Adaptive Statistical Designs

Janis Hardwick Quentin F. Stout
University of Michigan
ALATE – ALgorithms: Applied, Theoretical, Experimental
www.eecs.umich.edu/~alate/

- √ Adaptive designs allow investigators to adjust sampling decisions while the experiment is being carried out. Using them one can, for example: *optimize sampling costs*; *minimize experiment length*; and *reduce patient failures* during a clinical trial.
- √ Contrast this with fixed allocation designs in which sampling is specified prior to the beginning the experiment.

Scientific Problem: Despite their advantages, adaptive designs are rarely used, because they are far more complex to analyze and optimize. Analytic analyses are typically impossible.

Computational Challenge: Adaptive designs yield very large state spaces. Backward induction or dynamic programming is the classic approach taken for analysis and optimization. For example, a trial of N patients and T treatments, with merely binary outcomes, yields a state space of size $\Theta(N^{2T})$. The computational time grows similarly. This was previously thought to be computationally infeasible.

Progress: Using some new algorithms, careful implementation, and parallel computers, we have optimized and analyzed 3-treatment trials with $N = 200$ using only 16 processors of an IBM SP2 at the University of Michigan's Center for Parallel Computing. This problem is more than *100,000* times harder than the largest previous work (done on a Cray Y-MP). More generally, we can solve near-neighbor recurrences of arbitrary dimension using programs that are both time- and memory-efficient and *highly scalable*.

Example: Consider a clinical trial with 100 patients and two treatment options T_1 and T_2. The unknown success rates for T_1 and T_2 are p_1 and p_2. This plot illustrates the number of "extra" failures that occur if we use the optimal **adaptive design** rather than the omniscient design that assigns all patients to the superior treatment. The values plotted are referred to as the *expected successes lost* or the *regret* of the adaptive design. Note that the maximum value in the plot is less than 3.5. Had we used the standard fixed sampling design, (with half the patients on each treatment) $\longrightarrow$ the expected successes lost would be an order of magnitude greater with the maximum being **50** extra failures!