

2 Elements of classical computer science

Good references for the material of this section are [3], Chap. 3, [5], Secs. 2 and 3, [7], Sec. 6.1, and [11] as a nice readable account of complexity with some more details than we will treat (and need) here.

2.1 Bits of History

The inventor of the first programmable computer is probably Charles Babbage (1791-1871). He was interested in the automatic computation of mathematical tables and designed the mechanical “analytical engine” in the 1830s. The engine was to be controlled and programmed by punchcards (a technique already known from the automatic Jacquard loom), but was never actually built. Babbage’s unpublished notebooks were discovered in 1937 and the 31-digit accuracy “Difference Engine No. 2” was built to Babbage’s specifications in 1991. (Babbage was also Lucasian professor of mathematics in Cambridge, like Newton, Stokes, Dirac, and Hawking, and he invented the locomotive cowcatcher.)

The first computer programmer probably is Ada Augusta King, countess of Lovelace (1815-1852), daughter of the famous poet, Lord Byron, who devised a programme to compute Bernoulli numbers (recursively) with Babbage’s engine. From this example we learn that the practice of writing programmes for not-yet existing computers is older than the quantum age.

Another important figure from 19th century Britain is George Boole (1815-1864) who in 1847 published his ideas for formalizing logical operations by using operations like AND, OR, and NOT on binary numbers.

Alan Turing (1912-1954) invented the Turing machine in 1936 in the context of the decidability problem posed by David Hilbert: Is it always possible to decide whether a given mathematical statement is true or not? (It is not, and Turing’s machine helped show that.)

For further details on the history of computing consult the history section of the entry “Computers” in the Encyclopaedia Britannica (unfortunately no longer available for free on the web).

2.2 Boolean algebra and logic gates

Computers manipulate bit strings. Any transformation on n bits can be decomposed into one- and two-bit operations.

There are only two one-bit operations: identity and NOT.

A **logic gate** is a one-bit function of a two-bit argument:

$$(x, y) \longrightarrow f(x, y) \text{ where } x, y, f = 0 \text{ or } 1.$$

The four possible inputs 00, 01, 10, 11 can be mapped to two possible outputs 0 and 1 each; there are $2^4 = 16$ possible output strings of 4 symbols and thus there are 16 possible logic gates or “Boolean functions”. Note that these gates are *irreversible* since the output is shorter than the input.

Before we discuss these functions we will recall some elements of standard Boolean logic. Boolean logic knows two *truth values*, true and false, which we identify with 1 and 0, respectively, and the elementary operations NOT, OR, and AND from which all other operations and relations like IMPLIES or XOR can be constructed. The binary operations OR and AND are defined by their “truth tables”

x	y	x OR y	x AND y
0	0	0	0
0	1	1	0
1	0	1	0
1	1	1	1

Note that, for example

$$x \text{ XOR } y = (x \text{ OR } y) \text{ AND } \text{NOT } (x \text{ AND } y).$$

(XOR is often denoted by \oplus , because it is equivalent to addition modulo 2.) We now return to the 16 Boolean functions of two bits. We number them according to the four-bit output string as given in the above truth table, read from top to bottom and interpreted as a binary number. Example: AND outputs 0001=1 and OR outputs 0111=7. We can thus characterize each gate (function) by a number between 0 and 15 and discuss them in order:

- 0: The absurdity, for example $(x \text{ AND } y) \text{ AND } \text{NOT } (x \text{ AND } y)$.
- 1: $x \text{ AND } y$

- 2: $x \text{ AND } (\text{ NOT } y)$
- 4: $(\text{ NOT } x) \text{ AND } y$
- 8: $(\text{ NOT } x) \text{ AND } (\text{ NOT } y) = (x \text{ NOR } y)$
- 3: x , which can be written in a more complicated way: $x = x \text{ OR } (y \text{ AND } \text{ NOT } y)$
- 5: $y = \dots$ (see above)
- 9: $((\text{ NOT } x) \text{ AND } (\text{ NOT } y)) \text{ OR } (x \text{ AND } y) = \text{ NOT } (x \text{ XOR } y) = (x \text{ EQUALS } y)$

All others can be obtained by negating the above; notable are

- 15: The banality, for example $(x \text{ AND } y) \text{ OR } \text{ NOT } (x \text{ AND } y)$.
- 14: $\text{ NOT } (x \text{ AND } y) = x \text{ NAND } y$
- 13: $\text{ NOT } (x \text{ AND } (\text{ NOT } y)) = x \text{ IMPLIES } y$

We have thus seen that all logic gates can be constructed from the elementary Boolean operations. Furthermore, since

$$x \text{ OR } y = (\text{ NOT } x) \text{ NAND } (\text{ NOT } y),$$

we see that we *only* need NAND and NOT to achieve any desired classical logic gate.

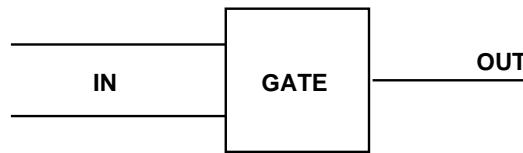


Figure 2.1: A logic gate.

If we picture logic gates as elements with two input lines and one output line and connect these logic gates by wires through which we feed the input, we get what is called the *network model of computation*.

2.2.1 Minimum set of irreversible gates

We note that

$$x \text{ NAND } y = \text{ NOT } (x \text{ AND } y) = (\text{ NOT } x) \text{ OR } (\text{ NOT } y) = 1 - xy \quad (x, y = 0, 1).$$

If we can copy x to another bit, we can use NAND to achieve NOT:

$$x \text{ NAND } x = 1 - x^2 = 1 - x = \text{ NOT } x$$

(where we have used $x^2 = x$ for $x = 0, 1$), or, if we can prepare a constant bit 1:

$$x \text{ NAND } 1 = 1 - x = \text{ NOT } x.$$

We can express AND and OR by NAND only:

$$(x \text{ NAND } y) \text{ NAND } (x \text{ NAND } y) = 1 - (1 - xy)^2 = 1 - (1 - 2xy + x^2y^2) = 1 - (1 - xy) = xy = x \text{ AND } y$$

and

$$(x \text{ NAND } x) \text{ NAND } (y \text{ NAND } y) = (\text{ NOT } x) \text{ NAND } (\text{ NOT } y) = 1 - (1 - x)(1 - y) = x + y - xy = x \text{ OR } y.$$

Thus NAND and COPY are a universal set of (irreversible) gates. It is also possible to use NOR in place of NAND. In fact, NAND and COPY can be performed by the same two-bit to two-bit gate, if we can prepare a bit in state 1. This is the NAND/NOT gate:

$$(x, y) \longrightarrow (1 - x, 1 - xy) = (\text{ NOT } x, x \text{ NAND } y).$$

The NOT and NAND functions are obviously achieved by ignoring one output bit. For $y=1$ we obtain COPY, combined with a NOT which can be inverted.

2.2.2 Minimum set of reversible gates

Although we know how to construct a universal set of irreversible gates there are good reasons to study the reversible alternative. Firstly, quantum gates *are* reversible, and secondly, it is interesting to discuss the possibility of reversible (dissipationless) computation.

A reversible computer evaluates an invertible n -bit function of n bits. Note that every irreversible function can be made reversible at the expense of additional bits:

$$x(n \text{ bits}) \longrightarrow f(m \text{ bits})$$

is replaced by

$$(x, n \text{ times } 0) \longrightarrow (x, f)(n + m \text{ bits}).$$

The possible *reversible* n -bit functions are the *permutations* of the 2^n possible bit strings; there are $(2^n)!$ such functions. For comparison, the number of *arbitrary* n -bit functions is $(2^n)^{(2^n)}$. The universal gate set for reversible (classical) computation will have to include three-bit gates, as discussed in [7]. The number of reversible 1-, 2-, and 3-bit gates is 2, 24, and 40320, respectively.

One of the more interesting reversible two-bit gates is the CNOT, also known as “reversible XOR”:

$$(x, y) \longrightarrow (x, x \text{ XOR } y),$$

that is

x	y	x	x XOR y
0	0	0	0
0	1	0	1
1	0	1	1
1	1	1	0

The reversibility is caused by the storage of x . The square of this gate is IDENTITY. XOR is often symbolized by \oplus . The CNOT gate can be used to copy a bit, because it maps

$$(x, 0) \longrightarrow (x, x).$$

The combination of three XOR gates in Fig. 2.2 achieves a SWAP:

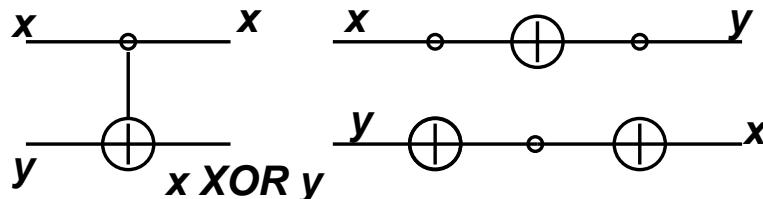


Figure 2.2: Left: Single CNOT gate. Right: SWAP gate

$$(x, y) \longrightarrow (x, x \text{ XOR } y) \longrightarrow ((x \text{ XOR } y) \text{ XOR } x, x \text{ XOR } y) \longrightarrow (y, y \text{ XOR } (x \text{ XOR } y)) = (y, x)$$

Thus the reversible XOR can be used to copy and move bits around. We will show now that the functionality of the universal NAND/NOT gate discussed above can be achieved by adding a three-bit gate to our toolbox, the *Toffoli gate* $\theta^{(3)}$, also known as controlled-controlled-NOT, which maps

$$(x, y, z) \longrightarrow (x, y, z \text{ XOR } xy).$$

This gate is universal, provided we can provide fixed input bits and ignore output bits:

- For $z = 1$ we have $(x, y, 1) \longrightarrow (x, y, 1 - xy) = (x, y, x \text{ NAND } y)$.
- For $x = 1$ we obtain $z \text{ XOR } y$ which can be used to copy, swap, etc.
- For $x = y = 1$ we obtain NOT.

Thus we can do any computation reversibly. In fact it is even possible to avoid the dissipative step of memory clearing (in principle): store all “garbage” which is generated during the reversible computation, copy the end result of the computation and then let the computation run backwards to clean up the garbage without dissipation. This may save some energy dissipation, but it has a price (as compared to reversible computation with final memory clearing):

- The time (number of steps) grows from T to roughly $2T$.
- Additional storage space growing roughly $\propto T$ is needed.

However, there are ways [7] to split the computation up in a number of steps which are inverted individually, so that the additional storage grows only $\propto \log T$, but in that case more computing time is needed.

Note that we have used the possibility to COPY bits without much thought. In quantum computation we will have to do without COPYing due to the “no cloning theorem”!

2.3 The Turing machine as a universal computer

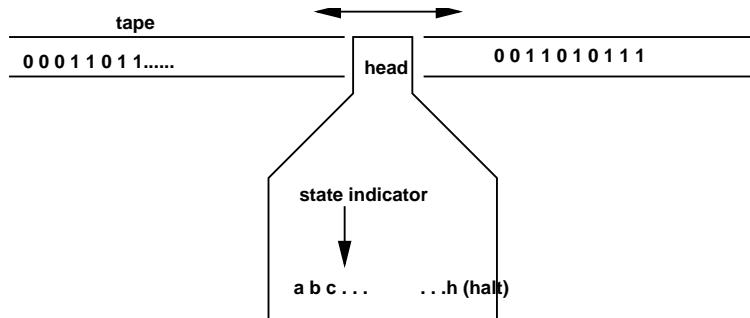


Figure 2.3: A Turing machine operating on a tape with binary symbols and possessing several internal states, including the *halt* state

The Turing machine acts on a tape (or string of symbols) as an input/output medium. It has a finite number of internal states. If the machine reads the binary symbol s from tape while being in state G , it will replace s by another symbol s' , change its state to G' and move the tape one step in direction d (left or right). The machine is completely specified by a *finite set of transition rules*

$$(s, G) \longrightarrow (s', G', d)$$

The machine has one special internal state, the “halt” state, in which the machine stops all further activity. On input the tape contains the “program” and “input data”, on output the result of the computation.

The (finite) set of transition rules for a given Turing machine T can be coded as a binary number $d[T]$ (the description of T). Let $T(x)$ be the output of T for a given input tape x . Turing showed that there exists a *universal Turing machine* U with

$$U(d[T], x) = T(x)$$

and the number of steps U needs to simulate each step of T is only a *polynomial* function of the length of $d[T]$. Thus we only have to supply the “description” $d[T]$ of T and the original input x on a tape to U and U will perform the same task as any machine T , with at most polynomial slowdown.

Other models of computation (for example the network model) are computationally equivalent to the Turing model: the same tasks can be performed with the same efficiency. Church [12] and Turing [13] stated the **Church-Turing thesis**: Every function which would naturally be regarded as computable can be computed by the universal Turing machine.

(“Computable functions” comprise an extremely broad range of tasks, such as text processing etc.) There is no proof of this thesis, but also no counterexample, despite many attempts to find one.

2.4 Complexity and complexity classes

We will not discuss many examples of problems from the different complexity classes. The article by Mertens [11] gives several examples in easy-to-read style.

Consider some task to be performed on an input (integer, for simplicity) number x , for example finding x^2 or determining if x is a prime. The number of bits needed to store x is

$$L = \log_2 x.$$

The *computational complexity* of the task characterizes how fast the number s of steps a Turing machine needs to solve the problem increases with L . Example: the method by which most of us have computed squares of “large” numbers in primary school has roughly

$$s \propto L^2$$

(if you identify s with the number of digits you have to write on your sheet of paper). This is a typical problem from class P : there is an algorithm for which s is a *polynomial* function of L . If s rises exponentially with L the problem is considered hard. (Note, however, that it is not possible to exclude the discovery of new algorithms which make problems tractable!)

It is often much easier to verify a solution than to find it; think of factorizing large numbers. The class NP consists of problems for which solutions can be verified in polynomial time. Of course P is contained in NP , but *it is not known if NP is actually larger than P* , basically because revolutionary algorithms may be discovered any day. NP means *non-deterministic polynomial*. A non-deterministic polynomial algorithm may **branch** into two paths which are both followed **in parallel**. Such a tree-like algorithm is able to perform an exponential number of calculational steps in polynomial time (at the expense of exponentially growing parallel computational capacity!). To verify a solution, however, one only has to follow “the right branch” of the tree and that is obviously possible in polynomial time.

Some problems may be *reduced* to other problems, that is, the solution of a problem P_1 may be used as a “step” or “subroutine” in an algorithm to solve another problem P_2 . Often it can be shown that P_2 may be solved by applying the subroutine P_1 a polynomial number of times; then P_2 is *polynomially reducible* to P_1 : $P_2 \leq P_1$ (Read: “ P_2 cannot be harder than P_1 .”) Some nice examples are provided by problems from graph theory, where one is looking for paths with certain properties through a given graph (or network), see [11]. A problem is called *NP-complete* if any NP problem can be reduced to it. Hundreds of NP-complete problems are known, one of the most famous being the *traveling salesman problem*. If somebody finds a polynomial solution for *any* NP-complete problem, then “P=NP” and one of the most fundamental problems of theoretical computer science is solved. This is, however, very unlikely, since many clever people have in vain tried to find such a solution. (It should be noted at this point that theoretical computer science bases its discussion of complexity classes on *worst case* complexity. In practical applications it is very often possible to find excellent approximations to, say, the traveling salesman problem within reasonable time.)

A famous example for a hard problem is the factoring problem already discussed in the preceding chapter. Some formulas, numbers, and references on this problem and its relation to cryptography can be found in section 3.2 of Steane’s review [5]. Since the invention of Shor’s [10] quantum factorization algorithm suspicions have grown that the factoring problem may be in class NPI (I for intermediate), that is, harder than P , but not NP-complete. If this class exists, $P \neq NP$.

Some functions may be not just hard to compute but *uncomputable* because the algorithm will never stop, or, nobody knows if it will ever stop. An example is the algorithm

while x is equal to the sum of two primes, add 2 to x , otherwise print x and halt , beginning at $x = 8$.

If this algorithm stops, we have found a counterexample to the famous Goldbach conjecture. Another famous unsolvable problem is the *halting problem*, which is stated very easily: Is there a general algorithm to decide if a Turing machine T with description (transition rules) $d[T]$ will stop on a certain input x ? There is a nice argument by Turing showing that such an algorithm does not exist. Suppose such an algorithm existed. Then it would be possible to make a Turing machine T_H which halts if and only if $T(d[T])$ (that is, T , fed its own description) does not halt:

$$T_H(d[T]) \text{ halts} \Leftrightarrow T_H(d[T]) \text{ does not halt}$$

(This is possible since the description $d[T]$ contains sufficient information about the way T works.) Now feed T_H the description of itself, that is, put $T = T_H$

$$T_H(d[T_H]) \text{ halts} \Leftrightarrow T_H(d[T_H]) \text{ does not halt.}$$

This contradiction shows that there is no algorithm that solves the halting problem. (Or that every Turing machine halts, fed its own description. That must be excluded somehow.) This is a nice recursive argument: let an algorithm find out something about its own structure. This kind of reasoning is typical of the field centered around Gödel’s incompleteness theorem; a very interesting literary piece of work centered about the ideas of recursiveness and self-reference (in mathematics and other fields of culture) is the book “Gödel, Escher, Bach” [14] by the physicist/computer scientist Douglas R. Hofstadter.