



WhatsUp: news from, for, through everyone

Antoine Boutet, Davide Frey, Rachid Guerraoui, Anne-Marie Kermarrec

► To cite this version:

Antoine Boutet, Davide Frey, Rachid Guerraoui, Anne-Marie Kermarrec. WhatsUp: news from, for, through everyone. 2010. inria-00468272

HAL Id: inria-00468272

<https://hal.inria.fr/inria-00468272>

Submitted on 30 Mar 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

WhatsUp: news from, for, through everyone

A. Boutet¹

Collaborators: D. Frey¹, R. Guerraoui² & A.-M. Kermarrec¹

¹ INRIA Rennes Bretagne-Atlantique - ² EPFL

EuroSys - April 13rd 2010

This work is part of the gossple project an ERC starting grant <http://gossple.fr>.

Abstract

WHATSUP (WUP) is a new form of electronic news which is decentralized and personalized. No big brother company controls the news, and no central server makes it vulnerable to failures, censorship or attacks. WUP does not rely on any explicit subscription scheme but leverages the interest expressed by the readers of the news to personalize its dissemination. Basically, WUP users express their interests simply by giving opinions about the news they read. This opinion, in turn, impacts the news these users receive in the future.

Underlying WUP lies BEEP, a biased epidemic dissemination protocol that delivers news to interested users in a fast manner and using little bandwidth, despite jamming and churn. BEEP is parameterized on the fly on a per user, per news and per dissemination hop basis. This protocol has two key characteristics: *orientation* and *amplification*. Every user forwards the news of interest to a randomly selected set of users with a preference towards those that have similar interests (*orientation*). The notion of interest does not rely on any explicit social network or subscription scheme, but rather on an implicit and dynamic overlay capturing the commonalities between users with respect to they are interested in. The size of the set of users to which a news is forwarded depends on the interest of the news and its origin (*amplification*).

1 Overview of WUP

Recent events in Iran and Haïti emphasized the importance of social networking through the Internet as the only means through which information can be disseminated, despite natural disasters and censorship. When e-mail, phone, and fax could not be used, news transited through Twitter. Yet, a system a la Twitter is clearly not a panacea. The company might be pressed by a government to stop its activities in certain areas. The system is hosted by a set of servers that can get attacked by some form of news bombing. In addition, the dissemination of the information is performed through an explicit network of friends (or followers) and a node can miss interesting information or receive spam.

WUP aims to bring news to people all over the world, from all over the world, in a decentralized and personalized manner. No government or big brother company should be able to filter the news and the information dissemination process should be robust to failures of computer and networking infrastructures.

Personalization here means that users should not be spammed with irrelevant news, yet they should receive news of high interest to them in a second. We believe that this notion of interest cannot be expressed with a subscription scheme based on topics or keywords, for the lack of any ontology and structure to classify news in the real world makes this simply impossible. Clearly, explicit social networks can be viewed as effective means to limit the scope of the dissemination of the news, i.e., one only receive news posted by friends. Yet, this again limits the kind of news to be received and declaring someone as friend one day does, by no means, qualify the news posted by that "friend" as highly relevant few months later (Facebook provides anecdotal evidences of that). On the other hand, relevant and important news could come from other people outside the circle of explicit friends.

In WUP, the interest of a user is approximated by which news that user found relevant in the past. Indeed, a WUP user is not only potentially active in the sense that she can post news, but also in the sense that she

can comment on the news by declaring it interesting or, in contrast, non-interesting. The more opinionated a user, the more personalized information it will get and the faster it will get it. More specifically, WUP maintains a dynamic overlay of implicit acquaintances, where the relationship is established when two users have expressed similar interests in the same news. The dissemination of the information is then performed using our biased epidemic protocol (BEEP) in a way that is robust to failures and attacks, as we explain below.

2 Overview of BEEP

At the heart of WUP lies an epidemic dissemination protocol. Such protocols are known to be simple to deploy and robust to dynamics. Basically, the idea is to disseminate information (news in our case) in a peer-to-peer manner, each peer forwarding it to a small set of randomly selected peers. Such a protocol cannot however be used as such to disseminate news in the real world because of its uniformity: all news and all users are treated equally and would in principle receive all news. Also, in its simple form, an epidemic dissemination protocol is very vulnerable to attacks as malicious users can jam the system with bogus news to overload the system, or news with non-acceptable content, scaring away potential users from using and contributing to it.

Instead, we use a biased epidemic protocol we call BEEP. Two ideas underly BEEP: *amplification* and *orientation*. Amplification means that the set of nodes to which a node forwards some information depends on (1) the interest expressed by the user at that node about the news and (2) the origin of the news. Clearly (1) is used to stop jams and rather propagate news that are meaningful to the user itself. In fact, (2) does the same job, in a complementary manner. The goal here is to propagate news coming from independent sources, leveraging the random nature of epidemic dissemination. More specifically, if some news come from nodes that are not well connected in the network (and possibly part of the same malicious coalition), it is more likely to be true. The idea of orientation is to bias the constitution of the set of nodes to which a user transmits the news at its disposal. This set is constituted from a large part to the implicit acquaintance of a node, complemented by a subset of nodes chosen randomly.

3 On-going and future work

Current work consists in fine-tuning the BEEP protocol and building the WUP system. The BEEP protocol requires first a network of acquaintances. This relies, in turn on a similarity metric allowing to capture the common interest in a news system. The random contacts of each peer are themselves provided by a random peer sampling service.

Second, the BEEP protocol must set the dissemination parameter to implement the amplification and orientation features of the protocol for that specific news, at that specific stage of the news dissemination, on that specific user. This typically means that the fanout (to how many peers this news should be disseminated) and the bias (to whom this news should be forwarded) have to be set at each dissemination stage.

A first prototype of WUP has been implemented and we are experimenting it. In particular, we are in the process of comparing the dissemination under WUP and traditional dissemination system with real traces from LiveJournal and Twitter. In addition, a real churn model has been used to stress our approach. Our preliminary results are encouraging and show that WUP delivers news of interest in fast manner and with little bandwidth and resists well to jamming.

Several challenging questions remain to be addressed:

- what is a relevant benchmark for the evaluation of WUP? we are currently simulating WUP to evaluate it in a large-scale system; we are also implementing WUP and evaluating it on PlanetLab to account for real network characteristics.
- What is the degree of anonymity and security that WUP should provide? We are currently analyzing the potential attacks on **wup** and the impact of misbehaviors and free-riding.
- What are the probabilistic reliability guarantees provided by BEEP?