AFOSR-TR-97-0654

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	15 Sept 1997	Final - 1 Oct 1994 - 28 Feb 1997

4. TITLE AND SUBTITLE
New Methods for Large Scale Local and Global Optimization

5. FUNDING NUMBERS
F49620-94-1-0101

6. AUTHOR(S)
Robert Schnabel and Richard Byrd

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
University of Colorado at Boulder
Boulder, CO 80309

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Office of Scientific Research/NM
110 Duncan Avenue, Suite 100
Bolling Air Force Base, DC 20332-0001

10. SPONSORING / MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for public release;
distribution unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

We have pursued all the topic areas described in the grant proposal during the period of this grant. These include global optimization methods for molecular cluster problems, global optimization methods for protein folding problems, smoothing methods for global optimization, other optimization topics from molecular chemistry, and related large-scale optimization topics. The largest amount of effort has gone into the construction of new global optimization methods for protein folding problems and the development of new smoothing approaches for global optimization. We have extended our stochastic/perturbation approach to large-scale global optimization to deal with proteins, and have had very good success on initial test problems. We have also developed a new, analytic smoothing approach, investigated some fundamental properties of smoothing, and successfully incorporated our smoothing approach into our global optimization algorithm. The combination of smoothing and our stochastic/perturbation approach is so far producing excellent results. A related accomplishment under this grant has been the successful application of our stochastic/perturbation approach to distance geometry problems from molecular chemistry. Finally, we have developed unified limited memory/truncated Newton methods for large-scale unconstrained optimization that seem to combine many of the advantages of each approach.

14. SUBJECT TERMS	15. NUMBER OF PAGES
Optimization, large-scale, molecular chemistry, global optimization	11
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	UL

Final Technical Report

U.S. Air Force Grant F49620-94-1-0101
Large-Scale Optimization Methods for Molecular Chemistry Problems

Richard Byrd, Robert Schnabel
Department of Computer Science
University of Colorado at Boulder
Boulder, CO 80309-0430

1 October 1994 - 28 February 1997

1. Summary

We have pursued all the topic areas described in the grant proposal during the period of this grant. These include global optimization methods for molecular cluster problems, global optimization methods for protein folding problems, smoothing methods for global optimization problems, other optimization topics from molecular chemistry, and related large-scale optimization topics. The largest amount of effort has gone into the construction of new global optimization methods for protein folding problems and the development of new smoothing approaches for global optimization. We have extended our stochastic/perturbation approach to large-scale global optimization to deal with proteins, and have had very good success on initial test problems. We have also developed a new, analytic smoothing approach, investigated some fundamental properties of smoothing, and successfully incorporated our smoothing approach into our global optimization algorithm. The combination of smoothing and our stochastic/perturbation approach is so far producing excellent results. A related accomplishment under this grant has been the successful application of our stochastic/perturbation approach to distance geometry problems from molecular chemistry. Finally, we have developed unified limited memory / truncated Newton methods for large-scale unconstrained optimization that seem to combine many of the advantages of each approach.

2. Global Optimization Methods for Molecular Cluster Problems

Under our previous AFOSR grant, we began the development of a global optimization approach for finding the lowest energy configurations of molecular structures. To do this, one must find the lowest (global) minimum of energy functions that generally have very many parameters and huge numbers of local minimizers. Therefore, these are very difficult global optimization problems. Previously developed general purpose global optimization algorithms were not nearly able to solve problems of this size.

Our approach is to develop a fairly general class of methods that are applicable to a broad class of partially separable large scale global optimization problems. The methods combine efficient stochastic global optimization techniques with several new, more deterministic perturbation techniques. The main phase of the approach is a series of steps that move from low minimizers to lower minimizers by selecting a small subset of the parameters that appears most fruitful to modify, and then applying a small-scale, stochastic global optimization method to modify these parameters with the remaining parameters temporarily fixed. This is followed by a local minimization algorithm applied to all the parameters. This

phase utilizes strategies that select the configurations to work on from among the best current configurations, in a way that balances the breadth of the search (trying configurations from several different origins) with the depth of the search (modifying the configurations with the current lowest energy values).

Our previous test results showed that this method successfully solves all Lennard-Jones problems with up to 76 atoms (the largest problem that we have attempted). This includes finding a new global minimizer for 75 atoms, and re-finding best minimizers for 66 and 72 atoms that had been found recently by other researchers. Overall, these results appear to be the best for the Lennard-Jones problem by any general or special purpose method. We had also worked on the water cluster problem, using the energy function of Coker and Watts. Early in this grant period we continued this research, considerably improving our results for water cluster with 20 and 21 minimizers (180 and 189 optimization parameters) through long runs on the Caltech Delta (using 64 processors), and finding considerably better configurations for 32 water molecules (288 parameters) than the previous research.

These algorithms formed much of the basis for the research described in the next two sections which is the core work performed under the support of this grant.

3. Global Optimization Methods for Protein Folding Problems

Our primary activity under this grant has been the development and testing of global optimization methods for the protein folding problem. This includes the basic methods discussed in this section and the smoothing approaches discussed in the following section. As with molecular clusters, the problem is to find the lowest energy configuration of a protein or other polymer. Again, these problems have many parameters and very many local minimizers, making them difficult global optimization problems. The main difference from molecular clusters is that proteins have a known chain structure that constrains the possible structures. Therefore algorithms need to be significantly different since they should only explore configurations that satisfy the chain structure. In particular, the types of moves of atoms or molecules that form the basis of our cluster algorithms are not appropriate for polymers. A related difference is that these problems are often parameterized in "internal" parameters that consist of dihedral (torsion) angles. This parameterization reduces that number of variables by more than a factor of ten, but removes the partial separability of the problem.

We developed a new global optimization method for solving polymer problems, adopted from our molecular cluster methods. Like the molecular cluster methods, it consists of two stages. The first, initialization stage finds a set of low local minimizers, and the second, improvement stage successively moves from low to even lower local minimizers. The initialization stage for polymers is entirely different than for clusters, due to the chain structure. It consists of a build-up process that inexpensively builds possible configurations one amino acid at a time. The improvement stage has the same high level structure as for clusters: select a configuration and some parameters to improve, perform a small-scale global optimization on this configuration with only the selected parameters varying, then perform full variable local minimizations from the best new configurations, and add the new configurations to the list of configurations. Beneath this level the key parts of the improvement stage are basically all new. Particularly important are the methodology for selecting which parameters to vary in the small global optimization, and the methodology for selecting the configurations to try to improve. The parameters that are varied are dihedral angles, so that the chain structure is maintained.

We have applied this algorithm to the protein polyalanine with 20, 30, 40, and 58 amino acids, and to the protein metenkephalin. Polyalanine is a commonly studied protein that is simple in that all amino acids are the same and have just a one-molecule residue. It is interesting because for the smaller sizes, the optimal structure is an alpha-helix, while for the 58 amino acid case, a bend structure consisting of two alpha-helical strands has slightly lower energy than the straight alpha-helix. Thus, the start of real folding behavior is exhibited in this case. Our algorithm has been very successful on these problems. The results have been very successful. For 20 amino-acids, the initialization phase already finds the alpha-helical global minimizer. For 30 and 40 amino-acids, the improvement phase finds the global minimizer, still an alpha helix. For 58 amino-acids, the improvement phase was able to find the best alpha-helical structure and an even lower energy bent structure, although with the incorporation of the smoothing techniques discussed in the next section, a slightly better bent structure was found. These results are noteworthy because we are not aware of other global optimization approaches that have been able to deal with polymer problems of this size. For metenkephalin, a small but more complex protein, we were able to find the global minimizer, but other researchers have succeeded on this protein as well. As discussed in section 4, the incorporation of smoothing has led to a considerably more powerful algorithm still. All of these computational results have been obtained with parallel algorithms run either on parallel computers or networks of workstations.

Much of the research our our global optimization approach for proteins concentrated on the heuristics for determining which configurations to try to improve, and which dihedral angles to use as parameters to the small-scale global optimization. We showed that a combination of breadth and depth was best in the choice of configurations, i.e. trying descendants of a variety of configurations from the initialization phase as well as trying to improve the best configurations found so far. We also showed for polyalanine that if one knows which parameters are farthest from their optimal values and chooses those for the parameters in the small-scale global optimization phase (an "ideal" strategy), then the algorithm works extremely well. Motivated by this, we investigated practical new strategies for choosing the parameters to vary, including some that use chemical knowledge such as Ramachandran plots or DSP programs for recognizing secondary structure. These new strategies significantly improved the ability of our method to find global optimal or near-optimal structures for polyalanine, and are motivating us to use chemical knowledge in this phase of our algorithm as we progress to more complex problems.

Another portion of our global optimization research for proteins has been tests of a "biasing" approach. Our work is derived from work of Head-Gordon, Gay, and Wright. In this work, penalty terms that favor predicted secondary structure are added to the energy function. For the protein polyalanine with 20, 30, 40, and 58 amino-acids, we have shown that biasing towards alpha-helix permits the initialization phase of our algorithm to already find the optimal, alpha-helical structure. This is a huge improvement in efficiency over algorithms without biasing. This research shows that if one has a prediction of secondary structure, it can be usefully incorporated in the global optimization method for determining the full folded structure. An interesting and unexpected observation in this research is that biasing is far more successful in Cartesian coordinates than in the reduced dihedral angle parameter space that we normally use. This leads to interesting tradeoffs since it is expensive to work in the full Cartesian space. Our ongoing research and that of Head-Gordon is investigating the use of biasing within the context of our full stochastic/perturbation global optimization algorithm.

4. Smoothing Methods for Global Optimization Problems

Our other main activity under this grant has been the development, testing and application of a new "smoothing" approach to use in the context of these global optimization problems. Smoothing is an approach that transforms the original multiple-minima problem to a smoother problem with one or a few minimizers, and then transforms back in a way that leads to the global minimizer. Our work in smoothing has had three very significant portions: development of a new, analytic smoothing function; study of fundamental properties of smoothing that have important practical implications; and successful incorporation of smoothing into our global optimization approach to protein folding. The rest of this section describes these portions of this research.

We first developed a new, simple, analytic way to smooth energy functions for molecular chemistry problems. It is based on replacing the terms in the energy function that have poles and make the global optimization problem difficult, generally Lennard-Jones and electrostatic terms, with analytic replacements that remove the pole and broaden the well of the minimizer. We have developed a simple two-parameter family of smoothed functions that progress continuously from the original functions to smoother and smoother variants. This approach offers an appealing alternative to the convolution or Taylor-series approaches used previously, which are expensive and have fundamental limitations due to the poles in the energy function. Our new approach appears to retain the advantages of the earlier approaches but be considerably cheaper and simpler to apply. In particular, we have shown for Lennard-Jones clusters that both the shapes and trajectory behaviors of our analytically smoothed functions are very similar to that produced by the smoothing approaches of Scheraga et al and of Coleman, Wu, and More'. We have also shown how to treat the Lennard-Jones and electrostatic terms together in an internally consistent way. This makes the approach immediately applicable to the empirical potential energy function for any protein.

Second, through experiments on molecular clusters, we have illustrated a fundamental issue with the smoothing approach: that "trajectory flips" appear to be very common. This essentially means that the best smoothed minimizer(s) don't generally lead back to the best unsmoothed minimizer(s). This shows that approaches that smooth back to one or a few minimizers and then trace back, which has been the common use of smoothing, will have significant limitations. In contrast, an approach such as ours that incorporates smoothing into a global optimization approach may have a much better chance of success. We have shown that this behavior is not just an artifact of our smoothing approach but is shared by other smoothing approaches.

Third, we have incorporated smoothing into our stochastic/perturbation approach to global optimization of proteins, with excellent results. The algorithm utilizes smoothing in two ways. In the initialization stage, the initialization is done entirely on a smoothed function, and then the smoothed minimizers are traced (in one step, so far) back to minimizers of the unsmoothed function. This stage alone already found the optimal alpha-helical shape for polyalanine with 20, 30, 40 and 58 amino-acids. This was a huge improvement in efficiency over the algorithm without smoothing. Then in the second, improvement stage, we use the best smoothed minimizers from the initialization phase to perform the improvement phase on a less smoothed function, then tracing the minimizers this produces back to the unsmoothed function (again, so far in just one step). To our surprise, this approach (automatically) located at a better bent global optimizer for the 58 amino-acid than we had found in several years of working on this problem without smoothing. To our knowledge this result is superior to any other work on this problem. Our algorithm using smoothing also found near-optimal solutions far more efficiently than the approach

without smoothing. We are very encouraged by these results and are expecting that smoothing will play a major role in our ongoing research on global optimization for protein folding

5. Other Optimization Topics from Molecular Chemistry

Another accomplishment in this research period has been the adaptation of our stochastic/perturbation global optimization approach to distance geometry problems. Distance geometry problems are a different type of global optimization problem that arise in molecular chemistry. They occur when one has measurements of some atom-atom distances in proteins (generally from NMR experiments), and wishes to reconstruct the molecular structure from these measurements. These problems are well-studied in the chemistry literature but hardly considered yet in the global optimization community. Motivated in part by some recent work of More' and Wu on these problems, we have adapted our stochastic/perturbation approach to solve them. These problems are naturally characterized in Cartesian coordinates and so are more akin to molecular cluster problems than protein folding problems from the viewpoint of our method. Adapting our algorithm required new heuristics for choosing subproblems; this could be done very naturally which is a nice validation of the generality of our approach. As in molecular cluster problems, the problems are partially separable which greatly aids our heuristics and the efficiency of our method. Our computational results on the problems tested by are excellent, far surpassing the results reported in previous work. The largest problems that we successfully solved had over 700 atoms, i.e. over 2000 optimization parameters. These are very large problems given the current state of the art in global optimization.

We also began work on a new, basic optimization topic of interest in molecular chemistry, the determination of saddle points. Saddle points with one negative eigenvalue determine transition states, which are of great interest to chemists. The chemistry community has constructed a variety of algorithms for locating these problems, but they appear somewhat ad-hoc and there does not seem to be any theoretical analysis for them. The problem really has three aspects: a global aspect of getting in the basin of attraction of a saddle point, a semi-local aspect of getting close to the saddle point one one is in its basin of attraction, and a local aspect of converging quickly to the saddle point once one is close enough to it. Our initial work concentrated on the local aspect. Chemistry methods for the local aspect have generally used the PSB update, because it is well defined for indefinite problems. However we expected that the scale-invariant properties of the SR1 and BFGS method should make them superior to the PSB for these problems if they could be implemented stably for them. We formulated a stable way to use BFGS and SR1 and have shown on a wide range of test problems that these updates, particularly the BFGS, are never appreciably worse than PSB and often far superior.

6. Related Large-Scale Optimization Topics

One optimization topic related to molecular chemistry problems that we have investigated under this grant is the convergence of the unconstrained optimization methods on "benign singular" problems. These problems occur in molecular cluster problems, distance geometry problems, and any other problems where the parameters are the locations of atoms or other objects in space. In the normal Cartesian parameterization, there are six excess degrees of freedom in such three-dimensional problems, and the Hessian matrix is singular everywhere. For this reason, such problems usually are re-parameterized to eliminate the (six) excess degrees of freedom, which involves some extra work and leads to an unnatural parameterization. However our computational results have shown that the unconstrained optimization

algorithms appear to converge quickly on the original parameterization even though the Hessian matrix is singular, and have attempted to show mathematically that this is the case. So far we have shown this for a set of cases that includes translations, and are working on establishing the more general result that will include rotations. If we are able to do this, it will eliminate the need for chemists to perform these re-parameterizations when working in Cartesian variables.

The second topic that we worked on is unified limited memory / truncated Newton methods for large-scale unconstrained optimization. These two methods are two of the most effective for general large-scale unconstrained optimization. The two methods are somewhat similar in that they do not depend on sparsity or special structure of the problem, and the required storage is linear in the dimension of the problem. However from the point of view of distribution of effort they are at opposite extremes. A truncated Newton method does more work at each iteration, approaching Newton's method in the limit, while limited memory methods do just a bit more work than a conjugate gradient method at each iteration. In comparative testing, limited memory methods are more efficient on highly nonquadratic problems that are not too ill-conditioned, while truncated Newton methods can handle ill-conditioning better, especially if the objective function is not too far from quadratic. Indeed, limited memory methods are sometimes simply unable to reach an accurate solution to some ill-conditioned problems which truncated Newton methods solve without difficulty.

The objective of our research has been to see if we can develop a combined method that combines the advantages of each approach. We have experimented with a simple hybrid method where one truncated Newton iteration is interspersed with about every ten limited memory BFGS iterations. We found that on many problems, the truncated Newton step is able to make more than enough progress relative to limited memory to make up for its greater cost. More interestingly, this framework allows us to save the second order information from the inner conjugate gradient steps of truncated Newton for use in the limited memory Hessian approximation at later steps. In experiments with such an information saving strategy we have found that performance of the method is significantly improved by the saving, at least for certain eigenvalue distributions of the Hessian. There seems to be a significant interplay between the number of inner conjugate gradient steps of the truncated Newton iteration and the number of limited memory iterations between truncated Newton iterations. This work has been done jointly with Professor Jorge Nocedal of Northwestern University.

7. Publications Resulting from this Grant

(1) A. Bouaricha and R. Schnabel, "TENSOLVE: A software package for solving systems of nonlinear equations and nonlinear least squares problems using tensor methods", to appear in *ACM Transactions on Mathematical Software.*

(2) R. Byrd, T. Derby, E. Eskow, K. Oldenkamp, and R. Schnabel, "A new stochastic/perturbation method for large-scale global optimization and its application to water cluster problems", *Large-Scale Optimization: State of the Art,* W. Hager, D. Hearn, and P. Pardalos, eds., Kluwer Academic Publishers, 1994, pp. 71-84.

(3) R. Byrd, E. Eskow, B. Oldenkamp, R. Schnabel, and A. van der Hoek, "A parallel global optimization method for solving molecular cluster and polymer conformation problems", Proceedings of the Seventh SIAM Conference on Parallel Processing for Scientific Computing, SIAM, 1995, pp. 72-77

(4) R. Byrd, E. Eskow, and R. Schnabel, "A new large-scale global optimization method and its application to Lennard-Jones problems", University of Colorado Technical Report CU-CS-630-92.

(5) R. Byrd, E. Eskow, R. Schnabel, C.-S. Shao, A. van der Hoek, and Z. Zou, "Global optimization methods for protein folding problems", Proceedings of the DIMACS Workshop on Global Minimization of Nonconvex Energy Functions: Molecular Conformation and Protein Folding, P. Pardalos, D. Shalloway, and G. Xue, eds., American Mathematical Society, 1995, pp. 29-39.

(6) R. Byrd, H. Fayez Khalfan, and R. Schnabel, "Analysis of a symmetric rank-one trust region method", *SIAM Journal on Optimization* 6, 1996, pp. 1025-1039.

(7) R. Byrd, P. Lu, J. Nocedal and C. Zhu, "A limited memory algorithm for bound constrained optimization," *SIAM Journal on Scientific Computing* 16, 1996, pp. 1190-1208.

(8) R. Byrd, J. Nocedal, and R. Schnabel, "Representations of quasi-Newton matrices and their use in limited memory methods", *Mathematical Programming* 63, 1994, pp, 129-156.

(9) R. Byrd, J. Nocedal and C. Zhu, "Towards a discrete Newton method with memory for large-scale optimization", in *Nonlinear Optimization and Applications,* G.Di Pillo and F. Giannesi, eds., Plenum, 1996.

(10) T. Chow, E. Eskow, and R. Schnabel, "A software package for unconstrained optimization using tensor methods," *ACM Transactions on Mathematical Software* 20, 1994, pp. 518-530.

(11) D. Feng and R. Schnabel, "Tensor methods for equality constrained optimization", *SIAM Journal on Optimization, August 1996.*

(12) D. Feng and R. Schnabel, "Local convergence analysis of tensor and SQP methods for singular constrained optimization", to appear in *SIAM Journal on Optimization.*

(13) R. Schnabel, "Parallel nonlinear optimization: limitations, opportunities, and challenges", *Algorithms for Continuous Optimization: The State of the Art,* E. Spedicato, ed., Kluwer Academic Publishers, 1994, pp. 531-559.

(14) R. Schnabel, "A view of the limitations, opportunities, and challenges in parallel nonlinear optimization", *Parallel Computing* 21, 1995, pp. 875-905.

(15) C.-S. Shao, R. Byrd, E. Eskow and R. Schnabel, "Global optimization for molecular cluster problems using a new smoothing approach", to appear in *Large Scale Optimization with Applications, Part III: Molecular Structure and Optimization*, L. Biegler, T. Coleman, A. Conn, and F. Santosa eds., Springer-Verlag.

(16) Y. Yuan and R. Byrd, "Non-quasi-Newton updates for unconstrained optimization", *Journal of Computational Mathematics* 13, 1995, pp. 95-107.

(17) Z. Zou, R. Byrd and R. Schnabel, "A stochastic/perturbation global optimization algorithm for distance geometry problems", to appear in *Journal on Global Optimization*.

(18) C. Zhu, R.H. Byrd, P. Lu and J. Nocedal, "L-BFGS-B -- Fortran subroutines for large-scale bound constrained optimization", to appear in *ACM Transactions on Mathematical Software*.

8. Professional Personnel Supported by Grant

R. Byrd, Co-Principal Investigator

R. Schnabel, Principal Investigator

Azmi Aqil, Research Assistant, 8/96 - 2/97

Mr. Azmi is a Ph.D. student in the Department of Computer Science who is doing his Ph.D. thesis research in global optimization. He has been the main person investigating the performance of smoothing approaches on protein folding problems during the period he has worked on the project. His Ph.D. thesis, expected to completed by summer 1998, will continue in this line of research.

Twan Ettes, Research Assistant, 8/95 - 7/96

Mr. Ettes worked on this project while a visiting graduate student in our department from Erasmus University in Rotterdam, Netherlands. Mr. Ettes performed much of the computational work in testing global optimization strategies for protein folding, especially for the polyalanine problem. He also contributed extensively to the development of new algorithmic strategies, and to the investigation of the parallel efficiency of the method. He completed his M.S. thesis under our supervision that describes this work.

Andre van der Hoek, Research Assistant, 12/93 - 7/94

Mr. van der Hoek worked on this project while a visiting graduate student in our department. He was the main person involved in the implementation and testing of the global optimization methods for polymer problems, and also made many contributions to their development. His M.S. thesis, at Erasmus University in the Netherlands, was on this topic and was directed by Prof. Schnabel.

Anna Szczyrba, Research Assistant, 9/96 - 2/97

Ms. Szczyrba is an M.S. student in the Department of Computer Science who may continue for her Ph.D. She has been conducting computational experiments with our stochastic/perturbation global optimization approach.

Chung-Shang Shao, Research Assistant, 12/93 - 8/95

Mr. Shao is a Ph.D. student in the Department of Computer Science at the University of Colorado at Boulder. He worked extensively on the smoothing approaches to global optimization, conducting much of the fundamental testing and making many important observations. He is now working on a thesis in parallel systems that is motivated by the global optimization work, and continues to participate in our group and to make research contributions.

Zhihong Zou, Research Assistant, 8/94 - 8/96

Mr. Zou is Ph.D. student in the Department of Computer Science at the University of Colorado at Boulder. His contributions included experimental research on global optimization methods for protein folding problems, taking the lead in investigating the application of our global optimization methodology to distance geometry problems, and investigating the computational and theoretical issues in new algorithms for locating saddle points.

9. Presentations about Grant Research at Technical Meetings

(1) R. Byrd, "Global Optimization of Molecular Structure by a Stochastic-Perturbation Method", invited talk, Fifth SIAM Conference on Optimization, Victoria, Canada, May 1996.

(2) R. Byrd, "Optimization Methods for Problems from Molecular Chemistry", invited talk, International Conference on Nonlinear Programming, Beijing, China, Sept. 1996.

(3) R. Byrd, "An Interior Point Method for General Nonlinear Programming using Trust Regions", invited talk, Oberwohlfach Conference on Mathematical Optimization, Oberwohlfach, Germany, Jan. 1997.

(4) E. Eskow, "Global Optimization Methods for Protein Folding Problems", invited talk, American Chemical Society Annual Meeting, Orlando, Florida, Aug. 1996.

(5) R. Schnabel, "An overview of parallel architectures, algorithms, and languages", invited talk, Workshop on Optimization and Parallel Computation, U.A.E. University, Al-Ain, United Arab Emirates, May, 1994.

(6) R. Schnabel, "Global optimization methods for molecular configuration problems", invited talk, Workshop on Optimization and Parallel Computation, U.A.E. University, Al-Ain, United Arab Emirates, May, 1994.

(7) R. Schnabel, "Global optimization methods for molecular configuration problems", invited minisymposium talk, SIAM National Meeting, San Diego, CA, July 1994.

(8) R. Schnabel, "Global optimization methods for molecular conformation problems", invited talk, Fifteenth International Symposium on Mathematical Programming, Ann Arbor, MI, August 1994.

(9) R. Schnabel, "Global Optimization for Protein Folding Problems", invited talk, DIMACS Workshop on Global Minimization of Nonconvex Energy Functions, New Brunswick, N.J., Mar. 1995.

(10) R. Schnabel, "Global Optimization Methods for Molecular Configuration Problems", invited talk, IMA Workshop on Large-Scale Optimization, Minneapolis, July 1995.

(11) R. Schnabel, "A Parallel Global Optimization Methods for Solving Molecular Cluster and Protein Conformation Problems", Seventh SIAM Conference on Parallel Processing for Scientific Computing, San Francisco, California, February 1995.

(12) R. Schnabel, "Global Optimization Methods for Molecular Configuration Problems", invited talk, SIAM National Meeting, Charlotte, NC, Oct. 1995.

(13) R. Schnabel, "Global Optimization, Parallel Optimization", invited talk, Linear Algebra in Optimization Workshop, Albi, France, Apr. 1996.

(14) R. Schnabel, "Convergence of Unconstrained Optimization Methods on Benign Singular Problems", Fifth SIAM Conference on Optimization, Victoria, May 1996.

R. Schnabel also presented colloquium talks on Global Optimization Methods for Molecular Configuration Problems during this period at:

Argonne National Laboratory
Cornell University
Colorado School of Mines
National Institutes of Health
Stanford University
University of Colorado
Wright Air Force Laboratory