

Designs for Ramp-Constrained Day-Ahead Auctions

Shmuel S. Oren

IEOR Dept.

UC Berkeley

Berkeley, CA 94720

oren@ieor.berkeley.edu

Andrew M. Ross

ISE Dept.

Lehigh University

Bethlehem, PA 18015

amr5@lehigh.edu *

Abstract

Some electric power markets allow bidders to specify constraints on ramp rates for increasing or decreasing power production. We show in a small example that a bidder could use an overly restrictive constraint to increase profits, and explore the cause by visualizing the feasible region from the linear program corresponding to the power auction. We propose two penalty approaches to discourage bidders from such a tactic: one based on duality theory of Linear Programming, the other based on social cost differences caused by ramp constraints. We evaluate the two approaches using a simplified scaled model of the California power system, with actual 2001 California demand data.

1. Introduction

Many restructured electricity systems rely on self-commitment of generation resources rather than on central unit commitment. This structure avoids some of the incentive-compatibility problems associated with more centralized systems such as the original UK system (prior to NETA), PJM, NYPP, New England pools which involve multi-dimensional auctions allowing bidders to specify technical constraints on the dispatch. Such auctions are often susceptible to manipulation allowing bidders the opportunity to profit by specifying deceiving technical constraints. Unfortunately, in systems that rely on self-commitment and clear the hourly day ahead market without consideration of intertemporal constraints on dispatch, mismatches between the ISO schedule and the capabilities of generators must be made up in the real-time balancing market. Not only is this an expensive solution, it shifts a perhaps unnecessary volume of energy transactions to the real time balancing market. Furthermore, although some generation

technologies hinder efficient scheduling due to their ramp constraints, and others assist with their rapid ramping capabilities, the rapid-ramping plants are not rewarded for the flexibility they bring to the system. In this paper we explore ways to allow bidders to specify ramp rate constraints while mitigating to some extent the possibilities for deceptive bidding. We are concerned with two effects of such bidding: an increase in overall cost, and inequity of outcomes.

We deal here with a day-ahead energy market without network considerations, and allow bidders to specify ramp rate constraints as part of their offers. The decision to turn a unit on is still left to the generator who will need to absorb the startup cost and can ensure its minimum generation level using zero or negative offer prices for that amount. In order to accommodate ramp constraints, the market operator would need to clear the markets for all the 24 hours simultaneously using an optimization algorithm. This problem, fortunately, is far less complicated than a unit commitment problem since it does not introduce discrete variables into the problem. If the objective is to minimize the social cost of the dispatch, then the market clearing problem with ramp constraints can be formulated as a Linear Programming (LP) problem and solved by standard algorithms and software. The solution gets more complex if the objective is to minimize total procurement cost, since that criterion introduces nonconvexities. It should be noted that the day-ahead electricity auction in the Spanish system allows bidders to specify ramp constraints as well as a floor on their total 24-hour revenue, which are incorporated into the market clearing formulation and solved using a heuristic algorithm [1].

One simple strategy for designing the auction to avoid manipulation is to allow ramp-rate constraints to change only once a month, or some other suitable time interval, since generator technologies do not change very rapidly. While this has little theoretical backing, its practical effect would hopefully be to prevent generators from rapidly responding to market conditions with false ramp rates, leaving them with little choice but to report their true ramp rate

*This work was supported by the Power Systems Engineering Research Center (PSERC) and by the Electric Power Research Institute.

constraints. However, it would also allow a company to lock in a misleading constraint and profit from it for a whole month. Another simple idea is to have a regulatory agency certify ramp rates, as the California regulations stated; however, this leads to only an upper bound on the rates. Bidders might be bidding into various markets, and so would need to split their ramp rates among them, giving them a valid reason to specify ramp rates lower than their certified values.

The problem data given in Table 1 help demonstrate the incentive compatibility problem that can arise when ramp rates can be specified and are accounted for in the dispatch but the spot market prices for energy are based on snap shot market clearing that does not account for intertemporal constraints affecting the dispatch. For illustrative purposes we will assume that each time period has a duration of 12 hours. Table 2 summarizes the auction results (for minimizing social cost) in the absence of ramp constraints.

Suppose now that Generator B specifies an intertemporal constraint requiring that his dispatch level at Peak is no greater than his dispatch level at Off-peak. Then the auction results under optimal social-cost dispatch are as given in Table 3. We note that the intertemporal constraint stipulated by Generator B caused an increase in social cost due to displacement of cheap energy with more expensive energy, while the resulting increase in market clearing prices increased the net profits of both generators A and B. While one could not fault Generator A for enjoying the windfall, it is clearly inappropriate for Generator B to reap extra profits by stipulating a constraint that impedes efficiency. Such a profit opportunity could motivate generators to misrepresent their ramping capability in order to drive up prices. Some market designs (e.g. the old UK system) attempt to prevent misrepresentation of constraint by barring a constrained generator from setting the clearing price. Our example demonstrates, however, that such a restriction still does not solve the problem since the constrained generator may force a more expensive unit into the dispatch and benefit from the higher clearing price set by that unit. The Spanish market design eliminates such perverse incentives by forcing generators to bear the dispatch consequences of their ramp constraints, which in our example would amount to forcing Generator B out. This rule solves the incentives problem but unfortunately, it may also unnecessarily increase the social cost of the dispatch.

2. Visualization

We wish to visualize the feasible region and optimal solution for our small example, to learn how Generator B can profit by specifying a ramp constraint. Our small example has 6 variables (one for each generator in each period) and 2 equality constraints (one for demand in each period). Us-

Table 1. Problem Data

	Off-peak	Peak
Demand	1 GW	3 GW
Gen. A offers	1 GW, \$10/MWh	1 GW, \$10/MWh
Gen. B offers	2 GW, \$15/MWh	2 GW, \$15/MWh
Gen. C offers	2 GW, \$25/MWh	2 GW, \$25/MWh

Table 2. Auction Results without ramp constraints

	Off-peak	Peak
Clearing Price	\$10/MWh	\$15/MWh
Gen. A	1 GW	1 GW
Gen. B	0	2 GW
Gen. C	0	0

Table 3. Auction Results with Gen. B ramp constraint

	Off-peak	Peak
Clearing Price	\$10/MWh	\$15/MWh
Gen. A	0 GW	1 GW
Gen. B	1	1 GW
Gen. C	0	1 GW

Table 4. Financial summary

	Unconstrained	Ramp-Constrained
Gen. A profit	\$60000	\$180000
Gen. B profit	0	\$120000
Gen. C profit	0	0
Social Cost	\$600000	\$780000
Acquisition Cost	\$660000	\$1080000

ing the equality constraints we can eliminate two variables. Furthermore, Generator A is always assigned 1 GW during Peak regardless of us including or excluding Generator B's ramp constraint, so we can treat Generator A during Peak as a fixed quantity rather than as a variable. This reduces us to 3 variables, which are more easily visualized than our original 6. Fig. 1 shows the feasible region without Generator B's ramp constraint; it is a triangular prism, and the optimal solution is marked with an extra circle. We start imposing the ramp constraint gradually in Fig. 2, where the increase from Off-peak to Peak must be under 1.5 GW for Generator B. We see that the new constraint plane has cut off two corners of the non-ramp-constrained region, pushing the optimal solution so Generator B gets more business during Off-peak, and Generator A gets less. Also, we must now take Generator C's higher price. Fig. 3 shows the feasible region (an irregular tetrahedron) once the full ramp constraint that B specified has been imposed; Generator A has been pushed out of the Off-peak solution entirely. Overall, we see the problem with allowing bidders to specify ramp constraints: it allows them to specify slanted constraint planes (as opposed to upper-bound constraints on power output, which are orthogonal to the axes and cannot push the solution away from other bidders).




Figure 1. Feasible Region without ramp constraints

When the feasible region is viewed this way (especially as in Fig. 2), it is natural to think of the effects of the ramp-rate constraint in terms of LP sensitivity theory. We will explore this idea further in the next section.




Figure 2. Feasible Region with Generator B increase limited to 1.5 GW

3. Penalty Systems

If we are to allow bidders to specify ramp constraints, we should ensure that our final dispatch satisfies all of the constraints, since we do not know which are true and which are misleading. That is, in trying to make our auction more incentive-compatible, we have little leeway in the dispatch. We do have some room to adjust the payments, though. One option, explored in [3], is the Vickrey-Clarke-Groves (VCG) auction. A VCG auction effectively pays companies to bid truthfully. It pays each offer the generation cost based on the revealed parameter and an additional premium that reflect the contribution of the offer to social welfare i.e., the difference in the optimal value of the objective function with and without the offer. While the VCG auction is incentive compatible and efficient, it may lead to revenue insufficiency, and is considered undesirable due to its radical departure from the uniform price philosophy underlying commodity markets. Instead, we propose to capture the essence of the VCG approach by starting with the usual uniform (market-clearing) payments as the benchmark and impose financial penalties on companies whose ramp constraints are active at the optimal solution. This would tend to reduce the acquisition costs, instead of increasing them as the VCG auction does. The following are desirable features of penalty systems:

1. Avoid under-penalizing:
 - reduce the incentive to specify misleading constraints
 - recover the increase in social cost from ramp constraints




Figure 3. Feasible Region with full Generator B ramp constraint

2. Avoid over-penalizing:
 - more than the corresponding profit increase
 - more than the corresponding social cost increase
 - so much that a bidder's costs are not recovered
3. Quickly computable
4. Transparent
5. Unaffected by multiple optimal dispatch solutions

To these ends, we offer two penalty system proposals:

1. A penalty based on LP sensitivity theory (dual variables and right-hand-side (RHS) ranges)
2. A penalty based on re-optimizing without each bidder's ramp constraints in turn.

These systems are focused on the problem of the increased social cost; they do not directly address the potential profit increase, though we will investigate their effects on this.

3.1. Duality-Based Penalties

The first option calculates penalties as follows: the full auction LP is solved, and dual variables for each active ramp rate constraint are calculated the usual way. Also, it is simple to extract for each ramp constraint a range in which the RHS can vary without changing the optimal LP basis. If a particular active ramp constraint has dual variable λ and allowable RHS change δ , we propose that a financial penalty of $\lambda \cdot \delta$ be imposed on the company that specified the ramp

constraint. In our example above, the dual variable for Generator B's ramp constraint is 120 \$/MW, and the allowable RHS range is 1 GW, for a penalty of \$120000. This reduced B's income from \$480000 to \$360000, exactly its value in the solution without the ramp constraint. In this case, Generator B has no incentive to specify a misleading ramp constraint, unless it is in league with Generator A (who benefits from B's ramp constraint and escapes any penalty payments).

Unfortunately, it is not hard to find examples where the penalty does not exactly compensate for the shifted profits. In a few cases, the penalty is too much; this tends to happen when the ramp constraint would be violated in only one period, but after adding the constraint for all periods (as it natural) there are two periods where it is binding, so penalties are charged for both periods. In other cases, the penalty is not enough, so that a company still profits by giving a misleading ramp constraint. This can happen when the ramp constraint chops off too many corners from the feasible region. That is, our penalty system is based on the idea of Fig. 2, where the ramp-constrained optimum is adjacent to the optimum without the ramp constraints. It is this adjacency that determines how large the RHS-ranges are. It is a matter of coincidence in the costs that this example works well even with the full ramp constraint, where the ramp-constrained optimum is not adjacent to the original optimum. To avoid this situation, we consider next dropping each bidder's ramp constraints in turn.

3.2. Cost-Difference Penalties

In this penalty system, we run the auction with all ramp constraints, and save that solution. For each bidder with active ramp constraints, we re-optimize having relaxed its ramp constraints. The difference in social cost between the two solutions is the penalty to the bidder. We use the social cost difference, rather than the acquisition cost difference, because social cost (as revealed by the offers) reflects the true cost of constraints and furthermore, the duality-based penalties are also based on social cost. Also, since our optimization minimizes the social cost, alternate optimal solutions can cause larger differences in the acquisition cost than bidders should be penalized for. This approach requires as many optimizations as there are bidders with active ramp constraints. However, each optimization starts with an easy known feasible point (the optimal solution with all ramp constraints imposed). Thus the calculation of penalties is not as computationally intensive as premium computation in a VCG auction, where each new optimization (with an offer removed) starts without a known feasible point.

By dropping ramp constraints and re-optimizing, our penalties are not restricted by other adjacent vertices of the

feasible region, so in most cases penalties will be at least as large as the duality-based penalties (the exceptions being when duality-based penalties over-penalize because of multiple active constraints). This makes the auction more incentive-compatible. However, it is possible that a penalty might be so large that the penalized generator will end up with a deficit. Indeed, the penalty might be even greater than the income itself. In either of these cases, we can assume that the penalized company would want to withdraw its offer. We would then exclude that company and re-optimize (the question remains whether to exclude all such companies simultaneously, or one at a time starting with the worst-off). Such exclusion, however, will drive up social and acquisition costs since eliminating an offer amounts to adding a constraint on the optimization. This happens with the duality-based penalty as well, but not as often, since duality-based penalties are typically smaller. In either case, there is a tradeoff between reducing acquisition cost by imposing penalties, and increasing cost by forcing out offers through excessive penalties. Another problem might occur if an offer that is forced out by a large penalty is needed for reliability reasons. If this happens on a continuing basis then a Reliability-Must-Run (RMR) contract could be enacted, but it would be harder to deal with if it happened only occasionally.

4. Results for CalECo, California 2001

We tested our suggested penalty approaches using generator data from the CalECo system of [4, 5] which represents a scaled abstraction of the California power system developed for the purpose of evaluating production simulation models. The details of the CalECo system are provided in the cited references and will be omitted here due to space limitation. The demand data, for 330 days of 2001, was collected from the California ISO web site by Matt Schneider. Because the maximum output of the CalECo generators (12500 MW total) is not on the same scale as the current California demand data, we scaled the demand data to create three data sets in which the scaled annual peak load (originally 41244 MW, on August 7 at 4pm) left reserves of 1.5%, 5%, and 15%. To isolate the effects of the ramp constraint mechanisms, we assumed that no bidders specified ramp constraints except for the “gas3” bidder, who specified a ramp constraint of either 150 MW/hour or 200 MW/hour (2.5 or 3.33 MW/minute), which is clearly more restrictive than a real gas-fueled generator would face. This constraint was imposed for every hour of every day.

For each day in each data set, the following was done:

1. Optimize social cost without any ramp constraints
2. Optimize social cost with all ramp constraints

3. Compute duality-based penalties
4. Compute cost-difference penalties

In all cases, the nuclear, QF, and coal bidders were run at full capacity in every period, so ramp constraints for them would have been irrelevant.

To get a feeling for the data set, Fig. 4 shows the social cost throughout the year for one of the three demand levels, with no ramp constraints imposed (the other two demand levels look almost the same, except for scaling). Notice that,




Figure 4. Daily social cost for one demand scenario

due to the scaled demand, the financial figures are much smaller than one would expect in reality. For this reason, we will focus our evaluation on percentage rather than absolute changes. In Fig. 5 we show the percent that profit for “gas3” increased when that bidder specified a ramp constraint. This figure is broken out into categories by demand level and ramp constraint value (150 or 200 MW/hour). The higher percentage increases tended to be on days without much original profit, though. Tables 5 and 6 summarize Fig. 5, showing the percent of days with a profit increase and the average percentage increases for these days (weighted by profit amounts). We focus only on the days with an increase because a bidder that uses this scheme would try to carefully choose when to specify the constraint, so as to avoid days on which specifying a ramp constraint would result in reduced profit.

Up to this point, we have described the effects of misleading ramp constraints. Now we turn to the effects of the proposed penalty schemes. Fig. 6 shows, in percentages, the net profit increases for the various scenarios, plotted against

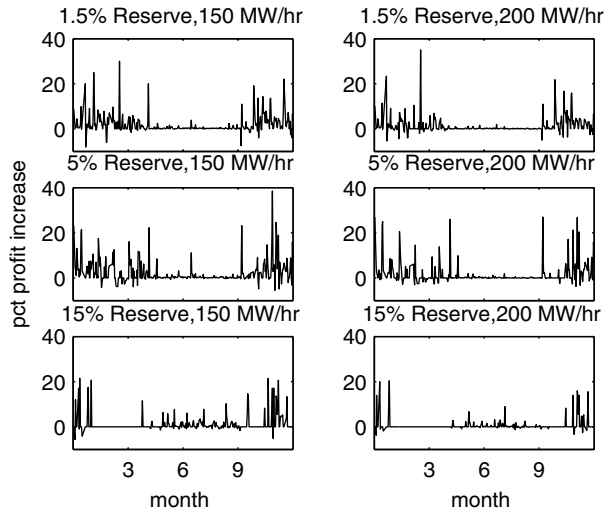


Figure 5. Daily profit increase (in percent) from no-constraint to constrained

Table 5. Percent of days with increased profits

Reserve	150 MW/hr	200 MW/hr
1.5 %	75%	53%
5 %	63%	38%
15 %	28%	16%

the gross increases, with penalties calculated from the LP dual. Fig. 7 similarly shows the results from the cost-based penalties. Ideally, these graphs would be scattered along a horizontal line through zero. Unfortunately, we see that this is not the case. Also, there are days where the gross increase is negative, and then a further penalty is applied. This is acceptable to the extent that our aim is to recover increases in social cost due to ramp constraints.

Fig. 8 shows, for duality-based penalties, the percentage of the increase in social cost that the penalties recovered, as the year goes on. We would hope that the values would concentrate around 100 percent. For the cost-based penalties, the amount recovered is always exactly 100 percent, when only one bidder specifies a ramp constraint. It is not clear that this would be true when more than one bidder has an active ramp constraint; the amount recovered might be less than or greater than the total social cost difference.

Table 6. Percent profit increase (when positive)

Reserve	150 MW/hr	200 MW/hr
1.5 %	0.45 %	0.45%
5 %	0.66 %	0.71%
15 %	0.76 %	0.67%




Figure 6. Net (after-penalty) vs. Gross percentage increase in profits, dual penalties

5. Variations on Spain's System

Spain's electricity market rules favor a heuristic solution procedure, rather than a process based on mathematical programming. The market rules for Spain's system include the following restriction [2, pg. 28]:

In any case, when the owner of a production unit which includes the rising/start-up or descending/stop load gradient condition in an electric power sale offer, the market operator shall assign the producer a lower quantity of power than the latter would have received if it had not included the cited condition.

This prevents a company from gaining business by specifying ramp constraints, which is our aim, without specifying penalty payments. Inspired by this rule, we consider some variations: when a ramp constraint is specified, the constrained dispatch must be compared to the unconstrained dispatch in one of these four ways:

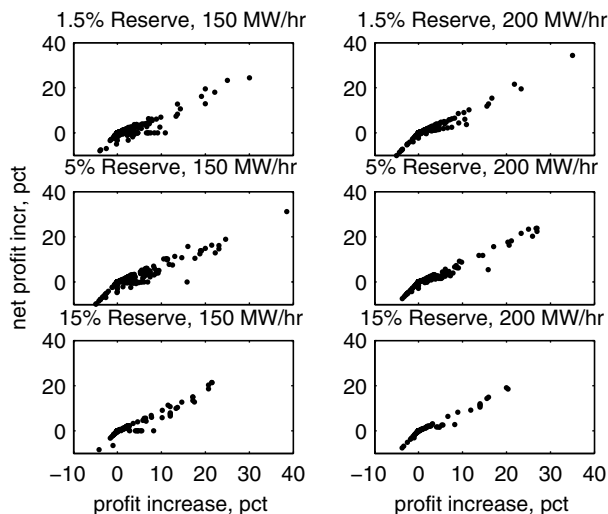


Figure 7. Net (after-penalty) vs. Gross percentage increase in profits, cost penalties




Figure 8. Percent of social cost increase that is recovered by dual-based penalties

1. The power quantity cannot increase in any hour, or
2. The sum of power dispatches over the day (i.e., energy dispatch) cannot increase, or
3. The income cannot increase in any hour, or
4. The total income for the day cannot increase.

The first two are fairly easy to implement as simple linear constraints once the initial dispatch (without ramp constraints) is obtained. The first becomes a set of 24 constraints, and the second becomes a single constraint. The third and fourth variations are much harder to implement, since the income in any period is the product of the market-clearing price and the dispatch quantity, and so is a nonlinear term. Furthermore, we have to add binary variables to the LP formulation to calculate the market-clearing price in this context. It is possible to eliminate the nonlinearity by noting that the market-clearing price must come from the set of offers, and creating a constraint for each combination of possibilities, but this becomes unwieldy very quickly. Overall, from the market perspective (ignoring implementation difficulties) it seems that restrictions on the total daily income make more sense than hourly incomes, due to cost differences between hours of the day. Also, income constraints seem better than MW allocation constraints, since the bottom line is profit rather than power generation (although in Spain power generation may have an indirect effect on profit due to stranded cost payments.)

While any of these four approaches sound fair, there are two other predicaments that should be considered: they can

make the problem infeasible, and they depend heavily on the initial solution. Electric power dispatch problems are notorious for having multiple optimal solutions. If the solution chosen as the initial one gives a particular company only a small allocation (of power, energy, or income), while another solution gives it a larger share, it seems unfair to restrict that company to the smaller of the two. To avoid this problem, though, we might have to optimize once for each company, trying to give it as big an allocation as possible while maintaining (near-) optimality. This would significantly increase the computational requirements of the auction process.

6. Conclusions

The duality-based penalties have the advantage of being easily computed from a single run of the auction LP, along with having the usual economic interpretations of dual variables. However, we have seen that they are not particularly effective in recovering social cost losses from generators due to strategic specification of ramp constraints. The cost-based penalties require many optimizations to be run, but they perform better on social cost recovery. Neither penalty system performs very well at removing the incentive for any particular bidder to specify misleading ramp constraints. The variations proposed for the Spanish approach are all subject to the problem of multiple optimal solutions, and the two that make the most sense financially are the most difficult to implement. Further research is needed into the problem of incentive-compatible auctions with ramp con-

straints. However, at this point it seems that in order to prevent gaming generators should be restricted in how often they can restate such constraints.

References

- [1] J. Contreras, O. Candiles, J. I. de la Fuente, and T. Gomez. Auction design in day-ahead electricity markets (republished). *IEEE Transactions on Power Systems*, 16(3):409–417, Aug 2001.
- [2] C. O. del Mercado Español de Electricidad, S.A. Electricity market activity rules, Apr 2001. Available: <http://www.omel.es/en/normativa/mreglasconadhesion.htm>.
- [3] B. F. Hobbs, M. H. Rothkopf, L. C. Hyde, and R. P. O'Neill. Evaluation of a truthful revelation auction in the context of energy markets with nonconcave benefits. *Journal of Regulatory Economics*, 18(1):5–32, Jul 2000.
- [4] R. Johnson, S. Oren, and A. Svoboda. Equity and efficiency of unit commitment in competitive electricity markets. *Utilities Policy*, 6(1):9–20, 1997.
- [5] C. Marnay and T. Strauss. Effectiveness of antithetic sampling and stratified sampling in Monte Carlo chronological production cost modeling. *IEEE Transactions on Power Systems*, 6(2):669–675, May 1991.