searches, and social network data including the social graph, location traces, user-generated content and mobile dialogues.
2. The application of behavioral, semantic and social analysis to these signals can educe the underlying interests or needs. Detected interests and needs have associated probabilities and weights, derived from the confidence level of the observations and correlations.
3. Recommending advertisements / offers (reactively, proactively or serendipitously) to the consumer from the closest suitable business, delivered at the most apposite time and location, and in a customized and personalized form, attempts to maximize the impact and relevance of the message, and subsequent probability of conversion.

2.1 Interest Signals

The concept of interest signals outlined here is similar to John Battelle's Database of Intentions [2], but extended with the addition of user-generated content, dialogues and inventory, in an attempt to give a more real-time and comprehensive view of the consumer's interests (see Figure 1). For GCMA to work these interest signals need to be: Modulated – transferred online as a result of real world consumer interactions; Propagated – shared and updated across various online interest repositories and sources; Demodulated – transferred as the underlying need or interest to an offline merchant.

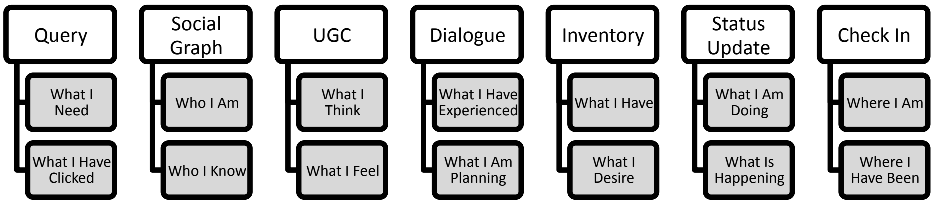

Fig. 1. Interest Repositories and Signals

These processes aim to remove the distinction between online and offline worlds, and as a result support a more flexible model of the consumer purchasing cycle. Figure 2 summarizes the overall path from interest signal to serving an advertisement.

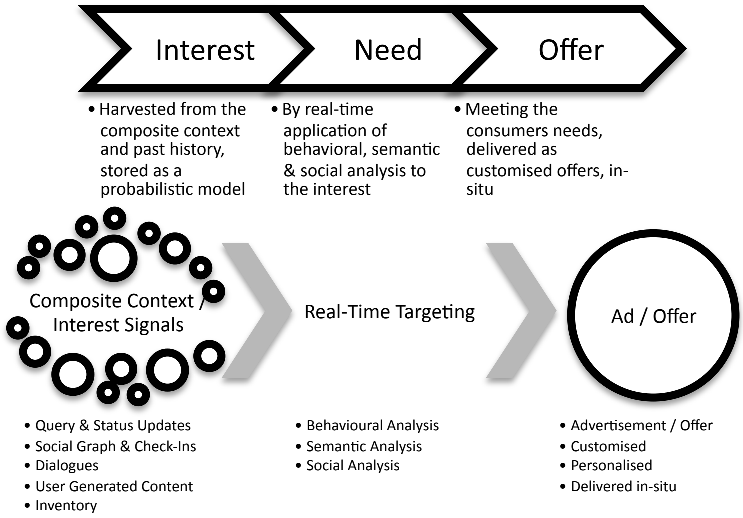

Fig. 2. Path from Interest to Advertisement

2.2 Composite Context

To implement GCMA, the context sensing abilities of the application need to be extensive, the scope of which is illustrated by the concept of Composite Context. For a context to be composite, it must encompass a number of physical and virtual elements – which in the case of GCMA effectively embody the consumer's geosocial data (see Figure 3). This approach bridges the online and offline worlds, and provides the foundations for a needs-based location based advertising platform.

In [1] Anderson *et al*. take the view that mining multiple sources, such social networks, user-generated content, messaging and location traces leads to a more nuanced and finer understanding the consumer's interests, individual or shared, and proposes a geo-social data management service to facilitate this. Liu *et al.* [8] also share our vision for an extensive context-aware architecture built on interest signals, and accordingly defines their "mission statement and research agenda."

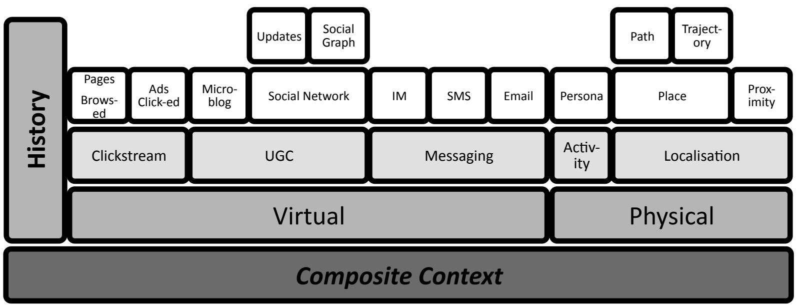

Fig. 3. Composite Context Structure

2.3 Technical Architecture

Figure 4 gives a high-level overview of the technical architecture for the GCMA proof-of-concept implementation. Note that in order to meet time and resource constraints only the most fundamental targeting approach was implemented (semantic), with simplified versions of Composite Context and User History. The construction of the interest model was performed manually, and advertisement selection and serving was simulated by passing the identified keywords to Google (see section 3.1).

The prototype GCMA application is essentially a *User Event Recorder*, in that it monitors a range of user actions and records them in strict sequence of occurrence to a hosted SQL database. The data is then used to build a model of the user's interests, which is in turn used to direct the selection and serving of advertisements to the user. The events and actions of interest (recorded from the point in time the application is launched) are: SMSs sent and received; http address of browsed pages; GPS coordinates / Cell ID; Tweet body; Foursquare check-in location, name, and categories.

The application records this data in the background, and apart from a launch screen and a configuration screen where the user's Twitter and Foursquare account information and recording polling rate is specified, no other GUI element is used. For each SMS, tweet and browsed web page, the topics and sentiments (including any implied actions or timescales) are determined using OpenCalais and OpenAmplify semantic analysis APIs.[1] The results of this analysis are stored in the database alongside the original communication.

3 User Study

To understand the potential value of our approach to advertising we conducted a field study of the application's use in-the-wild. The intrinsic linkage between context-aware systems and context makes it difficult to test them in a laboratory setting, as it is exigent to recreate the confluence of events and factors affecting the systems functionality, performance and usability [5]. Nevertheless, field studies are the best method known to researchers to mitigate these effects, and can be a rich source of quantitative and qualitative data.

For the study portion of this project we applied the guidelines, practices and principles for proof-of-concept within-subjects field studies as described in [10] [3] [4]. The field study helped us to collect a significant amount of real-world in-situ usage data; observe any issues or problems that participants may of experienced; and understand the impact and potential improvements for the system.

Each participant was required to perform three different advertising usage scenarios and assess the relevance of advertisements served, as well as answering other qualitative and quantitative surveys and interviews. We explained to the participants what tools and methods they had their disposal to enact the scenarios but left them to decide how best to use these to meet the objectives of the particular situation. The

---

1 http://opencalais.com/about (accessed 01/02/12)
http://openamplify.com/about (accessed 01/02/12)

order of the scenarios was randomized for each participant to minimize learning effects across subjects.

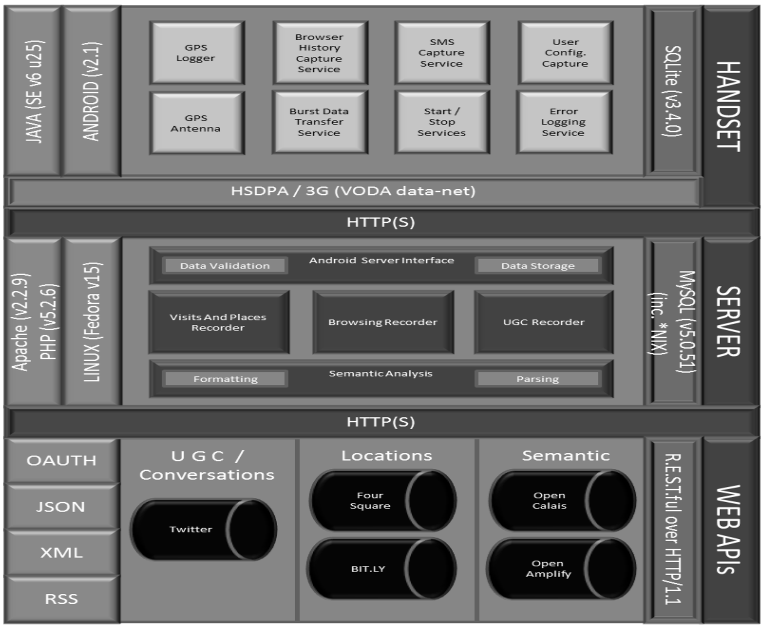

Fig. 4. Application Architecture for GCMA

3.1 Design

We placed an advertisement on a local website, asking for volunteers for a paid scientific study on mobile advertising. Due to the technology-centric nature of our study, we remunerated participants to avoid attracting only technophiles/early adopters who could skew or bias the proceedings. The advertisement explained that the completion /of the study was necessary for payment, the process and duration (see Figure 5), and nature of data collected, including privacy measures and intended use for the data.

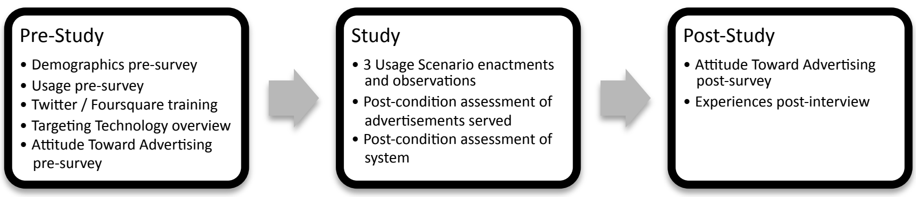

Fig. 5. Field Study Process Overview

Of the 11 respondents, only four (2M 2F) had the required background and best covered the anticipated and intended demographic for GCMA users. The chosen participants, possibly motivated by the belief they would benefit from the offers and discounts offered by the proposed system, ranged in age from 23 to 41 (median 31), and came from a variety of occupational and educational backgrounds. The advertisement outlined the required profile for a participant: user of PC as part of their daily activities; owner and user of smartphone; regular browser of internet from the desktop and smartphone; purchaser of goods or services from internet; and user of Twitter and/or Foursquare.

In order to gather data on in-situ use of GCMA we employed a within-subjects field experiment methodology. Each participant was given a Sony Xperia X10 Mini 3G Smartphone, running Android 2.1 and our application as a background service. The handsets had data connections, AGPS capability, and came preloaded with Twitter and Foursquare clients. The browser and SMS functionality was provided by OEM applications.

As part of the field study, each participant had to perform three advertising usage scenarios:

- Deciding on a particular bar for a night out (BAR);
- Choosing a restaurant (FOOD);
- Finding an apartment to rent (RENT).

Each usage scenario was designed to demonstrate a different aspect of GCMA, so that participants could compare its behaviour versus a traditional advertising system like Google. The BAR usage scenario was intended to illustrate the Conversational Analysis functionality of GCMA; the FOOD scenario aimed to show the Interest Model functionality; and finally, RENT attempted to demonstrate Important Places. These usage scenarios involved non-trivial interactions with our application and other supporting systems such as text messaging, Twitter, Foursquare and the default browser. The whole field study procedure lasted under 10 hours per participant, scheduled non-contiguously over two weeks to match participants' availability.

The execution of each usage scenarios began in the same way, with the participant entering their Twitter and Foursquare credentials into the GCMA application, and then setting it to record. The application recorded, with timestamps, each communication sent or received from the participant's handset and the GPS locations throughout the enactment of the scenario. This data was then post-processed to generate the advertisements that would have been generated had the participant been using: 1) a full implementation of GCMA, and 2) a traditional advertising system like Google.

On the day following the enactment, we sat with each participant and walked them through the scenarios, highlighting the points at which advertisements would have been served and the body of actual advertisements themselves. Each participant was shown two sets of advertisements per scenario, each corresponding to a different condition for the experiment. The experiment was designed so that the participants would not see the advertisements selected by the system until after they had performed all the required scenarios. This allowed us to separate the assessment of the system's usability from the assessment of its effectiveness, and thus circumventing any possible confounding effects due to well/badly matched advertisements served mid-process.

In addition, each participant completed a number of pre-study, post-study and post-condition surveys and interviews to enable us to collect further information. The two most important of these were the:

- Post-Condition Assessment of Advertisements Served Survey (PCAASS): For every advertisement served, the participant was asked to rate them for relevancy in relation the usage scenario they were enacting. This was completed for both conditions across all usage scenarios, thus allowing us to compare the participant's view of the different technologies for the different advertising situations
- Attitude Toward Advertising Pre- and Post-Study Survey (ATAPPSS): The Attitude Toward Advertising survey was a Likert-scale based survey that asked participants to rate their experience of advertising, in terms of value, trust, intention and any other notable facets. It was completed by the participants both at the beginning of field study, and then at the end, after being exposed to both conditions across all usage scenarios.

3.2 Results

All four participants completed the enactment of the usage scenarios, and the data collected by the application was of sufficient quantity and quality to allow us to construct models of their interests and subsequently match advertisements against them. Furthermore, the application logs showed that the enactment times were proportional to the complexity of the given scenarios and thus not confounded by any other factors, possibly such as confusion over how to use of the system or performance issues.

From the analysis of the PCAASS data for the two fully executed usage scenarios (the third was not performed due to technical problems on the day), although there was no homogenous trend across all participants, there were significant uplifts in relative rating between conditions for P3 and P4, with a lesser indication for P2 (see Figure 6). However, the aggregate net mean change by participant (for both conditions and across all usage scenarios), showed that there was a positive change for all participants, with the most significant increase recorded by P3.

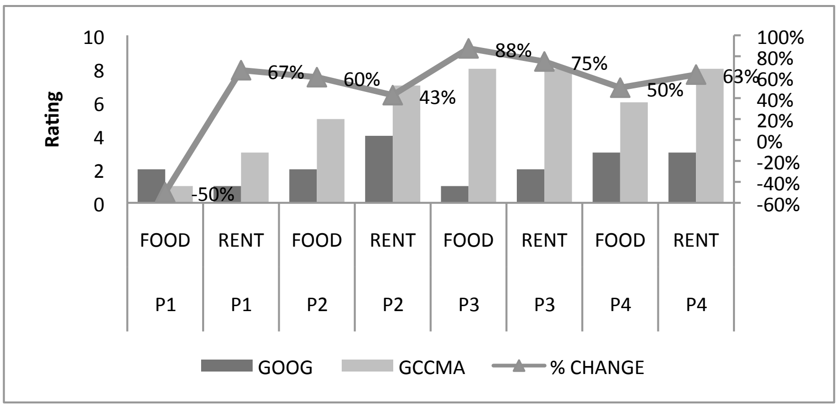

Fig. 6. Rating Of Advertisements Served By Condition across Usage Scenarios

Analysis of the ATAPPSS data reflected a stronger sentiment. This survey was a snapshot of the participants' attitudes toward mobile advertising in general, before and after having gone through the study. This allowed us to capture the change (if any) in participants' attitudes toward advertising attributable to having taken part in the field study and experiencing the different conditions and usage scenarios. From examining the means pre- and post- study, we saw a homogenous uplift across all participants, with significant changes in mean observable in the answers to questions V1, T2, T3 and I1 (see figure 7). [2]

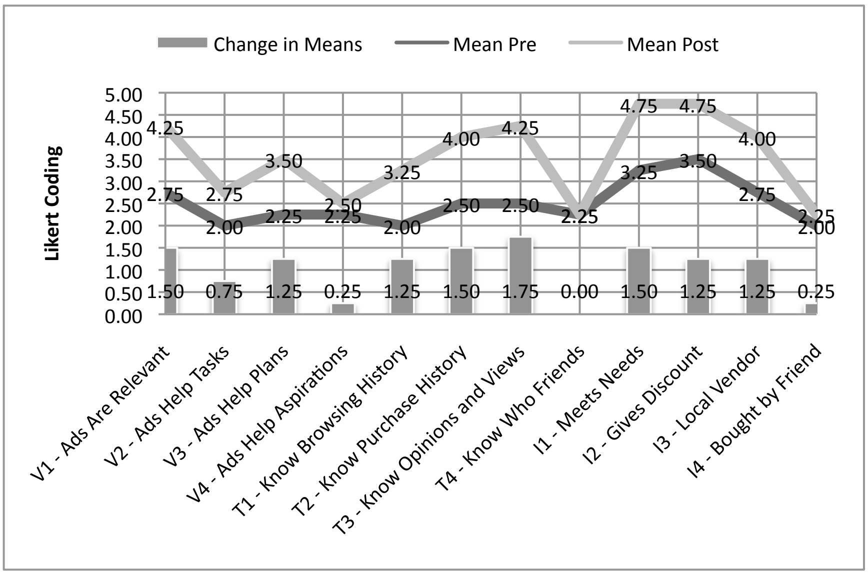

Fig. 7. Change in Means of Participant Responses for ATAPPSS

4 Summary and Discussion

We are mindful that our proof-of-concept is just one way to realise an alternative approach for advertising, which we have tested with a limited number of participants over a restricted range of usage situations. As such, we cannot claim that this approach and our prototype are universally applicable, but for the small field study group made up from our intended demographic, the results demonstrate a definite level of support for the research objective and hypothesis.

We have explored the feasibility and usefulness of geosocial composite context and interest signals for matching advertisements on mobile devices. We have built a proof-of-concept application and used it for a field study involving a small group of

[2] V1:"Advertising is relevant", T2:"I don't mind advertisers knowing what products or services I have purchased", T3:"I don't mind advertisers knowing about my opinions and views", and I1:"I am likely to purchase the goods or services when the advertisement matches my needs".

users enacting different advertising scenarios. After analysing the resulting system logs, survey data and observations, we saw a statistically significant increase in participants' opinion of advertising relevance and improvement in their attitude toward advertising overall. To our knowledge, we are unique in examining such a broad set of interest signals, and as such extend the scope of exploration into how signals affected mobile advertising effectiveness and potential benefits for both consumers and advertisers.

Whilst our study had limitations, including the lack of real-time feedback and delayed matching/serving of advertisements, we managed to provide our participants with a palpable taste of a possible future for advertising where advertisements could be relevant but still respect privacy, and be calm while simultaneously being engaging. Directions for future work including looking at different combinations of interest signals and time periods for building the interest model, increasing number of participants, exploring a wider range of usage scenarios, and a simulation-free implementation that could provide real-time functionality.

References

1. Paul Anderson et al. On managing social data for enabling socially-aware applications and services. In Proceedings of the 3rd Workshop on Social Network Systems, SNS '10, pages 7:1–7:6, New York, NY, USA, 2010. ACM.
2. John Battelle. Database of intentions chart - version 2, updated for commerce. Battelle Media, March 2010. http://tinyurl.com/5stw433 Last accessed 01/02/12
3. A.J. Bernheim Brush. Ubiquitous Computing Field Studies. In John Krumm (ed.), Ubiquitous Computing Fundamentals, 2009, pp. 161-202
4. A.J. Bernheim Brush et al. User Experiences with Activity-Based Navigation on Mobile Devices. In Proceedings of MobileHCI 2010, September 7 - 10, 2010, Lisbon, Portugal.
5. Anind K. Dey. Context Aware Computing. In John Krumm (ed.), Ubiquitous Computing Fundamentals, 2009, pp. 321-350
6. Dan Frommer. Ron Conway's confidential investment megatrend - O2O commerce. Business Insider, February 2011. http://tinyurl.com/47s9mp9 Last accessed 02/02/12
7. Shawndra Hill, Foster Provost, and Chris Volinsky. Network-based marketing: Identifying likely adopters via consumer networks. Statistical Science, 21(2):256–276, 2006.
8. Jie Liu Michel Goraczko, Aman Kansal, Dimitrios Lymberopoulos, Suman Nath, and Bodhi Priyantha. Subjective Sensing: Mission statement and a research agenda. Technical report, Microsoft Research, July 2009.
9. Miller McPherson, Lynn Smith-Lovin, and James M. Cook. Birds of a feather: Homophily in social networks. Annual Review of Sociology, 27:415–444, 2001.
10. Meredith Ringel Morris et al: A field study of knowledge workers' use of interactive horizontal displays. Workshop on Horizontal Interactive Human-Computer Systems, 2008, pp. 105-112.
11. Alex Rampell. Why online2offline commerce is a trillion dollar opportunity. Tech Crunch, August 2010. http://tinyurl.com/28fmpgz Last accessed 02/02/12
12. Erick Schonfeld. Forrester forecast: Online retail sales will grow to USD 250 billion by 2014. Tech Crunch, March 2010. http://tinyurl.com/yascwmq Last accessed 02/02/12
13. Erica Suyeon Kim and Steve Sangki Han. An analytical way to find influencers on social networks and validate their effects in disseminating social games. In Proceedings of the 2009 International Conference on Advances in Social Network Analysis and Mining, pp. 41–46, Washington, DC, USA, 2009. IEEE Computer Society.