Not All Adware is Badware: Towards Privacy-Aware Advertising

Hamed Haddadi, Saikat Guha, Paul Francis

Max Planck Institute for Software Systems (MPI-SWS), Germany

Abstract. Online advertising is a major economic force in the Internet today. A basic goal of any advertising system is to accurately target the ad to the recipient audience. While Internet technology brings the promise of extremely well-targeted ad placement, there have always been serious privacy concerns surrounding personalization. Today there is a constant battle between privacy advocates and advertisers, where advertisers try to push new personalization technologies, and privacy advocates try to stop them. As long as privacy advocates, however, are unable to propose an alternative personalization system that is private, this is a battle they are destined to lose. This paper presents the framework for such an alternative system, the Private Verifiable Advertising (Privad). We describe the privacy issues associated with today's advertising systems, describe Privad, and discuss its pros and cons and the challenges that remain.

1 Introduction

Online advertising is one of the key economic drivers of the modern Internet economy. It supports many web sites and web services, is the basis for such industry giants as Google and Yahoo!, and helps pay for data centers and ISPs. Online advertising has long been touted as a natural basis for highly targeted and personalized advertising. It is good for advertisers who can avoid losing money by delivering ads to uninterested consumers, and good for consumers, who do not get bombarded with ads that are not in line with their interests. Unfortunately, personalized online advertising, at least so far, comes at a price: individual privacy. In order to deliver ads that individuals are interested in, advertisers learn what individual's interests are. While a number of private advertising systems have been designed, these either exhibit too weak a notion of privacy, or are too expensive to deploy (see Section 4). This paper introduces what we believe can form the first viable strongly private personalized online advertising system.

Online advertising has a long and storied history. Because it has always violated individual privacy, its practice has always been controversial opposed by privacy advocates (where here we use the term broadly to encompass governments and concerned users as well as privacy advocacy groups like the Electronic Frontier Foundation (EFF)). Privacy advocates have had some successes: most adware has been shut down, through a combination of legal (individual and

class-action lawsuits) and technical means (virus detection software). More recently, privacy advocates have had some success in stopping or slowing trials of ISP-based advertising technologies such as proposed by the companies NebuAd [1] and Phorm [2].

Privacy advocates have also had some major failures, specifically in the now ubiquitous advertising models deployed by Google, Yahoo!, Microsoft, and others. In spite of the fact that each of these companies, including recently Google, maintain personalization information about users, privacy advocates have so far been pretty much powerless to put a stop to it. Arguably the reason for this is that these and other companies play such a central role in the Internet economy that the deal that privacy advocates offer, privacy OR advertising, is not acceptable to industry, governments, or most individuals. In other words, as long as privacy advocates do not offer up an alternative privacy-preserving approach to personalized advertising, we will be stuck with the status quo or worse.

This paper introduces a framework for that alternative model. Specifically, we propose a Private Verifiable Advertising (Privad), a personalized advertising model based on the following three principles:

1. That all personalization information is kept on the client computer, and all decisions about what ads to show are made *purely local* to the client computer. This is made possible by pushing many or all ads to all client computers in advance of showing each ad.
2. That all reports about which ads are shown provide *zero information* to any given system component that would otherwise allow that component to associate users with ads viewed.
3. That privacy advocates are able to *verify* that the first two principles are indeed maintained in practice by observing all messages that go in and out of the various system components. It is this verification principle in particular that makes Privad powerful and unique.

It is worth pointing out that Privad is, by any reasonable definition, adware: software runs on client computers that monitors user behavior, builds up a profile of some sort, and uses that to present targeted ads to the user. We argue that not all adware is badware, and that in fact a client-based personalization system (i.e. adware) is the only way to achieve privacy in advertising. "Good adware" is not an oxymoron.

The rest of this paper is structured as follows. In Section 2 we introduce the basics of current advertising systems and we discuss user privacy. Section 3 presents the Privad framework. In Section 4 we give an overview of related work on privacy and advertising. Finally in Section 5 we conclude the paper and speculate on potential future avenues for research and for the advertising industry.

---

[1] `http://www.nebuad.com`
[2] `http://www.phorm.com`

2   Advertising Basics

For this paper, we define four major components of advertising systems: advertisers, publishers, clients, and brokers. Advertisers wish to sell their products or services through ads. Publishers provider opportunities to view ads, for instance by providing space for ad banners. Clients are the computers that show publisher web pages and ads to users. Brokers bring together advertisers, publishers, and clients. They provide ads to users, gather statistics about what ads were shown on which publisher pages, collect money from the advertisers, and pay the publishers.

Brokers can serve ads with or without user personalization. For instance, Google provides ads that match the topic of the web page with the embedded banner, on the assumption that if the user is interested in the web page, then the user is also likely to be interested in the ad. In this limited case, there is no personalization (Google, until recently, did not profile the user per se). Even in this case, however, there is a potential privacy violation: Google and other brokers see many of the web pages each client visits. This information can easily be used to personalize clients and their users. Even if the broker does not directly obtain Personally Identifying Information (PII) such as names, addresses, social security numbers and so on, PII can often be obtain through tracking cookies or by correlating IP addresses with PII obtained elsewhere [6].

In spite of the fact that some advertising can be done without user personalization, increasingly this personalization is being done by major advertising companies in form of persistent cookies. These cookies store information about user's visits to different websites, the frequency of such visits and the click information. An example of use of such cookies is with Behavioral Targeting in Google AdSense[3] or the Microsoft advertisement program. Although these brokers use some of the collected information in order to personalize the ads shown, they insist the systems select ads based only on non-PII data. For example, they store page views, clicks and search terms used for ad personalization separately from the user contact information or other PII data that directly identifies him or her. They also provide the ability to opt-out of personalized ads. Google has also tried to alleviate the concerns grown over about profiling users by promising not to create sensitive interest categories like race or religion or cross-correlating the data with other information saved in Google accounts. However the consumer can only trust Google to adhere to its word. Similar levels of privacy concerns have grown over commercial software such as the Google Desktop[4], where personal information can be used to effect public search results for targeted advertising[5].

These methods all lead to erosion of user privacy by the historically dominant players in the field such as Google, Microsoft and Yahoo!, as well as by relative newcomers like FaceBook. This leaves those concerned about privacy in a frustrating situation. On one hand, they can not easily modify the opera-

---

[3] http://googleblog.blogspot.com/2009/03/making-ads-more-interesting.html
[4] http://desktop.google.com/
[5] http://www.theregister.co.uk/2004/10/15/google_desktop_privacy/

tional behavior of the big players, in part because there is little regulation of the industry, and in part because these players play such an important role in the Internet economy. On the other hand, they spend a great amount of effort trying to impede new players and technologies to enter the market, without offering any viable alternative solutions. Ironically, this only helps solidify the position of the dominant players.
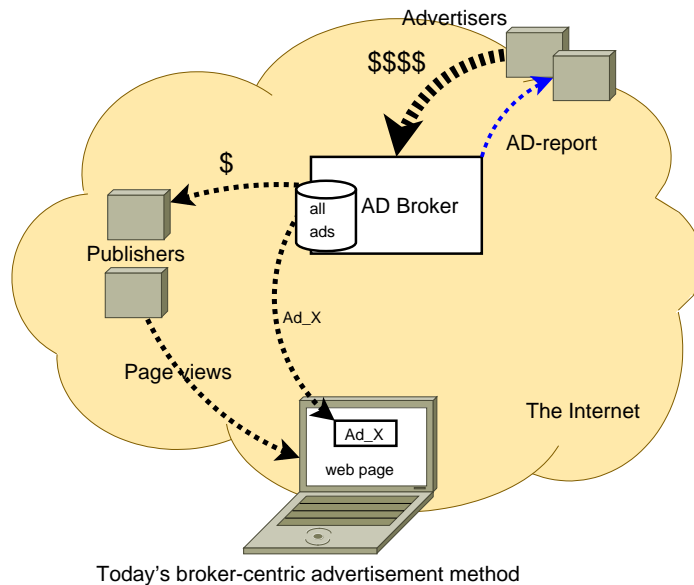

Today's broker-centric advertisement method

Fig. 1. Modern keyword-based advertising.

Figure 1 illustrates the current advertising model. The advertiser provides ads and bids (how much the advertiser is willing to pay for views and clicks of its ads) to the broker. When a publisher provides banner space to the client on a web page, a request goes to the broker asking to fill in the banner space with an appropriate ad. The broker makes the decision as to which ads to place based on a number of criteria such as the keywords for the web page, personalization information about the client, the keywords of the ad, and the bid associated with the ad. The broker provides the ad to the client, informs the advertiser of the ad view, and eventually settles the financial bookkeeping with advertisers and publishers. The broker also records user clicks on ads and possibly other events associated with ads (such as a product purchase), and reports to the advertiser. The advertiser uses the received information to run its ad campaign. That is, to decide how much to bid for ads, which keywords to associate ads with, which demographics to target, and so on.

3   The Privad Framework

Fig. 2. The Privad architecture.

Figure 2 illustrates the Privad model. The first thing to note here is that it preserves the economic framework of existing advertising systems. There are advertisers, publishers, clients, and a broker. Privad does introduce a fifth component, the untrusted proxy. The proxy is an entity that wishes to ensure that privacy is maintained, but nevertheless is itself not trusted with private information. As such, it could be operated by government agencies or by privacy advocacy groups (though possibly financially supported by the broker).

In Privad, no single entity obtains private information, including the proxy, and can therefore be untrusted. We do require that the proxy and the broker do not collude with each other. Proxy operation, however, can easily be monitored to insure with high probability that no collusion is taking place. Because the proxy is intrinsically interested in maintaining privacy, it should be willing to open itself to such monitoring.

Although Privad preserves the basic economic structure of current advertising systems, there are key differences. The main difference is that ads are served not by the broker, but by the client itself. This is done by providing the client with two things: a database of all or most ads, and a software agent that selects ads from the database to show the user, probably though not necessarily through user profiling.

In this model, when the client receives a webpage with banner space from a publisher, it itself fills in the banner with an appropriate ad. It generates an encrypted report identifying the ad and the publisher. The broker can decrypt the report, but not the proxy. The report is transmitted to the proxy, which now

knows that the client viewed or clicked on an ad, but cannot know which ad or which publisher. The proxy forwards the report to the broker, and in so doing hides the IP address of the client. Upon decrypting the report, the broker now knows that some client viewed a particular ad served by a particular publisher, but does not know which client. What's more, there is no information in the report that allows the broker to associated multiple reports with a given client. The broker then reports to the advertiser as with today's systems. A similar report is produced by the client for other ad events such as clicks, purchases, and so on.

As already stated, a critical feature of Privad is that its privacy attributes can be verified, thus eliminating the need for trust in any single component. This can be done by a monitor positioned in two places: at the client and at the proxy. Specifically, monitors can observe messages leaving the client, arriving at the proxy, and leaving the proxy. Of course, the monitor cannot verify that no private information is leaked from the client if it cannot view the contents of the encrypted report.
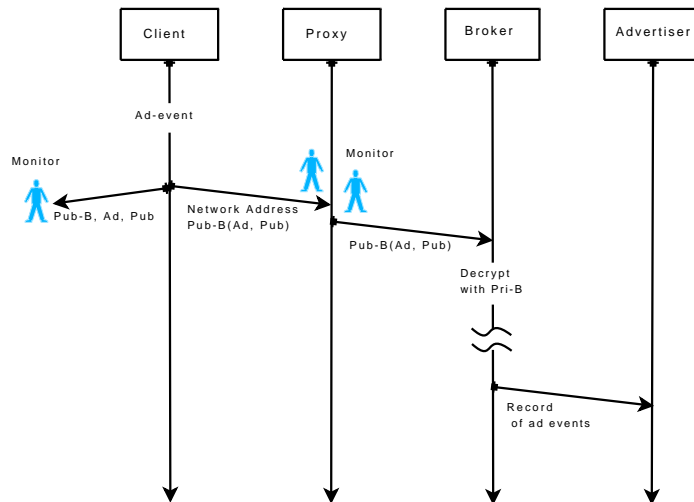

Fig. 3. Inter-party Encryption details.

Figure 3 gives a somewhat simplified illustration how this problem is solved. When the client produces a report, it encrypts the report with the public key of the broker, denoted $Pub - B$. At the same time, it gives a monitor positioned at the client $Pub - B$ as well as the contents of the report, the ad ID Ad and the publisher ID Pub. This monitor could be software running on behalf of a concerned user or privacy advocate. With this, the monitor is able to replicate the encryption step and produce an encrypted report. If the encrypted report produced by the monitor matches that transmitted by the client, then the mon-

itor is assured that the report contains only *Ad* and *Pub*, and not some private information.

In practice, it will be necessary to add salt to the encrypted part of the message to prevent a dictionary attack at the untrusted Proxy. If this salt is produced through a pseudo-random number generator, then the client cannot use the salt as a covert channel (i.e. pick "random" numbers that in fact embed personal user information). If the monitor provides the initial key, and knows the pseudo-random number generator algorithm, then it can validate that the salt is not a covert channel.

Likewise, monitors at the input and output of the proxy can verify that each report going into the proxy is matched by a report going out of the proxy, and that therefore the proxy did not convey the IP address to the broker.

Of course, these reports alone are not the only means by which the client could potentially communicate private information to the broker, or that the proxy could communicate IP addresses to the broker. Legitimate processes within client computers commonly transmit thousands of messages in the background to a wide variety of destinations around the Internet. The software agent could hide its own messages within this background noise. We admit that individual users would not be in a position to detect such messages. Rather, we envision that privacy advocates would install the software agent in clients that they control and carefully monitor all messages going in and out. Privacy advocates could then build up reasonable confidence that no software agents are leaking private information.

Likewise, a proxy intent on colluding with the broker could convey client IP addresses through background channels such as file transfers or even disks sent through physical mail. Proxy operation, however, is relatively simple, and so we would reasonably expect a monitor to be able to audit the operating procedures of a proxy to insure with high confidence that no collusion is taking place.

In the following, we provide additional detail and related challenges.

3.1   Profiling

User profiling has a long established history in the advertising industry. Starting from initial adware and banner advertisements and leading to today's cookie placement strategies and user traffic stream analysis, the brokers and advertisers have been continuously working on improving their penetration levels into more fine-grained groups of users. This has, unfortunately, been all at the expense of gathering a large amount of information from the users who, usually unknowingly, accept the long and complicated terms and conditions used by most providers. This has the effect of hiding the monitoring abilities of such cookies and adware software from users.

With Privad, Users would voluntarily install the software agent, probably bundled with some useful software package as historically done with adware (and, unfortunately, spyware). Because Privad is private, and hopefully has the imprimatur of privacy advocates and government agencies, it will not be necessary to obscure the profiling operation of Privad from users. Users would opt-in

with full knowledge. It should go without saying that the software agent would be disabled when users un-install the associated bundled software package.

Of course, it is necessary to keep the profiling information stored locally on the client safe from the prying eyes of other users of the client or other malware that may invade the client. This problem is similar in spirit to the privacy issues created by the browsing histories kept by browsers. The problem is in some respects potentially made worse by Privad, however, because the profile may contain such information as salary and specific shopping habits. To mitigate this, the profile can be encrypted as long as it is on disk, and only kept in the clear while in memory. Another option might be to allow users to flush the profile from time to time, though the ability to do this, or the frequency with which it can be done, has to be weighed against the loss in ad targeting.

3.2 Ad Database Dissemination

A major challenge in Privad is ad dissemination. Not only do potentially hundreds of thousands of ads need to be disseminated, the associated bids and targetting meta-data such as keywords need to be disseminated as well. If all ads are pushed to all clients, then nothing is learned about those clients. If this does not adequately scale, however, the broker must be selective in which ads it gives to which clients. To the extent that this selection requires specific knowledge about client profiles, Privad leaks private information.

Before discussing how we can mitigate these scaling issues, it is worth pointing out that substantial volumes of information can be broadcast to clients with existing broadband installations. For instance, thousands of ads can be downloaded into a client for the equivalent bandwidth of a few minutes of YouTube. To ease load on the proxy server, the download could take place as an ongoing BitTorrent. Existing peer-to-peer video broadcasting systems, for instance, have demonstrated the ability to disseminate many megabytes of video stream to thousands of clients within a few tens of seconds [7]. The client could locally discard any ads that are clearly unlikely to be relevant, thus reducing its memory requirements.

One way to reduce the volume of distribution, while preserving privacy, would be to distribute each ad to a randomly selected fraction of clients. The advertiser only requires that each ad is shown a certain number of times, not that the ad ultimately distributed to all clients. If the budget for an ad is relatively small, then the ad may correspondingly be distributed to a relatively smaller number of clients. Randomness could be achieved by exploiting the randomness inherent in a BitTorrent peer-to-peer style of distribution. For instance, each ad could have attached to it the number of clients that should store the ad. This number could be divided and decremented as the ad is shared with peers and stored, thus limited the distribution of the ad.

If random scoping is not adequate, then scoping based on certain relatively non-sensitive demographic information could be done. For instance, ads could be disseminated based on language or location.

While we have some confidence that the scaling issues associated with ad dissemination can be overcome, measurement and experimentation are ultimately required to support or undermine this confidence.

4 Related work

In this section we briefly present the related work on user profiling and private advertising. We also cover the case of advertising on online social networks as the amount of detail provided by users in their profiles puts these social networks in a strong position in relation to targeted advertising. For example in the case of FaceBook[6] there were a range of advertisement plans devised by the owners. Using the FaceBook interface, companies would be able to selectively target FaceBook's members in order to research the appeal of new products through a polling system called Engagement Ads and FaceBook Beacons. However in some cases the user's privacy concerns forced changes to such policies[7].

Recently there has been attempts by ISPs to monitor customers' traffic and offer them targeted advertisements. For example information such as web pages viewed and links clicked on, including advertisements, search terms, browsing trends and page click chains, response to advertisements and more importantly demographics such as language, the day of the week, time and domain locations are used by some of these cooperative advertisement business plans[8].

In UK, British Telecommunication initially announced that a controversial online advertisement system would be rolled out, but it stressed that any profiling will be done only with the knowledge and agreement of the customer. This announcement was motivated by the belief that systems such as Phorm[9] are the only way to keep ISPs afloat in the future. In the trial stages, BT had admitted that it secretly used customer data to test Phorm's ad targeting technology, and that it covered it up when customers raised questions over the suspicious redirects. BT faced legal action from customers who were not pleased that their web traffic was compromised.[10] However this has not yet stopped development of business ties between Phorm and a few major UK ISPs in order to monitor user traffic for targeted advertising.

We are not the first to design privacy-preserving advertising, though the academic literature contains surprisingly few examples. There are a number of patents that claim to be privacy preserving[11], but many of these are private only by virtue of claiming not to gather explicit PII per se. They can, however, associate user IP addresses with ad events such as views and clicks, and are therefore only weakly private.

---

[6] `http://www.facebook.com`

[7] `http://www.nytimes.com/external/readwriteweb/2009/02/02/`
`02readwriteweb-facebook_sells_your_data.html`

[8] `http://www.nebuad.com/privacy`

[9] `http://www.phorm.com/`

[10] `http://www.theregister.co.uk/2008/03/17/bt_phorm_lies/`

[11] US patents 6085229, 6182050, and 6370578.

Juels [4] designed a strongly private system as early as 2001. Like Privad, Juels' system assumes an agent running on the client. Unlike Privad, however, Juels made the assumption that pushing ads to the client was not feasible (and indeed in 2001 it certainly was not as feasible). Therefore, Juels proposes a system whereby the client requests a certain type of ad, and the broker supplies that ad. To preserve privacy, however, Juels places an untrusted mixnet between the client and the broker. Mixnets such as the one that Juels proposes are robust against collusion, and thus go further than Privad, which requires monitoring. However, mixnets are complex and heavyweight, and introduce delays into the ad serving process. As a result, it doesn't strike us as attractive an approach as Privad.

Kazienko and Adamski [5] propose AdROSA, a method for displaying web banners using usage mining based on processing of all user's HTTP request during visits to publisher's website. This method was suggested to solve the problem of automatic personalization of web banner advertisements with respect to user privacy and recent advertising policies. It is based on extracting knowledge from the web page content and historical user sessions as well as the current behavior of the online user, using data-mining techniques. Historical user sessions are stored in the form of vectors in the AdROSA database and they are clustered to obtain typical, aggregated user sessions. The site content of the publisher's web pages is automatically processed in a similar way. The system extracts terms from the HTML content of each page. However passing on such information to the advertisement broker may generate the same privacy concerns experienced by efforts of Phorm and NebuAd.

Similar issues as text advertisement have been seen in banner advertisements. Claessens *et al.* [2] highlight some of the privacy issues with current advertising model based on banner displays, which over time could be linked to individual users using server-side analysis of the cookies stored on user's system. In addition, there is a security threat from the publisher or users to cooperate in order to increase the click rate, hence claiming more money from the advertiser. They suggest the separation of profile managers from user by use of anonymous networks and web proxies. The profile management side can be done at the user end. They suggest a client-side banner selection solution will not be practical in terms of scalability of updates.

5   Future Directions and Conclusions

In this paper we have highlighted the privacy issues with current advertising schemes and outline a new system, Privad, based on a variant of adware systems. Specifically, Privad runs a software agent at the client that locally profiles the user and serves out ads, thus preserving the privacy of the users. Privad, however, operates within the current economic framework of advertising consisting of advertisers, publishers, and brokers. Ad events like view and clicks are reported to brokers and advertisers in such a way that no single system component needs

to be trusted individually, and that the privacy of the system can be verified with high confidence.

Although Privad is promising, there are a number of technical challenges ahead.

We have already mentioned the issue of scalable ad and bid dissemination, and we are continuing research in this area.

We believe that the Privad model engenders a wide range of advertisement selection approaches. At the simple extreme, the agent could simply monitor which search terms the user types into the popular search engines, and later match these with keywords associated with advertisements. On the other extreme, the agent could do sophisticated demographic profiling of the user, and advertisers could bid on fine-grained demographics. Taken to its logical limit, this approach could end up revealing private information in another way. For instance, imagine that profiling was so detailed that an advertiser could target an ad to the demographic "families with two or more children under the age of 10, living in large metropolitan cities, with income between $80,000 and $100,000". In this case, when a user clicks on such an ad, the advertiser immediately knows a great deal about the user. Overall, we need to strike a balance between the granularity of targeting and privacy. This is another ongoing topic of research.

There are claims that augmenting social networks with online markets places improves trust between transactions and increases user satisfaction. The advantages of a social-based system over a reputation-based system have been studied previously, showing that malicious users can still boost their transaction-based reputation, while they can not improve their social-based trustworthiness [1]. An interesting avenue for more targeted profiling is also to perform distributed cluster detection algorithms on profiles based on their interests. Such methods have also been developed in order to present better search results by linking to users' social network group [3]. Ironically, such an approach may overcome the privacy concerns with highly targeted ads. If a private reputation system can be devised, then all that is known by the advertiser is that a user who clicked on an ad has something in common with some other user who also clicked on the ad. The advertiser, however, doesn't know what that something is.

Another major problem in advertising systems today is click fraud. For Privad to be viable, it must be at least no more susceptible to click fraud than existing advertising systems, and ideally it improves in this area. While this is also ongoing research, we observe that since the agent can monitor user behavior, it is potentially in a good position to detect clickfraud. Making this work depends on being able to protect the agent software from being tampered with.

Finally, it is worth mentioning that our approach is not limited to web browsers. The system can serve ads in a variety of ways, for example in virtual reality games, Internet chat conversations, or embedded in other applications. This variety of channels makes our approach well suited for future evolution of advertising systems.

References

1. R. Bhattacharjee and A. Goel. Avoiding ballot stuffing in ebay-like reputation systems. In *P2PECON '05: Proceedings of the 2005 ACM SIGCOMM workshop on Economics of peer-to-peer systems*, pages 133–137, New York, NY, USA, 2005. ACM.
2. J. Claessens, C. Diaz, R. Faustinelli, and B. Preneel. A secure and privacy-preserving web banner system for targeted advertising, 2003.
3. K. P. Gummadi, A. Mislove, and P. Druschel. Exploiting Social Networks for Internet Search . In *Proc. 5th Workshop on Hot Topics in Networks*, pages 79–84, Irvine, CA, 2006.
4. A. Juels. Targeted advertising ... and privacy too. In *CT-RSA 2001: Proceedings of the 2001 Conference on Topics in Cryptology*, pages 408–424, London, UK, 2001. Springer-Verlag.
5. P. Kazienko and M. Adamski. Adrosa–adaptive personalization of web advertising. *Information Sciences*, 177(11):2269 – 2295, 2007.
6. B. Krishnamurthy and C. Wills. On the leakage of personally identifiable information via online social networks. In *WOSN '09: Proceedings of the second workshop on Online social networks, Barcelona, USA*.
7. X. Zhang, J. Liu, B. Li, and T. Yum. Coolstreaming/donet: A data-driven overlay network for efficient live media streaming. In *IEEE Infocom 2005*, Miami, 2005.