

In an *amortized analysis*, we average the time required to perform a sequence of data-structure operations over all the operations performed. With amortized analysis, we can show that the average cost of an operation is small, if we average over a sequence of operations, even though a single operation within the sequence might be expensive. Amortized analysis differs from average-case analysis in that probability is not involved; an amortized analysis guarantees the *average performance of each operation in the worst case*.

The first three sections of this chapter cover the three most common techniques used in amortized analysis. Section 17.1 starts with aggregate analysis, in which we determine an upper bound  $T(n)$  on the total cost of a sequence of  $n$  operations. The average cost per operation is then  $T(n)/n$ . We take the average cost as the amortized cost of each operation, so that all operations have the same amortized cost.

Section 17.2 covers the accounting method, in which we determine an amortized cost of each operation. When there is more than one type of operation, each type of operation may have a different amortized cost. The accounting method overcharges some operations early in the sequence, storing the overcharge as “prepaid credit” on specific objects in the data structure. Later in the sequence, the credit pays for operations that are charged less than they actually cost.

Section 17.3 discusses the potential method, which is like the accounting method in that we determine the amortized cost of each operation and may overcharge operations early on to compensate for undercharges later. The potential method maintains the credit as the “potential energy” of the data structure as a whole instead of associating the credit with individual objects within the data structure.

We shall use two examples to examine these three methods. One is a stack with the additional operation MULTIPOP, which pops several objects at once. The other is a binary counter that counts up from 0 by means of the single operation INCREMENT.

While reading this chapter, bear in mind that the charges assigned during an amortized analysis are for analysis purposes only. They need not—and should not—appear in the code. If, for example, we assign a credit to an object  $x$  when using the accounting method, we have no need to assign an appropriate amount to some attribute, such as  $x.credit$ , in the code.

When we perform an amortized analysis, we often gain insight into a particular data structure, and this insight can help us optimize the design. In Section 17.4, for example, we shall use the potential method to analyze a dynamically expanding and contracting table.

---

## 17.1 Aggregate analysis

In *aggregate analysis*, we show that for all  $n$ , a sequence of  $n$  operations takes *worst-case* time  $T(n)$  in total. In the worst case, the average cost, or *amortized cost*, per operation is therefore  $T(n)/n$ . Note that this amortized cost applies to each operation, even when there are several types of operations in the sequence. The other two methods we shall study in this chapter, the accounting method and the potential method, may assign different amortized costs to different types of operations.

### Stack operations


In our first example of aggregate analysis, we analyze stacks that have been augmented with a new operation. Section 10.1 presented the two fundamental stack operations, each of which takes  $O(1)$  time:

$PUSH(S, x)$  pushes object  $x$  onto stack  $S$ .

$POP(S)$  pops the top of stack  $S$  and returns the popped object. Calling  $POP$  on an empty stack generates an error.

Since each of these operations runs in  $O(1)$  time, let us consider the cost of each to be 1. The total cost of a sequence of  $n$   $PUSH$  and  $POP$  operations is therefore  $n$ , and the actual running time for  $n$  operations is therefore  $\Theta(n)$ .

Now we add the stack operation  $MULTIPOP(S, k)$ , which removes the  $k$  top objects of stack  $S$ , popping the entire stack if the stack contains fewer than  $k$  objects. Of course, we assume that  $k$  is positive; otherwise the  $MULTIPOP$  operation leaves the stack unchanged. In the following pseudocode, the operation  $STACK-EMPTY$  returns  $TRUE$  if there are no objects currently on the stack, and  $FALSE$  otherwise.



**Figure 17.1** The action of `MULTIPOP` on a stack  $S$ , shown initially in (a). The top 4 objects are popped by `MULTIPOP( $S$ , 4)`, whose result is shown in (b). The next operation is `MULTIPOP( $S$ , 7)`, which empties the stack—shown in (c)—since there were fewer than 7 objects remaining.

`MULTIPOP( $S$ ,  $k$ )`

```

1  while not STACK-EMPTY( $S$ ) and  $k > 0$ 
2      POP( $S$ )
3       $k = k - 1$ 

```

Figure 17.1 shows an example of `MULTIPOP`.

What is the running time of `MULTIPOP( $S$ ,  $k$ )` on a stack of  $s$  objects? The actual running time is linear in the number of `POP` operations actually executed, and thus we can analyze `MULTIPOP` in terms of the abstract costs of 1 each for `PUSH` and `POP`. The number of iterations of the **while** loop is the number  $\min(s, k)$  of objects popped off the stack. Each iteration of the loop makes one call to `POP` in line 2. Thus, the total cost of `MULTIPOP` is  $\min(s, k)$ , and the actual running time is a linear function of this cost.

Let us analyze a sequence of  $n$  `PUSH`, `POP`, and `MULTIPOP` operations on an initially empty stack. The worst-case cost of a `MULTIPOP` operation in the sequence is  $O(n)$ , since the stack size is at most  $n$ . The worst-case time of any stack operation is therefore  $O(n)$ , and hence a sequence of  $n$  operations costs  $O(n^2)$ , since we may have  $O(n)$  `MULTIPOP` operations costing  $O(n)$  each. Although this analysis is correct, the  $O(n^2)$  result, which we obtained by considering the worst-case cost of each operation individually, is not tight.

Using aggregate analysis, we can obtain a better upper bound that considers the entire sequence of  $n$  operations. In fact, although a single `MULTIPOP` operation can be expensive, any sequence of  $n$  `PUSH`, `POP`, and `MULTIPOP` operations on an initially empty stack can cost at most  $O(n)$ . Why? We can pop each object from the stack at most once for each time we have pushed it onto the stack. Therefore, the number of times that `POP` can be called on a nonempty stack, including calls within `MULTIPOP`, is at most the number of `PUSH` operations, which is at most  $n$ . For any value of  $n$ , any sequence of  $n$  `PUSH`, `POP`, and `MULTIPOP` operations takes a total of  $O(n)$  time. The average cost of an operation is  $O(n)/n = O(1)$ . In aggregate

analysis, we assign the amortized cost of each operation to be the average cost. In this example, therefore, all three stack operations have an amortized cost of  $O(1)$ .

We emphasize again that although we have just shown that the average cost, and hence the running time, of a stack operation is  $O(1)$ , we did not use probabilistic reasoning. We actually showed a *worst-case* bound of  $O(n)$  on a sequence of  $n$  operations. Dividing this total cost by  $n$  yielded the average cost per operation, or the amortized cost.

### Incrementing a binary counter

As another example of aggregate analysis, consider the problem of implementing a  $k$ -bit binary counter that counts upward from 0. We use an array  $A[0..k-1]$  of bits, where  $A.length = k$ , as the counter. A binary number  $x$  that is stored in the counter has its lowest-order bit in  $A[0]$  and its highest-order bit in  $A[k-1]$ , so that  $x = \sum_{i=0}^{k-1} A[i] \cdot 2^i$ . Initially,  $x = 0$ , and thus  $A[i] = 0$  for  $i = 0, 1, \dots, k-1$ . To add 1 (modulo  $2^k$ ) to the value in the counter, we use the following procedure.

```

INCREMENT( $A$ )
1   $i = 0$ 
2  while  $i < A.length$  and  $A[i] == 1$ 
3       $A[i] = 0$ 
4       $i = i + 1$ 
5  if  $i < A.length$ 
6       $A[i] = 1$ 

```

Figure 17.2 shows what happens to a binary counter as we increment it 16 times, starting with the initial value 0 and ending with the value 16. At the start of each iteration of the **while** loop in lines 2–4, we wish to add a 1 into position  $i$ . If  $A[i] = 1$ , then adding 1 flips the bit to 0 in position  $i$  and yields a carry of 1, to be added into position  $i + 1$  on the next iteration of the loop. Otherwise, the loop ends, and then, if  $i < k$ , we know that  $A[i] = 0$ , so that line 6 adds a 1 into position  $i$ , flipping the 0 to a 1. The cost of each INCREMENT operation is linear in the number of bits flipped.

As with the stack example, a cursory analysis yields a bound that is correct but not tight. A single execution of INCREMENT takes time  $\Theta(k)$  in the worst case, in which array  $A$  contains all 1s. Thus, a sequence of  $n$  INCREMENT operations on an initially zero counter takes time  $O(nk)$  in the worst case.

We can tighten our analysis to yield a worst-case cost of  $O(n)$  for a sequence of  $n$  INCREMENT operations by observing that not all bits flip each time INCREMENT is called. As Figure 17.2 shows,  $A[0]$  does flip each time INCREMENT is called. The next bit up,  $A[1]$ , flips only every other time: a sequence of  $n$  INCREMENT

Counter value	A[7]	A[6]	A[5]	A[4]	A[3]	A[2]	A[1]	A[0]	Total cost
0	0	0	0	0	0	0	0	0	0
1	0	0	0	0	0	0	0	1	1
2	0	0	0	0	0	0	1	0	3
3	0	0	0	0	0	0	1	1	4
4	0	0	0	0	0	1	0	0	7
5	0	0	0	0	0	1	0	1	8
6	0	0	0	0	0	1	1	0	10
7	0	0	0	0	0	1	1	1	11
8	0	0	0	0	1	0	0	0	15
9	0	0	0	0	1	0	0	1	16
10	0	0	0	0	1	0	1	0	18
11	0	0	0	0	1	0	1	1	19
12	0	0	0	0	1	1	0	0	22
13	0	0	0	0	1	1	0	1	23
14	0	0	0	0	1	1	1	0	25
15	0	0	0	0	1	1	1	1	26
16	0	0	0	1	0	0	0	0	31

**Figure 17.2** An 8-bit binary counter as its value goes from 0 to 16 by a sequence of 16 INCREMENT operations. Bits that flip to achieve the next value are shaded. The running cost for flipping bits is shown at the right. Notice that the total cost is always less than twice the total number of INCREMENT operations.

operations on an initially zero counter causes  $A[1]$  to flip  $\lfloor n/2 \rfloor$  times. Similarly, bit  $A[2]$  flips only every fourth time, or  $\lfloor n/4 \rfloor$  times in a sequence of  $n$  INCREMENT operations. In general, for  $i = 0, 1, \dots, k-1$ , bit  $A[i]$  flips  $\lfloor n/2^i \rfloor$  times in a sequence of  $n$  INCREMENT operations on an initially zero counter. For  $i \geq k$ , bit  $A[i]$  does not exist, and so it cannot flip. The total number of flips in the sequence is thus

$$\begin{aligned} \sum_{i=0}^{k-1} \left\lfloor \frac{n}{2^i} \right\rfloor &< n \sum_{i=0}^{\infty} \frac{1}{2^i} \\ &= 2n, \end{aligned}$$

by equation (A.6). The worst-case time for a sequence of  $n$  INCREMENT operations on an initially zero counter is therefore  $O(n)$ . The average cost of each operation, and therefore the amortized cost per operation, is  $O(n)/n = O(1)$ .

## **Reference :**

**Text Book**-T. H. Cormen, C. E. Leiserson , R. L. Rivest and C. Stein : Introduction to Algorithms, Third Edition ,The MIT Press Cambridge, Massachusetts London, England .