

The Anatomy of Sindbad: A Location-Aware Social Networking System

Mohamed Sarwat*, Jie Bao*, Ahmed Eldawy*, Justin J. Levandoski†, Amr Magdy*, and Mohamed F. Mokbel*

*Dept. of Computer Science and Engineering, University of Minnesota, Minneapolis, MN 55455

†Microsoft Research, Redmond, WA 98052-6399

*{sarwat,baojie,eldawy,amr,mokbel}@cs.umn.edu, †justin.levandoski@microsoft.com

ABSTRACT

This paper features Sindbad; a location-based social networking system. Sindbad supports three new services beyond traditional social networking services, namely, *location-aware news feed*, *location-aware recommender*, and *location-aware ranking*. These new services not only consider social relevance for its users, but they also consider spatial relevance. Since location-aware social networking systems have to deal with large number of users, large number of messages, and user mobility, efficiency and scalability are important issues. To this end, Sindbad encapsulates its three main services inside the query processing engine of PostgreSQL. Usage and internal functionality of Sindbad are implemented with PostgreSQL and Google Maps API. Both a web and android phone applications are built on top of Sindbad for better interaction with the system users.

Categories and Subject Descriptors

H.2.8 [Database Applications]: Spatial databases and GIS

General Terms

Design, Management, Performance, Algorithms

Keywords

Social Networking, Recommender Systems, News Feed, Spatial Rating, Spatial Message

1. INTRODUCTION

In this paper we feature Sindbad [10]: a location-based social networking system. Sindbad distinguishes itself from existing social networking systems (e.g., Facebook [4] and Twitter [12]) as it injects location-awareness within every aspect of social interaction and functionality in the system. For example, posted messages in Sindbad have inherent spatial extents (i.e., spatial location and spatial range) and users receive friend news feed based on their locations and the spatial extents of messages posted by their friends. The functionality of Sindbad is fundamentally different from current incarnations of location-based social networks (e.g., Facebook

Places [5], Foursquare [7]). These existing location-based social networks are strictly built for mobile devices, and only allow users to receive messages about the whereabouts of their friends (e.g., Foursquare “check-ins” that give an alert that “your friend Alice has checked in at restaurant A”). Sindbad, on the other hand, takes a broader approach that marries functionality of traditional social-networks with location-based social scenarios (e.g., friend news posts with spatial extents, location-influenced recommendations) [3]. Thus, Sindbad is appropriate for both traditional social networking scenarios (e.g., desktop-based applications) as well as location-based scenarios (e.g., mobile-based applications).

Users of Sindbad can entertain one or more of the following functionalities: (a) select their friend list as well as getting listed as friends to other users in a same way like traditional social network systems, (b) post (spatial) messages and/or rate (spatial) objects (e.g., restaurant or movies), which will be seen by their friends, (c) once a user logs on to Sindbad, the user will see an incoming *location-aware* news feed from the user friends. The news feed is selected based on both the user location and the spatial extents of the posted messages, and (d) get a *location-aware* recommendation about spatial items, e.g., restaurants, or non-spatial items, e.g., movies. The recommendation is based on the user location, the item location, and what are the items that the friends of the user have liked. In summary, Sindbad distinguishes itself from all previous systems in one or more of the following aspects:

1. Posted messages in Sindbad have spatial extents in which they are deemed relevant to only those friends who are located within the message spatial extent.
2. Sindbad allows its users to express their opinions by rating different items, e.g., restaurants, movies, or stores.
3. Sindbad is equipped with a *location-aware news feed* module that, for any user, efficiently retrieves the relevant messages from her friends based on the user location and the message spatial extents.
4. Sindbad is equipped with a *location-aware recommender* module that, for any user, efficiently suggests (spatial) items based on users locations, items locations, and previous ratings by user friends.
5. Sindbad is equipped with a *location-aware ranking* module that efficiently selects the top- k relevant objects produced from either the location-aware news feed or the recommender system modules. The ranking is based on both the spatial and social relevance.
6. A major part of Sindbad is built inside PostgreSQL; an open-source DBMS. Hence, Sindbad: (a) takes advantage of the scalability provided by the DBMS, and (b) is able to employ early pruning techniques inside the DBMS engine, which

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ACM SIGSPATIAL LBSN '12, November 6, 2012. Redondo Beach, CA, USA

Copyright 2012 ACM 978-1-4503-1698-9/12/11 ...\$15.00.




Figure 1: Sindbad System Architecture.

yields an efficient performance for the news feed, recommendation, and ranking functionalities.

7. Sindbad provides an RESTful [6] web API so that it would allow a wide variety of applications to easily communicate with Sindbad and make use of its unique features. Both a web and smart phone (i.e., Android) applications are implemented on top of Sindbad to enrich the system user experience.

The rest of the paper is organized as follows: Section 2 gives the system architecture of Sindbad. The three main components of Sindbad, namely, *location-aware news feed (GeoFeed)*, *location-aware recommender system (LARS)*, and *location-aware ranking*, are discussed in Sections 3, 4, and 5, respectively. Finally, section 6 concludes the paper.

2. SINDBAD ARCHITECTURE

Figure 1 depicts the Sindbad system architecture that consists of three main modules, namely, *location-aware news feed (GeoFeed)*, *location-aware ranking*, and *location-aware recommender (LARS)*, and three types of stored data, namely, spatial messages, user profiles, and spatial ratings. The communication between Sindbad and the outside world is held through RESTful web API interface (named *Sindbad API Functions* in Figure 1). The API functions facilitates building a wide variety of applications (e.g., web applications, smart phone applications) on top of Sindbad. As shown in Figure 1, Sindbad API functions can also be used to complement the functionality of existing social networking websites e.g., Facebook, and turn their news feed and recommendation to be location-aware. Sindbad can take five different types of input (i.e., through the API interface): profile updates, a new message, a new rating, a location-aware news feed query, and a location-aware recommender query. The actions taken by Sindbad for each input is described as follows:

Profile updates. As in typical social networking systems, Sindbad users can update their personal information, their friend list, or accept a friend invitation from others.

A new message. Users can post spatial messages to be seen by their friends, if relevant. A spatial message is represented by the tuple: (*MessageID*, *Content*, *Timestamp*, *Spatial*), where *MessageID* and *Content* represent the message identifier and contents, respectively, *Timestamp* is the time the message is generated, while *Spatial* in-

dicates the spatial range for which the message is effective. The message is deemed relevant to only those users who are located within its spatial range.

A new rating. Sindbad users can give location-aware (spatial) ratings to various items in a scale from one to five. Location-aware (spatial) ratings can take any of these three forms: (1) *Spatial ratings for non-spatial items*, represented as a four-tuple (*user*, *userLocation*, *rating*, *item*); for example, a user located at home rating a book, (2) *Non-spatial ratings for spatial items*, represented as a four-tuple (*user*, *rating*, *item*, *itemLocation*); for example, a user with unknown location rating a restaurant with an inherent location, and (3) *Spatial ratings for spatial items*, represented as a five-tuple (*user*, *userLocation*, *rating*, *item*, *itemLocation*); for example, a user at his/her office rating a restaurant with an inherent location.

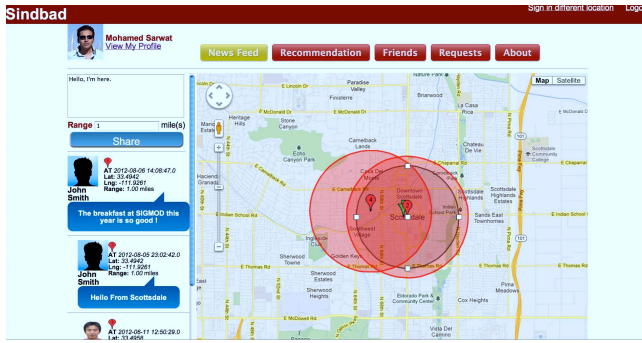
Location-aware news feed queries. Once a Sindbad user logs on to the system, a location-aware news feed query is fired to retrieve the relevant news feed, i.e., messages posted by the user's friends that have spatial extents covering the location of the requesting user. Details of the execution of the location-aware news feed query will be discussed in Section 3. The output of the *location-aware news feed* module (GeoFeed) will be processed further by the *location-aware ranking* module to get only the top-*k* news feed based on the spatial and social relevance, which will be returned to the user as the requested news feed. Details of the *location-aware ranking* module will be described in Section 5.

Location-aware recommendation queries. Sindbad users can request recommendations of either spatial items (e.g., restaurants, stores) or non-spatial items (e.g., movies) by explicitly issuing a location-aware recommendation query. The *location-aware recommender* module (LARS) suggests a set of items based on: (a) the user location (if available), (b) the item location (if available), and (c) ratings previously posted by either the user or the user's friends. Details of LARS will be discussed in Section 4. Similar to location-aware news feed queries, the output of LARS goes through the *location-aware ranking* module to select only the top-*k* items based on both spatial and social relevance.

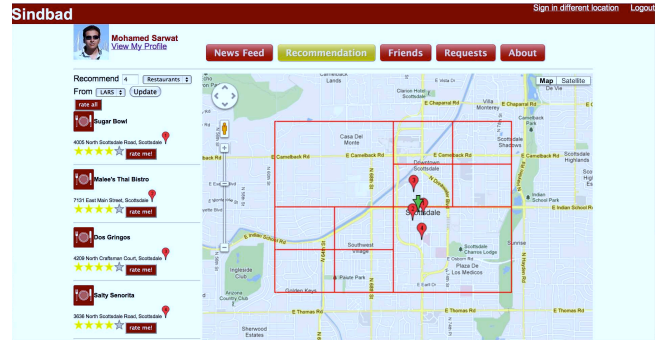
3. LOCATION-AWARE NEWS FEED

Motivation. Although news feed functionality is widely available in all social network systems [11], these systems select the relevant messages either based on the message timestamp or some importance criteria that ignores the spatial aspect of posted messages. Thus, users may miss several important messages that are spatially related to them. For example, when a traveling user logs on to a social network site, the user would like to get the news feed that match his/her new location, rather than receiving the most recent (non-spatial) news feed. The same concept applies for users who continuously log onto the system from the same location, yet have a large number of friends. It is of essence for such users to limit their news feed to the messages related to their location. Examples of the location-aware news feed returned by Sindbad include a message about local news, a comment about a local store, or a status message targeting friends in a certain area.

Contribution. The main idea of the *location-aware news feed* module (GeoFeed) is to abstract the location-aware news feed problem into one that evaluates a set of location-based point queries against each friend in a user's friend list that retrieves the set of messages issued that overlap with the querying user's location. The *location-aware news feed* is equipped with three different approaches for evaluating each location-based query: (1) *spatial pull approach*, in which the query is answered through exploiting a spatial index over the messages posted by the friend, (2) *spatial push approach*, in which the query simply retrieves the answer from



(a) Sindbad News Feed Service

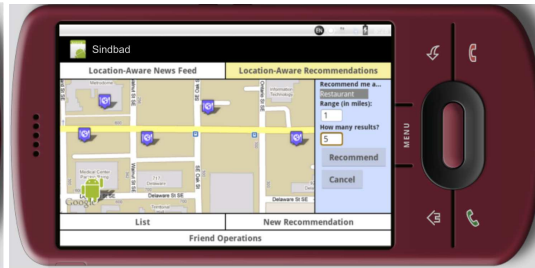


(b) Sindbad Recommendation Service

Figure 2: Sindbad Web Application Screenshots



(a) Location-Aware News Feed



(b) Location-Aware Recommender

Figure 3: Sindbad Android Phone Application Screenshots

a pre-computed materialized view maintained by the friend, and (3) *shared push approach*, in which the pre-computation and materialized view maintenance are shared among multiple users. Then, the main challenge of GeoFeed is to decide on when to use each of these three approaches for a query.

A better response time calls for using the *spatial push* approach for all location-aware news feed queries issued to Sindbad. In this case, all location-aware news feed are pre-computed. However, this approach results in tremendous system overhead since a massive number of materialized views must be maintained. On the other hand, favoring system overhead may result in executing more queries using the *spatial pull* approach as no views need to be maintained. However, this approach may result in a long query response time for users who have a large number of friends, since they will suffer a long delay when retrieving their news feed. Sindbad takes these factors into account when deciding on which approach to use to evaluate each query in a way that minimizes the system overhead and guarantees a certain user response time. Sindbad is equipped with an elegant decision model that decides upon using these approaches in a way that: (a) minimizes the system overhead for delivering the location-aware news feed, and (b) guarantees a certain response time for each user to obtain the requested location-aware news feed. More details about GeoFeed decision model are provided in [2].

Application Dynamics. Figures 2(a) and 3(a) give an example of how the user receives the location-aware news feed service via both the Sindbad web application and phone application. When the user logs on to Sindbad, the application displays the relevant news for the user on the map. Each message is associated with a circle representing the range of each message. The user may “change” locations by dragging and dropping the green arrow on the map, to see how the news feed will change accordingly. The user may also submit a geo-tagged message to the system. For instance, the user

shares a message "The pizza at ABC restaurant is awesome" with a range distance of one mile.

4. LOCATION-AWARE RECOMMENDER

Motivation. This section describes the *location-aware recommender* module [8, 9] of Sindbad. In general, recommender systems make use of community opinions to help users identify useful items from a considerably large search space (e.g., Amazon inventory, Netflix movies). The technique used by many of these widely deployed systems is collaborative filtering [1], which analyzes past community opinions to find correlations of similar users and items to suggest a set of personalized items to the querying user. Community opinions are expressed through explicit ratings represented by the triple (*user, rating, item*) that represents a user providing a numeric rating for an item (e.g., movie). Unfortunately, ratings represented by this triple ignore the fact that both users and (some) items are spatial in nature. Unlike traditional recommendation techniques that assume the (non-spatial) rating triple (*user, rating, item*), the *location-aware recommender* module (LARS) in Sindbad supports a taxonomy of three types of location-based ratings: (1) *Spatial ratings for non-spatial items*, represented as a four-tuple (*user, ulocation, rating, item*), where *ulocation* presents the user location, (2) *Non-spatial ratings for spatial items*, represented as a four-tuple (*user, rating, item, ilocation*), where *ilocation* presents the item location, and (3) *Spatial ratings for spatial items*, represented as a five-tuple (*user, ulocation, rating, item, ilocation*), where *ulocation* and *ilocation* present the user and item locations, respectively.

Contribution. Sindbad produces recommendations using *spatial user ratings for non-spatial items* by employing a user partitioning technique that exploits the user location embedded in the ratings. This technique uses an adaptive pyramid structure to partition ratings by their user location attribute into spatial regions of varying sizes at different hierarchies. Then, for a querying user located in a region \mathcal{R} , we apply an existing collaborative filtering technique [1]

that utilizes only the ratings located in \mathcal{R} . The challenge, however, is to determine whether all regions in the pyramid must be maintained in order to balance two contradicting factors: scalability and locality. Maintaining a large number of regions increases locality (i.e., recommendations unique to smaller spatial regions), yet adversely affects system scalability because each region requires storage and maintenance of a collaborative filtering data structure (i.e., model) necessary to produce recommendations. The pyramid dynamically adapts to find the right pyramid shape that minimizes storage overhead, i.e., increases scalability, without sacrificing locality. When deciding to merge or split a pyramid cell, we define a system parameter \mathcal{M} , a real number in the range [0,1] that defines a tradeoff between scalability gain and locality loss. LARS merges the child cells into the parent cell if:

$$(1 - \mathcal{M}) * scalability_gain > \mathcal{M} * locality_loss \quad (1)$$

Sindbad produces recommendations using *non-spatial user ratings for spatial items* by employing a travel penalty technique that accounts for item locations embedded in the ratings. This technique favors recommendation candidates the closer they are in travel distance to a querying user. The challenge here is to avoid computing the travel distance for all spatial items to produce the list of recommended items, as this will greatly consume system resources. Sindbad addresses this challenge by employing an efficient query processing framework capable of terminating early once it discovers that the list of recommended items cannot be altered by processing more candidates. Sindbad employs an efficient query processing framework capable of terminating early once it discovers that the list of recommended items cannot be altered by processing more candidates. Finally, to produce recommendations using *spatial ratings for spatial items*, Sindbad employs both the user partitioning and travel penalty techniques to address the user and item locations associated with the ratings. More details about LARS are given in [8].

Application Dynamics. Figures 2(b) and 3(a) gives an example of how the user interacts with the location-aware recommendation service through both the Sindbad web and phone applications. The user can ask Sindbad for recommendations (e.g., Restaurants in Scottsdale where SIGMOD takes place) by clicking on the recommendation tab of the web interface or by clicking on the location-aware recommendation button in the mobile app. The user then enters the type of object he is interested in (e.g., restaurant, theaters, stores) a spatial range in miles, and also the number of recommended items to be returned to him and then presses the *Recommend* button. The recommended items are then shown on the map.

5. LOCATION-AWARE RANKING

Sindbad users may have different preferences over messages from the location-based news feed or recommender system. For example, a traveling user may be more interested in messages that were issued close to her current locations. On the other hand, the stationary user maybe more interested in the most recently issued messages. Moreover, due to the large volume of messages submitted to Sindbad and the user's limited viewing capability (e.g., 40 messages for the web page and 20 messages for mobile application), Sindbad provides a *location-aware ranking* module that is responsible for ranking the results coming out of the *location-aware news feed* module and *location-aware recommender* module, based on the user's preferences.

GeoRank ranks the messages based on multiple domains, i.e., temporal domain and spatial domain. Each user can also specify a preference parameter ω to indicate her preference over the temporal

or spatial domain (e.g., $\omega = 1$ indicates the user cares only for the close messages, and $\omega = 0$ indicates the user cares only for the recent messages). The final rank score will be calculated based on the following equation:

$$\text{Rank Score} = \omega \times \text{Spatial Score} + (1 - \omega) \times \text{Temporal Score}$$

Instead of ranking all objects, Sindbad *location-aware ranking* module encapsulates the user's ranking preferences within the query processor to improve the response time for the user. Moreover, if the user is continuously online, Sindbad location-aware ranking module is responsible for keeping either the news feed or the recommended items correctly ranked (i.e., continuously evaluating the ranking function) as the user moves or if the user changes her preferences. The main idea is to reduce the number of friends the system querying for the relevant messages. Because we notice that for a top-k preferred location-based news feed query, it needs to retrieve the relevant messages from at most k friends. Thus, we track the highest ranked message scores from all the user friends to select a candidate set of friends to retrieve the relevant messages. Then, GeoRank queries friends by their highest ranked message scores and keeps updating the ranking boundary for top-k candidate messages. The query processing terminates early once the highest ranking score of the next querying friend is less than the ranking boundary to save the significant amount of redundant computations.

6. CONCLUSION AND FUTURE WORK

In the paper, we gave an overview of Sindbad system architecture. Sindbad is a location-aware social networking system that considers the spatial location as first class citizen within every aspects of social networking services. We introduced the three main modules that constitute Sindbad: (1) Location-Aware News: that delivers location-aware messages to the user, (2) Location-Aware Recommendation: that generates recommendations based on both the users and items spatial locations, and (3) Location-Aware Ranking: that incorporates the spatial location in ranking both the user news and recommendations. In the future, we plan to plug-in more social services inside Sindbad to enrich the user experience.

7. ACKNOWLEDGEMENT

This work is supported in part by the National Science Foundation under Grants IIS-0811998, IIS-0811935, CNS-0708604, IIS-0952977 and by a Microsoft Research Gift.

8. REFERENCES

- [1] G. Adomavicius and A. Tuzhilin. Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions. *TKDE*, 17(6):734–749, 2005.
- [2] J. Bao, M. F. Mokbel, and C.-Y. Chow. GeoFeed: A Location-Aware News Feed System. In *ICDE*, 2012.
- [3] C.-Y. Chow, J. Bao, and M. F. Mokbel. Towards Location-based Social Networking Services. In *ACM SIGSPATIAL-LBSN*, 2010.
- [4] Facebook. <http://www.facebook.com/>.
- [5] The Facebook Blog, "Facebook Places": <http://tinyurl.com/3aetfs3>.
- [6] R. T. Fielding and R. N. Taylor. Principled Design Of The Modern Web architecture. In *ICSE*, 2000.
- [7] Foursquare: <http://foursquare.com>.
- [8] J. J. Levandoski, M. Sarwat, A. Eldawy, and M. F. Mokbel. LARS: A Location-Aware Recommender System. In *ICDE*, 2012.
- [9] M. Sarwat. Recdb: Towards dbms support for online recommender systems. In *SIGMOD/PODS PhD Symposium*, pages 33–38, 2012.
- [10] M. Sarwat, J. Bao, A. Eldawy, J. J. Levandoski, A. Magdy, and M. F. Mokbel. Sindbad: a location-based social networking system. In *SIGMOD*, 2012.
- [11] A. Silberstein, J. Terrace, B. F. Cooper, and R. Ramakrishnan. Feeding Frenzy: Selectively Materializing User's Event Feed. In *SIGMOD*, 2010.
- [12] Twitter. <http://www.twitter.com/>.