



11-2009

An Empirical Study on the Benefit of Split Loads with the Pickup and Delivery Problem

Maciek Nowak

Loyola University Chicago, mnowak4@luc.edu

Chelsea C. White

Georgia Institute of Technology - Main Campus

Ozlem Ergun

Georgia Institute of Technology - Main Campus

Follow this and additional works at: https://ecommons.luc.edu/isom_facpubs



Part of the [Business Administration, Management, and Operations Commons](#), [Management Information Systems Commons](#), and the [Technology and Innovation Commons](#)

Author Manuscript

This is a pre-publication author manuscript of the final, published article.

Recommended Citation

Nowak, Maciek; White, Chelsea C.; and Ergun, Ozlem. An Empirical Study on the Benefit of Split Loads with the Pickup and Delivery Problem. *European Journal of Operational Research*, 198, 1: , 2009. Retrieved from Loyola eCommons, Information Systems and Operations Management: Faculty Publications & Other Works, <http://dx.doi.org/10.1016/j.ejor.2008.09.041>

This Article is brought to you for free and open access by the Faculty Publications and Other Works by Department at Loyola eCommons. It has been accepted for inclusion in Information Systems and Operations Management: Faculty Publications & Other Works by an authorized administrator of Loyola eCommons. For more information, please contact ecommons@luc.edu.



This work is licensed under a [Creative Commons Attribution-NonCommercial-No Derivative Works 3.0 License](#).
© 2008 Elsevier B.V.

An Empirical Study on the Benefit of Split Loads with the Pickup and Delivery Problem

Maciek Nowak ^{a,*}, Ozlem Ergun ^b, Chelsea C. White, III ^b

^a*Loyola University Chicago, Department of Information Systems and Operations Management, 1 E. Pearson Ave., Chicago, IL, 60611, USA, voice: 312-915-6113, fax: 312-915-7207*

^b*Georgia Institute of Technology, School of Industrial and Systems Engineering, 765 Ferst Drive, NW, Atlanta, Georgia, 30332, USA*

Abstract

Splitting loads such that the delivery of certain loads is completed in multiple trips rather than one trip has been shown to have benefit for both the classic Vehicle Routing Problem and the Pickup and Delivery Problem. However, the magnitude of the benefit may be affected by various problem characteristics. In this paper, we characterize those real world environments in which split loads are most likely to be beneficial. Based on practitioner interest, we determine how the benefit is affected by the mean load size and variance, number of origins relative to the number of destinations, the percentage of origin-destination pairs with a load requiring service, and the clustering of origin and destination locations. We find that the magnitude of benefit: is greatest for load sizes just over one half vehicle capacity as these loads can not be combined without splitting, while they are the easiest to combine on a vehicle with splitting; increases as the number of loads sharing an origin or destination increases because there are more potential load combinations to split at each stop; and increases as the average distance from an origin to a destination

increases because splitting loads reduces the trips from origins to destinations.

Key words: transportation, vehicle routing, split pickup and delivery

1 Introduction

2 Splitting loads such that the delivery of certain loads is completed in multi-
3 ple trips rather than one trip results in opportunities for a reduction in cost
4 and the number of vehicles used. Several studies have shown the benefit of
5 split deliveries for the classic Vehicle Routing Problem (VRP), in which a
6 vehicle operating out of a depot makes a series of deliveries on each route
7 ((Dror et al., 1994), (Frizzell and Giffin, 1995), (Sierksma and Tijssen, 1998),
8 (Archetti et al., 2006)). More recently, Nowak et al. (2008) quantified the ben-
9 efit for the Pickup and Delivery Problem (PDP), in which a vehicle picks up
10 a load from a specific origin and delivers it to its destination. They showed
11 theoretically that the optimal load size for splitting is just above one half of
12 a truckload and supported this result with empirical evidence. Furthermore,
13 a real world example was used to show that certain problem characteristics
14 may limit the benefit of split loads.

15 Although the theoretical results are of interest, practitioners have found the
16 results regarding the characteristics of the problem that have an effect on the
17 benefit of split loads to be of more use. The real world case presented in Nowak
18 et al. (2008) showed that these benefits are affected by the per stop cost asso-

* Corresponding author.

Email addresses: mnowak4@luc.edu (Maciek Nowak), oergun@isye.gatech.edu
(Ozlem Ergun), cwhite@isye.gatech.edu (Chelsea C. White, III).

19 ciated with each pickup or delivery, the size of the loads requiring service and
20 the number of common origins or destinations requiring service. In this paper,
21 we determine the degree to which these characteristics impact the benefit. We
22 focus on the latter two characteristics of real world environments, the size of
23 loads to be delivered and the distribution of flow over the network, while also
24 analyzing the geographic orientation of origins and destinations. Specifically,
25 we determine how the magnitude of benefit is affected by mean load size and
26 variance, the number of origins relative to the number of destinations, the
27 percentage of origin-destination pairs with a load requiring service, and the
28 clustering of origin and destination locations.

29 We find that the magnitude of benefit: is greatest for load sizes just over one
30 half vehicle capacity as these loads can not be combined without splitting,
31 while they are the easiest to combine on a vehicle with splitting; increases as
32 the number of loads sharing an origin or destination increases because there are
33 more potential load combinations to split at each stop; and increases as the
34 average distance from an origin to a destination increases because splitting
35 loads reduces the trips from origins to destinations. Through this analysis,
36 practitioners will find a guide describing those instances where splitting loads
37 is most beneficial, as well as instances where the additional effort associated
38 with load splitting is not justified.

39 This paper is organized as follows. Section 2 describes the experimental design
40 including the common traits shared by all problem instances tested. Section
41 3 presents the results from the tests on the mean load size and variance.
42 Section 4 describes the effect that the number of origins relative to the number
43 of destinations has on the magnitude of benefit from split loads. Section 5
44 discusses how benefit is affected by the number of loads to be serviced from

45 a common origin or to a common destination. Section 6 analyzes two types
46 of location clusters and how they influence the benefit. Section 7 summarizes
47 the results of the paper.

48 **2 Experimental Design**

49 Several sets of problem instances are generated using the different character-
50 istics to be tested, as described in the following sections. However, all of the
51 instances share several common traits. The majority of problem sets tested
52 have 50, 100 or 150 transportation requests, as these sizes are similar to the
53 problem sizes used in testing of the SDVRP (Dror and Trudeau, 1989, 1990).
54 Each transportation request contains the origin and destination location co-
55 ordinates and the fraction of a truckload to be delivered. X and Y coordinates
56 for the pickup and delivery locations and load sizes are randomly generated.
57 The locations are uniformly distributed over the range $[-40,40]$ for both X and
58 Y coordinates for the problems in Sections 3-5, while Section 6 analyzes dif-
59 ferent distributions for the locations. The load sizes are all less than or equal
60 to vehicle capacity, which is set at one, without loss of generality. This is done
61 to determine the load sizes that benefit most from splitting that can otherwise
62 be serviced by a vehicle in one trip without splitting. The case study discussed
63 by Nowak et al. (2008) presented tests run with load sizes greater than vehicle
64 capacity.

65 The heuristic developed in Nowak et al. (2008) is used to solve each problem
66 under two scenarios, both with and without split loads. This heuristic func-
67 tions by randomly generating a split load for a solution that initially has a
68 unique route dedicated to each load. After the split load is created, local search

69 techniques, such as route combination, load swapping and insertion, are used
70 to improve the solution. Additional splits are created and local improvements
71 made until no cost reduction is found. The use of a heuristic is justified as
72 both the PDP and PDPSL are NP-hard (Nowak et al., 2008). The cost of the
73 solution in each case is equivalent to the distance traveled by the vehicle. The
74 two costs are compared to determine the percentage cost reduction that is
75 found through the use of split loads. An additional constraint is implemented
76 limiting routes to 500 miles. While loads could still be combined on a vehicle
77 due to the relative proximity of stops, this prevents a vehicle from servicing
78 all of the loads on one route.

79 In order to evaluate the heuristic, it was tested on a set of eleven standard
80 TSPLIB problem instances, finding solutions within 5% of the best known cost
81 of seven instances and within 10% of all instances. Given that the heuristic
82 was designed to focus on the additional complexities associated with the PDP,
83 we find these results to be acceptable. The heuristic was coded in C and all
84 experiments were run on a 2.4-GHz Xeon processor with a 400-Mhz frontside
85 bus and 2 GB RAM.

86 **3 Mean Load Size and Variance**

87 Previous research has shown that the most benefit from split loads occurs
88 with load sizes just over one half vehicle capacity. The following theorem was
89 presented by Nowak et al. (2008):

90 **Theorem 1** *Given the origin and destination locations of a set of k loads, a*
91 *vehicle of capacity Q , and a very small value, ϵ , let $v(PDPSL)$ be the cost of*
92 *the optimal PDPSL solution to deliver these loads and $v(PDP)$ be the cost of*

93 *the optimal PDP solution. Then the ratio $v(PDP)/v(PDPSL)$ is maximized*
94 *when the loads are all of size $Q/2 + \epsilon$, as $k \rightarrow \infty$.*

95 This theorem was supported with tests run on a variety of load instances
96 uniformly generated over several load size ranges. While these results provide
97 some basic insight into those loads that are most likely to lead to a cost reduc-
98 tion through splitting, a more in depth look at mean load size and variance is
99 of interest. Classifying industries by the size of loads delivered is difficult for
100 most goods transported, so this analysis will not define those sectors that are
101 most likely to gain benefit from split loads based on load size. However, the
102 following tests may provide a simple guide for any company wishing to deter-
103 mine if split loads should be considered based on the load sizes they generally
104 transport. Archetti et al. (2008) studied the effect of mean load size and vari-
105 ance on split loads for the VRP, finding results similar to those reported here.
106 Prior to an analysis of load size variance, the benefit for various load sizes
107 with no variance is presented with a more defined picture than that found
108 in Nowak et al. (2008). Determining the benefit without variance provides a
109 baseline indicating the exact load sizes for which the most and least benefit
110 may be found over the range from zero to one truckload.

111 Problem instances with 50, 100 and 150 load requests of a common size are
112 generated with load sizes incremented by 0.05 truckloads (TL) in the range
113 $[0.05 - 0.95]$. Additional problems are tested for those sizes where significant
114 changes in benefit may be expected. These load sizes include 0.11, 0.21, 0.26,
115 0.33, 0.34 and 0.51. Three different sets of location coordinates are randomly
116 generated to test each load size. Figure 1 displays the reduction in cost with
117 split loads for each of the load sizes tested, where the results are very similar
118 for each number of load requests. Almost all benefit from the use of split loads

119 is eliminated for the sizes 0.1, 0.2, 0.25, 0.33, and 0.5, or Q/k for $k = 2, \dots$,
 120 where $Q = 1$ is vehicle capacity. These load sizes can easily be combined on
 121 a vehicle with no splitting required. Peaks in cost reduction are found for the
 122 load sizes 0.11, 0.21, 0.26, 0.34, and 0.51. When splitting is allowed, these
 123 load sizes may have as little as 0.01 TL split off to allow for loads to be placed
 124 on a vehicle simultaneously. Although these results show the load sizes that
 125 provide the most (and least) benefit with split loads, it is rare to find a set
 126 of circumstances in the real world where all loads to be transported are of a
 127 common size. Therefore, it is of interest to determine the effect that load size
 variance has on this benefit.

Fig. 1. Average Percentage Cost Reduction with Split Loads for Load Sizes in the
 Range [0.05 – 0.95]

128

129 New problem instances are randomly generated with means ranging from 0.05
 130 to 0.95 TL and with variances from 0.005 to 0.08. A beta distribution is used
 131 to generate the load sizes as this distribution is defined on the interval [0,1]
 132 and all loads for the problems described here are of a size less than or equal to
 133 one truckload. This distribution is parameterized by two non-negative shape
 134 parameters, typically denoted by α and β , which are estimated using the
 135 method of moments with the following two equations:

$$136 \quad \alpha = x \left(\frac{x(1-x)}{v} - 1 \right), \tag{1}$$

$$137 \quad \beta = (1-x) \left(\frac{x(1-x)}{v} - 1 \right), \tag{2}$$

138 where x is the desired sample mean and v is the desired sample variance.
 139 The load sizes are then generated using the beta distribution function from

140 the GNU Scientific Library. Two different problem instances are randomly
141 generated with load sizes corresponding to each mean and variance. Two sets
142 of location coordinates are also randomly generated, such that four problem
143 instances are tested for each mean and variance.

144 Figure 2 provides the cost reduction for each variance, overlaid by the results
145 displayed in Figure 1 for which there is no variance. Any variance has an
146 immediate effect on the benefit of split loads. For those load size means below
147 one half of a truckload the peaks and minima are virtually eliminated. For the
148 variance of 0.005 there are two slight dips, one at 0.45 TL and the other at 0.2
149 - 0.25 TL. The variation is small enough such that the majority of load sizes
150 are still in the range [0.4-0.5] for the mean of 0.45 and [0.15-0.3] for the mean
151 of 0.2 or 0.25. Loads in these ranges are easily combined on a vehicle with
152 no splitting required. The other variances tested display an almost constant
153 percentage of cost reduction for each mean up to 0.5 TL. Load size means just
154 above one half of a truckload still result in a greater cost reduction, even with
155 a variance as high as 0.04. However, the peak in cost reduction diminishes as
156 the variance increases. A greater variance for a mean load size above one half
157 vehicle capacity results in problem sets with more load sizes below one half
158 vehicle capacity, allowing for more loads to be combined on a vehicle without
159 splitting and a reduction in benefit. Similarly, a greater variance for a mean
160 load size below one half vehicle capacity results in problem sets with more load
161 sizes above one half vehicle capacity and more splitting required to combine
162 loads on a vehicle, with an increase in benefit. This is further illustrated in
163 the following table.

164 Table 1 presents the average cost reduction over all load size means for each
165 variance. Although the problem instances with no variance result in the widest

Fig. 2. Average Percentage Cost Reduction with Split Loads for each Load Size Mean and Variance

166 range of values for cost reduction, the average reduction is not significantly
167 greater than those instances with some variance. Most real world problems will
168 have varying load sizes and these results indicate that there is some benefit
169 associated with almost any mean and variance combination. However, there
170 is a clear difference between loads with a mean size above and below one half
171 a truckload. Table 1 also separates the average percentage cost reduction to
172 show this distinction. As is evident in Figure 2, there is a greater benefit for
173 load sizes with a mean greater than half a truckload at variances up to 0.04.

174 Generalizing these results for any real world case is difficult, as most industries
175 can not be classified by the load sizes in which their goods are transported.
176 However, this is a very important factor in determining if splitting loads will
177 provide a significant benefit. When the majority of load sizes are clustered
178 around one half of a truckload, split loads should provide an opportunity for
179 cost savings. Other load sizes may result in a benefit, but it would most likely
180 be reduced. The load sizes used for the remaining problem instances in this
181 study fall in the range $[0.51 - 0.60]$, as loads of this size result in the most
182 opportunity for benefit from split loads. Because of this, changes in benefit
are most visible as other problem characteristics are altered.

Table 1

Average Percentage Cost Reduction with Split Loads for the Tested Load Size
Variances

183

184 4 Number of Origins and Destinations

185 When the Pickup and Delivery Problem has only one origin or one destination
186 it is reduced to the Vehicle Routing Problem. As described earlier, the benefits
187 associated with using split loads with the VRP have been quantified. Relaxing
188 the VRP to allow for multiple origins and destinations raises the question of
189 how that benefit is affected by the ratio of the number of origins to the number
190 of destinations. This should allow for a comparison between industries with
191 heavy inbound or outbound flow and those with a mixed flow. Industries with
192 heavy inbound flow, where a large number of materials or parts are used to
193 make few products (ie, auto industry), should have loads leaving from many
194 origins with a few common destinations, while heavy outbound flow, where few
195 materials make many products (ie, chemical industry), should be characterized
196 by loads leaving from a few common origins to many destinations.

197 To determine the effect of the number of origins relative to number of desti-
198 nations on the magnitude of benefit, problem instances with various ratios are
199 tested. To minimize variability between problems the ratios are selected such
200 that each problem requires the delivery of a similar number of loads, 50, 100 or
201 150. The ratios of the seven 50 load problem sets tested are: (number of origins
202 : number of destinations) 25 : 2, 12 : 4, 10 : 5, 7 : 7, 5 : 10, 4 : 12, and 2 : 25.
203 The ratios of the nine 100 load problem sets tested are: 50 : 2, 25 : 4, 20 : 5, 14 :
204 7, 10 : 10, 7 : 14, 5 : 20, 4 : 25, and 2 : 50. The ratios of the nine 150 load prob-
205 lem sets tested are: 75 : 2, 37 : 4, 30 : 5, 15 : 10, 12 : 12, 10 : 15, 5 : 30, 4 : 37,
206 and 2 : 75. Six different sets of location coordinates and five different sets
207 of load sizes are randomly generated for each ratio, resulting in 30 problem
208 instances for each number of loads.

209 Figure 3 presents the cost reduction for each of the ratios. All ratios result
210 in a cost reduction between 25 and 34%. The most benefit is found in the
211 instances that most closely represent the VRP, with either two origins or two
212 destinations. Benefit is reduced as the ratio of the number of origins to the
213 number of destinations approaches one. This is because opportunities for load
214 splitting grow as the number of loads departing from or arriving to a common
215 location increases. In the instance with two origins and 150 loads, when the
216 vehicle arrives for a pick up there are 75 different loads to select from to create
217 a combination of split loads. Dropping off loads at only two destinations has
218 similar opportunities. The instances with ratios of 37 : 4 and 4 : 37 have
219 38 fewer loads leaving from or arriving to any origin or destination, thereby
220 resulting in the largest decline in cost reduction. Less variance is found between
221 the other ratios as the change in the number of loads available at each origin
or destination is not as great.

Fig. 3. Average Percentage Cost Reduction with Split Loads for Problem Instances
with a Varying Number of Origins and Destinations

222

223 These results indicate that split loads would be most beneficial in a situation
224 where many loads are departing from or arriving to a common location. As
225 with the industry example described earlier, this indicates that the most ben-
226 efit would be found in the supply chain for production processes with heavy
227 inbound flow or heavy outbound flow. These supply chains have many loads
228 sharing common origins or destinations that provide for the most potential
229 split load combinations. For the remainder of this paper we will report results
230 for the ratios 5 : 10, 5 : 20, 10 : 10 and 10 : 15. The results for other ratios are
231 similar, with overall cost reduction slightly increased or decreased dependent

232 on the ratio.

233 **5 Origin-Destination Pairs Requiring Service**

234 Every origin-destination pair has a load requiring service in each problem in-
235 stance generated above. However, as shown by Nowak et al. (2008), a real
236 world problem instance will likely not have this characteristic and this may
237 have an effect on the benefit of split loads. To evaluate this effect, the percent-
238 age of origin-destination pairs requiring service is reduced. Several problem
239 instances are generated with a varying percentage of origins and destinations
240 between which a load must be delivered. Each instance has 50, 100 or 150
241 origin-destination pairs and the percentages of these pairs requiring service is
242 100%, 80%, 60%, 40%, and 20%. Nine different sets of load sizes and three sets
243 of location coordinates are randomly generated for each percentage, resulting
244 in 27 problem instances for each number of loads requiring service. All load
245 sizes are in the range $[0.51 - 0.60]$. Problem instances are generated for the
246 ratios 5 : 10, 5 : 20, 10 : 10 and 10 : 15.

247 Figure 4 presents the cost reduction for the various instances. As the percent-
248 age of origin-destination pairs requiring service decreases, the cost reduction
249 decreases as well. This can be attributed to a similar factor that caused the
250 change in benefit as the ratio of origins to destinations approaches one. As the
251 percentage of origin-destination pairs requiring service is reduced, each origin
252 or destination has fewer loads to select from when creating a combination to
253 place on a vehicle. This is most evident with the 5 : 10 ratio problem instances
254 with 20% of pairs requiring service. Each origin has only one to three loads de-
255 parting, while each destination has one load arriving. With fewer opportunities

256 to split and combine loads onto a vehicle, the potential benefit is diminished.
257 This is an indicator that in those real world situations with many isolated
258 locations that have a limited number of loads requiring service, splitting loads
has limited benefit.

Fig. 4. Average Percentage Cost Reduction with Split Loads when the Percentage
of Origin-Destination Pairs with a Load Requiring Service Varies

259

260 **6 Origin and Destination Location**

261 The effect of the location of the origins and destinations on the benefit of
262 split loads has not been tested. Nowak et al. (2008) used locations uniformly
263 generated over the test area for the random problem instances. In this section,
264 several different location configurations that correspond to real world scenarios
265 are tested.

266 One common scenario that occurs in the real world is that of origins clus-
267 tered separately from destinations. In the auto industry, parts suppliers are
268 closely located while production facilities are also clustered together. There
269 is not much movement within these two clusters, with most shipments mov-
270 ing between the clusters. To evaluate the change in benefit associated with
271 clustering, several different problem instances are generated.

272 Location coordinates are generated in three different configurations, A, B and
273 C. For Configuration A, the X and Y coordinates for the origin are both ran-
274 domly generated in the range $(0, 30)$ while the destination coordinates are
275 generated in the range $(-30, 0)$, such that the two clusters are separate but

276 adjacent. The origin and destination clusters are spaced further apart in Con-
277 figurations B and C, where they are separated by a minimum of 30 and 60
278 units, respectively. Six sets of location coordinates and five sets of load sizes
279 are randomly generated for each configuration, resulting in 30 problem in-
280 stances. All load sizes are in the range $[0.51 - 0.60]$. Problem instances are
281 generated for the ratios 5 : 10, 5 : 20, 10 : 10 and 10 : 15.

282 Figure 5 presents the cost reduction for both the random and clustered prob-
283 lem instances. Clustered origins and destinations result in a significant increase
284 in cost reduction over randomly located origins and destinations. Splitting
285 loads leads to more trips from origin to origin or destination to destination
286 and fewer trips from an origin to a destination, as the vehicle picks up smaller
287 loads from several origins rather than picking up a large unsplit load from
288 one origin that is immediately transported to its destination. Clustering ori-
289 gins and destinations separately increases the average distance between origins
290 and destinations relative to the average distance between origins or between
291 destinations. Because splitting loads reduces the number of trips from ori-
292 gins to destinations, clustering leads to an increase in the potential benefit
293 from splitting. As seen with Configurations B and C, lengthening the dis-
294 tance between the clusters increases the average distance from an origin to a
destination, further increasing the potential benefit.

Fig. 5. Average Percentage Cost Reduction with Split Loads when Origins and
Destinations are Clustered

295

296 Another scenario tested separates the locations into three geographically di-
297 vided clusters. Each cluster consists of several origins and destinations. Loads
298 are delivered primarily within each cluster, with a few loads delivered between

299 clusters. These problem instances are similar to a real world scenario in which
300 loads are transported within several regions, with very few delivered between
301 regions, such as with a retail distribution network.

302 Three configurations are tested, each with a different number of origins and
303 destinations in the three clusters. These configurations are described in Ta-
304 ble 2. Problem instances are also generated with different restrictions on the
305 number of loads that required service between clusters: instances with no
306 loads requiring service between clusters, instances with approximately 50% of
307 all loads requiring service to be delivered between clusters, and instances with
308 up to 85% of all loads that may be delivered between clusters. Six sets of loca-
309 tion coordinates and three sets of load sizes are randomly generated for each
310 configuration and level of allowable inter-cluster load movement, resulting in
18 problem instances. All load sizes are in the range $[0.51 - 0.60]$.

Table 2

Number of Origins \times Number of Destinations for each Cluster within each Config-
uration

311

312 Table 3 presents the average percentage reduction in cost for each configu-
313 ration and level of allowable inter-cluster load movement. As the number of
314 inter-cluster moves increases, so does the reduction in cost. Just as with the
315 problem instances presented above, where the origins and destinations are
316 clustered separately, this is a result of an increase in the average distance that
317 must be traveled to deliver loads between origins and destinations. Without
318 inter-cluster moves the delivery of all loads occurs within the limited bound-
319 aries of a cluster. The average distance from an origin to a destination is the
320 same as the average distance between two origins or between two destinations.

321 As more inter-cluster moves are made, the average distance traveled by the
322 vehicle from an origin to a destination increases. Allowing split loads results
323 in a decrease in the number of moves between origins and destinations relative
to the number of moves between origins or between destinations.

Table 3

Average Percentage Cost Reduction with Split Loads for each Configuration with
Various Restrictions on Moves Between Clusters

324

325 This result further underlines the usefulness of split loads when loads must
326 be delivered over longer distances. Although benefit was found for instances
327 where loads were only transported within the clusters, the cost reduction was
328 markedly greater when loads were also delivered over the longer distances
329 between clusters.

330 Altering the number of origins relative to the number of destinations per
331 cluster also had an effect. Configuration A, which had an equal number of
332 origins and destinations for each cluster, showed the least amount of cost
333 reduction. As with the results found in Section 4, this configuration afforded
334 the least opportunity to generate multiple split load combinations at each
335 origin. When the number of origins and number of destinations in a cluster
336 were not equivalent, the cost reduction increased.

337 **7 Conclusions**

338 The benefit associated with split loads varies considerably with most problem
339 characteristics including load size, number of loads, and the configuration of
340 origins and destinations. By testing various problem instances, we have found

341 three primary factors that affect the benefit:

342 (1) Although some benefit was found with almost any mean load size and
343 variance, those loads larger than one half of vehicle capacity showed the
344 most potential, even with greater variances. These loads can not be com-
345 bined without splitting, while they are the easiest to combine on a vehicle
346 with splitting.

347 (2) As the number of loads available at a common location for pickup or
348 delivery increases, so does the potential benefit from split loads. This is
349 due to the increase in potential load combinations to split at each stop.
350 This was shown by changing the ratio of the number of origins relative
351 to the number of destinations, where the benefit decreased as the ratio
352 approached one, and by decreasing the percentage of origin-destination
353 pairs with a load requiring service, where the benefit decreased with the
354 percentage.

355 (3) Increasing the average distance from an origin to a destination relative
356 to the distance from origin to origin and destination to destination has a
357 positive effect on the benefit of split loads. Because splitting loads reduces
358 the number of trips from origins to destinations, clustering leads to an
359 increase in the potential benefit from splitting. Both clustered scenarios
360 supported this result, where clustering origins separately from destina-
361 tions increased the reduction in cost from split loads as the clusters were
362 spaced farther apart, while limiting the number of loads that could travel
363 between the three separated clusters decreased the cost reduction.

364 **References**

- 365 Archetti, C., Hertz, A., Speranza, M., 2006. A tabu search algorithm for the
366 split delivery vehicle routing problem. *Transportation Science* 40, 64–73.
- 367 Archetti, C., Savelsbergh, M. W. P., Speranza, M., 2008. To split or not to
368 split: That is the question, *Transportation Research Part E*. Forthcoming.
- 369 Dror, M., Laporte, G., Trudeau, P., 1994. Vehicle routing with split deliveries.
370 *Discrete Applied Mathematics* 50, 239–254.
- 371 Dror, M., Trudeau, P., 1989. Savings by split delivery routing. *Transportation*
372 *Science* 23, 141–145.
- 373 Dror, M., Trudeau, P., 1990. Split delivery routing. *Naval Research Logistics*
374 37, 383–402.
- 375 Frizzell, P. W., Giffin, J. W., 1995. The split delivery vehicle scheduling prob-
376 lem with time windows and grid network distance. *Computers and Opera-*
377 *tions Research* 22, 655–667.
- 378 Nowak, M. A., Ergun, O., White, III, C. C., 2008. Pickup and delivery with
379 split loads, *Transportation Science*. Forthcoming.
- 380 Sierksma, G., Tijssen, G. A., 1998. Routing helicopters for crew exchanges on
381 off-shore locations. *Annals of Operations Research* 76, 261–286.