



Chapter 21 - Data Structures for Disjoint Sets

Introduction to Algorithms, Third Edition

by Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest and Clifford Stein

The MIT Press © 2009 Citation

Recommend?

◀ Previous

Next ▶

21.1 Disjoint-Set Operations

A **disjoint-set data structure** maintains a collection $\mathcal{S} = \{S_1, S_2, \dots, S_k\}$ of disjoint dynamic sets. We identify each set by a **representative**, which is some member of the set. In some applications, it doesn't matter which member is used as the representative; we care only that if we ask for the representative of a dynamic set twice without modifying the set between the requests, we get the same answer both times. Other applications may require a prespecified rule for choosing the representative, such as choosing the smallest member in the set (assuming, of course, that the elements can be ordered).

As in the other dynamic-set implementations we have studied, we represent each element of a set by an object. Letting x denote an object, we wish to support the following operations:

MAKE-SET(x) creates a new set whose only member (and thus representative) is x . Since the sets are disjoint, we require that x not already be in some other set.

UNION(x, y) unites the dynamic sets that contain x and y , say S_x and S_y , into a new set that is the union of these two sets. We assume that the two sets are disjoint prior to the operation. The representative of the resulting set is any member of $S_x \cup S_y$, although many implementations of **UNION** specifically choose the representative of either S_x or S_y as the new representative. Since we require the sets in the collection to be disjoint, conceptually we destroy sets S_x and S_y , removing them from the collection \mathcal{S} . In practice, we often absorb the elements of one of the sets into the other set.

FIND-SET(x) returns a pointer to the representative of the (unique) set containing x .

Throughout this chapter, we shall analyze the running times of disjoint-set data structures in terms of two parameters: n , the number of **MAKE-SET** operations, and m , the total number of **MAKE-SET**, **UNION**, and **FIND-SET** operations. Since the sets are disjoint, each **UNION** operation reduces the number of sets by one. After $n - 1$ **UNION** operations, therefore, only one set remains. The number of **UNION** operations is thus at most $n - 1$. Note also that since the **MAKE-SET** operations are included in the total number of operations m , we have $m \geq n$. We assume that the n **MAKE-SET** operations are the first n operations performed.

An Application of Disjoint-Set Data Structures

One of the many applications of disjoint-set data structures arises in determining the connected components of an undirected graph (see [Section B.4](#)). [Figure 21.1\(a\)](#), for example, shows a graph with four connected components.