Computational Complexity

Salil Vadhan
School of Engineering & Applied Sciences
Harvard University

Synonyms

Complexity theory

Related concepts and keywords

Exponential time; O-notation; One-way function; Polynomial time; Security (Computational, Unconditional); Sub-exponential time;

Definition

Computational complexity theory is the study of the minimal resources needed to solve computational problems. In particular, it aims to distinguish between those problems that possess efficient algorithms (the "easy" problems) and those that are inherently intractable (the "hard" problems). Thus computational complexity provides a foundation for most of modern cryptography, where the aim is to design cryptosystems that are "easy to use" but "hard to break". (See security (computational, unconditional).)

Theory

Running Time. The most basic resource studied in computational complexity is *running time* — the number of basic "steps" taken by an algorithm. (Other resources, such as *space* (i.e., memory usage), are also studied, but they will not be discussed them here.) To make this precise, one needs to fix a model of computation (such as the Turing machine), but here it suffices to informally think of it as the number of "bit operations" when the input is given as a string of 0's and 1's. Typically, the running time is measured as a function of the *input length*. For numerical problems, it is assumed the input is represented in binary, so the length of an integer N is roughly $\log_2 N$. For example, the elementary-school method for adding two n-bit numbers

has running time proportional to n. (For each bit of the output, we add the corresponding input bits plus the carry.) More succinctly, it is said that addition can be solved in time "order n", denoted $O(n)$ (see O-notation). The elementary-school multiplication algorithm, on the other hand, can be seen to have running time $O(n^2)$. In these examples (and in much of complexity theory), the running time is measured in the *worst case*. That is, one measures the maximum running time over all inputs of length n.

Polynomial Time. Both the addition and multiplication algorithms are considered to be efficient, because their running time grows only mildly with the input length. More generally, polynomial time (running time $O(n^c)$ for a constant c), is typically adopted as the criterion of efficiency in computational complexity. The class of all computational problems possessing polynomial-time algorithms is denoted **P**. (Typically, **P** is defined as a class of *decision problems* (i.e. problems with a yes/no answer), but here no such restriction is made.) Thus ADDITION and MULTIPLICATION are in **P**, and more generally one thinks of **P** as identifying the "easy" computational problems. Even though not all polynomial-time algorithms are fast in practice, this criterion has the advantage of robustness: the class **P** seems to be independent of changes in computing technology. **P** is an example of a *complexity class* — a class of computational problems defined via some algorithmic constraint, in this case "polynomial time".

In contrast, algorithms that do not run in polynomial time are considered infeasible. For example, consider the *trial division* algorithms for integer factoring or primality testing (see primality test). For an n-bit number, trial division can take time up to $2^{n/2}$, which is exponential time rather than polynomial time in n. Thus, even for moderate values of n (e.g. $n = 200$) trial division of n-bit numbers is completely infeasible for present-day computers, whereas addition and multiplication can be done in a fraction of a second. Computational complexity, however, is not concerned with the efficiency of a particular algorithm (such as trial division), but rather whether a problem has *any* efficient algorithm at all. Indeed, for primality testing, there are polynomial-time algorithms known (see prime number), so PRIMALITY is in **P**. For integer factoring, on the other hand, the fastest known algorithm has running time greater than $2^{n^{1/3}}$, which is far from polynomial. Indeed, it is believed that FACTORING is not in **P**; the RSA and Rabin cryptosystems (see RSA public-key encryption, RSA digital signature scheme, Rabin cryptosystem,

Rabin digital signature scheme) rely on this conjecture. One of the ultimate goals of computational complexity is to rigorously prove such *lower bounds*, i.e. establish theorems stating that there is no polynomial-time algorithm for a given problem. (Unfortunately, to date, such theorems have been elusive, so cryptography continues to rest on conjectures, albeit widely believed ones. More on this below.)

Polynomial Security. Given the above association of "polynomial time" with feasible computation, the general goal of cryptography becomes to construct cryptographic protocols that have polynomial efficiency (i.e., can be executed in polynomial time) but super-polynomial security (i.e., cannot be broken in polynomial time). This guarantees that, for a sufficiently large setting of the *security parameter* (which roughly corresponds to the input length in complexity theory), "breaking" the protocol takes much more time than using the protocol. This is referred to as *asymptotic security.*

While polynomial time and asymptotic security are very useful for the theoretical development of the subject, more refined measures are needed to evaluate real-life implementations. Specifically, one needs to consider the complexity of using and breaking the system for fixed values of the input length, e.g. $n = 1000$, in terms of the actual time (e.g. in seconds) taken on current technology (as opposed to the "basic steps" taken on an abstract model of computation). Efforts in this direction are referred to as *concrete security.* Almost all results in computational complexity and cryptography, while usually stated asymptotically, can be interpreted in concrete terms. However, they are often not optimized for concrete security (where even constant factors hidden in <u>O-notation</u> are important).

Even with asymptotic security, it is sometimes preferable to demand that the gap between the efficiency and security of cryptographic protocols grows even more than polynomially fast. For example, instead of asking simply for super-polynomial security, one may ask for *subexponential security* (i.e. cannot be broken in time 2^{n^ϵ} for some constant $\epsilon > 0$; see <u>subexponential time</u>). Based on the current best known algorithms (see the <u>number field sieve for factoring</u>), it seems that FACTORING may have subexponential hardness and hence the cryptographic protocols based on its hardness may have subexponential security. Even better would be *exponential security*, meaning that the protocol cannot be broken in time $2^{\epsilon n}$ for some constant $\epsilon > 0$; see <u>exponential time</u>. (This refers to terminology in the

cryptography literature. In the computational complexity literature, 2^{n^ϵ} is typically referred to as exponential and $2^{\epsilon n}$ as strongly exponential.)

Complexity-Based Cryptography. As described above, a major aim of complexity theory is to identify problems that cannot be solved in polynomial time and a major aim of cryptography is to construct protocols that cannot be broken in polynomial time. These two goals are clearly well-matched. However, since proving lower bounds (at least for the kinds of problems arising in cryptography) seems beyond the reach of current techniques in complexity theory, an alternative approach is needed.

Present-day complexity-based cryptography therefore takes a *reductionist approach*: it attempts to relate the wide variety of complicated and subtle computational problems arising in cryptography (forging a signature, computing partial information about an encrypted message, etc.) to a few, simply stated assumptions about the complexity of various computational problems. For example, under the assumption that there is no polynomial-time algorithm for FACTORING (that succeeds on a significant fraction of composites of the form $n = pq$), it has been demonstrated (through a large body of research) that it is possible to construct algorithms for almost all cryptographic tasks of interest (e.g., asymmetric cryptosystems, digital signature schemes, secure multi-party computation, etc.). However, since the assumption that FACTORING is not in $\mathbf{P}$ is only a conjecture and could very well turn out to be false, it is not desirable to have all of modern cryptography to rest on this single assumption. Thus another major goal of complexity-based cryptography is to abstract the properties of computational problems that enable us to build cryptographic protocols from them. This way, even if one problem turns out to be in $\mathbf{P}$, any other problem satisfying those properties can be used without changing any of the theory. In other words, the aim is to base cryptography on assumptions that are as weak and general as possible.

Modern cryptography has had tremendous success with this reductionist approach. Indeed, it is now known how to base almost all basic cryptographic tasks on a few simple and general complexity assumptions (that do not rely on the intractability of a single computational problem, but may be realized by any of several candidate problems). Among other things, the text below discusses the notion of a *reduction* from complexity theory that is central to this reductionist approach, and the types of general assumptions, such as the existence of *one-way functions*, on which cryptography can be based.

Reductions. One of the most important notions in computational complexity, which has been inherited by cryptography, is that of a *reduction* between computational problems. A problem Π is said to reduce to problem Γ if Π can be solved in polynomial time given access to an "oracle" that solves Γ (i.e. a hypothetical black box that will solve Γ on instances of our choosing in a single time step). Intuitively, this captures the idea that problem Π is no harder than problem Γ. For a simple example, let us see that PRIMALITY reduces to FACTORING (without using the fact that PRIMALITY is in $\mathbf{P}$, which makes the reduction trivial). Suppose you have an oracle that, when fed any integer, returns its prime factorization in one time step. Then you could solve PRIMALITY in polynomial time as follows: on input N, feed the oracle with N, output "`prime`" if the only factor returned by the oracle is N itself, and output "`composite`" otherwise.

It is easy to see that if problem Π reduces to problem Γ, and $\Gamma \in \mathbf{P}$, then $\Pi \in \mathbf{P}$: if the oracle queries are substituted with the actual polynomial-algorithm for Γ, the result is a polynomial-time algorithm for Π. Turning this around, $\Pi \notin \mathbf{P}$ implies that $\Gamma \notin \mathbf{P}$. Thus, reductions give a way to use an assumption that one problem is intractable to deduce that other problems are intractable. Much work in cryptography is based on this paradigm: for example, one may take a complexity assumption such as "there is no polynomial-time algorithm for FACTORING" and use reductions to deduce statements such as "there is no polynomial-time algorithm for breaking encryption scheme X". (As discussed later, for cryptography, the formalizations of such statements and the notions of reduction in cryptography are more involved than suggested here.)

NP. Another important complexity class is $\mathbf{NP}$. Roughly speaking, this is the class of all computational problems for which solutions can be *verified* in polynomial time. ($\mathbf{NP}$ stands for *nondeterministic polynomial time*. Like $\mathbf{P}$, $\mathbf{NP}$ is typically defined as a class of decision problems, but again that constraint is not essential for our informal discussion.) For example, given that PRIMALITY is in $\mathbf{P}$, one can easily see that FACTORING is in $\mathbf{NP}$: to verify that a supposed prime factorization of a number N is correct, simply test each of the factors for primality and check that their product equals N. $\mathbf{NP}$ can be thought of as the class of "well-posed" search problems: it is not reasonable to search for something unless you can recognize when you have found it. Given this natural definition, it is not surprising that the class $\mathbf{NP}$

has taken on a fundamental position in computer science.

It is evident that $\mathbf{P} \subseteq \mathbf{NP}$, but whether or not $\mathbf{P} = \mathbf{NP}$ is considered to be one of the most important open problems in mathematics and computer science. It is widely believed that $\mathbf{P} \neq \mathbf{NP}$, indeed, FACTORING is one candidate for a problem in $\mathbf{NP} \setminus \mathbf{P}$. In addition to FACTORING, $\mathbf{NP}$ contains many other computational problems of great importance, from many disciplines, for which no polynomial-time algorithms are known.

The significance of $\mathbf{NP}$ as a complexity class is due in part to the $\mathbf{NP}$-*complete* problems. A computational problem Π is said to be $\mathbf{NP}$-complete if $\Pi \in \mathbf{NP}$ and every problem in $\mathbf{NP}$ reduces to Π. Thus the $\mathbf{NP}$-complete problems are the "hardest" problems in $\mathbf{NP}$, and are the most likely to be intractable. (Indeed, if even a single problem in $\mathbf{NP}$ is not in $\mathbf{P}$, then all the $\mathbf{NP}$-complete problems are not in $\mathbf{P}$.) Remarkably, thousands of natural computational problems have been shown to be $\mathbf{NP}$-complete. (See [2].) Thus, it is an appealing possibility to build cryptosystems out of $\mathbf{NP}$-complete problems, but unfortunately, $\mathbf{NP}$-completeness does not seem sufficient for cryptographic purposes (as discussed later).

Randomized Algorithms. Throughout cryptography, it is assumed that parties have the ability to make random choices; indeed this is how one models the notion of a secret key. Thus, it is natural to allow not just algorithms whose computation proceeds deterministically (as in the definition of $\mathbf{P}$), but also consider *randomized algorithms* — ones that may make random choices in their computation. (Thus, such algorithms are designed to be implemented with a physical source of randomness. See random bit generation (hardware).)

Such a randomized (or *probabilistic*) algorithm A is said to solve a given computational problem if on every input x, the algorithm outputs the correct answer with high probability (over its random choices). The error probability of such a randomized algorithm can be made arbitrarily small by running the algorithm many times. For examples of randomized algorithms, see the probabilistic primality tests in the entry on prime number. The class of computational problems having polynomial-time randomized algorithms is denoted $\mathbf{BPP}$ (which stands for "bounded-error probabilistic polynomial time"). A widely believed strengthening of the $\mathbf{P} \neq \mathbf{NP}$ conjecture is that $\mathbf{NP} \not\subseteq \mathbf{BPP}$.

P vs. NP and Cryptography. The assumption $\mathbf{P} \neq \mathbf{NP}$ (and even $\mathbf{NP} \not\subseteq \mathbf{BPP}$) is *necessary* for most of modern cryptography. For example, take any efficient encryption scheme and consider the following computational problem: given a ciphertext C, find the corresponding message M along with the key K and any randomization R used in the encryption process. This is an $\mathbf{NP}$ problem: the solution (M, K, R) can be verified by re-encrypting the message M using the key K and the randomization R and checking whether the result equals C. Thus, if $\mathbf{P} = \mathbf{NP}$, this problem can be solved in polynomial time, i.e. there is an efficient algorithm for breaking the encryption scheme. (Technically, to conclude that the cryptosystem is broken requires that the message M is uniquely determined by ciphertext C. This will be the case for most messages if the message length is greater than the key length. If the message length is less than or equal to the key length, then there exist encryption schemes that achieve information-theoretic security for a single encryption, e.g. the one-time pad, regardless of whether or not $\mathbf{P} = \mathbf{NP}$. See Shannon's model.)

However, the assumption $\mathbf{P} \neq \mathbf{NP}$ (or even $\mathbf{NP} \not\subseteq \mathbf{BPP}$) does not appear *sufficient* for cryptography. The main reason for this is that $\mathbf{P} \neq \mathbf{NP}$ refers to *worst-case complexity*. That is, the fact that a computational problem Π is not in $\mathbf{P}$ only means that for every polynomial-time algorithm A, there *exist* inputs on which A fails to solve Π. However, these "hard inputs" could conceivably be very rare and very hard to find. Intuitively, to make use of intractability (for the security of cryptosystems), one needs to be able to efficiently generate hard instances of an intractable computational problem.

One-way functions. The notion of a one-way function captures the kind of computational intractability needed in cryptography. Informally, a one-way function is a function f that is "easy to evaluate" but "hard to invert". That is, it is required that the function f can be computed in polynomial time, but given $y = f(x)$, it is intractable to recover x. The difficulty of inversion is required to hold even when the input x is chosen *at random*. Thus, one can efficiently generate hard instances of the problem "find a preimage of y", by selecting x at random and setting $y = f(x)$. (Note that this process actually generates a hard instance together with a solution; this is another way in which one-way functions are stronger than what follows from $\mathbf{P} \neq \mathbf{NP}$.) To formalize the definition, one needs the concept of a *negligible function*. A function $\epsilon : \mathbb{N} \to [0, 1]$ is negligible if for every constant c, there

is an n_0 such that $\epsilon(n) \leq 1/n^c$ for all $n \geq n_0$. That is, ϵ vanishes faster than the reciprocal of any polynomial. Then the definition is as follows:

Definition 1 (one-way function) *A one-to-one function f is* one-way *if it satisfies the following conditions.*

1. *(Easy to evaluate) f can be evaluated in polynomial time.*

2. *(Hard to invert) For every probabilistic polynomial-time algorithm A, there is a negligible function ϵ such that*

$$\Pr[A(f(X)) = X] \leq \epsilon(n),$$

where the probability is taken over selecting a input X of length n uniformly at random and the random choices of the algorithm A.

For simplicity, the definition above is restricted to one-to-one one-way functions. Without the one-to-one constraint, the definition should refer to the problem of finding *some* preimage of $f(X)$, i.e. require the probability that $A(f(X)) \in f^{-1}(f(X))$ is negligible. (For technical reasons, it is also required that f does not shrink its input too much, e.g. that the length of $|f(x)|$ and length of $|x|$ are polynomially related (in both directions.)

The length n of the input can be thought of as corresponding to the *security parameter* (or *key length*) in a cryptographic protocol using f. If f is one-way, it is guaranteed that by making n sufficiently large, inverting f takes much more time than evaluating f. However to know how large to set n in an implementation requires a concrete security analogue of the above definition, where the maximum success probability ϵ is specified for A with a particular running time on a particular input length n, and a particular model of computation.

The "inversion problem" is an **NP** problem (to verify that X is a preimage of Y, simply evaluate $f(X)$ and compare with Y). Thus, if **NP** $\subseteq$ **BPP** then one-way functions do not exist. However, the converse is an open problem, and proving it would be a major breakthrough in complexity theory. Fortunately, even though the existence of one-way functions does not appear to follow from **NP** $\not\subseteq$ **BPP**, there are a number of natural candidates for one-way functions.

Some Candidate One-Way Functions. These examples are described informally, and may not all match up perfectly with the simplified definition above. In particular, some are actually *collections of one-way functions* $\mathcal{F} = \{f_i : \mathcal{D}_i \to \mathcal{R}_i\}$, in the functions f_i are parameterized by an index i that is generated by some randomized algorithm. (Actually, one can convert a collection of one-way functions into a single one-way function, and conversely. See [4].)

1. (Multiplication) $f(p, q) = p \cdot q$, where p and q are primes of equal length. Inverting f is the FACTORING problem (see integer factoring, which indeed seems intractable even on random inputs of the form $p \cdot q$).

2. (Subset Sum) $f(x_1, \ldots, x_n, S) = (x_1, \ldots, x_n, \sum_{i \in S} x_i)$. Here each x_i is an n-bit integer and $S \subseteq [n]$. Inverting f is the SUBSET SUM problem (see knapsack cryptographic schemes). This problem is known to be **NP**-complete, but for the reasons discussed above, this does not provide convincing evidence that f is one way (nevertheless it seems to be so).

3. (The Discrete Log Collection) $f_{G,g}(x) = g^x$, where G is a cyclic group (e.g. $G = \mathbb{Z}_p^*$ for prime p), g is a generator of G, and $x \in \{1, \ldots, |G|-1\}$. Inverting $f_{G,g}$ is the DISCRETE LOGproblem (see discrete logarithm problem), which seems intractable. This (like the next two examples) is actually a collection of one-way functions, parametrized by the group G and generator g.

4. (The RSA Collection) $f_{n,e}(x) = x^e \bmod n$, where n is the product of two equal-length primes, e satisfies $\gcd(e, \phi(n)) = 1$, and $x \in \mathbb{Z}_n^*$, . Inverting $f_{n,e}$ is the RSA problem.

5. (Rabin's Collection (see Rabin cryptosystem, Rabin digital signature scheme) $f_n(x) = x^2 \bmod n$, where n is a composite and $x \in \mathbb{Z}_n^*$. Inverting f_n is known to be as hard as factoring n.

6. (Hash functions & block ciphers) Most cryptographic hash functions seem to be *finite* analogues of one-way functions with respect to *concrete* security. Similarly, one can obtain candidate one-way functions from block ciphers, say by defining $f(K)$ to be the block cipher applied to some fixed message using key K.

In a long sequence of works by many researchers, it has been shown that one-way functions are indeed the "right assumption" for complexity-based cryptography. On one hand, almost all tasks in cryptography imply the existence of one-way functions. Conversely (and more remarkably), many useful cryptographic tasks can be accomplished given any one-way function.

Theorem 1 *The existence of one-way functions is necessary and sufficient for each of the following:*

- *The existence of <u>commitment</u> schemes.*

- *The existence of <u>pseudo-random number generators</u>.*

- *The existence of <u>pseudorandom functions</u>.*

- *The existence of <u>symmetric cryptosystems</u>.*

- *The existence of <u>digital signature schemes</u>.*

These results are proven via the notion of reducibility mentioned above, albeit in much more sophisticated forms. For example, to show that the existence of one-way functions implies the existence of pseudorandom generators, one describes a general construction of a pseudorandom generator G from any one-way function f. To prove the correctness of this construction, one shows how to "reduce" the task of inverting the one-way function f to that of "distinguishing" the output of the pseudorandom generator G from a truly random sequence. That is, any polynomial-time algorithm that distinguishes the pseudorandom generator can be converted into a polynomial-time algorithm that inverts the one-way function. But if f is one-way, it cannot be inverted, implying that the pseudorandom generator is secure. These reductions are much more delicate than those arising in, say, the **NP**-completeness, because they involve non-traditional computational tasks (e.g., inversion, distinguishing) that must be analyzed in the average case (i.e. with respect to nonnegligible success probability).

The general constructions asserted in Theorem 1 are very involved and not efficient enough to be used in practice (though still polynomial time), so it should be interpreted only as a "plausibility result". However, from special cases of one-way functions, such as one-way permutations (see <u>one-way function</u>) or some of the specific candidate one-way functions mentioned earlier, much more efficient constructions are known.

Trapdoor Functions. For some tasks in cryptography, most notably public-key encryption (see public-key cryptography), one-way functions do not seem to suffice, and additional properties are used. One such property is the *trapdoor* property, which requires that the function *can* be easily inverted given certain "trapdoor information". What follows is not the full definition, but just a list of the main properties. (See also trapdoor one-way function.)

Definition 2 (trapdoor functions, informal) *A collection of one-to-one functions $\mathcal{F} = \{f_i : \mathcal{D}_i \to \mathcal{R}_i\}$ is a collection of* trapdoor functions *if*

1. *(Efficient generation) There is a probabilistic polynomial-time algorithm that, on input a security parameter n, generates a pair (i, t_i), where i is the index to a (random) function in the family and t_i is the associated "trapdoor information".*

2. *(Easy to evaluate) Given i and $x \in \mathcal{D}_i$, one can compute $f_i(x)$ in polynomial time.*

3. *(Hard to invert) There is no probabilistic polynomial-time algorithm that on input $(i, f_i(x))$ outputs x with nonnegligible probability. (Here, the probability is taken over i, $x \in \mathcal{D}_i$, and the coin tosses of the inverter.)*

4. *(Easy to invert with trapdoor) Given t_i and $f_i(x)$, one can compute x in polynomial time.*

Thus, trapdoor functions are *collections* of one-way functions with an additional trapdoor property (Item 4). The RSA and Rabin collections described earlier have the trapdoor property. Specifically, they can be inverted in polynomial time given the factorization of the modulus n.

One of the main applications of trapdoor functions is for the construction of public-key encryption schemes.

Theorem 2 *If trapdoor functions exist, then public-key encryption schemes exist.*

There are a number of other useful strengthenings of the notion of a one-way function, discussed elsewhere in this volume: claw-free permutations, collision-resistant hash functions (see collision resistance), and universal one-way hash functions.

Other Interactions with Cryptography. The interaction between computational complexity and cryptography has has been very fertile. The text above describes the role that computational complexity plays in cryptography. Conversely, several important concepts that originated in cryptography research have had a tremendous impact on computational complexity. Two notable examples are the notions of pseudo-random number generators and interactive proof systems.

Open Problems

Computational complexity has a vast collection of open problems. The discussion above touched upon three that are particularly relevant to cryptography:

- Does $\mathbf{P} = \mathbf{NP}$?

- Is FACTORING in $\mathbf{BPP}$?

- Does $\mathbf{NP} \not\subseteq \mathbf{BPP}$ imply the existence of one-way functions?

Acknowledgments.

References

[1] Sanjeev Arora and Boaz Barak. *Computational complexity.* Cambridge University Press, Cambridge, 2009. A modern approach.

[2] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* W.H. Freeman and Company, 1979.

[3] Oded Goldreich. *Computational complexity: a conceptual perspective.* Cambridge University Press, Cambridge, 2008.

[4] Oded Goldreich. *Foundations of cryptography.* Cambridge University Press, Cambridge, 2001. Basic tools.

[5] Oded Goldreich. *Foundations of cryptography. II.* Cambridge University Press, Cambridge, 2004. Basic Applications.

[6] Michael Sipser. *Introduction to the Theory of Computation.* PWS Publishing, 1997.