Recitation 14 — Hashing and Leftist Heaps

Parallel and Sequential Data Structures and Algorithms, 15-210 (Fall 2013)

December 4, 2013

Announcements:

- Homework 9 is due Friday. It's a lot of fun, but has several DP problems to think about, and may take you a while. Hopefully, you've started it.
- The final exam will be next Friday, December 13 at 8:30 AM. Watch for announcements about rooms.

Today's Agenda:

- Hashing Review
- Parallel Hashing
- Leftist Heaps
- Hashing Application

1 Hashing Review

We have a large space S_{keys} of keys (this might be infinite as in the case of the set of all possible strings, or simply large like all binary strings of length 1024) and a target range $\{0, \ldots, m-1\}$. We normally expect $|S_{keys}| \gg m$. A hash function h is a mapping from S_{keys} to $\{0, \ldots, m-1\}$.

One important application of hash functions is in implementing a dictionary data structure (hash table).

What are some properties we want out of a hash function?

- It should be a deterministic function—if we ask for the hash value of the same key twice, we should get the same value each time.
- It should distribute keys across buckets uniformly.

In general, there's a tension between wanting "random-like" behavior and wanting determinism. For now, we'll suppose we have a "good" hash function (without going too deep into what that means). For concreteness, though, we will make the following assumption:

> **The Simple Uniform Hashing Assumption.** Any given key $k \in K$ is equally likely to be hashed to one of the m possible values, independently of what values other keys might hash to.

Let's now address another problem with hashing, collision resolution.

Q: What does "load factor" refer to and why does it matter?

A: The ratio n/m. Literally, it is how full the hash table is. It is also an indicator of how often we should expect a collision.

Q: What are some techniques for dealing with a collision?

A: Separate chaining and open addressing.

1. *Separate chaining*: we build a so-called "chain"—a sequence or a list containing all the keys that hash to the same value; therefore, with this method, insertion, search, and deletion operations require traversing such a chain and have running time proportional to the length of the chain.
2. *Open addressing:* Fit each key to a slot in a single array. For example, when a key k is inserted, we first consider the hash location $h(k)$ and if that is occupied, we proceed to a different location according to some *probe sequence* until an empty slot is found. Searching is done in the same manner.

Formally, a probe sequence is a function $h(k,i)$ which decides the i-th alternative location for key k. For convenience we will define $h(k,0) = h(k)$. Thus, for a key k, we'd consider locations $h(k,0), h(k,1), h(k,2), \ldots$ in this order.

Q: Can you suggest some probe sequence patterns?

Linear probing is the simplest and most widely used probe sequence strategy, whereby $h(k,i) = [h(k) + i)] \bmod m$. That is, first it tries $h(k)$, then $h(k)+1$, then $h(k)+2$, so on so forth. This sequence wraps around if we reach the end of the array. This strategy has very good cache locality and performs really well in practice. In fact, it also has many nice theoretical performance guarantees, which we won't be able to cover here. Because of its simplicity, linear probing is often the method of choice in generic hash table implementations these days.

Alternatively, we have *quadratic probing* where $h(k,i) = [h(k)+i^2] \bmod m$, or *double hashing* where $h(k,i) = [h(k) + i \cdot g(k)] \bmod m$ utilizing a second hash function g. In theory, all these techniques perform within constant factors of each other in expectation and are ideal in different sets of conditions.

2 Parallel Hashing

Q: What do we mean by hashing in parallel?

A: In the parallel context, instead of inserting, finding or deleting one key at a time, each operation takes a set of keys and performs the operation on all of the keys in parallel. This becomes a problem when there are collisions, but with relatively few keys being used in parallel and a good hash function, these should be rare.

Q: How might we parallelize open addressing?

A: The idea is to use open addressing in multiple rounds. For `insert`, each round attempts to write the keys into the table at their appropriate hash position in parallel. For any key that cannot be written because another key is already there, the key continues for another round using its next probe location. Rounds repeat until it writes all the keys to the table.

In order to prevent writing to a position already occupied in the table, we introduce a variant of the `inject` function. The function

$$\texttt{injectCond}(IV,S) : (\texttt{int} \times \alpha)\ \texttt{seq} \times (\alpha\ \texttt{option})\ \texttt{seq} \to (\alpha\ \texttt{option})\ \texttt{seq}$$

takes a sequence of index-value pairs $\langle (i_1,v_1),\ldots,(i_n,v_n)\rangle$ and a target sequence S and conditionally writes each value v_j into location i_j of S. In particular it writes the value only if the location is set to NONE and there is no previous equal index in IV. That is, it conditionally writes the value for the *first* occurrence of an index; recall `inject` uses the *last* occurrence of an index.

For example, if $S = \langle \texttt{SOME}\ 5, \texttt{NONE}, \texttt{NONE}, \texttt{SOME}\ 42, \texttt{NONE}, \texttt{SOME}\ 28\rangle$, then
$\texttt{injectCond}(\langle(3,43),(1,97),(4,8),(1,35)\rangle, S) = \langle \texttt{SOME}\ 5, \texttt{SOME}\ 97, \texttt{NONE}, \texttt{SOME}\ 42, \texttt{SOME}\ 8, \texttt{SOME}\ 28\rangle$.

Let's write `insert`. Let T be our hash table and K be the set of keys we wish to insert.

```
1  function insert(T,K) =
2  let
3    function insert'(T,K,i) =
4      if |K| = 0 then T
5      else let
6          T' = injectCond({(h(k,i),k) : k ∈ K}, T)
7          K' = {k : k ∈ K | T[h(k,i)] ≠ k}
8      in
9        insert'(T',K',i+1)
10     end
11 in
12   insert'(T,K,0)
13 end
```

For round i (starting at $i = 0$), `insert` attempts to put each key k into the hash table at position $h(k,i)$, but only if the position is empty. To see whether it successfully wrote a key to the table, it reads the values written to the table and checks if they are the same as the keys. In this way it can filter out all the keys that it successfully wrote to the table. It repeats the process on any keys that did not get hashed on the next round using their next probe position $h(k,i+1)$. Rounds continue until every element was put in the hash table.

For example, suppose the table has the following entries before round i:

$T =$

0	1	2	3	4	5	6	7
	A		B			D	F

If $K = \langle E, C \rangle$ and $h(E,i)$ is 1 and $h(C,i)$ is 2, then $IV = \langle (1,E),(2,C) \rangle$ and `insert'` would fail to write E to index 1 but would succeed in writing C to index 2, resulting in the following table:

$T' =$

0	1	2	3	4	5	6	7
	A	C	B			D	F

It then repeats the process with $K' = \langle E \rangle$ and $i+1$.

Note that if T is implemented using an `stseq`, then parallel `insert` basically does the same work as the sequential version which adds the keys one by one. The difference is that the parallel version may add keys to the table in a different order than the sequential. For example, with linear probing, the parallel version adds C first using 1 probe and then adds E at index 4 using 4 probes:

$T_P =$

0	1	2	3	4	5	6	7
	A	C	B	E		D	F

Whereas, the sequential version might add E first using 2 probes, and then C using 3 probes:

$T_S =$

0	1	2	3	4	5	6	7
	A	E	B	C		D	F

Both make 5 probes in the table. Since we showed that, with suitable hash functions and load factors, the expected cost of insert is $O(1)$, the expected work for the parallel version is $O(|K|)$. In addition, if the hash table is large enough, the expected size of K decreases by a constant fraction in each round, so the span is $O(\log |K|)$.

3 Heaps

Recall from lecture that a priority queue data structure can be implemented using a heap. A min-heap is a tree satisfying the property that the key stored at every node is less than or equal to the keys of all of its

descendants. A max-heap is one where the key stored at every node is greater than or equal to the keys of all of its descendants.

A useful operation of heaps is the *meld* operation, which joins two heaps into one. An implementation of meld traverses the right spine of each tree, so we would like to guarantee that the right spine is not too long. In particular if our two trees are of size n and m, we would like to guarantee that the right spines of the two trees are of length $O(\log n)$ and $O(\log m)$, respectively. In this case, melding the two trees would require $O(\log n + \log m)$ work and span.

3.1 Leftist Heaps

A leftist heap is a type of heap which supports efficient meld operations. It has the property that the right spine of any subtree is at most as long as the subtree's left spine.

We define the *rank* of a node x as

$$\mathsf{rank}(x) = \#\text{ of nodes on the right spine of the subtree rooted at } x,$$

and more formally:

$$\mathsf{rank}(\mathsf{leaf}) = 0$$
$$\mathsf{rank}(\mathsf{node}(\_, \_, R)) = 1 + \mathsf{rank}(R)$$

Recall that all nodes of a leftist heap have the "leftist property". That is, if $L(x)$ and $R(x)$ are the left and right children of x, then we have:

Leftist Property: For all node x in a leftist heap, $\mathsf{rank}(L(x)) \geq \mathsf{rank}(R(x))$

This is why the tree is called leftist: for each node in the heap, the rank of the left child must be at least the rank of the right child. That way, the right spine of a leftist heap is kept short as most of the entries will amass on the left. We made this explicit in lecture by proving the following lemma.

Lemma 3.1. *In a leftist heap with n entries, the rank of the root node is at most* $\log_2(n+1)$.

In plain English, this lemma says *leftist heaps have a short right spine,* about $\log n$ in length.

The code for the meld operation (on min-heaps) from lecture is shown here:

```
1  datatype PQ = Leaf | Node of (int × key × PQ × PQ)

2  function rank Leaf = 0
3    | rank (Node(r, _, _, _)) = r

4  function makeLeftistNode (v, L, R) =
5    if (rank(L) < rank(R))
6    then Node(1 + rank(L), v, R, L)
7    else Node(1 + rank(R), v, L, R)

8  function meld (A, B) =
9    case (A, B) of
10     (_, Leaf) ⇒ A
11   | (Leaf, _) ⇒ B
12   | (Node(_, k_a, L_a, R_a), Node(_, k_b, L_b, R_b)) ⇒
13       case Key.compare(k_a, k_b) of
14         LESS ⇒ makeLeftistNode (k_a, L_a, meld(R_a, B))
15       | _ ⇒ makeLeftistNode (k_b, L_b, meld(A, R_b))
```

Q: Given a sequence $S = \langle 1, 2, 3, 4, 5, 6, 7, 8 \rangle$, what is the resulting leftist heap after applying the following code?

```
fun singleton v = Node(1, v, Leaf, Leaf)
Seq.reduce meld Leaf (Seq.map singleton S)
```

A:

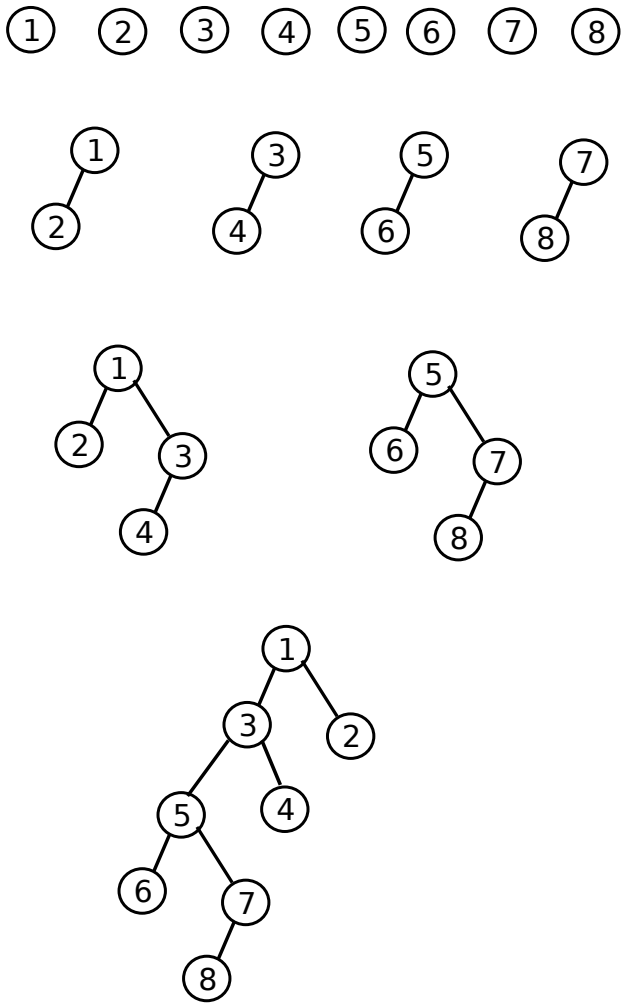

4 Example hashing application: removing duplicates

In class, one of the examples that showed good speedup was removing duplicates. That is, suppose we have a sequence of n elements possibly with some duplicate entries, and we want to remove the duplicates.

$$\langle \text{“}quux\text{”}, \text{“}foo\text{”}, \text{“}bar\text{”}, \text{“}foo\text{”}, \text{“}baz\text{”}, \text{“}bar\text{”} \rangle$$

Q: A brute force solution would look at all pairs, but this is clearly inefficient. Which pairs do we really need to look at? What's the relationship between two keys *having the same hash value* and *being equal*?

A: Being equal $\Rightarrow$ having the same hash index, but having the same hash index $\not\Rightarrow$ being equal.

Already, we have reduced our solution space by only comparing those keys that have the same hash.

Q: If we use separate chaining, how long do we expect a chain to be (at least intuitively)?

A: With the simple uniform hashing assumption, the probability that a key k hashes to a value t is $1/m$; therefore, the length of a chain is simply the number of keys, out of n keys, that hash to the same value t. We know by linearity of expectation that this is n/m.

This leads to the following algorithm: We're going to hash all the values in our sequence and compare only those values within buckets.

Aside (ignore if you wish): Note that if the chains have lengths $B_1, B_2, \ldots, B_m$. Comparing entries within chain i will take $O(B_i^2)$ work; therefore, *even if* $\mathbf{E}\left[B_i\right] = n/m$ may be constant, the expected work $\mathbf{E}\left[B_i^2\right]$ can be much larger. Analyzing this requires looking at the "second moment" of B_i.

On our example above, let's suppose "bar" and "baz" hash to the same index. We'd get

1	2	3
foo, foo	bar, baz, bar	quux

Q: How can we implement this in parallel?

A: Hmm... separate chaining is non-trivial to implement in parallel.

Fortunately, there is a way use open addressing to give $O(n)$ work and $O(\log n)$ span. Q: How might we use parallel open addressing to check for duplicates?

A: Each element attempts to write its value in the hash table. Let's say the initial input consists of n elements and we're using a table of size m. As will be apparent, we'll want m to be larger than n to guarantee efficiency.

Q: What happens after round 1 of insertion?

A: Some keys have been written to the array. If two keys are equal, only one of them got successfully hashed.

Q: What can we change about how we proceed to the next round in order to eliminate duplicates?

A: When collecting the keys to retry inserting in the next round (by comparing each element to the one at its hashed position in the array), exclude duplicates of elements that already got hashed.

Q: How does an element know if it is a duplicate?

A: Instead of just writing elements into the table, we write a pair (v, i) where v is the value and i is its index in the original sequence. If two elements are duplicates, only one of them will get written with its index. The other has the same key as the element written to the hash table but not the same index. That is, element S_j is a duplicate if what got written in $h(S_j)$ is (x, ℓ) where $x = S_j$ but $\ell \neq j$.

Q: Do all unique elements get written to the hash table?

A: No. Some may collide with elements in the hash table, so we repeat the process until there are no elements left. However, unlike hashing with open addressing, we can start the second round with an empty hash table after collecting the elements that successfully hashed, since these are definitely in our non-duplicate sequence.

In summary, using contraction, we proceed in rounds, where each round does the following:

1. For $i = 0, \ldots, |S| - 1$, each element S_i attempts to "write" the value (S_i, i) into location $h(S_i)$ in an array using `injectCond`.

2. We will divide S into unique elements (ACCEPT) and potentially unique elements (RETRY) as follows:

$$\begin{aligned}\mathsf{ACCEPT} &= \langle S_i \in S | T[h(S_i)] = (S_i, i)\rangle \\ \mathsf{RETRY} &= \langle S_i \in S | \#1(T[h(S_i)]) \neq S_i\rangle\end{aligned}$$

3. Recurse on RETRY, appending together all the ACCEPT's.

The ACCEPT elements are those that successfully wrote to the hash table, and the RETRY elements are those that attempted to but did not write to the hash table and are not a duplicate of an element in ACCEPT. *It is crucial that* RETRY *does not contain a duplicate of an element in* ACCEPT. Further, note that implicitly, there is the other group REJECT which is thrown away: this group is made up of the elements that are duplicates of what we already have in ACCEPT.

Why is this algorithm correct? It is easy to see that if a key k is present in S, we'll never throw it away until we include it in ACCEPT, so we only need to argue that we never put two copies of the same key in ACCEPT. Let's consider a round of this algorithm. If S_i and S_j are the same key, only one of them will be in ACCEPT, and furthermore, none of the entries of this key can be in RETRY.

We'll now analyze this algorithm: To bound work, we're interested in knowing how large RETRY is on an input sequence S of length n. Although the worst case might be bad, we're happy with expected-case behaviors. For this, it suffices to compute the probability that an entry S_i is included in RETRY. This happens *only if* the key S_i hashes to the same value as some other key S_j where $S_i \neq S_j$. Therefore, with the simple uniform hashing assumption, we have that for any i,

$$\begin{aligned}\mathbf{Pr}\left[S_i \in \mathsf{RETRY}\right] &\leq \mathbf{Pr}\left[\exists j \text{ s.t. } S_i \neq S_j \wedge h(S_i) = h(S_j)\right] \\ &\leq \sum_{j:S_j \neq S_i} \mathbf{Pr}\left[h(S_i) = h(S_j)\right] \\ &\leq n/m,\end{aligned}$$

where we have upper bounded the probability with a union bound.

If $m = 3n/2$, then $n/m = 2/3$, and by linearity of expectation, we have $|\mathsf{RETRY}| \leq 2n/3$. We have seen this recurrence pattern before; this gives that the total work is expected $O(n)$ because in expectation, this forms a geometrically decreasing sequence. Furthermore, we can show that the number iterations is $O(\log n)$ in expectation.