Station Repacking in the Incentive Auction

Neil Newman and Kevin Leyton-Brown
Department of Computer Science
University of British Columbia
{newmanne, kevinlb}@cs.ubc.ca

***Abstract*—Station repacking problems ask whether a given set of television stations can each be assigned a channel such that no two stations excessively interfere with each other. Solving repacking problems was critical to running the Incentive Auction, which reduced the number of channels available to television stations in the UHF band and needed to make many queries about which sets of stations could fit in the channels that remained in order to set prices. We built an auction simulator that generates realistic station repacking instances.**

I. Introduction

Over 13 months in 2016–17 the US Federal Communications Commission (FCC) conducted an "incentive auction" to repurpose radio spectrum from broadcast television to wireless internet [1]. The result of the auction was to remove 14 UHF-TV channels from broadcast use, sell 70 MHz of wireless internet licenses for $19.8 billion, and create 14 MHz of spectrum for unlicensed uses. With fewer UHF channels remaining for TV broadcast, the TV spectrum was also reorganized. Each station was either "repacked" in the leftover channels or voluntarily sold its broadcast rights; volunteers were compensated with $10.05 billion.

Our focus is the "reverse auction", the component of the incentive auction that determined which stations would go off the air and for what prices. Station repacking is central to the reverse auction, as it arises every time the FCC tries to lower a price—in practice, tens of thousands of times in a single auction. The auction rules are roughly as follows: First, all participating stations are given initial price quotes. Each station responds either that it agrees to sell its broadcast rights at the quoted price or else it declines to participate ("exits the auction"), meaning that it will be guaranteed some interference-free channel in the spectrum that will remain available to television stations after the auction's conclusion. The auction then repeatedly iterates over active bidders. Every time a bidder i is considered, a repacking problem is solved to determine if i can be feasibly repacked along with all exited stations. If such a feasible repacking exists, i is given a (geometrically) lower price quote and again has the options of accepting or exiting. Otherwise, i is *frozen*: its price stops descending and it is no longer active. The auction ends when all bidders are either frozen, exited, or receive price quotes of zero.

Since station repacking problems are NP-hard, and problems must be solved sequentially, and the auction executes on a real-time schedule allowing only so much time per problem, some problems will remain unsolved. The auction design is robust to such failures, treating them as proofs of infeasibility at the expense of raising the cost required to clear spectrum.

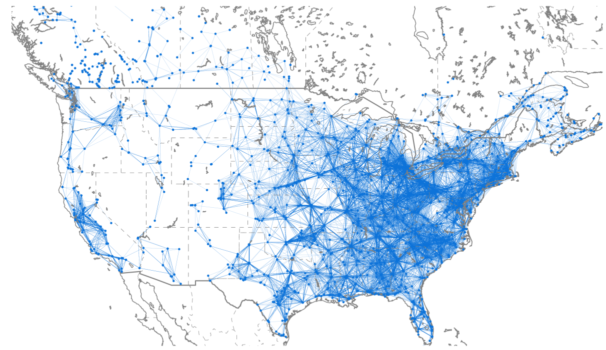

Fig. 1. Interference graph visualizing the constraint data (2 990 stations; 2 575 466 channel-specific interference constraints).

II. Station Repacking

We now give a formal definition of the station repacking problem. Each television station in the US and Canada $s \in \mathcal{S}$ was assigned to a channel $c_s \in \mathcal{C} \subseteq \mathbb{N}$ prior to the auction that ensured that it would not excessively interfere with other, nearby stations. The FCC determined pairs of channel assignments that would cause harmful interference based on a complex, grid-based physical simulation ("OET-69" [2]). Let $\mathcal{I} \subseteq (\mathcal{S} \times \mathcal{C})^2$ denote a set of *forbidden station–channel pairs* $\{(s,c),(s',c')\}$, each representing the proposition that stations s and s' may not concurrently be assigned to channels c and c', respectively. The effect of the auction was to remove some broadcasters from the airwaves and to reassign channels to the remaining stations from a reduced set. This reduced set was defined by a *clearing target*, corresponding to some channel $\overline{c} \in \mathcal{C}$ such that all stations are only eligible to be assigned channels from $\overline{\mathcal{C}} = \{c \in \mathcal{C} \mid c < \overline{c}\}$. Each station can only be assigned a channel on a subset of $\overline{\mathcal{C}}$, given by a *domain* function $\mathcal{D} : \mathcal{S} \rightarrow 2^{\overline{\mathcal{C}}}$ that maps from stations to these reduced sets. The *station repacking problem* is then the task of finding a repacking $\gamma : \mathcal{S} \rightarrow \overline{\mathcal{C}}$ that assigns each station a channel from its domain that satisfies the interference constraints: i.e., for which $\gamma(s) \in \mathcal{D}(s)$ for all $s \in S$, and $\gamma(s) = c \Rightarrow \gamma(s') \neq c'$ for all $\{(s,c),(s',c')\} \in \mathcal{I}$. A problem instance thus corresponds to a set of stations $S \subseteq \mathcal{S}$ and channels $C \subseteq \overline{\mathcal{C}}$ into which they must be packed, with domains $\mathcal{D}$ and constraints $\mathcal{I}$ implicitly being restricted to S and C; we call the resulting restrictions D and I.

We define the *interference graph* as an undirected graph in which there is one vertex per station and an edge exists between two vertices s and s' if the corresponding stations participate together in at least one interference constraint: i.e., if there exist $c, c' \in C$ such that $\{(s,c),(s',c')\} \in I$. Figure 1 shows the interference graph for this auction.

Lastly, we note that descending clock auctions repeatedly generate station repacking problems by adding a single station s^+ to a set S^- of provably repackable stations. This means that every station repacking problem $(S^- \cup \{s^+\}, C)$ comes with a partial assignment $\gamma^- : S^- \to C$ that we know is feasible on restricted station set S^-.

III. SAT Encoding

Given a station repacking problem (S, C) with domains D and interference constraints I, we create a Boolean variable $x_{s,c} \in \{\top, \bot\}$ for every station–channel pair $(s,c) \in S \times C$, representing the proposition that station s is assigned to channel c. We then create three kinds of clauses: (1) $\bigvee_{d \in D(s)} x_{s,d} \;\; \forall s \in S$ (each station is assigned at least one channel); (2) $\neg x_{s,c} \vee \neg x_{s,c'} \;\; \forall s \in S, \forall c, c' \neq c \in D(s)$ (each station is assigned at most one channel); (3) $\neg x_{s,c} \vee \neg x_{s',c'} \;\; \forall \{(s,c),(s',c')\} \in I$ (interference constraints are respected). Note that (2) is optional: if a station is assigned more than one channel, we can simply pick one channel to assign it from among these channels arbitrarily[1]. A SAT encoding of a problem involving all stations at a clearing target of 36 involved 73 187 variables and 2 917 866 clauses.

IV. Generating Repacking Problems

The exact repacking problems that arose in the auction are not publicly available. We therefore wrote our own reverse auction simulator and released it as open source software; it is available at http://www.cs.ubc.ca/labs/beta/Projects/SATFC/. Our simulator assumes that stations bid to myopically maximize their utility. Our simulator can sample from two value models. The first model[2] is from Doraszelski et al [3] and the second we created by interpreting the FCC's publicly released bid data as censored value observations (for more detail, see [4]).

The choice of solver for repacking problems can be customized in our simulator. Since an auction's trajectory can change based on whether or not a given problem is solved, using different solvers on auctions with otherwise the same strategy for bidders may lead to different auction trajectories and therefore generate different problems. The simulator's clearing target is controlled by a parameter. Further details on the simulator can be found in [4], [5].

[1]These constraints can be toggled on or off through a parameter in our software.

[2]The authors shared their fit model parameters with us, but requested we did not make them public. Therefore, in order to run simulations under this model a user must supply their own parameters.

V. Benchmark

We will now discuss the parameters of the simulator used to create our submitted benchmark. We ran 20 simulations at the 84 Mhz clearing target, corresponding to a maximum allowable channel of 36. We note that this is the amount of spectrum actually cleared by the Incentive Auction. We used Doraszelski et al's value model. As our solver, we used SATFC 2.3.1, the solver used by the FCC, with a 60 second cutoff. All of our experiments were run on Intel Xeon E5-2640 v2 processors on nodes with 96 GB of RAM. The simulations took between 3.52 and 4.02 hours to run (wall time), of which the majority of the time, between 2.52 and 3.07 hours, was spent within SATFC.

We sampled 10 000 "nontrivial" problems uniformly at random from all of the problems across all simulations to use as our dataset, where we defined nontrivial problems as those that could not be solved by greedily augmenting the previous assignment γ^-. This benchmark set consists of 9 482 feasible problems, 121 infeasible problems, and 397 problems that timed out at our 1 minute cutoff and for which we have not determined the feasibility. Constraints of type (2) are included in these problems[3]. While the submitted benchmark consists of only a few instances, the full set of instances are available as CNFs at http://cs.ubc.ca/labs/beta/Projects/SATFC. The previous solution γ^- and a MIP encoding are also available for each problem.

VI. Additional Benchmarks

We conclude this document with a pointer to three additional, more recent station repacking benchmarks we created that, due to extraneous circumstances, were not the ones submitted to the competition. To create these benchmarks, we reran smaller scale versions of three of the experiments in [4], storing all of the station repacking problems that arose in the auction simulations. Each experiment varied some facet of the incentive auction design in order to compare the resulting differences in the simulated economic outcomes. We briefly summarize the experiments below; we refer the reader to [4] for more details.

a) Scoring Rules: These experiments compared four different methods of setting the auction's station-specific opening prices ("scoring rules") in simulations with an 84 MHz clearing target.

b) Feasibility Checker: These experiments compared replacing SATFC with one of five different feasibility checkers to solve station repacking problems in simulations with an 84 MHz clearing target.

c) Clearing Procedure: These experiments compared four different initial clearing targets of the auction, holding the final clearing target fixed at 84 MHz.

[3]We note that this representation may not correspond directly to the CNFs that were solved in the simulations. SATFC is a parallel portfolio of several solvers, some of which use the other encoding, and some of which perform preprocessing on a constraint satisfaction problem representation prior to SAT conversion.

	SAT	UNSAT	Unknown
Scoring Rules	2969	3004	4027
Feasibility Checker	829	759	8412
Clearing Procedure	2438	3506	4056

TABLE I

FEASIBILITY DISTRIBUTION OF REPACKING PROBLEMS IN EACH OF OUR THREE ADDITIONAL BENCHMARKS.

As above, a 60 second cutoff was used for solving the station repacking problems in each simulation. For each experiment, we ran five simulations (each using a different sampled value profile) for each configuration for each of the two value models. This amounted to a total of 140 auction simulations, leaving us with 12.7 million problems, of which 4.4 million problems were "nontrivial". From each experiment, we sampled 10 000 problems uniformly at random. The feasibility distributions of each benchmark are summarized in Table I. The feasibility checker experiment has a higher incidence of timeouts since it uses solvers other than SATFC 2.3.1, which were not optimized for solving repacking problems.

A dictionary file is included along with each benchmark that relates each problem back to the experiment configuration that it originated from and contains information on runtime[4]. Constraints of type (2) are included in all of these problems. The previous solution γ^{-} is included for each problem. The benchmarks can be downloaded at http://cs.ubc.ca/labs/beta/Projects/SATFC.

ACKNOWLEDGEMENTS

We thank Pranay Jain for his assistance in creating the additional benchmarks.

REFERENCES

[1] K. Leyton-Brown, P. Milgrom, and I. Segal, "Economics and computer science of a radio spectrum reallocation," *Proceedings of the National Academy of Sciences*, p. 201701997, Jun. 2017. [Online]. Available: http://www.pnas.org/lookup/doi/10.1073/pnas.1701997114

[2] FCC, "Office of engineering and technology releases and seeks comment on updated OET-69 software," *FCC Public Notice*, vol. DA 13-138, February 2013.

[3] U. Doraszelski, K. Seim, M. Sinkinson, and P. Wang, "Ownership concentration and strategic supply reduction," National Bureau of Economic Research, Working Paper, http://www.nber.org/papers/w23034 23034, January 2017. [Online]. Available: http://www.nber.org/papers/w23034

[4] N. Newman, K. Leyton-Brown, P. Milgrom, and I. Segal, "Incentive auction design alternatives: A simulation study," in *EC '20: The 21st ACM Conference on Economics and Computation, Virtual Event, Hungary, July 13-17, 2020*. ACM, 2020, pp. 603–604. [Online]. Available: https://doi.org/10.1145/3391403.3399499

[5] N. Newman, A. Fréchette, and K. Leyton-Brown, "Deep optimization for spectrum repacking," *Communications of the ACM*, vol. 61, no. 1, pp. 97–104, 2017.

[4]For problems coming from the feasibility checker experiment, note that walltime is with respect to the solver used in the experiment, not necessarily SATFC.