Enabling Social Notifications at Mass Scale

Contact: Prof. Roman Vitenberg (romanvi@ifi.uio.no)

Motivation and background:
Electronic services nowadays are typically accompanied by discussion forums and similar social communities. To enhance the main functionality, services deliver social notifications to such communities or individual users. For example, the main music streaming service of Spotify is accompanied by an engine that delivers real-time notifications to each online user, as shown in the illustrations below. These notifications report when a user's friend comes online and goes offline, what music track a user's friend listens to, when a user's friend updates a shared content, such as a playlist, etc. Similar delivery of social notifications is provided alongside most real-time cooperative activities: video and audio streaming, new content uploaded by friends on Facebook, mobile and ad-hoc interaction, online gaming, large-scale human gatherings during the commute, Olympic games, exhibitions, as well as many others.

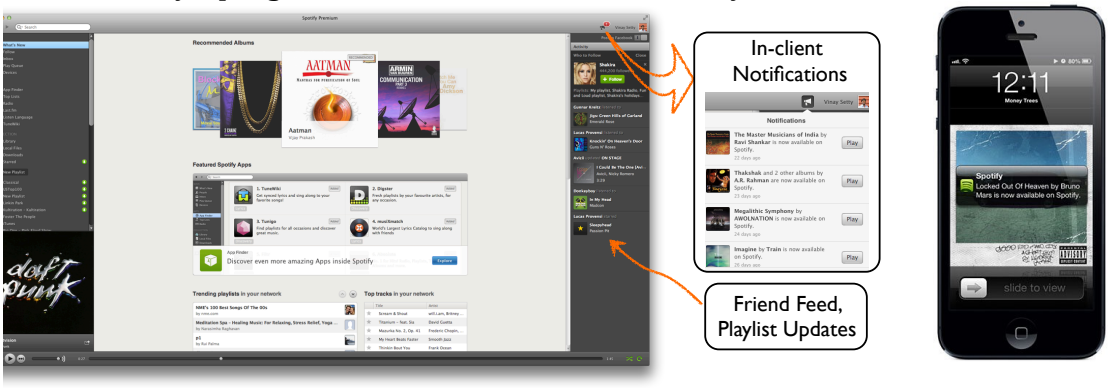

The volume of social notifications is quite staggering. For example, the notification service at Spotify delivers over a billion notifications every single day, with a total volume of over 2 terabytes [1]. The scale makes it imperative to apply big data science techniques of data analytics and machine learning.

However, the technology for deploying services for social notifications at mass scale is not in place yet. The first challenge is that it is non-trivial to measure and model quality of user experience thanks to delivered notifications [2]. Secondly, notification services have traditionally been deployed on in-house enterprise clusters. With the advent of cloud computing, a viable alternative of running such services in the cloud became available. An enterprise may choose between using a generic notification engine (such as Azure Service Bus or PubNub included in Microsoft Azure and Amazon EC2, respectively) and moving the deployment of its proprietary engine optimized for the application needs to the cloud. Unfortunately, neither alternative is satisfactory at present, as explained in [3]: generic notification engines do not take into account the specifics of social notifications whereas proprietary engines are not designed to operate atop cloud services. The questions of cloud resource allocation and cost become critical in this context, yet they have never been considered for social notification services.

[1] V. Setty, G. Kreitz, R. Vitenberg, M. van Steen, G. Urdaneta, S. Gimåker: The Hidden Pub/sub of Spotify, DEBS 2013.

[2] M. Uddin, V. Setty, Y. Zhao, R. Vitenberg, N. Venkatasubramanian: RichNote: Adaptive Selection and Delivery of Rich Media Notifications to Mobile Users, ICDCS 2016.
[3] V. Setty, R. Vitenberg, G. Kreitz, G. Urdaneta, M. van Steen: Cost-Effective Resource Allocation for Deploying Pub/Sub on Cloud, ICDCS 2014.

Main objective and summary:
Devise novel concepts and mechanisms for deploying social notification services. A particular emphasis is on the user satisfaction due to delivered notifications and on allocation of resources. We plan to use state-of-the-art data science and machine learning techniques to design models for user satisfaction. The models will be constructed by profiling large-scale user activities over time. Using these models, we will develop replication strategies that consider user satisfaction metrics in addition to scalability, efficiency and speed of notification delivery.

We plan to build on our long-term collaboration and past research with Spotify to carry out this project. While the main goal of this research is to help enterprises that move their operation to the cloud, the problem is also beneficial for minimizing resource consumption when using in-house deployment.

What we offer in a nutshell
- A strong research environment. Our students have won best paper and best demo awards[1] at several conferences. In addition to Spotify, our alumni are employed by IBM Research, Microsoft, and highly reputable academic institutions in Europe.
- The group of Networks and Distributed Systems[2] offers a work environment that is well equipped with the newest hard- and software technology. The research group has tight bonds with Simula Research Laboratory. Furthermore, we have well-established links to national and international research institutions. We conduct collaborative research projects that are funded by Norwegian research funds, and the European Community.
- Opportunities for research stays with our renowned collaborators worldwide, including University of Toronto, TU Munich, and UC Irvine.
- Well-paid PhD positions, in a country which has been ranked by the UN as having the highest standard of living in the world, and which is known for its unique scenic beauty. The work is in a smart futuristic building that has won multiple awards.

Suitable Background and Requirements:
- Applicants must have a degree in Computer Science, or in a related study, with excellent results. They must also be able to demonstrate interest in scientific research. The evaluation considers many aspects of excellence, as well as the personal drive and organizational skills. The ideal candidate for the position will have strong background in distributed computing.
- You may apply if you have not yet completed your degree, but expect to do so before the position starts.
- Knowledge of Norwegian is not a prerequisite for application. English is our working language for research.

[1] http://www.mn.uio.no/ifi/om/aktuelt/utmerkelser/acm-debs-poster-award-2014.html
[2] http://www.mn.uio.no/ifi/english/research/groups/nd/index.html