

Economic Congestion Relief Across Multiple Regions Requires Tradable Physical Flow-Gate Rights

Shmuel S. Oren, *Fellow, IEEE*, and Andrew M. Ross

Abstract—This paper is concerned with market-based protocols for relieving congestion caused by transactions outside the control area in which the congestion occurred. One approach, proposed by Cadwalader *et al.* is based on dual decomposition in which out of area congestion is “priced-out” and added to the optimal power flow (OPF) objective function of the control area operator while the prices are determined iteratively via nodal energy adjustment bids. The paper demonstrates through a simple three node example that even with “correct prices” on out-of-area congested interfaces, the augmented AC-OPF objective function of a control area operator might not be locally convex at the optimal solution and hence the control area’s optimal dispatch may violate the thermal constraints on out-of-area interfaces. That conclusion supports the alternative “flow-based” approach that enforces thermal limits more directly, which is consistent with North American Electric Reliability Council’s (NERC’s) FLOWBAT proposal for interzonal transmission load relief (TLR).

Index Terms—Duality, interconnected power systems, optimization methods, power distribution, power generation dispatch, power industry.

I. INTRODUCTION

THE proliferation of competition in the electric power industry in the U.S. and around the world and the rapid growth of interregional trading of electric power require the development of procedures for coordinating congestion management across multiple control areas. U.S. Federal Energy Regulatory Commission (FERC) Order 2000, mandating the formation of regional transmission organizations, will undoubtedly accelerate interregional trading and increase the burden of interregional coordination of transmission use. The North American Electric Reliability Council (NERC) is in the process of implementing a multistage approach to transmission load relief (TLR) protocols aimed at keeping the use of the grid within secure capacity limits. There is general agreement that such protocols should promote efficiency and hence NERC envisions a gradual transition from the current approach of administrative curtailments to a market based approach [1]. However, the rate at which such a transition should occur and the ultimate market mechanism for interregional coordination are subject to debate.

There are two main contenders for market based TLR protocols. One approach, often referred to as the link-based approach,

envisions a system of tradable physical transmission rights on the congestion prone links (flowgates). The impact of any bilateral transaction on these lines can be determined from the Power Transfer Distribution Factor (PTDF) matrix of the network. Under this scheme the TLR protocol requires that any transaction must be backed by flowgate rights in the amount of the flow it generate on each flowgate. The rights are traded in a market for transmission rights that operates in parallel with the energy market. Since counterflow on congested links enables more flow in the congested direction, counterflow should be regarded as virtual flowgate rights that can be traded as if they were physical rights. Revenues from the sales of counterflow rights will subsidize out of merit generation that helps relieving congestion. The theoretical foundation for this approach was developed by Chao and Peck [2] who have shown that the energy and transmission rights market will converge to the economic dispatch equilibrium and that the transmission rights prices on the congested links converge to the corresponding marginal values of additional capacity on these links under economic dispatch. Tabors [3] discusses an implementation scenario for a congestion management approach based on tradable physical rights that are initially awarded through an auction. An important feature of this approach is that feasibility of the flows is ensured by the number of physical rights issued for each line and the enforcement of the trading rules requiring that the appropriate rights back each transaction. The importance of this feature will become clear later in the paper. One should also point out that a rule reverting all unused rights to the system operator (within a reasonable time frame) would alleviate the classic concern about potential exercise of market power through withholding of rights. Alternatively, such withholding can be prevented by defining the transmission rights as financial rights with scheduling priority as the “Firm Transmission Rights” in California. Bundling financial rights with scheduling priority has the force of physical right but the scheduling priority can only be exercised through scheduling of a transaction and hence withholding is not possible. Further reasons to use the flow-gate approach are discussed in [4].

The other approach, due to Cadwalader *et al.* (CHHP) [5], is based on the nodal pricing paradigm in which transmission prices are calculated from nodal price differences for energy. The fundamental difference between this approach and the former is that here there is no direct trading of transmission rights and the price of transmission is determined by the energy market and adjusted through energy adjustment bids. The basic idea in this approach is to decompose the global optimal

Manuscript received November 11, 2000; revised July 17, 2001.

The authors are with the Industrial Engineering and Operations Research Department, University of California, Berkeley, CA 94720 USA (e-mail: oren@ieor.berkeley.edu; aross@ieor.berkeley.edu).

Publisher Item Identifier S 0885-8950(02)01018-0.

dispatch problem into subproblems corresponding to the different control areas. Each control area operator optimizes his dispatch by explicitly accounting for transmission constraints within his jurisdiction whereas constraints on lines he affects outside of his control area are “priced out” and accounted for as an added cost in his objective function. The prices for out-of-area transmission impact are exchanged among the control area operators through an iterative process in which energy prices and schedules are adjusted to account for the out-of-area transmission cost and each control area operator recalculates and reports the new transmission prices in his area based on the adjusted nodal prices. CHHP provide a very eloquent and thorough description of a market mechanism for implementing their proposed approach with a discussion of various practical implementation issues along with a detailed theoretical foundation. Unfortunately, as we shall see below, the CHHP approach has a basic theoretical flaw and if implemented without additional corrective measures it may result in dispatches that violate transmission constraints in neighboring regions. Such possible mishaps can occur when a control area operator uses an ac model (rather than a dc approximation) in optimizing the dispatch in its own control area. We present this effect in a simple three-node example under the assumption that the proposed market based TLR process finds the “correct” prices for transmission (i.e., the shadow prices that would result if the dispatch was optimized jointly). Our example demonstrates that even with correct prices for transmission the proposed TLR protocol can fail. One would expect that approximate prices that would result from a realistic adjustment process could make things worse.

II. REGIONAL DECOMPOSITION THROUGH PARTIAL DUALIZATION

The mathematical underpinning of the CHHP approach is the decomposition of the economic dispatch problem for the entire grid through a technique known as partial dualization or Lagrangian relaxation. The basic idea of this technique is to “price out” some of the constraints in a nonlinear constrained optimization problem and move them from the constraints set to the objective function. To illustrate this concept we consider a simple nonlinear optimization problem with two constraints

$$P1: \min_{\vec{x}} f(\vec{x})$$

$$\text{subject to: } g(\vec{x}) \leq 0 \quad h(\vec{x}) \leq 0.$$

Suppose that the vector \vec{x}^* and corresponding Lagrange multipliers (shadow prices) λ_g^* , λ_h^* represent a local solution to $P1$, i.e., they satisfy first and second order optimality conditions. We can then formulate a new problem $P2$ in which the second constraint is relaxed and moved to the objective function (“dualized”) as follows:

$$P2: \min_{\vec{x}} \{f(\vec{x}) + \lambda \cdot h(\vec{x})\}$$

$$\text{subject to: } g(\vec{x}) \leq 0.$$

It can be easily shown that if $(\vec{x}^*, \lambda_g^*, \lambda_h^*)$ satisfy first order necessary (Karush/Kuhn-Tucker) conditions for $P1$ then (\vec{x}^*, λ_g^*)

will also satisfy these conditions for $P2$ when $\lambda = \lambda_h^*$. Thus, if second order conditions are also met then (\vec{x}^*, λ_g^*) solve $P2$ for $\lambda = \lambda_h^*$. Satisfying second order conditions is assured when a dualized constraint is represented in a convex manner. In such happy circumstances it is possible to develop an iterative procedure that alternates between adjusting the multiplier λ in the objective function of $P2$ (dual iteration) and moving toward the solution (\vec{x}^*, λ_g^*) (primal iteration), as CHHP proposes.

Unfortunately, when a dualized constraint is not represented in a convex manner, second order necessary conditions for $P2$ (even with the correct multiplier $\lambda = \lambda_h^*$) may not be satisfied by (\vec{x}^*, λ_g^*) . In such a case, \vec{x}^* represents an inflection point of the objective function in $P2$.

CHHP have employed the partial dualization approach outlined above to decompose the total grid optimal power flow (OPF) problem into regional problems in which transmission constraints inside a region are considered explicitly while out-of-region constraints are priced (using market-based prices) and added to the regional objective function. This works if one assumes that all control area operators employ a DC approximation of the network in determining their own OPF and in representing out of area constraints. Indeed, the example used by CHHP is based on a DC model. However, their approach can be problematic if the control area operator employs more realistic AC-OPF models. This will be demonstrated in the subsequent sections by means of a simple three-node, two-region example.

CHHP argue that when the solution to first order conditions of the dualized problem is not a solution for the market equilibrium problem then “. . . The difficulties would extend beyond the mechanics of decomposition to call into question the existence of a competitive market equilibrium and might point to a greater role for more direct management of the grid and less reliance on markets.”¹ It is important to realize, however, that the inadequacy of the partial dualization approach which we address is not inherent in the AC-OPF problem but rather in the decomposition approach and the associated market mechanism. To demonstrate this point it is useful to consider the alternative decomposition of the OPF problem that Kim and Baldick proposed [6]. In that approach the network is decomposed into subgrids by “cutting off” transmission lines along the borders, inserting dummy busses at each end of the cut, and adding coupling constraints forcing the cuts to match. The cuts are achieved mathematically by duplicating the border variables characterizing the line flow (i.e., real power, reactive power, voltage and phase). The coupling constraints forcing the border variables to match on both sides of a cut are dualized and this decomposes the total grid OPF problem into separate regional OPFs. The multipliers on the coupling constraints that enter the regional OPF objective functions are adjusted iteratively until the border variables on both sides of a cut match. CHHP highlight the conceptual similarity of their proposal to the Kim and Baldick approach. Furthermore, they argue that the computational success of the

¹The possibility that partial dualization of transmission constraints in the AC case may result in market solutions that differ from the centralized OPF due to loss of local convexity in the dualized objective function was brought to the attention of the authors in a private communication but dismissed in a footnote which contained this quote.

Kim and Baldick algorithm supports their conjecture that their scheme should converge rapidly if price variables in other regions don't change much. However, the Kim and Baldick decomposition dualizes only linear constraints (simple equalities of the border variables); hence, adding the dualized constraints to the objective function does not affect its convexity. Then, first and second order optimality conditions for the dualized problem are satisfied by the solution to the original problem. Kim and Baldick use an augmented Lagrangian approach that adds to the dualized objective an additional quadratic penalty term on violations of the coupling constraint. Adding the penalty term is a standard method of convexifying the dualized objective function, thus preventing the problems discussed above and speeding up convergence [7]. Where Kim and Baldick augment the Lagrangian for linear equality constraints, the CHHP formulation uses nonlinear inequalities; the resulting problem does not lend itself to a distributed or market-oriented interpretation or implementation. Thus, we evaluate the CHHP system in its original form, without augmenting the Lagrangian. Reference [7] also proposes a convexification system for block-separable nonlinear inequality programs. Unfortunately, our problem is not block-separable, so we cannot take advantage of this procedure.

The remainder of this paper constructs a simple illustrative example that demonstrates in the context of a three node AC model how the CHHP decomposition scheme (which, as written, does not include an augmented Lagrangian) can lead to a severe violation of thermal constraints even when all out of area constraints are "priced" correctly.

III. AC OPTIMAL POWER FLOW MODEL

We will use the power flow model from Wu *et al.* [8]. We have n nodes, and at each one we set a phase angle θ . The flow between node i and node j is $Y_{ij} \sin(\theta_i - \theta_j)$, where Y_{ij} is called the "admittance" of the line (if there is no (i, j) line we set $Y_{ij} = 0$). The quantities q_i keep track of the net power flow at each node. If $q_i < 0$ then node i is a net consumer of power; otherwise it is a net generator of power. Associated with each node there is a cost (benefit) curve $C_i(q_i)$ expressed as a function of net output that is convex and increasing. Negative cost represents consumption benefit. Our objective is to minimize the total cost. Each power line has a limited capacity M_{ij} . All power lines are symmetric. Our basic formulation is

$$\begin{aligned} & \text{minimize} \quad \sum_i C_i(q_i) \\ & \vec{q}, \vec{\theta} \\ & \text{subject to} \\ & \forall i: \sum_j Y_{ij} \sin(\theta_i - \theta_j) = q_i \quad (\text{node flow}) \\ & \forall i, j: Y_{ij} \sin(\theta_i - \theta_j) \leq M_{ij} \quad (\text{thermal}). \end{aligned}$$

This model ignores reactive power and resistive line losses, but it captures the essence of the planning problem. Because only the differences between phases are used, there is an extra degree of freedom, so we can choose one node's phase and set it equal

to zero. We will find it more convenient to substitute out the q_i variables. We get this formulation

$$\begin{aligned} & \text{minimize} \quad \sum_i C_i \left(\sum_j Y_{ij} \sin(\theta_i - \theta_j) \right) \\ & \text{subject to} \\ & \forall i, j: Y_{ij} \sin(\theta_i - \theta_j) \leq M_{ij} \quad (\text{thermal}). \end{aligned} \quad (1)$$

Denote the optimal solution to this problem $\vec{\theta}_A$, with Lagrange multipliers $\vec{\lambda}_A \leq \vec{0}$. Even though the C_i are convex in their input q_i , here we see that the objective function is not globally convex in $\vec{\theta}$. Among other things, it has period 2π in each θ_i . However, it is locally convex near the feasible region, as pointed out by Chao and Peck [2]. The formulation (1) is equivalent to that used by CHHP, but they express the objective function and constraints in terms of the nodal injections/withdrawals q_i . While it is more intuitive to use q_i as decision variables, graphical illustration of the objective function and constraints set is more convenient in terms of θ_i , as we shall see below. There is no substantive difference between the two representations since the optimal injections/withdrawals can be readily calculate from the optimal phases through the relation $q_i = \sum_j Y_{ij} \sin(\theta_i - \theta_j)$.

A. Feasible Region

To visualize the feasible region for the power flow problem in the phase variables (θ_i) , we will use a three-node network as in Fig. 1. As mentioned above, we can set one of the phase angles to zero, so we will make $\theta_3 = 0$. We also take arcsines on each side of the thermal constraints

$$|\theta_i - \theta_j| \leq \arcsin(M_{ij}/Y_{ij}).$$

In Fig. 2, we plot the feasible regions represented by these constraints in the (θ_1, θ_2) plane. These combine to give a feasible region that will often be an irregular hexagon, but could be an irregular quadrilateral if one of the constraints is never binding. This shape repeats every 2π units in each direction, but we will ignore those other regions since they will result in the same power flows. For a network with n nodes, the feasible region is a polyhedral solid in $n - 1$ dimensions with faces that are at 45 and 90° angles to the coordinate axes.

B. Partial Dualization

In order to solve this problem across multiple regions, we partition the tie lines into two sets, which we call Mine and Other. We explicitly monitor the thermal constraints on the lines that are Mine, while we relax the constraints on the Other lines using the CHHP partial dualization approach. The new formulation is

$$\begin{aligned} & \text{minimize} \quad \sum_{\forall i} C_i \left(\sum_{\forall j} Y_{ij} \sin(\theta_i - \theta_j) \right) \\ & \quad - \sum_{(i,j) \in \text{Other}} (\vec{\lambda}_A)_{ij} \cdot (Y_{ij} \sin(\theta_i - \theta_j) - M_{ij}) \\ & \text{subject to} \\ & \forall (i, j) \in \text{Mine}: Y_{ij} \sin(\theta_i - \theta_j) \leq M_{ij} \quad (\text{thermal}) \end{aligned} \quad (2)$$

As shown, we know that any solution to (1) is a first-order stationary point of (2). But, a minimum point of (2) might not match the optimal point for (1), since the dualized constraints are not in a convex form. This difficulty is not inherent in the problem but rather in the implementation of partial dualization. For instance, the difficulty would disappear if instead of dualizing the constraints on line flow (following CHHP) we would dualize the equivalent constraints on the phases θ_i which do represent convex sets. However, that would change the interpretation of the Lagrange multipliers to prices on phase angles rather than on flows and this interpretation would not lend itself to a market-based implementation of the dual iteration.

IV. EXAMPLE

To illustrate the potential problems with partial dualization, we present a small three-node example in Fig. 1. We first consider the centralized OPF for the entire network. The two degrees of freedom in this problem are the phases θ_1 and θ_2 (in radians). Fig. 3 shows the level sets of the cost function and the feasible region for the two phase variables described by the inner hexagonal region. The optimal solution is at the lower left corner of the rectangular constraint, i.e., the smallest feasible values of θ_1 and θ_2 . The corresponding optimal injections (withdrawals) and resulting line flows are shown on the network diagram in Fig. 4. We now calculate the Lagrange multipliers corresponding to the optimal point and dualize the constraints for the (1, 3) and (2, 3) lines. This bypasses the iterative market process needed to converge on the correct price for transmission and assumes that the correct prices have been found. Pricing out the two power lines removes the rectangular constraints from the previous picture, leaving the diagonal strip as the feasible region. Fig. 5 shows the level sets of the objective function and the relaxed constraint set in the dualized problem. Since the optimal solution to the original problem is interior to the remaining constraint, it is a stationary point of the dualized objective function as expected. However, this is an inflection point (which can be identified by the cusp in the level sets) rather than a local minimum. The minimum of the dualized problem, on the other hand, has moved outside the old rectangle to the point shown by the arrow where the constraint boundary is tangent to the level set of the dualized objective function. This means that if the operator has no information about the capacity of the two transmission lines that are not under his control and is only given price information about the marginal cost of using these lines, he will choose to operate at the new optimal point that is outside the feasible region prescribed by the priced-out constraints. The implication of the new optimum (i.e., the partially dualized solution) with respect to injections (withdrawals) and flows are given on the network diagram in Fig. 6. The flows on lines (1, 3) and (2, 3) exceed their thermal limits by 60% and 92%, respectively.

V. PRICE ADJUSTMENT

In deriving the optimal solution to the partially dualized (relaxed) optimization problem we have used the “correct” Lagrange multipliers based on the optimal solution to (1). Nevertheless, one may still wonder whether further adjustments to




Fig. 1. Particular network for our counterexample.




Fig. 2. Feasible regions for the three thermal constraints.

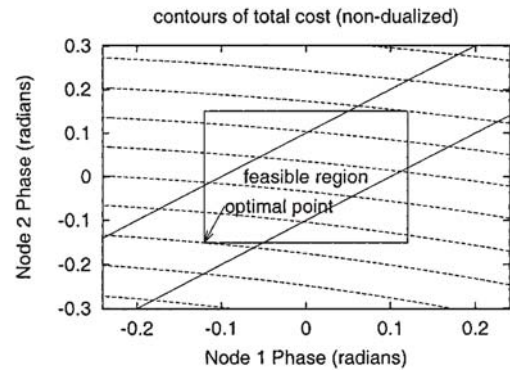


Fig. 3. Feasible region and contours of the true cost function.




Fig. 4. Optimal flow for our example.

these Lagrange multipliers can restore feasibility of the relaxed solution with respect to the violated thermal constraints. Such adjustments can be implemented in a market-based framework through two means. One is a direct adjustment to the link-based




Fig. 5. Contours of the partially dualized cost.

prices, which equal the Lagrange multipliers corresponding to the violated thermal constraints. The other is energy adjustment bids at the nodes, as proposed by CHHP. The energy adjustment bids can be obtained through a linear transformation of the adjustment to the Lagrange multipliers using the power transfer distribution factors (PTDF) corresponding to the local linearization of the constraints. Hence, without loss of generality we may examine the effect of direct adjustment to the Lagrange multipliers. In the context of Lagrangian relaxation approaches to optimization, which is what the CHHP proposal is trying to emulate through market processes, Lagrange multipliers corresponding to relaxed constraints are commonly adjusted by means of subgradient steps. A subgradient step amounts to incrementing each of the Lagrange multipliers corresponding to a violated relaxed constraint by an amount that is proportional to the violation. In mathematical terms the Lagrange multipliers are adjusted via the iteration

$$\vec{\lambda}^i = \vec{\lambda} + \alpha \cdot \vec{g}^+(\vec{x})$$

where \vec{g}^+ is the violation of the constraints (or zero if a constraint is not violated), evaluated at the current approximation to the solution—which in our case is the false optimum obtained from doing partial dualization with the correct Lagrange multipliers. In our example, since the constraint violations are large, the step size α can be very small and still have a noticeable effect.

In our example, only λ_{31} and λ_{32} are nonzero and so are the corresponding adjustments. Table I gives the values of the adjusted Lagrange multipliers corresponding to several values of α . The graphs in Figs. 7–10 illustrate the level sets of the dualized objective function resulting from using the modified Lagrange multipliers in pricing out the out-of-area thermal constraints. We also show in Table I the corresponding percentage changes to the Lagrange multipliers.

Note that the largest change is under one percent. Since, as indicated before, the corresponding adjustments to nodal energy prices are just linear transformations of the Lagrange multipliers adjustments the adjustments to the nodal energy prices will also be under 1%. The graphs illustrate that the dualized objective function is highly sensitive to changes in the Lagrange multipliers, and small adjustments to these multipliers (or equivalent nodal energy adjustment bids) will radically alter the shape of the objective function and the location of the optimal solution to the relaxed problem. In view of these sensitivities, even a computer implementation of the

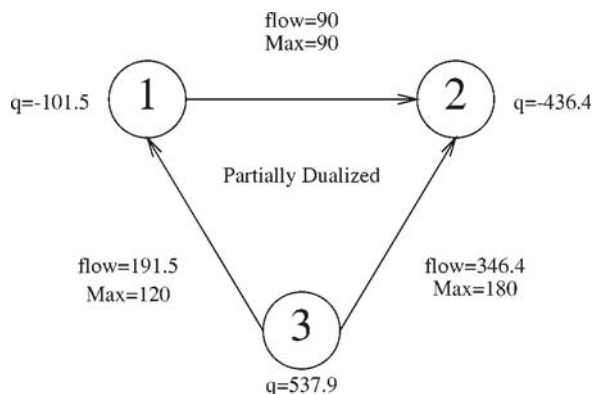


Fig. 6. “Optimal” (but actually infeasible) flows for the partially dualized problem.

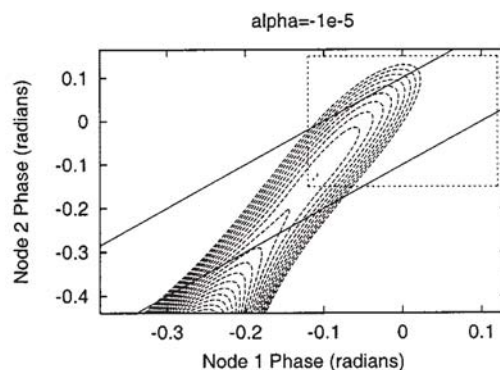


Fig. 7. Contours for $\alpha = -10^{-5}$ are almost unchanged from $\alpha = 0$. A local optimum is barely visible.




Fig. 8. Contours for $\alpha = -5 \cdot 10^{-5}$. Now a feasible local optimum is visible.

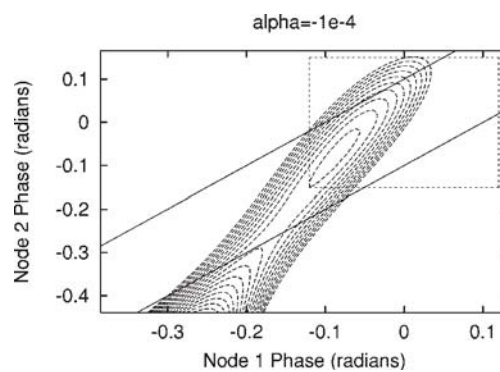


Fig. 9. Contours for $\alpha = -10^{-4}$. Optimum has moved away from true optimum.




Fig. 10. Contours for $\alpha = -10^{-3}$. Optimum has moved much farther away from true optimum.

TABLE I
PERCENT CHANGE IN λ FOR VARIOUS VALUES OF PENALTY α

α	$\bar{\lambda}^i$		percent change
0	λ_{31}	-7.2090	-
	λ_{32}	-34.7308	-
-10^{-5}	λ_{31}	-7.2097	0.00991
	λ_{32}	-34.7325	0.00479
$-5 \cdot 10^{-5}$	λ_{31}	-7.2126	0.0496
	λ_{32}	-34.7391	0.0240
-10^{-4}	λ_{31}	-7.2161	0.0991
	λ_{32}	-34.7474	0.0479
-10^{-3}	λ_{31}	-7.2805	0.991
	λ_{32}	-34.8972	0.479

Lagrangian relaxation approach with high numerical precision may face difficulties converging to a feasible optimal solution of the original problem. Achieving such convergence through a market-based implementation of the price adjustments is unconscionable.

VI. CONCLUSION

It is by now widely recognized that the physical realities of electric power systems limit the variety of market designs that will support decentralized operational paradigms and competition while respecting the secure operational limit of the system. Just as loop flow has been recognized as a phenomenon that must be contended with in the design of market-based congestion management protocols, nonlinearities due to the AC characteristics of the system must be recognized in the design of interregional TLR approaches. We have demonstrated that simple relaxation and pricing of thermal limits on out of region transmission lines can result in violation of such constraints. Furthermore, the solution is highly sensitive to the prices, so any attempt to correct such infeasibilities via market-based price adjustments will result in unpredictable outcomes. These observations raise serious questions about the viability of mechanisms that treat transmission constraints indirectly through adjustments in the energy market. This is a direct consequence of the basic conclusion indicated by our examples which is that price mechanisms without some form of direct quantity controls are insufficient to ensure feasible usage of transmission lines. By contrast, market designs based on direct trading of link-based physical transmission rights (or scheduling priorities) in

parallel with energy markets are not prone to such constraint violations. In such designs adherence to the secure limits is enforced by the number of rights issued whereas the market determines the value of these rights which are required to support energy trades. This link-based approach is in line with NERC's proposal for a gradual transition from a TLR protocol based on administrative curtailment of transactions to a market-based protocol in which flow-based rights on congested flowgates can be traded among competing users (the FLOWBAT approach).

Drawing policy conclusions based on stylized examples is always a precarious undertaking, although in the electric restructuring arena it is a way of life. Our example is by no means a representation of real electric power systems and admittedly was not easy to find. In fact, we identified this example using a systematic random search of cost functions and network parameters. It is also not clear whether in more complex networks the phenomenon we illustrated disappears or becomes more prevalent. Nevertheless, our counterexample identifies a flaw in the theory underlying the CHHP proposal and suggest caution in adopting such an approach. It also highlights a relative strength of the competing approach based on physical tradable flowgate rights, which is not subject to the same problem.

REFERENCES

- [1] NERC Transaction Reservation and Scheduling Self Directed Work Team, "Discussion paper on aligning transmission reservations and energy schedules to actual flows," NERC, Princeton, NJ, Nov. 1998. Distributed via the NERC process.
- [2] H. P. Chao and S. Peck, "A market mechanism for electric power transmission," *J. Reg. Econ.*, vol. 10, pp. 25–59, July 1996.
- [3] R. D. Tabors and R. Wilson, "Auctionable capacity rights and market-based pricing," unpublished, Apr. 21, 1998.
- [4] H. P. Chao, S. Peck, S. Oren, and R. Wilson, "Flow-based transmission rights and congestion management," *Electricity J.*, vol. 13, no. 8, pp. 38–58, Oct. 2000.
- [5] M. D. Cadwalader, S. M. Harvey, W. W. Hogan, and S. L. Pope. (1999, Oct.) Coordinating congestion relief across multiple regions. [Online]. Available: <http://www.ksg.harvard.edu/whogan>.
- [6] B. H. Kim and R. Baldick, "Coarse-grained distributed optimal power flow," *IEEE Trans. Power Syst.*, vol. 12, pp. 932–939, May 1997.
- [7] J. Eckstein, "Parallel alternating direction multiplier decomposition of convex programs," *J. Optim. Theory Applicat.*, vol. 80, no. 1, pp. 39–63, Jan. 1994.
- [8] F. Wu, P. Varaiya, P. Spiller, and S. Oren, "Folk theorems on transmission access: Proofs and counterexamples," *J. Reg. Econ.*, vol. 10, no. 1, pp. 5–23, July 1996.



Shmuel S. Oren (M'72–SM'99–F'01) received the B.Sc. and M.Sc. degrees in mechanical engineering from the Israel Institute of Technology (Technion), Haifa, and the M.S. and Ph.D. degrees in engineering economic systems from Stanford University, Stanford, CA, in 1965, 1969, and 1972, respectively.

Currently, he is Professor of Industrial Engineering and Operations Research at the University of California, Berkeley, and a former Chairman of that department. He is the Berkeley Site Director of PSerc, a multiuniversity Power Systems Engineering Research Center sponsored by the National Science Foundation and industry members. His academic research and consulting focus on planning scheduling and economic analysis of power systems, and in particular on issues concerning market design for competitive electricity. He published numerous papers in this area and has been a consultant on electricity restructuring issues to numerous public and private organizations in the United States and abroad.



Andrew M. Ross received the B.S. degree in mathematics from Harvey Mudd College, Claremont, CA, and the M.S. degree in operations research from the University of California, Berkeley, in 1996 and 1997, respectively. He is currently pursuing the Ph.D. degree in industrial engineering and operations research at the University of California, Berkeley

His teaching interests are in linear and nonlinear programming, while his research focuses on nonhomogeneous queueing systems that are common in the telecommunications industry.