

A Comparison of Ordered Greed to a First Fit Approach using the Bin Packing Problem

*Scott H. Douglas¹
Rochester Institute of Technology, Rochester, NY
February 23, 2005*

Abstract

The First Fit approach to solving the Bin Packing problem using a genetic algorithm (GA) does not always find the solution in an optimal timeframe. To this end, the Ordered Greed approach was created to solve problems with less computation. The Ordered Greed approach can be applied to many types of problems, including Bin Packing. This paper compares the Ordered Greed approach to the First Fit approach.

1. Introduction

The Bin Packing Problem attempts to place O objects of different weights into M bins of size T optimally. The optimization may attempt to minimize the number of bins used or to reduce the average waste of each bin. The Bin Packing problem has many applications ranging from optical drive storage allocation to truck packing. This paper uses a one dimensional bin packing system, but this could easily be adapted for use in N -dimensional space.

With the First Fit approach, a permutation of the O objects is created. The permutation is looped over adding elements to the next available bin. Once a bin is full, you begin adding elements to the next bin. You never go back to a bin once you've moved on.

With Ordered Greed, a permutation of the O objects is created. The permutation is looped over adding elements to the first bin it will fit into.

Since fitness evaluations are costly, they are a good metric of how a GA is performing. By running the Bin Packing problem with different mutation rates and population sizes, we hope to find sweet spots that minimize fitness evaluation calls and to provide empirical data for comparing approaches to solving permutation based GAs.

One caveat that should be noted is that the Ordered Greed fitness evaluation does more work than the First Fit algorithm. It takes more time to run because of this.

2. Software Design

2.1. Functional Description

The GA system must evolve a solution to the Bin Packing problem in the form of a permutation.

2.2. System Overview

The system implemented is similar to the system implemented in "Exploring Permutation-Based Mutation Operators with the N-Queens Problem" by Scott Douglas.

Chromosomes were permutation based instead of bit string based. Instead of an array of bits, the chromosomes consists of an array of integers. This makes crossover and mutation easier. In this system, it is impossible to morph a permutation element into another element. You can only shift elements. This ensures that the permutation is always valid.

Interfaces were created for defining custom crossover, fitness, and mutation functions. To use the system to solve a problem, simply provide a crossover function, a fitness function that measures the fitness of the permutation, and a mutation operator. Pass these parameters into a Population instance, which randomly generates the initial population. At this point, statistics about the population can be queried.

To advance to the next generation, call `nextGeneration()`. This function copies a certain percentage of individuals to the next generation based on fitness. For the remaining individuals, tournaments of two individuals are held to find two parents. Crossover is performed based on the crossover rate. This is done until the new generation is full. Mutation is employed based on the mutation rate parameter. The next generation process is repeated until the solution is found or a maximum number of fitness evaluations has been exceeded.

When gathering statistics, such as the average number of fitness evaluations required, the system runs the GA multiple times based on the repeat rate.

2.3. Initial Parameters

- Crossover Rate: 0.7

¹Send correspondence to `shd0326@cs.rit.edu`.

- Crossover Type: PMX
- Mutation Rate: 0.035 When Not Varied
- Max. Fitness Calls: 300,000
- Population Size: 50 When Not Varied
- Chromosome Length: 50
- Copy Rate: 0.1
- Repeat Rate: 5
- Measurements: Average

2.4. Encoding

Permutations of the range $[0, N - 1]$ are encoded as a series of N integers in a chromosome. Each integer represents an index into an array of object sizes that is held in the fitness function.

2.5. Fitness Evaluation

The GA system attempts to minimize the number of bins used.

3. Testing

In order to ensure the system was working properly before beginning experimentation, testing was performed on a simple data set with a known solution. The system was able to converge on correct solutions to the Bin Packing problem. The initial parameters were used.

4. Experimentation

Using the Bin Packing problem with the initial parameters, the mutation rate and population size were varied to determine sweet spots and to provide a comparison between the Ordered Greed and First Fit approaches. The results provide empirical evidence for comparing the approaches. The object set used for experimentation was N1C1W2_E.BPP from Dr. Klein's bin packing website at <http://www.wiwi.uni-jena.de/Entscheidung/binpp/index.htm>. This data set has a known optimal packing of thirty six bins.

5. Results

Refer to Figure 1 for a graph of the mutation rate varied while using both the First Fit and Ordered Greed approach. The Ordered Greed approach has no sweet spot because it solved the problem in the first generation every time. The number of fitness evaluations required is proportional to

the population size, which was fifty. The sweet spot for the First Fit approach was with a mutation rate of 0.026, which required 2,330 fitness evaluations. Any mutation rate value from the range $[0.015, 0.060]$ would provide a minimal number of fitness evaluations.

Refer to Figure 2 for a graph of the population size varied while using both the First Fit and Ordered Greed approach. Ordered Greed solved the problem within the first generation every time. Due to the system implementation, this caused the fitness evaluations required to be directly proportional the population size. The First Fit approach tended to require more fitness evaluations for larger population sizes. As the population size got larger, the variance between readings began to grow. For both approaches, the smallest population size gave the lowest number of fitness evaluations required.

6. Conclusion

Ordered Greed is far superior to the First Fit approach when working with the Bin Packing problem. Ordered Greed was able to solve the problem in the first generation for all the test cases that were used. The performance of the First Fit approach didn't come anywhere near this efficiency. The Ordered Greed approach was usually dependent on the population size. The mutation rate had no effect. The mutation rate greatly affected the First Fit approach. The mutation rate sweet spot of 0.026 identified in this paper agrees with results from the same author in earlier papers. Mutation is necessary for convergence, but too much mutation takes longer to converge.

Both the Ordered Greed and First Fit approach prefer smaller populations. This makes sense for the Ordered Greed approach since it always finds the solution almost immediately. For the First Fit approach, this was not intuitively obvious. This suggests that the bin packing problem works best with minimal genetic diversity.

The larger variance between successive readings in the population size graph makes sense. Since the experiment is repeated based on a constant repeat rate, the points picked are likely to be spread farther apart as the population size increases.

Due to the extent at which Ordered Greed reduced the number of fitness evaluations required, the extra time spent in the fitness evaluation step for Ordered Greed becomes inconsequential.

Future work could compare the Ordered Greed approach and the First Fit approach to other approaches, but at this point, it seems like it would be hard to beat Ordered Greed for the Bin Packing problem. These experiments could also be run in other problem spaces to see if the same results hold true.

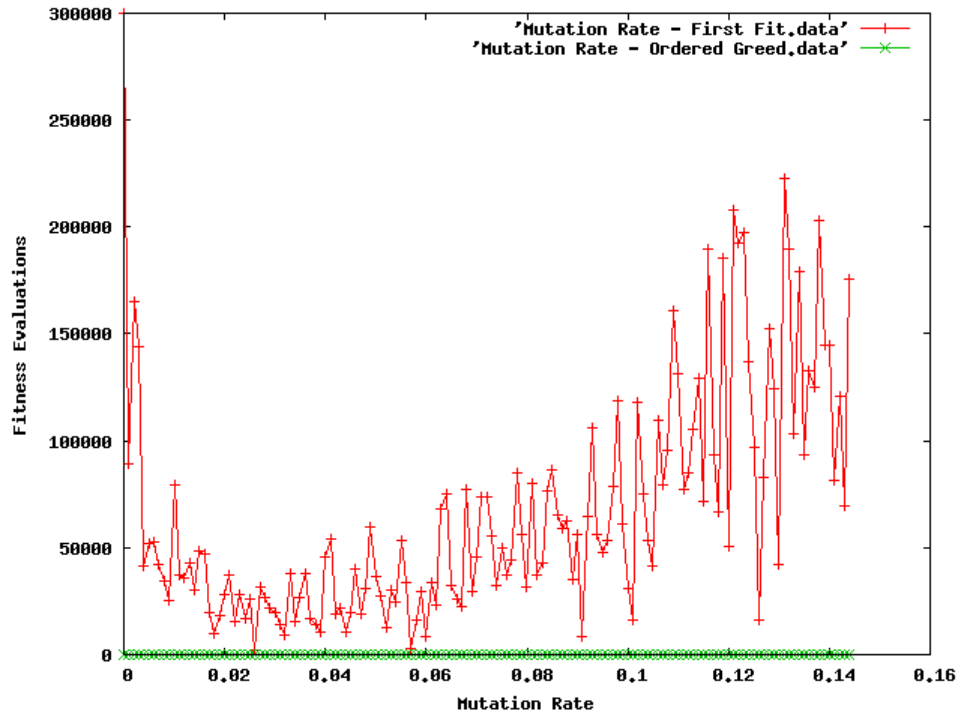


Figure 1: The mutation rate was varied to determine a comparison between the First Fit approach and the Ordered Greed approach.

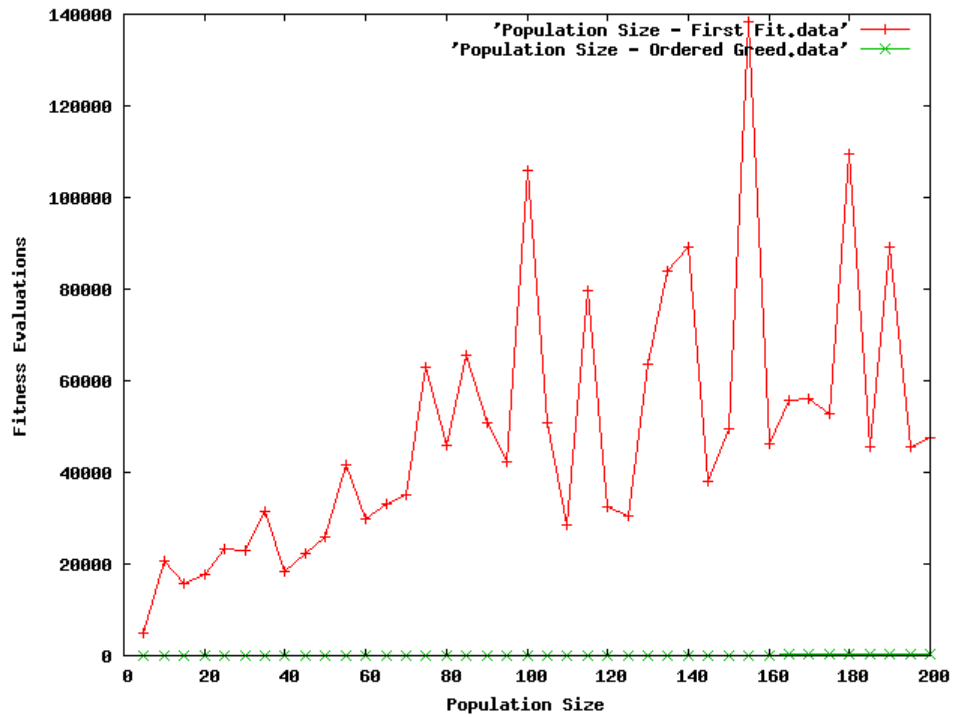


Figure 2: The population size was varied to determine a comparison between the First Fit approach and the Ordered Greed approach.