The energy performance coefficient – a robust indicator

Vilnis Vesma MA(Oxon) CEng MEI CMVP CEM

1. Introduction

The purpose of an energy performance indicator (EnPI) is to allow progress in energy management to be reported. An EnPI is a number whose value changes through time, summarising the relative improvement or deterioration in the energy efficiency of the building, process, enterprise, or other object to which it is applied.

Although it is commonly claimed that EnPIs are 'simple and easy to understand', I will argue that this is not generally true, that some traditional forms of EnPI are misleading, and that in some cases they cannot be calculated at all in a meaningful way. I will describe a methodology which overcomes these objections and provides an indicator which is far more meaningful and far more widely applicable, but remains simple to present. My aim (and it should be the aim of everyone reporting on energy performance) is to find an indicator which has these essential attributes:

1. It responds only to changes in energy performance, and is not affected by factors such as the weather, product throughputs, or other potentially distorting influences; and
2. The direction and magnitude of change are consistent with, and proportionate to, the change in energy performance.

I might add that it is highly beneficial if the chosen indicator can survive non-routine changes to the monitored object, such as the addition or removal of production capacity and the acquisition, disposal, or remodelling of buildings.

2. Examples of the problem

2.1 energy-intensive manufacturing

This example is taken from an energy-intensive process. If consists of monthly electricity consumption figures and corresponding production data (the figures used are reproduced in the appendix).

When monthly consumption is plotted against production output, as in Figure 1, they are seen to have a simple linear relationship. A best-fit straight line has been superimposed to represent typical performance.

I have chosen two particular months and we will compare their performance using specific energy ratios (SER) in MWh/tonne. This is a very common way of reporting energy performance.

	MWh	Tonne	SEC
January 2008	95.1	118.2	0.805
February 2009	45.0	51.7	0.870

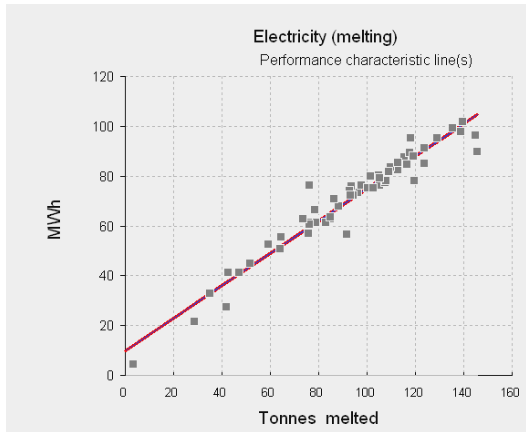

Figure 1

On this analysis, February 2009 seems to be worse, since more electricity was used per tonne of production (0.870 compared with 0.805 MWh/tonne).

In fact, when the two months are identified on the scatter diagram (Figure 2) January 2008—the higher point—can be seen to be the worse performer because consumption that month was significantly above the line of typical performance. February's consumption on the other hand was almost exactly what it should have been, despite returning a higher specific energy ratio. The reason for its higher SER is that fact that the 10 MWh of fixed monthly demand—represented by the intercept on the vertical axis—is divided by only 51.7 tonnes of production, rather than the 118.2 tonnes that applied in the January.

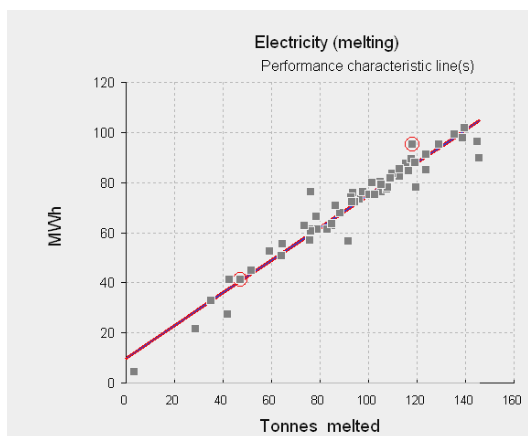

Figure 2

This simple numerical example illustrates the weakness of using a simple specific energy ratio as a performance indicator. In the built environment similar considerations will arise with seasonal weather variations. Although the use of annual totals may go a long way towards mitigating these distortions, firstly it does not completely eliminate the effect and secondly it introduces unacceptable delay in the evaluation.

Simple specific energy ratios have a further fatal weakness: they can only be computed if there is a single factor driving variation in consumption. In any other real-life scenario, they are impossible to compute at all.

2.2 Computer data centres

Computer data centres present an unusual problem. Their electricity uses are broadly threefold: energy delivered to the IT equipment in conditioned rooms; energy used for cooling; and overhead consumption for offices, lighting and general power. The amount of energy required for cooling depends primarily on the quantity delivered to the IT equipment in the conditioned space, since every kilowatt-hour of IT consumption represents a kilowatt-hour of heat to be pumped out. However, the weather usually has some influence, with cooling demand being higher in hot weather than in cold.

The industry-standard way of expressing data-centre energy efficiency is 'Power Utilisation Effectiveness' (PUE), the ratio of total electricity to that consumed in the IT equipment. Were there no cooling or overhead loads, PUE would be unity, but in practice a value of, say, 1.6 might be more commonly achievable. Three things can improve the PUE. One is cold weather (which reduces cooling demand); the second is an increase in IT load (which makes the overhead a smaller proportion of the total); and the third is energy-efficiency measures. Because PUE responds to all three influences, it cannot be trusted as a pure energy performance indicator.

3. Towards a solution

The problem with a specific energy ratio (e.g. MWh/tonne) when used as a measure of performance is that there is no unique target value which can be set. Even in a simple case like the one illustrated above in section 2.1, the target changes according to product throughput. In the data-centre case, as in many other real-life installations, there are two factors that need to be taken into account.

A solution can be inferred from ISO 50001 Section 4.6.1 where bullet-point (e) calls for evaluation of actual against ***expected*** consumption. The relationship between these two values is likely to be far more meaningful.

The straight-line relationship which we observed earlier can be used to illustrate the principle. From the data given in the appendix it is possible to derive a formula for expected consumption E:

E = 10 MWh per month + 0.65 MWh per tonne

Thus for example if the production output were 100 tonnes, the expected consumption would be 10 + (100 x 0.65) = 75 MWh, as illustrated in Figure 3.

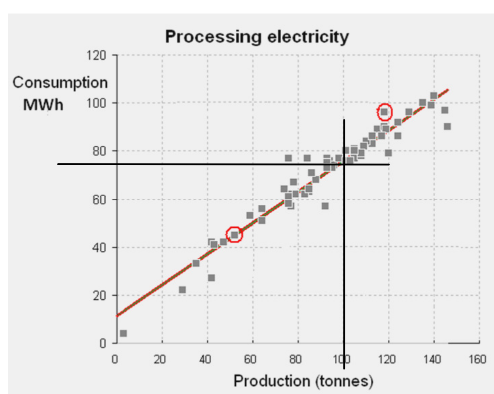

Figure 3

Likewise for a data centre, where there are two influences on cooling energy demand, one might have a formula for expected weekly cooling energy consumption E_c of this form:

$$E_c = a + b.E_{IT} + c.CDD$$

Where *a* represents the fixed weekly consumption, *b* and *c* are constants, E_{IT} is the energy delivered to IT equipment, and CDD is the number of cooling degree days. Both E_{IT} and CDD are variables that can be independently measured.

4. The Energy Performance Coefficient

The **ratio of actual to expected** energy consumption turns out to be a very well-behaved and versatile measure. I call it the Energy Performance Coefficient[1] or EnPC. A value of unity is neutral: it signifies that consumption is the same as would historically have been expected, given the current prevailing conditions. Greater than one signifies poor performance—higher consumption than expected—while a value less than 1.0 indicates improved energy performance.

This form of performance indicator is not affected by changes in the factors which routinely affect consumption, as long as we have taken them into account when calculating expected consumption. It only responds to changes in energy performance, which is one of the things we need to achieve.

Being unaffected by, for example, seasonal changes in the weather or levels of business activity, the EnPC is stable through the year when evaluated over shorter intervals. This makes it possible to monitor continuously against a fixed target, getting early warning of problems or confirmation of success. Any energy-using system, regardless of the number of factors known to influence its consumption, can be assigned a fixed target value for its performance indicator that will be the same over any interval: weekly, monthly, or annual.

EnPCs can be used in any application—buildings, industrial processes and vehicles—and can be calculated at any level: individual component; system; site; or enterprise. Moreover, it is possible to aggregate components' EnPCs to give a grand total for the installation as a whole, as follows.

[1] *This replaces the previous term 'Energy Intensity Coefficient'*

To take a paper mill as an example, one could track individual disaggregated EnPCs for each major subsystem:

Subsystem	Commodity	Driving factors	Method of calculating expected consumption
Paper Machine	Electricity	Gross production	Straight-line relationship
Paper Machine	Steam	Mass difference between stock feed and gross paper production including broke	Straight-line relationship
Cooling towers	Electricity	Outside air temperature and Vacuum-pump run hours	Multi-variate model
Broke plant	Electricity	Pulp output from broke plant	Straight-line relationship
Wood yard	Electricity	Timber processed	Straight-line relationship
Pulp plant	Electricity	Pulp produced (split between different grades)	Multi-variate model
Boilers	Total fuel	Steam produced	Straight-line relationship
Buildings	Heating fuel	Weather (expressed as degree days)	Straight-line relationship
Buildings	Electricity	None	Assumed constant

Each of these nine separately-metered consumption streams can have its own individual actual consumption a and expected consumption e which combine as follows to give an overall coefficient:

$$EnPC = \frac{a_1 + a_2 + ... + a_9}{e_1 + e_2 + ... + e_9}$$

This approach makes it relatively easy to deal with non-routine changes. Suppose, for example, that a new paper machine is added. It will increase demand for steam and electricity, but both requirements can be modelled and the formula for the coefficient will become:

$$EnC = \frac{a_1 + a_2 + ... + a_9 + a_{10} + a_{11}}{e_1 + e_2 + ... + e_9 + e_{10} + e_{11}}$$

If sub-units are removed, it is not even necessary to modify the formula. Both the a and e terms for the removed equipment fall to zero. The effect on the overall indicator is what you would hope: for example if the sub-unit that is removed was a poor performer, the EnPC will fall, indicating an improvement. If one removes something that had enjoyed improving energy performance, the coefficient will rise.

The procedure for dealing with non-routine changes is thus likely to be easy to define and document, and just as importantly, their effect on the overall coefficient is not disruptive. The inclusion of new plant operating as expected will simply tend to move the coefficient closer to unity.

6. Integration with existing management reporting

Although the Energy Performance Coefficient is superior to most of the alternative commonly accepted indicators (such as specific energy ratio, SER) it will always be difficult to introduce it, partly because doing so is an admission that previous reporting practices were wrong, and partly because the EnPC has a different numerical value unfamiliar to management. Fortunately, there is a compromise which enables the new methodology to be introduced in a subtle way. The procedure is as follows:

1. Choose a value for the existing indicator, S, which will be fixed as the 'baseline' value S_{base}.
2. In each subsequent reporting period, evaluate the current EnPC (call it $C_{current}$)
3. Report the product (S_{base} x $C_{current}$) as the *adjusted* performance indicator.

For example, suppose an enterprise has an initial crude specific energy ratio (SER) of 560 kWh/tonne, and that energy consumption is known to be affected by the weather as well as by variation in output. As explained in the body of this paper, both factors would be used to calculate expected consumption and thus the EnPC. Then:

- In year 0, the base year, the EnPC will be 1.0 (by definition).
- Suppose that in year 1 the EnPC was still 1.0 (there having been no improvement yet) but low throughput had pushed crude SER up to 577. This is misleading, because energy performance had not changed, and this is evident in the adjusted SER, which is simply 560 x 1.0 = 560 kWh/tonne
- Let us say that in year 2 they made some progress with energy performance and their EnPC fell to 0.9, while mild weather and high volumes had brought their crude SER down to 498 kWh/tonne. By the proposed method, the adjusted SER was 560 x 0.9 = 504 kWh/tonne, which does not look so good but is a more accurate reflection of their actual 10% improvement in energy performance.
- Finally in year 3 let us say that they were hit by bad weather and low volumes which took crude SER up to 530, despite further energy performance improvements which had taken their EnPC down to 0.8. The truth is reflected in the *adjusted* SER: 560 x 0.8 = 448 kWh/tonne.

So it can be seen that multiplying the **baseline indicator** by the **current EnPC** yields an **adjusted indicator** which looks similar to the enterprise's customary measure of energy performance, but which actually eliminates the distorting effects of all known external influences.

Acknowledgments

- Four projects have enabled the author to crystallise and refine the principles embodied in this paper. He wishes to acknowledge the following without whom it would not have been possible:
  - Chad Gilless, Global Energy Partners
    - Commissioned input to a white paper on energy baselines in industry
  - Jigar Shah, Institute for Industrial Productivity
    - Invited the author to advise the Chinese National Institute for Standardisation on performance evaluation in the oil refining and iron and steel industries
  - Michel de Laire, Chilean Energy Efficiency Agency
    - Commissioned the author to advise on performance evaluation in three ISO 50001 pilot projects
  - Boris Sučić, Institut Jožef Stefan, Ljubliana
    - Independently developed the same idea and persuaded the author to adopt his (superior) terminology “Energy Performance Coefficient”

Vilnis Vesma
Degree Days Direct Limited
Pound House
Market Square
Newent
Glos GL18 1PS

Vilnis@VESMA.COM

+44(0)1531 821350

Revision 3
5 June 2014

Appendix: data set used in the first example

These data have been taken from a real manufacturing plant, but for reasons of commercial confidentiality the numbers have been disguised. Both consumption and output figures have been multiplied by undisclosed factors and the sequence of months has been slightly changed.

For the purposes of analysis, July and August 2005 have been disregarded because of unreliable reporting of results.

None of these changes has a material effect on the conclusions drawn.

Month	Year	MWh	Tonne
4	2004	55.4	64.4
5	2004	50.7	64.1
6	2004	62.7	84.9
7	2004	4.3	3.4
8	2004	52.7	59.4
9	2004	61.4	83.1
10	2004	56.7	91.7
11	2004	61.2	79.0
12	2004	61.4	75.9
1	2005	60.6	76.2
2	2005	67.8	88.2
3	2005	74.8	93.9
4	2005	70.7	86.3
5	2005	79.7	101.4
6	2005	74.2	92.9
7	2005	0.0	0.0
8	2005	76.5	84.0
9	2005	73.2	96.4
10	2005	76.3	105.2
11	2005	76.1	97.7
12	2005	72.2	94.5
1	2006	75.9	93.4
2	2006	74.0	92.9
3	2006	82.6	112.7
4	2006	66.4	78.5
5	2006	75.3	100.5
6	2006	80.3	104.9
7	2006	27.2	41.7
8	2006	62.9	73.6
9	2006	72.3	93.3
10	2006	87.7	115.5
11	2006	91.1	123.8
12	2006	77.3	107.5
1	2007	83.5	109.7
2	2007	81.7	109.1
3	2007	85.3	113.0
4	2007	78.2	107.9
5	2007	89.3	117.7
6	2007	96.2	145.0
7	2007	41.3	42.2
8	2007	56.9	75.7
9	2007	79.0	105.3
10	2007	101.9	139.8
11	2007	99.3	135.5
12	2007	75.3	102.8
1	2008	95.1	118.2
2	2008	87.9	119.1
3	2008	84.7	116.7
4	2008	97.8	138.7
5	2008	95.3	128.9
6	2008	77.9	119.8
7	2008	32.7	35.2
8	2008	63.5	84.8
9	2008	85.0	123.8
10	2008	89.6	145.7
11	2008	76.1	76.3
12	2008	21.5	28.7
1	2009	41.1	42.8
2	2009	45.0	51.7
3	2009	41.3	47.1