# Have I Met You Before? Using Cross-Media Relations to Reduce SPIT

Kumiko Ono
Columbia University
New York, USA
kumiko@cs.columbia.edu

Henning Schulzrinne
Columbia University
New York, USA
hgs@cs.columbia.edu

## ABSTRACT

Most legitimate calls are from persons or organizations with strong social ties such as friends. Some legitimate calls, however, are from those with weak social ties such as a restaurant the callee booked a table on-line. Since a callee's contact list usually contains only the addresses of persons or organizations with strong social ties, filtering out unsolicited calls using the contact list is prone to false positives. To reduce these false positives, we first analyzed call logs and identified that legitimate calls are initiated from persons or organizations with weak social ties through transactions over the web or email exchanges. This paper proposes two approaches to label incoming calls by using cross-media relations to prior contact. One approach is that a potential caller offers the callee his contact addresses which might be used in future calls. Another is that a callee provides a potential caller with weakly-secret information. In order to be identified as someone the callee contacted before through other means, the caller can convey the information in future calls. The latter approach enables a callee to label incoming calls even without caller identifiers. Reducing false positives during filtering using our proposed approaches will contribute to the reduction in SPIT (SPam over Internet Telephony).

## Keywords

Unsolicited calls; SPIT prevention; cross-media relations; email; WWW; web; VoIP; SIP

## 1. INTRODUCTION

Unsolicited calls usually originate from persons or organizations, whom the callee does not know their contact addresses nor met before. Since an IP-based infrastructure is more vulnerable to unsolicited calls, as described in [1], people have recently been experiencing more SPIT calls. Most legitimate calls, by contrast, have caller identifiers (IDs) that the callee has seen before. Some legitimate calls, however, have unknown caller IDs. Examples of these legitimate calls include confirmations of appointments, reservations, or deliveries, and recorded notifications of flight delays or school closing on a snowy day. These legitimate calls are often mistakenly labeled as unsolicited calls at a SPIT filtering system since their caller IDs are not found on the callee's white list.

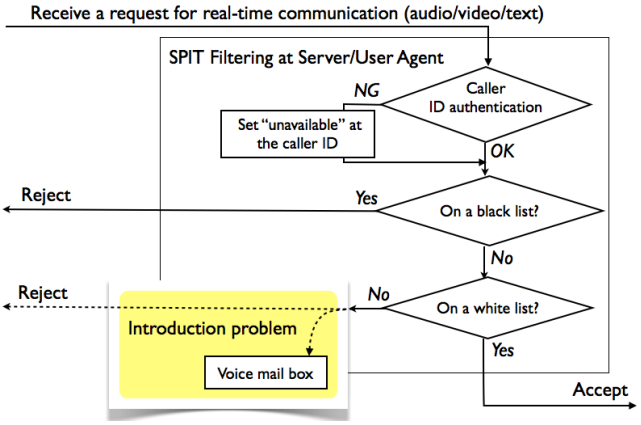


**Figure 1: Existing SPIT filter**

Figure 1 illustrates a typical SPIT filtering system, which authenticates caller IDs and looks them up on a black list or reject-list and a white list or accept-list. For a VoIP (Voice over IP) call using the SIP (Session Initiation Protocol) [2], the SIP Identity header [3] enables a callee to authenticate the caller ID. Some legitimate calls, however, are sent with "unavailable" caller IDs because the authentication of the caller IDs fails. For example, most international calls or calls through a VoIP–PSTN (Public Switched Telephone Network) gateway have no authenticated caller ID. These anonymous calls limit the effectiveness of labeling incoming calls based on the caller ID.

Generally, the callee's black list contains the contact addresses of undesired callers or links to a reputation service that gathers IDs of well-known malicious callers. Unfortunately, however, callers can pick a new caller ID for each call easily especially for VoIP calls. Thus, the effectiveness of the black list is limited.

On the other hand, the callee's white list contains the addresses from his contact list or address book, which is populated by contact addresses of people with strong ties in his social network [4] such as family members and friends. For


*IPTCOMM'09* July 7-8, 2009, Atlanta, Georgia, USA.

business use, the white list usually links to a directory service located on an LDAP (Lightweight Directory Access Protocol) [5] server. For either use, however, the white list does not usually include the addresses of persons or organizations with weak social ties [4] such as friends of a friend in an SNS (Social Network Service). When people who have with weak social ties place calls for the first time, their calls are filtered out since their caller IDs are not found in the callee's white list. This is called the "introduction problem." To mitigate this introduction problem, some systems forward these calls to a voice mail box, rather than reject them. However, this is not a desirable solution because it requires callee's time to check them and causes the delay of notifications. Thus, we need a better approach to label incoming calls from persons or organizations who have with weak social ties.

For this purpose, we analyze how legitimate calls from people with weak social ties are triggered. We then propose two mechanisms to label incoming calls by using cross-media relations between calls and previous contacts. For our first mechanism, a potential caller offers the callee his contact addresses which he might use in future calls. If the callee agrees, these contact addresses are added to his white list. We describe this mechanism further in Section 4.1. For our second mechanism, a callee provides a potential caller with weakly-secret information that the caller can use in future calls in order to be identified as someone the callee has had prior contact through other means, as outlined in Section 4.2. Section 5 describes a use case integrated with an SNS and Section 6 describes implementation to achieve these mechanisms. Finally, Section 7 concludes the paper.

## 2. RELATED WORK

Similar to preventing bulk unsolicited emails, spams, there is no panacea for preventing unsolicited calls; thus, a collection of solutions is needed. As described in [1], the solution space can be divided into two categories: one places procedural, computational, financial, and/or legal burdens on callers, and another labels incoming calls on the callee side. Our mechanisms using cross-media relations are used in conjunction with adding procedural burden on the caller side and enhancing the labeling mechanism on the callee side. Consent-based solutions in SIP [6] are also used in conjunction with the two categories. To grant a permission, whereas the consent-based solutions use additional SIP messages, our mechanisms reuse messages through other means than the SIP.

Most well-known solutions for labeling incoming calls are based on authenticated caller IDs as described in Section 1. Since maintaining static lists of caller IDs as white and black lists has the introduction problem, many approaches using social graphs have been proposed. To expand white lists using social networks, Ceglowski and Schachter [7] introduced address book sharing with privacy as an email attachment, while we [8] offered address book propagation within SIP messages. To update white lists based on communication history, Balasubramaniyan and his colleagues [9] introduced call credentials based on the call history of a caller. Dantu and Kolan [10] described learning systems based on unsolicited call traffic patterns in order to update reputation and black lists. Although their communication history limits to calls, Shacham and Schulzrinne [11] addressed using alternative communication channels, web transactions to collect potential caller IDs. This is the base work for our mechanisms using cross-media relations, which we expand to use email exchanges.

To label incoming calls without caller IDs, one of our labeling mechanisms is based on the destination address with subaddressing [12], which has already been deployed for emails. For calls, subaddressing in the userinfo of the SIP-URI is new, but the concept of extensions in the tel-URI is similar to call distribution at a PBX (Private Branch Exchange).

Relying on the observation that many unsolicited calls play prerecorded messages to decrease cost to the callers, Quittek and his colleagues [13] proposed Turing tests to detect human communication patterns. However, some legitimate calls from government agencies, credit card companies, or dealers among others are automated recorded messages. The SPIT detection system proposed by Mathieu and his colleagues [14] relies on a SPIT characteristic that unsolicited calls originate more error messages than legitimate calls. However, legitimate automated calls also share this property. Also, unless the originating carrier cooperates, it may be difficult to measure the outgoing call volume. This information may also be considered privacy sensitive or a business secret.

## 3. LEGITIMATE CALLS FROM WEAK SOCIAL TIES

Our quick survey gives a rough sense of how often people are experiencing unsolicited calls, how well-maintained contact lists are effective in labeling legitimate calls, and how legitimate calls from weak ties are initiated. In this survey, we gathered call records of 246 calls from eight cell phones and 136 calls from four landline phones from our colleagues at our lab. We also asked the participants about their relationship to legitimate callers whose IDs were not found on their contact lists.

Figure 2 indicates a significant difference in the proportions of unsolicited calls between cell and landline phones. Whereas only six percent of the incoming calls on cell phones were unsolicited calls, 52 percent of those on landline phones were unsolicited. We suspect that this difference was caused by the FTC (Federal Trade Commission) regulations that prohibit telemarketing calls to cell phones [15]. Even though we can reduce unsolicited telemarketing calls using national "Do Not Call" registry service [15], the effect is unfortunately limited. This is because their jurisdiction is limited over domestic telemarketers, not over international ones nor calls using VoIP. Also, some telemarketers appear to be flouting the law. Thus, we still need a technical mechanism to help a callee decide whether to accept incoming calls.

Figure 2 also illustrates a difference in the proportions of legitimate calls with known caller IDs. A larger proportion, 78 percent, of the calls for cell phones carried known caller IDs, which were found on the contact list, compared to 18 percent for landline phones. Since people usually maintain their contact lists on cell phones better than landlines, the result shows how well-maintained contact lists are useful to label incoming calls.

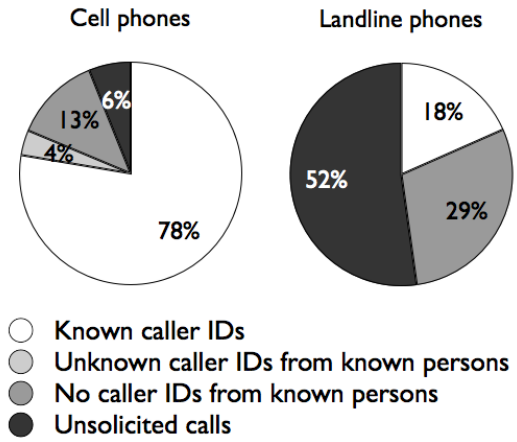


**Figure 2: Incoming calls: cell phones vs. landline phones**

Figure 2 also indicates that 17 or 29 percent of the incoming calls were legitimate, but with unknown or unavailable caller IDs. By asking the participants, we found that all these legitimate calls with unknown caller IDs were related to transactions over the web or email exchanges. For example, they were confirmation calls from the restaurants which the callee made on-line reservation, or notification calls of flight changes from the airline on which the callee booked flights. On the other hand, the legitimate calls with no caller IDs were international calls or calls through VoIP-PSTN gateways from people with strong ties. There were no calls from legitimate callers whom the callee has had no prior contact with. We summarize that even if we had collected a larger data set, most legitimate calls from people with weak ties would still have had previous contacts with the callees. This suggests that we need a new mechanism to label incoming calls beyond using caller IDs, and the solution could be use a piece of information related to the previous contacts.

From these indications, therefore, we set our goal to enhance a SPIT filtering system covering calls from persons or organizations with weak social ties. Our approach is to use a piece of information related to previous contacts between the callee and the caller, in addition to using caller IDs.

# 4. USING CROSS-MEDIA RELATIONS

Legitimate calls from persons or organizations with weak ties, as analyzed in Section 3, are usually preceded by previous contacts between the callee and caller through transactions over the web or email exchanges. Focusing on these previous contacts, we propose that both parties exchange additional information which can be used in future calls as an indication of a legitimate previous contact. We call this piece of information a "cross-media relation." Our approach is to expand filter conditions for incoming calls by using the cross-media relations as illustrated in Figure 3, which is also applicable to other real-time communication requests. We distinguish two types of cross-media relations: contact addresses offered by potential callers and weakly-secret information provided by a callee. The following outlines the mechanisms using each type of cross-media relations and shows our proposed filtering system.

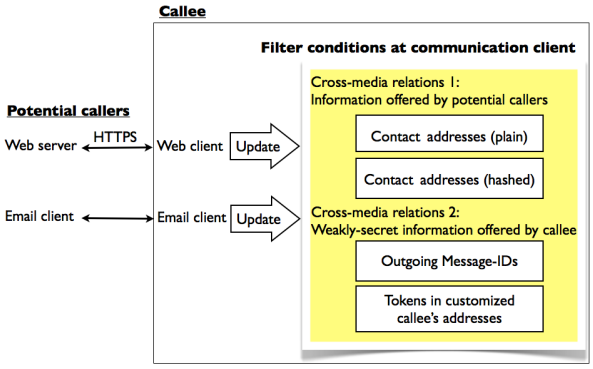


**Figure 3: Overview of proposed mechanisms**

## 4.1 Contact Addresses of Potential Callers

In general, the more contact addresses we can obtain from potential callers, the more incoming calls we can label, since a typical filter system uses the caller IDs as shown in Figure 1. Thus, persons or organizations that a callee contacts through the web transactions or emails offer their contact addresses which they might use in future calls with the callee.

Depending on the contact mechanism, callers use a different method to convey their contact addresses. In a web transaction, i.e., an HTTP transaction shown in Figure 4, the contact addresses, e.g., `sip:operator@book.airline.com`, are conveyed in a new HTTP header, Correspondence-URIs [11] or an HTML META tag, HTTP-EQUIV [16] in the response from the potential caller. In an email exchange shown in Figure 5, the contact addresses are contained in a vCard [17] attached to an email message sent from a potential caller.

After the callee receives the contact addresses of a potential caller, he adds them to his white list. To prevent misuse, the callee should be prompted for confirmation before updating his white list only for secure HTTP (HTTPS) [18] transactions.

The format of the contact address can be either plain text or a hash of it. Hashed contact addresses are suitable if the potential caller prefers concealing his routable address for privacy or operation reasons. For example, in an SNS, when a subscriber prefers not to publish his routable contact address, he can instead publish his hashed contact address for the limited purpose of filtering calls.

The mechanism to use this type of cross-media relations is appropriate in a case where the previous contact was one-to-one correspondence between the callee and the potential caller. However, we cannot apply this mechanism in several cases such as previous contact for membership in an association. In these cases, the callee should deliver weakly-secret information to potential callers.

## 4.2 Weakly-Secret Information

Weakly-secret information provided by a callee can also a type of the cross-media relations. Potential callers can use this information in future calls to be identified as someone with whom the callee has had prior contact through other

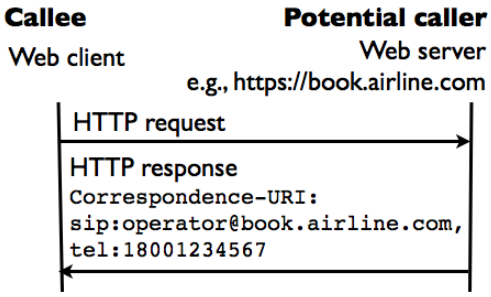


**Figure 4: HTTP message exchange where a potential caller delivers his contact addresses**


**Figure 5: Email message exchange where a potential caller delivers his contact addresses**

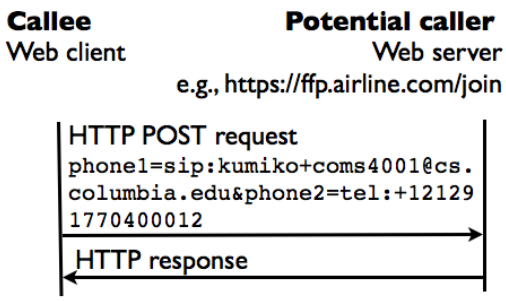


**Figure 6: HTTP message exchange where a callee delivers weakly-secret information**

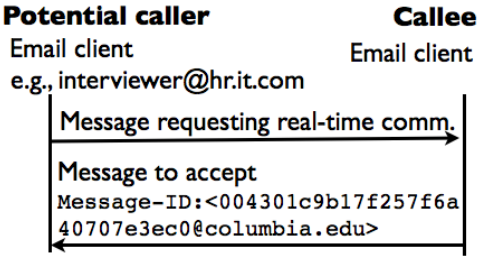


**Figure 7: Email message exchange where a callee delivers weakly-secret information**

means. This mechanism is useful in the following cases. One is where the previous contact was one-to-many correspondence between the callee and the potential callers. For example, when joining an association, the callee is unwilling to receive all the contact addresses of the potential callers in the association. Another case is where potential callers might use a different or no authenticated caller ID, due to the type of communication medium or service such as two-stage dialing for international calls.

Depending on the communication medium of the previous contact, a callee provides a potential caller with a different type of information. A customized contact address containing a random component or a token can be used when a callee fills out contact information on a web site, as shown in Figure 6, or in a vCard attached to an email message. The random component or token can be automatically generated in correspondence to the URL (Uniform Resource Locator) [19], or manually specified. In the examples in Figure 6, a token, `coms4001`, in the SIP-URI is set between the user name and the domain name preceded with +, in the same way as the email addressing practice called "sub-addressing" [12]. For tel-URI [20], a token, `0012`, follows the E.164 number like an extension. To convey this information in a later call, the caller just needs to set the destination address to the customized contact address.

Specifically in an email exchange, as shown in Figure 7, the message identifier of an email from the callee can be used. A potential caller first sends a message to the callee requesting a real-time communication. Only if the callee accepts the request, he will respond to it by email containing his contact address. As a result, the message identifier of the response email, which is set in the Message-ID [21] header, can be used as weakly-secret information to prove the acceptance from the callee. Thus, the message identifiers of outbound emails or SIP calls can be included by the potential caller in a later call, even if he uses a different caller ID or type of communication medium.

To convey the message identifier in a SIP call, the caller should set a SIP header extension, References [22] to its value. Although we need to define a new parameter of the References header, which currently limits the parameter to call identifiers, we assume that the References header is used for this purpose.

For message security, we should use an appropriate mechanism for each communication protocol. That is, in web transactions, we use HTTP over TLS, i.e, secure HTTP (HTTPS) for message confidentiality, its integrity, and the authentication of the web server. Also for emails, we use TLS (Transport Layer Security) [23] for all the hops from a client to the other. We also leverage anti-spam email mechanisms when receiving emails.

## 4.3 Proposed Filtering Process

Figure 8 depicts a new filtering process for incoming calls, modifying and adding conditionals using the cross-media relations. If the caller ID of the incoming call is not found on a black list, then the SPIT filter tries to find it on a white list. The white list contains contact addresses either in the plain text or hashed format. To increase the effectiveness of the white list, people within an organization can share a common white list stored on a server and can receive binary responses to query whether or not the caller ID is found on the list.

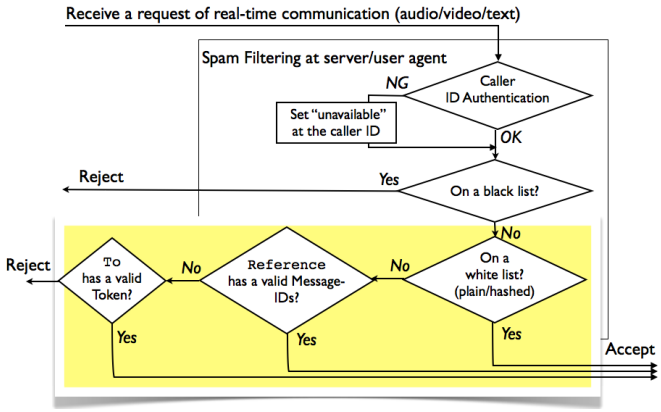


**Figure 8: SPIT filter using cross-media relations**

If the caller ID of the incoming call is not found on the white list, the new filtering process tests on two new conditionals. The first one is whether it contains a valid Message-ID value in the References header. The second is whether it contains a valid token in the destination address, i.e., in the To header. The validity can be determined by looking up on the filter conditions of message identifiers and tokens. If the test succeeds in either condition, the call request can be accepted.

# 5. A USE CASE: INTEGRATION WITH AN SNS

We describe how we can apply our proposed mechanisms in a typical existing SNS, since an SNS is the most popular and effective service for subscribers to maintain relationships with both strong and weak ties and to initiate real-time communications with each other. There are three typical services in a typical SNS: the subscription, the invitation of new friends to expand their own social network, the notification of the state updates of friends, and message exchanges among them. The following describes how subscribers can extract cross-media relations in each service.

A newcomer starts to subscribe to an SNS through a transaction over the web although the invitation to the subscription may be sent by email. In this web transaction, the newcomer fills in a sign-up form including his name and contact addresses by email and/or phone. By submitting the sign-up form, he can send a token in his customized phone contact address. Then, he can save the token corresponding to the URL of the web site as a filter condition. When he receives an incoming communication request destined for his contact address with the token, he can identify the caller as one of the subscribers in the SNS and decide whether to accept the request.

Next, a subscriber can expand his social network in the SNS by inviting his friends to add to his network or being invited by his friends to be added in their networks. Such an invitation is generally delivered by email asking the invitee to respond at the SNS web site. If the invitee accepts the invitation, he can receive his friend contact address in the HTTP response. The format of the contact address, in plain text or hashed, depends on the preference of the owner of the contact address. By adding the contact address to his white list, the invitee can prepare to label calls or text messages from the friend.

When notifying the status updates of a subscriber's social network, the notification message can contain the list of hashed contact addresses of the friends of his friend. Generally, users prefer concealing their own contact addresses from friends in the second degree. Thus, the hashed format of their contact addresses, rather than plain text format, is appropriate.

Among the members in his social network, a subscriber often exchanges messages through the SNS server. When he wants to talk with a girl of his friend, but does not know her contact address, he needs to send a message asking her contact address for a real-time communication. If he can receive a message response showing the acceptance and her contact address, he can send a call request with the Message-ID of the response. Therefore, she can identify him as a person corresponding to the previous message. Reversely, if he is asked and accepts her request, he needs to save the Message-ID of the response in order to label a later call from her.

Thus, in these services in an SNS, a subscriber can extract cross-media relations, prepare to label incoming calls or other real-time communication requests, and identify the caller or requester as a specific subscriber or one of the subscribers.

# 6. IMPLEMENTATION TO HAVE

As we outline below, our technique requires minimal changes. The following are required implementation for the SPIT filtering system using cross-media relations, a SIP proxy server, a caller and a callee. For each end user, we describe what kind of functions need to be in a SIP User Agent (UA), a web browser or server, and an email client.

## 6.1 Implementation of SPIT Filtering System

The SPIT filtering system using cross-media relations can be located in an inbound SIP proxy server or a SIP UAS (User Agent Server). As described in Section 4.3, the SPIT filtering system stores accept-lists of caller IDs with the hash algorithm, message IDs and tokens within customized contact addresses corresponding to the caller information related to prior contacts such as the URI for a transaction over the web. When receiving an incoming call, it reads the originator address in the addr-spec parameter of the From header, the destination address in the addr-spec parameter in the To headers, and the referred message identifier in the refer parameter of the References header. It then tries to find the values on the lists. If the call is accepted, the filtering system may provide corresponding caller information in order to ask the user's final decision whether to accept it. When the filtering system is located in a SIP proxy server, such caller information is conveyed in the refer parameter of the References header.

## 6.2 Implementation in a SIP Proxy Server

A SIP proxy server is required to implement any additional functions if it follows the email subaddressing practice described in [12]. As an inbound SIP proxy server, it allows subaddressing in the userinfo of the SIP-URI in the To header

and Request-URI. When the server determines the destination user name, the server ignores the string after the plus separator. Without any additional functions, a SIP proxy forwards the SIP References header.

## 6.3 Implementation at the Callee

When receiving a call, a SIP UAS shows the call information including the destination address in the To header and the refer parameter of the References header to help the callee's decision whether or not to accept it.

In a web browser, an add-on program supports following functions: generating and storing a token when a user is filling in a sign-up form, and extracting and storing contact addresses when receiving an HTTP response. When generating a token for a sign-up form, the add-on program stores the token corresponding to the timestamp of the generation and the URL to which the sign-up form sent in an HTTP request over TLS. When extracting contact addresses from the HTTP response over TLS, if the user agrees, the add-on program adds their contact addresses to his white list. Each contact address is stored with the hash algorithm if the address is in the hashed format, the timestamp of the response received, and the URL from which the response came.

For an email client, an IMAP (Internet Message Access Protocol) [24] client dedicated to our proposed mechanisms is needed. This is because the required functions run periodically without any user interaction, as long as the client can fetch saved outgoing and incoming legitimate emails. This IMAP client supports following functions: extracting and storing contact addresses in a vCard from incoming email messages and extracting and storing the message identifiers in the Message-ID headers from outgoing emails. The extracted information is not always needed to synchronize with saved outgoing emails. When a user deletes an outgoing email, the corresponding message identifier may remain for be looked up by a filtering system for a certain period of time. The message identifiers are stored as the filter conditions, corresponding to the timestamp of the email sent and the destination address in the To header.

## 6.4 Implementation at a Caller

A SIP UAC (User Agent Client) sets the References header extension in an outgoing request to convey the referred message identifier manually set by a caller. When the destination address is in the tel-URI, a SIP UAC may set the post dial parameter to the additional digits of the customized contact address of a callee.

A web server supports an HTTP header extension or HTML META tag when responding to an accepted sign-up form with the contact address which will be used in future calls.

Since attaching vCards to emails has been widely deployed, an email client does not need any additional functions at a caller.

# 7. CONCLUSIONS

To label incoming calls, we proposed to use cross-media relations between calls and previous contacts between caller and callee via transactions over the web or through email exchanges. These cross-media relations are expressed in two types of information. First, relations can be as potential callers' contact addresses in either plain text or hashed format, and second, they can be expressed as weakly-secret information in the callee's customized contact address or the message identifier of the callee's outgoing email. By enhancing existing filter conditions, our proposed mechanisms enable a callee to label incoming calls, not only from persons or organizations with weak ties, but also from callers who change their caller IDs. As a result, we expect to avoid most false positives that occur during filtering such calls represent in our survey 17 percent of the incoming calls for cell phones and 29 percent of the incoming calls for landlines.

In addition to the effect of reducing false positives, we expect to observe two secondary effects from enhanced filtering. One of these secondary effects is the ability to trace after the delivery of the customized contact address including the weakly-secret information. The weakly-secret information will identify the person who sold the contact address when the caller sells the contact address of a callee to a third party. Another secondary effect of our proposed filtering system is increased security as a result of maintaining hashed contact addresses, which are unroutable, as a filter condition. Without collecting contact addresses in plain text format, our system can be less vulnerable to viruses which spread using gathered routable addresses as target addresses.

The work described in this paper is the first step in ongoing efforts to integrate anti-SPIT mechanisms with anti-spam techniques. We plan to deploy the prototype with our proposed mechanisms in order to examine their effectiveness and usability.

# 8. REFERENCES


[1] J. Rosenberg and C. Jennings. The Session Initiation Protocol (SIP) and Spam. RFC 5039, IETF, January 2008.

[2] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, M. Handley, and E. Schooler. SIP: Session Initiation Protocol. RFC 3261, IETF, June 2002.

[3] J. Peterson and C. Jennings. Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP). RFC 4474, IETF, August 2006.

[4] M.S. Granovetter. The Strength of Weak Ties. *Amer. J. of Sociology*, 78:1360–80, May 1973.

[5] J. Sermersheim. Lightweight Directory Access Protocol (LDAP): The Protocol. RFC 4511, IETF, June 2006.

[6] J. Rosenberg, G. Camarillo, and D. Willis. A Framework for Consent-Based Communications in the Session Initiation Protocol (SIP). RFC 5360, IETF, October 2008.

[7] M. Ceglowski and J. Schachter. LOAF. `http://loaf.cantbedone.org`, 2004.

[8] K. Ono and H. Schulzrinne. Trust Path Discovery. Internet-draft, IETF, June 2006. `http://tools.ietf.org/html/draft-ono-trust-path-discovery-02`.

[9] V. Balasubramaniyan, M. Ahamad, and H. Park.

CallRank: Combating SPIT Using Call Duration, Social Networks and Global Reputation. In *CEAS 2007 - Fourth Conference on Email and Anti-Spam*, August 2007.

[10] R. Dantu and P. Kolan. Detecting Spam in VoIP Networks. In *Steps to Reducing Unwanted Traffic on the Internet Workshop*, pages 31–37. USENIX association, July 2005.

[11] R. Shacham and H. Schulzrinne. HTTP Header for Future Correspondence Addresses. Internet-draft, IETF, May 2007. `http://tools.ietf.org/html/draft-shacham-http-corr-uris-00.txt`.

[12] K. Murchison. Sieve Email Filtering: Subaddress Extension. RFC 5233, IETF, January 2008.

[13] J. Quittek, S. Niccolini, S. Tartarelli, M. Stiemerling, M. Brunner, and T. Ewald. Detecting SPIT Calls by Checking Human Communication Patterns. In *IEEE International Conference on Communications (ICC '07)*, pages 1979–1984, June 2007.

[14] B. Mathieu, Y. Gourhant, and Q. Loudier. SPIT mitigation by a network level anti SPIT entity. In *Third annual security workshop (VSW'06)*. ACM Press, June 2006.

[15] Federal Trade Commission. National Do-Not-Call Registry. C.R.F. Part 310, Telemarketing Sales Rule, June 2003. `http://www.ftc.gov/donotcall`.

[16] D. Raggett, A.L. Hors, and I. Jacobs. HTML 4.01 Specification. `http://www.w3.org/TR/REC-html40/`, December 1999.

[17] F. Dawson and T. Howes. vCard MIME Directory Profile. RFC 2426, IETF, September 1998.

[18] E. Rescorla. HTTP Over TLS. RFC 2818, IETF, May 2000.

[19] T. Berners-Lee, R. Fielding, and L. Masinter. Uniform Resource Identifier (URI): Generic Syntax. RFC 3986, IETF, January 2005.

[20] H. Schulzrinne. The tel URI for Telephone Numbers. RFC 3966, IETF, December 2004.

[21] P. Resnick. Internet Message Format. RFC 5322, IETF, October 2008.

[22] D. Worley. The References Header for the SIP. Internet-draft, IETF, February 2009. `http://www.ietf.org/internet-drafts/draft-worley-references-02.txt`.

[23] T. Dierks and E. Rescorla. The Transport Layer Security (TLS) Protocol Version 1.2. RFC 5246, IETF, August 2008.

[24] M. Crispin. INTERNET MESSAGE ACCESS PROTOCOL - VERSION 4rev1. RFC 3501, IETF, March 2003.