

Lecture 7 — Collect, Sets, and Tables

Parallel and Sequential Data Structures and Algorithms, 15-210 (Fall 2012)

Lectured by Guy Blelloch — September 18, 2012

Material in this lecture:

- The collect operation.
- Sets and Tables

Challenge: Given sequences of integers A and B of length n , and an initial value X_0 , calculate $X_{i+1} = A_i \cdot X_i + B_i$ for all $0 \leq i < n$. You need to do it in $O(n)$ work and $O(\log n)$ span.

1 Collect

In many applications it is useful to collect all items that share a common key. For example we might want to collect students by course, documents by word, or sales by date. More specifically let's say we had a sequence of pairs each consisting of a student's name and a course they are taking, such as

```
val Data = ⟨("jack sprat", "15-210"),
            ("jack sprat", "15-213"),
            ("mary contrary", "15-210"),
            ("mary contrary", "15-251"),
            ("mary contrary", "15-213"),
            ("peter piper", "15-150"),
            ("peter piper", "15-251"),
            ...⟩
```

and we want to collect all entries by course number so we have a list of everyone taking each course. Collecting values together based on a key is very common in processing databases. In relational database languages such as SQL it is referred to as "Group by". More generally it has many applications and furthermore it is naturally parallel.

We will use the function `collect` for this purpose, and it is part of the sequence library. Its interface is:

$$\text{collect} : (\alpha \times \alpha \rightarrow \text{order}) \rightarrow (\alpha \times \beta) \text{ seq} \rightarrow (\alpha \times \beta \text{ seq}) \text{ seq}$$

The first argument is a function for comparing keys of type α , and must define a total order over the keys. The second argument is a sequence of key-value pairs. The `collect` function collects all values that share the same key together into a sequence. If we wanted to collect the entries of `Data` given above by course number we could do the following:

[†]Lecture notes by Guy E Blelloch, Margaret Reid-Miller, and Kanat Tangwongsan.

```

val collectStrings = collect String.compare
val rosters = collectStrings( $\langle (n,c) : (c,n) \in \text{Data} \rangle$ )

```

This would give something like:

```

val rosters =  $\langle$  ("15-150",  $\langle$  "peter piper", ...  $\rangle$ )
              ("15-210",  $\langle$  "jack sprat", "mary contrary", ...  $\rangle$ )
              ("15-213",  $\langle$  "jack sprat", ...  $\rangle$ )
              ("15-251",  $\langle$  "mary contrary", "peter piper"  $\rangle$ )
              ...  $\rangle$ 

```

We use a map ($\langle (n,c) : (c,n) \in \text{Data} \rangle$) to put the course number in the first position in the tuple since that is the position used to collect on.

Collect can be implemented by sorting the keys based on the given comparison function, and then partitioning the resulting sequence. In particular, the sort will move all the equal keys so they are adjacent. A partition function can then identify the positions where the keys change values, and pull out all pairs between each change. Doing this partitioning can be done relatively easily by filtering out the indices where the value changes. The dominant cost of `collect` is therefore the cost of the sort. Assuming the comparison has complexity bounded above by W_c work and S_c span then the costs of `collect` are $O(W_c n \log n)$ work and $O(S_c \log^2 n)$ span. It is also possible to implement a version of `collect` that runs in linear work using hashing. But hashing would require that a hash function is also supplied and would not return the keys in sorted order. Later in the lecture we discuss tables which also have a `collect` function. However tables are specialized to the key type and therefore neither a comparison nor a hash function need to be passed as arguments.

1.1 Using Collect in Map Reduce

Some of you have probably heard of the map-reduce paradigm first developed by Google for programming certain data intensive parallel tasks. It is now widely used within Google as well as by many others to process large data sets on large clusters of machines—sometimes up to tens of thousands of machines in large data centers. The map-reduce paradigm is often used to analyze various statistical data over very large collections of documents, or over large log files that track the activity at web sites. Outside Google the most widely used implementation is the Apache Hadoop implementation, which has a free license (you can install it at home). The paradigm is different from the `mapReduce` function you might have seen in 15-150 which just involved a map then a reduce. The map-reduce paradigm actually involves a map followed by a collect followed by a bunch of reduces, and therefore might be better called the map-collect-reduces.

The map-reduce paradigm processes a collection of documents based on a map function f_m and a reduce function f_r supplied by the user. The f_m function must take a document as input and generate a sequence of key-value pairs as output. This function is mapped over all the documents. All key-value pairs across all documents are then collected based on the key. Finally the f_r function is applied to each of the keys along with its sequence of associated values to reduce to a single value.

In ML the types for map function f_m and reduce function f_r are the following:

$$\begin{aligned} f_m &: (\text{document} \rightarrow (\text{key} \times \alpha) \text{ seq}) \\ f_r &: (\text{key} \times (\alpha \text{ seq}) \rightarrow \beta) \end{aligned}$$

In most implementations of map-reduce the document is a string (just the contents of a file) and the key is also a string. Typically the α and β types are limited to certain types. Also, in most implementations both the f_m and f_r functions are sequential functions. Parallelism comes about since the f_m function is mapped over the documents in parallel, and the f_r function is mapped over the keys with associated values in parallel.

In ML map reduce can be implemented as follows

```

1 fun mapCollectReduce fm fr docs =
2   let
3     val pairs = flatten ⟨ fm(s) : s ∈ docs ⟩
4     val groups = collect String.compare pairs
5   in
6     ⟨ fr(g) : g ∈ groups ⟩
7   end

```

The function `flatten` simply flattens a nested sequence into a flat sequence, e.g.:

$$\begin{aligned} &\text{flatten} \langle \langle a, b, c \rangle, \langle d, e \rangle \rangle \\ \Rightarrow &\langle a, b, c, d, e \rangle \end{aligned}$$

As an example application of the paradigm, suppose we have a collection of documents, and we want to know how often every word appears across all documents. This can be done with the following f_m and f_r functions.

```

fun fm(doc) = ⟨ (w, 1) : tokens doc ⟩
fun fr(w, s) = (w, reduce + 0 s)

```

Here `tokens` is a function that takes a string and breaks it into tokens by removing whitespace and returning a sequence of strings between whitespace.

Now we can apply `mapCollectReduce` to generate a `countWords` function, and apply this to an example case.

```

val countWords = mapCollectReduce fm fr

countWords ⟨ "this is a document",
             "this is is another document",
             "a last document" ⟩
⇒ ⟨ ("a", 2), ("another", 1), ("document", 3), ("is", 3), ("last", 1), ("this", 2) ⟩

```

2 An Abstract Data Type for Sets

Sets undoubtedly play an important role in mathematics and are often needed in the implementation of various algorithms. Whereas a sequence is an ordered collection, a set is its *unordered* counterpart.

Most programming languages either support sets directly (e.g., Python and Ruby) or have libraries that support them (e.g., in the C++ STL library and Java collections framework). They sometimes have more than one implementation of sets. For example, Java has sets based on hash tables and balanced trees. Unsurprisingly, the set interface in different libraries and languages differ in subtle ways. So, when using one of these interfaces you should always read the documentation carefully.

We now define an abstract data type for sets. The definition follows the mathematical definition of sets from set theory and is purely functional. In particular, when updating a set (e.g. with insert or delete) it returns a new set rather than modifying the old set.

Definition 2.1. For a universe of elements \mathbb{U} (e.g. the integers or strings), the SET abstract data type is a type \mathbb{S} representing the power set of \mathbb{U} (i.e., all subsets of \mathbb{U}) along with the following functions:

<code>empty</code>	: \mathbb{S}	= \emptyset
<code>size(S)</code>	: $\mathbb{S} \rightarrow \mathbb{Z}_{\geq 0}$	= $ S $
<code>singleton(e)</code>	: $\mathbb{U} \rightarrow \mathbb{S}$	= $\{e\}$
<code>filter(f, S)</code>	: $((\mathbb{U} \rightarrow \{\text{T}, \text{F}\}) \times \mathbb{S}) \rightarrow \mathbb{S}$	= $\{s \in S \mid f(s)\}$
<code>find(S, e)</code>	: $\mathbb{S} \times \mathbb{U} \rightarrow \{\text{T}, \text{F}\}$	= $ \{s \in S \mid s = e\} = 1$
<code>insert(S, e)</code>	: $\mathbb{S} \times \mathbb{U} \rightarrow \mathbb{S}$	= $S \cup \{e\}$
<code>delete(S, e)</code>	: $\mathbb{S} \times \mathbb{U} \rightarrow \mathbb{S}$	= $S \setminus \{e\}$
<code>intersection(S_1, S_2)</code>	: $\mathbb{S} \times \mathbb{S} \rightarrow \mathbb{S}$	= $S_1 \cap S_2$
<code>union(S_1, S_2)</code>	: $\mathbb{S} \times \mathbb{S} \rightarrow \mathbb{S}$	= $S_1 \cup S_2$
<code>difference(S_1, S_2)</code>	: $\mathbb{S} \times \mathbb{S} \rightarrow \mathbb{S}$	= $S_1 \setminus S_2$

where $\mathbb{Z}_{\geq 0}$ are the non-negative integers.

We write this definition to be generic and not specific to Standard ML. In our library, the type \mathbb{S} is called `set` and the type \mathbb{U} is called `key`, the arguments are not necessarily in the same order, and some of the functions are curried. For example, the interface for `find` is `find : set → key → bool`. Please refer to the documents for details. In the pseudocode, we will give in class and in the notes we will use standard set notation as in the right hand column of the table above.

A Note about map. You may notice that the interface does not contain a `map` function. If we try to generalize the notion of `map` from sequences, a `map` function does not make sense in the context of a set: if we interpret `map` to take in a collection, apply some function to each element and return a collection of the same structure. Consider a function that always returns 0. Mapping this over a set would return all zeros, which would then be collapsed into a singleton set, containing exactly 0. Therefore, such a `map` would allow reducing the set of arbitrary size to a singleton, which doesn't match the `map` paradigm (which traditionally preserves the structure and size).

2.1 Cost Model

So far, we have laid out a semantic interface, but before we can put it to use, we need to worry about the cost specification. The most common efficient ways to implement sets are either using hashing or balanced trees. They have various tradeoffs in cost. For simplicity, we'll consider a cost model based on a balanced-tree implementation. We will cover how to implement these set operations when we talk about balanced trees later in the course. For now, a good intuition to have is that we use a comparison function to keep the elements in sorted order in a balanced tree. Since this requires comparisons inside the various set operations, the cost of these comparisons affects the work and span. For this, we'll assume that compare has C_w work and C_s span.

We have the following cost specification:

	<i>Work</i>	<i>Span</i>
size(S)	$O(1)$	$O(1)$
singleton(e)		
filter(f, S)	$O\left(\sum_{e \in S} W(f(e))\right)$	$O\left(\log S + \max_{e \in S} S(f(e))\right)$
find(S, e)		
insert(S, e)	$O(C_w \cdot \log S)$	$O(C_s \cdot \log S)$
delete(S, e)		
intersection(S_1, S_2)		
union(S_1, S_2)	$O\left(C_w \cdot m \cdot \log\left(1 + \frac{n}{m}\right)\right)$	$O\left(C_s \cdot \log(n + m)\right)$
difference(S_1, S_2)		

where $n = \max(|S_1|, |S_2|)$ and $m = \min(|S_1|, |S_2|)$.

The work bounds for `intersection`, `union`, and `difference` deserve further discussion—at a glance they might seem a bit funky. These turn out to be both the theoretical upper and lower bounds for any comparison-based implementation of sets. We will get to this later.

Notice that when the two sets have the same length ($n = m$), the work is simply

$$O(C_w \cdot m \cdot \log(1 + 1)) = O(C_w \cdot n).$$

This should not be surprising, because it corresponds to the cost of merging two approximately equal length sequences (effectively what these operations have to do).

Moreover, you should also observe that in the case that one of the sets is a singleton, then the work is $O(\log n)$.

Buy in Bulk and Save. On inspection, the functions `intersection`, `union`, and `difference` are simply the “parallel” counterparts of the functions `find`, `insert`, and `delete`, to wit:

- `intersection` — search for multiple elements instead of one.
- `union` — insert multiple elements.
- `difference` — delete multiple elements.

In fact, it is easy to implement `find`, `insert`, and `delete` in terms of the others.

```
find(S,e) = size(intersection(S, singleton(e))) = 1
insert(S,e) = union(S, singleton(e))
delete(S,e) = difference(S, singleton(e))
```

Since `intersection`, `union`, and `difference` can operate on multiple elements they are well suited for parallelism, while `find`, `insert`, and `delete` have no parallelism. Consequently, in designing parallel algorithms it is good to think about how to use `intersection`, `union`, and `difference` instead of `find`, `insert`, and `delete` if possible. For example, one way to convert a sequence to a set would be to insert the elements one by one, which can be coded as

```
fun fromSeq(S) = Seq.iter (fn (S',e) => Set.insert e S') Set.empty S
```

However, the above is sequential. To do it in parallel we could instead do

```
fun fromSeq(S) = Seq.reduce Set.union Set.empty (Seq.map Set.singleton S)
```

Exercise 1. What is the work and span of the first version of `fromSeq`.

Exercise 2. Show that on a sequence of length n the second version of `fromSeq` does $O(C_w n \log n)$ work and $O(\log^2 n)$ span.

3 Tables: Associating Each Element With A Value

Suppose we want to extend sets so that each element is associated with a payload. A table is an abstract data type that stores for each key data associated with it. The table ADT supplies operations for finding the value associated with a key, for inserting new key-value pairs, and for deleting keys and their associated value. Tables are also called dictionaries, associative arrays, maps, mappings, and functions (in set theory). Given our focus on parallelism, the interface we will discuss also supplies “parallel” operations that allow the user to insert multiple key-value pairs, to delete multiple keys, and to find the values associated with multiple keys.

In this class, the notation we are going to be using is

$$\{(k_1 \mapsto v_1), (k_2 \mapsto v_2), \dots, (k_n \mapsto v_n)\},$$

where we have *keys* and *values*—and each key k_i is associated with the value v_i . Mathematically, a table is simply a set of pairs and can therefore be written as a set of ordered pairs $\{(k_1, v_1), (k_2, v_2), \dots, (k_n, v_n)\}$. Our notation choice is largely to better identify when tables are being used.

As with sets, tables are commonly used in many applications. Most languages have tables either built in (e.g. dictionaries in Python, Perl, and Ruby), or have libraries to support them (e.g. `map` in the C++ STL library and the Java collections framework). We note that the interfaces for these languages and libraries have common features but typically differ in some important ways, so be

warned. Most do not support the “parallel” operations we discuss. Again, here we will define tables mathematically in terms of set theory before committing to a particular language.

Formally, a table is set of key-value pairs where each key appears only once in the set. Such sets are called *functions* in set theory since they map each key to a single value. We will avoid this terminology so that we don’t confuse it with functions in a programming language. However, note that the `(find T)` in the interface is precisely the “function” defined by the table `T`. In fact it is a *partial function* since the table might not contain all keys and therefore the function might not be defined on all inputs. Here is the definition of a table.

Definition 3.1. For a universe of keys \mathbb{K} , and a universe of values \mathbb{V} , the TABLE abstract data type is a type \mathbb{T} representing the power set of $\mathbb{K} \times \mathbb{V}$ restricted so that each key appears at most once (i.e., any set of key-value pairs where a key appears just once) along with the following functions:

$$\begin{array}{lll}
\text{empty} & : \mathbb{T} & = \emptyset \\
\text{size}(T) & : \mathbb{T} \rightarrow \mathbb{Z}_{\geq 0} & = |T| \\
\text{singleton}(k, v) & : \mathbb{K} \times \mathbb{V} \rightarrow \mathbb{T} & = \{(k, v)\} \\
\text{filter}(f, T) & : ((\mathbb{V} \rightarrow \{\text{F}, \text{T}\}) \times \mathbb{T}) \rightarrow \mathbb{T} & = \{(k, v) \in T \mid f(v)\} \\
\text{map}(f, T) & : ((\mathbb{K} \times \mathbb{V} \rightarrow \mathbb{V}) \times \mathbb{T}) \rightarrow \mathbb{T} & = \{(k, f(k, v)) \mid (k, v) \in T\} \\
\text{insert}(f, T, (k, v)) & : (\mathbb{V} \times \mathbb{V} \rightarrow \mathbb{V}) \times \mathbb{T} \times (\mathbb{K} \times \mathbb{V}) \rightarrow \mathbb{T} & = \\
& \quad \forall k \in \mathbb{K}, \begin{cases} (k, f(v, v')) & (k, v') \in T \\ (k, v) & (k, v') \notin T \end{cases} \\
\text{delete}(T, k) & : \mathbb{T} \times \mathbb{K} \rightarrow \mathbb{T} & = \{(k', v) \in T \mid k \neq k'\} \\
\text{find}(T, k) & : \mathbb{T} \times \mathbb{K} \rightarrow (\mathbb{V} \cup \perp) & = \begin{cases} v & (k, v) \in T \\ \perp & \text{otherwise} \end{cases} \\
\text{merge}(f, T_1, T_2) & : (\mathbb{V} \times \mathbb{V} \rightarrow \mathbb{V}) \times \mathbb{T} \times \mathbb{T} \rightarrow \mathbb{T} & = \\
& \quad \forall k \in \mathbb{K}, \begin{cases} (k, f(v_1, v_2)) & (k, v_1) \in T_1 \wedge (k, v_2) \in T_2 \\ (k, v_1) & (k, v_1) \in T_1 \\ (k, v_2) & (k, v_2) \in T_2 \end{cases} \\
\text{extract}(T, S) & : \mathbb{T} \times \mathbb{S} \rightarrow \mathbb{T} & = \{(k, v) \in T \mid k \in S\} \\
\text{erase}(T, S) & : \mathbb{T} \times \mathbb{S} \rightarrow \mathbb{T} & = \{(k, v) \in T \mid k \notin S\}
\end{array}$$

where \mathbb{S} is the power set of \mathbb{K} (i.e., any set of keys) and $\mathbb{Z}_{\geq 0}$ are the non-negative integers.

Distinct from sets, the `find` function does not return a Boolean, but instead it returns the value associated with the key `k`. As it may not find the key in the table, its result may be bottom (\perp). For this reason, in the Table library, the interface for `find` is `find : 'a table → key → 'a option`, where `'a` is the type of the values.

Unlike sets, when we insert an element, we can’t simply ignore that element if it is already present—their values might be different. For this reason, the `insert` function takes a function `f : V × V → V` as an argument. The purpose of `f` is to specify what to do if the key being inserted already exists in the table; `f` is applied to the two values. This function might simply return either its first or second argument, or it can be used, for example, to add the new value to the old one. The parallel counterpart of `find` is the `merge` function, which takes a similar function since it also has to consider the case that an element appears in both tables.

We also introduce new pseudocode notation for `map` and `filter` on tables:

$$\{(k \mapsto f(v)) \mid (k \mapsto v) \in T\}$$

is equivalent to $\text{map}(f, T)$ and

$$\{(k \mapsto v) \in T \mid p(v)\}$$

is equivalent to $\text{filter}(p, T)$.

The costs of the table operations are very similar to sets.

	<i>Work</i>	<i>Span</i>
$\text{size}(T)$ $\text{singleton}(k, v)$	$O(1)$	$O(1)$
$\text{filter}(f, T)$	$O\left(\sum_{(k,v) \in T} W(f(v))\right)$	$O\left(\log T + \max_{(k,v) \in T} S(f(v))\right)$
$\text{map}(f, T)$	$O\left(\sum_{(k,v) \in T} W(f(k, v))\right)$	$O\left(\max_{(k,v) \in T} S(f(k, v))\right)$
$\text{find}(S, k)$ $\text{insert}(T, (k, v))$ $\text{delete}(T, k)$	$O(C_w \log T)$	$O(C_s \log T)$
$\text{extract}(T_1, T_2)$ $\text{merge}(T_1, T_2)$ $\text{erase}(T_1, T_2)$	$O(C_w m \log(1 + \frac{n}{m}))$	$O(C_s \log(n + m))$

where $n = \max(|T_1|, |T_2|)$ and $m = \min(|T_1|, |T_2|)$.

As with sets there is a symmetry between the three operations `extract`, `merge`, and `erase`, and the three operations `find`, `insert`, and `delete`, respectively, where the prior three are effectively “parallel” versions of the earlier three. The `extract` operation can be used to find a set of values in a table, returning just the table entries corresponding to elements in the set. The `merge` operation can add multiple values to a table in parallel by merging two tables. The `erase` operation can delete multiple values from a table in parallel.

We note that, in the SML Table library we supply, the functions are polymorphic (accept any type) over the values but not the keys. In particular the signature starts as:

```
signature TABLE =
sig
  type 'a table
  type 'a t = 'a table
  structure Key : EQKEY
  type key = Key.t
  structure Seq : SEQUENCE
  type a seq = 'a Seq.seq
  type set = unit table
  ...
  val find : 'a table -> key -> 'a option
  ...
end
```

The `'a` in `'a table` refers to the type of the value. The key type is fixed to be `key`. Therefore there

are separate table structures for different keys (e.g. `IntTable`, `StringTable`). The reason to do this is because all the operations depend on the key type since keys need to be compared for a tree implementation, or hashed for a hash table implementation. Also note that the signature defines `set` to be a `unit table`. Indeed a set is just a special case of a table where there are no values.

In the SML Table library, we supply a `collect` operation that takes a sequence of key-value pairs and produces a table that maps every key in S to all the values associated with it in S , gathering all the values with the same key together in a sequence. This is equivalent to using a `sequence collect` followed by a `Table.fromSeq`. Alternatively, it can be implemented as

```
1 fun collect(S) =  
2   let  
3     val S' = {k ↦ ⟨v⟩ : (k,v) ∈ S}  
4   in  
5     Seq.reduce (Table.merge Seq.append) {} S'  
6   end
```

Exercise 3. *Figure out what this code does.*