LENS: LEveraging Social Networking and trust against Spam

Sufian Hameed*, Pan Hui†, Xiaoming Fu*, Nishanth Sastry‡
* University of Göttingen, Germany †Deutsche Telekom Labs, Germany ‡University of Cambridge

I. Introduction

Spam emails is still an open problem largely outnumbering legitimate ones. In 2010, 89%[1] of the emails were spams (262 billion spam messages daily) and the projections show that spam will incur a cost of $338 billion[2] by 2013.

The common state-of-the-art strategy used today only filter spam from the user's inbox (i.e. recipient's edge), but the spam already travels the network, and provokes non-negligible cost to network operators in terms of bandwidth and infrastructure. On the other hand, content-based filtering [2] has turned spam problem into false +ve and -ve one.

There have been innumerable attempts to solve the problem of spam, including, recently, solutions that exploit trust embedded in social networks to create solutions without false positives [4], [5]. RE: [4] introduced the idea that recipients can trust senders in their immediate social neighbourhood, and gave a zero false-positive mechanism for vetting emails sent by their immediate social circle i.e. direct friends and friends of friends (FoF). However, email coming outside this circle still had to be tested by noisy and unreliable spam filters.

In this work, we aim to create a system that, like RE:, can be deployed individually by small groups of users, but allows for a reach greater than FoF. In order to accomplish this, we create a per-recipient ego-centric view of the entire social network of email users. Anyone who is a friend or FoF can email the recipient directly. To enable legitimate senders who are farther away, the recipient enlists a set of trusted users, called Gate Keepers (GKs) at various hop counts away from himself. Each GK is allowed to vouch for new senders in his immediate social circle by issuing them unforgeable vouchers. Unless a GK vouches for the emails of potential senders from outside the social circle of a particular recipient, those e-mails are prevented from transmission. In this way *LENS* drastically reduces the consumption of Internet bandwidth by spam to control messages only.

II. LENS Architecture

LENS (Fig.1) is designed in a modular fashion and comprise of four main components; 1) community formation, 2) trust management, 3) spam report handler and 4) GK selection. All the components of *LENS* run on a Mail Server (MS). Each email user explicitly control his community (friends and FoF) and can give feedback by reporting spam emails. All the remaining functionality of *LENS* is handled transparently by the MS. In this section we briefly describe the first three components followed by elaborate discussion on GK selection.

[1]http://royal.pingdom.com/2011/01/19/email-spam-statistics/
[2]http://www.redcondor.com/company/

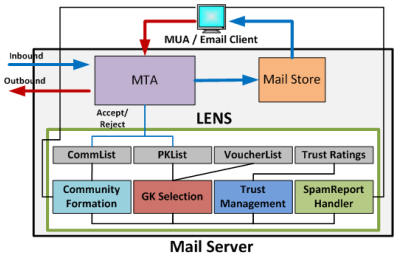

Fig. 1. *LENS* Architecture

A. Community Formation

The formation of a social community is a simple two step process i.e. *addition of friends and FoF.* For FoF addition The idea of FoF friend lists are not exchange among the friends. Instead a user can suggest two of his friends to add each other as FoF. By design, community formation is a selective process to preserve privacy.

B. Trust Management

This component is responsible for maintaining a system wide trust rating (TR) of the users and use them to determine the user type i.e. *legitimate, new and illegitimate*. Calculation of TR use mechanisms that are in principle similar to MailRank [3].

C. Spam Report Handler

This component handles spam reports and weights them according to the TR of the user to identify potential spammer.

D. GK Selection

GK serve as a means to vouch for legitimate users outside the community of the recipient for communication. To maintain a reliable trust structure, a GK is only authorized to vouch for the nodes in his own community. The GK selection consists of three stages as follows. *LENS* covers all the communication scenarios for legitimate emails (no legitimate email is stopped from transmission).

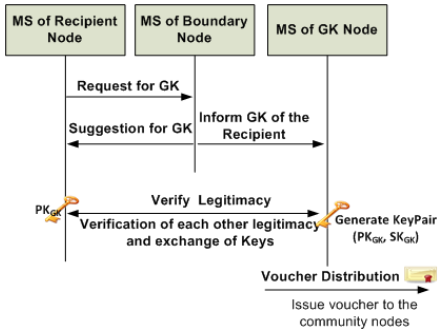

Fig. 2. GK selection and voucher distribution

Stage 1: GK selection in adjacent communities

The GK selection for a recipient (RN), in adjacent communities consist of three step.

i) Request: RN's MS requests all the boundary nodes(BNs) i.e. FoF in the RN's community to send their suggestion for good GKs.
ii) Suggestion: The MS of the BN will suggest a locally optimal (highest degree) friend of the BN to the RN as a GK. The MS of the BN will also inform the BN's friends about the recipient.
iii) Verification of Legitimacy: This step, with the help of trust management, ensures that the GK is legitimate. As a result of this step, a RSA based public key (PK) and secret key (SK) is generated for the GK. PK is shared with the RN and the SK is use to issue vouchers to entire community members of the GK. These members will use the issued vouchers if they need to communicate with the RN (see Fig. 2). All the users within a social radius of 5 would be able to send emails to the recipient with an assurance of being free from spam. Users having a social distance greater than 5 are covered in stage 2 below.

Stage 2: GK selection beyond adjacent communities

After completion of stage 1, RN's MS *sends a request to the selected GK's MSs* to help them look for GKs from their adjacent communities. As a result of this request, the GKs will use their BN to find new locally optimal GKs and *send their suggestions back to the RN*. Finally, *the RN's MS will verify legitimacy of the new set of GKs from social level 6* and extend reachability of the RN to level 8. The GK selection for higher levels must also consider the small world property of social networks [6], in order to *avoid random walks* on the social graph.

Stage 3: GK selection for new communication

If a user wants to send an email to a recipient (for the first time), who is not only outside its community but there is also no GK for the recipient within its community, *LENS* will perform the following two steps.
i) Announcement: announce the sender to the RN that wants to communicate.
ii) Verification of Legitimacy: start the legitimacy verification process to prove that the sender is and not a spammer. As a result of this process, the RN will add the sender as his GK.

This process is only performed once at the start of a new communication. After the sender is verified as a GK, not only the sender but his entire community can send email to the recipient.

E. Email Processing with LENS

LENS filters emails (based on recipient's community information or sender's voucher from an authorized GK) at the SMTP time (RCPT TO: command). In this way *LENS* is able to make acceptance/rejection decision before actually receiving the email data.

III. EVALUATIONS

For a scalable and effective solution, the goal of GK selection process is to select *minimum GK* for *maximum coverage*. For our evaluations we use Facebook [7] dataset consisting of 3.1 million users with over 23 million edges and an average of 15.2 friends per user. We randomly selected 4000 nodes with the community size between 100 and 1500 and tested them for stage 1 (since average path of the dataset is no more than 5 hops) of GK selection.

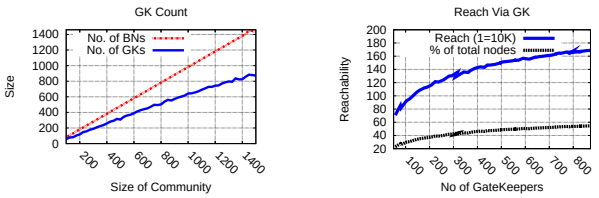

(a) No. of GKs (b) Reachability via GKs

Fig. 3. Evaluations on Facebook dataset

Fig. 3(a) presents the number of GKs selected for a recipient. The number of selected GKs ranges between 56 and 871. Fig. 3(b) shows the number of users that can reach a particular recipient with the help of GKs. With the GKs selected above, the reachability of the recipient is ranging from 700K to 1.7 million i.e. 22% to 55% of the total network, and remains above 40% most of the time.

We also run trace driven evaluation of *LENS* using real email traces of Enron [1] (1,136,760 messages exchanged between 52,747 users). *LENS* is able to receive *all* incoming emails in the email traces with an average of just 31 (0.06% of users) GKs per recipient.

We have implemented *LENS* prototype in Postfix/MailAvenger[3]. Our initial stress tests show that the overheads imposed by the additional processing are tolerably small. *LENS* is significantly less compute intensive (Up to 75% less CPU and 9% less memory) than current solutions like SpamAssassin.

IV. SUMMARY

In this paper we introduce *LENS*, a novel spam protection system based on the *social networking paradigm*, which mitigates spam beyond recipient's social circles with the help of trusted users, called GKs. Unless a GK vouches for the emails of potential senders from outside the social circle of a particular recipient, those e-mails are prevented from transmission. In this way *LENS* drastically reduces the consumption of Internet bandwidth by spam. *LENS* covers all the communication scenarios for legitimate emails (no legitimate email is stopped from transmission).

Our initial evaluation results, based on Facebook traces, demonstrate that reliable email delivery from millions of potential users is possible using GKs in the order of hundreds. The results of email trace driven evaluations shows that with the application of *LENS* we can effectively filter and accept all the legitimate inbound emails. Initial evaluations of the system prototype reveals that *LENS* remains lightweight and performs significantly better than SpamAssasin [2]. More experiments are being performed to evaluate system perform under various scenarios and study further potential improvements.

REFERENCES

[1] Enron email dataset. http://www.cs.cmu.edu/ enron/.
[2] Spamassassin. http://spamassassin.apache.org/.
[3] Paul Alexandru Chirita, Jörg Diederich, and Wolfgang Nejdl. Mailrank: using ranking for spam detection. In *Proc. of CIKM*.
[4] S. Garriss, M. Kaminsky, M. J. Freedman, B. Karp, D. Mazi'eres, and H. Yu. Re: Reliable email. In *Proc. of NSDI*, 2006.
[5] A. Mislove, A. Post, P. Druschel, and KP Gummadi. Ostra: Leveraging trust to thwart unwanted communication. In *Proc. of NSDI*, 2008.
[6] J. Travers and S. Milgram. An experimental study of the small world problem. *Sociometry*, 1969.
[7] Christo Wilson, Bryce Boe, Alessandra Sala, Krishna P.N. Puttaswamy, and Ben Y. Zhao. User interactions in social networks and their implications. In *Proc. of EuroSys*, 2009.

[3] http://www.mailavenger.org/