

=====

VRP SOLVER - ReadMe.txt

=====

Version 1.3

VRP Solver and any attached documentation are

(c) 2003 Lawrence V. Snyder

Department of Industrial and Systems Engineering

Lehigh University

Bethlehem, PA, USA

1. PURPOSE OF VRP SOLVER

VRP Solver implements an adaptation of the Clarke-Wright savings algorithm for vehicle routing problems. It takes input from a text file listing each customer's location (latitude and longitude) and demand. It builds vehicle routes that visit every city exactly once and that obey user-specified vehicle volume and distance limits. Results are displayed in graphical (map) form and in text form.

2. ALGORITHM IMPLEMENTED

The Clarke-Wright savings algorithm is a well-known algorithm in vehicle routing and is described in various papers and texts. The algorithm implemented by VRP Solver expands the Clarke-Wright algorithm in the following ways:

- Randomization: Instead of choosing the best pairing of routes at each step, choose one of the k best pairings, chosen randomly. Repeat several times and choose the best overall solution.
- Improvement Heuristics: After an initial solution is built, various improvement heuristics are performed. These include the well-known 2-opt and Or-opt operations (the Or-opt uses group sizes of 1, 2, and 3), as well as a swap operation in which two customers on different routes may be removed from their routes and inserted into the opposite route.

3. USING VRP SOLVER

The main screen in VRP Solver allows you to load and view the data, set options, and run the model. To load data, type the path and filename of the data file into the box labeled "Load Data From File:", or click the "Browse..." button to search for the file. After typing the filename or filling it using the Browse button, you must load the data by clicking the "Load" button. If the data are successfully read, they will appear in chart form in the box in the middle of the window. See section 5 for details on the data file.

At the bottom of the main screen are two boxes in which you can set the parameters of the problem. The "Truck Capacity" box specifies the maximum allowable volume a vehicle may carry (in units of demand). The "Truck Distance Limit" box specifies the maximum allowable distance a truck may travel in its round trip. If any customer's demand exceeds the truck capacity or is farther from the depot than 1/2 the truck distance limit, it will be placed on a route of its own.

The "Distances" button displays a window showing the distance matrix among pairs of cities in the data and allowing you to choose how distances are computed. See section 6 for more details.

The "Options" button displays a dialog box in which you can fine-tune the algorithm. See section 7 for more details.

To run the model, click the "Run Model" button. The progress bar at the bottom of the screen will indicate how far along the process is.

4. SOLUTION SCREEN

Once the model has finished running, VRP Solver will display the solution in a new window that has two tabs. The "Map" tab displays the solution in graphical form, with each city represented by a circle and routes drawn as lines connecting the customers. Each route is represented by a separate color, up to a maximum of 12 routes; routes after the twelfth are drawn in black. Check or uncheck the "Display Labels" box to turn the node labels on or off. Click "Save Bitmap..." to save a picture of the map in bitmap format.

The "Text" tab displays the solution in text form. For each route, VRP Solver lists the stops, in order, as well as the demand at each stop and the distance to each stop from the previous one. The total volume and distance for each route are also listed. After the route listings, the window indicates the total solution time, the time spent building the initial solution and improving it, and the number of each type of improvement that was performed. There is no facility to save the text, but you can copy and paste it into a text editor and save it from there.

At the bottom of the screen, when either tab is active, is a brief summary of the solution indicating the number of trucks used, the total distance, and the solution time (in seconds).

5. DATA FILE FORMAT

The data file must be in the following format: The first line contains the x-coordinate and the y-coordinate of the depot, in that order, separated by spaces or tabs. Each successive line contains the x-coordinate, y-coordinate, and demand of a single customer, separated by spaces or tabs. If distances are to be loaded from a file instead of computed based on coordinates, the x- and y-coordinates may be set to 0 or any other value, but they may not be left blank. Note that the depot has no demand. (A sample data file can be found in the attached file "data20.txt".)

6. DISTANCES

When you click the "Distances" button on the main form, you will see a window displaying the distance between each pair of cities in the data set. At the top of the form, you can choose whether to compute distances or load them from a file. If you want distances to be computed, you can choose whether to use Euclidean or great circle distances. If you want to load them from a file, enter the file name where indicated or click the Browse button to choose a file. You can also request VRP Solver to round all distances to the nearest integer by checking the obvious checkbox. The default (when you first load a data set) is Euclidean distances and no rounding.

If you change an option, the distances will not automatically be recomputed. To compute distances, you must click the "Compute Distances" button. The text below the Preferences box indicates whether the distance matrix as displayed reflects your current settings, or whether the distances need to be recomputed.

You can save the current distance matrix to a text file by clicking the "Save Distances" button.

When loading distances from a matrix, be aware of the following:

- Distances must be arranged in a grid containing the distance between each pair of cities, including the depot, in the same order in which they are listed on the main form. Successive entries must be separated by spaces, tabs, or carriage returns.
- VRP Solver does not check to make sure that the distances in the file are symmetric (that is, that the distance from i to j is the same as the distance from j to i), but the algorithms used by VRP Solver are not designed for asymmetric distances and may produce unpredictable results if the distances are asymmetric. Therefore, you should make sure that the distance file contains symmetric distances.

7. OPTIONS

The "Options" button on the main form produces a dialog box containing several options for the algorithm. The "Randomization Depth" and "Randomization Iter" options specify the depth and number of iterations for the randomization procedure used while building the initial solution. Randomization works as follows: In the first iteration, the best pairing of routes is chosen at each step. In the second iteration, one of the 2 best pairings is chosen at random at each step, and 'iter' solutions are built in this way. In the third iteration, one of the 3 best pairings is chosen at random at each step, and 'iter' solutions are built in this way. This continues until solutions are built choosing one of the 'depth' best pairings at each step. In total the procedure produces 'depth' x 'iter' solutions (possibly containing duplicates). The best solution is selected and improvement heuristics are run on it. Note that setting depth = iter = 1 corresponds to turning off the randomization process and running the standard Clarke-Wright algorithm to find an initial solution. In general, the smaller these parameters are, the quicker the algorithm will run but the poorer the solutions will be.

The next set of options allows you to choose which improvement heuristics to run (see Section 2). The final set of options concerns how Or-opt is performed. If the box labeled "Search all routes during Or-opt [3]" is checked, the algorithm performs a variation on the standard Or-opt (which considers moving a group of 3

cities to another position on the same route) in which the group of cities may be moved to a different route if it's cheaper to do so than to re-insert them on the same route. If the box is unchecked, the standard Or-opt procedure is performed. The next two options are the same, but applying to the 2- and 1-city Or-opt operations. In general, checking more options on this form will result in better solutions but longer run-times.

8. DISCLAIMER

THE AUTHOR OF THIS SOFTWARE MAKES NO CLAIMS, EXPRESSED OR IMPLIED, ABOUT THE PERFORMANCE OF THIS SOFTWARE, INCLUDING, BUT NOT LIMITED TO, THE STABILITY OF THE SOFTWARE OR THE QUALITY OF THE SOLUTIONS IT RETURNS. The author shall not be held liable for any damage or injury that results from the use of this software, including, but not limited to, damage to computer software or hardware and excess costs due to poor solutions implemented.

9. DISTRIBUTION

VRP Solver may be freely reproduced and distributed provided that this ReadMe.txt file is distributed with it, un-altered. The most recent version of VRP Solver may be downloaded from users.iems.northwestern.edu/~lsnyder/software.html.

10. CONTACT INFORMATION

Lawrence V. Snyder
Assistant Professor
Department of Industrial and Systems Engineering
Lehigh University
200 West Packer Avenue, Mohler Lab
Bethlehem, PA 18015
ivs2@lehigh.edu
<http://users.iems.northwestern.edu/~lsnyder>