

Combinatorial optimization problems typically require every possible solution to be evaluated to ensure finding the optimal solution. Since such exhaustive searches are often impractical, there is now a vast body of heuristic algorithms for them. Among the algorithms are those based on metaphors borrowed from other areas of science. The idea is that key elements of physical processes can be used abstractly to form the basis of an optimization algorithm. This article presents a broad overview of several metaphor-based algorithms, including the widely-used genetic and simulated annealing algorithms.

Possibly the most famous combinatorial optimization problem is the *Traveling Salesman Problem*: Given a set of points, either nodes on a graph or cities on a map, the objective is to find the shortest possible tour that visits every single point exactly once and then returns to its starting point. Simple heuristics for choosing a tour, such as iteratively picking the closest city to the current one, do not necessarily yield optimal tours. At the same time, an exhaustive search of all the possible tours to find the best one can be computationally impractical.

There are $(n-1)!/2$ possible tours for any n -city traveling-salesman problem. Thus, while there are only 2520 possible solutions (tours) for a small 8-city problem, there are 653,837,184,000 possible solutions for a 16-city problem. Even if a computer could evaluate a single tour in one-millionth of a second, it would require almost eight straight days to evaluate all of the possible solutions to a 16-city problem. Furthermore, this extensive computation would be of little use in another instance (with different data) of the same problem.

Another example of a combinatorial optimization problem is the *Location-Allocation Problem*. Given a set of facilities, each of which can serve a certain number of nodes on a graph, the objective is to place the facilities on the graph so that the average distance between each node and the facility that serves it is minimized. That is, where should the police department build stations (of possibly different sizes) so that the residents are as close as possible to the stations that serve them? Figure 1 (pg 22) shows an example of a location-allocation problem and its solution. The square nodes represent facilities serving nodes in their respective regions.

Strictly speaking, any problem that has a large set of discrete solutions and a cost function that rates those solutions relative to one another is a combinatorial optimization problem. In combinatorial optimization problems, the number of solutions grows exponentially as the size of the problem grows. This is what makes them so challenging.

For traveling salesman problems, for example, the size of the problem is the number of cities to visit. The solution set is the set of all possible tours through those cities. The cost function for a traveling salesman problem is the length of the tour.

Meanwhile, for location-allocation problems, the size of the problem is affected by both the number of nodes in the graph and the number of facilities. The solution set is the number of possible ways that the facilities can be placed on the graph crossed with the number of possible ways that the nodes can be assigned to facilities. The cost function is the average distance from each node to its corresponding facility.

The engineer's job is to find algorithms that can reliably produce results that are close to optimal in a reasonable amount of time. We will look at four successful attempts to apply metaphors from other sciences to develop algorithms for combinatorial optimization problems. We will also look at less successful results using the same approach.

Evolution and the Genetic Algorithm

As their name suggests, genetic algorithms attempt to mimic the process of biological evolution in developing a solution. The algorithm begins by randomly generating a reasonably large set of candidate solutions (the initial "population"). Then, each solution is rated for its "fitness" according to some metric that is appropriate for the problem.

For example, an attempt to solve a traveling salesman problem would start by generating a set of possible tours and then calculating each tour's length. The shorter the tour length, the higher the fitness of the solution.

After all of the solutions' fitnesses are calculated, a new set of solutions is produced by randomly picking solutions (with some replication, to keep the population more or less constant) out of the original set with a preference for the solutions with higher fitness. This is the "natural selection" of biolo-

Nature's algorithms

Social and natural metaphors
in algorithm design

Joseph Carnahan
and
Rahul Simha

© Photo Disc

gy, weeding out the weaker members of the population of solutions.

Next, members of the new set of solutions are paired. Parts of their solutions are swapped within each pair, followed by an occasional “mutation” to change a solution slightly. In theory, this swapping reflects the way that DNA is combined in sexual reproduction. Then, this pattern of evaluation, reproduction, swapping and mutation is repeated until a sufficiently satisfactory (low-cost) solution appears to dominate the rest of the population.

Genetic algorithms have been shown to perform fairly well in a wide variety of problems. The largest difficulties arise when there is no obvious way to “mate” solutions. Also, since genetic algorithms are a general-purpose problem solver, they can often be outperformed by more specialized algorithms designed for that particular problem.

Metallurgy & the simulated annealing algorithm

Simulated annealing is another general-purpose optimization problem solver. Its metaphor comes from metallurgy instead of biology. When molten metal cools too quickly, many desirable qualities, such as strength or magnetism, are lost or degraded. In contrast, when molten metal is cooled slowly by a process called annealing, these properties are retained or enhanced. This idea of slow cooling has been applied in the abstract to solving combinatorial




Fig. 1 A solution to an arbitrary location-allocation problem

Box A

```
S = initial_state; // The current "working" solution
minS = S; // The best solution so far
T = initial-temperature;
while (T > final-temperature)
  S' = Generate_New_State (S);
  if (Cost(S') < C(S))
    S = S';
  else if (Random() < exp((Cost(S) - Cost(S')) / T))
    S = S'; // even though Cost(S') > Cost(S)
  else
    // do nothing - stay in same state
  end if;
  if Cost(S) < Cost(minS)
    minS = S;
  end if;
  T=Get_New_Temperature();
end while;
output minS, Cost(minS);
```

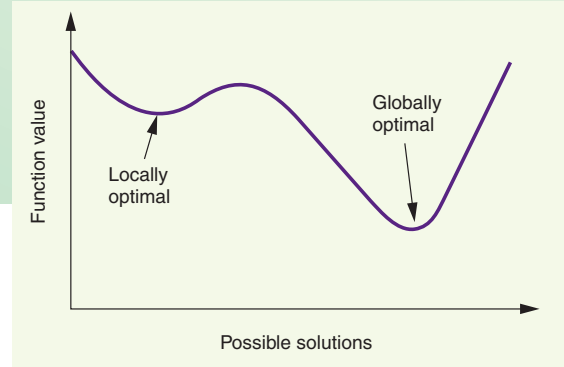


Fig. 2 Local and global optima

optimization problems.

The simulated annealing algorithm is most easily understood as an improvement on a very simple heuristic, called “greedy local search.” It operates as follows:

1. Start with any proposed solution.
2. Change the solution slightly. (For example, swap two cities in a traveling salesman tour).
3. Check to see if the new solution is better than the original. If not, undo the change.
4. Keep going back to step 2 until it is apparent that any small change will only make the solution worse.

Greedy local search has a major weakness: what if a dramatically better solution can be reached by making two or more consecutive changes? This algorithm finds a result that is only locally optimal: it is better than any of the results that are close to it, but it may not be the best result overall. The best result overall, which is what the algorithm is seeking, is called the “globally optimal” result. Figure 2 shows both a locally optimal solution and the globally optimal solution to an arbitrary function. Note that a genetic algorithm without mutation is also likely to get “stuck” in a local optimum.

Simulated annealing attempts to improve on greedy local search by occasionally taking a risk and accepting a worse solution. Specifically, when a solution less optimal than the current solution is proposed, the algorithm starts by computing ΔC , the difference between the cost of current solution and the cost of the proposed one. Then, the algorithm generates a random number r

(the function $\text{Random}()$ below) in the range $0 < r < 1$. If r is less than $e^{-\Delta C/T}$ then the proposed change is accepted. Here, T is the metaphoric “temperature,” that is iteratively decreased to make this probability of acceptance low at low temperatures. An outline of the algorithm is shown in Box A.

At the search’s beginning, the temperature is high and the algorithm is very likely to explore new solutions. This holds true even if those new solutions are slightly less desirable than the current one. However, as the search progresses, the temperature is slowly lowered and the algorithm is less and less likely to accept a switch to a locally less optimal solution.

Theoretically, it is possible to show that the temperature should be lowered in proportion to $(1/\log(n))$. However, at that rate, the cooling is too slow for the algorithm to terminate within a reasonable amount of time. Realistically, the temperature is usually lowered at a linear or exponential rate. Even with these faster cooling schemes, simulated annealing’s compromise between exploring all the possibilities and homing in on a locally optimal result has proven to be a very effective.

Physics and the location-allocation problem

The previous two examples are general purpose algorithms designed to handle

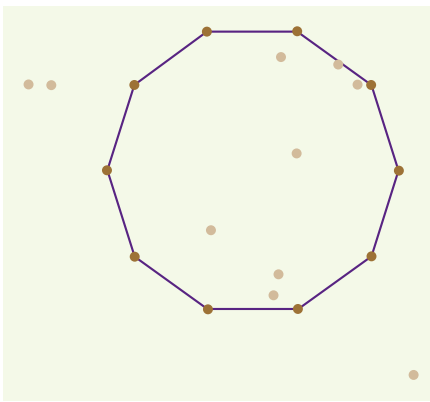


Fig. 3 An initial layout for the 10 city string-particle model (The brown points are the cities. The beige points are the moving particles.)

any combinatorial optimization problem. However, the characteristics of a particular problem often lend themselves to a more specialized heuristic. Here, too, physical metaphors have played a role in inspiring optimization algorithms.

To solve the location-allocation problem discussed in the beginning, an approach called the Simulated N-Body Algorithm has been adapted from the study of moving particles in physics. The facilities that need to be placed on the graph are modeled as particles. The mass of each particle is proportional to the capacity of the facility that it represents. Particles are allowed to move from node to node around the graph.

There is a repulsive force between each pair of particles that is inversely proportional to the distance between the particles. This means that to minimize the potential energy of the system, the particles will tend to spread out from one another. Also, to keep the particles from all being pushed to the outer edge of the graph, a repulsive force from an artificial surrounding “boundary” is added.

To fix the particles into a low-energy configuration, a process of aggregation and decomposition is used. Initially, every particle is aggregated into one large particle with the combined mass of all of its member particles. Then, the smallest of all of the particles is removed from the large aggregate particle. Next, the large particle and the small one are moved into whatever positions will give them the lowest potential energy. The smallest particle is fixed into place, and then the process is repeated by removing the next-smallest particle from the large aggregate particle. From smallest to largest, the particles are all eventually attached to various nodes. The configuration of particles that results from this approach turns

out to be a reasonably good configuration of facility locations.

Math economics and computer networks

A familiar pricing mechanism from economics has been used as a metaphor for an algorithm targeted at some optimization problems in computer networks and distributed systems. Although not strictly a combinatorial optimization algorithm, the algorithm can be applied to certain classes of continuous (real-variable) optimization problems. They also work for combinatorial optimization problems that can be approximated by transformation to a continuous optimization problem.

While the details may vary, the idea is to identify certain parts of the system as “commodities” and other parts of the system as “consumers” competing for these commodities. In a real economy, such competition is settled by prices: Each consumer buys as much as they can afford, buying more of whichever commodities they need more than others. The prices, in turn, are influenced by the demands of the consumers. Commodities that are in high demand become more expensive.

This system of pricing prevents any one consumer from purchasing all the valuable commodities. Yet, it still allows consumers to buy more or less of a given commodity, depending on how important that commodity is to them. In effect, the consumer is maximizing a utility function. This is a function that weights the importance of each commodity to the consumer.

A similar process can be used for a distributed system: Each consumer has a utility function and a fixed amount of money with which to bid. Whether the commodity is processor time, access to

a networked printer, or anything else, the bidding process establishes an equilibrium that takes into account the demands of each of the consumers.

For example, consider the situation where two computer terminals wish to share two database files. The files can be stored at either location, but accessing a part of the file stored at the other location is much slower than accessing it from the locally stored one. So, a decision must be made regarding how much of each file should be stored at each terminal.

Now, suppose that information has been gathered about the probability that an access at each location will require either file. Based on this probability and based on the amount of each file already stored locally, a utility function can be defined for each terminal. It is the probability that any given access at that terminal can be handled locally.

Each file is given an arbitrary starting price and each terminal is given a fixed amount of simulated money with which to bid. Maximizing the utility functions within the constraints of the amount of money that they have, each terminal asks for a certain portion of each file. After that, the prices of each commodity are adjusted in proportion to how much each of them was requested. Based on the new prices, the two terminals recompute their utility functions and make a new bid for a share of the two files. This iterative process (called *tatonnement* in microeconomics) is repeated until the prices have stopped changing and some sort of equilibrium has been achieved.

At any point in the bidding process, if either terminal starts requesting enormous amounts of either file, the price of that file will jump up. With only a fixed amount of money to bid with, the terminal will be forced to ask for less of the now expensive file. Similarly, if either file is accessed less often than the other one, its price will drop. This encourages both terminals to bid for more of it to increase the likelihood that database queries can be satisfied locally.

Particle physics, elasticity & the traveling salesman

The paper, “Emergent Collective Computational Abilities in Interacting Particle Systems,” by Z. Zhang, S. Bai, and G. Li, is another example of an attempt to apply a metaphor from particle physics to solving a combinatorial optimization problem. Specifically, the




Fig. 4 An in-progress snapshot of the 10 city string-particle model

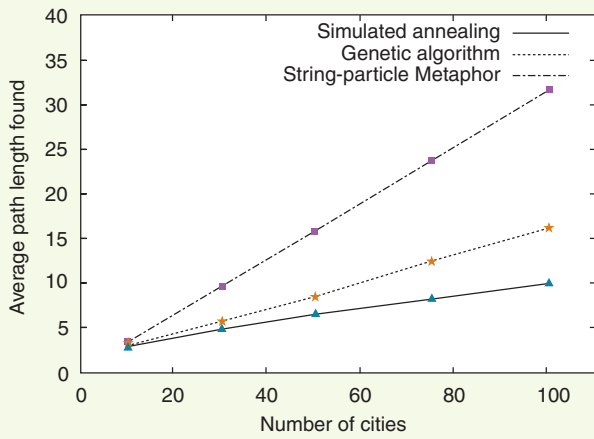


Fig. 5 Comparing algorithms for solving traveling salesman problems

goal is to solve traveling salesman problems with a simulation of charged particles connected by elastic bands.

The algorithm operates as follows: The cities that the traveling salesman seeks to visit are modeled as N fixed negatively-charged particles. Then, N freely moving, positively-charged particles are placed in a circle around the center of mass of the cities. The radius of this circle is equal to the average distance of the cities from their collective center of mass. The moving particles are connected by a loop of elastic string that applies a force proportional to the distance between the particles. Figure 3 shows an example of the particles in their initial configuration for a 10 city traveling salesman problem.

To find a solution, the particles are set free to move under the influence of their electric charges and the force of the elastic string. Figure 4 shows the particles for the example problem after they have begun moving. Whenever a moving particle collides with a city, it locks into place there and stops. When all of the particles stop moving, the elastic string now traces a tour through all of the cities.

This string-particle metaphor shows some promise as an algorithm, since the force of the elastic bands will try to minimize the distance between adjacent particles. However, some empirical testing has shown that other methods produce better solutions to traveling salesman problems.

Figure 5 is a graph that compares the costs of the traveling salesman solutions with the string-elasticity model to solutions found by two other metaphor-based algorithms. Clearly, simulated annealing produces the lowest-cost tours, followed by genetic algorithms.

(In fact, genetic algorithms do not perform well for traveling salesman problems because most schemes for mating two different tours actually reduce the average fitness of the two tours.) Also, the particle physics metaphor does not work well for the traveling salesman problem.

Discussion

As we have seen, a variety of algorithms are inspired by processes and systems studied in other sciences. These metaphor-based algorithms often suc-

ceed by introducing a completely new approach to a familiar problem. This new idea may or may not be better than any previous ones. But, the novelty itself has propelled several of these metaphor-based algorithms to popularity.

However, the catchy sound of a new metaphor does not necessarily reflect on the algorithm based on that metaphor. It certainly gives no guarantees about the quality of a particular implementation of that metaphor-based algorithm. In particular, genetic algorithms seem very intriguing since as human beings, we ourselves are the result of an evolutionary process. Naturally, there is some appeal in the assertion, “The algorithm that created me is a good algorithm.”

Personal appeal aside, there has been a fair amount of empirical testing of several metaphor-based algorithms. Particularly for combinatorial optimization problems, metaphor-based algorithms have provided useful alternatives.

In addition, metaphor-based algorithms are often fairly simple to implement. It is usually quicker to apply simulated annealing than to custom design a heuristic for a particular problem. However, experience shows that specially designed heuristics can often outperform general purpose heuristics such as annealing. Also, some metaphor-based algorithms lend themselves very well to parallel processing. Whereas, considerable parallelization effort may be required for a custom heuristic.

Finally, we note that genetic algorithms can be confused with two other areas of research, namely DNA computing and evolutionary algorithms. The former involves storing data in strands of actual DNA and using enzymes to break down and recombine the DNA to

process the data. The latter is the idea of applying genetic algorithms to the process of algorithm development itself. In this case, the population of solutions is actually replaced by a population of algorithms. They are allowed to compete and recombine to produce new algorithms. Also, the term “emergent behavior” is sometimes used to describe macro-behavior obtained from interacting components. This leads to yet another biology-inspired optimization algorithm, the neural network.

Read more about it

- D.E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. AddisonWesley Pub. Co., Reading, Mass., 1989.

- E. Aarts and J. Korst. *Simulated Annealing and Boltzmann Machines*. John-Wiley Sons, 1989.

- R. Simha, W. Cai and V. Spitzkovsky. Simulated N-Body: A New Particle Physics-Based Heuristic for a Euclidean Location-Allocation Problem. *J. Heuristics*, Vol. 7, No. 1, 2001, pp. 23-36.

- J.F. Kurose and R. Simha. A Microeconomic Approach to Optimal Resource Allocation in Distributed Computer Systems. *IEEE Transactions on Computers*, May 1989.

- D.F. Ferguson, C. Nikolaou, J. Sairamesh and Y. Yemini. *Economic Models for Allocating Resources in Computer Systems, in Market based Control of Distributed Systems*, Ed. Scott Clearwater, World Scientific Press, 1996.

- X. Yao (Ed.). *Evolutionary Computation: Theory and Applications*. World Scientific Pub., 1999.

- L.M. Adleman. Molecular computation of solutions to combinatorial problems. *Science*, 266:1021, 1994.

- S.S. Haykin. *Neural Networks: A Comprehensive Foundation*. Prentice-Hall, 1998.

About the authors

Joseph Carnahan graduated from the College of William and Mary, Virginia in 2000. He is currently a scientist at the Advanced Computational Technology Group at the Naval Surface Warfare Center, Dahlgren Division.

Rahul Simha is an Associate Professor of Computer Science at The George Washington University, Washington DC, 20052. Email: simha@seas.gwu.edu.

This work has been supported in part by grant CDA-9712718 from the National Science Foundation.