COCOS'02, A Workshop on
Global Constrained Optimization and Constraint Satisfaction, Held October 2–4, 2002, Sophia–Antipolis

R. Baker Kearfott (`rbk@louisiana.edu`)
University of Louisiana at Lafayette

The first international workshop on Global Constrained Optimization and Constraint Satisfaction (COCS'02) was held in Sophia Antipolis, a center of France's high-technology industry, on October 2–4, 2002. An excerpt from the "Workshop Objectives" page, found at the web page (COCONUT, 2002), is:

> "This workshop focuses on complete solving techniques for continuous constraint satisfaction and optimization problems that provide all solutions with full rigor. Less rigorous solution techniques are not excluded, since they may be part of complete relevant techniques. Complete solution techniques guarantee that all the constraints – e.g. security or tolerance criteria – are satisfied and the global optima identified."

The conference was part of the "COCONUT" (COntinuous CONstraints – Updating the Technology) project, funded by the European Union,

> "to integrate the currently available techniques from mathematical programming, constraint programming, and interval analysis into a single discipline, to get algorithms for global optimization and continuous constraint satisfaction problems that outperform the current generation of algorithms based on using only techniques from one or two of the traditions."

COCONUT researchers are from a consortium consisting of the University of Vienna, University of Coimbra, Darmstadt Technical University, Université Catholique de Louvain, University of Nantes, the Swiss Federal Institute of Technology, and the ILOG company.

The conference program is also posted at the web address (COCONUT, 2002).

1. Overall Themes

The conference organizers and scientific committee did an excellent job of bringing together prominent individuals from

1. the interval global optimization community,

2. a community of researchers doing deterministic global optimization without emphasis on rigor (using convexity and approximating or bounding by convex functions),

3. the interval constraint propagation community,

4. the semidefinite and DC programming communities,

5. and the applications community.

The overall workshop emphasis was on complete, practical software packages for global optimization. It became apparent early in the conference that each of these four research communities can benefit significantly from interaction with the other.

1.1. Interval Global Optimization

Participants in the workshop representing the interval global optimization community included Jürgen Garloff, Christian Jansson, Arnold Neumaier, Bill Walster, and me (Kearfott) (although a number of participants in the conference could represent more than one of these groups). The guiding ideas for this group are branch and bound methods combined with the Moore–Skelboe algorithm (Skelboe, 1974); Hansen first used the term "interval global optimization" for this class of algorithms. These algorithms involve interval estimation of the range of objective functions and interval Newton methods. Familiar to many of us, these methods appear in (Ratschek and Rokne, 1988), (Neumaier, 1990), (Hansen, 1992), (Kearfott, 1996), and in numerous journal articles and conference proceedings. In this research school, all algorithms contain rigorous interval computations, and the emphasis has been on obtaining more efficient algorithm variants with better interval extensions of the range of the function, gradient, and (for constrained problems) constraints. Although rigorous, most algorithms produced within this research community have in the past been too slow for typical industrial-sized problems.

1.2. Non-Rigorous Deterministic Global Optimization

Chris Floudas and Nick Sahinidis have led in the production of practical commercial software for global optimization without emphasis on rigor. Algorithms in this category have the same basic overall idea and, at the top level, use similar techniques to algorithms centered on rigor and interval analysis. In both rigorous interval-based algorithms and these non-rigorous algorithms, an upper bound on the global optimum of the objective function is used in conjunction with a lower bound on the objective function over a region, to eliminate subregions that do not contain the optimum. (If the lower bound over a region is greater than the upper bound on the global optimum, then the region can be eliminated.) However, efforts in this area have focussed on developing models of the objective (and constraints) which are accurate approximations and also for which the bounds can be easily computed; the philosophy has been that rigor, although useful (and necessary in some cases), is not absolutely necessary for many problems. Both Floudas and Sahinidis make use of convexity in the underestimating and overestimating models, allowing them to easily determine unique optima of these models. The resulting software packages have been successful for a number of larger, practical problems. A knowledgeable person examining both Floudas' and Sahinidis' presentations sees that many underestimating techniques can be made rigorous without a large performance penalty. In fact, Jansson's (Jansson, 2002) and Neumaier and Shcherbina's (Neumaier and Shcherbina, 2002) work point out one way in places where Floudas and Sahinidis use sparse linear programming packages to sort out constrained problems; this needs to be examined further.

1.3. Interval Constraint Propagation

The interval constraint propagation community, as the interval global optimization community and the non-rigorous deterministic global optimization community, also use subdivision and branch and bound methods. Like the interval global optimization community, constraint propagation people assume rigor is necessary. However, the emphasis is not on interval Newton methods to eliminate subregions, but on use of constraints and side conditions to narrow and eliminate subregions. In particular, a system of constraints can be considered of the form

$$c_i(x_1, \ldots, x_n) = 0, \quad 1 \le i \le m. \tag{1}$$

If new, narrower bounds on, say, x_j are made available, then the c_i can be, in principle, solved for the x_k, $k \ne j$, to try to use the narrower

bounds on x_j to obtain narrower bounds on x_k. The entire process can be iterated across i, j, and k until it becomes stationary. Some of the c_i can be of the form

$$\varphi(x_1, \ldots x_n) - \boldsymbol{v} = 0,$$

where φ is an objective to be minimized, and $\boldsymbol{v}$ represents bounds on the objective function: If better upper and lower bounds on the objective function are obtained (say, by exploiting the analytic description of φ), then better bounds on the independent variables can be computed. Similarly, additional constraints can be obtained by setting the gradient of the Lagrangian equal to zero. Also, derived constraints can be obtained by parsing the expressions involving the original objective, gradients, and constraints and assigning variables to the intermediate results in the computations; see (Kearfott, 1991) or (Kearfott, 1996, Chapter 7) for simple examples of this. Also, see the work of Neumaier and Schichl on directed acyclic graphs (DAG's), presented at the workshop and elsewhere.

Although the basic idea behind constraint propagation is simple, there are numerous variants, and many issues greatly affect the practicality of actual implementations. To sort out these issues, the constraint propagation community has focussed on computer languages for constraint propagation (within a more general class of languages termed "logic programming"), along with development of a theoretical framework and terminology for specific ways of using the constraints. This framework has led to recent advances in the practicality of constraint propagation techniques, and to a number of techniques that go beyond or complement the naive implementations that outsiders from the interval community have produced. However, some interval researchers are somewhat daunted by the need to learn this terminology. This workshop opened lines of communication and ameliorated this problem.

Simple examples show that constraint propagation alone (without some kind of interval Newton method, preconditioning, or other system solver) does not allow us to avoid subdivision altogether in highly-coupled systems of constraints. In fact, constraint propagation can be viewed as a kind of nonlinear Gauss–Seidel method. However, as Luc Jaulin convincingly pointed out at the workshop, constraint propagation may be the *only* method of handling large regions during a branch-and-bound process when the dimension is also large. (Nonetheless, constraint propagation alone cannot be expected to reduce regions to points except in special cases.)

Frédéric Benhamou, at the workshop and on the scientific committee, has been cited frequently in the early literature on constraint

propagation. A milestone includes Pascal van Hentenryck's "Numerica" software and corresponding book (Van Hentenryck et al., 1997).

1.4. DC Programming and Semidefinite Programming

The DC programming community decomposes objectives or constraints into Differences of Complex functions, while the semidefinite programming community approximates (or overestimates or underestimates) objective or constraints by convex functions. The convexity or the difference-of-convex structure is then used to advantage. DC programming is used in numerous places in the material discussed at the workshop, such as in Floudas' "α-BB" technique. Using semidefinite programming, Jean Lasserre presented a Matlab-based package "GloptiPoly" for global optimization over polynomials. GloptiPoly is available free of charge from the web page (Lasserre, 2002a). The underlying theory for GloptiPoly is presented in (Lasserre, 2001) and (Lasserre, 2002b).

1.5. Applications

The fifth group brought together at this workshop consisted of people with applications. To a large extent, these people have written their own computer codes, taking advantage of particular structure within their application and of whatever techniques are available. Several impressive successes with global and with verified techniques should encourage academic researchers in validated global optimization.

2. Highlights (A Personal View)

In the first invited lecture, Christodoulos Floudas presented a synopsis of work he has directed over the past fifteen years towards global optimization problems arising in the chemical engineering industry and elsewhere. A cornerstone of the methods is construction of underestimators to objectives and constraints by decomposing expressions into linear terms, quadratic terms, and more general convex terms. The "α BB" approach is then used to replace the general terms by convex underestimators. Floudas outlined several practical problems, including some with relatively large numbers of variables, for which the Floudas' approach, but not others, succeeded. Much of this can be found in Floudas' book (Floudas, 2000).

Nikolaos Sahinidis, giving the invited lecture on the next day, presented his "BARON" global optimization software, available commercially through GAMS. Although the techniques used are different (and

have been derived separately), BARON is similar to the work of Floudas et al in the sense that it is not completely rigorous (although both Floudas and Sahinidis use interval arithmetic in places), but is successful for similar types of relatively large-scale problems. See (Tawarmalani and Sahinidis, 2002) for an explanation of most of the techniques Sahinidis presented at the workshop.

The verified global optimization community at large can probably learn much from study of both Floudas' and Sahinidis' techniques.

Christian Jansson gave a clever technique, mentioned above and described, simultaneously but independently, in the references (Jansson, 2002) and (Neumaier and Shcherbina, 2002), for computing rigorous bounds on the optimum value for linear programs. The technique requires only $\mathcal{O}\left(n^2\right)$ operations (where n is the number of variables), and can take full advantage of sparsity. Although not giving bounds on the optimizing variables, the technique can potentially be combined with convex underestimating techniques of Floudas, Sahinidis and others to make these latter techniques rigorous.

Luc Jaulin presented an impressive constraint propagation success, in which he narrowed interval bounds from, essentially, the entire range of real numbers, to points or relatively narrow intervals, without subdividing, for a problem with thousands of variables. This system of constraints was obtained by discretizing a system for satellite location. Although there was some controversy about whether this was really a global or a local problem, and although constraint propagation alone was not able to reduce all coordinates completely to points, the example gave convincing evidence that at least some large global optimization problems not amenable to other techniques can be solved with rigorous constraint propagation.

Daniel Johnson outlined the method he and his colleagues developed for feasibility scheduling of operations for an entire oil refinery. The work both illustrates that interval techniques can be highly successful for large practical problems (including over 10,000 variables and constraints) and underscores the need to take advantage of specific problem structure in some cases.

A number of participants presented software packages, both free and otherwise, utilizing various programming languages. Axel Meeraus, the founder and head of the GAMS corporation, was present and participated in a panel discussion on useful directions for future work, while Michael Bussieck from GAMS gave a presentation on the system GAMS uses to test optimization packages before inclusion in the GAMS library. (GAMS, General Algebraic Modeling System, `www.gams.com`, is a commercial enterprise whose product features access to a number of solvers for nonlinear programming, unified through a common mod-

elling language. Software developers can find the testing framework at `www.gamsworld.org`.)

Other unique, clever, and useful innovations were presented but not mentioned here. A proceedings with complete papers will be published.

3. Future Directions

The workshop made clear that more careful comparison is necessary. In particular, a number of software packages, using different techniques exist, and the authors have presented results illustrating success for various problems. However, a systematic comparison of these packages on a common problem set does not presently exist in the open literature. In the mean time, techniques proliferate, and it is unclear which of these are superior, which are complementary, etc. Only careful testing will lead to software that is uniformly superior to that presently available.

Arnold Neumaier announced an open source initiative, in which global optimization software from the COCONUT project, along with a carefully defined framework, submission protocol and style guidelines, will be openly available, probably by December, 2003. Global optimization and constraint propagation researchers will be able to submit coordinate-narrowing and other techniques for inclusion into this system.

4. Acknowledgements

The author of this report wishes to thank the organizers, Christian Bliek and Djamila Sam-Harold for the administrative work that went into the workshop, and to thank the entire scientific committee for the excellent combination of topics and researchers.

I wish to warmly thank Arnold Neumaier, who went through several iterations of proofreading, and guided me to better understanding of certain aspects of the material.

Just because I have not mentioned a particular workshop participants here does not lessen their contributions. The reader is urged to explore the list of accepted papers and the workshop program at (COCONUT, 2002).

References

COCONUT: 2002, 'COCOS'02, The First International Workshop on Global Constrained Optimization and Constraint Satisfaction'. `http://liawww.epfl.ch/Cocos02/cocos02.html`.

Floudas, C. A.: 2000, *Deterministic Global Optimization: Theory, Algorithms and Applications*. Dordrecht, Netherlands: Kluwer.

Hansen, E. R.: 1992, *Global Optimization Using Interval Analysis*. New York: Marcel Dekker, Inc.

Jansson, C.: 2002, 'Rigorous Lower and Upper Bounds in Linear Programming'. Technical report, TU Hamburg-Harburg. `http://www.ti3/tu-harburg.de/paper/jansson/verification.ps`.

Kearfott, R. B.: 1991, 'Decomposition of Arithmetic Expressions to Improve the Behavior of Interval Iteration for Nonlinear Systems'. *Computing* **47**(2), 169–191.

Kearfott, R. B.: 1996, *Rigorous Global Search: Continuous Problems*. Dordrecht, Netherlands: Kluwer.

Lasserre, J.: 2002a, 'GloptiPoly - Global Optimization over Polynomials with Matlab and SeDuMi'. `http://www.laas.fr/~henrion/software/gloptipoly/gloptipoly.html`.

Lasserre, J. B.: 2001, 'Global Optimization with Polynomials and the Problem of Moments'. *SIAM J. Optim.* **11**(3).

Lasserre, J. B.: 2002b, 'An Explicit Equivalent Positive Semidefinite Program for Nonlinear 0-1 Programs'. *SIAM J. Optim.* **12**(3), 756–769.

Neumaier, A.: 1990, *Interval Methods for Systems of Equations*. Cambridge, England: Cambridge University Press.

Neumaier, A. and O. Shcherbina: 2002, 'Safe Bounds in Linear and Mixed-Integer Programming'. Technical report. `http://www.mat.univie.ac.at/~neum/papers.html`.

Ratschek, H. and J. Rokne: 1988, *New Computer Methods for Global Optimization*. New York: Wiley.

Skelboe, S.: 1974, 'Computation of Rational Interval Functions'. *BIT* **14**, 87–95.

Tawarmalani, M. and N. V. Sahinidis: 2002, *Convexification and Global Optimization in Continuous and Mixed-Integer Nonlinear Programming: Theory, Algorithms, Software, and Applications*. Dordrecht, Netherlands: Kluwer.

Van Hentenryck, P., L. Michel, and Y. Deville: 1997, *Numerica: A Modeling Language for Global Optimization*. Cambridge, MA: MIT Press.