Electricity Market Models: Lessons from Australasian Experience

E. Grant Read and Deb Chattopadhyay
University of Canterbury, Christchurch NZ, and
PHB- Asia Pacific, Wellington NZ[1]
g.read@mang.canterbury.ac.nz

Abstract

Market clearing models based on Linear Programming are being adopted in a number of electricity markets in different countries, with New Zealand and Australia being among the pioneers in this respect. This represents an exciting development for OR, worldwide, not only because these models are controlling a sector of vital national importance, but because large sums of money are being traded, virtually in real time, on the basis of model results. Here we reflect on experience with these models, draw out some lessons, and discuss the implications of these developments for OR practice in the sector, and elsewhere.

1 Introduction

Electricity market clearing models based on Linear Programming have now been adopted in several jurisdictions, and are being developed for others. In Australia, the Victorian gas market now operates on a very similar basis. New Zealand and Australia were among the pioneers in this respect, with the New Zealand market now having been operational for over 3 years. These models are controlling a sector of vital national importance, and coordinating the activities of a large number of participants on a real-time basis. Large sums of money are also being traded, virtually in real time, on the basis of model results. Thus, these models are significant OR applications in their own right, and by any standard. But they may also have wider interest as representative examples of a trend towards further development of what we refer to as “dual OR”. As such they represent a potentially exciting development for OR, worldwide.

The concept of “dual OR” is considered in more detail by Read [13], but the general concept involves optimisation of the environment within which decision-making occurs, rather than direct optimisation of particular decisions. In particular, electricity market models take offers and bids submitted by participants and “clear the market” to produce both (primal) dispatch instructions, and (dual) prices at which electricity is traded. Indeed, some authorities have advocated dispensing entirely with traditional centralised control in favour of relying on price driven control in all time frames. At this point in

[1] The authors have acted, via PHB, in a model certification role for the Electricity Market Company (formerly EMCO, now “the Marketplace Company”, M-co), in New Zealand, and the National Electricity Market Management Company (NEMMCO) in Australia. The views expressed here are purely personal, though, and do not necessarily represent those of M-co, NEMMCO, or PHB.

time, though, all markets solve a primal dispatch optimisation problem much the same as might have been solved in a traditional centralised control environment, and that (dispatch) solution is actually implemented in real time. In the Australian market, for example, physical dispatch is restricted to the 5 minute time frame, with partial re-optimisation every 30 seconds. The prices derived from the dual solution corresponding to this optimal "market clearing" dispatch are used to provide market coordination in all other time frames.

We claim that this type of application to support market operations is distinctly different from a traditional, primal oriented, optimisation even though, the dispatch formulation itself may not be very different. Traditional policy modeling often had a market orientation, too, as in MARKAL, MFFS and PIES (Fishbone and Abilock [4], Murphy et al [9], Manne et al [8]). Such models often incorporated dual oriented decomposition techniques, and the hypothetical prices they produced were often more important than the primal "solutions". But these models were only intended to inform decision-making in an off-line fashion. By way of contrast, modern electricity market models are very much 'on-line', and the prices they produce are not at all hypothetical, being the primary instrument whereby market participants receive both the information they need to make decisions which are consistent with those of other market participants, but the commercial incentives to do so.

Here we briefly describe the market models employed in Australia and New Zealand, and reflect on experience with them to date. Our primary goal, though, is to draw out some lessons about the significance of those models, the sometimes subtle formulation issues which arise in this context, and the implications of these developments for the way in which OR is practised in the sector, and perhaps elsewhere. Thus this paper draws on, and develops some of the themes originally canvassed by Read [11], prior to market start, and alluded to again by Read [12].

2 OR Practice in the Electricity Sector

Read [11], writing before commencement of the LP based market in New Zealand, identified and described six areas of new, or increased, OR activity, as shown in Table 1. Since that time, these area have developed as shown in Table 2:

Functional area	Illustrative example	Typical Client	Comment
Market Policy	Company and market structure	Government	New focus
Market Coordination	Energy market clearing	Market institution	New function
Market Support	Dispatch based spot pricing	Market institution	New function
Market Intelligence	Long term market equilibrium	Firm	New to this sector
Market Interaction	Reservoir release strategy	Firm	New to this sector
Cost Minimisation	River chain optimisation	Firm	New impetus

Table 1: Developing areas of OR Activity (from Read (1996))

Functional area	Developments
Market Policy	Reduced as activity as markets become established, but still important to resolve issues of market power etc.
Market Coordination & Support	Now routine functions, combined in the electricity market models described here
Market Intelligence	Now a routine function in many firms and consultancies
Market Interaction	Subject of considerable commercial and research activity, especially on bidding strategies
Cost Minimisation	Still a strong focus in competitive electricity firms

Table 2: Subsequent Developments in OR Activity

This paper does not attempt to provide a detailed update on all of the areas listed above, but focuses on electricity market models designed to perform market coordination/support functions. But some general observations may be in order.

First, the application of OR techniques to electricity market problems is by no means unique to Australasia, with Bunn [2] and Littlechild [6] reporting on similar OR activity associated with the early years of the UK market. As might be expected, though, the intense policy/research activity associated with market start, and with pioneering new approaches to electricity sector problems has abated where markets have now become established, although other parts of the world are just now experiencing strong growth in these areas, as they begin the reform process. But there are now more participants in the industry, and each must understand the LP based market clearing process, form its own judgements about industry trends, and optimise its own response. Thus it is not surprising that the industry continues to absorb a steady stream of OR graduates, and to sponsor research in the area.

One very significant trend is the increasing internationalisation of OR/economic consultancy in this sector. In order to obtain independent advice many firms must look to overseas consultancies, and many now have international affiliations due to ownership changes. On the other hand, international firms operating here have come to appreciate the value which local expertise can contribute, and there is a strong demand for that expertise in other parts of the world where market developments are less advanced. Thus, in part, the traditional sharing of ideas within a rather insular local OR community has been reduced, but replaced by a much broader network of international relationships.

In particular, local experience is making a significant international contribution with respect to the design of electricity markets, and of electricity market models. The remainder of this paper reviews experience with such models, and discusses critical issues related to their design, and audit requirements.

3 Description of Market Models

The general principle of the electricity market model is to maximise the benefit from electricity trading in the spot market, subject to satisfying the operating constraints. Within this broad principle, the specifics of market models are governed by a set of "market rules". In New Zealand, Australia, and more recently in parts of North America, electricity markets are based on LP based models. The general principle may best be described using the following simple diagram (Figure 1). The electricity market "sellers" (generators), and "buyers" (loads) at each node (or zone) submit a set of offers and bids. The market model aims at "clearing", or selecting, the best set of generation offers and load bids to maximise the benefit from trading, which is shown by the shaded area in Figure 1. In this clearing process, the physical constraints on electricity generation, consumption, transmission, and finally system security aspects must be observed. The market clearing process produces a generation/load dispatch schedule, and also a set of nodal (or zonal) spot prices.

The market clearing process is dynamic, and is performed very nearly in real time (every five minutes in the Australian market, and every thirty minutes in the New Zealand market). The real-time market dispatch and pricing process is preceded by a look-ahead pre-dispatch plan, produced by running a very similar optimisation, typically performed a day ahead.

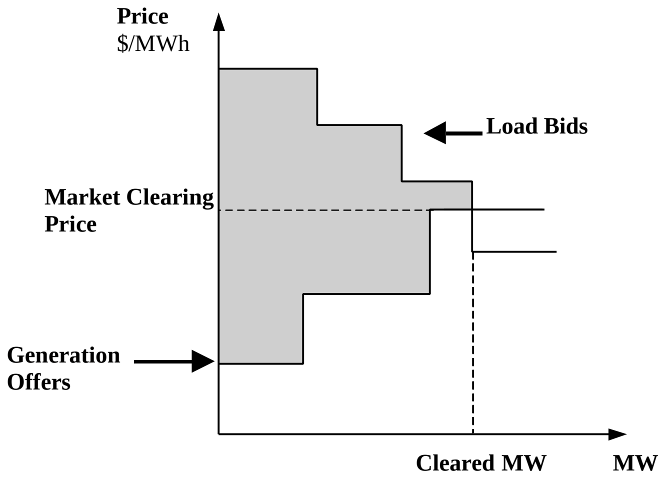

Figure 1: Overview of the market clearing process

The dispatch/pricing model is complex and may embody detailed representation of the physical and commercial realities of the power system. The formulation of such LP based dispatch/pricing models evolved over a decade starting with the pioneering work of Fred Schweppe of MIT [16]. Alvey *et al* [1] and Ma *et al* [7] describe the state-of-the-art market models that have been implemented for the New Zealand and Australian electricity markets respectively. The LP model essentially has three sets of decision variables namely, generation, transmission, and reserve. The common elements of the LP formulations may be summarised as:

Maximise Market Benefit

Subject to:

1. Zonal/Nodal energy balance constraints,
2. Zonal/Nodal reserve requirement constraints,
3. Tie-line capacity limit constraints for energy trading in normal condition,
4. Energy and reserve capacity-coupling constraints for generators,
5. Constraints on reserve response set by a piece-wise linear approximation of the governor characteristics,
6. Constraints on transfer of reserve and energy across tie-lines limited by contingency flow limit,
7. Generic network security constraints,
8. Constraints on up/down ramping ability of generators

Constraints 1-8 above try to capture the various technical and commercial requirements in the market. A detailed discussion on the specifics of each is beyond the scope of the current paper. The LP essentially attempts to approximate the complex interaction among the decision variables namely, generation, transmission, and reserve with a view to achieving a reasonable trade off between computational complexity *vs.* accuracy. Constraints 1-2 ensure that the ultimate goal of supplying adequate energy and reserve quantities are fulfilled. Constraints 3-8 represent various physical limitations in the system on provision of generation/transmission/reserve. Apart from these common elements, there is often a whole range of other physical and commercial issues specific to each individual system that are represented in the market LP. Such issues relate to, in the Australian context for example, approximate representation of unit commitment decisions for fast start generators/loads, energy availability for energy constrained units, etc.

The key "primal" outputs from these models are the generation dispatch instructions, which indirectly determine other interrelated outputs such as reserve, line flows, etc. The key "dual" outputs are the shadow prices on the energy and reserve balance constraints i.e., (1)-(2), although these prices essentially reflect the aggregate effect of all binding constraints.

4 Experience with Market Models

Perhaps the most important observation that can be made about electricity market models, after 3 years of operation in New Zealand, and 2 in Australia, is that they work. Day in and day out, every 30 minutes in New Zealand, and every 5 minutes in Australia, they produce dispatch solutions which work, in the sense that they are not only optimal, with respect to the economic considerations implicit in participant's offers and bids, but physically implementable. The lights have not gone out, contrary to some predictions, and nor has the sector been blighted by economic or financial disaster. Instead participants have come to rely on these models as a matter of course.

Which brings us to our second observation, which is that these models are far less controversial, once established in practice, than many would have predicted, prior to implementation. Certainly, there are aspects of the models which various participants would like to change, and disputes do arise from time to time. But these are actually few and far between (and certainly not the major source of consulting revenues some might have hoped for!!). By and large the principles behind the models, the model formulations, and the specific outcomes have been accepted by participants, and their bankers around the world, as the basis on which this sector now does business.

Finally, it seems that the dispatch schedules produced by these models are often better, in some respects at least, than those produced by more traditional dispatch processes. And why is this surprising? Because, theoretically, it should not be possible to produce better coordination in a decentralised environment than in a centralised environment where a single organisation has access to all the data, and can optimise all of the sectors resources without reference to anyone else. Thus, the goal of these electricity market models was never to improve upon traditional dispatch processes, but to create a market environment which allowed and "incentivised" participants to improve their performance in terms of productive and dynamic (investment) efficiency, while trying to preserve as much of the benefits of traditional coordination (allocative efficiency) as possible. In reality, it appears that, while centralised "public service" monopolies have been ideally placed to optimise operations they have often not, in fact, done so. Many of the techniques developed in universities, and published in the literature, have not actually been employed in practice, perhaps because the monopolist does not understand their potential and/or is not strongly motivated to do so. Developing a market certainly can reduce some aspects of dispatch efficiency when participants are motivated to misrepresent their true situation for commercial gain. But it can also improve some aspects, because independent commercial participants are highly motivated to ensure that the central coordinator maximises the value to be obtained from their particular plant.

5 Formulation of Market Models

A number of formulation issues have had to be dealt with in the process of developing and applying these market models. Some are simply extensions of traditional approaches to power system optimisation. These, include the inclusion of spinning reserve requirements, as in Drayton and Read [3], and the addition of a carefully graded set of penalties to ensure that, if necessary, constraints are violated in a carefully defined priority order.

Other issues relate to finding practical workarounds for the annoying limitations implicit in the use of convex optimisation. A prime example would be the fact that the piece-wise linear approximation normally applied to represent line losses becomes non-convex, and produces wildly inappropriate results, when nodal price become negative, which they legitimately can in electricity networks. Conversely, though, and leaving aside any computational difficulties, a convex representation is much preferred, because otherwise the economic theory of market clearing becomes much more complex. Difficulties have also arisen with market requirements to ensure pro-rata loading of tied offers, for example, because the optimisation is actually indifferent between any loading combination. Small penalties for deviations from pro-rata loading can be added where such ties are identifiable a priori, but this becomes virtually impossible as the formulation becomes more complex.

But, conceptually, the most important issue to emerge is the absolute necessity of a “clean” primal/dual formulation which produces both dispatch and pricing schedules which exactly, and always, match the market rules. There are a number of aspects to this, but the central point applies not only to general sloppiness with respect to formulation issues, but also to many traditional formulation “tricks” designed to ensure reliable and fast convergence to the primal solution. With most models, anything that guarantees the primal optimum is acceptable. But much more care must be taken with market models. Real examples include:

- Addition of backstop constraints to make doubly sure that unreasonable solutions will not be produced. If such constraints ever become active, they not only produce dispatch solutions but pricing effects which are inexplicable in terms of the market rules;

- Inclusion of equivalent, and redundant, versions of essentially the same constraints or variables, perhaps for notational convenience. This will not affect the primal solution, but can provide quite misleading prices, particularly when the price of a commodity is taken from the shadow price of a constraint without recognising that there may also be a shadow price on another equivalent constraint;

- Setting variables, directly or via constraints, to values known to be “correct”. This effectively shuts off any pricing signals relating to variation in those variables. In the limit, if we knew the optimal solution, we could set up constraints which left that as the only feasible point. This would produce a perfectly acceptable primal solution, but an utterly meaningless dual, since any set of prices would seem consistent with this “solution”;

- Utilising observed values of variables to simplify constraints, and perhaps pricing relationships, when those same variables are included elsewhere in the formulation. This means that the optimisation adjusts those variables without taking account of the full impact which such adjustment will really have on the solution, affecting pricing, and producing a solution which will not be optimal, and may not be feasible.

- Inclusion of network constraints in a “traditional” form, without consideration of pricing implications. Substitution is necessary to remove a natural redundancy in the equations governing flows in electricity networks. But a very large number of alternative substitutions may be made, each producing an apparently different set of constraints, and correspondingly different set of prices. Only one of these forms, and not necessarily the most obvious one, may produce prices corresponding to the market rules.

- Setting individual constraint violation penalties without regard to their combined effects. In some cases the system may be under such stress that multiple constraint violations become necessary. Where constraints are in conflict, these penalties effectively cancel, while sometimes a whole family of constraints simply reflects one underlying concern, perhaps representing a single non-linear constraint. This makes it difficult to guarantee the desired priority order for violation, and can cause unintended pricing effects, too.

6 Certification of Market Models

One distinctive OR activity associated with electricity market models is the need to certify the models prior to real-life application. One would like to think that all OR models were thoroughly tested prior to first use, but experience suggests this is, in fact, the exception, rather than the rule. Presumably, mission critical models are subjected to a very thorough testing, particularly where safety is at stake. But such testing is undertaken in-house, or by privately contracted specialists. By way of contrast, market models must undergo a very public certification process, not just to the satisfaction of the market/system operators, but of all market participants, plus their financial advisors and partners, and any relevant regulatory authorities. Ring et al [15] presented initial experience of this kind of certification process, which consists of two separate, but interdependent tasks:

First, the mathematical formulation must be reviewed to determine that it properly reflects the market rules agreed to by the market participants and regulatory authorities. In reality, this is a dual task, involving identification of those features of the formulation which have to be changed in order to become consistent with the rules, as well as those aspects of the rules which will have to be changed in order to be consistent with the capabilities of convex optimisation. This task can be considerably simplified by adopting the over-riding principle that, where possible, the rules should be interpreted in ways which are consistent with the fundamental choice to pursue an optimisation based market design. Still, establishing mutual consistency between rules and formulation is a significant task. In the Australian case, for example, the rules (NEMMCO Code, version 2.2, Chapter 3) currently consist of 113 pages of legal style clause and sub-clause, while the formulation document consists of 90 pages of equations and explanatory text. It should be understood that both documents represent the outcome of several years of development, and several successive attempts to deal with the conflicting requirements, concerns and viewpoints of various parties. Thus there are often consistency issues to address, and problems arising where, for example, detailed rules have been drafted to cover both dispatch and pricing, rather than simply allowing the model to set prices that are consistent with the primal formulation. It should not surprise anyone to hear that lawyers do not always have a good eye for primal/dual consistency or convexity, for example, or that those used to developing control-room software do not always see the commercial implications of their formulation choices. Achieving consistency can be a time-consuming process, given the lead-times involved in altering software, and particularly in gaining agreement of all parties to any rule changes. The review role also involves a significant element of facilitating a mutual education process, particularly by explaining the implications of modelling choices in terms which can be understood by both the engineering and commercial/legal fraternities, thus enabling an informed debate.

Second, the delivered software must be very thoroughly tested with a view to certifying that it correctly, and robustly, implements the agreed formulation. To date, such testing has been conducted on an entirely “black box” basis, with the testers having no direct access to the code (in this case written in AIMMS) which implements the formulation. This arrangement was originally instituted for reasons of commercial confidentiality, and creates obvious difficulties. But experience suggests that this kind of testing picks up errors which would never be apparent to someone following the apparent, but perhaps misleading, logic of the implemented formulation. It also adds discipline to the process since the tester must insist on precise documentation, rather than accepting the implemented code as a substitute. On the other hand, it should be recognised that, with hundreds of equation sets to test, it is simply not possible to test every possible combination. Thus the tester must rely on performing a thorough test on the basic model, and on each “feature” area separately, then supplement these tests with testing cases involving all plausible combinations of features. These combination cases also play a key role in “regression testing”, designed to make sure that any new version of the model continues to produce the same answers as previous versions, when applied to the same test set. In all, each new model release is subjected to 600-700 tests, and this entire test set must often be repeated several times, as bugs are found, and the model revised.

This testing process could doubtless be improved, by greater automation, for example, but it has so far proved effective in severely limiting the number of unexpected problems encountered in practice. Conversely, the number of errors found in testing indicates that the process makes a worthwhile contribution, despite the significant time and cost involved. But none of the above directly addresses the most fundamental requirement of the software… that it be absolutely reliable in producing physically feasible, and economically sensible dispatch outcomes. In part, this is supposed to be ensured by having those parties whose vital interests are at stake, and who have relevant expertise, involved in setting the rules, and then ensuring that the formulation and software be consistent with those rules. But the problem is complex, and the review/testing process has thrown up issues which require reconsideration from a practical operational perspective. Finally, though, all of the above is ultimately dependent on the reliability of a considerable data processing and communication network, both incoming and outgoing, and on data checks, manipulations, and interpretations performed outside of the model. All of these systems must be subject to its own testing processes, but particular care is required with respect to the model interface. It should be obvious that no amount of testing, or theoretical review, will avail if the data actually entered into the model does not correspond to what it is assumed to represent in the documentation, or have the assumed characteristics, such as non negativity, for example. On the other hand, the model must deal with circumstances in which it is presented with "unusual" data, such as negative offer prices, or when inconsistent data is the best available. For example, it takes around 10 seconds for the SCADA system to complete one sweep, updating all system data, but system conditions can change quite significantly in 10 seconds, if a breakdown occurs for example, so that data collected at the beginning and end of a sweep are inconsistent. Considerable effort is expended, during the testing process, to ensure that the model at least behaves predictably and robustly when presented with "inappropriate" data. But "garbage in" will still produce "garbage out".

7 Conclusions

The current generation of electricity market models represent an important class of OR application, and present some new challenges and opportunities to the OR community. As noted, the most important lesson from experience to date is that these models work, to the general satisfaction of a wide variety of participants, and their advisors drawn from many backgrounds. This diversity, of itself, presents a significant challenge to the OR practitioner, who must not only take responsibility for ensuring that "optimisation" is performed appropriately, but simultaneously deal with the technical/operational, data/communications and legal/commercial/political implications of any proposal, with a very strong overlay of economic policy. These requirements are focussed on the review/testing role described above, and on the prior role of designing market rules, and market software. Success requires a broad perspective which accepts a diversity of contributions, along with clear and appropriate communications, but tempered by a requirement to rigorously challenge and test both the intent and the delivered outcome of each aspect of any proposal. We suggest that these same lessons will also apply in other similar contexts.

References

[1] T. Alvey, D. Goodwin, X. Ma, D. Streiffert, and D.I. Sun " A Security-Constrained Bid-Clearing System for New Zealand Wholesale Electricity Market," *IEEE Transactions on Power Systems*, Paper # PE-923-PWRS-2-06-1997, July, 1997.

[2] D.W. Bunn "Evaluating the effects of privatising electricity" *Journal of Operations Research Society*, 1994, vol. 45, pp. 367-375.

[3] G.R.Drayton-Bright and E.G. Read "Using LP to Form a Market for Spinning Reserve", ORSNZ Proceedings, 1996, p119-124.

[4] L.G. Fishbone and H. Abilock "MARKAL: A Linear Programming Model for Energy Systems Analysis: Technical description of the BNL version" *International Journal of Energy Research*, Vol. 5, 1981, pp. 353-375.

[5] W.W. Hogan, E.G. Read, and B.J. Ring "Using Mathematical Programming for Electricity Spot Pricing", *International Transactions in Operations Research*, vol. 3, no.3-4, 1996, pp. 243-253

[6] S. C. Littlechild "Operational research and regulation: theory and practice", *Journal of Operations Research Society*, vol. 47, 1996, pp. 601-611.

[7] X. Ma, D. Sun and K. Cheung "Energy and Reserve Dispatch in a Multi-Zone Electricity Market" IEEE Transactions on Power Systems, Vol. 11, No. 3, August, 1999, pp. 913-919.

[8] A.S. Manne, R.G. Richels, J.P. Weyant "Energy Policy Modelling: A survey" *Operations Research*, Vol. 27, No. 1, 1979, pp. 1-36.

[9] F.H. Murphy *et al*, "Modelling and Forecasting Energy Markets with the Intermediate Future Forecasting System" *Operations Research*, Vol. 36, No.3, 1988, pp. 406-420.

[10] National Electricity Market Management Company, *Australian National Electricity Code*, Version 2.2, Australia. (http://electricity.net.au/codrule.htm)

[11] E.G. Read "OR Modelling for a Deregulated Electricity Sector", *International Transactions in Operations Research*, vol. 3, no. 2, 1996, p. 129-138.

[12] E.G. Read "OR in the Evolving Electricity Sector: the Australasian Experience", *Keynote Speech*, *APORS Triennial Conference*, Melbourne, 1997.

[13] E.G Read "Dual OR: a Partial Paradigm for the Future", forthcoming in *ORSNZ Proceedings*, 1999.

[14] E.G. Read and B.J. Ring *Dispatch Based Pricing*, Report submitted to Transpower New Zealand Limited, 1995.

[15] B.J. Ring, E.G. Read, and A.J. Turner "The application of OR Techniques in the Certification of Software Used to Clear Electricity Markets" presented to *APORS Triennial Conference*, Melbourne, 1997

[16] F. Schweppe, M.C. Caramanis, R.D. Tabors, and R.E. Bohn, *Spot Pricing of Electricity*, Kluwer Academic Publishers, 1988.