Acoustic Benchmark Validation of GRASP

Donald R. DelBalzo and Erik R. Rike
Neptune Sciences Div. of Planning Systems, Inc.
40201 Hwy 190 E. Slidell, LA 70461

Abstract*–*A genetic algorithm is used in non-homogeneous and anisotropic environments to nearly optimize sonar search tracks. The optimization metric is maximum cumulative detection probability against a target with specified characteristics (acoustic and tactical) during a fixed time period. This application for search planning is named GRASP, for Genetic Range-dependent Algorithm for Search Planning. A validation of GRASP solutions in various ocean environments is shown under benchmark conditions, i.e., fairly simple environments and target distributions. Directional, range-dependent sonar performance is estimated from Parabolic Equation calculations of transmission loss. The search tracks produced by the genetic algorithm are generally intuitive; they usually remain in high detection areas. When track solutions are counter-intuitive, we explain unexpected behavior (e.g., zigzag turns, tracks offset from symmetric features, and occasional departures from high detection areas) in terms of details in the acoustic field.

1 Introduction

Search planning has been the subject of active research, both academic and operational, since at least 1946. Theoretically optimal effort allocations have been obtained for search against both stationary targets [1] and moving targets [2], under the idealized assumption that search effort is an abstract quantity that may be allocated arbitrarily over the search region, subject only to a constraint on the total effort expended at any instant. Such solutions do not reflect the practical constraints that an acoustic searcher can expend effort only in the vicinity of his current location, can be in only one place at one time, and must expend time to move from place to place. There has also been much research on the discrete searcher path problem [3-5 and associated references]. In that problem, the search takes place in discrete space and time, the searcher and target are constrained to move at most one grid cell at each time step, and the searcher's effort is uniformly distributed over his current cell. Virtually all research on either the effort-allocation or the discrete-path problem has assumed that return on search effort is governed by the classic exponential search formula [1] and that targets do not react to the search.

The continuous search-path problem is much harder than the effort-allocation and discrete-path problems. We are unaware of any algorithm that produces provably optimal solutions in any but the most trivial cases of the continuous-path problem. The Genetic Range-dependent Algorithm for Search Planning (GRASP) is an application of a genetic algorithm (GA), which generates search tracks that maximize Cumulative Detection Probability (CDP), in order to solve the continuous-path problem. Initial GRASP results for simple definite-range performance reproduced intuitively optimal solutions in a modest number of generations, proving its computational efficiency [6].

Simulations were performed in complicated, non-homogeneous environments [7] where the "best answer" can be approximated intuitively. That work assumed a range-dependent propagation over flattened coarse-sand ridges 10 nm wide and 200m below the surface (high detection range), rising over a silty-clay basin 1700m deep (low detection range). Examples (double ridge and square annulus) are shown in Fig. 1. Those results showed consistent qualitative agreement between intuitive and GRASP tracks. This new work builds on the results of [6–7] by expanded testing in the same environments.

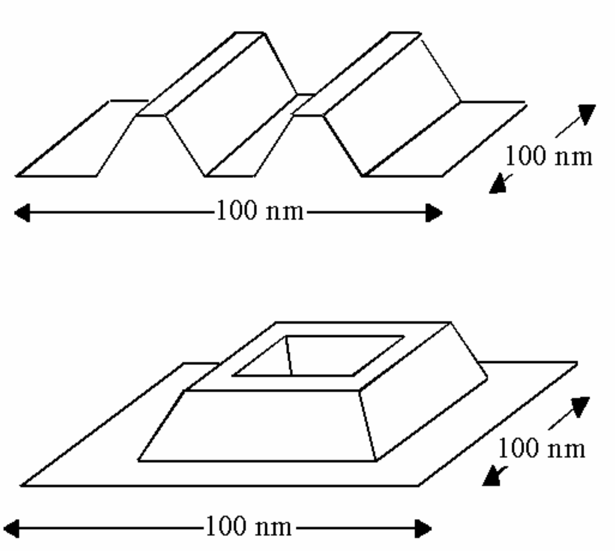

Fig. 1. Double ridge (above) and Square Annulus (below). The vertical scale is severely exaggerated.

2 Genetic algorithms

Genetic algorithms are an attempt to solve otherwise intractable problems by mimicking the process of evolution. A good introduction to the subject may be found in [8]. A somewhat simplistic view of evolution is that individuals in a population can be described by a chromosome, a reproductive mechanism (either sexual or asexual), mutation, and natural selection (survival of the fittest). A chromosome is a sequence of genes, each of which describes some aspect of the individual's structure. The chromosome as a whole completely determines the individual's characteristics, and in particular, its fitness.

Reproduction may occur either sexually or asexually. In the former, offspring are produced by interchanging portions of their parents' chromosomes. In the latter, each child is (except for mutation) an exact replica of the single parent. Most complex organisms reproduce sexually, and there is considerable evidence that even organisms that normally reproduce asexually occasionally engage in sexual reproduction. Either type of reproduction involves the possibility of mutation, in which some genes are randomly altered before being passed to the child. Natural selection tends to eliminate the least fit individuals in favor of the most fit, leaving the most fit to pass on their genes to the next generation.

Some genes, or combinations of genes, may increase the individual's fitness relative to the population as a whole. Some may reduce the individual's fitness, while others may have little or no effect. A species that reproduces entirely asexually can only improve by fortuitous, random mutations. Sexual reproduction offers a much greater chance of improving the species. When two parents exchange genes, the offspring have a much richer gene pool from which to draw. With luck, some child will inherit a combination of the genes that made each parent sufficiently fit to survive and reproduce, and become more fit than either parent. (Of course, some other child may inherit all of the bad genes, but that child will likely not survive.) By either method of reproduction, the odds favor the survival of good genes and the elimination of bad genes.

A GA mimics the process of evolution within a computer. The chromosome is an encoding of a trial solution to the optimization problem. The fitness function is the problem's objective function (or some monotone transformation thereof). Natural selection is mimicked by choosing individuals to reproduce with probability proportional to their fitness. Sexual reproduction is accomplished by exchanging segments of the parents' chromosomes. Asexual reproduction is a simple cloning process. Mutation is applied to the offspring by randomly perturbing some aspects (genes) of the trial solution. Many GAs use a combination of sexual and asexual reproduction.

3 Problem statement

We consider the problem of designing a fixed-speed search path, in continuous space and time, for a fixed time period, *T*, through a closed, connected, two-dimensional region with the intention of detecting either a fixed or moving target confined to the same region. The objective function is the CDP at time *T*. We assume that the searcher has a fixed 10-kt speed (but no other motion constraints) and that we can compute the CDP of a given searcher path against the distribution of target tracks. In GRASP, we evaluate the CDP by Monte Carlo simulation. Specifically, we simulate a large number of targets, and compute overall CDP as the average of the CDP against individual targets. To simplify this benchmark analysis we arbitrarily choose the Monte Carlo targets to be stationary and distributed uniformly on the grid. We assume that the CDP function is governed by a lambda-sigma fluctuation model [6].

A variety of benchmark cases were run. The geometries and starting locations are listed below and diagrammed in Fig.2.

1) Square Annulus:
    1: center of side of annulus
    2: corner of annulus
    3: corner of search area
    4: center of side of search area
2) Circular Annulus:
    1: center of annulus ridge
    2: corner of search area
    3: center of side of search area
3) Double Ridge
    1: on ridge 10 nm from lower end
    2: center of side of search area

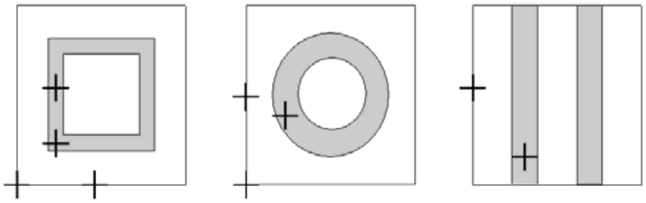

Fig. 2. Ridge shapes used in this benchmark analysis and nine starting positions for search plans.

4 Previous results

This section is a review of the major results from [6–7]. When an intuitive solution is obvious, GRASP paths spatially correlate with intuition. However, GRASP paths are rarely the straight paths that intuition favors, yet they almost always outperform intuition. This, in turn, improves intuition. Sometimes the shortest path to success is not a straight line. Analysis of benchmark trials reveals three competing principles, which are only strictly true for stationary targets and definite-range detection functions.

1) *Sharp turns introduce losses.* Sharp turns always introduce an area of overlapping sweep area. See Fig. 3. This overlap maximizes (to 100%) for a U-turn, but soft changes in direction tend to be low-cost.

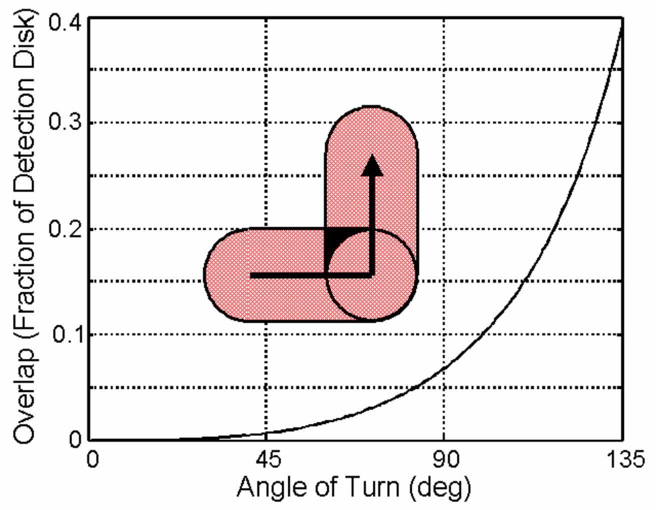

Fig. 3. Sharp turns (with radius of curvature smaller than the sweep radius) sweep some areas redundantly.

2) *Anisotropy implies a preferred direction of travel.* In a homogeneous but anisotropic region (*i.e.*, sweep width in a given direction is constant from point to point, but sweep widths in different directions vary at any given point), the preferred travel direction is perpendicular to the direction of maximum sweep width. See Fig. 4. This obviously maximizes sweep area, and therefore CDP. This means a straight line is optimal in any such region if search time will be exhausted before a boundary is encountered.

3) *Inhomogeneities and boundaries often favor non-linear paths.* A single, straight search path is no longer possible if the preferred-direction dimension is small compared to the search time available, and may no longer be optimal if the search area is inhomogeneous. An "optimal" path will negotiate the best possible trade-off between the first two principles, which usually involves some wiggling. Such will typically be the case on the ridges: detection range will be greatest when the searcher is on the ridge and looking along it, and so the preferred direction of travel will be across the ridge (and therefore quite short). Wiggling is the inevitable result, but wiggling with the most gentle serpentine motion possible. It is precisely such paths that GRASP tends toward even in the simple environments studied during these benchmarking trials.

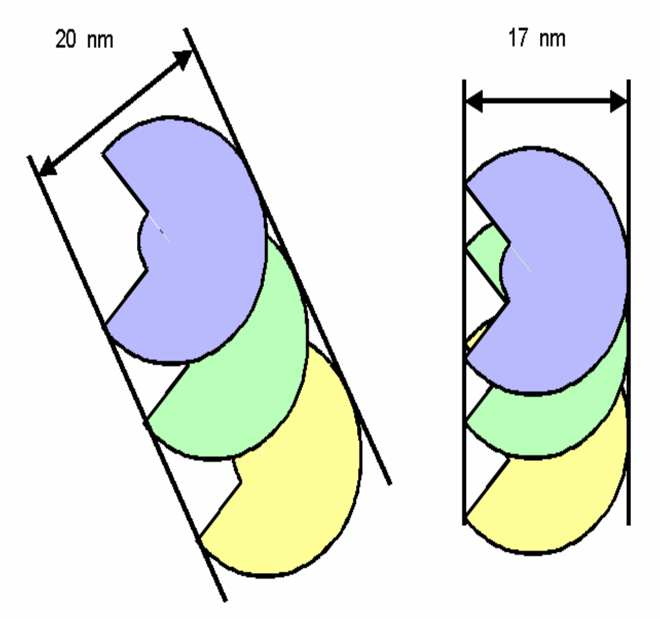

Fig. 4. An anisotropic pattern sweeps out different areas depending on the direction of travel.

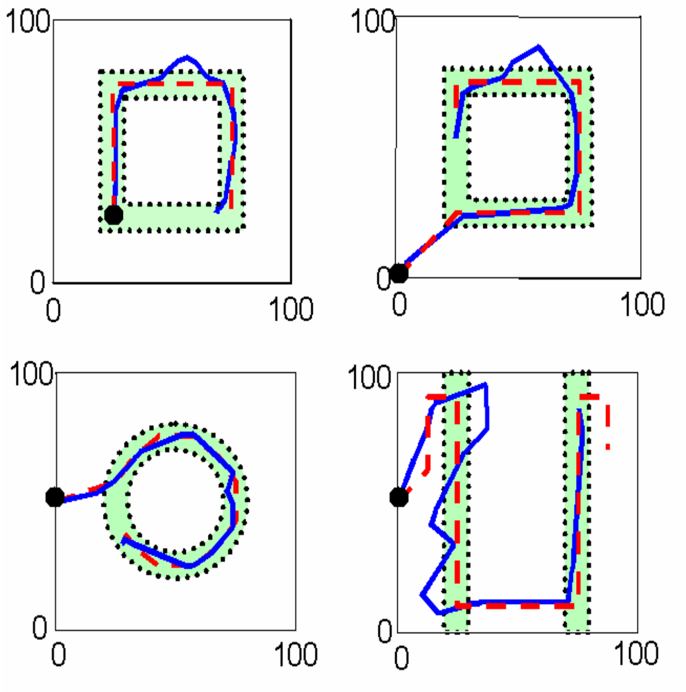

Fig. 5. Example GRASP benchmark solutions in 100 x 100 nm search areas. Dotted lines are the ridge tops, solid lines are GRASP tracks, and dashed lines are intuitive tracks. At upper left is the 15-hr Square Annulus case, at upper right is the 20-hr Square Annulus case, at lower left is the 16.5-hr Circular Annulus case, and at lower right is the 30-hr Double Ridge case.

5 Results

As stated in the introduction, we hoped that GRASP would produce tracks nearly as good as intuitive tracks. Fig. 5 shows four example solutions (taken from the twenty trials for this work) provided by GRASP compared to intuitive solutions. A solid circle indicates the starting point of each track. GRASP solutions generally follow the intuitive solutions, *i.e.*, tracks initially move to the ridge (if they do not start on it), and then follow the ridge. The resulting CDPs for GRASP were consistently greater than those for the intuitive solutions, however, despite the turns and occasional departures from the ridge. The numerical improvements are shown in Tables 1-3 for the nine geometry/starting point combinations and for several different total search times. The improvements, quantified by mean differences, are 1.4% (square annulus), 0.8% (circular annulus), and 1.3% (double ridge).

GRASP turns in Fig. 5 tend to be gentle, and there is wiggling from side to side on the ridges. However, these solutions also tend to spend some time off the ridge. The 30-hr search in the double ridge case (lower right) is a dramatic example, but this behavior is likely caused by a search time that is long compared to that necessary to just cover the ridge. The top 2 examples, however, show tracks leaving the ridge when clearly the search times are short enough that intuition would suggest staying on the ridge. This counter-intuitive behavior is explained below.

Table 1
Square Annulus

Start Location	Time (hr)	GA	User/ Intuitive	Differ-ence
1	15.0	25.51	24.12	1.4%
	17.5	28.87	27.42	1.4%
2	17.5	28.70	27.58	1.1%
3	20.0	29.74	28.15	1.6%
	22.0	32.24	30.99	1.3%
4	20.0	30.13	29.02	1.1%
	21.0	31.44	29.96	1.5%
	22.0	32.40	30.95	1.5%
Mean difference is 1.4%				

Table 2
Circular Annulus

Start Location	Time (hr)	GA	User/ Intuitive	Differ-ence
1	13.5	23.00	22.15	0.8%
	17.5	26.78	25.58	1.2%
2	19.0	26.22	26.20	0.0%
3	16.5	25.11	23.87	1.2%
Mean difference is 0.8%				

Table 3
Double Ridge

Start Location	Time (hr)	GA	User/ Intuitive	Differ-ence
1	8.0	14.26	13.12	1.1%
	18.5	28.29	26.35	1.9%
	20.0	29.94	28.54	1.4%
	21.0	31.23	30.00	1.2%
2	19.5	26.66	25.82	0.8%
	23.0	30.99	29.74	1.3%
	30.0	38.09	36.69	1.4%
	40.0	45.53	44.39	1.1%
Mean difference is 1.3%				

Fig. 6 shows the size of the detection areas for greater than 50% probability of detection (PD) for areas on a north-south ridge (away from the corners). At the west edge of the ridge (left), the detection area is roughly a semicircle with poor detection downhill to the West. At the center of the ridge (right), detection is symmetric both along (N-S) and perpendicular (E-W) to the axis of the ridge, but has twice the detection range along the ridge as perpendicular to it. For paths that travel exactly parallel to the ridge axis (*i.e.*, N-S), the greatest sweep width is achieved when the track is offset about halfway (2.5 nm) from the center of the ridge (middle panel in Fig. 6). The sweep widths (measured E-W) for the 3 examples in Fig. 6 are 14.76 nm (west edge of ridge), 14.88 nm (halfway between center and edge), and 12.71 nm (center of ridge).

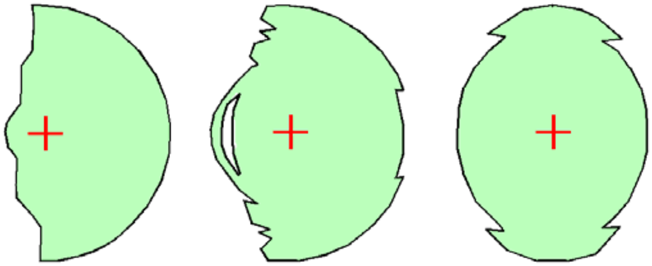

Fig. 6. Areas with >50% PD for ridges oriented N-S. Cross-hairs indicate searcher location. Left: west edge of ridge. Center: halfway to center of ridge. Right: center of ridge.

Fig. 7 illustrates that (using these benchmark calculations) sinusoidal or zigzag motion on the ridge can result in a greater effective sweep width than straight-line motion directly along the ridge. This is an extension of the result shown in Fig. 4. Increasing effective sweep width with wiggling is a good tactic; however, it is more important to increase sweep rate, the rate at which non-overlapping area is added, which argues for *gentle* wiggling. For example, for a zigzag path from the center of the ridge to about 2.5 nm offset and back, with end points 35 nm apart, the actual track distance is 1.0% longer than a straight-line track. If the effective sweep width is more than 1.0% wider, then wiggling is beneficial. Hence, the effective sweep width need be only about 0.36 nm greater than the 14.88 nm of the straight-line offset path. The width of the dashed vertical lines in Fig. 7 is 15.40 nm, or 0.52 nm wider than the sweep width of the offset path. Hence, wiggling need only include about 68% of this extra area to give a greater sweep rate. GRASP solutions that maximize CDP will automatically include the correct number and degree of turns.

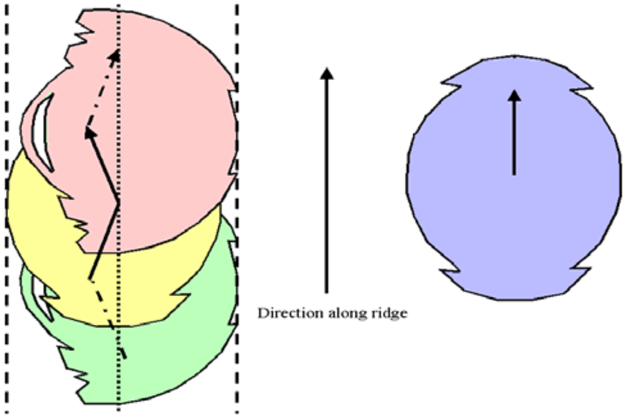

Fig. 7. Wiggling (exaggerated) on the ridge.

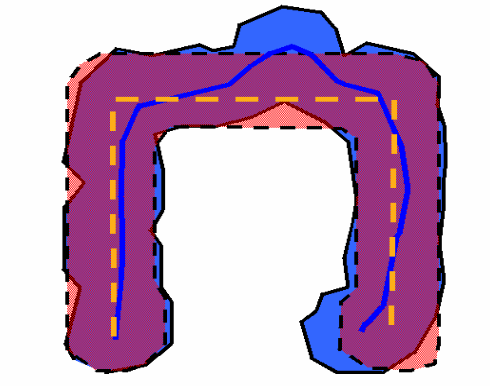

Fig. 8. Area swept by GRASP (solid line and border) and intuitive (dashed line and border) search paths for the 15-hour Square Annulus case. For comparison, the edge of the ridge is indented about 1.5 nm from the dashed border.

The final issue addressed here is the occasional departure from high detection areas, as shown in the top two cases in Fig. 5. Fig. 8 shows the area swept for GRASP (solid line and border) and the area swept for the intuitive path (dashed line and border) during the 15-hour square annulus search. Since detection range along the ridge is large, very little is lost by GRASP momentarily leaving the ridge. Specifically, the extra area swept (enclosed within solid border but outside dashed border) is about two or three times the area lost (enclosed within dashed border but outside solid border).

6 Conclusions

GRASP search tracks are in general qualitative agreement with intuitive tracks for several benchmark acoustic cases. We have shown that certain counter-intuitive detailed behavior (wiggling and occasionally leaving the high detection areas) can be advantageous. In fact, GRASP tracks outperform intuitive tracks in all cases examined here, based on cumulative detection probability. The improved performance results from exploiting the anisotropic nature of the benchmark acoustic fields. Work is continuing to enhance the capability of GRASP.

References

[1] Koopman, B.O. *Search and Screening*. Operations Evaluation Group Report 56, Navy Department, 1946. Reprinted (with additions) as *Search and Screening: General Principles with Historical Applications.* Pergamon Press, 1980.

[2] Brown, S.S. "Optimal search for a moving target in discrete space and time," *Operations Research*, Vol 28, No 6, 1980, pp. 1275–1289.

[3] Washburn, A.R. "Probability density of a moving particle," *Operations Research*, Vol 17, 1967, pp. 861–871.

[4] Washburn, A.R. *Search and Detection.* Institute for Operations Research and Management Sciences, 1996.

[5] Washburn, A.R. "Branch and bound methods for a search problem," *Naval Research Logistics*, Vol 45, 1998, pp. 243–257.

[6] Kierstead, D.P., and DelBalzo, D.R. "A genetic algorithm approach for planning search paths in complicated environments," *Military Operations Research* **8**(2) 45:59 (April 2003).

[7] DelBalzo, D.R., Wagstaff, M.D., Jr., Rike, E.R., and Hemsteter, K.P. "Environmental – acoustic impact on optimum sonar search," *Proceedings* of the 17th International Conference on Acoustics, Rome, 2001.

[8] Goldberg, D.E. *Genetic Algorithms in Search, Optimization & Machine Learning*. Addison-Wesley, 1989.