A Study of Shape Penalties in Vehicle Routing

Charles Gretton and Philip Kilby
Optimisation Research Group, NICTA and A.N.U., Australia

1 Introduction

Following recent studies of *visual attractiveness* in vehicle routing (Hollis and Green, 2012; Kim et al., 2008; Poot et al., 2002; Tang and Miller-Hooks, 2006; Zhou et al., 2006), we investigate the inclusion of shape and compactness penalties in computing solutions to the Vehicle Routing Problem (VRP) using the Adaptive Large Neighbourhood Search (ALNS) proposed by Ropke and Pisinger (2006). Visually attractive routes are sought predominantly for two reasons. First, our clients are reluctant to implement solutions that exhibit overlapping routes, or unacceptable shape. Second, as Hollis and Green (2012) have lately observed, the visual compactness of routes is indicative of the operational robustness of solutions. We are the first to investigate the concept of *bending energy* from Young et al. (1974) as a solution penalty in this setting. We are also the first to investigate a search that leverages the geographic center of every route encountered during search.

2 Penalties of Unattractiveness

We propose two penalties that relate to compact and nicely shaped routes. For compactness we use a penalty similar to *Measure Two* developed by Tang and Miller-Hooks (2006). Intuitively, a route has a centre of gravity and an optimal solution will minimise the distances between that centre and its visits. Optimal solutions will typically be compact and exhibit little inter-route overlap. Formally, where i and j denote visits, let $d(i, j)$ be the symmetric distance between i and j. Given a route R, a visit i is a good candidate for the geographical centre of R if it has a relatively low score $V(i, R) = \sum_{j \in R} d(i, j)$. The *median* of route R is a visit $C_R \in R$ with the lowest V score, in other words $C_R = \arg\min_{i \in R} V(i, R)$.

Tang and Miller-Hooks (2006) studied conditions of zero overlap between the convex hulls induced by routes in the euclidean case, giving a proof that if all visits are closest to the median of the route they participate in, then the hulls induced by those routes do not overlap. Here, we propose a solution penalty equal to the **sum of distances of assigned visits to their route medians**. This is a slight departure from Tang and Miller-Hooks (2006), who optimise a normalised sum and treat routes individually. In our

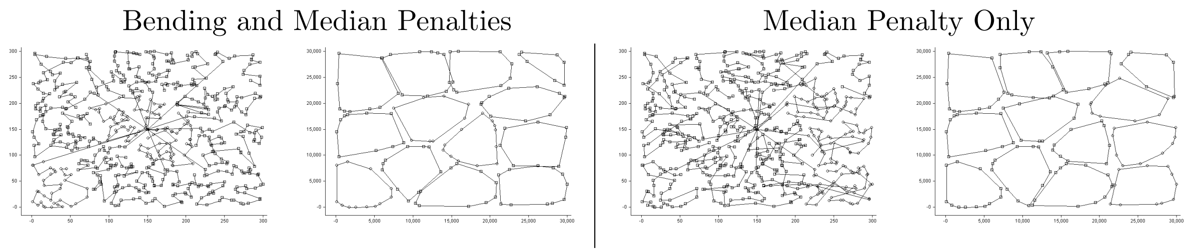

Figure 1: Depicts solutions to an instance with 600 requests distributed uniformly. Under each heading routes are depicted on the left, and the convex hulls induced by those routes on the right.

setting normalisation can be problematic if vehicles have different capacities, and in early iterations of search when incumbent routes comprise a significantly different number of visits.

Our second penalty requires that solutions minimise the total *bending energy* of routes. That quantity measures the amount of work required to create a shape using a linear thin-shelled medium. It was proposed by Young et al. (1974) to measure shape equivalence in computer vision, and has since found myriad application in that setting. A non-trivial contour with minimum bending energy is a circle. To measure the energy of a polygonal shape, one might induce a regular curve familiar to differential geometry. A more recent and satisfactory definition for the bending energy of polygons, called *visual curvature*, was proposed by Liu et al. (2008). Following those authors, we can adopt a fairly straightforward scale invariant penalty of the bending energy of a route, equal to the **sum of turn angles**. Solutions with routes having low bending energy therefore have non-jagged transitions between visits, and are visually attractive.

3 Experimental Observations

We experiment using instances of the VRP with time windows by Solomon and Desrosiers (1988), and extended Solomon instances by Gehring and Homberger (1999) with 100 to 1000 customers. Our evaluation also includes some real-world instances. We use Indigo (Kilby and Verden, 2011), a solver that implements an ALNS. Routes are built by inserting visits one at a time. When considering an insertion, penalties (adjusted by a multiplier for each penalty type) are added to the objective. For example, *minimum insertion cost* includes the penalty terms, and *regret* is the difference between best and second best insertion routes, where each cost includes the penalties that would be associated with each insertion. All experimental results reported in this abstract were obtained using 10^4 iterations of ALNS.

Our experiments yield a number of broad observations. Using the median penalty only, the area intersected by convex hulls induced by routes is reduced by 43% on average, relative to solutions found without shape penalties. Here, total solution distance increases by 12% on average, and only by 9% when we also use the bending energy penalty. In that

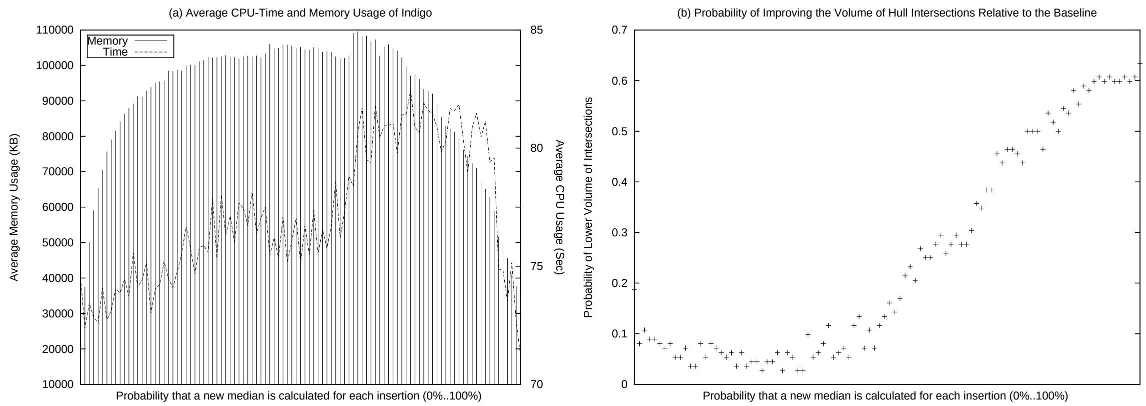

Figure 2: For different values of θ we measure: (a) average memory and CPU usage, and (b) the probability that the area of intersecting convex hulls induced by solution routes is less relative to a *baseline* solution computed without bending or median penalties.

case the good intersection statistic reported above changes marginally; Decreasing by at most 2 percentage points. We have illustrated in Figure 1 how optimising additionally for bending energy improves solution attractiveness.

Calculating route medians is computationally expensive, and each time a visit is inserted in a route we should calculate the penalty according to the (possibly) new median. We investigated the effect of varying how often route medians are actually calculated: 0% means the median remains unchanged from the original seed, through to 100% where the true median is calculated for each insertion. In between, we have a probability θ that the penalty is computed according to the true median, rather than according to the previous median computed for the route.

Our implementation caches partial computations of route medians for runtime efficiency. The size of that cache increases in proportion to the diversity of routes for which medians are computed during search. With that in mind, we can interpret Figure 2a as highlighting the degree to which search is focused as θ approaches either 0 or 100. The computational burden of search is low as θ approaches 0 because very few median calculations are performed. Here, the search is focused on compact routes around *seed* visits. As θ approaches 100 that burden decreases because the search quickly becomes focused on a relatively small set of compact solutions. Here, due to caching very few expensive median calculations are performed. Figure 2b shows that as θ increases the guidance associated with the median penalty is likely to yield more compact solutions.

4 Conclusions

We introduce the idea of using bending energy to guide attractive-route construction. We also use the notion of a route median—a visit that is geographically central to the route—

to penalise a route according to how distant the visits are from the median. Together those penalties guide ALNS to visually attractive solutions. We have some empirical evidence that bending energy helps identify optimised and attractive routes early in search. We also investigate the effect of changing the frequency with which true (versus inherited) route medians are used in evaluating penalties during search. We have seen empirically that increased frequency leads to more attractive solutions which feature compact routes with comparatively little inter-route overlap.

References

H. Gehring and J. Homberger. A parallel hybrid evolutionary metaheuristic for the vehicle routing problem with time windows. Short Course on Evolutionary Algorithms in Engineering and Computer Science, 1999.

B. Hollis and P. Green. Real-Life Vehicle Routing with Time Windows for Visual Attractiveness and Operational Robustness. *Asia-Pacific Journal of Operational Research*, 29(4):1–29, 2012.

P. Kilby and A. Verden. Flexible routing combing constraint programming, large neighbourhood search, and feature-based insertion. Paper presented at the 2nd Workshop on Artificial Intelligence and Logistics, 2011.

B.I. Kim, S. Kim, and S. Sahoo. Waste collection vehicle routing problem with time windows using multi-objective genetic algorithms. Paper presented at the 3rd IASTED International Conference on Computational Intelligence, 2007.

H. Liu, L. Latecki, and W. Liu. A Unified Curvature Definition for Regular, Polygonal, and Digital Planar Curves. *International Journal of Computer Vision*, 80(1):104–124, 2008.

A. Poot, G. Kant, and A. Wagelmans. A Savings Based Method for Real-Life Vehicle Routing Problems. *Journal of the Operational Research Society*, 53(1):57–68, 2002.

S. Ropke and D. Pisinger. An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows. *Transportation Science*, 40(4):455–472, 2006.

M. Solomon and J. Desrosiers. Time window constrained routing and scheduling problems. *Transportation Science*, 22(1):1–12, 1988.

H. Tang and E. Miller-Hooks. Interactive Heuristic for Practical Vehicle Routing Problem with Solution Shape Constraints. *Transportation Research record: Journal of the Transportation Research Board*, 1964:9-18, 2006.

I. Young, J. Walker and J. Bowie. An Analysis Technique for Biological Shape. I. *Journal of Information and Control*, 25(4):357–370, 1974.

C. Zhou, Y. Tan, L. Liao and Y. Liu. Solving the Multi-vehicle Pick-up and Delivery Problem with Time Widows by New Construction Heuristic. Paper presented at the 6th International Conference on Intelligent Systems Design and Applications, 2006.