5 August 2010

Comments on the TLAF Workshop of 26 July 2010

A Report for Viridian Energy Ltd

NERA
Economic Consulting

Confidential Draft

Author

Graham Shuttleworth

NERA Economic Consulting
15 Stratford Place
London W1C 1BE
United Kingdom
Tel: +44 20 7659 8500
Fax: +44 20 7659 8501
www.nera.com

Contents

Confidential Draft

Executive Summary

[optional – to be added later if necessary]

1. Introduction

Viridian has asked me to comment on certain aspects of the presentations made at the recent workshop on TLAFs. I attended the workshop, which was held in Dundalk on 26 July, as an observer. The following points sprang to my attention.

Several speakers claimed that the current system was "broken", and some even suggested in passing that *any* alternative scheme "had to" be an improvement. However, no-one presented any evidence to back up this claim. This part of the discussion was prompted largely by the presentation of ESB,[1] but in fact ESB provided no evidence of any flaw in the method of calculating TLAFs. ESB's statements about the method are partly incorrect and do not in any case amount to any criticism of the economics behind the use of marginal loss factors. Indeed, the analysis only confirms that the TLAF correctly measures the costs that generators in Cork impose on the system (and hence on other system users). It even contradicts the view that TLAFs are unpredictable or volatile in the short term.

ESB's key slides on this topic therefore represent little more than a complaint that ESB's plant is suffering a loss of revenue due to its disadvantageous location. That complaint does not provide any grounds for moving away from marginal loss factors, let alone for adopting a uniform TLAF, by reference to the high level objectives applicable to the RA's decision.

Instead, it would have been helpful if the workshop had devoted more time to examining the real source of concern about TLAFs, which is the variation in the TLAF for a particular point over the long term. Although the Irish Wind Energy Association hinted at such an issue in the slide entitled "Volatility Example", the discussion did not properly separate out variation between time periods, variation due to short term changes in operating conditions, and variation due to long term investment and demand growth.

Below, I consider the evidence presented by ESB and explain why it does not present any reasoned criticism of the current system. I also consider the problem of long term variation in TLAFs and possible responses that the Regulatory Authorities could and should have considered before reaching a decision.

---

1 Available at: http://www.allislandproject.org/en/transmission_decision_documents.aspx?article=52d81fe7-c75e-4f67-8fc1-7b87d57cad37&mode=author

2. Calculation of TLAFs

One part of ESB's presentation plots out the total system losses arising from different levels of output at the ESB ***and BGE*** plants (c. 850 MW in total) located in Cork – and hence the marginal losses associated with each increment of output. This analysis *confirmed* that the TLAF applied to those plants is a reasonable estimate of the transmission losses caused by the final increment of output (from either plant) when they are running at full output (i.e. most of the time). ESB claimed that the "actual" transmission losses caused by other increments of the plants' total output implied different (and more advantageous) estimates of the TLAF. However, that finding does not have any basis and does not indicate any problem with the current system.

2.1. Marginal Costs in Principle

There is no reason to expect that the marginal losses attributable to the marginal increment of output would ever equal the marginal losses attributable to any other ("inframarginal") increment of output, or the average TLAF over a wider range. In general, a marginal cost is the cost of producing the marginal unit and it does not usually equal the "inframarginal" cost of producing other units, or the average cost of producing all units taken together. ESB is therefore setting up a false standard for the appraisal of TLAFs.

For comparison, consider energy prices. The market price of energy in the SEM (and other markets) is the System Marginal Price, which is set equal to the incremental cost of the marginal generator. This market price is recognised as a fair representation of the value of electricity and as the price that encourages efficient decisions by producers and consumers alike. No-one is arguing for ESB's "principle" to be extended to energy pricing, such that the price for different units of output would be derived from market conditions assuming different levels of output.

2.2. Marginal Costs and Efficiency

The objective of efficient despatch is promoted by setting the TLAF equal to the marginal cost imposed by the last increment of output. It makes each system user responsible for the costs (i.e. the transmission losses) that they impose on the system by their decisions to produce that incremental output. Charging system users the marginal cost of their actions encourages them to make efficient decisions about the level of available capacity and the price of output offered to the system operator.

The principle at work here is the desire to encourage efficient decisions. System users should therefore bear the costs imposed on society (or on the electricity system) by their decisions. In the current context, the TLAF represents the estimated cost imposed on the transmission system by each generator's decision to commit the last increment of their plant's availability and output (in a set of expected scenarios).

The workshop could have discussed whether the scenarios used to calculate the TLAFs accurately reflect the actual pattern of output – e.g. whether the two generators actually run at full load. The workshop (and the RAs' Proposed Decision Paper) could also have discussed what decision is being affected by TLAFs and therefore whether to calculate incremental transmission losses for

(1) the decision to make a small increment of capacity (i.e. output) available for despatch; or

(2) the decision to make a larger increment of capacity (i.e. output) available for despatch; or

(3) the decision to commit a whole plant's capacity.

Taking incremental transmission losses over a larger increment of output (say 400 MW) would (according to ESB's presentation) reduce or stabilise the TLAF for generators in Cork. The presentation shows that the incremental losses ("ΔLosses") are a smaller proportion of incremental output ("ΔOutput") for an incremental output of 400 MW than for an incremental output of 20 MW. The TLAF calculated for such a large increment would represent the marginal cost of the decision to commit the whole generator, and would therefore encourage efficient decisions about plant commitment. Decisions about the precise level of output might be less efficient, but the difference between the two TLAFs is unlikely to reduce the efficiency of despatch by very much (at least for baseload generators).

2.3. Invalidity of ESB Example

ESB's presentation does not in fact criticise the principle of calculating marginal losses attributable to a generator's decision. The graphic shows the marginal losses attributable to each increment of c. 850 MW of capacity and also argues that the incremental losses due to the commitment of all this capacity is very small (1.75MW). ESB's presentation therefore *accepts* the desirability of assigning incremental losses to individual generators.

However, ESB's graphic estimates the incremental transmission losses caused by the combined output of *two* plants owned by two *different* companies. There is no rationale for adopting this approach, either as a method of calculating TLAFs or as a standard for appraising TLAFs calculated by other means. At no time will any single player in the market be faced with a decision about committing both plants – unless ESB and BGE explicitly collude over the operation of their Cork plants. As a result, there is no need to present anyone with the costs of such a decision.

The maximum capacity affected by any single decision is the capacity of one plant (c. 400 MW), as discussed above. ESB's presentation shows that the decision to commit the last 400 MW of capacity imposes additional transmission losses on the system, and that a TLAF substantially less than one will therefore encourage efficient decisions. There is no reason to adopt a zero TLAF (let alone uniform TLAFs) because of the combined impact of two generators.

One related proposal would be impossible to implement. It was suggested at the workshop that different increments of output should be assigned different TLAFs, based on the incremental losses at different levels of output, as shown in the ESB presentation. Leaving aside the difficult of carrying out this calculation for each generator within the Single Electricity Market, it would be impossible to say which of the increments in ESB's graphic belonged to the ESB plant and which to the BGE plant. Both generators would wish to claim the "first" 400 MW, which impose negative losses, and to avoid the "second" 400 MW, which impose positive losses. Any allocation of tranches between the two generators would be entirely arbitrary. The only rational approach is to recognise that, once they have been built, *both* generators are marginal and should be assigned the same loss factor. The TLAF for both plants would be less than one, whether it was estimated for incremental output of 20 MW or 400 MW.

If there were concerns about the long-term trends in TLAFs, it would be necessary to consider ways to provide some kind of hedging or two-part tariff, such that the TLAF for any individual generator was fixed by contract for a specific level of forecast output, and set equal to marginal losses for any difference between actual and forecast output. I consider these options in more detail in section 4 below.

2.4. Conclusions

Nothing in ESB's presentation of the TLAFs attributable to different tranches of capacity calls into question the principle of assigning losses to generators on the basis of the incremental losses that they impose upon the system. Indeed, ESB's position appears to be that the generators at Cork should bear little or no liability for losses precisely because losses are the same with and without these generators, i.e. because incremental losses associated with these generators equal zero.

Setting the liability for losses (or any other charges) equal to marginal costs encourages efficient decision-making. However, it is not efficient or rational to consider the decision to commit both generators together (unless their owners collude). One might wish to consider the incremental losses caused by committing one generator (rather than just the last increment of its output), but the resulting TLAF would (according to ESB's analysis) still be substantially less than one.

ESB may have wished to raise a point about long term trends in TLAFs, but did not. In any case, I consider that point in section 4 below.

3. Plausibility of TLAFs

Part of ESB's presentation questions the result that the two plants in Cork should be held responsible for marginal losses equal in some cases to total losses on the system (or rather, on the ROI part of the system). This argument appears to have been picked up in passing in the Proposed Decision Paper, which says on page 10 that "It has been argued that the marginal cost approach could lead to overly punitive losses being attributed to particular generators on the island." However, this part of the presentation provides no evidence of any failure or flaw in the system, since it is a possible outcome of a rational method (as described above) and because total transmission losses (particularly total losses within part of the system) do not represent any standard by which to judge marginal loss factors.

3.1. Total and Marginal Losses

If the transmission system is broadly in balance and load flows are minor, total transmission losses will be a small proportion of total generation. However, *incremental* output can still have a large impact on total losses – up to about 10% of the incremental output – if it is located far from a centre of (net) demand, as in the case of Cork. If there is a lot of capacity at that location – and the two plants in Cork amount to over 10% of total output on many occasions – then the losses allocated to that capacity will also be large relative to total losses. Such outcomes are consistent with incentives for efficient decisions and do not indicate any flaw in the calculation. Instead, complaints about this outcome merely indicate a dislike for the scale of the charge imposed on a particular user due to the location of their capacity.

A simple example may help to illustrate the case. Assume that, in one hour, total output in the SEM is 8,000 MW and that total losses equal 120 MW (1.5% of output). There is no real reason to divide the Irish electricity system at the border, but suppose that 72 MW of losses (60% of the total) arise within the ROI. When operating in baseload mode, the combined capacity of the ESB and BGE plants is about 850 MW – i.e. over 10% of total output. To assign these generators losses of 72 MW (i.e. equal to total losses in the ROI), they require a TLAF of only about 9%. Such a figure is not implausible, if the plants are located far from a centre of net load, because (as the ESB presentation shows) that is the scale of the incremental losses that their output creates.

For comparison, consider the effect of setting energy prices equal to system marginal costs. At times, the revenue earned by a large generator may exceed total generation costs within the Republic of Ireland. However, that provides no basis for criticising the pricing rule.

3.2. Conclusion

The ESB presentation therefore provides no indication that the current method is fundamentally flawed. Neither the sensitivity of the TLAF to the capacity in Cork, nor the values of the Cork TLAF relative to total losses indicate any problem related to the objectives driving the RAs' decision.

ESB might have suggested the possibility of using a larger increment to estimate the TLAF, as a better reflection of the decision actually facing the individual companies. However, there is no economic rationale for calculating a TLF for incremental output equal to the combined capacity (850 MW) owned by two, non-colluding companies. (See section 2.)

4. Volatility and Trends

BGE's presentation claimed that the current system is "volatile", "unpredictable", "penalising for investors" and (therefore) "unreliable as an investment signal". The ESB presentation did not consider the question of volatility, i.e. unpredictable variation in TLFs over time. However, the BGE and ESB presentations both implicitly raised a separate point, which was not brought out fully in the workshop, as to whether investors can hedge against future changes in TLAFs. This point was raised more clearly by representatives of the wind sector. Rather than suggesting a need to depart from the principle of calculating TLAFs from marginal losses, it suggests a need to offer system users some method of hedging against changes in TLAFs from year to year.

4.1. Background

The ESB presentation showed that the TLAF in Cork would be advantageous (i.e. would imply negative marginal losses) if there was only one plant in that location, because total system losses decline for each increment of output up to about 400 MW. Given output around that level, system losses are minimised and marginal losses are about zero.

Construction of the second plant must have changed the pattern of flows on the system, so that Cork switched from being an importing or balanced area (with an advantageous or zero TLAF) to being an exporting area (with a disadvantageous TLAF). The disadvantageous TLAF applies to both generators in this location.

4.2. Predictability of TLAFs for New Entrants

The BGE presentation suggests that the system should not penalise investors – and BGE in particular – for their apparent failure to foresee a decline in TLAFs caused by their investment. However, I would expect such an outcome to be broadly predictable using only simple information (on local output and consumption). ESB's ability to model losses for this presentation confirms that it can be done. Thus, the movement in TLAFs for Cork is not really an example of random volatility, but part of a predictable effect or trend which new entrants could have modelled or foreseen.

None of the RAs' objectives call for individual investors to be protected against the consequences of their own poor decision-making, particularly if the protection comes at the expense of investors who chose to locate plant in advantageous locations. In any case, the disadvantageous TLAF facing BGE may have been offset by a benefit in other locational costs, such as a lower cost of land, labour or other inputs. There is no reason to compensate an investor for a locational cost, whilst allowing the investor to capture a locational benefit.

4.3. Implications for Existing Investors

The wind generators (and ESB) have apparently suffered a deterioration in their TLAFs over time, because of decisions taken by later investors (including, probably, BGE's decision to locate plant in a disadvantageous location). The IWEA presentation showed a worsening of a TLAF between 2007 and 2010, which represents an unhedgeable risk for investors, and which has led to investors demanding that future projects provide a risk premium (as a cushion against similar risks). This risk merits further consideration, since it is affecting investment incentives. The problem for wind generators could be solved relatively simply, by making feed-in tariffs apply to output at the station gate, rather than output as delivered to the SEM (after adjustment for losses). In general, though, the RAs have simply opted for uniform TLAFs without considering whether other solutions would solve this problem more efficiently. I consider alternative methods of hedging in section 5.

4.4. Conclusion

The ESB presentation did not explicitly discuss any concept of volatility, since it analyses the predictable effect of increasing output from particular plants, not random effects. ESB's own work shows that the effect of constructing the BGE was susceptible to modelling, such that its effect on TLAFs was predictable. I am not aware of any reason why the RAs would wish to protect investors from the (adverse) consequences of their own decisions, particular adverse consequences that may have been offset by other benefits.

The impact of new entrants on the TLAFs facing existing investors is consistent with the need for short term efficient despatch. However, disquiet at the effects suggests a need for investors to be able to hedge against future changes in TLAFs. The RAs do not appear either to have identified this specific problem or to have considered alternative solutions to it.

5. Alternative Methods of Long-Term Hedging

Given the importance of deteriorating TLAFs for existing investors, the RAs should have considered the possibility of insulating past investors from the financial impact of decisions by later investors, even if generators remain exposed to TLAFs "at the margin". Such a system would exploit the efficient characteristics of a "two-part tariff" or "contractual property rights", which apply one rate to a fixed level of "inframarginal" output (to hedge risks) and the marginal rate to variations in output around that level (to preserve incentives for efficient decisions).

5.1. Hedging the TLAF

The aim of such schemes would be to protect investors as if they had signed a contract to fix their TLAF for a period of time, just as one can sign a contract that fixes the energy price for a limited period. The fixed TLAF in a contract for years 1 to n might be the TLAF actually applicable in year 0, or it could be defined by a formula that shifts the applicable value gradually from the year 0 TLAF to the current year's actual TLAF over a period of n years.[2] Other options are available and should be considered.

5.2. Preserving Incentives for Efficiency

As with an efficiently defined two-part tariff, it would be desirable for this hedge to apply to a fixed level of output, so that the current year's TLAF could apply to all decisions to vary output around that level. That fixed output might be defined by the outputs used when estimating TLAFs, for instance, or it might be a (declining) share of a plant's capacity or historical output.

Thus, consider a case where a generator locates in an area in which the TLAF is 1.00 and secures that rate for 200 MWh, but where the TLAF falls to 0.95 on construction of the generator, whose actual output varies between 300 MWh and 400 MWh. For the purpose of the SEM, the settlement mechanism would operate as shown in the table below.

To begin with, the generator would be credited with output at the current TLAF of 0.95. This would result in the generator contributing 15 MWh or 20 MWh for losses, depending on whether its output was 300 MWh or 400 MWh.

Next, the hedge would compensate the generator for the difference between a current TLAF of 0.95 and a hedged TLAF of 1.00, for a volume of 200 MWh, i.e. a credit of 10 MWh in both output scenarios.

The generator would be then credited for deliveries equal to the sum of its TLAF-adjusted output and the hedging adjustment, as shown in the last line of the table.

---

[2] If the scheme applies for n years, then in each year **y** (= 1 to **n**), the applicable TLAF would be given by the following formula:

*Contract TLAF for year y = (Year 0 TLAF * (**n-y**)/**n**) + (Current year TLAF * **y/n**).*

Applying this contract TLAF to a <u>fixed</u> volume (e.g. the output in the scenarios used to estimate TLAFs), whilst applying the current year TLAF to the difference between this volume and actual output, would ensure that current year TLAFs continued to provide an incentive for efficient decisions about output.

**Table 5.1:
The Impact of Hedging TLAFs**

Item		Output Scenario: Low	High	Difference	Row No.
Generation:					
Metered Output at the Station Gate	MWh	300	400	100	(1)
Current TLAF	*multiplier*	0.95	0.95		(2)
TLAF-Adjusted Output	MWh	285	380	95	(3)=(1)x(2)
Hedging					
Hedged Volume	MWh	200	200	0	(4)
Hedged TLAF	*multiplier*	1.00	1.00		(5)
Current TLAF	*multiplier*	0.95	0.95		(6)
Hedging Correction	*multiplier*	0.05	0.05		(7)=(5)-(6)
Hedging Adjustment	MWh	10	10	0	(8)=(4)x(8)
Settlement:					
Accredited Deliveries into SEM	MWh	295	390	95	(9)=(3)+(8)

Crucially, because the hedge covers a fixed volume, the generator is credited with an extra 95 MWh for increasing its output from 300 MWh to 400 MWh, just as it would have been if it faced the current TLAF without any hedge. As a result, the hedge does not harm incentives to generate efficiently.

5.3. Precedents

The creation of such hedging contracts in the SEM would not be the first case of arrangements intended to protect investors against volatility or trends in “real-time” transmission prices. At least three electricity markets in the United States offer retailers and other traders the opportunity to acquire contracts for hedging against variations in the difference between energy prices in different locations, i.e. the real-time cost of transmission.[3]

§ The PJM market offers multi-year contracts known as Financial Transmission Rights (FTRs); retailers acquire FTRs on the basis of an “endowment” of rights based on their historic peak demand and others may do so by taking part in auctions; in the latest round of auctions, the PJM sold off FTRs for the years 2011-2014;[4]

§ New England ISO offers annual FTRs, both by auction and in proportion to an endowment of retailers’ rights;[5]

---

3 These arrangements focus on retailers as the major drivers of investment in generation capacity; other markets might view generation as independent of retailers. The arrangements need to be centralised in markets where energy prices differ by location, because only the body receiving the revenue from real-time charges for transmission can offer a hedge against variations in the cost of real-time charges for transmission (without increasing their risk profile). Other players would merely be taking over an unhedgeable risk from generators.

4 http://pjm.com/markets-and-operations/ftr.aspx

5 http://www.iso-ne.com/support/faq/ftr/index.html

§ The New York Independent System Operator has just completed a filing for a system of annual "Transmission Congestion Contracts" assigned by reference to historic load flows and also for non-historic load flows.[6]

5.4. Implications for the Decision-Making Process

Appraisal of such hedging schemes might show them to be desirable by the standard of many the high-level objectives, but difficult to design or to implement in practice. Consideration of such a scheme in the abstract might nevertheless affect the evaluation of other options.

The RAs rejected some options (e.g. compression, averaging, etc) by comparison with TLAFs based on annual estimates of marginal losses and their impact on short term efficiency in dispatch. However, these same options might prove to be a close proxy for longer term TLAF contracts, and to offer longer term efficiency benefits. That would imply they had some merit that had been overlooked so far.

5.5. Conclusion

The presentations at the workshop indicated a concern over long-term trends in TLAFs, which are not adequately described by the word "volatility" nor necessarily solved most efficiently by imposing some arbitrary but stable level (such as uniform losses). There exist several possible methods of providing a longer term hedge. These methods offer both new options to be evaluated and a different standard by which to assess existing options. Despite the importance of this concern, the RAs do not seem to have considered such hedging schemes, either as distinct options or as a potential frame of reference for evaluating the existing options.

---

[6] See FERC Docket ER07-521-009, Order issued 15 July 2010 at: http://www.ferc.gov/whats-new/comm-meet/2010/071510/E-12.pdf.

NERA
Economic Consulting

NERA Economic Consulting
15 Stratford Place
London W1C 1BE
United Kingdom
Tel: +44 20 7659 8500
Fax: +44 20 7659 8501
www.nera.com

NERA UK Limited, registered in England and Wales, No 3974527
Registered Office: 15 Stratford Place, London W1C 1BE