

Current programs to discover causal relationships between genetic markers and phenotypes take a long time to run, and the accuracy of the results often depends on how long the researcher is willing to let the program run. Using parallel computing to decrease run time mitigates the effect of the trade-off between speed and accuracy. I implemented two parallelizations of a program that uses a Markov chain Monte Carlo search to find groups of genetic markers associated with a phenotype. The first parallelization runs multiple Markov chains in parallel, while the second runs a single chain but calculates transition probabilities in parallel. While the second method did not speed up the program, the first method offered a speedup of more than two times the original run time when run on a four core machine, and the speedup is likely to increase as the size of the data set increases. The speedup makes the program more practical for use with larger data sets, allowing researchers to consider more genetic markers and combinations of genetic markers in searching for genetic causes of particular phenotypes.