

Please use the \AuthorTitle command defined in defs.tex instead of \Author Please use the \AuthorTitle command defined in defs.tex instead of \Title

Skip lists: A randomized dictionary

The exposition is based on the following sources, which are all recommended reading:

1. Pugh: Skip lists: a probabilistic alternative to balanced treed. Proceedings WADS, LNCS 382, 1989, pp. 437-449
2. Sedgewick: Algorithmen in C++, 2002, Pearsons, (Chapter 13.5)
3. Lecture Script from Michiel Smid, University of the Saarland.

Introduction

We consider the so called *dictionary problem*. Given a set S of real numbers, store them in a data structure such that the following three operations can be performed efficiently:

- **Search(x):** Given the real number x , report the maximal element of $S \cup \{-\infty\}$ that is at most equal to x .
- **Insert(x):** Given a real number x insert it into the data structure.
- **Delete(x):** Given a real number x delete it into the data structure.

Introduction ⁽²⁾

The standard data structures for this problem is the balanced binary tree. It supports all the above operations in worst case time $O(\log n)$ and uses $O(n)$ space.

Well known classes of balanced trees are for example AVL-trees, $BB[\alpha]$ -trees and red-black-trees. In order to maintain their worst case time behaviour all those data structures need more or less elaborate *rebalancing* operations which make an implementation non-trivial, which in turn leads not to the best practical run times.

Introduction ⁽³⁾

We will introduce in this lecture an alternative, *randomized* data structure, the *skip list*. It uses *in expectation* linear space and supports the above dictionary operation in *expected* time $O(\log n)$ with *high probability*.

Why do we do this? We will see that the data structure is conceptually much simpler and more elegant than balanced trees. Nevertheless we will exchange a worst case runtime against an expected run time.

Introduction ⁽⁴⁾

However, the analysis will show that skip lists behave very well and are very fast in practice (the difference is similar to the deterministic merge sort and the randomized quicksort algorithm).

The goal of this lecture is to a) introduce you to the data structure, b) show you how to analyze the randomized run time, and c) to introduce you to *tail estimates* using *Chernoff bounds*.

Skip lists

Throughout the lecture we assume that we can generate random, independent bits in unit time. Let S be a set of n real numbers. Then we construct a sequence of sets S_1, S_2, \dots as follows:

1. For each element $x \in S$, flip a coin until zero comes up.
2. For each $i \geq 1$, S_i is the set of elements in S for which we flipped the coin at least i times.

Let h be the number of sets that are constructed. Then it is clear that

$$\emptyset = S_h \subseteq S_{h-1} \subseteq S_{h-2} \subseteq \dots \subseteq S_2 \subseteq S_1 = S$$

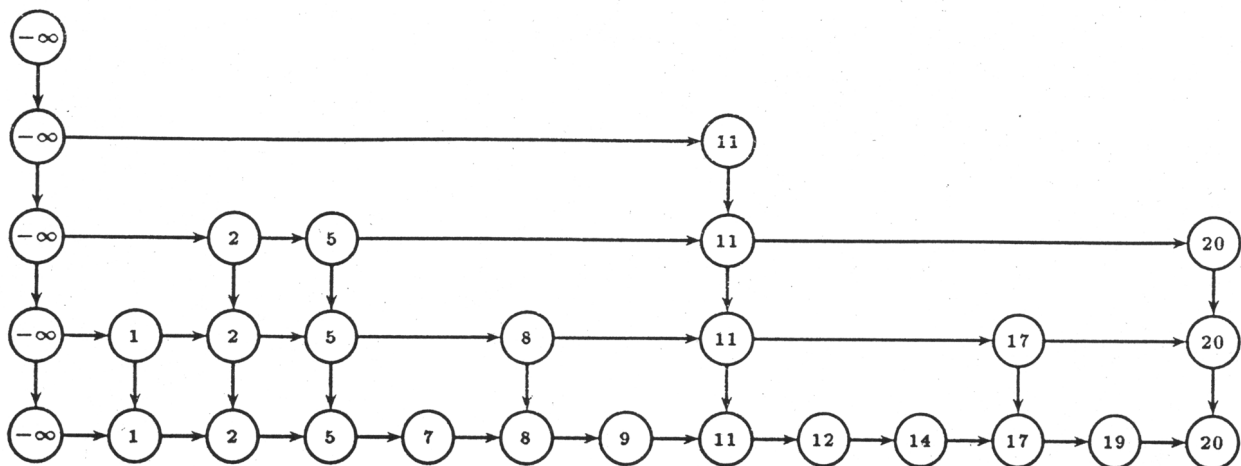
Skip lists ⁽²⁾

The skip list for S consists then of the following:

1. For each $1 \leq i \leq h$, the elements of $S_i \cup \{-\infty\}$ are stored in a sorted linked list L_i .
2. For each $1 < i \leq h$, there is a pointer from each $x \in L_i$ to its occurrence in L_{i-1} .

Skip lists ⁽³⁾

Here is an example. Suppose $S = \{1, 2, 5, 7, 8, 9, 11, 12, 14, 17, 19, 20\}$. Flipping coins might lead to $S_1 = S$, $S_2 = \{1, 2, 5, 8, 11, 17, 20\}$, $S_3 = \{2, 5, 11, 20\}$, $S_4 = \{11\}$, and $S_5 = \emptyset$.



Searching skip lists

We can now implement the search for x as follows:

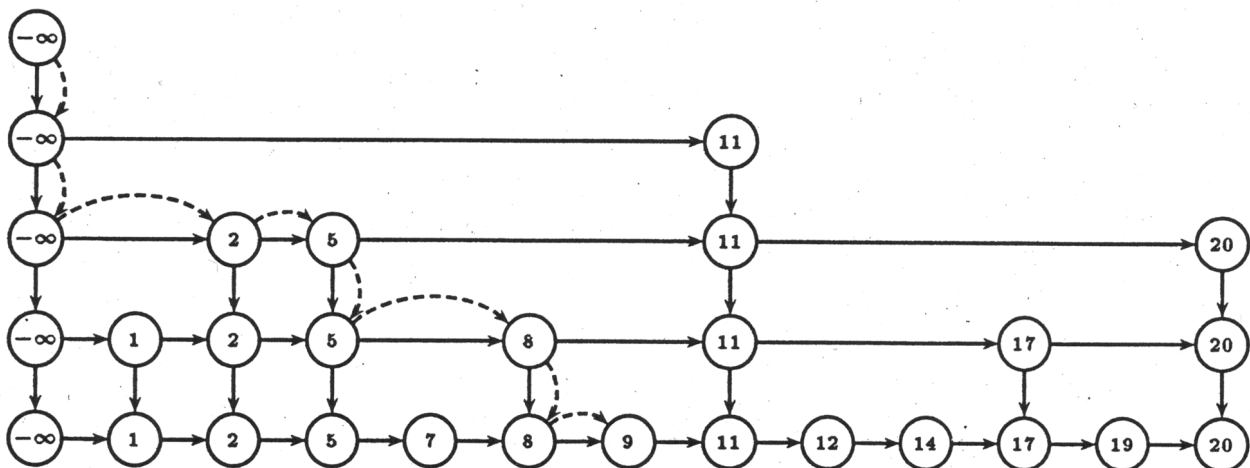
1. Let y_h be the only element in L_h .
2. For $i = h, h-1, \dots, 2$
 - (a) Follow the pointer from y_i in L_i to its occurrence in L_{i-1} .

- (b) Starting in y_{i-1} , walk to the right along L_{i-1} , until an element is reached that is larger than x or the end of L_{i-1} is reached. Let now y_{i-1} be the last encountered element in L_{i-1} that is at most equal to x .

3. Output y_1 .

Searching skip lists ⁽²⁾

The following figure illustrates the search step (we search for element 10):

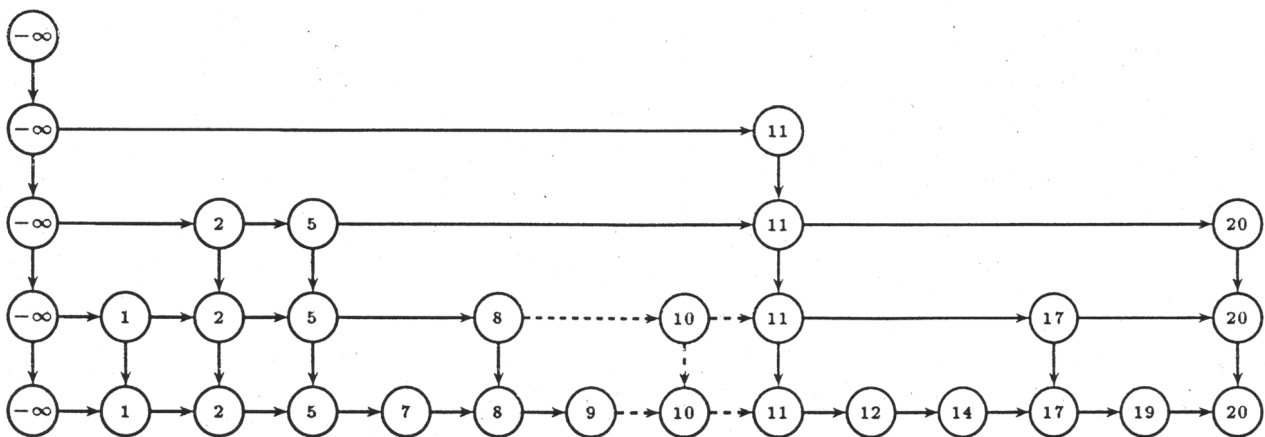


It is not hard to imagine how the insertion and deletion operations work on a skip list.

Inserting into a skip list

For inserting an element x into the dictionary we proceed as follows:

1. Run the search algorithm for x . Let y_1, y_2, \dots, y_h be the elements of L_1, L_2, \dots, L_h that are computed while searching. If $x = y_1$, the $x \in S$ and nothing has to be done. Hence assume that $x \neq y_1$.
2. Flip a coin until a zero comes up. Let l be the number of coin flips.
3. For each $1 \leq i \leq \min(l, h)$, add x to the list L_i immediately after y_i .
4. If $l \geq h$, then create new lists L_{h+1}, \dots, L_{l+1} storing the sets $S_{h+1} \cup \{-\infty\}$, where each set contains x except for S_{l+1} which is empty.
5. For each $1 < i \leq l$, give x in L_i a pointer to its occurrence in L_{i-1} .
6. If $l \geq h$, then for each $h+1 \leq i \leq l+1$, give $-\infty$ in L_i a pointer to its occurrence in L_{i-1} .
7. Set $h = \max(h, l+q)$.



Deleting from a skip list

1. Run the search algorithm for x . Let y_1, y_2, \dots, y_h be the elements of L_1, L_2, \dots, L_h that are computed while searching. If $x \neq y_1$, the $x \notin S$ and nothing has to be done. Hence assume that $x = y_1$.
2. For each $1 \leq i \leq h$ such that $x = y_i$ delete y_i from the list L_i .
3. For $i = h, h-1, \dots$: if L_{i-1} only stores $-\infty$, delete the list L_i and set $h = h-1$.

Why are skiplists efficient? The intuition

We have seen, that mostly what we do in skip lists is to search. The rebalancing is done by throwing a coin a few times and making local changes along the search path.

How expensive is the search? It is the sum over all traversed path length at each level. We expect there to be $\approx \log n$ levels. At each level we travel to the right. However for a fixed level we do not expect to do this long, since this would imply, that all the elements are *not* in the level above.

Hence we expect to spend a constant amount of time at each level which would add up to a total search time of $O(\log n)$. We will now prove this more formally.

Why are skip lists efficient? The proofs

The size of a skip list and the running times of the search and update algorithms are random variables. We will prove that their expected values are bound by $O(n)$ and $O(\log n)$ respectively.

Recall that h denotes the number of sets S_i that result from our probabilistic construction. How can we derive and upper bound for h ?

Let x be an element of S and $h(x)$ be the number of sets S_i that contain x . Then $h(x)$ is a random variable distributed according to a geometric distribution with $p = 1/2$. Hence $\Pr(H(x) = k) = (1/2)^k$ and $E(h(x)) = 2$. That means if we look at a specific element we only expect it to be in S_1 and S_2 .

Why are skip lists efficient? The proofs ⁽²⁾

Clearly $h = 1 + \max\{h(x) : x \in S\}$. From $E(h(x)) = 2$ for any $x \in S$, however, we cannot conclude that the expected value of h is three.

We can estimate $E(h)$ as follows. Again consider a fixed $x \in S$. It follows that for any $k \geq 1$, $h(x) \geq k$ if and only if the first $k - 1$ coin flips produced a one. That is $Pr(h(x) \geq k) = (1/2)^{k-1}$. In addition it is clear that $h \geq k + 1$ if and only if there is a $x \in S$ such that $h(x) \geq k$. Hence

$$Pr(h \geq k + 1) \leq n \cdot Pr(h(x) \geq k) = \frac{n}{2^{k-1}}$$

Why are skip lists efficient? The proofs ⁽³⁾

This estimate does not make sense for $k < 1 + \log n$. For those values of k we can use the trivial upper bound $Pr(h \geq k + 1) \leq 1$. Then $E(h)$ equals:

$$\sum_{k=0}^{\infty} Pr(h \geq k + 1) = \sum_{k=0}^{\lceil \log n \rceil} Pr(h \geq k + 1) + \sum_{k=1+\lceil \log n \rceil}^{\infty} Pr(h \geq k + 1).$$

(exercise: proof the first equality, that is $E(X) = \sum_{k=1}^{\infty} Pr(X \geq k)$ for a random variable X that takes values $\{0, 1, 2, \dots\}$.)

Why are skip lists efficient? The proofs ⁽⁴⁾

The first summation on the right hand side is at most $1 + \lceil \log n \rceil$. The second sum can be bounded from above by:

$$\sum_{k=1+\lceil \log n \rceil}^{\infty} \frac{n}{2^{k-1}} = n(1/2)^{\lceil \log n \rceil - 1} \leq n(1/2)^{\log n - 1} \leq 2.$$

Hence we have proved that $E(h) \leq 3 + \lceil \log n \rceil$.

Why are skip lists efficient? The proofs ⁽⁵⁾

The expected size of a skip list can easily be computed. Let M denote the total size of the sets S_1, S_2, \dots, S_h . Then $M = \sum_{x \in S} h(x)$ and by linearity of expectation:

$$E(M) = \sum_{x \in S} E(h(x)) = \sum_{x \in S} 2 = 2n.$$

If M' denotes the total number of nodes in a skip list, then M' is equal to M plus h . Hence

$$E(M') = E(M + h) = E(M) + E(h) \leq 2n + 3 + \lceil \log n \rceil.$$

Why are skip lists efficient? The proofs ⁽⁶⁾

What is left to do is to estimate the search costs.

Let x be a real number and let C_i denote the number of elements in the list L_i that are inspected when searching for x (We do not count the element of L_i at which the algorithm starts walking to the right. Hence, C_i counts comparisons between x and elements of S .) The search cost is then proportional to $\sum_{i=1}^h (1 + C_i)$.

Again we cannot use linearity of expectation since h is a random variable. Again the trick is to fix an integer A and analyze the search cost up to a level A and above level A separately (and differently).

Why are skip lists efficient? The proofs (7)

We first estimate the search level above A , i.e., the total costs in the lists $L_{A+1}, L_{A+2}, \dots, L_h$. Since the cost is at most equal to the total size of these lists, its expected value is at most equal to the expected value of $M_A := \sum_{i=A+1}^h |L_i|$.

How do we estimate this value? We first note that the lists L_i , $A+1 \leq i \leq h$, form a skip list for S_{A+1} . Hence we have:

$$E(M_A) = \sum_{k=0}^n E(M_A \mid |S_{A+1}| = k) \cdot \Pr(|S_{A+1}| = k)$$

where $E(M_A \mid |S_{A+1}| = k)$ is the expected size of a skip list with k elements. We have already seen that this is $O(k)$.

Why are skip lists efficient? The proofs (8)

Hence we only need to compute $\Pr(|S_{A+1}| = k)$. Since $|S_{A+1}| = k$ if and only if out of the n elements of S exactly k reach the level $A+1$, we have:

$$\Pr(|S_{A+1}| = k) = \binom{n}{k} \left(\frac{1}{2}\right)^{Ak} \left(1 - \left(\frac{1}{2}\right)^A\right)^{n-k}.$$

Setting $p = \frac{1}{2}^A$, we infer that the expected value of M_A is proportional to:

$$\begin{aligned} \sum_{k=0}^n k \cdot \binom{n}{k} p^k (1-p)^{n-k} &= \sum_{k=1}^n n \cdot \binom{n-1}{k-1} p^k (1-p)^{n-k} \\ &= n \cdot p \sum_{k=0}^{n-1} \binom{n-1}{k} p^k (1-p)^{n-1-k} \\ &= n \cdot p (p + (1-p))^{n-1} \\ &= n \cdot p \end{aligned}$$

Hence the expected search cost above level A is bounded by $O(n/2^A)$.

Why are skip lists efficient? The proofs (9)

Next we estimate the expected search cost in the lists L_1, L_2, \dots, L_A . Recall that C_i is the number of elements searched when searching for x . We use again conditional expectation. Let $l_i(x)$ be the number of elements in L_i that are at most equal to x . Then

$$E(C_i) = \sum_{k=1}^n E(C_i \mid l_i(x) = k) \cdot \Pr(l_i(x) = k).$$

Why are skip lists efficient? The proofs (10)

Assume that $l_i(k) = k$. Also assume that there is an element in L_i that is larger than x . Then $C_i = j$ if and only if the largest $j-1$ elements of L_i that are at most equal to x do not appear in L_{i+1} , but the element that immediately precedes these $j-1$ elements does appear in L_{i+1} .

Why are skip lists efficient? The proofs (11)

Hence

$$\Pr(C_i = j \mid l_i(x) = k) \leq \left(\frac{1}{2}\right)^{j-1}, 0 \leq j \leq k.$$

This inequality also holds if x is at least equal to the maximal element of L_i . From this we obtain:

$$\begin{aligned} E(C_i \mid l_i(x) = k) &= \sum_{j=0}^k j \cdot \Pr(C_i = j \mid l_i(x) = k) \\ &\leq \sum_{j=0}^k \frac{j}{2^{j-1}} \\ &\leq 4. \end{aligned}$$

(exercise.)

Why are skip lists efficient? The proofs (12)

This, in turn implies that

$$E(C_i) \leq \sum_{k=1}^n 4 \cdot \Pr(l_i(x) = k) = 4$$

It follows that the expected search cost up to level A is proportional to:

$$E\left(\sum_{i=1}^A (1 + C_i)\right) = \sum_{k=1}^A (1 + E(C_i)) \leq 5A$$

Summarizing we have shown that the expected search time for element x is bounded by:

$$O\left(\frac{n}{2^A} + A\right).$$

Setting A to $\log n$ we obtain the required bound of $O(\log n)$.

Tail estimates: Chernoff bounds

So far we proved bounds on the expected size, search time and update time for a skip list. In this section we consider so called *tail estimates*.

That is, we estimate the probability that the actual search time deviates significantly from its expected value. For example assume for a moment that the constant in the $O(\log n)$ term for the search time is one. Then we want to estimate the probability that the actual search time is at least $t \cdot \log n$.

We could derive an estimate using *Markov's inequality*.

Tail estimates: Chernoff bounds (2)

Lemma 1. Let X be a random variable that takes non-negative values, and let μ be the expected value of X . Then for any $t > 0$, $\Pr(X \geq t\mu) \leq \frac{1}{t}$.

Proof: Let $s = t\mu$. Then

$$\begin{aligned} \mu &= \sum_x x \cdot \Pr(X = x) \\ &\geq \sum_{x \geq s} x \cdot \Pr(X = x) \\ &\geq \sum_{x \geq s} s \cdot \Pr(X = x) \\ &= s \cdot \Pr(X \geq s) \end{aligned}$$

Tail estimates: Chernoff bounds ⁽³⁾

Hence the probability that the actual search time is at least $t \cdot \log n$ is less than or equal to $1/t$.

This is not very impressive. The probability that the search time is more than 100 times its expected value is at most $1/100$. So if this bounds was tight one search in a hundred takes more than 100 times the time of the average search.

Tail estimates: Chernoff bounds ⁽⁴⁾

In this section we will see that Chernoff bounds give a much tighter estimate. We will prove that the probability that the search time exceeds $t \cdot \log n$ is less than or equal to $\approx n^{-t/8}$ for $t \geq 5$.

Hence in a skip list of 1000 elements, the probability that the search time is more than 100 times its expected value is 10^{-38} which in practice means, it will never occur. (Even for $t = 50$ the bound is still 10^{-19} , and for $t = 10$ the probability is still only $\approx 2 \cdot 10^{-4}$).

Tail estimates: Chernoff bounds ⁽⁵⁾

Markov's inequality holds for any non-negative random variable. The Chernoff technique applies to random variables X that can be written as the sum $\sum_{i=1}^n X_i$ of mutually independent random variables X_i .

(Variables are called (mutually) independent if their joint density function is the product of the individual density functions. Beware that mutual independence is different than pairwise independence! (exercise)).

In such cases much better bounds can be obtained.

Tail estimates: Chernoff bounds ⁽⁶⁾

So let $X_1, X_2, X_3, \dots, X_n$ be a sequence of mutually independent random variables and let $X = \sum_{i=1}^n X_i$.

The *moment generating function (mgf)* for a (discrete) random variable Y is defined as

$$m_Y(\lambda) = E(e^{\lambda Y}) = \sum_y e^{\lambda y} \cdot \Pr(Y = y)$$

As the name suggests the function is used to easily generate the moments of the random variable Y . Clearly $m(0) = 1$ and it is easy to show that $\mu = m'(0)$ and $\sigma^2 = m''(0) - \mu^2$ (exercise).

Tail estimates: Chernoff bounds ⁽⁷⁾

In the case of X , which is a sum of n independent variables, $m_X(\lambda) = \prod_{i=1}^n m_{X_i}(\lambda)$ and of course the mean value of X is the derivate of the mgf at position 0 which is simply the product of all means of the X_i .

Or written down:

$$E(e^{\lambda X}) = E(e^{\lambda(X_1 + \dots + X_n)}) = \prod_{i=1}^n E(e^{\lambda X_i}).$$

Tail estimates: Chernoff bounds ⁽⁸⁾

Now let $s > 0$ and $\lambda > 0$. Since $X \geq s$ if and only if $e^{\lambda X} \geq e^{\lambda s}$, we have $\Pr(X \geq s) = \Pr(e^{\lambda X} \geq e^{\lambda s})$. By applying Markov's inequality to the non negative random variable $e^{\lambda X}$, we get

$$\Pr(X \geq s) = \Pr(e^{\lambda X} \geq e^{\lambda s}) \leq e^{-\lambda s} \cdot E(e^{\lambda X}).$$

This yields:

$$Pr(X \geq s) \leq e^{-\lambda s} \cdot \prod_{i=1}^n E(e^{\lambda X_i}), \text{ for } s > 0 \text{ and } \lambda > 0.$$

This is the basic inequality we work with. To estimate $Pr(X \geq s)$ we need bound on $E(e^{\lambda X_i})$. Of course those bounds depend on the probability distribution of X_i . We will now illustrate the technique using the *geometric distribution* with parameter $p = 1/2$.

Tail estimates: Chernoff bounds (9)

Let T be the number of flips we need until a one comes up in a series of coin flips. Then $Pr(T = k) = (1/2)^k$ for $k \geq 1$ and $E(T) = 2$. Now assume we are interested in T_n which is the number of flips we need until we obtain a one exactly n times (i.e. $T = T_1$).

Tail estimates: Chernoff bounds (10)

If we define the random variables X_i as the number of flips between the $(i-1)$ -st (excluding) and the i -th one (including), then X_i is distributed according to a geometric distribution. (This property is also called the memoryless property of the geometric or exponential distribution).

Then $T_n = \sum_{i=1}^n X_i$, where each X_i is distributed according to a geometric distribution and the expected value is $E(T_n) = 2n$, and Markov's inequality gives $Pr(T_n \geq (2+t)n) \leq \frac{2}{2+t}$.

Tail estimates: Chernoff bounds (11)

For $0 < \lambda < \log 2$ we have

$$E(e^{\lambda X_i}) = \sum_{k=1}^{\infty} e^{\lambda k} \cdot Pr(X_i = k) = \sum_{k=1}^{\infty} (e^{\lambda}/2)^k = \frac{e^{\lambda}}{2 - e^{\lambda}}$$

We now apply our basic inequality with $s = (2+t)n$, where $t > 0$ and get

$$Pr(T_n \geq (2+t)n) \leq e^{-\lambda(2+t)n} \left(\frac{e^{\lambda}}{2 - e^{\lambda}} \right)^n = \left(\frac{e^{-\lambda(1+t)}}{2 - e^{\lambda}} \right)^n$$

Now we choose λ such that the term on the right hand side is minimized (exercise) and find $\lambda = \log(1 + \frac{t}{2+t})$.

Tail estimates: Chernoff bounds (12)

Hence we have

$$Pr(T_n \geq (2+t)n) \leq (1+t/2)^n \left(1 - \frac{t}{2+2t} \right)^{(1+t)n}.$$

Since $1 - x \leq e^{-x}$ for all x , we have

$$\left(1 - \frac{t}{2+2t} \right)^{1+t} \leq (e^{-\frac{t}{2+2t}})^{1+t} = e^{-t/2}.$$

Moreover, $1 + t/2 \leq e^{t/4}$ for $t \geq 3$. This proves that for $t \geq 3$

$$Pr(T_n \geq (2+t)n) \leq e^{tn/4} \cdot e^{-tn/2} = e^{-tn/4}.$$

Compare this with the bound obtained from Markov's inequality (which was $\frac{2}{2+t}$)!

Tail estimates: Chernoff bounds (13)

We can subsume our finding in the following theorem:

Theorem 2. Let X_1, X_2, \dots, X_n be mutually independent random variables and assume that each X_i is distributed according to a geometric distribution. Let $T_n = \sum X_i$, then $E(T_n) = 2n$ and for any $t \geq 3$ holds:

$$\Pr(T_n \geq (2+t)n) \leq e^{tn/4} \cdot e^{-tn/2} = e^{-tn/4}.$$

Corollary 3. Let $c \geq 1$ be a constant and let m be a positive integer. Further let $n = c \cdot \ln m$. Then for any $s \geq 5$ it holds

$$\Pr(T_{\lceil n \rceil} \geq sn) \leq m^{-\frac{(s-2)c}{4}}.$$

Tail estimates: Skip lists

We can now use the above results to prove tail estimates for the size, search time and update time for a skip list. Consider again a skip list for n elements. Let M denote the total size of the sets S_1, \dots, S_h and let M' denote the total number of nodes of the skip list. Then $M' = M + h$. We have seen that $E(M') = E(M + h) \leq 2n + 3 + \lceil \log n \rceil$.

Now we are interested in the probability that M' is at least equal to $(2+t)n$. Clearly if $M' \geq (2+t)n$, then $h \geq \frac{tn}{2}$ or $M \geq (2 + \frac{t}{2})n$. As a result,

$$\Pr(M' \geq (2+t)n) \leq \Pr(h \geq \frac{tn}{2}) + \Pr(M \geq (2 + \frac{t}{2})n).$$

Tail estimates: Skip lists (2)

In the beginning of the lecture we estimated already the first probability on the right hand side and have:

$$\Pr(h \geq \frac{tn}{2}) \leq \frac{n}{2^{\frac{tn}{2}-2}} = e^{\ln n + 2 \ln 2 - \frac{tn}{2} \ln 2} \leq e^{-\frac{tn}{8}}.$$

It remains to bound $\Pr(M \geq (2 + \frac{t}{2})n)$.

Remember that $h(x)$ was defined as the number of sets S_i that contain x . Then the random variables $h(x)$ for $x \in S$ are all mutually independent and have a geometric distribution. Since $M = \sum_{x \in S} h(x)$ we can apply above theorem for geometric distributions and get:

$$\Pr(M \geq (2 + \frac{t}{2})n) \leq e^{-\frac{tn}{8}}.$$

That means it is extremely unlikely that a skip list deviates much from its expected size.

Tail estimates: Skip lists (3)

The analysis of the search time is a bit more complicated. Let $x \in \mathbb{R}$ and let $c \geq 1$ be a constant. We want to estimate the probability that the searching for x takes more than $c \cdot \log n$ steps. We analyze this by considering the cost up to level $\lceil c \cdot \log n \rceil$ and above level $\lceil c \cdot \log n \rceil$ separately.

Let A denote the number of nodes traversed in levels $1 + \lceil c \cdot \log n \rceil, 2 + \lceil c \cdot \log n \rceil, \dots, h$ when searching for x . We proved before that the expected value of A is bounded by $O(\frac{n}{2^{\lceil c \cdot \log n \rceil}}) = O(\frac{1}{n^{c \ln 2 - 1}})$. Hence Markov's inequality yields:

$$\Pr(A \geq 1) \leq E(A) = O(\frac{1}{n^{c \ln 2 - 1}}).$$

Tail estimates: Skip lists ⁽⁴⁾

Now let B be the number of nodes traversed in the lower levels. We want to express B now as the sum of mutually independent random variables each distributed according to a geometric distribution.

In order to do this we give an slightly different but equivalent construction of a skip list.

Let $S_1 = S = \langle s_{11} < s_{12} < \dots < s_{1n_1} \rangle$, where $n_1 = n$. Moreover, let $l_1(x)$ be the number of elements of S_1 that are at most equal to x .

Let $i \geq 1$ and assume that we have constructed $S_i = \langle s_{i1} < s_{i2} < \dots < s_{in_i} \rangle$ with $n_i = |S_i|$. Let $l_i(x) = |\{y \in S_i : y \leq x\}|$.

Tail estimates: Skip lists ⁽⁵⁾

We do the following:

1. We flip our coin $n_i - l_i(x)$ times.
2. We flip the coin $l_i(x)$ times.
3. If the previous $l_i(x)$ flips only produced zeros, then we flip the coin until we get a one.

Let f_i be the total number of flips made during these stages, and denote the outcome by $F_{i1}, F_{i2}, \dots, F_{if_i}$. Note that $f_i \geq n_i$. We proceed as follows:

Tail estimates: Skip lists ⁽⁶⁾

1. Define $S_{i+1} = \{s_{ij} : 1 \leq j \leq n_i \text{ and } F_{ij} = 1\}$, $n_{i+1} = |S_{i+1}|$ and $l_{i+1}(x) = |\{y \in S_{i+1} : y \leq x\}|$.
2. We define the random variable X_i as the number of flips after stage 1 until the first 1 comes up.
3. We define the random variable C_i as the minimum of $1 + l_i(x)$ and X_i .

The construction stops as soon as the empty set S_{i+1} has been constructed.

Tail estimates: Skip lists ⁽⁷⁾

Hence we have now

$$\emptyset = S_h \subset S_{h-1} \subseteq S_{h-2} \subseteq \dots \subseteq S_2 \subseteq S_1 = S.$$

and random variables X_1, X_2, \dots, X_{h-1} and C_1, C_2, \dots, C_{h-1} . This construction defines a probability distribution on skip lists, which is the same as that of our previous two constructions.

Tail estimates: Skip lists ⁽⁸⁾

Lets go back to the analysis of the search time. For convenience, we define $C_i = 0$ for all $i \geq h$. Recall that B is the number of nodes traversed in the levels $1, 2, \dots, \lceil c \cdot \log n \rceil$. Hence we have $B = \sum_{i=1}^{\lceil c \cdot \log n \rceil} C_i$. The means we have written B as the sum of random variables.

However, since C_i depends on $l_i(x)$, which depends on n_i which in turn depends on n_{i-1} , we see that the variables C_1, C_2, \dots, C_{h-1} are not mutually independent.

Tail estimates: Skip lists ⁽⁹⁾

On the other hand, it is easy to see that $C_i \leq X_i$ for all $1 \leq i \leq h-1$. Moreover, the random variables X_i , $1 \leq i \leq h-1$ are mutually independent and each one is distributed according to a geometric distribution. For $i \geq h$, let X_i also denote a random variable distributed according to a geometric distribution. Since

$$B = \sum_{i=1}^{\lceil c \cdot \log n \rceil} C_i \leq \sum_{i=1}^{\lceil c \cdot \log n \rceil} X_i,$$

we infer

$$P(B \geq s \cdot c \cdot \log n) \leq Pr\left(\sum_{i=1}^{\lceil c \cdot \log n \rceil} X_i \geq s \cdot c \cdot \log n\right)$$

The above Corollary immediately gives:

$$P(B \geq s \cdot c \cdot \log n) \leq n^{\frac{-(s-2)c}{4}} \leq \frac{1}{n^{\frac{sc}{8}}},$$

for $c \geq 1$ and $s \geq 5$.

Tail estimates: Skip lists (10)

Now we can give the tail estimate for the search time. Let T denote the total number of nodes traversed when searching for x . Then $T = A + B$. Moreover, if $T \geq 1 + 5c \cdot \log n$, then $A \geq 1$ or $B \geq 5c \cdot \log n$. Hence

$$Pr(T \geq 1 + 5c \cdot \log n) \leq Pr(A \geq 1) + Pr(B \geq 5c \cdot \log n).$$

Our results for the two probabilities on the right hand side imply for $c \geq 1$,

$$Pr(T \geq 1 + 5c \cdot \log n) = O(n^{1-c \cdot \ln 2} + n^{-5c/8}).$$

Taking $c = t/5$, where $t \geq 5$ we get

$$Pr(T \geq 1 + t \log n) = O(n^{1-(t \cdot \ln 2)/5} + n^{-t/8})$$