

OR Successes Run the Gamut, From Concrete to Kidneys

June 25, 2004

Barry A. Cipra

When Karla Hoffman talks about concrete applications of operations research, she means it literally. During the last three years, Hoffman and her PhD student Martin Durbin at George Mason University have developed an OR model for Virginia Concrete, the nation's seventh-largest concrete company. In a talk titled "The Dance of the Thirty-Ton Trucks," Hoffman described their work at a conference of the Institute for Operations Research and the Management Sciences, April 25-27, in Cambridge, Massachusetts.

The INFORMS conference, "Applying Science to the Art of Business," was held in conjunction with the 50th-anniversary celebration of the MIT Operations Research Center. More than 400 participants attended talks and tutorials on subjects ranging from airline scheduling to healthcare delivery. Finalists for the 2004 Franz Edelman prize presented papers on OR applications for companies including John Deere and Royal Philips Electronics. This year's winner was a team from Motorola, whose entry, an automatic tool for conducting online negotiations with suppliers, has saved the company \$600 million over the last three years. The competitors also included an international team of American and Russian nuclear disarmament experts who have developed a multi-attribute utility model for the disposal or reprocessing of plutonium from the two nations' nuclear arsenals. The seven finalist papers will appear next January in the INFORMS journal *Interfaces*.

But back to those thirty-ton trucks.

Optimal Choreography

In her talk (subtitled "Demand Dispatching in a Dynamic Environment"), Hoffman outlined one of the chief challenges faced by vendors of concrete: Once wet concrete has been loaded into one of the familiar trucks you see around construction sites, that truck has approximately two hours to get where it's going and dump its load; otherwise, the concrete will begin to harden inside the truck. In an ideal world, a planner can work out each day's schedule a day (or more) ahead, tell the drivers when they'll need to arrive in the morning and where they'll be delivering, and fit in last-minute orders wherever possible. But the construction industry is far from ideal. Orders get cancelled or delayed, weather gets in the way, and foremen often grossly underestimate the amount of concrete they actually need. (Once a job is started, it has to be finished, no matter how many loads of concrete it takes—you can't split a job in two unless you want a big crack in your foundation.) Indeed, Hoffman points out, the industry norm is modification of 95% of all orders in the course of a day. Ninety-five percent.

That puts enormous strain on dispatchers, who routinely burn out after a year or two on the job. Hoffman and Durbin, who works for Decisive Analytics Corporation in Arlington, Virginia, were brought in to develop computerized tools that would facilitate the dispatchers' job. "The idea was not to take the human out of the loop, but to give the humans the tools they needed to do their job well and with less stress," Hoffman explains. "At first all we wanted to do was mimic what the dispatchers and schedulers did, to be sure we understood the problem. Then we could go about changing the way that they did business."

One of the key changes was to convert the routing of trucks from a simple two-step to an intricate ballet. Virginia Concrete had approximately 120 trucks assigned to 10 ready-mix concrete plants. Their original system tied each truck to a particular plant: After a job, each truck would return to its "home" plant for

In This Section

[SIAM News](#)

[Current Issue](#)

[Subscribe](#)

[Archives](#)

[All Conferences and Events](#)

[Announcements and Online News](#)

[Professional Opportunities](#)

its next load. (Jobs were assigned to plants based on geographic proximity.) Durbin and Hoffman saw that they could achieve some savings simply by dropping this constraint and allowing trucks to go farther afield, based on a day's evolving schedule. (A related constraint was not dropped: At the end of the day, trucks return to their home plants, so that the drivers can get back to their cars.)

Hoffman and Durbin formulated the optimization as a time-space network flow problem. The nodes of the network are specific locations (e.g., at the plant, on the road, at the site) at specific times, updated at intervals as short as three minutes (that being the time it typically takes to load a concrete truck). The edges are trucks. In the initial, overnight planning stage, the solution includes a fleet of "phantom" trucks—nonexistent vehicles reflecting the planner's anticipation of cancelled orders, whose mathematical purpose is to guarantee a nonempty feasible set. During the day, of course, the phantoms must be replaced with real (shall we say "concrete"?) trucks. This is done in part by "slipping" and "stretching": either delaying the start of a job or, for multitruck jobs, dispatching trucks at, say, 15- rather than 10-minute intervals.

In addition to the obvious constraints, such as the amount of time it takes a truck to get from plant to job (a variable that can be monitored in real time by GPS, which Virginia Concrete had installed on its fleet), the optimization also allows for such constraints as the designation of "premium" customers, who should be given first priority (e.g., exempted from slip and stretch). "Like all businesses, there are some very large customers that make up an awful lot of the profit of the business," Hoffman says. "Even if you're using an automated system, you want to make sure that they get their concrete on time as asked for."

Virginia Concrete has been using the optimization software since 2002. (During initial testing, dispatchers let the algorithm make all the scheduling decisions, to see how well it worked on its own. They are now allowed to override the computer's solutions.) The company has been able to increase the amount of concrete delivered per driver by 26%. Overall, "the company is so excited about this that they've taken us to other people in the industry, saying 'you've got to get into this too,'" Hoffman says.

From her perspective, "the exciting part of it is that the technology is likely to be useful for lots of other applications." Part of the fun of OR, she says, is the exposure to new fields. Like an ornithologist in a nature preserve, Hoffman now notices concrete trucks wherever she goes. "I've learned more about concrete than I ever thought I would need to know."

OR in the OR

Directed graphs can save lives. A little hyperbolic, perhaps, but that's the gist of Alvin Roth's analysis of a new proposal for arranging kidney transplants, a life-saving operation that is performed nearly 15,000 times a year in the United States—but with a waiting list that adds 23,000 new patients each year. Roth, an economist at Harvard University, outlined the proposal in a talk at the INFORMS conference.

Demand for kidney transplants in the U.S. far outstrips supply. The United Network for Organ Sharing (UNOS) currently has more than 55,000 patients on its waitlist (in 2002, the number was 54,844, compared with 22,063 in 1992). The waitlist grows by about 3500 per year. The rest of the equation, $23,000 - 15,000 = 3500 + D$, is accounted for by approximately 3500 deaths and 1000 patients who become so ill they no longer qualify for transplants. If more people became organ donors (or if organ donation became an "opt-out" box on driver's license applications), many thousands more lives could be saved each year.

Most transplanted kidneys come from donors who no longer need their organs,

namely cadavers. But because it's possible to live with just one kidney, a growing number come from living donors, usually close relatives of the patients. In 2002, for example, 8534 cadaveric and 6233 live-donor transplants were performed (up from 7202 and 2535 in 1992). It's the live donors Roth has an eye on.

An economist who takes an interest in organ transplants may sound suspicious, but Roth is quick to stress that he is not talking about the buying or selling of organs—which is a felony under U.S. law. Nevertheless, he says, there are aspects of live-donor transplants that fall within the scope of economic analysis. Economists are interested mainly in creating efficient markets-mechanisms for the exchange of goods that lead to results that cannot be improved for some without harming others. Money is merely a convenient (but in this case forbidden) lubricant.

If all kidneys were equal, there would be no problem: A patient with a willing donor would simply get that kidney, and patients without donors would wait their turn for a cadaver to turn up. But not every kidney is suitable for every patient. To begin with, the blood types need to match. A patient with type-O blood can accept only a type-O kidney. Type-A or -B patients can accept type-O as well as their own types. Finally, type-AB patients can accept any type. These strictures put a premium on type-O donors, and place type-O patients at a disadvantage.

In addition, the patient and donor must be compatible with respect to a set of six kinds of proteins, collectively called HLA. Some transplants are ruled out because the patient has antibodies against one of the prospective donor's HLA proteins. The proteins don't otherwise need to match, but, in general, the more that do the better (which is one reason close relatives are good donor candidates).

A potential donor who is not a match for his or her patient/friend normally goes home, and the patient goes on the waitlist for a cadaver kidney. Occasionally, however, a doctor or hospital realizes that two patient-donor couples, while neither is a match individually, are a good match for each other. In that case a "paired exchange" can be arranged: Patient A gets donor B's kidney, and patient B gets donor A's. This can be advantageous even when the individual couples do match: The exchanged couples might be an even better match. Oddly, though, there is as yet no organized registry for paired exchanges; most potential donors effectively disappear if they are not a match for the intended recipients.

If paired exchanges make sense, why not cyclic exchanges? In fact, Roth says, the Johns Hopkins Comprehensive Transplant Center in Baltimore carried out just such a procedure last year among three couples. But Roth envisions even more: a registry of patient-donor couples and an algorithm that would optimally assign donors to patients.

The basic idea is for the patients (actually their doctors) to list the available donors in order of preference (which, for simplicity, is assumed to be strict rank order—i.e., without ties). The algorithm, known as the "top trading cycles" (TTC) method, has each patient point to his or her top choice, and each donor point to his or her patient/friend. The resulting directed graph is guaranteed to have at least one cycle, and because each node points to exactly one other node, the cycle or cycles created are unambiguous—no directed edge can belong to two cycles. TTC sends these cycles to the operating room and then repeats itself with the remaining patient-donor couples.

TTC is an old algorithm. It was used in a 1974 analysis of a housing market by Lloyd Shapley and Herbert Scarf, who credited the algorithm to David Gale. In the Shapley-Scarf analysis, the patients were homeowners, and the donors their

current homes, which they wished to exchange. The beauty of TTC is that it is efficient and strategy-proof. In each round, that is, no one has any incentive to point to anyone but his or her top choice, and the final result cannot be improved for anyone without someone else suffering.

But there is a fly in the ointment: Patients may rank some donors, including their own, as completely unacceptable. (This isn't much of a concern in a housing-market analysis, where it can be assumed that all homeowners will at least settle for their current abodes.) In other words, patients may prefer a spot on the waitlist for cadavers over some (or even all) of the donors. This introduces complications, both mathematical and ethical.

The ethical complication concerns incentives for the patient-donor couple and fairness to people already on the waitlist. If a patient who prefers to enter the waitlist is simply added at the end, as seems fair to those already signed up, his or her donor has no real incentive to donate a kidney to anyone else; the couple's doctor could not now ethically recommend that the donor undergo the operation. It may seem sensible, on the other hand, to simply insert the patient in the position of the person near the top of the waitlist who receives a live-donor kidney as a result of the exchange. (This has already been considered in the case of "indirect" pair exchanges, for a single, incompatible patient-donor couple.) But ethical problems arise here, too. In particular, type-O patients on the waitlist are disadvantaged if another type-O patient is placed ahead of them. The medical profession as a whole must see to the welfare of all patients, not just those with friends who are willing to help them out.

The mathematical challenge is one of ensuring efficiency while remaining strategy-proof. Roth and colleagues Tayfun Sönmez and M. Utku Ünver of Koç University in Istanbul, Turkey, have addressed these mathematical problems and some of the ethical aspects in an approach they call the "top trading cycles and chains" (TTCC) algorithm. Their analysis appeared in the May 2004 issue of the *Quarterly Journal of Economics*.

In TTCC, each patient's preferences include the waitlist along with the donors. The algorithm begins just like TTC—patients point to their top choices, any cycles that appear are removed, and the algorithm re-runs with the remaining patient-donor couples. But at some point, the algorithm runs out of cycles: The directed graph is a tree rooted at the waitlist.

If there is no branching in this tree, the solution is fairly simple. Donations are simply along each chain, with the patient at one end going on the waitlist and the donor kidney at the other end going to someone already on the waitlist. (To be sure, it's still necessary to consider the order in which the chains are prioritized, since this affects patients' placement on the waitlist.

If the tree branches, TTCC must apply a "chain-selection rule." The rule begins by picking a chain. It must then specify whether the donor kidney at the tail end should go immediately to someone on the waitlist, or if it should remain available for the next round of the algorithm. (Being at the tail end doesn't necessarily mean that no one wants that kidney, even in the current round—the selection rule can terminate a chain wherever it wants. For example, it might pick the chain starting at the patient who's been given highest priority for getting a kidney.)

Roth and colleagues have shown that any chain-selection rule that leaves the tail-end donor kidney available for the next round is efficient. They have also found some rules that are strategy-proof. One such rule is to randomly (or by some priority rule) pick a "chain" of one patient who is pointing either at the waitlist or at an available tail-end kidney. Another, mentioned earlier, is to start with the highest-priority patient. One rule that is not strategy-proof is the longest-chain rule, in which the aim is to make as many patient-donor

assignments as possible in each round. This tempts patients to point to less-preferred donors in an early round in order to forge a long chain, for fear that they'll wind up on the waitlist if they don't.

(Roth et al. point to one other strategy-proof rule that, while not efficient, has ethical appeal for the problem of type-O patients already on the waitlist. In this rule patients with type-O donors are given the highest priority. As long as type-O donors are available, the chain starts with the highest-priority patient and his or her type-O donor goes straight to a person on the waitlist. When the type-O donors are gone, the rule leaves the tail-end kidney available for the next round.)

Even patients who are promised high priority on the waitlist may consider that option not worth the risk to their donors. Patients can express this in their preferences by ranking their (incompatible) donors *above* entering the waitlist. If the algorithm ever produces such a cycle, the corresponding patient-donor couple simply drop out. (They can try again with a different pool, or the patient can enter the waiting list without any extra priority.) Roth and colleagues have simulated a kidney market in which only 40% of patients rank the waitlist above their own (incompatible) donors, using statistics on the prevalence of blood types and HLA characteristics. (They also use statistics on patient and donor ages, which play a role in determining the desirability of a kidney.)

The results are promising. In general, only 55% of patient-donor couples are compatible. In simulated markets with 100 couples, simple pair exchange alone increases the percentage of patients who receive a live kidney to about 73.5%. (With 300 couples, the percentage goes up to 75%. In general, the larger the market, the better the results.) With the TTCC algorithm, the rate goes to nearly 90%. Of the 10 (out of 100) patients who don't get a living-donor kidney, six (and their donors) drop out, and four go onto the waiting list.

The (simulated) exchanges also increase the number of HLA matches, which improves patients' prospects. Roth believes that the TTCC algorithm would also ultimately benefit even donorless patients on the waitlist, by making the waitlist shorter. After all, if some 6000 donor transplants are currently done each year, there are probably another 5000 incompatible patient-donor couples. If, instead of letting these donors disappear and placing the patients at the end of the waitlist, a market runs with all 11,000 couples, the simulations suggest that about 10,000 of the patients will receive live-donor kidneys and, of the remainder, only about 400 will receive priority on the waitlist. Although ethical aspects would have to be dealt with in starting up such a system (donorless patients already on the waitlist would not benefit from a sudden influx of high-priority patients from chains in the market), the increased number of live-donor transplants would cause the waitlist to grow more slowly (it could conceivably even shrink), so that in the future all patients would move through it more quickly than they would otherwise.

Barring advances in medical technology or social changes in people's attitudes about organ donation, there will always be a waitlist for kidney transplants and some percentage of patients will die waiting. But if some relatively simple mathematics can save some of those lives, Roth for one thinks it's worth a try. Anyone with a friend or family member undergoing dialysis is likely to agree.

Barry A. Cipra is a mathematician and writer based in Northfield, Minnesota.

- [See more news from this issue](#)

Phone: +1-215-382-9800 | FAX: +1-215-386-7999

[site map](#) | [privacy policy](#) | webmaster@siam.org | [suggestions](#)

Banner art adapted from a figure by Hinke M. Osinga and Bernd Krauskopf (University of Auckland, NZ).