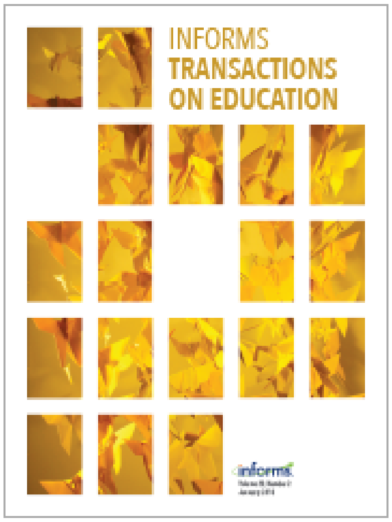

INFORMS Transactions on Education

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Getting Beyond the First Result of Solving a Vehicle Routing Problem

John F. Wellington, Stephen A. Lewis

informs
http://pubsonline.informs.org/journal/ited

INFORMS TRANSACTIONS ON EDUCATION
Vol. 22, No. 1, September 2021, pp. 9–27
ISSN 1532-0545 (online)

Getting Beyond the First Result of Solving a Vehicle Routing Problem

John F. Wellington,[a] **Stephen A. Lewis**[a]

[a] Purdue University Fort Wayne, Campton, New Hampshire 03223

Contact: wellingj@pfw.edu, https://orcid.org/0000-0001-8865-5340 (JFW); mongrelworks.org@gmail.com (SAL)

Received: June 2, 2020
Revised: November 7, 2020; January 28, 2021, March 1, 2021
Accepted: March 2, 2021
Published Online in Articles in Advance: July 6, 2021

https://doi.org/10.1287/ited.2021.0255

Abstract. The simple vehicle routing problem (VRP) is a common topic of discussion in introductory operations research/management science courses. The VRP can be framed in a variety of ways, and it can be difficult to solve to optimality. For solution purposes, introductory textbooks demonstrate how Excel's Evolutionary Solver (ES) add-in produces a routing. The ES utilizes a genetic algorithm with a heuristic stopping rule to produce a routing that is not guaranteed to be optimal. Beyond pointing out that search controls, such as maximum execution time, may be extended and followed by restart(s) of ES, textbook treatments do not offer alternative ways to continue the search for a possibly better routing. In this paper, a suite of ways is presented in which students may investigate beyond what ES produces or any other optimality-uncertain VRP solution method. The suite includes perturbation methods and other ways that function within an Excel spreadsheet environment that is popular with students and textbook writers. Because there is no demonstrable feature that confirms optimality, the student problem Solver must settle for a 'best found' result as unsettling as it may be. The incertitude is addressed.

Supplemental Material: The supplemental material is available at https://doi.org/10.1287/ited.2021.0255.

Keywords: Excel spreadsheet modeling • traveling salesman problem • vehicle routing problem

1. Introduction

The vehicle routing problem (VRP) is a classic operations research (OR) problem whose solution identifies the optimal way to sequence the touring of K number of vehicles among n number of locations for delivery or other purposes. Its solution resolves which site is visited first, the one thereafter, etc., for each vehicle so that fleet travel distance is as small as possible. Generally, the VRP is concerned with routing $K > 1$ vehicles, and the related traveling salesman problem (TSP) addresses a $K = 1$ routing situation. In the classic form of the sequencing, both routing problems require that all sites are visited, each site is visited by just one vehicle, and each vehicle leaves the starting point (depot) and does not return there until all its assigned sites are visited. Additionally, the VRP may require vehicle carrying capacities be observed; customer delivery requirements at each location be met; and vehicle/driver assignments with regard to number of stops and distance traveled be balanced, that is, as much alike as possible. For discussion of balance, see Carter and Ragsdale (2002), Pollaris et al. (2015), and Vidal et al. (2020). See Toth and Vigo (2014) and Vidal et al. (2020) for other considerations and variants of the VRP. The challenge presented by both routing problems is their combinatorial nature. It accounts for some solution methods that do not guarantee optimal routing but are expeditious in producing what is perceived as a good result. With some VRPs, long execution times may move students to halt the procedure prior to termination. Consequently, the problem solving brings students face to face with an optimality-uncertain result. Students may choose to terminate the problem solving with the best found first result or continue the investigation in search of a better routing. A scenario of this kind could arise with a VRP that is the subject of an end-of-chapter textbook problem, a course project, a case study, or a client-sponsored project. For students who would like to continue the investigation, introductory textbook narratives do not offer much beyond pointing out how the procedure can be restarted under different execution controls. This article is concerned with demonstrating how students could proceed differently within an Excel spreadsheet environment. We note at this point that Excel-based solution implementations are popular treatments in introductory textbooks.

From teaching experience, the VRP is a popular topic with students. Its popularity may be due to student perception that it is a common problem faced by business operations, and it is easy to visualize. Students quickly grasp the combinatorial challenge of sequencing the delivery of product to customers and other applications of the VRP. Demonstration of how the number of possible routings grows with n is generally convincing evidence for students that, except for small n, exhaustive enumeration of all routings with as many as $n!$ possibilities is not a practical solution method in most cases; see Ragsdale (2015). Nonetheless, textbooks have end-of-chapter exercises with routing problems of $n = 30$ and 45 sites; see Winston and Albright (2009). To address routing problems of such size, introductory textbooks include demonstrations of how Excel's 'alldifferent' constraint feature and its Evolutionary Solver (ES) add-in produce a solution; see Ragsdale (2018) and Winston and Albright (2009). Although the demonstrations are generally for the $K = 1$ TSP, it is shown how the spreadsheet methodology can be extended to the $K > 1$ VRP. Because of the combinatorial aspects of the VRP, solution times for ES may be long even for modest n. Accordingly, spreadsheet applications for both the TSP and VRP include user-specified controls, such as limiting its execution time. When limits are encountered, the procedure pauses and prompts the student user to continue the search or to terminate. A student may extend the limits more than once to prolong the search and increase the likelihood of finding a better solution.

Alternatively, a suite of ways is presented that is rooted in perturbing the best available routing produced by ES, another Solver, or by other means not guaranteed to terminate with the optimal routing. Although the suite does not guarantee the optimal solution or a routing better than what was available at the start, it does provide the student with a method and guidance in how to proceed when the impetus to find something better so moves the student. It also helps the student to deal with the incertitude/ambiguity of concluding the problem solving with the best of what was found. Spreadsheet implementation of the proposed suite makes use of Excel features well known to students.

The following are aspects of the Excel-based textbook treatments of the routing problem that provide context for the manner in which the proposed suite is operationalized.

1. Screenshots of the Excel worksheets presented in textbooks show the routing solution in a range of contiguous spreadsheet cells whose contents, when read serially, compose the sequence of site visits. The cell contents are location identifiers/indices generally of the form 1, …, n with one identifier per cell. Concatenation of those cell contents using commas for delimiters produces what is referred to as the string representation of a routing solution. For the $K = 1$ and $n = 3$ routing problem, one possibility is the string 0,2,1,3,0 resulting from the concatenation of the contents of five adjacent cells, for example, A1:A5. The string indicates that the sole touring vehicle leaves the depot denoted by 0; visits locations 2,1,3 in that order; and then and only then returns to the depot. Note too that reading the string right to left provides an alternative route. Representing a routing sequence as a string of location identifiers that denotes the order of site visits is adopted in this article. See Toth and Vigo (2014), Winston and Albright (2009), Ragsdale (2015), Hillier et al. (2017), and Taja (2018) for its practice in textbooks.

2. Introductory textbook treatments of the routing problem do not generally include a math programming formulation. The visual appeal of the string form and the ease of its translation to the contents of spreadsheet cells for solution purposes may account for this. Textbook demonstrations show how Excel functions and features familiar to most students are utilized in solving a routing problem. They make use of well-known Excel built-in functions, such as SUMPRODUCT(), INDEX(), and OFFSET(), whose arguments are cell references to the location identifiers of a routing string; see Ragsdale's (2018, chapter 8) treatment for a TSP. The demonstrations point out how Excel's 'alldifferent' and Solver add-in features produce alternative routing solutions. Implementation of the proposed suite also proceeds in a spreadsheet environment. However, it does not make use of a Solver add-in or Excel's 'alldifferent' feature. It perturbs rather than permutes location identifiers. An Excel one-way data table produces perturbations of a given route and, in turn, alternatives for evaluation. Doing so does not depend upon a formal mathematical modeling of the routing problem. However, for the interested reader, a math model is presented in Appendix A.

3. Some Excel Solver add-ins that are used for problem solving demonstrations in introductory OR/management science (MS) textbooks include a genetic algorithm (GA). ES is a popular adaptation for the vehicle routing problem. It makes use of a GA to identify mutations of a routing sequence. Because the GA includes a heuristic, the ES solution is not guaranteed to be optimal. Because the proposed suite does not exhaustively search the VRP solution space, it too does not guarantee an optimal result.

4. The spreadsheet demonstrations in the textbooks do not include a seamless way to proceed from the one vehicle routing problem to multiple vehicles, that is, to the $K > 1$ VRP. However, the Excel workbook 'VRP Example.xlsx' that is included in the supplemental material does.

5. Given the introductory textbook preference for Excel-based solution methods, any proposed extension of

the problem solving to be useful should function within an Excel spreadsheet environment.

The methods of the proposed suite begin with perturbing an available solution and proceed by doing the same to successive routes. One type of perturbation is the positional interchange of a selected location identifier with each of the other location indices in a given routing. Hillier et al. (2017) refer to a pair exchange of this kind as a subtour reversal. To illustrate, suppose the route 1,2,5,7,3,6,4,10,9,8 for a VRP with $K = 2$ and $n = 10$ is the output of ES or another route-producing scheme that does not guarantee the optimal result. As such, it would be a candidate for perturbation investigation. In this routing string, notation of departure from and return to the depot (= 0) is omitted with the understanding that the tour for each vehicle begins and ends there. The underlining serves to distinguish the assignments for the two vehicles. Suppose investigation is directed to positional interchanges of location index 2 with the other indices 1,3,4,5, … ,10. Examples are 2,1,5,7,3,6,4,10,9,8 resulting from the positional interchange of indices 2 and 1; 1,3,5,7,2,6,4,10,9,8 for indices 2 and 3; 1,4,5,7,3,6,2,10,9,8 for indices 2 and 4; and 1,5,2,7,3,6,4,10,9,8 for indices 2 and 5. In this way, alternative routings are made available for evaluation. In the Excel workbook that is included in the supplemental material accompanying this article, the positional interchanges are achieved using a one-way Excel data table. A user-operated Excel spin button triggers population of successive one-way data tables displaying the perturbed routings and their evaluative measures.

The suite offers simple ways for students to search for successively better routings. Each way begins with an optimality-uncertain routing solution that may be the output of a formal modeling of the routing problem and its translation to solution-producing software; Excel's ES or another Solver; a heuristic, such as the nearest neighbor rule (NNR) (see Toth and Vigo 2014); or a scheme for randomly sequencing the delivery stops for each vehicle. It may also begin with an opportunistic routing. The suite is not intended to replace introductory textbook treatments of the routing problem. Rather, because it proceeds from an available solution, it should be looked upon as a complement to those treatments. The suite of perturbations begins where optimality-uncertain solution methods terminate.

Concurrent execution of ES and the suite is a novel way to make use of the possibly long time spent waiting for ES or another Solver to terminate. Either or both procedures could begin with the same routing such as that produced by following the NNR. As a heuristic, the NNR does not guarantee the optimum, but it is a relatively easy route to produce. In this way, upon conclusion of ES, two routings would be available for comparative purposes. If the Solver result is the better of the two, the student could make it the subject of a new round of investigation by the methods of the suite.

Included in the suite is an addendum for use with textbook Solvers. It produces either the best-found feasible routing or the best-found near-feasible routing when the former is not encountered in the search.

The suite is new to the pedagogical VRP literature. It is intended to move and assist students in the search for a possibly better routing solution within the Excel problem-solving environment addressed in their textbooks.

As noted, when the problem-solving software, such as Excel's ES, makes use of a heuristic, the concluding result is a 'good' solution but not necessarily the optimal routing. In fact, if the optimal solution is encountered, it would be unknown as such to the user. There is no optimality signature that identifies the best route. The most that can be said about the terminating route is that it is the 'best found.' Successive runs of ES under different user execution controls may result in no improvement and may lead students to conclude that further investigation is unnecessary or fruitless; that is, nothing better can be found. In this regard, Baker and Camm (2005, p. 5) state,

> Some students seem vulnerable to a psychological trap when they use the evolutionary Solver. A series of runs under different parameter settings that produces the same solution may induce students to conclude prematurely that they have found an optimum. As teachers, we need to develop better guidelines for implementing the evolutionary Solver so that the chances of reaching an optimum improve.

FrontlineSystems (2021, p. 1), the developer of Excel's ES, cautions, "When you use the Evolutionary method, you may find—like other users of genetic and evolutionary algorithms—that you spend a lot of time running and re-running Solver, trying to find better solutions. This is an inescapable consequence of using a method that makes few or no assumptions about the nature of the problem. You can never be sure whether you've found the best solution, or what the payoff might be of running the evolutionary algorithm for a longer time." After the experience of learning to formally frame allocation models that lend themselves to solution techniques producing an optimal result, these aspects of the concluding VRP solution are new and unsettling to some students.

The suite of demonstrations in this article are ways to show students how to deal with the optimality uncertainty of a Solver result or the outcome obtained by other means that do not guarantee the optimum. The demonstrations show how to explore beyond what textbook software automatically produces. By design, the demonstrations in this article address a routing problem with practical requirements. It is shown that they are replicable in the same Excel environment as ES.

The hands-on approach of the proposed problem solving makes the student an active participant in searching for what could be better. It provides the student with a way to deal with the incertitude of the concluding result particularly if it arises from an automatic software application. If the outcome of the suite is a better solution, the suite has served its purpose. The suite enables students to discover for themselves a good solution. The student leaves the problem-solving arena with knowledge that reasonable effort was expended in the search. In this way, the student moves beyond nominal problem solving.

2. Relevant Literature

The vehicle routing literature is extensive and spans many disciplines in which the framework of the VRP or the TSP applies. A review of solution methods is found in Toth and Vigo (2014) and Laporte (2009) for the VRP and in Gutin and Punnen (2002) for the TSP. Braekers et al. (2016), Pillac et al. (2013), and Eksioglu et al. (2009) provide comprehensive taxonomies of the VRP literature. A variety of solution methodologies appears in this literature. They include exact and approximate solution methods. Exact methods relate to a formal optimization model of the routing problem; require translation to solution-producing software, such as MATLAB, Visual Basic for Applications (VBA), C, GAMS, AMPL, CPLEX or another modeling language; and produce the optimal route. An approximation method, as the name implies, produces a result with some performance expectation expressed in terms of the solution's proximity to the unknown optimal fleet distance. Otherwise, heuristic-based methods are available. They produce a good solution but provide no guarantee of its optimality or proximity thereto. They typically guide the sequencing of delivery locations one at a time in an emerging itinerary. Some heuristic methods are relatively straightforward and easy to teach and implement. See Toth and Vigo (2014) for detailed discussions of these topics. Ragsdale (2015) and Winston and Albright (2009) acknowledge the use of heuristics in the Excel-based Solvers addressed in their textbooks and the optimality uncertainty of the results.

The pedagogical literature of relevance for this article includes popular OR/MS introductory textbooks, such as Anderson et al. (2019), Hillier et al. (2017), Hillier and Hillier (2019), Ragsdale (2015, 2018), Winston and Albright (2009, 2012), and Taja (2018). In these sources, the introductory routing problem is usually a TSP. As previously noted, solution treatments include use of evolutionary methods implemented within an Excel spreadsheet environment using a proprietary Solver. Excel is a popular application for demonstrating how to solve optimization problems of the kind addressed in introductory OR/MS textbooks. In regard to this trend, consider the following remark from Baker and Camm (2005, p. 1): "Instead of requiring students to use unfamiliar stand-alone software that most would no longer use after graduation, the spreadsheet allows OR to be taught in an environment students have already accepted and will continue to use." Ragsdale (2001) and Hesse and Scerno (2009) also advance the use of spreadsheets in teaching students how to solve OR/MS problems. Erdogan (2017, p. 62) makes the case for Excel-based solution tools as "the standard software for small to medium scale quantitative analysis for businesses." As such, demonstration of how the proposed suite functions within Excel is consistent with the problem-solving pedagogy adopted in contemporary introductory textbooks. The practice may also be consistent with Erdogan's (2017) remark concerning the software environment some students may encounter in their early career placements.

It should be noted that Excel's ES is not the only spreadsheet Solver. In their solution treatments of OR problems, Hillier and Hillier (2019) use Excel's Analytic Solver. Open-source Excel-based Solver tools include Erdogan's (2017) VRP spreadsheet Solver. It uses VBA code that the author states can be understood and modified by medium-level programmers Erdogan (2017). Mason (2013, p. 45) provides SolverStudio, an "Excel add-in that combines the power of modelling languages including AMPL, CPLEX, GAMS, PuLP, GMPL, and Gurobi's Python environment with the familiarity and ease of use of Excel" that can be applied to the VRP. To the extent that Solvers such as these terminate with optimality-uncertain solutions, the proposed suite of treatments provides students with means for exploring beyond their results.

Case studies have become a popular feature in introductory textbooks. They allow narratives to go beyond simple well-framed routing problems of first discussion to unstructured routing situations with larger n and K. Hillier et al. (2017), Hillier and Hillier (2019), and Winston and Albright (2009), among others, include cases. Mini cases are also common in contemporary textbooks; see Hillier and Hillier (2019) for examples. More complicated cases dealing with the delivery of mobile healthcare services or the scheduling of tourist visits to historic and cultural sites appear in the VRP literature. See Erdogan (2017) for treatment of these routing situations as well as Milburn et al. (2018), Drake et al. (2011), and Koksalan and Salman (2003) for others. At a minimum, cases provide students with a view to the breadth of applications of VRP modeling and to practical routing considerations and how they are accommodated in problem solving. As noted by Milburn et al. (2017, p. 75), The primary objective of the case is to "provide students with hands-on experience developing and applying solution techniques for an unstructured vehicle routing problem." The 'hands-on' remark is noteworthy. It is central to the theme of this article that active student involvement in VRP solution discovery is good pedagogy whether it is experienced in a case study, an end-of-chapter exercise, a course project, or a client-sponsored project.

3. Method: The Proposed Suite

The proposed suite consists of ways to search for successively better routings by perturbing an available route and its successors in an orderly manner to generate alternative routings for evaluation. The symmetric VRP is the common form of the problem treated in introductory textbooks and is the form of the problem addressed in this article. Consequently, the distance between any two locations i and j denoted as d_{ij} is identical to d_{ji}, i, j = 0,1 … ,n, $i \neq j$. A schematic of the suite's Excel implementation appears in Table B.1 of the appendix.

3.1. Using ES to Produce a Routing

It is assumed that the reader has familiarity with using Excel's ES or another Solver that can produce a VRP solution to serve as input to the methods of the proposed suite. Ragsdale (2018) provides a good introduction to using Excel's ES to produce a solution for a simple VRP.

3.2. Using Perturbation Methods to Search for a Better Routing

In the search for better routings, consider how the indices in a routing of interest may be perturbed and, in turn, produce different routing sequences for evaluation. The perturbations discussed in this article are referred to as (i) type 1, in which each location index i (= 1, … , n) one at a time is repositioned as the jth (1st, 2nd, … , nth, left orientation) location index in a routing sequence, and (ii) type 2 in which index i is positionally interchanged with each of the other indices j (j = 1, … , n, $j \neq i$) one at a time. Table 1 displays a few perturbations of type 1 in which index 1 is repositioned in the 1st, … , 4th positions of the routing of row 0. Note the staircase appearance of the bolded index 1.

Table 2 shows a few type 2 perturbations in which index 1 is first positionally interchanged with index 2 and then with indices 3, 4, and 5 one at a time in that order. Note how the bolded original position of index 1 becomes repopulated among the other indices.

If the better of the best of the exhaustive types 1 and 2 perturbations is superior to the previous best found result, it becomes the new best found routing, and the investigation proceeds in like manner based on the new result. Perturbations of types 1 and 2 are operationalized using Excel one-way data tables.

Table 1. Illustration of a Few Type 1 Perturbations

Row	Perturbed routing
0	0,5,16,8,4,18,12,6,10,13,9,17,19,2,11,14,15,7,**1**,3,0
1	0,**1**,5,16,8,4,18,12,6,10,13,9,17,19,2,11,14,15,7,3,0
2	0,5,**1**,16,8,4,18,12,6,10,13,9,17,19,2,11,14,15,7,3,0
3	0,5,16,**1**,8,4,18,12,6,10,13,9,17,19,2,11,14,15,7,3,0
4	0,5,16,8,**1**,4,18,12,6,10,13,9,17,19,2,11,14,15,7,3,0

3.3. Using Penalty Functions with ES

As previously noted, there are routing situations that have specific requirements to be observed. Such a routing situation is referred to as a constrained vehicle routing problem and denoted by CVRP.

The requirements of a CVRP impact the efficiency of methods, such as Excel's ES, that utilize a GA. Initial routing assignments for GAs are assigned randomly. Accordingly, they may violate the stipulated restrictions of the CVRP. The GA may continue to 'guess' subsequent assignments (feasible or not) as a routing 'evolves.' Thus, GAs, especially general purpose GAs, are less efficient under restricted optimization, such as the CVRP.

Although some of the issues with GAs are well known among GA researchers, they may not be known to those teaching introductory OR/MS/operations management courses. For example, even with CVRPs of modest n and c (number of constraints), students may experience long execution times using software with an imbedded GA. And Excel's ES may terminate without a feasible result. Alternatively, consider the following. Suppose conformance with and violation of routing restrictions were accommodated within the objective function in the form of penalty terms, see Winston and Albright (2009). Conformance would have zero penalty and violation a large positive value. A routing, feasible or not, would result. If the terminating route is infeasible, it would have the smallest magnitude of constraint violation among those found. It would be known by the magnitude of the concluding fleet distance value. For demonstration purposes, suppose vehicle carrying capacities are limited to 700 per vehicle for a VRP of interest with K = 4 touring vehicles. Further, consider a routing with vehicle loads 698, 670, 693, and 701 calculated and displayed in Excel cells J6:J9, respectively, versus the routing with loads 648, 650, 663, and 801. Between the two, the first routing would be preferred with one constraint violation of magnitude one. However, because both are infeasible solutions, the preferred solution may not be returned by ES and, hence, not known to the problem Solver. However, if '=10000*COUNTIF(J6:J9,">700")' is included as a term in the objective function addressed by ES, and assuming no feasible routing is found, ES would return a solution with some indication of constraint violation(s) of magnitude 10,000 or more. The penalty term produces a count of the number of constraint violations weighted by 10,000. Note that the penalty term would be the same for both solutions. The preferred solution with less constraint violation would not necessarily emerge. Alternatively, consider the penalty function of the Excel form '=10000*MAX(MAX(J6:J9)-700,0),' where the magnitude

Table 2. Illustration of a Few Type 2 Perturbations

Row	Indices interchanged	Perturbed routing
0	—	0,5,16,8,4,18,12,6,10,13,9,17,19,2,11,14,15,7,**1**,3,0
1	1 and 2	0,5,16,8,4,18,12,6,10,13,9,17,19,**1**,11,14,15,7,**2**,3,0
2	1 and 3	0,5,16,8,4,18,12,6,10,13,9,17,19,2,11,14,15,7,**3**,**1**,0
3	1 and 4	0,5,16,8,**1**,18,12,6,10,13,9,17,19,2,11,14,15,7,**4**,3,0
4	1 and 5	0,**1**,16,8,4,18,12,6,10,13,9,17,19,2,11,14,15,7,**5**,3,0

of the worst constraint violation is weighted heaviest. Given objective function minimization, the routing with the lesser violation would emerge. The solution with vehicle loads 698 (J6), 670 (J7), 693 (J8), and 701 (J9) would have a constraint violation of magnitude one, whereas the routing with 648 (J6), 650 (J7), 663 (J8), and 801 (J9) would have a violation of magnitude 101. If a feasible route is not encountered, the penalty term(s) would likely assist the Solver in returning the best found infeasible route versus termination with the message that no feasible solution was found. See formal treatment of the penalty function approach in Appendix A.

3.4. The Step-by-Step Procedure

The student may proceed in the following manner using the methods of the suite:

Step 0. Find an initial routing. Set it to the best found. Go to Step 1.

Step 1. Perturb indices of the current best found routing using the type 1 method. Record the best *m* routing solutions by fleet distance and conformance to routing requirements. Each routing so treated should have a fleet distance value smaller than the value for the best found route designated as such at this point in the investigation. Go to Step 2.

Step 2. Perturb indices of the current best found routing using the type 2 method. Record the best *m* routing solutions by fleet distance and conformance to routing requirements. Each routing so treated should have a fleet distance value smaller than the value for the best found route designated as such at this point in the investigation. Go to Step 3.

Step 3. Collectively rank the 2*m* number of results identified in Steps 1 and 2 in ascending order by fleet distance. They constitute the list of perturbation candidates to be tested for permanent implementation. Then, in the order of the list, evaluate for fleet distance and conformance to routing restrictions the implementation of each perturbation candidate one at a time. Do this by editing accordingly the current best found routing. When an implementation produces a feasible routing with a distance measure better than the currently designated best found, leave the implementation permanent and denote the result as the new best found route; otherwise, undo the editing. Go on to the next candidate. When all candidates of the current list have been tested and implementations when warranted made permanent, go to Step 4.

Step 4. Continue? If yes, go to Step 1; otherwise, Stop.

In the manner of Steps 1–4, the procedure produces successively better routes or concludes with nothing better than the starting route of Step 0. Step 3 identifies serial perturbations of both types that would otherwise have been unknown.

The following are notes that relate to the procedure. First, no firm rule is offered for assigning a value to *m*. A value of five or six may be sufficient for discovery purposes. The treatment of *m* should be flexible enough to allow inclusion of more (or fewer) than *m* attractive candidates in Steps 1 and 2 as discovery proceeds. Limiting the number to a preset *m* is intended to expedite Steps 1–3. Second, the outcomes of Steps 1 and 2 are feasible routings. In addition, they would not be noted as best unless their fleet distance values were better than the value for the best found routing in effect at the onset of Step 1. When evaluated in Steps 1 and 2, each perturbation is treated as if it were the only change to the best found routing in effect at the time of examination. However, in Step 3, because of successive changes that preceded the testing of a perturbation candidate, its implementation as described may no longer be possible, or its testing may produce an infeasible or distance-inferior result. Third, depending on the best result at the conclusion of Step 3, the student may decide to repeat Steps 1–3, that is, perform another iteration of the same. If a significant improvement in fleet distance resulted, possibly an even better routing may result from another iteration. When no candidate or one with marginal improvement in fleet distance is the concluding outcome of Step 3, the student should think about stopping the investigation. These are the considerations to be addressed in Step 4. Fourth, the outcomes of Steps 1 and 2 need to be manually recorded to perform the ranking required at the beginning of Step 3. For a candidate identified in Step 1, recording the index to be repositioned, the locus (between what two indices) of its relocation, and the corresponding fleet distance value suffice. For a Step 2 candidate, recording the two indices to be positionally interchanged and the corresponding fleet distance value suffice. It is not necessary to record the resulting routing string with each instance of Steps 1 and 2. As a final note, in Step

3, once the repositioning or interchange of a location index is made permanent, any remaining perturbation candidates involving the location index should be struck from the current candidate list. This should be practiced until a new iteration is initiated, that is, when the procedure returns to Step 1. The illustration of Section 4 and the narrative of Section 5 make these points clear.

Discussion of how Steps 1–3 are operationalized in the proffered Excel workbook 'VRP Example.xlsx' appears in Appendix B.

4. Illustration

The suite of ways in which a solution to a VRP of interest may be investigated for improvement is demonstrated in this section.

4.1. The VRP of Illustration

The VRP selected for demonstration was taken from instance set B of Augerat et al. (1995) and is referred to by the label B-n50-k7. In the notation of this paper, the routing problem is a CVRP with $n = 49$ delivery sites, one depot, and $K = 7$ touring vehicles. The published requirements of the problem are displayed in Table 3. Although balanced vehicle assignments were not a requirement for the B-n50-k7, they were for this illustration.

The VRP B-n50-k7 was selected for the following reasons. Its dimensionality ($K = 7$, $n = 49$) allows a rich treatment of ways to seek a better solution. Additionally, it includes practical routing considerations. In regard to the latter, one can reasonably argue that delivery resources are limited, and restrictions, such as vehicle carrying capacities, customer delivery requirements, and balanced vehicle/driver delivery assignments, are practical considerations and should be accommodated in the problem solving. Balance is reflective of good resource stewardship. It was framed for demonstration purposes as the complete use of the delivery fleet and a uniform number of delivery stops for each vehicle. In addition, a bound to the extreme in assigned driver/vehicle travel distances was included. It was framed as the allowed maximal difference in travel distance between any two vehicles (MDTD2V) and operationalized as the difference between the maximum and minimum of assigned driving distances. Restricting MDTD2V was one way to flatten the variation among assigned vehicle driving distances and, in turn, to effect better utilization of driving personnel and vehicle usage.

In summary, the routing problem of illustration is a CVRP to identify the smallest fleet distance routing for $K = 7$ vehicles touring $n = 49$ sites subject to six specific considerations. They include observing uniform vehicle carrying capacities of 100 units by number, volume, or weight; meeting customer delivery requirements as specified in published B-n50-k7 and restated in Table 3; making full use of the available fleet of K (= 7) vehicles; prohibiting subtours in which a vehicle would return to the depot before visiting all locations uniquely assigned to it; and enforcing balanced assignments reflected in the MDTD2V of 100 and uniform ($n/K = 49/7$) number of delivery stops per vehicle. Addressing this scenario was challenging. The number ($n = 49$) of sites to be visited is large. For ES and its GA, the solution time was slowed by the need to satisfy the declared restrictions. However, the VRP as described has practical considerations and serves as a good introduction to VRPs the student may encounter in a course assignment, a course project, or a client-sponsored project. It is a manageable problem for students.

4.2. The Demonstrations

The reported best found routing for B-n50-k7 with a fleet distance of 741 is displayed in row 1 of Table 4. It requires seven vehicles, between 4 and 11 delivery stops per vehicle, and MDTD2V of 66. In pursuit of balanced assignments, the routing sequence as given was edited to show seven ($n/K = 49/7$) delivery stops for each of the seven vehicles. The first seven locations of the reported B-n50-k7 routing were assigned to vehicle 1, the next seven to vehicle 2, etc., to produce the routing of row 2, Table 4, with fleet distance 1,050. Note that the vehicle capacity and the MDTD2V restrictions were exceeded. The scenarios of rows 1 and 2 provide some expectation of what could possibly be found from further investigation. The best found fleet distance for the fully constrained CVRP must be at least 741 and perhaps no larger than the 1,050 value of row 2. Under the requirements of Table 3 and balanced delivery assignments, Excel's ES produced the routing of row 3 with fleet distance 2,390. The ES framing of the problem includes two terms in the objective function to penalize violations of vehicle load capacities and balanced delivery assignments. In executing ES, the number of location

Table 3. Delivery Requirements for the Illustration

Location	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
Delivery requirement	21	8	11	7	21	5	13	10	9	20	7	12	23	2	4	14	12	3
Location	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36
Delivery requirement	5	13	5	12	2	3	18	24	4	63	19	2	9	4	9	23	6	3
Location	37	38	39	40	41	42	43	44	45	46	47	48	49					
Delivery requirement	12	7	17	22	26	14	9	2	16	24	4	19	11					

Note. Load capacity per vehicle is 100.

Table 4. Perturbation Results for VRP B-n50-k7 Beginning with an ES Routing

Row	Action	Routing result[a] (number of vehicle delivery stops)[b]	Fleet distance (vehicle distances) MDTD2V[c]	Vehicle loads
1	Reported best found routing under the requirements of Table 3 only for CVRP B-n50-k7	47,31,41,7,37,11,20,8,5,36,45,21,35,4,29,46, 9,38,22,33,34,49,44,19,17,48,2,6,32,28,40,15, 30,25,18,14,43,24,27,16,42,10,3,23,1,13,26, 39,12,(4,8,8,8,4,11,6)	741, 66 (85,151,102,104, 85,99,115)	52,87,93,85, 93,100,99
2	Edited routing of row 1 to reflect the balanced delivery requirement of uniform number (= 7) of delivery stops per vehicle	47,31,41,7,37,11,20,8,5,36,45,21,35,4, 29,46,9,38,22,33,34,49,44,19,17,48,2,6, 32,28,40,15,30,25,18,14,43,24,27,16,42,10, 3,23,1,13,26,39,12,	1,050, 114 (208,171,152, 104, 152,94,169)	84,68,103,62, 116,66,110
3	Routing produced by ES with objective function terms penalizing vehicle loads in excess of 100 and MDTD2V exceeding 100	31,30,17,40,34,37,5,2,4,36,13,9,22,29, 6,45,3,26,24,43,41,8,38,15,49,46,21,20, 39,42,25,19,1,12,47,7,28,44,23,33,27,11, 10,32,14,18,16,48,35,	2,390, 97 (289,298,347, 386,332,373,365)	101,81,94,74, 91,100,68
4	The best found feasible routing produced through perturbation analyses that began with routing of row 3	40,44,19,17,48,49,34,6,21,1,13,26,12,33, 35,4,29,9,38,22,7,37,11,45,36,5,8,20, 3,42,10,16,27,41,47,15,2,32,28,14,24,43, 30,25,18,31,39,46,23,	1,007, 90 (102,191,112, 174,101,170,157)	94,99,73,82, 93,93,75
5	The best found near-feasible routing produced through perturbation analyses that began with routing of row 3	40,44,19,17,48,49,34,6,21,1,13,26,12,33, 4,35,29,38,9,22,31,37,11,45,36,5,8,20, 3,42,10,16,27,41,47,15,2,32,28,14,24,43, 30,25,18,7,39,46,23,	978, 109 (102,191,82,174, 101,170,158)	94,99,69,82, 93,93,79

[a]Alternating underscores distinguish location assignments for vehicles. In doing so, notation (= 0) of departure/return to the depot for each vehicle is omitted.
[b]Unless noted otherwise, the number of site visits per vehicle is seven.
[c]Maximal difference in travel distance between any two vehicles (MDTD2V).

stops per vehicle was kept at seven to conform to balanced delivery requirements. The ES result had one instance of constraint violation; see row 3 of Table 4, where the assigned load to vehicle 1 was 101, the minimal violation. At termination, Excel's ES reported that it could not find a better solution. Consequently, it was the best found infeasible routing at this point in the investigation. The search then turned to examining what the two proposed perturbation methods could produce in improvement to the routing of row 3. It concluded with the best found feasible routing displayed in row 4 and the best found near-feasible routing in row 5 with respective fleet distances of 1,007 and 978. The infeasibility of the latter was the MDTD2V of 109. Because no further improvements were found using the perturbation methods, the investigation that began with the ES solution of row 3 concluded with the results of rows 4 and 5. The perturbation specifics of row 4 are detailed in the supplemental material.

The suite of investigations can begin with any solution. Table 5 shows the results of the investigation that began with the routing produced under the NNR. The routing resulted from sequencing the successively closest locations beginning with departure from the depot and ending at the same. Visual inspections of the table of pair distances between all locations sufficed in composing the sequencing. Row 1 displays the NNR routing with fleet distance of 1,009. The routing violates the MDTD2V restriction (= 100) and uniform number (= 7) of delivery stops per vehicle. Row 2 is the outcome of editing the NNR routing sequence to show seven delivery stops for each of seven vehicles. The first seven locations of the routing of row 1 were assigned to vehicle 1, the next 7 to vehicle 2, etc., to produce the routing with fleet distance 934 and three instances of constraint violation. The perturbation investigation began with the routing of row 2. Each row 3–11 of Table 5 displays the routing at the end of Step 3 of the four-step procedure introduced in Section 3.4 and the perturbations that produced it. Although the routing of row 3 presented one instance (= 139) of constraint violation, it was entertained with a 'wait and see' approach to resolving the violation in subsequent perturbations. There was no guarantee that this approach would yield feasible solutions. However, it did occur with the routings of rows 4–11. The investigation concluded with the feasible routing reported in row 11 with the fleet distance value of 843. The latter is a reduction of 91 (= 934 − 843) in fleet distance or 9.74% improvement. The number of implemented perturbations reported under the column label 'Action' in rows 3–11 varies between one and six. Each row 3–11 relates to one iteration

Table 5. Perturbation Results for VRP B-n50-k7 Beginning with an NNR Routing

Row	Action	Routing result[a] (number of vehicle delivery stops)[b,c]	Fleet distance (vehicle distances) MDTD2V[d]	Vehicle loads
1	The routing produced under the NNR restricted by the requirements of Table 3 only	47,33,22,12,38,9,23,29,35,4,46,13,26,39,1,31,41,7,30, 25,18,14,43,24,27,16,10,42,3,34,49,40,15,6,28,32, 2,48,17,44,19,36,45,21,5,20,11,37,8, (10,4,8,8,4,4,10,1)	1,009, 147 (77,110,172, 140,68,97,215,130)	87,88,94,98,42,94,96,10
2	Edited routing of row 1 to reflect the balanced delivery requirement of uniform number of delivery stops per vehicle	47,33,22,12,38,9,23,29,35,4,46,13,26,39,1,31,41,7, 30,25,18,14,43,24,27,16,10,42,3,34,49,40,15,6, 28,32,2,48,17,44,19,36,45,21,5,20,11,37,8,	934, 110 (65,120,164, 94,162,175,154)	55,120,92,66,139,53,84
3	In the routing of row 2, repositioned index 46 between indices 9 and 23	47,33,22,12,38,9,**46**,23,29,35,4,13,26,39, 1,31,41,7,30,25,18, 14,43,24,27,16,10,42,3,34,49,40,15,6,28, 32,2,48,17,44,19,36,45,21,5,20,11,37,8,	967, 81 (98,120,164, 94,162,175,154)	77,98,92,66,139,53,84
4	In routing of row 3, interchanged positions of indices 1 and 9 then repositioned index 48 between indices 49 and 40	47,33,22,12,38,**1**,46,23,29,35,4,13,26,39, **9**,31,41,7,30,25,18,14,43,24,27,16,10,42,3,34, 49,**48**,40,15,6,28,32,2,17,44,19,36,45,21,5,20,11,37,8,	936, 81 (112,120,117, 94,164,175,154)	89,98,80,66,95,97,84
5	In routing of row 4, repositioned index 9 between indices 6 and 28	47,33,22,12,38,1,46,23,29,35,4,13,26,39,31, 41,7,30,25,18,14,43,24,27,16,10,42,3,34,49, 48,40,15,6,**9**,28,32,2,17,44,19,36,45,21,5,20,11,37,8,	916, 82 (112,120,111, 93,151,175,154)	89,98,73,75,93,97,84
6	In the routing of row 5, repositioned index 8 between indices 5 and 20 then index 2 between indices 19 and 36 and then the following exchanges between index pairs: 9,47; 38,46; 14,25; 19,44	**9**,33,22,12,**46**,1,**38**,23,29,35,4,13,26,39, 31,41,7,30,**14**,18,**25**,43,24,27,16,10,42,3, 34,49,48,40,15,6,**47**,28,32,17,**19**,**44**,**2**,36, 45,21,5,**8**,20,11,37,	897, 77 (118,120,109, 93,139,170,148)	94,98,73,75,88,97,84
7	In the routing of row 6, repositioned index 9 between indices 1 and 38; then made the following exchanges between index pairs: 6,36 then 17,44	33,22,12,46,1,**9**,38,23,29,35,4,13,26,39, 31,41,7,30,14,18,25,43,24,27,16,10,42,3,34, 49,48,40,15,**36**,47,28,32,**44**,19,**17**,2,**6**,45,21, 5,8,20,11,37,	877, 100 (108,120,109, 93,193,106,148)	94,98,73,75,86,99,84
8	In the routing of row 7, made the following exchanges between index pairs: 47,45 then 48,15	33,22,12,46,1,9,38,23,29,35,4,13,26,39, 31,41,7,30,14,18,25,43,24,27,16,10,42,3, 34,49,**15**,40,**48**,36,**45**,28,32,44,19,17,2,6, **47**,21,5,8,20,11,37,	866, 74 (108,120,109, 93,167,106,163)	94,98,73,75,98,99,72
9	In the routing of row 8, repositioned index 23 between indices 39 and 31; repositioned index 47 to the rightmost position; then repositioned index 1 between indices 4 and 13	33,22,12,46,9,38,29,35,4,**1**,13,26,39,**23**, 31,41,7,30,14,18,25,43,24,27,16,10,42,3, 34,49,15,40,48,36,45,28,32,44,19,17,2,6, 21,5,8,20,11,37,**47**,	859, 74 (107,119,109, 93,167,106,158)	92,100,73,75,98,99,72
10	In the routing of row 9, interchanged positions of indices 47 and 29	33,22,12,46,9,38,**47**,35,4,1,13,26,39,23, 31,41,7,30,14,18,25,43,24,27,16,10,42,3, 34,49,15,40,48,36,45,28,32,44,19,17,2,6, 21,5,8,20,11,37,**29**,	844, 74 (96,119,109, 93,167,106,154)	77,100,73,75,98,99,87
11	In the routing of row 10, repositioned index 47 to leftmost position	47,33,22,12,46,9,38,35,4,1,13,26,39,23, 31,41,7,30,14,18,25,43,24,27,16,10,42,3, 34,49,15,40,48,36,45,28,32,44,19,17,2,6, 21,5,8,20,11,37,29,	843, 74 (95,119,109, 93,167,106,154)	77,100,73,75,98,99,87

[a]Alternating underscores distinguish location assignments for vehicles. In doing so, notation (= 0) of departure/return to the depot for each vehicle is omitted.
[b]Unless noted otherwise, the number of site visits per vehicle is seven.
[c]Bolding denotes indices affected by the cited perturbations.
[d]Maximal difference in travel distance between any two vehicles (MDTD2V).

of the four-step procedure. For illustration purposes, the details related to the Action items reported in rows 6–8 of Table 5 are presented in Appendix C. After implementation of the Action item of row 11, the investigation stopped with the best found routing noted there and its distance value of 843. No type 1 or 2 perturbation thereafter offered improvement.

Students will ask about producing and evaluating randomly generated routings. They feel doing so provides opportunity to escape the possibility of being stalled at a local optima in the VRP solution space. One way to proceed is to sectionalize the best found routing and to randomize the order of the indices in a given section. A motivation to do so is the manner in which the indices are assigned one at a time, that is, the early, middle, and latter assignments. The early assignments to an emerging route strongly affect the outcome, that is, fleet distance. In this regard, consider the manner of assignments produced under the NNR. Early assignments, by design, connect short-distance links leaving longer distance links for later consideration. Randomization of the elements of a substring of early assignments may counter the myopia. Just as the methods of perturbation provide the means for investigating possible improvements to the routings produced by ES and the NNR, it may do the same for the best found routing produced by randomization.

4.3. Discussion

The demonstrations of Tables 4 and 5 show how the student who chooses to investigate beyond an automatically produced optimality-uncertain VRP solution may proceed. They also point to the utility of applying more than one method to solving a VRP of interest. As noted, the routing under the NNR is comparatively easy to produce as an initial solution. As suggested, while a Solver is in execution, perturbation investigations may proceed starting with the NNR result. If the Solver result is disappointing, the perturbation result may provide an alternative as it did in the illustration of Section 4.2. Doing so is a different approach to problem solving. Of course, it requires more problem-solving time and student involvement. However, the outcome, as it was for the problem of illustration, may be worthwhile. Although discovery of a better routing beyond the starting result is not guaranteed, the opportunity to explore in the demonstrated ways should not be overlooked.

5. Teaching Suggestions and Experiences

The following discussion is intended to facilitate operationalizing the four-step procedure of Section 3.4 for instructional purposes. To assist, the following are available:

1. The Excel workbook 'VRP Example.xlsx' that is set up and ready to use for the VRP of Section 4. Worksheet 'VRP Model' of the workbook is the locus of operation for each step of the procedure of Section 3.4. The workbook is available among the supplemental material accompanying this article. The setup includes the starting route of row 5, Table 5, that was entered in cells BR1, CC1, and CR1. Their purpose is explained.

2. Table B.1 of Appendix B is a schematic that points out the specifics (locus of operation, cell locations of required input data, an assisting navigation tool, and output information) for each step of the procedure.

3. Appendix C is a narrative of a walkthrough of two complete iterations of the VRP example of Section 4 as displayed in rows 6 - 8 of Table 5.

A suggested layout for recording information generated in performing Steps 1–3 of the procedure appears

Table 6. Suggested Layout for Recording Information

Column 1			Column 2			Column 3			
Ranked best five outcomes of type 1 perturbations obtained from Step 1			Ranked best five outcomes of type 2 perturbations obtained from Step 2			Candidate list—Step 3 Ranked collective outcomes of Steps 1 and 2			
Item	Perturbation[a]	Fleet distance value	Item	Perturbation[b]	Fleet distance value	Order	Item	Perturbation	Fleet distance value upon implementation[c]
1	5,**8**,20	910	6	8,11	913	1	1	5,**8**,20	910
2	19,**2**,36	912	7	9,47	913	2	2	19,**2**,36	?
3	17,**19**,44	914	8	11,8	913	3	6	8,11	?
4	14,**25**,43	914	9	20,8	913	4	7	9,47	?
5	46,**38**,23	914	10	38,46	913	5	8	11,8	?
						6	9	20,8	?
						7	10	38,46	?
						8	3	17,**19**,44	?
						9	4	14,**25**,43	?
						10	5	46,**38**,23	?

[a]Repositioned the bolded location index between the two indicated indices.
[b]Interchanged positions of the indicated index pairs.
[c]? denotes that fleet distance value is not available until the perturbation is tested.

as Table 6. The layout shows, by example, what and how to record the information produced by stepping through the procedure. The candidate list of column 3 of Table 6 determines the order in which the items/perturbations of columns 1 and 2 are successively tested for implementation consideration. Note the manner in which the list is composed. To further assist, the following narrative is presented. It includes partial screen shots of worksheet 'VRP Model' that relate to stepping through the procedure. In the discussion, cell references relate to worksheet 'VRP Model,' blue-filled denotes required user input, and the starting route taken from row 5, Table 5, has the distance value of 916. In the screen shots, note the cell locations of the operations for each step.

Figure 1 is the screenshot of the area of worksheet 'VRP Model' in which type 1 (Step 1) perturbations (repositioning a given location index everywhere in the routing string) are investigated. To proceed, enter the routing of interest in blue-filled cell BR1, and in blue-filled cell BP3 the first location index (1 is suggested) to be investigated by repositioning. Each possible repositioning of the location index and the corresponding fleet distance value appears in cells BT6 and below and cells BU6 and below, respectively, partially shown in Figure 1. Information regarding the best repositioning of the index in cell BP3 and the corresponding fleet distance value appears, respectively, in cells BQ5 and BQ4. 'Best' refers to the repositioning among all feasible possibilities with the smallest fleet distance value. The up arrow of the spin button in the vicinity of cell BQ3 increments the entry of cell BP3 by 1 and the down arrow decrements it by 1. Any valid location index (1, … , 49) may be entered manually in cell BP3. During the course of the investigation, when changing the location index in cell BP3, the student should record the perturbation and the associated distance value for just the successively better routes as encountered. After examining the repositioning of the last location index, the best five or so outcomes among the recorded results were ranked and labeled in the manner of the three subheadings of column 1 of Table 6. This completes Step 1 of the four-step procedure. Before using this section of the worksheet, make sure the entries (desired number of location stops for each vehicle) in blue-filled cells CJ6 and below relate to the routing of cell BR1.

Figure 2 is the screenshot of the area of worksheet 'VRP Model' in which type 2 (Step 2) perturbations (positional interchanges of a given location index with all others) are investigated. To proceed, enter the routing of interest in blue-filled cell CR1 and in blue-filled cell CP3 the first location index (one is suggested) to be positionally interchanged with each of the other location indices. Any valid location index (1, … , 49) may be entered manually in cell CP3. Cells CR6 and below and CS6 and below display the perturbed routing possibilities and the corresponding fleet distance value for each, partially shown in Figure 2. Information regarding the best index to interchange with the entry of cell CP3 appears in cell CP4 and the corresponding fleet distance value in cell CQ4. 'Best' refers to the positional interchange among all feasible possibilities that results in the smallest fleet distance value. The spin button near cell CQ3 assists in incrementing/decrementing the location index in cell CP3. When changing the location index in cell CP3, the student should record the perturbation and the associated fleet distance value for only the successively better routings so encountered. After examining the last interchange, the best five or so recorded outcomes are ranked and labeled in the manner of the three subheadings of column 2, Table 6. This completes Step 2 of the four-step procedure. Before using this section of the worksheet, make sure the entries (desired number of

Figure 1. Screenshot of the Area of Worksheet 'VRP Model' for Investigating Type 1 Perturbations

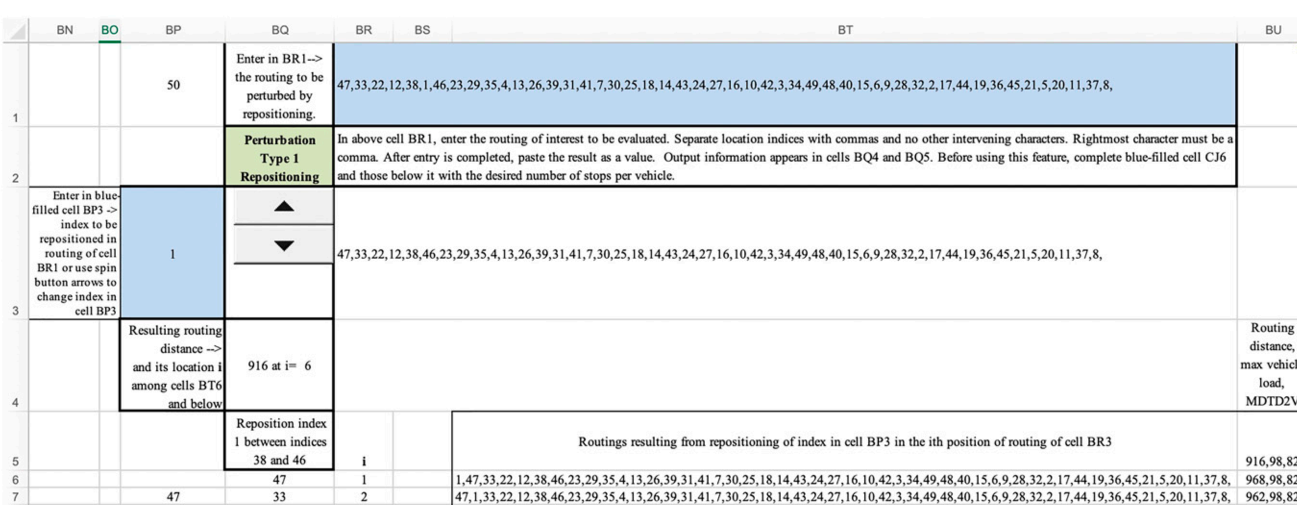

Figure 2. Screenshot of the Area of Worksheet 'VRP Model' for Investigating Type 2 Perturbations

	CO	CP	CQ	CR	CS
1		Enter in cell CR1 --> the routing to be perturbed by interchanging.		47,33,22,12,38,1,46,23,29,35,4,13,26,39,31,41,7,30,25,18,14,43,24,27,16,10,42,3,34,49,48,40,15,6,9,28,32,2,17,44,19,36,45,21,5,20,11,37,8,	50
2		**Perturbation Type 2 Interchanging index pairs**		In above cell CR1, enter the routing of interest to be evaluated. Separate location indices with commas and no other intervening characters. Rightmost character must be a comma. After entry is completed, paste the result as a value. Output information appears in cells CP4 and CQ4. Before using this feature, complete blue-filled cell CJ6 and below with the desired number of stops per vehicle.	
3	Enter in cell CP3 --> index to be interchanged or use spin button arrows to change the index in cell CP3.	1	▲ ▼	,47,33,22,12,38,1,46,23,29,35,4,13,26,39,31,41,7,30,25,18,14,43,24,27,16,10,42,3,34,49,48,40,15,6,9,28,32,2,17,44,19,36,45,21,5,20,11,37,8,	
4	Best index i --> to interchange with index in cell CP3	46	916	<-- Routing distance resulting from interchange of indices displayed in cells CP3 and CP4	Routing distance, max vehicle load, MDTD2V
5	**i**			Routings resulting from interchange of index in cell CP3 with indices i in cells CO6 and below … one at a time	916,98,82
6	1	N/A	N/A	N/A	0,0,0
7	2	,47,33,22	,47,33,22,12	47,33,22,12,38,2,46,23,29,35,4,13,26,39,31,41,7,30,25,18,14,43,24,27,16,10,42,3,34,49,48,40,15,6,9,28,32,1,17,44,19,36,45,21,5,20,11,37,8,	1218,110,262
8	3	,47,33,22	,47,33,22,12	47,33,22,12,38,3,46,23,29,35,4,13,26,39,31,41,7,30,25,18,14,43,24,27,16,10,42,1,34,49,48,40,15,6,9,28,32,2,17,44,19,36,45,21,5,20,11,37,8,	1056,98,64

location stops for each vehicle) of blue-filled cells CJ6 and below relate to the routing of cell CR1; see Figure 3. Upon completion of Step 2, the information needed for composing the candidate list is available, that is, moving to Step 3 and testing.

Figure 3 displays the section of worksheet 'VRP Model' to test (Step 3) the perturbations of the candidate list. Each candidate is tested in the order of the list. If the candidate list of column 3, Table 6, applied, students would begin by editing the best available routing that is maintained in cell CC1 according to the perturbation at the top of the list. For this situation, the first candidate to be tested is item 1, calling for the repositioning of location index 8 between indices 5 and 20. The editing in cell CC1 to this effect resulted in the fleet distance value of 910 in cell CK4, not shown in Figure 3. It is better than the fleet distance value of 916 that was known at the beginning of the testing as shown in cell CK4 of Figure 3. Point out to students

Figure 3. Screenshot of the Area of Worksheet 'VRP Model' for Evaluating a Routing

	CB	CC	CD	CE	CF	CG	CH	CI	CJ	CK	CL	CM	CN
1	50	47,33,22,12,38,1,46,23,29,35,4,13,26,39,31,41,7,30,25,18,14,43,24,27,16,10,42,3,34,49,48,40,15,6,9,28,32,2,17,44,19,36,45,21,5,20,11,37,8,											
2	**Evaluating a routing of interest**	Enter the routing to be evaluated in blue-filled cell CC1 with location indices separated by commas and no other intervening characters such as blanks. The last character in the string should be a comma. The desired number of location stops per vehicle must be entered in blue-filled cells CJ6 and below. Vehicles not utilized in the routing should have blank entries in the corresponding cells of CJ6 and below. The fleet distance value appears in cell CK4 and other relevant information in cells CL4, CK6 and below, CL6 and below, and cell CM6. For details about distances traveled by each vehicle and their loads,see cells CG6 and below and cells CH6 and below with vehicle reference in cells CE6 and below.											
3									No. of fleet location stops	Fleet travel distance	Fleet delivery load		
4									49	916	609		
5		Comma	Position	Vehicle	Location index	Distance	Vehicle load	Vehicle id	Number of location stops for each vehicle	Vehicle travel distance	Vehicle delivery load	Max difference between any two vehicle travel distances MDTD2V	
6		3	1	1	47	21	4	1	7	112	89	82	
7		6	2	1	33	4	9	2	7	120	98		
8		9	3	1	22	2	12	3	7	111	73		
9		12	4	1	12	2	12	4	7	93	75		
10		15	5	1	38	2	7	5	7	151	93		
11		17	6	1	1	26	21	6	7	175	97		
12		20	7	1	46	55	24	7	7	154	84		

that the displays in cells CK6:CK12, CL6:CL12, and CM6 relate to the feasibility of the edited routing. The displays in cells CK6:CK12 identify the travel distances for each vehicle, and their summative balancing feature (MDTD2V) appears in cell CM6. The latter should be inspected for its compliance with the bound (blue-filled cell P2) set to 100 for the illustration. The displays in cells CL6:CL12 show the vehicle loads associated with the routing. They should be examined for conformance to the uniform vehicle capacity of 100 (blue-filled cell L2) set for the illustration. Visual inspection suffices for assessing conformance. If feasibility is confirmed and the fleet distance value of cell CK4 is better than the value for the best found routing at the time of testing, as it is in this case, the perturbation in cell CC1 is left permanent. The routing there is noted as the current best found. The associated fleet distance value becomes the value of comparison in subsequent candidate list testing. If the fleet distance value did not test well, the edit to cell CC1 should be undone using Excel's 'undo' feature. Testing would go on to the next (second) candidate in the list, that is, item 2 (19,2,36) calling for the repositioning of location index 2 between indices 19 and 36 in cell CC1.

Recall the following. When evaluated in Steps 1 and 2, each perturbation is treated as if it were the only change to the best found routing in effect at the time of examination. However, in Step 3, because of successive changes that preceded the testing of a perturbation candidate, its implementation as described may no longer be possible. A candidate so affected should not be tested and should be noted with a strikethrough in the list. Also, testing may produce an infeasible or distance-inferior result. Until these inspections are completed, disposition of items in the candidate list cannot be made, hence, the reason for the question marks in column 3 of Table 6. The candidate list indicates the order in which the perturbations should be tested for permanent implementation and not the expected fleet distance values. This should be pointed out to students. Also, bring to the students' attention that, in going through the candidate list, they are identifying compound serial perturbations that would otherwise not be evident. This is an important focus to keep in mind when navigating the detail.

Testing the complete candidate list is referred to as one iteration of the four-step procedure. It may be looked upon as culling the list. Throughout the testing, the entry of cell CC1 is maintained as the best found routing. Demonstration of how to proceed completely through selective iterations (rows) of Table 5 is given in Appendix C.

At the end of the exercise, after the demonstration of one iteration, allow some time for discussion of how to proceed next, keeping in mind what to recommend as a final routing. This should include discussion of the ambiguity/incertitude associated with the final recommendation. Ask who among the students would continue the investigation and why they would do so. That discussion has good learning value. It should reflect much of what we promote in contemporary decision making, that is, teamwork, consensus building, creative use of technology, and dealing with the ambiguity/incertitude that accompanies a decision or recommendation. The document 'In class VRP exercise.pdf' among the supplemental material has more detail on conducting the exercise.

It may also be instructive to point out to students the span and locus of the changes that one iteration of the perturbation procedure can bring about. For an example, see the bolded indices of the routing of row 6, column 3, of Table 5. To search for possibly better routes, we need to 'shake up' the subject route. This is done by perturbing its location indices.

It is the experience of the authors that, given the means and guidance, students will investigate beyond the first routing produced by the automatic means of ES, another Solver, or by other means. In these situations, the instructor may be looked upon as a coach whose task is to move students to explore how the routing at hand can be perturbed in the search for something better. It becomes an exercise in coming to grips and resolution with the inherent uncertainty that VRP problem solving presents.

Teaching the methods of perturbation may appear daunting to the instructor. However, there is learning value in taking students through even one iteration of the four-step procedure of Section 3.4, for example, any of rows 3–11 of Table 5. It shows students how the problem solving may continue when their intellectual curiosity or an imperative (course assignment or project, client-sponsored project, etc.) drives doing so. The exercise demonstrates to students how the ambiguity of the final recommended routing solution may be addressed. As VRP problem Solvers, they do not have to be limited to the output of the automatic methods of textbook software. This is different problem solving. In the opinion of the authors, it is a good learning experience.

6. Summary, Limitations, and Future Development

In this article, it was shown how the VRP solution space becomes accessible to students and, in turn, to discovery of alternative routing solutions. The demonstrated techniques for doing so may be looked upon as complements to Excel-based Solvers such as ES and other methods that employ solution methodologies that do not guarantee the optimal routing. Proceeding in the demonstrated ways makes students active in solution discovery and not passive bystanders waiting for ES and its little understood proprietary algorithm to conclude and learning little more than what the vigil of waiting produces. We, as teachers, should move students beyond nominal problem solving of this kind. This is true for heuristic-based VRP solution methods in which ambiguity/inexactness is characteristic of the result. Enabling students to explore beyond the solutions of introductory textbook software is a way to address the incertitude of 'what could be better.'

Although the Excel workbook that is included in the supplemental material can accommodate up to $n = 49$ delivery sites and $K = 10$ vehicles, it is not intended as a template. It is offered to facilitate reproducing the results presented in the illustrations and, in turn, to promote student investigations. Instructors and students are welcome to adapt the worksheets to meet their problem-solving needs.

The authors are working to further facilitate types 1 and 2 investigations within worksheet 'VRP Model' in ways that are understood by most students.

Acknowledgments

The authors acknowledge the significant contributions that the editor and an anonymous associate editor have made in helping to shape this paper, especially Section 3.3. Their efforts are much appreciated.

Appendix A

The formalization of the VRP treated in this paper is adapted from Kulkarni and Bhave (1985) and appears as follows. For n number of sites to be visited by K number of vehicles, the objective is to

$$\text{minimize} \sum_{k=1}^{K}\sum_{i=0}^{n}\sum_{j=1}^{n} d_{ij}x_{ijk} \tag{A.1}$$

subject to

$$\sum_{i=0}^{n}\sum_{k=1}^{K} x_{ijk} = 1 \quad j = 1,\ldots,n, \tag{A.2}$$

$$\sum_{j=0}^{n}\sum_{k=1}^{K} x_{ijk} = 1 \quad i = 1,\ldots,n, \tag{A.3}$$

$$\sum_{i=0}^{n} x_{ihk} - \sum_{j=0}^{n} x_{hjk} = 0 \quad k = 1,\ldots,K \quad h = 0,1,\ldots,n, \tag{A.4}$$

$$\sum_{i=0}^{n} Q_i \sum_{j=0}^{n} x_{ijk} \le VCap_k \quad k = 1,\ldots,K, \tag{A.5}$$

$$\max_{k=1,\ldots,K}\left(\sum_{i=0}^{n}\sum_{j=0}^{n} d_{ij}\ x_{ijk}\right) - \min_{k=1,\ldots,K}\left(\sum_{i=0}^{n}\sum_{j=0}^{n} d_{ij}\ x_{ijk}\right) \le VBal, \tag{A.6}$$

$$\sum_{i=1}^{n}\sum_{j=1}^{n} x_{ijk} = NVStops_k \quad k = 1,\ldots,K, \tag{A.7}$$

$$\sum_{j=1}^{n} x_{0jk} \le 1 \quad k = 1,\ldots,K, \tag{A.8}$$

$$\sum_{i=1}^{n} x_{i0k} \le 1 \quad k = 1,\ldots,K, \tag{A.9}$$

$$\text{subtour elimination constraints,} \tag{A.10}$$

$$x_{ijk} = 0 \text{ or } 1 \quad \text{for all } i,j,k, \tag{A.11}$$

where $x_{ijk} = 1$ (0) if vehicle k travels (not) between sites i and j and $d_{ijk} = d_{ij} = d_{ji}$ is the distance traveled between locations i and j by vehicle k. The Q_i is the required delivery quantity for site i (= 1, …, n); $VCap_k$ is the load/carrying capacity of vehicle k (= 1, …, K); $VBal$ is the maximal allowed difference in travel distance between any two vehicles (MDTD2V); and $NVStops_k$ is the number of vehicle site visits assigned to vehicle k (= 1, …, K). The (A.1) represents the objective of minimizing fleet travel distance; (A.2) and (A.3) ensure that each site is visited by one vehicle; (A.4) is the conservation-of-flow requirement ensuring that any vehicle visiting a site leaves the same; (A.5) limits assignment of customer delivery requirements for any vehicle to the vehicle's carrying capacity; (A.6) and (A.7) enforce the requirements of balanced vehicle/driver assignments; (A.8) and (A.9) call for the K number of vehicles to return to the depot (= 0); (A.10) are restrictions that ensure each vehicle does not return to the depot before visiting each assigned site. For specification of the latter, see Kulkarni and Bhave (1985), who also noted some redundancies in (A.2)–(A.9).

For the illustration of Section 4 in which $n = 49$ and $K = 7$, the following apply. The $Q_1, \ldots, Q_{n=49}$ are as given in Table 3; $VCap_1 = \ldots = VCap_{K=7} = 100$; $VBal$ = MDTD2V = 100; and $NVStops_1 = \ldots = NVStops_{K=7} = n/K = 49/7$.

Accommodation of the penalty terms for constraint violation in the objective function may be formalized as

$$\text{Minimize} \sum_{k=1}^{K}\sum_{i=0}^{n}\sum_{j=0}^{n} d_{ij}\, x_{ijk} + P_1 V_1 + P_2 V_2, \tag{A.12}$$

where

$$V_1 = \max\left(\max_{k=1,\ldots,K}\left(\sum_{i=1}^{n} Q_i \sum_{j=0}^{n} x_{ijk} - VCap_k\right), 0\right), \tag{A.13}$$

$$V_2 = \max\left(\left(\max_{k=1,\ldots,K}\sum_{i=0}^{n}\sum_{j=1}^{n} d_{ij}x_{ijk} - \min_{k=1,\ldots,K}\sum_{i=0}^{n}\sum_{j=1}^{n} d_{ij}x_{ijk}\right) - VBal, 0\right), \tag{A.14}$$

and P_1 is the penalty associated with V_1 (exceeding vehicle load capacity) and P_2 is the penalty for V_2 (exceeding MDTD2V). The inner parenthetical expression of (A.13), if nonpositive, is, in absolute value, the amount of unused capacity for vehicle k; if positive, it is the amount by which a vehicle load exceeds its capacity, and the penalty P_1 applies. The inner parenthetical expression of (A.14) is the maximal difference between the largest and smallest assigned travel distances among all K vehicles. When the outer expression is positive, MDTD2V is exceeded, and the penalty P_2 applies. In situations in which no penalties apply, $P_1 = P_2 = 0$; otherwise, the penalty is assigned a distinctive positive value. In Section 4, the routing derived from ES was obtained using $P_1 = 10{,}000$ and $P_2 = 20{,}000$.

For any routing solution obtained from (A.1)–(A.11), there is equivalency in what the $X_{ij} = 1$ denote and the corresponding string form. It is pointed out with an example. For a VRP with $K = 1$ and $n = 3$, the routing sequence 0,2,1,3,0 in string notation is equivalent to $X_{0,2} = X_{2,1} = X_{1,3} = X_{3,0} = 1$ and $X_{i,j} = 0$ otherwise. When presented in this manner, the string representation is seen as a concatenation of the subscripts of the $X_{i,j}= 1$. For the routing situation, there are as many as $n(n + 1)$ $X_{i,j}$ decision variables in the modeling (A.1)–(A.11). When the location identifier 0 is omitted, there are n location references in a routing string.

Appendix B

Table B.1 is a schematic that describes how to use worksheet 'VRP Model' that appears in Excel workbook 'VRP Example.xlsx.' The workbook is included in the supplemental material. All cell references in Table B.1 relate to worksheet 'VRP Model.'

Table B.1. How to Use Worksheet 'VRP Model'

Category of information	For solution using Evolutionary Solver	For using the general purpose routing evaluator	For performing type 1 perturbations Repositioning evaluator	For performing type 2 perturbations Interchanging Index pairs evaluator
Where in worksheet 'VRP Model'	In the vicinity of green-filled cell E2	In the vicinity of green-filled cell CB2	In the vicinity of green-filled cell BQ2	In the vicinity of green-filled cell CP2
Input: Applies to all	Distance between each location pair referenced by location identifiers/indices $0,1,\ldots,n$ must be entered in blue-filled cells in the vicinity of cell L5. Row and column labels there assist data entry. Because the symmetric VRP is the problem of interest, the upper diagonal of pair distance values suffices. Enter among blue-filled cells in the range H2:AB2 input information as described in nearby cells.			
Input: Location of user supplied routing to be perturbed or analyzed	Optional: Enter starting route if available in cells D6 and below; otherwise, those cells should be blank.	Enter routing to be evaluated in blue-filled cell CC1. See note in cell CC2.	Enter routing to be perturbed by repositioning in blue-filled cell BR1. See note in cell BR2.	Enter routing to be perturbed by interchanging index pairs in blue-filled cell CR1. See note in cell CR2.
Input: Location of user supplied first location index to be perturbed	N/A	N/A	Enter in blue-filled cell BP3.	Enter in blue-filled cell CP3.
Input: Location of other necessary user-supplied information	See blue-filled cells: H2, J2, L2, P2, X2, AB2; cells A6 and below; cells H6 and below. Enter there the requisite data as described in nearby cells. See the MSWord document 'Setup for Solver.docx' in supplemental material.	Enter in blue-filled cells CJ6 and below the desired number of delivery stops for each vehicle. See note in cell CC2.	Enter in blue-filled cells CJ6 and below the desired number of delivery stops for each vehicle. See note in cell CC2.	Enter in blue-filled cells CJ6 and below the desired number of delivery stops for each vehicle. See note in cell CC2.
Location of spin button utility	N/A	N/A	In the vicinity of cell BQ3	In the vicinity of cell CQ3
Output: Location of the output routing and its fleet distance and related values.	Routing appears in cells D6 and below; fleet distance value in cell I4 and related routing information in cells I6 and below, J6 and below, cell K6, and cells J4:L4.	Best fleet distance value appears in cell CK4 and related routing values in cells CK6 and below, CL6 and below, and in cell CM6	Best fleet distance value and notation of corresponding routing appear in cells BQ4 and BQ5.	Best fleet distance value is displayed in cell CQ4 and corresponding index pairs in cells CP3 and CP4.

Appendix C

The complete candidate list that resulted from Steps 1–3 of the procedure outlined in Section 3.4 is not shown for the scenarios of rows 3–11 of Table 5. Only those candidates whose implementation offered successively feasible improvements in fleet distance are reported there. Table C.1 provides the details of culling the candidate list for the scenarios of rows 6–8 of Table 5. For each, the candidate list appears in column 5 of Table C.1. Some candidates could not be implemented as described given permanently implemented perturbations that preceded their testing. Some produced inferior fleet distance values. These instances are noted with strikethroughs in Table C.1. The candidate perturbations that do not have a strikethrough were permanently implemented as encountered in the testing. For these instances, the resulting fleet distances values are noted.

The testing reported in the row labeled 6 began with the first perturbation item 1 (5,8,20) of the candidate list in row 6, column 5, of Table C.1. It tested successfully with a fleet distance value of 910. Consequently, its implementation was left permanent. The second perturbation (item 2: 19,2,36) in the list also tested successfully with fleet distance 906, and its implementation was left permanent. Next, because the third perturbation (item 3: ,8,21) and the fourth perturbation (item 9: 8,11) addressed location index 8 that was previously implemented as item 1 (5,8,20), they lost their contextual references and were not tested and consequently struck as shown. The fifth perturbation (item 10: 9,47) tested successfully with fleet

Table C.1. Details of the Results of Rows 6–8 of Table 5

Row number in Table 5	Fleet distance beginning route	Step 1 outcomes best of repositioning (item reference, how bolded index was repositioned, distance value)			Step 2 outcomes best of interchanging index pairs (item reference, index pair interchanged, distance value)			Step 3 outcomes candidate list of ranked best outcomes (item reference, perturbation, resulting fleet distance value)			Fleet distance value and corresponding routing at end of Step 3[b]
6	916	1	5,**8**,20	910	9	9 8,11	913	1	5,8,20	910	897
	47,33,22,12,38,1,46,	2	19,**2**,36	912	10	10 9,47	913	2	19,2,36	906	**9**,33,22,12,4**6**,1,**38**,
	23,29,35,4,13,26,39,	3	,8,**21**[a]	912	11	11 11,8	913	3	~~,8,21~~[a]	903	23,29,35,4,13,26,39,
	31,41,7,30,25,18,14,	4	17,**19**,44	914	12	12 20,8	913	9	~~8,11~~	900	31,41,7,30,**14**,18,**25,**
	43,24,27,16,10,42,3,	5	14,**25**,43	914	13	13 38,46	913	10	9,47	898	43,24,27,16,10,42,3,
	34,49,48,40,15,6,9,	6	,8,**37**[a]	914	14	14 14,25	914	11	~~11,8~~	897	34,49,48,40,15,6,**47,**
	28,32,2,17,44,19,36,	7	46,**38**,23	914	15	15 19,44	914	12	~~20,8~~		28,32,17,1**9**,**44**,**2**,36,
	45,21,5,20,11,37,8,	8	19,**44**,36	914				13	38,46		45,21,5,**8**,20,11,37,
								14	14,25		
								15	19,44		
								4	~~17,19,44~~		
								5	~~14,25,43~~		
								6	~~,8,37~~[a]		
								7	~~46,38,23~~		
								8	~~19,44,36~~		
7	897	1	1,**9**,38	887	5	6,36	889	1	1,9,38	887	877
	9,33,22,12,46,1,38,	2	38,**33**,23	894	6	9,38	894	5	6,36	879	33,22,12,46,1,**9**,38,
	23,29,35,4,13,26,39,	3	36,**21**,45	896	7	33,38	894	2	~~38,33,23~~	~~880~~	23,29,35,4,13,26,39,
	31,41,7,30,14,18,25,	4	9,**38**,33	896	8	17,44	895	6	~~9,38~~	~~885~~	31,41,7,30,14,18,25,
	43,24,27,16,10,42,3,				9	21,36	896	7	~~33,38~~	877	43,24,27,16,10,42,3,
	34,49,48,40,15,6,47,				10	22,9	896	8	17,44		34,49,48,40,15,**36**,47,
	28,32,17,19,44,2,36,							3	~~36,21,45~~		28,32,**44**,19,**17**,2,**6**,
	45,21,5,8,20,11,37,							4	~~9,38,33~~		45,21,5,8,20,11,37,
								9	~~21,36~~		
								10	~~22,9~~		
8	877	1	15,**48**,36	876	4	47,45	868	4	47,45	868	866
	33,22,12,46,1,9,38,	2	21,**45**,5	876	5	36,21	870	5	~~36,21~~	579	33,22,12,46,1,9,38,
	23,29,35,4,13,26,39,	3	6,**21**,45	876	6	48,15	875	6	48,15	866	23,29,35,4,13,26,39,
	31,41,7,30,14,18,25,							1	~~15,48,36~~		31,41,7,30,14,18,25,
	43,24,27,16,10,42,3,							2	~~21,45,5~~		43,24,27,16,10,42,3,
	34,49,48,40,15,36,47,							3	~~6,21,45~~		34,49,**15**,40,**48**,36,**45**,
	28,32,44,19,17,2,6,										28,32,44,19,17,2,6,
	45,21,5,8,20,11,37,										**47**,21,5,8,20,11,37,

[a]Denotes that the bolded index was repositioned to the rightmost position of the routing string.
[b]Bolding denotes indices affected by implemented perturbations.

Figure C.1. Appearance of Worksheet 'VRP Model' at the Start of the Investigations of Row Labeled 6 in Table C.1

	CB	CC	CD	CE	CF	CG	CH	CI	CJ	CK	CL	CM	CN
1	50	47,33,22,12,38,1,46,23,29,35,4,13,26,39,31,41,7,30,25,18,14,43,24,27,16,10,42,3,34,49,48,40,15,6,9,28,32,2,17,44,19,36,45,21,5,20,11,37,8,											
2	Evaluating a routing of interest	Enter the routing to be evaluated in blue-filled cell CC1 with location indices separated by commas and no other intervening characters such as blanks. The last character in the string should be a comma. The desired number of location stops per vehicle must be entered in blue-filled cells CJ6 and below. Vehicles not utilized in the routing should have blank entries in the corresponding cells of CJ6 and below. The fleet distance value appears in cell CK4 and other relevant information in cells CL4, CK6 and below, CL6 and below, and cell CM6. For details about distances traveled by each vehicle and their loads,see cells CG6 and below and cells CH6 and below with vehicle reference in cells CE6 and below.											
3									No. of fleet location stops	Fleet travel distance	Fleet delivery load		
4									49	916	609		
5		Comma	Position	Vehicle	Location index	Distance	Vehicle load	Vehicle id	Number of location stops for each vehicle	Vehicle travel distance	Vehicle delivery load	Max difference between any two vehicle travel distances MDTD2V	
6		3	1	1	47	21	4	1	7	112	89	82	
7		6	2	1	33	4	9	2	7	120	98		
8		9	3	1	22	2	12	3	7	111	73		
9		12	4	1	12	2	12	4	7	93	75		
10		15	5	1	38	2	7	5	7	151	93		
11		17	6	1	1	26	21	6	7	175	97		
12		20	7	1	46	55	24	7	7	154	84		

distance 903, so its implementation was left permanent. The sixth (item 11: 11,8) and seventh (item 12: 20,8) in the list were next for testing. Because each involved location index 8 that was previously repositioned as item 1 (5,8,20), these perturbations as described lost their contextual references and were not tested. They were struck in the list. Next, the eighth perturbation (item 13: 38,46), the ninth (item 14: 14,25), and the 10th (item 15: 19,44) were tested separately in that order and produced successively better fleet distance values of 900, 898, and 897, respectively. Accordingly, they were permanently implemented. The next perturbation scheduled for testing was item 4 calling for the repositioning 17, 19, 44 that due to preceeding implementations was configured as such at this point. Items referenced as 5 - 8 were scheduled for testing next in that order. However, each could not be configured as described due to preceeding implementations of items 14, 1, 13, and 15. This concluded the testing of row 6. In summary, items 1, 2, 10, and 13–15 tested well as encountered in the investigations and were permanently implemented. The last unstruck result in column 5 row 6, identified the best found routing at the conclusion of this testing. It served as the subject routing for type 1 (Step 1) and type 2 (Step 2) perturbation investigations to be examined in the next iteration of the procedure, that is, row 7.

In the row labeled 7 in Table C.1, the outcomes of the next iteration are displayed. The best outcomes of Steps 1 and 2 are noted there as well as the resulting list of perturbation candidates to be tested in column 5. The first two perturbations (items 1 and 5) at the top of the candidate list tested well individually in successive order with fleet distances of 887 and 879, respectively, and were permanently implemented as encountered. The third through fifth perturbations in the list, that is, item 2 (38,33,23), item 6 (9,38), and item 7 (33,38), were struck from the list for the following reasons: item 2 because its testing resulted in the fleet distance value of 880 that was inferior to the 879 value in effect at the time of its testing, item 6 because it involved location index 9 that was previously repositioned as item 1, and item 7 because its testing resulted in the fleet distance value 885 that was inferior to the 879 value in effect at the time of its testing. The sixth perturbation (item 8: 17,44) tested successfully and was permanently implemented with the new best fleet distance value of 877. The seventh perturbation (item 3: 36,21,45) was next for testing. Items referenced as 3, 4, 9, and 10 were scheduled next for testing in that order. They could not be implemented as described and were struck in the list. The eighth perturbation (item 4: 9,38,33) was next for testing. Because of the repositioning of index 9, index 38 could not be repositioned between indices 9 and 33. It was struck in the list. The ninth (item 9: 21,36) and 10th (item 10: 22,9) perturbations in the list were next for testing in that order. Because location index 36 was previously interchanged (see item 5: 6,36) and location index 9 was previously repositioned (see item 1: 1,9,38), these perturbations were not tested. They were struck in the candidate list. This concluded the testing. In summary, items 1, 5, and 8 tested well and were permanently implemented as encountered in the testing. The best found routing at this point is displayed in the last column of the row labeled 7 with fleet distance value of 877. It served as the subject routing for type 1 (Step 1) and type 2 (Step 2) perturbations of the next iteration.

The following is additional detail regarding the preceding narrative. Figure C.1 is a partial screen shot of the area of worksheet 'VRP Model' as it appeared at the start of the investigation reported in the row labeled 6 of Table C.1 with fleet distance value of 916. Cell CC1 in Figure C.1 displays

Figure C.2. Appearance of Worksheet 'VRP Model' at the Conclusion of the Investigations of Row Labeled 6 in Table C.1

	CB	CC	CD	CE	CF	CG	CH	CI	CJ	CK	CL	CM	CN
1	50	9,33,22,12,46,1,38,23,29,35,4,13,26,39,31,41,7,30,14,18,25,43,24,27,16,10,42,3,34,49,48,40,15,6,47,28,32,17,19,44,2,36,45,21,5,8,20,11,37,											
2	**Evaluating a routing of interest**	Enter the routing to be evaluated in blue-filled cell CC1 with location indices separated by commas and no other intervening characters such as blanks. The last character in the string should be a comma. The desired number of location stops per vehicle must be entered in blue-filled cells CJ6 and below. Vehicles not utilized in the routing should have blank entries in the corresponding cells of CJ6 and below. The fleet distance value appears in cell CK4 and other relevant information in cells CL4, CK6 and below, CL6 and below, and cell CM6. For details about distances traveled by each vehicle and their loads,see cells CG6 and below and cells CH6 and below with vehicle reference in cells CE6 and below.											
3									No. of fleet location stops	Fleet travel distance	Fleet delivery load		
4									49	897	609		
5		Comma	Position	Vehicle	Location index	Distance	Vehicle load	Vehicle id	Number of location stops for each vehicle	Vehicle travel distance	Vehicle delivery load	Max difference between any two vehicle travel distances MDTD2V	
6		2	1	1	9	28	9	1	7	118	94	77	
7		5	2	1	33	6	9	2	7	120	98		
8		8	3	1	22	2	12	3	7	109	73		
9		11	4	1	12	2	12	4	7	93	75		
10		14	5	1	46	19	24	5	7	139	88		
11		16	6	1	1	8	21	6	7	170	97		
12		19	7	1	38	53	7	7	7	148	84		

the best available routing at the start of this testing. Items of the candidate list of row 6 were tested in blue-filled cell CC1 one at a time beginning with perturbation item 1 (5,8,20/910) at the top of the list. Each item was tested by making the indicated edit/perturbation in cell CC1 of the worksheet. If the edit (perturbation under investigation) resulted in a feasible solution with fleet distance value (cell CK4) less than the value for the best found routing at that point of the investigation, the edit was left permanent in cell CC1, and the routing there was denoted as the best found. Otherwise, the edit was undone. For feasibility conformance, cells CL6:CL12 (vehicle loads) and cell CM6 (MDTD2V) were inspected. After these assessments and disposition of the edit, the investigation moved to the next item in the candidate list, editing cell CC1 accordingly, comparing the result to the current best found route and depending on the outcome of the comparison leaving the edit permanent or undoing it.

Figure C.2 is a partial screen shot of the worksheet area at the conclusion of the testing described in the row labeled 6 in Table C.1. In Table C.1, the second and the last column entries of each row were presented for the sake of illustration. In practice, using worksheet 'VRP Model,' only the best available routing is maintained in cell CC1. It changes as testing warrants. Once the candidate list is composed, testing moves quickly. In composing the lists of types 1 and 2 testing candidates, usually five to eight perturbation instances sufficed for each.

References

Anderson D, Sweeney D, Williams T, Camm J, Cochoran J, Fry M, Ohlmann J (2019) *An Introduction to Management Science: Quantitative Approaches to Decision Making*, 15th ed. (Cengage, Boston).

Augerat P, Belenguer J, Benavent E, Corberan A, Naddef D, Rinaldi G (1995) Computational results with a branch and cut code for the capacitated vehicle routing problem. Technical Report 949-M, Universite Joseph Fourier, Grenoble, France.

Baker K, Camm J (2005) On the use of integer programming vs. evolutionary Solver in spreadsheet optimization. *INFORMS Trans. Ed.* 5(3):1–7.

Braekers K, Ramaekers K, Van Nieuwenhuyse I (2016) The vehicle routing problem: State of the art classification and review. *Comput. Indust. Engrg.* 99:300–313.

Carter A, Ragsdale C (2002) Scheduling newspaper advertising inserts using genetic algorithms. *OMEGA: Internat. J. Management Sci.* 30(6):415–421.

Drake M, Griffin M, Swann J (2011) Case article: Keeping logistics under wraps. *INFORMS Trans. Ed.* 11(2):57–62.

Eksioglu B, Vural A, Reisman A (2009) The vehicle routing problem: A taxonomic review. *Comput. Indust. Engrg.* 57(4):1472–1483.

Erdogan G (2017) An open source spreadsheet Solver for vehicle routing problems. *Comput. Oper. Res.* 87:62–72.

FrontlineSystems (2020) Accessed May 5, 2021, https://www.solver.com/excel-solver-evolutionary-solving-method-stopping-conditions.

Gutin G, Punnen A, eds. (2002) *The Traveling Salesman Problem and Its Variations* (Kluwer Academic Publishers, Dordrecht, Netherlands).

Hesse R, Scerno D (2009) How electronic spreadsheets changed the world. *INFORMS J. Appl. Analytics* 39(2):159–167.

Hillier F, Hillier M (2019) *Introduction to Management Science: A Modelling and Case Studies Approach with Spreadsheet*, 6th ed. (McGraw Hill Education, New York).

Hillier F, Lieberman G, Nag B, Basu P (2017) *Introduction to Operations Research*, 10th ed. (McGraw Hill Education, Noida, India).

Koksalan M, Salman F (2003) Beer in the classroom: A case study of location and distribution decisions. *INFORMS Trans. Ed.* 4(1):65–77.

Kulkarni R, Bhave P (1985) Integer programming formulations of vehicle routing problems. *Eur.n J. Oper. Res.* 20(1):58–67.

Laporte G (2009) Fifty years of vehicle routing. *Transportation Sci.* 43(4):408–416.

Mason A (2013) SolverStudio: A new tool for better optimisation and simulation modeling in excel. *INFORMS Trans. Ed.* 14(1):45–52.

Milburn A, Kirac E, Hadianniasar M (2017) Case article-growing pains: A case study of large-scale vehicle routing. *INFORMS Trans. Ed.* 17(2):75–80.

Pillac V, Gendreau M, Guéret C, Medaglia A (2013) A review of dynamic vehicle routing problems. *Eur. J. Oper. Res.* 225(1):1–11.

Pollaris H, Braekers K, Caris A, Janssens G, Limbourg S (2015) Vehicle routing problems with loading constraints: State-of-the art and future directions. *OR Spectrum* 37:297–330.

Ragsdale C (2001) Teaching management science with spreadsheets: From decision models to decision support. *INFORMS Trans. Ed.* 1(2):68–74.

Ragsdale C (2015) *Spreadsheet Modeling and Decision Analysis*, 7th ed. (South-Western Cengage Learning, Mason, OH).

Ragsdale C (2018) *Spreadsheet Modeling and Decision Analysis: A Practical Introduction to Business Analytics*, 8th ed. (South-Western Cengage Learning, Mason, OH).

Taja H (2018) *Operations Research: An Introduction*, 10th ed. (Pearson Education, Uttar Pradesh, India).

Toth P, Vigo D (2014) *Vehicle Routing: Problems, Methods, and Applications*, 2nd ed., MOS-SIAM Series on Optimization, vol. 18 (SIAM, Philadelphia).

Vidal T, Laporte G, Matl P (2020) A concise guide to existing and emerging vehicle routing variants. *Eur. J. Oper. Res.* 286(2): 401–416.

Winston W, Albright S (2009) *Practical Management Science*, rev. 3rd ed. (South-Western Cengage Learning, Mason, OH).

Winston W, Albright S (2012) *Practical Management Science*, 4th ed. (South-Western Cengage Learning, Mason, OH).