

Scalable Discovery of Hidden Emails from Large Folders

Giuseppe Carenini, Raymond T. Ng, Xiaodong Zhou
Department of Computer Science
University of British Columbia, Canada
{carenini, rng, xdzhou}@cs.ubc.ca

ABSTRACT

The popularity of email has triggered researchers to look for ways to help users better organize the enormous amount of information stored in their email folders. One challenge that has not been studied extensively in text mining is the identification and reconstruction of hidden emails. A hidden email is an original email that has been quoted in at least one email in a folder, but does not present itself in the same folder. It may have been (un)intentionally deleted or may never have been received. The discovery and reconstruction of hidden emails is critical for many applications including email classification, summarization and forensics. This paper proposes a framework for reconstructing hidden emails using the embedded quotations found in messages further down the thread hierarchy. We evaluate the robustness and scalability of our framework by using the Enron public email corpus. Our experiments show that hidden emails exist widely in that corpus and also that our optimization techniques are effective in processing large email folders.

Categories and Subject Descriptors: H.2.8 [Database Management]: Database Applications - Data Mining

General Terms: Algorithms, Performance

Keywords: text mining, hidden email, forensics

1. INTRODUCTION AND MOTIVATION

The popularity of email has triggered researchers to look for ways to help users better organize and use their mail folders, e.g., classification [7], task management [4] and user interface [9]. One key difference between emails and other types of documents is the threaded nature of emails. According to one study, over 60% of emails are threaded [5].

In this paper, we study a problem largely overlooked in text mining - the problem of discovering and reconstructing hidden emails. A hidden email is an original email that has been quoted in subsequent emails kept in a user's folder, but is not itself present in the folder(s). Anyone who has ever

managed folders is accustomed to the tedium of manually shunting messages between folders, as well as deciding which messages to keep and which to delete. Accidental deletion often occurs. Hidden emails also occur when new recipients are included in an existing thread. Whether the original email was deleted or never existed, it still may be found in the quotation of subsequent emails.

Beyond helping users to better manage their ever growing folders, hidden email reconstruction also find many forensic applications. For instance, emails may be deliberately deleted by a malicious user to avoid certain information to be revealed. In this case, investigators will need to reconstruct parts of the deleted emails that may be quoted in other emails stored in other people's folders.

The problem this paper attempts to solve is: how hidden emails can be discovered and reconstructed in a robust and efficient way using the embedded quotations found in messages further down the thread hierarchy.

As a concrete case study, we report in this paper the results of applying our framework to the Enron email corpus. To the best of our knowledge, it is the biggest public accessible email dataset, which has been of great value for research in many areas, including email classification. The two key purposes of using this dataset are to verify the robustness of our method on real data, and to ascertain its scalability.

While we discuss related work in Section 2, we describe the overall framework to regenerate hidden emails by a running example in Section 3. Although this basic framework can deliver the functionality, it cannot deliver the efficiency, nor the robustness, to deal with large real folders. In Section 4, we develop two ways to optimize the performance of the algorithms. Both are based on the use of word-indexing. They are designed to scale up to large folders and long emails.

In Section 5, we report a comprehensive case study based on the Enron dataset. We examine the prevalence of hidden emails and fragments to assess the importance of the hidden email reconstruction problem. We evaluate the robustness of our algorithm when applied to real folders. Last, but not least, we evaluate the success of our optimizations in scaling up to the needs of large folders and long emails.

2. RELATED WORK

From a research perspective, email and newsgroups differ from traditional documents in many aspects. A key difference is that there is a high level of hierarchical and referential relationship among emails in any folder, i.e., document threading. This relationship has caught the attention of many researchers. In [6], Lam et al. propose to sum-

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.

KDD'05, August 21–24, 2005, Chicago, Illinois, USA.

Copyright 2005 ACM 1-59593-135-X/05/0008 ...\$5.00.

Algorithm HiddenEmailFinder

Input: email folder MF , and reference email folders RF_1, \dots, RF_k
 Output: a set of bulletized hidden emails in MF

1. For each email $M \in MF$, extract all the quoted fragments.
2. For each quoted fragment F , match F against all emails in MF as well as those in RF_1, \dots, RF_k . In particular, identify the LCS between F and M for every email $M \in MF$ or $M \in RF_i$ for some i . Depending on the length of the LCS, F may be declared hidden or not.
3. Find possible overlaps between hidden fragments, split them if necessary, and create a precedence graph G .
4. Decompose G into its weakly connected components.
5. For each weakly connected component GC , do:
 - (a) Process GC with Algorithm *graph2email* as described in [1]. If the graph is complete and strict, output the reconstructed hidden emails. Otherwise, use the heuristics to deal with non-strictness and/or incompleteness. Output the reconstructed hidden emails.

Figure 1: A Skeleton of HiddenEmailFinder

marize a set of emails based on their threading hierarchy. They mention the existence of deleted emails in the hierarchy, which creates complications not present in newsgroups. However, they do not study how to regenerate those deleted emails. Newman [8] indicates the possibility of orphaned quotations and warns that applications such as classification and summarization would be adversely affected as a result, but does not explore the issue further.

Carvalho et al. [2] studied the problem of signature and quotation detection within an email. Their work can help us to identify quotations.

In [1], we present a preliminary report on the hidden email discovery and regeneration problem using a small synthetic dataset. Section 3 describes the overall framework. However, this framework does not scale up to large real datasets and long emails. Thus, the key contributions of the work presented here is to devise robust and scalable algorithms, as well as to evaluate them rigorously with real email folders.

Our research on the reconstruction of missing emails is also relevant to the area of document forensics, where document reconstruction from fragments is crucial. For instance, Shanmugasundaram et al. in [10] propose the reconstruction of totally ordered digital evidence from randomly scattered fragments. With respect to our goal of reconstructing the hidden email, as well as in document forensics, a total order is not always possible. Forcing one where none exists may be incorrect and even misleading. We believe that a partial order representation, i.e., the bulletized model, constitutes a reasonable solution that adequately compromises between accuracy and completeness concerns.

3. A BASIC FRAMEWORK FOR RECONSTRUCTING HIDDEN EMAILS

For any given email folder, some emails may contain quotations from original messages that neither exist in the same folder nor in a set of reference folders; the originals may have been deleted or were never received at all. Each of those quotations is considered a *hidden fragment*, as part of a *hidden email*. Several hidden fragments may all originate from the same hidden email, and a hidden fragment may be quoted in multiple emails. Our goal is to reconstruct hidden emails from the hidden fragments by finding the relative ordering


of the fragments and, where possible, piecing them together.

Figure 1 shows the overall algorithm called HiddenEmailFinder, which involves identifying hidden fragments (Step 1&2), creating a *precedence graph* to represent hidden fragments and their relative orders (Step 3), reconstructing hidden emails based on a *bulletized email model* (Step 4&5).

HiddenEmailFinder summarizes our framework of reconstructing hidden emails. Due to the lack of space, the details of the framework, e.g., formal definitions, theorems and algorithms, can be found in [1], and are not discussed here. In the following we only briefly introduce it by one example.

Subject: Re: Midterm Details > a) I need to meet with a > faculty recruit. > c) Warren and Qiang, can > you go SOWK directly? No problem. > f) I will bring classlists. Do they need to sign on the list? - Warren	Subject: Re: Midterm Details > a) I need to meet with a > faculty recruit. > b) Don, can you go LSK directly? Sure. > d) I will bring the exams with me. I can help you carry it. > f) I will bring classlists. Is there a seating plan? Don	Subject: Re: Midterm Details > a) I need to meet with a > faculty recruit. > b) Don, can you go LSK directly? Don, I'll go with you too. > e) The exam takes 48 minutes. > f) I will bring classlists. Do we have a seating plan as last term? Kevin
E ₁	E ₂	E ₃

(a) Folder MF



(b) The precedence graph and the hidden email

Figure 2: Example for HiddenEmailFinder

Figure 2(a) shows a folder MF containing three emails E_1, E_2, E_3 , each of which quotes from the same hidden email which is absent from MF . In Step 1 of HiddenEmailFinder, we extract the quoted fragments of every email in MF . For the ease of representation we use the first item number in each fragment to represent it. For example, E_1 has two quoted fragments $[ac], [f]$, and the rest are non-quoted fragments. In Step 2, for each quoted fragment F we need to identify whether or not it is a hidden fragment. In order to adapt to the complications of quoted fragments, e.g., deletion, insertion and forwarded messages, we use Longest Common Substring (LCS) to identify a match (more details are given in Section 4.2). Let τ represent the LCS of two fragments F, F' , where F' is a non-quoted fragment. Hence, $F = [F_1; \tau; F_2], F' = [F'_1; \tau; F'_2]$. If τ is longer than a given threshold $minLen$, τ is not considered as a hidden fragment and is removed from F . F is replaced by F_1 and F_2 . In this example, if we set $minLen$ to 10, the only hidden fragments identified at the end of Step 2 are: $[ac], [f], [ab], [d]$ and $[ef]$.

Since multiple emails may quote the same hidden email, we need to group hidden fragments quoting the same hidden email and piece them together in an appropriate order. In Step 3, we build a precedence graph to achieve this goal. We assume that if the overlapping of two hidden fragments is sufficiently long, the fragments originate from the same hidden email. Similar to Step 2, let τ represent the LCS of two hidden fragments F, F' and $F = [F_1; \tau; F_2], F' = [F'_1; \tau; F'_2]$.

If τ is longer than $minLen$, i.e., F and F' overlap long enough, F and F' are replaced by new hidden fragments $\{F_1, F_2, F'_1, F'_2, \tau\}$. For example, when we compare hidden fragments $[ac]$ with $[ab]$, $[ab]$ is replaced by two hidden fragments $[a]$ and $[b]$ and similarly $[ac]$ is replaced by $[a]$ and $[c]$. In this way, we end up with 6 distinct hidden fragments, $[a], [b], [c], [d], [e], [f]$, in Step 3.

Next, we build a precedence graph $G = (V, E)$ as follows. Each node in V represents a distinct hidden fragment and each edge $(F, F') \in E$ shows that there exists at least one email M containing both F and F' in the same textual order. Figure 2(b) shows the precedence graph built for the hidden fragments in Figure 2(a).

In Step 4, the graph is decomposed into its weakly connected components (only one in the example). Because a user can quote emails freely and arbitrarily, each component is not guaranteed to be a simple chain of nodes. As shown in Figure 2(b), the textual ordering of $b\&c$ and $d\&e$ is undetermined. We call such nodes *incompatible nodes*. In order to reconstruct and display the hidden email, we need to represent those partial ordering.

However, users usually read a document sequentially and are not accustomed to reading graphical representations of document fragments. We propose the bulletized email model to solve this problem. We use *bullets* to represent incompatible nodes and *offsets* to represent nested relationships among bulletized items. Figure 2(b) shows the bulletized email reconstructed in Step 5. In the regenerated hidden email, fragments $[b], [c]$ and $[d], [e]$ are bulletized to represent the partial ordering between them respectively. $[d], [e]$ are also indented with offsets which show that $[b]$ precedes both $[d]$ and $[e]$.

In [1], we give precise conditions under which a weakly connected precedence graph can be exactly represented by one bulletized email. Moreover, we also give heuristics to deal with situations in which those conditions are not satisfied.

4. OPTIMIZATIONS FOR LARGE FOLDERS AND LONG EMAILS

Even though HiddenEmailFinder delivers all the required functionalities, a preliminary experimental evaluation revealed two bottlenecks in the hidden fragment identification steps (Step 2 in Figure 1) when dealing with large folders and long emails. The first bottleneck is due to the large number of matches that may need to be performed between quoted fragments and other emails in the folders, while the second bottleneck is due to how efficiently LCS matching is performed (as discussed in Section 3.1 in [1]). Below we describe two optimizations to overcome these bottlenecks.

4.1 Email Filtering by Indexing

Step 2 of HiddenEmailFinder requires that a quoted fragment F be matched against every single email M in the primary folder MF , as well as with every email in the reference folders RF_1, \dots, RF_k . The reference folders are useful because a fragment may be hidden from MF simply because the user filed the original email into another folder.

The matching in Step 2 stops when either a match is found (in which case the quoted fragment F is not a hidden fragment), or a match is not found anywhere (in which case F is considered hidden). When the folders are small, a straight-

Algorithm EmailFiltering

Input: a word index, a frequent word list FW , a quoted fragment F
Output: a list of email ids possibly matching F

1. Tokenize F to a set of words w , and remove all the stop-words.
2. For each w not in the list FW , use the word index to identify L_w .
3. Return the unioned list, i.e., $\cup_{w \in F \wedge w \notin FW} L_w$.

Figure 3: A Skeleton of Algorithm EmailFiltering

forward string comparison is acceptable. But this does not scale up to large folders.

The first optimization is to use a word index. Each index entry is of the form: $\langle w, L_w \rangle$, where w is a word in the email corpus, and L_w is a list of ids of emails containing at least one occurrence of w . For example, the word “available” may have the following entry in the word index: $\langle available, \langle id = 17, id = 278 \rangle \rangle$. Like information retrieval systems, the word index does not contain high frequency closed-class terms (i.e., stop-words) such as the definite article “the”. In general, the index is created by one complete pass over all the emails in the corpus.

Given the word index and a quoted fragment F to be matched, the quoted fragment is first tokenized to words with all the stop-words removed. Then for each word w , the word index is used to return the list L_w . To support LCS matching between a quoted fragment and an email, a match is allowed even if not all the words are found in the email. Thus, we take the union of the lists, i.e., $\cup_{w \in F} L_w$. This filtering process guarantees no false dismissals in the sense that only emails in the unioned list can ever match F (unless F is made up of stop-words only!).

Figure 3 shows a skeleton of this process. It incorporates an additional optimization to reduce the size of the unioned list $\cup_{w \in F} L_w$. Specifically, it further excludes the most frequent open-class words (i.e., non-stop-words) in the corpus. Hereafter, we denote this list of words as FW . By doing so, we only obtain the unioned list $\cup_{w \in F \wedge w \notin FW} L_w$. We define the length of FW as *frequent word threshold* ft , i.e., the top- ft most frequent words are kept in FW . In Section 5.4, we show that the choice of frequent word threshold has a great impact on the runtime.

4.2 LCS-Anchoring by Indexing

Let us now review the role of LCS in Step 2 of algorithm HiddenEmailFinder. Suppose that the original email is a sequence of fragments $OM = \langle F_1, F_2, \dots, F_5 \rangle$. When the user quotes this email, the user might perform various actions to this sequence, as she can edit the fragments as free text. She can quote the exact sequence verbatim; or she can delete the beginning and/or the end parts (e.g., $QF_1 = \langle F_2, F_3, F_4 \rangle$). In a more sophisticated setting, she may quote $QF_2 = \langle F_2, F_4 \rangle$ to reduce the length. Furthermore, she may copy another fragment F_6 from another email to form $QF_3 = \langle F_2, F_6, F_4 \rangle$.

So given a quoted fragment, the task is to match it against other emails. In the case of QF_1 , a simple substring searching is sufficient to determine that QF_1 originates from OM . However, substring searching is not able to handle QF_2 and QF_3 . In contrast, LCS matching can correctly handle QF_1, QF_2 and QF_3 . Here the question is whether a simple substring matching is sufficient for real data. In our experimentation with the Enron dataset, we find that many quoted fragments contain quotations from more than one

Algorithm LCS-Anchoring

Input: the word index, the frequent word list FW , a quoted fragment F , and an email M

Output: the LCS between F and M

1. Tokenize F to a set of words w , removing the stop-words and keeping only those not in FW .
2. If M does not appear in any of the lists L_w for all the remaining w 's, return the empty string.
3. Otherwise, for each such w ,
 - (a) for each anchor position pos_i
 - i. Align w at pos_i in M and F .
 - ii. Expand the matched substring forward and backward as much as possible.
4. Return the longest matched substring in the nested loop.

Figure 4: A Skeleton of Algorithm LCS-Anchoring

email. Thus, to maximize the robustness of HiddenEmailFinder, it is necessary to use LCS.

The problem with LCS is that its complexity is quadratic in the length of the fragment and the email. For long emails and/or quotations, implementing LCS naively is not scalable. We propose to extend the word index from the email filtering step to tackle this problem. In particular, for each email in the list L_w , we also record the positions at which the word w occurs in the corresponding email, i.e., each entry in L_w is of the form $\langle id, \{pos_1, \dots, pos_k\} \rangle$. For example, the word “available” may have the following index entry: $\langle available, \langle \langle id = 17, pos = \{89, 3475\} \rangle, \langle id = 278, pos = \{190, 345, 3805\} \rangle \rangle$.

Then given a quoted fragment F , as before, F is tokenized to words. For each word w , and each email M in L_w , we can use the list $\{pos_1, \dots, pos_k\}$ as “anchors” to facilitate the matching between F and M . For example, let us say that F contains the word “available.” Then position 89 in email 17 is used as an anchor to match up the word “available” in F and email 17. By expanding forward and backward from the anchor position as much as possible, the longest common substring with respect to the anchor position is formed. Similar anchoring and expansion occurs at position 3475 in email 17, and the three specified positions in email 278. If the quoted fragment is tokenized to multiple words, the above process is conducted for each word w , and the longest common substring is picked. Figure 4 gives a skeleton of this optimization step called LCS-anchoring. This optimization is intended to be used in Step 2 of HiddenEmailFinder. It can also be used in Step 3 to optimize the identification of possible overlaps between hidden fragments.

5. THE ENRON CASE STUDY

5.1 The Data and the Setup

The Enron email dataset was made public by the US Federal Energy Regulatory Commission during the ex-Enron investigation. A detailed description can be found at www-2.cs.cmu.edu/~enron. To the best of our knowledge, this is the largest public accessible email dataset. This dataset contains about half a million messages belonging to 150 users and 3500 folders with all attachments deleted. Many analyses and preprocessing studies have been done on the Enron dataset. For example, the SIAM’05 Workshop on Link Analysis, Spam Detection and Anti-terrorism published several indexes of the Enron dataset. In our experiments, we use

their word indexes instead of building our own. The word index contains 160,203 unique words. Recall that whenever we refer to a frequent word threshold (ft), we mean that the top- ft words are considered too frequent to be used in EmailFiltering or LCS-Anchoring. Below we vary ft from 1000 to 80,000, corresponding to about 0.6% and 50% respectively.

For most of the results reported below, we focus on the inbox folders of the users. Of the 150 users, 137 have an inbox folder. The number of emails in those folders ranges from 3 to 1466. The average and median number of emails are 327 and 223 respectively.

5.2 Prevalence of Hidden Fragments

For each user in the Enron dataset, we identify all the hidden fragments in the inbox folder. Figure 5(a) shows the number of emails that contains at least one hidden fragment. Due to lack of space, we only show the largest 50 inbox folders sorted by ascending folder size (ranging from 338 to 1466 emails). As can be seen from this figure, there are 5 inbox folders with more than 300 emails containing at least one hidden fragment.

While Figure 5(a) shows the absolute values, Figure 5(b) displays the percentage of emails containing at least one hidden fragment (i.e., relative to the folder size). Because percentages may not make sense for small folders, we exclude folders with less than 50 emails. The x-axis of the graph shows the percentage, ranging from 0% to 60%. The y-axis shows the number of users with the given percentage. It is interesting to see that about half of the users are within the range of 15% to 30%.

The reader may wonder whether hidden fragments just represent a phenomenon of the user being diligent in filing her emails into an appropriate folder. To examine this effect, we check other folders of the same user. Hereafter, we refer to a hidden fragment as “global” if it is a fragment that cannot be found in *all* the folders of the user. We refer to a hidden fragment as “local” if it is a hidden fragment within the (inbox) folder, but is otherwise found in some other folder of the user. Let us denote the numbers of global and local hidden fragments by n_g and n_l respectively. We define the *recollection rate* as the ratio of $n_l/(n_l + n_g)$. That is to say, the closer the ratio is to 1, the smaller is the number of global hidden fragments.

Figure 5(c) shows a histogram of the recollection rates for all the users. It is interesting to see that most users have a recollection rate of less than 15%. That is to say, there is less than 15% of hidden fragments that can be found in the other folders of the user. Thus, hidden fragments do not seem to be simply a phenomenon of the user filing the emails to other folders; they are truly missing from the user’s folders.

While it is clear that hidden fragments are prevalent in the Enron corpus, the immediate question here is how general this phenomenon is for a “typical” real user. Let us review how the Enron dataset was prepared. As reported in [3], emails were deliberately deleted from the first published version of the Enron dataset on the users’ request for privacy and redaction purposes. It is estimated that about 8% of the emails were deleted to form the current version. Consider the following two aspects:

- First, the deleted emails are believed to be more personal in nature. It is reasonable to assume that they




Figure 5: Prevalence of Hidden Emails

were less likely to be quoted in many Enron business emails.

- Second, as will be discussed later in Section 5.3, the average number of reconstructed hidden emails per user is about 60. Given that the average inbox folder size is 327 emails, if the 8% of the deleted emails were evenly distributed in the inbox folder of each user, this would correspond to 26 emails in the folder. The gap between 26 and 60 is significant.

The question of whether hidden fragments are prevalent in a typical user’s folder is hard to be answered definitively. But for the two reasons discussed above, the abundance of (global) hidden fragments we found in the Enron corpus may well generalize to other real datasets.

There is actually another interesting point to make here. Emails were deleted partly to protect privacy. However, some of these deleted emails may be recoverable from emails in other folders. Thus, if there is an initial set of emails to be protected, the framework that we develop here can help to strengthen the protection by identifying other emails quoting the initial set.

5.3 Robustness

A key parameter in HiddenEmailFinder(Step 2&3) is the choice of $minLen$, the threshold for considering an overlap significant in identifying hidden fragments and in creating the precedence graph. A larger $minLen$ can avoid incorrect matching of common sentences and quotations, and hence the resulting precedence graph is less connected. On the other hand, it may also miss short quotations, and create more hidden fragments. The following table shows the average and median number of reconstructed hidden emails with respect to $minLen$ across all the inbox folders. As $minLen$ increases from 40 to 160, there is only a very slight increase in the number of reconstructed hidden emails. This shows that our algorithmic framework is stable with respect to the choice of $minLen$.

minLen	40	60	80	100	120	160
average	61.5	63.5	64.8	65.4	65.8	66.6
median	38	38.5	40	40	41	41

Recall from Step 3 of HiddenEmailFinder that when we build the precedence graph, if two hidden fragments overlap sufficiently, we assume that they quote the same original message. However, this may include some false positives by some common phrases contained in two unrelated emails.

One significant example is the usage of signature files of users, public email systems, etc. In the Enron dataset, there is also the copyright claim. We need to avoid such cases to be counted as overlap in the generation of the precedence graph.

One solution is to automatically identify and exclude signature files and copyright claims before matching quoted fragments. This is itself a research problem [2]. In our experimentation, we did manual detection on the Enron users and recognized that a large majority of the signature files and copyright claims were enclosed in two identical lines of asterisks. Thus, we programmed a simple regular expression filter to remove all such matches. They were excluded from all the results reported here.

5.4 Effectiveness of Optimizations

The intent is to observe the tradeoff between runtime efficiency and output quality. We design our experiments in the following way. In the first round of experiments, we only apply the heuristic of EmailFiltering and change the frequent word threshold (ft). We vary ft from 500 to 80,000. In the second round of experiments, we apply both EmailFiltering and LCS-Anchoring. For both rounds, we record the runtime and the number of reconstructed hidden emails for all the inbox folders.

Figure 6(a) shows the median runtime performance. The x-axis is drawn in log scale of the frequent word threshold ft . Let us first focus on the curve applying only the EmailFiltering algorithm. The basic, unoptimized version of HiddenEmailFinder corresponds to the case when $ft = 0$. The median runtime for this case is about 10 minutes, which is not shown in this figure. As the value of ft increases, the runtime improves by as much as 2 orders of magnitude, down to less than 10 seconds for $ft = 10,000$.

The second curve in Figure 6(a) shows the additional gain in efficiency when LCS-Anchoring is applied on top of EmailFiltering. The gap between the two curves shows that there is a definite bonus in applying LCS-Anchoring. The gain becomes smaller as ft increases because EmailFiltering alone has already eliminated a lot of emails required for matching, thereby reducing the number of times that LCS-Anchoring is performed.

The question is whether the significant gain in efficiency is achieved through reduced quality. Figure 6(b) shows that the number of reconstructed hidden emails when ft changes from 1000 to 80,000. As ft increases from 1000 to 80,000, the number of reconstructed hidden emails increases very slightly, reflecting the reduced connectivity of the prece-

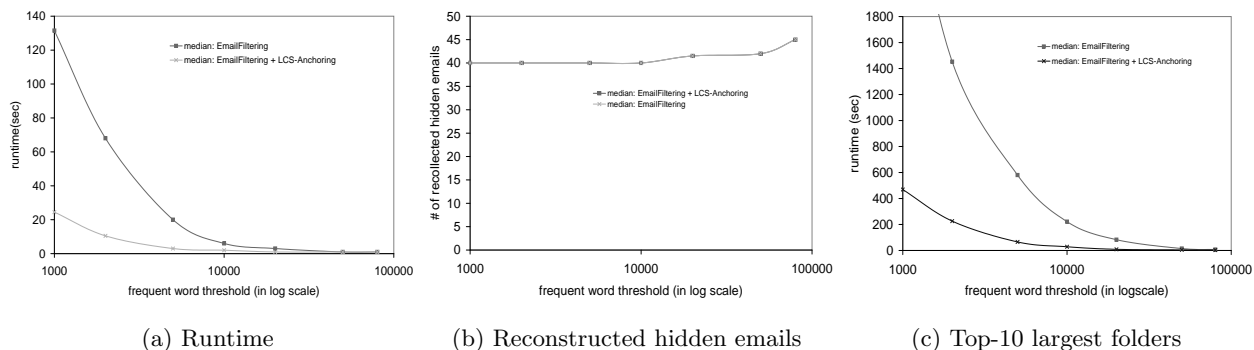


Figure 6: Effectiveness of Optimizations

dence graph. Given that the two curves in Figure 6(b) almost completely coincide, it is clear that both EmailFiltering and LCS-Anchoring can bring about a gain in efficiency without causing a degradation in the output quality.

Figure 6(a) does not include the average runtime because there is a large discrepancy between folders on how long it takes to process them. Figure 6(c) shows the extreme case of the top-10 largest folders. Among these top-10 folders, the median folder contains 1,152 emails, with 37 emails each longer than 1,000 words. Large folders and long emails take significantly more time than the smaller ones. The two curves in the figure show the median runtime across the 10 folders when EmailFiltering alone and when EmailFiltering and LCS-Anchoring are applied. Like in Figure 6(a), it is clear that both techniques are effective. But unlike in Figure 6(a), this time the gap is far more significant when LCS-Anchoring is applied. This convincingly shows the importance of LCS-Anchoring for long emails and large folders.

6. CONCLUSION

This paper studies the problem of reconstructing hidden emails using the embedded quotations found in messages further down the thread hierarchy. We optimize the basic HiddenEmailFinder algorithm to deal with large folders and long emails. The two optimizations are based on word indexing to reduce the number of emails that need to be matched and to reduce the amount of effort required to find the LCS between the fragment and the email under consideration. As a side benefit, given an initial set of emails to be protected, HiddenEmailFinder may be used to strengthen the protection by identifying other emails quoting the initial set.

Another key contribution of this paper is the Enron case study. From our experimentation, many valuable lessons are learned. First, we observe that global hidden fragments are prevalent in the Enron corpus. This prevalence may well generalize to other real datasets. Second, we show that our framework is robust in dealing with real folders. Last but not least, we show that both the EmailFiltering and the LCS-Anchoring techniques are effective in providing scalability to large folders and long emails.

Our future plans include applying natural language understanding techniques to make even more intelligent decisions about piecing fragments together and representing them to the user. This work is an integral part of a larger project on text mining. The goal is to help users better manage their emails. We plan to develop email summarization and classification tools. These tools will be based on various kinds

of graphs extracted from the emails, in a style similar to the precedence graph that HiddenEmailFinder generates.

7. REFERENCES

- [1] Giuseppe Carenini, Raymond Ng, Xiaodong Zhou and Ed Zwart. Discovery and Regeneration of Hidden Emails. *ACM Symposium on Applied Computing(SAC)*, Santa Fe, New Mexico, USA, March, 2005, pp. 503–510.
- [2] Vitor R. Carvalho and William W. Cohen. Learning to Extract Signature and Reply Lines from Email. *First Conferences on Emails and Anti-spam*, Mountain View, CA, USA, July, 2004.
- [3] US Federal Energy Regulatory Commission. <http://www.ferc.gov/industries/electric/indus-act/wem/pa02-2/info-release.asp>
- [4] Gwizdka, J, Chignell. M.H. Individual Differences and Task-based User Interface Evaluation: A Case Study of Pending Tasks in Email. *Interacting with Computers, Elsevier Science*, Minneapolis, Minnesota, USA , 2004, v.16(4) pp. 550–551.
- [5] Bryan Klimt and Yiming Yang. The Enron corpus: a new dataset for email classification research. *European Conference on Machine Learning (ECML)*, Italy, 2004, pp. 217–226.
- [6] Derek Lam, Steven L. Rohall, Chris Schmandt and Mia K. Stern. Exploiting E-mail structure to improve summarization. *CSCW'02 Poster Session*, New Orleans, Louisiana, United States, 2002,
- [7] Ani Nenkova and Amit Bagga. Email classification for contact centers. *ACM Symposium on Applied Computing(SAC)*, Melbourne, Florida, 2003, pp. 789–792.
- [8] Paula S. Newman. Exploring discussion lists: steps and directions. *Proceedings of the second ACM/IEEE-CS joint conference on Digital libraries*, Portland, Oregon, USA, 2002, pp. 126–134.
- [9] Steven L. Rohall. Reinventing email. *CSCW'02 Workshop: Redesigning email for the 21st century*, Portland, Oregon, USA, November, 2002.
- [10] Kulesh Shanmugasundaram and Nasir D. Memon. Automatic Reassembly of Document Fragments via Context Based Statistical Models *ACSAC*, 2003, pp. 152–159.