Technical Perspective
Moving Spectrum

By David C. Parkes

HOW CAN YOU transform the use of radio spectrum from low-value use such as over-the-air television to high-value use such as wireless broadband? A "big bang" response was proposed in the early 2000s. The idea was to use a market in which owners of spectrum could sell their spectrum for new uses by others. But this was plagued with difficulty.

Imagine TV stations as owners of cars in a parking lot. Each wireless company wants to buy a large, contiguous space in the lot. Doing so will create great value. But to succeed, it must buy space from each of a number of stations. And each station will try to seek a huge amount of money for this space! The solution to this "holdout problem" came in a new regulation that gave the Federal Communications Commission (FCC) the right to move TV stations to new frequencies. This creates competition. As long as enough stations sell spaces *anywhere* in the lot then remaining cars can be moved around to create space.

The FCC's "incentive auction" proceeds in stages, each stage associated with a *clearing target*. To illustrate, suppose the target is to sell half of the parking lot. Each stage proceeds in two interrelated parts. First, buyers compete to set prices for the large, contiguous spaces that would be created if half the lot were sold. Buyers drop out over time and the auction continues until every buyer who remains can be allocated space within the clearing target. Second, sellers place lower and lower bids in a "reverse auction" to sell their individual spaces. They drop out over time (refusing to sell their space) and the auction continues until every seller who remains is needed to be able to move around the sellers who refuse to sell in order to meet the clearing target. This repeats, each stage with a progressively smaller target, until the revenue generated is at least that of the cost of the reverse auction.

The incentive auction presents a critical algorithmic challenge—that of solving a *repacking problem* when operating the reverse auction, and doing so quickly and at scale. The real problem is more complicated than repacking a parking lot. It is as if all the parking slots were shaped like Tetris pieces, and furthermore my psychedelic-colored car cannot go right next to your white car (in reality, there are complex, physics-based interference constraints that dictate what is feasible). Consider a seller who is willing to sell at the current price. The repacking problem asks: *Is there a feasible assignment of the available spectrum (the spectrum not targeted to buyers) to this seller and all the other stations who have already chosen not to sell?* If the answer is NO (or not provably YES) then we must finalize the price to this seller. If the answer is YES then we can lower the price and give the station the choice of taking the next price or walking away.

This scenario provides the backdrop for the breathtaking research contribution presented in Newman et al. For the auction to run in a reasonable amount of time, the FCC wanted to run two rounds of price updates per day, giving a few hours to process a round. Bids had to be processed sequentially. Taking into account the fraction of instances that finish very quickly, this required solving each repacking instance within approximately one minute. State-of-the-art algorithms were able to solve only 80% of hard instances within this time cut-off. The solver in Newman et al. can solve 88% of hard instances in *under a second* and 96% within the cutoff. The authors estimate the impact of these improvements could represent a $1 billion gain to the U.S. government in lower prices and thus reduced cost in the reverse auction.

How did they do this? The approach is that of *data-driven algorithm design*. This allows decisions about the precise design of an algo-rithm, and in this case the design of a portfolio of different, complementary algorithms that run in parallel, to be made in a principled, data-driven way. The role of an algorithm designer becomes that of designing configurable algorithms as well as using creativity to come up with domain-specific heuristics that could prove useful. In addition, the designer provides a realistic distribution of inputs on which to optimize performance. Having exposed 191 parameters in the design space, each nested as many as four-levels deep, the authors call their approach "deep optimization."

There are larger lessons here. First, through outstanding computer science in developing a customized solution to this difficult and important problem, the researchers provided confidence that these large-scale, NP-hard problems could be solved at scale. Second, this success speaks to the importance of sustained research—this is the culmination of a concerted research effort over more than a decade into data-driven algorithm design. This project itself is a tour de force in "big bang" computer science, with an open source simulator, the inclusion of 20 state-of-the-art SAT solvers, and days of compute time to generate their results.

Eventually, the FCC's Incentive auction generated $19.8 billion of revenue, $10.1 billion of which flowed to broadcasters. It moved 84MHz of spectrum to highest and best use with all trades voluntary. This is one of the biggest successes to date for the Economics and Computer Science research agenda. [C]

David C. Parkes is the George F. Colony Professor of Computer Science, and Co-Director, Harvard Data Science Initiative in the Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA.