Note: Comparison between the congestion management model and the congestion relief market

Introduction

This note considers the design and policy implications of the congestion relief market (CRM) model and compares it to the congestion management model (CMM). The CRM was originally developed by Edify Energy. A modified design was published by the Clean Energy Council (CEC) in its submission to the consultation paper of the Energy Security Board (ESB):

CEC the modified CRM model
CEC CRM formulation and simple example models
CEC CRM demonstration model 2 node model.xlsx
CEC CRM demonstration model 4 node model.xlsx

The note examines:

1. The key differences between the CRM and the CMM
2. The workability of the CRM
3. The key design choices needed to align the CRM with the objectives of access reform
4. Issues that might arise in implementing the CRM

The ESB will circulate future working papers to discuss potential solutions to the issues identified.

Overview of the CRM model

CRM description

The CRM is based on two dispatch runs. Whilst the CRM designers refer to these two runs as "energy dispatch" and "CRM dispatch", this paper refers to *access dispatch* and *physical dispatch*, respectively. This terminology change provides a more intuitive description of each dispatch's role in the CRM, and aligns with CMM terminology to facilitate comparisons between the two models.

The vision of the CRM designers is that the two dispatches will be functionally identical to each other and to the existing dispatch by the NEM dispatch engine (NEMDE). Only the bid inputs will differ, with generators submitting separate sets of offers for the two dispatches.

Access dispatch determines the quantity on which generators are paid RRP. In the CMM, this is referred to as the access quantity. Physical dispatch determines the quantity that the generator must physically produce. Any differences between the two dispatches are paid at the locational marginal price (LMP) for the generator:[1]

CRM$ = A x RRP + (G-A) x LMP

Where:

CRM$ = energy[2] payment to generator from AEMO settlement under CRM
A = level of access, determined by access dispatch
G = actual generation, determined by physical dispatch

[1] Ignoring loss factors for simplicity
[2] Note that there might need to be a similar settlement formula for FCAS. This is discussed later in the paper.

Comparison between the CRM and CMM

As noted, CMM also determines an access quantity and uses the same settlement formula:

CMM$ = A x RRP + (G-A) x LMP

Where:

CMM$ = energy payment to generator from AEMO settlement
A = level of access, determined by the CMM access allocation
G = actual generation, determined by physical dispatch

Unlike the CRM, the CMM does not use a dispatch run to determine access quantities but instead does this algorithmically. However, the formulae above show that the two processes are analogous, and the access dispatch in CRM can be thought of as a sophisticated access allocation method.

The main constraint on the quantity of access allocation is the sufficiency of settlement payments to generators from retailers:[3]

Σ CMM$ ≤ D x RRP

Where:

D = total retailer demand

Roughly speaking, this can be ensured by requiring that the access allocated represents a feasible dispatch: that each generator could in principle be dispatched at its access level without violating any constraints. In the CRM this is done explicitly; the access dispatch ensures feasibility by requiring that every constraint is complied with.

The CMM takes a shortcut to ensuring feasibility. Because the settlement payments are based on prices (LMP and RRP) and, because these prices depend only on those constraints that are binding in physical dispatch[4], the CMM only ensures that the access allocation does not violate these binding constraints. Given that there might only be a handful, at most, of transmission constraints binding in a region in any dispatch interval (DI), this is a faster way of allocating access than solving a complete dispatch problem.

However, this shortcut means that the CMM access allocation might not actually be physically feasible. In fact, it does not even represent a complete or coherent dispatch. Firstly, generators who do not participate in *any* binding constraint will not be allocated any access.[5] Secondly, generators who participate in multiple binding constraints will be allocated a different amount of access for each constraint, whereas in a feasible dispatch only a single dispatch target can exist for each generator. Notwithstanding this, even if the access allocation represented a complete and coherent dispatch, the method only ensures that physically-congested constraints are not violated. So it might inadvertently overload other constraints which are not physically-congested.[6]

[3] Ignoring inter-regional flows for simplicity. These are discussed later.
[4] Alternatively, this can be thought of in terms of the CMM charging for congestion and distributing congestion rebates. Where there is no congestion, there are no charges or rebates.
[5] Generators unconstrained in dispatch have LMP = RRP, and are simply paid RRP on their output (no need for access to be allocated).
[6] For example, there may be 100 transmission constraints in dispatch, only one of which becomes binding in physical dispatch, causing congestion. The CMM access allocation is a kind of re-dispatch, but it doesn't check

It does not matter in the CMM that the dispatch is not feasible, because it is just a financial arrangement. It is not proposed to actually dispatch generators in accordance with their access. However, for CRM this feasibility is important, because the opt-out mechanism allows generators to demand that they are physically dispatched at their access level.[7]

Similarities between CRM and European uniform-price auction models

The CRM bears some similarities to market models where financial markets are settled on an unconstrained[8] uniform-price auction and unders-and-overs between the financial cleared volume and physical output are paid at local prices, specific to each generator's location. Such market designs are the norm in Europe, for example, with all transmission constraints within each country removed in the unconstrained auction, thus permitting a uniform clearing price in each country.

There are two key differences between these European markets and the CRM:

1. The European uniform-price auctions typically clear ahead of time (eg day-ahead) whereas the CRM's access dispatch clears in real-time;[9]
2. The CRM's access dispatch is *not* unconstrained; rather it uses the same transmission constraints as physical dispatch.

Despite these differences, the operation of European markets demonstrates the feasibility of this type of model and the kind of behaviours and outcomes that can result.

Conclusions

The CRM and CMM are similar in that both allocate access to generators and then pay RRP on that access and LMP on differences between access and physical output. The difference is that the CRM bases access allocation on a feasible dispatch, whereas the CMM access allocation is algorithmic and only ensures feasibility on physically binding constraints.

The CRM designers have not framed it in this way, but rather see it as an initial dispatch based on today's arrangements followed by a re-dispatch to more efficiently manage congestion. In this sense, it has some similarities (and some key differences) to European electricity markets. Either framing is legitimate and can offer different insights.

---

all 100 constraints. It only checks the one that was binding physically. It is quite possible that the "re-dispatch" would violate one or more of the remaining 99 constraints.

[7] Discussed further below.

[8] In the sense of no transmission constraints

[9] Although there have in the past been market designs where the unconstrained dispatch is cleared in real-time, akin to the access dispatch in the CRM. The Ontario market used to have this design, and such a design was planned for South Korea but never implemented.

Access dispatch

Overview

In a dispatch run, NEMDE calculates the dispatch targets for energy and FCAS for each generating unit. The inputs into NEMDE are constraints (relating to demand, transmission and generation) and bids (generator offers). NEMDE then finds a dispatch which is:

- *Feasible*, in that no constraints are violated; and
- *Optimal*, in that it has the lowest dispatch cost of all feasible dispatches

Currently, NEMDE calculates physical dispatch based on generator bids. AEMO sends out dispatch targets to generators, who are required to operate their generating units to meet the targets.

Under the CRM, this dispatch process continues broadly unchanged.[10] However, there is now an additional dispatch process: access dispatch. This is purely a financial exercise; it doesn't change physical outcomes,[11] which continue to be based on physical dispatch.

Practicalities of Access Dispatch

At a high-level, the access dispatch is like any other dispatch, with an optimal dispatch based on bids. Bids are likely to be similar to those in today's physical dispatch so the outcomes should be similar.

But potential issues arise due to:

- Running a dispatch that is divorced from physical reality
- Having somewhat different bidding incentives from today's dispatch.

The physical issues are discussed in this section and the bidding incentives in another section below.

Today's physical dispatch is tied to the physical state of the system in several ways:

- Ramp rates limits are predicated on the metered output of generators at the start of the DI.
- Thermal transmission constraints are formulated as "feedback constraints" which include metered line flows on the RHS of the constraint.
- Dispatch non-conformance can be identified by comparing metered output to dispatch targets, so bids that do not reflect physical capability will be indirectly identified and penalised.

Access dispatch is liable to be different to physical dispatch – and hence physical reality – and so these ties cannot be applied to access dispatch. For example, consider a generator with a ramp rate, up or down, of 5MW a minute or 25MW over a dispatch interval. If the metered output at the start of the DI is 200MW, then NEMDE can only physically dispatch that generator to values between 175MW at 225MW for the end of the DI. But it would not seem sensible to apply the same limits to access dispatch, which might have a different starting point.

There are likely to be ways around this: for example, ramp rate limits in access dispatch could be based on the dispatch target from the prior access dispatch. Pre-dispatch is similarly a fictitious dispatch, not tied to physical reality. Constraints and ramp rate limits are formulated differently to reflect this. AEMO also runs "what-if" dispatches during intervention periods which are, similarly,

---

[10] Differences arise in how it is coupled with the access dispatch, as discussed below
[11] At least, not directly.

fictitious dispatches not directly tied to physical dispatch. Nevertheless, there could be some design and implementation costs associated with developing these workarounds.

That leaves the question of ensuring physical feasibility. For example, thermal generators are typically slow to start and have a minimum generation (mingen) operating limit, below which they cannot remain stable. In today's dispatch, they must notify when they will be on-line and then ensure that they are always dispatched to at least their mingen when on-line: typically by bidding -$1000 for this output level.

There are a number of design questions to consider for the CRM:

- Should bids into access dispatch be similarly restricted?
- Would the on-line period have to be the same as in physical dispatch?
- If so, how would compliance against this requirement be monitored, given that there is no "dispatch non-conformance" to signal non-compliance?

Special Dispatch Runs and Re-runs

Currently, there are various situations where additional dispatch runs are required, commonly where specific physical constraints on operation are unable to be modelled by the NEMDE functionality in a single run. These additional runs will need to continue for the physical dispatch under the CRM model. It is worth examining whether these additional runs are also required for access dispatch, since avoiding this might significantly reduce the cost of implementing and operating this model.

FCAS

If access dispatch is formulated the same as physical dispatch then it must incorporate FCAS as well as energy. That would suggest two sets of FCAS bids, leading to two sets of dispatch prices and dispatch targets, with FCAS settlement being analogous to energy settlement in combining the two: eg

$$FCAS\$ = P_A \times Q_A + P_G \times (Q_G - Q_A)$$

Where

FCAS$ = total payment for FCAS provision
P_A = FCAS price in access dispatch
Q_A = FCAS quantity dispatched in access dispatch
P_G = FCAS price in physical dispatch[12]
Q_G = FCAS quantity dispatched in physical dispatch

This settlement would be applied separately for each of the 8 FCAS services. There may need to be some corresponding adjustment to FCAS cost-recovery algebra.

An alternative approach would be not to dispatch FCAS in the access dispatch: ie set FCAS requirements to zero. Whilst this would then not strictly be a feasible dispatch, this would not normally matter for the purposes of settlement adequacy.[13] All FCAS would then be dispatched in

[12] Note that, since there are rarely intra-regional FCAS constraints, there will generally be a single FCAS dispatch price for each region. Thus the issue of whether to charge a regional price or a local price does not arise in the way that it does for energy.
[13] This approach might leave it unclear how to formulate "islanding constraints", in which FCAS requirements are incorporated into transmission constraints.

the physical dispatch, with a single price and simple settlement for FCAS provided just like today. However, under this alternative approach, generators opting out of the CRM (ie locking their physical dispatch to their access dispatch) may find it more difficult to provide FCAS.

Settlement Residue and Interconnector Clamping

Because the access dispatch is feasible, CRM settlement will generally result in a settlement surplus[14], that is:

Aggregate payments to generator ≤ aggregate payments from retailers

However, this ignores payments to interconnectors: referred to in the current market design as inter-regional settlement residues (IRSRs). Currently, the IRSR[15] depends upon the price difference between the regions that the interconnector flows between:

$IRSR\$ = IC \times (RRP_M – RRP_X)$

Where:

IC = physical flow from the exporting region to the importing region
RRP_X = RRP in the exporting region
RRP_M = RRP in the importing region
As usual, losses are ignored for simplicity

When there is a *counterprice* flow ($RRP_M < RRP_X$) the IRSR is negative. Put another way, there is an overall settlement deficit between generator payments and retailer receipts. To manage this problem today, AEMO places clamping constraints on interconnectors when these deficits are large, forcing the interconnector flows to be dispatched back to zero and thus setting the IRSR to zero.

IRSR arises under the CRM and CMM designs, but it depends on access, not on physical dispatch:

$IRSR\$ = A_{IC} \times (RRP_M – RRP_X)$

Where:

A_{IC} is the access allocated to the relevant interconnector

Under CMM, it is anticipated that the method for allocating access to interconnectors will be designed to avoid settlement deficits arising. Similarly, a method is likely required for the CRM. Given that access allocation in the CRM is based on access dispatch, any method to prevent deficits must apply a corresponding constraint to access dispatch. Such constraints will be similar to the current interconnector clamping constraints. But, unlike today's market, the clamping will apply to access dispatch, not to physical dispatch. So, as with the CMM, physical counterprice flows will be permitted, and will naturally arise when these are consistent with an optimal dispatch outcome.

Conclusions

Because access dispatch is designed to be functionally identical to physical dispatch in the CRM, it should be feasible to calculate it using the same systems and processes that are used in today's

[14] If exactly all the constraints binding physical dispatch also bind in access dispatch, then settlement will balance exactly. However, this is generally not the case, and those constraints that bind in physical dispatch but *don't* bind in access dispatch generate a surplus, because the congestion charges collected on physical dispatch are not fully returned to generators in the rebates calculated by access dispatch.
[15] Note that IRSR payments are currently sold to market participants and traders through the settlement residue auction. Since IRSRs continue to be generated in the CMM and CRM, this process could continue.

dispatch or pre-dispatch. However, a question arises whether this identity is strictly necessary. It may be possible and appropriate to simplify the access dispatch somewhat to reduce costs whilst preserving its effectiveness in achieving the objectives of the CRM model.

Coupling between Access and Physical Dispatches

Coupling Constraints

An important feature of the CRM is the ability of individual generators to *opt out* of LMP exposure, or otherwise to explicitly manage this exposure. Since LMP is paid on the difference between the two dispatches (ie *G* minus *A*), the dispatch quantities need to be identical for generators who opt out. This is done by setting additional constraints in the physical dispatch, of the form:

G = A

Alternatively, to limit its LMP exposure, a generator can bid upper and lower bounds (eg +/- 10MW) to the difference, with physical dispatch then incorporating the associated constraints:

G ≥ A – lower bound
G ≤ A + upper bound

This paper refers to these constraints as *coupling constraints*, in that they couple the outcomes of the two dispatches.

In principle, every generator could opt out in this way,[16] thus making the two dispatches identical. In this situation, it would be critical to ensure that access dispatch is feasible so that the identical physical dispatch is also feasible and does not endanger system security.[17]

These coupling constraints are how a generator notifies AEMO: "I want to be physically dispatched at or close to my access level." Generators already notify AEMO of their desired dispatch level through their bids. Generators could similarly bid so that their dispatch tracks their access level, but two difficulties arise. Firstly, because access allocation is only calculated in real-time, there is no opportunity to construct the required bids in advance. Secondly, a generator may not have the physical capability to track its changing access level in dispatch: eg if the access quantity suddenly jumps up from one DI to the next, and ramp rate limits prevent generator output following suit.

The CRM addresses these difficulties by:

- Formulating the coupling constraints in a way that does not require advanced knowledge of access quantity; and
- Ensuring, through the access dispatch, that access represents a physically feasible dispatch: for each generator and for the system as a whole.

[16] Although unlikely in practice, because generators will be financially rewarded for taking LMP exposure.
[17] In practice, the formulation of physical dispatch would not permit system security to be endangered by financial considerations. NEMDE would choose to violate the coupling constraints before violating any security constraints. But that would make the opt-out ineffective under these conditions.

Therefore, the access dispatch design in CRM[18] follows inevitably from the opt-out requirement. It would be difficult or impossible to opt out if different access allocation methods were used, such as in the CMM.

Co-optimisation between Access and Physical Dispatches

As well as coupling the two dispatches, the CRM also *co-optimises* the two dispatches. It does this by formulating a single objective function which is the sum of the two dispatch costs:

Objective Function = Cost of access dispatch + cost of physical dispatch

This means that the two dispatches are solved at the same time, as a single optimisation problem, rather than as two separate optimisation problems. It is analogous to how today's NEMDE co-optimises the dispatch of energy and FCAS:

NEMDE objective function = cost of energy dispatch + cost of FCAS dispatch

Energy and FCAS dispatch are coupled,[19] so this co-optimisation means that energy dispatch will in part reflect FCAS bids and vice versa. This is an important feature of NEMDE which reduces the total cost of dispatch.

In the CRM, only the physical dispatch incurs actual costs. The access dispatch is simply a financial arrangement and does not cause any costs per se.[20] So the purpose of this co-optimisation in the CRM context is unclear.

Bidding Incentives created by Co-optimisation

A design principle that is ubiquitous in market models that employ a combination of access and LMP,[21] is that the method for allocating access is *independent* of the bids that are submitted into physical dispatch. That is to avoid bids for physical access being tailored towards getting a better access allocation, rather than simply to optimise the financial position in physical dispatch. Such mixed incentives are liable to worsen the efficiency of physical dispatch. Today's market design does not conform with this principle, leading to the -$1000 bidding that distorts physical dispatch.

In a co-optimisation, all bids can affect everything. For example, in co-optimisation of energy and FCAS today, energy bids can affect FCAS dispatch and vice versa. Similarly, with co-optimisation of the two dispatches in the CRM, physical bids can affect access, and this clearly violates the above design principle.

That could create a new form of "race to the floor" bidding in the CRM, with generators bidding minus $1000 into *both* dispatches ("minus $2000" across the co-optimised model) and coupling the two dispatches.[22] A co-optimised dispatch will dispatch this generator in preference (other things being equal) to a generator bidding -$1000 into access dispatch, and bidding cost-reflectively into

---

[18] That is to say, the access allocation must be based on a feasible dispatch. However, it need not be based on a dispatch that generators actively bid into. For example, bids could instead be determined using an administrative process. But the philosophy of the CRM is that access dispatch looks like today's dispatch, so this *does* require active bidding.

[19] Through the "trapezium constraints" that define FCAS capability as a function of energy output.

[20] Of course, it affects the financial position of individual participants, but these are just wealth transfers that net out overall.

[21] Including related models that use financial transmission rights as a proxy for access.

[22] It is not entirely clear whether, under the CRM design, an opt-out generator would need to submit bids into physical dispatch at all, since those bids don't seem to play any useful role. But, in this case, it is not clear what the co-optimisation is likely to achieve, since there are no longer two sets of bids to co-optimise between.

physical dispatch. It seems plausible – although this has not been investigated thoroughly – that, as a result, every in-merit constrained generator might end up bidding "minus $2000" into the CRM, just as such generators today bid at -$1000. [23] The outcome would be that physical dispatch remains distorted and the CRM fails to achieve the ESB's access objectives.

The non-co-optimised alternative

This problem of distorted bidding incentives arises solely because of the co-optimisation between the two dispatches. This co-optimisation does not appear necessary to achieve the CRM objectives. Furthermore, it will likely add to implementation costs by requiring the development of a new dispatch engine that can co-optimise two dispatches in this way. It would seem preferable to have the two dispatches done separately and sequentially as follows:

1. Determine the access dispatch based on access bids; then
2. Determine the physical dispatch based on physical bids and any coupling constraints.

So, in effect, NEMDE would be run twice at the start of each dispatch interval to solve the two dispatch problems. This alternative approach would seem to achieve the CRM objectives, whilst avoiding the problem of having access dependent on physical bids and associated distortions to physical dispatch.

The ESB is seeking feedback from the TWG (and public consultation) on the proposed design before submitting the options to AEMO, which will determine the implementation risks and mitigation plan.

Conclusions

In the CRM design, the two dispatches are coupled in two ways:

- Through coupling constraints that a generator can submit in order to opt out of the CRM;
- Through the two dispatches being co-optimised within a single optimisation problem.

The first element is an important feature, creating the opt-out facility which is an integral part of the CRM philosophy. The second feature seems superfluous and could damage dispatch efficiency by encouraging a continuation of -$1000 bidding in physical dispatch.

The proposed non co-optimised alternative does not appear to create threshold implementation risks but AEMO will be best placed to confirm its system integration plan and impact on solve times.

[23] Generators bidding "-$2000" in this way would need to opt-out of the CRM, and so give up any profitable trading opportunities in this market. However, this would likely be more than made up by the extra access dispatch (on which RRP is paid) obtained by this bidding strategy.

Pricing and bidding behaviour

Overview

Under the CMM and CRM, generators are paid RRP on their access dispatch, which is analogous to generators being paid RRP on their physical dispatch today. In the CRM, bidding incentives[24] will likely encourage constrained, in-merit generators to bid at -$1000 into access dispatch to maximise their access to RRP. This means that access dispatch is likely to have the winner-takes-all (WTA) characteristic seen in today's dispatch.[25] This can be seen as a strength or a weakness. On the upside, it means that generators' access to RRP will be similar, other things being equal, to today's market, reducing the financial disruption when the CRM is introduced. On the downside, WTA dispatch might be seen as an accidental and undesirable feature of today's market design which access reform should seek to eliminate or restrict. The ESB has developed alternative access allocation methods in the CMM which remove or mitigate this WTA characteristic. The CRM design does not offer these allocation alternatives, at least as currently specified.

Arbitrage opportunities

The CRM gives generators the opportunity to bid separately into two different markets. Such a situation will typically present arbitrage opportunities for profit-maximising traders; they will seek to sell into the higher-priced market and buy from the lower-priced market.

Recall the formula for generator settlement under CRM:

CRM$ = A x RRP + (G-A) x LMP

Where:

A = level of access dispatch
G = level of physical dispatch

Generators can sell a portion of their physical output into the access market at RRP and the remainder into the physical market at LMP. For example, a generator dispatched at 100MW might sell:

- 100MW at RRP and 0MW at LMP
- 0MW at RRP and 100MW at LMP
- 50MW at RRP and 50MW at LMP
- 130MW at RRP and -30MW at LMP
- Etc

An arbitrage strategy would see the generator aiming to sell into the market with the higher price: ie

- If LMP<RRP, sell maximum output into access dispatch.
- If LMP>RRP, sell maximum output into physical dispatch.

[24] Although there are also some new incentives, discussed below.

[25] Meaning that generators will be dispatched in order of participation in the constraint: generators with low participation factors are likely to be fully dispatched, whereas generators with high participation factors are liable to be fully constrained off.

For the purposes of this arbitrage analysis, we can categorise each generator in terms of how its output affects congestion:[26]

- *It worsens it*: the generator has *positive participation* and has LMP < RRP
- *It relieves it*: the generator has *negative participation* and has LMP > RRP
- *It doesn't affect it:* the generator has *zero participation* and has LMP = RRP

Opportunities for arbitrage for the first two types are discussed below.[27] Arbitrage by the third type is discussed in the following section.

It is useful to further subdivide generators into:

- *In-merit*: operating costs are below RRP
- *Out-of-merit (OOM)*: operating costs are above RRP

A positive participation generator has RRP > LMP and will aim to maximise its access dispatch by bidding -$1000. This occurs today, but only for in-merit generators. In the CRM, *all* positive participation generators – both in-merit and OOM - will bid -$1000, irrespective of cost. Unlike today, access dispatch does not create any actual operating cost for the generator; it is purely financial. Only physical dispatch causes operating costs, and an OOM generator need only ensure that it is not physically dispatched.[28]

A negative participation generator has RRP < LMP and will aim to *minimise* its access dispatch by bidding +$15,000.[29] This occurs today, but only for OOM generators; if an in-merit generator did this it would deny itself a profitable opportunity to run against RRP. But, in the CRM, *all* negative participation generators will bid +$15,000, irrespective of cost, because the opportunity remains to be physically dispatched and profit from their output.

Thus, although the access dispatch is set up to be equivalent to today's dispatch, bidding incentives will be different for two categories of generator, as summarised in table 1, below.

Congestion type	*Merit position*	*Bid in Today's dispatch*	*Bid in access dispatch*
Positive participation	*In-merit*	-$1000	-$1000
Positive participation	*OOM*	at cost	-$1000
Negative participation	*In-merit*	At cost	+$15,000
Negative participation	*OOM*	+$15,000	+$15,000

Table 1: arbitrage bidding in access dispatch

This raises the question as to whether generators should be freely permitted to bid in this way in access dispatch, recognising that this will change outcomes relative to today's dispatch, or whether bidding should be regulated to ensure that access dispatch better reflects today's bidding and dispatch and, if so, how would this be done. It would be straightforward to prohibit generators that are bidding out-of-merit in physical dispatch from bidding -$1000 into access dispatch. However,

[26] At a point in time. Over time, the location of congestion will vary and the generator's impact on it will vary accordingly.
[27] The picture is complicated where there are two or more points of congestion, so for simplicity of exposition this paper assumes that it is just at one point: ie a single transmission constraint is binding in physical dispatch.
[28] For example, by bidding at cost into physical dispatch.
[29] it could, alternatively, bid unavailable, but we will assume that it is not permitted to do this unless also bidding unavailable in physical dispatch, which it will not wish to do.

out-of-merit generators could easily get around that by rebidding "in-merit" into physical dispatch[31] So the regulations would need some way to infer a generator's in-merit status other than from its bid e.g. based on its bidding during periods without congestion.

Similar questions also arise for the CMM design, albeit expressed in terms of the access allocation method rather than bidding incentives. These questions will need to be answered for the respective designs to be finalised.

RRP

In both access models, RRP is the price paid to generators for the access they have been allocated, as well as the price that retailers pay. But this begs the question as to how RRP is calculated.

Today's RRP is calculated as part of the physical dispatch process. It is technically the LMP at the regional reference node (RRN). But there are two dispatches in the CRM, each of which calculates an RRP. There is a design choice about which RRP to use to settle the market. Other market designs take different approaches to determine the RRP or its equivalent:

- In European markets, the "RRP" is set by the unconstrained schedule, analogous to using the RRP from access dispatch.
- The CMM uses the RRP from physical dispatch.

The CRM design uses access RRP, which reflects its philosophy of setting up the access dispatch as a continuation of today's market design, to which the physical dispatch is appended. But it would equally be possible to use physical RRP, to bring the CRM design closer to the CMM.

The choice of RRP has implications in three areas:

- settlement adequacy
- arbitrage opportunities
- gaming opportunities.

These are considered in turn below.

Settlement adequacy

As discussed earlier the IRSR on each interconnector is governed by the formula:

$$IRSR\$ = IC \times (RRP_M - RRP_X)$$

This becomes negative when there is a counterprice flow (ie $RRP_M < RRP_X$), so the interconnector must be clamped in access dispatch in this situation. But which RRPs should we use? If access RRP is used, the RRPs come from access dispatch, which is analogous to today's clamping. If physical RRP is used, the RRPs come from physical dispatch. In that case, we are taking the prices from one dispatch and the quantities (flows) from the other. This does not seem to present any practical difficulties; either way, the clamping constraint must be applied before the RRPs are known for certain, just as is the case today. But it creates a rather peculiar, hybrid definition of "counterprice flow", which might have ramifications for market outcomes more generally: for example, by creating volatility or oscillations in dispatch and pricing.

---

[31] So long as they ensure their bids are still higher than LMP, they will not be physically dispatched. So they could rebid anywhere between LMP and RRP.

Arbitrage opportunities

Zero participation generators have LMP = RRP. Or, more precisely have LMP = physical RRP, since LMPs are determined by physical dispatch. If RRP is based on access RRP, then the arbitrage discussed above for constrained generators becomes, for zero-participation generators, an arbitrage between access RRP and dispatch RRP. Bidding incentives will be analogous to those identified above:

- If access RRP > physical RRP, sell maximum output into access dispatch
- If access RRP < physical RRP, sell maximum output into physical dispatch

This is not to suggest that zero-participation generators would be bidding at -$1000 or +$15,000 as constrained generators would. The bidding behaviour would likely be more nuanced: raising or lowering bids somewhat, depending on the expected RRP difference.

Arbitrage activity always closes the price differences between the two arbitraged markets, by raising demand – and hence price – in the lower-priced market. This arbitraging by zero-participation generators is likely to have similar effects. If this is effective at closing the gap between the two RRPs, it might not matter which is used in the CRM. Or, put another way, one might as well choose to use physical RRP[32] which would avoid creating this arbitrage opportunity in the first place.

Gaming opportunities

If RRP is based on physical RRP, then generators at the RRN[33] are paid:

$$\text{CRM\$} = A \times RRP_{phys} + (G\text{-}A) \times LMP_{RRN}$$

$$= A \times RRP_{phys} + (G\text{-}A) \times RRP_{phys}$$

$$= G \times RRP_{phys}$$

That is, they are indifferent to their access dispatch level and so indifferent to how they bid into access dispatch. They could bid at cost, at the MPF or at the MPC; their settlement payment would be identical. But whilst the access bid does not affect the generator itself, it could affect other generators: for example, if all RRN generators bid at the MPF, this would reduce the access dispatch of constrained generators and so reduce their access to RRP.

A portfolio generator, owning and bidding both constrained and RRN generators, would be liable to bid its RRN generator in a way that increased the access dispatch of its constrained generator. That could give an unfair access advantage to that portfolio generator over another generator that does not own RRN generators: eg a wind-farm owner.

This issue could be avoided by settling energy at the access RRP[34], meaning that the bidding of RRN generators would be disciplined by its potential impact on RRP, similar to today. Or there might be design mechanisms that could be introduced to either regulate access bidding of RRN generators, or to mitigate its impacts on access dispatch.

Dispatch variances

A dispatch *variance* is the difference between a generator's output and its dispatch target. Variances are inevitable in physical dispatch, including under the CMM and CRM.

[32] Noting, however, the earlier concern about the impact on interconnector clamping dynamics.
[33] Or, more generally, generators that are unconstrained and so have the same LMP as at the RRN.
[34] Although note that there are concerns around that option also, as discussed in the previous section.

In the CMM, generators are paid LMP at the margin for their physical, metered output. This means that dispatch variances are automatically paid at LMP. The CRM designers propose that variances are paid at RRP, although the alternative option of paying at LMP is acknowledged. The chosen option is consistent with the vision of the CRM as a continuation of today's market (under which variances are paid at RRP) with a congestion relief market appended. It also means that generators who opt out of the CRM to avoid LMP exposure aren't inadvertently exposed to LMP as a result of dispatch variances.

Variances will generally be small, because dispatch non-conformance is tightly regulated. But it is worth examining why this is the case. Constrained-off generators paid RRP for variances would commercially prefer, if they were permitted, to generate at a higher level than their dispatch targets. But such variances could overload the relevant constraint and so must be prevented.

Now consider the situation under LMP settlement. A generator will be constrained off only if its LMP is at or below its offer price, making it unprofitable to deviate from its dispatch target. So, paying LMP on variances would substantially reduce the impact of any dispatch non-conformance; paying RRP on them would leave the current situation unchanged.

Another reason to manage dispatch non-conformance is to reduce the amount of regulation FCAS required. But an incentive to do this is already in place through the causer-pays mechanism. A further improvement to this incentive is being considered through the "double-sided causer-pays" design, where generators would be rewarded for variances if these assist with frequency regulation. Existing conformance rules would limit the opportunity for generators to these new signals. If LMP settlement were used for dispatch variances, these rules could be relaxed and frequency regulation further improved.

Whilst opt-out generators would prefer not to be exposed to LMP on dispatch variances, the fact is that they are already exposed to FCAS prices on these variances, through the application of causer pays. It is not clear that paying LMP on variances under the CRM would substantially add to these existing risks.

Conclusions

The existence of two separate markets, with separate bids, in the CRM creates opportunities to arbitrage between the two. It should be considered whether such arbitrage would be inconsistent with access objectives and, if so, how this could be managed.

Similar issues arise in the CMM, which would need to be managed administratively, through the design of the access allocation method.

Overall conclusions

The list of design elements and issues discussed in this paper are summarised in table 2 below.

There are some material design choices affecting both the CRM and CMM including:

- Wealth transfers to out-of-merit generators or in-merit negative participation generators
- Managing LMP exposure

The CRM creates additional complexity in its design considerations and implementation because of the number and type of systems that it affects (bidding, pre-dispatch and dispatch) compared to affecting settlement systems in the CMM.

Material items affecting only the CRM include:

- Choice of RRP based on physical or access dispatch
- Co-optimisation (or not) between access and physical dispatch
- Complexity of access dispatch including ancillary markets.

Readers should note that the list of 'material' items is limited to the points raised in this working paper.

Element	CMM design	CRM design	Issue	Potential solution
		Access Allocation		
Complexity of Access allocation	Simple Algorithms	Access dispatch	Access dispatch is potentially complex	Consider simplifying CRM access dispatch
FCAS	Not relevant	FCAS included in access dispatch	More complex FCAS pricing and settlement	Consider simplification options
Settlement Residue	Algorithm allocates access to interconnectors, no negative residue	IRSR on interconnectors in access dispatch	Counterprice flows in access dispatch will cause settlement deficits	Clamp interconnectors in access dispatch but not physical dispatch
		Coupling Access and Physical Dispatch		
Managing LMP exposure	Through bidding around access forecast	Through coupling constraints and bidding	CMM and CRM lack access forecasts	Consider access forecasts for CRM and CMM
Co-optimise access and physical dispatches	Not an issue	Through combined objective function	CRM co-optimisation encourages disorderly bidding in physical dispatch	Remove co-optimisation and solve as two separate dispatches
		Pricing and Behaviour		
Out-of-merit positive-participation generators	Aim to exclude from access allocation	Incentivised to bid into access dispatch	Possible wealth transfer to these generators in CMM and CRM	Consider excluding from access allocation in CMM. Consider bidding rules to manage this in CRM.

Element	CMM design	CRM design	Issue	Potential solution
In-merit, negative-participation generators	Policy not decided	Incentivised to bid out of access dispatch	Possible wealth transfer to these generators in CMM and CRM	Consider bidding rules to manage this in CRM.
RRP settlement	RRP from physical dispatch	RRP from access dispatch	CRM design encourages arbitrage between the two RRPs	Consider using physical RRP for CRM
Dispatch variances	Paid at LMP	Paid at RRP	LMP provides better incentives for managing variances	CRM should pay variances at LMP

Table 2: design comparisons, issues and options