



On the search path length of random binary skip graphs

Philippe Duchon, Larchevêque Hubert

► To cite this version:

Philippe Duchon, Larchevêque Hubert. On the search path length of random binary skip graphs. ANALCO 2010, Golin, Mordecai and Sedgewick, Robert, Jan 2010, Austin, United States. pp.1-8. inria-00547935

HAL Id: inria-00547935

<https://inria.hal.science/inria-00547935>

Submitted on 17 Dec 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On the search path length of random binary skip graphs

Philippe Duchon, Hubert Larchevêque

December 17, 2009

Abstract

In this paper we consider the skip graph data structure, a load balancing alternative to skip lists, designed to perform better in a distributed environment.

We extend previous results of Devroye on skip lists, and prove that the maximum length of a search path in a random binary skip graph of size n is of order $\log n$ with high probability.

1 Introduction

A skip list [9, 8] is an ordered data structure based on a succession of linked lists with geometrically decreasing numbers of items. Skip lists come into deterministic [8] and randomized [9] flavours. The deterministic versions have guaranteed properties whereas randomized skip lists only offer high probability performance. Devroye studied precisely the height of a random skip list in [6]. This height H_n is the maximum length of a search path for any key from the top of the skip list. Devroye proved that this height H_n is of order $\log n$ (see below for a more precise statement).

The skip graph, introduced by Aspnes and Shah in [2, 3], is a variant of the skip list, designed to perform better in a distributed environment. In a skip graph, the whole data structure can be distributed among a large number of nodes, and the structure provides good load balancing and fault tolerance properties. During past years interesting variants of skip graphs have been studied, like skip nets [?], skip webs [1] or rainbow skip graphs [7], and some theoretical results are known about its structure [4, 3].

In a skip graph of size n (see below for more precise definitions), any key can be searched for in n different (but not independent) skip lists, and this results in as many search paths. We define the *height* H'_n of a random skip graph as the maximum length of all search paths in the skip graph. As a first step towards a more detailed analysis of random skip graphs, the goal of this paper is to give an upper bound on the length of the maximal search path in a skip graph, by extending to skip graphs previous results of Devroye on skip lists. Our main result is the following:

THEOREM 1.1. *For any positive ϵ , the height H'_n of a skip graph built on a list of n elements is such that $\mathbb{P}(H'_n \leq (2c + \epsilon) \log_2 n) = 1 - o(1)$, where c is the unique solution of the equation*

$$x - 1 - x \log_2(x) + (x - 1) \log_2(x - 1) = 0.$$

It can easily be argued that the height H'_n of a skip graph $O(\log^2 n)$, since the number of levels of a skip graph is of order $\log n$ w.h.p, and the number of elements at a given level on a given search path can be of order $\log n$. Our result ensures that the height is in fact of order $\log n$, ensuring good expected execution time for some algorithms based on skip graphs, such as the prefix sum computation of an ordered list of elements.

For example, consider a set of n processors, each having a given attribute. Let this set be ordered (following the processor's identifiers if necessary). In the *prefix sum computation problem*, each processor

must compute the sum of attributes of all preceding processors. This problem often arises in parallel computing since it is a useful block for many algorithms in this area (for a more detailed survey on applications using prefix sums computation, see [?]).

This problem has also been specifically studied by the authors and coauthors in [5], where the prefix sums are to be computed in a skip graph. The number of steps needed for the algorithm presented in [5] is the height of the skip graph, H'_n , thus is of probabilistic order $\log n$ instead of $\log^2 n$ as originally claimed.

The present paper is organized as follows. In Section 2 we describe the models we use for skip lists and skip graphs in the rest of the paper. In Section 3 we extend an upper bound result of Devroye for random skip lists with parameter $p = 1/2$. In Section 4 we use this result to get a high probability upper bound on the height of a skip graph, and finally, in Section 5 we complete our theoretical results by a brief empirical study.

2 Models and notations

2.1 Skip lists A skip list is defined by a decreasing sequence $S_0 \supseteq S_1 \supseteq \dots \supseteq S_L$ of finite sets of keys from a totally ordered universe.

In a skip list, each of the S_i is represented as an ordered (singly or doubly) linked list; each element of S_i also stores a pointer to the corresponding element of S_{i-1} . The maximum level L of a skip list $\mathcal{S}$ will be noted $L(\mathcal{S})$.

In a *random* skip list of parameter p for a keyset $S_0 = \{x_1 \leq x_2 \leq \dots \leq x_n\}$, each key $x \in S_0$ is independently assigned a geometric random variable of parameter $1 - p$, G_x . The key x appears in the G_x first sets $S_0, \dots, S_{G_x-1}$.

Search paths in skip lists The *search path* for a key x in a skip list $\mathcal{S}$, denoted by $P_{\mathcal{S}}(x)$ represents the set of cells one has to examine while searching for it. Its length is denoted by $|P_{\mathcal{S}}(x)|$.

For each level $0 \leq i \leq L(\mathcal{S})$ and each key $x \in S_i$, the cell in S_i for key $x = x_m$ is represented by a node at coordinates (m, i) . We add sentinel nodes $(0, i)$ and $(n + 1, i)$ at each level. Thus, there are pointers from each node (m, i) to $(m, i - 1)$ and (m', i) where m' is the rank in S_0 of $\text{succ}_{S_i}(x_m)$, *i.e.*, $\text{succ}_{S_i}(x_m) = x_{m'}$.

The search path for a key x in a skip list $\mathcal{S}$ starts at $(0, L(\mathcal{S}))$ and ends at the first node of S_0 with an attached key greater than or equal to x . If the i -th node is (m, l) , the next node (m', l') is defined as follows :

- if $\text{succ}_{S_l}(x_m) > x$, $(m', l') = (m, l - 1)$ (unless $l = 0$, in which case the next node is $(m + 1, 0)$ and the path ends);
- otherwise, $l' = l$ and m' is the rank in S_0 of $\text{succ}_{S_l}(x_m)$.

The maximum length of a search path in a skip list $\mathcal{S}$, called its height, is denoted by $H(\mathcal{S})$. Following Devroye, we also denote by E_i the length of the end of the search path for x_i , starting with the first node with positive abscissa.

2.2 Skip graphs In a binary skip graph for a keyset $S_0 = \{x_1 \leq \dots \leq x_n\}$, each key x is assigned a (potentially infinite) uniform random binary word w_x . For each finite word w such that at least one key has w as a prefix of its binary word, consider the linked list S^w of such keys (the list S^ϵ is just the linked list S_0 of all keys in the skip graph). With probability 1, each word w_x has a shortest finite prefix that separates it from all other words w_{x_i} ; the length of this prefix is the *level* of x in the skip graph, noted $L(x)$.

As a result, the sequence of lists $S_0(x) = S^\epsilon, S_1(x), S_{L(x)}(x)$, in which key x appears, make up a skip list (this skip list is not distributed exactly as a random skip list as defined previously; the exact distribution will be described later).

Figure 1 depicts an instance of skip graph built on a set of 6 elements.




Figure 1: A skip graph instance built on a set of 6 elements

Aspnes and Shah proved in [2, 3] that the expected maximum level of a skip graph is $O(\log n)$. They also proved that the expected search time in a skip graph was in $O(\log n)$.

Search paths in skip graphs In contrast to the skip list, search paths in a skip graph $\mathcal{G}$ are defined by a starting key x and a target key y (in a distributed setting where each key corresponds to a single processing node, or where each node is assigned a small number of keys, x would be a key assigned to the node initiating the search). The search path $P_{\mathcal{G}}(x, y)$ is simply the search path for y in the skip list

$\mathcal{S}(x) = (S_0(x), S_1(x), \dots, S_{L(x)}(x))$.

The *height* of a skip graph $\mathcal{G}$, denoted by $H'(\mathcal{G})$, is the maximum, over all pairs of keys (x, y) , of the lengths of search paths in $\mathcal{G}$.

Remark: the careful reader will have noticed that our description of the search path in skip graphs is not coherent with that of [2, 3], where skip graphs are described using doubly linked lists and searching goes backwards when $y < x$, while we use a forward search in a singly linked list. This makes no difference, since searching forward in the “keys smaller than x ” portion of a skip list is equivalent to

searching backwards in the reversed sublist, and both the skip list and its reversed list have the same distribution.

3 Upper bounds on the height of a skip list

To prove Theorem 1.1, we first prove the following one, which is an extension of Devroye's results in [6]:

THEOREM 3.1. *For any $\alpha \geq 1$ and any $c > \alpha c_p$, where c_p is the unique positive solution of the equation*

$$x - 1 + (x - 1) \log_{1/p}(x - 1) - x \log_{1/p} x = 0,$$

the height H_n of a random skip list (with geometric parameter p) of n elements is such that $\mathbb{P}(H_n > c \log_{1/p} n) = o(n^{1-\alpha})$,

Remark: The case $\alpha = 1$ of Theorem 3.1, together with a corresponding lower bound, is proved in [6].

Proof. [Proof of Theorem 3.1]

Following Devroye's proof scheme, we bound the probability of a large height: for any k and l ,

$$\mathbb{P}\{H_n > k\} \leq \mathbb{P}\{L(\mathcal{S}) > l\} + n\mathbb{P}\{E_n > k, L(\mathcal{S}) \leq l\} \quad (3.1)$$

$$\leq \mathbb{P}\{L(\mathcal{S}) > l\} + n\mathbb{P}\left\{\sum_{j=0}^l N_j > k\right\}, \quad (3.2)$$

where N_j is the number of nodes at level j in the search path for x_n .

We first pick a value of l such that the probability $\mathbb{P}(L(\mathcal{S}) > l)$ is $o(n^{1-\alpha})$: since the number of levels of the skip list is the maximum of n independent geometric random variables with parameter p ,

$$\mathbb{P}(L(\mathcal{S}) > l) \leq n \cdot p^l,$$

so that it is sufficient to have $l = \alpha \log_{1/p} n + \omega(1)$. We pick

$$l = \alpha \log_{1/p} n + \sqrt{\log_{1/p} n}.$$

We will pick $k = \lceil \theta \log_{1/p} n \rceil$, θ to be suitably chosen so that

$$\mathbb{P}\{E_n > k, L(\mathcal{S}) \leq l\} = o(n^{-\alpha}).$$

Let N_j be the number of nodes at level j on the search path $P_{\mathcal{S}}(x_n)$. The N_j 's are collectively stochastically smaller than i.i.d. geometric p random variables: $\mathbb{P}\{N'_j = i\} = (1-p)^i p$, $i \geq 1$. Thus, the second term in (3.2) is upper bounded as follows:

$$\mathbb{P}\left\{\sum_{j=0}^l N_j > k\right\} \leq \mathbb{P}\left\{\sum_{j=0}^l N'_j > k\right\}$$

We now use the classical Chernoff-Hoeffding exponential bounding method: for any $0 < t < \ln 1/(1-p)$,

$$\begin{aligned} \mathbb{P}\left\{\sum_{j=0}^l N'_j > k\right\} &\leq \frac{\mathbb{E}\left(e^{t \sum_{j=0}^l N'_j}\right)}{e^{tk}} \\ &= \frac{\left(\mathbb{E}(e^{t N'_0})\right)^{l+1}}{e^{tk}} \\ &= e^{-tk} \left(\frac{pe^t}{1 - (1-p)e^t}\right)^{l+1}. \end{aligned}$$

Provided $\theta > \alpha$, we can take t such that $e^t = \frac{k-l}{k(1-p)}$. This gives us, with $u = \frac{l}{k}$:

$$\begin{aligned} \mathbb{P} \left\{ \sum_{j=0}^l N_j > k \right\} &\leq \left(\frac{k-l}{k(1-p)} \right)^{-k} \left(\frac{p(k-l)}{(1-p)l} \right)^{l+1} \\ &\leq \frac{p(1-u)}{(1-p)u} \left(\left(\frac{p}{u} \right)^u \left(\frac{1-p}{1-u} \right)^{1-u} \right)^k \end{aligned}$$

Note that $u \rightarrow \frac{\alpha}{\theta}$ as $n \rightarrow \infty$. Thus, we have

$$\begin{aligned} n \mathbb{P} \left\{ \sum_{j=0}^l N_j > k \right\} &\leq n \frac{p(1-u)}{(1-p)u} \left(\left(\frac{p}{u} \right)^u \left(\frac{1-p}{1-u} \right)^{1-u} \right)^k \\ &= \frac{p(1-u)}{(1-p)u} n^{1+\theta(-u \log_{1/p} u + u \log_{1/p} p - (1-u) \log_{1/p} (1-u) + (1-u) \log_{1/p} (1-p)) + o(1)} \end{aligned}$$

For the upper bound to be $o(n^{1-\alpha})$, it is sufficient to have (replacing u by its limit),

$$(\theta - \alpha) \log_{1/p} \left(\frac{\theta - \alpha}{1 - p} \right) - \theta \log_{1/p} \theta + \alpha \log_{1/p} \alpha > 0. \quad (3.3)$$

Setting $f_{\alpha,p}(x) = (x - \alpha) \log_{1/p} \left(\frac{x - \alpha}{1 - p} \right) - x \log_{1/p} x + \alpha \log_{1/p} \alpha$, we find that $f_{\alpha,p}(\alpha x) = \alpha f_{1,p}(x)$, and easily check that $f_{1,p}$ has a unique positive zero and is positive for $x > c_p$. Thus, (3.3) is equivalent to $\theta > \alpha c_p$.

4 The relationship between random skip lists and random skip graphs

In order to properly analyze random skip graphs, it is important to understand what the set $\mathcal{S}(x_i)$ of lists in which a given element x_i appears looks like.

These lists make up a skip list that looks a lot like a random skip list of size n (with geometric parameter $1/2$), but the probability distribution is not quite the same (for example in Figure 2 is depicted the skip list corresponding to node 5 in the skip graph of Figure 1). In fact, the list “seen” by element x_i is distributed exactly as a “normal” random skip list (with parameter $p = 1/2$) of size $n - 1$, into which one then inserts an additional element, with final rank i , and with a height one larger than the maximum height of the skip list. To see that this is true, consider a fixed element x_i and its attached binary word w_{x_i} . Conditioned on $w_{x_i} = w$, for each other element x_j , the length (increased by 1) of the longest common prefix between w and w_{x_j} is geometrically distributed with parameter $1/2$, and these lengths are independent; thus, the number of lists in $\mathcal{S}(x_i)$ in which x_j appears is geometrically distributed.

Now consider a pair $(\mathcal{S}, \mathcal{S}')$ of skip lists, where $\mathcal{S}$ is an arbitrary skip list of size $n - 1$ and $\mathcal{S}'$, as described above, is obtained from $\mathcal{S}$ by inserting an n -th element x_i , in position i , with a height one larger than the maximum height h of elements in $\mathcal{S}$.

LEMMA 4.1. *For any key x , the lengths of the search paths for x in $\mathcal{S}$ and $\mathcal{S}'$ satisfy*

$$|P_{\mathcal{S}'}(x)| \leq 2 + |P_{\mathcal{S}}(x)|.$$




Figure 2: The skip list “seen” by node 5 in the skip graph of Figure 1

Proof. Let us consider the possible relative positions of x and x_i . To make notations simpler, we assign the half-integer position $i - 1/2$ to x_i in $\mathcal{S}'$, so that larger keys keep the same position in $\mathcal{S}$ and $\mathcal{S}'$.

If $x < x_i$, then the search paths in $\mathcal{S}$ and $\mathcal{S}'$ are exactly the same, except that the path in $\mathcal{S}'$ starts one level higher, and we have

$$|P_{\mathcal{S}'}(x)| = 1 + |P_{\mathcal{S}}(x)|.$$

If $x = x_i$, the search path for x in $\mathcal{S}'$ is reduced to one node per level, which is certainly at most 1 larger than any search path in $\mathcal{S}$.

If $x > x_i$, consider the first node (i', k) with rank $i' \geq i$ in the search path $P_{\mathcal{S}}(x)$. Then the search path in $\mathcal{S}'$ is simply

$$(0, L(\mathcal{S}) + 1), (i - 1/2, L(\mathcal{S}) + 1), (i - 1/2, L(\mathcal{S})), \dots, (i - 1/2, k)$$

followed by the end of the search path in $\mathcal{S}$, starting at (i', k) . Again, the initial segment of this search path uses only one node per level between levels k and $L(\mathcal{S})$ (the search path in $\mathcal{S}$ uses at least as many), and two nodes at level $L(\mathcal{S}) + 1$, which account for the additional $+2$ term in the lemma.

With this lemma and Theorem 3.1, we are now ready to prove our main result.

Proof. [Proof of Theorem 1.1] Let $\mathcal{G}$ be a random skip graph on n keys. For $H'(\mathcal{G})$ to be larger than k , at least one of the n skip lists $\mathcal{S}(x_i)$ has to have height larger than k . If we set $k \geq c \log_2(n)$ with

$c > 2c_{1/2}$, Lemma 4.1 and Theorem 3.1 ensure that, for each $1 \leq i \leq n$,

$$\mathbb{P}\{H(\mathcal{S}(x_i)) > k\} = o(1/n),$$

(uniformly in i), so that by a union bound, we have

$$\mathbb{P}\{H'(\mathcal{G}) > k\} = o(1).$$

5 Experimental evaluation

In the previous section we obtained a high probability upper bound on the maximum height of a skip-graph by a constant c times $\log n$, where n is the number of elements of the skip-graph. In this section we use simulations to estimate how tight this upper bound is.

Figure 3 shows the empiric distribution functions for the normalized heights of both skip lists and skip graphs of different sizes; for each size the simulations involved 100 random skip lists and skip graphs.




Figure 3: Empiric repartition functions of the normalized height of the structure (skip list or skip graph)

For a random skip list of parameter $p = 1/2$, the constant $c_{1/2} \simeq 4.403$ corresponds to the “top” of the curves in Figure 3 even though the size is still relatively small. By contrast, the top of the curves related to skip graphs seems to correspond to a constant closer to 5.5 than the upper bound of $2c_{1/2} \simeq 8.807$ implied by our theoretical result.

For skip lists, the upper bound is matched by a corresponding lower bound [6], so that the height of a skip list H_n , normalized by $\log_2 n$, tends to $c_{1/2}$ in probability. In the present paper, we have no corresponding lower bound (other than that for skip lists) for skip graphs, and the simulation results seem to indicate that our upper bound is not tight; in any case, the union bound technique we used is unlikely to give much stronger results, and a more careful analysis of the relationship between the different skip lists that appear in a skip graph is required.

6 Conclusion

In this paper we gave an upper bound on the height of a skip graph, proving that with high probability, the height is no larger than twice that of a skip list of the same size.

We observe empirically that this upper bound seems not to be tight. Simulation results seem compatible with the conjecture that $H'_n/\log_2(n)$ converges in probability to a constant, though such a constant is likely to be smaller than our upper bound implies. At the moment, more work is required to obtain such a potential result.

References

- [1] L. Arge, D. Eppstein, and M.T. Goodrich. Skip-webs: efficient distributed data structures for multi-dimensional data sets. *Proceedings of the twenty-fourth annual ACM SIGACT-SIGOPS symposium on Principles of distributed computing*, pages 69–76, 2005.
- [2] J. Aspnes and G. Shah. Skip graphs. *Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 384–393, 2003.
- [3] James Aspnes and Gauri Shah. Skip graphs. *ACM Transactions on Algorithms*, 3(4):37, November 2007.
- [4] James Aspnes and Udi Wieder. The expansion and mixing time of skip graphs with applications. In *SPAA '05: Proceedings of the seventeenth annual ACM symposium on Parallelism in algorithms and architectures*, pages 126–134, New York, NY, USA, 2005. ACM.
- [5] O. Beaumont, N. Bonichon, P. Duchon, and H. Larcheveque. Distributed approximation algorithm for resource clustering. In LNCS Series Springer, editor, *Proceedings of SIROCCO'08*, volume 5058, pages 61–73, 2008.
- [6] Luc Devroye. A limit theory for random skip lists. *The Annals of Applied Probability*, Vol. 2, No. 3, pages 597–609, Aug. 1992.
- [7] Michael T. Goodrich, Michael J. Nelson, and Jonathan Z. Sun. The rainbow skip graph: a fault-tolerant constant-degree distributed data structure. In *SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 384–393, New York, NY, USA, 2006. ACM.
- [8] J.I. Munro, T. Papadakis, and R. Sedgewick. Deterministic skip lists. *Proceedings of the third annual ACM-SIAM symposium on Discrete algorithms*, pages 367–375, 1992.
- [9] W. Pugh. Skip Lists: A Probabilistic Alternative to Balanced Trees. *Communications of the ACM*, 33(6), 1990.