

Introduction

This issue of **Quick Takes** describes how the IESO determines prices and schedules simultaneously for both the energy and operating reserve (OR) markets. In order to do so, this document will discuss both joint optimization and system marginal cost pricing. While joint optimization (also known as “co-optimization”) leads to the maximum gain from trade for all participants in the market, it can also lead to what may appear to be the non-intuitive scheduling of individual resources. This paper will explain why.

Please note that the examples and discussion in this paper are greatly simplified in order to convey the principles involved as clearly as possible. This paper is written for a general audience, and therefore, does not attempt to be a comprehensive explanation of the functioning of the algorithm.¹

Background

The IESO administers real-time markets for energy and three classes of operating reserve (10-minute spinning, 10-minute non-spinning and 30-minute non-spinning).² Any offer of OR must be accompanied by an energy bid or offer for an equal or greater number of megawatts.

There are two methods for scheduling energy and OR resources in order to satisfy demand:

- Sequential optimization, where resources are first used to satisfy the demand in the energy market, with any remaining resources used to satisfy the demand in the OR markets.
- Joint optimization, where the offers in all markets are evaluated at the same time in order to satisfy demand in such a way that the economic gain from trade for all participants is maximized by taking advantage of the relative differences in energy and OR offers.

The IESO-administered markets use the second method as it utilizes the available resources in the most cost effective manner for the market as a whole.

¹ For more detailed information concerning the DSO, please refer to: Market Rules Chapter 7 and Appendix 7.5 which are available on the [Rules and Manuals](#) web page.

² Operating reserve is standby power that the IESO can activate in order to meet a contingency on the grid. For additional information on operating reserve, please see the *Introduction to Ontario’s Physical Markets* available on the IESO’s [Marketplace Training](#) web pages.

Economic Gain from Trade

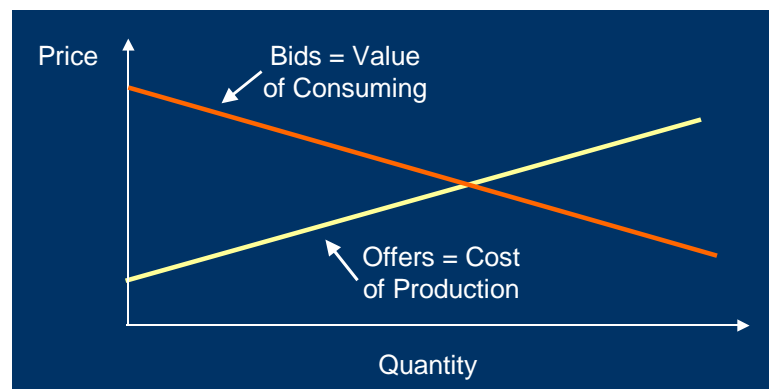
As mentioned above, joint optimization is used for scheduling energy and OR because it maximizes the economic gain from trade across the market. Economic gain from trade for the market is the difference between the value to consumers of the electricity consumed and the cost to suppliers of producing that electricity.

The value placed on electricity differs from one consumer to another, and even from one time to another for the same consumer. Similarly, the cost of producing electricity differs from time to time and from supplier to supplier. Given this, how does the IESO quantify the value of electricity to consumers or the cost to suppliers?

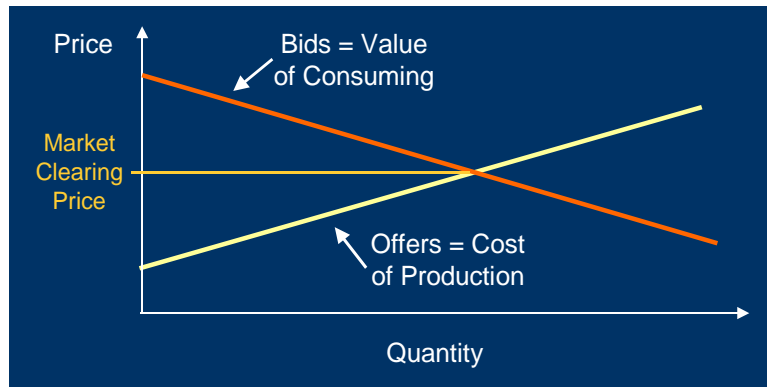
The computer software at the heart of the pricing and scheduling process is called the Dispatch Scheduling Optimizer (DSO) algorithm. In calculating prices, the DSO algorithm assumes that:

- Bids to purchase electricity represent the value consumers place on electricity, and
- Offers to supply electricity represent the cost to suppliers of producing electricity

The IESO receives bids from dispatchable loads indicating the maximum price they are willing to pay for electricity. It also receives offers to supply electricity which indicate the lowest price a supplier is willing to accept for their production. These bids and offers are then stacked economically into bid and offer curves. This can be illustrated simply as follows:



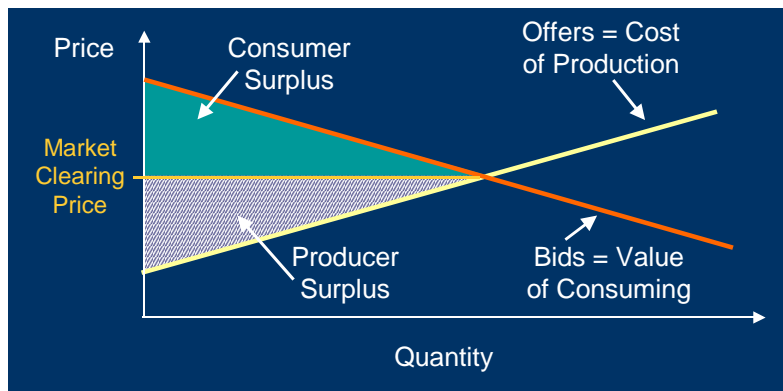
As price increases, the quantity of electricity that consumers are willing to consume decreases because fewer and fewer consumers have valued electricity at that level. As price decreases, the quantity of electricity that suppliers are willing to produce decreases. The optimal price is one where the willingness of consumers to consume equals the willingness of suppliers to supply. This point is attained where the two lines intersect, as in the chart below:



Setting the price here maximizes the economic gain from trade by allowing for the highest possible consumer surplus and the highest possible producer surplus:

- Consumer surplus is the difference between the value of electricity to the consumer (as reflected in their bids) and what they must pay for electricity.
- Producer surplus is the difference between the cost of producing electricity (as reflected in their offers) and the revenue received for the electricity.

Consumer and producer surplus can be illustrated as follows:



The shaded area below the market clearing price represents producer surplus while the shaded area above the market clearing price represents consumer surplus. By setting price at the intersection of the two curves, the entire area between the bid and offer curves is available to the market as a surplus. As a result, the economic gain from trade for the market as a whole has been maximized.

The point of intersection between the bid and offer curves can only be directly calculated where the price tolerance of all participants is known. On the producer side, this is not an issue as all producers provide the IESO with their price tolerance.³ However, the great majority of consumers

³ Non-dispatchable generation (i.e., self-scheduling and intermittent generators) provide a price at which they are likely to reduce their output to zero MW.

within the IESO-administered markets do not submit bids because they are non-dispatchable.⁴ How then can we ensure the maximization of the economic gain from trade for the market as a whole when we don't know the value consumers place on energy?

The DSO maximizes the economic gain from trade by finding the combination of resources that satisfies the demand of non-dispatchable consumers across both the energy and OR markets at the lowest possible cost of production based on the offers submitted. The DSO can use demand to represent the value non-dispatchable consumers place on electricity because their consumption by definition is insensitive to price – they will consume no matter what the price. Therefore, by finding the least cost solution that satisfies non-dispatchable demand, the DSO is maximizing the economic gain from trade for these participants. The DSO also evaluates the bids received from dispatchable consumers and determines any additional resources required to satisfy all dispatchable load willing to pay. Based on the offers made by the resources thus selected by the DSO to satisfy all demand across the energy and OR markets, market clearing prices are determined using a system marginal cost-based pricing method.

System Marginal Cost Pricing

Prices in the IESO real-time energy and OR markets are determined by the Dispatch Scheduling Optimizer (DSO) linear program using a system marginal cost method. This means that prices are most often based on the incremental cost to the market of the product (exceptions include situations of shortfall). The cost to the market of the last megawatt (MW) of demand served is the same as the additional cost to the market as a whole of satisfying an additional increment of demand, except at the point of transition from one offer price to another. In the following examples, offer price transition points are avoided and the additional cost to the market as a whole of satisfying one more MW of demand is used as the marginal price in order to present the principles involved as clearly as possible.

Joint Optimization

Joint optimization allows the DSO to trade-off resources between the energy and OR markets in order to find the schedule that meets the required demand while minimizing the cost of production and, thereby, maximizing the economic gain from trade. We can illustrate this with the following example. Assume:

- Dispatchable energy demand of 200 MW bid at \$1,000/MW
- The requirement for OR is 100 MW
- There is only one class of OR

⁴ Non-dispatchable consumers withdraw electricity without direction from the IESO and pay the wholesale market price.

- There are three generators offering both energy and operating reserve.⁵ Their offers are as follows:
 - Gen 1 offered 200 MW of energy at \$25/MW and 100 MW of OR at \$1/MW
 - Gen 2 offered 200 MW of energy at \$26/MW and 100 MW of OR at \$4/MW
 - Gen 3 offered 100 MW of energy at \$40/MW and 100 MW of OR at \$5/MW

Let's contrast the outcome of joint optimization in this example with that of a sequential optimization solution.⁶ If the DSO performed sequential optimization, it would minimize cost in the energy market first and then use any remaining resources to satisfy demand in the OR market:

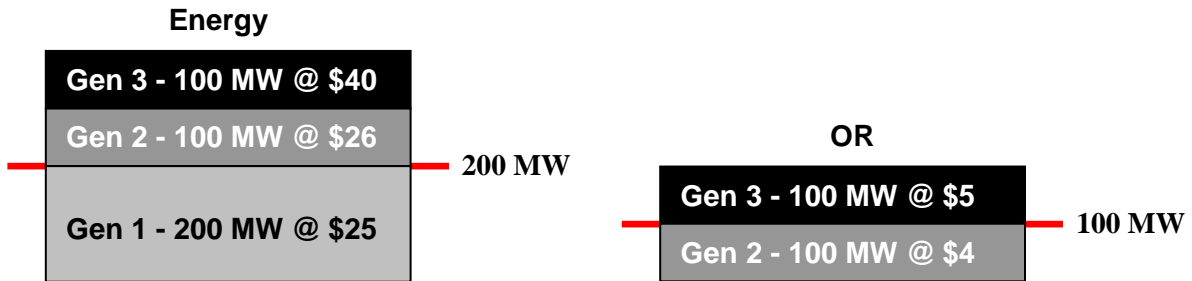


Figure 1 - Sequential Optimization

In this solution, all of Gen 1's 200 MW offer at \$25 was accepted in the energy market as it had the lowest price. This means that none of its 100 MW offer in the OR market can be used since all that Gen 1 can supply at any one time, based on its offer, is 200 MW. As a result, we cannot use any of Gen 1's \$1/MW OR offer. Instead, we have to use Gen 2's more expensive \$4/MW OR offer. Given this, what is the total cost of production to the market as a whole of this solution?

$$\begin{array}{ccc}
 \text{Energy} & & \text{OR} \\
 \underbrace{\hspace{1.5cm}} & & \underbrace{\hspace{1.5cm}} \\
 = 200 \text{ MW} \times \$25 + 100 \text{ MW} \times \$4 = \mathbf{\$5,400} \\
 \underbrace{\hspace{1.5cm}} & & \underbrace{\hspace{1.5cm}} \\
 \text{Gen 1} & & \text{Gen 2}
 \end{array}$$

Next, let's see what the total cost of production would be if we used a joint optimization solution. Joint optimization allows the DSO to examine all of the possible ways of satisfying demand in order to find the solution with the lowest possible production cost to the market. By doing so, it maximizes the gain from trade for the market as a whole:

⁵ In this Quick Take, it is assumed that all offers are for the maximum capacity of the unit.

⁶ To ease understanding, supply offers set the price in all examples in this Quick Take. However, dispatchable loads can also set the price.




Figure 2 - Joint Optimization Solution

The total cost of production in this solution is:

$$\begin{array}{c}
 \text{Energy} \qquad \qquad \qquad \text{OR} \\
 \underbrace{\hspace{10em}} \qquad \underbrace{\hspace{5em}} \\
 = 100 \text{ MW} \times \$25 + 100 \text{ MW} \times \$26 + 100 \text{ MW} \times \$1 = \mathbf{\$5,200} \\
 \underbrace{\hspace{3em}} \qquad \underbrace{\hspace{3em}} \qquad \underbrace{\hspace{3em}} \\
 \text{Gen 1} \qquad \qquad \text{Gen 2} \qquad \qquad \text{Gen 1}
 \end{array}$$

This is lower than the \$5,400 cost of production if the resources had been sequentially optimized because it takes advantage of the relative costs across both the energy and OR markets. Therefore, this solution results in a greater economic gain from trade for the market at a whole.

Now that the resources have been scheduled, how does the DSO set market clearing prices?

Joint Optimization and Market Clearing Prices

As discussed above, prices are based on the total cost of satisfying the **next** MW of demand. This process is fairly straightforward when only one market has to be considered. However, the DSO uses joint optimization whereby resources are traded-off between markets to find the optimal combination of resources that satisfies demand at the lowest possible cost. Therefore, we have to consider what the effect of satisfying one additional MW of demand in one market is on the other markets. This is so because by satisfying one more MW of demand in one market, we are making that MW unavailable to satisfy demand in the other markets. For example, if we schedule one additional MW in the energy market, we make that MW unavailable for use in the OR markets.

For example, assume:

- Non-dispatchable energy demand is 18 MW
- The requirement for operating reserve is 5 MW
- There is only one class of operating reserve
- There are three generators in the market with the following offers:
 - Gen 1 offered 10 MW of energy at \$5/MW and 10 MW of OR at \$5/MW
 - Gen 2 offered 10 MW of energy at \$8/MW and 10 MW of OR at \$1/MW
 - Gen 3 offered 10 MW of energy at \$10/MW and 10 MW of OR at \$4/MW

The optimal solution is shown in figure 3:



Figure 3 - Joint Optimization Solution

In this solution, Generator 2 is used to provide energy and OR rather than just energy, based on the comparative offer prices for OR and energy. In other words, the total cost for the two combined markets is reduced by taking advantage of Generator 2’s relatively low OR offer price.

In this economic solution, the total cost of production is:

$$\underbrace{10 \text{ MW} \times \$5}_{\text{Gen 1}} + \underbrace{5 \text{ MW} \times \$8}_{\text{Gen 2}} + \underbrace{3 \text{ MW} \times \$10}_{\text{Gen 3}} + \underbrace{5 \text{ MW} \times \$1}_{\text{Gen 2}} = \$125$$

Now let’s consider what the market clearing prices for the energy and OR markets are in this example. Remember that prices are set at the total cost to the markets of satisfying one more MW of demand. Using joint optimization, prices are determined based on the total cost across both the energy and OR markets of an incremental MW of demand.

To satisfy one additional MW of energy market demand most cost effectively, we would increase the schedule of Gen 3 by one MW. Doing so would mean that the cost of satisfying 19 MW of energy demand instead of 18 MW would be:

$$\underbrace{10 \text{ MW} \times \$5}_{\text{Gen 1}} + \underbrace{5 \text{ MW} \times \$8}_{\text{Gen 2}} + \underbrace{4 \text{ MW} \times \$10}_{\text{Gen 3}} + \underbrace{5 \text{ MW} \times \$1}_{\text{Gen 2}} = \$135$$

This is \$10 higher than the total cost of providing 18 MW of energy (which was \$125), so the incremental cost of an additional MW of energy demand is \$10. The energy market clearing price is therefore \$10.

One might think that the most cost effective solution would have been to increase the schedule of Gen 2 by one MW because its energy offer price was \$8 versus Gen 3’s \$10 offer. This is not the case because the effect on the OR market must also be considered. Using Gen 2 to satisfy the additional energy market MW would require the DSO to reduce the MWs available from Gen 2 in the OR market by one MW. In order to continue to meet the OR market demand of 5 MWs, the DSO would have to take another MW from Gen 3. Gen 3 had offered OR at \$4/MW, which is

substantially higher than Gen 2's OR offer of \$1/MW. In this scenario, the cost to the market would be:

$$\begin{array}{c}
 \text{Energy} \qquad \qquad \qquad \text{OR} \\
 \underbrace{10 \text{ MW} \times \$5 + 6 \text{ MW} \times \$8 + 3 \text{ MW} \times \$10}_{\text{Gen 1} \quad \text{Gen 2} \quad \text{Gen 3}} + \underbrace{4 \text{ MW} \times \$1 + 1 \text{ MW} \times \$4}_{\text{Gen 2} \quad \text{Gen 3}} = \$136
 \end{array}$$

This is \$1 more expensive in total than using Gen 3 to satisfy the additional MW of energy demand as calculated above. Therefore, using Gen 2 to satisfy an additional MW of demand in the energy market would actually be a higher cost solution.

What is the price for OR in this example? The solution to this question illustrates one of the effects of using joint optimization - the market clearing price can actually be set at a price that does not correspond to any single offer received. Remember that the market clearing price is set based on the incremental **cost** to provide the product. Because we have to consider the cost across markets, the incremental cost may or may not be the same as any offer received. Let's work through it:

If the OR requirement in the above example went up by 1 MW, that MW would most economically come from Gen 2 who had made an offer of \$1/MW. However, based on its offers, Gen 2 can only produce 10 MW of total output. Increasing its schedule in the OR market by 1 MW would require reducing its schedule in the energy market by the same 1 MW. In order to meet demand in the energy market, one additional MW of Gen 3 output would be required. Gen 3's energy offer, however, was more expensive than Gen 2's (\$10 versus \$8). Therefore, the true cost to the market as a whole of satisfying one more MW of OR is not \$1 (the cost of Gen 2's OR offer), but rather something that reflects the fact that more expensive supply is also now required in the energy market. In this example, the total cost to the market of satisfying one additional MW of OR demand would be:

$$\begin{array}{c}
 \text{Energy} \qquad \qquad \qquad \text{OR} \\
 \underbrace{10 \text{ MW} \times \$5 + 4 \text{ MW} \times \$8 + 4 \text{ MW} \times \$10}_{\text{Gen 1} \quad \text{Gen 2} \quad \text{Gen 3}} + \underbrace{6 \text{ MW} \times \$1}_{\text{Gen 2}} = \$128
 \end{array}$$

This is \$3 higher than the optimal solution for 5 MW of OR of \$125 shown on page 6. Therefore, the OR market clearing price in this example is \$3. This is not a price offered by any one of the suppliers. However, it is the true incremental cost to the market as a whole of serving one more MW of OR demand.

Summary - Impact of Joint Optimization on Price and Schedules

The joint optimization process leads to the scheduling of resources across the energy and operating reserve markets based on an “optimal solution”. The optimal solution is the one that maximizes the economic gain from trade across all markets. Maximizing the economic gain from trade results in scheduling the offered resources in such a way that demand can be satisfied at the lowest possible cost of production. The result of the process of trading off resources between the energy and OR markets can be market clearing prices that do not correspond to any single offer submitted. In the example above, the incremental impact on the energy market led to a clearing price of \$3/MW for OR, even though no participant made a \$3 OR offer.

This process of trading off resources can also result in market participant schedules that may not appear intuitive at first. In the above example, Gen 2 had offered energy into the market at \$8. The clearing price for energy was \$10 and yet Gen 2 was not scheduled for its full energy offer. Because of their relatively low OR offer price of \$1, Gen 2 was scheduled for 5 MW of OR. This scheduling difference had no negative impact on Gen 2, however, as their operating profit in both markets was the same (\$10 Energy MCP - \$8 Energy offer price or \$3 OR MCP - \$1 OR offer price).⁷ This is another aspect of the joint optimization function: It will not schedule participants for the advantage of the market if doing so negatively impacts the participant themselves.⁸

Additional Resources

Market Rules Chapter 7, sections 4 and 8 and Appendix 7.5 available on the [Rules and Manuals](#) page

[Introduction to Ontario’s Physical Markets](#)

For additional information, please contact the IESO at:

Toll Free: 1-888-448-7777

Tel: 905-403-6900

Fax: 905-403-6921

customer.relations@ieso.ca

⁷ As both the constrained and unconstrained modes of the algorithm use joint optimization, no congestion management settlement credits (CMSC) would be payable. Please see the [Introduction to Ontario’s Physical Markets](#) for additional information on CMSC.

⁸ System constraints or other factors may lead to scheduling between the markets that may lead to reduced operating profit.