Towards an Optimal Approach to Soft Constraint Problems
(Extended Abstract)

Martine Ceberio and Vladik Kreinovich
Department of Computer Science
University of Texas at El Paso, 500 W. University
El Paso, TX 79968, USA, {mceberio,vladik}@cs.utep.edu

Why constraints? In many areas of science and engineering, we are interested in solving design and control problems. Usually, in these problems, the users describe several *constraints* that the desired design or control must satisfy, and our objective is to find a design (correspondingly, a control) that satisfies all these constraints.

In mathematical terms, a design or a control can be usually represented by the values of the relevant numerical parameters $x = (x_1, \ldots, x_n)$. For example, an airplane design can be described in terms of the geometric parameters of the plane, the thickness of the plates that form the airplane's skin, the weight and power of the engine, etc. A typical constraint describes a limitation on some characteristics of this design: e.g., the airplane's speed must exceed a certain threshold, its fuel use must not exceed a certain amount, and the overall cost must be within given limits. Each of the corresponding characteristics y (speed, fuel use, etc.) can be uniquely determined by the design $y = f(x_1, \ldots, x_n)$ for some computable function f. Thus, each constraint can be described as either an inequality of the type $f(x_1, \ldots, x_n) \le y_0$ or $f(x_1, \ldots, x_n) \ge y_0$ – or as an equality $f(x_1, \ldots, x_n) = y_0$ (for example, if we want to design the fastest airplane within a fixed cost).

In many real-life situations, the constraints are consistent, i.e., there exist designs that satisfy all these constraints. In this case, it is desirable to find a design that satisfies a given finite set of inequality/equality-type constraints. The corresponding problem is called the problem of *constraint satisfaction*.

Often, in such situations, there are many different designs that satisfy the given constraints. In this case, it is desirable to select one of these designs. Users can often describe their preference in terms of an *objective function* $g(x_1, \ldots, x_n)$ whose value should be made as large as possible. In such situations, we are interested in maximizing the given function $g(x_1, \ldots, x_n)$ under the given constraints. This problem is called a *constrained optimization problem*.

In general, constraint satisfaction and constraint optimization problems are NP-hard; see, e.g., [3, 6]. In practice, however, there exist many efficient tools for solving these problems, including numerous efficient tools that provide validated solution to these problems.

What are soft constraints? In many practical situations, if we formulate all the users' desired as constraints, we often end up with an inconsistent set of constraints. For example, a user may want to design a plane that is as fast and as fuel-efficient as the existing Airbus or Boeing planes, but that will decrease the noise level to 0.

The reason why constraints are inconsistent is that while some of these constraints are absolute requirements (e.g., safety constraints for a plain), other onstraints are simply recommendations, desires, that a user wants to be implemented if possible – but that can be dismissed if it is not possible to satiafy them. Such "not required" constraints are called *soft constraints*.

Soft constraints are an important research topic; see, e.g., the proceedings of the latest conference [4] and references therein. Our own work in soft constraints is described, e.g., in [1, 2].

Priority approach to soft constraints: a brief description. The main idea behind this approach is that if we cannot satisfy all the constraints, we should at least satisfy as many constraints as possible.

One of the natural approaches to soft constraints is therefore to ask the user to *prioritize* their constraints, from the absolutely required to the less required. Once the user sorted all his/her constraints from the most required to the least required, into a sequence $C_1 \succ C_2 \succ \ldots \succ C_n$, we try to find the largest possible value $k = k_{\rm opt}$ for which all the constraints $C_1, C_2, \ldots, C_k$ are still consistent.

Priority approach to soft constraints: a computational question. The existing constraint satisfaction tools enable us, given constraints, either to find a design that satisfies these constraints, or to conclude that the given constraints are inconsistent.

There are many different ways how we can use one of these tools to solve soft constraint problems. For example, we can sequentially apply this tool to constraints sets $\{C_1\}$, $\{C_1, C_2\}$, $\ldots$, until we find the first value l for which the corresponding set is inconsistent. Then, the previous value $k_{\rm opt} = l - 1$ is the desired largest k, and the corresponding design is the desired one.

A (potential) problem with this approach is that when the number of constraints is large, the constraint satisfaction tools take a long time to run, so if we have to run the tool for many different sets, we make the process even slower.

Alternatively, we can, e.g., use bisection to find the desired value $k_{\rm opt}$. At each stage of this iterative method, we have an interval $[k^-, k^+]$ that is guaranteed to contain this desired value, i.e., for which the system of the first k^- constraints is consistent while the system of the first k^+ constraints is inconsistent.

Initially, $k^- = 0$ (if we have no constraints, then, of course, the problem is consistent) and $k^+ = n$ (we cannot satisfy all constraints: this is the definition of the soft constraint problem).

Then, sequentially, we test whether the midpoint $k_m \stackrel{\rm def}{=} \lfloor (k^- + k^+)/2 \rfloor$ of the interval is consistent or not, and, depending on the result of this test, replace the original interval with a half-size one: $[k^-, k_m]$ or $[k_m, k^+]$. On each iteration, the interval decreases in half, so after $\log_2(n)$ iteration, we get the interval of width 1 – i.e., we get the desired value $k_{\rm opt}$. The advantage of this technique is that we need fewer iterations to find k, but the disadvantage is that some of these iterations may require analyzing much more constraints – n as opposed to $k \ll n$ – and thus, may take much longer.

Other methods of finding the optimal k are possible. The question is: which of the methods is optimal?

Of course, different methods may be optimal for different cases. What we would like to do is find the methods for which some reasonably defined "worst-case" computation time is optimal. Let us formulate this problem in precise terms.

Toward formalizing this question. We would like to describe the best strategy of finding $k_{\rm opt}$. In order to do that, we need to explain what we mean by a strategy.

All we can do is check, for several values k_i, whether the corresponding constraint sets are consistent. If we know that constraints sets are consistent for the values $k_1, \ldots, k_m$ and inconsistent for $k = k'_1, \ldots, k'_p$, then this knowledge is equivalent to knowing that the constraints are consistent for $k^- \stackrel{\rm def}{=} \max(k_1, \ldots, k_m)$ and inconsistent for $k^+ \stackrel{\rm def}{=} \min(k'_1, \ldots, k'_p)$.

In other words, at each stage of the computations, we know that k is within an interval $[k^-, k^+]$ – i.e., that:

- the set of the first k^- inequalities is consistent, while
- the set of the first k^+ inequalities is inconsistent.

If $k^+ = k^- + 1$, then we know that k^- is the desired largest value for which constraints are consistent. If $k^+ > k^- + 1$, then, to continue looking for $k_{\rm opt}$, we must select a value within the interval $[k^-, k^+]$ that we will check next. This value can, in general, depend on the number of a step.

So, in general, we can define the *method* as follows:

Definition 1. *A* method *is a mapping that maps each pair* (I, s)*, where:*

- I *is an integer-valued interval* $[k^-, k^+]$*, where* $0 \le k^- < k^+ \le n$ *and* $k^+ > k^- + 1$*, and*
- s *is a positive integer (= number of step)*

into an integer $k_{\rm next}$ *from the open interval* (k^-, k^+)*.*

One can easily see that each method generates a strategy that eventually leads to the desired value $k_{\rm opt}$: once we check whether the system of constraints is consistent or not, we get a new interval $[k^-, k_{\rm next}]$ or $[k_{\rm next}, k^+]$. For example:

- the sequential search S corresponds to the method in which we select $k_{\rm next} = k^- + 1$;
- bisection B corresponds to the method in which we select $k_{\rm next} = \lfloor (k^- + k^+)/2 \rfloor$.

How can we estimate the worst-case complexity of each method? As we have mentioned, the problem of constraint satisfaction is known to be NP-hard. This means, crudely speaking, that the computation time grows exponentially with the number of constraints. For most NP-hard problems such propositional satisfiability, the actual worst-case complexity of the known algorithms grows as 2^n. For Boolean-type constraints, it is therefore reasonable to assume that the computational complexity of checking the consistency of the system of k constraints is proportional to 2^k.

Similarly, in a more general case of finite constraints in which each variable has $p \ge 2$ different values, we need, in the worst case, p^k checks to check all possible values of these constraint variables. It is therefore reasonable to assume that in general, the computational complexity of checking the consistency of the system of k constraints grows as p^k for some $p \ge 2$.

For every method M and for every value k of $k_{\rm opt}$, we can now define the overall time $T_M(k)$ that this method spends on this case as the sum of the values $p^{k'}$ for all k' for which this method checks consistency.

For each value of $k_{\rm opt}$, the ideal case is when someone's intuition informs us of the correct value $k_{\rm opt}$. In this case, we do not need to check for too many different values k, all we need to do is check whether indeed the system of first k constraints is consistent while the system of the first $k+1$ constraints is inconsistent. In this case, the overall expenses are equal to $p^k + p^{k+1}$.

In real life, no such intuition is available, so we have to test more values k and thus, spend more time. The smaller the "overhead" in comparison with the ideal case, the better. We can therefore define the *overhead* $O_p(M)$ of a method M as

$$O_p(M) \stackrel{\rm def}{=} \max_k \frac{T_M(k)}{p^k + p^{k+1}}.$$

Let $p \ge 2$ be fixed. We say that a method M is *optimal* for this p if out of all methods, it has the smallest possible overhead: $O_p(M_{\rm opt}) = \min\limits_M O_p(M)$. Now, we are ready to describe our main result:

Theorem. *For every* $p \ge 2$*, the sequential search method* S *is optimal.*

Proof: main idea. For sequential search, if the actual value of $k_{\rm opt}$ is k, we check consistency of sets consisting of 1, 2, ..., k, and $k+1$ constraints. As a result, the overall time is equal to

$$T_S(k) = p + p^2 + \ldots + p^{k+1}.$$

Here,

$$T_M(k) = p^{k+1} \cdot (1 + p^{-1} + p^{-2} + \ldots + p^{-k}) < p^{k+1} \cdot (1 + p^{-1} + p^{-2} + \ldots) = \frac{p^{k+1}}{(1 - p^{-1})}.$$

Since $p^k + p^{k+1} = p^{k+1} \cdot (1 + p^{-1})$, we conclude that

$$O_S(p) < \frac{1}{(1 - p^{-1}) \cdot (1 + p^{-1})}.$$

Let us now show that every other method has a larger overhead. Indeed, the specific feature of S is that in S, out of all possible integers between k^- and k^+, we always select $k_{\rm next} = k^- + 1$ as the next value to check. This means that in every other method $M \neq S$, there exists an interval in which we select a value $k_{\rm next} > k^- + 2$. In this case, if the actual value $k_{\rm opt}$ is equal to the corresponding k^-, then in this method, we check both k and $\geq k+2$. Later on, we still need to check the value $k+1$ – to make sure that k is indeed the largest consistent value. Thus, for this k, the method M spends at least time $T_M(k) \geq p^k + p^{k+1} + p^{k+2}$. This lower bound can be described as $T_M(k) \geq p^{k+1} \cdot (p + 1 + p^{-1})$, hence

$$O_M(k) = \frac{T_M(k)}{p^{k+1} \cdot (1 + p^{-1})} \geq \frac{p + 1 + p^{-1}}{1 + p^{-1}}.$$

To complete our proof, we must show that

$$\frac{p + 1 + p^{-1}}{1 + p^{-1}} \geq \frac{1}{(1 - p^{-1}) \cdot (1 + p^{-1})}.$$

Multiplying both sides of this inequality by the denominator of the right-hand side, we get an equivalent inequality $p - p^{-2} \geq 1$, i.e., equivalently, that $p^3 - p^2 - 1 \geq 0$. One can show that the equation $p^3 - p - 1 = 0$ has a solution $p_0 \approx 1.47$, and that for $p \geq p_0$, its left-hand side is an increasing function – since its derivative is $3p^2 - 1 \geq 3p_0^2 - 1 > 0$. Thus, $p^3 - p - 1 > 0$ for all $p \geq p_0$ – in particular, for all $p \geq 2$. Q.E.D.

Comment. From the purely mathematical viewpoint, the above problem is very similar to the following planning problem: the existing AI-based planners either find a plan of given length k or conclude that such a plan is impossible. Based on such a planner, what is the best way to find the shortest plan?

A partial solution to this problem is given in the paper [5]; in this paper, we also consider the cases when $p < 2$ and when, instead of optimizing the worst-case overhead, we optimize the average-case overhead. It is desirable to extend our soft constraint results to similar cases.

Acknowledgments. This work was supported in part by NASA under cooperative agreement NCC5-209, by NSF grants EAR-0112968, EAR-0225670, and EIA-0321328, and by NIH grant 3T34GM008048-20S1. This research was partly done when M.C. was a Visiting Researcher at NII, Tokyo, Japan.

The authors want to thank Seetharami R. Seelam for his help.

References

[1] F. Benhamou and M. Ceberio, "Soft Constraints: A Unifying Framework applied to Continuous Soft Constraints", *Proceedings of the ERCIM/CoLogNET Workshop*, 2003.

[2] M. Ceberio, *Under- and over-constrained numerical CSP: symbolic tools and soft constraints*, Ph.D. Dissertation, Nantes, France, 2003.

[3] V. Kreinovich, A. Lakeyev, J. Rohn, and P. Kahl, *Computational complexity and feasibility of data processing and interval computations*, Kluwer, Dordrecht, 1997.

[4] Proceedings of the 6th International Workshop on Preferences and Soft Constraints, Held in conjunction with 10th International Conference on Principles and Practice of Constraint Programming CP'2004, September 27–October 1, 2004, Toronto, Canada.

[5] R. A. Trejo, J. Galloway, C. Sachar, V. Kreinovich, C. Baral, and L. C. Tuan, "From Planning to Searching for the Shortest Plan: An Optimal Transition", *International Journal of Uncertainty, Fuzziness, Knowledge-Based Systems (IJUFKS)*, 2001, Vol. 9, No. 6, pp. 827–838.

[6] S. A. Vavasis, *Nonlinear Optimization: Complexity Issues*, Oxford University Press, New York, 1991.