The Evolution of Tide Data Management

Introduction

CherSoft specialises in developing marine software, particularly software related to tides.

In this article we will review the process of collecting and using tide data and how it is currently managed.

Then we'll talk a bit about the history and why I think that has influenced tide data management.

After that, we'll look at why tide data needs to be managed differently and how it could be done.

Finally we will look at some of the opportunities that will be opened up by these improvements.

Why are tides important?

First, let's have a look at a few reasons why tides are so important.

Navigation

When can a supertanker like this transit the Malacca Strait? The strait is only 23m deep at its shallowest. The navigator has to time it just right to have enough water under the keel.

Long term sea level monitoring

How long will the Thames barrier be effective for? It was closed 4 times in the 1980s. It was closed 35 times in the 1990s. It was closed 75 times in the first decade of this century. Things are clearly changing – the SE of England is sinking due to post glacial rebound and the sea level is rising. We need to keep track of this.

Extreme event detection and risk analysis

On this diagram, the blue line shows the height of tide on 28 Feb 2010 at La Rochelle. The green line is the predicted tide. Hurricane Xynthia caused a surge of 1.5m exactly at the time of a high tide, causing massive flooding. We need the tide information to do proper risk analysis for these rare events.

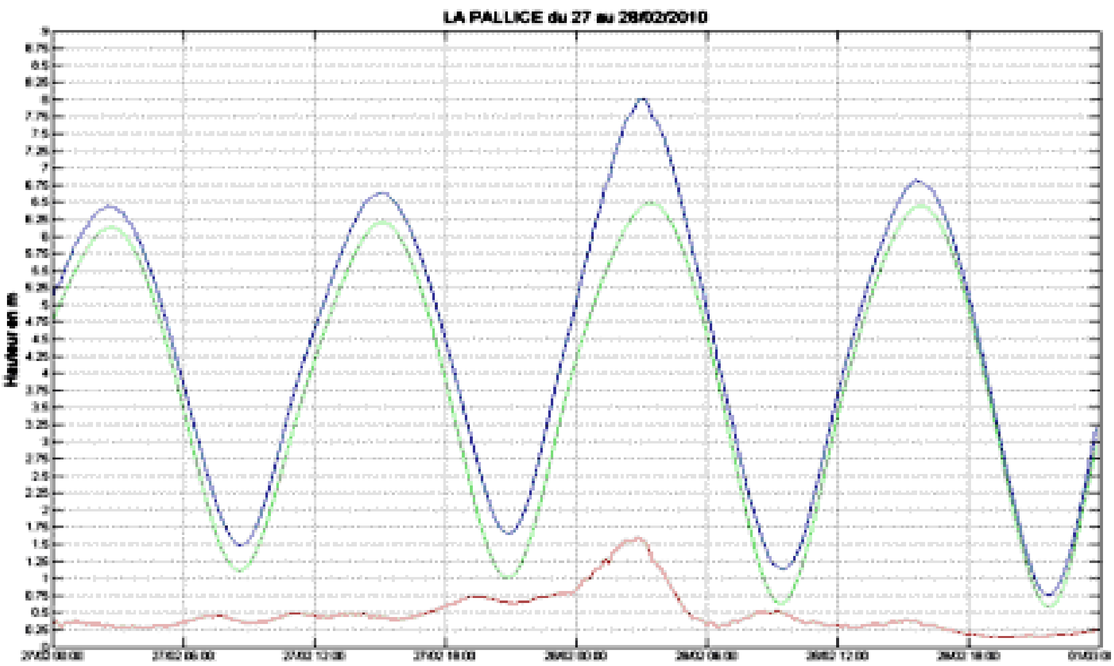

Planning tide power installations

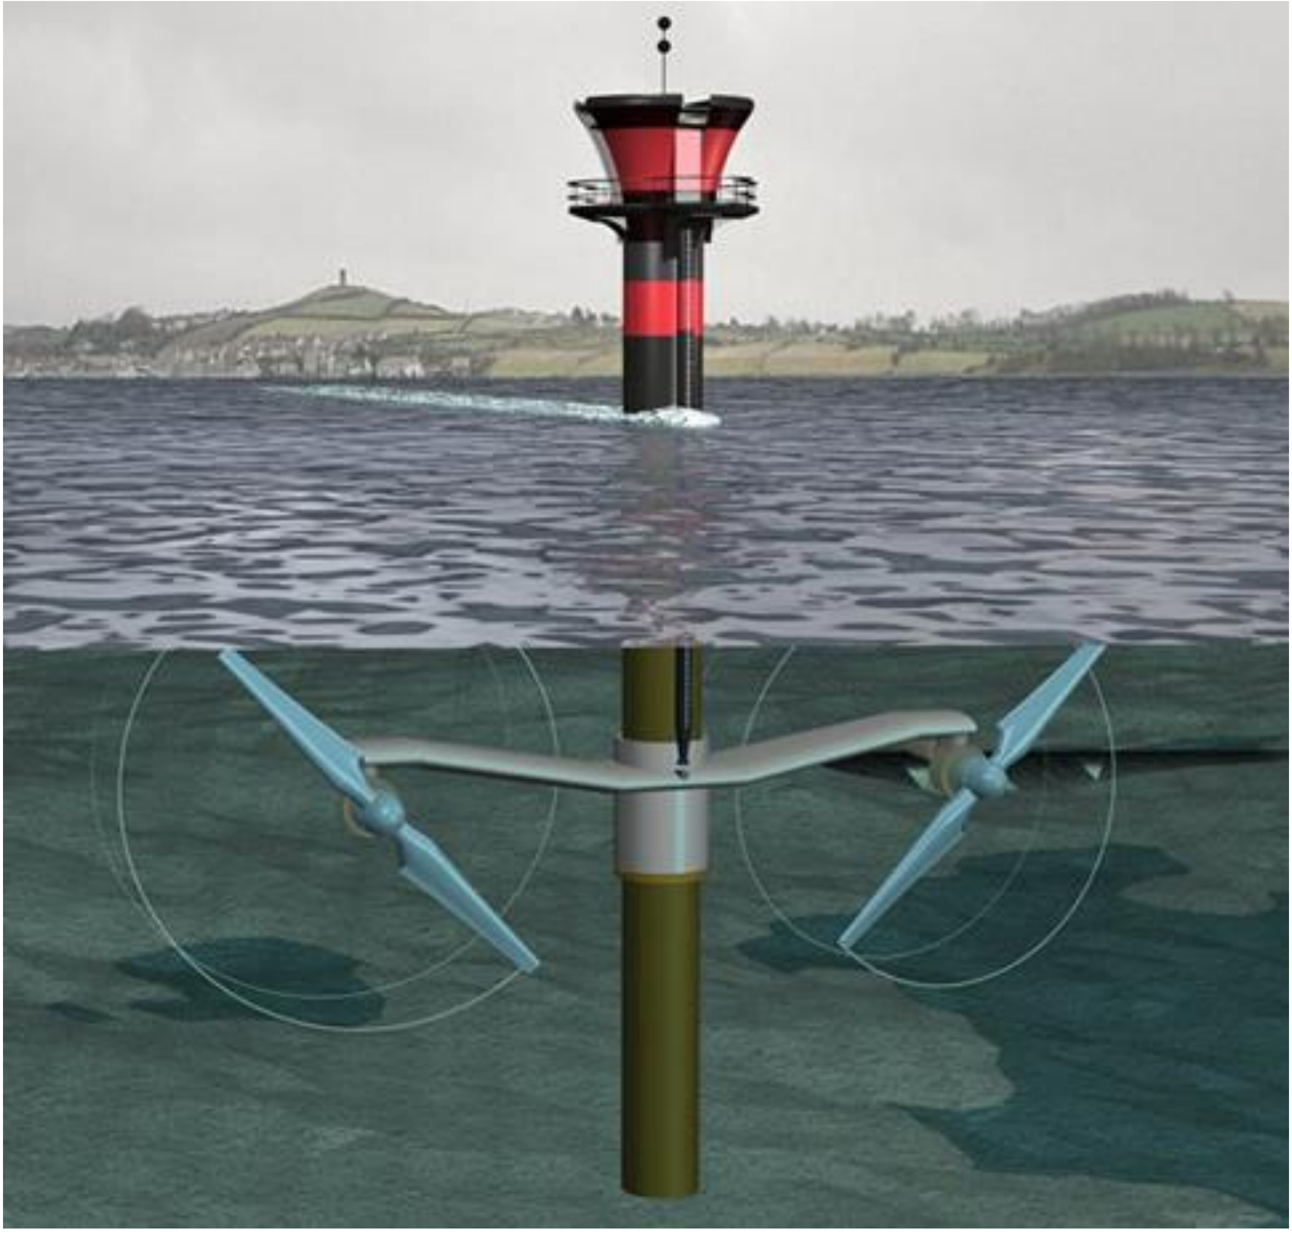

There are now serious attempts to extract tide energy and generate electricity. I know there have been trials and experiments for years, but most of those were shown on Tomorrow's World which was the kiss of death for any invention. And the ideas weren't any good. Now I think there are worthwhile designs coming off the drawing boards.

There are 3.5 terra watts of tide power. A few percent of this would be very useful. The latest ideas are basically underwater windmills, extracting kinetic energy. This needs detailed and accurate tide height and tidal stream information.

As with any kinetic energy extractor, the power output of the machine is proportional to the cube of the speed of the fluid. This is unfortunate, because it means very low energy yield at below average speeds. An extra tenth of a knot of tidal stream can make all the difference to a generator's viability. So you want to put your watermill in a narrow, shallow constricted seaway to get the maximum power output. That is a pity because we tend to use such areas as shipping lanes!

Monitoring and Predicting Mean Sea Level changes

Sea Level change is critical for numerous low lying areas, such as the Maldives. This is Malé. You might have expected a bounty advert island, but this is their main city. Most of the population lives between 0.8 and 2m above sea level.

In 2007, 55 of the islands were flooded by a high tide.

And as mean sea level is projected to rise by up to half a metre over next century, events like this are bound to become more common.

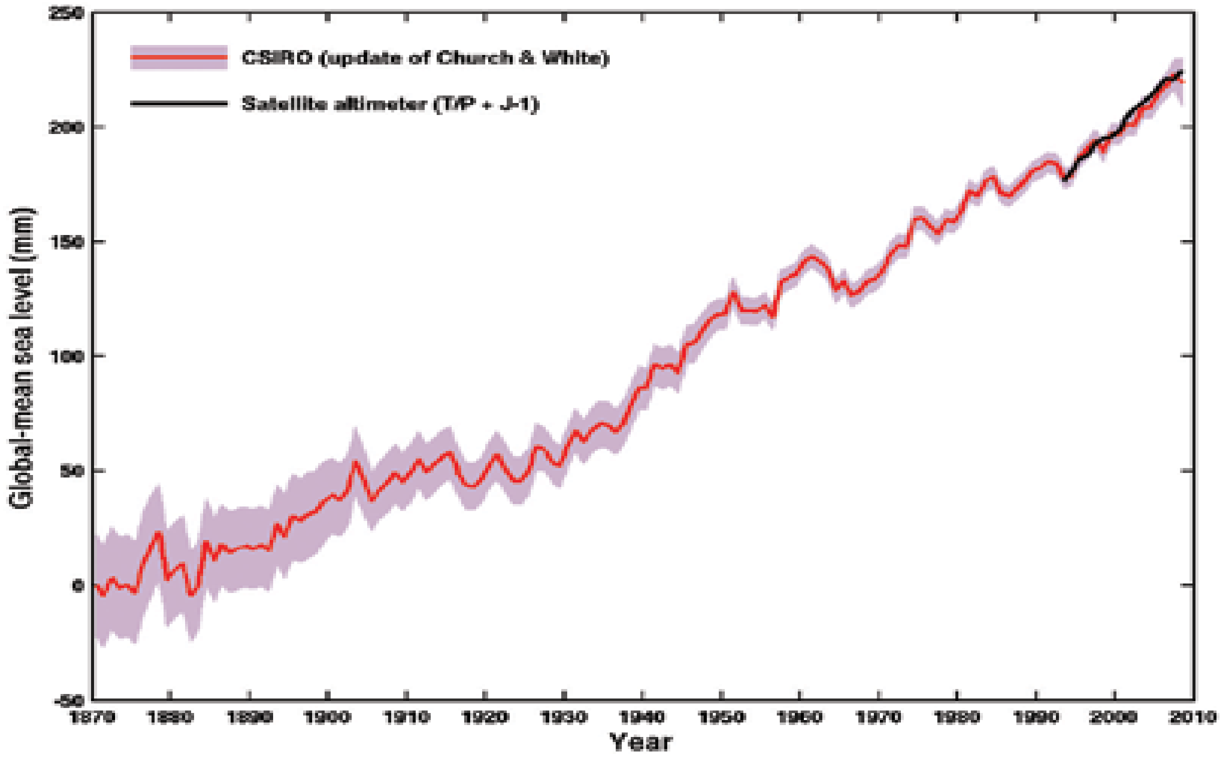

This graph shows the trend in mean sea level over the last 140 years – a 200 mm rise. The data in this graph has the potential to prove or disprove global warming. The attempts to model, and eventually predict, the effects of global warming, or global chaos as it should be called, make our old tidal records as valuable as the new ones, if not more so.

While tide gauge records only cover the last few hundred years, there are indirect measurements which can give us a much longer perspective. On this graph, the blue line is from ESSO measurements derived from oil exploration. The red line is from a completely separate study based on surface geological interpretation.

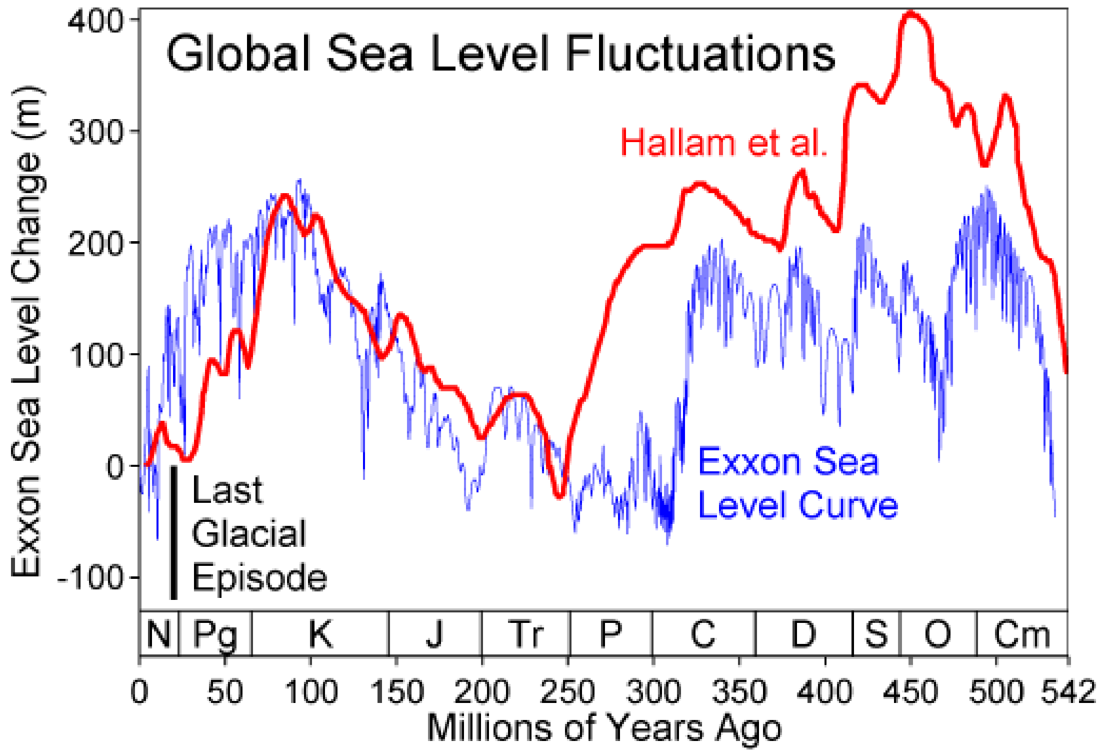

Notice the scales – we are looking at hundreds of metres of sea level change over the last 500 million years.

For comparison, melting all the ice caps would add 80m and the end of the last ice age raised levels by 120m as indicated by the black bar.

Incidentally, thermal expansion of the oceans due to increasing water temperature is actually a larger effect than melting the ice caps. They were roughly equal in the 20th century but thermal expansion will dominate this century.

Now; all changes that have been recorded so far using all the data we know of from tide gauges are just 0.2m in total. That is just one thousandth of the change over the last 500 million years. It is too small to even show up on this graph.

Given that we know that even a one metre of sea level rise will have a huge impact, and seeing the scale of level changes in geological time, it is obvious that we are trying to measure signals well down in the noise. We are going to have to wring every last bit of meaning out of the measurements.

The existing tide management process

It is worth a quick summary of the existing process for collecting and using tide data. We need to know the steps involved in tide processing before suggesting improvements. My aim here is to identify the processing steps rather than to worry about the fine details of each piece of equipment.

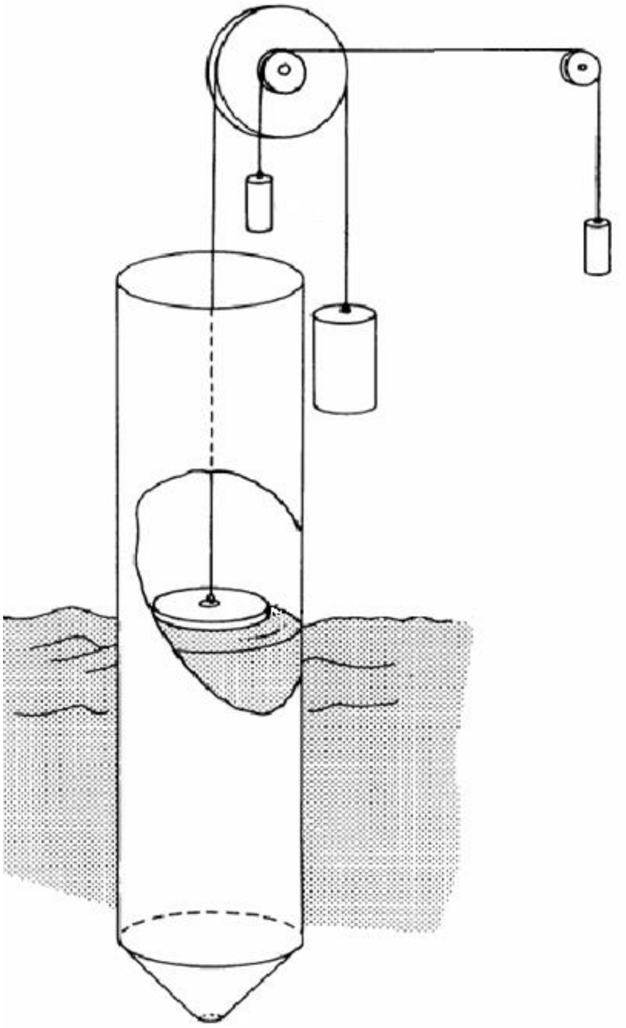

Starting with the oldest first, on the left is a tide staff - the oldest type of gauge. It is still an important calibration and reference tool. Its big problem is that there is no way to read it automatically.

This tide staff is on the outside of a stilling well – a tube housing a float which drives a recorder via wires and pulleys. The tube is a filter which gets rid of wave action for the benefit of the measurements and to protect the float.

Ever since the float gauge was invented, people have striven for a gauge with no moving parts.

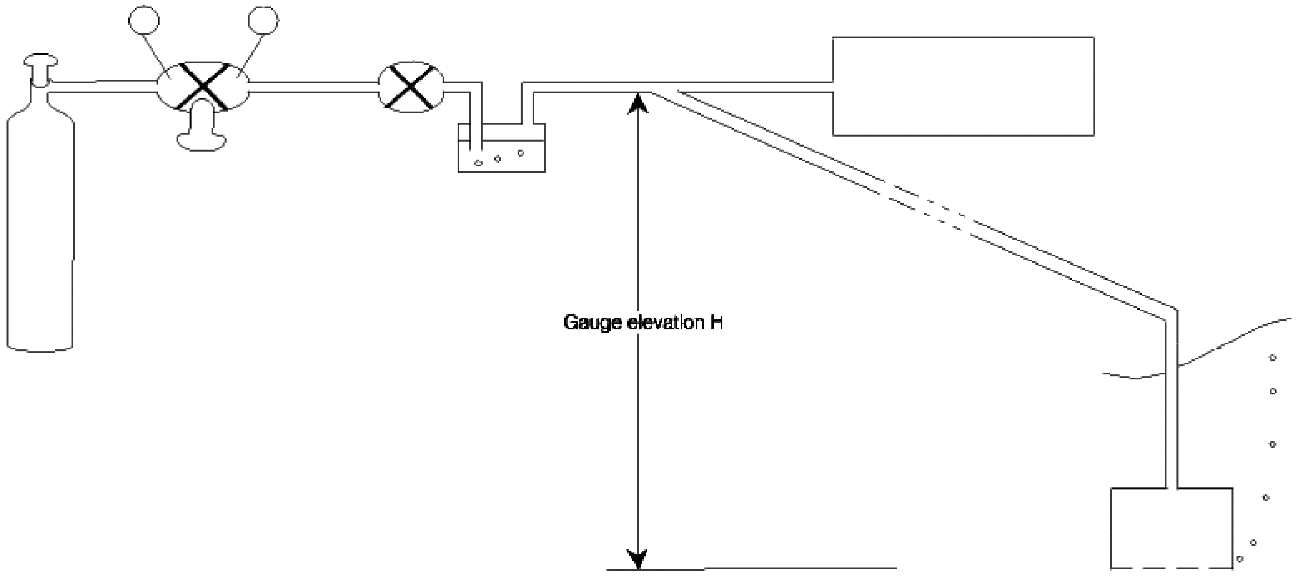

The first successful one was the bubbler pressure gauge (pictured above) which is based on Archimedes principle.

Compressed air is used to blow a stream of bubbles out of the open end of a thin tube fixed above the sea bed. Providing the flow rate is low enough, the pressure at the land end of the tube is the same as the pressure at the underwater open end; and that is proportional to the height of the water column above it.

A bubbler gauge is useful when there is no vertical wall to mount a gauge on. The only part exposed to the sea is the thin pipe which can run horizontally for hundreds of metres if necessary. Very useful, but the pipe does tend to get damaged by anchors, fishing etc.

As with any pressure gauge, this is an indirect measurement – atmospheric pressure and water density are needed to calculate the height, and water density depends on water temperature and salinity. More data to collect, check and store; more scope for error.

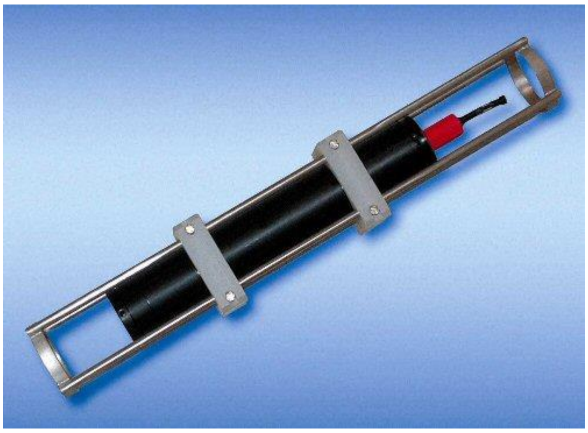

Later still, we have some modern gauges. On the left is a direct pressure sensor gauge made by Valeport. This is particularly useful offshore – one can be simply dropped over the side and left on the sea bed for a month or two happily recording away. The gauge is then recovered and its data extracted.

On the right is a radar gauge made by OTT. It performs a more direct measurement affected only by the speed of light though the air between the gauge and the sea surface.

The lower lifetime cost of the 'no moving parts' gauges and, with radar, the 'no contact' gauges makes it practical to have many more permanent gauges.

Data Loggers

Once measured, the data has to be stored. A data logger does this. It is a small item nowadays – the one below is memory stick sized and can store 6 months of tide height. Many gauges build the logging function in.

There is an important extra step here: Filtering. The gauge is measuring water height, but that includes waves, the wash from ships and oscillations from other sources. We want to ignore all this noise. Fortunately most of it is at a higher frequency than any real tidal effect, so it can be removed using a low pass filter – averaging in other words. Typically, each recorded value is the average over 2 to 4 minutes. And the logger normally stores one value every 6 or 10 minutes.

The older types of gauge did the filtering mechanically – a float gauge is housed in a stilling well, as I mentioned earlier. The stilling well is a large diameter pipe with only a small hole connecting it to the sea. Typically the hole is only one fiftieth of the pipe diameter – so only one two thousand five hundredth of the cross sectional area. Therefore the flow is limited and so is the rate of change of level within the pipe.

Other types of data need to be measured and logged as well – particularly for a pressure sensor where we need atmospheric pressure, water temperature and salinity to calculate the water height. Wind speed and direction are also important factors worth recording.

Then, having captured the data, we need a computer.

This computer is a 1960s classic, the IBM 360, and it really does look like a computer should look.

But what do we actually use the computer for? Currently, almost always, we use it just for storage and for calculations. The computer is used to store the inputs, do the sums and store the results. Typically, the inputs and results are stored in hundreds of files in hundreds more directories and sub directories.

Storage

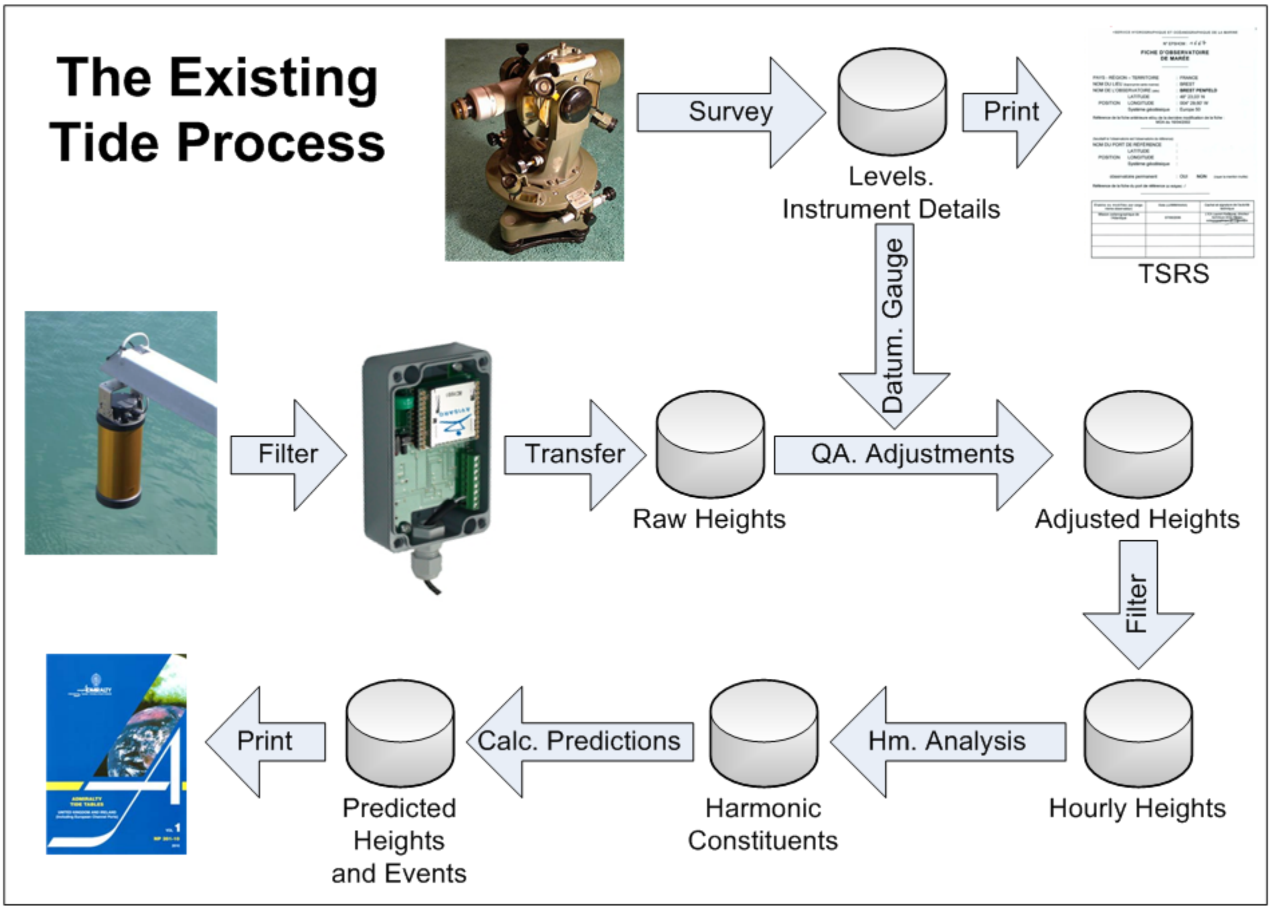

The computer disk is like a large filing cabinet. There are lots of steps in the tide process. Separate files are used to store the results of each step. These are the drums in this diagram.

Raw heights (in the centre of the diagram) are received either continuously or periodically from a gauge. Weekly collection by modem is common.

The heights need to be checked – that's the arrow called "QA Adjustments". There are many tests that exploit the specific nature of tide data to check that it is sensible – that the clock is accurate, that there are no steps, flats or spikes. Bad data needs to be rejected.

The raw heights need to be read in conjunction with the vertical datum. This comes from the survey information for the station – we can store that on the computer too. You can see the vertical datum information in the drum marked "Instrument Details" at the top of the diagram.

Then, when enough data has been accumulated, maybe every month or three, the data is filtered again to produce hourly heights, in the bottom right of the diagram.

The checked raw height and hourly heights might be sent to an archiving service, such as the British Oceanographic Data Centre, the BODC.

A few times a year, the hourly heights will be analysed (the "Hm. Analysis" arrow) to generate a new set of harmonic constituents. These will be compared with the previously generated ones. Any big difference would signify that something had gone wrong.

Then in the "Calc. Predictions" arrow, the harmonic constituents are used to generate new predictions.

Finally, in the bottom left of the diagram, the predictions might be formatted for printing for an annual tide table book, or published on a website.

Notice, that at each stage, more data files are produced which have to be stored. And whoever is doing this work has to be able to find them again later – probably months later as different stages of this process are performed at different frequencies – anything from daily or weekly to annually.

Also notice that here we have only considered one port. There may be tens or hundreds of other ports to process as well.

We can see that there is a major office management task here. Before looking at some of the issues with this, it is interesting to look at some of the history – how it was done before electronic computers. Maybe we will learn something from how they organised their offices.

A dip into history

"it will be HW when the full or new moon bears SSE"

Before the renaissance, non-harmonic methods were used. The people who created Drake's Dial relied on casual observation and folk lore accumulated over the centuries. Drakes Dial relies on spring tides at a place always being at the same time of day, and the time of day being related to the moon's bearing. This allowed a dial calculator to give rough HW times, as moon bearings to people who could barely read and who had had no clocks. For example, "it will be HW when the full or new moon bears SSE".

Of course, we also had real time tide displays over 100 years ago.

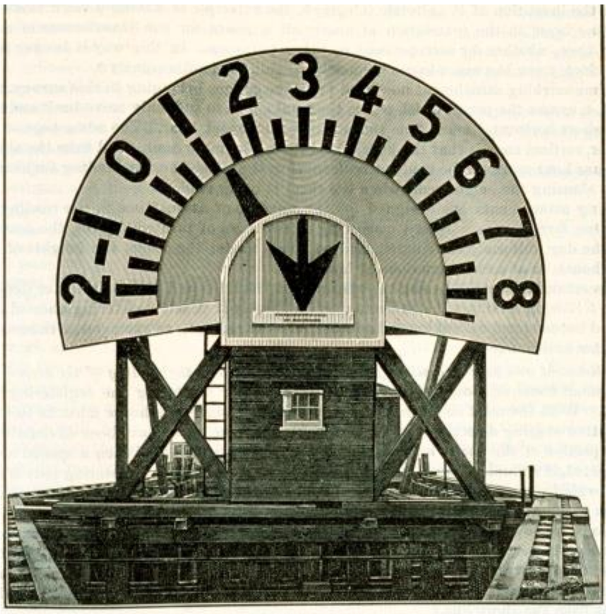

Fig. 5.—Tidal indicator, Delaware River, Delaware.

The wooden station on the left, used in 1897 in New York is one of the earliest examples of a real-time, tide-measuring device. When entering or leaving port, mariners would view it through binoculars. The pointer shows the present level of the water. The vertical arrow shows whether the tide is rising or falling. It is about 20 feet high.

The big advance in quantitative analysis came from Newton's realisation that gravity causes tides. Bernoulli added an understanding of fluid flow. Laplace developed tidal equations which are still in use today. William Thomson (Lord Kelvin) performed the first systematic harmonic analysis of tides.

People like Thomson in Liverpool and William Ferrel in the US invented mechanical solutions to the calculation problems. Some of them are so brilliant that I just have to mention them here.

Drum tide recorders are well known and quite simple, obvious devices. They record 14 days of tide on a single drum that rotates once every 24 hours. Brunel knocked this one up in a spare moment from building bridges, tunnels, railway lines, and the odd ship or two.

But how about the first digital tide recorder?

This is from 1966. It is entirely electro mechanical and generated punched paper tape. These machines were introduced by NOAA and used all over America. I don't think they foresaw the computer revolution – this machine looks like it is built to last forever.

The mechanical tide calculator is also well known. All it has to do is add together sine waves of various amplitudes and phases to produce a predicted tide height.

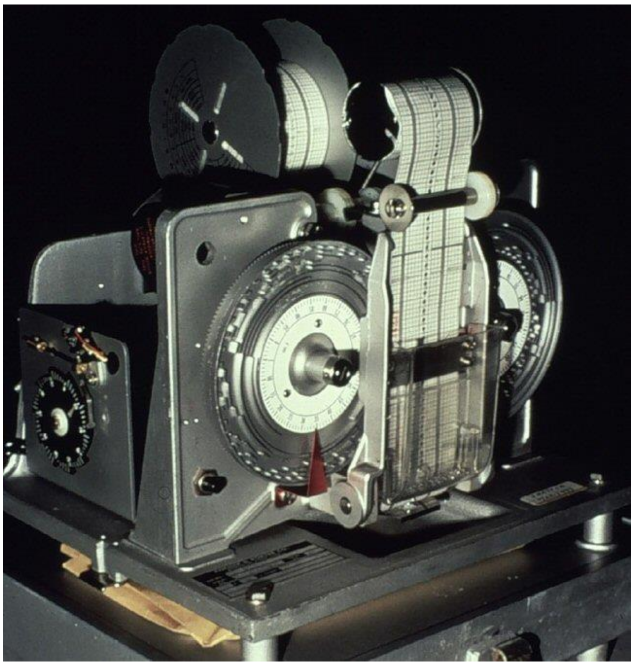

This is William Ferrel's Tide Prediction Machine Number 2, which he had built in the 1880s in the US. It could sum 39 constituents and weighed in at 1 ½ tons.

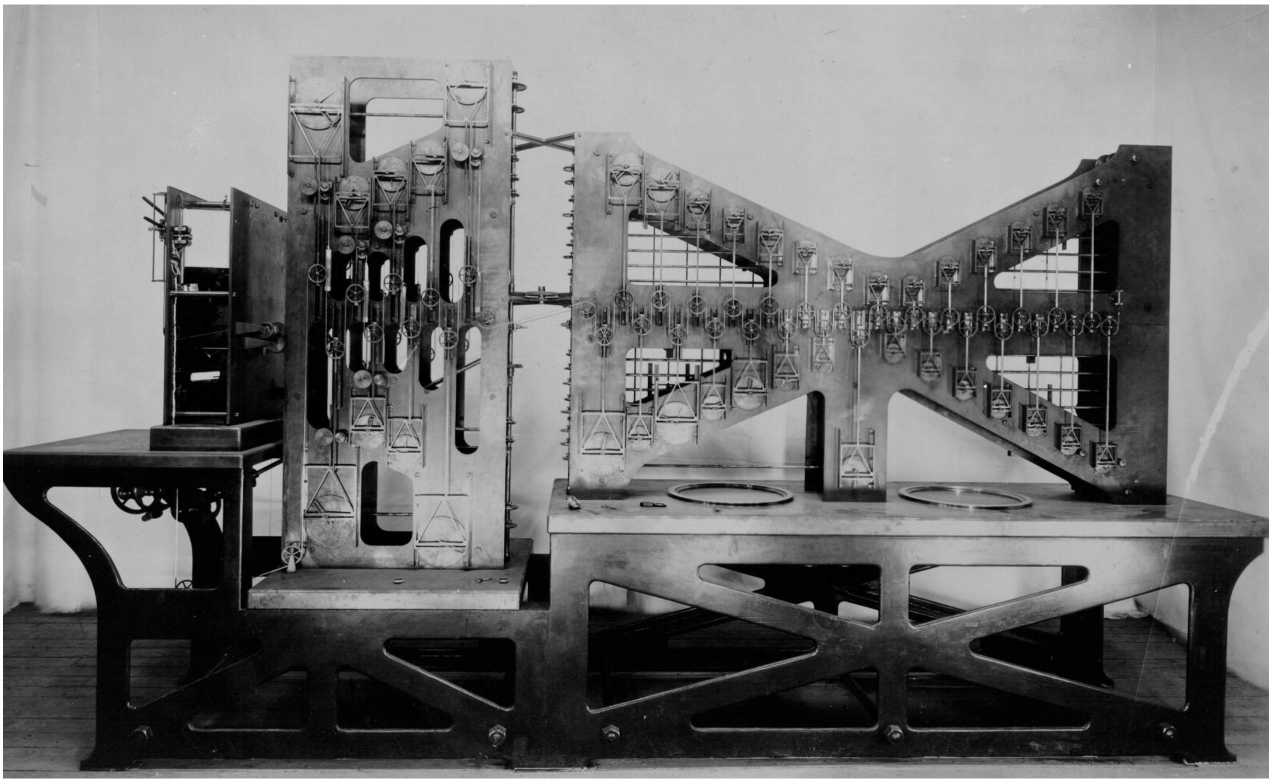

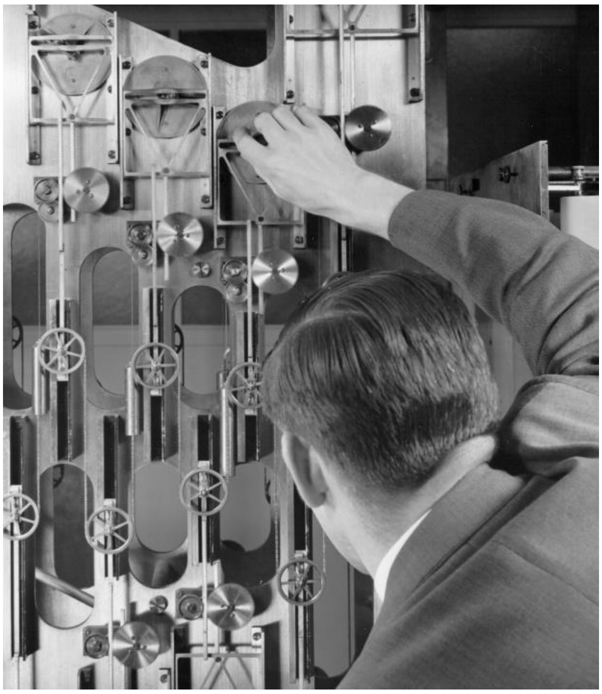

On the left you can see the wire that performs the adding function. Incidentally, it was found that this wire was the source of most of the inaccuracies in these machines. The best solution was a nickel tape – the thin dimension gave flexibility so that the tape conformed to the wheels. The width gave it sufficient cross section to have enough strength to avoid stretch.

On the right, you can see how the mechanism converts rotation into a sine function using a pin in a slot.

This machine sums harmonic constituents to generate predictions; but analysing the tide height to obtain the harmonic constituents in the first place is a much more interesting problem. This isn't simple adding anymore. It is calculus. It requires integration.

The basic process of mechanical integration already existed in a machine called a Planimeter which calculates, for example, the area of a field on a map. Its central component is a continuously variable gear – that is the cone in the picture. The output is the area within the curve followed by the stylus. Thomson's genius was that he realised he could do a similar thing with a tide height graph as input. The output would be the harmonic constants.

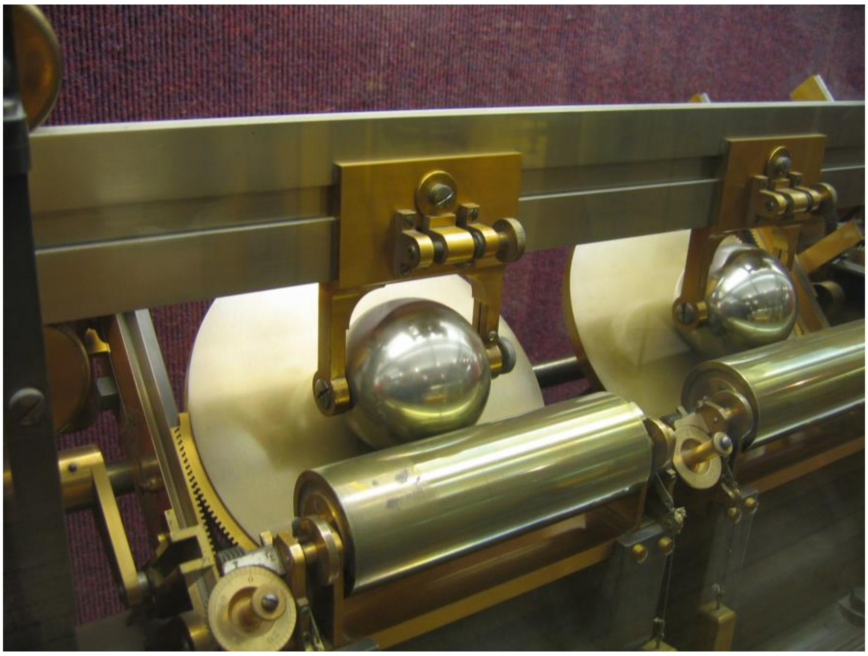

The result was a mechanical integrator –the roller, disk and sphere tide analyser which he invented in 1876. This is a fundamental leap from the tide predictor. It is not a simple calculator. It can solve second order differential equations. The same idea was used for mechanical gun laying until electronic computers took over.

We can only image the effort in making these one-off machines to the bleeding edge tolerances required to get accurate results. What I take from this history lesson is that all the effort had to be focused on the calculations. The calculations were everything.

Harmonic analysis by purely mechanical means was a triumph. But I think their focus on calculation stayed with the later generation as they started to use electronic computers. To be fair, the early electronic computers weren't easy to use either; but now things are different.

Let's now look as some issues with the tide process, keeping this history in mind as we go.

Problems with the Tide Process

In the past, a huge amount of manual work was required to analyse a single port and then generate the predictions. We saw the tide analyser machine. Think how this was used – someone had to manually track the tide curve with a stylus to get the height information into the machine. One slip and you had to start again. Or you could go very slowly to avoid mistakes. Either way, it took ages.

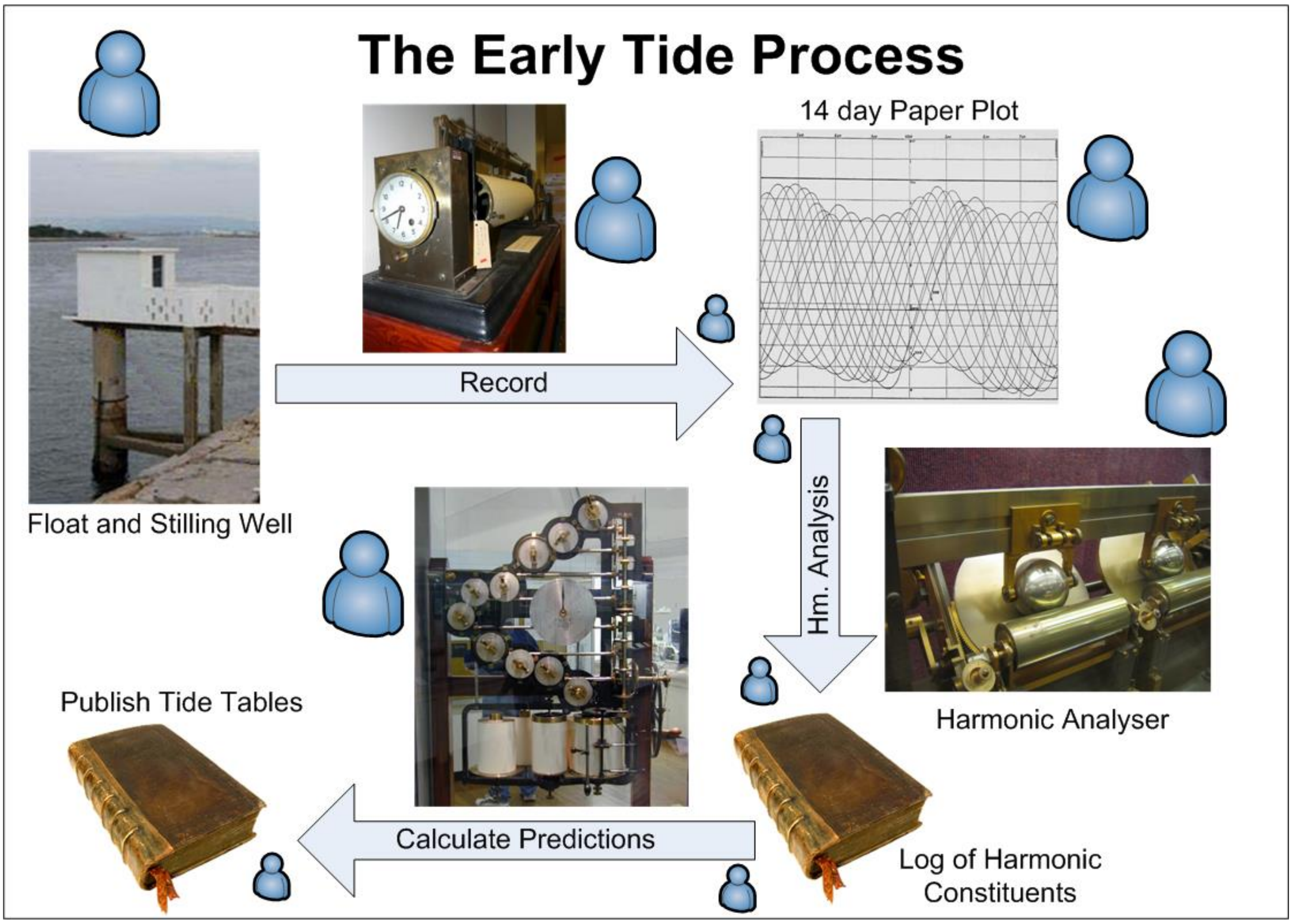

Here we can see the skilled steps in the process, which I've shown with the larger blue people. Just keeping these mechanical wonders working was skilled. Keeping the float and wires working freely. Changing the paper on the drum recorder whilst ensuring that its clock is correct. Transcribing the data from the drum paper. Checking the data. Choosing the harmonics for the analysis.

Setting up a machine was done by hand. Recording the results was too. Each of the points where I have put a little blue person is a simple manual step. Every single one had to be correct.

It took 4 hours to run a year's worth of predictions for a port, but much longer than this to set up the machine. Joseph Proudman and Arthur Doodson, again in Liverpool, used to provide predictions for 160 ports. I think this was the maximum number of ports the predicting machine could do flat out!

Clearly this was a performance limited technology. As soon as electronic computers were half way to being practical, they were on to it.

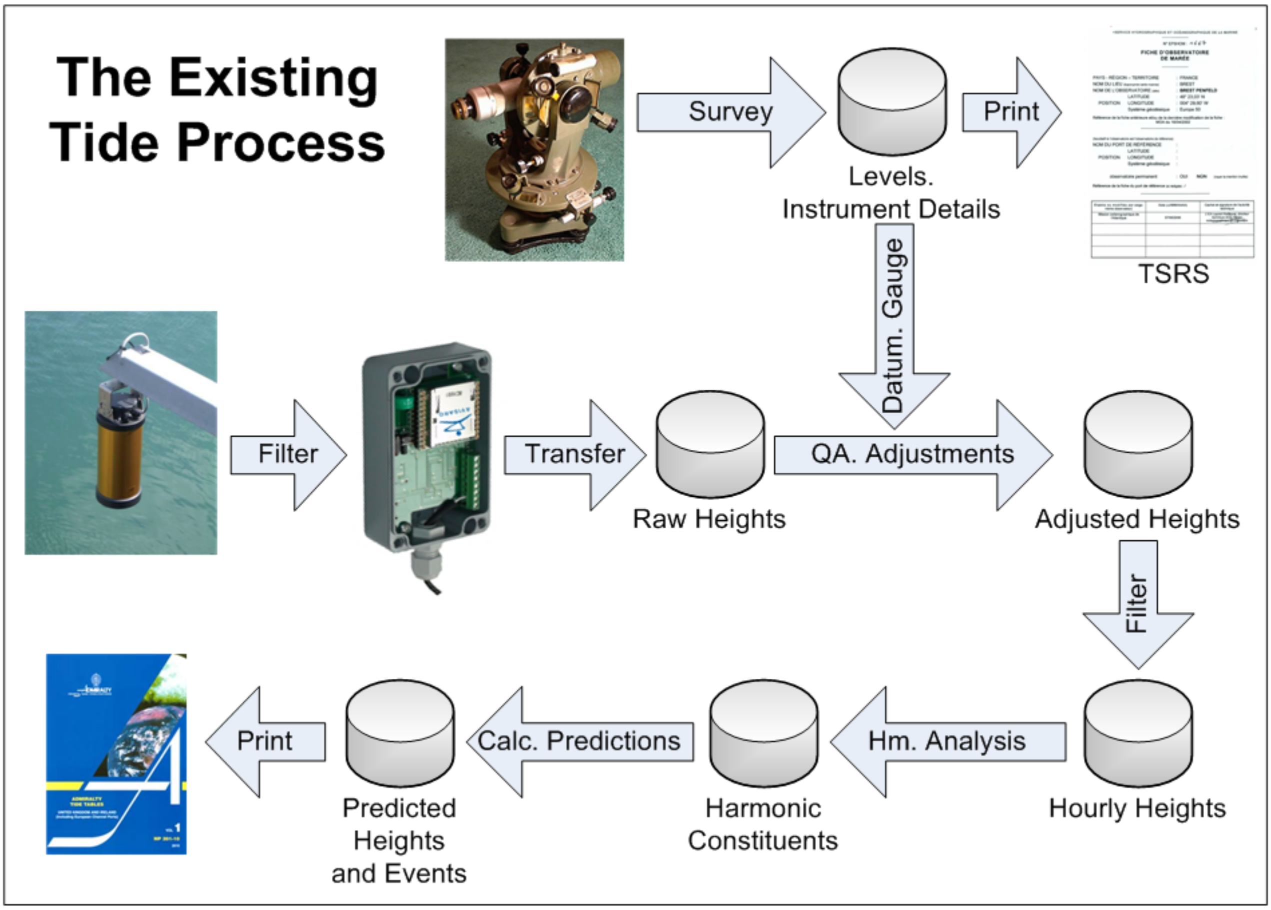

Here, again, is the computer based process that is common now. This is very similar to the earlier process, but with computers instead of mechanical machines to do the calculations, and computer files instead of paper to store the results.

Computers solved the performance problem; and they were accurate too –they did not introduce additional errors beyond those in the data.

There are a couple of big improvements that aren't shown on the diagram.

1. Recording the measurements digitally.
2. Computerised typesetting.

Both were major leaps in reliability and were a tremendous saving of labour. They meant that more ports could be processed.

CHERSOFT

Being a computer early adopter must also have been a major effort. But it was also a great success. But notice one thing; the scope of what was automated was exactly the same before and after electronic computers.

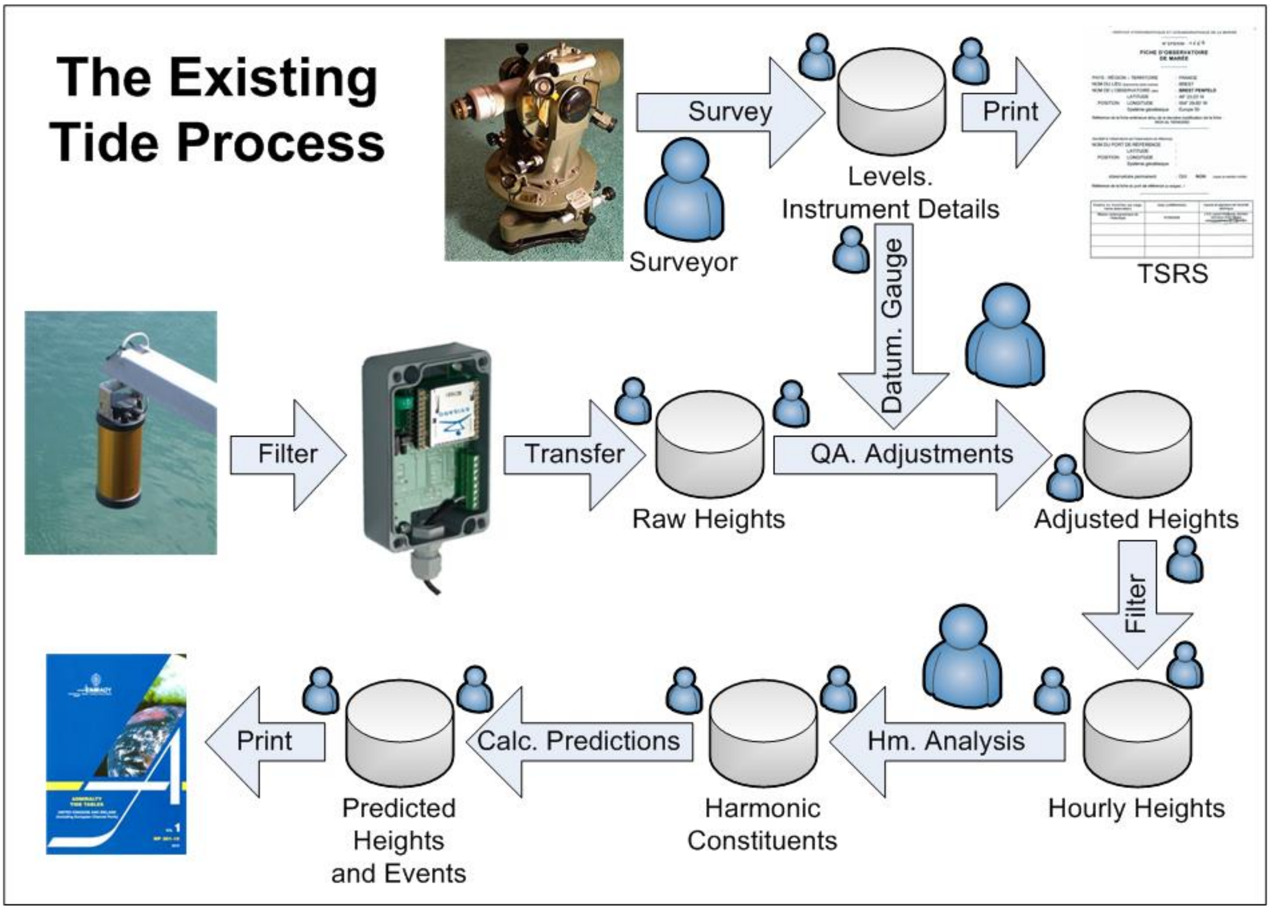

The skilled steps still exist –surveying, inspecting and checking the data, choosing the harmonic analysis strategy. People are still doing this work. Again, skilled steps are marked by the big blue people on the diagram. The simple steps are still there as well – represented by little blue people. These steps include choosing the correct files, putting the results in the correct place and remembering where you put them. There are even more steps here than with the earlier process. People are still doing these simple but critical steps manually.

When electronics replaced mechanics, the old organisation and procedures lived on. Manual organisation of the processes was necessary with the earliest electronic computers and just continued even when the computers got much better.

Gradually, computers got faster. But, more work appeared to fill any potential vacuum; there was great demand for tide predictions at more ports. The Admiralty Tide Tables got fatter and fatter. Eventually it was split into 4 separate books.

This is how most tide offices operate today. The data is stored as computer files. The computer does the calculations, but people still do all of the organising.

Things are happening that will make this way of working harder and harder.

For example,

1. The number of tide stations is increasing. There's more stuff to keep track of.
2. These stations are more likely to have permanent gauges, so the data from them is continuous and the quantity ever increasing. There's even more stuff to keep track of.
3. We keep more raw data. Soon it will be 1 minute samples instead of 6 or 10.
4. There are more uses for the data. There are more different types of publication – books, websites, exports to other systems, even the need for continuous, just in time, publication. This all adds up to wanting to extract more data more often. This can interfere with the ongoing work of updating the tide records because we normally only want to publish fully checked and approved data, not what people are currently working on. The two can be kept separate manually, but it gets harder the more data there is.
5. Tide offices need to do more with less – they have fewer and busier staff.

So what can we do about it?

A better process

The simplest improvement is to just display the data better while checking it

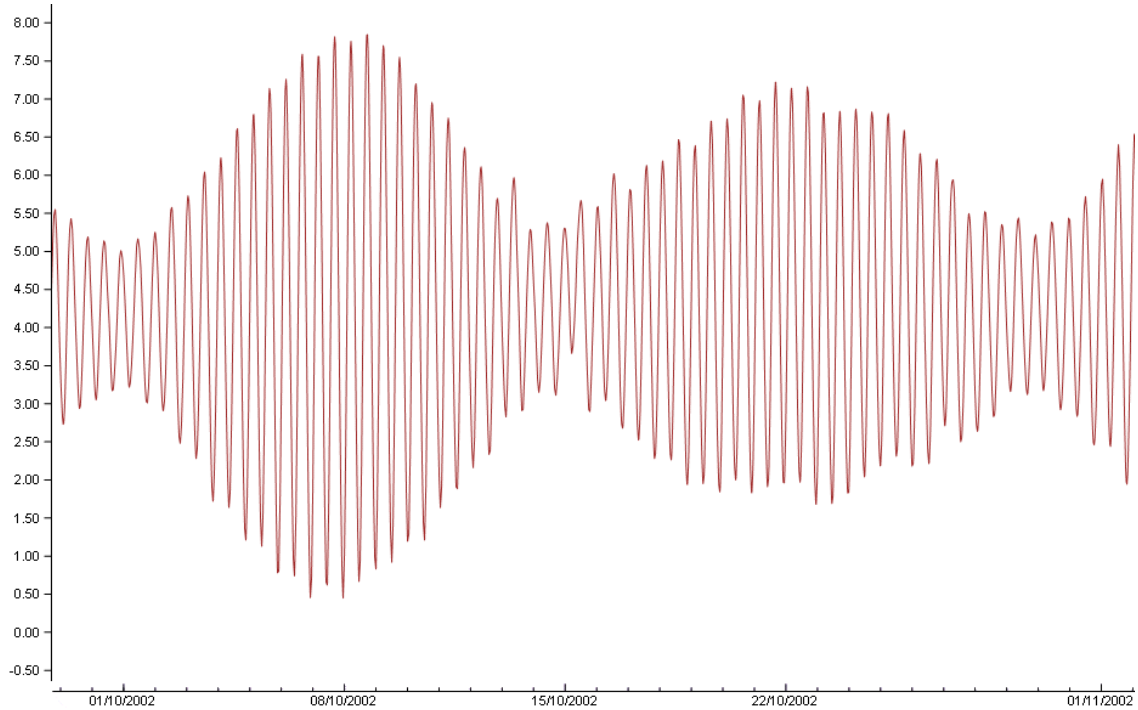

Here is a simple graph which was perfect when you only had a month or two of tide data. But it doesn't work so well for a hundred or a thousand times as much data. Many people are still working from printouts like this. Now we have loads of data; what does 100 years worth look like?

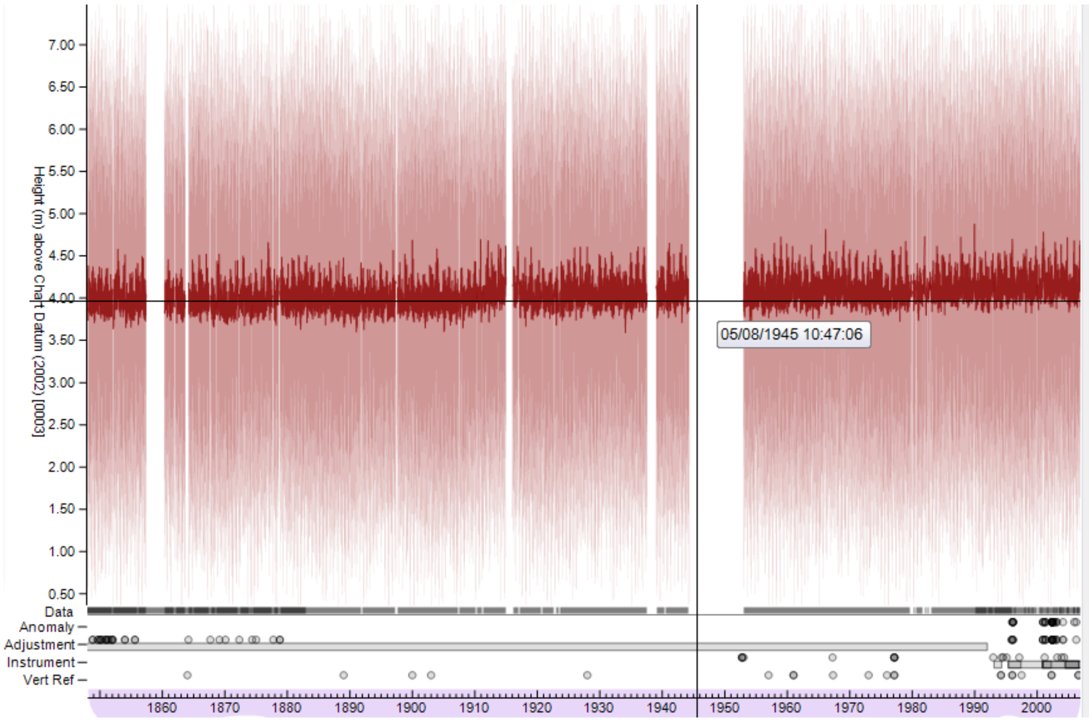

Here is 140 years of data for Brest. The light red is the hourly height data. The dark red is daily mean sea level. The gaps are mainly wars.

It is a bit of a sideways jump, but think, for a moment, about the performance of video games. Their ability to handle vast and complex scenes at fantastic speeds. That same computer can just as easily display hundreds of years of tide data. And zoom in and out of it to see the detail at the flick of a mouse.

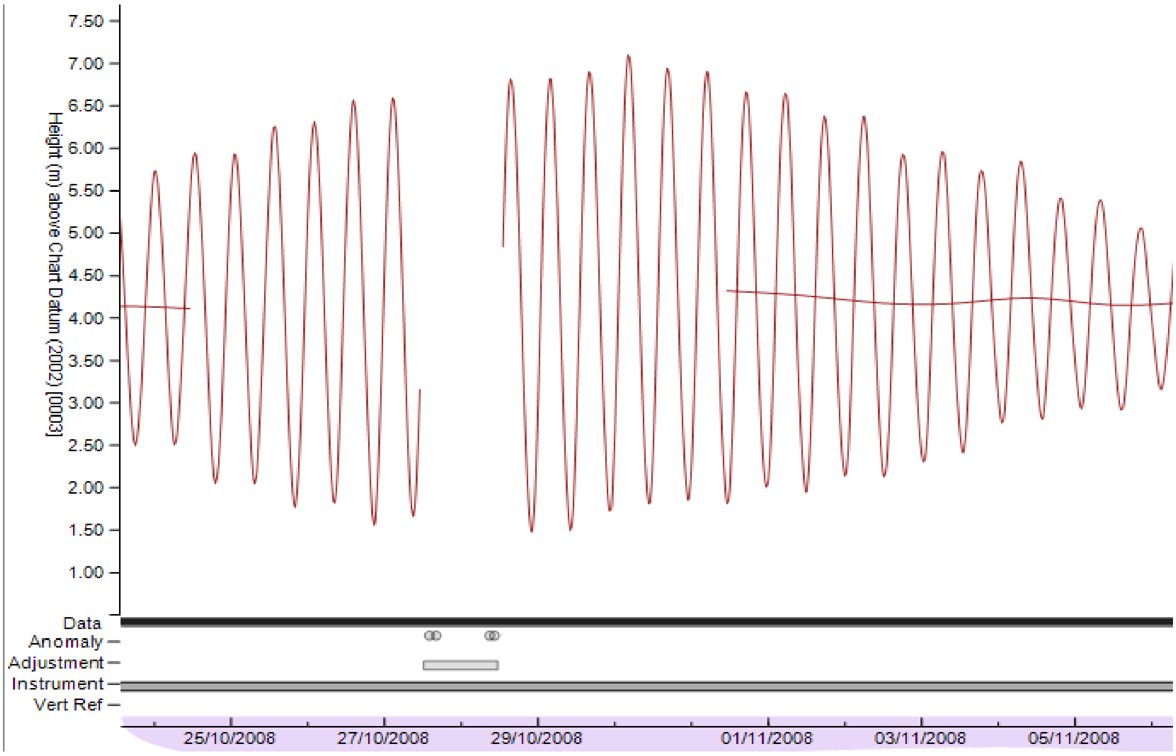

Here is 14 days. It’s easy to see the anomaly in the graph, where the instrument hasn’t been recording for some reason.

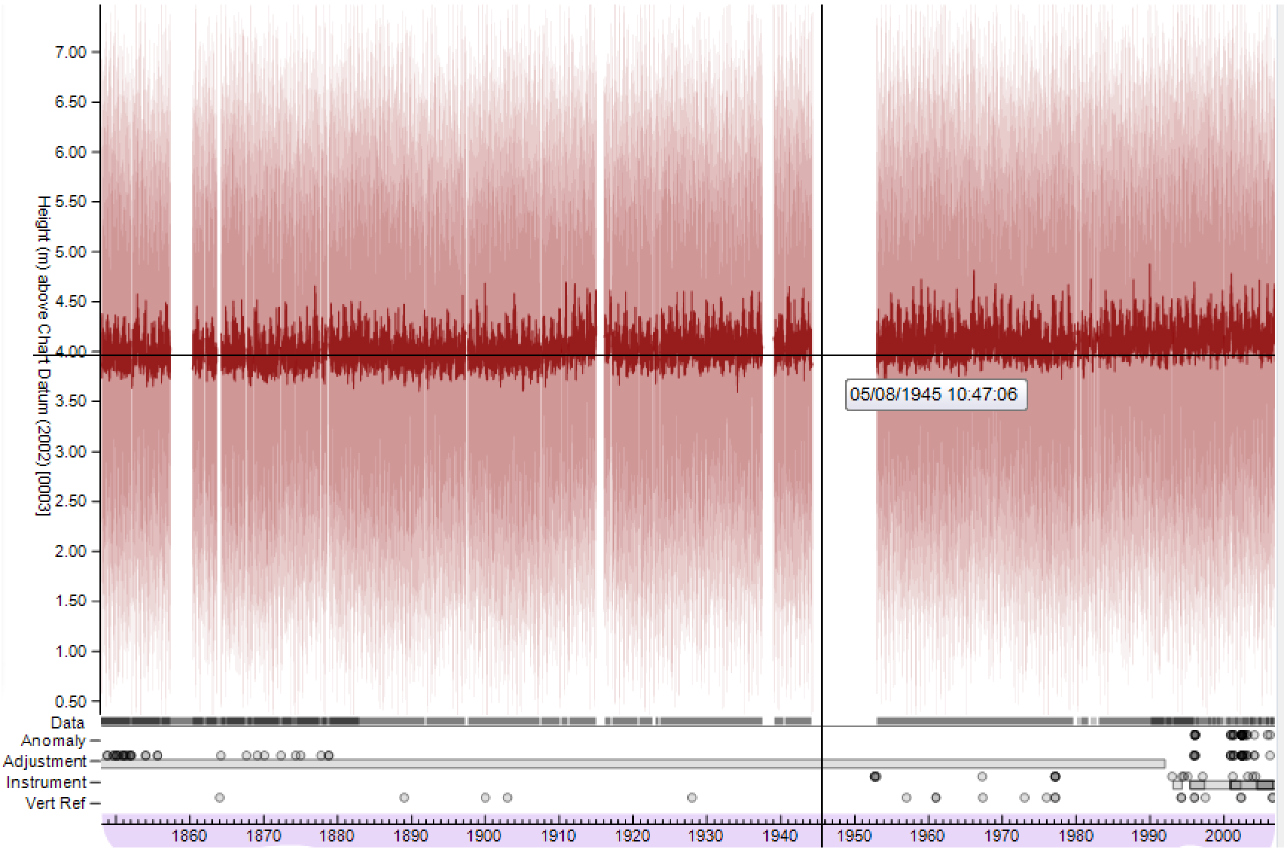

With a graph that covers a huge time span like this one of 140 years of data, we need help to find the interesting bits. Imagine now that the computer knows the date of every event at the tide station – every survey, gauge replacement, calibration. It could show these on the graph – after all, they are the most likely times for errors to be introduced into the height record. Each one could be a dot on the graph.

Now we can calculate the residual and mark all the interesting bits of that on the graph.

We can let the hydrographer edit the graph directly, and add comments and adjustments. Most of the graph will be fine – the aim is to automatically highlight the few places worth checking. Now we have a useful tool even when we have many years worth of data.

Static printouts of tide graphs simply don't give a sufficiently directed view of the data.

Next we consider data storage. In most existing tide offices, as we have seen, people have to know their way round large directory structures of tide data. Each port might be a directory. It probably contains several subdirectories for different types of data – raw height, hourly height, checked, unchecked etc. Anyone working with the data has to choose the correct set of files for the tide station they are working on. This is a nightmare task for any human, but a suitably structured computer system can handle it easily and get it right every time.

This is simple stuff – even a computer could do it. There is no reason to have to choose the correct height files for a particular port manually. With permanent gauges there are likely to be tens or hundreds of files per station, and tens or hundreds of stations. People should be focusing on the skilled jobs – the checking and analysis decisions.

There will also be metadata - the tide station record sheets - for each port. If you are lucky, these will be on line, but that often just means stored as a file – a Word document or PDF. If you are unlucky they will still be in a manila folder in a filing cabinet.

There is more to the metadata than just being able to find it. For example, without the vertical datum we cannot fully understand the height data. So a really useful tide management system would store the survey results and important events, such as tide gauge replacement and calibration, in a structured way. Of course the historic metadata is also needed, so that we have all the relevant information when looking at historic measurements as well as current ones. Structure means, for instance, that the system knows the date of each instrument change. It allows the computer system to link things together, making the whole more accessible to us.

A better process for tide data moves beyond using the computer for analysis and calculation. It uses the computer as a tool to manage the data. This is widely practiced in other areas – think of a factory, or a supermarket supply chain.

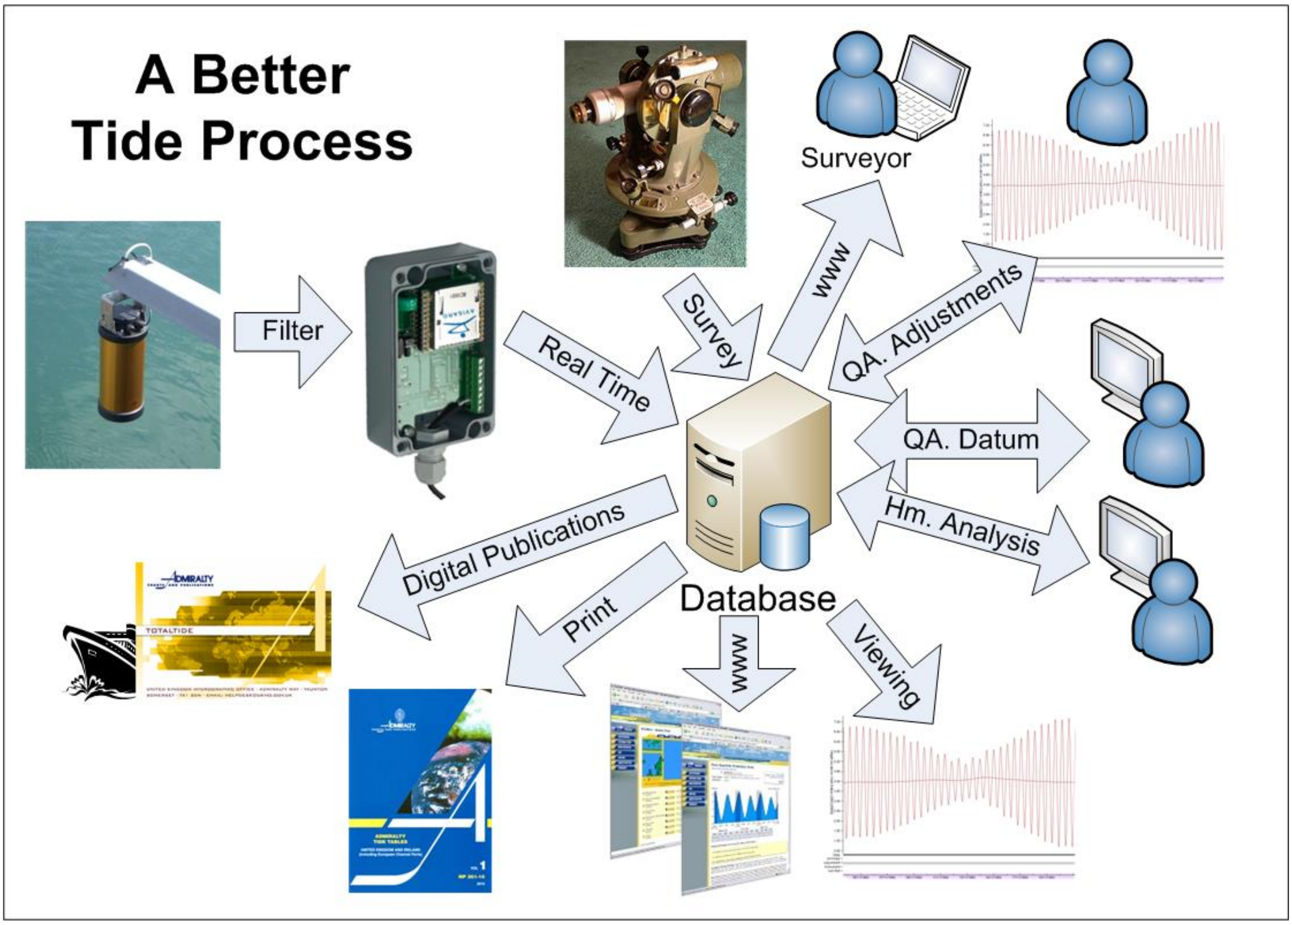

So we put a database at the centre of operations. There are still skilled tasks for people to perform such as checking and analysis, but the database is used do the mechanical simple things like retrieving the correct data for a particular station.

The joy of such a system is that, not only can it do the drudge work automatically and accurately, but it can separate the work in progress from the approved, publishable data without the user having to keep track of individual files themselves.

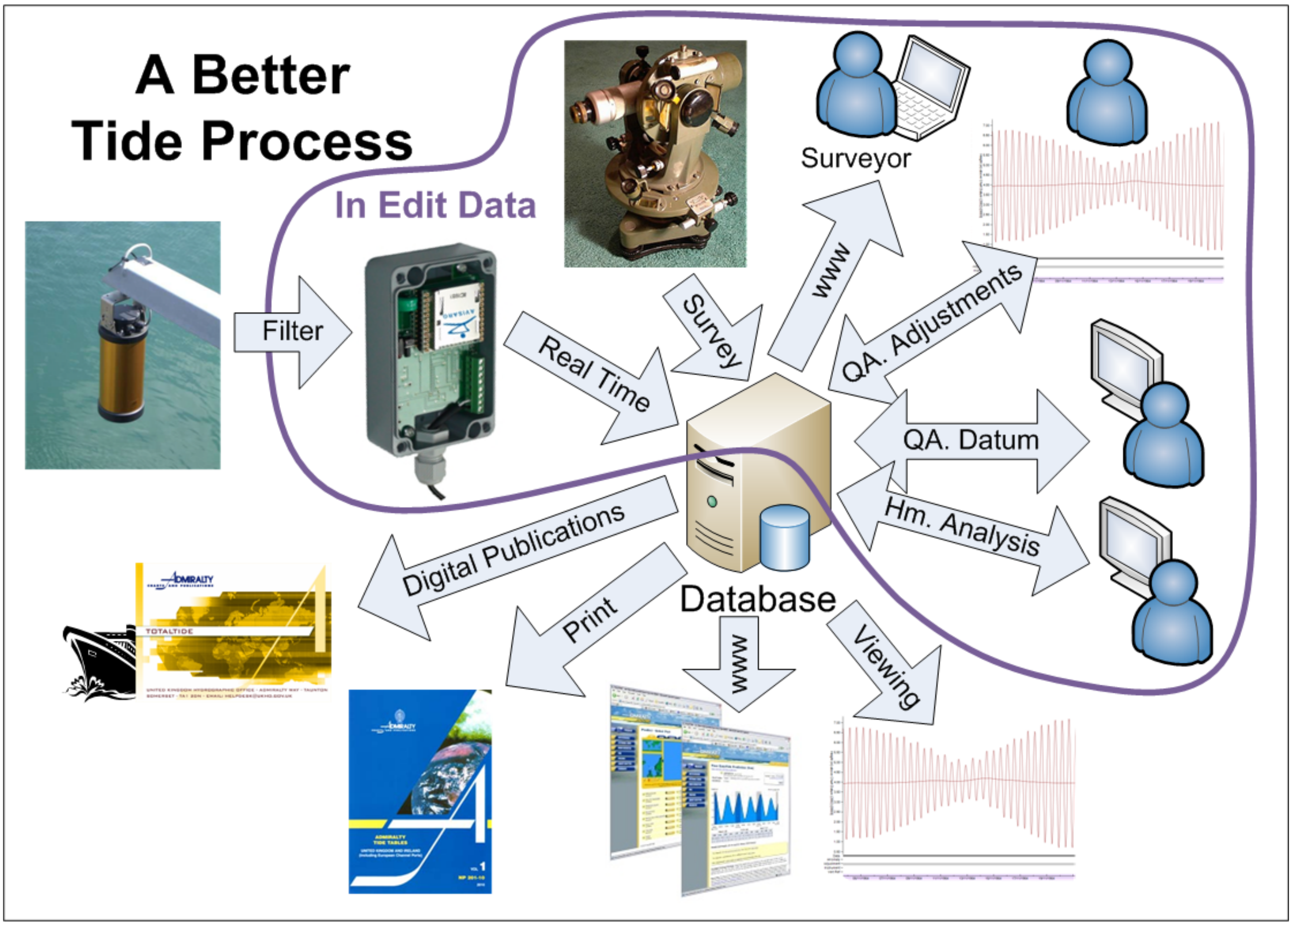

Here we've outlined the work in progress with a purple line on the diagram. This is the data that is being worked on by the hydrographer. It takes time. She doesn't want to publish it until the work is finished. Traditionally, keeping this data separate involves manually taking a copy of the files to work on. This is tedious and error prone. Tedium is what computers are best at. It is simple for a computer program to do this and to always do it accurately.

At some point, the hydrographer decides that the work on a tide station is complete, she approves it, and it is then allowed to be published.

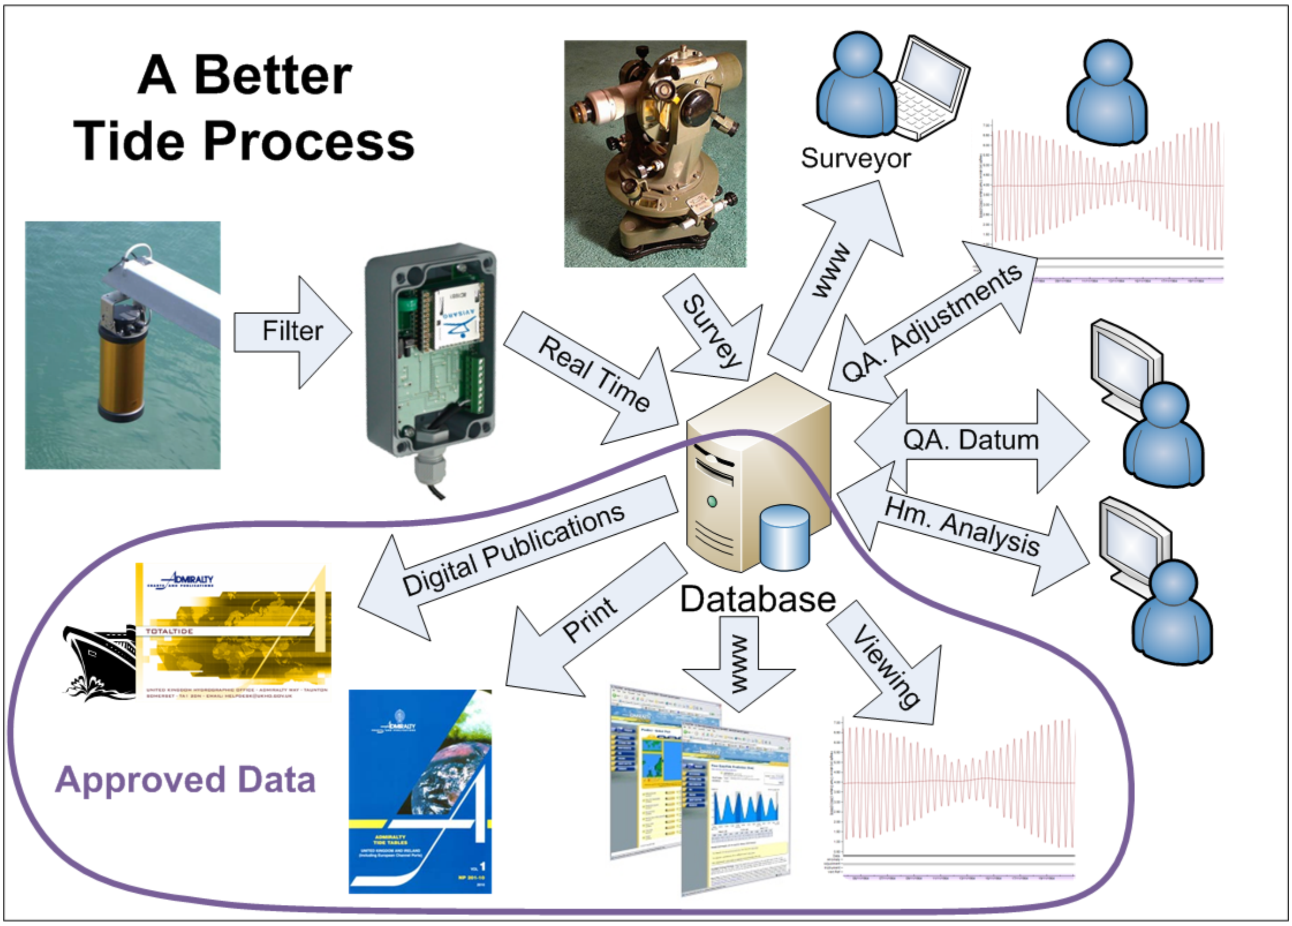

Here is the publishable data. Because the computer keeps it separate from the work in progress, it can be exported at any time, whether to a tide table book, a website, or BODC.

Given structured metadata, as we have described, the data exports can be more complete and meaningful. For instance, the entire vertical datum history could be exported in a way that is intelligible to other computer systems.

We suggest that the safest and most efficient way to handle tide data is to put all the measurements and the metadata into a database. In database size terms, this is trivial. But the rewards are that the hydrographer can concentrate on the hydrography and leave the organisation, retrieval and backup of the data to the computer.

Once all the information is in a structured database, new ways of working become so much easier

1. Multiple publications can be supported without interrupting the checking and analysis. Even continuously updated publications are easy because work in progress is separate from publishable data.
2. Distributed working is possible – someone on site, inspecting a gauge, can connect directly to the current version of the tide station record. They can update it with the latest measurements there and then without any concern about imminent publication.

3. The system can answer more sophisticated questions. Before all we could do was ask for the location of the files for a particular port. Now we can ask for the tide heights for a port for a specific date range, or we could ask for all the tide stations where a particular gauge was used, or ask for the tide height at all the tide stations in Europe at 3 PM on Christmas day.
4. By actively managing the data, the system will be able to cope with the future when there is too much data to ever be able to look at, and check, manually. Think of bathymetry. The surveyors can't even store all the side scan sonar and LIDAR data any more. They just generate a model from it and throw the raw scan data away. We need the same level of confidence in our data organisation tools for tides.

Active Data

This is a big step. It is a move away from passive files to what we at CherSoft refer to as Active Data. Active data allows you to ask questions in a domain specific way, rather than a computer specific way. So in the case of tides, we can ask tide orientated questions instead of searching for files in directories.

For example

"Get me the tide data for Brest for 1990 to 2000"

Not

"Get the names of the files in the Brest folder that will contain the data from 1800 to the present day, which I can then manually search through to find the right time period"

Future Directions

Currently, accurate tidal prediction is just for point positions – where the gauges are located.

But we now want to know the tide everywhere – continuously, not just at tide stations. This will improve sea level modelling; it will help decisions on tide power generation such as the tide windmills; it will improve navigation.

We can interpolate tide gauge measurements to make a continuous estimate

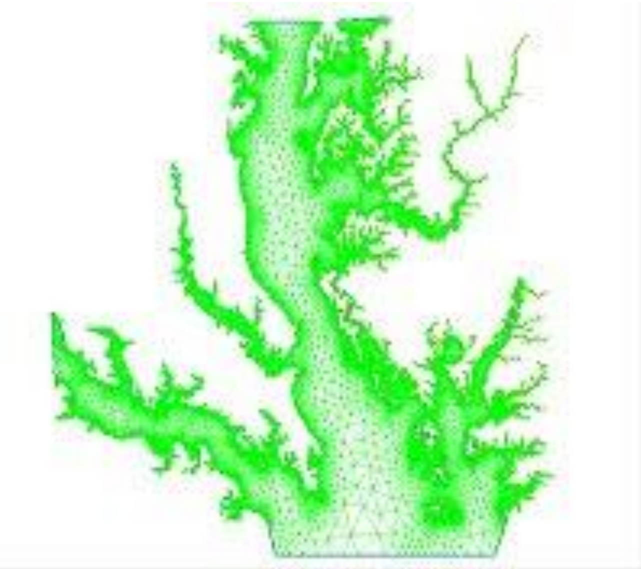

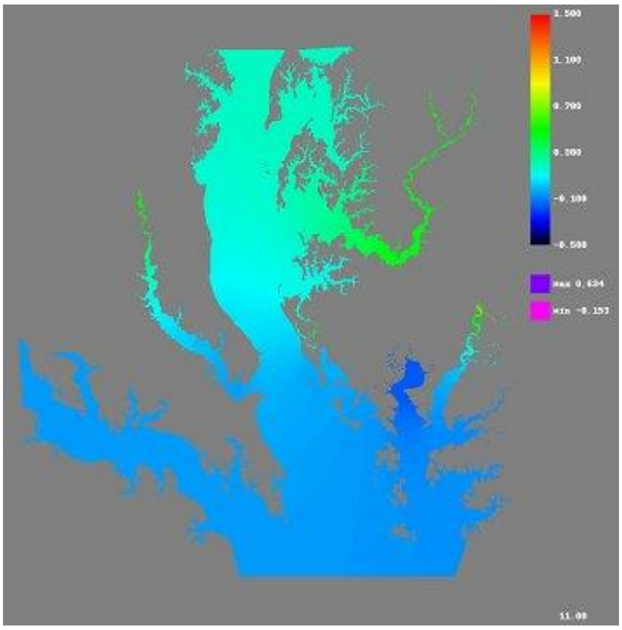

NOAA have developed a system called TCARI which provides continuous tide height from just tide gauge data. It works by interpolating each component of the water height separately using a mesh. A separate analysis is done for the vertical datum, each harmonic component and the residual.

The left hand picture shows the mesh. The triangles are smallest near the coast where the detail is wanted.

The right hand picture shows the interpolation of one component – the vertical datum. There is an adjustment of +0.5 metres in the dark blue areas and -0.5 metres in the green areas. The result is a smoother, more continuous prediction than earlier methods.

Another approach to continuous tide predictions is to use satellite altimetry data.

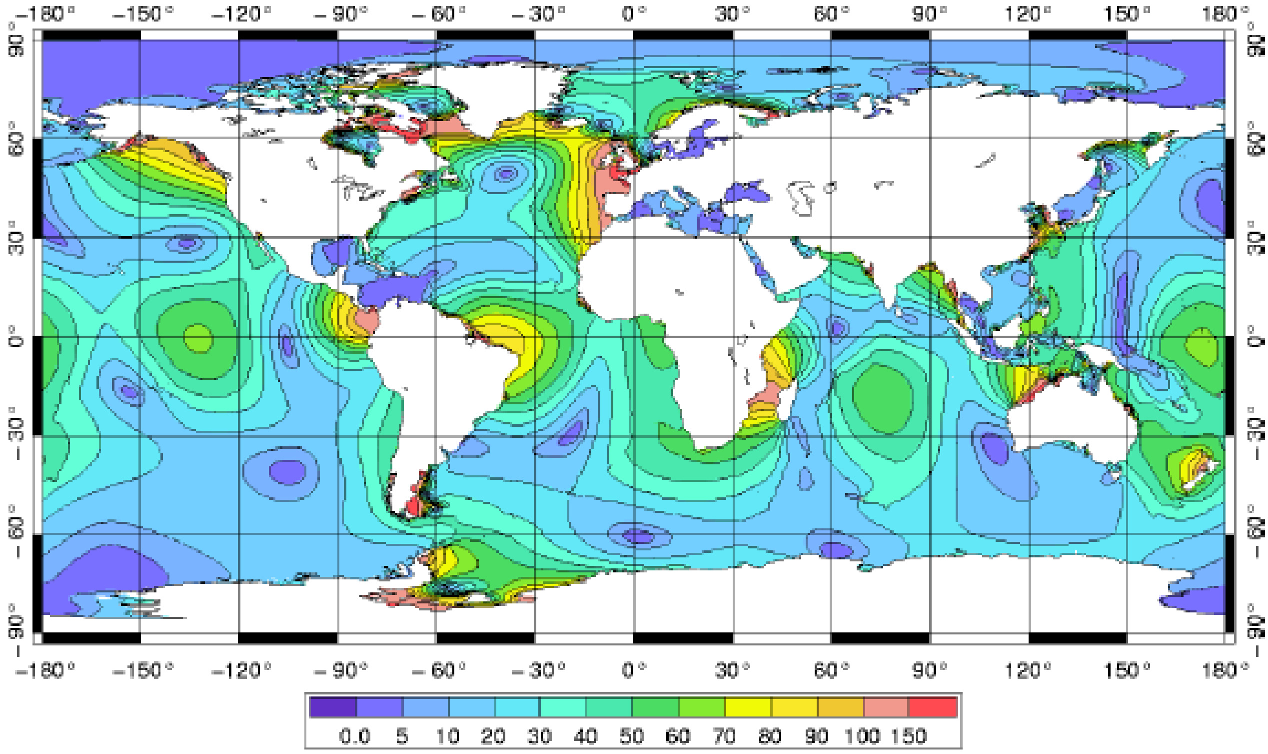

Here, the value of the M2 harmonic constituent has been calculated just from satellite data. The other major harmonic constituents have also been determined. This is worldwide and continuous, but not very accurate.

The real win is to integrate tide gauge data with satellite altimetry. This has the potential to provide predictions with tide gauge accuracy at gauges and continuity between them. But the calculations involve huge amounts of data, from the satellite, as well as lots of data from all the tide stations at once. It is much easier to provide the answers to the sorts of questions these systems require using an Active Data system than directories of loose files.

In the longer term, the individual tide stations will be subsumed into a spatially continuous prediction system. Look at what the Met Office has done with weather calculations – now obviously they are lot less accurate than we are with tides - but they have moved from making predictions at individual sites to using that information to contribute to a global model incorporating lots of other types of data, from radar, satellite etc.

This is the Jason 2 satellite, launched in 2008. It is collecting the most accurate altimetry data ever. They are aiming for 2.5 cm height accuracy.

Now take the same idea, combining Satellite and Tide gauge data and do it just in time to provide real time predictions continuously over the whole coastal area. Wouldn't the environment agency and the tide and wave power people love this, and pay for it? Not to mention the existing ships, ports and leisure users. The next step is to observe how much GPS positioning has improved.

Real Time Kinematic (RTK) GPS can resolve 1 cm horizontally and, more importantly for us, 3cm vertically. It is rather expensive, but it is electronics and computing. It will get cheaper. In 10 years time it will probably be competitive with DGPS now.

Ships could report this height information in real time via AIS. Suddenly we have thousands more tide stations. It is complicated data, as the ships are moving and their accuracy is not quite as good as a fixed tide station. But it will get better. Feed that information into the model and send the results straight back to the vessel. Now the ship can know exactly what this particular tide will be when they dock in 2 hours time. This sort of certainty will make docking and other marine operations more precise, easier, safer and, in the long run, cheaper.

All these advances need ready access to both tide data and some of the metadata, particularly the vertical datum. A managed database solution will serve this much better than directories of loose files. There will soon be so much data that manual checking will be impractical – the system has to ensure it is correct by design. Active Data does this.

Our Conclusions

Having reviewed the uses of tide predictions and the tide prediction process, we can see the importance of accurate and long term tide data.

We have seen that many organisations made fantastic progress when the technology was tough – examples are automatic data recording in Victorian times and electronic computer solutions when a computer occupied a whole room or even several.

But it seems that this is where progress stopped. Many organisations have tried and trusted calculation programs, but they still do all the data administration and organisation manually. The computer disk is just a giant filling cabinet and the drawers are still being opened and closed by hand.

We suggest that the time is right for a technology leap to a more complete computerised solution for tides where

- Tide data and metadata are managed together in a structured way: as Active Data.
- The hydrographer can concentrate on hydrography and leave the organisation and storage aspects to the software.
- We use modern visualisation techniques to speed up checking and validation, and get away from the static printed graphs dictated by 1970s computing.

The infrastructure will then be ready to handle the much larger volumes of data created by the increased numbers of permanent gauges and higher capture rates.

Then the tide data will be ready for the big leap after that – to continuous predictions based on the merging of satellite altimetry, tide gauge and vessel data, eventually in real time. This will eliminate a big chuck of uncertainty from all users of the sea whether they are involved in shipping, energy extraction or costal protection.

We know that very long term tide records are important. Who knows what other uses for this data will crop up? After all, the Victorian and Edwardian recorders had no inkling of the use to which their data is now being put.

We don't know all the possible uses for our data in the future, but I suggest we can help by getting organised now.