the stress state parameterization described above, consisting of the variables $(\phi, \delta, \Omega, \psi)$. Here the borehole azimuth, δ, and the borehole deviation, Ω, are the first two Euler angles, and ψ is the third Euler angle. A N-dimensional optimizer is then initialized with the GAs results to find the best fitting stress state.

Genetic algorithms are an attractive approach to solving hard, nonlinear problems in which the forward calculation is straightforward but more traditional techniques might fail. Some of the advantages of GAs are that they efficiently search the problem space, do not require a good starting model, and do not get trapped in local minima (unlike gradient search methods).

GAs operate on a population of models. The models are typically binary coded, just as floating point numbers are encoded in a computer. An individual binary-encoded model is termed a chromosome. The first generation of chromosomes is randomly generated. Each chromosome has associated with it an "objective value," which is a problem-specific measure of how well the chromosome solves the problem. From the objective value is calculated a "fitness value": the higher the fitness, the better the chromosome. In our work, each chromosome represents a distinct regional stress state, and the objective value is the weighted one-norm misfit between the observed and theoretical **IJK** breakout angles. Since smaller one-norm misfits correspond to higher levels of fitness, we use the following equation to relate the two:

$$f_i = \frac{\bar{o} + \sigma_o - o_i}{\sigma_o}, \tag{12}$$

where $\bar{o}$ is the population's mean objective value, σ_o is the standard deviation of the objective values, o_i is the ith chromosome's objective value, and f_i is the ith fitness value [*Holland*, 1975].

After the fitness values have been calculated, chromosomes are selected to "mate" to create the next generation of chromosomes. Chromosomes with higher fitness values proportionately mate more often. Mating between two chromosomes is performed by randomly exchanging part of the binary patterns of both parent chromosomes. This operation is known as "crossover." Crossover is performed only roughly 60% of the time between two chromosomes. In the other 40% of matings the two chromosomes are carried directory into the next generation without crossover.

The final operation of the GA is mutation, whereby a small fraction of the bits of a chromosome are flipped. This process introduces variability into the population and allows broader searching of the solution space. The next generation of chromosomes has now been created, and the cycle begins anew with the calculation of the population's objective values.

We run our GAs with populations of several hundred chromosomes for several hundred generations. Since GAs do not guarantee that they find the optimal solution to a problem, we take the particular stress state corresponding to the minimum misfit observed for all of the GA runs completed as the starting point for a general N-dimensional optimizer routine. We use the Powell optimizer as described and coded by *Brent* [1973]. This particular optimizer does not require the use of derivatives of the function it is trying to optimize, a feature that makes it attractive for the borehole breakout problem.

Experience with various breakout data sets has shown that starting the Powell optimizer with a purely random stress state does not lead to a global minimum of the one-norm misfit, even for theoretically generated breakout data. The GA thus supplies an important first step in efficiently searching the space of stress states for good starting stress states for the Powell optimizer. In practice we run roughly several hundred separate GAs, each with different random numbers, to make sure we find the global one-norm misfit minimum. Even though we use the Powell optimizer after the GA, the misfit is not a smooth function of the stress state, since different GA and Powell runs find different local minima. We could have chosen to start a GA search around the results of the Powell search, but we choose not to do so, since such a search would probably not improve the results all that much.

Confidence Limits

For the results of the stress state optimization to be of any worth the confidence limits in the results must be determined. We determine the confidence limits on the best fitting stress state, using a modification of *Gephart and Forsyth*'s [1984] technique for defining confidence limits on stress state determinations from focal mechanisms.

We want to find the 95% confidence limit on the model we obtained that minimized the one-norm misfit (equation (11)). We do this by assuming that by performing this one experiment (an experiment in the sense that oil wells were drilled and borehole breakout orientations were measured) we know something about the ensemble of one-norm misfits of the best fitting models that would result if we were able to perform the experiment again as many times as we liked. From this assumption we can say that our particular misfit lies