ISE

Industrial and Systems Engineering

Can We Prevent the Gaming of Ramp Constraints?

Shmuel S. Oren
UC Berkeley

Andrew M. Ross
Lehigh University

Report No. 03T-001

LEHIGH University

200 West Packer Avenue
Bethlehem, PA 18015

Can we prevent the gaming of ramp constraints?

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 *

January 24, 2003

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 three penalty approaches to discourage bidders from such a tactic: two based on duality theory of Linear Programming, the other based on social cost differences caused by ramp constraints. We evaluate the 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. Similarly, FERC's proposed Standard Market Design allows ramp constraints to be specified. 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

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

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

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 constraints 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.

Table 3: Auction Results with Gen. B ramp constraint

	Off-peak	Peak
Clearing Price	$10/MWh	$25/MWh
Gen. A	0 GW	1 GW
Gen. B	1	1 GW
Gen. C	0	1 GW

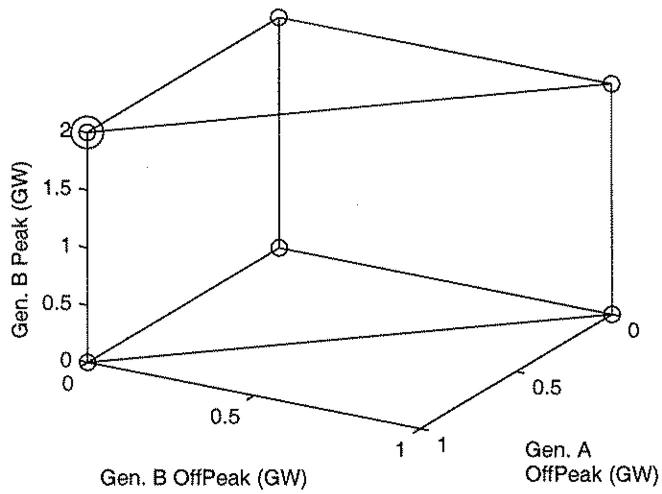

Figure 1: Feasible Region without ramp constraints

eration 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
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

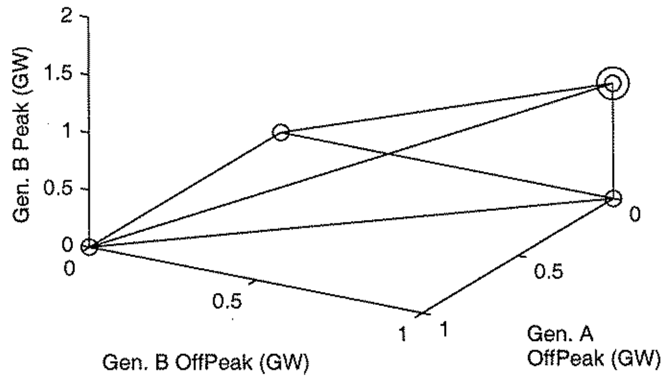

Figure 3: Feasible Region with full Generator B ramp constraint

unless it is in league with Generator A (who benefits from B's ramp constraint and escapes any penalty payments).

In Figure 4, we use the viewpoint of parametric linear programming to show the three proposed penalty systems; for now we will focus only on the first system. The figure shows how the optimal social cost would decrease as a generator's maximum ramp rate (in just one hour of the day) is increased. Suppose the current ramp rate is at the point marked a—the corresponding optimal social cost is marked with a dot. The slope of the cost curve is the dual variable λ for the ramp constraint. If the ramp rate increased to point b, the optimal basis for the LP would change, and hence the cost curve changes slope. The distance $\delta = (b - a)$ corresponds to the allowable RHS change. Thus, the PP1 penalty $\lambda \cdot \delta$ corresponds to the vertical distance shown by arrow 1.

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.

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 it 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 could happen with the PP1 penalties as well, but not as often, since PP1 penalties are typically smaller than PP2 penalties. In either case, there is a trade-off 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.

3.3 Penalty Proposal 3

Our third proposed penalty (PP3) is again based on duality theory, just like PP1. However, it is more drastic: instead of multiplying the dual variable λ by the allowable change in the RHS before the basis changes, we multiply it by the RHS change needed to relax the ramp constraint completely. That is, if the dispatch without ramp constraints had a maximum dispatched ramp rate of c (referring to 4), and the ramp-constrained dispatch has a ramp rate of a, we multiply λ by $(c - a)$ to get the penalty for this ramp constraint. On the figure, we see that this corresponds to the vertical distance between the original optimum point and the cross-hatch marked on the dotted line. Due to the convexity of the cost curve, this point is lower than the optimal social cost with the constraint relaxed. This vertical distance is shown by arrow 3.

We calculate the value of c by looking at the ramp rates implied for each hour transition by the non-ramp-constrained dispatch. This means that PP3 requires two optimizations (one ramp-constrained, the other not) whereas PP1 requires only one. However, this is still faster than PP2, which requires one optimization for each bidder. A potential downside is that PP3 is more vulnerable than PP1 or PP2 to multiple optimal dispatches, since the size of the penalty depends on c, which might vary from one optimal solution to another.

An interesting note is that the value of $c - a$ is not always positive. This can happen when the unconstrained dispatch had a large one-period jump that is then spread out by the ramp constraints to cover two or more periods. We could either ignore the penalty for that ramp constraint in those cases (in effect, let the penalty be $\lambda \cdot \max(0, c - a)$), or continue using the formula $\lambda \cdot (c - a)$. We have chosen to continue with the simpler formula, since in one case the resulting total penalty was equal to the PP2 penalty. This is an interesting result that might deserve further study, but must be left for another time.

To summarize our proposed penalty systems, we expect (from theory and from Fig. 4) that the penalty sizes will usually fall into the order PP1 $\leq$ PP2 $\leq$ PP3, where the equality cases are a common occurrence. However, we have mentioned some toy cases where PP1 is larger than expected, or where a penalty is so large that a bidder

4.1 Effects of Ramp Constraints

Before we evaluate the penalty systems, we will explore the behavior of the simulated system without the penalties. To get a feeling for the data set, Fig. 5 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

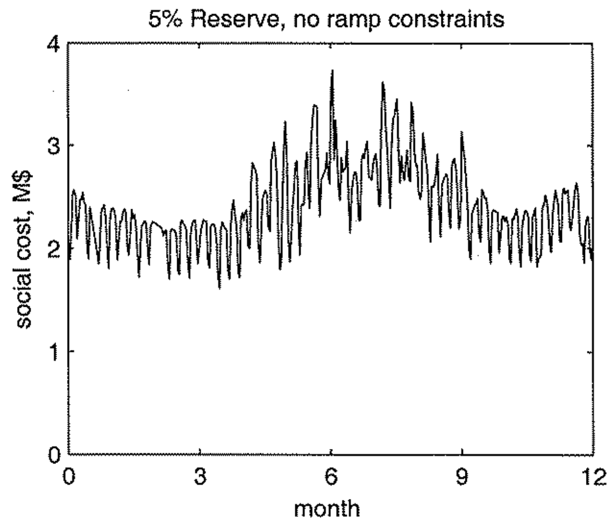

Figure 5: Daily social cost for one demand scenario

that, 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. 6 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 6 and 7 summarize Fig. 6, showing the percent of days with a profit increase and the 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

4.2 Effects of Penalty Systems

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. 7 shows, in percentages, the net profit increases (after PP1 penalties) for the various scenarios, plotted against the gross increases. Fig. 8 similarly shows the results for PP2, and Fig. 9 for PP3. Ideally, these graphs would be scattered along a horizontal line through zero (the net

Table 7: 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%

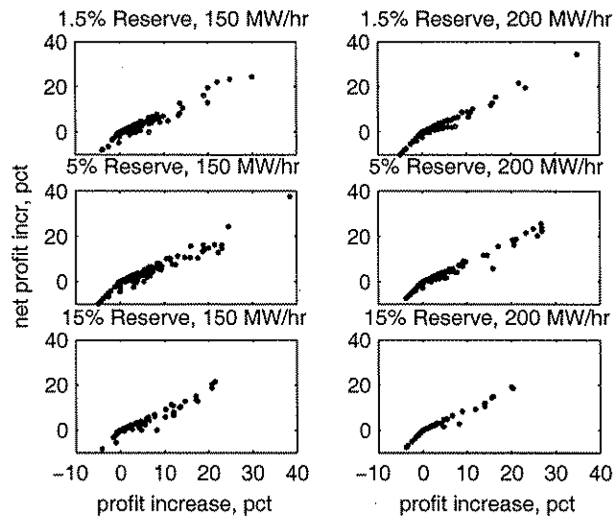

Figure 7: Net (after-PP1-penalty) vs. Gross percentage increase in profits

would prefer to back out) were not an issue in our simulated data set—it did not happen in any of our 6 scenarios (ramp rates and demand levels) as it did in a few toy examples. Also, our expectation that PP1 $\leq$ PP2 $\leq$ PP3 was true in most cases. It was always true that PP2 $\leq$ PP3, and Table 8 summarizes how many days had anomalous results for PP1. We see more anomalous results when reserves are tight, and (in all but one case) more anomalies when the ramp constraints are more restrictive.

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

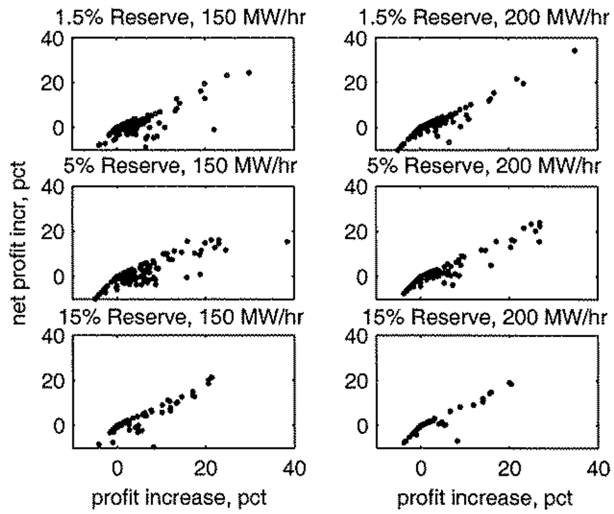

Figure 9: Net (after-PP3-penalty) vs. Gross percentage increase in profits

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.

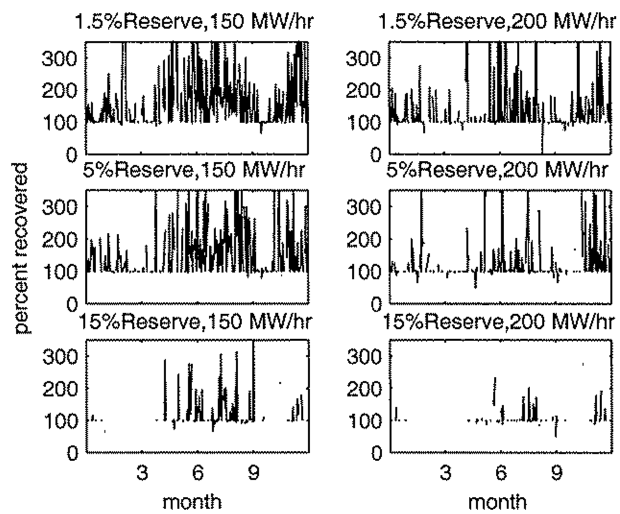

Figure 11: Percent of social cost increase that PP3 recovers

[3] Benjamin F. Hobbs, Michael H. Rothkopf, Laurel C. Hyde, and Richard 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.B. Johnson, S.S. Oren, and A.J. Svoboda. Equity and efficiency of unit commitment in competitive electricity markets. *Utilities Policy*, 6(1):9–20, 1997.

[5] Chris Marnay and Todd 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.