Uniform Price Auctions and Optimal Power Flow

Ray D. Zimmerman

February 12, 2010*

MATPOWER *Technical Note 1*

*First draft December 20, 2007

Contents

1 Notation

i	index of generator or dispatchable load, also used to index the corresponding offers or bids
$k(i)$	bus number corresponding to location of generator or dispatchable load i
p_G^i, p_D^i	real power produced by generator i, consumed by dispatchable load i, respectively
λ_p^k	nodal prices for real power at bus k as computed by the OPF
o_p^i, b_p^i	price of single-block real power offer or bid i, corresponding to generator or dispatchable load i, respectively
$o_p^{i,\mathrm{LA}}, b_p^{i,\mathrm{LA}}$	price of last (partially or fully) accepted block for real power offer and bid i, defined to be $-\infty$ and ∞, respectively, if all blocks are rejected
$o_p^{i,\mathrm{FR}}, b_p^{i,\mathrm{FR}}$	price of first fully rejected block for real power offer and bid i, defined to be ∞ and $-\infty$, respectively, if no blocks are fully rejected
q	replace all instances of p above with q for equivalent reactive power quantities

2 Uniform Price Auctions

2.1 Single-Sided Auctions

Consider a single-sided uniform price auction for the supply of an inelastic demand. Offers are ranked in increasing order and accepted beginning with the least expensive and continuing until the demand is satisfied. The uniform price is then set equal to either the *last accepted offer* (LAO) or the *first rejected offer* (FRO). For a continuous quantity commodity with block offers, except for the special case where the quantity clears on a block boundary, there is a marginal block which is partially accepted, as shown in Figure 1. This block is taken as the last accepted block and its price, the *marginal price*, corresponds to the incremental cost of additional demand. In this case, the first rejected price is taken to be the price of the first fully rejected block. Whether the uniform price is set by the LAO or the FRO, it will be greater than or equal to all accepted offers and therefore acceptable to all of the selected suppliers.

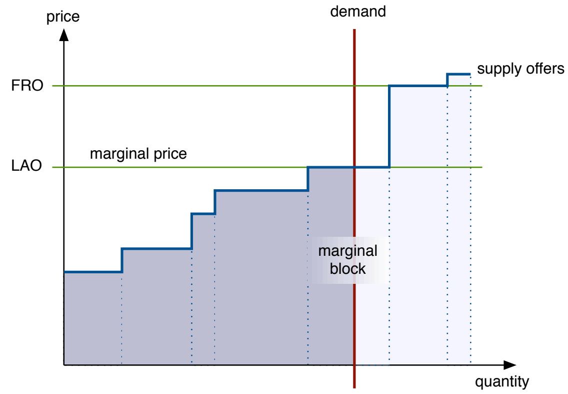

Figure 1: Single-sided Seller Auctions

Similarly, in a single-sided auction with bids to procure an inelastic supply, bids are ranked in decreasing order and accepted beginning with the highest and continuing until the supply is exhausted. Once again there are two pricing options, *last accepted bid* (LAB) and *first rejected bid* (FRB), as shown in Figure 2. In this case, the marginal price will be equal to the last accepted bid. Either case results in a uniform price that is less than or equal to all accepted bids and therefore acceptable to all of the selected buyers.

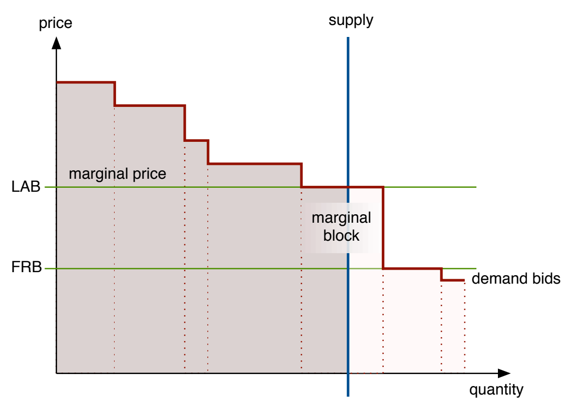

Figure 2: Single-sided Buyer Auctions

2.2 Two-Sided Auctions

In two-sided auctions, offers and bids are ranked as above, and an equal quantity of each is accepted, beginning with the highest bids and lowest offers, until supply or demand is exhausted or the offer price exceeds the bid price. Again, except for the special case where the quantity clears on a block boundary, there will be a partially accepted block. This marginal block could be either a bid or an offer. In general, the last accepted bid will be greater than the last accepted offer, leaving a *bid-offer gap*.

A uniform price set to anything within this bid-offer gap will be satisfactory for all buyers and sellers. That is it will be less than or equal to all accepted bids and greater than or equal to all accepted offers. While the FRO and FRB are not guaranteed to yield prices in this range, and are therefore not suitable for two-sided auctions, there are a number of other valid pricing rules in addition to LAO and LAB. For example, the *marginal* or *first price* rule chooses LAO or LAB, depending which one is the marginal unit. A *split-the-difference* rule uses the mid-point of the bid-offer gap by taking the average of LAO and LAB. And, somewhat analogous to the first rejected single-sided auctions, a *second price* rule could be defined as follows. If the marginal unit is a bid, set the price to the maximum of the LAO or FRB. If it is an offer, set the price to the minimum of the LAB or FRO.

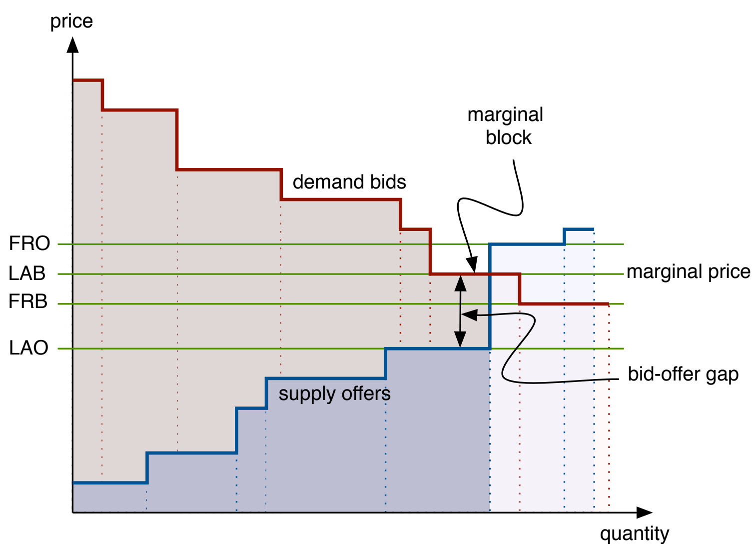

Figure 3: Double-sided Auctions

3 Auctions using DC Optimal Power Flow

This flexibility in the choice of uniform price can be generalized to auctions that are solved by "smart markets" that take into account externalities, such as an optimal power flow (OPF) used to solve a power market subject to network constraints.

When a DC OPF is used to solve a real power market with block offers/bids, it produces a set of *nodal prices* λ_p^k for real power. These prices correspond to the incremental cost (benefit) of additional demand (supply) at each node. In an uncongested transmission network, these prices will be uniform and they will be equal to the price of the marginal or last accepted unit. If the market is a one-sided auction with inelastic demand, the λ_p^k can be used directly as the uniform price, yielding LAO pricing. The price of the first rejected block could also be used directly to set a uniform price, just as in a standard auction. Similarly, for two-sided markets, using the λ_p^k directly as the uniform price corresponds to the *first-price* auction described above, where the price is determined by the marginal unit. Alternatively, the uniform price can be set to the last accepted offer, last accepted bid or anything else in between.

With congestion, however the nodal prices λ_p^k are no longer uniform, but vary based on location. The differences betweeen these nodal prices represent the cost of transmission between locations. The prices are still determined by the marginal offer(s)/bid(s), but since they vary by location, finding the equivalent of a first rejected price or of a bid-offer gap at each node is no longer straightforward.

The nodal prices represent the marginal value of power at a location and can be used to compute exchange rates for normalizing location specific prices, offers or bids to a reference location r. Since the marginal value of a single unit of real power at bus k is λ_p^k and the same unit has a value of λ_p^r at bus r, any price, offer or bid at bus k can be converted to the equivalent at bus r by multiplying by an exchange rate of λ_p^r/λ_p^k.

Given the OPF solution which specifies the nodal prices along with which offers and bids are accepted and which are rejected, these offers and bids can be normalized to a reference location r, then rank ordered as in a standard auction and the normalized uniform price can be chosen directly according to the desired pricing rule (first-price, LAO, LAB, FRO, FRB, split-the-difference, second-price). However, this uniform price only applies at the reference bus r. The equivalent price at bus k is found by multiplying by the appropriate exchange rate λ_p^k/λ_p^r. Notice that this results in a simple scaling of all nodal prices by some factor χ. For example, if the normalized uniform price at bus r is u_p^r, then the uniform price at each bus k is simply

$$u_p^k = \left(\frac{u_p^r}{\lambda_p^r}\right)\lambda_p^k = \chi\lambda_p^k \tag{1}$$

With standard uniform price auctions each of the pricing rules mentioned results in prices that are acceptable to all participants. Similarly, a scale factor χ, appropriately determined by the method above, will result in nodal prices that are satisfactory for all accepted units.[1] It is also important to note that these prices are still consistent with the solution found by the OPF. This is equivalent to multiplying the objective function by χ, in effect solving the same problem with the cost expressed in a different currency with an exchange rate of χ relative to the original currency. The optimal allocations are not affected by such scaling.

The exchange rate corresponding to the last accepted offer χ^{LAO} is equal to the scale factor obtained by scaling prices[2] down until some accepted offer becomes equal to its nodal price. For last accepted bid, it corresponds to scaling prices up until some accepted bid is equal to its nodal price. In the case of first rejected offer (bid), prices are scaled up (down) until a rejected offer (bid) is equal to its nodal price.

Since the nodal prices produced by the OPF marginal prices corresponding to the first-price rule, the corresponding exchange rate χ^{first} is always equal to 1. The exchange rates corresponding to the other uniform pricing rules mentioned can be found by the following formulas:

$$\chi^{\mathrm{LAO}} = \max_i \frac{o_p^{i,\mathrm{LA}}}{\lambda_p^{k(i)}} \tag{2}$$

$$\chi^{\mathrm{LAB}} = \min_i \frac{b_p^{i,\mathrm{LA}}}{\lambda_p^{k(i)}} \tag{3}$$

$$\chi^{\mathrm{FRO}} = \min_i \frac{o_p^{i,\mathrm{FR}}}{\lambda_p^{k(i)}} \tag{4}$$

$$\chi^{\mathrm{FRB}} = \max_i \frac{b_p^{i,\mathrm{FR}}}{\lambda_p^{k(i)}} \tag{5}$$

$$\chi^{\mathrm{split}} = \frac{\chi^{\mathrm{LAO}} + \chi^{\mathrm{LAB}}}{2} \tag{6}$$

$$\chi^{\mathrm{first}} = 1 \tag{7}$$

[1] In the case of FRO and FRB, this is only true for single-sided auctions, so they are not appropriate for two-sided auctions.

[2] Here prices are assumed to be positive.

$$\chi^{\text{second}} = \begin{cases} \min(\chi^{\text{FRO}}, \chi^{\text{LAB}}) & \text{if an offer is marginal} \\ \max(\chi^{\text{FRB}}, \chi^{\text{LAO}}) & \text{if a bid is marginal} \end{cases} \quad (8)$$

The fact that $\chi^{\text{first}} = 1$ implies that at least one of χ^{LAO} and χ^{LAB} will always be equal to 1 as well. If there is at least one marginal offer, then $\chi^{\text{LAO}} = 1$ and if there is at least one marginal bid $\chi^{\text{LAB}} = 1$. In many cases, there will be multiple marginal units. If both an offer and a bid are marginal, there is no bid-offer gap and $\chi^{\text{LAO}} = \chi^{\text{LAB}} = 1$. The table below summarizes the characteristics of the various exchange rates according to which kind of marginal blocks are present in the OPF solution.

	marginal offer	marginal bid	both marginal
χ^{LAO}	1	≤ 1	1
χ^{LAB}	≥ 1	1	1
χ^{FRO}	≥ 1		≥ 1
χ^{FRB}		≤ 1	≤ 1
χ^{split}	≥ 1	≤ 1	1
χ^{first}	1	1	1
χ^{second}	≥ 1	≤ 1	1

4 Extension to AC OPF and Reactive Power

Similarly, an AC OPF can be used to solve a "smart market" for real power. In the absence of dispatchable loads with constant power factor constraints, the exchange rates for the various pricing rules can be calculated in exactly the same way as for the DC OPF. Even without congestion, the AC OPF typically produces non-uniform nodal prices due to losses. Congestion in the form of binding line flow or voltage limits only increases the nodal price differences.

However, as soon as constant power factor dispatchable loads are included, new complications are introduced due to the coupling of the real and reactive power. Even if the objective function deals only with real power,[3] the AC OPF still yields a set of nodal prices for reactive power λ_q^k as well as for real power λ_p^k.

4.1 Bundling Real and Reactive Power

If a dispatchable load i is modeled with a constant power factor, so that the ratio of reactive to real demand is a constant

[3] That is, there is no market for reactive power and reactive power supply is "free".

$$\kappa_i = \frac{q_D^i}{p_D^i} \tag{9}$$

then the real and reactive power consumption of this load can be thought of as a single "bundled" commodity. The value of this commodity can be expressed on a per MW (or per MVAr) basis. If this load is located at bus k and λ_p^k and λ_q^k are the prices of real and reactive power, respectively, then the value v of the bundled power is

$$\begin{aligned} v &= \lambda_p^k \cdot p_D^i + \lambda_q^k \cdot q_D^i && (10) \\ &= \lambda_p^k \cdot p_D^i + \lambda_q^k \cdot \kappa_i \, p_D^i && (11) \\ &= (\lambda_p^k + \kappa_i \lambda_q^k) \, p_D^i && (12) \end{aligned}$$

In other words the per MW price of the bundled commodity is $\lambda_p^k + \kappa_i \lambda_q^k$. Similarly, the per MVAr price is $\lambda_p^k / \kappa_i + \lambda_q^k$.

This implicit bundling implies that the following single-block bids are equivalent for such a dispatchable load, where the third case defines a continuum of equivalent bids that lie on the line connecting the first two.

	b_p^i	b_q^i
1	α	0
2	0	α/κ_i
3	$\gamma\alpha$	$(1-\gamma)\alpha/\kappa_i$

Each indicates a willingness to pay up to the same amount for a given quantity of the bundled product, namely α per MW consumed.

4.2 Coupled Auctions

It is important to take this bundling into account when finding the last accepted and first rejected bids and the bid-offer gap. Consider the first case from the table above, where $b_p^i = \alpha$ and $b_q^i = 0$. Ordinarily, a bid is accepted if it is higher than the price. But due to the bundling, the bid will be rejected if it is in the range $\lambda_p^k < \alpha < \lambda_p^k + \kappa_i \lambda_q^k$. In other words, in order for the bid to be accepted α must exceed, not only the real power price, but the bundled per MW price. More generally, the bid will be accepted if the bundled bid exceeds the bundled price, it will be marginal when the bundled bid is equal to the bundled price, and it will be rejected when the bundled bid is less than the bundled price. Note that the "bundling" is determined by κ_i and is therefore specific to the individual load. Loads with different power factors at the same bus will have different bundled prices.

Consequently, the exchange rates in (3) and (5) must be modified to use the bundled bids and bundled nodal prices,

$$\chi^{\mathrm{LAB}} = \min_i \frac{b_{pq}^{i,\mathrm{LA}}}{\lambda_{pq}^{k(i)}} \tag{13}$$

$$\chi^{\mathrm{FRB}} = \max_i \frac{b_{pq}^{i,\mathrm{FR}}}{\lambda_{pq}^{k(i)}} \tag{14}$$

where the bundled price is simply $\lambda_{pq}^{k(i)} = \lambda_p^{k(i)} + \kappa_i \lambda_q^{k(i)}$. For LAB, the bundled bid is likewise straightforward. But for FRB, it gets more complicated due to the subtlety of defining the first rejected block of a bundled commodity. If $p_D^{i,\mathrm{FR}}$ and $q_D^{i,\mathrm{FR}}$ are used to denote the quantities of real and reactive power at which the respective first rejected bid block begins, then the last accepted and first rejected bundled bid prices used in (13) and (14) can be defined as

$$b_{pq}^{i,\mathrm{LA}} = b_p^{i,\mathrm{LA}} + \kappa_i b_q^{i,\mathrm{LA}} \tag{15}$$

$$b_{pq}^{i,\mathrm{FR}} = \begin{cases} b_p^{i,\mathrm{FR}} + \kappa_i b_q^{i,\mathrm{LA}} & \text{if } \kappa_i p_D^{i,\mathrm{FR}} \leq q_D^{i,\mathrm{FR}} \\ b_p^{i,\mathrm{FR}} + \kappa_i b_q^{i,\mathrm{FR}} & \text{if } \kappa_i p_D^{i,\mathrm{FR}} = q_D^{i,\mathrm{FR}} \\ b_p^{i,\mathrm{LA}} + \kappa_i b_q^{i,\mathrm{FR}} & \text{if } \kappa_i p_D^{i,\mathrm{FR}} \geq q_D^{i,\mathrm{FR}} \end{cases} \tag{16}$$

For a load with negative VAr consumption, the bundled per MW bid incorporates any corresponding offer for reactive power. In this case, let $q_D^{i,\mathrm{FR}}$ denote the quantity of reactive power demand at which the first rejected offer block begins. Then both κ_i and $q_D^{i,\mathrm{FR}}$ are negative, and

$$b_{pq}^{i,\mathrm{LA}} = b_p^{i,\mathrm{LA}} + \kappa_i o_q^{i,\mathrm{LA}} \tag{17}$$

$$b_{pq}^{i,\mathrm{FR}} = \begin{cases} b_p^{i,\mathrm{FR}} + \kappa_i o_q^{i,\mathrm{LA}} & \text{if } -\kappa_i p_D^{i,\mathrm{FR}} \leq -q_D^{i,\mathrm{FR}} \\ b_p^{i,\mathrm{FR}} + \kappa_i o_q^{i,\mathrm{FR}} & \text{if } -\kappa_i p_D^{i,\mathrm{FR}} = -q_D^{i,\mathrm{FR}} \\ b_p^{i,\mathrm{LA}} + \kappa_i o_q^{i,\mathrm{FR}} & \text{if } -\kappa_i p_D^{i,\mathrm{FR}} \geq -q_D^{i,\mathrm{FR}} \end{cases} \tag{18}$$

An important implication of this coupling of real and reactive in the determination of these exchange rates is that it is not valid to treat real and reactive power markets as separate, uncoupled auctions, using different pricing rules for real power and reactive power. In other words, choosing LAO for real power and FRO for reactive power is not legitimate when there are constant power factor dispatchable loads in the system. Furthermore, since the real and reactive power are coupled by the network itself, not just by these loads, it is not consistent to use one exchange rate to adjust real power prices and another to adjust reactive power prices. The resulting prices would be inconsistent with the original OPF solution.

In a general case, an AC OPF based auction may include bids and offers for both real and reactive power. In the MATPOWER implementation, a generator may have an offer for real power and an independent offer and/or bid for reactive power, depending upon the reactive power output capability of the unit. A dispatchable load may have a bid for real power coupled with either a bid or an offer for reactive power, depending on the fixed power factor. Any scaling of prices to implement a particular pricing rule must take into account both real and reactive offers of generators, possible uncoupled reactive bids from generators and the coupled nature of the bids from the dispatchable loads. Combining all of these factors yields the following generalized formulas for the exchange rates first presented in equations (2) through (5).

$$\chi^{\text{LAO}} = \max\left(\max_i \frac{o_p^{i,\text{LA}}}{\lambda_p^{k(i)}}, \max_i \frac{o_q^{i,\text{LA}}}{\lambda_q^{k(i)}}\right) \tag{19}$$

$$\chi^{\text{LAB}} = \min\left(\min_i \frac{b_{pq}^{i,\text{LA}}}{\lambda_{pq}^{k(i)}}, \min_i \frac{b_q^{i,\text{LA}}}{\lambda_q^{k(i)}}\right) \tag{20}$$

$$\chi^{\text{FRO}} = \min\left(\min_i \frac{o_p^{i,\text{FR}}}{\lambda_p^{k(i)}}, \min_i \frac{o_q^{i,\text{FR}}}{\lambda_q^{k(i)}}\right) \tag{21}$$

$$\chi^{\text{FRB}} = \max\left(\max_i \frac{b_{pq}^{i,\text{FR}}}{\lambda_{pq}^{k(i)}}, \max_i \frac{b_q^{i,\text{FR}}}{\lambda_q^{k(i)}}\right) \tag{22}$$

5 Practical Considerations

5.1 Numerical Issues

The exchange rates χ are all derived from ratios of bids/offers and nodal prices. In particular, of the set of ratios considered, the one closest to 1 sets the exchange rate. Clearly, numerical issues arise as the numerator and denominator of this defining term become small. Any numerical errors in the prices λ, arising from the finite precision of the OPF solution, are magnified when dividing by them if they are small.

Suppose the defining term is one with $|\lambda| < \delta$ for some small threshold δ. If we label this term a, the price λ_a and the corresponding exchange rate, χ_a. Now, if this term is eliminated, some other term b will define the new exchange rate χ_b. The only possible problem which could arise is a slight violation of the bid or offer corresponding to term a. At most, this violation will be on the order of $|\chi_b - \chi_a|\delta$. This will typically also be small if δ is small.

So for some small value of δ, it is safe to eliminate any terms with $|\lambda| < \delta$. For the extreme case where this eliminates all of the terms, essentially none of the prices are accurate enough to generate a meaningful exchange rate, so the exchange rate χ can be explicitly set to 1.

5.2 Binding Minimum Generation Limits

If generator i has a non-zero minimum generation limit $P_{\min}$, it is possible for the OPF solution to result in a nodal price $\lambda_p^{k(i)}$ that is less than the offer o_p^i. In this case, the shadow price on the minimum generation limit $\mu^i_{P_{\min}}$ will be equal to the difference.

$$\mu^i_{P_{\min}} = o_p^i - \lambda_p^{k(i)} \tag{23}$$

Using the $\lambda_p^{k(i)}$ as the price, corresponding to a first-price auction, will not be acceptable to the seller, since it is less than the offer o_p^i. In order to satisfy the seller's offer, the price must be the sum of $\lambda_p^{k(i)}$ and $\mu^i_{P_{\min}}$.

Now, suppose another pricing rule is used, with a corresponding exchange rate of $\chi \geq 1$. If χ scales the objective function of the OPF, it would result in a scaling of both the nodal prices and all shadow prices. So the corresponding adjusted price paid to generator i would be

$$\chi o_p^i = \chi(\lambda_p^{k(i)} + \mu^i_{P_{\min}}) \tag{24}$$

Clearly this is still not acceptable for the case where $\chi \leq 1$, since the price would still be less than the offer.

As a practical compromise, the following procedure is used to compute prices.

1. Determine the exchange rate χ for the desired pricing rule.
2. Set nodal prices to the scaled values $\chi\lambda_p^{k(i)}$.
3. For any generators with a binding minimum generation limit, if $\chi\lambda_p^{k(i)} < o_p^i$, set the price equal to the offer o_p^i.

5.3 Application to Emergency Imports

In the context of a single-sided market with inelastic demand where supply withholding is permitted, it is possible that a given set of offers will be insufficient to meet demand, resulting in an infeasible OPF problem. In the context of electricity market experiments, where the market solver is treated as a "black box", it is not acceptable to have a market that does not solve.

One way to overcome the infeasibility is to introduce emergency imports, implemented as generators with very high costs, whose generation will only be used if all other offered capacity has been exhausted. This has the potentially unintended consequence of producing very high prices when there is a shortage, set by the high-priced imports rather than the lower-priced offers from the sellers.

The following technique can be used to allow for emergency imports without having them set the price when they are needed. Treat each import as a fixed real power injection (fixed negative load) paired with an equally sized high-priced dispatchable load. Under normal circumstances the two will exactly cancel one another. However, when there is a supply shortage, the dispatchable load will be cut back, resulting in a net injection at that location. So far, this is completely equivalent to using a high priced import generator. However, since it is the bid for this dispatchable load that is marginal, there may be a large bid-offer gap. Choosing LAO pricing will drop the price to the bottom of that gap, where the most expensive offer is the one determining the market price.

It should be noted that this technique does not guarantee that prices will not exceed the highest offer. In these highly stressed cases, it is not uncommon to have some offered capacity which is unusable due to other binding network constraints. This results in a marginal generator offer along with the marginal import bid, eliminating the bid-offer gap and producing import level prices in portions of the network.

5.4 Visualization

Graphs showing offer and bid stacks can be very useful for visualizing the auction outcomes. In OPF-based markets, where there are externalities imposed by the network, it is not particularly useful to display the sorted set of raw bids and offers. Specifically, the accepted and rejected units will not typically be neatly separated from one another as in a simple auction.

However, if the offers and bids are normalizing to a single reference location r, the normalized offers and bids *can* be ranked and visualized in the standard fashion. For an offer o_p^i, the normalized offer can be written as

$$\bar{o}_p^i = \left(\frac{\lambda_p^r}{\lambda_p^{k(i)}} \right) o_p^i \tag{25}$$

It can be instructive to plot the offers as the original offer plus a "transmission adjustment" as shown in Figure 4, where the dashed black line shows the original offers and the solid black line is the adjusted offers. This adjustment is simply the difference between the normalized offer and the original offer.

$$\bar{o}_p^i - o_p^i = \left(\frac{\lambda_p^r}{\lambda_p^{k(i)}} - 1 \right) o_p^i \tag{26}$$

If the reference location r is chosen as the one with the largest value of λ_p then these adjustments will all be positive and appear as stacked on top of the original offers.

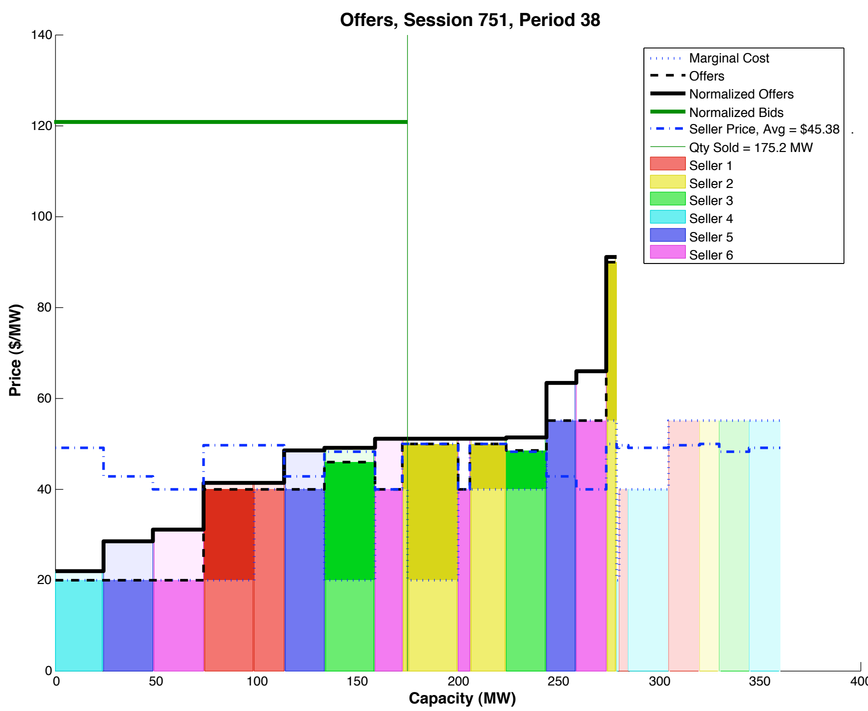

Figure 4: Example Adjusted Offer Stack.

The uniform price determined by the chosen pricing rule at the reference bus r

is then multiplied by $\lambda_p^{k(i)}/\lambda_p^r$ for each block to get the corresponding clearing price, plotted as the blue dash-dotted line in Figure 4.

Visualizing the demand side requires bundling the real and reactive portions together into a single per MW bid and then normalizing that bid to the reference location. For example, the normalized, bundled bid for load i would be

$$\bar{b}_{pq}^i = \left(\frac{\lambda_p^r}{\lambda_p^{k(i)}} \right) \left(b_p^i + \kappa_i b_q^i \right) \tag{27}$$

5.5 MATPOWER Implementation

The smart market auction implementation in MATPOWER includes all of the uniform pricing rules mentioned for real power. For joint real and reactive power auctions, it implements only the discriminative price auction (pay-as-offer/bid) and the first price auction (based directly on nodal prices from OPF).

In versions of MATPOWER prior to 4.0, the pricing rules for the various auction types were implemented incorrectly by *shifting* prices rather than scaling them, resulting in prices that are not consistent with the OPF solution. That is, using the resulting prices to replace the offers and bids and resolving the OPF would result in a different dispatch. Version 4.0 and later have corrected this problem.