Enhancing Electrical Supply by Pumped Storage in Tidal Lagoons

David J.C. MacKay
Cavendish Laboratory, University of Cambridge
`mackay@mrao.cam.ac.uk`

March 13, 2007 – Draft 1.8 – first published 5/3/07

Summary

The principle that the net energy delivered by a tidal pool can be increased by pumping extra water into the pool at high tide or by pumping extra water out of the pool at low tide is well known in the industry. On paper, pumping can potentially enhance the net power delivered by a factor of about four. However, pumping seems generally to be viewed as a minor optional extra, delivering only a modest power enhancement. Two possible reasons why pumping is not emphasized in tidal designs are that increasing the vertical water range introduces additional costs (for example, higher walls), and that alternating between pumping and generating worsens the intermittency-of-supply problem from which simple tide pools suffer.

The intermittency-of-supply problem also causes problems for wind. How can we switch to wind power if the wind might stop blowing for two days at a time? Chemical or kinetic-energy storage systems are an economical way to smooth out the fluctuations of wind power on a time-scale of minutes, but what about hours and days?

Perhaps a shift of perspective on tidal lagoons is helpful. I sketch designs for a large pumped-storage system located at sea-level with a dual purpose: first, it can turn power that is poorly matched to demand into high-value demand-following power; and second, it can simultaneously serve as a tidal power station. Large designs with a capacity of several gigawatts are the most economical.

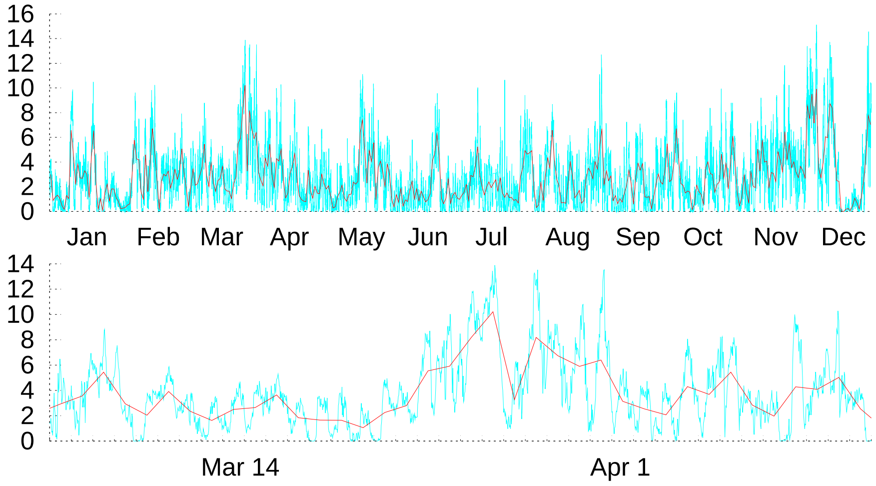

Figure 1. Cambridge mean wind speed in metres per second, daily (heavy line), and half-hourly (light line) during 2006. The lower figure shows detail from the upper. Thanks to Digital Technology Group, Computer laboratory, Cambridge This weather station is on the roof of the Gates building, roughly 10 m high. Wind speeds at a height of 50 m are usually about 25% bigger.

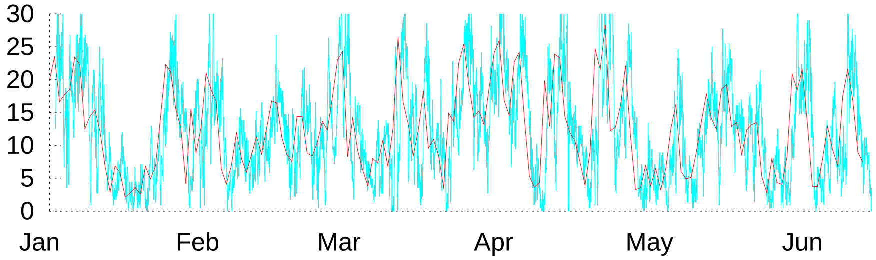

Figure 2. Cairngorm mean wind speed in metres per second, daily (heavy line), and half-hourly (light line), during six months of 2006. Thanks to Heriot–Watt University Physics Department.

Storage and wind

Offshore wind farms deliver, on average, about 3 W per m^2 of sea-floor area (or $3\,\text{MW/km}^2$, if you prefer).

Imagine that Britain had '30 GW' of wind farms – fifteen times as much as today. I put quotes round '30 GW' to emphasize that the nominal capacity of wind farms is much bigger than the average power delivered. The standard 'capacity factor' in the UK wind industry seems to be 1/3, so 30 GW of wind farms would be expected to deliver, on average, 10 GW.

Winds fluctuate (figures 1, 2). So this average of 10 GW would be delivered burstily: 30 GW one hour, and 0 GW the next, on one day; and perhaps 0 GW all day on the following day. How can such bursty power be made useful to society?

The default approach is to build back-up stations using some other sort of power – most likely fossil fuel – which sit idle when the wind blows, and are switched on when it does not, or when demand peaks. Another approach would be to manage demand – using smart electric-car chargers, for example, which use electricity when it is cheap; or running the Aluminium plant and the water-purification factory only when the wind blows. A third approach is storage. The storage required to deliver 10 GW for 24 hours is

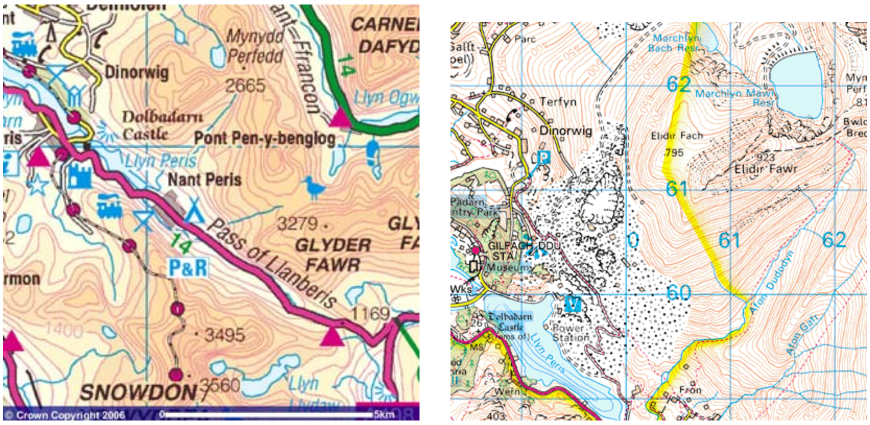

Figure 3. Dinorwig, in the Snowdonia National Park. The left map is a 10 km by 10 km area. In the right map the blue grid is made of 1 km squares. Dinorwig is the home of a 9 GWh storage system, using Marchlyn Mawr (615E, 620N) and Llyn Peris (590E, 598N) as its upper and lower reservoirs. Images produced from Ordnance Survey's Get-a-map service `www.ordnancesurvey.co.uk/getamap`. Images reproduced with permission of Ordnance Survey. © Crown Copyright 2006

240 GWh – twenty-six times as big as the 9 GWh of Dinorwig.

The Scottish island of Fair Isle (population 70, area 5.6 km^2) has pioneered several of these technologies. To solve the demand-management problem, Fairisle has for over 25 years had *two* electricity networks that distribute power from two wind turbines and, if necessary, a diesel electric generator. Standard electricity service is provided on one network, and electric heating is delivered by a second set of cables. The electric heating is mainly served by excess electricity from the turbines that would otherwise have had to be dumped. Remote frequency-sensitive programmable relays control individual water heaters and storage heaters in the individual buildings of the community. In fact there's up to six frequency channels per household, so the system behaves like seven networks. Fair Isle also successfully trialled a kinetic energy storage system (a flywheel) to store energy during oscillations of wind strength (with a period of 12 to 20 seconds).

Designs for multi-purpose storage/tidal systems

Key ideas for an energy-enhancing pumped-storage system:

1. It is said that connecting large numbers of wind turbines to the national electricity grid could lead to instabilities. We thus propose decoupling wind turbines from the grid, *plugging them directly into pumped storage systems instead.* The wind-to-pump connection could be a flexible grid with much wider tolerances than the national network.

2. The pumped storage system is located in a region with large

tides. Water is pumped to and from the sea in such a way that (a) the power delivered can respond to the grid's demand, eliminating problems of intermittency; and (b) we get more power out than we put in. (Yes, I mean that the energy delivered when generating *exceeds* the energy received – in contrast to Dinorwig, which has a round-trip efficiency of about 75%.)

3. When the demand for pumped storage is low (during a few calm days, say), the facility can also function as a stand-alone tidal power station. By using multiple lagoons, it's possible to turn the intrinsically intermittent tidal power into always-on, demand-following capacity.

4. The facility could also buy electricity from the national grid for pumped storage, just like Dinorwig.

In sum, it's a storage system that is more than 100% efficient. It's a storage system that can also produce its own power when it's not needed for storage. Or, it's a tidal facility that still provides a valuable function even when the tides are small.

Rough models

Let's assume a tidal range of $2h = 4\,\mathrm{m}$ throughout. I'll also assume that hydroelectric generators have an efficiency of 90% and that pumps have an efficiency of 85%. (These figures are based on the pumped storage system at Dinorwig, whose round-trip efficiency is about 75%. I am not sure what the best figures are for low-head tidal turbines. In their paper based on La Rance, Shaw and Watson [2003a] assume pumping efficiencies up to 66%, with best efficiency at large head, and generating efficiency 80%.)

Let's start by finding some benchmarks for energy production.

Production on ebb and flow (no pumping, no demand-following)

THE POWER OF AN ARTIFICIAL TIDE POOL. To estimate the power of an artificial tide pool, imagine that it's filled rapidly at high tide, and emptied rapidly at low tide. Power is generated in both directions. The change in potential energy of the water, each six hours, is mgh, where h is the change in height of the centre of mass of the water, which is half the range (figure 4). The mass per unit

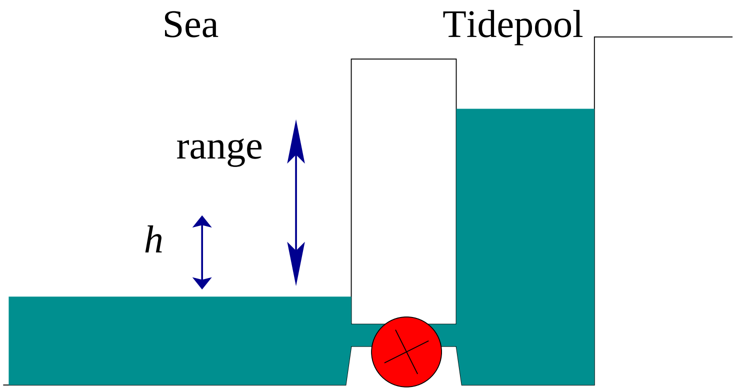

Figure 4. An artificial tide pool. The pool was filled at high tide, and now it's low tide. We let the water out through the electricity generator to turn the water's potential energy into electricity.

land-area covered by tide-pool is $\rho \times (2h)$, where ρ is the density of water ($1000\,\mathrm{kg/m^3}$). So the power per unit area delivered by a tide pool is

$$\frac{2\rho h g h}{6\,\text{hours}}$$

Plugging in $h = 2\,\mathrm{m}$, we find

$$\text{Power per unit area of tide-pool} = 3.6\,\mathrm{W/m^2}.$$

Allowing for an efficiency of 90% for conversion of this power to electricity, we get

$$\text{Power per unit area of tide-pool} = 3.3\,\mathrm{W^{(e)}/m^2}.$$

(Or $3.3\,\mathrm{MW/km^2}$.)

Tidal pools with pumping

The pumping trick artificially increases the amplitude of the tides in the tidal pool so as the amplify the power obtained. The energy cost of pumping in extra water at high tide is repaid with interest when the same water is let out at low tide; similarly, extra water can be pumped out at low tide, then let back in at high tide. Let's work out the theoretical limit for this technology.

I'll assume that generation has an efficiency of $\epsilon_{\mathrm{g}} = 0.9$ and that pumping has an efficiency of $\epsilon_{\mathrm{p}} = 0.85$.

Let the tidal range be $2h$. I'll assume that the prices of buying and selling electricity are the same at high tide and low tide, so that the optimal height boost b to which the pool is pumped above high water is given by (marginal cost of more pumping = marginal return of water):

$$b/\epsilon_{\mathrm{p}} = \epsilon_{\mathrm{g}}(b + 2h)$$

Defining the round-trip efficiency $\epsilon = \epsilon_{\mathrm{g}}\epsilon_{\mathrm{p}}$, we have

$$b = 2h\frac{\epsilon}{1-\epsilon}$$

For example, with a tidal range of $2h = 4\,\mathrm{m}$, and a round-trip efficiency of $\epsilon = 76\%$, the optimal boost is $b = 13\,\mathrm{m}$.

Let's assume the complementary trick is used at low tide. (This requires that the basin have a vertical range of 30 m!) The delivered power per unit area is then

$$\left(\frac{1}{2}\rho g \epsilon_{\mathrm{g}}(b+2h)^2 - \frac{1}{2}\rho g \frac{1}{\epsilon_{\mathrm{p}}} b^2\right) \Big/ T,$$

where T is the time from high tide to low tide. We can express this as the power without pumping, scaled up by a boost factor

$$\left(\frac{1}{1-\epsilon}\right),$$

which is a factor of about 4.

Tidal amplitude h (m)	Optimal boost height b (m)	Power with pumping (W/m^2)	Power without pumping (W/m^2)
0.5	3.3	0.9	0.2
1.0	6.5	3.5	0.8
2.0	**13**	**14**	**3.3**
3.0	20	31	7.4
4.0	26	56	13

Unfortunately, this pumping trick will rarely be exploited to the full because of the economics of basin construction: full exploitation of pumping requires the total height of the pool to be roughly 4 times the tidal range, and increases the delivered power by a factor of 4. But the material in a sea-wall of height H scales as H^2, so presumably the cost of constructing a wall four times as high will be more than four times as great. Extra cash would probably be better spent on enlarging a tidal pool horizontally rather than vertically.

The pumping trick can nevertheless be used for free whenever the natural tides are smaller than the maximum tidal range. The next table gives the power delivered if the boost height is set to h, that is, the range in the pool is just double the external range.

Tidal amplitude h (m)	Boost height b (m)	Power with pumping (W/m^2)	Power without pumping (W/m^2)
0.5	0.5	0.4	0.2
1.0	1.0	1.6	0.8
2.0	**2.0**	**6.3**	**3.3**
3.0	3.0	14	7.4
4.0	4.0	25	13

A doubling of vertical range is plausible at neap tides, since neap tides are typically about half as high as spring tides. Pumping the pool at neaps so that the full springs range is used thus allows neap tides to deliver roughly twice as much power as they would offer without pumping. So a system with pumping would show two-weekly variations in power of just a factor of 2 instead of 4.

These benchmarks – **3.3** W/m^2 without pumping and **6.3** W/m^2 with pumping – assume that power is delivered and demanded at exactly the optimal times, and that there is no limit to the flow rate of water in the system. Such a system is highly intermittent and spikey. We now examine a more reasonable, smooth, but still intermittent, benchmark.

An intermittent solution that alternates between steady pumping and steady generating

Figure 5 shows a pumping and generating schedule where the system is always active; it spends exactly half the time pumping (at constant power) and half generating (at constant power). The system alternately sucks 7 W/m^2 from the electricity grid (for three hours) and delivers 20 W/m^2 (for three hours). The net energy contribution is thus 6.5 W/m^2. The range required is about 10 m – slightly more than double the tidal range of 4 m.

Multiple-lagoon solutions

Using multiple pools – for example, a high pool and a low pool – doesn't increase the deliverable power, but does increase the flexibility of when power can be delivered, thus enhancing the value of a facility. A two-pool facility is 'always on', and would be able to provide the same sort of valuable service as the Dinorwig station.

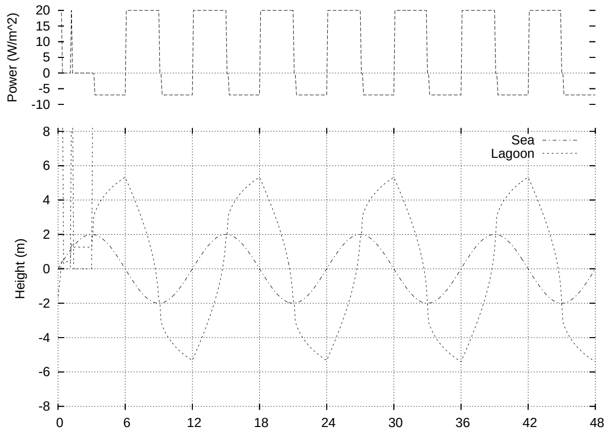

Figure 5. A bursty tidal power option using one lagoon at sea-level. The tidal range is $2h = 4\,\mathrm{m}$. The system alternately sucks $7\,\mathrm{W/m^2}$ from the electricity grid (for three hours) and delivers $20\,\mathrm{W/m^2}$ (for three hours). The net energy contribution is thus $6.5\,\mathrm{W/m^2}$.

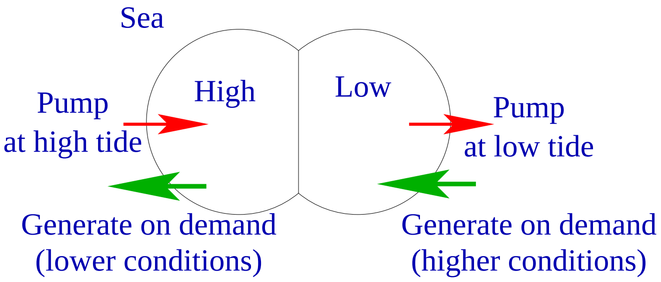

Figure 6. Design assumed: one, two, or three lagoons are located at sea-level. While one lagoon is being pumped full or pumped empty, the other lagoon may be delivering steady, demand-following power to the grid. Pumping may be powered by bursty sources such as wind, by spare power from the grid (say, in the future, from nuclear power stations), or by the facility itself, using one lagoon's power to pump the other lagoon to a greater height.

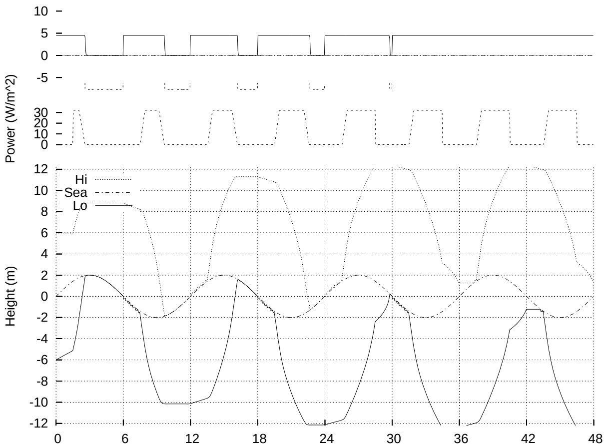

Figure 7. A two-lagoon system with no power input from the grid or wind. The tidal range is $2h = 4\,\mathrm{m}$. Self-pumping takes place with a power of $32\,\mathrm{W/m^2}$. After an initial set-up period of a couple of periods, the system delivers a steady $4.5\,\mathrm{W/m^2}$.
Top graph: solid line – power delivered to grid.
Second graph: self-pumping power.

A two-pool facility can do its own pumping.

Thus a tidal station can turn the intermittent tidal power into demand-following power. As an extreme simple case, let's assume that demand is absolutely steady. Not realistic, but a challenging target for a renewable source to deliver!

Figure 7 shows a possible schedule for a two-lagoon system. In contrast to the single-lagoon schedule, where pumping periods and generating periods alternate, each lasting 3 hours (with switches from pumping to generating at high tide and low tide), here generating happens all the time and pumping lasts for three hours around each high tide and three hours around each low tide. One lagoon's water level is always above sea-level; the other's is always below.

In this figure, the pumping into or out of one lagoon is entirely funded by the energy in the other lagoon. No energy is required from the grid. After an initial set-up period of a couple of periods, the system delivers a steady $4.5\,\mathrm{W/m^2}$. The range is about $25\,\mathrm{m}$ (about six times the tidal range).

The same facility can simultaneously be used for pumped storage.

For simplicity and clarity, I again assume that the demand is steady. I also assume in the computations that the power being stored is steady, but the system would work equally well if the incoming power fluctuated around its average value on a timescale of minutes or one or two hours.

Figure 8 shows the result of using the same schedule, doing self-pumping for three hours around each high and low tide, plus pumping $5.5\,\mathrm{W/m^2}$ of 'bursty' wind power into the appropriate la-

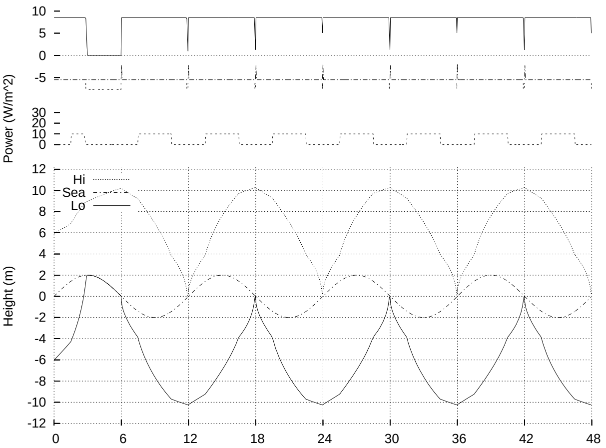

Figure 8. A two-lagoon system receiving $5.5\,\mathrm{W/m^2}$ of bursty wind power and delivering $8.5\,\mathrm{W/m^2}$ of steady power. The tidal range is $2h = 4\,\mathrm{m}$. Self-pumping takes place with a power of $10\,\mathrm{W/m^2}$.
Top graph: solid line – power delivered to grid; dashed line – average power received from intermittent source, e.g. wind.
Second graph: self-pumping power.

goon all the time. The system is generating a steady $8.5\,\mathrm{W/m^2}$. The range is roughly 20 m (five times the tidal range).

The facility could also be used for pumped storage alone.

Perhaps pumps and generators are a valuable resource and none are available for the self-pumping trick. Figure 9 shows results for two incoming power conditions. In (a), $5.5\,\mathrm{W/m^2}$ of 'bursty' wind power is turned into $7.5\,\mathrm{W/m^2}$ of steady power; the range between the high pool's maximum and the low pool's minimum is 16 m. In (b) $18\,\mathrm{W/m^2}$ of 'bursty' wind power is turned into $19\,\mathrm{W/m^2}$ of steady power; the range required is about 26 m.

Other designs, future work

The next design I would like to explore uses three lagoons. The two-lagoon solution (self-pumping) doesn't deliver as much power per unit area, and required larger vertical amplitudes, than the one-lagoon solution with externally-funded pumping. I expect that there are various three- or four-lagoon solutions in which one or two of the lagoons follow trajectories like that of the one-lagoon solution, with most or all of the required pumping funded internally.

An obvious piece of further work is to explore the economics of realistic daily supply and demand inputs. It's possible that the economically optimal pumping and generating strategy might sometimes be to exploit just one high tide and one low tide for pumping each night, and generate at appropriately selected times in the day.

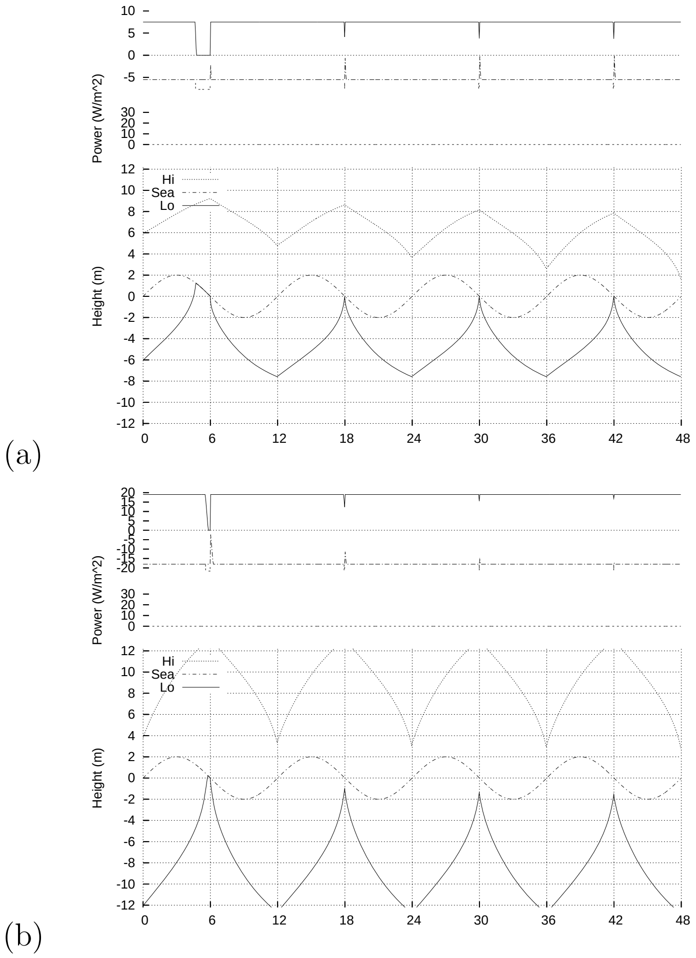

Figure 9. In these simulations, no self-pumping takes place. (a) A two-lagoon system receiving $5.5\,\mathrm{W/m^2}$ of bursty wind power and delivering $7.5\,\mathrm{W/m^2}$ of steady power. (b) A two-lagoon system receiving $18\,\mathrm{W/m^2}$ of bursty wind power and delivering $19\,\mathrm{W/m^2}$ of steady power. The tidal range is $2h = 4\,\mathrm{m}$. Top graph: solid line – power delivered to grid; dashed line – average power received from intermittent source, e.g. wind.

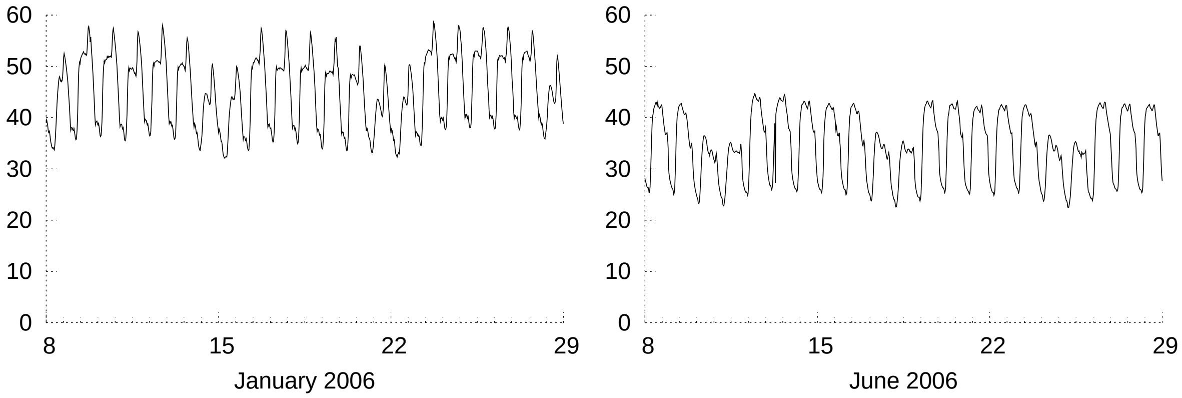

Figure 10. Electricity demand in Great Britain (in GW) during three winter weeks and three summer weeks of 2006.

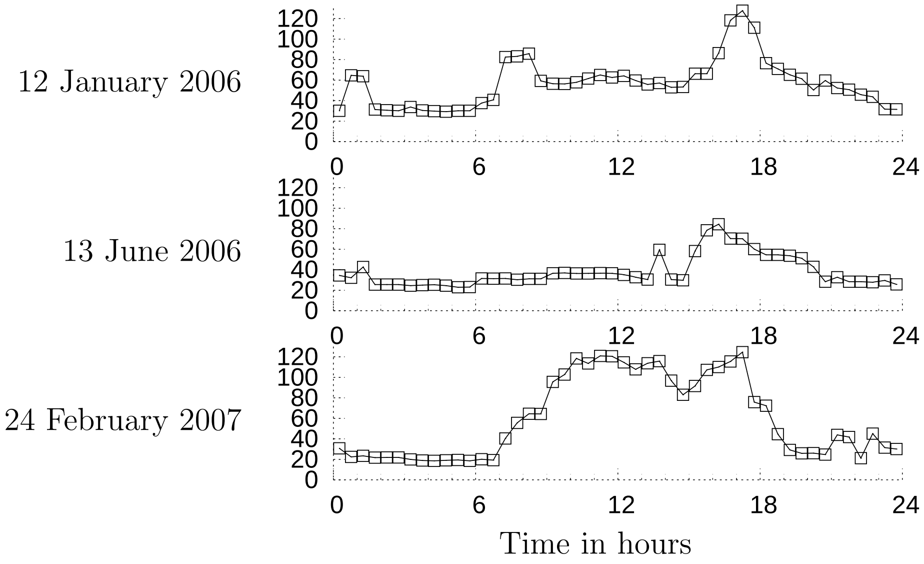

Figure 11. Electricity prices in Great Britain (in £ per MWh) on three days in 2006 and 2007.

To reduce costs associated with high walls we could look for lopsided schedules where lagoons are pumped down to lower extremes and not pumped up so high.

Criticisms: I've ignored the true dependence of generating and pumping efficiency on head. I've assumed the sea is an inexhaustible source or sink of water at the current sea-level. Once the system reaches a sufficiently large size, its sucking and blowing will have a significant effect on local sea-level.

I've not taken account of the cost of turbines, assuming that we can install whatever pumping and generating capacity these schedules call for.

Discussion

Some of these ranges are enormous. Where could such a system be put? What would it cost, and what would it be worth?

One simple observation is that the value delivered scales as the area of the lagoons, but the dominant part of the cost – the walls – scales as the circumference. Very large systems are thus favoured by simple economics.

Let's pick a benchmark size. How about 10 km × 10 km?

- A plain old intermittent tide-pool of this size (3.3 MW/km^2) would deliver 330 MW on average (assuming, as usual, a 4 m tidal range).

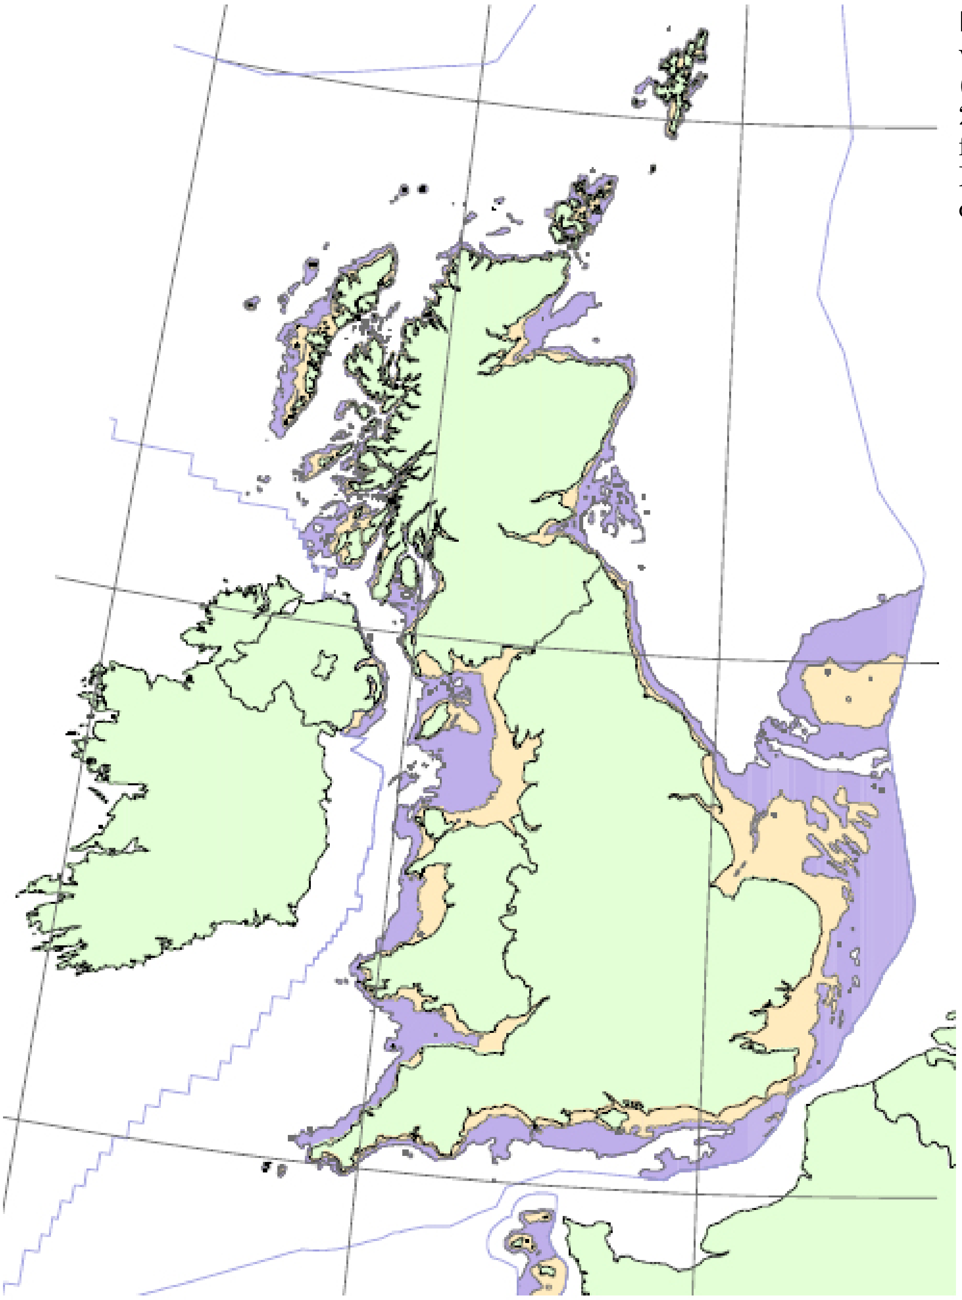

Figure 12. UK territorial waters with depth less than 25 m (yellow) and depth between 25 m and 50 m (magenta). Data from DTI Atlas of Renewable Marine Resources. © Crown copyright.

- With pumping to 5 m above and below mean sea-level, a single lagoon would deliver a net power of 650 MW.
- The two-lagoon solution that does its own pumping would deliver a **steady** 450 MW.
- It could also be used as a pumped storage system for intermittent or unwanted electricity. For example, it could turn 550 MW (average) of bursty wind power into 750 MW of steady power. (A round-trip efficiency of 135%!) Or it could turn 1.8 GW of bursty wind power into 1.9 GW of steady power. A round-trip efficiency of 105% compares favourably with Dinorwig's 75%.
- For comparison of pumped storage capacity with Dinorwig – a profitable power station with a capacity to store 9 GWh – these tide pools would have a capacity of about 20 GWh (assuming a height change of 13 m). So this facility would be worth two Dinorwigs. Indeed, it would be worth more, since it would be better than 100% efficient, in contrast to Dinorwig's 75%.

We need at least half of this water to have a depth of about 13 m below mean sea-level. There's lots of shallow water around Britain. We need the tidal range to be as large as possible, too. An ideal location might be an area of shallow sea surrounding a small island, where the pumping facilities could be built. Alternatively the high pool could be built on land. Offshore lagoons have many advantages, as advocated by Tidal Electric limited, and Friends of the Earth Friends of the Earth Cymru [2004]. I think the best two locations in the British Isles are, on the East Coast, The Wash (where the mean spring tidal range exceeds 7 m) and, on the West Coast, anywhere in the Irish Sea from the Mersey to the mouth of Morecambe Bay (where there have been proposals to build a 12-mile bridge with built-in tidal and wind power). The mean spring tidal range here is 7–8 m. Morecambe Bay already has a gas field, so there is a precedent for energy exploitation. From wikipedia: 'A lease has been granted for the development of two wind turbine sites in Morecambe Bay, one at Walney Island and the other at Cleveleys. Together these will have around 50 turbines.' The Wash would be big enough to fit one 10 km by 10 km tidal facility, but perhaps not more. The Irish sea is bigger. Both locations have a tidal range bigger than I assumed, so the potential power is bigger – perhaps about twice as big, on average.

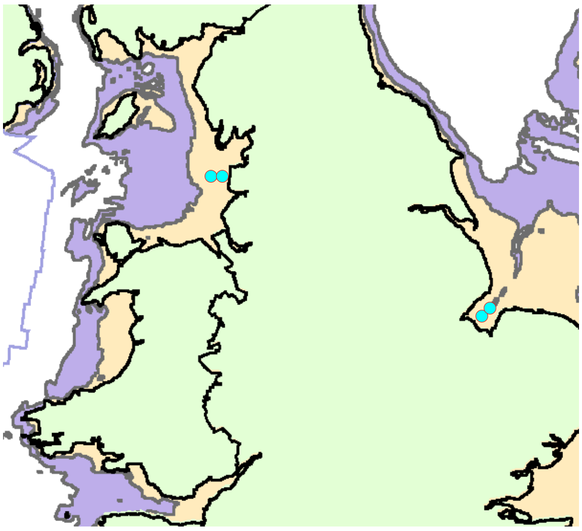

Figure 13. Two locations with plentiful shallow water and big tides: The Wash, and the Irish Sea (marked by two circles each). UK territorial waters with depth less than 25 m (yellow) and depth between 25 m and 50 m (magenta). Data from DTI Atlas of Renewable Marine Resources. © Crown copyright.

What about the cost? Two circular lagoons enclosing 100 km^2 would require 50 km of walls. The low pool's walls would be in water of depth about 13 m. And (for the most ambitious schedules described here) we need the high pool to have a wall of height 13 m above sea-level. Let's look at some costs from Tidal Electric limited. Their plan for a small tidal lagoon in Swansea Bay (where the mean tidal range varies between 4.1 m and 8.5 m) involved 9 km of walls and would cost either £82 million (according to Tidal Electric, AEA Technology, and W.S. Atkins Engineering) or £234 million (according to critics of Tidal Electric's scheme). The cost of the wall was estimated to be £49 million or £114 million respectively. Taking the larger of these two figures, the cost per km of wall is £13 million. This wall was of height 16 m from sea bed to crest. The walls I was imagining above would be slightly higher or perhaps twice as high (if the high pool is built in water of the same depth as the low pool). The wall, using this technology, would thus cost at least £0.65 billion. Perhaps costs could be reduced by alternative wall construction methods. And I think the wall heights could be trimmed quite a lot without spoiling the results sketched here. Doubling the wall's cost to allow for all the other stuff, I'll propose £1.3 billion as the cost for a 10 km × 10 km two-lagoon system.

Dinorwig cost £0.4 billion in 1980 money, so £1.3 billion for a facility superior to two Dinorwigs sounds a reasonable deal to me. Another way of expressing the value of the facility is to take what people currently spend on wind turbines – for example, £500 million on the '650 MW' Lewis wind farm, plus £375 million on the Lewis–Mainland electricity connection: an expenditure of about £0.9 billion on roughly 220 MW (average) of intermittent power. Scaling this up, 550 MW of bursty wind power seems to be valued at £2.25 billion. The pumped storage solution presented in figure 9(a), requiring walls of height 9 m above mean sealevel, would turn this 550 MW of bursty wind power into 750 MW of steady demand-following power. It seems plausible to me that this service would be worth the estimated cost of £1.3 billion.

If the cost needs to be reduced, we simply make the system bigger. For example, we multiply the area by four (to 20 km × 20 km) and double the length of all the walls. The estimated cost roughly doubles (to £2.6 billion, say), but the storage quadruples to 40 GWh (more than four Dinorwigs). As a source of tidal power, this quadrupled station could deliver a steady 1.8 GW all day and all night, and could serve peak demand.

Cost comparison with vanadium flow batteries

For comparison, VRB power systems have provided a 12 MWh energy storage system for the Sorne Hill windfarm in Ireland (currently '32 MW', increasing to '39 MW'). I think VRB stands for vanadium redox battery. This storage system is a big 'flow battery', a vanadium-based redox regenerative fuel cell, with a couple of tanks full of vanadium in different chemical states. This storage system can smooth the output of its windfarm on a time-scale of minutes, but the longest time for which it could deliver one third of the 'capacity' (during a lull in the wind) is one hour. The same company installed a 1.1 MWh system on Tasmania. It can deliver 200 kW for four hours, 300 kW for 5 minutes and 400 kW for 10 seconds.

A 1.5 MWh vanadium system costing \$480 000 occupies 70 m^2 with a mass of 107 tonnes. Its efficiency is 70–75%, round-trip.

Scaling this up, and translating into British, a 10 GWh system using vanadium would cost £1.64 billion; a 20 GWh system would cost £3.3 billion. The tidal-pumped-storage system thus looks competitive with the storage technology currently used for large windfarms. [Scaling up the Vanadium technology to 10 or 20 GWh might have a noticeable effect on the world Vanadium market, but it is probably feasible. Current worldwide production of Vanadium is 40 000 tonnes per year. A 10 GWh system, assuming 1-molar Vanadium solution, would contain 36 000 tonnes of Vanadium – one year's worth of current production. Vanadium is currently produced as a by-product of other processes, and the total world Vanadium resource is estimated to be 63 million tonnes.]

Compressed-air storage

Compressed air storage is said to be significantly less expensive than other large-scale storage options [Denholm et al., 2005]. “Energy is stored by compressing air in an airtight underground storage cavern. To extract stored energy, compressed air is drawn from the storage vessel, heated, and then expanded through a high-pressure turbine, which captures some of the energy in the compressed air. The air is then mixed with fuel and combusted, with the exhaust expanded through a low-pressure gas turbine. The turbines are connected to an electrical generator. Turbine exhaust heat and gas burners are used to preheat cavern air entering the turbines. CAES can be considered a hybrid generation/storage system because it requires combustion in the gas turbine. The storage benefit of precompressed air is the elimination of the turbine input compressor stage, which uses approximately 60% of the mechanical energy produced by a standard combustion turbine. By utilization of precompressed air, CAES effectively ”stores” the mechanical energy that would be required to turn the input compressor and uses nearly all of the turbine mechanical energy to drive the electric generator.”

1 kWh of electricity generated by the CAES turbine requires 4649 kJ of fuel (1.3 kWh) plus 0.735 kWh of compressor electricity. This is said to be five times more efficient than the most efficient plain fossil combustion technology.

Further details including a life-cycle analysis are in the paper [Denholm et al., 2005].

I haven’t found a figure for the cost of such a storage system.

Additional opportunities

A pair of lagoons in the sea with 13 m-high walls and electrical plumbing installed would be a good place to locate wind turbines. The turbines would be offshore, which is good, but erection and maintenance of turbines on the walls would be much easier and cheaper than for regular offshore turbines. 100 m diameter turbines (with ‘capacity’ 3.5 MW) could be placed every 500 m – 100 turbines in total, with a ‘capacity’ of 350 MW. A good combination: wind, pumped storage, and tidal energy, all enhancing each other.

Perhaps to kill four birds with one stone, we could sequester carbon too: the walls could be built out of artificial limestone, or coal!

Acknowledgements

I thank Stephen Salter, Denis Mollison, Adrian Wrigley, Tim Jervis, Marcus Frean, and Trevor Whittaker for helpful discussions.

References

J. A. Baines, V. G. Newman, I. W. Hanna, T. H. Douglas, W. J. Carlyle, I. L. Jones, D. M. Eaton, and G. Zeronian. Dinorwig pumped storage scheme. *Institution of Civil Engineers Proceedings pt. 1*, 74:635–718, November 1983.

J. A. Baines, V. G. Newman, I. W. Hanna, T. H. Douglas, W. J. Carlyle, I. L. Jones, D. M. Eaton, and G. Zeronian. Dinorwig pumped storage scheme. *Institution of Civil Engineers Proceedings pt. 1*, 80:493–536, April 1986.

C. Baker, J. Walbancke, and P. Leach. Tidal lagoon power generation scheme in swansea bay, 2006. URL `http://www.dti.gov.uk/files/file30617.pdf`. A report on behalf of the Department of Trade and Industry and the Welsh Development Agency.

B. D. T. . R. S. C. A two-basin tidal power scheme for the severn estuary. In *Conference on new approaches to tidal power*, 1982.

P. Denholm, G. L. Kulcinski, and T. Holloway. Emissions and energy efficiency assessment of baseload wind energy systems. *Environ Sci Technol*, 39(6):1903–1911, March 2005. ISSN 0013-936X. URL `http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=pubmed&dopt=Abstract&list_uids=15819254`.

Department of Trade and Industry. DTI atlas of UK marine renewable energy resources, December 2004.

Friends of the Earth Cymru. A severn barrage or tidal lagoons?, January 2004. URL `http://www.foe.co.uk/resource/briefings/severn_barrage_lagoons.pdf`.

T. J. Hammons. Tidal power. *Proceedings of the IEEE*, 8(3):419–433, March 1993.

Z. Kowalik. Tide distribution and tapping into tidal energy. *Oceanologia*, 46(3):291–331, 2004.

S. H. Salter. Possible under-estimation of the UK tidal resource, 2005. URL `http://www.dti.gov.uk/files/file31313.pdf`. Submission for DTI Energy Review Question 2.

T. L. Shaw and M. J. Watson. The effects of pumping on the energy potential of a tidal power barrage. *Engineering Sustainability*, 156(2):111–117, June 2003a. ISSN 1478-4637. doi: 10.1680/ensu.156.2.111.37018. URL `http://www.atypon-link.com/TELF/doi/abs/10.1680/ensu.156.2.111.37018?cookieSet=1&journalCode=ensu`.

T. L. Shaw and M. J. Watson. Flexible power generation from a severn barrage, April 2006. URL `http://www.dti.gov.uk/files/file31332.pdf`.

T. L. Shaw and M. J. Watson. Flexible power generation from the tides. *Engineering Sustainability*, 156(2):119–123, 2003b. ISSN 1478-4629.

S. J. Taylor. The severn barrage – definition study for a new appraisal of the project, 2002. URL `http://www.dti.gov.uk/files/file15363.pdf`. ETSU REPORT NO. T/09/00212/00/REP.

Power translation chart

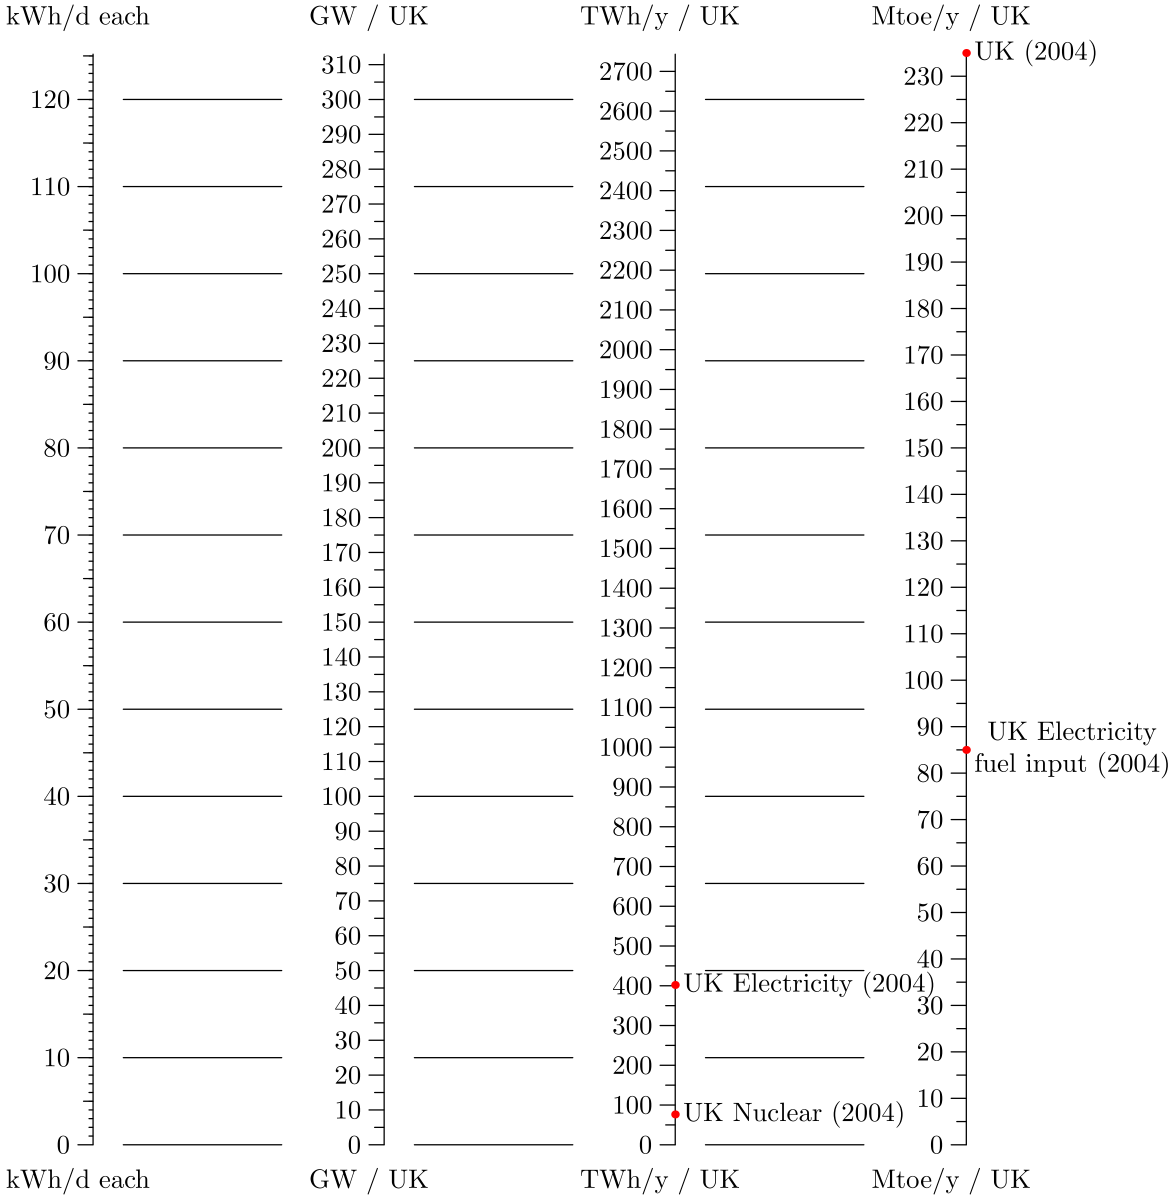

1 kWh/d	the same as 1/24 kW
GW	often used for ‘**capacity**’ (peak output)
TWh/y	often used for average output
1 Mtoe	‘one million tonnes of oil equivalent’

‘UK’ = 60 million people

USA: 300 kWh/d each
Europe: 120 kWh/d each