Imbalance Reserves: Supply, Demand, and Sufficiency

Formerly, "A Metric for Imbalance Reserves Requirements, Supply, and Adequacy"
Monday, June 18, 2012
Michael Schilmoeller, NWPCC

This paper provides a framework for quantifying the requirement for and supply of imbalance reserves. It also presents a means of measuring any insufficiency of supply.

The following is the outline of the paper. The section **Introduction** (page 3) sets out background, examples, objectives, and problems to familiarize the reader with the subject. The introduction gives a preliminary list of definitions and assumptions that will simplify subsequent arguments. At the end of the paper, the consequences of alternative assumptions are discussed.

The paper next characterizes of reserve supplies. The section **The Economics of Reserve Supply** (page 5) briefly demonstrates why the value of services increases with ramp rate. This will be important to thinking about how reserve capacity is restored when conditions permit. The section **Imbalance Reserve Supply** (page 6) shows that the amount of reserves available to a system is sensitive to how a balancing authority manages its resources. Ensembles of machines under coordinated control, such as under automatic generation control (AGC), can typically provide much more rapid response than can any smaller set of generators. The paper shows that the imbalance supply can be equated to a particular geometric object, the ***convex hull***. The convex hull shows the amount of time a resource ensemble can provide various ramp rates for imbalance excursions. It also provides a partial ordering on resource ensembles that makes it meaningful to speak of "larger" and "smaller" ensembles. The ensemble also can be put into a unique form, which this paper calls ***the standard form***.

Discussion then turns to **Imbalance Reserve Requirements** (page 15). The paper uses INC events to explain concepts. It addresses DEC events only near the end of the section (page 25).[1] Characterizing imbalance requirements is done in stages. First, the exposition considers the simple case where requirements are strictly increasing, that is, there is no recovery of capacity (page 16). Requirements are identified at this point with the machine ensembles necessary to meet the requirements. The ***cumulative ramp duration curve*** or ***CRDC*** facilitates the comparison of the imbalance requirement with reserve supply. An example compares the strictly increasing requirement to both adequate and inadequate supplies. The argument then considers the ***recovery of*** capacity dispatched to meet imbalances (page 19). This requires careful description of the initial conditions at the beginning of a response. The combination of the initial conditions with the response itself will be referred to as a "path." (Paths may have trivial initial conditions.) Paths reflect the net deployment of resources from the beginning of an excursion

[1] In what follows, an INC is a positive excursion from the zero level of balancing reserves, such as those shown in Figure 1; a DEC is a negative excursion from the zero level. A response, for either INCs or DECs is an increase in the magnitude of the excursion. A response during an INC is a requirement for more generation resource relative to load; a response during a DEC is a requirement for less generation resource compared to load. A recovery is the opposite of a response. A recovery is a reduction in the magnitude of the excursion.

through the end of any particular response.

The total system imbalance requirement over a particular history of excursions is determined by the minimum resource ensemble necessary to meet ***all of the paths*** encountered in the sample history. While it is attractive to think that providing for the largest excursion or the largest ramp rate would suffice, in fact no single path or subset of paths provides sufficient information. A simple history of responses will double-count ramp rate requirements, because recoveries will likely restore some ramping capability. Alternatively, the initial conditions present at the beginning of a response do not reflect all earlier requirements. The paper shows (page 22) the convex hull on the union of vertices of the path CRDCs is a means to calculate the requirement over all paths. This convex hull for demand has unique description in terms of the minimum resources required to meet the demand. That description, in fact, is the desired requirements characterization (page 24).

A somewhat digressive section entitled **The Requirements Algorithm** (page 27) presents the results of a computer program that implements the ideas presented in this paper. The section walks the reader through the calculation of a path, for example. The calculations use the imbalance requirement example that introduces the paper (Figure 1 on page 3). It presents the convex hull and standard form (Figure 30, page 30) for that example requirement. These should help to make the new concepts clearer and more familiar.

The section **Measuring the Sufficiency of Imbalance Reserves** (page 31) returns to the comparison of supply and demand for imbalance reserves using convex hulls. The comparison indicates whether a given supply meets all path requirements. Attached to the paper is a proof that the resources will be sufficient to meet the requirements if and only if the demand hull is contained within (or is “below”) the resource supply hull.

Up to this point, convex hulls have been created from constituent time series and resource capabilities. Building on example data from the preceding section, the section traces the sequence in the opposite direction. The section provides an example of going from the convex hull and resources in standard form back to the specific sequence of events that gave rise to insufficiencies (page 35).

The final section of the paper, **Our Assumptions Revisited** (page 35), reconsiders the introductory assumptions. Our current knowledge is limited regarding machine constraints during response and recovery. It is prudent to consider the consequences of alternative assumptions for this model. The section finds that minor bookkeeping changes would permit us to manage exceptions to the weaker assumptions. The section concludes with ideas, not directly related to the assumptions, for extending and improving the approaches presented throughout this paper.

An attachment (**Formalities**, page 43) contains rigorous proofs of some of the more technical claims. It is not intended for the general readership. Other attachments include the input data and code used in preparing the example in the section **The Requirements Algorithm** and a Table of Contents for this paper.

Introduction

In power engineering, imbalance requirements are defined as excursions from a baseline (zero) left over after providing load-following and scheduled exchanges. The following six-day example of five-minute observations (Figure 1) is from the Bonneville Power Administration's (BPA) website[2]. While resources generally are not dispatched exactly to the requirement, this history will provide us with an ideal case, which can be modified as practice dictates.

The principal requirement is that sufficient resources with adequate ramp rates are available to meet the minute-to-minute or even second-to-second variations that present themselves. These requirements depend on the particular sequence of events. That is, resources can recover from excursions when the sign of the requirement changes. A requirement that is punctuated with periods of recovery is consequently quite different than one that is not. This is true even if the magnitude of the excursion and the maximum ramp rates are identical.

Trying to characterize supply of reserves is equally challenging. Different combinations of resources can provide the same service. Two 1-MW units, each with 1 MW/min ramp rates can substitute for a 2-MW unit with a 2 MW/min ramp rate. These may not be economic equivalents, however.

Multiple units with low ramp rate, deployed at the same time, can replace a single unit with high ramp rate. Any unit with higher ramp rate can meet lower ramp rate requirements, of course. Some substitutions that we think might work, however, are not equivalent. We cannot substitute

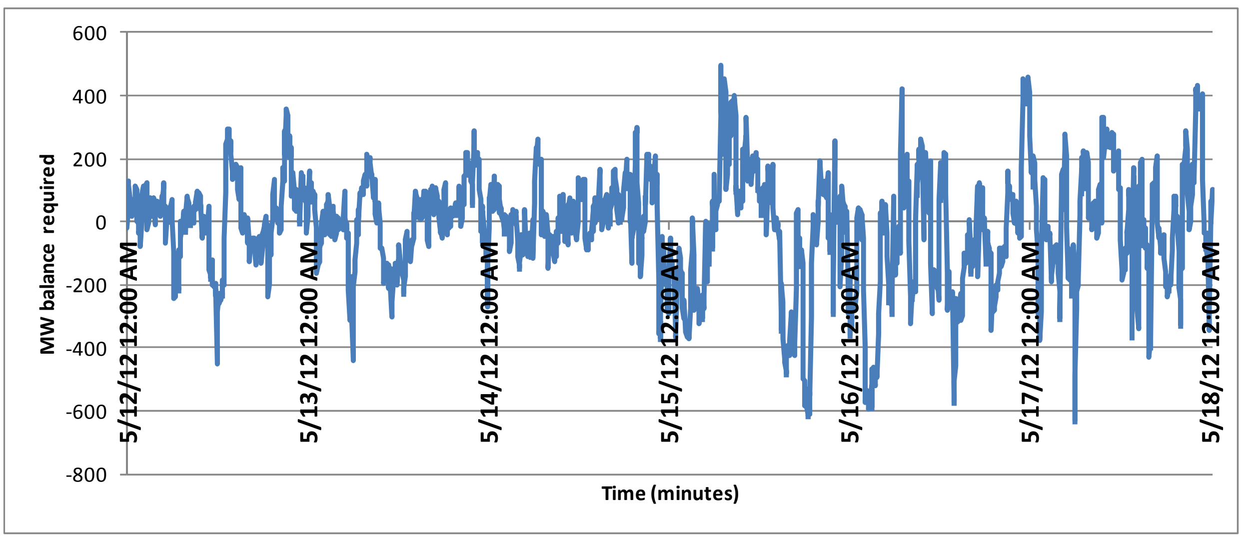

Figure 1: Imbalance Reserve Requirements

one 2 MW unit with 1 MW/min for two 1 MW units with 1 MW/min ramp rate. The second will ramp twice as fast as the first.

Other complications are the question of the protocols for deploying and restoring units. For example, once units have been deployed, what is the best way to recover their capacities to

[2] http://transmission.bpa.gov/Business/Operations/Wind/reserves.txt

prepare for the next, unforeseeable excursion? The answers to such questions will depend on machine properties and on policies about managing the risk of not meeting load.

The objectives of this paper will be as follows.

1. Create the simplest possible model that quantifies the supply and demand of reserves necessary for imbalance requirements.
2. Show how the model indicates the insufficiency of supply.
3. Address the assumptions of the simple model and suggest changes to the basic model that would make it suitable for more general situations.

Before presenting the engineering argument, it is helpful to state certain assumptions. These assumptions are not critical to the arguments that follow. However, fixing the rules for how resources are operated will make following the argument much easier.

1. Perfect foresight is not assumed. In particular, resources return to a state wherefrom they can be redeployed as soon as possible. This will be referred to below as the "standby" condition for the resource. This condition is typically – but not always – where they can be deployed to provide equally expected up- and down-balancing demands.[3]

2. The supply curve for response services has increasing marginal cost. The figures presented here use the assumption that high-ramp rate resources and services are more expensive ($/kW and $/MWh) than lower ramp-rate resources. Machines capable of providing higher ramp rates may also have limited availability. (The paper discusses imbalance reserve economics at greater length in the next section.) As a consequence of these considerations, higher-ramp rate units are restored to standby before lower ramp rate units.

3. Each time the requirement for balancing reserves crosses zero, all resources are returned to their standby condition. This is not hard to achieve. It is really a consequence of the first two assumptions. As the system requirement returns to zero, the highest ramp-rate units are returned to ready, then the second-highest ramp rate units are returned to ready, and so forth. When the zero line is crossed, the total displacement must net to zero. The protocol therefore requires that even the lowest-ramp rate units must be returned to their standby condition.

4. Units can recover capacity at least as fast as they respond. This is not obviously unreasonable for INC reserves, because it takes power to increase torque, but torque can be removed in an instant. So far, it appears to comport with practice as well, at least

[3] The concepts in this paper do not use optimization as it is commonly understood. For many situations, optimization and imperfect foresight are fundamentally antithetical concepts. The ideas here do use a special kind of optimization, however, that lends itself well to imperfect foresight, namely the "greedy algorithm." While it may not be recognized by that name, the greedy algorithm is ubiquitous in the power industry. For example, costs are minimized by dispatching the least expensive ($/MWh) resources, than the next most expensive resource, and so forth. This myopic form of optimization works for a surprising number of problems, for example, in mathematical graph theory, statistics, economics, and engineering.

anecdotally.[4] For DEC reserves, we expect the converse, that recovery – which increases generation – has a more constrained ramp rate. DEC response would be at least as rapid as recovery.

The assumptions regarding recovery are perhaps the weakest. There may be thermal or other constraints on some kinds of machines that limit how fast machines can recover their capacity, for example. The paper returns to reconsider this assumption in the section **Our Assumptions Revisited** on page 35.

Also note that while the terms "capacity" and "power" have different meanings, either one may be applied to a given example, depending on the context. Capacity is the ability to provide power; it determines the maximum available power or rate of energy delivery from a generator. The paper is attempting to find the power requirement for imbalance and by implication the units, ramp rates, and capacities necessary to meet those requirements. As we discuss an increasing maximum power requirement associated with different situations, therefore, we may use the expression "increasing capacity."

We begin by addressing the situation where a known set of power generation resources have a given amount of imbalance reserve capacity. A background on the economics of these resources, however, will inform that exploration.

The Economics of Reserve Supply

It will be convenient to have an economic "merit order" for imbalance reserves. The merit is given by ramp rate. The convention will be that higher ramp-rate power generation will have higher gross value than lower ramp-rate generation.

This merit order has support in an arbitrage argument. A MW of capacity with a higher ramp rate can provide all of the imbalance service that a MW of a unit with lower ramp rate can provide. That is, a MW of capacity with a higher ramp rate can always be deployed more slowly. Moreover, a MW of capacity with a higher ramp rate can participate in transactions that would not be available to a MW of capacity with a lower ramp rate. A MW of capacity with a higher ramp rate must therefore have value ($/MW) at least as great as that a MW of a unit with lower ramp rate. Its availability for higher-ramp rate transactions would make the value strictly greater.

The same is true of the value of the energy ($/MWh). Should a MW of capacity with a higher ramp rate be deployed, when a MW of a unit with lower ramp rate is available and would suffice? There is an opportunity cost using higher ramp-rate capability prematurely. We therefore expect a MW of capacity with a higher ramp rate must have an energy value ($/MWh) at least as great as that a MW of a unit with lower ramp rate. Its availability for additional transactions would again make the value strictly greater.

[4] Kevin Nordt, Director of Power Management for Grant Public Utility District, in a telephone conversation, June 5, 2012. Also note that a recently released IEEE standard 1310-2012 deals with precisely this question. The title of the standard is "IEEE Recommended Practice for Thermal Cycle Testing of Form-Wound Stator Bars and Coils for Large Rotating Machines"

We can illustrate the cost relationship with the supply curve in Figure 2. This is a conceptual illustration showing increasing value with ramp rate capability. It is not intended to represent any particular generation, so no increments are shown on the axes. It basically shows that value in $/MW is strictly increasing. While the illustration has a convex curve, any increasing function is possible, including one that is concave. The only constraint is that slope remains positive.

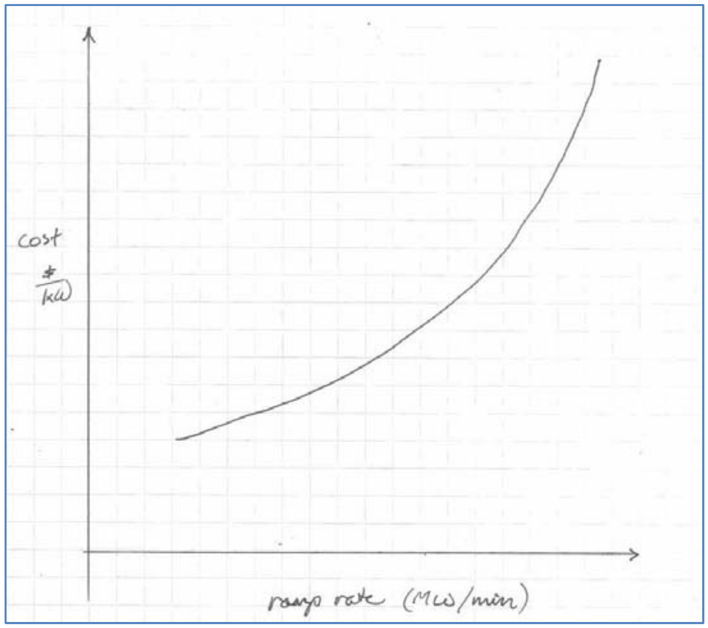

Figure 2: Ramp rate supply curve

Some may protest the use of a ramp rate on the horizontal axis for a supply curve, which typically shows increasing quantities. They would probably not protest the use of capacity (MW) on the horizontal axis. Capacity, however, stands with respect to energy as ramp rate stands with respect to capacity. The first is simply the rate of change of the other. We have seen, in fact, that the number and variety of transactions available to a MW of capacity increases as the ramp rate increases. We explore this observation further in a moment.

Because it has higher value, the recovery of higher ramp-rate resources should precede that of lower ramp rate resources if at all possible. This is our Assumption 2 on page 4.

Note also that all of the arguments above speak of gross value of imbalance reserves. The net value of imbalance reserves will depend on the cost to provide the reserves, which the paper does not address. Suffice it to say that costs can vary widely and resource operators may refrain from using or offering capacity into the market if costs exceed the gross value.

Imbalance Reserve Supply

Given a list of resources and their capability to provide INC and DEC reserves, what is their combined ability to provide imbalance reserves? The answer to this question depends on how the resources are deployed. This paper will use the term "protocol" to refer to how a system operators use their resources.

Consider the situation where, by custom or convenience, balancing authority operators rely on a single source of reserves to meet much of the reserve requirement. This occurs, for example, in situations where hydrogeneration at mid-Columbia is the principal source of load following and imbalance reserves. If the main source of imbalance reserves becomes exhausted, of course, the operators would then turn to other sources as the available equipment ramp rates and costs dictate. We refer to this as "sequential" deployment.

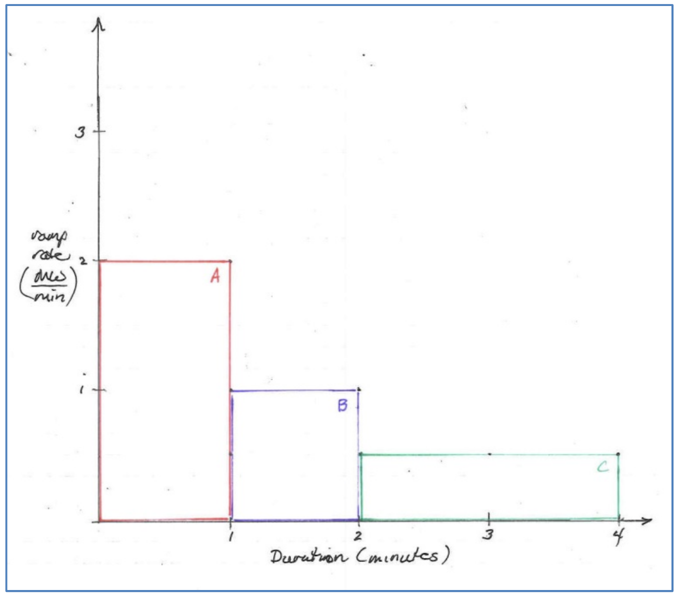

Figure 3: Sequential Ramping

Alternatively, resources could be controlled by AGC, which provides the opportunity for simultaneous dispatch.[5] We refer to this as "simultaneous" deployment. In this case, the potential system ramp rate of the ensemble of the resources would be higher because the system ramp rate would be the sum of the unit ramp rates. A few pictures will illustrate the difference.

Figure 3 shows the first situation, where an operator relies on one unit at a time to meet imbalance requirements. The vertical axis shows the ramp rate in MW/min; the horizontal axis indicates how many minutes this ramp rate can be sustained. Unit A, on the left-hand side of the graph has a ramp rate capability of 2 MW per minute and can sustain this ramping for one minute. Of course, this means the capacity of the unit is 2 MW, which is also the area of the rectangle representing Unit A. Similarly, Unit B, which lies to the immediate right of Unit A, has a ramp rate capability of 1 MW per minute and a capacity of is 1 MW. Unit C, which lies to the right of Unit B, has a ramp rate capability of one half MW per minute but also has a capacity of 1 MW. This means that Unit C can ramp twice as long as Unit B.

Contrast the first situation with the second, where the same three units can be dispatched simultaneously. This is illustrated in Figure 4. Here, the units are sorted by the duration over which they can provide their maximum ramp rate, from shortest to longest. Unit A has the same duration as Unit B, so their stacking order is arbitrary. Now the system is capable of ramping at 3.5 MW per minute, but only for one minute. After one minute, dispatch capability falls to one-half MW per minute. Note that when sorted by duration in this manner, the ***total system*** ramp rate is automatically sorted by ramp rate.

When illustrated this way, the advantages[6] of the simultaneous ensemble in Figure 4 are not obvious. What is missing from the picture is the ability of a high-ramp rate unit to substitute for a low-ramp rate unit. That is, there is nothing that prevents the simultaneous ensemble to dispatch exactly like the sequential ensemble in Figure 3. One way that this could be accomplished, using proportional dispatch, appears in Figure 5. (This is one of a myriad of possible schemes.) What the figure illustrates is the simultaneous contribution of the three units on the left-hand side of Figure 4, all reduced to ramp rates below their maximum rates. For the

[5] This situation is not far-fetched. It is how utilities with coal-fired and nuclear power plants can meet imbalance requirements. It also appears to closely resemble a protocol, Reliability-Based Control (RBC), that the National Electric Reliability Council (NERC) is currently evaluating for possible implementation. (See http://www.nerc.com/filez/standards/Reliability-Based_Control_Project_2007-18.html .)

[6] We need to acknowledge that this discussion completely sidesteps the issue of equipping power plants to provide simultaneous service. That equipment is not without cost. Implicitly, we are assuming the cost of communication and control equipment is negligible compared to the value of the service.

first minute, they are all ramping at 2/3.5 of their maximum rate. In the second minute, they are ramping at 1/3.5 of their maximum rate. In the third minute, they are ramping at 0.5/3.5 of their maximum rate. In the fourth minute, the capacities of Units A and B are exhausted, and the remaining capacity in Unit C is deployed at its maximum rate.

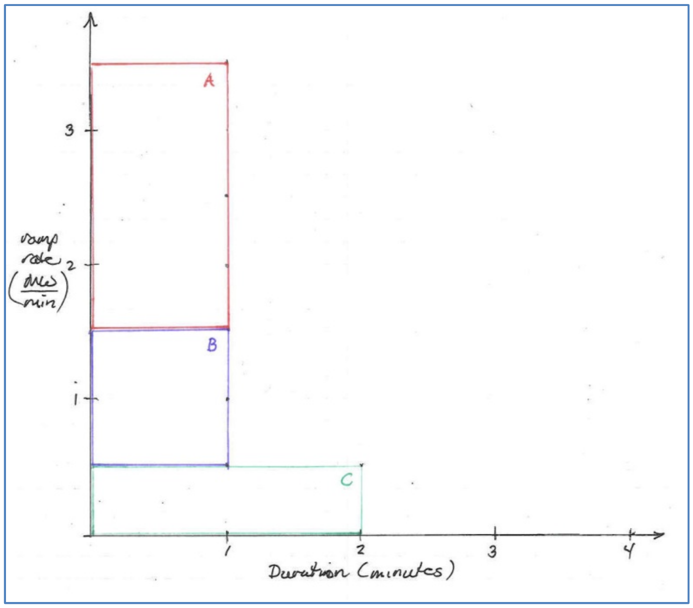

Figure 4: Simultaneous Ramping

There are two things about Figure 5 that suggest it is not unrealistic. First, at no time does the ramp rate of any unit exceed it maximum. Second, the picture respects the capacities of each unit. To see this, note that in each of the first three minutes, the Unit A ramp rate is 4/7 of the total. This corresponds to its share in the first minute of Figure 4. Likewise, Unit B is 2/7 of the total, and Unit C is 1/7 of the total. Recall that the area of the blocks corresponds to capacity. In the first three minutes of Figure 5, the total capacity deployed is 7/2 = 3.5 MW. This means 2 MW = 4/7*7/2 of Unit A is deployed, 1 MW = 2/7*7/2 of Unit B is deployed, and 1/2 MW = 1/7*7/2 of Unit C is deployed. One-half (1/2) MW of Unit C is available for minute 4 in Figure 5.

So how do we go about capturing, in a figure, this ability of a faster ramp-rate unit to substitute for a slower ramp-rate unit? Clearly, the two features we need to recognize are ramp rate and capacity.

The Convex Hull (Part I)

Given a protocol for deployment, like those illustrated in Figure 3 and Figure 4, the system capabilities in the figures can be sorted by their ramp rates. The result is called the ***ramp duration curve (RDC)***. The system capability reflects the individual resource capabilities and the system protocol, as we have seen. The particular resources and protocol, however, cannot be determined by this system

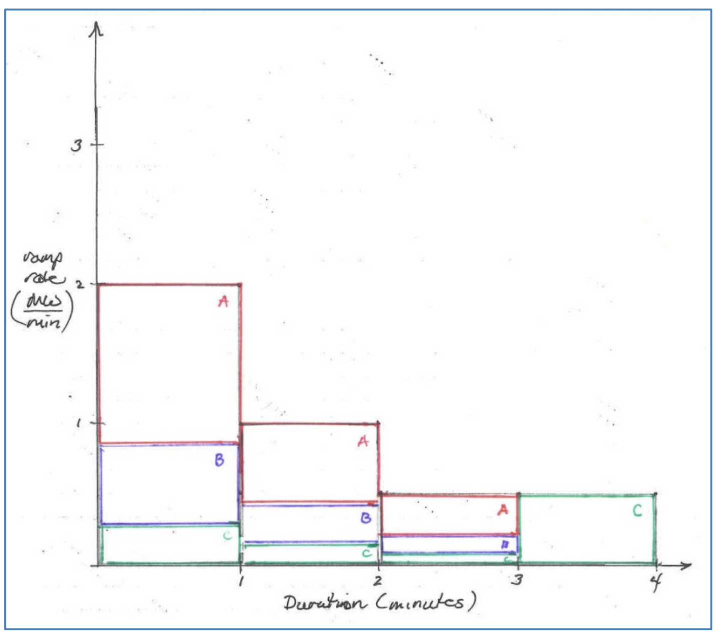

Figure 5: Substitution

capability. Figure 6 shows the system capabilities associated with the last two examples. Again, these particular examples represent the same resources, but with distinct protocols.

It may not be obvious that we need to do this, because resources in the preceding examples have

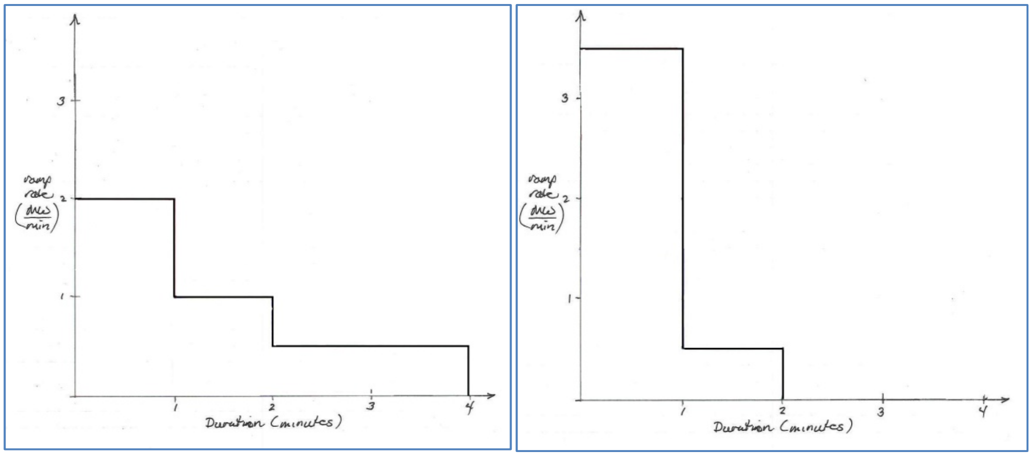

Figure 6: System capabilities for sequential (left) and simultaneous (right) deployment

already been sorted. However, there is nothing in principle to prevent operators from deploying resources with differing ramp rates in any particular sequence. Blocks corresponding to Units A and B in Figure 3 might have been reversed, for example, for various reasons.

Sorting the system capabilities by their ramp rates, however, is the first step in making two distinct sets of resources and protocols comparable. It is straight-forward, for example, to compare the ramp rates of the fastest units in the two sets. This is clearly a fundamental difference between systems.

Sorting the blocks, however, achieves another important end. Sorting makes easier to determine which substitutions are permitted. Any block can substitute for any other block to their right. If the ramp rate of the left-most block is not needed to meet a particular requirement, its capacity is available to meet a requirement with lower ramp rate. This kind of asymmetric substitution is a key feature of imbalance resources.

In adding the potential capacities left to right, what we are constructing is a ***cumulative ramp duration curve (CRDC)*** for capacity as a function of time (potential duration). The slope of each line segment in the CRDC indicates the ramp rate of the corresponding RDC block; the height indicates the cumulative potential capacity available. Figure 7 shows the CRDC corresponding to the resources and protocol in Figure 3. The dashed lines are not part of the CRDC but are present merely to guide the eye.

A notable feature of the CRDC is that slopes of the line segments are positive and strictly decreasing. This makes the area under the curve convex in the mathematical sense of the word. A straight line can be drawn between any two points on the boundary and will lie under (or interior to) the curve itself.

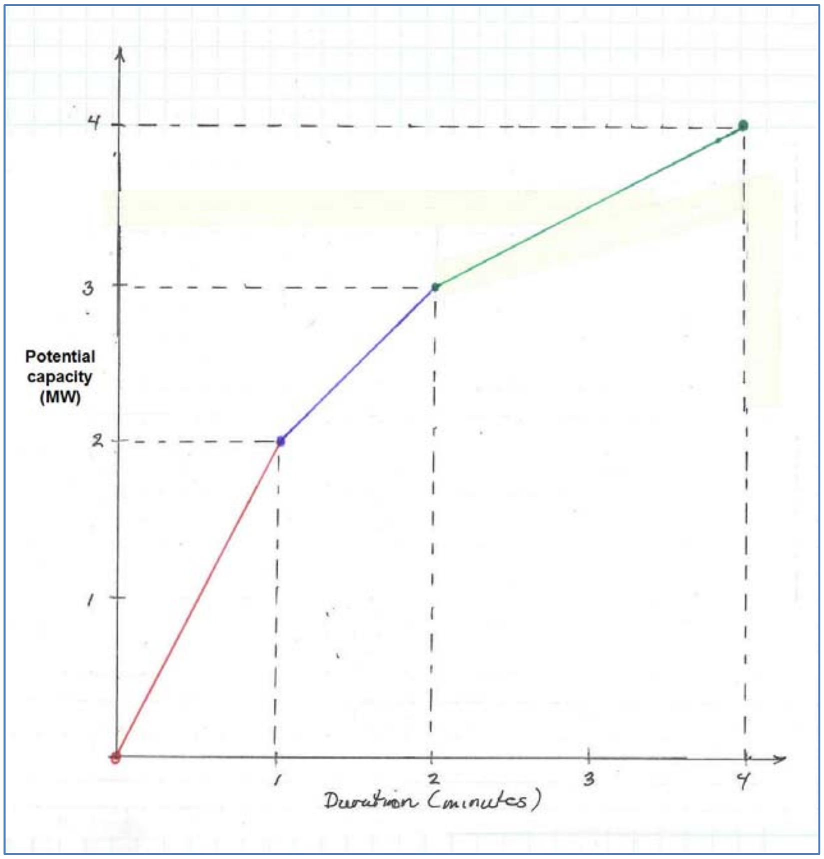

Figure 7: Convex hull for sequential case

We convert the CRDC into a ***convex hull*** by adding two points. The mathematical definition of a convex hull requires that it be a closed polygon. We satisfy the definition by adding two points to those four points defining the convex hull in Figure 7. One point is placed at (X, 4) and one at (X, 0) on the horizontal axis. The value of X is chosen "sufficiently large." What this means will be clear in a few more paragraphs.

The addition of these two points, however, is not simply one of aesthetics. It is physically the case that the total system capacity can never exceed – in this example – 4 MW, irrespective of the duration of interest. Adding these points also makes it easier to communicate how two systems of resources compare, as we will now see.

In Figure 8, we have added the convex hull for the *simultaneous* system described by Figure 4. Note that the convex hull for the sequentially dispatched system never rises above the convex hull for the simultaneous system and lies below it until the end of the 4-minute duration. Here is where the value of the convex hull as a comparative tool is evident. At a glance, we can tell that the simultaneous system is in a sense superior to the sequential system. We have already discussed the reasons for this. The graph makes it evident that any requirement the sequential system can meet, the simultaneous system can meet as well. But the converse is not true.

In Figure 8, we have also added the two points mentioned in the previous paragraph that make the curves into convex polygons. The values of the points are (4, 4) and (4, 0). The value 4 for the variable X in the preceding paragraph makes it clear what it means for the sequential hull to "lie below" the simultaneous hull. Later, as other comparisons of hulls become necessary, the X value for the two added points will be shifted to the right as necessary to facilitate comparisons.

Ordering Resource Ensembles

There is a natural order for resource ensembles, which Figure 8 has introduced. If the top of the hull of one ensemble lies over the hull of another, we will say the first is "***larger than***" or "***greater than***" the second.[7] If the hulls cross at any point, we will say that they are not comparable or that one will be ***insufficient*** to meet the requirements met by the other. The reasoning for these statements follows.

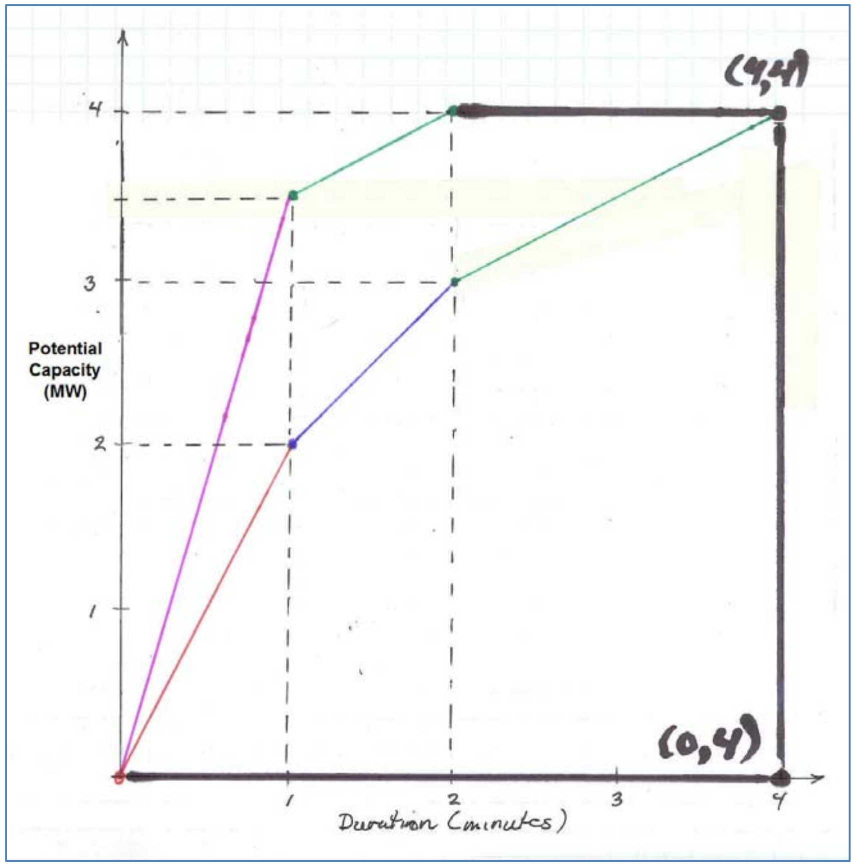

Figure 8: Comparing the convex hulls of two resource ensembles

Figure 7 contains line segments that have interpretation as vectors. That is, they can be added and subtracted in meaningful and useful ways. For example, look at the first two edges (left to right) of a supply hull. As vectors, their sum is a vector from the origin to the right-hand endpoint (2, 3) of the second edge. It corresponds to deploying the total capacity of the two resources over their total duration at a constant rate. This just is a proportional dispatch of the two resources. The constant rate of deployment for the combined deployment is the average of their ramp rates. The same is true of any set of edges in this figure.

If the hull of one ensemble lies over another, it means that any requirement that the subordinate ensemble can meet, the superior ensemble can also meet by combining resource deployment as necessary. This is true because segment of the subordinate hull, there is a segment between points on the superior hull that has the same slope. This means we can meet the requirement.[8] This is another way of viewing the events in Figure 5.

If the hulls cross, however, each resource ensemble will have some requirement that it can meet that the other ensemble cannot meet. To see this, think about a simple ramp requirement defined

[7] Characterizing these hulls as polygons, it might be more proper to say that one hull contains the other.
[8] The detailed argument appears in the attachment **Formalities** on page 44.

by a vector from the origin to the first vertex following a crossing and lying above the hull of the other ensemble. This ramp cannot be achieved by the other ensemble.

These observations are formalized in the next propositions.

Proposition 1: The capacity and duration of the sum of a set of vectors (hull edges) in a convex hull is the sum of capacity and duration, respectively, of its constituents.

Proposition 2: The ramp rate of the sum of a set of vectors (hull edges) in a convex hull is the average of its constituents' ramp rates with respect to their duration.

Proposition 3: If an ensemble has a convex hull that lies above that of a second ensemble, the first ensemble can meet the same requirements for ramping and capacity as can the second ensemble.

Commuting Ramping Events and Capability

To create the convex hull, it was necessary to sort the RDC resource blocks. It is reasonable to ask if any information is lost when the blocks are rearranged by sorting. Does the resulting object still have meaning? An analogy will help explain why sorting does not distort the physics or economics of the situation.

On page 6, it was observed that capacity stands with respect to energy as ramp rate stands with respect to capacity. This is true because one is the rate of change of the other. If ***R*** is the ramp rate, ***C*** is the capacity and ***E*** is the energy of a resource, this can be expressed concisely as follows.

$$R = \frac{\partial C}{\partial t} \text{ and } C = \frac{\partial E}{\partial t}$$

If power plants have adequate capacity to meet peak requirements and sufficient fuel to supply those plants, the specific order of dispatch has no bearing on their sufficiency.[9] The same is true of imbalance reserves. As long as we have adequate ramp rate to meet our requirement and sufficient capacity for each, the specific order of the ramps has no bearing on sufficiency.

As we mention in the introduction, recovery of unit capacity changes the "fuel use" of imbalance reserves. Consequently, we need to take care to exclude that situation. It will be useful to record this observation as a particular observation:

Proposition 4: For strictly increasing imbalance reserve requirements, the order of occurrence of requirements with different ramp rates has no bearing on resource sufficiency.

The analogy of imbalance reserves to energy systems also suggests a particular way of representing an ensemble of resources under a given protocol. When dispatching power plants to meet load requirements, the convention is to use "base load" units, intermediate units, and peaking units. Operators do not switch plants on and off according to the capacity of the load

[9] This is the reason load duration curves work in reliability analysis.

requirement. Clearly that would be unnecessary, complicated, and quite expensive. What we are doing with sequential deployment of reserves, however, is not that much different.

The Standard Form

Pursuing the line of thinking, consider a protocol for reserves that resulted in the ensemble shown in Figure 9. Note that this example has no periods of recovery, which would appear as negative values in this graph. For the purpose of this example, sorting the requirements will not provide any additional insight, so we will elect not to do so. In this illustration, we have intentionally left the effective protocol ambiguous. That is, Figure 9 tells us nothing about how many units are combining their ramp rates, if any, to achieve these results.

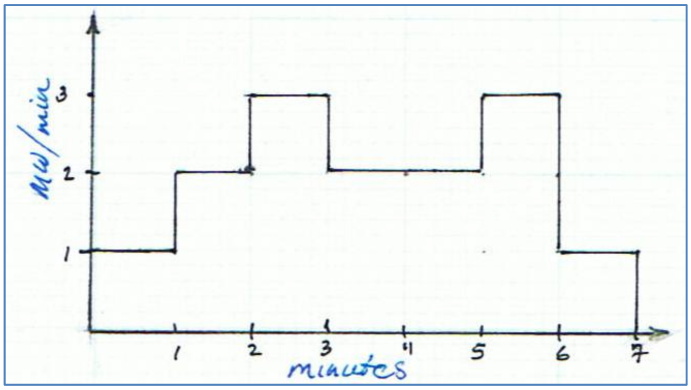

Figure 9: Ramp rates versus time

Now, if this were a profile of energy loads instead of ramp rates, we would conceive of plants, dispatched according to merit order, to meet the requirement. Power plants near the top of the dispatch stack would operate the least number of hours. That is, we would conceive of base load, intermediate, and peaking units along the lines illustrated in Figure 10.

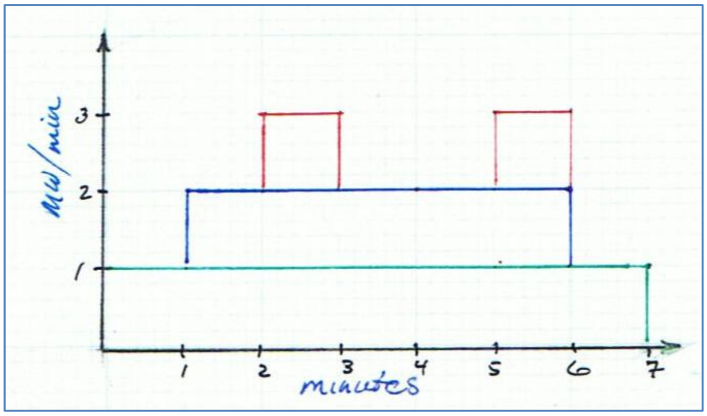

Figure 10: Decomposition of requirement

This arrangement also makes sense for resources providing imbalance reserves. Distinguishing services that contribute to meeting peaks better facilitates the valuation of the services. Ordering system ramping capability by duration (capacity factor) also sorts by system ramp rate. This means the representation better reflects the order in which operators deploy and recover the services. Operators naturally deploy resources according to their merit order to minimize cost. This is done without perfect foresight, confining deployment to the minimum system ramp rate resource ensemble. This practice is therefore in keeping with our assumption that operators cannot anticipate size or direction of the next requirement.

The real advantage of this particular RDC representation, however, that it is unique and bears a direct relation to the convex hull. Each face of the hull is associated with a particular increment of ramping capability. These increments typically do not have the same size. The only requirement is that their number of minutes of availability be declining as we move up the stack. Consolidating these to indicate the duration of service available from the ensemble produces

Figure 11. We will refer to this format the "***standard form***."[10] In this example, the standard form is as follows. All three units have 1 MW/min ramp rate. The "base load" unit has 7 minutes duration, or a capacity of 7 MW. The "intermediate unit" has a capacity of 5 MW, and the "peaking" unit has a capacity of 2 MW. Again, capacities are identified with the area of the rectangles.

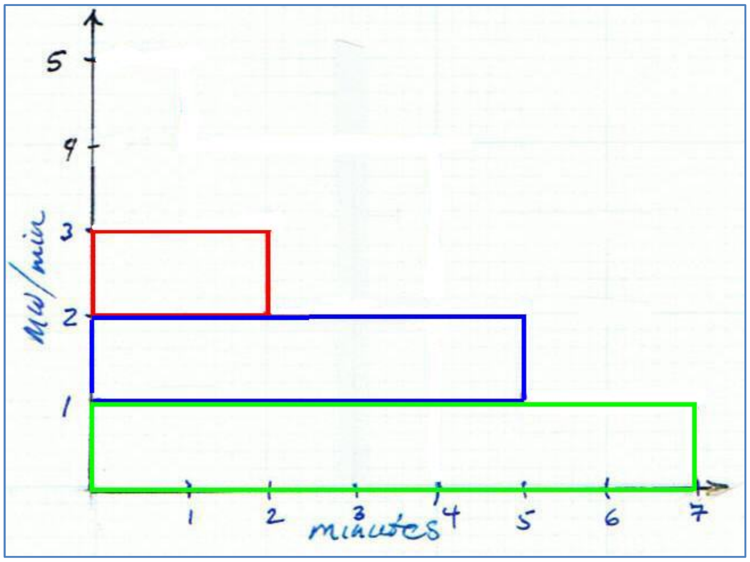

Figure 11: Standard Form for resources

To summarize the lessons of this section, we have learned that

1. The gross value of ramping services will have value ($/kW) that increases with the ramp rate. This stems from an asymmetric substitution among such resources.
2. The amount of imbalance reserves of various ramp rates depends in a sensitive fashion on the way that the resources are deployed, i.e., their protocols.
3. The RDC and convex hull for imbalance resources represents both the ramp rate and capacity available from an ensemble of resources.
4. The value of the convex hull for imbalance resources is its ability to capture and represent the value that stems from the asymmetric substitution of the resources. That is, it reflects the fact that a MW of higher-ramp rate capacity can serve the requirements of a lower-ramp rate MW, but not conversely.
5. If an ensemble has a convex hull that lies above that of a second ensemble, the first ensemble can meet the same requirements for ramping and capacity as can the second (**Proposition 3**).
6. Capacity stands with respect to energy as ramp rate stands with respect to capacity, or

$$R = \frac{\partial C}{\partial t} \text{ and } C = \frac{\partial E}{\partial t}$$

[10] In an earlier paper, we gave it the misleading and ultimately inaccurate name of "the least-cost" form.

7. For strictly increasing imbalance reserve requirements, the order of occurrence of requirements with different ramp rates has no bearing on resource sufficiency (**Proposition 4**).
8. The standard form of a resource ensemble has the advantage of a) being unique, b) corresponding to the geometry of the convex hull of the system, and c) reflecting the increasing marginal value of ramp rate.

These observations are also preparation for the next section. They will be applied in an alternative form to characterize reserve requirements.

Imbalance Reserve Requirements

This section develops an approach to quantifying imbalance reserve requirements. As we shall see, this approach has the advantage of making sources of and requirements for imbalance reserve comparable.

The approach is best developed in stages, however, because of the distinct behaviors we expect of machines and operators during excursions and the recoveries from excursions. Initial discussion is limited to examples of INC excursions. The first stage of examination will study a sequence of increasing requirements. The second stage will explore requirements where recovery occurs. The initial conditions enable us to capture initial machine loadings where there is partial recovery of capability. A response with initial conditions is called a "path" in this paper. Finally, the paper presents a technique for determining resource requirements over an arbitrary number of such paths. The paper shows this larger resource requirement characterizes the total requirement uniquely. Finally, the paper discusses considerations in evaluating DEC excursions.

The following example illustrates some of the difficulty in evaluating imbalance reserve requirements. Figure 12 shows a fragment of a hypothetical INC excursion. Let us say that we are keeping track of excursions to create some statistical description of this particular requirement. Counting the number of and ramp rate associated with each segment of the seven-minute snippet in Figure 12, we might construct the following table.

	Count	Minute obs
MW/min		
-1	2	3,4
+1	3	1, 5, 7
+2	2	2,6

The columns (left to right) are the ramp rates (MW/min), the number of times the ramp rate appears in Figure 12 ("Count") and the minutes in which they occur ("Minute obs").

What would these statistics tell us about our requirements? Very little, unfortunately – when requirements decline (a recovery for resources) in minutes 3 and 4, some resource would need to be backed down. It would not be unreasonable, given the assumptions laid out on page 4, that an operator would want to restore the resource that had provided the highest ramp rate. The highest ramp rate occurred in the second minute. That capacity, however, would then be available to serve the excursion that occurred in the sixth minute. That would suggest that 3 MW of 1 MW/min and 2 MW of 2 MW/min service should suffice to meet this requirement.

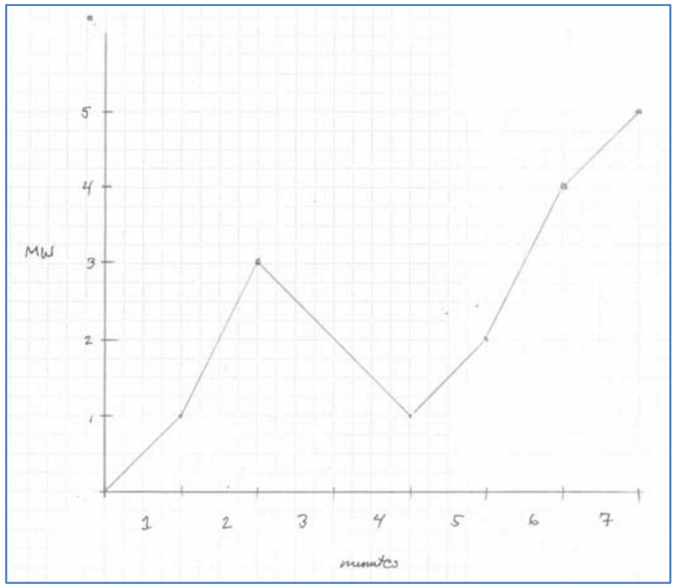

Figure 12: INC excursion

But what would the outcome have been if the recovery in minutes 3 and 4 did not occur until the end of this period? (After all, the statistics in the preceding table do not specify any particular order.) Then clearly not 2 MW but 4 MW of 2 MW/min would have been necessary. Placing the two periods of 2 MW/min ramping adjacent to one another provides no opportunity for recovery. This illustrates, of course, that the specific chronology of events is important.

Clearly our assumptions about what happens during a recovery are critical to any conclusions we might hope to obtain. However, in this section we will assume that recoveries are not constrained and we can recover units at least as fast as they can ramp. (This is assumption 4 on page 4.) We suspect, however, this may not be the case in particular situations. Specifically, when recovery occurs during a DEC event, the generation must increase which is exactly the constraint we are addressing in the INC event during a response.

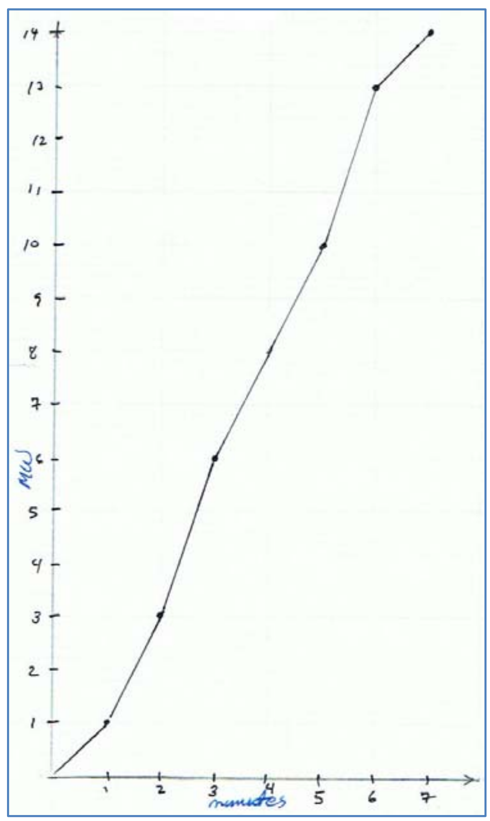

Figure 13: Increasing requirements

The assumptions about recoveries merit particular scrutiny. The final section of the paper returns to consider how alternative assumptions about recovery, deployment, and other features of behavior would affect the approach discussed here.

On a separate topic, note that we have begun to introduce a convention used repeatedly in the rest of this paper. Imbalance requirements will be identified with the "minimal" set of resources necessary to meet the requirements. What is meant by a "minimal set of resources" will become more evident with a few examples. We have already referred to the fact, however, that specific services should suffice to meet the requirement in Figure 12. This implies some set of machines and a protocol capable of providing that service. Be prepared for this paper to alternate between these two representations of requirements.

Strictly Increasing Requirements

Consider now the simple case of a sequence of increasing requirements. (See Figure 13) The key aspects of this sequence are the ramp rates, the slopes of the line segments in the figure, and the total amount of capacity required. Both of these features are captured if we graph the ramp rates versus time, as in Figure 14.

If we relate the requirements in Figure 14 to sufficient resources to meet the requirements, we can also view Figure 14 as the capability of some ensemble of resources. As was the case for resources, we can sort the requirements by their ramp rate requirements. The result is the ***ramp duration curve (RDC)*** of the requirement, and it appears in Figure 15.

We can now ask whether another ensemble of resources would be sufficient to meet this requirement. For example, assume there is an ensemble that could provide 5 MW/min ramping for one minute (5 MW) and 4 MW/min for three additional minutes (12 MW). Would that ensemble suffice? If we naively compare ramp rates, it is evident that the highest ramping requirements are covered, at least for four minutes. But what happens after that? This is where substitution again becomes important.

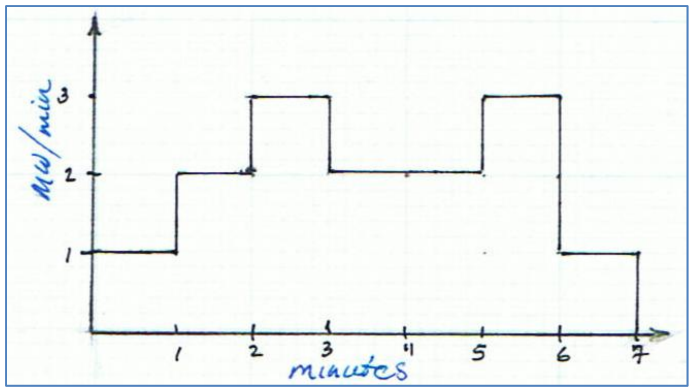

Figure 14: Ramp rates versus time

Figure 16 compares the requirement described in Figure 14 with the candidate ensemble of resources, described in the preceding paragraph. To compare the highest ramp rate requirement with the resources with the highest ramp rate resources, both are in RDC form. This permits direct comparison of the ramp rates. As pointed out, the candidate supply curve falls below the requirement's RDC curve at four minutes.

What is true about this situation, however, is that the candidate ensemble of resources ***is sufficient***. The total capacity provided by the high-ramp rate service exceeds the total capacity required. That is, the area under the blue line in Figure 16 is larger than the area under the orange line. The fact that the blue line falls below the orange line does not matter. The faster-ramp rate units can simply be deployed at a lower ramp rate. Observe, however, that an ensemble that provides 4MW/min for only two additional minutes (8 MW) is not! It lacks sufficient capacity.

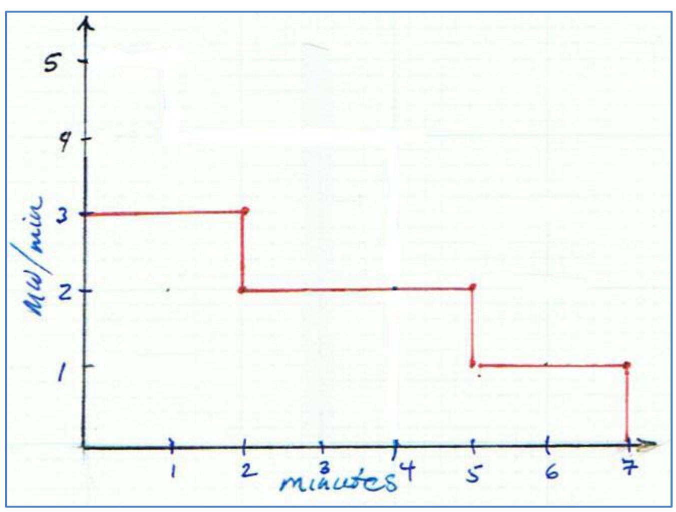

Figure 15: Requirements in RDC form

The preceding example is an early example of the fact that a system can be constrained not by its ramp rate, but by its capacity. This is analogous to energy dispatch systems that are not capacity constrained but fuel constrained.

The Comparison of Supply and Demand

A direct, graphical comparison of the preceding example is possible if we integrate or add up the areas in Figure 16 under the blue and orange lines, moving from left to right, as we did for the resources in the previous section. This operation creates the graph in Figure 17. The solid orange line is the ***cumulative ramp duration curve (CRDC)*** of the requirement. The reader will recognize the solid blue line as the convex hull for candidate resource ensemble.[11]

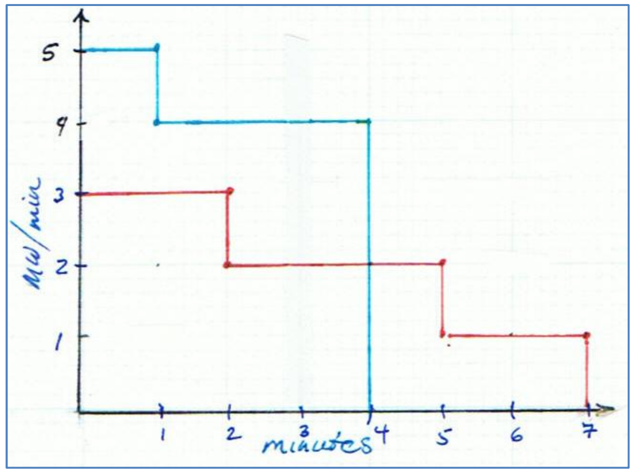

Figure 16: Comparisons

Figure 17 is a simple illustration of one of the fundamental findings of this paper. In this case, the finding says that a given ensemble of resources is sufficient to meet reserve requirements if and only if the convex hull of the supply is above the CRDC of the requirements. The result is more general and is proved in the attachment, **Formalities**. For the simple path CRDC we have here, however, the argument is simple. It is identical to that which we used for resources. If the CRDC lies above the supply hull at some point, no combination of ensemble resources can provide sufficient ramp rate for the number of minutes required. Otherwise, they can.

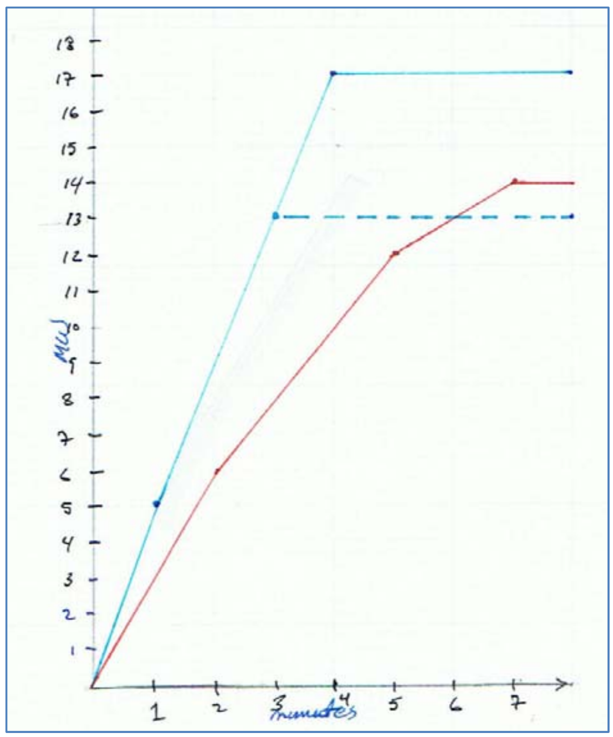

Figure 17: Convex hull for supply and CRDC of demand

The dashed blue line is the result for the inadequate ensemble consisting of 5 MW of 5 MW/min ramping and only 8 MW of 4MW/min. It is now evident that the inadequate candidate does not have enough capacity to meet the requirement. The line shows the inadequate ensemble can only provide 13 MW of service in total and the requirement will eventually exceed the supply.

[11] What we are calling the convex hull on the supply is essentially the CRDC form. When we get further along with requirements, we will distinguish the CRDC of a particular path from the convex hull on all paths. At that point, describing supply and demand as comparable objects will simply the discussion.

So far, we have considered a simple case where requirements are monotonically increasing and units are never allowed to recover. Of course, the status of machines during a response will depend on the condition for the machines at the beginning of the response period. The starting conditions, in turn, are the result of any recovery of unit capacity made in prior periods. The next section discusses this calculation.

Requirement Paths

As we have seen, it is necessary to incorporate the chronological history into a characterization of requirements. We will eventually see an example (page 33) where meeting the largest capacity and largest ramp rate requirement of Figure 1 does not guarantee sufficiency. Effectively, any approach to characterizing requirements will have to consider all of the paths of requirements encountered in an excursion. Our next example will help illustrate this.

Figure 18 is an example of a requirement path during an INC event. Up to the three minute point, the reserves resources are providing response. Between the three and five-minute marks, the reserves are in recovery. From five minutes to eight minutes, the reserves are again in response.

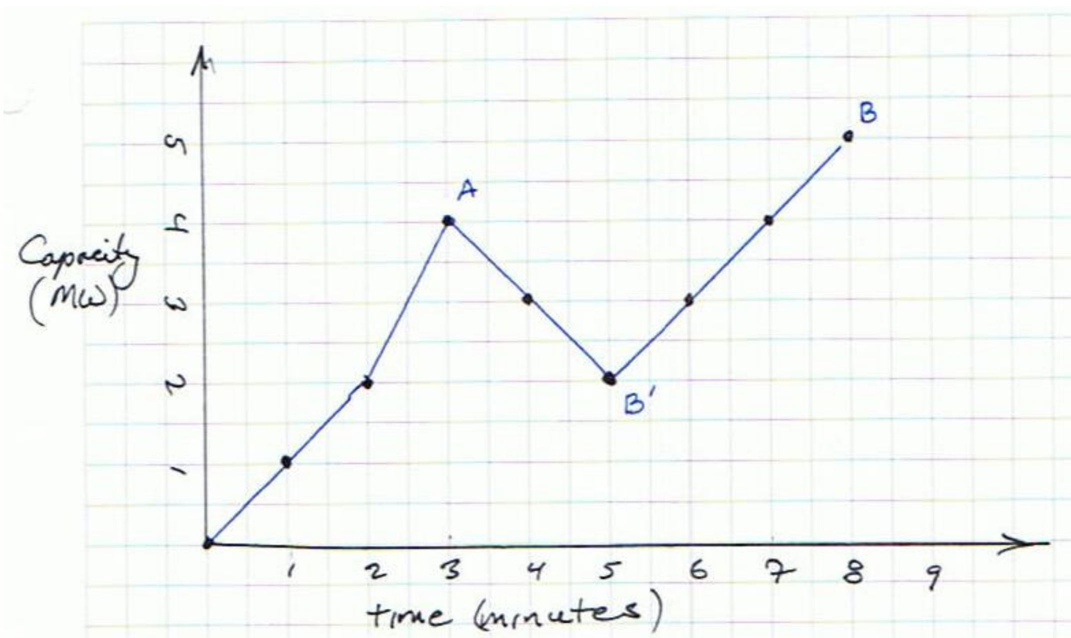

Figure 18: An excursion with recovery

There are, however, effectively two paths illustrated in Figure 18 that need to be considered. These are labeled “A” and “B” in Figure 19. Path A is something we have already explored in some detail in the previous section. It consists of a monotonically increasing requirement. An ensemble with 2 MW of 2 MW/min and 2 MW of 1MW/min ramping service suffices.

Path B, on the other hand, reflects the history from the beginning of the excursion, but it also includes a recovery. During the recovery, the 2 MW/min unit would presumably be returned to standby condition. Figure 20 shows the deployment and recovery of this unit. The recovery ramp rate is only 1 MW/min. This is half the response rate of the 2 MW/min unit. We assume that the 2 MW/min unit would have no difficulty reducing its loading at this rate. Recall from the assumptions on page 4 that the higher ramp rate unit will be restored to standby condition before any lower ramp rate units.

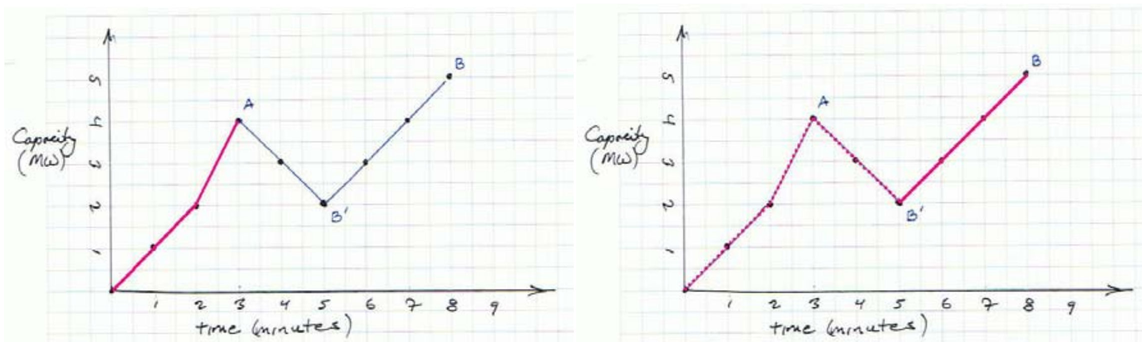

Figure 19: Two paths

At the beginning of the response phase in path B that begins at minute 5, labeled B′, only the 1 MW/min unit is ***not*** at standby. An initial condition for path B exists at this point B′. It is the result of the path indicated by the dotted line in Figure 19 preceding B′. After minute 5, the ramp never exceeds 1 MW/min. The period between point B′ and the point B comprises the response of path B. The requirement to meet path B, the combination of the initial condition and the response, nets to 5 MW of 1 MW/min resource.

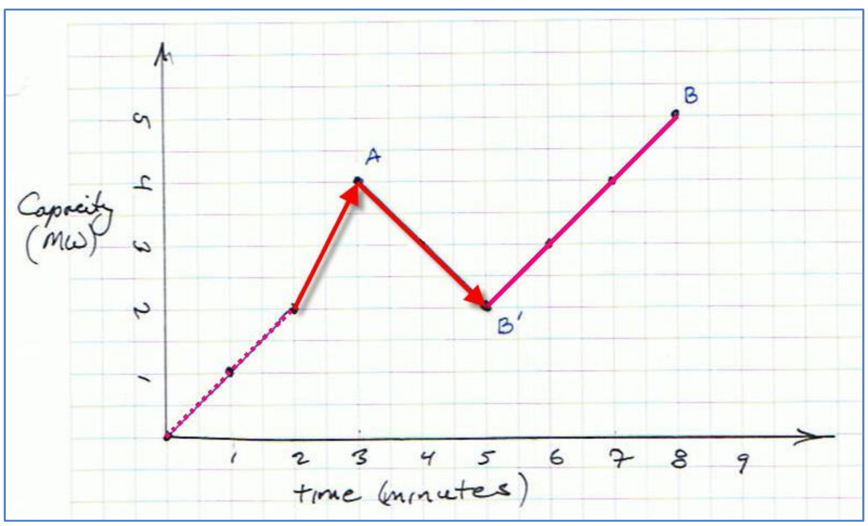

Figure 20: Returning 2 MW/min unit to standby

Figure 21 shows the CRDCs for Path A and Path B. We notice that the resources that would be sufficient to meet Path A would lack the capacity to satisfy Path B.

Figure 21 makes it evident that our treatment of paths reduces the excursion history to an initial condition, a unit loading. Moreover, the path summarized as an initial condition and a response ***is a strictly increasing requirement***. Because it is a strictly increasing requirement, the conclusions of the preceding section are applicable.

The path concept has significant implications that the reader should understand and feel comfortable using. For example, initial conditions are treated the same way as response events, in the sense that they may be rearranged, combined, and sorted with the response events. Paths are also independent of each other in the sense that the order in which they are evaluated has no effect on questions of sufficiency. The following briefly explores these practices.

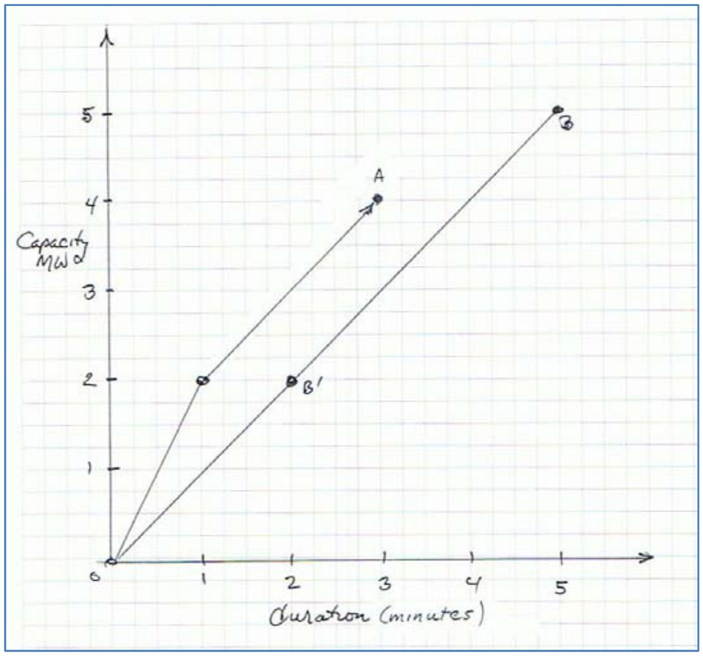

Figure 21: Requirements for Path A and Path B

Convincing oneself that replacing a history of behavior back to the beginning of an excursion with an initial condition, that is to say, an initial loading, requires a thought experiment. We are attempting to inventory the use of resources to meet a requirement. Reconsider Paths A and B in Figure 19. How much of which resources have we used? Starting with Path A, we have used 2 MW of the 2 MW/min resources and 2 MW of the 1 MW/min resources. The recovery returns the capacity of the 2 MW/min resources, but clearly we would have needed that resource to complete Path A. How much additional resource will we use to complete Path B? After recovery, an additional 3 MW of 1 MW/min will suffice. Now, in evaluating the total requirement, pay attention to which requirements these two paths have in common – common requirements should properly be counted only once – and which requirements differ. By marking requirements as initial conditions, we are merely recognizing that the resource loadings comprising the initial condition are shared in common with another path. The methods we will use to evaluate the joint requirement will count contributions toward such requirements ***as additions*** only if they differ from requirements that have already been counted.

By the same token, each path may be compared in any order with any other path. This is true irrespective of whether the path arises from the same excursion or a different excursion. Why does this make sense? The reason is similar to the one we just considered. If we are to identify the resources required to meet all requirements, the exercise remains one of tracking which requirements have already been counted and which requirements are different. If each path has an initial condition to indicate requirements in common with other paths, those requirements are not double-counted. The books will balance, so to speak.

The reader should also convince themselves that the treatment of recoveries makes sense, given our current assumptions. Recoveries have ***as their initial condition*** only the loading of the machines at the end of the preceding path. The unloading of units certainly does not contribute to requirements. The question then becomes which units will unload. If nothing prevents high ramp rate units from being the first to unload, economics suggests doing so is beneficial. Of course, assumptions about recovery can change this situation, as we will see in the section **Our Assumptions Revisited**.

Do not confuse the concept of a ***path*** with the events from the beginning of an excursion through the end of a response. Readers who do so will undoubtedly become frustrated. The path is a mathematical construct. A path has an initial condition, a particular capacity loading, which ***reflects or encodes*** the trace of events back to the beginning of an excursion. That is quite different from the events themselves. It is this property that makes it possible to change the order of path evaluations without affecting sufficiency estimates. It is also the key to understanding why the procedures for recovering imbalance capacity are less significant to the model than it might initially appear.

Minimally Sufficient Resources

We have considered paths with both trivial and non-trivial initial conditions. This is a good point at which to summarize in a proposition about the resources to which we are equating requirements:

Proposition 5: The CRDC of a path is equivalent to the minimal set of sufficient resources.

This should not be hard to see. As discussed, **Proposition 4** says we can reorder the requirements and place them in the form of a CRDC. The resources ensemble corresponding to and having the same shape as the CRDC must be minimal. Any other sufficient resource ensemble must have a convex hull that lie strictly above the CRDC curve at some point and therefore will be larger.

Returning now to the discussion of Figure 21 we ask, what kind of resource ensemble is necessary to satisfy both path A ***and*** path B? One contains a requirement for the faster ramping resource; the other contains a requirement for more capacity. The way to combine those requirements is by use of the convex hull for requirements.

Satisfying a Collection of Paths: The Convex Hull (Part II)

Consider the convex hull for requirements defined by the smallest polygon containing the vertices in Figure 21.[12] Figure 22 illustrates the top of this convex hull. Consider the minimal sufficient set of resources that would satisfy the requirements corresponding to this convex hull. We identify those resources with the demand hull itself. That this set of resources satisfies the requirements of both Path A and Path B is not evident. It would require a description of precisely how the resources corresponding to the hull are deployed in both cases. The method must be general enough that the specifics of the hull and the CRDC paths are not material. That description is the proof of the fundamental theorem of this paper, which is presented in the section **Formalities** (page 43).

[12] As explained on page 9, our convex hull has a requirement beyond those customarily stipulated for this object, that the hull have two vertices not present here that guarantee that the top of the hull consist of non-decreasing line segments.

The convex hull performs the tricky accounting described in the previous section. Not only will it count the common path requirements only once, but it counts unique path requirements only if they represent 1) higher ramp rates, or 2) greater amounts of capacities not already provided by higher ramp rates. This may not be evident to the reader, however, without significantly more experience with the structures.
We will not divert the discussion at this point to go through the theorem. If our reasoning is correct, however, it should be possible to illustrate how the specific requirement in Figure 18 is met with the resource ensemble described here. This is a good illustration of the principles discussed in this section and is achieved using the two figures below.

First, reduce the convex hull in Figure 22 into the RDC of the constituent resources. The standard form decomposition is laid out in Figure 23. The reader can verify the calculation by adding capacity increases from left to right.

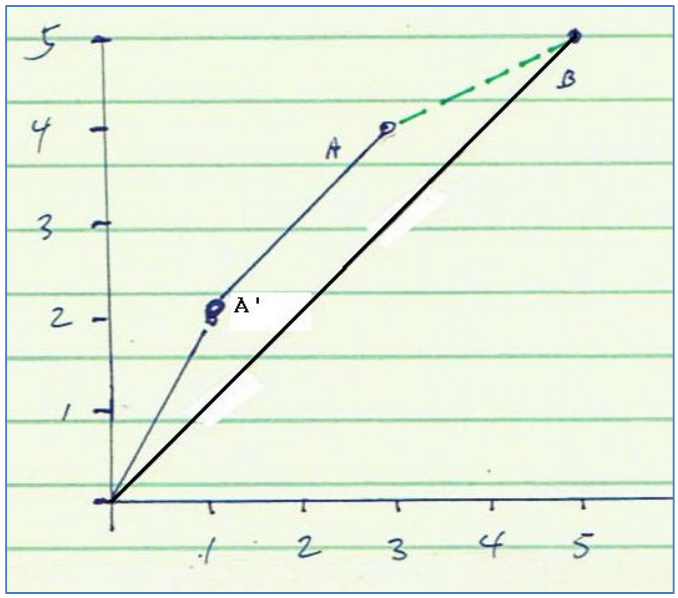

Figure 22: Hull for Paths A and B

Now, trace the dispatch of our resource ensemble to the requirements in Figure 18. To aid the discussion, the ½ MW/min resource at the bottom of the resource stack in Figure 23 and illustrated in green is called Resource I. The next resource in the stack, another ½ MW/min resource with smaller capacity is drawn in blue and is called Resource II. The 1 MW/min resource at the top of the stack is Resource III.

In Figure 24, the two resources with smaller ramp rate, Resources I and II, are deployed to provide the requisite 1 MW/min for the first two minutes. Resource III joins these two in the third minute to provide a total ramp of 2 MW/min. During recovery, first Resource III is restored to standby and then 1 MW of Resource II. (Since Resources II and III have the same ramp rate, this choice is entirely arbitrary.) At five minutes, the 1 MW/min ramp begins. One and a half MW of Resource I are depleted, leaving only 1 MW. Resource I can therefore contribute ½ MW/min for two additional minutes. Resource II, however, is only partially restored and starts out at minute 5 with ½ MW of loading. It therefore also has 1 MW (two minutes) of contribution left. In the last minute, Resource III must again be placed in service to complete the path.

The previous discussion illustrates some important techniques. This example is not intended as a proof. That task is relegated to the section **Formalities**. Nevertheless, we can begin to see how it is possible to move back and forth between the convex hull and the original description of the

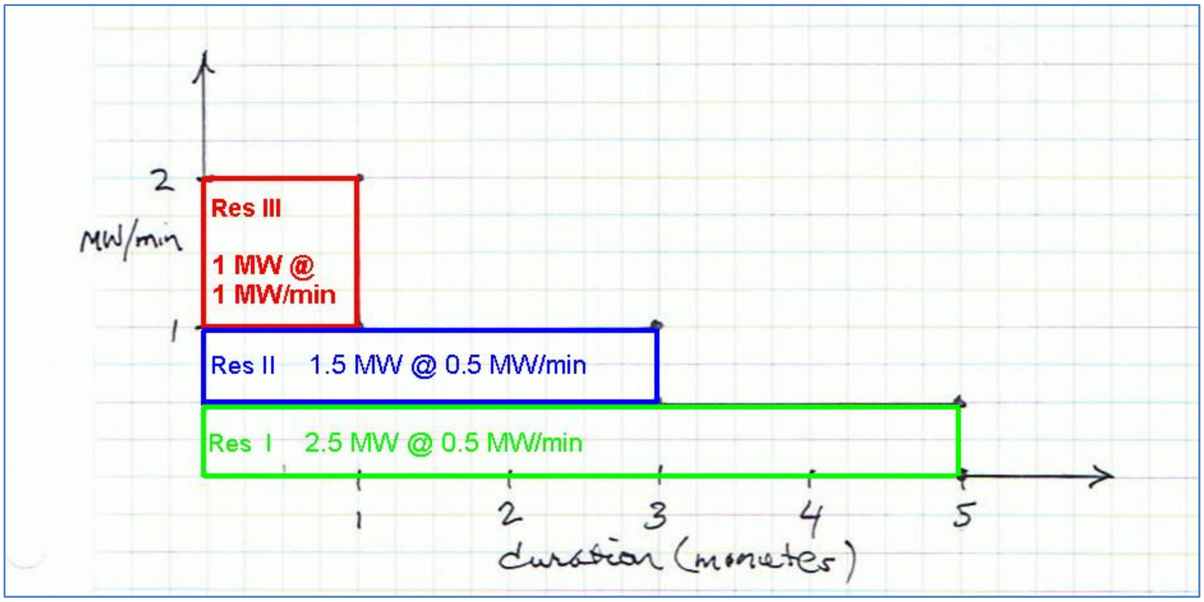

Figure 23: Decomposition to standard form

system. The section **Measuring the Sufficiency of Imbalance Reserves** will return with additional examples.

The Characterization of Imbalance Reserves

A complete characterization of INC reserves extends the two-path case in the preceding section to an unlimited number of paths. The construction of the convex hull for an arbitrary number of such paths can be performed from the complete collection of requirement vertices or incrementally, adding one path at a time. The result will be the same.

The incremental construction, for example, would take the convex hull on two paths, which is effectively their combined CRDC, and augment it with the third CRDC. The process of augmentation is the same as for two CRDCs. The joint CRDC for the three paths is the convex hull on their vertices. The process can be continued as many times as necessary.

The proof the main theorem for this paper takes the other tack. It compares the supply hull with the demand hull on all the paths and gives the technique for meeting any path selected. Since the choice of path is arbitrary, it effectively solves for all the paths simultaneously.

There will typically be many paths for each excursion and many excursions. Again, because we assume the machines are returned to standby mode at the beginning of each excursion, paths from distinct excursions have no effect on one another and may be comingled. That is, their comingling will have no effect on the final set of paths.

The algorithm for creating the measuring or characterizing INC requirement should therefore be clear. In brief:

1. Find the vertices for each path in the collection of INC excursions. The initial condition for the path will be trivial (no loading of any units) if the path is the first path of any excursion (INC event). Where the path's initial condition is not trivial, determine the initial condition given the rules for restoring imbalance reserve capacity. Our assumption 4 on page 4 holds that any capacity recovered will restore reserve capacity according to the resources' ramp rates. The resource with the highest ramp rate is restored first. Alternative practices are discussed in the section **The Recovery Assumptions** on page 38.

2. Find the convex hull on the collection of path vertices, observing the requirement of non-decreasing capacity stipulated on page 10.

3. Derive the standard form of resources from the convex hull. These resources have unique representation and completely characterize the requirement.

The discussion in this section has been careful to restrict examples to INC excursions. The topic of DEC reserves begins in the next section.

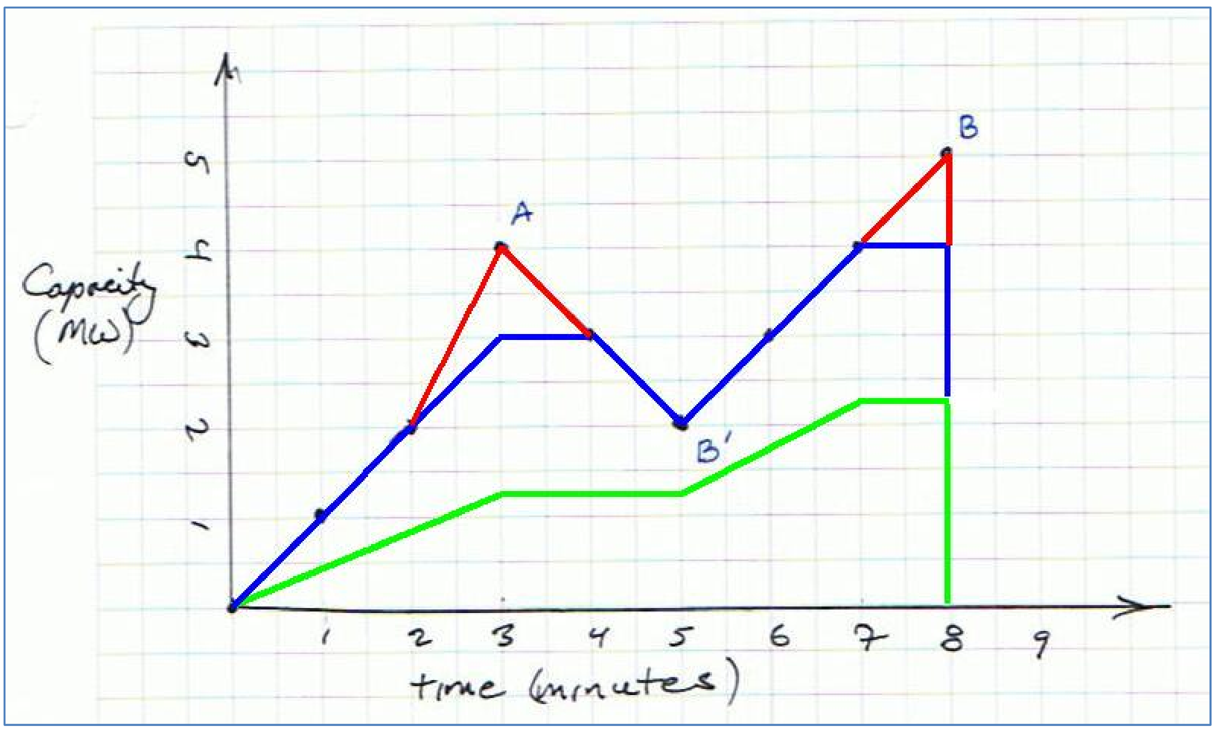

Figure 24: Hypothetical dispatch of ramping resources

The Treatment of DEC Reserves

If DEC excursions are similar to INC excursions, apart from the sign of the excursion, the same calculations discussed above would apply. There are reasons, however, why they could be quite different.

After adjusting for the sign change, their behavior could be the mirror image of INC excursions in the following sense. Recoveries, rather than responses, might be ramp rate constrained. It takes power – which is fundamentally limited – to increase torque, whereas decreasing torque and output can be interrupted virtually instantaneously. (When this is inadvertent, we call it a forced outage.) This suggests that DEC recovery ramp rates would be constrained because they entail increasing generation.

This situation is explored in greater detail in the section **Our Assumptions Revisited**. To see why this might ***not*** be a fundamental problem for us, however, consider the following interpretation. Shifting down ("re-biasing") the zero to the lowest DEC excursion as in Figure 25 and reducing base generation the same amount nets to the levels in Figure 1. Then all of the excursions become INC excursions, of course. Taking this interpretation down to the resource level requires more thought. The capacity of resources that provide both INC and DEC reserves would probably be re-biased to the bottom of their DEC range. The most significant changes, conceptually, is the rule for restoring capacity during a recovery. The rule would endeavor to keep the resource in standby at some mid-capacity value in this case. We return to develop these ideas further in the last section of the paper.

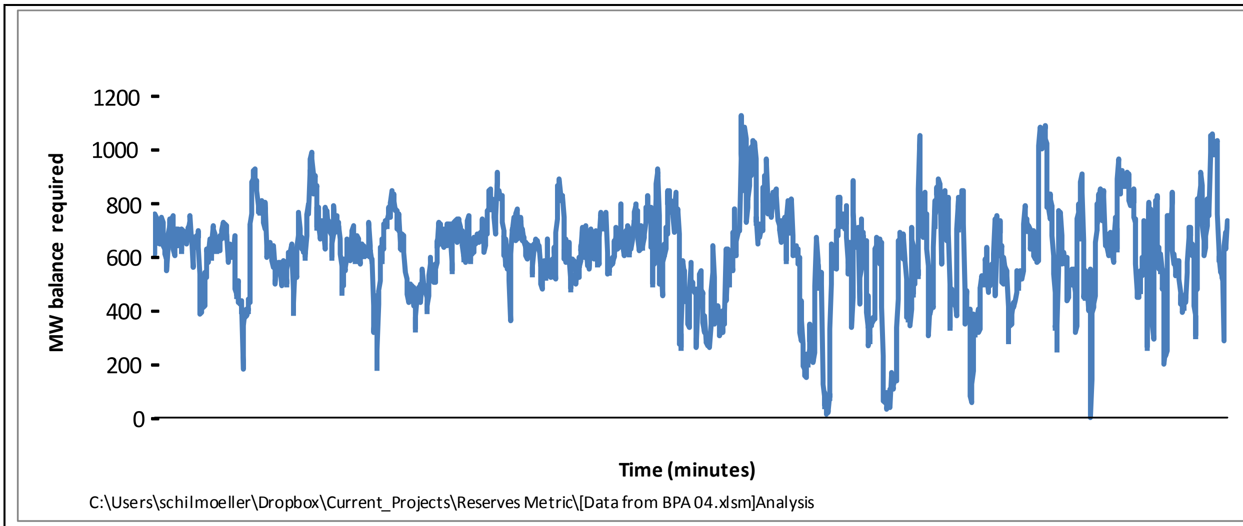

Figure 25: Biasing the data

This completes the discussion of characterizing imbalance reserve requirements. To summarize this section, we have observed the following.

1. Simple statistics on reserve requirements are not sufficient to capture the necessary elements of behavior. The specific chronology of imbalance is critical to this characterization.

2. What happens during resource recovery after response is important to our understanding of requirements.
3. A simple response can be characterized in terms of the RDC of the resources necessary to meet the requirement. This RDC form of resources, in turn, lends itself to the construction of a convex hull for demand.
4. A given ensemble of resources is sufficient to meet reserve requirements if and only if the convex hull of the supply is above the convex hull of the requirements. The proof of this assertion appears in the section **Formalities.**
5. The requirement for any INC excursion is determined by the convex hull on the collection of paths that make up the excursion. Each path has representation as the RDC of resources requisite to meet the path's requirements. The RDC of a path corresponds to the minimal set of sufficient resources (**Proposition 2**). The vertices in the union of path CRDC vertices determine the convex hull of the collective requirement.
6. The characterization of an imbalance requirement is given uniquely by the standard form of the resources defined by the convex hull of the collective requirement.
7. The preceding constructions provide a means to move back and forth between the convex hull and the original description of requirements.

A sound theoretical construct should make it easy to create tools to assist us in addressing problems and questions. A mathematical proof, at a certain level, is indistinguishable from a computer subroutine. Computer simulation provides a means of verifying theory. It also facilitates the exploration of assumptions, the creation of insights, and the extension of ideas to new and unanticipated realms.

The next section outlines a computer application that renders the data in Figure 1 into the forms introduced in this paper. Several extensions and interpretations follow. The application also will facilitate implementation of these ideas when certain key assumptions change. Such potential changes are the subject of a subsequent section, **Our** Assumptions Revisited.

The Requirements Algorithm

A computer program written in VBA/Excel uses the concepts presented in this paper to evaluate the requirement illustrated in Figure 1. It has computed the paths, the convex hull for the requirements, and the associated standard form resources. Taking a brief digression from development of concepts and seeing a few of them in action may help fix the ideas more firmly in the mind's eye.

The data for the computer application appears in Figure 26. The columns, from left to right, are the time, specified in five minute intervals; the balance requirement (MW); a label that specifies whether the record belongs to an INC event or a DEC event; the implied ramp rate in MW per five-minute interval; and a label indicating whether generation is increasing or decreasing. For ease of visualizing the data, cells with increases in generation are pink and those with decreases are green. The entire data set is embedded as an Excel worksheet in an attachment at the end of this document.

The increases and decreases in generation are what define the paths. For INC excursions, response is in pink. The first record in each excursion must be a response. These require special treatment, because typically they will not be full five minutes in duration. For example, the first response listed, 58 MW/5 min. is about 3.1 minutes. This is interpolated from the requirement before the INC event, -22 MW, the requirement after the INC event, 36 MW. This means the corresponding capacity would be 40.0 = 3.1/5.0*58 MW.

If the INC response is not the first response of the excursion, other treatment is necessary. The algorithm must calculate the initial condition by subtracting the recovery energy from the previous response. It must add the loaded unit capacities to the requirement for the current response and resort the collection by ramp rate.

Figure 27, which show the first three paths of many hundreds, will make this evident. The figure contains the first paths encountered in the INC event after they have been sorted by ramp rate. The first path, as mentioned above, starts out with zero or trivial initial loading condition. The records for 5/12/2012 12:05:00 AM and 5/12/2012 12:10:00 AM are the first path data. The units must be sorted by ramp rate and the cumulative duration computed. The report properly indicates 5 minutes of 91MW/5 min ramping followed by 3.1 minutes of 58 MW/5 min ramping, as we just determined. The calculations show the cumulative amount of capacity as a function of cumulative duration, as required for the calculation of the convex hull of this path.

Date/Time	Bal Res	Inc/Dec	Ramp (MW/min)	Ramp INC/DEC
5/12/2012 0:00	-22	DEC	0	NC
5/12/2012 0:05	36	INC	58	INC
5/12/2012 0:10	127	INC	91	INC
5/12/2012 0:15	117	INC	-10	DEC
5/12/2012 0:20	69	INC	-48	DEC
5/12/2012 0:25	58	INC	-11	DEC
5/12/2012 0:30	21	INC	-37	DEC
5/12/2012 0:35	48	INC	27	INC
5/12/2012 0:40	18	INC	-30	DEC
5/12/2012 0:45	71	INC	53	INC
5/12/2012 0:50	36	INC	-35	DEC
5/12/2012 0:55	30	INC	-6	DEC
5/12/2012 1:00	68	INC	38	INC
5/12/2012 1:05	114	INC	46	INC
5/12/2012 1:10	101	INC	-13	DEC
5/12/2012 1:15	8	INC	-93	DEC
5/12/2012 1:20	1	INC	-7	DEC
5/12/2012 1:25	-6	DEC	-7	DEC
5/12/2012 1:30	-22	DEC	-16	DEC
5/12/2012 1:35	-37	DEC	-15	DEC
5/12/2012 1:40	-18	DEC	19	INC
5/12/2012 1:45	-81	DEC	-63	DEC
5/12/2012 1:50	3	INC	84	INC
5/12/2012 1:55	54	INC	51	INC
5/12/2012 2:00	88	INC	34	INC

Figure 26: Input Data

The second path in Figure 27 is more interesting. This response is preceded by a recovery of 106 MW, as shown in Figure 26. This capacity offsets the 91 MW deployment of the fastest resource and all but 20.96 MW of the 58 MW/5 min requirement. There is only five minutes of this 27 MW/5 min requirement (ending 5/12/2012 12:35:00 AM). That requirement is added to the initial condition to create the second path.

The third path (ending 5/12/2012 12:45:00 AM) lasts only 5 minutes and follows a 30 MW/5 min recovery. The 30 MW is removed from the previous path, offsetting the 58 MW/5 min capacity requirements and 9.04 MW of the 27 MW/5 min loading. This leaves 3.33 minutes or 18 MW of the 27 MW/5 min requirement remaining. The full 5 minutes of the 53 MW/5 min response remains, as usual. After sorting by ramp rate, the result is the third path shown.

The paths in Figure 27 are formatted as they are for easy graphing. Each will result in a curve that starts at the origin and traces the convex hull representation for the path.

INC RESPONSE Period	INC RESPONSE MW/Min Sorted	INC RESPONSE Cum MW/Min
0.00	0.00	0.00
5.00	91.00	91.00
8.10	58.00	126.96
0.00	0.00	0.00
1.81	58.00	20.96
6.81	27.00	47.96
0.00	0.00	0.00
5.00	53.00	53.00
8.33	27.00	70.96

Figure 27: Inc responses

The next report illustrating the output of the computer application appears in Figure 28. This is description of the convex hull on all the vertices, that is, the convex hull on the collection of requirement paths. The algorithm for computing the convex hull is very simple, because the problem is two dimensional and only a small number of vertices are involved.

INC RESPONSE Period	INC RESPONSE HULL MW/Min
0.0	0.0
2.9	256.0
8.4	419.0
12.5	449.0
17.5	474.0
27.5	488.0
32.5	493.0

Figure 28: Convex hull

Once the convex hull is computed, the next task is to reduce it to standard form. The application does this in two steps. The first step is to interpret the convex hull using a sequential protocol and sort by ramp rate. This results in the sequence of resources in Figure 29. Then the resources are interpreted in standard form, i.e., simultaneous deployment. The process is simply one of taking differences of ramp rates and summing up the cumulative durations in minutes. This process produces Figure 30.

Note that the capacity provided by all the resources in Figure 30, 493 MW, meets the highest INC encountered in Figure 1. That event occurred at 7:00 AM on 5/15/2012. The total ramp rate provided by all the resources, 87 MW/min, 435 MW over 5 minutes, also matches the largest ramp rate. That event occurred between 10:00 PM and 10:05 PM that same evening. This interval, however, contains a transition from DEC to INC. The INC portion lasted for 2.9 minutes.

The duration of each resource in Figure 30 is obtained by dividing the capacity of each resource by its respective ramp rate. The duration of the first resource in Figure 30, 168.3 MW at 57.2

MW/min, is the shortest, at about 2.9 minutes. If all of the resources are deployed simultaneously, the ramp rate of the system is 87 MW. The combined system is therefore capable of providing 256 MW = 87 MW/min for 2.9 minutes. This corresponds to the first point on the demand hull to the right of the origin in Figure 31. It also corresponds to the requirement between 10:00 PM and 10:05 PM on 5/15/2012 described in the previous paragraph. The other points in Figure 30 can be confirmed in a similar fashion.

INC RESPONSE Periods	INC RESPONSE CUM MW/Min
2.9	87.0
4.1	7.3
5.0	5.0
5.0	1.0
5.5	29.8
10.0	1.4

Figure 29: Convex hull increments

Figure 31 brings the preceding representations together. It illustrates the INC paths with blue lines and convex hull for the INC requirements with the red line with markers. The application used over 1,700 five-minute samples. These samples are for the reserve requirements from 5/12/2012 12:00 A.M. to 5/18/2012 12:20 A.M. The application runs in less than a second.

It should be clear that accommodating revised recovery assumptions only requires modifying a single function that specifies the beginning state of machines at the initial condition of each path. This is purely a bookkeeping issue. The paths are the fundamental building blocks, and it does not matter how the initial conditions for a given path are constructed. Initial conditions have no ramp rate, per se, but only the associated ramp rate of units that have been loaded. Those effects are captured by the definition of a path.

INC RESPONSE RAMP RATE (MW/MIN)	INC RESPONSE RESERVE MW
57.2	168.3
22.5	189.0
3.6	63.0
2.3	29.2
1.0	32.5
0.4	11.0
87.0	493.0

Figure 30: Standard Form

The program can construct the convex hull and standard form representation of the DEC requirements, as well. Using the assumption that DEC recovery ramp rates are at least as great as response ramp rates, the program produced Figure 36 on page 40. The assumption is reconsidered in the context of that figure.

The program also tracks the INC and DEC recovery paths and creates a convex hull for them as well. An assumption implemented in the current algorithm is that units can recover from an excursion at least as fast as they respond. The convex hulls for the recovery paths confirm that, for this sample, recovery ramp rate never exceeded the response ramp rate requirements.

The logic is as an embedded pdf file in an attachment at the end of this paper. For the purpose of creating a concise and clear representation, the error handling subroutines and other features of the code are suppressed. Double-click on the embedded pdf to open and edit the 36-page document.

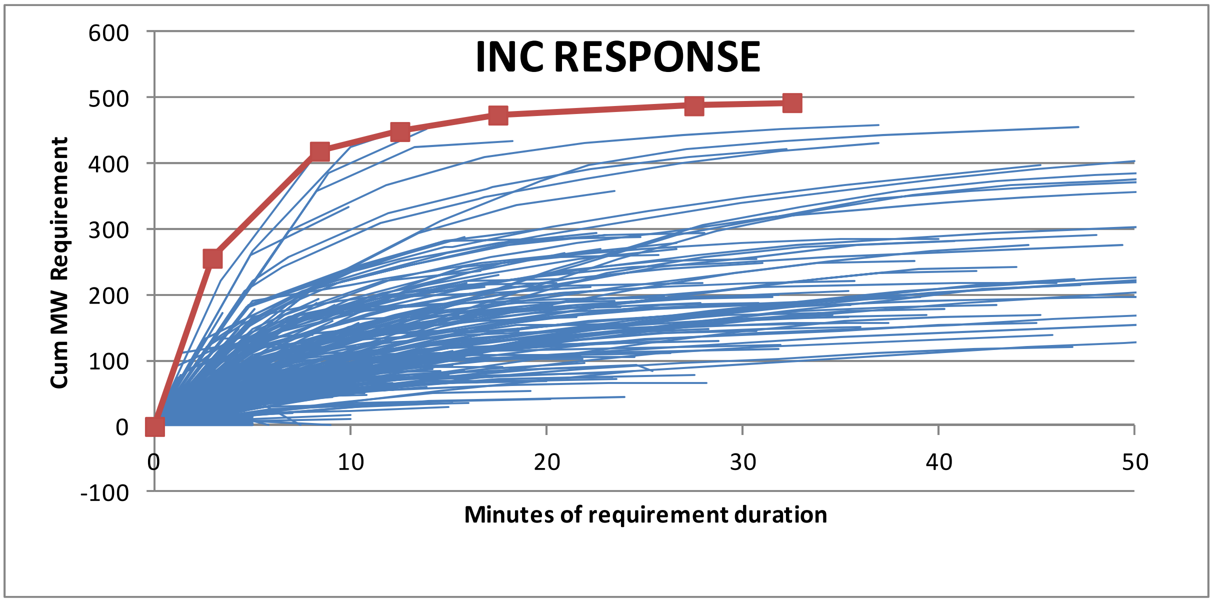

Figure 31: Path profiles and the demand hull

This brief section made the following points.

1. Each path implicitly reflects its initial conditions. The initial conditions are the result of recovering capacity in earlier periods. The section stepped through the calculations for the first three paths in the first INC event.
2. The convex hull for system imbalance demand from the vertices of all paths.
3. The standard form of system demand has attributes, such as total capacity and total ramp rate, which can be compared directly to the chronological input data.
4. While the section presents INC responses requirements, the computer applications simultaneously produces DEC response requirements and both INC and DEC recovery requirements.

Having discussed sufficiency in the context of supply and individual path requirements, we turn to the subject of the sufficiency of system requirement. This is the next section.

Measuring the Sufficiency of Imbalance Reserves

A remarkable fact is that the supply and demand convex hull geometries tell us immediately whether a given supply is adequate. The supply will be adequate if and only if the convex hull of the supply is above the convex hull of the demand along its entirety.[13]

[13] See page 17 and the attachment entitled **Formalities.**

A Metric for Imbalance Reserve Sufficiency

Comparing convex hulls is not a simple task, because they cannot be compared directly, which is to say they cannot be placed in any strict order. That is the purpose of a metric or distance measure.[14]

A metric by definition reduces comparisons to a single number, like percent loss-of-load probability (LOLP), cost, or TailVaR90 risk. A space of convex hulls could be assigned a metric that characterized their "distance" from one another, but it may have little meaning.

An alternative approach is to reduce the difference between a supply hull and a demand hull to a single number is a way that is useful and meaningful. One approach is to measure the extent to which the demand exceeds the supply. Figure 32, for example, shows one such situation. The area of the triangle bounded by the supply hull on the bottom and the demand hull on the top has units of energy (capacity x time). This area of the triangle is the amount this imbalance supply falls short of the ideal (minimal and sufficient) imbalance supply for this requirement.

Note that this area does <u>not</u> measure the shortfall energy or capacity of any particular path or set of paths comprising the system requirement. The vertex of the requirements hull is associated with at least one path, but the path itself typically will not resemble the hull near the vertex. Care must be taken in describing the shortfall. It should be evident, however, that the system will be adequate if and only if this metric value is zero.

The advantage of using the area of the exceedance is that it captures the relative size of the difference. An alternative might be merely to count, say, the number of vertices that are excluded. Counts like this generally do not capture the magnitude of the problem.

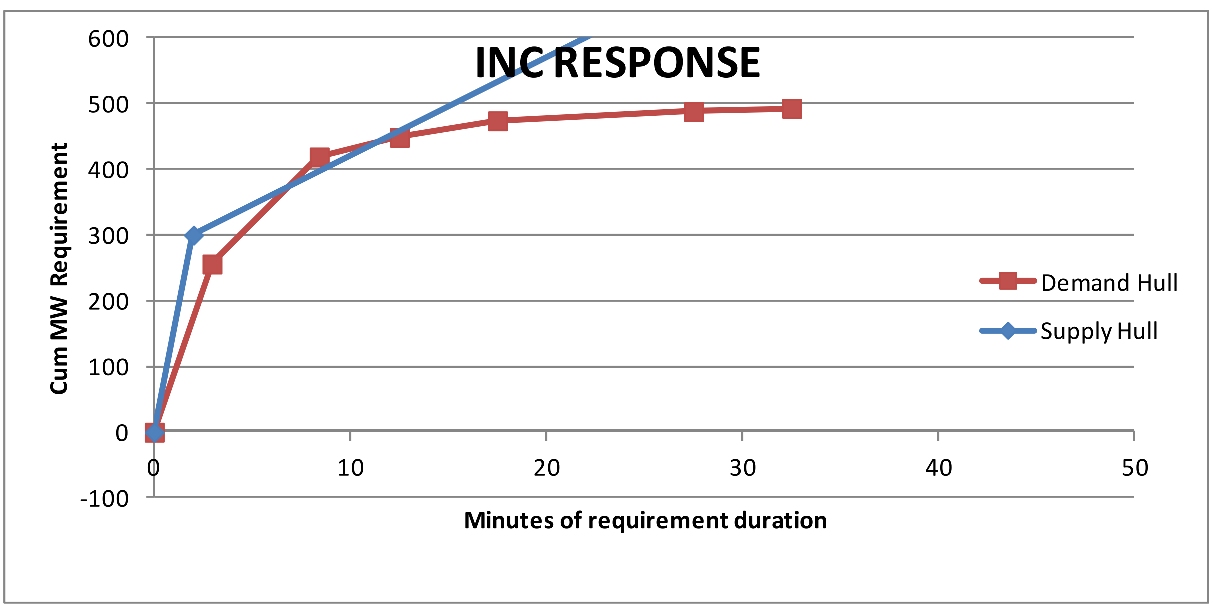

Figure 32: Supply and Demand Hulls

[14] The technical requirements for a metric appear, for example, in http://en.wikipedia.org/wiki/Metric_(mathematics)

The Probabilistic Nature of Supply and Demand

This paper uses a particular sample of the requirement to determine the reserve requirements. However, any sample of requirements, such as that illustrated in Figure 1, is subject to uncertainty. That is, there is an expectation that the magnitudes and ramp rates in the sample are not the largest that might occur. Consequently, the convex hull for demand, and perhaps supply as well, are fundamentally probabilistic objects.

This is an area of research that would benefit from more work. One untested concept is to assign likelihood to ramps of various rates and durations and to calculate the expected energy of excursions relative to the ideal imbalance supply. This approach would distinguish insufficiency events according to their size, as discussed in the previous paragraph.

Isolating and Illustrating Insufficiency

The data from the computer analysis described in the previous section permits us to present a useful kind of resource sufficiency analysis technique. Figure 32 illustrates a system that is <u>not</u> adequate. The supply is 300 MW of a 150 MW/min resource and over 300 MW of 15 MW/min resources. The combination of capacities exceeds the maximum excursion of 493 MW and the maximum ramp rate of 87 MW/min. Nevertheless comparison the supply and demand hulls suggests the supply is inadequate to meet the requirement. The supply graph falls below the demand graph at 8.4 minutes.

Up to this point, the convex hulls have been constructed from the time series of requirements and the specifics of available resources. It is informative to see a case where the development goes in the opposite direction, from the hulls to the time series and resources.

The design of the supply hull in Figure 32 is based entirely on the geometry of the demand hull. The two line segments in the supply hull were simply chosen to fall below the demand hull at a particular vertex of the demand hull. As we will show, it is nevertheless straight-forward to isolate the path (or paths) that could not be satisfied. Moreover, we will reproduce the chronology of resource deployment to show specifically how the resources fail to satisfy the requirement.

To isolate the event associated with shortfall in Figure 32, recall that each vertex in a path and in the demand hull corresponds to a particular requirements path. The capacity requirement associated with the vertex, 419 MW, distinguishes it. Sifting through the paths with this magnitude and matching the path ramp rates on either side of the vertex made it easy to find the event. The requirement is circled in Figure 33, and Figure 34 "zooms in" on the time period so the event can be seen more clearly. The individual five-minute observations over the hour of interest are evident in Figure 34.

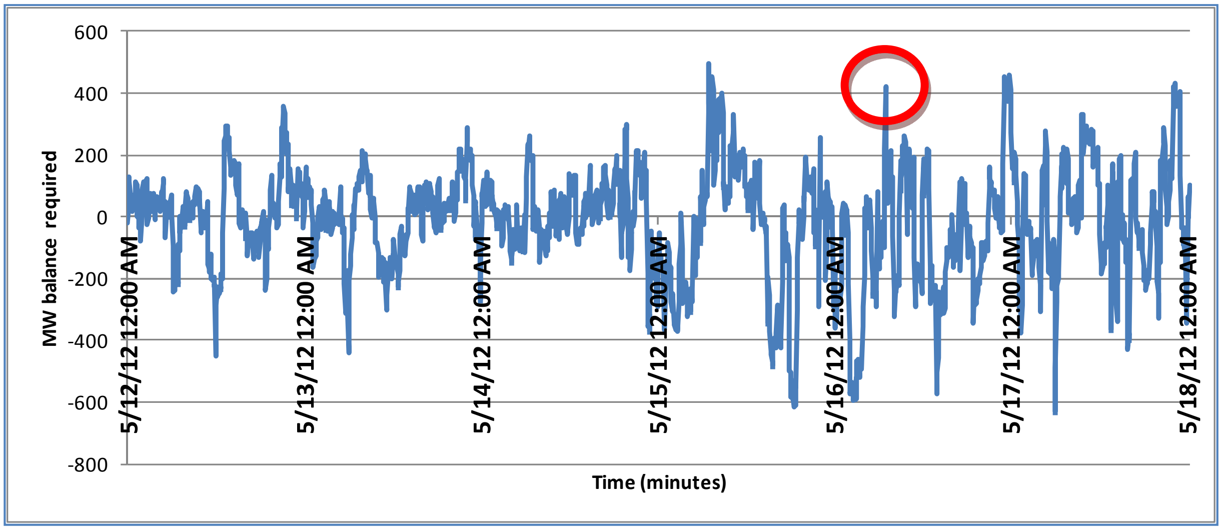

Figure 33: The associated event

The two resources in standard form that correspond to the supply hull are 300 MW of 15 MW/min service and two minutes (270 MW) of 135 MW/min service. The 135 MW/min service is 150 MW/min service less the 15 MW/min service. The lower ramp-rate service is shown in green in Figure 35; the higher ramp-rate service is pink. The shortfall is circled in red.

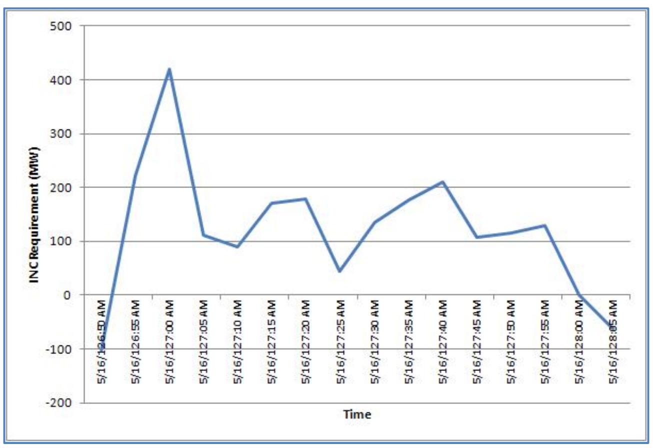

Figure 34: Finding the exception

The shortfall corresponds to a fuel insufficiency in conventional generation planning. That is, the ramp rate of the most responsive unit far exceeds the highest ramp rate observed in any path. Likewise, the total capacity far exceeds the requirement of any path. It is the capacity associated with the unit providing 135 MW/min that is not sufficient. The unit could not continue to ramp up its deployment because the unit ran out of capacity.

In this section, we

1. Used the first few records of input data to reproduce the convex hull of several paths.
2. Examined the sequential and standard form of resources describing the convex hull for requirements.
3. Created a hypothetical supply of imbalance reserves that, in principle, would be insufficient according to the convex hull comparison. We then confirmed its insufficiency by finding the specific insufficiency event in Figure 1.

This paper has relied on several assumptions (see page 3) that we now examine.

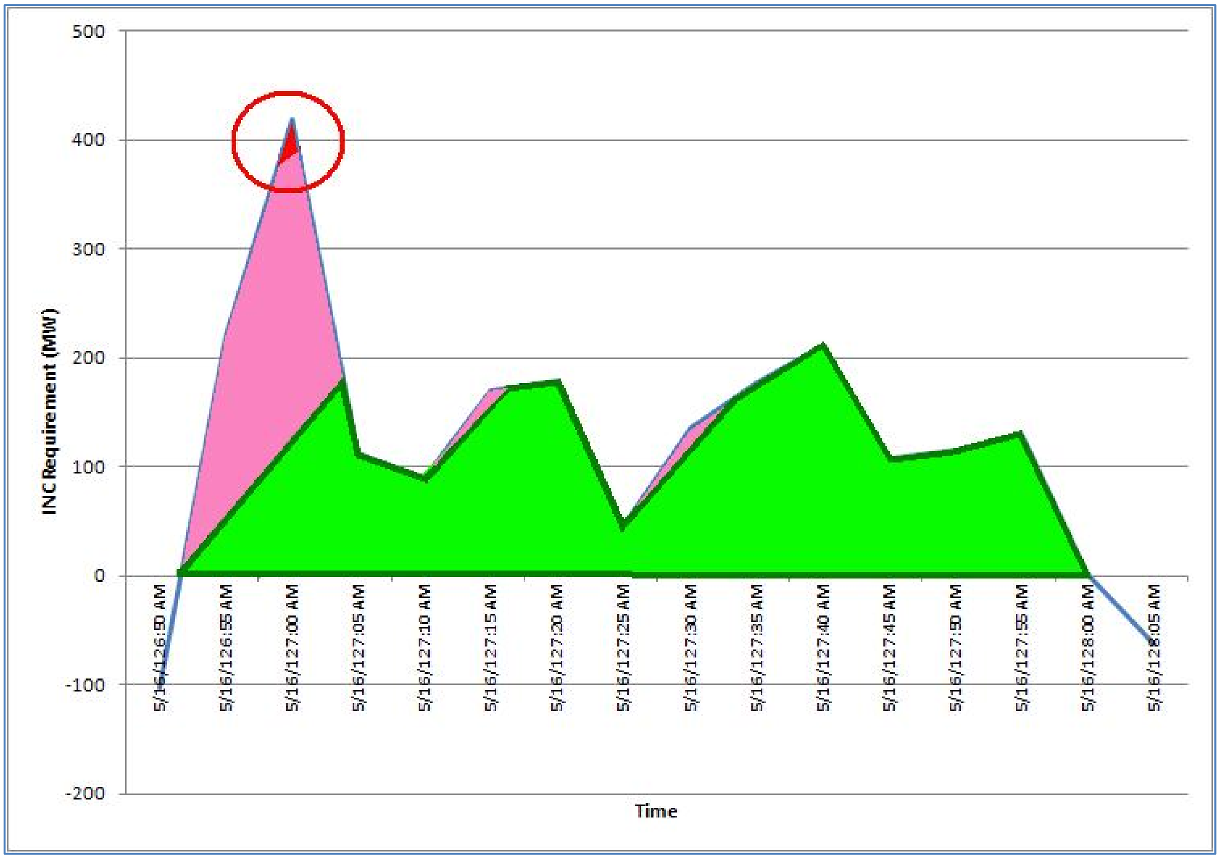

Figure 35: The exception

Our Assumptions Revisited

The weakest of our assumptions fall into several categories. Before discussing these weaker assumptions in detail, the beginning of this section will try to put the value of this paper and its concepts into context. After discussing the assumptions, the section will also outline some outstanding tasks and opportunities.

The Accuracy and Completeness of the Model

While the paper has stated relatively narrow objectives, we should keep in mind the context out of which these objectives were created. We are currently in a time of extensive renewable development and it is not unreasonable to assume this development will continue for the foreseeable future. The Region has struggled with questions about how best to manage the new resources, which are variable. These tend to be wind and to a lesser extent solar generation. That variability has predictable consequences for imbalance reserves.

There has been much good work done around the region on the engineering requirements for systems with variable resources. These typically focus on the likelihood of violating the NERC

CPS2 standards[15], which stipulate constraints on the frequency or likelihood of not meeting capacity requirements.

Little, however, has been forthcoming about the economics of alternatives to manage variable generation. In fact, little has been done to quantify supply and demand of imbalance reserves.[16] Part of the challenge in evaluating and comparing supply and demand has been a lack of consensus on how that can be accomplished. Anecdotal information suggests there is dissatisfaction with analyses using cost minimization techniques such as Lagrangian relaxation. The disappointment stems from two sources. One is the inability of commercial models in general to represent the main source of ancillary services in the Pacific Northwest, hydrogeneration. The other is the overly optimistic results produced by optimization models. The perfect foresight assumption appears to be at the root of this difficulty.

What the region needs at this point, therefore, are tools to guide high-level policy about acquiring and allocating resources to integrate variable generation. Lacking necessary tools and complete solutions, a few, relatively simple conceptual models – such as those presented in this paper – can have significant value.

The value of many models lies in what they ignore. Models like Ohm's law are gross approximations that, nevertheless, are useful for many applications.

Models for higher-level policy do not need the same level of ***accuracy*** as models for operation. There are circumstances, in fact, where it makes sense to intentionally ignore details that distinguish one operator or balancing authority from another. Sometimes the value of characterizing supply and demand is more to provide a relative measure rather than an absolute number. Often what we seek is not so much the most precise representation, but instead the most general representation. The virtue of such a representation can stem from its ***lack of*** detail. This may permits decision-makers to make the broadest possible comparisons.

Not only are precision and accuracy of any representation relevant, but ***the intended use*** of the results, as well. The purpose of the techniques presented in this paper is to make alternative requirements and sources of imbalance reserves comparable. That task is different from answering the question of what amounts of particular machines guarantee a sufficient and robust regional or utility system. The latter, for example, must take up questions about the accuracy of samples. This likely gets into statistical questions. The reliability and availability of the resources to satisfy such requirements would be another consideration.

The point here is that the criteria for evaluating the classification for imbalance reserves across utilities and sub-regions might be very different from those for guiding operations or utility-specific planning. We might call this the IJAMS[17] principle.

[15] See, for reference, item M2 on page 3 of http://www.nerc.com/files/BAL-001-0_1a.pdf
[16] Note, however, that the Oregon Public Utility Commission is driving progress in this area. (Order 12-013, UM 1461, Sec II. D. Integrated Resource Planning Flexible Resources Guidelines, http://apps.puc.state.or.us/orders/2012ords/12-013.pdf)
[17] It's just a model, silly.

General Assumptions

With that preamble, the following are some more general concerns that have been or might be raised about this approach to measuring imbalance reserve supply and demand. The section will then turn to more specific questions about resource recovery and the nature of DEC reserves.

1. The sources of data used for the requirements sample (e.g., Figure 1 on page 3) may not reflect the true requirement for services, because operators do not deploy resources according to a calculated requirement.

   RESPONSE

   This is a valid issue. "…Every movement in Figure 1 may be a sufficient condition for keeping the system whole, but it is not a necessary condition." There are many reasons why this might be the case. Operators may find it counterproductive to react too quickly to requirements. Overly rapid reaction can create or exacerbate balancing problems. Moreover, standards such as the NERC CPS2 require that area control error (ACE) be corrected within a ten-minute period. ACE itself is cumulative value and does not need minute to minute, much less second to second response, per se. It is common for operators to "lean on the system" for seconds to minutes while steering the balance of supply and demand.

   It would seem that the appropriate accommodation, however, is to determine what kinds of adjustments are appropriate and customary, and to make these to the sample data. While the resulting values would likely differ significantly, the over-arching need to characterize requirements certainly remains.

   One consideration in implementing any such accommodation is the scope of such practices. The introduction of this section discussed this in some detail, so it does not need to be repeated here.

   Another consideration is that observations of requirement taken only every five minutes are in effect already filtered. There is no information about shorter excursions that take place. Some thought probably needs to be given to such matters as sample "aliasing"[18] and filtering in order to make sample observations a better match system operation.

2. Some excursions from zero have only a single observation, a DEC or an INC. How is that handled?

   RESPONSE

   For this reason and to obtain reasonable result in general, it is necessary to interpolate the event time at first and last intervals with a zero-crossing. Sometimes the first and last intervals share a single observation. If ***t*** is the cross-over time in minutes after the

[18] Sample aliasing occurs when, by virtue of periodic sampling at discrete points in time, artificial signals are created. See for example, http://en.wikipedia.org/wiki/Aliasing

beginning of the five-minute interval, y_1 is the requirement (MW) at the beginning of the interval and y_2 the requirement at the end of the interval,

$$t = 5y_1 \,/(y_1 - y_2)$$

This is true irrespective of whether the transition is from INC to DEC (y_2 is negative) for from DEC to INC (y_1 is negative). This interpolation for duration is used to arrive at the duration of both parts, INC and DEC, of the interval.

The Recovery Assumptions

There are several potential problems related to the assumption that resources can recover at least as quickly as they respond. The following list contains some issue-specific response. Following the list is a more general discussion of the related challenges and potential solutions.

3. What happens to initial conditions of a path if units cannot recover as quickly as we assume?

4. We assume that resources return to standby when the zero line is crossed. If recovery is constrained, how would that affect this assumption and the conclusions?

5. Should recoveries from DEC excursions be handled the same way as recoveries from INC excursions? Should DEC responses be handled the same way as INC responses?

   The issue here is that INC responses increase generation, while DEC responses decrease generation. There may be fundamental physical reasons why machines would perform differently in these two cases. In fact, an INC response may look a lot more like a DEC recovery, and vice versa.

RESPONSE

Before getting into particulars, here are a few general observations. First, we probably need to solicit feedback from utilities and operators before investing any significant effort in fine-tuning the assumptions or computer tools. This is probably pretty obvious.

Second, note that the units in question are units associated with the convex hull. In the paper, the recovery assumption is invoked to compute the initial condition for a path. The units that would actually recover, however, are not those associated with any given path. They are units associated with the convex hull. Typically, these are much more rapid-responding than is necessary for any particular path. While it is by no means assured, we would expect a unit with more rapid response would also be capable of more rapid recovery. Recall also that it is the recovery of rapid-response units that is most important.

One easy check on the recovery assumption that should be performed is the examination of the convex hull on the recovery events. The basic algorithm currently performs a simple calculation of the capacity recovered during any recovery period and offsets the capacity of loaded units, starting with the highest ramp-rate unit. Previous incarnations of the code, however, computed the paths associated with the recoveries in a fashion identical to what is currently done for responses. Because the ramp-rate of the recoveries is of interest, their "initial conditions" are immaterial. Comparing the convex hulls for recoveries to those of responses, however, indicates whether recovery ramp rates will be a problem. If so, this situation would merit investigation. Perhaps the hull of the combined modes (response and recovery) would comprise a better measure of need. If not, the question of recovery rates may be moot.

Having said this, we expect it is still rather easy to accommodate recovery constraints for particular units, groups of units, or all units. The use of the recovery assumption in this model (conceptual and computational) is to arrive at estimates of initial machine loadings at the beginning of each path. How units recover capacity is therefore not a fundamental issue. It merely affects the amount of capacity that will eventually be required for each kind of ramp rate. The ramp rates themselves should be unaffected by the recovery assumptions. None of the proofs appearing in this paper rely on specific assumptions about the recovery of units.

Estimating initial machine loadings at the beginning of each path can be as simple or complex as necessary to obtain realistic results. The current model simply adds up the total recovery capacity and allocates it back to the deployed resources, most rapid-response unit first. Alternatively, the model could have employed a detailed accounting of unit-specific recovery constraints. The constraints could reflect the state of each machine and its circumstances. Again, none of this has any bearing on the methods presented here.

The preceding paragraph addresses primarily item 3 above. The other questions require more exploration.

Item 4 above deals with the related question of whether excursions can be isolated as the paper assumes, and the resources can always be returned to standby with a zero-crossing. This raises an interesting issue. Is the problem merely that the recovery lags the zero-crossing? It would then seem to be the case that units have more capacity available for subsequent deployment than we had assumed. If a unit is still loaded from an INC response when transitioning to a DEC response, this means there is actually more DEC capacity available than assumed. That is, generation would start out above the standby level. Consequently, the assumptions that they return to standby is conservative in the sense that requirements may be somewhat over-stated. Thus this might not be a fundamental concern.

If the problem is not merely one of recovery lags, it suggests some asymmetry of response and recovery. It is possible, for example, that every time the zero is crossed, the recovery problem grows. If the zero point were somehow floating up in some absolute sense, the units would not be able to recover because zero would be higher each time. This ratcheting need for capacity would appear to lead to runaway situation. It seems that what we are describing, however, is really something like a load-following situation. A source of load-following generation is the solution to this hypothetical situation, and the ratcheting problem would disappear.

Item 5 appears in this list because of its relationship to recovery. However, it is a more complex and extensive issue than what happens during the recovery. DEC events have the potential for behaving in a manner very different from INCs for a number of good reasons. Consequently, it has its own section, which now follows.

The Treatment of DEC Reserves

In the current model, DEC responses are treated like INC responses, and similarly for recoveries. Because different kinds of resources participate in these kinds of reserves, separate bookkeeping appears appropriate. The assumption that all units return to standby condition during a zero-crossing facilitates this disaggregation.

9. If DEC recoveries increase generation, wouldn't they resemble INC responses with regard to ramp rate constraints? If DEC responses decrease generation, wouldn't by resemble INC recoveries? What implications does this have for the model?

10. If different resources participate in DEC responses, how does the model keep track of that?

RESPONSE

The observation is that it takes power – which is fundamentally limited – to increase torque, whereas decreasing torque and output can be interrupted virtually instantaneously. (When this is inadvertent, we call it a forced outage.) This suggests that DEC recovery ramp rates would be constrained and responses ramp rates would be unconstrained.

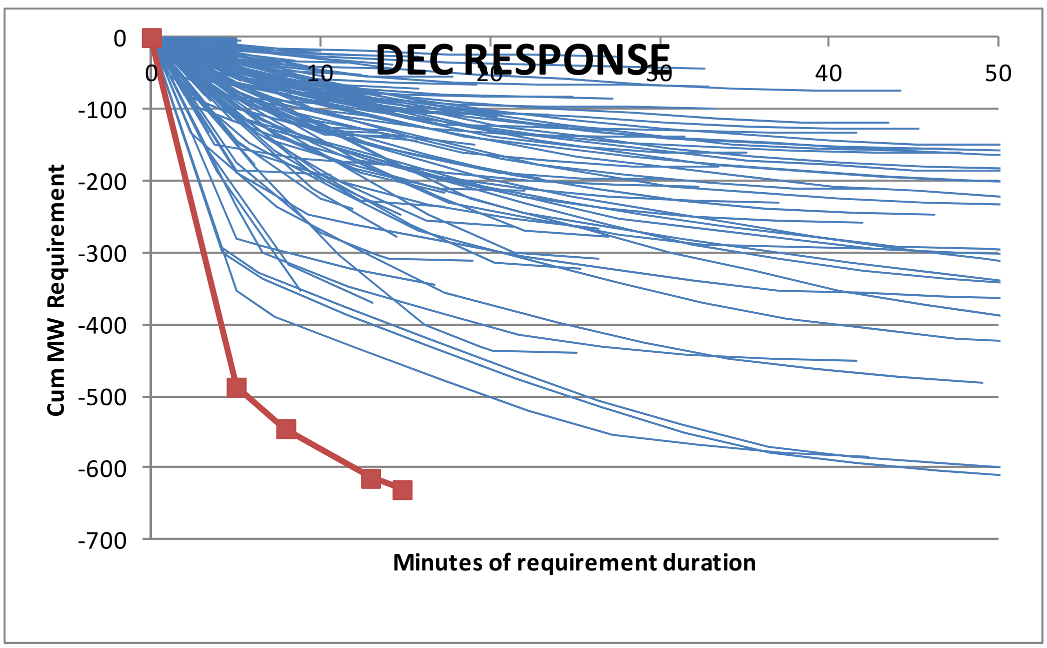

Figure 36: Convex hull for DEC requirements

There is a way to address these concerns, however. It may, in fact, be more natural in the sense that it dispenses with the admittedly arbitrary zero on the imbalance graph.
Consider the interpretation of the imbalance requirements in Figure 1 if the zero line were merely shifted down (“re-biased”) to the lowest DEC excursion. This is illustrated in Figure 25. Then all of the excursions become INC excursions. Taking this interpretation down to the resource level requires more thought. Capacity of resources that provide both INC and DEC reserves must be re-biased to the bottom of their DEC range. The most significant changes, perhaps, is the rule for restoring capacity during a recovery. The rule would endeavor to keep the resource in standby at some mid-capacity value in this case.

The model could deal with the restoration of capacity during the newly redefined recoveries precisely as described earlier. Again, this has no affect on the proofs or fundamental concepts of this paper. The only change would be to the bookkeeping of unloading units to determine the initial conditions of the newly defined responses.

Where the interesting twist arises is in the specification of resources. What happens when a resource is available only during DEC events (or only during INC events)? The answer lies in the capacity of the units. Clearly, DEC units that could not participate in INC events would have capacity constrained to the magnitude of the DEC events. INC resources that could not participate in DEC events would have availability that would be state-related, i.e., related to the state of loading.

This is not as difficult as it sounds. To achieve this, it suffices to parse the requirements into two levels corresponding the DEC and INC levels. Now responses to INC events are informed with additional initial conditions that reflect movement from the bottom of the DEC range. In all other respects, including the treatment of recovery rate constraints, the treatment would be identical. Only the capacity requirements of the INC paths would require interpretation.

Other Potential Areas for Work

While not assumptions, this paper has made observations that raise questions and highlighted opportunities that are not pursued here.

- The analogy of imbalance reserve systems to energy systems captured in the relationships

  $$R = \frac{\partial C}{\partial t} \text{ and } C = \frac{\partial E}{\partial t}$$

  suggests that energy systems with “recoverable” capacity, that is, batteries and other storage technology, are the correct analogy for imbalance reserves. Models, techniques, and insights about those technologies, to the extent they exist, could well be directly applicable to evaluating imbalance reserves.

- The ***standard form*** for system imbalance was the selected RDC to represent the convex hull on the system requirement. It was selected because of its potential for better addressing economic issues. The advantages, however, are not well developed in this paper.

Specifically, Figure 6 on page 9 has two examples of system requirements in RDC form. The convex hull and the rest of the paper do not require that these requirements be reduced any further. It is the integral (sum) of these levels that define the convex hull. It therefore does not matter how we choose to think them. We can consider them as a sequence of ramping resources as we show in Figure 3 (page 7) or as a set of resources deployed simultaneously as in Figure 4. The ***standard form*** adopts the latter. Both representations, however, can be used as unique descriptors for the hull. Both could be used to distinguish between the two systems in Figure 6.

The original motivation of the ***standard form*** was two-fold. First, because of the advantages inherent in simultaneous deployment, it seemed to be a better representation of the direction the industry will eventually move. Second, and more fundamentally, the representation partitions the system requirements (and resources) according to marginal gross value. This would appear to be more useful in addressing the economic questions that are the motivation for this work.

Economic efficiency calls for costs and benefits to be reflected back to their sources by means of suitable prices. The paper argues that increments of ramp rate capability, as illustrated by the increments in the standard form, have increasing gross value per megawatt. Placing requirements and supply into standard form, therefore, should produce a kind of supply curve for the resources. Such a curve could serve as a useful tool in setting those prices and in making the appropriate allocations of resources and costs.

- The paper (section **Measuring the Sufficiency of Imbalance Reserves** on page 31) refers to the outstanding task of developing a statistical description for sufficiency using the concepts in this paper. This should not require original mathematics and would significantly enhance the usefulness of the results.

We are hopeful that others will find these issues interesting and take up the task of answering the questions inherent in these unexplored and unshared areas.

■

Formalities

The objective of this work is to find the imbalance resources that satisfy a given imbalance requirement. Imbalance reserve requirements over some period may be thought of as a time series of observations, consisting of observed average ramp rates over short periods. Many different time series representing reserve requirements, however, potentially have the same resource requirement. Consequently, we seek a classification scheme for reserve requirements that permits us to determine whether a given ensemble of the resources is sufficient to meet a given time series of requirements. This implies that any scheme permit us to directly compare requirements and resources.

Classification is a bijective map between equivalence classes and distinct invariants of each class. In our case, the equivalence classes contain time series (or histories) of imbalance capacity requirement which are equivalent if they are satisfied by the same minimal set of resources, i.e., have resource ensembles with the same ***standard form***.[19] The invariant of an equivalence class is the vector defining the convex hulls or, equivalently, the associated resources in standard form.

Invoking convexity in the following arguments assumes something in addition to the standard assumptions for a convex hull. The CRDC that makes up the "top" of our convex hull has the other significant feature that it is ***non-decreasing***. This is a consequence of the fact that the curve represents an accumulation of positive values. The right-most segment typically is horizontal, which is associated with the trivial (zero) ramp. In the following, we are always working in the context of the top of each hull.

Proposition 1: The capacity and duration of the sum of a set of vectors (hull edges) in a convex hull is the sum of capacity and duration, respectively, of its constituents.

Proposition 2: The ramp rate of the sum of a set of vectors (hull edges) in a convex hull is the average of its constituents' ramp rates with respect to their duration.

Proposition 3: If an ensemble has a convex hull that lies above that of a second ensemble, the first ensemble can meet the same requirements for ramping and capacity as can the second ensemble.

Proposition 4: For strictly increasing imbalance reserve requirements, the order of occurrence of requirements with different ramp rates has no bearing on resource sufficiency.

Proposition 5: The CRDC of a path is equivalent to the minimal set of sufficient resources.

The arguments for the preceding statements appear in the text. (See pages 12, 12, 12, 12, and 22, respectively). Note that **Proposition 3** is not reference below and is in fact a corollary of the **Theorem**.

[19] See the text for a definition of these terms.

With these preliminaries, we present the fundamental result of this paper and its proof. In the following, the ***weak sense*** of "above" and "below" means "at or above" and "at or below." If the term is not qualified, the strict sense of the term is implied. The term ***system imbalance requirement*** is used to represent the requirement over all paths that make up a given requirement history.

Theorem The imbalance supply is sufficient to meet a system imbalance requirement if and only if the convex hull of supply lies above (weak sense) that of the demand hull.

⇨ Contrapositive argument: We show that if a point on the demand hull lies above the supply hull, then the supply is not sufficient. Because both the supply and demand hulls are convex, it suffices to consider a vertex of the demand hull lying above the supply hull. Any vertex of the demand hull must be associated with one or more requirement paths. Consider any particular path containing this vertex. By **Proposition 5**, this supply is insufficient for that path requirement because the supply lies below the minimum sufficient ensemble.

Consequently, a sufficient supply implies the supply hull lies above the demand hull.

🡐 Induction on the edges of an arbitrary path CRDC, from right (least ramp rate) to left, starting with the first (right-most) non-trivial edge of the CRDC. The points in Figure 37 below may help the reader follow the next steps.

Select an arbitrary path CRDC contained in the demand hull. If the supply hull lies weakly above the right-most endpoint of the demand hull, it lies weakly above the right-most endpoint of the path CRDC. The convexity of both the supply hull and the path CRDC guarantees that a diagonal line exists (*A* to *B*) with the same slope as the edge of the path CRDC with endpoints on the supply hull. Place the right endpoint (*B*) of diagonal over the right endpoint of the path CRDC edge. Considered as a vector from the left endpoint (*A*) to the right endpoint (*B*), we know from **Proposition 2** that the diagonal has an average ramp rate equal to that of the path CRDC edge. The path CRDC edge, which represents a constant ramp requirement, can therefore be satisfied by the constituents of the diagonal. As observed in **Proposition 4**, the particular chronology of the requirements has no bearing on the adequacy of this supply.

We will refer to the capacity of the supply hull between the endpoints of the diagonal (*A, E, F, B*) as the "deployed capacity." The difference (*A* to *C*) between the diagonal vector (*A* to *B*) and the portion of the diagonal vector over the path CRDC edge (*C* to *B*) represents the remaining capacity of the deployed resources after meeting the path CRDC's first edge. The effective supply hull available to meet subsequent requirements becomes the curve (0, D, A, C). Because the path CRDC is convex, the average ramp rate of the remaining deployed capacity will be less than the ramp rate of the next path CRDC edge, which lies to the left of the subject path CRDC edge. The diagonal lies weakly above the path CRDC at a point (*C*) above the left endpoint of the path CRDC edge. Thus we satisfy the initial conditions for the next step of this induction. We can continue with the next edge of the path CRDC in this fashion from right to left until all the edges of the path CRDC edges are satisfied.

Because the choice of path CRDC was arbitrary, the statement is established for all path CRDCs. Consequently, if the supply hull lies above the demand hull, the supply is sufficient.

QED

Note that the proof of sufficiency in this theorem is constructive. Constructive proofs are often quite powerful. In this case, we describe which resources will be deployed not only for any individual path in the collection represented by the demand hull, but for the demand hull itself. (The demand hull observes the same convexity requirement as the path CRDC.) This provides a near-chronological rendition of the deployment of resources in any path. This proof therefore enables analysts to isolate the particular paths and events that give rise to any resources insufficiency. There are many corollaries that we could write. Below are examples.

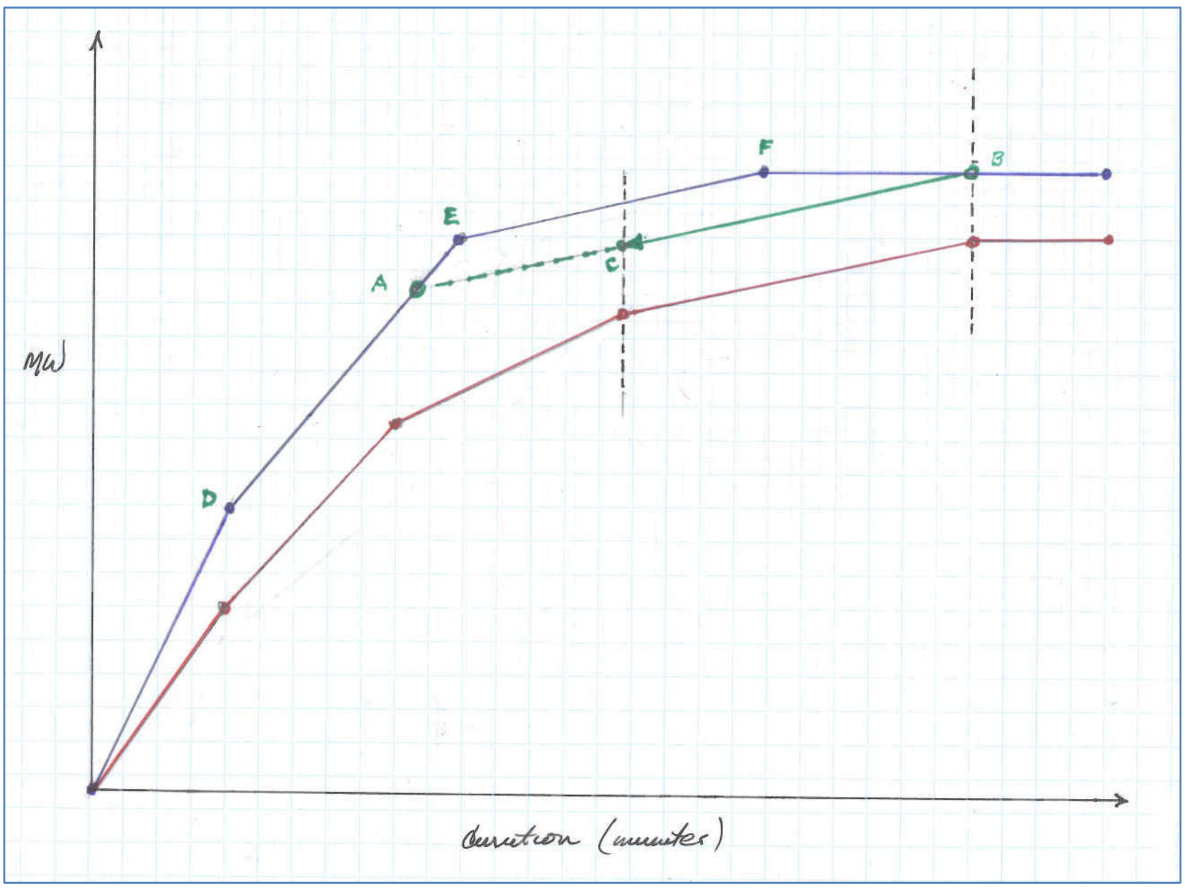

Figure 37: Induction proof, step 1

Corollary If a supply hull dominates a second supply hull, the first supply hull can provide all the imbalance services as the second.

Corollary The convex hull of the vertices of a set of requirement paths determines the minimal, sufficient set (standard form) of resources that satisfies all the paths.

Sufficiency: Each requirement path lies below (weak sense) the supply hull defined by the hull on the union of the path vertices. By the **Theorem**, the supply must be sufficient for each path.

Minimality: By the definition of the convex hull, it is minimal. Any smaller hull would not contain some vertex and, by the proof of necessity in the **Theorem**, would not be sufficient.

QED

■

Data Used for Creating the Examples in This Document

Below is a worksheet of the data presented in Figure 1 and used in the examples presented in this document. It is embedded Excel 2007 and can be opened by right-clicking on the image below. Opening might be more manageable than editing (the default invoked by double-clicking) or converting the object.

Source: http://transmission.bpa.gov/Business/Operations/Wind/reserves.txt						
BPA Balancing Reserves Deployed						
at 5-minute intervals, last 7 days						
Dates: 12May2012 - 19May2012 (last updated 18May2012 00:21:50) Pacific Time						
Based on 5-min MW readings from the BPA SCADA system for points 108044, 108						
BPA/Technical Operations (TOT-OpInfo@bpa.gov				row		
			row			
min (dec):	-630		1528	min ramp:	-487	
max (inc):	493.0		1144	max ramp:	435.0	
Date/Time	Bal Res	INC Max	DEC Max	Inc/Dec	Ramp (MW/min)	Ramp INC/DEC
5/12/2012 0:00	-22	726	-923	DEC	**0**	NC
5/12/2012 0:05	36	726	-923	INC	58	INC
5/12/2012 0:10	127	726	-923	INC	91	INC
5/12/2012 0:15	117	726	-923	INC	-10	DEC
5/12/2012 0:20	69	726	-923	INC	-48	DEC
5/12/2012 0:25	58	726	-923	INC	-11	DEC
5/12/2012 0:30	21	726	-923	INC	-37	DEC
5/12/2012 0:35	48	726	-923	INC	27	INC
5/12/2012 0:40	18	726	-923	INC	-30	DEC

VBA/Excel Code for Creating the Paths, Convex Hull, and Standard Forms

The following is an embedded pdf with the principal class module for performing the calculations presented here. It is approximately 36 pages long. It might be most manageable if opened rather than edited or converted. (These are mouse right-click options.)

```
Sheet1.cls                                          Sheet1                          Printed:  12-Jun-12 00:06
File Date:  12-Jun-12 00:04                     <Declarations>

VERSION 1.0 CLASS
BEGIN
    MultiUse = -1  'True
END
Attribute VB_Name = "Sheet1"
Attribute VB_GlobalNameSpace = False
Attribute VB_Creatable = False
Attribute VB_PredeclaredId = True
Attribute VB_Exposed = True
Option Explicit
'Required NWPCC ErrorHandling Library


Const lINC_RESPONSE = 0
Const lINC_RECOVERY = 1
Const lDEC_RECOVERY = 2
Const lDEC_RESPONSE = 3
'Used for vInput
Const lMINUTES As Long = 0
Const lORIGINALROW As Long = 1
Const lRAMPRATE As Long = 2
'Used for vResults
Const iVALUE As Integer = 0
Const iORIGINAL_ORDER As Integer = 1
Const iORIGINAL_ROW As Integer = 2
Const iMINUTE As Integer = 3

Const lCOL_INPUT_BALANCE_MW = 2
Const lCOL_INPUT_EXCURSION_TYPE = 5     'INC or DEC
Const lCOL_INPUT_RAMP_MW_PER_PERIOD = 6
Const lCOL_INPUT_RAMP_DIRECTION = 7 _
  'Response or recovery, but this indicates whether generation is increasing (response for
'for INCs and recovery for DECS) or decreaing (recovery for INCs and response for DECs)

Const lCOL_INC_RESPONSE_PER = 1
Const lCOL_INC_RESPONSE_SORTED = 2
Const lCOL_INC_RESPONSE_CUM = 3
Const lCOL_INC_RECOVERY_PER = 4
Const lCOL_INC_RECOVERY_SORTED = 5
Const lCOL_INC_RECOVERY_CUM = 6

Const lCOL_DEC_RESPONSE_PER = 7
Const lCOL_DEC_RESPONSE_SORTED = 8
Michael Schilmoeller                          Page  1                    NW Power and Conservation Council
```

Table of Contents