

**Constraint Programming**

White Paper

**Constraint Programming**

White Paper



August, 2011  
[www.mu-sigma.com](http://www.mu-sigma.com)

# Constraint programming

Integer programming (IP) and mixed integer programming (MIP) have served the operations research (OR) community well for many years. They have been traditionally used to help organizations make extremely strategic decisions such as, “Should I open a new refinery at location X or location Y?” to very tactical, yet complex, questions such as, “Which routes should I operate to minimize my vehicle routing costs?” or “How should a baseball season be scheduled such that all teams have an equal number of home games?”

Thanks to the advances in IP/MIP solver technology and availability of cheap computing power, the optimization community has been able to solve most of these questions to near optimality. However, many other questions remain intractable for IP/MIP formulations because of their computational complexity. Optimization of these problems by nature requires a multi-faceted approach.

Some classes of optimization problems, such as linear programming models, can be solved by an all-purpose method, but most problems need individual attention. Typically, the secret to solving a problem is to take advantage of its particular structure. The result is a proliferation of optimization methods. This imposes an obvious burden on anyone who seeks the right method and software to solve a given problem. Also, the various disciplines (such as IP, constraint programming, heuristic optimization, etc.) continue to move in largely separate orbits. This not only imposes the inconvenience of becoming familiar with multiple solvers, but it passes up the advantages of integrated problem solving.

Recent research has shown that there is much to be gained by exploiting the complementary strengths of different approaches to optimization. Mathematical programmers are experts at relaxation techniques and polyhedral analysis. Constraint programming (CP) is distinguished by its inference techniques and modeling power. Rather than choose between these, one would like to have them both available to attack a given problem. Some problems submit to a single approach, but others benefit from the more flexible modeling and orders-of-magnitude speedup in computation that can occur when ideas from different fields are combined. The advantages of integrated methods are being demonstrated in terms of the several problems such as sports tournament scheduling, just in time (JIT) job scheduling and chemical process optimization that could not be solved easily earlier but can now be solved in a reasonable timeframe using integrated methods.

Functionality	CP	IP
Sequencing and Scheduling	→	↓
Ability to handle messy constraints	→	↓
Succinctness of model	→	↓
Constraints convey problem structure?	↑	↓
Ability to handle very large number of variables?	↓	→
Robustness of state of the art software	↓	↑
Solution gap easily available?	↓	↑

## Need for Integrating IP and CP

CP techniques are good at “inference tasks” that involve reducing the domain of the variables. However, they lack the support of relaxations that IP has. Moreover, IP methodology also benefits from the extensive use of duality theory. Each technique has its strong and weak points. It is up to the O.R. practitioner to skillfully combine these techniques in a way that they complement each other.

Traditionally, scheduling and sequencing problems in the manufacturing sector have been solved using IP. However, a glance at Pochet’s [Pochet & Wolsey, 2006] standard monograph on IP for scheduling indicates that these problems seldom have a structure that is favorable. In most cases, the complex business constraints so essential in real-life applications obscure the nice structure of the bare problem. CP has special constraints that perfectly map to these business constraints, and therefore there are advantages not only in solving the problem but also in modeling it.

Let’s take an example from the airline industry where there are a certain number of planes arriving at an airport at the same time. Each of these planes has to go to a different gate. Representing this constraint using IP requires some skill, as we have to look at each pair of flights and then say that of the pair only one flight can be assigned to a particular gate. However, CP has an “all different” constraint that is designed to tackle such situations. It simply says that each of these planes has to go to a different gate – this is as close to plain English as possible.

CP on the other hand suffers from limitations in terms of the number of variables it can handle. Large numbers of variables present in a constraint make propagation more difficult. This is not only a limitation of the current software that is used for CP but also a limitation of the methodology in some sense. IP methodology has the ability to handle literally billions of variables through the use of techniques like column generation. IP techniques also benefit from the concepts of duality and LP relaxations that help provide best bounds on the solution and thus help us track how far from the goal one is.

Another benefit that IP has over CP currently is its legacy: IP techniques have been around for many more years than CP, and this has led to the development of robust IP solvers that can solve industry-level problems rapidly. However, with O.R. practitioners’ interest in CP growing, constraint-based solvers will soon catch up.

## How IP and CP combine to produce a more powerful approach

Aside from computational advantages, an integrated approach provides a richer modeling environment that can result in simpler models and less debugging. The full repertoire of global constraints used in CP is potentially available, as are nonlinear expressions used in continuous global optimization. Frequent use of metaconstraints not only simplifies the model but also allows the solver to exploit problem structure. Effective integrated modeling draws from a large library of metaconstraints and presupposes that the user has some familiarity with this library. For instance, the library may contain a constraint that defines piecewise linear costs, a constraint that requires flow balance in a network, a constraint that prevents scheduled jobs from overlapping and so forth. Each constraint is written with parameters that specify the shape of the function, the structure of the network or the processing times of the jobs. When sitting down to formulate a model, the user would browse the library for constraints that appear to relate to the problem at hand. Integrated modeling therefore places the burden of identifying problem structure on the user, and in so doing it takes full advantage of the human capacity for pattern recognition. Users identify highly structured subsets of constraints, which allow the solver to apply the best-known analysis of these structures to solve the problem

efficiently. In addition, only certain metaconstraints tend to occur in a given problem domain. This means that only a relevant portion of the library must be presented to a practitioner in that domain.

## Practical Applications

IP and CP have already been used in an integrated fashion to solve a wide variety of practical problems. In some cases the results have been dramatic. Problems that could not be solved by IP or CP alone could be solved using an integrated approach. In other cases there has been a significant reduction in run-time or a huge improvement in solution quality due to an integrated approach.

One of the first applications of this integrated approach was to schedule Major League Baseball games in the United States. The business constraints on the sequence of home and away matches made the problem intractable to a purely IP-based approach. So the problem was split into a “master problem” and a “sub-problem.” IP techniques were used to solve the master problem while CP was used for the sub-problem. This combined strategy helped to solve the problem involving eight teams for the first time.

Another interesting application is automatic recording of broadcast content. Such a recorder tries to match a given user profile with the information submitted by the different TV channels. For example, a user may be interested in thrillers, the more recent the better. The digital video recorder is supposed to record programs such that the user’s satisfaction is maximized. As the number of channels may be enormous (more than 100 digital channels are possible), a service that automatically provides an individual selection is highly appreciated. In this context, two restrictions have to be met. First, the storage capacity is limited, and second, only one video can be recorded at a time. An integrated approach using relaxed solutions from IP along with CP is up to 10 times faster than MIP, which in turn outperforms pure CP.

Scheduling is one area where CP is known to be stronger than IP. In traditional job-scheduling problems, there is a penalty only for finishing the job later than the scheduled end time, not for finishing it earlier. However in just in time (JIT) scheduling, jobs that finish early are also penalized. For this problem, CP consistently outperforms IP. However, combined use of IP-based relaxations and CP for a class of JIT scheduling problems enabled 67 out of 90 instances to be solved while CP alone could solve only 12.

In another application for call-center scheduling, utilizing CP for the master problem and LP for the sub-problem solved twice as many instances as the original MIP formulation. Extremely encouraging results were obtained for a Polypropylene batch-scheduling problem. In this case the roles were reversed, and MIP was used for the master problem while CP was used for the sub-problem. This approach solved a previously insoluble problem in 10 minutes.

From these practical examples it is apparent that coupling the powerful constraint propagation approach of CP with relaxations from MIP is fundamentally changing the way practitioners use O.R. to solve problems. One of the key areas where the integrated approach is adding value is in scheduling problems in hospitals. In particular, the Nurse Scheduling Problem is characterized by the complexity of its business constraints. In this case, apart from benefits in solution run-time, CP cuts down the model formulation time significantly. Before CP arrived, modeling this problem in MIP was considered a very difficult task requiring advanced skills and experience in building models. The CP formulations for this case are quite compact and do not require the services of very highly skilled O.R. analysts. Ease in modeling also translates into faster development, thereby reducing the to-market time for O.R. solutions.

Soon, O.R. practitioners across industries will start using CP in conjunction with IP/MIP to solve a variety of high-impact optimization problems such as assortment optimization, sales-force territory alignment and SKU rationalization. A new methodology, constraint integer programming (CIP), which integrates IP and CP, is also emerging.

## **Is the Additional Complexity Worth It?**

One of the key limitations in using the integrated framework is that the burden of integrating the solvers to some extent still resides upon the user. Though current commercial packages provide for scripting functionality to connect one type of solver to another, there is a lack of a truly integrated language that by looking at the formulations would decide by itself the best way to decompose and solve the problem. Development of such an integrated modeling framework is the need of the hour.

O.R. is more than anything else a practical approach to problem solving. We already know that O.R. problems cannot be solved by sitting in hallowed academic portals without knowing what happens on the ground. We must steer clear of over-reliance on any particular technique because each technique has its limitations and we might lose out on what the other technique has to offer. It is clear that the future belongs to those who smartly use a combination of techniques that complement each other, and CIP is already showing us the way.