Exploration of Enhancements to Dispatch Methodology and Processes (SE-61)

Solution Identification Paper

Public

Introduction

Dispatchable facilities have raised concerns about dispatch instructions (e.g. the number of dispatch instructions and dispatch reversals) and their impact on plant operational requirements and operating efficiencies. In a previous paper[1], the IESO identified four problems contributing to generator dispatch issues. Those problems were categorized under four headings: Modelling, Cost Accounting, Forecast Error and Fixed One Hour Bid Window.

The purpose of this paper is to present a description of all plausible solutions that address at least one of the identified problems. While a full range of possible solutions should be presented, those that are not feasible due to implementation costs (including the time required to implement), market structure, inter-jurisdictional integration complexities, or government policy constraints should be discounted or ruled out when appropriate.

The solution identification section of this paper is organized according to the four categorized problem headings. Under each heading the problem statement and source of inefficiency are restated. Following this a series of solutions are presented. A description of each solution is provided as well as an initial assessment of feasibility. For solutions not ruled out and expected to involve considerable change to the market design, subsequent reports will be required, gradually providing more design detail, market amendments as well as proposed changes to IT/settlement systems.

For those solutions of highest priority, and where more than one possible solution to a particular problem exist and where further quantification of the costs and benefits of each solution is required, the IESO will pass them to the next stage of the project life cycle for a more formal CBA to be conducted.

The IESO has issued this report of plausible solutions to summarize much of the discussion around dispatch issues. Despite the temporary suspension of new SE-61 activity, stakeholders are welcome to submit feedback, including their view on the relative priority of the various solutions. Based on this report and stakeholder feedback and with a better appreciation of the magnitude of change and the time involved to implement the solutions, the IESO will review the prioritization of the solutions relative to other issues in 2011.

Solution Identification

Problem 1: Modelling

The optimal dispatch instruction should take into account the physical limitations of a dispatchable unit, but each instruction only accounts for the operating characteristics that are modelled in the DSO.

A dispatch instruction that can't be met by a dispatchable unit may be an inefficient market outcome if the total cost of the dispatch instruction and accompanying compensating action, such as AGC, operating reserve activation (ORA) or an out of market action, is greater than the cost of an alternative dispatch that could be met.

[1] The paper can be found on the IESO's website at: http://www.ieso.ca/imoweb/consult/consult_se61.asp.

Solution: Inertia modelling

In discussing dispatch issues, generators specifically raised the lack of modelling in the DSO as a likely contributor to undesirable dispatch instructions. Inertia Modelling is a new technique that introduces additional constraints on a unit's ramp capability into the DSO. Specifically, inertia modelling places constraints on a unit's ramp rate as well as on a unit's rate of change in ramp rate. Conceptually, inertia modelling can be understood with an effective analogy relating a unit's operating characteristics to kinematic properties. Consider the MW output of a generating unit to be equivalent to position, the ramp rate of a unit equivalent to speed and the rate change in ramp rate equivalent to acceleration. Using a relationship between position, speed and acceleration, a unit's MW output at the end of the next five minute interval can be computed based on its current MW output, its available ramp rate and the rate of change in ramp rate. It is also possible to linearize the relationship for inclusion in the optimization process.

At present, the DSO uses a simplified model for ramp rate and an implicit model for the rate of change in ramp rate. The DSO accounts for ramp limitations by allowing participants to submit ramp rates at five different breakpoints. The piecewise laminations are an approximation for a continuous ramp curve. The DSO also uses an approximation when modelling acceleration. This is most noticeable when a unit is dispatched up and then immediately down, using the max ramp up and down rates. In this scenario, a maximum acceleration is implicitly assumed based on the difference between the maximum ramp up rate and the maximum ramp down rate. For example, when a unit receives a dispatch instruction that makes use of the full ramp up capability, say 8MW/min, followed by a dispatch instruction that makes use of the full ramp down capability, say 5MW/min, the DSO is implicitly assuming an acceleration capability of 13MW/min[2]. It may be that the supplied ramp rates are achievable when a unit is dispatched in a constant direction, but they may not be applicable during a reversal. If the simplifications and assumptions used by the DSO do not accurately reflect a unit's operating characteristics, a set of dispatch instructions may be inefficient.

Inertia modelling would add greater complexity to the DSO but also greater accuracy in the way dispatch instructions reflect unit capabilities. To account for acceleration limitations a new set of equations would be developed and incorporated into the linear programming model. During the optimization process, the DSO would ensure a set of dispatch instructions do not violate unit acceleration limits. If limits were placed on acceleration, a unit would likely receive a more gradual set of dispatches to achieve a reversal.

A possible implication of incorporating an explicit constraint on acceleration is the apparent removal of ramp depth[2] and dispatch flexibility from the market. However, if the constraint reflected actual capabilities then no ramp depth or flexibility is lost. If units were previously unable to follow their dispatch instructions, then the presumed ramp depth without any consideration for acceleration was fictitious. The inclusion of dispatch limiting operating parameters will allow the DSO to formulate dispatch instructions that will more accurately reflect operating capabilities. The result would be the addition of dispatch accuracy, and not the removal of dispatch flexibility.

Based on a proof of concept model, IESO research has demonstrated the viability of Inertia Modelling. However, to implement Inertia Modelling, the IESO estimates a project scope similar to that of the Multi Interval Optimization (MIO) design. To deliver the MIO scope as originally planned, 60 weeks of labour and $2.5M[3] was committed to the project. Actual implementation was extended by several weeks due to the

---

[2] Limits on unit acceleration will affect the available ramp depth in the market, as a unit's available ramp rate next interval is a function of the unit's current ramp rate and acceleration capability.

[3] $2.5M in 2004 dollars.

additional time required in stakeholdering user requirements. For what would be a significant undertaking, clear evidence would need to demonstrate that the current assumptions and simplifications used by the DSO are insufficient before additional model complexity could be justified.

Solution: Period of steady operations

A less refined technique of modelling generator dispatch capabilities is requiring a period of steady operation before reversing a generator's dispatch trajectory. During the MIO implementation, a functionality called period of steady operation was designed but never activated. The design was intended for slow moving generating facilities to ensure that units would not reverse direction without a minimum period of steady operation. After the minimum period of steady operation, the unit is available to be normally dispatched. The steady state period is defined in terms of a number of 5 minute intervals, with the minimum number of intervals as 0 and the maximum number of intervals as 2.

The proposed design considered a thermal unit in steady operation if the magnitude of the schedule change between the last two intervals was less than an adjustable variable times the unit offer ramp rate. If the magnitude of the schedule change was higher, the unit was considered ramping up or down. The minimum and maximum output of the unit was set based on its ramp status as follows[4]:

- If ramping up, the minimum output will be set equal to the last dispatch and the maximum output will be based on its full offer ramp up capability.
- If ramping down, the maximum output will be set equal to the last dispatch and the minimum output will be based on its full ramp down capability.

The period of steady operation is an approximation of a unit's ability to reverse its dispatch trajectory, i.e. move from ramp up to ramp down. The lack of consideration for unit specific dispatch characteristics may still result in dispatch instructions that cannot be followed or it may result in the removal of ramp depth and dispatch flexibility from the market. For example, a generator whose unit capability falls in between available options will need to choose between dispatch scenarios where it cannot follow dispatch instructions or where dispatch instructions do not make full use of its capabilities.

At the time of the MIO project, the period of steady operation was designed[5], developed, tested and implemented on the Market Information System (MIS) but never utilized in production. Given the existing design, incorporating a period of steady operation into the DSO would be relatively straightforward. However, before implementation could take place a thorough analysis of the impact of the change on market participant dispatches would be needed. When the period of steady operation was first discussed during the MIO project, significant concern was expressed about the ripple effect on dispatch instructions. It was theorized that if the marginal unit was unable to move because of a constraint, then a sub marginal unit or units would have to move. The move of the sub marginal units, could then impact their availability to move in subsequent intervals. The end result would be a number of resources off their economic base points and an overall increase in the total number of dispatches. For the period of steady operations to be considered a viable alternative, the limited benefits would have to outweigh potential negative outcomes.

---

[4] http://www.ieso.ca/imoweb/pubs/consult/mep/MIO_WG_2003Aug29_FuncRequirements.pdf

[5] The period of steady operation was written into Chapter 7 of the market rules as a required input for the dispatch scheduling and pricing process.

Solution: Load smoothing

Unit operating characteristics can also be considered in the development of the load forecast profile. Instead of including additional model complexity into the DSO, unit operating characteristics can be implicitly accounted for in the 10 minute ahead load forecast. In the current market design, marginal resources are not expected to follow the ever changing variations in real time market demand. Instead, units with specialized automatic generation control (AGC) equipment shadow demand from second to second, and marginal resources follow a slower moving curve. The slower moving demand curve is a formulation of the Load Predictor (LP) forecast tool. The LP provides the load forecast for each 5-minute interval in the MIO horizon. It is based on a historical load pattern, load telemetry and weather conditions. Instead of formulating a best estimate for actual load, the LP uses a smoothed forecast. The smoothed forecast is less accurate but also less volatile. An argument could be made that depending on the marginal resource, the degree of smoothing is insufficient.

The LP smoothes the demand forecast by using a line of best fit technique. To determine the 10 minute ahead dispatch forecast, the LP linearizes the set of three past actual values and three most recent predictions for future intervals; it then interpolates to determine an estimate of the current actual load. From the current estimate of load a future load estimate can be calculated. To increase the degree of smoothing, additional intervals can be included in the line of best fit technique, i.e. using actual load from four previous intervals. The consequence of dispatching to a smoothed forecast is a loss of accuracy that must be compensated for with AGC. As the degree of smoothing increases, so too does the forecast error. If the forecast error increases so too must the amount of AGC.

The implementation of a smoother forecast would be relatively straightforward. The DSO design contains a parameter that allows the degree of smoothing to be manually adjusted with relative ease. Although implementation is straightforward, significant analysis on the loss of forecast accuracy would be necessary prior to any change. The impact on market efficiency would also need to be studied, in particular, any additional requirements or costs in AGC. Generally, the IESO considers the addition of less transparent costs, such as AGC, to be a negative outcome for the province.

The relationship between the different resource types and the degree of smoothing would also need to be understood. Smoothing would be most directly felt by the marginal unit, which is not a constant, yet the degree of smoothing is. Forecast smoothing is a broad stroke solution that does not consider generator specific characteristics and at best could achieve a degree of smoothing that, on average, could be met by the marginal unit.

Solution: Load decomposition

Load breakdown is another method that can be used to smooth the load forecast. The IESO dispatches generation to meet the 10 minute ahead load forecast, so as the forecast fluctuates, so too will unit dispatch instructions. Load breakdown is a technique used to create a number of different forecasts that are smoother and easier to follow. The process entails decomposing the load forecast into its various signals using spectral analysis. Spectral analysis is a mathematical technique that is used to extract smooth sinusoidal curves from complex oscillating curves. The extracted curves have known amplitudes and frequencies. When the extracted curves are added together, they reform the original signal.

With this technique, the total forecast is divided into smaller segments, some that change faster than others. For example, consider the decomposition of the load forecast into two sinusoid waves and a high frequency noise wave. The first forecast, or wave, might account for 60% (on average) of total load and have a very gradual change from one interval to the next. The second forecast might account for 35% (on average) of total load and change more quickly than the previous wave, but still in a predictable way. A third forecast might account for 5% (on average) of total demand and change very rapidly in a less predictable way. For any given interval, the sum of all three forecasts would total the expected load.

The decomposition process described above would generate a set of sinusoid forecasts that each has a predictable frequency and amplitude. Generators could use the known properties of the different forecasts to self select which forecasts its units are capable of following. Presumably, the rate of change of the different forecasts and the operating characteristics of a unit, would dictate which forecast they select to follow. For example, slower moving nuclear facilities would respond to a lower frequency forecast and faster moving hydro facilities would respond to a higher frequency forecast. This method implicitly accounts for the operating capabilities of different generating units by allowing generators to self select the optimal dispatch trajectory their units should follow.

As another approximate solution, load breakdown does not address generator specific dispatch capabilities. The decomposed wave forms may or may not reflect the operating capabilities of generating units. This method creates forecasts for an average unit, rather than customizing dispatch instructions to each unit's abilities. For example, a unit may find its self underutilized in following a particular forecast, yet the next available forecast may be beyond its capabilities. This solution is likely to improve the dispatch for some units, but not all.

A second challenge is to ensure generators are properly incented to follow the appropriate forecast. To meet total load sufficient resources would need to be available to respond to each of the different forecasts. Proper incentive could be achieved if each forecast represented a different energy product, with separate offers and clearing prices. Such a proposal is unlikely to achieve consensus among market participants. Due to the complexities of designing multiple energy products and the consensus needed from market participants, the IESO does not consider this solution feasible.

Solution: Dispatch decomposition

Another means of avoiding additional complexity in the DSO is to implicitly account for unit operating characteristics in the formulation of dispatch instructions. A set of dispatch instructions could be designed to cater to the different operating needs of different generators. This could be achieved by creating a new energy product. Presently, the IESO uses energy from three sources to meet demand: AGC, real-time dispatch and intertie resources. AGC resources are dispatched on a second to second basis, real-time dispatch is done on a five minute basis and imports are dispatched hourly. One could argue that the Ontario market lacks an energy product for resources that operate optimally in the range between real-time dispatch and imports. Such a product has often been referred to as a 15 minute dispatch.

Despite the frequent reference to a 15 minute dispatch, very few market features have been published or even discussed. Multiple designs may exist, each with different market features. For illustrative purposes, one such design is described here. To implement a 15 minute dispatch the IESO would have to ensure sufficient dispatch flexibility remained to meet demand. Such a requirement would necessitate a portion of the generation fleet to continue to operate on a 5 minute basis. The required capacity could be obtained through a monthly auction. These units would receive the auction clearing price as a capacity payment for making their units available for dispatch on a 5 minute basis. To participate in the auction, these units would have a requirement to make available a minimum dispatch range and ramp rate. If a unit was not successful in the auction or did not participate in the auction it would by default be designated a 15 minute resource.

The successful participants from the auction would then be dispatched on a 5 minute basis, as today. The remainder of resources would be dispatched on a 15 minute basis. This would require a second MIO to run in parallel with the existing MIO, but on a 15 minute basis. The Ontario market would now have two constrained runs. The 15 minute MIO would dispatch units based on a 15 minute trajectory, while the 5 minute MIO would dispatch units on a 5 minute trajectory. When the 15 minute and 5 minute MIO runs overlap, resources would be dispatched concurrently. For the two intervals of the 5 minute dispatch that do not overlap with the 15 minute dispatch, the DSO would assume a fixed output from the 15 minute units (interpolating between the two most recent 15 minute dispatches) and dispatch the 5 minute movers to meet the variation in demand. The 5 minute resources would have to be closely managed so that they are not dispatched to their full capability, but have available capacity to manoeuvre up and down. The process described above would affect only the real-time constrained sequence. The real-time unconstrained (i.e. HOEP) would not be affected.

Prior to implementing the scenario described above the IESO would have to resolve a number of issues. The IESO would want to understand the implications of altering the constrained sequence and not the unconstrained, and in particular, the effects on CMSC. Additionally, a detailed study would be needed to understand the capacity and ramp needs of the 5 minute dispatch product. Apart from the necessary pre-design work, a project scope for the implementation of a second MIO sequence can reasonably assumed to be equal to or less than one third of the original MIO project. A 15 minute dispatch would be a significant undertaking, requiring substantial IESO and stakeholder resources.

Solution: Pseudo units in real-time

As new generation technologies enter the market, new modelling techniques will be required. With the substantial arrival of combined-cycle generating plants (CCP) new modelling techniques are required to accurately reflect the various configurations of combustion turbines (CT) and steam turbines (ST). Without considering the multiple operating states of a CCP, the DSO can produce unrealistic schedules and reduce the overall efficiency of day-ahead commitments and real-time dispatches.

As part of the Enhanced Day-Ahead Commitment (EDAC) project, the IESO is implementing the Pseudo-Unit model (PSU)[6] to provide increased functionality in the DSO. The PSU model will represent a CCP with one or more CT and a single ST as one or more pseudo units, each comprised of a single CT and its associated portion of the ST capacity. The changes will only affect the day-ahead commitment process. Generators are still required to submit offers for their physical units in real-time. To fully reflect the dispatch capabilities of a CCP the relationship between CT and ST must also be modeled in real-time.

A real-time model would apply the same logic used in EDAC, where a pseudo unit is comprised of a single CT and its associated portion of the ST capacity. The model parameters, offers and settlement would be consistent with the EDAC design, but incorporated within the real-time constrained algorithm. In real-time, the PSU would be dispatched separately by the DSO (and separate from the physical units) resulting in dispatches that realistically represent the capabilities of the units. Generators would no longer be required to submit offers for their physical units, as their pseudo unit offers submitted day-ahead could be used for dispatch.

Although the PSU model has been designed for the day-ahead timeframe, it will still require a significant effort to implement into the IESO's real-time operations. In addition to the design work, pseudo-units in

[6] For a complete description of the Pseudo-Unit model see: http://www.ieso.ca/imoweb/pubs/consult/se21-edac/se21-edac-20081128-Pseudo_Units.pdf.

real-time will require amending the Market Rules. Considering the IESO's experience with pseudo-units in the day-ahead timeframe, it is realistic and feasible to pursue implementation in real-time.

Problem 2: Cost Accounting

A set of dispatch instructions should reflect the lowest cost solution to meet demand, but a solution is not the lowest cost because all relevant costs are not accounted for during the optimization process.

If all relevant costs are not considered in issuing dispatch instructions then the DSO may dispatch in a more costly manner than otherwise needed.

Solution: V cost curve

Adding a cost for ramp into the DSO is one way to account for the avoidable costs of dispatching a generator. In a competitive market, a unit's offer reflects their marginal cost, or the additional costs required to produce an additional MW. A unit's offer includes all marginal costs for a specified level of output, namely the incremental energy cost, variable O&M and any opportunity cost of not running. The costs included in an offer are incurred when a generator responds to a dispatch, and are known in advance of the dispatch. Other costs, that are also marginal, are more difficult to include in an offer because they are particular to a type of dispatch or they are incurred over the longer term. Costs that are typically not included in offers, or are difficult to include, are fuel inefficiency costs, certain wear and tear costs and environmental costs. These costs are related to the dispatch instructions a unit receives, but can be difficult to predict and for that reason would have to be averaged if they were included in an offer price. For example, certain wear and tear costs are immediate costs, but are only determined after a dispatch instruction has been received. A unit may incur additional wear and tear cost if it receives a dispatch reversal but not if it receives a consistent set of dispatch instructions in the same direction. These costs are avoidable, depending on how the unit is dispatched. For these costs to be considered avoidable by the DSO they must be explicit from a unit's offer price, and accounted for during the optimization process.

One method of explicitly recognizing the cost of ramp, or manoeuvring, is known as a V cost curve. This method considers a per MW cost for moving a unit away from its operating point. Because the cost is a per MW cost, the cost associated with moving a unit up or down increases linearly from the operating point, resulting in a V shaped cost curve. When formulating dispatch instructions, the DSO would consider the manoeuvring cost in addition to a unit's offer.

In practice, the extra cost of movement is achieved by adding a per MW cost to the offer curve above the operating point and subtracting a per MW cost from the offer curve below the operating point. From the perspective of the DSO, this has the effect of making an additional MW appear more costly and the savings from a reduction of a MW appear less. It is anticipated that the DSO would then only dispatch units when the market savings outweigh the manoeuvring costs. In contrast to today, where no cost to moving units is assumed, and thus any lower cost combination of offers will always be dispatched, regardless of the magnitude of total market savings.

The ability of the V cost curve to reduce the number of dispatches has its limits. Many dispatches that exist today would continue to persist. Except for instances where the DSO dispatches to optimize for changes in system constraints, an explicit ramp cost is unlikely to reduce the number of dispatches but rather shift the dispatch instruction to the new lowest cost generator, i.e. the generator with the lowest total offer and ramp cost. Dispatches for changes in imports/exports, unplanned outages, changes in bids/offers, ramp constraints and changes in demand will still be needed.

The modified offer curves are applicable to the constrained schedule only. The V cost curve would ensure the correct, or lowest cost, generator is dispatched but would not impact the HOEP. Instead, the disparity between the constrained and unconstrained schedule is likely to impact congestion management settlement credits (CMSC). Careful study would be needed to understand the impact on price transparency, especially if HOEP were left unchanged and CMSC grew. As well as the necessary research, the addition of a V cost curve to the constrained schedule would require a re-coding of the DSO, albeit a relatively modest one. While the change is minimal, the analysis required to estimate the impacts of the change may be complex.

Solution: Extended inertia modelling

Extended inertia modelling is another technique that applies explicit costs to the manoeuvring of a unit. Like the inertia modelling technique described previously, constraints are placed on unit ramp and acceleration, but in addition, generator specific costs are applied to each of these movements. This methodology recognizes that a unit's costs are not limited to, or easily incorporated into, its offer price and that moving a unit is more costly than not. For example, if a unit that offers $20 for 100MW has a ramp cost of $4/MW/min and an acceleration cost of $1/MW/min^2 and is dispatched from a flat 50MW to 55MW, a total cost of $105 will be incurred. The total cost for the additional 5MW is based on three components: $100 from the offer price, $4 from the 1MW/min ramp, and 1$ from the 1MW/min^2 acceleration. With the additional information the DSO will not only minimize the total cost of production based on offers but also based on the cost of moving units. The more detailed information will allow the DSO to optimize the total amount of unit manoeuvring over all resources. This could result in a unit with a $104 offer and no movement costs being dispatched in advance of the unit described above.

Extended inertia modelling would build upon the design elements of inertia modelling. Given the natural extension, it is expected that extended inertia modelling would add cost, complexity and time to the previously described scope of inertia modelling. Perhaps of greatest concern to the IESO are the implications for market pricing and settlement. Extended modelling could be present in the constrained algorithm, like the V cost curve, or it could be present in both the constrained and unconstrained. In either case, there would be profound impacts for either the hourly Ontario energy price, congestion management settlement credits or uplift charges. These are but a few of the many implications that would have to be considered. The implementation of extended inertial modelling would represent a fundamental change in Ontario's electricity market, the likes of which has not been seen by neighbouring jurisdictions. For such a significant undertaking stakeholders and the IESO would need to believe that the costs and complexities of implementation would be outweighed by any improvement in dispatch efficiency.

Problem 3: Forecast Error

A set of dispatch instructions should reflect the lowest cost solution to meet demand, but each solution is less than optimal because forecast error is not minimized.

If the DSO dispatch includes a forecast error, then an inefficient mix of resources was dispatched, which resulted in a higher cost than needed.

Solution: Load predictor update

Any reduction in load forecast error will result in a more efficient dispatch. The Load Predictor, used by the IESO to predict intra hour load, has a built in bias that can result in forecast errors. When a bias is introduced into a forecast used for dispatch, the DSO will schedule an inappropriate level of generation. An over forecast will result in more generation scheduled than necessary and an under forecast will result in too little generation scheduled. Reducing forecast error will reduce the size and number of inefficient dispatches.

The LP forecasts total load using the line of best fit technique described previously. The DSO, however, requires a forecast for total non dispatchable load (NDL), not total load. The current methodology used by the DSO to forecast NDL can create instances of dispatch volatility. The DSO estimates NDL based on the forecast for total system demand less a "snapshot" of dispatchable load (DL). As the level of DL fluctuates, so too will the forecast for NDL. The IESO has observed instances of a 100MW or more reduction in DL over a period of 5 minutes. A 100MW reduction in DL can result in an immediate increase of 100MW in forecast NDL, a counter intuitive result. The volatility in the forecast is translated into volatility in the dispatch instructions for generators.

The forecast error within the LP can be reduced by altering the way NDL is forecast. Instead of using the methodology described previously, the LP can forecast NDL directly. Fluctuations in DL would no longer impact the NDL forecast. The result would be a smoother forecast and a smoother set of dispatch instructions.

The proposed change would involve a redesign of the LP tool. To complete the redesign, new variables would need to be developed and several existing variables would have to be redefined. The overall forecast methodology, however, would remain unchanged. The IESO is currently completing a Cost Benefit Analysis on the Load Predictor update. Should the benefits outweigh the costs the IESO will consider the change relative to other priorities.

Solution: Self-scheduling generation forecast update

Forecast errors can also be present in self-scheduling generator forecasts. Although the IESO does not send dispatch instructions to self-scheduling generators on a five minute basis, instructions are still formulated by the DSO to account for the energy production from these units. The current method for dispatching a self-scheduling generator is to assume their current production will persist 10 minutes into the future. In this way, the dispatch is a forecast of future energy production, often referred to as persistence forecasting. However, if a generator's submitted schedule changes between hours, persistence forecasting is substituted for a procedural dispatch. In the event of a schedule change, the unit will be ramped to their new schedule using a fixed pre-determined profile in the real-time constrained schedule. The DSO will perform ramping to prevent a sudden jump during the hour crossing. The cross-hour ramp starts at the end of the hour and is completed in fifteen minutes. After the three intervals, the DSO will return to persistence forecasting.

The ramping method described above works well under the assumption that generators ramp to their new schedules in a way consistent with the current design. However, the IESO has observed that wind generators, a subset of self-scheduling generators, rarely if ever transition between schedules over a three interval period. At the time the ramp feature was developed wind would not have been a consideration due to the low levels of penetration. Now, with 1,085MW of installed capacity, wind forecast errors present reliability concerns. With wind resources, it may be that persistence forecasting is a better predictor of output at all times, even across hours where schedules change. With increased amounts of variable generation expected in a future resource mix, any forecast error associated with an assumed ramp is likely to grow.

The IESO is currently investigating the potential improvement in self-scheduling generator forecast accuracy by moving to persistence forecasting at all times. All self-scheduling generators will be examined to determine if the benefits apply to all other resources types, and not just variable generators. Pending the findings of the analysis, the IESO would conduct a Cost Benefit Analysis to determine if moving to persistence forecasting at all times presents an overall net benefit to the province of Ontario.

Problem 4: Fixed One Hour Bid Window

A set of dispatch instructions should reflect the lowest cost solution to meet demand, but each solution is limited by the fixed one hour bid window that may prevent the optimal dispatch.

If a known mix of resources could theoretically meet demand at a lower cost than the dispatch but is limited by market design features, the dispatch may be inefficient.

Solution: More frequent dispatch of intertie resources

The economic assessment of intertie resources on a more frequent basis than the current hourly process would allow for a lower cost dispatch solution. Ideally, intertie resources would be dispatched on a 5 minute basis, but any dispatch frequency greater than the current hourly dispatch would enhance the efficiency of the IESO administered market.

Intertie resources are currently scheduled in the hour ahead pre-dispatch and are fixed for the entire dispatch hour. This process can lead to the intertie resource being uneconomic for certain intervals in the dispatch hour. On an hourly granularity intertie resources are on average economic, but this does not imply that the resource is economic for every interval during the dispatch hour. If the intertie resource could be scheduled around the uneconomic intervals, a lower cost solution could be achieved.

Avoiding the uneconomic intervals would also reduce the unnecessary dispatch of domestic resources that most commonly occurs across the top of the hour. Consider, for example, an hour in which demand is increasing. In real-time, intertie resources are placed at the bottom of the energy stack and displace the energy of the marginal unit from the interval preceding the dispatch hour. Although the domestic resource might have an offer price that is less than the offer of the intertie resource, the domestic resource will be dispatched down to accommodate the entry of the intertie resource. As demand increases through the dispatch hour the domestic resource will be reloaded. If the intertie resource could be scheduled around intervals in which it is not economic there would be no need to dispatch down the domestic resource, instead, the intertie resource would be loaded when its offer is economic.

Complicating the more frequent assessment of intertie resources is the increased coordination efforts needed between jurisdictions. A reasonable balance between increased efficiency and increased coordination efforts would need to be struck. Also, a more frequent assessment of intertie resources would not resolve the numerous seams issues that exist between the Ontario market and neighbouring jurisdictions. Most notably, the uniquely Ontario unconstrained energy pricing method and the locational marginal pricing method used by many neighbouring jurisdictions. The Regional Electricity Trading: Opportunities and Challenges for Ontario paper discuss this issue among the many others[7].

The more frequent dispatch of intertie resources represents a real opportunity to enhance the efficiency of the Ontario wholesale electricity market. However, the IESO would require the cooperation of its neighbouring system operators to coordinate such an effort. As well, significant development work would be necessary to design a process around the numerous seams issues. Despite the many known challenges it may be worth IESO and stakeholder effort to investigate a high level design.

---

[7] The full report can be found on the IESO website at: http://www.ieso.ca/imoweb/pubs/marketreports/omo/2009/trade.pdf.

Solution: More frequent offer submissions

Allowing generators to submit offers on a more frequent basis than the current hourly requirement may improve dispatch efficiency. Generators would continue to receive dispatch instructions every five minutes but their offers would be valid for time periods shorter than one hour, such as every 30 minutes or every 15 minutes. Ideally, generators would submit offers that reflect the value of their energy for each 5 minute interval, but any bid window less than hourly would be an improvement. For most generators, more frequent offer submissions may have little value, as the marginal cost of energy production is unlikely to vary on a 5 minute basis. However, energy limited resources may find it advantageous to submit their offers on a more frequent basis. It is most efficient for energy limited resources to supply energy during peak periods when energy is most needed and when prices are highest. If a hydro facility has 36 intervals of stored energy, the most efficient use of that energy is during the 36 highest priced intervals in which the resource is infra marginal. Allowing generators to offer at the interval granularity would allow energy limited resources to precisely target the delivery of their energy to the highest priced intervals, rather than the highest priced hours.

More frequent offer submissions could result in a more efficient dispatch if participants receive reliable price signals from the market and respond accordingly. In order for energy limited resources to deploy their energy efficiently, that is, offer their energy during peak periods, they must be able to predict higher priced intervals. However, if resources are allowed to submit their offers too frequently, price volatility may be introduced into the market and distort price signals. As unit offers change, the DSO will re-dispatch units to minimize the total cost of meeting demand. If ramp constraints are present, the market clearing price will be higher than the simple stacking of energy offers would suggest. If the DSO is frequently ramp constrained, prices may become volatile and resources may not be able to extract actionable information from the market price.

Whether or not price volatility would actually be introduced into the market as a result of frequent offer changes is unclear at present, as this would depend on market participant behaviour. It is also unclear that if price volatility was introduced into the market, energy limited resources would be unable to reliably deploy their energy during peak periods, thus negating the potential benefits of more frequent offer submissions. These issues would have to be investigated prior to any design work.

Reducing the fixed one hour bid window would represent a fundamental market design change. The change may smooth the transitions which are currently massed on the hour but it may also increase intra hour ramping. One would have to believe that energy limited resources would offer marginally during peak periods in order to eliminate the ramping down (up) of units to accommodate the sudden arrival (departure) of energy limited resources. A great deal of pre market design work would be necessary to understand the implications of more frequent offer submissions. Clearly identifiable efficiency gains would be needed to warrant what would be a significant market design change.