Universidad Carlos III de Madrid
e-Archivo

Institutional Repository

A First Step Towards User Assisted Online Social Networks

Michal Kryczka
IMDEA Networks
michal.kryczka@imdea.org

Ruben Cuevas
Univ. Carlos III de Madrid
rcuevas@it.uc3m.es

Carmen Guerrero
Univ. Carlos III de Madrid
guerrero@it.uc3m.es

Eiko Yoneki
University of Cambridge
eiko.yoneki@cl.cam.ac.uk

Arturo Azcorra
Univ. Carlos III de Madrid
azcorra@it.uc3m.es

ABSTRACT

The current Online Social Networks' infrastructure is composed by thousands of servers distributed across data-centers spread over several geographical locations. These servers store *all* the users' information (profile, contacts, contents, etc). Such an infrastructure incurs high operational and maintenance costs. Furthermore, this may threaten the scalability, the reliability, the availability and the privacy of the offered service. On the other hand this centralized approach gives to the OSN provider full control over a huge amount of valuable information. This information constitutes the basis of the OSN provider's business.

Most of the storage capacity is dedicated to store the user's content (e.g. photos, videos, etc). We believe that OSN provider does not have strong incentive to dedicate a large part of its infrastructure to store majority part of this content.

In this position paper we introduce the concept of *user assisted Online Social Network* (uaOSN). This novel architecture seeks to distribute the storage load associated to the content (e.g. photos, videos, etc) among the OSN's users. Thus the OSN provider keeps the control on the relevant information while reducing the operational and maintenance costs. We discuss the benefits that this proposal may produce for both, the OSN provider and the users. We also discuss the technical aspects to be considered and compare this solution to other distributed approaches.

Categories and Subject Descriptors

E.1 [**Data Structures**]: Distributed data structures; H.2.4 [**Systems**]: Distributed databases

General Terms

Management

Keywords

Distributed Online Social Networking

1. INTRODUCTION

In the last years Online Social Networks (OSNs) have irrupted in the Internet becoming the most popular applications. For instance, Facebook counts with more than 350M registered users and MySpace has 260M of users. This tremendous success has attracted the interest of the research community that is giving the first steps into the understanding and the improvement of the OSNs.

The current infrastructure of OSNs is composed by dozens of thousands of servers spread across different datacenters in different geographical locations around the globe. The OSNs store *everything* in this infrastructure: user's profile, user's social contacts, user's content, etc. This gives full control to the OSN provider over precious information about millions of users. This is the basis of its business. On the downside, such a centralized infrastructure has several problems:

- Poor scalability: The larger the number of users and the higher activity of them, the larger the centralized infrastructure has to be.

- High cost: In order to maintain a large number of users and the derived data, the OSN provider needs to set up a large number of servers and datacenters facilities. It is well known that these facilities are becoming more and more expensive due to cooling, electricity and traffic among other factors.

- Single Point of failure: It can be easily a victim of DoS attacks [1]. In addition, storing data exclusively in a single location can lead to information loses [2].

In this position paper we introduce the concept of *user assisted Online Social Networks (uaOSNs)*. We propose a decentralized storage architecture where the content which is not strictly needed to be possessed by the provider is stored by the users. Since the main functionality of an OSN is allowing the users to exchange information and interact among them, this model seems to be suitable. uaOSN permits to keep (at least) the same level of functionality while limiting the provider's operational and maintenance cost. Furthermore, uaOSN does not affect the business model of OSN providers, because only the least relevant information is distributed.

The rest of the paper is organised as follows. Section 2 presents more in detail difficulties of the current centralised approach of OSN. In Section 3 we introduce the architecture of uaOSN. We present conceptual and technical description of new model and we outline the possible benefits it can bring. Section 4 provides brief information about related work before concluding in Section 5.

2. LIMITATIONS OF CENTRALIZED OSN

In this section we present the limitations of the current centralised approach of OSNs. Furthermore, we highlight two main areas in which user assisted approach could bring benefits: (i) maintenance and operational costs reduction and (ii) availability and reliability improvement. We focus on Facebook since it is the main player on OSN market, however our claims can be extended to most of OSN providers.

2.1 Maintenance cost control

One of the main feature of Facebook is to allow users to upload their photos and videos and share them with their friends. As for November 2009, Facebook stores 20 billion of unique photos in 4 resolutions which gives 80 billions of stored photos [3]. Furthermore, 2 billions of photos and 14 millions of videos are uploaded to the site each month. Such a situation obviously generates tremendous costs connected with storing the content and threads the scalability of the system. In April 2008, Facebook possessed 10,000 servers [4]. But the growing number of users and enormous amount of data to store made it necessary to borrow $100 million in May 2008 and next $100 million in March 2009 in order to purchase more servers resulting with 30,000 servers as for October 2009 [5].

However it is not only a matter of purchasing the servers but also of maintaining them. The company needs to lease or to build data centers with the necessary facilities. In addition, cooling and electricity costs are an important issue in current data centers [6]. It is estimated that Facebook spends around $1 million per month only for energy consumption (data from October 2008) [7]. Moreover the average cost of cooling data centers is similar to the cost of powering its servers [6]. In summary Facebook would need at least $2 million per month just for cooling and powering its servers.

On the other hand, OSN generates a huge amount of traffic. Facebook has 200 billion page view monthly and 2 billion pieces of content are uploaded every week. Users spend more than 8 billion minutes per day on Facebook generating 3.9 trillion feed actions and 1 billion chat messages per day [3]. This potentially produces large cost of transit traffic. To overcome this problem Facebook relies on Akamai CDN. Although this cuts the traffic bill it still represents an important cost.

We believe that a distributed storage approach like uaOSN can alleviate the storage load of the private OSN infrastructure, thus the number of required server can be smaller. Therefore the operational and maintenance costs are reduced. In addition, a distributed approach can partially or even totally substitute CDN function.

2.2 Reliability, Availability and Security

The next aspect in which distributing OSNs can bring benefits is reliability and availability. In January 2008 Watch-Mouse (a company which deals with website performance monitoring) carried out a set of tests to evaluate the reliability and availability of 104 social networking sites [8]. They measured the waiting time experienced by site visitors. Almost 50% of webpage result with score more than 1000 (where 500 means good performance, 1000 bad one and 1500 means serious problems for the user). The worst OSN in case of availability was the most popular Facebook which scored 6629 points.

Moreover, the OSNs have been a target of network attacks. For instance, in August 2009 Facebook and Twitter were the victims of Denial of Service attacks, which resulted in a degraded service for some of the Facebook users and putting offline Twitter [1].

Another Facebook's availability problem occurred in October 2009 when some of the Facebook accounts were inaccessible during several days and, what is worse, some of the content of about 150k users were lost [2].

Using a distributed architecture into OSN can mitigate the aforementioned problems. On one hand it may prevent DoS attacks since there is no single point of attack. However, it can occur that we will have several weak points of attack instead. The safety of whole system will increase but some additional system should be implemented to prevent attacks on the content of particular user. Furthermore it is more unlikely to suffer from data loses as content is replicated in multiple locations. On the other hand by distributing the content in an intelligent way (i.e. locality strategies) we can also improve the reliability and the availability and reduce the access time.

3. USER ASSISTED ONLINE SOCIAL NETWORKS (UAOSNS)

In this Section we describe the model of user assisted Online Social Networks. In more detail we present: (i) the concept of uaOSNs, (ii) the technical aspects of the solution, (iii) the enhancements over centralised approach and (iv) possible business solutions.

3.1 Conceptual description

In traditional Online Social Networks (like Facebook, Okrut etc.) all users content is stored in the facility of the OSN provider. This creates costs (data storage, data transfer, etc.) which raise proportionally with the number of OSN users and their increasing activity.

The main idea behind uaOSNs is to distribute the content among the facilities of users. We claim that this architecture allows to significantly reduce the operational cost of OSNs providers. The type of the content which could be distributed among the users has to be carefully chosen to not affect the business model of the OSN provider. We believe that distributing data like photos or videos of the users fulfils this criteria while largely reducing the storage overhead of the OSNs provider. Figure 1 depicts the idea of uaOSN. In traditional OSNs all queries are addressed to the service access point of the OSN provider and then to the server which stores the content. In uaOSN query is addressed to the provider which informs user about place of the storage of the content. The variation where some of the queries are directly addressed between the users is also possible.

Moreover, the distributed storage of the content should not affect the user experience. This means that the functionality of the OSN (especially availability of the content
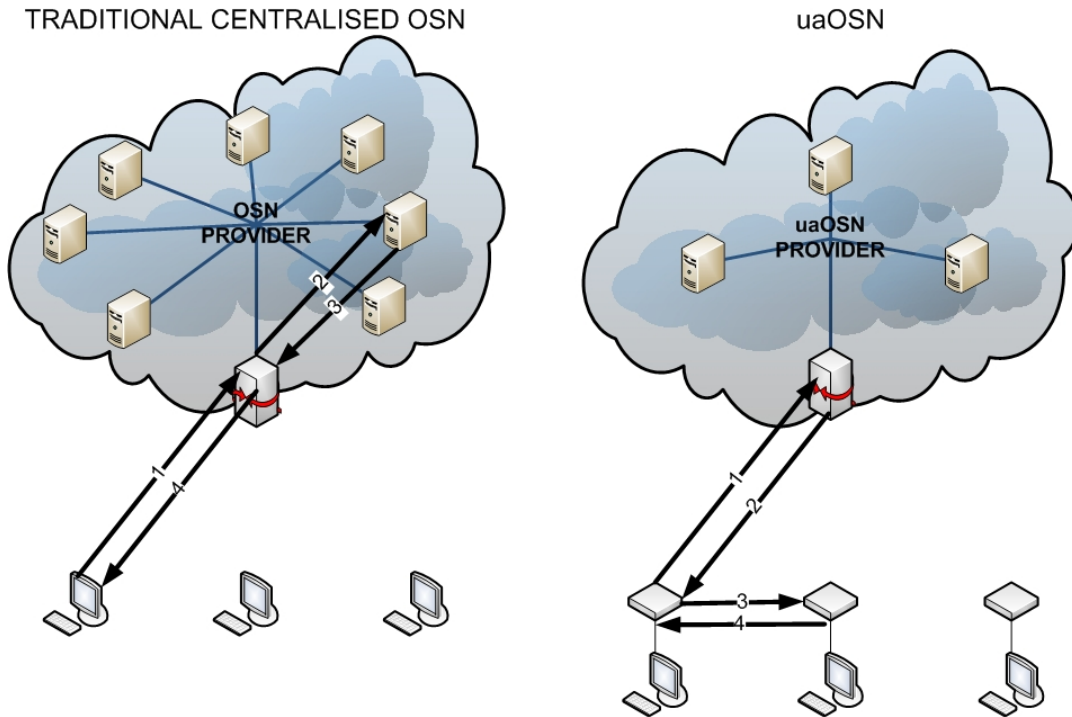
Figure 1: Comparison of traditional OSN and uaOSN

and content privacy) should remain the same from the point of view of the user. Furthermore we believe that user experience can even be improved. For instance we can achieve a higher availability due to locality replication techniques and also higher reliability because of a data replication scheme.

3.2 Technical description

In this section we describe the basic model of uaOSN. However, some modifications are possible according to the needs of the particular OSN provider and service.

In uaOSN architecture the storage overhead of the OSN provider is largely reduced while keeping the control over relevant information. The uaOSN provider stores the basic user information such as user profile and social graph. The content of the user (like photos, videos, wall messages, private messages etc.) are stored in a distributed fashion. However, according to the provider preferences and needs (reliability or censorial purposes), some additional data can be chosen to be stored in a central manner. For example, Facebook portal generates home webpage for each user which contains last activity of the user's friends. The information needed to generate such a page for a particular user could be stored in the provider's facility.

3.2.1 Storage placing

In this section we discuss several distributed storage strategies.

Firstly, the content can be directly stored on the desktop machine of the users. The obvious advantage of this solution is its relatively low cost of deployment. The OSN provider should prepare proper software which is installed and run on the user's machine. On the downside this approach suffers from availability issues. Desktop machines are usually not powered 24 hours per day thus they produce lots of churn.

Additionally they are not reliable. In order to overcome these issues replication techniques needs to be implemented. It produces data transfer overhead and storage overhead (a large number of concurrent replicas need to be maintained).

A second approach is to use the set-top boxes/residential routers to store the OSN data. These devices are expected to be equipped with a hard drive in the close future [9]. Due to the fact that routers are usually not turned off by the users, the availability of the data is much higher in this case. On the downside, this solution, is more costly than previous one in terms of deployment.

A third possibility is using a paid storage services (like Amazon S3 [10]). The obvious advantage of this solution is high availability of the data. However, this architecture forces users to pay for storage space. Furthermore, the granularity of locality techniques is affected (for instance Amazon S3 service only offers three location: US, US West and EU).

Finally, all of the previously described approaches can be mixed and used in parallel.

3.2.2 Replication scheme

Independently of the storage facility used (desktop machine or set-top box), their reliability cannot be compared with server-class machines placed in server farms. Additionally there is a large portion of users that can not participate in the distributed storage infrastructure (e.g. mobile users). This implies that to assure the same availability of the content as with a centralized OSN, we need to apply smart replication scheme.

We need to consider several aspects when choosing a strategy for content replication. First aspect is locality. Content should be replicated in a node close to the consumer. This improves availability (access delay) and reduces the tran-

sit traffic. Assuming that all the queries are addressed to uaOSN provider, this one is able to identify the location of the consumer of given content. Based on historical data and frequency of the requests provider can decide the optimum location of the replica in terms of accessibility. The performance of locality techniques improves with the size of the distributed system. That is why they are appropriate for large-scale distributed systems such as BitTorrent [11][12]. We believe that they could be successfully applied in OSNs like Facebook for example, which comes with 350M of the users.

In second aspect we consider social graph. The node should prefer set-top boxes of friends to store its content. For this purpose system should have information regarding to the owner of each set-top box. This solution has several advantages. If the friend needs a data about the user, it can occur it is already stored in its set-top box. Moreover, this solution relax the privacy constraints: the content is stored in the facility of the user who anyway has access to it.

In case when data is replicated in a stranger's node, it needs to be ciphered for privacy purpose thus a security scheme is a must. This scheme should easily allow to add/revoke users, to distribute keys to friends, etc and should not produce significant overhead. To facilitate this scheme, the OSN provider could play the role of trusted third-party (issuer of the keys). Because of the nature of data exchange in OSN (one person allows to see the content to the group of friends) security mechanisms designed for multicast could be adopted [13][14].

3.2.3   uaOSN provider role

In uaOSN architecture provider works as *content indexer*. It is aware and participate in the decision of placing the content. In a most conservative approach all the queries are addressed to the OSN private server, that redirect the user to the actual location of the data. This give full control to the provider on the OSN activity. The level of availability of this approach is equivalent to the one offered by centralised approaches. A possible variation, is to address some of the queries to the OSN central entity whereas some others are directly exchanged between users.

It is worth noting that since we have central indexing service our solution is robust to user mobility (change of IP address) and churn (when a router disconnects and connects to the system after a while, it is likely to do it with a different IP address). In uaOSN, if this occurs the device informs the provider about the new IP address and then the provider updates the index.

For privacy reasons the uaOSN provider acts as a trusted third-party and key issuer. This allows the provider to perform censorship tasks.

The rest of the current functions of the provider are not affected by the usage of our proposal.

3.3   Business model

The key point of the uaOSN architecture is the usage of home user set-top boxes/routers with hard drive facility. This device could be partially or fully subsidized by the OSN provider. As a result, the user could get a useful device for an attractive price while OSN provider could benefit from using it for content storage. Another way to convince users to use uaOSN software in both, set-top boxes and desktops, is introducing a loyalty program. User could get points

for installing uaOSN software (buying set-top box) and then for allowing the provider using it (usage time, shared storage space, size of exchanged traffic, etc.). These points could be then used in the provider's portal (additional storage limit, additional functionality, buying virtual things, etc.)

A second approach is that proposed in [9]. In this case the ISP is the owner of set-top box. The uaOSN provider can rent the service of certain number of set-top boxes from different ISPs.

The first model has higher deployment cost but allows uaOSN to construct its own distributed infrastructure. Whereas the second option has no deployment cost but the uaOSN highly depends on ISPs. Finally we would like to highlight that both approaches can be funded from the savings on data center infrastructure and maintenance described in Section 2.

3.4   Summary of improvements

We claim above that using uaOSNs can bring benefits to both: OSN providers and end users. Next we summarize this benefits. uaOSN provider benefits are:

- *reducing infrastructure cost*: since the major part of the storage overhead (photos, videos, etc) is distributed, providers can reduce the number of private servers dedicated to data storage. This also reduces those cost connected with maintaining data centers (energy, cooling, personnel, etc).

- *reducing content delivery costs*: on one hand using appropriate locality replication strategies reduces the amount of traffic crossing transit links. On the other hand the uaOSN offers to the provider the possibility of creating its own distributed storage infrastructure, thus removing CDN service, or at least mitigating dependency on it.

- *robustness to DoS attacks* The distributed nature of uaOSNs makes them more robust to DoS attacks than its centralised counterpart.

uaOSN users benefits are:

- *higher reliability*: uaOSN uses external and distributed replication of the content. This avoid the possible information loses produced in a centralised storage scheme.

- *better access to the content*: the uaOSN uses locality as first criteria for replication algorithm. Then, it is likely that a given user gather a content replica that is close to it, thus reducing the delay to access the content.

- *higher availability of the service*: In the described case where part of the queries are decentralised to be issued directly between the users, the availability of the service improves since part of the service works even if the central server is inaccessible.

4.   RELATED WORK

In this Section we present other distributed OSN approaches and discuss their differences with our uaOSN.

Vis-a-Vis [15] is a framework for decentralized OSN. It is based on the concept of Virtual Individual Server (VIS), this is a personal virtual machine used for managing and

storing the personal content of the user. The user's VIS runs within some utility computing infrastructure. Furthermore, the VIS nodes form overlay networks which reflect the users' social connections. The whole system is based on a multi-tier DHT structure. The top level, called meta-group, is used to advertise and look for public groups and all VISs belong to it. Other groups are created also in a DHT-manner and they reflect the social groups from the OSN.

A second work [16] presents a system that deals with users' security and privacy. The system contains three main entities: the *matryoshkas*, a trusted identification service, and a peer to peer substrate. Each user has an associated matryoshka which is a logical infrastructure of concentric rings where the user is located in the center (core). The innermost ring is formed by those nodes that are trusted by the user, the next one is composed by nodes trusted by the nodes from the inner ring, and so on. The data of the user is stored in the first ring and can be accessed by other users in a privacy-preserving manner. The trusted identification service is responsible for authentication and the peer to peer substrate (e.g. a DHT) provides the basic mechanisms for accessing to user's data.

PeerSoN [17] is another proposal for P2P social networking. It has a similar goal as Vis-a-Vis. In order to address the problem of privacy PeerSoN resigns from a central entity. The mechanism relies on encrypting the users data. Furthermore the system uses a procedure for key distribution and revocation. This architecture makes impossible for any third entity to mine the data or infer either the users behaviour or users relationships.

An interesting concept is presented in [18]. It proposes creating a personal containers (blog, photos) in the cloud. This content can be accessed by different applications which the user is participating e.g. OSNs or video streaming services, etc. Therefore this is a meta-service that could be used by any of the previous proposals.

It is worth noting that all the described systems are primarily focused on privacy preserving. Their aim is to address the concerns related to the fact that OSN providers concentrate a huge amount of personal data of millions of user in one central place. This includes: (*i*) granting full copyrights to the OSN provider over the content which is upload, (*ii*) the censorship performed by provider and (*iii*) possible data disclosures, etc. All the described systems are non-commercial and self-organized social networks. However, these schemes present some withdraws. The most important one is the availability of the data that can be stored in either storage facilities (e.g. Amazon S3 severs) or the users' PCs. The first approach makes the content fully available but incurs in economical costs for the users. Whereas, the second approach requires replication techniques to guarantee availability what incurs in network costs.

On the other hand, our approach does not exclude the OSN provider from the picture. Instead of decentralizing the full system, we propose to distribute only the storage of the *heavy* content (not relevant for the business model of the OSN provider) that produces high infrastructural and operational costs. In other words, our approach allows to significantly cut the *opex* of the company without loosing any functionality. It also gives benefits to the users such as a better availability and reliability of the service. Furthermore uaOSN is also less complex than the previous systems since many important operations (creating new user, log-

ging, granting and revoking rights) can be easily performed with the help of a trusted central entity.

5. CONCLUSION

We have discussed the problem raised from the central nature of current OSNs. We have presented the concept of uaOSNs in which irrelevant content for the OSN provider (e.g. photos or videos) is distributed among the OSN users. We have shown that this architecture can bring improvements both to the OSN providers and end users. Finally we have compared our proposal to other distributed solutions clarifying the differences.

As future work, we aim to validate the claims presented in this position papers with real traces from commercial OSNs. Furthermore We will also work in the implementation of an uaOSN prototype.

Acknowledgment

6. REFERENCES

[1] [Online]. Available: http://news.cnet.com/8301-27080_3-10305200-245.html?tag=mncol

[2] [Online]. Available: http://news.cnet.com/8301-13577_3-10373349-36.html

[3] [Online]. Available: http://www.jacobsschool.ucsd.edu/news/news_events/event.sfe?id=930

[4] [Online]. Available: http://www.datacenterknowledge.com/archives/2008/04/23/facebook-now-running-10000-web-servers/

[5] [Online]. Available: http://www.datacenterknowledge.com/archives/2009/10/13/facebook-now-has-30000-servers/

[6] V. Valancius, N. Laoutaris, L. Massoulie, C. Diot, and P. Rodriguez, "Greening the internet with nano data centers," *ACM Conext 2009, Rome*, December 2009.

[7] [Online]. Available: http://www.datacenterknowledge.com/archives/2008/10/31/facebook-1-million-a-month-in-power-costs/

[8] [Online]. Available: http://www.watchmouse.com/en/press/Social-networking-sites-slow-and-inaccessible.html

[9] N. Laoutaris, P. Rodriguez, and L. Massoulie, "ECHOS: edge capacity hosting overlays of nano data centers," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 1, pp. 51–54, 2008.

[10] [Online]. Available: http://aws.amazon.com/s3/

[11] D. R. Choffnes and F. E. Bustamante, "Taming the torrent: a practical approach to reducing cross-isp traffic in peer-to-peer systems," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 4, pp. 363–374, 2008.

[12] R. Cuevas, N. Laoutaris, X. Yang, G. Siganos, and P. Rodriguez, "Deep diving into bittorrent locality," *Technical Report, available from: http://arxiv.org/abs/0907.3874*, 2009.

[13] M. Ning-bo, H. Yu-Pu, and O. Hai-wen, "Broadcast encryption scheme based on RSA," *The Journal of China Universities of Posts and Telecommunications*, February 2009.

[14] D. Boneh and M. Hamburg, "Generalized identity based and broadcast encryption schemes," *ASIACRYPT 2008*, 2008.

[15] A. Shakimov, H. Lim, L. P. Cox, and R. Caceres, "Vis-a-Vis: online social networking via virtual individual servers," *Duke University Technical Report TR-2008-05*, October 2008.

[16] L. A. Cutillo, R. Molva, and T. Strufe, "Privacy preserving social networking through decentralization," *Proceedings of WONS 2009, The Sixth International Conference on Wireless On-demand Network Systems and Services, Snowbird, Utah, USA*, February 2009.

[17] S. Buchegger, D. Schioberg, L.-H. Vu, and A. Datta, "PeerSon: P2P social networking - early experiences and insights," *SocialNets 2009, The 2nd Workshop on Social Network Systems, Nuernberg, Germany*, March 2009.

[18] A. Madhavapeddy, R. Mortier, J. Crowcroft, and S. Hand, "Multiscale not multicore: Efficient heterogeneous cloud computing," *ACM-BCS Visions of Computer Science 2010*, 2010.