

ISO White Paper

Inter-Regional Interchange Scheduling (IRIS)
Analysis and Options

January 5, 2011

ISO NEW ENGLAND AND NEW YORK ISO

INTER-REGIONAL INTERCHANGE SCHEDULING: ANALYSIS AND OPTIONS

Principal Authors:

Matthew White, Ph.D., *Senior Economist*, ISO New England
Robert Pike, *Director, Market Design*, New York ISO

Assisted By:

Christopher Brown, *Energy Market Products*, New York ISO
Ronald Coutu, *Business Architecture & Technology*, ISO New England
Ben Ewing, *Market Analysis & Settlements*, ISO New England
Shaun Johnson, *Energy Market Products*, New York ISO
Cheryl Mendrala, *System Operations Support*, ISO New England

ISO White Papers are intended to facilitate discussion of technical market, planning, and operational issues among stakeholders and ISO staff. This white paper reflects analysis and options developed jointly by the New York ISO and ISO New England. In the event a representation herein differs from provisions of an ISO's tariff, the tariff shall prevail.

ISO New England
One Sullivan Road
Holyoke, MA 01040
www.iso-ne.com

New York ISO
10 Krey Boulevard
Rensselaer, NY 12144
www.nyiso.com

TABLE OF CONTENTS

Executive Summary	iv
I. Inter-Regional Trade	
A. Interconnected Markets	I-1
B. Physical and Financial Obligations	I-3
II. Problems: Evidence and Causes	
A. Economic Inefficiencies	II-1
B. Cost/Benefit Consequences	II-12
C. Analysis of Root Causes	II-15
III. Solution Option A: Tie Optimization	
A. Concept and Design Principles	III-1
B. Economic Framework	III-4
C. Higher Frequency Scheduling	III-6
D. What About Congestion?	III-10
IV. Solution Option B: Coordinated Transactions	
A. Concept and Design Principles	IV-1
B. Economic Framework	IV-2
C. Mechanics: HFS and Submission	IV-6
D. What About Congestion?	IV-7
V. Day-Ahead Markets and Interactions	
A. Day-Ahead Market Issues	V-1
B. Day-Ahead External Transactions	V-2
C. Settlement Methods	V-4
VI. Hedging Price Risk at the Interface	VI-1

EXECUTIVE SUMMARY

In July 2010, ISO New England Inc. (ISO-NE) and the New York Independent System Operator (NYISO) commenced a joint project to evaluate the economic and operational performance of energy interchange on their interconnected transmission network. This long-term project has two phases. Phase I, from 2010-2013, seeks to improve economic coordination between the two regions' electricity markets. Phase II, from 2012-2014, will focus on coordinated congestion management and network modeling.

This White Paper is the initial Phase I report. It evaluates the performance of the current inter-regional interchange system, describes alternative market procedures that could improve this performance, and provides preliminary economic benefit estimates from these improvements. The purpose of this White Paper is to facilitate stakeholder discussion of these options, and develop consensus recommendations that NYISO and ISO-NE can refine and implement.

THE STATUS QUO

To enable physical trade of power requires an extensive set of market rules and procedures. The market monitor for NYISO and the (external) market monitor for ISO-NE, Potomac Economics, has expressed concern that the current rules governing inter-regional trade yield frequent price disparities between regions.¹ Unless the transmission network is congested, these price disparities imply low-cost generation is used too little and high-cost generation is used too much. That runs counter to the ISOs' shared objective of meeting demand at the minimum production cost.

The current inter-regional interchange system between New York and New England does not realize all of the potential benefits from trade between regions. Analysis in this report indicates that improved scheduling would produce significant benefits. Ideally, power should flow from the region with lower costs to the region with higher costs. However, the region with lower costs switches back and forth frequently—often reversing within each day. The current scheduling system cannot react quickly to these changes. As a result, at the primary transmission interface between NYISO and ISO-NE power flows in the wrong direction—from the high-priced region to the low-priced region—more than 4000 hours per year.

¹ Potomac Economics, Ltd., External Market Monitor for ISO New England Inc., [2009 Assessment of the Electricity Markets in New England](#) ("New England Assessment"); Potomac Economics, Ltd., Market Monitor for the New York ISO, [2009 State of the Markets Report for New York ISO](#) ("New York Report").

In addition, data indicate that during the remaining hours of the year there is ample transmission capacity available to move additional power from the lower-cost region to the higher-cost region. As a result, production costs would be lower if the existing transmission interconnections were scheduled efficiently.

Potomac Economics estimated the economic benefits that could have been achieved if the transmission interconnections between New York and New England were scheduled efficiently. Comparing the status quo to an efficiently scheduled system, the estimated total production cost of meeting demand in the two regions (combined) would have been lower by a cumulative \$77 million from 2006 through 2010. These production cost savings accrue to both regions.

The cost reductions would also produce lower locational marginal prices (LMPs) in each region. Potomac Economics estimates that if the transmission interface had been efficiently scheduled, loads' total energy expenditures in the two regions would have been lower by a cumulative \$784 million from 2006 through 2010. Each region's energy expenditures would be significantly lower in every year examined, with magnitudes that vary by year with fuel costs and system conditions.

SOLUTION OPTIONS

To solve the problem of inefficient tie schedules between ISO-NE and NYISO, the two ISOs established a joint design team to develop solution options and recommendations. This White Paper presents conceptual designs for two solutions: (A) Tie Optimization and (B) Coordinated Transaction Scheduling. Either of these two options would have lower production costs than the status quo and result in significant savings for load.

While other options were examined, such as maintaining the current system with increased scheduling frequency, the two solution options in this report provide the greatest potential efficiency improvement. Each option directly targets the root causes of the inefficiencies in the current inter-regional scheduling system. In addition, each solution option adheres to several key design principles:

- **MARKET-BASED SOLUTIONS.** The solution options both use competitive, market-based offers to determine the real-time schedule of energy interchange between their interconnected transmission networks.
- **ALL SETTLEMENTS AT LMP.** All scheduled energy flows between regions are priced at the LMP. This facilitates market transparency and correctly prices congestion.
- **ISOs HAVE NO FINANCIAL POSITION IN MARKETS.** The ISOs do not directly participate in the markets and do not buy or sell power. The ISOs continue to act as independent settlement administrators for the payments to and from market participants.

Both Tie Optimization and CTS employ several common elements:

- Higher frequency schedule changes across external interfaces;
- Elimination of charges/credits on external transactions that deter trade;
- Financial instruments (FTR/TCC) to hedge price risk at external interfaces.

To implement these elements, the two solution options share many operational and settlement details. However, they differ in the information they require of market participants. Conceptually, Coordinated Transaction Scheduling (CTS) is more like the current inter-regional trading system: CTS retains a role for external transaction offers to help determine real-time interface schedules between regions. In contrast, the Tie Optimization option is like the least-cost economic dispatch system used internally for each ISO's energy market: It relies on the bid-based supply offers from generators and demand resources to determine real-time LMPs and transmission flows within and between the two ISOs' networks.

SOLUTION OPTION A: TIE OPTIMIZATION

The core concept of Tie Optimization is for the ISOs to optimize their external transmission links in the same way, or as closely as possible, that the ISOs optimize transmission internally. This achieves the lowest possible production cost and efficiently uses the existing transmission infrastructure.

The concept that underlies Tie Optimization is not new. It is the same bid-based, security-constrained least cost dispatch logic that underlies the wholesale energy market administered by each ISO. This competitive market design applies to all internal nodes and internal transmission facilities today. Tie Optimization simply extends this standard market design to cover the pool transmission facilities that interconnect the two ISOs.

Operationally, Tie Optimization coordinates real-time energy dispatch across both ISOs' control areas through the exchange of load and offer data every fifteen minutes. This is made possible because of advances in communications and information technology, which allow the ISOs to implement a (near) joint energy dispatch without merging control rooms. We describe this system, called High Frequency Scheduling (HFS), in detail in Part III.

A subset of each ISOs' market participants actively trade energy across the interface today. For them, important considerations are (1) hedging (congestion) price risk at the interface, and (2) fulfilling existing contractual obligations that involve scheduling between ISO regions. To address (1), the ISOs anticipate developing financial products (TCCs/FTRs) that would provide greater hedging ability at the interface than exists today. To address (2), the ISOs would revise certain ISO-specified scheduling obligations to conform to the Tie Optimization system, simplifying the current scheduling

requirements,² and work with other market participants to handle existing contractual scheduling obligations under the new system.

SOLUTION OPTION B: COORDINATED TRANSACTION SCHEDULING

The second solution option is a package of external transaction enhancements called *coordinated transaction scheduling* (CTS). Like Tie Optimization, CTS employs higher frequency scheduling (HFS) and eliminates charges/credits on external transactions that deter trade. In contrast to today's inter-regional scheduling system, CTS features:

- A simplified bid format, called an *interface bid*, for real-time scheduling;
- Coordinated acceptance of interface bids by the ISOs, using an improved clearing rule.

Like the current external transaction system, the ISOs would use two sets of market-based offers under CTS to set real-time external interface schedules: (1) participants' external transaction offers to buy and sell across an interface, cleared against (2) the real-time generation supply stacks in each region. However, the structure of the external transaction bid format, and how it clears, differs between CTS and the existing inter-regional trading system. An *interface bid* is a unified transaction to buy and sell power simultaneously on each side of the interface. This bid structure is designed to resolve one of the root causes of the current system's inefficiencies, ensuring that transactions determining real-time flows result in a net interface schedule that moves power from the lower-cost region to the higher-cost region.

As with Tie Optimization, CTS would enable market participants that actively trade energy across the interface today to (1) hedge (congestion) price risk across the interface, and (2) fulfill existing contractual obligations that involve scheduling between ISO regions. To address (1), the ISOs anticipate developing financial products (TCCs/FTRs) that are compatible with CTS and enable greater hedging ability across the interface than exists today. To address (2), the ISOs would revise certain ISO-specified scheduling obligations to conform to the CTS system and work with market participants to handle existing contractual scheduling obligations under the new system.

RECOMMENDATIONS

The ISOs recommend the Tie Optimization option because it is the more efficient solution. This emphasis on efficiency reflects the ISOs' shared philosophy of using competitive wholesale markets to meet power demand at the minimum production

² In particular, the scheduling obligations imposed on a participant with a Capacity Supply Obligation to ISO New England associated with a resource physically located in the New York control area.

cost. The CTS system is presented as an option because it has the potential to provide significant efficiency improvements over the current system.

Tie Optimization and CTS will both enable the real-time net tie schedule to be adjusted frequently (every 15 minutes) in response to changing system conditions. This HFS procedure represents an important solution to a second root cause of inefficient schedules today: The inflexibility of current market rules to change tie schedules more frequently than hourly, in a power system where the location of the lowest-cost resource can change every dispatch interval.

The CTS system is not expected to produce as complete a price convergence between regions as Tie Optimization. The profit margin that market participants require to accept real-time price risk between regions when trading power will result in a price difference between New York and New England. That difference means the CTS system will tend to produce less efficient schedules, and higher production costs, than Tie Optimization.

With HFS and the improved clearing rule, it is possible the CTS system might yield price convergence that is nearly as efficient as Tie Optimization. The efficiency loss with CTS is difficult to quantify prospectively because the CTS bid format is a new product without clear parallel in other electricity markets today. The ISOs are actively engaged, with the assistance of Potomac Economics, in an effort to gauge whether production costs would be materially higher with CTS than with Tie Optimization.

I. INTER-REGIONAL TRADE

A. INTERCONNECTED MARKETS

New York ISO and ISO New England (“the ISOs”) are private, non-profit, regional transmission organizations that serve New York and the New England states, respectively. Each ISO operates the bulk electric power system in its region, and is responsible for the short-term reliability of the power grid. The ISOs seek to carry out these functions in an efficient, cost effective manner. To meet this objective, each ISO administers competitive wholesale markets for electricity, capacity, and related services.

Although the wholesale electricity markets operated by NYISO and ISO-NE are administered separately, the markets are interdependent. Each hour, power flows from one region to the other based on the buying and selling activities of market participants. Across the primary transmission interface between New York and New England, annual scheduled power flows are close to balanced: In 2009, 44 percent (1.6 TWh) of total interchange flowed toward New England, and 56% (1.9 TWh) flowed toward New York.

The fundamental drivers of this inter-regional trade are the variations in electricity demand in each region from one hour to the next, and generation cost differences between regions. In a modern power grid, the location of the lowest-cost generation facility able to serve demand can—and does—change from moment to moment. Accordingly, both the volume and the direction of trade between New York and New England typically change over the course of each day.

PHYSICAL INTERCONNECTIONS

The power transmission networks operated by NYISO and ISO-NE are interconnected at their border. Table I-1 lists the nine major transmission links between New York and New England. In total, these links are capable of transferring approximately 1800 megawatts (MW) of power between New York and New England under normal operating conditions. To put this in perspective, 1800 MW is approximately 12% of New England’s average power consumption in 2009, and a similar percentage, 10%, of New York’s. Thus, in any particular hour, each region could meet a significant portion of its power consumption with imports from the other.³

³ These interconnections also serve important reliability purposes, and NYISO and ISO-NE have longstanding arrangements to that effect.

Table I-1. Transmission Interconnections Between New York and New England

Interface	Network Path (From / To)	Type	Rating (kV)	Capacity (MW)
NYN	Alps NY / Berkshire MA	AC	345	Eastbound: 1400 Westbound: 1200
	Pleasant Valley NY / Long Mtn CT	AC	345	
	Rotterdam NY / Bear Swamp MA	AC	230	
	Hoosick NY / Bennington VT	AC	115	
	Whitehall NY / Blissville VT	AC	115	
	Plattsburgh NY / Grand Isle VT	AC	115	
	Smithfield NY / Falls Village CT	AC	69	
NNC	Northport NY / Norwalk CT	AC	138	100
CSC	Shoreham NY / New Haven CT	DC		330

Note: Capacity is nominal total transmission capability under normal operating conditions.

Physical transmission links between different regional grid operators are known as *interties*, or simply *ties* for short. The set of seven alternating current (AC) ties labeled “NYN” are called the New York North Interface.⁴ The NYN interface comprises the majority of the transmission capacity between the two regions, and accordingly is the focus of this report. All lines in this interface are known as “pool transmission facilities,” and the ISOs are responsible for scheduling all power flows across these transmission links.

The two remaining ties in Table I-1 operate differently than the NYN interface. The Cross Sound Cable (CSC) (NYISO: NPX-CSC) and the Northport-Norwalk (NNC) (NYISO: NPX-1385) link run between Connecticut and New York underneath the Long Island Sound. The former is a direct current line, and the latter is a controllable AC line (via phase angle regulators). These lines are used nearly all hours they are in service to deliver power to Long Island.⁵

Unlike the eight AC transmission links (NYN and NNC), the CSC is not a pool transmission facility but has merchant transmission status under FERC regulations. Parties that seek to trade electricity across it must use a different reservation system than the ISOs’ process for pool transmission facilities. This report does not evaluate, nor propose any changes to, the reservation and scheduling system governing merchant transmission facilities.

⁴ NYISO refers to the NYN Interface as the NYSIO-ISONNE interface. This report uses NYN hereafter for brevity.
⁵ See Potomac Economics, *2009 Assessment of the Electricity Markets in New England* at 16.

B. PHYSICAL AND FINANCIAL OBLIGATIONS

To enable physical trade of power requires an extensive set of market rules and procedures. These rules not only affect the physical flow of power, but also influence the cost of operating the power grid and the prices consumers ultimately pay.

Even for market participants that transact energy on a daily basis, the economic logic of the current trading system between New York and New England can be complex. In this section, we distill these market rules and procedures into their key elements and their economic purpose.

THE MAIN ELEMENTS

There are four main elements to the current inter-regional trading system:

1. **REQUESTS.** Market participants submit requests to buy or sell power at the ‘border’ separately to each ISO (e.g., a request to buy on the New England side and to sell on the New York side);
2. **ACCEPTANCE.** Each ISO independently clears the requests on its side, based primarily on economic comparisons to other requests and to the ISO’s generation supply stack;
3. **DELIVERY.** During the delivery period, each ISO dispatches internal generation so the total physical flow of power between regions matches (as closely as possible) the aggregate quantity of offers accepted by both ISOs;
4. **SETTLEMENT.** Market participants with accepted requests incur a financial obligation, as discussed below.

Each step is performed for both the day-ahead forward market and hour ahead (“real-time”) trading, except that the delivery stage in Step 3 is omitted day-ahead.

There are many procedural details associated with each of these four steps, and they differ between the ISOs. Rather than elaborate on these details here, it is useful to consider the physical and financial obligations this transaction system entails, and the roles of the ISO and the market participant in making it work.

PHYSICAL AND FINANCIAL OBLIGATIONS: WHO DOES WHAT?

Market participants refer to accepted offers to buy or sell across the interface between ISOs as *external transactions*. An external transaction is a binding financial arrangement between each ISO and the market participant. Although settlements are performed separately by each ISO, the market participant’s net gain or loss on the transaction is

ultimately this: It pays, or is paid, the difference in locational marginal prices (LMP) between each region, plus various fees. The difference is calculated for the delivery hours to which the transaction applies, and for each megawatt transacted.

For example: If the ISOs clear a market participant's offer to buy (export) from New York and sell (import) into New England in an upcoming hour, and the LMP is \$40 per MWh in NYISO and \$50 per MWh in ISO-NE during that hour, the participant's net gain is $\$50 - \$40 = \$10$ for each MWh transacted (before fees).

It is important to note that transaction requests are submitted and accepted in advance of when the power flows. (So-called "real-time" transaction offers are due at least an hour in advance, for example.) This means that when a transaction is accepted, there is uncertainty about the LMPs at which it will settle. If, continuing the example, the real-time LMP in ISO-NE (where the participant sold) turns out to be \$30 per MWh, then the participant has a net loss of $\$30 - \$40 = -\$10$ per MWh (before fees). Thus a market participant can make or lose money on the external transaction.

Now consider physical delivery. Although real-time schedule requests are nominally 'physical' energy trades, the physical delivery obligation applies only to the two ISOs. That is, a market participant arranging an external transaction between New York and New England incurs no *physical* obligation by doing so. Excepting capacity market products, a participant submitting an external transaction need not supply generation to 'match' its buy or sell request—or have any physical assets at all.⁶ If a market participant's external transaction offer is accepted, the market participant's only obligation is financial.

The distinction between financial and physical obligations is an important feature of market design for inter-regional trade. Market participants with external transactions do not dispatch resources to satisfy the interface schedule between NYISO and ISO-NE under the existing trading system, and they will not do so under either of the options presented in this report. Instead, the ISOs will continue to bear the obligation to determine the physical tie schedule by aggregating market offers, and to execute it using least-cost dispatch on each side of the border.

THE ECONOMIC PURPOSE

What economic purpose is served by allowing a market participant to buy power in one region and sell it in the other? Fundamentally, this is a mechanism to converge the LMPs between two adjacent power networks. Here's why: After all transactions are submitted, the ISOs aggregate the accepted transactions into a net tie schedule across each interface. The net tie schedule determines the direction and magnitude of power

⁶ This contrasts with capacity imports, under which a market participant incurs a physical capacity availability obligation.

flows between regions that the ISOs seek to maintain over the upcoming hour. When the net tie schedule sends power from the lower-cost to the higher-cost region, the ISOs displace an expensive generation facility in one region with a less-expensive power source in the other region. This enables the ISOs to meet demand at lower total production cost, a central ISO objective.

Consistent with this objective, the ISOs' current market rules are designed to enable market participants to earn a profit when they buy and sell in a way that converges LMPs between regions. In general, market participants take a financial loss when they do the reverse (cause price divergence). This aligns private incentives for trade with the public's interest in minimizing total production costs.

While market participants have the incentive to execute efficient trades under the current system, the current system does not produce optimal results. This occurs because of shortcomings in the current trading system's design that impede price convergence. The inability of the current scheduling system to effectively converge prices is amply evident in the data that show price convergence remains limited. This is documented in detail in Part II.

Fortunately, the current trading system is not the only way to schedule power flows between regions. A key feature of today's system is that the two ISOs use market-based information to set the net interface schedule between regions. However, there are many alternative ways to use market-based information to determine interface schedules. The ISOs believe that two specific alternatives deserve careful consideration, as they would perform significantly better than the current system and overcome many of the documented problems. These alternatives are presented in Parts III and IV.

II. PROBLEMS: EVIDENCE AND CAUSES

Since at least 2003, market monitor Potomac Economics has expressed concern that the current inter-regional trade system fails to converge prices between ISOs.⁷ The ISOs share the concern. Both the market monitor and the ISOs have recognized that inter-regional trade improvements must be prioritized in relation to other important market improvements.

Regional price differences that persist over many years are a symptom of underlying problems. Unless the transmission network is congested, price differences imply low-cost generation is used too little and high-cost generation is used too much. This runs counter to the ISOs' shared objective of meeting total demand at the lowest possible cost. It also implies loads are paying higher energy prices than necessary.

In this section, we examine data that reveal the extent of the problem. We then analyze the root causes that point to underlying flaws in the current trading system.

A. ECONOMIC INEFFICIENCIES

To evaluate how well the current trading system serves its economic purpose, we first examine price and transmission data. Two central conclusions emerge. First, in most hours of the year, there is ample transmission capacity available to move additional power from the lower-cost ISO to the higher-cost ISO.

Second, the region with lower costs switches back and forth frequently—often reversing within each day. At the primary border between NYISO and ISO-NE, each region has the lower price about equally often over the course of the year. As a result, production costs in each region would be reduced if the transmission ties between them were scheduled more efficiently.

⁷ David Patton, *Coordinated Interchange Recommendations*, March 13, 2003 (Presentation to New England RTO Working Group).

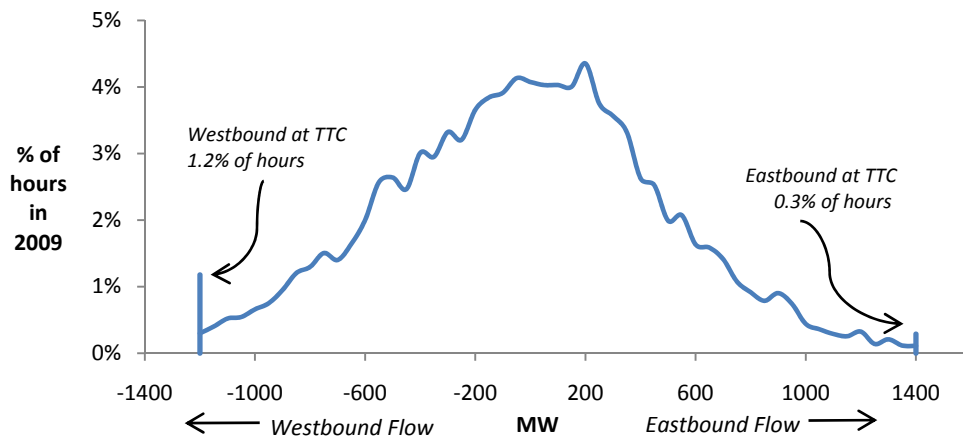


Figure II-1. Scheduled flows on the (NYN) interface between NYISO and ISO-NE in 2009.

TIE UTILIZATION

Is there unused transmission capacity across the interface between New York and New England? In brief, yes. This is evident in transmission operational data.

Figure II-1 shows the frequency distribution of hourly power flows across the primary interface (NYN) between NYISO and ISO-NE for 2009. In the figure, the height of the curve indicates the fraction of hours in 2009 in which directional flows equaled a particular MW level. For example, scheduled flows were 200 MW eastbound approximately 4% of the year.

Figure II-1 indicates that the primary interface (NYN) typically operates at far less than its total transmission capacity (TTC). For example, over seventy-five percent of the time scheduled and actual flows across the interface are between -600 MW and +700 MW (these values are half the normal westbound and eastbound normal capacities, respectively). So more than three-fourths of the year, the transmission interface operates at less than half its capacity.

We have also examined the same data for 2007 and 2008, and the results are substantially identical. The pattern shown in Figure II-1 has been stable for several years.

TIE CONGESTION

Congestion is Rare. The figure shows the primary NY-NE transmission interface is constrained very few hours per year. The height of the solid (blue) line at -1200 MW means the interface is at or near its nominal westbound total transfer capacity (TTC)

about 1.2% of the year; similarly, the interface is at or near its nominal eastbound TTC of 1400 MW only 3/10ths of one percent per year.⁸

As a technical matter, the interface's operating capacity is occasionally de-rated for reliability reasons. The underlying transmission data indicate the interface was constrained at a de-rated MW limit approximately 1 percent of the hours in 2009.⁹

In sum, transmission capacity constraints did not bind on this interface 97% of the hours in 2009. As Figure II-1 makes clear, there is ample transmission capacity to move additional power between NYISO and ISO-NE across their primary interface.

IMPORTS/EXPORTS

Balanced Imports and Exports. A second point to note in Figure II-1 is that the frequency distribution of flows is centered at about zero. The actual average hourly flow across this interface in 2009 was 42 MW westbound. This means that across their largest interface, neither New York nor New England is predominantly an exporter to the other on an annual basis.

In addition, there is little parallel path (loop) flow across the AC transmission system between NYISO and ISO-NE. This is an important difference from the loop-flow transmission coordination issues that NYISO faces on its interfaces with other ISOs, and that are being actively addressed in NYISO's Broader Regional Markets initiative.¹⁰

INEFFICIENT TIE UNDER-SCHEDULING

It is important to emphasize that, standing alone, the data in Figure II-1 do not imply the transmission system is scheduled inefficiently. It indicates there is ample capacity to move additional power—but it does not show whether it would be economic to move additional power.

To determine whether the interface is inefficiently utilized, additional data are needed. Next we examine scheduled tie flows during hours when one ISO is operating higher-cost generation than the other, and the transmission interface between them is unconstrained.

⁸ In this report we characterize an interface as congested, or operating 'at or near' its TTC, if scheduled flows are at or within 2% of the TTC applicable at the time the schedule is finalized for the upcoming operating hour. Key conclusions in this report are not sensitive to expanding the 2% tolerance (this is evident in the low incidence of flows at levels near east- and westbound nominal TTC in Figure II-1).

⁹ There are other constraints on changes in transmission flows. In particular, pool ramp constraints bound on the interface (on either the NYISO or ISO-NE side) approximately 1000 hours in 2009. These constraints do not limit the *level* of scheduled and actual flows, but restrict *changes* in flows from one hour to the next.

¹⁰ New York ISO, *Report on Broader Regional Markets*, FERC Docket ER08-1281 (January 12, 2010).

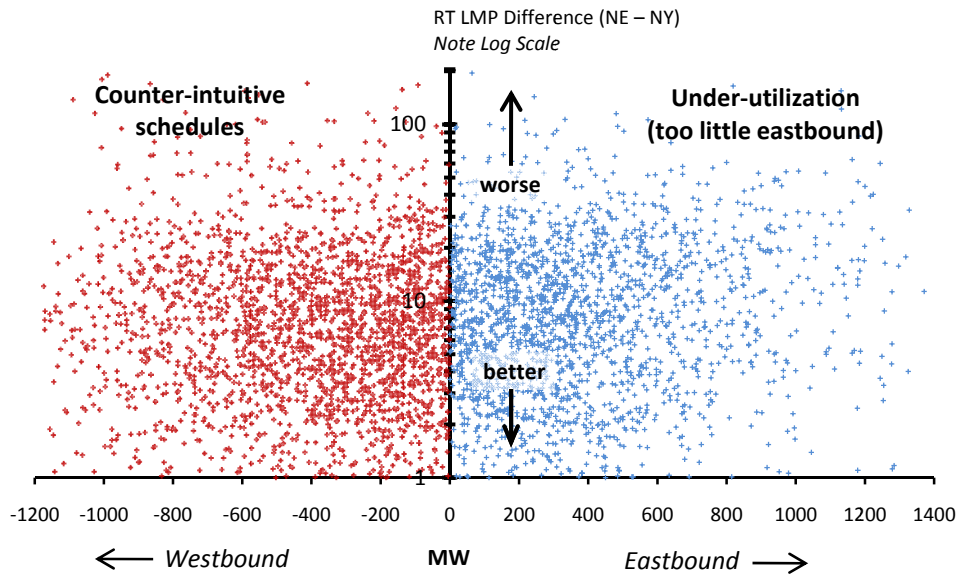


Figure II-2(a). Scheduled net tie flows and LMP differences across the primary interface (NYN) during hours in 2009 when (1) New England's LMP exceeds New York's LMP and (2) the interface TTC is not binding. Note the price difference on the vertical axis is in logarithmic scale. Prices are real-time hourly integrated LMPs for the external proxy buses (Sandy Pond and Roseton price nodes).

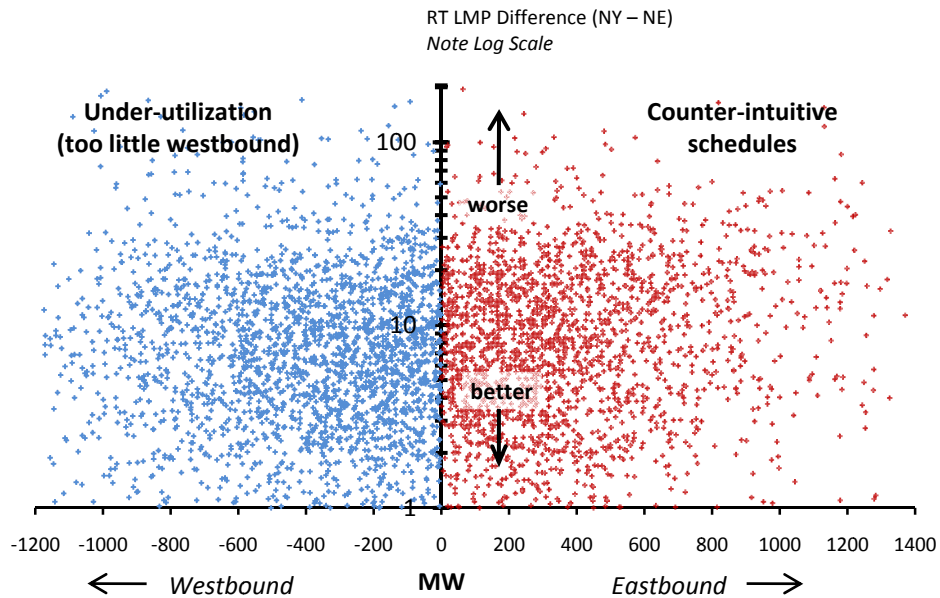


Figure II-2(b). Scheduled net tie flows and LMP differences across the primary interface (NYN) during hours in 2009 when the high-price region is reversed: (1) New York's LMP exceeds New England's LMP and (2) the interface TTC is not binding. The price difference on the vertical axis is reversed from (a), and in logarithmic scale.

We present separate charts for two situations: Hours when New England has higher prices, and hours when New York has higher prices. The first situation is shown in Figure II-2(a) (*top previous page*). Each dot in the scatter-plot represents an hour during 2009 when two conditions apply: (1) New England's RT LMP exceeds New York's RT LMP, and (2) the interface TTC is not binding.¹¹ This occurs about half of the year. The horizontal axis is the scheduled tie flow across the primary interface (NYN), and the vertical axis is the LMP *difference* between the two regions that hour (NE minus NY). Note the vertical axis is in logarithmic scale.

In the lower graph, Figure II-2(b) shows the situation during hours when (1) New York's RT LMP exceeds New England's RT LMP, and (2) the interface TTC is not binding. This occurs during the other half of the year. Note the vertical axis in the lower graph shows (the logarithm of) the price difference subtracted the other way, or NY minus NE.

The right-hand side of the top graph, Figure II-2(a), shows the hours when New England's price is higher and scheduled flows are eastbound (blue dots). This occurs 2141 hours, or about 24%, of the year in 2009. This is the economically correct direction when New England has higher prices: It displaces higher-cost generation with lower-cost generation from New York.

In the lower graph, Figure II-2(b), the left-hand panel shows hours when New York's price is higher and scheduled flows are westbound. This is the economically correct direction when New York has higher prices, as it substitutes lower-cost generation from New England for higher-cost generation New York.

UNDER-UTILIZATION

Even though trade is the economically correct direction about half the time, the data in the top-right and lower-left panels (blue points) imply the NYN interface is inefficiently under-utilized during these hours. Although flows are in the correct direction, the tie is unconstrained and the regions' LMPs remain far apart. The average price difference in these hours is \$11.82 per MWh, and price differences exceeding \$20 per MWh are not uncommon.

From the perspective of the ISOs that seek to operate the power grid in a least-cost manner, too little power is flowing in the correct direction more than 4000 hours per year. That is a true economic loss: One ISO is running higher-cost generation on the margin than the other, and society could save much of this difference if the transmission network was scheduled efficiently.

¹¹ All price data presented in this section are the 'border', or proxy bus, real-time hourly integrated LMPs for NYISO and ISO-NE (Sandy Pond and Roseton price nodes).

The economic costs of under-scheduling have a useful visual interpretation. Consider Figure II-3. Here the black curve represents the generation supply offer stack in ISO 1, which is exporting. The blue curve represents the generation supply offer stack of importing ISO 2, shown here in *descending* bid-cost order. (In practice these supply curves have stair-step shapes, but they are drawn here as smooth curves for clarity).

In Figure II-3, the (red) shaded triangle represents the excess production costs incurred because the tie is under-utilized. If exporting ISO 1 increased production by the under-utilized amount, and importing ISO 2 reduced production accordingly, total costs would decrease. The savings is the difference in their supply curves over the range of this schedule change.

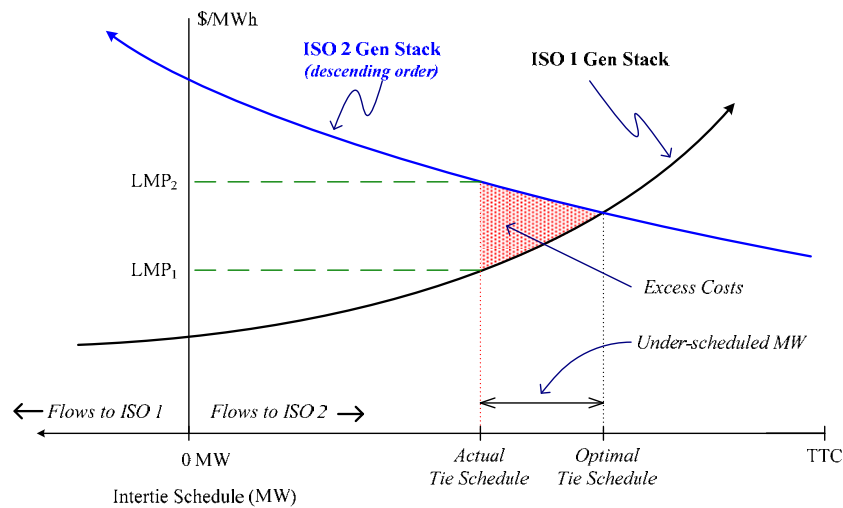


Figure II-3. The excess costs of tie under-utilization

In Figure II-3, which real ISO represents which curve? It can vary by hour. The top-right panel in Figure II-2(a) corresponds to the situation depicted in Figure II-3 with NYISO as ISO 1 (exporting) and ISO-NE as ISO 2 (importing). The bottom-left panel in Figure II-2(b) applies with the flows reversed, that is, for the situation in Figure II-3 with NYISO as ISO 2 (importing) and ISO-NE as ISO 1 (exporting).

Regardless of which direction the flows go, excess costs exist whenever the tie is under-utilized. About half the time it occurs, ISO-NE incurs the excess costs; the other half of the time, NYISO incurs the excess costs. Both ISOs' systems would operate more efficiently if these excess costs were reduced with more efficient scheduling.

COUNTER-INTUITIVE FLOW

The left-hand panel (red dots) in Figure II-2(a) and right-hand panel in II-2(b) illustrate a phenomenon known as *counter-intuitive flow* (or “wrong-way” flow). It happens a lot: Nearly half of the time that New England has higher-cost generation on the margin than New York, the net scheduled flow is *westbound* into New York. We see the same problem of counter-intuitive flow similarly often in Figure II-2(b), when New York has higher-cost generation on the margin.

IMPACT

Counter-intuitive flow has material adverse consequences. It means the net tie schedule is causing the exporting ISO to *increase* production from high-cost generation at the margin, and the importing ISO to *decrease* production from low-cost generation. This economically perverse outcome raises the total costs of serving demand, relative to a lower level of flows (that is, flows closer to zero). Total production costs would fall if flows were reduced (toward zero), up until the point where it caused the ISOs’ re-dispatch to produce equal LMPs.

INTERPRETATION

To understand the economic impact of counter-intuitive flow, consider Figure II-4. This shows the supply curves drawn to illustrate a situation with counter-intuitive flow.

Here the actual tie schedule *exceeds* the optimal tie schedule. This forces the exporting ISO to incur higher generation costs (at the margin) than the importing ISO. The flows are “counter-intuitive” in this situation because the exporting ISO’s price (LMP 1) exceeds the importing ISO’s price (LMP 2); that is, flows are from the high-cost to the low-cost region *at the margin*.




Figure II-4. Counter-intuitive flow creates excess costs

When counter-intuitive flows occur, there are excess costs. In principle, the solution to costly counter-intuitive flow is to reduce the net tie schedule (toward zero). If the tie schedule was reduced to the point where the ISO's LMPs are equal—the optimal tie schedule—then total costs would fall by the dollar amount represented by the (red) shaded triangle in Figure II-4.

WRONG WAY CASE

One other situation merits note. It is also possible for counter-intuitive flows to occur if the actual tie schedule is on the *opposite side of zero* from the optimal tie schedule. This can produce large excess costs, relative to the optimal tie schedule. In Figure II-2(a) and (b), the hours shown in the top-left and bottom-right panels (the red dots) include both the situation illustrated in Figure II-4, and any hours in which the schedule has the wrong direction entirely. Together, these counter-intuitive flows occur nearly half the year in 2009.

WHY IT OCCURS

In practice, counter-intuitive flow can occur for a number of different reasons. One problem is changes in system conditions that can lead LMPs to reverse between the time transactions are submitted and accepted (typically an hour before execution) and the time real-time prices are determined (almost two hours later). The inability of market participants or the ISOs to alter the net tie schedule in response to real-time price reversals can result in counter-intuitive flow.

The second problem is that market rules currently allow a market participant to submit schedule requests to intentionally 'buy high, sell low' across the interface. This directly produces counter-intuitive flows, unless a greater number of megawatts are accepted (or "clear") in the opposite direction.¹² We discuss this problem in greater detail in section II.C, below.

A third problem arises due to the uncertainty associated with real-time uplift (NCPC/BPGC) charges. If a market participant schedules a transaction day-ahead between regions, but the region with the lower costs reverses the next day operating day, the market participant may choose not to deviate from its day-ahead schedule in order to avoid real-time deviation (balancing) charges. This leads to counter-intuitive flow in real-time.

¹² Market participants point out they may intentionally submit counter-intuitive (high-to-low price) schedule requests to fulfill other contractual obligations. For example, ISO-NE's capacity import rules can require this, as may terms of certain renewable energy credits.

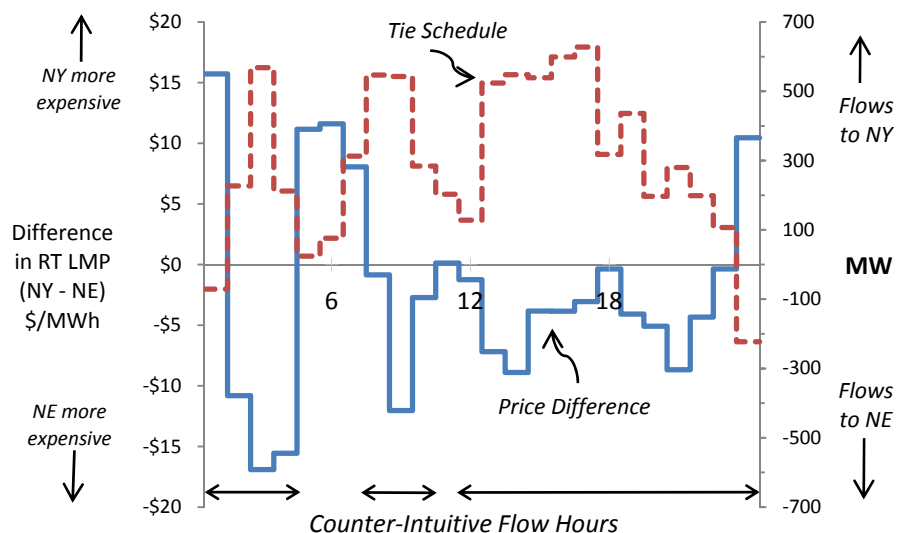


Figure II-5. Counter-intuitive schedules continue for hours on end (June 16, 2009)

EXAMPLE

Counter-intuitive flows are not only associated with rapid price reversals across the interface. They can, and do, persist for hours on end. Figure II-5 (*next page*) provides an example for one day (June 16, 2009). This was an ordinary operating day for both ISOs, with LMPs rising in their usual pattern that summer from overnight in the \$10-20/MWh range to afternoon highs in the \$40-50/MWh range.

The figure shows that from 2 AM until 11 PM, the net tie schedule (NYN) was westbound and ranged from near zero to 650 MW. However, for nearly all of these hours the LMP was *higher* on the New England side of the border. The counter-intuitive flow continues for hour after hour—20 hours in all that day. No participant submitted potentially profitable external transactions in the opposite direction and closed the price spread (possibly due to the risk of prices inverting, or due to the ISOs' transaction fees on external transactions).

This one-day example is by no means unique. As the prevalence of points in the red portions of Figure II-2(a) and (b) suggest, counter-intuitive flow occurs nearly half of the hours per year. There are many days, such as in Figure II-5, when counter-intuitive flow continues for long periods at a time.

A final point to note from Figures II-2(a) and (b) is that the fraction of hours with counter-intuitive flows is similar whether New England is the higher cost region (top graph) or New York is the higher cost region (lower graph). That means both regions stand to benefit from lower production costs under a system that produces more efficient scheduling and reduces counter-intuitive flow.

BALANCE OF PRICES

As the economic supply-curve depictions in Figures II-3 and II-4 suggest, the gains from more efficient scheduling depend on how often there are large price differences on either side of the border. The scatter-plots in Figure II-2(a) and (b) suggest this occurs a lot. Here we take a closer look at how often prices differ between the New York and New England sides of the interface.

PRICE DIFFERENCES

Figure II-6 shows an annual duration curve for the *difference* in hourly real-time locational marginal prices between NYISO and ISO-NE in 2009. Reading the duration curve in the usual way, it indicates how often the NYISO price exceeds the ISO-NE price by a given amount (or more). For example, 20% of the time, NYISO's real-time price exceeded ISO-NE's by about \$7 or more per MWh. This specific duration curve is for the price difference across the primary interface (NYN), and excludes hours in which the interface transmission capacity limit is binding (3% of hours).

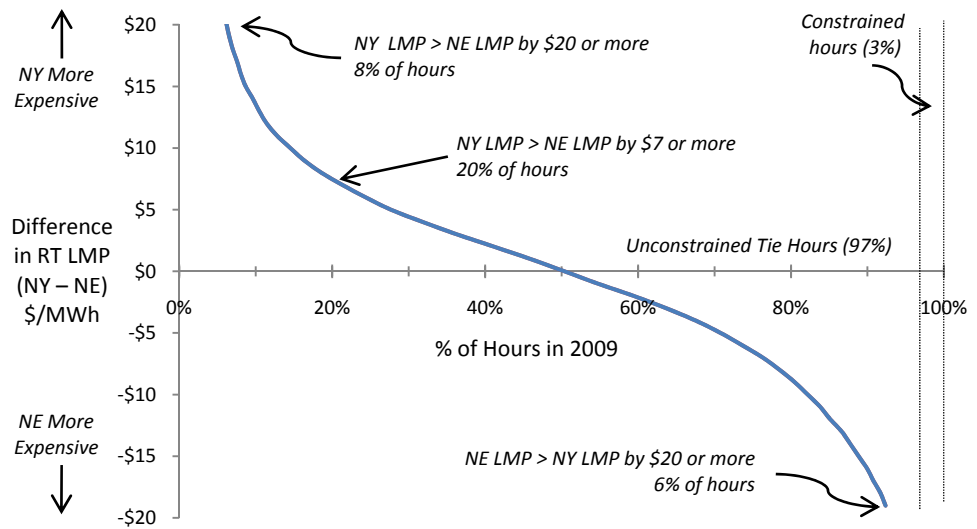


Figure II-6. Annual duration curve for the LMP difference across the NYN interface, 2009

IMPLICATIONS

Three observations are important. First, there are wide price differences in most hours. The price difference exceeds \$5 per MWh (in absolute value) more than half of the year, and exceeds \$10 per MWh (in absolute value) nearly one-third of the year. These are hours in which there is transmission capacity available to schedule additional transfers across the interface.

Second, the duration curve crosses the horizontal axis at the fifty percent mark. This means each region had the higher price at the primary interface (NYN) equally often.

Third, large price spreads occur with similar frequency in both directions. The figure reveals that in 8% of the hours in 2009, the NYISO price exceeded ISO-NE's price by \$20 per MWh or more. At the opposite end of the duration curve, the data show that in 6% of the hours the ISO-NE price exceeded NYISO's price by \$20 per MWh or more. The general symmetry of the price duration curve indicates the fraction of time in which one ISO has higher prices than the other is similar each way.

Taken together, these observations indicate that if an efficient interface scheduling system had been in place in 2009, the cost savings would accrue to both regions. The simulation results described in section II.B (below) confirm this.

VOLATILITY

Are inefficient tie schedules caused by volatile real-time price differences from hour to hour—and therefore hard for market participants to predict in advance? Yes, in part. The volatility of real-time price differences is evident using statistical indexes based on changes in the real-time price spread.

Figure II-7 plots the *50/50 Volatility Index* for the hourly real-time price difference between NYISO and ISO-NE at their border (the proxy bus price spread). Interpreting volatility indexes can be complicated, but the idea is simple: Larger values of the volatility index mean greater hour-to-hour changes in the price spread, in either direction.




Figure II-7. Volatility of the hourly price spread between NY and NE over time.

This volatility index calculates the median change, from one hour to the next, of the (absolute) LMP spread across the NYN interface. The median is calculated over the last 30 days (720 hourly observations) on a rolling basis. For example, an index value of 10 means that half the time during the last 30 days, the price spread changed from one hour to the next by \$10/MWh or more (either way); the other half of the time, the price spread changed by \$10/MWh or less.

There are several key observations from Figure II-7. First and foremost, real-time price differences between regions can change greatly from one hour to the next. That contributes to tie under-scheduling and counter-intuitive flow. The current external transaction scheduling system produces inefficient net tie schedules—in effect, ‘mistakes’—partly because system conditions change faster than the net tie schedule does.

This implies that to improve efficiency significantly, any tie scheduling system will need to update the net tie schedule at a higher frequency to react to changing system conditions. Higher frequency scheduling is an integral component of the solution options described in Parts III and IV.

Second, there is no clear evidence that price differences between the ISOs are becoming more volatile over the years. The high year is 2008, when fuel prices set record highs; the low year is 2009, when both fuel prices and power demand were low. This suggests that the benefits from a more efficient tie scheduling system would likely vary from year to year, depending on the underlying cost and demand drivers that determine how often the ISOs are operating on the steep portions of their supply curves.

In summary, the data on LMPs and transmission schedules examined here present a compelling case that the current external transaction system could be improved to serve its central economic purpose: To converge prices at the interface. That would reduce total system costs, and the prices paid by loads, in each region. We consider the economic benefits that may be realized from a more efficient scheduling system next.

B. COST/BENEFIT CONSEQUENCES

The prevalence of tie under-utilization and counter-intuitive flow means NYISO and ISO-NE are incurring higher production costs than necessary. How much would these costs be reduced if the interface between the two regions was scheduled efficiently?

To answer this question, NYISO and ISO-NE requested that Potomac Economics analyze each ISO’s market and transmission data through 2010. The analysis also examines the changes in average LMPs in each region, and changes in loads’ energy market

expenditures, relative to the status quo. A detailed report will be forthcoming from Potomac Economics in early 2011; here we summarize prior estimates from Potomac Economics' *Annual Assessment* reports for New England and New York.¹³

Potomac's methodology employs a simulation model of production costs in each region, including the major transmission constraints internal to each ISO and between them. Using actual generation market offers and demand data, the model simulates the hourly generation re-dispatch that would converge prices, or bind the interface if that occurs first, between the two regions. It then calculates the change in production costs associated with this re-dispatch. Conceptually, this corresponds to estimating the size of the red excess cost triangles represented in Figures II-3 and II-4.¹⁴

One important cautionary note is in order. This method estimates "ideal" scheduling across the interface, as if the ISO's had a crystal ball revealing the next hour's LMP differences when setting the interface schedule between them. In reality, any feasible system will be subject to contingencies and unforeseen events that may limit its efficiency gains to less than this ideal. Accordingly, the results below are best viewed as the *potential* cost savings from an optimized tie scheduling system.

MAIN FINDINGS

Table II-1 summarizes the main findings. The estimated cumulative reduction in total production costs over the last five years totals \$77 million dollars for the two regions combined. The annual figures vary modestly from year to year, being higher when fuel (principally natural gas) prices are high and lower when fuel costs decline, as occurred in 2009 and 2010.

Table II-1. Estimated Impact of Efficient Interface Scheduling

Year	Reduction in Production Costs	Reduction in Loads' Expenditures		
		Total	New England	New York
2010*	10	186	147	38
2009	10	125	68	57
2008	19	154	25	129
2007	21	199	22	177
2006	17	120	61	59
Cumulative	77	784	323	460

Note: 2010 data through October.

Sources: Potomac Economics, *Annual Assessment of the New England Markets* and *Annual Assessment of the New York Markets* for 2009 and updates.

¹³ *Op cit.*, note 1, sections III and IV, respectively.

¹⁴ Additional details can be found in Potomac Economics, *Annual Assessment*, §III (*op cit.*), and Potomac Economics, *Answers to Questions from Market Participants Regarding the 2008 State of the Market Report* (June 24, 2009).

The estimated reductions in energy expenditures by loads are an order of magnitude larger than the reductions in production costs. This occurs because increasing tie utilization by 200 MW reduces the production costs associated with that 200 MW, but it affects the LMPs for *all* load (unless limited by transmission constraints)—which may be 20,000 MW in each region at the time.

The estimated cumulative reduction in energy expenditures by loads if the interface was efficiently scheduled over the last five years totals over three quarters of a billion dollars. These savings vary by year and by region. During normal operating years, such as 2006 and 2009, the expenditure reductions to loads are quite similar in each region. During a year in which one ISO or the other has a greater number of price spikes, operating reserve shortages, or other adverse system conditions—as occurred in 2010, 2007, and 2006—the distribution of benefits skews toward the ISO experiencing more adverse system conditions. In this respect, efficient tie scheduling provides a degree of insurance to each region’s loads that reduces the impact of price spikes.

Importantly, the data indicate that both regions’ loads would experience lower costs in all years with efficient tie scheduling. The primary reason is that, in any particular hour, the lower-cost ISO tends to be operating on a flat portion of its supply curve, and the higher-cost ISO tends to be operating on a steeper portion of its supply curve. Sending additional megawatts from the low-cost to high-cost region tends to drop the price in the high-cost region a lot, but raise price in the low-cost region only a little. Because the region with lower costs varies from day to day and within each day, both New York and New England experience decreases in average LMPs on an annual basis.

Table II-2 summarizes this phenomenon using the data for 2010 (through October). It shows that with efficient tie scheduling that equalizes LMPs (up to transmission limits), the average LMP when New England is importing would fall by \$7.32 per MWh, but when exporting would rise only about half as much, \$3.75 per MWh. The results are similar for New York: Average LMP when it is importing would fall by \$7.09, but when exporting would rise by only \$3.89.

Table II-2. LMP Impacts of Efficient Interface Scheduling

Average Change in Hourly (Real-time) LMP, in \$ per MWh		
	When Importing	When Exporting
New England	- 7.32	+ 3.75
New York	- 7.09	+ 3.89

Notes. 2010 data through October. *Source:* Potomac Economics

Two final results from Potomac’s analysis merit note. First, the estimated change in total scheduled power flows between New York and New England with efficient tie schedules is modest: An increase of between 1 and 10%, in each direction, in most years. Second, the frequency of congestion across the primary interface remains only slight, increasing by 2-3 percentage points per year (up from the current 3 percent; see Figure II-1). This implies congestion is likely to remain infrequent across the primary interface between New York and New England.

SUMMARY

Taken altogether, the data support two main conclusions. First, the current external transaction system produces demonstrably inefficient outcomes. These inefficient outcomes cause excess production costs in each region, averaging in the low tens of millions of dollars annually. Second, loads pay for the inefficiencies of the current external transaction system. Their total energy expenditures would be on the order of one to two hundred million dollars lower annually—or perhaps half a million dollars *per day* lower—if the real-time inter-regional interchange system produced efficient tie schedules.

C. ANALYSIS OF ROOT CAUSES

The economic inefficiencies we observe with the current inter-regional trading system can be traced to three root causes. These root causes are important for understanding what problems must be solved, and why the solution options presented in Parts III and IV will produce substantially more efficient tie schedules than today’s system.

The three root causes are:

1. *Latency Delay.* The time delay between when the tie is scheduled and when power flows, during which time system conditions and LMPs may change.
2. *Non-economic Clearing.* The ISOs make decisions about which tie schedule requests to accept without economic coordination, producing inefficient schedules.
3. *Transaction Costs.* The fees and charges levied by each ISO on external transactions serve as a disincentive to engage in trade, impeding price convergence and raising total system costs.

Each is addressed below.

THE LATENCY PROBLEM

“Latency” is the time delay between when (1) an ISO determines whether or not an external transaction request is economic (clears) and accepts it, and (2) the transaction’s execution is complete. Under the current inter-regional trading system, the latency

delay is nearly two hours for both NYISO and ISO-NE. Each ISO determines whether or not a “real-time” external transaction offer clears about an hour before the corresponding delivery hour, and this quantity remains fixed for the full delivery hour. This combination of hourly lead-time and the hourly duration-time are referred to an “hourly scheduling” system, although the latency delay is double that.

CONSEQUENCES

Latency delay causes excess costs for the system as a whole. Power system conditions can change from minute to minute, altering each ISO’s bid-based marginal generation cost. That can make a transaction that appeared economic when cleared an hour earlier actually uneconomic during the delivery hour.

The consequences of these *ex post* uneconomic transactions are evident in the data (section II.A) in two ways:

- *Tie under-utilization.* If the importing region’s LMP rises relative to the exporting region’s LMP after all transactions are accepted, imports are more valuable in real-time than the market anticipated. Even though power is flowing the ‘right’ way, *not enough* power is flowing the right way to minimize total system costs (*c.f.* Figure II-3).
- *Counter-intuitive flow.* If the importing region’s LMP falls relative to the exporting region’s LMP after all transactions are accepted, counter-intuitive flows may result. That inefficiently displaces low-cost generation in one ISO with high-cost generation in the other ISO at the margin, *increasing* total costs for the two regions overall (*c.f.* Figure II-4).

Latency delay is an economic problem because price information available to the ISO and to market participants more than an hour in advance is imperfect. This is evident in the volatility index (Figure II-6) above. The LMP data indicate that the region with the lower LMP (at the border) switched from one hour to the next over one-third of the time in 2009.

IMPLICATIONS

From the perspective of a market participant submitting an external transaction request, the latency problem creates financial risk. If the price difference between regions changes after an external transaction is accepted, the market participant can end up “buying high and selling low,” losing money on each megawatt scheduled. This risk poses a deterrent to submitting external transactions in the first place, exacerbating the tie under-scheduling problem.¹⁵

¹⁵ Under certain conditions, each ISO compensates a market participant for financial losses on external transaction requests that clear an hour in advance but that incur a loss at real-time prices. In theory this may reduce the tie under-scheduling problem, but it will exacerbate the counter-intuitive flow. It also increases total ‘uplift’ charges ultimately paid by other market participants.

From the standpoint of economically sound market design, the best solution is an inter-regional interchange system that minimizes latency delay and the excess costs it creates. More concretely, resolving the latency problem requires:

1. Reducing the time lag between when the ISOs determine the aggregate net tie schedule and when the power actually flows; and
2. Enabling more frequent updating of the aggregate net tie schedule to reflect changes in system conditions and energy prices in real-time.

Satisfying each of these goals is a central design objective of both major solution options presented in Parts III and IV.

NON-ECONOMIC CLEARING

Under the current system, market participants submit separate external transaction requests to each ISO (*e.g.*, an offer to buy, or export, from ISO-NE is submitted only to ISO-NE, and the matching offer to sell, or import, in NYISO is submitted only to NYISO). There is no economic coordination between the ISOs when they set the aggregate net tie schedule. This absence of economic coordination when external transaction requests are accepted produces inefficient tie schedules, and raises total system costs. We explain how next.

CONCEPT

In an economically sound market design, an external transaction should be accepted, and increment the net tie schedule, if condition (A) is true:

- (A) The importing region's expected LMP exceeds the exporting region's expected LMP.

(Assuming sufficient transmission capacity). This condition ensures low-cost generation displaces high-cost generation, lowering total production costs.

The current external transaction system does not verify condition (A), however. Instead, it checks two different conditions:

- (B) The offer to buy (export) exceeds the exporting region's expected LMP;
- (C) The offer to sell (import) is less than the importing region's expected LMP.

The ISO receiving the export-side schedule request checks (B), and the ISO receiving the import-side schedule request checks (C). If—and only if—both conditions are satisfied (“check out”), then the external request is cleared “to flow”. These cleared transactions determine the net interface schedule during the requested delivery hour(s).

WHAT GOES WRONG?

The problem here is that conditions (B) and (C) do not imply condition (A). This means that transactions can, and are, routinely scheduled to flow that do not reduce total system costs. In fact, cleared transactions that do not satisfy condition (A) *raise* total system costs—to the detriment of everyone that buys power.

This problem is known as *non-economic clearing* because the necessary condition for a transaction to be economic—that is, condition (A)—is not checked by the ISOs before external transactions are cleared to flow.

What goes wrong in practice? The participant’s offer to buy (export) from one ISO may clear at a *higher* price, while the offer to sell (import) in the other ISO clears at a *lower* price. Unless a greater number of megawatts are accepted (clear) in the opposite direction, this raises total costs for the system as a whole: High-cost generation in the exporting region is displacing low-cost generation in the importing region.

Note that this is a fundamentally different problem from latency. Here, there is no latency delay at all: The transaction is scheduled in the wrong way from the start. The root of the market design problem is that today’s inter-regional interchange system checks (B) and (C), instead of checking all three conditions, before accepting transactions that determine the physical tie schedule.

PREVALENCE

The prevalence of this problem is indicated in Figures II-8(a) and II-8(b) (*next page*). The vertical axis shows the difference in the two ISO’s *scheduling prices*, each of which is used by the ISO to accept its ‘side’ of the external transaction request (approximately 45 minutes before the delivery hour).¹⁶ That acceptance process corresponds to the ISO checking condition (B) or (C) (depending whether it is exporting or importing, respectively). The horizontal axis is the net tie schedule across the primary interface (NYN), with each dot a separate hour. Any point in the top-left or lower-right panel (red dots) of the two figures has a higher expected LMP in the exporting ISO, and therefore fails condition (A).

There are a lot of hours that fail condition (A). The data in Figures II-8(a) and (b) cover all hours from July through December 2009. Of these, the ISOs are scheduling power from the high-to-low cost region (at the margin), raising expected total production costs, 44 percent of the time.¹⁷ This is not because of latency; if the ISOs had these data in real time, they would have been able to see that the interchange scheduling system had produced an inefficient net tie schedule 45 minutes in advance.

¹⁶ NYISO’s scheduling price is generated by its forward-looking real-time commitment algorithm at (approximately) 45 minutes before the hour. ISO-NE’s scheduling price is determined separately, but at the same time, by software tools that evaluate the generation bid stack and expected dispatch rate during the delivery hour. Each ISO’s scheduling price can be interpreted as the ISO’s internal forecast of its LMP (at the border) during the next delivery hour.

¹⁷ This phenomenon is most common when a market participant submits a price-based bid on the NYISO side of the border, and a “fixed bid” on the ISO-NE side of the border (a “fixed bid” is an offer to pay up \$1000 per MWh to export, or to accept any price above zero to import). Any such bid-pair that clears on the NYISO side will also clear in ISO-NE, even if condition (A) fails.

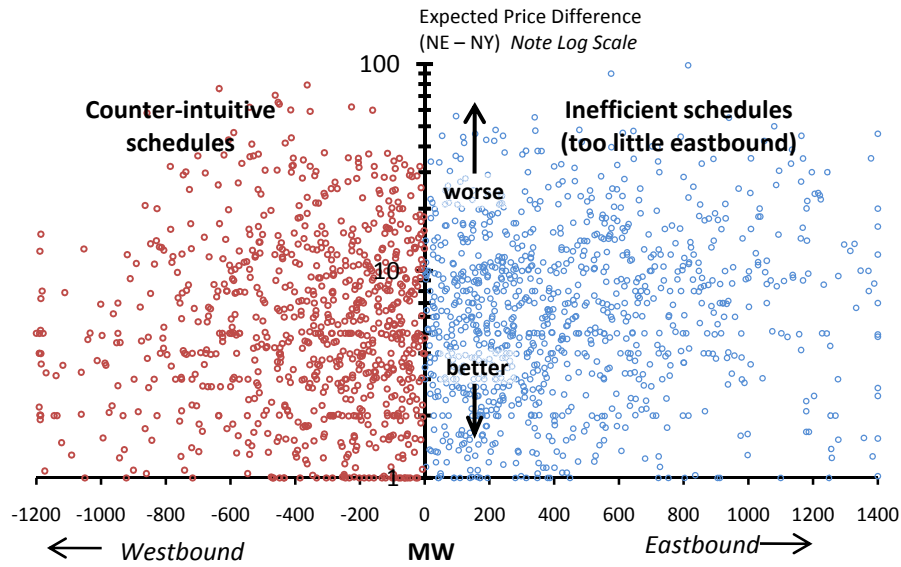


Figure II-8(a). Scheduled tie flow and the *expected* price difference across the primary interface (NYN) when New England's expected LMP exceeds New York's expected LMP during the upcoming schedule hour. Hourly data, July through December 2009. Note the price difference on the vertical axis is in logarithmic scale.

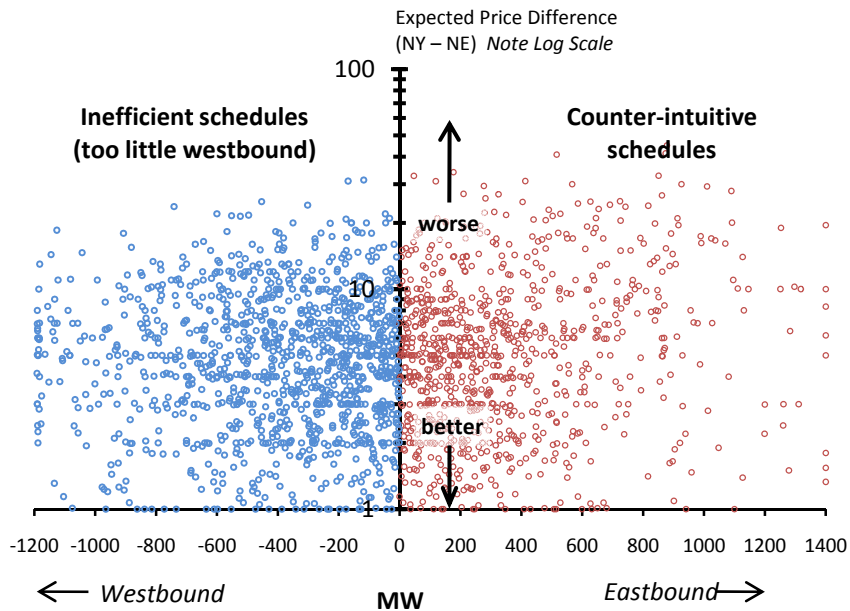


Figure II-8(b). Scheduled tie flows and *expected* price differences across the primary interface (NYN) when New York's expected LMP exceeds New England's expected LMP during the upcoming schedule hour. Hourly data, July through December 2009. Note the price difference is reversed from Figure II-8(a), and in logarithmic scale.

The upper-right and lower-left panels (blue dots) reveal an additional, different scheduling inefficiency. These panels show hours when the net interface schedule is in the direction that is expected to be economically correct (meaning, it satisfies condition (A)). The figure makes clear, however, that too little power is being scheduled to converge the LMPs most of these hours. That means total system costs are unnecessarily high: Trade fails to displace high-cost generation with additional lower-cost generation available from the exporting region. Of the remaining 56 percent of the hours, this situation occurs in almost all of them.

IMPLICATIONS

From the standpoint of economically sound market design, an efficient inter-regional interchange system should yield an interface schedule in the direction that satisfies condition (A). This would correct a fundamental market design flaw of the current inter-regional trading system that yields excess costs.

Correcting this fundamental flaw is a design objective satisfied by both of the major solution options presented in Parts III and IV.

TRANSACTION COSTS

The third root cause of inefficient net tie schedules is the transaction costs that market participants pay when they schedule external transactions. Both ISO-NE and NYISO impose a number of different fees and charges on external transactions. These include:

- An allocation to external transactions of general ‘uplift’ costs incurred by the ISO, on a per-megawatt basis.¹⁸
- Financial Impact Charges imposed by NYISO on import/export requests that fail the “check-out” process for reasons within the market participant’s control.
- ISO scheduling fees paid by market participants on a per-megawatt hour basis, which for external transactions are levied by *both* ISOs.

In practice, the most consequential of these charges is likely the first. Real-time uplift averages a few dollars per megawatt on an annual basis, and is levied differently in NYISO and ISO-NE. However, it is highly variable from day to day, the charges are difficult to predict in advance, and there is no practical means for a market participant to hedge against them.

CONSEQUENCES

The allocation of these fees to external transactions reduces trade between regions and adversely impacts total production costs. The reason is straightforward: To cover the expected fees levied by the ISOs, a market participant will submit an external

¹⁸ In NYISO, generators receive Bid Production Cost Guarantees (BPGC); in ISO-NE, this is known as Net Commitment Period Compensation (NCPC). To minimize unfamiliar acronyms, we use the generic term ‘uplift’ here.

transaction request when the expected price difference between regions exceeds the expected total fees. Put in financial terms, the participant will incorporate a margin, or “bid-ask spread”, between the prices at which it is willing to buy and sell across the interface.

This rational behavior prevents price convergence between regions. System costs will be higher than necessary because the ties will tend to be under-utilized, relative to external transaction volumes if there were no per-megawatt transaction fees. In economic terms, transaction fees on external transactions act like a ‘tax’ that deters trade and distorts production costs upward.

How large is the distortion due to transaction costs? It is difficult to assess precisely, because the largest component—real-time uplift charges—is highly variable and creates risk that deters trade. On a monthly basis, real-time uplift charges average about \$1 per MWh in NYISO and average between \$1 and \$4 per MWh in ISO-NE. ISO scheduling fees add another \$1-2 per MWh or so (from both ISO combined). Together, these suggest a risk-neutral market participant would not find it profitable to schedule transactions that would drive prices between regions closer than perhaps \$5 per MWh, maybe more.

IMPLICATIONS

Transaction fees can be allocated in a number of different ways. Allocating them in a way that prevents price convergence between regions means one region will tend to have higher-cost generation on the margin when lower-cost power is available and ample transmission capacity between them. This raises the cost of meeting total power demand.

While the allocation of uplift to external transactions in the real-time market ‘saves’ other market participants from paying these fees, this allocation may be a Pyrrhic victory that costs loads in the end. An example suggests why. An allocation of \$5 per MWh on 1000 MW of external transactions in a particular hour means other market participants avoided \$5,000 in ISO uplift charges. An efficient inter-regional interchange system would converge prices between regions, reducing the price spread by \$5 per MWh.

What is that convergence potentially worth to loads? In electricity markets, when prices converge between regions the importing region’s LMP will typically fall more than the exporting region’s LMP rises. (See again Table II-2). For example, the LMP may fall by \$3 per MWh in the importing region, and rise by \$2 per MWh in the exporting region, to eliminate the \$5 per MWh price spread between regions. Assuming both regions have equal loads of (say) 20 GWh at the time, the *net* savings to loads from eliminating the external transaction fees in this scenario would be at least \$15,000 for that hour (20 GWh × (\$3 – \$2)/MWh, less a portion of the \$5,000 in re-allocated fees).

The bottom line is that allocating uplift and other transaction fees to external transactions impedes price convergence between regions and raises system production costs. The magnitude of the benefits to loads from eliminating an uplift allocation to external transactions, in the form of lower LMPs due to greater inter-regional competition, could well outweigh the re-allocated costs.

III. SOLUTION OPTION A: TIE OPTIMIZATION

To solve the problem of inefficient interface schedules between ISO-NE and NYISO, the two ISOs established a joint design team to develop solution options and recommendations. In this section and the next, we present conceptual designs for two solution options: (A) Tie Optimization and (B) Coordinated Transaction Scheduling.

Either of these two options would be a major improvement over the status quo. Either option would eliminate much of the inefficiencies that result in tie under-scheduling and counter-intuitive flows. Relative to the current system, either option would yield lower production costs for both regions and significant savings for loads.

The two options we present are mutually exclusive, meaning the ISOs cannot implement both simultaneously. This is an operational constraint: The two systems employ different underlying mechanisms for using market-based information to determine the net interface schedule between regions. Of the two, the ISOs recommend the Tie Optimization option because it is the more efficient solution. This emphasis on efficiency reflects the ISOs' shared philosophy of using competitive wholesale markets to meet power demand at the lowest possible cost.

We describe the Tie Optimization option next. The Coordinated Transaction Scheduling option is described in Part IV.

A. CONCEPT AND DESIGN PRINCIPLES

Although market design details can be complex, the core idea of Tie Optimization is simple. It is this:

The ISOs manage the transmission ties between them in the same way, or as close as possible to, the ISOs manage transmission internally.

In an important sense, Tie Optimization is not a new design. It is the same bid-based, security-constrained least cost dispatch logic that underlies the wholesale energy market administered by each ISO. This competitive market design applies to all internal nodes and internal transmission facilities today. Tie Optimization simply extends this market design to cover the pool transmission facilities that interconnect the ISOs.

In practice, this means three things:

- Optimizing physical transmission flows to minimize the total costs of meeting inter-regional demand, using market-based resource supply offers;
- Providing voluntary financial instruments to help market participants hedge price risk at the interface, for those who wish to do so;
- Eliminating procedural obligations that can require market participants to submit schedule requests in inefficient ways under the current system.

The first point addresses the inefficiently higher costs that result from the way physical transmission flows are scheduled across the interface today. Tie Optimization coordinates real-time energy dispatch across both ISOs' control areas to minimize total production costs. This is made possible because of advances in communications and information technology in recent years that enable the ISOs to implement a (near) joint energy dispatch without merging control rooms. We describe how this process works in more detail below (sections III.B and III.C).

The second and third elements in the list above provide market participants who currently "move power" across the interface with new options to meet their business needs. Only a subset of market participants actively engage in external transactions today, and their reasons for doing so are varied. Accordingly, the ISOs are interested in discussions with market participants about how to best accommodate their varied needs, in ways that are consistent with efficient operation of the inter-regional transmission network.

CORE DESIGN PRINCIPLES

The ISOs developed the Tie Optimization option as a solution to the inefficiencies documented in Part II. The development of this option is guided by four core market design principles.

1. **EFFICIENCY.** Tie Optimization is designed to schedule the transmission interface in the most efficient way, serving demand in both regions at the lowest possible production cost.
2. **MARKET-BASED TRANSMISSION FLOWS.** The ISOs use competitive, market-based supply offers from market participants to determine the power system's dispatch. Least-cost dispatch determines all transmission schedules, within and between the two ISOs.
3. **ALL SETTLEMENTS AT LMP.** All energy flows across the interface are priced at the interface LMP. This facilitates market transparency, properly reveals and prices real-time congestion, and sends correct price signals to all market participants about the value of energy flowing between regions. Under either reform

option, there are no uplift debits or credits associated with inter-regional transmission flows.

4. ISOs HAVE NO FINANCIAL POSITION IN MARKETS. The ISOs do not directly participate in markets, and do not buy nor sell. Loads pay LMP for all power consumed, wherever the generation is located; generators receive the LMP at their locations, regardless of where the power goes. The ISOs will continue to act as independent settlement administrators for the payments to and from market participants. The same basic settlement procedures are used to settle cash flows on each ISO's internal energy market today.

RELATION TO VRD DISCUSSIONS (2003-2005)

Back in the 2003-2005 period, ISO-NE, NYISO, and their stakeholders engaged in a discussion of "virtual regional dispatch" (VRD). VRD did not take the form of a fully developed market design proposal; discussions centered on (somewhat amorphous) proposals to more closely integrate power scheduling between ISOs.¹⁹

In 2005, the two ISOs took the concrete step of arranging an Inter-Hour Transaction Scheduling (ITS) experiment. The ITS experiment was operational in purpose, not economic. Reports indicate there was little impact on market prices, but no systematic economic analysis was conducted.²⁰ Neither the ISOs, nor their stakeholders, viewed further development as a priority given other pressing market development projects at the time.²¹

The financial arrangements under which the experiment was conducted raised concerns among some market participants. The experiment proceeded under special tariff provisions allowing the net costs to flow through to an ISO's operating expense.²² Importantly, that will not occur under Tie Optimization. Instead, with Tie Optimization load pays generation for energy that flows across the interface in the same way that load pays generation for energy flows *within* an ISO's footprint. The ISOs act as settlement administrators for payments from loads to generation; the ISOs do not buy or sell power and take no financial position in the energy market.

¹⁹ See *Northeast ISOs Seams Resolution Report*, Item P18, October 8, 2010 (http://www.nyiso.com/public/webdocs/documents/regulatory/seams_issues/Seams_Closed.pdf#page=18).

²⁰ See *Intra-Hour Transaction Scheduling Between NYISO and ISO-NE: Report on the Pilot Tests Conducted on April 20-21, 2005* (http://www.iso-ne.com/committees/comm_wkgrps/mrks_comm/mrks/mtrls/2005/nov142005/a2_pilot1_report_2005_1_1_08clean.doc)

²¹ See *ITS: How to Progress*, Joint Presentation to NYISO MSWG and ISO-NE Markets Committee (November 14, 2005) (http://www.iso-ne.com/committees/comm_wkgrps/mrks_comm/mrks/mtrls/2005/nov142005/a4_its_how_to_progress_2005_11_11.ppt)

²² *Op cit.*, p. 3 (note 2).




Figure III-1. The tie is scheduled efficiently when LMPs are equal

B. ECONOMIC FRAMEWORK

A central difference between Tie Optimization and the current external transaction system is *which* market-based information the ISOs use to determine transmission schedules between regions. This section describes the economic logic of real-time Tie Optimization and its market foundations. Tools and process to enable a market participant to manage price risk are distinct solution elements, addressed in sections V and VI.

Although operational details of how the ISOs manage transmission can be intricate, the economic logic of Tie Optimization is conceptually simple. Consider Figure III-1. At a specific point in time, imagine the black curve represents the generation supply offer stack in one ISO (in this case, NYISO). This stack characterizes NYISO’s incremental cost of delivering power at its side of the (NYN) interface. Similarly, the blue curve represents the generation supply offer stack of the other ISO (in this case, ISO-NE), shown here in *descending* cost order.²³

The optimal tie schedule is the level of tie flow, in MW, that equates LMPs on each side of the border. In the figure, this is the point where the two supply stacks cross (labeled “Optimal Tie Schedule”). By setting the net aggregate tie flow at the level that equates

²³ More precisely, the supply curve is NYISO’s incremental dispatch cost at the NYN interface proxy bus. Technically, each supply curve shown will incorporate all generation shift factors changes and internal transmission constraints that may bind and affect the proxy bus price over any segment of the range of admissible tie flows.

each ISO's LMP at the border, costs are minimized: There is no other way to allocate total generation, across all facilities in both regions, that yields lower total production costs.

As drawn in Figure III-1, at zero tie flow ISO-NE has higher costs than NYISO. Thus the optimal tie schedule is eastbound. The logic applies similarly in hours when NYISO has higher costs than ISO-NE (at zero tie flow). See Figure III-2. In this situation, the optimal tie schedule is westbound. Again, the optimal schedule is where the supply stacks cross.

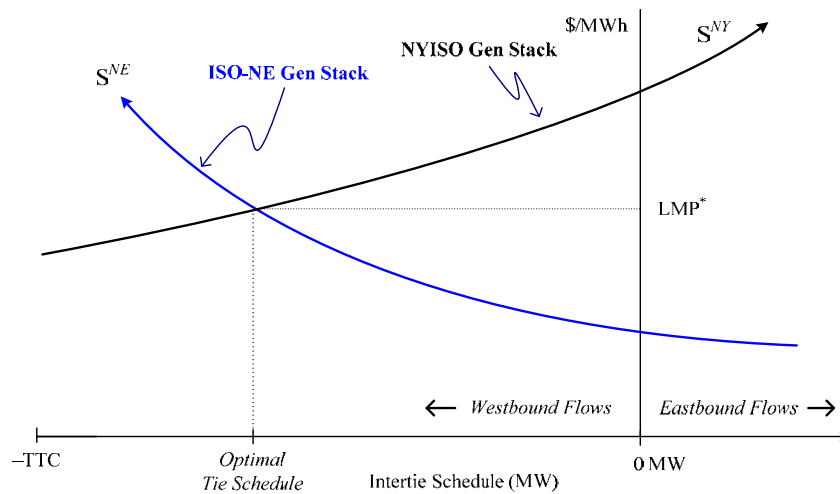


Figure III-2. Tie Optimization when the optimal schedule is westbound

The ISOs are able to determine this using their existing network models and market offer data by exchanging their exact interface (proxy bus) dispatch stack information in real-time. In effect, Tie Optimization amounts to pooling information on each ISOs' incremental bid-cost of supply, and decremental bid-cost of reduced supply, across the transmission interface. By pooling this information, the ISOs can determine how much to dispatch up and down on each side of the border to equate LMPs over the next dispatch interval. This prevents one ISO from operating higher cost generation (at the margin) than the other, whenever there is transmission capacity available to move power in the low-to-high cost direction.

WHO PAYS WHOM?

Figures III-1 and III-2 also reveal the basic settlement logic. For the moment, suppose there is no congestion nor losses. Load on the import side (whichever ISO that may be) pays the LMP at its locations. Generation on the export side is paid the LMP at its locations. If the LMPs are equal at the border, the (energy component of) LMP paid by load and paid to generation is the same.

To provide the appropriate cash flows, the ISOs—acting as settlement administrators—transfer specific funds coming from loads in the importing ISO to generation in the exporting ISO. The transfer is equal to the border LMP (or LMP* in Figures III-1 and III-2), times the tie schedule megawatts. In essence, the ISOs act as a joint settlement administrator for the appropriate payments from loads to generation across the interface.

Like internal dispatch today, there are additional settlement elements that arise due to congestion and marginal losses: LMPs at the border will not always equal internal node LMPs, nor always equal each other (e.g., if the tie is congested). Nevertheless, the basic settlement logic at the interface generalizes in the same way it is today applied to internal settlements for real-time congestion and marginal energy losses.

In sum, the central rationale for Tie Optimization is the same reason we use bid-based central dispatch today. A coordinated dispatch approach using both regions' market-based generation offers will produce the most efficient solution, minimizing both under-scheduling and counter-intuitive flows between regions.

C. HIGHER FREQUENCY SCHEDULING

To minimize latency delays (see section II.C), it is desirable to set the tie schedule as close to real-time as possible, and to update it as frequently as system conditions change. This section describes the basic timing and information flows between ISOs in order to “cross the stacks” and implement higher frequency schedule (HFS) changes under Tie Optimization. HFS applies to real-time operations and to price formation in the real-time energy markets. (Day-ahead markets and interactions are addressed in Part V).

A substantively identical timing of information flows between the ISOs characterizes the Coordinated Transaction Scheduling option described in Part IV.

OVERVIEW

The ISOs propose to update the net interface schedule across each (AC) interface between ISO-NE and NYISO approximately 5-to-10 minutes before the energy flows. These schedules are (nominally) fixed for 15 minutes between each update. Within each 15 minute interval, each ISO will perform its internal dispatch.

The choice of 15 minute intervals is determined by current technology and operational considerations, not economic theory. With additional advances in technology, it may be possible to further shorten this interval. The key constraint is that Tie Optimization requires “pre-scheduling” passes of the unit dispatch system by each ISO, and an exchange of detailed solution information between ISOs after each pass. The resulting

optimized net tie schedule is then incorporated into each ISO's next energy dispatch solution sent to generators.

The "pre-scheduling" process determining the net interface schedules is a look-ahead system, producing a binding net tie schedule for the upcoming tie schedule interval and 'advisory' net tie schedules for subsequent intervals. The look-ahead advisory net tie schedules preserve one important feature of today's hour-ahead tie scheduling system: It provides the system operators with information (expected net tie flows) that can be important for evaluating the near-term system trajectory (up to 60 minutes out) and making operational decisions over the next hour.

TIMELINE

The key steps determining the aggregate net tie schedule across a particular interface are best described using a timeline. This timeline is indicative, as exact times can vary during real-time operations and current dispatch timing conventions differ slightly between NYISO and ISO-NE.

For concreteness, we walk through the key steps and explanations in detail here. Times are in minutes.

T-20 *Step Pre-Schedule.* ISO-NE performs a set of "pre-scheduling" unit-dispatch system evaluations to evaluate the (bid-based) cost of incremental and decremental energy at the interface proxy bus at time T, T+15, T+30 and T+45.

This (parallel processed) evaluation determines the complete proxy-bus supply stack over the [-1200MW, +1400MW] transmission capacity of the interface. ISO-NE operational constraints on tie flows during these scheduling intervals may limit this range, and would be incorporated into the evaluation.

In the context of Figures III-1 and III-2, step *Pre-Schedule* determines ISO-NE's (blue) supply curve at the ISO-NE/NYISO border.

T-17 ISO-NE completes its supply stack evaluation. It passes the interface dispatch-rate schedule (proxy bus supply stack) to NYISO. Accompanying this information are any constraints indicated by ISO-NE operators governing interface flows over the upcoming schedule intervals.

T-15 *Step TieOpt.* NYISO integrates the ISO-NE interface dispatch-rate schedule into its RTD (real-time dispatch) optimization. The ISO-NE proxy-bus dispatch schedule is incorporated into RTD as the incremental cost incurred (by ISO-NE) to provide additional power across the interface into NYISO, and decremental cost avoided (by ISO-NE) by additional power flows across the interface into ISO-NE.

The RTD optimization determines each ISO's scheduled interface flow target for the 15 minute period starting at time T and 'advisory' tie schedule targets for the 15 minute periods starting at T+15, T+30 and T+45. If there are no binding constraints on the interface, each target equates NYISO's expected RT LMP at with ISO-NE's expected LMP.

In the context of Figure III-1 and III-2, step *TieOpt* implicitly determines NYISO's (black) supply curve at the border, and explicitly determines the quantity where the supply curves intersect (the optimized tie flow).

- T-11** NYISO completes RTD optimization. It passes the optimized tie schedule MW to ISO-NE. Incorporated into the optimized tie schedules (for T, T+15, T+30 and T+45) are any constraints indicated by NYISO operators governing interface flows expected over the upcoming schedule intervals, as well as all constraints received from ISO-NE.
- T-10** *Step RTD.* Each ISO performs its internal (real-time) dispatch, taking the optimized tie schedule MW for T as an input. The ramp profile is executed from T-5 to T+5. (This is the identical process the ISOs use today when they run RTD/UDS at 10 minutes before the top of the hour, since today's hourly tie schedule changes are executed over a nominal ramp interval from T-5 to T+5).
- T-5** *Step Pre-Schedule update.* ISO-NE performs *Step Pre-Schedule* again, using updated system information as of T-55, to update its proxy-bus supply stack for T+15, T+30, T+45 and T+60.
- Step TieOpt update.* NYISO performs *Step TieOpt* again, using updated system information as of T and the results of ISO-NE's *Step Pre-Schedule* update. The *TieOpt* update produces new tie schedule targets for T+15, T+30, T+45 and T+60, sent to each ISO.
- T+5** *Step RTD.* Each ISO performs its internal (real-time) dispatch as usual, now incorporating the updated optimized tie schedule target for T+15. The ramp profile is executed from T+10 to T+20.
- T+10** *Step Pre-Schedule update* commences by ISO-NE.
- T+15** *Step TieOpt* update commences by NYISO, producing new tie schedule targets for T+30, T+45,
- T+20** *Step RTD*, taking the optimized tie schedule for T+30 as a constraint, with ramp profile executed from T+25 to T+35.

And so forth.

A few observations here are useful. First, the real-time dispatch steps that occur every five minutes are operationally the same as internal dispatch procedures today. It is the *inputs* that feed each ISO's internal real-time dispatch that change.

Tie Optimization 'feeds' the existing real-time dispatch a tie schedule target that is not fixed for an hour, as it is today. Instead, Tie Optimization feeds the existing real-time

dispatch process with a sequence of tie schedule levels that are updated every 15 minutes. Each update also provides system operators with revised ‘look-ahead’ schedule information for each 15 minute interval, over the next 60 minutes.

In effect, with Tie Optimization the ISOs treat every 15-minute real-time point (e.g., the :00 minute, :15 minute, :30 minute, ...) in the same way they dispatch today to meet the ‘top of the hour’ ramp window for hourly schedule changes. By doing so every 15 minutes, instead of only once per hour, the system produces a sequence of back-to-back ramp windows that is intended to smoothly adjust from one interval to the next.

This adjustment of the tie schedules is the key point. It enables the system to respond, in a more economically efficient way, to changes in the location of the lowest marginal-cost generation in either region. In effect, the dispatch process would use the market-based offers from all dispatchable resources—across both regions—to continually adjust the tie schedules, in near real-time, to re-balance production in the least-cost way.

ADDITIONAL OPERATIONAL OBSERVATIONS

In the HFS timeline, the Pre-Scheduling step is always performed by ISO-NE, and the Tie Optimization step is always performed by NYISO five minutes later. This sequencing of ISO actions builds on the comparative advantages of each ISO’s existing technology infrastructure. ISO-NE’s unit dispatch system accommodates a parallel-processing configuration that enables it to evaluate its supply stack over a range of candidate tie schedule levels in approximately the same amount of time needed to obtain a single unit dispatch solution. Thus, it is computationally efficient for ISO-NE to perform step *Pre-Schedule*. NYISO’s current real-time dispatch produces generator set points for the next time interval (5 minutes), as well as ‘advisory’ generator set points for 15, 30, and 45 minutes out. Thus, by performing step *TieOpt*, NYISO can produce advisory tie schedules for these same future intervals within an existing software optimization system.

There are a number of additional implementation decisions with HFS. We briefly note several of these implementation decisions that may be of interest at this stage. The scope of analysis for these issues is more appropriate to an operational design proposal than this report’s conceptual design analysis. Some of these implementation decisions may depend on pre-deployment performance testing and evaluation.

- *Proxy-Bus Granularity.* HFS may warrant evaluation of the appropriate granularity of the proxy-bus representations in network models and optimization, relative to current single-node proxy bus element for the NYN interface and separate proxy bus element for NNC (1385 line) used by each ISO.
- *SAR.* Because it may increase utilization of the tie lines significantly, Tie Optimization may impact the shared activation of reserves (SAR) between NYISO and ISO-NE. In principle, changes to the availability of reserves across

the border can be properly valued by the HFS algorithm, since both NYISO and ISO-NE have market-based reserve prices. This would result in the tie being scheduled to a lower level of capacity when the value of reserves across it exceeds the marginal value of energy across it.

- *Pool ramp constraints.* Priority rules may be needed to satisfy pool ramp limits across multiple interfaces simultaneously, if the scheduling ‘lead time’ varies across interface. This may be primarily an issue for the first-pass ISO (i.e., ISO-NE). For example, ISO-NE may need procedures to allocate pool ramp limits between the NYISO/ISO-NE, HQ/ISO-NE, and NBSO/ISO-NE interfaces during the top-of-the-hour ramp window.

D. WHAT ABOUT CONGESTION?

Tie Optimization seeks to “cross the stacks” to equate each ISO’s LMP on either side of the interface. However, at times there may be binding transmission constraints that prevent this. Although the economic logic of congestion pricing is the same whether it occurs across internal or across external transmission links, the external ties are administratively different. That presents two questions: First, if interface transmission constraints bind, how will the real-time congestion price across the interface be set? And second, who gets (and who pays) the congestion revenue? We address each in turn.

REAL-TIME CONGESTION PRICES

Real-time congestion prices across an external interface under Tie Optimization mirror congestion pricing internally. Consider Figure III-3. As before, the black curve represents the generation supply offer stack in one ISO (in this case, NYISO), indicating its incremental cost of delivering power at its side of the (NYN) interface. Similarly, the blue curve represents the generation supply offer stack in the other ISO (in this case, ISO-NE), shown again in *descending* cost order.

In this situation, the point where the supply stacks cross exceeds the total transfer capability (TTC) of the interface. The optimal tie schedule is therefore limited to the TTC, and the HFS Tie Optimization process will set the tie schedule to that value.

The binding transmission constraint across the interface produces price separation between markets. Each ISO’s LMP calculator will set the real-time LMP on its side of the interface as shown in Figure III-3. The cost of congestion across the tie is the difference between the importing and exporting LMPs in each region. In Figure III-3, the real-time congestion price is the difference $LMP^{NE} - LMP^{NY}$.

The logic applies just the same if the optimal tie schedule is transmission-constrained westbound in hours when New England has lower costs than New York.

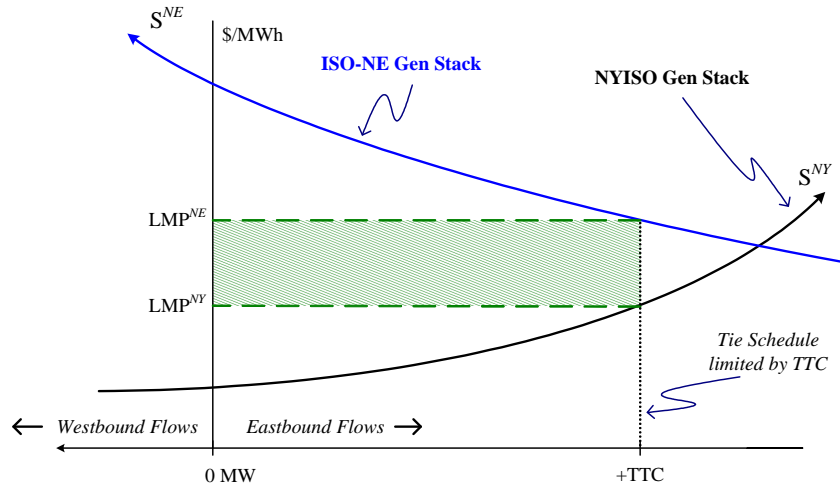


Figure III-3. Tie Optimization produces correct real-time congestion prices if total transmission capacity (TTC) limits tie flows between regions..

WHO PAYS WHAT?

There are a number of different ways to settle under the Tie Optimization solution option. For the sake of clarity, let's imagine—for the moment—that all energy is transacted in the real-time market. (This avoids the need to track deviations from day-ahead market positions, an issue considered in detail in Part V).

As usual, all payments are made between loads and generation, with the ISOs acting as settlement administrators. Load on the import side (in this case, ISO-NE) pays the LMP at its internal locations. Absent internal congestion (and ignoring losses), this is the same price paid to generation in the importing ISO. However, in the importing ISO there is less generation than load—the difference being the imported megawatts. That means the importing ISO receives more energy market revenue from its loads than it pays out to internal generation.

Where does the excess revenue go? Two places. Part of it must be paid to generation in the exporting ISO. As usual, generation in the exporting ISO is paid the LMP at its location. This is LMP^{NY} in Figure III-3 (absent internal congestion and ignoring losses). Acting as joint settlement administrators, the two ISOs transfer funds debited from importing-side loads to the credit of generation in the exporting ISO. This transfer equals the tie flow MW times the LMP on the exporting side of the interface (LMP^{NY} in Figure III-3).

Then there's the second part: When there is congestion, like in Figure III-3, there is a surplus left over. The amount left over is the *congestion rent*. In the figure, this is the shaded green area. It equals the congestion price across the interface times the (constrained) tie schedule MW.

What happens with the congestion rent? In practice, the congestion rent accrues (primarily) in the day-ahead market, where most energy is transacted, not in the real-time market. This provides the revenue stream necessary to fund risk-management products (FTR/TCC) across the transmission interface between regions.

Precisely how day-ahead congestion revenue accrues across an external interface requires a discussion of day-ahead markets and how they interact with congestion pricing of real-time flows. Day-ahead markets and settlements are similar under both the Tie Optimization and CTS solution options. Thus, we defer the details to Part V below, and explain the CTS option for real-time scheduling first.

The important point to note here is that the day-ahead settlement process does not change the basic economics of congestion pricing, or that loads must pay generation the appropriate LMP for power exported across the interface to serve them. The ISOs' role as joint settlement administrators for power flows between regions ensures the appropriate payments are made, and congestion revenue is accrued properly.

IV. SOLUTION OPTION B: COORDINATED TRANSACTION SCHEDULING

A. CONCEPT AND DESIGN PRINCIPLES

The second solution option is a package of external transaction enhancements called *coordinated transaction scheduling* (CTS). The CTS has four major elements:

- High frequency scheduling (HFS) across external interfaces;
- Elimination of charges/credits on external transactions that deter trade;
- A simplified bid format, called an *interface bid*, for real-time scheduling; and
- Coordinated acceptance of bids by the ISOs, using an improved clearing rule.

The first two of these elements are substantively the same as the Tie Optimization option. The second two of these elements differ.

The core philosophy of the CTS is that the ISOs will determine the schedule across the (AC) interfaces between NYISO and ISO-NE using external transaction offers. Like the current external transaction system, under CTS the ISOs will use two sets of market-based offers to set the tie schedule: (1) participants' external transaction offers to buy and sell across the interface, cleared against (2) the real-time generation supply stacks in each region. However, the structure of the external transaction bid format, and how it clears, differs between CTS and the existing inter-regional trading system. These differences are designed to solve the root causes of the current system's inefficiencies, as explained in section II.C.

The CTS solution option differs from the Tie Optimization solution option in a broad way. Conceptually, the CTS is more like the current external transaction system than Tie Optimization: CTS retains a role for external transaction offers to determine tie schedules and real-time LMPs. In contrast, the Tie Optimization option looks like the least-cost economic dispatch process used internally by each ISO: It relies on only the bid-based supply offers from generators (along with load information) to determine real-time LMPs and all transmission flows.

CORE DESIGN PRINCIPLES

The core design principles governing CTS are the same as those for Tie Optimization. The differences arise in how these design principles are applied.

In brief, the principles are four:

1. **IMPROVED EFFICIENCY.** The CTS is designed to operate the transmission interface more efficiently than the current inter-regional trading system, meeting the total power demand of both regions at lower cost than today.
2. **MARKET-BASED TRANSMISSION FLOWS.** The ISOs use competitive, market-based supply offers from market participants to determine the power system's dispatch and net external tie schedule.
3. **ALL SETTLEMENTS AT LMP.** All energy flows between regions are priced at LMP. This facilitates market transparency and prices congestion across the interface. Under either reform option, there are no uplift debits or credits associated with inter-regional transmission flows.
4. **ISOs HAVE NO FINANCIAL POSITION IN MARKETS.** The ISOs do not directly participate in markets, and do not buy nor sell. Loads pay LMP for all power consumed, wherever the generation is located; generators receive the LMP at their locations, regardless of where the power goes. The ISOs will continue to act as independent settlement administrators for the payments to and from market participants.

B. ECONOMIC FRAMEWORK

Two central elements of the CTS are the simplified bid format for real-time external transactions, and the clearing rule used to determine the net tie schedule with these bids. We address each in turn.

THE INTERFACE BID FORMAT

An *interface bid* is a unified transaction to buy and sell power simultaneously on each side of the interface. It settles at the real-time price difference. For example, an interface bid across the NYN interface is an offer to buy at the proxy bus on one side of this interface, and sell on the other.

Mechanically, an interface bid consists of three numbers: A price, a direction, and a quantity. The price indicates the minimum expected price difference between the two nodes that the participant is willing to accept. The direction indicates at which proxy bus the participant wants to buy, and at which it wants to sell. The quantity indicates

how many megawatts the participant is willing to transact. If the interface bid is accepted (i.e., clears), the participant is paid the real-time price difference between the two nodes.

Like external transactions today, a market participant that submits an interface bid incurs no physical delivery obligation. All interface bids settle financially. From the ISOs' perspective, interface bids are simply a market-based device with which the ISOs determine the real-time net tie schedule between ISO-NE and NYISO.

DETERMINING THE TIE SCHEDULE

Under CTS, a market participant submits an interface bid to either ISO. All bids are pooled by the two ISOs, who apply a coordinated acceptance (clearing) process that determines the net tie schedule.

The economic logic of this process is simplest to convey graphically. Consider Figure IV-1. At a specific point in time, imagine the black curve represents the generation supply offer stack in one ISO (in this case, NYISO). This stack characterizes the incremental cost of delivering power at its side of the (NYN) interface.²⁴ Similarly, the blue curve represents the generation supply offer stack of the other ISO (in this case, ISO-NE), shown here in *descending* cost order.

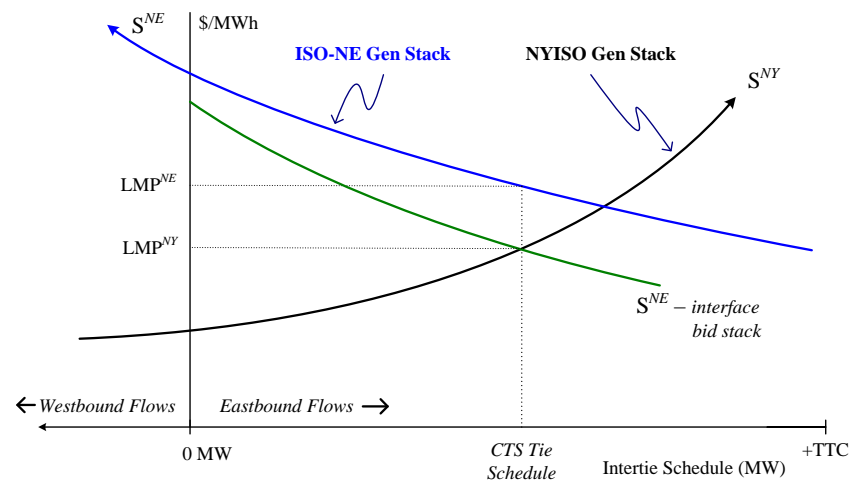


Figure IV-1. Determining the tie schedule using interface bids

²⁴ More precisely, the supply curve is NYISO's incremental dispatch cost at the NYN interface proxy bus. Like the Tie Optimization option, each supply curve shown will incorporate all generation shift factors changes and internal transmission constraints that may bind and affect the proxy bus price over any segment of the range of admissible tie flows.

Unlike Tie Optimization, the ISOs will not automatically set the net tie schedule where the ISOs' generation supply stacks intersect. Instead, they will modify the supply stacks by simultaneously clearing the interface bids. The clearing rule is: An interface bid is accepted if the offered price is less than the expected LMP difference across the interface at the time the ISOs set the interface schedule.

How would the ISOs implement this rule? Mechanically, it takes three steps. First, the ISOs must assemble the total *interface bid stack*. Imagine taking the set of all interface bids indicating the same direction (say, eastbound). These bids are stacked, from lowest to highest price, to create their own 'supply curve'.

Second, we modify the expected generation supply curve of ISO-NE by *subtracting* the (eastbound) interface bid supply curve. In Figure IV-1, the result is depicted as the descending (green) curve. The CTS interface schedule is then set where the green curve intersects the NYISO's expected generation supply stack. All interface bids to the left of the CTS tie schedule are accepted; all interface bids to the right of the CTS tie schedule are not.

As drawn in Figure IV-1, at zero tie flow ISO-NE has higher costs than NYISO. Thus the tie schedule is eastbound. The logic applies similarly in hours when NYISO has higher costs than ISO-NE at zero flow. See Figure IV-2. Here the ISO's assemble the interface bid stack for all offers in the *westbound* direction, and *add* it to ISO-NE's generation supply stack. Again, the CTS tie schedule is set where the marginal interface bid equals the difference in LMPs across the interface.

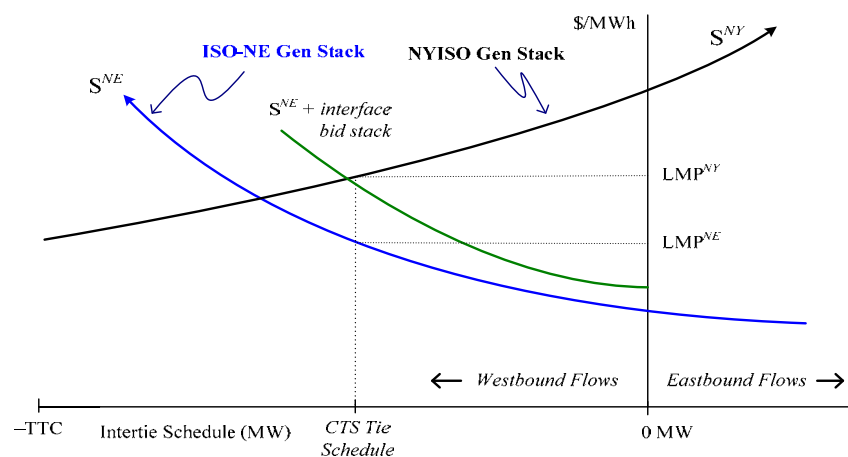


Figure IV-2. Determining a westbound CTS tie schedule

Figures IV-1 and IV-2 reveal the basic settlement logic. For the moment, suppose there is no congestion, nor losses. Load on the import side (whichever ISO that may be) pays the LMP at its locations. In the importing region, internal generation is paid the same LMP. However, there is less generation than load in the importing ISO, so the importing ISO receives more revenue from load than it pays to its internal generation.

Where does the excess revenue go? It is paid to the market participants whose interface bids cleared. In this sense, participants with accepted interface bids are selling power to the importing region's loads.

The exporting ISO's settlement logic is symmetric. All generation is paid the LMP at its location. However, there is more generation than load in the exporting ISO. The market participants whose interface bids cleared cover the difference: They pay the exporting region's LMP for each MW cleared at the interface. In this sense, participants with accepted interface bids are buying power supplied by the exporting region's generation.

Like internal dispatch today, there are additional settlement elements that arise due to congestion and marginal losses: LMPs at the border will not always equal internal node LMPs, and the tie may be congested. Nevertheless, the basic settlement logic across the interface generalizes in straightforward way. We go through the settlement details, including day-ahead market interactions, in Part V.

EFFICIENCY

Figures IV-1 and IV-2 show situations without transmission capacity limitations, and reveal that CTS does not equate expected locational marginal prices on each side of the interface. The expected LMP in the importing region is higher than the LMP in the exporting region (unless the last accepted interface bid price is zero). As a result, the CTS interface schedule megawatts will be less than the interface schedule set by Tie Optimization (section III.B) when interface capacity is not limited. In theory, this means the CTS process is less efficient than Tie Optimization: If expected prices do not converge completely, one ISO continues to operate higher-cost generation that could be displaced by lower-cost generation from the other region—but is not.

A key property of the CTS bid format and clearing rule is that it solves the *non-economic clearing* problem that is a root cause of scheduling inefficiencies today (section II.C). The clearing rule ensures that, if a transaction is accepted, the expected price in the importing ISO is higher than the expected price in the exporting ISO. In the terminology of section II.C, the CTS clearing rule always satisfies condition (A). That ensures the interface schedule will never be set to flow in the wrong direction at the time the schedule is evaluated.

There is a similarity and a difference between the CTS option and the Tie Optimization in this regard. In Section II.C (Figure II-8), we showed that the current trading system frequently produces counter-intuitive schedules, and frequently exhibits under-utilization, at the time the tie is scheduled. Both the CTS option and the Tie

Optimization option completely solve the *first* of these two problems: By ensuring the importing region's LMP is at least as large as the exporting region's LMP, both options preclude interface schedules that are counter-intuitive (cost-increasing) at the time scheduled.

However, the two options differ with respect to tie-underutilization. Tie Optimization converges expected LMPs when the schedules are set, while the CTS leaves the importing ISO's LMP higher than the exporting ISO's LMP (unless the marginal interface bid is zero). That means the interface would be under-scheduled relative to the (efficient) level that minimizes total production costs.

C. MECHANICS: HFS AND SUBMISSION TIMING

Both the Tie Optimization and the CTS solution options share the same higher-frequency scheduling (HFS) system. Under CTS and Tie Optimization, the ISOs require the same information on one another's current bid stack at the border, as indicated in Figures IV-1 and IV-2, to determine the tie schedule. Mechanically, the only modification to HFS between Tie Optimization and CTS is that the interface bid stack is added to, or subtracted from, ISO-NE's generation bid stack before the interface schedule is calculated.

The ISOs anticipate that the precise timing of information flows between NYISO and ISO-NE to implement HFS would be the same under either option. Accordingly, we refer the reader to section III.C for the indicative timeline and individual steps.

INTERFACE BID SUBMISSION

As noted above, interface bids are pooled and conveyed to both ISOs. To do so, the ISOs anticipate developing a common bid submission platform for market participants submitting interface bids. The platform would provide a one-stop, fully-automated bid submission and validation tool, eliminating today's cumbersome submission procedures. In addition, the common submission platform would eliminate 'fail to check-out' outcomes, an inefficient source of financial risk for market participants submitting external transaction requests today.

One important element under CTS is the timing of interface bid submission. Current external transaction rules require schedule offer requests to be submitted 75 minutes before the delivery hour on the NYISO side, and 60 minutes on the ISO-NE side. This will need to be a single point in time, since all bids would be submitted to a common portal. The ISOs anticipate the bid time under CTS will remain 75 minutes in advance of the start of the delivery period to which the interface bid applies. This is to accommodate the look-ahead information needs of the ISOs dispatch and commitment systems, which

assess changes in physical tie schedules 75 minutes forward as an input into generation dispatch and real-time commitment optimization.

D. WHAT ABOUT CONGESTION?

At times, there may be binding transmission constraints that prevent the CTS clearing rule from setting the LMP difference across the interface equal to the marginal interface bid. The CTS addresses this using a modified congestion pricing rule that differs from internal congestion pricing.

CTS CONGESTION COST

Under CTS, real-time congestion pricing across the interface has an additional element not applicable to pricing across *internal* transmission links. The additional element is the marginal interface bid. The payment made to each accepted interface bidder reduces the expected congestion revenue that accrues under CTS, relative to congestion pricing under Tie Optimization or between internal network nodes.

To illustrate the difference, consider Figure IV-3. As before, the black curve represents the generation supply offer stack in one ISO (in this case, NYISO), indicating its incremental cost of delivering energy to its side of the interface. Similarly, the blue curve represents the generation supply offer stack in the other ISO (in this case, ISO-NE) on its side of the interface, shown again in *descending* cost order.

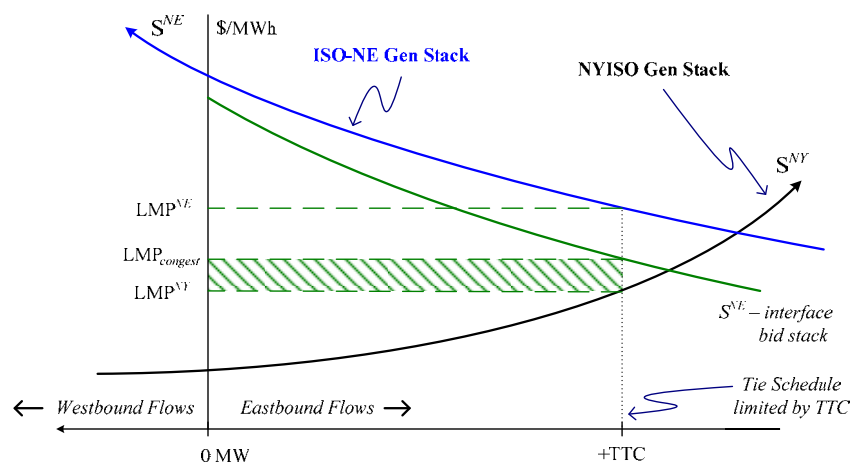


Figure IV-3. Payments to interface bids reduce the amount of congestion revenue collected if the tie's transmission capacity is binding.

CONGESTION RENT

In this situation, the total transfer capability (TTC) of the interface is less than the tie schedule that CTS would normally set—where the last accepted interface bid equals the difference between the two ISO’s generation stacks. The tie schedule is therefore limited to the TTC, and the HFS process will set the tie schedule to that value.

The binding transmission constraint across the tie ensures price separation between markets. For the importing region, the importing ISO’s LMP calculator will set the real-time LMP on its side of the interface equal to its marginal cost of energy delivered to the interface. This is shown as LMP^{NE} in Figure IV-3. At all internal nodes in the importing ISO (here, ISO-NE), this is the real-time LMP paid by loads and paid to generation (absent internal congestion and losses).

For the exporting region, there is a lower real-time (proxy bus) LMP equal its marginal cost of energy delivered to the interface. This is shown as LMP^{NY} in Figure IV-3. In the absence of any internal congestion (or losses), this is the real-time LMP paid by loads and paid to generation at all internal nodes in the exporting ISO (here, NYISO).

In Figure IV-3, there is a congestion price for energy delivered across the constrained interface. The congestion price is where the modified congestion price rule applies. As drawn in Figure IV-3, the importing ISO credits each accepted interface bid at a price equal to its real-time internal LMP, or LMP^{NE} in Figure IV-3. The exporting ISO debits each accepted interface bid at a price equal to its real-time LMP plus the congestion price, or $LMP_{congest}$ in Figure IV-3.

In this way, the congestion rent equals the shaded (green) box in Figure IV-3. The (real-time) congestion price under CTS is the difference between each ISO’s marginal cost of delivering energy to the interface, less the price of the last accepted interface bid.

What happens with the congestion revenue collected under CTS? In implementation, the ISOs propose to use a modified settlement procedure from that indicated in Figure IV-3. The modification produces the same congestion rent and payments to interface bids. However, the modification allocates the congestion revenue in equal measure to the existing congestion revenue funds administered by each ISO. The precise settlement method is described in Part V, next.

In addition, under both Tie Optimization and CTS congestion revenue would accrue (primarily) in the day-ahead market—where most energy is transacted—not in the real-time market. This provides the revenue stream necessary for each ISO to fund risk-management products (TCC/FTR) across transmission interfaces between regions. We describe the structure and funding of each ISO’s TCC/FTRs for the interface in Part VI.

V. DAY-AHEAD MARKETS AND INTERACTIONS

The Tie Optimization and the CTS options apply to the determination of inter-regional energy interchange in real-time. This report does not propose to apply either option to the day-ahead markets. Each ISO will continue to operate separate day-ahead markets, with separately-scheduled external transactions. However, there are important interactions between each ISO's day-ahead market and the real-time prices and interface schedules under an economic coordination system like Tie Optimization or CTS.

In this section we explain how each ISO's day-ahead market interacts with the real-time market under the Tie Optimization and CTS options, including the treatment of day-ahead external transaction offers and settlement procedures.

A. DAY-AHEAD MARKET ISSUES

There are two reasons why the solution options presented above focus on real-time, as opposed to day-ahead, interface schedules. First, to solve the problems documented in section II, the interface schedule must be coordinated efficiently in real-time. If the scheduling process does not produce price convergence in real-time (up to congestion), the ISOs are incurring higher production costs than necessary.

Second, close coordination of ISO-NE's and NYISO's day-ahead markets has less certain incremental benefits. At present, the ISOs' day-ahead markets are *asynchronous*: They operate during non-overlapping time periods prior to the operating day. In NYISO's day-ahead market, bids are due at 5 AM and the market clears by 11 AM prior to the operating day; in ISO-NE's day-ahead market, bids are due by noon and the market clears by 4 PM.

Asynchronous markets are difficult to coordinate. The problem is that NYISO's earlier-closing day-ahead market cannot directly incorporate information about the price of power available from ISO-NE, which won't be known for several hours. Similarly, after the price of power in ISO-NE is known later that day, participants cannot revise offers in NYISO's day-ahead market because that market is closed. These timing issues mean it is not possible to construct a clearing algorithm that determines an interface schedule to equate day-ahead locational marginal prices between regions.

Nevertheless, it is possible to achieve a level of coordination between day-ahead markets based on market participants' expectations about real-time prices. These expectations are incorporated into day-ahead prices by allowing participants in each day-ahead market to buy or sell at the interface proxy bus between NYISO and ISO-NE.

B. DAY-AHEAD EXTERNAL TRANSACTIONS

There are two ways in which a market participant could offer to buy or sell at the interface between regions in the day-ahead market: External transactions and virtual transactions. The two are treated the same in the day-ahead market.

An external transaction offer in the day-ahead market is an offer either to buy, or to sell, that settles at the day-ahead locational marginal price at the external interface. Like today, an accepted external transaction is a binding financial arrangement between a market participant and the ISO. The external transaction does not create a physical obligation for the market participant.

With minor modifications to today's practices, both the Tie Optimization and the CTS options can accommodate external transaction offers in each ISO's day-ahead energy market. The minor modifications relate to how day-ahead external transactions are linked and settle in to the real-time market.

LINKING DA AND RT EXTERNAL TRANSACTIONS

Under CTS, an external transaction that clears an ISO's day-ahead market can be linked to a specific interface bid in the same direction submitted by the same market participant. If the interface bid also clears, then the external transaction would be deemed "to flow" in real-time. That means a participant submitting an external transaction that clears in each ISO's day-ahead market, and that clears (for the same MW) as an interface bid in real-time, would have no balancing (deviation) charges associated with real-time settlement.²⁵

Note, however, that interface bids can only clear in the economically-correct direction. That is a central objective of the CTS design: By only clearing interface bids in the economically correct direction, it solves one of the root causes of today's inefficiencies (see *Non-Economic Clearing* in section II.C). The point to note is that day-ahead external

²⁵ In NYISO, credits and debits associated with differences between a market participant's day-ahead and real-time market positions (cleared MW) are called *balancing charges*; in ISO-NE, they are commonly called *deviation charges*. We use both terms synonymously.

transaction offers need to be in the economically-correct direction to avoid balancing (deviation) charges at real-time settlement.

The same role for day-ahead external transactions can be achieved under Tie Optimization. Under that design option, a market participant submitting a day-ahead external transaction may request that the ISOs deem the transaction “to flow” in real time. Like today, a day-ahead external transaction that is deemed to flow in real time under Tie Optimization would not incur balancing (deviation) charges associated with real-time settlement. A day-ahead external transaction that is not deemed “to flow” in real time would incur a balancing (deviation) credit/debit, settled at the real-time LMP.

Under Tie Optimization, the market participant’s day-ahead external transactions need to clear each ISO’s day-ahead market (for the same MW) in order to avoid balancing (deviation) charges associated with real-time settlement. In addition, as with CTS, day-ahead external transactions need to clear in the economically-correct direction to be deemed to flow in real time.²⁶

RATIONALE FOR PRESERVING EXTERNAL TRANSACTIONS

Why provide day-ahead external transaction functionality, in a world in which the real-time net interface schedule is optimized with either CTS or Tie Optimization? There are three main reasons.

First and foremost, external transaction offers can help improve price formation at the external interface in the (separate) day-ahead markets administered by each ISO. If there are more day-ahead offers to buy or sell at the external interface, then there is more information going into the determination of day-ahead prices, day-ahead quantities, and the market’s prediction of the economically-correct power flow direction across the interface. Information on the economically-correct power flow direction for the next day, and its approximate magnitude, helps the ISOs to make day-ahead generator commitment decisions in the most cost-effective way.

Second, external transactions enable a market participant to lock-in the prices associated with buying and selling energy across the interface on a day-ahead basis. That allows the participant to avoid being exposed to the volatility of real-time price differences across the interface. Of course, those price differences should be much lower and much less volatile than they have been in the past, under either of the two solution options presented here.

²⁶ Under both CTS and Tie Optimization, if real-time TTC is less than the total MW of day-ahead cleared external transactions, (some) day-ahead transactions would not be deemed to flow and would incur balancing charges at real-time settlement.

Last, external transaction functionality enables a market participant with an existing contractual arrangement that obligates it to schedule an external transaction between regions to continue to satisfy that obligation. Such obligations can arise for a number of reasons outside the purview of the ISOs, such as the external scheduling requirements associated with some states' renewable energy certificate (REC) or renewable portfolio standard (RPS) programs.

C. INTERFACE SETTLEMENT METHODS

There are several methods that could be used to settle day-ahead and real-time energy schedules across the interfaces between NYISO and ISO-NE under Tie Optimization or CTS. In this section, we step through one of these possible settlement methods. Although settlements methods are (necessarily) detailed, they also convey useful information about how many of the different components of the two solution options described above work, and interact, with the day-ahead markets.

The settlement details also indicate how congestion revenue accrues across the interface is treated under CTS or Tie Optimization. That is necessary for the ISOs to be able to fund TCC/FTR products that can be used by market participants to hedge energy price risk at the interface.

The settlement method described here has several desirable properties:

- It enables (expected) congestion revenue across the interface to accrue in the day-ahead markets, even though the two ISOs' day-ahead markets continue to operate asynchronously;
- The day-ahead congestion revenue can fund a TCC/FTR that hedges day-ahead LMP differences at the interface.
- It enables the day-ahead markets to produce information about the expected real-time power transfers between regions, which can help improve each ISO's unit commitment decisions prior to the operating day.

The settlement method discussed here is applicable, with minor differences, to either the Tie Optimization option or to the CTS option for real-time scheduling. We consider the Tie Optimization case first. The extension to the CTS case involves an additional element to accommodate compensation to participants submitting real-time interface bids, explained subsequently.

PROXY BUS PRICES

Under standard market design, the settlement of scheduled energy flows across an external interface is based on *external proxy bus* prices. An ISO's external proxy bus

price is the bid-based incremental (dispatch) cost of delivering another megawatt to the ISO's own side of interface.

Note an ISO's external proxy bus price includes the cost of *internal* congestion to the interface, and the cost of energy losses to the interface. The external proxy bus LMP would be the same as the internal LMPs in the ISO's network if there is no internal congestion or losses.

In real-time, there is a *sending* (exporting) ISO and a *receiving* (importing) ISO across the interface. Figure V-1 illustrates the external proxy bus prices set by the sending and receiving ISOs when the external interface transmission capacity is binding in real-time. In this event, there is price separation between markets.

The difference in each ISO's external proxy bus prices equals the price of congestion across the interface in real-time. This is shown as the height of the shaded (green) area in Figure V-1. As discussed next, the congestion price is divided into two equal components, which serve to apportion any real-time congestion revenue equally to each ISO's congestion revenue fund.

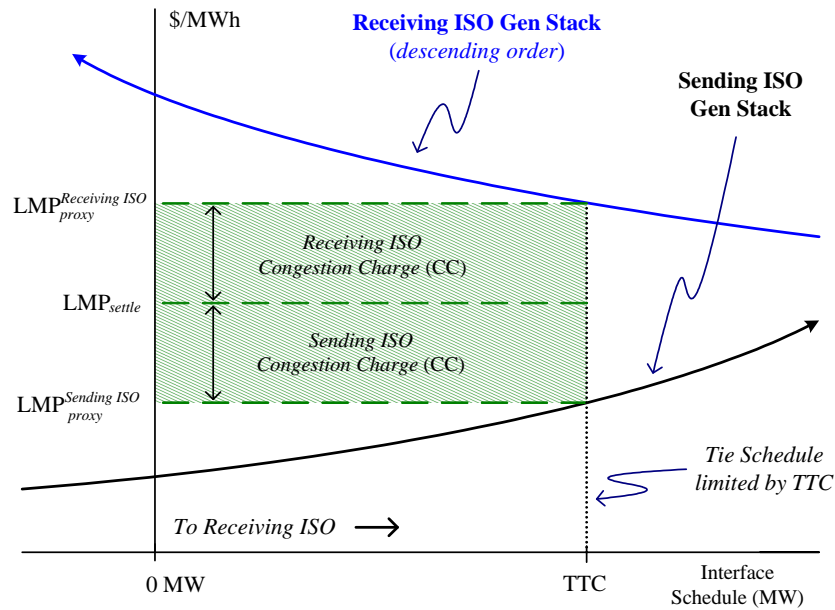


Figure V-1. Real-time external proxy bus LMPs, interface settlement LMP, and interface congestion charges (CC) under Tie Optimization.

SETTLEMENT RULES: REAL-TIME

Under either Tie Optimization or CTS, all load and generation pays, or is paid, the LMP at its location. For the receiving ISO, more revenue is paid by internal load than must be paid to internal generation. This difference is attributable to the imported megawatts. Under Tie Optimization, the two ISOs—acting as joint settlement administrators—must transfer funds from loads in the importing region to compensate generators in the exporting region for the imported megawatts. Under CTS, a similar transfer of funds between receiving ISO loads and sending ISO generation occurs, but through debits and credits to the market participants with accepted interface bids.

The primary complication that external interface flows pose for settlements is that when congestion occurs across an external interface, the congestion revenue must be allocated—explicitly or implicitly—between regions. While there are a number of different means to do so, one of the simplest, most transparent, and equitable methods is based on the concept of an *interface settlement LMP*.

Under Tie Optimization, the real-time interface settlement price is simply the midpoint between the sending and receiving ISOs' real-time external proxy bus prices:

$$\text{Settlement LMP} = \frac{1}{2} \times (\text{Sending ISO RT proxy bus LMP} + \text{Receiving ISO RT proxy bus LMP})$$

The interface settlement LMP is shown as LMP_{settle} in Figure V-1. If there is no price separation between markets, then each ISO's external proxy bus LMP will equal the settlement LMP.

The settlement and proxy bus LMPs provide for a simple method to settle the real-time energy flows across the interface, and to allocate equally any real-time congestion revenue that accrues across the external interface. Under Tie Optimization, a transfer is performed by the ISOs at the interface settlement price using an interface settlement account:

- The receiving ISO credits an interface settlement account, in an amount equal to the optimized tie schedule MW times the settlement LMP;
- The sending ISO debits the interface settlement account, in an amount equal to the optimized tie schedule MW times the settlement LMP.

The interface settlement account always nets to zero.

This method is simplest in the case where there is no price separation within or between each region. In that case, the interface settlement price equals the LMP paid by loads at its location and equals the LMP paid to generation at its location.

When there is price separation between regions, this settlement process divides the total (real-time) cost of congestion across an external interface into two equal parts. In Figure V-1, these parts are the Sending ISO Congestion Charge and the Receiving ISO

Congestion Charge. In this way, each ISO's congestion revenue fund accrues an equal amount of the total real-time congestion revenue across the external interface.

In sum, this settlement method has the following properties:

- All generation is paid the LMP at their locations, and all load pays the LMP at their locations, as usual under standard market design;
- The difference in real-time proxy prices at the external interface transparently conveys the correct real-time cost of congestion (i.e., the marginal opportunity cost of limited transmission capacity between regions);
- Any real-time congestion revenue across the external interface accrues in equal measure to the congestion revenue funds administered (separately) by each ISO.

SIMPLE EXAMPLES

Imagine day-ahead cleared load equals real-time load in each region, and there are no external transactions cleared in the day-ahead markets. Assume no internal congestion within each ISO, and no energy losses.

Scenario A. Suppose the RT LMP in each ISO is \$50 / MWh, and Tie Optimization sets a real-time schedule of 1000 MW across an unconstrained external interface. Generation in the sending ISO has a positive deviation from day-ahead of 1000 MW in real-time, and generation in the receiving ISO has a negative deviation from day-ahead of 1000 MW in real time. In real-time settlement, the receiving ISO debits generators with the negative deviation that hour a total of \$50,000 (= \$50 / MWh x 1000 MW). It transfers this amount to the sending ISO, and the sending ISO credits internal generators with the positive deviation the same amount.

Scenario B. Suppose the RT LMP is \$50 / MWh in the sending ISO, \$70 / MWh in the receiving ISO, and Tie Optimization sets a real-time schedule of 1200 MW that congests the external interface. Generation in the sending ISO has a positive deviation from day-ahead of 1200 MW, and generation in the receiving ISO has a negative deviation from day-ahead of 1200 MW. In real-time settlements:

- The receiving ISO debits internal generation at its LMP of \$70 / MWh, in an amount of \$84,000 (= \$70 / MWh x 1200 MW);
- The interface settlement price is \$60 / MWh, so the receiving ISO transfers \$72,000 (= \$60 / MWh x 1200 MW) to the sending ISO;
- The sending ISO credits internal generation at its LMP of \$50 / MWh, in an amount of \$60,000 (= \$50 / MWh x 1200 MW);

- Each ISO accrues a credit to its congestion revenue fund of \$12,000, equal to each ISO's interface congestion charge of \$10 / MWh times the 1200 MW interface schedule.

The important points to note about this example are two. First, market participants with real-time balancing charges (deviations) from day-ahead positions are credited or debited at the LMP at their location, as usual under standard market design.

Second, regardless of the magnitude of the price separation between regions, each ISO has the same real-time congestion price across the external interface. If the numbers assumed in Scenario B were instead applied between two locations separated by a congested interface *within* a single ISO, the congestion price across the internal interface would be \$20 / MWh. When it applies across an *external* interface, each ISO applies an interface congestion price of half as much. By doing so, the total congestion price applied by the two ISOs—here \$20 / MWh—is the economically-correct congestion price under standard market design.

CTS INTERFACE SETTLEMENT PRICES

The principles underlying the settlement method outlined above apply to both Tie Optimization and CTS. However, there are some important differences. These differences are necessary to accommodate settlement of the interface bids under CTS.

Under CTS, there are *two* real-time interface settlement prices. Each ISO's interface settlement price is equal the ISO's external proxy bus price adjusted for the cost of congestion across the interface:

- Receiving ISO's settlement LMP = Receiving ISO's RT proxy bus LMP
– Scheduled Congestion Charge
- Sending ISO's settlement LMP = Sending ISO's RT proxy bus LMP
+ Scheduled Congestion Charge

To interpret the settlement LMPs and Scheduled Congestion Charges, a graph may help. Figure V-2 (*next page*) shows how the cleared interface bids affect real-time settlement under CTS. As before, the external proxy bus prices represent each ISO's bid-based incremental (dispatch) cost of delivering another megawatt to the ISO's *own* side of interface.

The supply curve in the lower axes in Figure V-2 is the interface bid stack (only interface bids in the economically correct direction are shown). When the external interface is congested in real time, the last accepted interface bid is the marginal interface bid (MIB). In the upper axes, the marginal interface bid is the height of the (unshaded) box between the sending and receiving ISO settlement prices. Each scheduled congestion charge (SCC) per MW is the height of the green box between the ISO's proxy bus LMP and its settlement LMP.

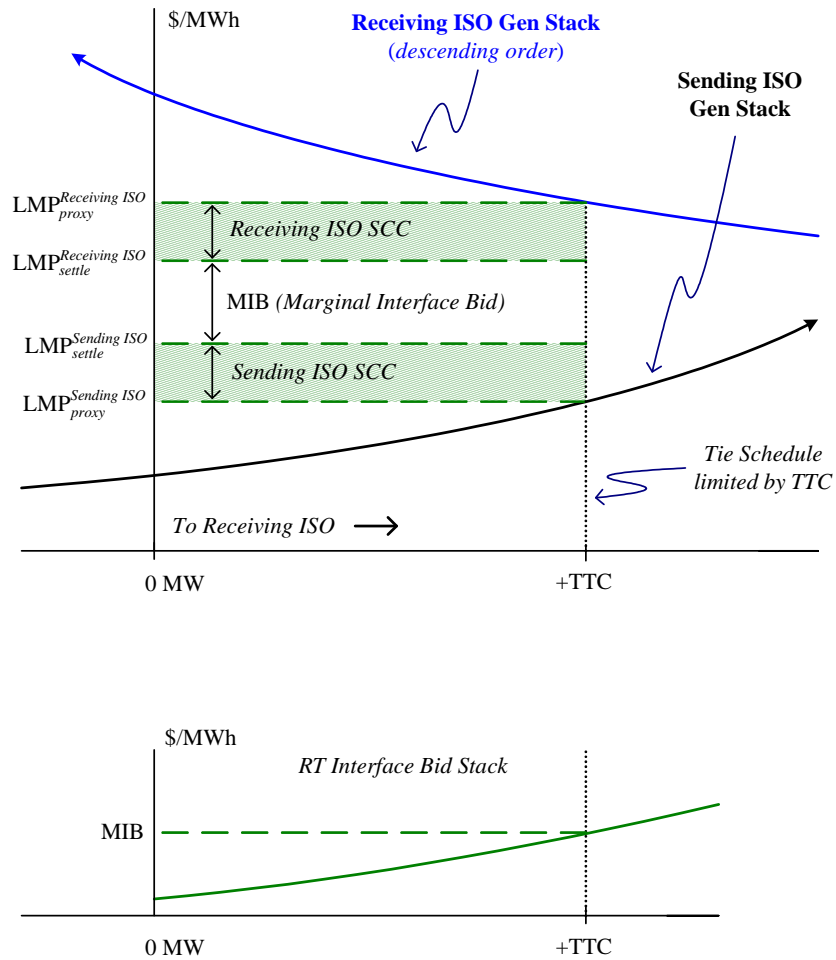


Figure V-2. Scheduled congestion charges (SCC) and marginal interface bid (MIB) under CTS when the interface is congested in real-time.

These settlement and proxy bus LMPs provide for a method to settle the real-time energy flows across the interface, and to allocate any (real-time) congestion revenue between markets. The real-time settlement rules for interface bids are two:

- The receiving ISO credits participants submitting cleared interface bids, in an amount equal to the cleared interface bid MW times the receiving ISO settlement LMP;
- The sending ISO debits participants submitting cleared interface bids, in an amount equal to the cleared interface bid MW times the sending ISO settlement LMP.

All participants with cleared interface bids are charged a uniform price by the receiving ISO, and receive a uniform price from the sending ISO. That means the net gain to a participant with a cleared interface bid is the MIB price times the quantity (in MW) of cleared interface bids it submitted.

Note, however, that if a participant with a cleared interface bid also has a cleared day-ahead external transaction for the same MW, then the two transactions will offset in real-time settlement. The interface bid settlement rules generate net cash flows in real-time settlements only for the *deviation* (imbalance) in cleared quantities between the interface bid and the day-ahead transaction.²⁷

CTS SCHEDULED CONGESTION CHARGES

Figure V-2 shows a situation where there is congestion in real-time. There is one aspect of the SCC that applies to the CTS option and is not readily conveyed by Figure V-2.

While the external proxy bus prices are real-time LMPs, the SCC is determined when the tie is *scheduled*. Under HFS, the interface is scheduled before the real-time prices are known. The *expected* real-time external proxy bus LMPs at the time the tie is scheduled are called as the ISOs' *scheduling prices*. Under CTS, the SSC is calculated by subtracting the marginal interface bid (MIB) from the difference in scheduling prices:

$$\text{SSC} = (\text{Receiving ISO's scheduling price} - \text{Sending ISO's scheduling price}) - \text{MIB}$$

If there are no changes in system conditions between when the tie is scheduled and when real-time internal LMPs are determined, then the scheduling price and the RT proxy bus LMP will be the same.

²⁷ This is achieved by crediting (or debiting) the external transaction at the real-time proxy price. That is a net zero settlement if the interface bid and day-ahead external transaction have an identical cleared MW, and non-zero for MW deviations.

Why set the congestion charge, and thus the sending ISO's proxy price, based on *scheduling* prices? The reason is that it affects who bears latency risk. Even under HFS, there exists a (15 minute) latency delay between when the tie schedule is set and when the power actually flows. If, for instance, the sending ISO experiences a contingency that raises its real-time internal LMP above the price expected when the tie is scheduled, the sending ISO debits the interface bidder the sending ISO's actual real-time external proxy bus LMP less a fixed SCC charge, times the cleared interface bid MW. In this way, interface bidders (that do not have a matching day-ahead position) bear some of the financial risk associated with latency delays under HFS.

REAL-TIME SETTLEMENT OF DAY-AHEAD EXTERNAL TRANSACTIONS

The interface settlement rules for Tie Optimization and for CTS generalize to handle real-time settlement of external and virtual transactions at external proxy buses. The key point is that for any day-ahead external or virtual transaction at the external proxy bus that incurs real-time balancing (deviation) charges, the real-time settlements are priced at the external proxy bus LMP plus (or minus) the interface congestion charge. That is, external transactions are priced at the interface settlement LMP.

Although numerical examples become detailed, the settlement rules are conceptually straightforward. Under both Tie Optimization and CTS:

- A cleared day-ahead virtual transaction at an external proxy bus is settled in real-time at the ISO's interface settlement LMP.

Under Tie Optimization:

- A cleared day-ahead external transaction that is not deemed to flow in real time is settled in real-time at the interface settlement LMP;
- A cleared day-ahead external transaction that is deemed to flow in real time has no balancing (deviation) charges associated with real-time settlement.

Under CTS, the rules are similar but a day-ahead external transaction requires a corresponding cleared interface bid in order to avoid balancing (deviation) charges in real-time settlement:

- A cleared day-ahead external transaction that does not have a matching cleared interface bid for the same MW is settled in real-time at the ISO's interface settlement LMP;
- A cleared day-ahead external transaction that has a matching cleared interface bid for the same MW has no balancing (deviation) charges associated with real-time settlement.

SETTLEMENT RULES: DAY-AHEAD

The settlement rules for external and virtual transactions in the day-ahead market at the proxy bus follow standard market design. These rules operate in the same way under CTS and Tie Optimization. Since the day-ahead markets operate asynchronously, day-ahead settlement, day-ahead external proxy bus pricing, and day-ahead external interface congestion charges are applied by each ISO separately.

The main day-ahead settlement rules for external (and virtual) transactions at an external proxy bus pricing point are these:

- A day-ahead external transaction or virtual transaction offer to buy (export) at the proxy bus clears if the offer price exceeds the day-ahead proxy bus price; and
- A day-ahead external transaction or virtual transaction offer to sell (import) at the proxy bus clears if the offer price is less than the day-ahead proxy bus price.
- The proxy bus price equates the total MW of supply and demand that clear at the proxy bus (unless that level of MW exceeds total transmission capability).²⁸

Real-time charges and credits are based only on deviations from day-ahead positions. In this way, a market participant submitting a day-ahead external transaction can (largely) avoid exposure to real-time prices across the interface.

ECONOMIC IMPLICATIONS

The basic settlement method outlined above has a number of important economic implications.

First, it enables congestion revenue associated with real-time interface transmission constraints to accrue in the day-ahead market, to the extent that day-ahead prices anticipate it. That enables a traditional FTR/TCC instrument to be funded from day-ahead congestion revenue.

Second, the method is applicable to both the CTS option and the Tie Optimization option (which amounts to treating the interface bids as priced at zero). Thus, to some degree, the decision over whether the ISO's should pursue the CTS option or the Tie Optimization option for the process to determine the real-time net tie schedule can be divorced from the decision over how to set proxy-bus prices and settle the associated cash flows.

²⁸ If total transmission capacity binds day-ahead, then the proxy bus price must be modified by the shadow cost of the transmission constraint. We omit the mathematical details.

Third, the settlement logic reveals that CTS and Tie Optimization differ in the extent to which they enable a TCC/FTR holder to hedge the total (internal) LMP difference between regions. To see why, compare the total congestion surplus (the shaded green areas) in Figures V-1 and V-2: The congestion surplus is larger with Tie Optimization. It is equal to the interface schedule megawatts times the *sum* of the SCCs + MIB shown in Figure V-2. In effect, the profit margin of a participant with a cleared interface bid (the value of MIB) reduces the congestion revenue that would otherwise ultimately accrue to congestion revenue rights holders under Tie Optimization.

VI. HEDGING PRICE RISK AT THE INTERFACE

A subset of market participants actively trade power between regions, or have long-term contracts tied to assets in both New York and New England. For these participants, the existing external transaction system provides limited means to hedge against the risk of price differences between regions. As Figure II-7 indicates, real-time price differences across the interface are historically volatile.

Although Tie Optimization and CTS are both designed to reduce price differences between regions, and should therefore reduce this volatility substantially, congestion cannot be whisked away by market design. A hedging mechanism, such as TCC/FTRs, at the interface may be valuable to market participants.

EXAMPLE

Suppose a market participant with a generator located in New York has a year-long supply arrangement with a municipal distribution utility in New England. Under this arrangement the municipal utility pays the market participant a fixed price per MWh, in exchange for which the participant remits to ISO-NE the day-ahead energy market charges that the municipal utility incurs at its location. (This type of financial arrangement is common, often implemented as an internal bilateral transaction administered by ISO-NE's settlement department).

Under this arrangement, the market participant is subject to four distinct price risks: Energy price risk at the generator; internal congestion in NYISO between the generator's location and the NE/NY border; internal congestion in ISO-NE between the municipal utility's location and the NE/NY border; and price risk across the interface itself.

Existing TCC/FTR instruments enable the market participant to (largely) hedge against the internal congestion price risk within each ISO. However, there is no ready means to hedge the price risk between regions. If the LMP in New England rises relative to the LMP in New York, the market participant ends up with a high bill from ISO-NE for the municipal utility's load obligation, and low revenue at the generator's location in New York with which to cover it.

IMPLICATIONS

The high (historical) volatility of the price difference across the interface suggests a market participant in this position faces a significant financial exposure on its fixed-price contract with the municipal utility. The inability to hedge a significant price risk will inhibit a participant's willingness to enter into a fixed price contract in the first place—or

increase the fixed price the participant is willing to offer in a long-term contract. That reduces the competitiveness of markets for long-term, fixed-price power contracts.

In effect, even if this (hypothetical) municipal utility never deals with the world of TCCs/FTRs directly, or conducts any inter-regional energy trading *per se*, the opportunities it can find to acquire stable-price long-term power supplies in today's market environment are affected by whether hedging instruments across ISOs' external interfaces are well designed.

DESIGN ISSUES

The natural instrument with which to enable market participants to manage price risk across the interface is a TCC/FTR product. Accordingly, both of the real-time interface scheduling options in this report provide a means to fund this financial product.

Economic coordination of real-time power flows enables the two ISOs to set congestion prices at their external (proxy-bus) pricing points in real-time. A settlement system such as that outlined in section V.C will accrue the congestion revenue, on a day-ahead basis, during settlements. This day-ahead congestion revenue can be paid out to the buyer of the TCC/FTR, partially reducing its price risk across the interface.

While the market design for a TCC/FTR at the interface is mostly standard (in the sense of mirroring how a nodal TCC/FTR works between points internal to an ISO's network), there are a few additional wrinkles. These arise because each ISO separately administers the TCC/FTRs between its internal locations and the external interface.

Two TCC/FTRs

The settlement system described above (section V.C) has the property that both ISOs accrue congestion revenue across the interface. This provides for an appealing means to structure a TCC/FTR to pay out this congestion revenue: Each ISO issues a *separate* TCC/FTR to the external interface. The TCC/FTR pays the difference between the (day-ahead) price at an internal location and the ISO's (day-ahead) interface settlement price. An ISO's day-ahead interface settlement price is the ISO's day-ahead price at the external proxy bus plus (or minus) the ISO's day-ahead interface congestion price.

If day-ahead market prices are perfect predictors of real-time prices, then the day-ahead congestion revenue of a TCC/FTR to the external interface will equal the ISO's congestion charge as shown in Figure V-1 (under Tie Optimization) or schedule congestion charge as shown in Figure V-2 (under CTS), plus any internal congestion charge to the external proxy bus. In effect, if day-ahead prices match real-time prices, then the TCC/FTR to the external interface pays its bearer the sum of any congestion cost to the external proxy bus *plus* one-half the total cost of congestion *across* the external interface.

There are some advantages and disadvantages to structuring TCC/FTRs at the external interface as separate products administered by each ISO. From the perspective of a market participant that may wish to acquire the product, it provides flexibility for a

TCC/FTR buyer to manage congestion price risk in each market separately, consistent with the individual day-ahead market executions. However, a TCC/FTR buyer that wishes to acquire insurance against price separation between a location within New York and a location within New England will need to acquire two TCC/FTRs, and participate in the TCC/FTR auctions administered by two ISOs. With this structure, each ISO's TCC/FTR to the external interface remits the day-ahead congestion revenue collected by the ISO's own day-ahead market to parties that acquired its TCC/FTR across the interface.

From an administrative standpoint, there are a number of advantages. As with any TCC/FTR, the ISOs need to cover financial assurance for the instrument, estimate revenue insufficiency risk (that is, how many megawatts of TCCs/FTRs to auction), and decide how far in advance to issue the TCCs/FTRs. Both NYISO and ISO-NE have well established, and somewhat different, administrative procedures and tariff provisions to carry out these functions. Structuring TCC/FTRs as separate products to the external interface will enable each ISO to continue to use its existing process for these essential functions.

REVENUE ALLOCATION

From the perspective of the revenue-rights holders in each region, this separate administration of TCC/FTR products "to the interface" addresses an important issue: How the total revenue from the sale of TCCs/FTRs at the interface accrues to two different ISOs' auction revenue rights holders. Under the two-TCC/FTR structure, each ISO would credit the auction-revenue rights holders in its region alone for the proceeds from the interface TCC/FTR auction instrument it administers. The allocation of auction revenue from TCC/FTR instruments to the external interface would follow each ISO's existing tariff provisions for the allocation of TCC/FTR auction revenue.

CONGESTION RESIDUALS

The third issue to consider is the allocation of real-time congestion residuals (in New England, congestion residuals are often simply called *real-time congestion*). Real-time congestion residuals can arise in two ways. One way occurs if (a) the day-ahead markets do not predict congestion (the day-ahead cleared quantity at an ISO's external proxy bus is less than total transmission capacity across the interface), but (b) in the real-time market total cleared quantities are greater, and the interface transmission constraint binds in real-time. In this case, there is a positive contribution to total congestion revenue from participants' real-time deviations from day-ahead quantities.

Alternatively, congestion residuals can occur when transmission constraints bind in real-time operations at a lower MW rating than the transmission capacity value used to clear the day-ahead market. In this situation, the real-time congestion residual is typically negative.

Both NYISO and ISO-NE have established procedures and tariff language governing the allocation of real-time congestion residuals.²⁹ The only new element here is how to split it between each ISO's market participants when it accrues across the interface, instead of across internal transmission constraints. The settlement method outlined in section V.C addresses this issue in a direct and equitable way. Because each ISO's real-time interface settlement LMP is calculated to always produce equal real-time interface congestion charges by the sending and receiving ISO, any real-time congestion revenue (whether positive or negative) will always be accrued in equal measure to each ISO's congestion revenue fund.

BACK TO CTS AND TIE OPTIMIZATION

It is important to observe that the foregoing TCC/FTR design issues are equally germane to both the CTS option and the Tie Optimization option. Thus, to a significant degree, the decision of whether the ISOs should pursue the CTS option or the Tie Optimization option to schedule real-time net tie flows is distinct from decisions on TCC/FTR auction revenue and instrument design.

A final, more subtle, economic issue should be pointed out in connection with hedging price risk across the interface. This is the fact that price differences between New York and New England are primarily not due to congestion. As shown in Figure II-1, congestion occurs only a few percent of the hours per year (at the NYN interface). Yet the price differences across this interface, as shown in Figure II-6, are non-zero nearly *every hour of the year*. The risk that market participants may wish to hedge between regions is not primarily caused by congestion.

Nevertheless, the two solution options presented here would enable a TCC/FTR, constructed in a standard way, to provide a hedge against price risk for day-ahead price separation at the interface. In addition, a settlement system like that outlined above (section V.C) would enable a market participant to transact on a day-ahead basis at each ISO's external proxy bus, then avoid (most causes of) real-time balancing (deviation) charges under either Tie Optimization or under CTS. For these reasons, the combination of a TCC/FTR instrument across the interface and either Tie Optimization or CTS for real-time net tie scheduling would help market participants to hedge price risk across the interface in a way they cannot do so today.

²⁹ Neither of the real-time net tie scheduling solution options presented in this paper contemplate changes to these procedures.

