

Partition Algorithms

Ralph Freese

March 4, 1997

These are algorithms for partitions on the set $\{0, 1, \dots, n-1\}$. We represent partitions abstractly as forests, i.e., a collection of trees, one tree for each block of the partition. We only need the parent information about the tree so we represent the partition as a vector V with $V[i]$ the parent of i unless i has no parent (and so is a root), in which case $V[i]$ is negative the size of the block with i . In this scheme the least partition would be represented by the vector $\langle -1, -1, \dots, -1 \rangle$ and the greatest partition could be represented in many ways including the vector $\langle -n, 0, \dots, 0 \rangle$. [2] contains an elementary discussion of this type of representation of partitions.

We say that a vector representing a partition is in *normal form* if the root of each block is the least element of that block and the parent of each nonroot is its root. This form is unique, i.e., two vectors represent the same partition if and only if they have the same normal form. The examples above are in normal form. Algorithm 1 gives a simple recursive procedure for finding the root of any element i . Note that i and j are in the same block if and only if $root(i) = root(j)$.

```
1 procedure root( $i, V$ )  
2    $j \leftarrow V[i]$   
3   if  $j < 0$  then return( $i$ )  
4   else return(root( $j$ )) endif  
5 endprocedure
```

Algorithm 1: Finding the root

The running time for *root* is proportional to the depth of i in its tree, so we would like to keep the depth of the forest small. Algorithm 2 finds the root and at the same time modifies V so that the parent of i is its root without increasing the order of magnitude of the running time.

In many applications you want to build up a partition by starting with the least partition and repeatedly join blocks together. Algorithm 3 does this.

Note that Algorithm 3 always joins the smaller block onto the larger block. This assures us that the resulting partition will have depth at most $\log_2 n$ as the next theorem shows.

```

1 procedure root(i, V)
2   j ← V[i]
3   if j < 0 then return(i)
4   else V[i] ← root(j); return(V[i]) endif
5 endprocedure

```

Algorithm 2: Finding the root and compressing *V*

```

1 procedure join-blocks(i, j, V)
2   ri ← root(i, V); rj ← root(j, V);
3   if ri ≠ rj then
4     si ← −V[ri]; sj ← −V[rj]
5     if si < sj then
6       V[i] ← rj; V[j] ← −(si + sj)
7     else
8       V[j] ← ri; V[i] ← −(si + sj)
9     endif
10    endif
11    return(V)
12 endprocedure

```

Algorithm 3: Join two blocks together

Theorem 1 *If Algorithm 3 is applied any number of times starting with the least partition, the depth of the resulting partition will never exceed $\log_2 n$.*

Proof: Let *i* be a fixed node. Note an application of *join-blocks* increases the depth of *i* by at most 1 and, if this occurs, the size of the block with *i* is at least doubled. Thus the depth of *i* can be increased (by 1) from its original value of 0 at most $\log_2 n$ times.

This result shows that the time required to run the *join-blocks*-procedure *m* times is $O(m \log_2 n)$. In [1] Tarjan has shown that, if we use the *root* operation given in Algorithm 2, the time required is $O(m\alpha(m))$, where α is the pseudo-inverse of the Ackermann function. The Ackermann function is extremely fast growing and so α grows very slowly; in fact, $\alpha(m) \leq 4$ unless *m* is at least

$$2^{2^{2^{\cdots^2}}}$$

with 65536 2's.

By Theorem 1 we may assume that all our (representations of) partitions have depth at most $\log_2 n$. The *rank* of a partition (in the partition lattice Π_n of an *n* element set) is $n - k$, where *k* is the number of blocks. The join of two partitions *U* and *V* can be found by executing *join-blocks*(*i*, *U*[*i*], *V*) for each *i* which is not a root of *U*. This can be done in time $O(\text{rank}(\text{U}) \log_2 n)$

and so in time $O(n \log_2 n)$. (Actually, such an algorithm should make a copy of V so the original V is not modified.)

It is relatively easy to write $O(n \log_2 n)$ time procedures for putting V into normal form and for testing if $V \leq U$ in Π_n . Finding an $O(n \log_2 n)$ time algorithm for the meet of two partitions is a little more difficult. Algorithm 4 does this. In this algorithm, HT is a hash table. (In place of a hash table, one could use a balanced tree or some other data structure described in texts on algorithms and data structures.)

```

1 procedure meet( $V_1, V_2$ )
2    $n \leftarrow \text{size}(V_1)$ 
3   for  $i \in \mathbb{Z}$  with  $0 \leq i < n$  do
4      $r_1 \leftarrow \text{root}(i, V_1); r_2 \leftarrow \text{root}(i, V_2)$ 
5     if HT[ $r_1, r_2$ ] is defined then
6        $r \leftarrow \text{HT}[r_1, r_2]$ 
7        $V[r] \leftarrow V[r] - 1$ 
8        $V[i] \leftarrow r$ 
9     else
10      HT[ $r_1, r_2$ ]  $\leftarrow i$ 
11       $V[i] \leftarrow -1$ 
12    endif
13  endfor
14  return( $V$ )
15 endprocedure

```

Algorithm 4: Meet of two partitions

References

- [1] R. E. Tarjan, *Efficiency of a good but not linear set union algorithm*, J. Assoc. Comput. Mach. **22** (1975), 215–225.
- [2] M. A. Weiss, *Data Structures and Algorithm Analysis*, Benjamin Cummings, Redwood City, California, 1994, Second Edition.