

Compressed Bloom Filters

[Extended Abstract]

Michael Mitzenmacher^{*}
Harvard University
33 Oxford St.
Cambridge, MA 02138
michaelm@eecs.harvard.edu

ABSTRACT

A Bloom filter is a simple space-efficient randomized data structure for representing a set in order to support membership queries. Although Bloom filters allow false positives, for many applications the space savings outweigh this drawback when the probability of an error is sufficiently low. We introduce *compressed Bloom filters*, which improve performance when the Bloom filter is passed as a message, and its transmission size is a limiting factor. For example, Bloom filters have been suggested as a means for sharing Web cache information. In this setting, proxies do not share the exact contents of their caches, but instead periodically broadcast Bloom filters representing their cache. By using compressed Bloom filters, proxies can reduce the number of bits broadcast, the false positive rate, and/or the amount of computation per lookup. The cost is the processing time for compression and decompression, which can use simple arithmetic coding, and more memory use at the proxies, which utilize the larger uncompressed form of the Bloom filter.

1. INTRODUCTION

Bloom filters are an excellent data structure for succinctly representing a set in order to support membership queries [2]. We describe them in detail in Section 2.1; here, we simply note that the data structure is randomized (in that it uses randomly selected hash functions), and hence has some probability of giving a *false positive*; that is, it may incorrectly return that an element is in a set when it is not. For many applications, the probability of a false positive can be made sufficiently small and the space savings are significant enough that Bloom filters are useful.

In fact, Bloom filters have a great deal of potential for distributed protocols where systems need to share information

^{*}Supported in part by NSF CAREER Grant CCR-9983832, an Alfred P. Sloan Research Fellowship, and a Rosen Fund award.

about what data they have available. For example, Fan et al. describe how Bloom filters can be used for Web cache sharing [6]. To reduce message traffic, proxies do not transfer URL lists corresponding to the exact contents of their caches, but instead periodically broadcast Bloom filters that represent the contents of their cache. If a proxy wishes to determine if another proxy has a page in its cache, it checks the appropriate Bloom filter. In the case of a false positive, a proxy may request a page from another proxy, only to find that proxy does not actually have that page cached. In that case, some additional delay has been incurred. The small chance of a false positive introduced by using a Bloom filter is greatly outweighed by the significant reduction in network traffic achieved by using the succinct Bloom filter instead of sending the full list of cache contents. Bloom filters have also been suggested for other distributed protocols, e.g. [5, 8].

Our paper is based on the following insight: in this situation, the Bloom filter plays a dual role. It is both a data structure being used at the proxies, and a message being passed between them. When we use the Bloom filter as a data structure, we may tune its parameters for optimal performance as a data structure; that is, we minimize the probability of a false positive for a given memory size and number of items. Indeed, this is the approach taken in the analysis of [6]. If this data structure is also being passed around as a message, however, then we introduce another performance measure we may wish to optimize for: transmission size. Transmission size may be of greater importance when the amount of network traffic is a concern but there is memory available at the endpoint machines. This is especially true in distributed systems where information must be transmitted repeatedly from one endpoint machine to many others. For example, in the Web cache sharing system described above, the required memory at each proxy is linear in the number of proxies, while the total message traffic rate is quadratic in the number of proxies, assuming point-to-point communication is used. Moreover, the amount of memory required at the endpoint machines is fixed for the life of the system, where the traffic is additive over the life of the system.

Transmission size can be affected by using compression. In this paper, we show how compressing a Bloom filter can lead to improved performance. By using compressed Bloom filters, protocols reduce the number of bits broadcast, the false positive rate, and/or the amount of computation per lookup.

The tradeoff costs are the increased processing requirement for compression and decompression and larger memory requirements at the endpoint machines, who may use a larger original uncompressed form of the Bloom filter in order to achieve improved transmission size.

We start by defining the problem as an optimization problem, which we solve using some simplifying assumptions. We then consider practical issues, including effective compression schemes and actual performance. We recommend arithmetic coding [9], a simple compression scheme well-suited to this situation with fast implementations.

Our work underscores an important general principle for distributed algorithms: when using a data structure as a message, one should consider the parameters of the data structure with both of these roles in mind. If transmission size is important, tuning the parameters so that compression can be used effectively may yield dividends.

2. COMPRESSED BLOOM FILTERS: THEORY

2.1 Bloom filters

We begin by introducing Bloom filters, following the framework and analysis of [6].

A Bloom filter for representing a set $S = \{s_1, s_2, \dots, s_n\}$ of n elements is described by an array of m bits, initially all set to 0. A Bloom filter uses k independent hash functions h_1, \dots, h_k with range $\{0, \dots, m-1\}$. We make the natural assumption that these hash functions map each item in the universe to a random number uniform over the range $\{0, \dots, m-1\}$ for mathematical convenience. For each element $s \in S$, the bits $h_i(s)$ are set to 1 for $1 \leq i \leq k$. A location can be set to 1 multiple times, but only the first change has an effect. To check if an item x is in S , we check whether all $h_i(x)$ are set to 1. If not, then clearly x is not a member of S . If all $h_i(x)$ are set to 1, we assume that x is in S , although we are wrong with some probability. Hence a Bloom filter may yield a *false positive*, where it suggests that an element x is in S even though it is not. For many applications, this is acceptable as long as the probability of a false positive is sufficiently small.

The probability of a false positive for an element not in the set, or the *false positive rate*, can be calculated in a straightforward fashion, given our assumption that hash functions are perfectly random. After all the elements of S are hashed into the Bloom filter, the probability that a specific bit is still 0 is

$$\left(1 - \frac{1}{m}\right)^{kn} \approx e^{-kn/m}.$$

We let $p = e^{-kn/m}$. The probability of a false positive is then

$$\left(1 - \left(1 - \frac{1}{m}\right)^{kn}\right)^k \approx \left(1 - e^{-kn/m}\right)^k = (1-p)^k.$$

We let $f = \left(1 - e^{-kn/m}\right)^k = (1-p)^k$. Note that we use the asymptotic approximations p and f to represent respec-

tively the probability a bit in the Bloom filter is 0 and the probability of a false positive from now on for convenience.

Although it is clear from the above discussion, it is worth noting that there are *three* fundamental performance metrics for Bloom filters that can be traded off: computation time (corresponding to the number of hash functions k), size (corresponding to the array size m), and the probability of error (corresponding to the false positive rate f).

Suppose we are given m and n and we wish to optimize the number of hash functions k to minimize the false positive rate f . There are two competing forces: using more hash functions gives us more chances to find a 0 bit for an element that is not a member of S , but using fewer hash functions increases the fraction of 0 bits in the array. The optimal number of hash functions that minimizes f as a function of k is easily found taking the derivative. More conveniently, note that f equals $\exp(k \ln(1 - e^{-kn/m}))$. Let $g = k \ln(1 - e^{-kn/m})$. Minimizing the false positive rate f is equivalent to minimizing g with respect to k . We find

$$\frac{dg}{dk} = \ln\left(1 - e^{-\frac{kn}{m}}\right) + \frac{kn}{m} \frac{e^{-\frac{kn}{m}}}{1 - e^{-\frac{kn}{m}}}.$$

It is easy to check that the derivative is 0 when $k = (\ln 2) \cdot (m/n)$; further efforts reveal that this is a global minimum. In this case the false positive rate f is $(1/2)^k = (0.6185)^{m/n}$. In practice, of course, k must be an integer, and smaller k might be preferred since they reduce the amount of computation necessary.

Note that Bloom filters are highly effective even if $m = cn$ for a small constant c , such as $c = 8$. An alternative approach if more bits are available is to simply hash each item into $\Theta(\log n)$ bits and send a list of hash values. Bloom filters can allow significantly fewer bits to be sent with better false positive rates.

2.2 Compressed Bloom filters

Our optimization above of the number of hash functions k is based on the assumption that we wish to minimize the failure of a false positive as a function of the array size m and the number of objects n . This is the correct optimization if we consider the Bloom filter as an object residing in memory. In the Web cache application, however, the Bloom filter is not just an object that resides in memory, but an object that must be transferred between proxies. This fact suggests that we may not want to optimize the number of hash functions for m and n , but instead optimize the number of hash functions for the size of the data that needs to be sent, or the *transmission size*. The transmission size, however, need not be m ; we might be able to compress the bit array. Therefore we choose our parameters to minimize the failure probability after using compression.

Let us consider the standard uncompressed Bloom filter, which is optimized for $k = (\ln 2) \cdot (m/n)$. Can we gain anything by compressing the resulting bit array? If we choose the above value for k , under our assumption of good random hash functions, the bit array appears to be a random string of m 0's and 1's, with each entry being 0 or 1 inde-

pendently with probability $1/2$.¹ Hence compression does not gain anything for this choice of k .

Suppose, however, we instead choose k so that each of the entries in the m bit array is 1 with probability $1/3$. Then we can take advantage of this fact to compress the m bit array and reduce the transmission size. After transmission, the bit array is decompressed for actual use. Note that the uncompressed Bloom filter size is still m bits.² While this choice of k is not optimal for the uncompressed size m , if our goal is to optimize for the transmission size, using compression may yield a better result. The question is whether this compression gains us anything, or if we would have been better off simply using a smaller number of bits in our array and optimizing for that size.

To contrast with the original Bloom filter discussion, we note that for compressed Bloom filters there are now *four* fundamental performance metrics for Bloom filters that can be traded off. Besides computation time (corresponding to the number of hash functions k) and the probability of error (corresponding to the false positive rate f), there are two other metrics: the *uncompressed filter size* that the Bloom filter has in the proxy memory, which we continue to denote by the number of array bits m ; and the *transmission size* corresponding to its size after compression, which we denote by z . Our starting point is the consideration that in many situations the transmission size may be more important than the uncompressed filter size.

We may establish the problem as an optimization problem as follows. Let $z \geq n$ be the desired compressed size. Recall that each bit in the bit array is 0 with probability p ; we treat the bits as independent. Also, as a mathematically convenient approximation, we assume that we have an optimal compressor. That is, we assume that our m bit filter can be compressed down to only $mH(p)$ bits, where $H(p) = -p \log_2 p - (1-p) \log_2 (1-p)$ is the entropy function. Our compressor therefore uses the optimal $H(p)$ bits on average for each bit in the original string. We consider the practical implications more carefully subsequently. We note just that near-optimal compressors exist; arithmetic coding, for example, requires on average less than $H(p) + \epsilon$ bits per character for any $\epsilon > 0$ given suitably large strings.

Our optimization problem is as follows: given n and z , choose m and k to minimize f subject to $mH(p) \leq z$. One possibility is to choose $m = z$ and $k = (\ln 2) \cdot (m/n)$ so that $p = 1/2$; this is the original optimized Bloom filter. Hence we can guarantee that $f \leq (0.6185)^{z/n}$.

¹Technically, this is not precisely true, since bits are not completely independent: the fact that one bit was set to 1 affects the probability of other bits being set to 1. Asymptotically (and in practice) this effect is negligible; see, for example, [1]. Henceforth we make the simplifying assumption of independence of the bit values in the array.

²A compression scheme that also provided random access might allow us to compute on the compressed Bloom filter; however, achieving random access, efficiency, and good compression simultaneously is generally not possible. While this is a possible point for future study, we assume here that computation on the Bloom filter is done after decompression.

We can, however, do better. Indeed, in theory this choice of k is the *worst* choice possible once we allow compression. To see this, let us parameterize so that $k = \alpha m/n$. Then $p = e^{-\alpha}$, and we must choose m and α to minimize $f = (1 - e^{-\alpha})^{\alpha m/n}$ subject to $m = z/H(e^{-\alpha})$ (we may without loss of generality choose m as large as possible). Equivalently, we have $f = (1 - e^{-\alpha})^{\alpha z/(nH(e^{-\alpha}))}$. Since z and n are fixed with $z \geq n$, we wish to minimize

$$\beta = f^{n/z} = (1 - e^{-\alpha})^{\alpha/H(e^{-\alpha})}.$$

Let $\alpha = -\ln x$. Then

$$\begin{aligned} \beta &= (1 - x)^{-\ln x/H(x)} \\ &= \exp\left(\frac{-\ln(x) \cdot \ln(1 - x)}{(-\log_2 e)(x \ln x + (1 - x) \ln(1 - x))}\right). \end{aligned}$$

The value of β is maximized when the exponent is maximized, or equivalently when the term

$$\gamma = \left(\frac{x}{\ln(1 - x)} + \frac{1 - x}{\ln x}\right)$$

is minimized. Note that

$$\frac{d\gamma}{dx} = \frac{1}{\ln(1 - x)} - \frac{1}{\ln x} + \frac{x}{(1 - x) \ln^2(1 - x)} - \frac{1 - x}{x \ln^2(x)}.$$

The value of $\frac{d\gamma}{dx}$ is clearly 0 when $x = 1/2$, and using symmetry it is easy to check that $\frac{d\gamma}{dx}$ is negative for $x < 1/2$ and positive for $x > 1/2$. Hence the maximum probability of a false positive using a compressed Bloom filter occurs when $x = 1/2$, corresponding to $\alpha = \ln 2$ or equivalently $k = (\ln 2) \cdot (m/n)$.

We emphasize the point again: the number of hash functions that minimizes the false positive rate without compression in fact maximizes the false positive rate with compression. Said in another way, in our idealized setting using compression always decreases the false positive rate.

The argument above also shows that γ is minimized and hence β and f are minimized in one of the limiting situations as x goes to 0 or 1, corresponding to α going to infinity or 0. In each case, using for example the expansion $\ln(1 - x) \approx -x - x^2/2 - x^3/3 - \dots$, we find that γ goes to -1 . Hence β goes to $1/2$ in both limiting cases, and we can in theory achieve a false positive rate arbitrarily close to $(0.5)^{z/n}$ by letting the number of hash functions go to 0 or infinity.

In practice we are significantly more constrained. For example, we must use at least one hash function, and the number of hash functions must be an integer. Note, however, that the theory shows we may achieve improved performance by taking $k < \ln 2 \cdot (m/n)$ for the compressed Bloom filter. This has the additional benefit that a compressed Bloom filter uses fewer hash functions and hence requires less computation per lookup. Further practical considerations are discussed in the next section.

Note that the optimization framework developed above is not the only one possible. For example, one could instead fix the desired false positive rate f and optimize for the final compressed size z . To compare in this situation, note that

in the limit as the number of hash functions goes to zero the compressed Bloom filter has a false positive rate tending to $(0.5)^{z/n}$ while the standard Bloom filter has a false positive rate tending to $(0.5)^{(m \ln 2)/n}$. Hence the best possible compressed Bloom filter achieving the same false positive rate as the standard Bloom filter would have $z = m \ln 2$, a savings in size of roughly 30%. Again, this is significantly better than what can be realized in practice.

The primary point of this theoretical analysis is to demonstrate that compression is a viable means of improving performance, in terms of reducing the false positive rate for a desired compressed size, or for reduced the transmission size for a fixed false positive rate. Indeed, because the compressed Bloom filter allows us another performance metric, it provides more flexibility than the standard original Bloom filter. An additional benefit is the compressed Bloom filters uses a smaller number of hash functions, so that lookups are more efficient. Based on this theory, we now consider implementation details and specific examples.

3. COMPRESSED BLOOM FILTERS: PRACTICE

Our theoretical analysis avoided several issues that are important for a real implementation:

- Restrictions on m : While the size z of the compressed Bloom filter may be of primary importance, limitations on the size m of the uncompressed Bloom filter also constrain the possibilities. For example, while theoretically we can do well using one hash function and compressing, achieving a false positive rate of ϵ with one hash function requires $m \approx n/\epsilon$, which may be too large for real applications.

Also, it may be desirable to have m be a power of two for various computations. We do not restrict ourselves to powers of two here.

- Compression overhead: Compression schemes do not achieve optimal performance; all compression schemes have some associated overhead. Hence the gain from the compressed Bloom filter must overcome the associated overhead costs.
- Compression variability: Of perhaps greater practical importance is that if there is an absolute maximum packet size, we generally want that the compressed array does not exceed the packet size. Otherwise, we may have to send multiple packets, leading to increased network traffic. (In some situations, we may want to bound the number of packets used to send the compressed Bloom filter; the idea is the same.) Compression performance varies depending on the input; moreover, if the number of elements n in the set S cannot be exactly determined in advance, a misprediction of n could yield insufficient compression.
- Hashing performance: Depending on the data and the hash functions chosen, real hash functions may behave differently from the analysis above.

The issue of achieving good hashing performance on arbitrary data sets is outside the scope of this paper, and

we do not consider it further except to raise the following points. First, in practice we suspect that using standard universal families of hash functions [4, 10] or MD5 (used in [6]) will be suitable. Second, in situations where hashing performance is not sufficiently random, we expect that compressed Bloom filters will still generally outperform the uncompressed Bloom filter. The point is that if the false positive rate of a compressed Bloom filter is increased because of weak hash functions, we would expect the false positive rate of the uncompressed Bloom filter to increase as well; moreover, since compressed Bloom filters use fewer hash functions, we expect the effect will be worse for the uncompressed filter. For compressed Bloom filters, however, there is the additional problem that weak hash functions may yield bit arrays that do not compress as much as expected. The choice of parameters may therefore have to be tuned for the particular data type.

For compression issues, arithmetic coding provides a flexible compression mechanism for achieving near-optimal performance with low variability. Loosely speaking, for a random m bit string where the bit values are independent and each bit is 0 with probability p and 1 with probability $1-p$, arithmetic coding compresses the string to near $mH(p)$ bits with high probability, with the deviation from the average having a Chernoff-like bound. For more information on arithmetic coding, we refer the reader to [9, 11]. For more precise statements and details regarding the low variability of arithmetic coding, we refer the reader to [7]. More details will appear in the final version.

Given this compression scheme, we suggest the following approach. Choose a maximum desired uncompressed size m . Then design a compressed Bloom filter using the above theory using a slightly smaller compressed size than desired; for example, if the goal is that the compressed size be z , design the structure so that the compressed size is $0.99z$. This provides room for some variability in compression; the amount of room necessary depends on m . A similar effect may be achieved by slightly overestimating n . If our uncompressed filter is more than half full of zeroes, then if we have fewer than expected elements in the set, our filter will tend have even more zeroes than expected, and hence will compress better. With this design, the compressed filter should be within the desired size range with high probability.

To deal with cases that still do not compress adequately, we suggest using multiple filter types. Each filter type t is associated with an array of size m , a set of hash functions, and a decompression scheme. These types are agreed on ahead of time. A few bits in the header can be used to represent the filter type. If one of the filter types is the standard Bloom filter (no compression) then the set can always be sent appropriately using at least one of the types. In most cases two types—compressed and uncompressed—would be sufficient.

3.1 Examples

We test the theoretical framework above by examining a few specific examples of the performance improvements possible using compressed Bloom filters. We consider cases where eight and sixteen bits are used in the compressed Bloom filter for each element; this corresponds to configurations examined in [6].

Suppose we wish to use at most eight bits per set element in our transmission with a Bloom filter; that is, $z/n = m/n = 8$. Then using the optimal number of hash functions $k = 6$ yields a false positive rate of 0.0216. For $k = 5$, the false positive probability is only 0.0217, so this might be preferable in practice. If we are willing to allow 14 array bits for the uncompressed Bloom filter per set element, then we can reduce the false positive rate by almost 20% to 0.0177 and reduce the number of hash functions to two while keeping the (theoretical) transmitted bits per element z/n below 8, as shown in Table 1.

It is also interesting to compare the standard Bloom and the compressed Bloom filter pictorially in this case where $z/n = 8$. In Figure 1 we show the false positive rate as a function of the number of hash functions k based on the theoretical analysis of Sections 2.1 and 2.2, where we allow k to behave as a continuous variable. Note that as the theory predicts the optimized uncompressed filter actually yields the largest false positive rate once we introduce compression.

We tested the compressed Bloom filter via simulation. We repeated the following experiment 100,000 times. A Bloom filter for $n = 10,000$ elements and $m = 140,000$ bits was created, with each element being hashed to two positions chosen independently and uniformly at random in the bit array. The resulting array was then compressed using a publicly available arithmetic coding compressor based on the work of Moffat, Neal, and Witten [3, 9].³ Using $z = mH(p)$ suggests that the compressed size should be near 9,904 bytes; to meet the bound of 8 bits per element requires the compressed size not exceed 10,000 bytes. Over the 100,000 trials we found the average compressed array size to be 9,920 bytes, including all overhead; the standard deviation was 11.375 bytes; and the maximum compressed array size was only 9,971 bytes, giving us several bytes of room to spare. For larger m and n , we would expect even greater concentration of the compressed size around its mean; for smaller m and n , the variance would be a larger fraction of the compressed size. We believe example provides good insight into what is achievable in real situations.

Theoretically we can do even better by using just one hash function, although this greatly increases the number of array bits per element, as seen in Table 1.

Similarly, considering the specific case of a Bloom filter where $z/n = m/n = 16$, we would use eleven hash functions to achieve an optimal false positive rate of 0.000459. As eleven hash functions seems somewhat large, we note that we could reduce the number of hash functions used without applying compression, but using only six hash functions more than doubles f to 0.000935. Table 2 summarizes the improvements available using compressed Bloom filters. If we allow 28 array bits per element, our false positive rate falls about 30% while using only four hash functions. If we allow 48 array bits per element, our false positive rate falls over 50% using only three hash functions. We simulated

³We note that this is an adaptive compressor, which bases its prediction of the next bit based on the bits seen thus far. Technically is slightly suboptimal for our purposes, since we generally know the probability distribution of the bits ahead of time. In practice the difference is quite small.

the case with $n = 10,000$ elements, $m = 480,000$ bits, and $k = 3$ hash functions using 100,000 trials. The theoretical considerations above suggest the compressed size will be 19,787 bytes. Over our simulation trials, the average compressed array size was 19,805 bytes, including all overhead; the standard deviation was 14.386 bytes; and the maximum compressed array size was only 19,865 bytes, well below the 20,000 bytes available.

We have also tested the case where $z/n = m/n = 4$ against using $m/n = 7$, or seven array bits per table. We expect this case may prove less useful in practical situations because the false positive rate is so high. In this case using the standard Bloom filter with the optimal three hash functions yields a false positive rate of 0.147; using $m/n = 7$ and one hash function gives a false positive rate of 0.133. Again, we performed 100,000 random experiments with $n = 10,000$. The largest compressed filter required 4,998 bytes, just shy of the 5,000 byte limit.

As previously mentioned, we may also consider the optimization problem in another light: we may try to maintain the same false positive ratio while minimizing the transmission size. In Tables 4 and 5 we offer examples based on this scenario. Our results yield transmission size decreases in the range of roughly 5-15% for systems of reasonable size. Here again our simulations bear out our theoretical analysis. For example, using $n = 10,000$ elements, $m = 126,000$ bits, and $k = 2$ hash functions over 100,000 trials, we find the average compressed filter required 9,493 bytes, closely matching the theoretical prediction. The largest filter over the 100,000 trials required 9,539 bytes.

4. EXTENSIONS

In the Web cache sharing setting, the proxies periodically broadcast updates to their cache contents. As described in [6], these updates can either be new Bloom filters or representations of the changes between the updated filter and the old filter. The difference, or *delta*, between the updated and old filter can be represented by the exclusive-or of the corresponding bit arrays of size m , which can then be compressed using arithmetic coding as above. For example, one may decide that updates should be broadcast whenever 5% of the underlying array bits have changed; in this case, the compressed size of the delta would be roughly $mH(0.05)$. Hence one may wish to optimize the array size for a target size of the compressed delta and allow the one-time cost of longer initial messages to establish a base Bloom filter at the beginning. Compressing the delta can clearly yield substantial gains in this setting.

Also in [6], the authors also describe an extension to a Bloom filter, where instead of using a bit array the Bloom filter array uses a small number of bits per entry to keep counts. The j th entry is incremented for each hash function h_i and each element x represented by the filter such that $h_i(x) = j$. The *counting Bloom filter* is useful when items can be deleted from the filter; when an item x is deleted, one can decrement the value at location $h_i(x)$ in the array for each of the k hash functions, i.e. for $1 \leq i \leq k$. We emphasize that these counting Bloom filters are not passed as messages in [6]; they are only used locally.

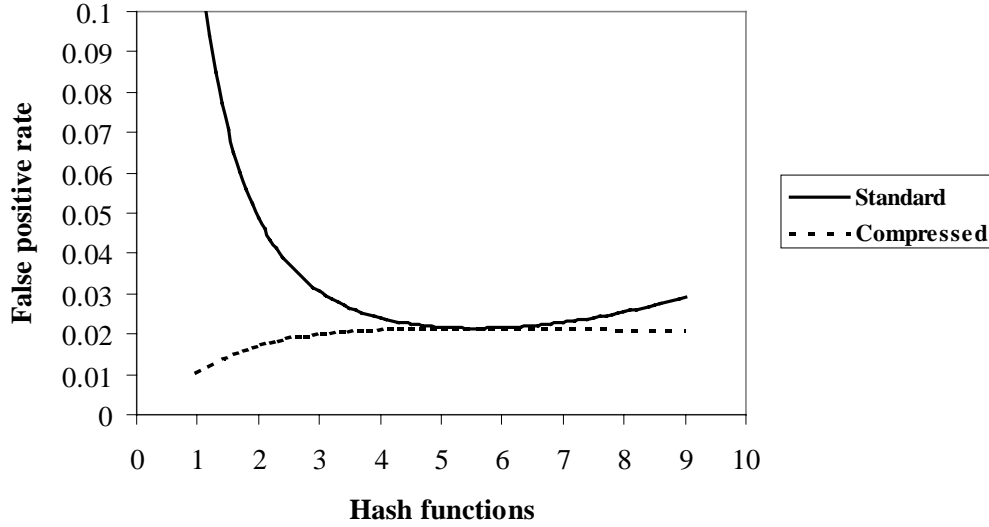


Figure 1: The false positive rate as a function of the number of hash functions for compressed and standard Bloom filters using 8 bits per element.

Array bits per element	m/n	8	14	92
Transmission bits per element	z/n	8	7.923	7.923
Hash functions	k	6	2	1
False positive rate	f	0.0216	0.0177	0.0108

Table 1: At most eight bits per item (compressed).

Array bits per element	m/n	16	28	48
Transmission bits per element	z/n	16	15.846	15.829
Hash functions	k	11	4	3
False positive rate	f	0.000459	0.000314	0.000222

Table 2: At most sixteen bits per item (compressed).

Array bits per element	m/n	4	7
Transmission bits per element	z/n	4	3.962
Hash functions	k	3	1
False positive rate	f	0.147	0.133

Table 3: At most four bits per item (compressed).

Array bits per element	m/n	8	12.6	46
Transmission bits per element	z/n	8	7.582	6.891
Hash functions	k	6	2	1
False positive rate	f	0.0216	0.0216	0.0215

Table 4: Maintaining a false positive rate around 0.02.

Array bits per element	m/n	16	37.5	93
Transmission bits per element	z/n	16	14.666	13.815
Hash functions	k	11	3	2
False positive rate	f	0.000459	0.000454	0.000453

Table 5: Maintaining a false positive rate around 0.00045.

We note in passing that if one wanted to pass counting Bloom filters as messages, compression would yield substantial gains. The entropy per array entry would be much smaller than the number of bits used per entry, since large counts would be extremely unlikely. Our optimization approach for finding appropriate parameters can be extended to this situation, and arithmetic coding remains highly effective.

5. CONCLUSIONS

We have shown that using compression can improve Bloom filter performance, in the sense that we can achieve a smaller false positive rate as a function of compressed size over a Bloom filter that does not use compression. This improvement is useful in situations where the Bloom filter is being used as a message as well as a data structure in a distributed protocol. In this case, it may be significantly more important to optimize the transmission size rather than the data structure size. An interesting feature of our work is that the use of compression affects how one should tune the parameters of the data structure. It would be interesting to find other useful examples of data structures that can be tuned effectively in a different manner when being compressed.

While we have not yet implemented compressed Bloom filters in the context of a full working system for an application such as distributed Web caching, we hope to do so in the near future. We expect that significant performance improvement will occur even after the minor costs of compression and decompression time are factored in.

Acknowledgments

The author would like to thank Andrei Broder for introducing him to Bloom filters and for helpful discussions.

6. REFERENCES

- [1] M. Adler, S. Chakrabarti, M. Mitzenmacher, L. Rasmussen. Parallel randomized load balancing. *Random Structures and Algorithms*, 13(2), pages 159-188, 1998.
- [2] B. Bloom. Space/time tradeoffs in in hash coding with allowable errors. *Communications of the ACM*, 13(7):422-426, July 1970.
- [3] J. Carpinelli, W. Salomonsen, A. Moffat, R. Neal, and I. H. Witten. Source code for arithmetic coding, Version 1. Available at http://www.cs.mu.oz.au/~alistair/arith_coder/. March 1995.
- [4] J. L. Carter and M. N. Wegman. Universal classes of hash functions. *Journal of Computer and System Sciences*, 18, pages 143-154, 1979.
- [5] S. Czerwinski, B. Zhao, T. Hodes, A. Joseph, and R. Katz. An architecture for a secure service discovery service. In *Proceedings of the Fifth Annual International Conference on Mobile Computing and Networks (MobiCOM '99)*, 1999.
- [6] L. Fan, P. Cao, J. Almeida, and A. Broder. Summary cache: A scalable wide-area Web cache sharing protocol. In *Proceeding of SIGCOMM '98*, 1998.
- [7] P. G. Howard and J. Vitter. Analysis of arithmetic coding for data compression. *Information Processing and Management*, vol 28. no. 6, pages 749-763, 1992.
- [8] J. Kubiawicz, D. Bindel, Y. Chen, S. Czerwinski, P. Eaton, D. Geels, R. Gummadi, S. Rhea, H. Weatherspoon, W. Weimer, C. Wells, and B. Zhao. OceanStore: An architecture for global-scale persistent storage. In *Proceedings of ASPLOS 2000*, 2000.
- [9] A. Moffat, R. Neal, and I. H. Witten. Arithmetic coding revisited. *ACM Transactions on Information Systems*, 16(3):256-294, July 1998.
- [10] M. V. Ramakrishna. Practical performance of Bloom filters and parallel free-text searching. *Communications of the ACM*, 32(10):1237-1239, October 1989.
- [11] I. H. Witten, A. Moffat, and T. Bell. **Managing Gigabytes (2nd Edition)**. Morgan Kaufmann, San Francisco, 1999.

Extended version available as Technical Report 1361, Computer Sciences Department, Univ. of Wisconsin-Madison, February 1999.