Chapter 1: Computational Models, Branching Programs and Circuits Anup Rao March 28, 2022

OUR GOAL IN THIS COURSE IS to mathematically capture the concept of computation. In the broadest sense, what is computation? What is the best way to model computation? How can we reason about whether something can be efficiently computed or not? What are the resources that are worth measuring with regards to computational efficiency? Complexity theory is about addressing these questions.

During this class we shall explore many mathematical concepts that will help us deal with these topics. Along the way, we will naturally be led to studying some very interesting questions. Here is a sample of some of the questions we will discuss:

- 1. Can every function be computed?
- 2. Can every true statement be proved?
- 3. Can every slow algorithm be replaced with a faster algorithm that computes the same function?
- 4. Can every slow sequential algorithm be replaced with a faster parallel algorithm?
- 5. Is it possible to check the correctness of a computation that took a long time using much less time?

Throughout, we will use clever definitions and ideas from mathematics to reason about these questions.

Before we begin, we need to define what we mean by computation.

Computation

A FIRST ATTEMPT at defining a computation might be to say that it is a process that *manipulates information* in some way, or has an *inputoutput* behavior. For example, when reading this sentence, your brain takes the information encoded graphically and translates that information into letter, words and ideas, and so performs a computation.

At this point you may feel that we have made the model too general to be useful, so it might be useful to explain what is *not* computation. A useful mathematical abstraction that captures some of the things we have discussed is the abstraction of *functions*. Given two sets *D*, *R*, a function

$$f: D \to R$$

assigns a value $f(x) \in R$ to every element of $x \in D$. So, if we think of D as the set of all possible images, and R as the set of all sentences, f can be defined to be the function that maps the picture of a sentence to the actual sentence. This abstraction misses something that is inherent about physical computational processes: computations are local. At any point, the state of two parts of the brain that are far away from each other cannot affect each other.

Informally, a computational process is a process that manipulates information in some *local* or restricted way. Interesting computational processes are ones that manage to have a complicated global effect through incremental local steps. There are many such computational processes, and we will not be able to talk about all of them in detail in this course. The aim of this course is to show you, in broad strokes, how you can start to reason about such computational processes and their complexity.

Computational Complexity

How CAN WE DISTINGUISH functions that are easy to compute from functions that are hard? What makes some things easy and other things hard?

In order to tackle this kind of question, we first need mathematical models that captures exactly what the process we are interested in can do cheaply, and what takes more effort. We would like our models to be general enough that they capture most real computational processes, and simple enough that we can ask and understand easy questions about them.

A crucial issue is how the information being manipulated is encoded. For example, if numbers are encoded using their prime factorization (both in the input and output), then it is slightly easier for us to multiply two 100 digit numbers than if they are encoded using their digits. Under this representation, addition of numbers, and comparing two numbers becomes harder.

So, given a function $f : D \rightarrow R$, we would like to be able to quantify how difficult it is to compute this function. We shall make two immediate simplifications.

• We shall identify the the input domain with the set of binary strings of arbitrary length {0,1}*, or the set of strings of some fixed length {0,1}ⁿ. Since every countable set can be mapped to

For example, we are very good at reading text, but multiplying 100 digit numbers takes us considerably more time, even though the amount of information is contained in a picture is much more than the information contained in a 100 digit number.

Addition is hard when numbers are represented in terms of their factorization because we would have to factor the sum to bring it back in this representation. We do not know of any efficient algorithms for factoring numbers.

Figure 1: A finite state automaton that accepts strings that have a 1 in some odd location.



the first set, and every finite set can be mapped to the second, this does not lose too much generality.

• We shall often restrict the output domain to be $R = \{0, 1\}$. This does lose some generality, but it will turn out that most of our ideas will easily translate to the situation when the output domain is bigger. Further, for most examples of functions to bigger domains that are hard to compute, we shall be able to easily find corresponding boolean functions that are hard to compute.

Next we give several examples of computational models, and discuss some strengths and weaknesses of each of them.

Finite Automata

A FINITE AUTOMATON CAN BE USED to compute functions f: $\{0,1\}^* \rightarrow \{0,1\}$. An example is shown in Figure 1. One starts at the start state (labeled *S*), and reads the input bit by bit, transitioning on the states. The output is 1 if and only if we ever hit the accept state (labeled *A*).

The weaknesses of this model:

- The set of functions *f* : {0,1}* → {0,1} that can be computed by any finite automaton is not very general. Even more damning—finite automata cannot compute functions that we can efficiently compute in practice.
- It does not give a way to measure the complexity of computing functions. We can count the number of states in the automaton, but since this is always a constant independent of *n*, this measure

Only functions that represent *regular* languages can be computed by a finite automaton.

Try to prove that no finite automaton can compute the function whose output is 1 if and only if the input is a palindrome. of complexity doesn't scale with the input size, and so does not map to what we are trying to capture — the level of difficulty involved in carryout the computation.

Finite automata are a *uniform* model of computation. This means that there is a single description of the process that can be used to carry out the computation no matter how long the length of the input is.

Branching Programs

BRANCHING PROGRAMS ARE VERY SIMILAR to finiate automata. This is a model that computes functions $f : \{0, 1\}^n \rightarrow \{0, 1\}$.

A branching program is a directed acyclic graph where every vertex is labeled by either a variable or an output value. Every vertex that is labeled by a variable has exactly two edges coming out of it, one labeled 1 and the other labeled 0. The vertices are partitioned into disjoint sets L_0, \ldots, L_k , and all edges go from some set L_i to L_{i+1} . There is a special designated start node in the graph. To carry out the computation, we start at the start state, and follow the indicated path by reading the variable that labels the state that we are on in each step. When we hit a node labeled by an output value, the output of the computation is that output value.

There are two measures of the complexity of the program. The *width* is the size of the largest layer. The length is the number of layers.

This model has the advantage that it can actually compute all functions:

Fact 1. Every function $f : \{0,1\}^n \to \{0,1\}$ can be computed by a branching program with width 2^n , and length n + 1.

Proof Consider the branching program that is a rooted tree, where every node at level *i* reads the variable x_i . The program simply remembers the entire input. Each of the 2^n inputs *x* gets mapped to a distinct leaf, so that leaf can be labeled by the value f(x) to compute f.

Branching programs will be used later in the course as a way to capture computations that use a small amount of space (or memory). The main drawback of branching programs is that there are algorithms which run very quickly on computers in practice that we don't know how to model as small branching programs. So, the branching program complexity doesn't seem to capture everything we want to capture about efficient computation. In class, we discussed how every finite automaton for a function $f : \{0,1\}^* \rightarrow \{0,1\}$ can be used to get a branching program of size O(n) that computes f on all the inputs of length n.



Figure 2: A branching program that computes whether or not the number of 1's in the input is even.

Branching programs define a *non-uniform* model of computation. The program only tells you how carry out computation on an input of length *n*. To talk about the asymptotic complexity of computing a function $f : \{0,1\}^* \rightarrow \{0,1\}$ that is defined on strings of all lengths, we need to talk about families of branching programs, and discuss the complexity of the programs as *n* gets larger and larger.

Communication Complexity

In the model of communication complexity, there are two parties, Alice and Bob. Alice has an input $x \in \{0,1\}^n$, and Bob has an input $y \in \{0,1\}^n$. They wish to compute a function f(x,y), and in order to do so, they communicate. The complexity measure is the total number of bits they must exchange to compute f.

Now let us be more rigorous. A communication protocol for computing the function f is a binary tree, where every internal vertex in the tree v is associated with Alice or Bob, and is labeled by a function $m_v : \{0,1\}^n \rightarrow \{0,1\}$. The leaves of the binary tree are labeled by bits. To execute the protocol, the players start at the root of the tree. If the root r is associated with Alice, then Alice sends Bob the bit $m_r(x)$. If the root is associated with Bob, then Bob sends Alice $m_r(y)$. If this message is 0, the players move to the vertex that is the left child of the root, and if it is 1, they move to the right. In this way, the players communicate and move down the tree, until they hit a leaf. At this point they output the value labeling the leaf.

The communication complexity of such a protocol is the depth of the tree. In this way, we have managed to come up with a model of computation that captures conversations between two people!

Boolean Circuits

A *boolean circuit* computing a function $f : \{0,1\}^n \to \{0,1\}$ is a directed acyclic graph with the following properties. Every vertex (also called a gate) has at most 2 edges coming in to it. If there are 0 edges coming in, then the vertex is labeled with an input variable x_i , or the constants 0 or 1. Otherwise, the vertex is labeled with one of

Communication complexity is extremely powerful as a method for proving lower bounds. For example, if you want to prove that a chip for computing a given function must have large area, you can cut the chip into two parts in such a way that the number of wires crossing the cut is small. Then think of the left side of the cut as "Alice" and the right side as "Bob". If the communication complexity is large, then there must be a lot of wires crossing the cut, and this gives a lower bound on the area of the chip. It is hard to imagine a model of computation that cannot be broken up into two communicating parts in this way!





the boolean operators \land , \lor , \neg , and computes the specified operation on the bits that come in along the incoming edges. One of the gates in the circuit is designated the output node. This is the node whose value is the output of the circuit.

When every gate has out-degree at most 1, the circuit is called a *formula*. In the case of a formula, the graph of the circuit looks like a tree after edges have been converted into undirected edges.

A circuit can also be viewed as a program in a simple programming language, where every line is an assignment. For example, the circuit in Figure 3 is equivalent to this program:

- 1. $y_1 = x_1 \wedge x_2$
- 2. $y_2 = x_3 \lor x_4$
- 3. $y_3 = \neg x_5$
- 4. $y_4 = x_5 \wedge x_6$
- 5. $y_5 = y_1 \lor y_2$
- 6. $y_6 = x_5 \lor y_4$
- 7. $y_7 = y_5 \wedge y_3$
- 8. $y_8 = y_7 \lor y_6$

There are two major quantities we can measure to capture the complexity of a circuit:

Definition 2. The size of the circuit is the number of gates in the circuit.

Since every gate in the circuit has at most 2 incoming edges, the size of the circuit is proportional to the number of edges in the graph that defines the circuit:



Figure 4: Recursive construction of a circuit for *f*.

Fact 3. *The size of the circuit is the same as the number of edges in the circuit, up to a factor of 2.*

We can also measure the *depth* of the circuit:

Definition 4. *The* depth *of the circuit is the length of the longest input to output path.*

The depth complexity is a measure of how much parallel time it takes to compute the function.

Just like branching programs, boolean circuits can complete *every* function:

Theorem 5. Every function $f : \{0,1\}^n \to \{0,1\}$ can be computed by a circuit of size at most $O(2^n)$.

Proof We construct the circuit recursively. When n = 1, there is clearly a constant sized circuit that computes *f*, since *f* must be either a constant, x_1 or $\neg x_1$.

For n > 1, let f_0 denote the function on n - 1 bits given by $f_0(x) = f(x,0)$, and $f_1(x) = f(x,1)$. Then by induction we can compute f_0, f_1 recursively, and combine them using the value of the last bit to obtain f, as in Figure 4. When $x_n = 1$, the circuit outputs $f_1(x_1, ..., x_{n-1})$, and when $x_n = 0$, the circuit outputs $f_0(x_1, ..., x_{n-1})$.

If S_n is the size of the resulting circuit when the underlying function takes an n bit input, we have proved that

$$S_n \le 2S_{n-1} + 5.$$

Expanding this recurrence, and using the fact that $S_1 \leq 5$, we get that

$$S_n \le \sum_{i=1}^n 2^i 5 = 5 \cdot (2^{n+1} - 2) < 10 \cdot 2^n,$$

where here we used the formula for computing the sum of a geometric series. \blacksquare

The above theorem is not the best result we know about this subject. In fact, we know:

Theorem 6. Every function $f : \{0,1\}^n \to \{0,1\}$ can be computed by a circuit of size at most $O(2^n/n)$.

You will be asked to prove this on your homework.