
Reducing E-Discovery Cost by Filtering Included Emails

Tsuen-Wan “Johnny” Ngan
Symantec Research Labs

Abstract

As business activities becoming more digitalized, electronic information is often produced as vital evidence during civil litigation. The process of discovering information as evidence is getting increasingly expensive as the volume of data explodes. This surging demand calls for a solution to reduce the cost associated with discovery. In this paper, we propose filtering included emails as a means to reduce the volume of review data, and describe efficient algorithms to identify those emails. Experiments show that this can reduce the number of emails to be reviewed by 20% in a corporate email corpse.

1 Introduction

Email has become an indispensable part of communication for enterprises of any size, both internally and externally. It has long been considered as an official means of communication. Nevertheless, its speed and convenience come with a hefty price of management and storage overhead.

The problem is highlighted with the introduction of legislation such as the Health Insurance Portability and Accountability Act (HIPAA) in 1996 and the Sarbanes-Oxley Act (SOx) in 2002, both of which have significantly increased the importance of managing and storing all information within enterprises. Email, in particular, has become one of the most important content types that needs to be retained. This is further complicated by the fact that retention rules could vary based on users, locations, file types, and whatnot. In fact, for larger enterprises, the number of items to be retained could easily be multi-billions.

The management of this sheer volume of data has become a formidable task. Off-the-shelf storage products and solutions often provide basic processing features, such as whole-item de-duplication. Also known as single-instance storage, it allows the same item retrieved from different

locations and users to be stored only once. Other common functions are usually based on email headers, allowing boolean queries to select and filter emails. However, these limited features cannot meet the new demands for *e-discovery*.

E-discovery, or electronic discovery, refers to discovery of evidence from electronic data in civil litigation. Unlike traditional discovery from paper information, the uniqueness of electronic information stems from its intangible form, volume, transience, and persistence. Furthermore, electronic information is usually accompanied by metadata, which is rarely present in paper information.

E-discovery is often time-consuming and expensive. Litigations sometimes require humongous amount of emails to be reviewed to extract evidence. As accuracy is the primary goal, this process is usually done manually by attorneys that charged by the hours or data size. A survey conducted in 2006 [Fulbright, 2006] on senior corporate counsels showed that companies are facing growing numbers of lawsuits, and fewer are escaping them, even as arbitration numbers also grow. The cost of this active environment for dispute is very high: Companies with \$1 billion or more in annual revenues spent an average of \$31.5 million on all their legal matters, and 40% of the study samples had at least one lawsuit filed against them in which \$20 million or more was at issue. Several existing products aim to facilitate and streamline the e-discovery process [Symantec, 2008, EMC, 2008, ZANTAZ, 2008], but little is done on reducing the amount of emails to be reviewed by analyzing email content.

This paper considers practical ways to alleviate the burden on the manual review process. The key insight is that during e-discovery, any email that is completely included by another email in the repository can be safely ignored, without risking to lose evidence. To that end, we propose an efficient paragraph-based algorithm for statically and dynamically finding all email inclusions in an repository. We also experimented the algorithms against the Enron corpus and a mailing list trace.

The rest of this paper is organized as follows. Section 2 describes the process of e-discovery and explains the advantages for filtering included emails. Section 3 depicts our paragraph-based comparison algorithm to quickly find inclusions. The algorithm is tested on two data sources in Section 4, with discussions in Section 5. Finally, Section 6 covers related work and Section 7 concludes.

2 The E-discovery Process

Due to the potential of litigations, companies are required to preserve electronic information. E-discovery is the subject of amendments to the Federal Rules of Civil Procedure [Supreme Court, 2006], effective since December 2006. The amendment is concerned with a company's duty to preserve and produce electronically stored information in the face of litigation or pending litigation. Parties involved in civil court proceedings must provide a list of all electronically stored information that might be relevant to the case. Involved parties must also discuss the forms in which this information should be produced, and the party who requests the information may specify the form or forms in which electronically stored information is to be produced.

Since e-discovery is relatively new, there is no well-established standards and procedures for enterprises to follow. E-discovery typically involves a number of phases, and various industry bodies have proposed standards as to how each phase should be satisfied. One of those industry bodies, called the Electronic Discovery Reference Model project (EDRM) [EDRM Group, 2008], is backed by many technology and service vendors in the discovery market. This model, laying down e-discovery standards and guidelines, contains the following steps:

1. **Identification.** The scope, breadth, and depth of electronically stored information that might be pursued during discovery is first assessed and determined.
2. **Preservation.** The stored information is protected against destruction and alternations.
3. **Collection.** The stored information is gathered from various sources, including backup tapes, drives, and portable storage devices.
4. **Processing.** The overall set of data is reduced by setting aside duplications and data that are irrelevant due to their type, origin, or date.
5. **Review.** The collected information is evaluated for relevance.
6. **Analysis.** The discovery materials are analyzed to determine relevant summary information.
7. **Production.** The information is delivered to various recipients on various media.
8. **Presentation.** The information is finally presented at depositions, hearings, and trials.

Just a reminder to make sure you get Tim a list of books/regions that are currently designated to be traded. We need to change trader names etc. so make sure he gets us this information as soon as possible.

Stacey

Also, will you get Tim to give direction on what traders will be allowed to trade what regions.

Thanks,
Stacey

—Original Message—

From: White, Stacey W.
Sent: Tuesday, January 29, 2002 3:45 PM
To: Postlethwaite, John
Subject: Reminder

Just a reminder to make sure you get Tim a list of books/regions that are currently designated to be traded. We need to change trader names etc. so make sure he gets us this information as soon as possible.

Stacey

Figure 1: A typical business email and a follow-up email that includes the original one.

As the volume of data decreases down the steps, the relevance remains in the volume rises. Among these steps, review and analysis are usually the most expensive, since a large amount of data still needs to be manually processed by attorneys. Moreover, they are specific to a particular litigation, and need to be redone for each new litigation. Hence, any optimization should be done on minimizing the volume of data involved in these steps.

Consider an email thread with multiple emails in a corporate environment. When an email is replied or forwarded, it is common practice to quote the content of the original email fully in verbatim. From a reviewing perspective, it suffices to read only the last email in the thread, which already contains all other emails. Thus, all but the last email can be filtered without losing evidence. An example is shown in Figure 1.

Of course, a response email may or may not quote its parent email in entirety, and email threads are not necessarily linear. It is therefore necessary to compare email content to determine which emails can be safely filtered.

This design not only can reduce the number of emails that need manual review, it can also group emails that are partially included together, even when other threading heuristics [Yeh & Harnly, 2006] fail. A reviewer can then read these related emails together, making the review process more efficient.

3 Finding Email Inclusions

This section describes our proposed algorithms for finding inclusions.

3.1 Defining Inclusions

Conceptually, an email is “included” when its content is completely present in other emails. But this definition with respect to e-discovery is unclear. We need a definition so that the meaning of an included email would be captured, but at the same time not too restrictive that it would be rendered useless.

Recall the purpose of finding inclusion is to allow reviewers to skip certain emails, with the guarantee that they would not miss any evidence. To that end, inclusion must be defined conservatively. The same sentence or paragraph in different emails may carry different meanings. For instance, it would be incorrect to conclude that a reply of “Yes” can be ignored if there are some unrelated emails containing “Yes.” Thus, we require an email to be completely included by another email.

What should the smallest unit for comparison be? To conserve its meaning, the minimal basic unit should be sentences. But note that paragraphs are usually unmodified after quotation. By expanding the basic unit to paragraphs, it can greatly reduce the number of items for comparisons without affecting much on accuracy.

Of course, this paragraph-based definition may result in false negatives (failure to find an inclusion). For example, when a paragraph is broken into two in a reply. For the purpose of reducing volume of review data, these false negatives would only prevent some emails from being filtered, but would not remove potential evidence. We consider this as an acceptable tradeoff for faster searching.

3.2 Paragraph-Based Inclusion Check

We now present a high-level design of our paragraph-based inclusion check. Each email is represented by a set of digests, where each digest represents a paragraph. Fundamentally, inclusion check can be done with two tests:

- *Subset test*: Given email E , determine if E is contained by any member of a set of emails.
- *Superset test*: Given email E , determine if E contains any members of a set of emails.

With a large set of emails, the superset test is much more expensive than the subset test. Imagine an inverted index is built that maps paragraphs to emails that contain them. For the subset test, it suffices to use *any* paragraph in E to select candidate emails. However, for the superset test, it is necessary to consider *all* paragraphs in E .

```
-----Original Message-----
From: Allen, Phillip K.
Sent: Friday, December 07, 2001 5:14 AM
To: Dunton, Heather
Subject: RE: West Position

> At 04:18 PM 10/1/2001 -0500, you wrote:
```

Figure 2: Two examples of software-generated quotation text in email content.

Because of the difference in performance, the two tests are suitable for different use cases. For an existing email archive, all emails are present at the very beginning. To find inclusions, it suffices to iterate the subset test over each email once. However, for a live email archiving system, when a new email arrives, it is necessary to check inclusion in both directions.¹ Therefore, a tool that is tightly integrated to such a system would need to run both tests on each new email. This paper addresses the latter, more computational-intensive problem.

3.3 Implementation Details

The rest of this section describes details of our proposed algorithm. The algorithm consists of preprocessing individual emails and finding inclusions between them.

3.3.1 Preprocessing Individual Emails

The preprocessing converts each email into a set of digests. It involves three steps: (1) Removing extraneous text; (2) dividing email into paragraphs; and (3) hashing each paragraph. Since some of these steps depend on the language and its usage practice, for simplicity we assume all emails are in English.

The first step is to remove extraneous text. Extraneous text is defined as parts in an email that are not inputted by human users. This includes email headers and quotation text in content that are automatically generated by email client software. Examples of the latter are shown in Figure 2.

The remaining email content is then divided into paragraphs by locating paragraph separators. For text emails, a paragraph separator could be a line with no alphanumeric characters or a paragraph separator character defined in Unicode. For html emails, it could be a paragraph tag (<p/>) or consecutive line break tags (
).

The last step is to compute a digest for each paragraph. Since we are only interested in text and not its format, only alphanumeric characters are extracted from

¹Note that while new emails are unlikely to be included by existing emails in normal circumstances, emails may not arrive in chronological order. For example, they may be collected from different sources at different time.

each paragraph and hashed. This ensures that automatic alternations like insertion of quoting symbols (e.g., $>$) and line breaks would not affect the generated digests. A universal hash function [Carter & Wegman, 1979] like UMAC [Black et al., 1999] is preferred due to the use of Bloom filters below, but any commonly used ones like MD5 or SHA-1 would be sufficient.

3.3.2 Finding Inclusions

Each email is already converted to a set of digests. The problem is reduced to running both subset test and superset test efficiently.

If the superset test is performed by iterating over all paragraphs and check all emails that contain them, the process would be very slow, as some paragraphs appear in a large number of emails (see Section 4.1.4). Every time when a new email containing a popular paragraph arrives, it would be necessary to iterate over all emails that contain that popular paragraph. To expedite the process, we need to handle popular paragraphs separately.

Popular Paragraphs

The algorithm partitions paragraphs into *popular* and *unpopular* based on a threshold of occurrences. Under this partitioning, emails are further classified into two sets: (1) Emails with only popular paragraphs (called \mathcal{P}); and (2) Emails with at least one unpopular paragraph (called \mathcal{Q}). Two inverted indices are built, mapping paragraphs to emails in \mathcal{P} and \mathcal{Q} .

For a paragraph to be popular, it must appear in a large number of emails. Many of them may be unrelated to each other. Examples include signatures and greetings gratitude. Since emails usually contain other pieces of useful information, very few emails should fall into \mathcal{P} .

To search for inclusion candidates, the algorithm iterates over all paragraphs in the email, but handles popular and unpopular paragraphs differently. For each unpopular paragraph, since they are only present in \mathcal{Q} , all emails that contain it in \mathcal{Q} are added to the candidate set. However, for each popular paragraph, it suffices to add only emails in \mathcal{P} . This is because any email in \mathcal{Q} would also contain unpopular paragraph, so if it is an inclusion, it would be already selected. Thus, we avoid the worst case, namely, iterate over all emails that contain a popular paragraph.

Once the set of candidate emails are selected, it is still necessary to verify the inclusion. The set of candidate emails can still be very large, since any email that share any common paragraph is in the set. Here we use Bloom filters to filter out most false positives.

Bloom Filter

Bloom filter [Bloom, 1970] is a space-efficient probabilis-

tic data structure used to test set membership. It maintains a bit vector initialized to all 0. Each element added to the filter will set some bits to 1. A set membership test is done by checking if all the bits corresponding to the given element are set to 1. There are usually more elements than the number of bits available, so Bloom filters may have false positives but not false negatives. We extend this data structure to test for subset between emails.

Note that some emails contain a large number of paragraphs (see Section 4.1.2). To avoid many emails from setting their Bloom filters to entirely 1 and ineffective, we correspond each paragraph to only one bit.

Once we have the Bloom filters for two emails, we can quickly test for inclusions. If email A contains all the paragraphs in email B , all the bits that are set in B 's Bloom filter would also be set in that of A . Hence, the bitwise AND of the two Bloom filters should be the same as B 's Bloom filter.

Of course, similar to set membership test, this inclusion test can still have false positives. After this test is passed, it is imperative to compare the sets of digests from the two emails to conclude with certainty that they are indeed inclusion.

Blacklisting Uninteresting Paragraphs

A possible optimization is to ignore popular paragraphs that do not contribute much in content. Examples include company names, greetings, gratitude, signatures, advertisements (from webmails or electronic devices), boilerplates, and subscription information for mailing lists. While technically their presence can affect inclusion decisions, these paragraphs can be ignored without affecting filter quality.

4 Experiments

In this section, we examine the proposed algorithm against two datasets. We also investigate the effectiveness of removing included emails and the performance improvement from using Bloom filters.

4.1 Datasets

To understand the performance of the algorithm under different inputs, we use two datasets: The *Enron email corpus* and a *mailing list dataset* we created. Here we describe these datasets with some comparisons and analyses.

The Enron email corpus contains data from about 150 users, mostly senior management of the now-defunct company Enron Corporation [Klimt & Yang, 2004]. It was originally made available by the Federal Energy Regulatory Commission during its investigation. The data does not include attachments, and some emails have been deleted "as part of a redaction effort due to requests from affected em-

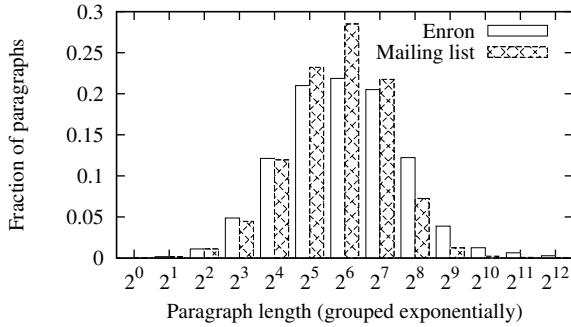


Figure 3: The length distribution of paragraphs, measured in alphanumeric characters, grouped exponentially. The mailing list dataset has more middle-sized paragraphs while the Enron corpus has more longer paragraphs.

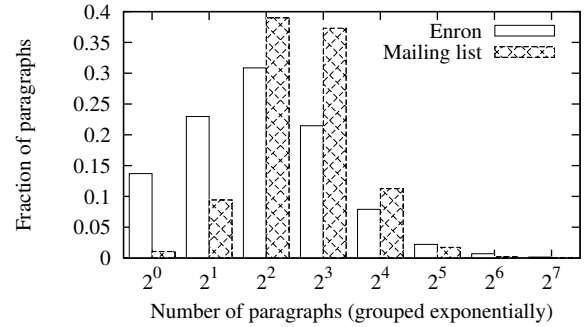


Figure 4: The distribution of the number of paragraphs per email, grouped exponentially. The Enron corpus has more shorter emails, likely due to short personal notes and sending attachments between colleagues.

ployees.” This serves as a corporate email trace.

The mailing lists dataset is collected from four mailing lists at December 2007 through their archive websites: The Cygwin project mailing list archive [Cygwin, 2007], general Gentoo user support and discussion mailing list [Gentoo, 2007], MySQL general discussion [MySQL, 2007], and qmail mailing list [qmail, 2007]. This serves as a general email discussion trace.

A summary of the two datasets is shown in Table 1. The bottom part of the table shows statistics after de-duplication, which is done by comparing MD5 digests on alphanumeric characters of each email. Less than half of the emails remained after de-duplication in the Enron corpus, but only 5% are removed in the mailing list dataset. This is because the Enron corpus includes “Sent” folders, and some emails are sent to multiple recipients.

Even though the two datasets are from different sources, below we show that they possess similar properties.

4.1.1 Paragraph Length

Figure 3 shows the length distribution of distinct paragraphs, measured in the number of alphanumeric characters. It shows that both datasets have similar fraction of shorter paragraphs. However, the mailing list dataset has more middle-sized paragraphs, whereas the Enron corpus has more longer paragraphs.

The former is probably because all the mailing lists are related to programming and software and often contain moderately long code snippets. The latter can be attributed to two reasons: (1) Some emails in the corpus were converted from other formats to text, but during the conversion paragraph margins were not properly retained; and (2) some emails contain very long XML documents that are stored in one large paragraph.

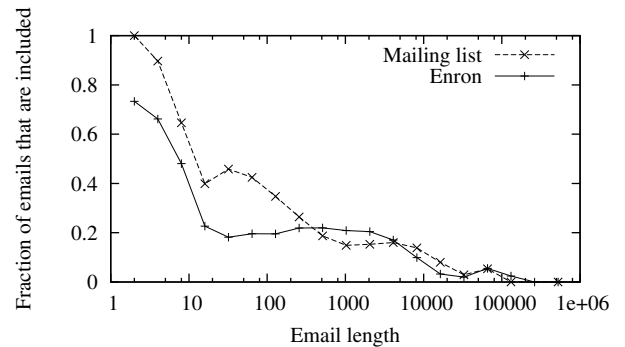


Figure 5: The correlation between number of paragraphs and the probability that it is included.

4.1.2 Paragraphs Per Email

Figure 4 shows the number of paragraphs per email. In both datasets, most of the mails contain a small number of paragraphs. The main difference between the two datasets is that the Enron corpus has more shorter emails, likely due to short personal notes and sending attachments between colleagues. On the other hand, postings on mailing lists usually contain at least a few paragraphs to describe or answer the questions, or to include code snippets.

4.1.3 Size and Likelihood of Inclusion

The size of an email, measured in either alphanumeric characters or number of paragraphs, has a negative correlation to the probability that it is included. Figure 5 and 6 shows that most included emails are relatively short (less than 100 alphanumeric characters or 10 paragraphs).

4.1.4 Paragraph Appearances

Figure 7 shows the number of appearances of distinct paragraphs over each entire repository. For both datasets, the frequency of appearance goes down logarithmically even as each group grows exponentially in size. Both curves have

Table 1: The two datasets and their detailed properties. The bottom part shows statistics after de-duplications.

	Enron	Mailing list
Data type	Corporate emails	Mailing list discussions
Number of emails	517,431	486,869
Total size (excluding headers)	961MB	680MB
Average size per email	1,858 bytes	1,397 bytes
Average number of paragraphs per email	8.32	9.62
Number of distinct paragraphs	1,020,319	1,916,744
Number of emails after removing duplications	248,517	464,766
Fraction of remaining emails	48.03%	95.46%
Total size of remaining emails	473MB	656MB

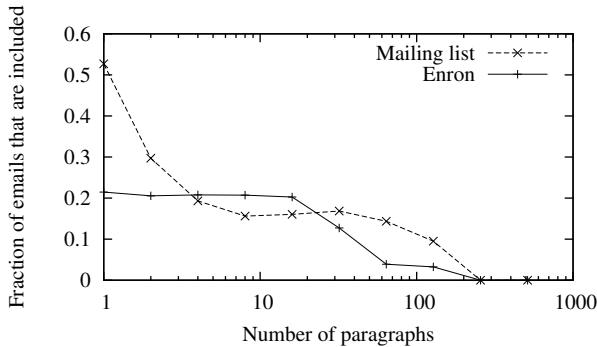


Figure 6: The correlation between email length and the probability that it is included.

very long tails: The most common paragraph appeared in the Enron corpse 23,052 times, whereas that for the mailing list dataset is 54,316.

Figure 8 shows the data in cumulative distribution curves. Most paragraphs appeared only for a few times. For example, for both datasets, less than 1% of the paragraphs appeared more than 10 times. Thus, a small popular threshold can be used to select a small fraction of popular paragraphs from the rest.

As some paragraphs appear in many emails, a possible application of paragraph-based analysis is to perform de-duplication at the paragraph level. Specifically, paragraphs that are identical can be stored only once, with duplicated paragraphs pointing to a common copy. We performed an informal measurement on this feature. Even after whole-item de-duplication and without compression, this feature can reduce storage space by an additional 22–28%.

4.2 Effectiveness of Filtering Included Emails

The goal of finding included emails is to reduce the volume of data to be reviewed. We first determine how well this idea may work in practice by measure the amount of emails that are completely included and can thus be ignored during e-discovery. Table 2 shows the fraction of included emails. In both datasets, even after de-duplication, this amounts to

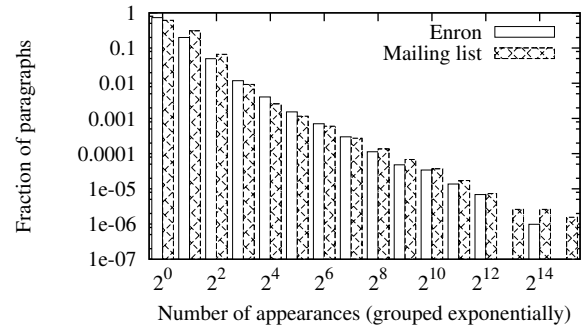


Figure 7: The distribution of appearance of paragraphs, grouped exponentially. For both datasets, the frequency goes down logarithmically even as each range expands exponentially.

Table 2: Fraction of included emails that can be removed before reviewing.

	Enron	Mailing list
No. of nonidentical items	248,517	464,766
Included items	50,787	87,544
Fraction of included items	20.4%	18.8%

around 20% of the emails, all of which can be removed without affecting the accuracy of reviewing.

4.3 Algorithm Performance

We built a prototype of the algorithm to get a handle of its performance. The prototype is ran on a Windows XP desktop with a Pentium-4 3.4 GHz CPU. The prototype uses Berkeley DB version 4.6.19 as a large hashtable for inverted indices. The prototype runs entirely in memory. The program is ran with and without a blacklist, which is manually built to include 41 popular but unimportant paragraphs.

Table 3 shows the running speed of the prototype averaged over the entire datasets. Note that these numbers do not including the time spent on reading emails from disk. The table shows that the algorithm can process multi-MB per

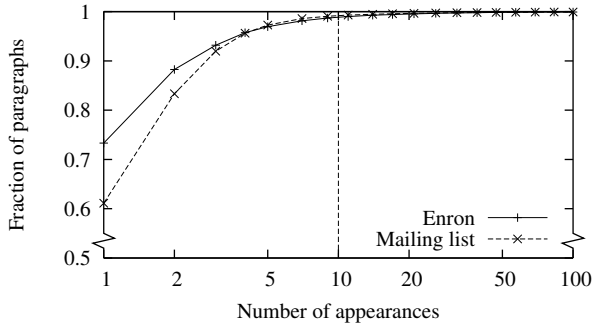


Figure 8: The cumulative distribution of appearances of paragraphs. The line shows a cutoff threshold of 10.

Table 3: Average running speed of the prototype algorithm. These numbers do not include the disk I/O time for reading emails from disk.

	Enron	Mailing list
Without blacklist	2.01 MB/s	2.48 MB/s
With blacklist	2.64 MB/s	3.99 MB/s

second, fast enough to run in a real-time system. Also, when a blacklist is used to filter unimportant paragraphs, the algorithm can run 31–61% faster. The prototype also seems to scale reasonably well over the datasets tested, as the last 1% of the emails are only 40–50% slower than the first 1%.

The rest of this section shows two more micro-benchmarks on two parameters used.

4.3.1 Bloom Filter

Bloom filter is used to filter out emails that are definitely not inclusions. With a Bloom filter size of 128 bits, its false positive rate is 2% for the Enron corpus and 0.9% for the mailing list dataset. This shows that a Bloom filter can be very effective even at a size much smaller than an average email. It also allows us to filter out most candidates without even looking at any of their digests.

Although Bloom filter is very effective, it is unclear how large it should be, as it is nontrivial to analytically derive the false positive rate and optimal Bloom filter size. The traditional analysis [Broder & Mitzenmacher, 2002] is for set membership test and is not applicable to our subset test. Moreover, unrelated emails often share common paragraphs. The conditional probabilities in the analysis cannot be assumed to be independent. Instead, we study the trade-off by experimentally measuring the false positive rate at different Bloom filter sizes.

The result is shown in Figure 9. Clearly, for both datasets, the false positive rate is reduced by roughly the same factor every time the Bloom filter size is doubled. This result is

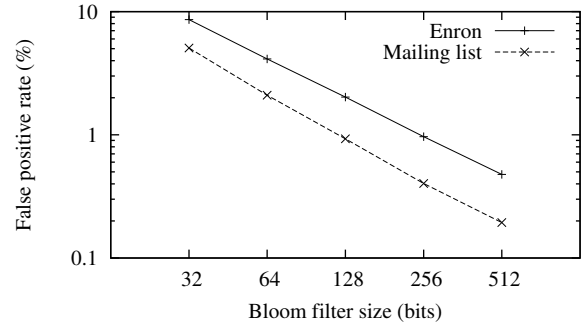


Figure 9: Effectiveness of Bloom filter for subset comparisons at different filter sizes. Similar to the original use of Bloom filters, doubling the size of the filter reduces the false positives by a factor.

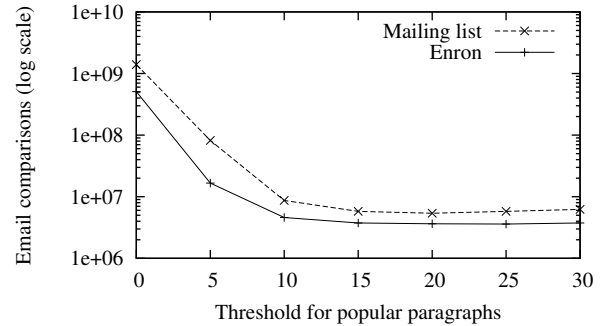


Figure 10: Filtering cost with different popular thresholds. A zero threshold represents the case when there is no separate popular paragraph set. Observe that a reasonable threshold reduces the filtering cost by a factor of 200.

similar to using Bloom filters for set membership test.

4.3.2 Popularity Threshold

Lastly, Figure 10 shows the cost of finding inclusions using different popular thresholds. The cost is measured in terms of the number of full set comparisons (i.e., the number of times the Bloom filter test is passed) needed. For both datasets, the number of comparisons reduced by two orders of magnitude when a popular threshold is used and is increased to 10, but remained roughly at the same level when the threshold is further increased.

This shows the division of popular and unpopular paragraphs is very effective. With this division, popular paragraphs are usually not used to search for inclusion candidates. This eliminates the worst case of iterating thousands of emails just because they happen to share a common popular paragraph. However, when the threshold is too small, too many emails are classified as emails with only popular paragraphs. This forces the algorithm to search with popular paragraphs more frequently.

5 Discussions

5.1 Minimum Paragraph Length

As shown in Section 4.1.4, some paragraphs appear in many emails. Among the most common paragraphs, some of them are very short. We decided against imposing an artificial minimum paragraph length. This is because short paragraphs, like a simple “No,” can still convey important meaning. If they are removed, it might cause some important emails to be considered included and ignored, which is a risk that cannot be taken.

5.2 Email Threads

Email threads can be used to guide searching for inclusions, but are not used here. First, they are rarely complete with different email clients. Second, users may copy and paste content, which cannot be tracked by threads. Moreover, the algorithm is fast enough that it is feasible to run on the entire email repository anyway.

Conversely, inclusion can help to derive email threads. However, this benefit is also limited, as emails in a thread may or may not include each other.

5.3 Selecting Alphanumerical Characters

One interesting question is whether choosing only alphanumerical characters for hashing is reasonable. For instance, this does not work well against source code. However, note that e-discovery usually concern what did senior management know and discuss, so the focus is on textual data. Other domain-specific contents should be handled by specialized tools or experts.

5.4 Language Dependency

While our algorithms for finding inclusions are independent of the language used by the emails, the preprocessing needs to be aware of the language, as it has to know how to divide email content into paragraphs and remove formatting characters. Also, to remove extraneous text, it may require the preprocessor to recognize quotation text in the language used.

6 Related Work

An alternative to finding inclusion is to group emails into email threads, perhaps through header information or other heuristics [Yeh & Harnly, 2006]. While this would provide reviewers the same benefit of reviewing the same discussion at a time, further comparisons are still needed to filter out included emails. Moreover, these heuristic-based algorithms may not be accurate, and are ineffective for emails

that are inherently not in the same thread, for example when text is copied from one email and pasted into another.

Email signatures do not contain useful information, but they may appear in many unrelated emails and cause slow-down. A simple optimization is to remove these signatures before running our algorithms using a signature extraction algorithm [Carvalho & Cohen, 2004].

Fingerprinting techniques, e.g., Rabin fingerprinting [Rabin, 1981], coupled with fingerprint selection algorithms like Winnowing [Schleimer et al., 2003], can similarly generate digests for finding emails that share common portions. The same comparison algorithm can then be applied for finding inclusions. Our paragraph-based generation is more suitable for this application since it exploits the fact that paragraphs are usually unmodified in responses. Our algorithm guarantees to generate a digest for each paragraph, no matter how short it is, and only generates one digest for very long paragraph, which reduces the number of comparisons but increases accuracy.

There are other machine-learning research that work on email content, including automatic clustering [Huang et al., 2004] and classification [Surendran et al., 2005] and other techniques for filtering spams [Joachims, 1999, Sahami et al., 1998]. These techniques are unsuitable for e-discovery, as enterprises cannot risk missing any crucial evidence from their false negatives.

7 Conclusions

In this paper, we used inclusions to filter emails for e-discovery. We proposed a paragraph-based comparison to find inclusions, and algorithms for performing comparisons efficiently. Our experiment showed a 20% reduction of number of emails to be reviewed in an e-discovery process.

Our algorithm divide paragraphs into popular and unpopular. One interesting observation is that this division is a special case; paragraphs can be divided into a larger number of groups based on popularity, and likewise emails can be divided based on their least popular paragraphs. The benefit from a finer granularity of division is left as future work.

Acknowledgements

We would like to thank Guy Bunker for the original idea of filtering included emails to reduce volume of review data, and Tzi-cker Chiueh for helpful discussions and comments.

References

[Black et al., 1999] Black, J., Halevi, S., Krawczyk, H., Krovetz, T., & Rogaway, P. (1999). UMAC: Fast and

- secure message authentication. *Advances in Cryptology — CRYPTO '99*. Santa Barbara, CA.
- [Bloom, 1970] Bloom, B. H. (1970). Space/time trade-offs in hash coding with allowable errors. *Communications of the ACM*, 13, 422–426.
- [Broder & Mitzenmacher, 2002] Broder, A., & Mitzenmacher, M. (2002). Network applications of Bloom filters: A survey. *Proceedings of the 40th Annual Allerton Conference on Communication, Control, and Computing*. Urbana, IL.
- [Carter & Wegman, 1979] Carter, J. L., & Wegman, M. N. (1979). Universal classes of hash functions. *Journal of Computer and System Sciences*, 18, 143–154.
- [Carvalho & Cohen, 2004] Carvalho, V. R., & Cohen, W. W. (2004). Learning to extract signature and reply lines from email. *Proceedings of the 1st Conference on Email and Anti-Spam (CEAS)*. Mountain View, CA.
- [Cygwin, 2007] Cygwin (2007). The Cygwin project mailing list archives. Retrieved from <http://cygwin.com/ml/cygwin/>.
- [EDRM Group, 2008] EDRM Group (2008). The electronic discovery reference model (EDRM) projects. <http://edrm.net/>.
- [EMC, 2008] EMC (2008). EMC EmailXtender family. <http://www.emc.com/products/family/email-xtender-family.htm>.
- [Fulbright, 2006] Fulbright (2006). *Third annual litigation trends survey findings*. Fulbright & Jaworski L.L.P.
- [Gentoo, 2007] Gentoo (2007). General Gentoo user support and discussion mailing list. Retrieved from <http://archives.gentoo.org/gentoo-user/>.
- [Huang et al., 2004] Huang, Y., Govindaraju, D., Mitchell, T., de Carvalho, V. R., & Cohen, W. (2004). Inferring ongoing activities of workstation users by clustering email. *Proceedings of the 1st Conference on Email and Anti-Spam (CEAS)*. Mountain View, CA.
- [Joachims, 1999] Joachims, T. (1999). Transductive inference for text classification using support vector machines. *Proceedings of 16th International Conference on Machine Learning*. Bled, Slovenia.
- [Klimt & Yang, 2004] Klimt, B., & Yang, Y. (2004). Introducing the Enron corpus. *Proceedings of the 1st Conference on Email and Anti-Spam (CEAS)*. Mountain View, CA. Data available at <http://www.cs.cmu.edu/~enron/>.
- [MySQL, 2007] MySQL (2007). MySQL general discussion. Retrieved from <http://lists.mysql.com/mysql/>.
- [qmail, 2007] qmail (2007). qmail mailing list. Retrieved from <http://www.ornl.gov/lists/mailling-lists/qmail/>.
- [Rabin, 1981] Rabin, M. O. (1981). *Fingerprinting by random polynomials* (Technical Report TR-15-81). Center for Research in Computing Technology, Harvard University.
- [Sahami et al., 1998] Sahami, M., Dumais, S., Heckerman, D., & Horvitz, E. (1998). A Bayesian approach to filtering junk e-mail. *Proceedings of the AAAI Workshop on Learning for Text Categorization*. Madison, Wisconsin.
- [Schleimer et al., 2003] Schleimer, S., Wilkerson, D. S., & Aiken, A. (2003). Winnowing: Local algorithms for document fingerprinting. *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD)*. San Diego, CA.
- [Supreme Court, 2006] Supreme Court (2006). *Amendments to the Federal rules of civil procedure*.
- [Surendran et al., 2005] Surendran, A. C., Platt, J. C., & Renshaw, E. (2005). Automatic discovery of personal topics to organize email. *Proceedings of the 2nd Conference on Email and Anti-Spam (CEAS)*. Palo Alto, CA.
- [Symantec, 2008] Symantec (2008). Symantec Enterprise Vault. <http://www.symantec.com/ev/>.
- [Yeh & Harnly, 2006] Yeh, J.-Y., & Harnly, A. (2006). Email thread reassembly using similarity matching. *Proceedings of the 3rd Conference on Email and Anti-Spam (CEAS)*. Mountain View, CA.
- [ZANTAZ, 2008] ZANTAZ (2008). ZANTAZ Enterprise Archive Solution. <http://www.zantaz.com/products/eas.php>.