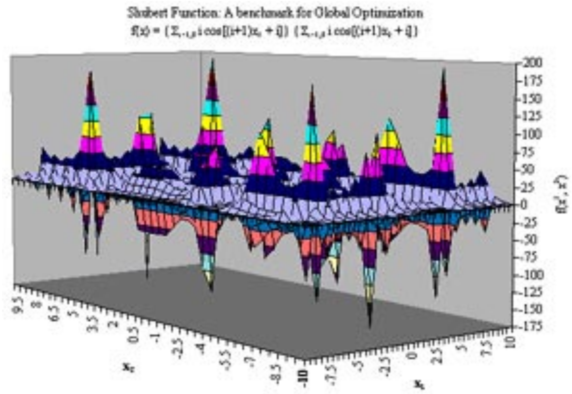

A DOE SUCCESS: *TRUST*
A Computational Tool for Global Optimization

The Project: The Terminal Repeller Unconstrained Subenergy Tunneling (TRUST) tool for Global Optimization is designed to find the optimal solution to a problem for which many sub-optimal solutions are possible. Global optimization occurs in a variety of applications in engineering, manufacturing, robotics, navigation, and all branches of science. The goal may be to resolve manipulator redundancy or to minimize: travel distance or time, cost, used resources, waste, risk, etc. In some applications, the goal may be to maximize capacity, efficiency, industrial throughput, overall quality, and/or profit, etc.

Consider a mobile manipulator performing a task. The control parameters are accelerations that change the velocities and values of the actuators that determine the position and orientation of the vehicle and the manipulator. For a vehicle with four steerable drive wheels, the actuators change the steering angle and rotation of each of the wheels. For a seven-degree-of-freedom manipulator with a one-degree-of-freedom gripper, eight actuators control the system. The global optimization is to determine the actuator accelerations that will perform a task in minimum time by moving the mobile manipulator in a cluttered environment while avoiding obstacles, minimizing the torque on the joints, and avoiding joint limits and velocity limits.

In mathematical terms, the global optimization problem can be stated as follows: given a function of many variables, find its global (lowest) minimum among all the other local (less deep) minima. The difficulty of the global optimization problem is threefold:

> Local criteria (such as the annulment of derivatives) do NOT differentiate between local and global minima. The only way to uncover the global minimum is exhaustive search and comparison of all the possible local minima outcomes.
>
> Exhaustive search becomes exponentially expensive as the dimensionality of the problem grows. Hence, such a search becomes actually infeasible. Unfortunately, most real life problems fall into this category.
>
> Finally, in many practical situations the value of the function to be minimized is not given analytically, but is computed using complex computer codes that are extremely expensive to run. This makes the number of "function evaluations" the paramount criterion in comparing the efficiency of optimization algorithms.

These inherent difficulties render current methods quite tedious, time-consuming, and thereby costly. Furthermore even when implemented on the most powerful parallel computing hardware, conventional optimizers may require hours to find a solution, and the solution reached is often far from optimal.

The Impact: TRUST overcomes these drawbacks by combining several new, original concepts, namely: Pijavskij cones, subenergy tunneling, and terminal repellers. To understand these rather technical concepts, we resort to an analogy.

In most instances, finding the global minimum of a very complicated function can be compared to finding the needle in the haystack. The search procedure would be enormously sped up if one knew in advance in what regions of the haystack NOT to bother to search. Suppose one finds himself in a certain point of the stack. The Pijavskij cones are like cones of shade thrown from that point into the stack. Based on minimal a priori information about the function, they tell us there is no need to search in those shaded regions. With each search, the shaded region increases, thereby eliminating larger and larger parts from the search space. This enhancement increases the efficiency of the search phase of the algorithm.

Suppose now that after eliminating "barren regions" the search point finds itself in the vicinity of a (local) minimum. This means that one is on the slope of a pit and one "rolls down the slope" until the bottom (minimum) is reached. The local minimum is signaled by a zero gradient (slope). But once the slope is zero, the fall has ended and the point is "trapped" at the bottom of the pit. To escape from the local minimum a tunneling transformation is applied, whereby all the “relief” the function may have above that minimum is flattened out, while the other values of the function (lower than that minimum) are left unchanged. Then a repelling mechanism is turned on whereby the search point is pushed away from the current position on a flat surface and is set free to explore and discover other (necessarily lower) minima. Once a new (lower) local minimum has been reached the procedure is applied anew. By the very design of the algorithm, the last minimum found is the global one.

TRUST allows the formulation of an efficient and reasonable stopping criterion to ensure proper balance between cost and accuracy. This stopping criterion, first designed at ORNL/CESAR for machine learning, has been successfully carried over to large-scale global optimization problems. As a result, TRUST has achieved a dramatic improvement in terms of time (cost) and accuracy not only for rigorous mathematical benchmarks, but—most importantly—over leading methods currently used in subsurface imaging, multi-sensor integration, and redundancy resolution in robot manipulation.

Overall, TRUST's efficiency is 5 times higher than that of its best competitor and 45 times higher than that of its weakest competitors (from a selected group including the best-performing known competitors) on standard non-convex optimization benchmarks.

The TRUST methodology received a 1998 *R&D 100* award.

For more information contact:

J. Barhen, ORNL
barhenj@ornl.gov