[PD-01-001] Linear Programming and GIS

Abstract

Linear programming is a set of methods for finding optimal solutions to mathematical models composed of a set of linear functions. Many spatial location problems can be structured as linear programs. However, even modest-sized problem instances can be very difficult to solve due to the combinatorial complexity of the problems and the associated computational expense that they incur. Geographic Information Systems software does not typically incorporate formal linear programming functionality, and instead commonly uses heuristic solution procedures to generate near-optimal solutions quickly. There is growing interest in integrating the spatial analytic tools incorporated in Geographic Information Systems with the solution power of linear programming software to generate guaranteed optimal solutions to spatial location problems.

Keywords: combinatorial compexity, decision making, flows and networks, linear programming, location allocation modelling, logistics, mathematical modeling, network analysis, optimization, optimization and location-allocation modeling, transportation

Author & citation

Curtin, K. and Turner, D. (2019). Linear Programming and GIS. The Geographic Information Science & Technology Body of Knowledge (4th Quarter 2019 Edition), John P. Wilson (ed.). DOI: 10.22224/gistbok/2019.4.1

This Topic is also available in the following editions:

DiBiase, D., DeMers, M., Johnson, A., Kemp, K., Luck, A. T., Plewe, B., and Wentz, E. (2006). Linear Programming. The Geographic Information Science & Technology Body of Knowledge. Washington, DC: Association of American Geographers.

Explanation

1. Definitions
2. What is Linear Programming?
3. Solving Optimization Problems in GIS
4. Integrating Linear Programming and GIS
5. Future Developments

1. Definitions

Optimization: the calculation of the maximum or minimum output of a mathematical function.

Linear function: a polynomial function of degree at most 1.

Combinatorial Complexity: a property of a system in which an ordering or selection of combinations of inter-related problem elements or decision variable values is desired; when

there are a large number of such value combinations, a problem instance is said to be highly combinatorially complex.

Heuristic: a procedure for calculating an approximate solution to a problem that otherwise might be intractable or prohibitively expensive to solve optimally.

Linear programming: a structured method for optimizing the output of a linear objective function where the solution to that function is constrained by a set of bounding linear functions.

Solver: a software application or library designed to solve complex mathematical problems like those found in linear programming.

2. What is Linear Programming?

Linear programming (also known as LP or linear optimization) is the process used to solve a linear program, a mathematical model for finding optimal solutions to problems composed of a set of linear functions. A linear program consists of a linear objective function, which represents the goal that one wishes to optimize, and a set of bounding linear constraint functions that define the limits of the optimization function's feasible solution space. In a spatial context, the goal of the objective function could be minimizing cost or distance travelled, maximizing the spatial coverage of a population, or finding the shortest route among a set of stops, among many other possible objectives. Example constraints might include a limit on the number of facilities to locate, a maximum distance a person would need to travel to a service facility to be considered spatially covered, or a capacity on the number of customers a facility could serve. The values of the decision variables defined in these linear functions are calculated via LP to produce an optimized objective function.

Some of the classic examples of LP problems are spatial (Church & Sorenson 1996, Tong & Murray 2012). The traveling salesman problem (TSP), for example, posits a salesman who must visit a set of cities on a network and return home while taking the shortest possible path. As an illustration of the complexity of problems like the TSP, consider that the number of permutations in the potential combinatorial solution set approaches the factorial of the number of cities, (n-1)!. A TSP problem with 11 cities has 3.6 million potential solution permutations. A TSP problem with 21 cities has 2.4 * 1017. It is very difficult – if not impossible in most computational settings – to evaluate each of those possible paths for the salesman in order to determine which provides the shortest path. Therefore, more effective solution procedures must be implemented.

Solution procedures for problems formulated as LPs include both optimal and heuristic procedures. Many optimal LP solution procedures are derived from the Simplex algorithm (Dantzig 1948). In principle, the Simplex algorithm leverages the fact that the potential solution space for many LP problems can be represented by a convex, multidimensional polytope bounded by the constraint functions. Simplex operates by inspecting only the corner point solutions of this convex hull rather than searching the entire feasible region and all the combinations of decision variable values that it contains. There are many variants of, and extensions to, the Simplex Method tailored to particular problem structures and sizes. These methods are typically instantiated in solution software that is designed specifically to solve optimization problems (see a discussion of some of these below)

(Gearhart et al. 2013). However, even robust LP solution software cannot match the demands of highly combinatorially complex problems. In these cases, the problems can be addressed with heuristic solution procedures.

Heuristic solution procedures are typically much faster than optimal methods, but are not guaranteed to produce an optimal solution. These methods are used when the computational or time constraints for solving a problem make it necessary. One risk with heuristics, however, is that there is no way of knowing how close a solution generated by a heuristic solution procedure is to optimal. Typically, heuristics are evaluated by comparing the solutions they provide against optimal solutions for problem instances that are small enough to solve optimally with another method. The assumption is that if the heuristic performs well compared to optimal on small problems, it will likely perform well on larger problem instances where the optimal solution is unknowable. For larger problem instances, different heuristics can be compared against each other to see which provides better solutions consistently. Note that the same procedure may guarantee an optimal solution for one type of problem but not for another, where it is a heuristic. An example is a greedy procedure which produces optimal solutions for the minimal spanning tree problem, but typically fails to find optimal solutions for the Traveling Salesman Problem (and many others).

There are a wide range of heuristics that can be brought to bear on spatial optimization problems (Lei et al. 2016). These include interchange algorithms, genetic algorithms, simulated annealing, and others. One widely used family of heuristics is known as Tabu search (Glover 1989). This method is based on an interchange search – that is, an initial feasible solution (one which does not violate any constraints) is determined, and then individual decisions (e.g. a facility location choice) is exchanged (or swapped) for a decision that was not in that initial solution. If the exchange leads to an improved objective function value then the swap is accepted and additional interchanges are explored. Interchange methods continue until no better solution can be determined. Because interchange methods are prone to becoming trapped in local optima, Tabu search was developed to allow the heuristic more flexibility to search the solution space. Essentially, Tabu search allows an exchange to occur even if the objective function value is worsened, on the chance that a solution closer to the global optimum can be subsequently found.

Given this review, it is clear that there is a significant set of spatial optimization problems, and that these problems can be extremely difficult to solve optimally. It is not surprising that geographic information scientists would seek to use their toolsets to address these difficult yet compelling spatial problems.

3. Solving Optimization Problems in GIS

Geographic Information Systems typically have very limited capacity for solving spatial optimization problems with LP. Likewise, LP solvers and other related specialized software have evolved largely as entirely separate software packages from GIS software. However, GIS does support graph and topological data structures, which are historically ripe mediums for spatial optimization analysis. Therefore, there are commonalities between the research areas, yet significant differences in capabilities.

The optimization problems that are most frequently addressed in GIS tend to operate on

network data models. The specific objectives often concern minimal cost routing, including the shortest path problem, the TSP, determining network distance origin-destination matrices, and similar network optimization problems. Although the developers of commercial GIS software do not always reveal their proprietary solution procedures, it is well known that some form of Tabu search is a common approach to solving these problems in the GIS context. Tests of GIS solution procedures on problem instances where the optimal solution is known have shown the level to which the GIS can produce sub-optimal solutions (Curtin et al. 2014).

For other facility location and allocation problems, the network distances are an input to the optimal or near-optimal location of facilities where the objective is covering (set covering or maximal covering) or minimizing the distance to serve a demand (Hakimi 1964, Church & ReVelle 1974). These transportation related resource allocation tools can be used to – for example – place fire stations in such a way that the population covered within an acceptable service time is maximized. The problem of minimizing distance to demands is widely termed the P-median problem. More explicitly, the P-median seeks to locate a user-determined number of facilities (P) on a network, such that the total demand weighted distance from demand locations to their nearest facility is minimized. Figure 1 shows an example of a P-Median solution generated through a heuristic solution procedure on a network in a GIS.

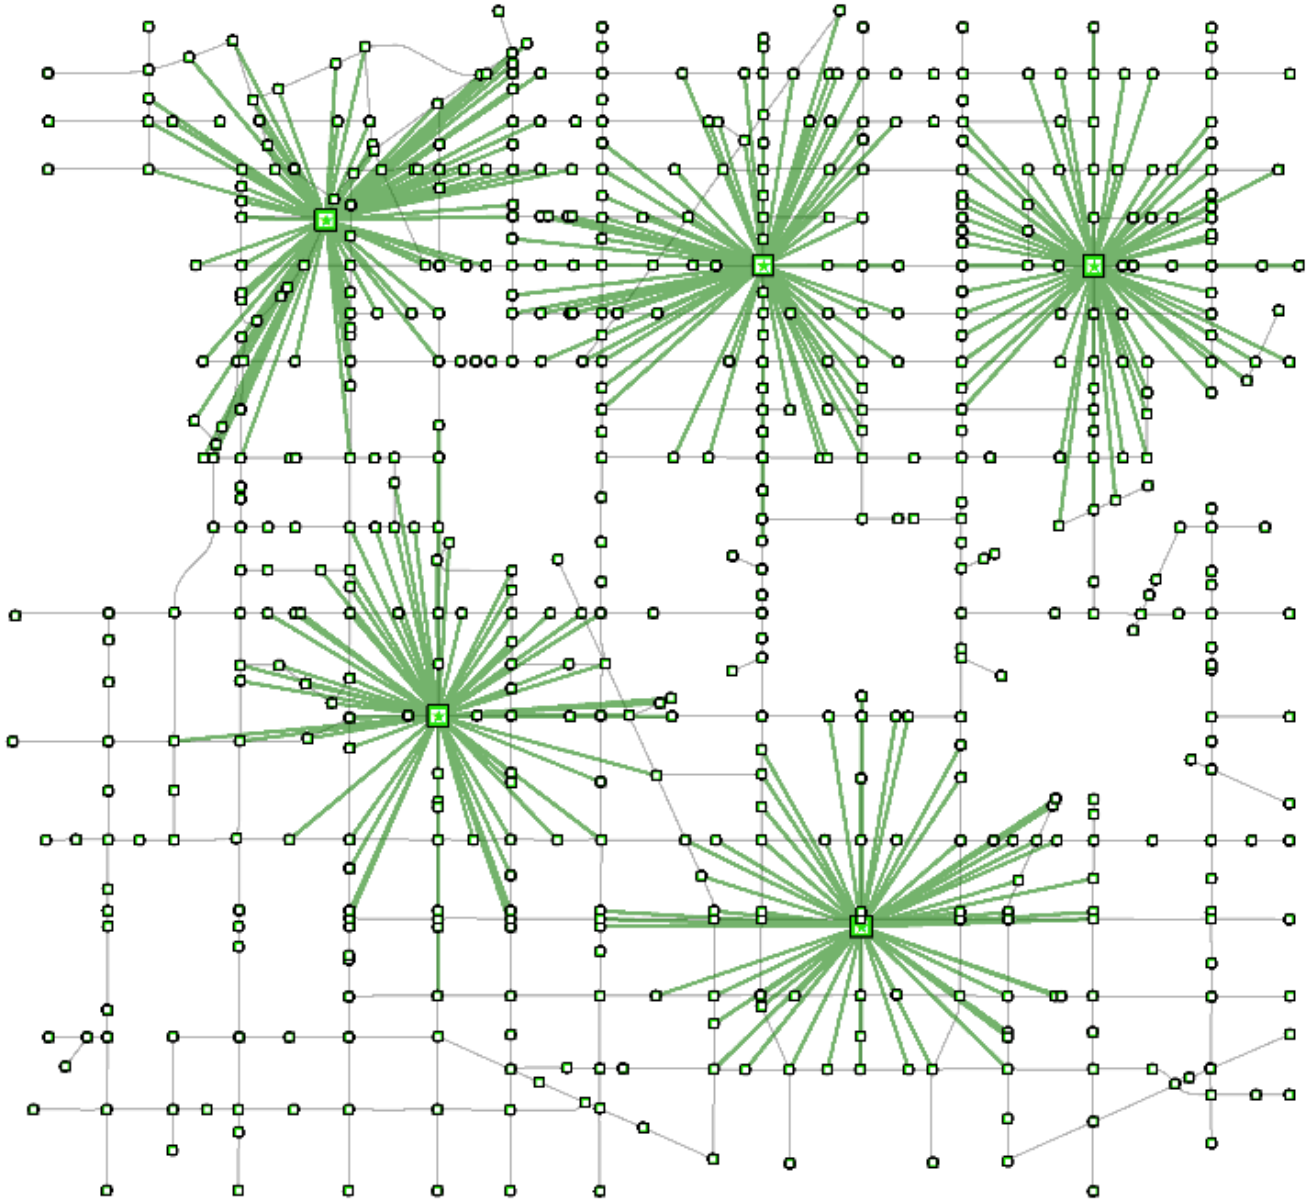


Figure 1: Five p-median site selections that maximally cover nodes given a defined network distance. Note that while network distances are used in the solution, green lines are used to graphically indicate the allocations. Source: authors.

Not all spatial optimization problems employing linear programming involve graph or network structures, however. In particular, raster based (grid) data structures can also be employed in the context of spatial optimization (Cova & Church 2000, Ligmann-Zielinska et al. 2008).

The greatest limitation currently in the ability to employ GIS for a wider variety of spatial optimization problems is that the solution procedures are designed to address a single problem. There is a rich literature in spatial optimization that contains many thousands of interesting problems and problem variants. New models are being developed and published all of the time. In contrast, there are only a handful of problems that can be addressed by off-the-shelf GIS software. While the problems that have been implemented are among the most important and have proven to be extremely useful in a broad range of application

areas, the inability to customize problems is a severe limitation. For this reason, there is increasing interest in the integration of GIS software with LP modeling and solution software.

4. Integrating Linear Programming and GIS

With the increasing acceptance of modular, decoupled workflow in software design, there is a greater expectation that software developed independently should be amenable to integration. That is, they should be able to accept inputs from, produce output for, and become more tightly integrated with each other. The idea is that integrating disparate sets of methods instantiated in software can allow researchers to approach and solve problems that neither set of methods could address in isolation. This is the motivation behind integrating GIS with LP solution software.

An integrated GIS/LP system can take advantage of the capabilities of each knowledge area. Typical GIS environments provide underlying database functionality that allows the input, storage, and manipulation of both spatial and aspatial data. Perhaps most importantly, there are a range of spatial analytic functions that can produce the necessary inputs for a spatial optimization problem (e.g. demand values from census geography, origin-destination matrices from transportation networks). These inputs to a model can be transferred to LP solution software in a loosely or tightly coupled system. At one end of that range, output files can be generated from the GIS and saved in a format that can be read by the LP solver, then subsequently submitted to the solver in a separate process. Alternatively, there are methods of tying the two software platforms together with so that the inputs are transferred directly from the GIS to the LP solver. Once that transfer is made, typical LP solution software has functionality that permits the formulation of any virtually any spatial optimization problem that can be structured as a linear program. Many of the commonly used solvers have a variety of pre-processing and solution procedures. Some of these can take advantage of either the structure of the problem itself or the nature of the data for a problem instance to reduce the problem size and increase the likelihood of determining an optimal solution. Finally, in a GIS/LP system the variable values can be transferred from the LP solver back to the GIS to take advantage of its cartographic capabilities, which allow a spatial inspection of the solution. Figure 2 shows an example of an optimal logistics delivery plan that was derived from LP solver, and which can be visualized and queried in the GIS environment.

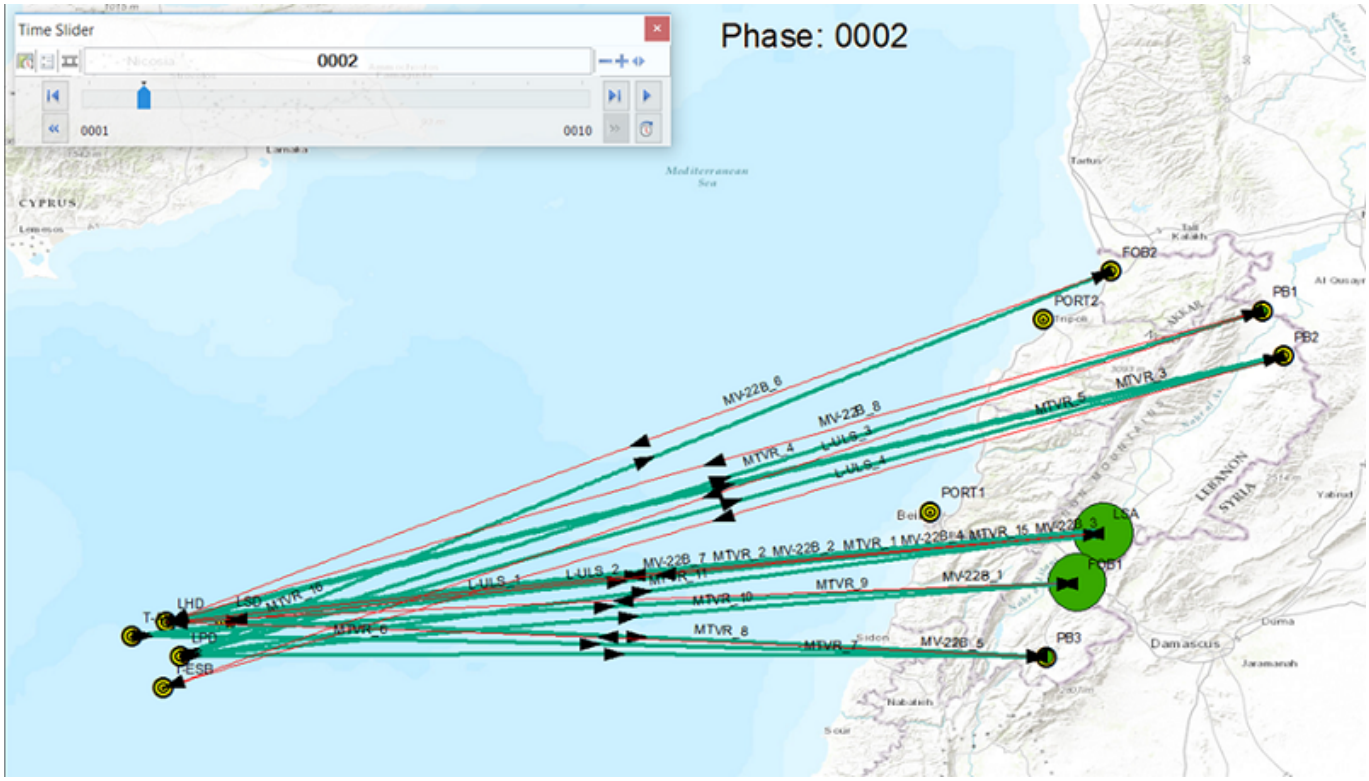


Figure 2. GIS display of an optimal logistics delivery solution, from ships at sea to ports and inland facilities. Source: authors.

Once this transition from GIS to LP and back to GIS is complete, the true advantage of the integration becomes apparent. The GIS could not have generated the solution without the LP. The LP could not generate a cartographic display or provide the environment for spatial query. Most importantly, however, the GIS can now be used to perform additional spatial analysis of the solution (e.g. clustering of chosen facilities or analysis of unmet demand characteristics). These analyses can be used to modify the linear program, driving forward the research process into new areas of inquiry.

5. Future Developments

It is likely that advances in the solution of LPs in the context of GIS will proceed along several lines. It may be that GIS applications will add functionality that either can compute the solution of a broader range of LP problems or incorporate a greater variety of solution procedures (either optimal, heuristic, or both). Conversely, LP solution software may adopt spatial analytic tools and techniques for users addressing problems with a spatial nature. It is probable that advances will also be made in the integration of the systems and the ease of translation and transition from one system to the other. It may be that techniques are developed to help a user determine if a problem is tractable for a particular solution procedure, and based on this assessment the user will be able to make informed decisions about the best solution process. Finally, there are major ongoing developments in cloud computing and high-performance computing that have the capability to dramatically expand the size of problems that can be addressed. All of these advances will influence the ways in which GIS and linear programming work together.

References

Church, R. L., & ReVelle, C. S. (1974). The maximal covering location problem. Papers of the Regional Science Association, 32, 101–118.

Church, R. L., & Sorensen, P. (1996). Integrating normative location models into GIS: problems and prospects with the p-median model. In P. Longley & M. Batty (Eds.), Spatial Analysis: Modelling in a GIS Environment (p. 400). New York: Wiley.

Cova, T. J., & Church, R. L. (2000). Contiguity Constraints for Single-Region Site Search Problems. Geographical Analysis, 32(4), 306–329.

Curtin, K.M., G. Voicu, M.T. Rice, and A. Stefanidis (2014). A Comparative Analysis of Traveling Salesman Solutions from Geographic Information Systems. Transactions in GIS, 18 (2), 286-301.

Dantzig, G. (1948). Programming in a Linear Structure, U.S. Air Force Comptroller, USAF, Washington, D.C.

Gearhart, J. L., Adair, K. L., Durfee, J. D., Jones, K. A., Martin, N., & Detry, R. J. (2013). Comparison of open-source linear programming solvers. (No. SAND2013-8847). Sandia National Lab. (SNL-NM), Albuquerque, NM (United States).

Glover, F. (1989). Tabu Search – Part 1. Journal on Computing, 1(2), 190–206.

Hakimi, S. L. (1964). Optimum location of switching centers and the absolute centers and medians of a graph. Operations Research 12, 450–459.

Lei, T. L., Church, R. L., & Lei, Z. (2016). A unified approach for location-allocation analysis: integrating GIS, distributed computing and spatial optimization. International Journal of Geographical Information Science, 30(3), 515–534.

Ligmann-Zielinska, A., Church, R. L., & Jankowski, P. (2008). Spatial optimization as a generative technique for sustainable multiobjective land-use allocation. International Journal of Geographical Information Science, 22(6), 601–622.

Tong, D., & Murray, A. T. (2012). Spatial Optimization in Geography. Annals of the Association of American Geographers, 102(6), 1290–1309.