

Is System Control Entirely by Price Feasible?

Fernando L. Alvarado, *Fellow, IEEE*
alvarado@engr.wisc.edu

Abstract— This paper revisits the possibility of controlling the power system entirely by means of price signals. It expands on notions introduced in an earlier paper and addresses several unresolved issues: problems with linear cost structures, response delays, varying costs, market power, and stability problems caused by market/system interactions. The results suggest that control by price can, in fact, be made to work with some caveats.

Index Terms—Power system control, pricing, market design.

I. INTRODUCTION

Ever since Fred Schweppe and his co-workers published their seminal paper on homeostatic control of a power system [1-4], the notion of using prices to control the system has been used in a number of ways. An important extension of the work of Schweppe was provided by William Hogan, who in 1992 introduced the concept of contract networks as a practical extension to these earlier notions because it permitted the establishment of property rights within networks and allowed (approximately) efficient prices to be determined from a dispatch that was influenced by the judgment of human operators [5]. More recently, Glavitsch and Alvarado [6] illustrated how (at least in principle) an operator could use prices to control congestion in the power system even under conditions where no information was explicitly shared by the generators with the system operator. The work by Glavitsch and Alvarado not only used prices (and prices alone) to resolve the problem of managing congestion, but further established in a theoretical setting that the system operator (who in this work was also in charge of “clearing” a real time market) could “post” prices for every node location that attained the desired objective of attaining optimal system dispatch without the need for any bids. Even after a serious disturbance, an operator could, in theory, post prices that would result in a new system equilibrium that would not only be optimal but also resolve the congestion. This was possible under the assumption that every generator would choose to operate anytime the price offered was above its marginal cost of production. Furthermore (and significantly) this work illustrated how the operator could infer and anticipate the behavior that any particular price pattern would elicit from generators prior to issuing and posting prices. This was, of course, subject to several clearly stated assumptions about the costs (there were assumed to be quadratic) and the behavior of the generators (costs were fixed over time and no market power was ever exercised).

This work was supported by the CERTS consortium under DOE Interagency Agreement DE-AI-99EE35075 with NSF.

F. L. Alvarado is a Professor at the University of Wisconsin and a Senior Consultant at Christensen Associates.

This paper revisits several concepts from [6], and addresses some dangling issues left unresolved by this earlier work:

- The requirement that cost functions be quadratic. Linear functions, although seemingly simpler, complicate the control problem because their all-on all-off characteristics. Linear costs would render control by prices jumpy at best, seemingly erratic under more extreme conditions, and completely unfeasible in some cases.
- The issue of response dynamics. Even if we assume that posting a price elicits a predictable response, attaining the new equilibrium takes time and the delays in achieving the transition can create serious operational difficulties which may include the excitation of unstable electromechanical system modes [7-8] as a result of the interaction between prices and system response characteristics.
- The possibility that generator costs may change with time faster than the operator can track them (non-stationary costs). In other words, the assumption that the operator can infer marginal costs from behavior observation relies on the assumption that costs do not change over time. However, in energy-constrained situations (such as hydro systems) or in cases where fuel costs are volatile, such assumption may be invalid.
- The possibility that generators may attempt to exercise market power and fail to respond even when the price should ordinarily induce a desired behavior [9].

One additional topic addressed in this paper is the possibility of using price signals for controlling all aspects of system operation, including such items as reactive power injection, reserve provision and other necessary system quantities. For example, real time prices may be posted for reactive power injection (and consumption), prices may also be posted in real time related to reserve requirements (although these would be a bit harder to monitor and measure than energy prices), and a price component associated with frequency (the original component in homeostatic control) may also be posted. These prices would not only vary over the course of a day depending on system conditions, but would vary by location based on system losses and congestion conditions.

II. LOCAL MARGINAL PRICING OVERVIEW

A Locational Marginal Price (LMP) at a given point in time and at a given system location is nothing more than the cheapest way by which one can deliver one MW of electricity to a particular node while from the available generators while respecting all the constraints and system limits in effect. The locational marginal prices themselves can be calculated in a variety of ways:

1. The system can be operated optimally “before” the one MW increase and “after” the 1 MW increase. The additional cost of operating the system optimally after delivering the additional MW is the LMP. This method of determining LMPs is, of course, highly impractical.
2. The LMPs can be obtained from a knowledge of “sensitivity factors” (sensitivity of constraining flows to injections) for each marginal generator and a simple calculation that establishes the cost of increasing 1 MW of production at a given location while holding the offending flow(s) unchanged. A set of two equations in two unknowns (for the case of a single constraining element) is solved for each desired LMP. It is necessary to know what elements have constrained and where are the marginal generators.
3. The LMPs can be obtained simply as the Lagrange multipliers associated with the nodal injection equations during the solution of the underlying optimization problem. This method has the virtue that it is not necessary to know where are the marginal units or which are the constraining elements ahead of time. Of course, this is not always the desired context.
4. The LMPs can also be obtained from a “transposed Jacobian” solution at a given operating point, with the limiting equations replacing the “original Jacobian” rows and columns. This is, in effect, the “adjoint network” approach for network analysis [11].

Regardless of how obtained, the LMPs create a pattern in the network that establishes the marginal cost of electricity at any system location. Although the prices may vary a great deal as a function of location, all four methods above should give the exact same LMP values provided the same conditions are used.

Once the LMPs are known, they can be used for a variety of purposes:

1. In calculations of price settlements for energy consumed or delivered, assuming that an LMP context is the agreed-upon market design.
2. In establishing price differentials among the nodes in order to settle transmission property rights settlements to the holders of Financial Transmission Rights (FTRs).
3. Or, as in this paper, as a means for posting prices in the network for the purpose of influencing the operation of the system, by taking advantage of the natural business sense of all market participants, who will see it as advantageous to operate anytime the income from operating exceeds the costs of operation.

III. CONTROL BY PRICE – KNOWN COSTS

Control by price in our context means that to increase power production at a location, you do not send a “raise” pulse to the generator. Instead, you increase the posted price - and wait for the generator(s) to respond. To increase production everywhere, you increase the price everywhere - and wait. To reduce production, you lower the price. In a more extended implementation of control by price, to get reactive power production and regulate voltage, you post a price for reactive power. To get “reserves,” you post a price for the reserves.

We assume, of course, that posting of a price at a node elicits a (non-mandatory) response on the part of every generator. Every generator (and every load, for that matter) will be free to choose whether to increase or decrease its output (or increase or decrease its consumption). If the price posted is the Locational Marginal Price and every generator responds according to their cost, optimal operation should ensue.

While such price responsiveness may not be as simple and transparent as we assume it to be, there is sufficient evidence from actual systems (PJM and New York) to suggest that the system is, indeed, responsive to price.

In this section we summarize the findings of [6]. Two situations are considered. If the costs of operation for each generator in the network are known to the system operator, then the optimal operating point can be readily obtained by solving an Optimum Power Flow problem. This solution will, in fact, determine the prices that can be sent to the market to attain the desired optimum point *under the assumption that the costs of every generator are not only known, but quadratic functions of production level*.

The key concept of [6] is that we can convert a “congested optimization” problem into an “uncongested optimization” problem with an identical solution, and we can do so without having to know what the costs of the various generators are, provided we can assume that the costs of the generators are quadratic and invariant. The concept is best illustrated first considering a case where the costs are known. Let the generation costs for generator i be:

$$C_i = a_i + b_i P_i + \frac{1}{2} c_i P_i^2 \quad (1)$$

For a lossless system with a given demand, the power balance equation is:

$$P_1 + P_2 + \dots + P_{n_g} = P_D \quad (2)$$

The Lagrangian for this optimization problem is:

$$L = \sum_{i=1}^{n_g} \left(a_i + b_i P_i + \frac{1}{2} c_i P_i^2 \right) + \lambda \left(P_1 + P_2 + \dots + P_{n_g} - P_D \right) \quad (3)$$

And the optimality conditions are:

$$\begin{aligned} b_1 + c_1 P_1 &= \lambda \\ b_2 + c_2 P_2 &= \lambda \\ &\vdots \\ b_{n_g} + c_{n_g} P_{n_g} &= \lambda \\ P_1 + P_2 + \dots + P_{n_g} &= P_D \end{aligned} \quad (4)$$

Or, in matrix form:

$$\left[\begin{array}{ccc|c} c_1 & & & -1 \\ & c_2 & & -1 \\ & & \ddots & \\ & & & c_{n_g} & -1 \\ \hline 1 & 1 & 1 & \lambda \end{array} \right] \begin{bmatrix} P_1 \\ P_2 \\ \vdots \\ P_{n_g} \end{bmatrix} = \begin{bmatrix} -b_1 \\ -b_2 \\ \vdots \\ -b_{n_g} \\ P_D \end{bmatrix} \quad (5)$$

Using “Matlab notation” this becomes:

$$\begin{aligned} \mathbf{diag}(\mathbf{c}) \cdot \mathbf{P} - \mathbf{ones}(n_g, 1) \cdot \lambda &= -\mathbf{b} \\ \mathbf{ones}(1, n_g) \cdot \mathbf{P} &= P_D \end{aligned} \quad (6)$$

If a line congests, the resulting congestion condition can be expressed as one more constraint, which in Matlab notation becomes:

$$\mathbf{S} \cdot \mathbf{P} = \mathbf{p}^{\max} \quad (7)$$

This introduces an additional Lagrange multipliers, μ . The new expanded equations at the solution point are:

$$\begin{aligned} \mathbf{diag}(\mathbf{c}) \cdot \mathbf{P} - \mathbf{ones}(n_g, 1) \cdot \lambda + \mathbf{S}^T \mu &= -\mathbf{b} \\ \mathbf{ones}(1, n_g) \cdot \mathbf{P} &= P_D \\ \mathbf{S} \cdot \mathbf{P} &= \mathbf{p}^{\max} \end{aligned} \quad (8)$$

At the heart of the method is the premise that the same solution \mathbf{P}^* can be attained by price alone if a vector $\boldsymbol{\beta}$ is added to the vector \mathbf{b} . This ‘‘congestion price adjustment’’ vector can be determined as follows:

$$\boldsymbol{\beta} = -\mathbf{diag}(\mathbf{c}) \cdot \mathbf{P}^* + \mathbf{ones}(n_g, 1) \cdot \lambda - \mathbf{b} \quad (9)$$

Once this vector is determined, it is used as a price signal that, in effect, modifies \mathbf{b} . The result is an uncongested problem that has the same solution as the congested problem:

$$\begin{aligned} \mathbf{diag}(\mathbf{c}) \cdot \mathbf{P} - \mathbf{ones}(n_g, 1) \cdot \lambda &= -(\mathbf{b} + \boldsymbol{\beta}) \\ \mathbf{ones}(1, n_g) \cdot \mathbf{P} &= P_D \end{aligned} \quad (10)$$

We illustrate the process with a numeric example. Consider the system in Figure 1. If this problem is solved ignoring the flow constraint, the solution obtained is as given.




Figure 1: Solution as an unconstrained problem.

Consider now the calculation and issuing of the price signal $\boldsymbol{\beta}$ and the solution of the subsequent optimization problem (whether the solution is attained ‘‘centrally’’ or by self-dispatch by the generators). The result is illustrated in Figure 2.




Figure 2: Unconstrained solution with locational price adjustments.

The implication of this result is the market participants responding only to posted price signals would converge to an optimal dispatch without the need to be aware of any congestion relief efforts on the part of the system dispatcher.

IV. CONTROL BY PRICE – UNKNOWN COSTS

The more interesting situation is when the costs of every generator are unknown to the system operator. If one can assume that the costs are not only quadratic, but also invariant with time, then it is possible to infer these costs from an observation of the response of generators to market prices. Reference [6] illustrates how a sequence of market observations are sufficient to establish enough information about the behavior of generators to price signals to be able to predict their behavior under any other price signal sent to them. If this is true, it then becomes feasible to operate the system by issuing the ‘‘correct’’ price signals necessary to induced the desired optimal behavior under any other set of conditions, including conditions that result in congestion of one or more lines. These price signals turn out to be none other than the locational marginal prices, although in the context used in the present paper these are not so much locational marginal prices as they are signals to control generator output. Refer to [6] for details on the precise sequence of observations necessary to infer the generator cost characteristics to such a level of detail that this becomes feasible. Suffice it to say that this is feasible.

V. PIECEWISE-LINEAR COSTS

When the cost characteristics of a given generator are purely linear (constant marginal costs), any response to a price above the marginal cost of the unit will tend toward maximizing the output of the unit.¹ For large units this can result in situations where a steady-state solution simply cannot be attained by price alone.

We illustrate the nature of this problem with a simple example. Consider the system in Figure 3. The optimal solution for this linear cost case is also illustrated.

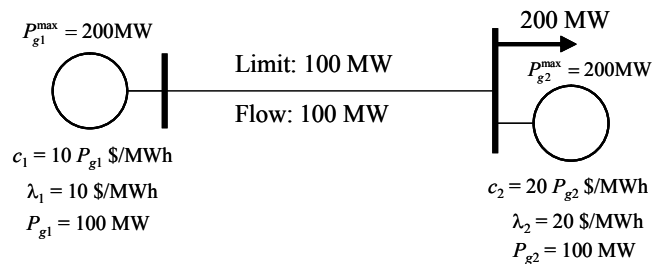



Figure 3: Linear cost case, optimal solution.

If we attempt to reach this optimal operating point by price alone, we fail. Figure 4 illustrates two possible situations, neither of which succeeds in the attempt.



(a) Case 1: This dispatch is not feasible: the line overloads.

¹ Declining marginal costs are even more cumbersome than linear marginal costs, but they are not considered here.




Figure 4: Neither price pattern leads to optimality.

VI. RESPONSE DYNAMICS

The analysis in [6] assumed only steady-state operation. The transition from any given operating state to an optimal operating state took place immediately. In practice there are sometimes significant delays within the system. There are delays in determining and posting of prices. There are also delays due to ramping rate limitations of generators.

To illustrate the nature of the possible problems associated with delays, consider the example from Figure 1, but where the load has just suddenly jumped from 100 MW to 200 MW (as a result, for example, of the loss of 100 MW of local generation, not shown). The pre-disturbance optimal dispatch corresponds to $P_{g1}=100$, no overload occurs, and both prices are equal to 10 \\$/MWh. The ultimate optimal dispatch, attainable by pricing alone, is illustrated in Figure 2. After the disturbance and before any pricing response is attained the generation/load balance is attained by the generator automatic frequency response (AGC) characteristics. Assume both generators have equal AGC characteristics. The response after the disturbance will be $P_{g1}=150$ and $P_{g2}=50$. This results in an overload which will hopefully be corrected by the response to prices. However, the occurrence of the overload may prompt the operator to post an even higher price in order to elicit the participation of some other generator that, although more expensive, may have better response characteristics. Of course, this will distort the eventual steady state, and will require that prices once again be adjusted after the transitional period is over. The ultimate result can be a sequence of over and under-corrections to account for the slowness of some generating units.


VII. WHAT TO DO?

On the surface, it would appear that control by price is doomed if either delays or linear (or worse yet, declining) marginal costs are the norm. However, the problem can be at least partially resolved as follows:


- Losses have a tendency to be vary as quadratic function of any single injection. A price signal can be sent that reflects the quadratic nature of the losses and the (usually) diminishing marginal value of injection at any single location as a function of supply. Assume, for example, that the losses associated with injections at location 1 for Figure 2 relative to the same injection at location 2 were $\text{Losses}=0.025P_{g1}^2$, then a price signal that reflected the losses would include an additive price term equal to

$\Delta\pi = 0.05P_{g1}$ which would modify the corresponding value of β . This would transform the problem in Figure 2 to the problem in Figure 1, and eliminate the issues associated with linear cost structures.

- Assume that the response time capability of each unit is 10 MW/min (linear). A price signal above the cost would elicit a “ramp up” response at this rate. A price signal below this rate would elicit a lowering of the output of a unit, at the same rate. If prices were to be adjusted every two minutes, it would be possible to attain a “near optimal” limit cycle response by an alternating sequence of prices, as illustrated in Figure 5.



(a) Price as a function of time period.



(b) Output of generators and line flow.

Figure 5: Quasi-steady state quasi-optimal response attained by time-domain price modulation.

VIII. NON-STATIONARY COSTS

Part of the way in which [6] established that the system could be steered to an optimal operating point was because the operator was able to infer generator marginal costs. That, coupled with the assumptions of a quadratic cost, no market power and an instantaneous response capability, leads to the ability to steer the system to almost any desired new condition. However, in many cases the costs of operation (whether they are actual fuel costs or whether they are opportunity costs) can vary over time. Such an effect would diminish the ability of an operator to properly predict response to a posted price and increase the difficulty associated with “control by price.”

The solution to difficulties associated with non-stationary costs is two-fold:

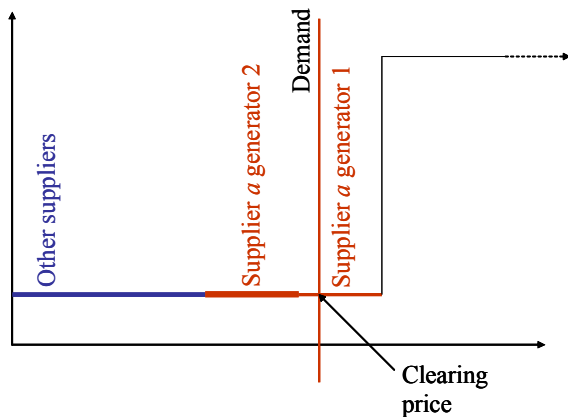
- First, a more complete model of the generator cost structure can be established and parameters to this model can be fit by market observation. Parameters that may matter in the assessment of generator response include minimum up or down times, startup and shutdown costs,

and status of the generator total energy or emissions limitations. A more complete assessment of generator parameters by market observation alone is beyond the intent of this paper.

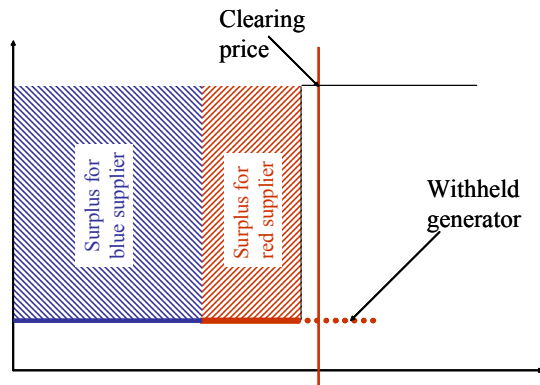
- Second, the operator may rely on feedback observations. Anytime prices are issued, it is possible to ascertain who is responding to a price pattern change and take action according to the observed degree of response. In other words, an operator would not attempt to steer the system to a new optimal point, only steer the system in the direction of a better operating point. In a sense, this is precisely what Locational Marginal Pricing is about!

IX. MARKET POWER ISSUES

In all the foregoing there has been the implicit assumption that there is no market power either possible or being exercised by any generator. In a market where generators bid prices and the operator clears the market, the exercise of market power (or the attempt to exercise market power) could manifest itself as higher price bids on the part of generators or alternatively as the withholding of generators from production. In the present framework, where there are no bids, only posted prices, market power would manifest itself as the refusal of generators from participating in production, even under conditions where the prices posted are above the marginal costs of a generator. In this section we explore some circumstances under which this would be likely to happen.



(a) Initial clearing price. No profits (no surplus).



(b) Clearing price if supplier *a* withholds one unit.

Figure 6: Prices rise and supplier surplus increases as a result of one supplier withholding output.

The most important thing to recognize is that market power is difficult to detect. Market power refers to the ability to raise prices significantly above the efficient economic equilibrium by either raising prices (economic withholding) or by withholding quantity (capacity withholding). In the case of control by price, only capacity withholding makes sense.

To see how capacity withholding can benefit a generator, consider the case of an entity owning two generating units, and privy to the information that if one of its units were to be withheld this would most likely result in a significant increase in the market price associated with the other unit. This is illustrated in Figure 6.

Detecting market power is not simple. There are essentially two approaches: (a) Simulate the system and try to predict what the price should be in the absence power, or (b) try to predict the behavior of profit-maximizing market participants that are price-takers and verify if the behavior of actual participants is consistent with this behavior. Approach (a) is quite difficult to make work. Only approach (b) seems to be practical [10]. Thus, for purposes of market monitoring it is essential to be able to anticipate optimal dispatch and bidding policy on the part of every generator. Such an effort needs to be an integral part of any “control by price” system.

X. VOLTAGES, RESERVES AND STABILITY

Consider the extension of the notion of controlling a system by price alone to the control of quantities other than energy alone. We discuss two such quantities: voltage and reserves.

The control of voltage in the system is done primarily by means of reactive power injection. Under conditions where voltages are essentially within acceptable ranges, the only significant value of reactive power is as a means of reducing system losses. However, reactive power injections can have significant value under limiting conditions. The two features that distinguishes reactive power in a control by price system are (a) its volatility is likely to be quite high, (b) its time response characteristics are likely to be much faster, and (c) it often is “lumpy” rather than continuous (this is specially the case for switched shunt system elements).

Reserves correspond, by definition, to unused capacity in generating units that can be made available on short notice upon the occurrence of system events and outages. Creating a reserve market can be done exactly as the main market on energy (it can have locational components), but it also requires that speed of response be specified and that the manner in which the activation of these reserves is to take place (perhaps the only needed signal is a drop in system frequency, but more commonly it can be expected to be explicit signals sent to the reserve generators).

The final topic that deserves mention is stability. A perfectly stable system, when connected to a perfectly rational and stable market can give rise, under some conditions, to electromechanical stabilities. Thus, care must be exercised when including price as part of an electromechanical system feedback loop. Instabilities (including electromechanical instabilities) can develop. For further details, see [7-8].

XI. CONTROL BY PRICE IN PRACTICE

One can say that an LMP system (such as the LMP system in effect in PJM or in New York) is accomplishing what this paper has described. Indeed, except for some aspects pertaining to the lack of signals for voltage support or the disregard of piecewise-linear effects or dynamic response issues, this is precisely what is being done. Real world operation seems to suggest that some of the concerns raised in this paper are not a major impediment to control by price. Both the PJM and the New York systems seem to be functioning relatively well. A different perspective on the same issue, however, would suggest that many of the present (albeit minor) problems and difficulties associated with the operation of these actual systems are, in fact, due to the very issues raised in this paper. A more thorough understanding of the nature of these problems (particularly market power issues) will eventually lead to better and smoother market designs and better operation of electricity markets.

XII. CONCLUSIONS

Control by price can be viewed as a natural and logical extension of Locational Marginal Pricing. In principle, it leads to the same optimal operating point, with the added advantage that control by price is more compatible with true free markets. Two of the main limitations of control by price (namely linear or declining marginal cost structures) can be addressed, at least in principle, with a proper design of the price signals that are sent to the market. Limitations associated with variability of costs will require better modeling of generator cost characteristics, a topic beyond the scope of this paper. Market power issues remain a large concern, but the issue is no different in a "control by price" environment than it is in a more conventional market design. Finally, the entire market and electrical control system must be analyzed as a single joint system in order to establish the stability of the feedback controls. Failure to do so can lead to operational problems. With respect to the possible extension of control by price, there seem to be no fundamental limitations about the possibility of extending control by price to also encompass voltage control and even reserves, provided appropriate signals are given in a timely manner. Market deployment of this concept will, however, require time and careful design of the precise requirements.

XIII. ACKNOWLEDGMENT

The author gratefully acknowledges Rajesh Rajaraman for discussions that led to the improvement of the paper.

XIV. REFERENCES

- [1] F.C. Schweppe, R.D. Tabors, J.L. Kirtley Jr, H.R. Outhred, F.H. Pickel and A.J. Cox, "Homeostatic utility control," *IEEE Transactions on Power Apparatus and Systems*, Vol.PAS-99, No.3, May-June 1980, pp.1151-1163.
- [2] M.C. Caramanis, R.E. Bohn, F.C. Schweppe, "Optimal Spot Pricing, Theory and Practice," *IEEE Transactions on Power Apparatus and Systems*, Vol. 109, No. 9, September 1982, pp. 3234-3245.
- [3] F.C. Schweppe, M.C. Caramanis and R.D. Tabors, "Evaluation of spot price based electricity rates," *IEEE Transactions on Power Apparatus and Systems*, Vol.PAS-104, No.7, July 1985, pp.1644-1655.

- [4] F. Schweppe, M. Caramanis, R. Tabors and R. Bohn, *Spot pricing of electricity* (Kluwer Academic) 1987.
- [5] W.W. Hogan, "Contract Networks for Electric Power Transmission," *Journal of Regulatory Economics*, Vol. 4, 1992, pp. 211-242.
- [6] H. Glavitsch and F.L. Alvarado, "Management of multiple congested conditions in unbundled operation of a power system," *IEEE Transactions on Power Systems*, Vol. 13, No. 3, August 1998, pp. 1013-1019.
- [7] F.L. Alvarado, "The Stability of Power System Markets," *IEEE Transactions on Power Systems*, Vol. 14, No. 2, May 1999, pp. 505-511.
- [8] F.L. Alvarado, J. Meng, C. DeMarco and W. Mota, "Stability Analysis of Interconnected Power Systems Coupled with Market Dynamics," *IEEE Transactions on Power Systems*, Vol. 16, No. 4, November 2001 pp. 695-701.
- [9] S.M. Harvey and W.W. Hogan, "Market Power and Withholding," December 2001. Available from ksghome.harvard.edu/~whogan.
- [10] R. Rajaraman and F. L. Alvarado, "Testing for market power in hydro-dominant regions," VIII SEPOPE, Brasilia, Brazil, May 2002.
- [11] S.A. Director and R.A. Rohrer, "The Generalized Adjoint Network and Network Sensitivities," *IEEE Transactions on Circuit Theory*, Vol. 16, No. 3, August 1969, pp. 317-323.

XV. BIOGRAPHIES



Fernando L. Alvarado (F'1998) was born in Lima, Peru, on September 15, 1945. He is a Professor of Electrical and Computer Engineering at the University of Wisconsin and a Senior Consultant at Christensen Associates. He earned a Ph.D. in Electrical and Computer Engineering from the University of Michigan in 1972, a Masters from Clarkson University in 1969 and a PE and BEE degree from the National University of Engineering in Lima, Peru. He is a fellow of IEEE. He is the Vice-Chairman of the IEEE Energy Policy Committee. He recently served as a member on a panel of experts for the study of a National Interest Transmission Grid for the US, reporting results to the Secretary of Energy in December 2001. He participated in a Presidential Task Force on Critical Infrastructures during 1997 and 1998. He has hundreds of journal publications, articles, book chapters, reports and conference presentations. He is recognized for his work on the integration of economics and electric power networks. He has developed methods for efficient trading in power networks, including a method for hedging against price uncertainty (with R. Rajaraman). He has done extensive work for CIGRE on the topic of ancillary services. He is a well-known software expert, with emphasis on large scale computation, sparse matrices and least squares methods and co-author of several software packages.