

15-819: Foundations of Quantitative Program Analysis
Lecture 1: Traditional Complexity Analysis

Jan Hoffmann

September 3, 2019

1 Introduction

In the first part of this course, we are interested in determining the resource usage of a program as an easily-understood function of the inputs. Such a resource analysis is a one of the most classic problems of computer science. It has been for instance studied by Donald Knuth in his book *The Art of Computer Programming*. Today, resource analysis is often performed without a formally defined cost model (or cost semantics) and without aid by computers. In particular, bounds are derived without automation or mechanical correctness checks. Moreover, the derived bounds often hold asymptotically only, that is, for large inputs and with other multiplicative constants. That such an analysis is meaningful at all in practice seems to be largely a coincident since it is applied to programs that have a uniform behavior. However, from a mathematical point of view, asymptotic analysis seems rather pointless since it applies only to *large* inputs where large could mean larger than the number of atoms in the universe. In many applications, we are interested in comparing algorithms with the same algorithmic behavior or within a given range of inputs. In these cases, asymptotic bounds are not helpful. Moreover, we have to precisely determine the constant factors anyway in a formal asymptotic analysis.

For these reasons, we will calculate and prove *non-asymptotic* bounds with concrete constant factors in this course. Interestingly, Donald Knuth does rarely use Big-O notation in *The Art of Computer Programming*. He formulates algorithms in a machine language for the MIX architecture and plays close attention to concrete and best possible values of constants in the analyses. We will operate at a much higher-level functional programming language and with a formally defined *cost semantics*. Our focus is then on automatically deriving bounds that come with efficiently checkable certificates. However, before we come to that, we want to remind ourselves how resource analysis is performed without these techniques.

In this lecture, we review *recurrence relations* and *asymptotic analysis*, which are some of the most popular ways to perform resource analysis. While recurrence relations are a powerful and flexible mechanism, they also have shortcomings. For example, an analysis with recurrence relations becomes cumbersome in the presence of more complex data structures (lists of trees, etc.) and when analyzing sequences of operations (or function calls), which are both present in many programs. For analyzing the worst-case resource usage of a sequence of operations, there is this the powerful idea of *amortized analysis*, will stay on our minds throughout the semester.

2 Recurrence Relations

A popular (and somewhat systematic) way of performing resource analysis is to use recurrence relations. Most commonly, recurrence relations are used for manual complexity analysis but we will also look at them from the perspective of automatic and mechanized resource analysis.

First, we should define what a recurrence relation is. The textbook *Introduction to Algorithms* contains the following definition.

A recurrence is an equation or inequality that describes a function in terms of its value on smaller inputs.

One could argue that this definition is a bit too broad because an input does not necessarily have a unique order to which *smaller* can refer. The aforementioned book contains moreover only recurrence relations where the inputs are natural numbers. The definition from Carnegie Mellons's algorithms course (15-451/651) reads as follows.

A recurrence relation is a description of the running time on an input of size n as a function of n and the running time on inputs of smaller sizes.

Here, we have the mention of a size but it is not clear if multiple arguments are permitted and if they all need to be abstracted as sizes. The exact definition of the term recurrence relation is probably not that important. For the purpose of this course, we will use the following.

Definition. A recurrence relation is a recursive function T of type $\mathbb{N}^k \rightarrow \mathbb{N}$.

The idea is that the arguments of $T : \mathbb{N}^k \rightarrow \mathbb{N}$ represent different sizes or measures that related in some way to the inputs of the program that is the subject of the resource analysis. Resource analysis using recurrence relations proceeds in two steps. First, we derive a recurrence relation that describes the resource usage of the program. Second, we find a closed (i.e., non-recursive) form of the recurrence. When approaching the problem from a mechanization and automation perspective, we realize that both steps are non-trivial.

To start with, it is sometimes not straightforward to find the right size parameters for the analysis but an automatic approach could settle for the most common ones like the size or height of a data structure. The derivation of the recurrence is also not trivial; particularly if we also want a proof that it correctly describes the resource usage of the program. One of the difficulties is that it is often not immediately obvious what the size of the data structures in the program is as a function of the inputs. For example, data could be returned by a complex auxiliary function before being passed to a recursive call. Finally, solving (i.e., finding closed forms for) recurrences is a difficult problem and undecidable in general. If we are looking for exact (non-asymptotic) solutions, we can only rely on a few techniques from calculus (see below) and mainly have to guess and substitute solutions or solution templates. As we will see, the situation looks slightly better for asymptotic solutions and this is one of the reasons why asymptotic resource analysis is popular.

2.1 The Substitution Method

We will now perform a simple resource analysis using recurrence relations to illustrate both steps.

1. Deriving a recurrence relation $T(\vec{n})$ from the description of the algorithm (or the program).
2. Deriving a closed form from $T(\vec{n})$, that is, a non-recursive easily-understood function $f(n)$ such that $f(n) = T(n)$.

To find such a function $f(n)$, we will use the *substitution method*, which is one of the simplest and most general ways of solving recurrences.

Consider the factorial function *fac* that implement in OCaml below.

```
let rec fac n =
  if n > 1 then
    n * (fac (n-1))
  else
    1
```

Assume we are interested in the number of multiplications that are performed by *fac*. The first question that arises is how to abstract an integer with a natural number. In this case we can see right away that the cost for negative inputs is 0 and focus on positive integers which are simply abstracted by the values. We can express the cost with the following recurrence.

$$\begin{aligned} T_{\text{fac}}(0) &= 0 \\ T_{\text{fac}}(1) &= 0 \\ T_{\text{fac}}(n) &= T_{\text{fac}}(n-1) + 1 \quad \text{if } n > 1 \end{aligned}$$

If $n = 0$ or $n = 1$ then the program does not execute any multiplications. If $n > 1$ then $\text{fac}(n)$ performance a multiplication plus the multiplications performed by $\text{fac}(n-1)$. We can prove by induction on n that the recurrence accurately captures the cost of the function.

We now solve the recurrence relation with the *substitution method*. We guess the solution $f(n)$ for $T_{\text{fac}}(n)$ is linear, that is, $f(n) = c_1 n + c_0$ for constants c_1 and c_0 . To verify that $f(n)$ is a correct solution and to determine the constants, we substitute $f(n)$ into the recurrence relations. We obtain

$$\begin{aligned} 0 \cdot c_1 + c_0 &= 0 \\ 1 \cdot c_1 + c_0 &= 0 \end{aligned}$$

which implies $c_0 = c_1 = 0$. However, this doesn't satisfy the remaining equality $T_{\text{fac}}(n) = T_{\text{fac}}(n-1) + 1$ which implies that there is no linear solution for the recurrence relation.

We could now start looking for non-linear solution but the recurrence is so simple that we can see that the problem originates from the base cases. So we simply drop $T_{\text{fac}}(0) = 0$ and only focus on positive numbers. While this seems to be an insignificant detail in our manual analysis, such issues make it difficult to automatically derive the right recurrence relations from code.

We again substitute our guess (or *ansatz*) into the remaining two equations.

$$\begin{aligned} c_1 + c_0 &= 0 \\ c_1(n+1) + c_0 &= c_1 n + c_0 + 1 \end{aligned}$$

Now we can solve for c_1 and c_0 and obtain $c_1 = 1$ and $c_0 = -1$. A sanity check shows that $f(n) = n - 1$ is indeed a solution of the recurrence.

Unfortunately, coming up with a candidate for a solution is not always straightforward.

2.2 The Recursion-Tree Method

A powerful method for finding an asymptotic solution (or candidate solution) for a recurrence for particular divide and conquer algorithms is the *recursion tree method* (or master method). Recall the following definitions for asymptotic notation.

Definition 1. Let $f, g : \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ be two functions from natural number to non-negative reals. We write

- $f(x) \in O(g(x))$ if there exists $C > 0$ and $N \in \mathbb{N}$ such that $f(x) \leq C \cdot g(x)$ for all $x \geq N$
- $f(x) \in \Omega(g(x))$ if there exists $C > 0$ and $N \in \mathbb{N}$ such that $f(x) \geq C \cdot g(x)$ for all $x \geq N$
- $f(x) \in \Theta(g(x))$ if $f(x) \in \Omega(g(x))$ and $f(x) \in O(g(x))$

You may notice that these definitions only apply to functions with one argument. In the homework, you will see that they cannot be directly generalized to multiple arguments.

Consider a divide and conquer algorithm that divides a problem of size n in a sub-problems of size n/b , which are recursively solved. The solutions of the sub-problems are then combined to a solution of the original problem of size n . Assume that the combination of the sub-solutions and the division into sub-problems needs time cn^k . The run time of such an algorithm can then be expressed as a recurrence relation $T(n)$ as follows where a, b, c , and k are positive constants.

$$\begin{aligned} T(1) &= c \\ T(n) &= aT\left(\frac{n}{b}\right) + cn^k \quad \text{if } n > 1 \end{aligned}$$

Example 1. Consider for example the divide and conquer algorithm merge sort. Merge sort divides a problem of size n into 2 sub-problems of size $\frac{n}{2}$. Division and merging takes time cn for some constant c . We thus have $a = b = 2$, $k = 1$ and $T_{\text{merge}} = 2T\left(\frac{n}{2}\right) + cn$.

The recurrence relation $T(n)$ can be solved systematically using the recursion-tree method. To see how consider the recursion tree that arises when unrolling the recurrence $T(n)$ in Figure 1. We will derive a closed form for $T(n)$ by summing up the cost level by level. The cost for the top




Figure 1: Recursion tree for $T(n) = aT(\frac{n}{b}) + cn^k$.

level is cn^k , the cost for the second level is $ac(\frac{n}{b})^k$, and the cost of the third level is $a^2c(\frac{n}{b^2})^k$. In total, the sum of costs is

$$T(n) = cn^k \left(1 + \frac{a}{b^k} + \left(\frac{a}{b^k}\right)^2 + \left(\frac{a}{b^k}\right)^3 + \dots + \left(\frac{a}{b^k}\right)^{\log_b n} \right) \quad (1)$$

If we define $r = \frac{a}{b^k}$ then we have

$$T(n) = cn^k (1 + r + r^2 + r^3 + \dots + r^{\log_b n}) \quad (2)$$

We now distinguish three cases.

Case $r < 1$. Then $1 + r + r^2 + r^3 + \dots + r^{\log_b n}$ are the first $\log_b n$ terms of a geometric series. So

$$1 + r + r^2 + r^3 + \dots + r^{\log_b n} < 1 + r + r^2 + r^3 + \dots = \frac{1}{1-r}$$

and $cn^k \leq T(n) \leq \frac{cn^k}{1-r}$. Thus $T(n) \in \Theta(n^k)$ since r and c are constants.

Case $r = 1$ Then $1 + r + r^2 + r^3 + \dots + r^{\log_b n} = \log_b n + 1$ and

$$T(n) = cn^k (\log_b n + 1) \in \Theta(n^k \log n).$$

Case $r > 1$ Then we have

$$1 + r + r^2 + r^3 + \dots + r^{\log_b n} = \frac{1 - r^{\log_b n + 1}}{1 - r} = \frac{r^{\log_b n + 1} - 1}{r - 1} = \Theta(r^{\log_b n}).$$

Thus

$$\begin{aligned} T(n) &\in \Theta(n^k (\frac{a}{b^k})^{\log_b n}) \\ &= \Theta(a^{\log_b n}) && \text{since } b^{k \cdot \log_b n} = n^k \\ &= \Theta(n^{\log_b a}) && \text{since } a^{\log_b n} = b^{(\log_b a)(\log_b n)} = n^{\log_b a} \end{aligned}$$

In summary, the so called *master method* for solving recurrences for divide and conquer algorithms is stated by the following theorem.

Theorem 1 (Master Theorem). Let $T(n) = aT(\frac{n}{b}) + cn^k$ be a recurrence relations and let $a, b, c, k > 0$ be positive constants. Then

$$T(n) \in \begin{cases} \Theta(n^k) & \text{if } \frac{a}{b^k} < 1 \\ \Theta(n^k \log n) & \text{if } \frac{a}{b^k} = 1 \\ \Theta(n^{\log_b a}) & \text{if } \frac{a}{b^k} > 1 \end{cases}$$

There are many variants and generalizations of the recursion tree method and more powerful “master theorems” in which cn^k is replaced with an arbitrary function $f(n)$. However, they are all limited to divide and conquer algorithms, derive asymptotic bounds only, have only one parameter, and cannot be applied if the “split factors” a and b are not constants.

Example 2. Consider again merge sort with recurrence $T_{\text{merge}}(n) = 2T(\frac{n}{2}) + cn$ and $a = b = 2$, and $k = 1$. Therefore $\frac{a}{b^k} = 1$ and by Theorem 1 it follows that $T_{\text{merge}}(n) = \Theta(n \log n)$.

So how useful is the recursion tree method for finding an *ansatz* for exactly solving a recurrence relation? It is only partially helpful. Do not be deceived by the use of Θ in the theorem. If $T(n) \in \Theta(n^k)$ does not mean that $T(n) = cn^k$ for some c . It’s probably worth trying an *ansatz* with a polynomial of degree k . However, arbitrary functions like $n^k + n \log n$ are members of $\Theta(n^k)$. So there’s no guarantee that this approach will succeed.

2.3 Linear Recurrence Relations

One of the few techniques for obtaining exact solutions to recurrence relations is based on calculating roots of characteristic polynomials. It applies to *homogeneous linear recurrence relations with constant coefficients*. These are recurrence relations of the form

$$T(n) = c_1 T(n-1) + \dots + c_d T(n-d)$$

for constants c_i with $c_d \neq 0$. We call d the degree of the recurrence. To uniquely define the value of the recurrence T we have to define the initial or base values $T(0), \dots, T(d-1)$.

We can find a closed form for a homogeneous linear recurrence relation by finding the roots of the characteristic polynomial

$$P(x) = x^d - c_1 x^{d-1} - \dots - c_d$$

which has d roots r_1, \dots, r_d . There are two cases. If the roots are pairwise distinct then

$$T(n) = k_1 r_1^n + \dots + k_d r_d^n$$

for constants k_i that can be determined using the initial values $T(0), \dots, T(d-1)$.

If the roots are not pairwise distinct then we add additional terms to the products in the solution that are determined by the number of identical roots that occurred in the formula already:

$$T(n) = k_1 n^{u_1} r_1^n + \dots + k_d n^{u_d} r_d^n$$

Here, u_i is number of identical roots r_j ($j < i$ and $r_j = r_i$) that appeared before r_i in the formula. Note that this is a generalization of the formula for pairwise distinct roots.

Non-Homogeneous Recurrence Relations Recurrence relations that arise in resource analysis usually not homogeneous because they have a cost component in addition to the cost of the base cases. The form of such *non-homogeneous linear recurrence relations with constant coefficients* is given by

$$T(n) = c_1 T(n-1) + \dots + c_d T(n-d) + P(n)$$

where $P(n)$ is a function in n . If $P(n)$ is a polynomial then we can reduce the problem of finding a solution to the problem of finding a solution for a homogeneous linear recurrence relation.

This is best illustrated with an example. Let us consider the split function for merge sort in Figure 2, which recursively splits a list into two lists of (almost) equal length. To this end, split removes the first two list elements, recursively splits the remaining list, and adds one of the removed elements to each of the returned lists. If we count the number of cons operations, the cost can be described by the following recurrence relation.

$$\begin{aligned} T_{\text{split}}(0) &= 0 \\ T_{\text{split}}(1) &= 1 \\ T_{\text{split}}(n+2) &= 2 + T_{\text{split}}(n) \end{aligned}$$

```

let rec split l =
  match l with
  | [] → ([], [])
  | x1::xs →
    match xs with
    | [] → ([x1], [])
    | x2::xs' →
      let (l1,l2) = split xs' in
      (x1::l1,x2::l2)

```

Figure 2: The split function for merge sort.

To convert the recurrence into a homogeneous recurrence that has the same closed-form solution, we use the equality $T_{\text{split}}(n+3) = 2 + T_{\text{split}}(n+1)$ to subtract $T_{\text{split}}(n+3)$ from both sides of the original recurrence. We obtain $T_{\text{split}}(n+2) - T_{\text{split}}(n+3) = 2 + T_{\text{split}}(n) - 2 - T_{\text{split}}(n+1)$ and thus

$$T_{\text{split}}(n) = T_{\text{split}}(n-1) + T_{\text{split}}(n-2) - T_{\text{split}}(n-3)$$

The characteristic polynomial for $T_{\text{split}}(n)$ is

$$P_{\text{split}}(x) = x^3 - x^2 - x + 1$$

The roots of the characteristic polynomial are 1 and -1 . The root 1 appears twice as we can write $P_{\text{split}}(x)$ as $(x-r)^2q(x)$ for $r=1$ and a polynomial $q(x)$, namely $q(x) = x+1$.

Using the previously discussed formula we have

$$\begin{aligned} T_{\text{split}}(n) &= k_1(-1)^n + k_21^n + k_3n1^n \\ &= k_1(-1)^n + k_2 + k_3n \end{aligned}$$

for yet undetermined constants k_i . Since $T_{\text{split}}(0) = 0$ we have $k_1 + k_2 = 0$. Since $T_{\text{split}}(1) = 1$ we have $k_2 + k_3 - k_1 = 1$ and since $T_{\text{split}}(2) = 1$ we have $k_1 + k_2 + 2k_3 = 2$. Therefore $k_1 = k_2 = 0$, $k_3 = 1$, and $T_{\text{split}}(n) = n$.

We can also experiment with other cost. If we for example define $T_{\text{split}}(1) = 0$ then we get the more interesting result $T_{\text{split}}(n) = 0.5(-1)^n - 0.5 + n$.

2.4 Example: Quick Sort

Hoare's quicksort algorithm is an example of a divide and conquer algorithm for which the master method cannot be applied. Figure 3 shows an implementation of quicksort in OCaml.

Our goal is to use recurrence relations to analyze the worst-case complexity of quicksort. The cost model we consider is the number of function calls. The first challenge in the analysis is that quicksort is higher-order function: It takes the comparison function le as its first argument. This is not uncommon, even in languages like C which do not have first class functions such as OCaml. In the C standard library, one of quicksort's arguments is a pointer to a comparison function. We will ignore the issue of higher order functions for now and will just assume that all comparison functions have a constant cost. The second challenge in the analysis is that quicksort is a curried function: The call `quicksort gt` consumes a constant amount of resources and returns a function closure. We will also ignore the issue of curried functions and will for now assume that quicksort is always applied to both of its arguments.

Before we can analyze quicksort, we have to analyze the helper functions `append` and `partition`. The recurrence relations that we obtain from the code are

$$\begin{aligned} T_{\text{app}}(0, m) &= 0 \\ T_{\text{app}}(n, m) &= T_{\text{app}}(n-1, m) + 1 \quad \text{if } n > 1 \\ \\ T_{\text{par}}(0) &= 0 \\ T_{\text{par}}(n) &= T_{\text{par}}(n-1) + 2 \quad \text{if } n > 1 \end{aligned}$$

```

let rec append l1 l2 =
  match l1 with
  | [] → l2
  | x::xs → x::(append xs l2)

let rec partition f l =
  match l with
  | [] → ([], [])
  | x::xs →
    let (cs,bs) = partition f xs in
    if f x then
      (cs,x::bs)
    else
      (x::cs,bs)

let rec quicksort le = function
| [] → []
| x::xs →
  let ys, zs = partition (le x) xs in
  append (quicksort le ys) (x :: (quicksort le zs))

```

Figure 3: An implementation of quicksort in OCaml.

For *append* (T_{app}), n is the length of the first list in the argument and m is the length of the second list in the argument. For *portioning* (T_{par}), n is the length of the second argument. Similar as for, T_{fac} we can derive $T_{\text{app}}(n, m) = n$ and $T_{\text{par}}(n) = 2n$.

We proceed to derive a recurrence relation T_{qs} for the function *quicksort*. When looking at the code, the question arises what the size of the lists ys and zs in the recursive calls for *quicksort* is. In fact, if $|\ell|$ denotes the length of a list ℓ , we have $|ys| + |zs| = |xs|$. So we need to prove by induction on xs that for every function f if $(ys, zs) = \text{partition } f xs$ then $|ys| + |zs| = |xs|$.

Before, we can state the recurrence for *quicksort*, we have to show that $|\text{quicksort } le xs| = |xs|$. However, we omit this step for brevity.

$$\begin{aligned}
 T_{\text{qs}}(0) &= 0 \\
 T_{\text{qs}}(n+1) &= \max_{0 \leq i \leq n} (T_{\text{qs}}(i) + T_{\text{qs}}(n-i-1) + i) + 2n + 4
 \end{aligned}$$

Let us first assume that the array is always split in the middle, that is, $i = \lfloor n/2 \rfloor$. We then have

$$T_{\text{qs}}(n) = 2T_{\text{qs}}(n/2) + \Theta(n)$$

and can use case $a/b^k = 1$ of the master method to infer $T_{\text{qs}}(n) \in \Theta(n \log n)$.

Let us now assume that the list is always split in the most imbalanced way, that is, $i = 0$. Then

$$\begin{aligned}
 T_{\text{qs}}(n) &= T_{\text{qs}}(n-1) + T_{\text{qs}}(0) + \Theta(n) \\
 &= T_{\text{qs}}(n-1) + \Theta(n)
 \end{aligned}$$

In this case, We can show with the substitution method that $T_{\text{qs}}(n) \in \Theta(n^2)$.

We conjecture that in general $T_{\text{qs}}(n) \leq cn^2$ for a constant c and use the substitution method to prove it.

$$\begin{aligned}
 T_{\text{qs}}(n) &\leq \max_{0 \leq i \leq n-1} (ci^2 + c(n-i-1)^2) + \Theta(n) \\
 &= \max_{0 \leq i \leq n-1} c(i^2 + (n-i-1)^2) + \Theta(n)
 \end{aligned}$$

For a fixed n , the polynomial $i^2 + (n-i-1)^2$ takes maximums in the range $0 \leq i \leq n-1$ for $i = 0$ and $i = n-1$. To verify this claim observe that the second derivation with respect to i is positive. Thus we have

$$\max_{0 \leq i \leq n-1} i^2 + (n-i-1)^2 \leq (n-1)^2 = n^2 - 2n + 1$$

and conclude

$$T_{qs} \leq cn^2 - 2cn + c + \Theta(n) \in O(n^2).$$

We have now performed only an asymptotic analysis. To determine the exact constants is even more tedious. It is not even clear how an exact solution to the recurrence could look like since, as we have seen, the resource usage varies based on the choices of the pivot element. However, we have abstracted away the elements of the input list and it seems impossible to recover from that abstraction. So we stop here as the previous calculations already illustrate the difficulties with automating and mechanizing recurrence solving.

3 Amortized Analysis with the Potential Method

Amortized analysis is a technique to derive a worst-case resource-usage bound for a sequence of operations (or function calls). For many data structures, the amount of resources that an operation consumes can vary substantially depending on the state of the data structure. Nevertheless, high cost (e.g., reorganizing the data structure) will occur with some predictable frequency; they *amortize* over time. Summing up the worst-case costs of operations in such a sequence of operations would lead to gross over-approximations of the cost. In these cases, amortized analysis provides a much more precise approach.

Remark. *You will sometimes hear that amortized analysis is a way of determining the amortized or even average cost of an operation. However, the goal of an amortized analysis is to determine the worst-case cost of a sequence of operation. As a means to an end, we sometimes define the amortized cost of an operation but this is not the final result of the analysis.*

Potential Functions To amortize the cost of different operations, we introduce a potential function

$$\Phi : \text{State} \rightarrow \mathbb{Q}_{\geq 0}.$$

The idea is that an operation $o : \text{State} \rightarrow \text{State}$ that is executed in state S can use the potential $\Phi(S)$ to pay for the cost of the operation. More specifically, we say that the *amortized cost* of operation o with respect to Φ is

$$\text{acost}(o) = \max_{S \in \text{State}} \text{cost}(o) + \Phi(o(S)) - \Phi(S)$$

where $\text{cost}(o)$ is the actual resource consumption of o .

Theorem 2. *Given starting state S_0 and operations o_1, \dots, o_n , we have*

$$\sum_{1 \leq i \leq n} \text{cost}(o_i) \leq \sum_{1 \leq i \leq n} \text{acost}(o_i) + \Phi(S_0).$$

Proof. Let $S_i = o_i(S_{i-1})$ be the intermediate states. Then

$$\begin{aligned} \sum_{1 \leq i \leq n} \text{acost}(o_i) + \Phi(S_0) &\geq \Phi(S_0) + \sum_{1 \leq i \leq n} \text{cost}(o_i) + \Phi(S_i) - \Phi(S_{i-1}) \\ &= \Phi(S_0) + \Phi(S_n) - \Phi(S_0) + \sum_{1 \leq i \leq n} \text{cost}(o_i) \\ &= \Phi(S_n) + \sum_{1 \leq i \leq n} \text{cost}(o_i) \\ &\geq \sum_{1 \leq i \leq n} \text{cost}(o_i) \end{aligned}$$

□

The challenge in performing an amortized analysis is to choose the right potential function Φ . The goal is to chose Φ so that the amortized cost $\text{cost}(o) + \Phi(o(S)) - \Phi(S)$ is identical for every state S .

Example: Stack Given is an implementation of a stack with operations *push* and *pop*. We assume that $\text{cost}(\text{push}(S,x)) = 1$ and $\text{cost}(\text{pop}(S)) = \min(|S|, 1)$. Consequently, the cost of a sequence of n *push* and *pop* operations is less or equal to n .

Consider the additional operation *multipop*(S,k) that pops k elements from the stack. If *multipop*(S,k) is executed for a stack S with less than k elements then all elements of S are popped. Similarly, if *pop*(S) is applied to an empty stack S then S is unchanged and some error message is returned. The cost of *multipop*(S,k) is $\min(|S|, k)$.

What is the worst-case cost of a sequence of n *push*, *pop*, and *multipop* operations? In a conservative analysis we would first establish the worst-case cost of each operation. So we would argue that the worst-case cost for *multipop* is $n - 1$ since there are $n - 1$ elements on the stack in the worst-case. Consequently, the cost of n operations can only be bounded by n^2 operations. This is clearly a very loose bound since every element on the stack can only be popped once.

Using the potential method, we can perform a much more precise analysis. For a stack S we define the potential

$$\Phi(S) = |S|$$

to be the height of the stack. The amortized cost of push is then

$$\text{acost}(\text{push}) = 2$$

since $\Phi(\text{push}(S, x)) - \Phi(S) = 1$ for every stack S and $\text{cost}(\text{push}(S)) = 1$. Furthermore, we have $\Phi(\text{pop}(S)) - \Phi(S) = -\min(|S|, 1)$ and $\text{cost}(\text{pop}(S)) = \min(|S|, 1)$ for every S and thus

$$\text{acost}(\text{pop}) = 0.$$

Similarly, $\Phi(\text{multipop}(S, k)) - \Phi(S) = -\min(|S|, k)$ and $\text{cost}(\text{multipop}(S, k)) = \min(|S|, k)$ for every S and thus

$$\text{acost}(\text{multipop}) = 0.$$

Therefore it follows from Theorem 2 that the cost of n operations is at most $2n$. It is also easy to see that the cost is actually bounded by $2m$, where m is the number *push* operations in the sequence.

Example: Binary Counter Consider a binary counter $b = b_k, \dots, b_0$ which is implemented using a list or an array of bits of fixed length. The binary counter only has an operation *inc*(b) which increments b by 1. The cost of *inc*(b) is defined as number of bits that have to flipped for the update. We observe that at most one 0 is flipped to a 1 in an increment. However, multiple 1s can be flipped to 0s. Let $|b|_1$ be the number of 1s in the counter. Then $|b|_1 + 1$ is an upper bound on the number of bits that are flipped by one increment.

We are interested in the cost of n increment operations. To derive a worst-case bound, we use the potential method of amortized analysis. We define the potential

$$\Phi(b) = |b|_1$$

To determine the *amortized cost* of *acost*(*inc*) assume that the operation *inc*(b) modifies t_i bits. As discussed earlier, $t_i - 1$ flips are from 1 to 0 and at most one flip (zero can happen during overflow) is from 0 to 1. So we have

$$\begin{aligned} \text{cost}(\text{inc}(b)) + \Phi(\text{inc}(b)) - \Phi(b) &= t_i + \Phi(\text{inc}(b)) - \Phi(b) \\ &\leq t_i + (\Phi(b) - (t_i - 1) + 1) - \Phi(b) \\ &= 2 \end{aligned}$$

It follows that $\Phi(b^0) + 2n$ is an upper bound on the number of flips in a sequence of n increment operations. Here b^0 is the initial counter.

Example: Dynamic Table Another standard example of amortized analysis is a dynamic table T that is implemented with an array. Assume first that we only have an insert operation $insert(T, x)$, which inserts a new element x into the table. We simply insert the new elements successively until the array is full. In this case, we allocate a new array of double the size of the current array and copy the elements over. Our cost model in this example is to count the number of insertions. So the worst-case cost of one $insert(T, x)$ is $|T| + 1$. However, we are again interested in a sequence of n insert operations.

Let $|T|_{\text{elm}}$ be the number of stored elements in T . We define the potential function

$$\Phi(T) = 2|T|_{\text{elm}} - |T|$$

If we always start with an empty array then the load factor is never smaller than $1/2$ and $\Phi(T) \geq 0$. If we would allow arbitrary load factors at the beginning we can simply define $\Phi(T) = \max(0, 2|T|_{\text{elm}} - |T|)$.

We now determine the amortized cost. There are two cases. Assume first that $|T|_{\text{elm}} = |T|$. Then

$$\begin{aligned} \text{cost}(insert(T, x)) + \Phi(insert(T, x)) - \Phi(T) &= |T| + 1 + \Phi(insert(T, x)) - (2|T|_{\text{elm}} - |T|) \\ &= |T| + 1 + \Phi(insert(T, x)) - |T| \\ &= |T| + 1 + (2(|T| + 1) - 2|T|) - |T| \\ &= 3 \end{aligned}$$

Assume now that $|T|_{\text{elm}} < |T|$. Then

$$\begin{aligned} \text{cost}(insert(T, x)) + \Phi(insert(T, x)) - \Phi(T) &= 1 + \Phi(insert(T, x)) - (2|T|_{\text{elm}} - |T|) \\ &= 1 + (2|T|_{\text{elm}} + 2 - |T|) - (2|T|_{\text{elm}} - |T|) \\ &= 3 \end{aligned}$$

As a result, we have $acost(insert) = 3$.

We now add another operation $remove(T, x)$ that deletes the element x from the table T . Similar, the expansion now also want to contract the table if the load factor is equal to a constant $c \geq 1/2$. You might have the intuition to pick $c = 1/2$. However, this is not a good choice since a sequence of alternating inserts and deletes can cause quadratic cost in the length of the sequence.

A better choice is $c = 1/4$. Now we define the potential to be

$$\Phi(T) = \begin{cases} 2|T|_{\text{elm}} - |T| & \text{if } |T|_{\text{elm}}/|T| \geq 1/2 \\ |T|/2 - |T|_{\text{elm}} & \text{if } |T|_{\text{elm}}/|T| < 1/2 \end{cases}$$

The amortized costs are then $acost(insert) = 3$ and $acost(remove) = 3$.