


Skip list

Skip List		
Type	List	
Invented	1989	
Invented by	W. Pugh	
Time complexity in big O notation		
	Average	Worst case
Space	$O(n)$	$O(n \log n)^{[1]}$
Search	$O(\log n)$	$O(n)$
Insert	$O(\log n)$	$O(n)$
Delete	$O(\log n)$	$O(n)$

Part of a series on
Probabilistic data structures
Bloom filter · Quotient filter · Skip list
Random trees
Random binary tree · Treap Rapidly exploring random tree
Related
Randomized algorithm
Computer science Portal

A **skip list** is a data structure for storing a sorted list of items using a hierarchy of linked lists that connect increasingly sparse subsequences of the items. These auxiliary lists allow item lookup with efficiency comparable to balanced binary search trees (that is, with number of probes proportional to $\log n$ instead of n).



Each link of the sparser lists skips over many items of the full list in one step, hence the structure's name. These forward links may be added in a randomized way with a geometric / negative binomial distribution. Insert, search and delete operations are performed in logarithmic expected time. The links may also be added in a non-probabilistic way so as to guarantee amortized (rather than merely expected) logarithmic cost.^[2]

Description

A skip list is built in layers. The bottom layer is an ordinary ordered linked list. Each higher layer acts as an "express lane" for the lists below, where an element in layer i appears in layer $i+1$ with some fixed probability p (two commonly used values for p are $1/2$ or $1/4$). On average, each element appears in $1/(1-p)$ lists, and the tallest element (usually a special head element at the front of the skip list) in $\log_{1/p} n$ lists.

A search for a target element begins at the head element in the top list, and proceeds horizontally until the current element is greater than or equal to the target. If the current element is equal to the target, it has been found. If the current element is greater than the target, or the search reaches the end of the linked list, the procedure is repeated after returning to the previous element and dropping down vertically to the next lower list. The expected number of steps in each linked list is at most $1/p$, which can be seen by tracing the search path backwards from the target until reaching an element that appears in the next higher list or reaching the beginning of the current list. Therefore, the total *expected* cost of a search is $(\log_{1/p} n)/p$, which is $\mathcal{O}(\log n)$ when p is a constant. By choosing different values of p , it is possible to trade search costs against storage costs.

Implementation details

The elements used for a skip list can contain more than one pointer since they can participate in more than one list.

Insertions and deletions are implemented much like the corresponding linked-list operations, except that "tall" elements must be inserted into or deleted from more than one linked list.

$\mathcal{O}(n)$ operations, which force us to visit every node in ascending order (such as printing the entire list), provide the opportunity to perform a behind-the-scenes derandomization of the level structure of the skip-list in an optimal way, bringing the skip list to $\mathcal{O}(\log n)$ search time. (Choose the level of the i 'th finite node to be 1 plus the number of times we can repeatedly divide i by 2 before it becomes odd. Also, $i=0$ for the negative infinity header as we have the usual special case of choosing the highest possible level for negative and/or positive infinite nodes.) However this also allows someone to know where all of the higher-than-level 1 nodes are and delete them.

Alternatively, we could make the level structure quasi-random in the following way:

```

make all nodes level 1
j ← 1
while the number of nodes at level j > 1 do
  for each i'th node at level j do
    if i is odd
      if i is not the last node at level j
        randomly choose whether to promote it to level j+1
      else
        do not promote
      end if
    else if i is even and node i-1 was not promoted
      promote it to level j+1
    end if
  repeat
  j ← j + 1
repeat

```

Like the derandomized version, quasi-randomization is only done when there is some other reason to be running a $\mathcal{O}(n)$ operation (which visits every node).

The advantage of this quasi-randomness is that it doesn't give away nearly as much level-structure related information to an adversarial user as the de-randomized one. This is desirable because an adversarial user who is able to tell which nodes are not at the lowest level can pessimize performance by simply deleting higher-level nodes. The search performance is still guaranteed to be logarithmic.

It would be tempting to make the following "optimization": In the part which says "Next, for each i'th...", forget about doing a coin-flip for each even-odd pair. Just flip a coin once to decide whether to promote only the even ones or only the odd ones. Instead of $\mathcal{O}(n \log n)$ coin flips, there would only be $\mathcal{O}(\log n)$ of them. Unfortunately, this gives the adversarial user a 50/50 chance of being correct upon guessing that all of the even numbered nodes (among the ones at level 1 or higher) are higher than level one. This is despite the property that he has a very low probability of guessing that a particular node is at level N for some integer N .

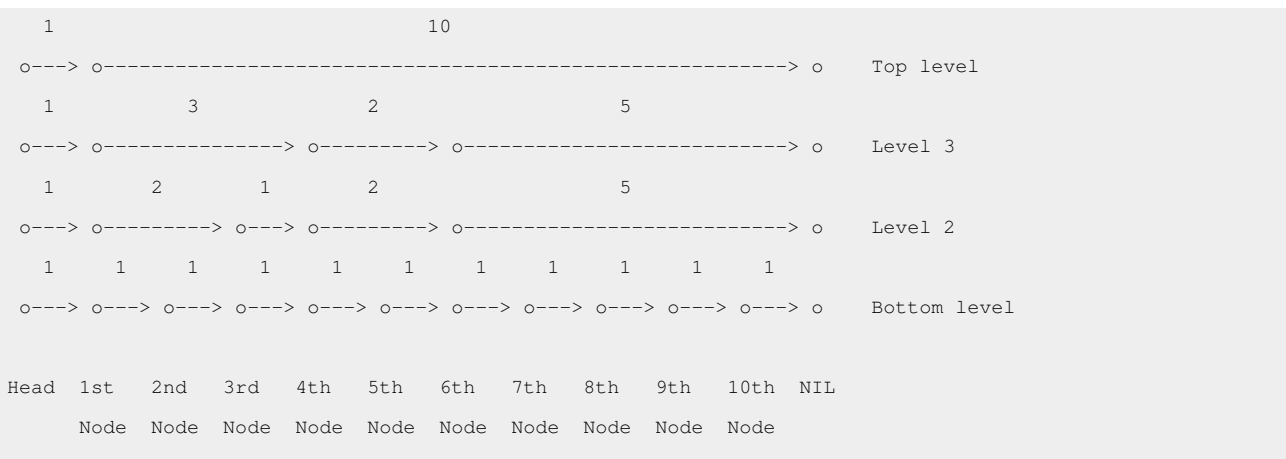
A skip list does not provide the same absolute worst-case performance guarantees as more traditional balanced tree data structures, because it is always possible (though with very low probability) that the coin-flips used to build the skip list will produce a badly balanced structure. However, they work well in practice, and the randomized balancing scheme has been argued to be easier to implement than the deterministic balancing schemes used in balanced binary search trees. Skip lists are also useful in parallel computing, where insertions can be done in different parts of the skip list in parallel without any global rebalancing of the data structure. Such parallelism can be especially advantageous for resource discovery in an ad-hoc Wireless network because a randomized skip list can be made robust to the loss of any single node.

There has been some evidence that skip lists have worse real-world performance and space requirements than B trees due to memory locality and other issues.^[3]

Indexable skiplist

As described above, a skiplist is capable of fast $\mathcal{O}(\log n)$ insertion and removal of values from a sorted sequence, but it has only slow $\mathcal{O}(n)$ lookups of values at a given position in the sequence (i.e. return the 500th value); however, with a minor modification the speed of random access indexed lookups can be improved to $\mathcal{O}(\log n)$. For every link, also store the width of the link. The width is defined as the number of bottom layer links being traversed by each of the higher layer "express lane" links.

For example, here are the widths of the links in the example at the top of the page:



Notice that the width of a higher level link is the sum of the component links below it (i.e. the width 10 link spans the links of widths 3, 2 and 5 immediately below it). Consequently, the sum of all widths is the same on every level ($10 + 1 = 1 + 3 + 2 + 5 = 1 + 2 + 1 + 2 + 5$).

To index the skiplist and find the i 'th value, traverse the skiplist while counting down the widths of each traversed link. Descend a level whenever the upcoming width would be too large.

For example, to find the node in the fifth position (Node 5), traverse a link of width 1 at the top level. Now four more steps are needed but the next width on this level is ten which is too large, so drop one level. Traverse one link of width 3. Since another step of width 2 would be too far, drop down to the bottom level. Now traverse the final link of width 1 to reach the target running total of 5 (1+3+1).

```

function lookupByPositionIndex(i)
  node ← head
  i ← i + 1                                # don't count the head as a step
  for level from top to bottom do
    while i ≥ node.width[level] do # if next step is not too far
      i ← i - node.width[level] # subtract the current width
      node ← node.next[level]   # traverse forward at the current level
    repeat
  repeat
  return node.value
end function

```

This method of implementing indexing is detailed in Section 3.4 Linear List Operations in "A skip list cookbook" by William Pugh ^[4].

History

Skip lists were first described in 1990 by William Pugh.

To quote the author:

Skip lists are a probabilistic data structure that seem likely to supplant balanced trees as the implementation method of choice for many applications. Skip list algorithms have the same asymptotic expected time bounds as balanced trees and are simpler, faster and use less space.

Usages

List of applications and frameworks that use skip lists:

- Cyrus IMAP server offers a "skiplist" backend DB implementation (source file ^[5])
- QMap ^[6] (up to Qt 4) template class of Qt that provides a dictionary.
- Redis, an ANSI-C open-source persistent key/value store for Posix systems, uses skip lists in its implementation of ordered sets.
- nessDB ^[7], a very fast key-value embedded Database Storage Engine (Using log-structured-merge (LSM) trees), uses skip lists for its memtable.
- skipdb ^[8] is an open-source database format using ordered key/value pairs.
- ConcurrentSkipListSet ^[9] and ConcurrentSkipListMap ^[10] in the Java 1.6 API.
- leveldb ^[11], a fast key-value storage library written at Google that provides an ordered mapping from string keys to string values
- Skip lists are used for efficient statistical computations ^[12] of running medians ^[13] (also known as moving medians).

Skip lists are also used in distributed applications (where the nodes represent physical computers, and pointers represent network connections) and for implementing highly scalable concurrent priority queues with less lock contention, ^[14] or even without locking, as well lockless concurrent dictionaries. There are also several US patents for using skip lists to implement (lockless) priority queues and concurrent dictionaries. ^[citation needed]

See Also

- Bloom filter
- Skip graph
- Skip trees, an alternative data structure to Skip lists in a concurrent approach: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.47.514>
- Skip tree graphs: <http://www0.cs.ucl.ac.uk/staff/a.gonzalezbeltran/pubs/icc2007.pdf>, <http://www0.cs.ucl.ac.uk/staff/a.gonzalezbeltran/pubs/AGB-comcom08.pdf>

References

- [1] <http://www.cs.uwaterloo.ca/research/tr/1993/28/root2side.pdf>
- [2] Deterministic skip lists (<http://www.ic.unicamp.br/~celio/peer2peer/skip-net-graph/deterministic-skip-lists-munro.pdf>)
- [3] http://resnet.uoregon.edu/~gurney_j/jmpc/skiplist.html
- [4] <http://cg.scs.carleton.ca/~morin/teaching/5408/refs/p90b.pdf>
- [5] http://git.cyrusimap.org/cyrus-imapd/tree/lib/cyrusdb_skiplist.c
- [6] <http://qt-project.org/doc/qt-4.8/qmap.html#details>
- [7] <https://github.com/shuttler/nessDB>
- [8] <http://www.dekorte.com/projects/opensource/skipdb/>
- [9] <http://download.oracle.com/javase/6/docs/api/java/util/concurrent/ConcurrentSkipListSet.html>
- [10] <http://download.oracle.com/javase/6/docs/api/java/util/concurrent/ConcurrentSkipListMap.html>
- [11] <https://code.google.com/p/leveldb/>
- [12] <http://code.activestate.com/recipes/576930/>
- [13] https://en.wikipedia.org/wiki/Moving_average#Moving_median
- [14] Skiplist-based concurrent priority queues (<http://dx.doi.org/10.1109/IPDPS.2000.845994>)

External links

- "Skip list" entry (<http://nist.gov/dads/HTML/skiplist.html>) in the Dictionary of Algorithms and Data Structures
- Skip Lists: A Linked List with Self-Balancing BST-Like Properties ([http://msdn.microsoft.com/en-us/library/ms379573\(VS.80\).aspx#datastructures20_4_topic4](http://msdn.microsoft.com/en-us/library/ms379573(VS.80).aspx#datastructures20_4_topic4)) on MSDN in C# 2.0
- SkipDB, a BerkeleyDB-style database implemented using skip lists. (<http://dekorte.com/projects/opensource/SkipDB/>)
- Skip Lists lecture (MIT OpenCourseWare: Introduction to Algorithms) (http://videlectures.net/mit6046jf05_demaine lec12/)
- Open Data Structures - Chapter 4 - Skiplists (http://opendatastructures.org/versions/edition-0.1e/ods-java/4_Skiplists.html)

Demo applets

- Skip List Applet (<http://people.ksp.sk/~kuko/bak/index.html>) by Kubo Kovac
- Thomas Wenger's demo applet on skiplists (<http://iamwww.unibe.ch/~wenger/DA/SkipList/>)

Implementations

- A generic Skip List in C++ (<http://codingplayground.blogspot.com/2009/01/generic-skip-list-skiplist.html>) by Antonio Gulli
- Algorithm::SkipList, implementation in Perl on CPAN (<https://metacpan.org/module/Algorithm::SkipList>)
- John Shipman's implementation in Python (<http://infohost.nmt.edu/tcc/help/lang/python/examples/pyskip/>)
- Raymond Hettinger's implementation in Python (<http://code.activestate.com/recipes/576930/>)
- A Lua port of John Shipman's Python version (http://love2d.org/wiki/Skip_list)
- Java Implementation with index based access (<https://gist.github.com/dmx2010/5426422>)
- ConcurrentSkipListSet documentation for Java 6 (<http://java.sun.com/javase/6/docs/api/java/util/concurrent/ConcurrentSkipListSet.html>) (and sourcecode (<http://www.docjar.com/html/api/java/util/>

[concurrent/ConcurrentSkipListSet.java.html](#))

Article Sources and Contributors

Skip list *Source:* <http://en.wikipedia.org/w/index.php?oldid=576041586> *Contributors:* A3 nm, Alan Dawrst, Almkglor, Altenmann, Andreas Kaufmann, Antaeus Feldspar, Bondolo, Bovineone, Braddunbar, CRGreathouse, Carlsotr, Cedar101, Cereblio, ChadmcDaniel, Charles Matthews, Cybercobra, Dcoetzee, Denisarona, DevynCI, Dmx2010, Doradus, Dysprosia, Fredrik, Funandtrvl, Gene91, Intrgr, Ivan Štambuk, JaGa, Jim10701, Jorge Stolfi, Jrockway, Kukolar, Laurienne Bell, Menahem.fuchs, MinorContributor, Mrwojo, Musically ut, Nczempin, Nkour, Noldoaran, Nsfmc, Overred, Patmorin, Pet3ris, Piet Delpoort, Populus, PuercoPop, Puraltruism, R. S. Shaw, Rdhettinger, RhaneKom, Rpk512, Sanchom, Silly rabbit, Stevenj, Svick, Two Bananas, Viebel, Vishalvishnoi, Waffleguy4, Wavelength, Wojciech mula, Xcez-be, Zr2d2, 175 anonymous edits

Image Sources, Licenses and Contributors

Image:Skip list.svg *Source:* http://en.wikipedia.org/w/index.php?title=File:Skip_list.svg *License:* Public Domain *Contributors:* Wojciech Mula

License

Creative Commons Attribution-Share Alike 3.0
[//creativecommons.org/licenses/by-sa/3.0/](http://creativecommons.org/licenses/by-sa/3.0/)
