УДК

R. Clark, S. Callejas, L. Medina, D. Robinson, N. Torres, V. Kreinovich
University of Texas at El Paso, El Paso, Texas, USA
M. Zakharevich
SeeCure Systems, Inc., San Jose, California, USA

ARE PRACTICAL NP-HARD PROBLEMS REALLY HARD?

Many practical problems are NP-hard. In most practical problems, we need to select the best possible alternative.

Each possible alternative can be described by the values of the corresponding parameters. The relative quality of different alternatives can be described by an objective function. In these terms, finding the best possible alternative means finding the values of the parameters for which the objective function attains the largest possible value – i.e., attains its global maximum.

Such global optimization problems are, in general, NP-hard – which means that unless P = NP (which most computer scientists do not believe to be possible), no feasible algorithm is possible that would solve all particular cases of this problem.

NP-hardness was also proven for many specific classes of optimization problems. In this sense, many practical problems are NP-hard.

But are these problems really hard? Sometimes, practical problems are indeed hard to solve, but in many cases, they are reasonably easily solved. How can we explain this seeming contradiction between the theoretical result – that everything is hard – and the practical observation – that most problems are feasibly solvable?

Our explanation: idea. The usual definitions underlying the notion of NP-hardness implicitly assume that we solve the problem “from scratch”: we have the formulation of the problem, we know nothing else, and we need to come up with a solution.

Such situations really happen, especially when we go into the unknown – when we plan a spaceflight into an area that has not yet been explored, when we design a completely new type of device, etc.

However, such “from scratch” situations are rare. For example, in control, the situation usually changes gradually, so by the time we need a control decision, we have, at our disposal, optimal or close-to-optimal solutions to very similar optimization problems that we have solves in the past. In such situations, we do not need to start from scratch: we can start with the previous solution and modify it a little bit.

Let us show that this indeed explains why many practical problems are feasible.

An important fact that we use in our explanation. An important fact that we will use in our explanation is that in almost all cases – almost all in some reasonable sense – the global maximum is attained at exactly one location.

This fact is, actually, a technically complex result using different probability distributions on the class of all objective functions. However, as we will show, this result also has a commonsense explanation that does not require us to deal with not-very-intuitive technical details of the corresponding mathematics.

Commonsense explanation of the above important fact. Let D denote the domain on which the objective function f(x) is defined. Let us divide this domain into two disjoint parts D_1 and D_2. Let m_1 and m_2 denote the maxima of the objective function f(x) on each of these domains. From the purely mathematical viewpoint, it is possible that these maxima are equal. However, from the commonsense viewpoint, the probability that the two values corresponding to two different subdomains should be equal is zero – e.g., if we measure the heights of two different mountains, it would be improbable that we end up with exactly the same number. Maybe we will get close numbers, but more accurate measurements will show the difference.

We can thus dismiss such improbable cases when the objective function attains the exact same maximum on two different subdomains. In this case, if we have a location x_M at which f(x) attains its global maximum M, and we divide the domain D into two parts one of which contains this location, the maximum on the remaining part should be different from M – and thus, smaller than M.

This indeed implies that the maximum is attained in only one location.

Explanation continues. The objective function corresponding to the next moment of time is close to the objective function corresponding to the previous moment of time – for which we already know the solution.

The closer the moments of time, the closer the objective functions. It is possible to show that if a new function g(x) is sufficiently close to the function f(x) that attains its global maximum at a single point x_M, then the global maximum of g(x) will also be attained close to x_M.

It is known that for smooth objective functions – and objective functions are usually smooth -- finding the maximum in a small vicinity of a point is a feasible task.

Conclusion. Thus, we have indeed explained why many practical problems are indeed feasible.