THE HYDRAULICS OF TIDE MILLS

David H Jones

This paper reports the beginning of a study of the operation of a traditional tide mill. The main aims of the study are to compare the hydraulic power and working regimes of known tide mill sites and to seek connections between types of tide mill and their regions of application.

As several large tide power plants have been proposed in recent times, and one actually built, the basic theory of tide power must have been studied, and no doubt published. I have not yet consulted such work as I do not have easy access to it at the moment, although it must be searched eventually. This is unlikely to be important, because a modern study would obviously assume a plant using turbines. This removes a limitation felt by the tide mill using a waterwheel; a difference which so modifies the working conditions as to necessitate a new study based on first principles.

Essential, that is, if a theoretical study is worth making at all. It is tempting to apply modern analysis to an old design in the belief that it is a contribution to technical history, but it may well be no such thing. For example, the computer has made possible full stress analyses of complex structures, and every so often these methods are applied to some historic example. The usual result is to prove that the structure is stable; a result which could be obtained with far less effort by noticing that it still stands. Such work may be an exercise in statics, but it is not history.

Nevertheless, modern methods are of value in historical studies, but they must be approached as history. The purpose of the analysis is to answer questions, and it is essential that such questions should have historical significance. This is true of most research; if the information exists at all, it is frequently more difficult to ask the significant question than to answer them. For the tide mill, technical considerations are likely to throw light on such questions as:

a) how did the capacity of a tide power site compare with freshwater sites in the same region?
b) what was required of a tide power site, and how widely were they available (both exploited and unexploited)?
c) where design options were different combinations of output and working period, what choices were made?
d) accounts exist of "improved" tide mill designs which were not adopted: were they soundly based, and what were they intended to achieve?
e) could they have been improved by the use of mill designs in use at the time of their construction, but not applied to tide mills?
f) if so, why were they not applied?
g) when different regions used very different designs, did their characteristics match local conditions, or were they adopted or non-technical reasons?

What does not seem profitable is to compare the performance a tide mill with what might be obtainable from a modern plant designed for the same site. As the tide mill was not a major inspiration for invention, we should only consider the machines which were available to their builders. A millwright with a commission to build (or rebuild) a tide mill might improve the details of the waterwheel, but he could hardly have been driven to invent the turbine. The following analysis begins with assumptions which could only be approached by a turbine, but this is only for convenience; they are discarded before attempting to draw conclusions.

Available Power

For an intermittently-operating machine like a tide mill, power needs a definition. The capacity of the mill depends on the mean power, averaged over one tide cycle, or 24 hours, depending on whether it is worked in one or two shifts. This is also true of some freshwater mills, but the tide mill is far more extreme. The peak power, occuring during part (perhaps most) of the working period, depends on the capacity of the plant and the working period duration. The potential output of any water power site is the product of water flow and fall, but for a tide mill, these factors are not easy to identify. If we postpone consideration of those cases where the supply is augmented by fresh water, the average (12 hour) flow is determined by the pond; the sea can be regarded as an infinite source and sink, uninfluenced by the presence of the mill. The fall is more complicated. It obviously cannot exceed the tide range, and is likely to be much less. As the flow is drawn from the pond, its level drops and the working fall is reduced. The mean fall is half the tide range, but this is not generally useful. If, as is very likely, the pond sides and bottom slope substantially, its capacity will not be proportional to the height of water in it.

The potential of a site depends primarily on the amount of energy which can be stored in the pond, and this is a function of tide range, pond volume, and the shape of the pond cross-section. Where this data is available, the available power at that site can be calculated, although predicting the performance of a practical tide mill is difficult and subject to many uncertainties. As a start to understanding the subject, a rough analysis has been made, based on the following assuptions:

1 The sea is an infinite source and sink of water; i.e, the local tide is not affected by the presence of the mill.

2 The tide height is a sinusoidal function of time.

3 The tide range also varies sinusoidally, with a period of 30 days.

4 Longer period variations, though present, may be ignored.

5 The sluices are adequate to allow the pond to fill to High Water Level.

6 The waterwheel cannot run with any backwater.

7 The wheel is operated with constant power input throughout the working period.

8 The wheel - and the mill - can operate over the speed range demanded by the variation of head as the pond empties.

9 The wheel is run at with a power input which will empty the pond (to the minimum workable head) at the time the rising tide reaches the tailrace.

10 The wheel can accept sufficient to allow condition 9 to be met.

Most of these assumptions are fully justified, in that they cause little error. However, nos. 5 and 6 are clearly not true, and must be re-examined if the analysis is refined and used on data from real sites. No. 8 is suspect, and they analysis may have to be revised to eliminate it. As for Nos. 9 and 10, they imply that the size of the pond is one of the main factors which determines the output of a tide mill, and a major aim of the analysis is to establish whether or not this is so.

Such a simple analysis can only deal with idealised cases which could not be fully realised in practice, but they can be be used to establish the ultimate limit of performance.

Waterwheel Characteristics

Many waterwheel designs will only operate efficiently - if at all - if the headrace level remains within a narrow range. Hardly any can accommodate large level changes on the tailrace side and these exceptions were not normally applied to tide mills.

If we initially consider only British tide mills, we find that nearly all of them use some form of undershot wheel - either traditional or Poncelet. As the Poncelet wheel is an undershot wheel with improved detail design, their characteristics are similar. These wheels will operate with substantial changes in headrace level, but their optimum speed varies with it. On the tailrace side, there is a maximum level above which the wheel is backwatered and either stopped or its performance impaired. The wheel is obviously unaffected by lower tailrace levels, but can obtain no benefit from the extra fall.

The siting of such a waterwheel clearly involves considerable compromise if it is to be tide powered. If it is raised, the working period is increased as the tailrace will be clear longer, but the power is reduced by the smaller fall. On neap tides, the initial headrace level may be too low to run the wheel at all. If it is lowered, the power will be increased, but the working period will be reduced, while on neap tides the tailrace may not clear at all. Even the power increase when working may prove small, as the increased fall, especially on spring tides, may be more than the wheel can use effectively.

Design Aims

Although the obvious starting point was to consider the amount of power available at a given site, and show how it could be maximised, it is now clear that this is only one of a number of design options, and not necessarily the most advantageous. At the most basic level, the object of the design was to build a mill which could do as much work as was required, as economically as possible. Here "economically" means making minimum demand for the scarcest resource, whether this be land, building material, millwork, construction labour, or operating labour.

If land is in demand and the pond area either reclaimable or has other uses as water, minimising the pond size - which is equivalent to maximising the power - is important. However, if a natural pond of generous size is available with no other claims on its use, spreading the power over a long working period may be more desirable.

Practical Mill Design

The builder of a traditional tide mill would have to make two main decisions; siting the dam, and choosing the waterheel level. The dam site would have been influenced by so many factors other than the technical performance of the mill that it seems only worth considering it for specific cases.

Although the tides round the British Isles can range as high as 7 m, the range at a typical tide mill site is very much less. There are several reasons for this. Sites are typically in creeks and estuaries, where the tide range is low. They are no doubt chosen because they offer natural or easily-built ponds, and for the protection they offer against storm and wave damage. Another reason is that a waterwheel could not utilise a high tide range anyway. A fall of more than about 3 m required an overshot or a backshot wheel, which could not accommodate a wide range of intake levels, and a tide mill depended on using a water supply which dropped to a small fraction of its initial level as the pond emptied.

The optimum design seems to be: a large plant (a broad waterwheel, or two wheels) with the wheels set so that they are just clear of backwater at low water on a neap tide. (At normal tides they will have a little more power, and a much longer working period). It was probably desired not to lose too much work on neaps (backwater or pond shape)

There is no doubt that fill-and-empty operation, with either a reversing wheel or two opposite-facing wheels, would have given both more power and longer working time, but they were hardly used anywhere. Perhaps capacity was not the issue; after all, how many regularly worked two shifts?

The lifting waterwheel (moulin pendant; Panstermühle) would not have been useful in those forms; they were designed to let excess water run to waste, not use it efficiently at all levels.

One- or two-shift operation is an important issue, as it leads to an uncertainty over output of a factor of two. It is quite impossible to resolve this, as there is little doubt that practically all millers would have worked either one or two shifts according to demand.

(The mathematical analysis of the energy content of ponds of various cross-sectional shapes and its variation throughout the monthly cycle, and the variation of working period have been omitted to save space, as a later stage of this work will shortly be published by TIMS, Trans. Sixth Symp. Molinology, 1985.)

Discussion

JARVIS I find this quite an intriguing exercise. To people who have given it any thought at all it must be obvious that there is a limitation somewhere, but where are the limitations? I think you are pointing the way. I would like to ask, have you considered any particular site, to see how near or otherwise the optimum situation has been achieved. I appreciate there are difficulties, as the sites as they are today may be considerably changed from what they were when the tide mills were operating, even if you know the height of the dam.

JONES Yes, I plan to do that. So far the only data I have is for Eling. A little has come from it; it is clear that they took something like the optimum as I understood it. They certainly had two wheels there, and they certainly attempted to get a lot of water out. It will only just operate on the neap tides, if at all; they were prepared to lose a little working time, though not much. They seemed more concerned with having enough plant to do a lot of work in a limited period rather than extend the working period. I haven't got much further than that. Of course, there are two other complications; one is that all this assumes there is no freshwater feed to the pond. This is true of some sites, but not all. In fact, we have the whole range from no freshwater input to those wholly dependant on freshwater, where the tide is merely an obstruction. Eling does have a significant freshwater input.

PLUNKETT With the lagoon-type model you have no complications; it is governed entirely by the tide. There are no other parameters, except the mean sea level at the site.

JONES The other complication also applies at Eling; it is in a very anomalous tide region. The assumption of a sinusoidal variation is an approximation anywhere, but Southampton is an extreme departure from it. However, it does not affect the total power available, but only the working period, which is shortened. I find it convenient to consider approximate sinusoids in terms of their harmonic content, so I analysed the Southampton tide record and found about 30% each of 2nd and 3rd harmonic. The other components - up to the 20th - had considerable effect on the shape of the curve, but their amplitudes were very small. The phase of the 2nd harmonic is such that it broadens the top considerably, and narrows the bottom, which is the working period. It might have been more convenient on that site to have turned the mill round and run while filling the pond instead of emptying it. That would have lost the benefit of the freshwater, of course.

Beyond that, it should be possible to get levels from a number of existing tide mill sites, even where the mill has gone. It only needs a very few levels. The tide levels haven't changed. The big uncertainty is that the pond capacity may have been reduced by silting. And were they originally dredged or not? If they were, presumably when the mills were operating successfully, the ponds were kept dredged. In decline, they would be allowed to silt up.

JARVIS Another question on the operation of waterwheels - undershot or Poncelet - is not the power related to speed and fall? Would there not be an advantage, if you were trying to optimise, in having a variable gear ratio, changed at some point in the tide range?

JONES That is a possibility. This question of speed is the other complication I was referring to. The simple analysis assumes that the wheel runs with reasonable efficiency over the whole range as the pond empties. That may be far from the truth. For an impulse wheel there is an optimum speed which is a strong function of the head. The practical efficiency with a widely-varying head is likely to be very poor, which raises the question of whether anything could be done about it. It could be, because we know of an example where they did adapt the waterwheel for this condition, at Woodbridge. There we have a three-piece gate, so that it first operates as a breast wheel, and when the level gets too low, it can run as a lower breast wheel, and finally as an undershot. It seems a very sound design; simple, easy to maintain, and easy to operate, and I don't see why it was not more widely used.

JARVIS This is somewhat like the windmill situation, where provided you get enough power it doesn't matter how much wind you use. Given enough water, it doesn't matter how efficiently you use it.

JONES Yes, the aim would have been to build a mill which could meet the demand throughout the year. A tide mill has the advantage of a reliable minimum power; if there was sometimes more than was required, it was ignored.

TURNER What of the Portugese tide mills, which used a number of horizontal wheels; perhaps 8 wheels on some of them, with 16 pairs of stones. Does that mean that they used all 8 at the peak level diffference, gradually reducing to perhaps one or two at the lowest level? They are going to use a vast amount of water at certain times, and if they closed off all but one they would only use very little.

JONES I agree the Portugese design is a very interesting one, and I would like to include it when I get enough information. The fact that they have a large number of wheels is just what we would expect, because they can run them all when they can run at all. This is the tide mill problem; all the work is concentrated into a short period, so they must work several times faster for part of the time. To me, the puzzling thing about the Portugese tide mills is the choice of wheel. Not just that they use horizontal wheels, but that they seem to have chosen the wrong type of horizontal wheel. The Portugese had experience of both vertical and horizontal wheels, and perhaps of the "pit type" wheel (roue à cuve) which was specially designed to suit the large-flow situation, yet they chose a spout-fed wheel. Perhaps there are other factors we haven't thought of yet.

PLUNKETT Perhaps it depends on the tidal range and conditions of working. There is one English tide mill with a turbine conversion of 1904 (I think) which was claimed to work in several feet of tidal water, which it wouldn't do when it had a waterwheel. There is a letter in Armfield's sales catalogue from the owner of the mill, saying what an improvement it was.

A BRYAN What about the use of a floating mill in a tidal race. Where they ever used in this country?

JONES I have never heard of a floating mill being used in a tidal situation in any country, though I had realised it was a possibility.

JARVIS London had two, at different times.

JONES If they were above London Bridge, as I believe they were, their situation would not have been tidal.

P BRYAN What about a moulin pendant in a tidal situation?

JONES At first sight, it seems a very useful thing to try, but I don't think it is really suitable. The purpose of the moulin pendant was to extract a small amount of power from a large river with a variable level, and it was always running water to waste. Indeed, at most sites, at least in their original form, there was no damming, and most of the water simply bypassed the mill. Except when the wheel was at its lowest, a lot of water passed under the wheel, contributing nothing. With a tide mill, anything other than the traditional undershot waterwheel would only be useful if it economised water. A lifting wheel would only do that if the wheel pit bottom also lifted with it. Such a design was proposed, and something similar was built, but it seems that the problems of building and maintaining outwieghed any advantage. I believe the East Greenwich tide mill used this arrangement; also, one of the wheels in House Mill, Bromley-by-Bow, was a lifting wheel, although it is now fixed.

PLUNKETT Any more questions? No.

JONES In that case I would like to ask one. Can anyone suggest any other directions in which this study could be extended?

PLUNKETT That might be better considered after my contribution, which is closely related, and deals with one or two questions originally raised by this work.