The pullometer challenge – part 1

Before Christmas (p.1207) I announced that Richard Major of Oswestry had developed a 'pullometer' – a device for displaying the force used during ringing – that satisfied the first part of the challenge I set a year ago, and that I had presented Richard with a cash prize. Here is the full story.

The story so far

Twenty years ago people discussed the idea of measuring and displaying the force applied to the rope. A few people had worked on it but nothing was available for widespread use. When I announced the Pullometer Challenge a year ago (RW 9 Dec 2016), I explained its training potential and described several approaches, all of which seemed to have pros and cons.

Five of the people who responded to my December 2016 article were currently active: two using devices to measure changes in the ringer's weight, two using motion sensors on the headstock and one using a device to measure rope tension directly. All five sent me promising preliminary results.

At the time if I had been asked to guess which would succeed it would probably have been weight measurement, because if the drawbacks could be overcome – spurious body movement effects and not being synchronised to the bell's movement – it would be easier to buy and install than a system needing fixed hardware. I felt most sceptical about direct measurement of rope tension.

During the spring various developers sent me preliminary results. I always tried to provide helpful feedback, for example about how to relate the force waveforms to what the bell was doing (when it was over the balance or at peak swing). If I offered substantive advice to one, which I felt might give an unfair advantage, then I copied it to the other developers.

By the summer, only Richard Major in Oswestry was sending me results, and in October he told me that the system would soon be ready for evaluation. So I booked a trip at the end of November to try it out for myself.

Trying the system

My main interest – apart from having a go – was to see if I could 'read' the resulting display, and make sense of it in terms of what I had been doing on the end of the rope. I wanted a realistic test where I was focusing on the ringing rather than just experimenting to see what happened, so I asked for the test bell to be connected to a simulator as well as the pullometer. I rang a few rounds and then pitched into a course of Plain Bob Minor. I was ringing quite fast and my rope was too long, so I had to work quite hard to strike properly.

The force was displayed as a vertical bar in real time but I tended not to look at that and focused on ringing the bell. When I stopped the data was analysed and displayed. I was pleased that with no prior experience of the system I quickly identified the difference between handstrokes and backstrokes (which were not marked in the version then in use). I could also see where I applied force to change speed while hunting and dodging. I spotted a few extra places where I had applied a bit of more force, which I assumed were when the bell didn't quite do what I expected, making me put in extra effort to stay on track.

Figure 1 shows the main display of force against time. The alternating yellow/buff areas show handstroke/backstroke and the white 'crack' down the middle of each separates the rising and falling stroke, making it easier to see how much force was applied before and after the direction reversed – the division between checking ad pulling. Both these aids to readability were added after my visit. The ticks at the top also show where the bell reverses direction, and the pink bars show when it is beyond the balance point (which in this extract happens on the over-blows of two dodges, and in 2nds place coming off the front).

The system lets you overlay two separate whole pulls for comparison. This will be useful to compare a learner's handling with that of someone more expert, for example a short snatch versus a long steady (and probably lighter) action. It was demonstrated to me by comparing different bits of my ringing, Figure 2.

To make the compison easier the curves are 'fitted together' by slightly adjusting the two time scales. The blue time marks at the bottom (that go with the blue curve A) are slightly different from the green time marks at the top (that go with the green curve B). The (orange) reversal points for both curves are aligned.

This example compares two whole pulls hunting down through 5-4. With 'perfect handling' I would expect to use a bit more force than normal to move the bell into 5ths, and then let it 'roll down the hill' using rather less effort through 4ths. But the blue curve (my first two blows in the method) show that I used more force to get my backstroke down into 4ths. Obviously I didn't quite get the speed change right on the first blow – something I have often observed in learners. The green curve shows the same point in the following lead, by which time I had settled down and where I handled it better, using a bit more force at the handstroke to change speed, followed by a more relaxed backstroke in 4ths.

Such subtlety is more advanced than I expect will be needed in many teaching situations, where the comparisons will relate to gross behaviour like over pulling, snatching, and so on. But I am encouraged that even with an experienced ringer on the rope it is still possible to 'read' the displays (especially with no prior experience of interpreting force measurements) to understand what happened between ringer and bell,

As well as force against time, the system can show force (vertical) against bell position (horizontal), Figure 3. The balance points (TDC = Top Dead Centre) are shown by thick black markers at left and right (handstroke & backstroke respectively) with BDC (Bottom Dead Centre) in the middle. The 'cloud' of curves show every stroke rung, with blue moving from handstroke to backstroke and red moving from backstroke to handstroke. I have added arrows showing the direction of movement. You can highlight an individual whole pull (in this example when I was lying two blows in 6th place).

It is immediately apparent that the bell never went over the balance at backstroke but it did quite often at handstroke, where I tended to use larger forces.

The humps before and after the handstroke aren't due to my handling – they represent the force needed to accelerate the rope when it rapidly reverses direction as the garter hole passes the pulley.

Another interesting feature is the consistent difference between rising to and descending from both strokes. More force is almost always applied on the way down. Thinking about the underlying physics, I realised that the net area inside the loops represents energy lost in the system – due to rope friction, etc, and maybe due to handling technique. It will be interesting to see whether that feature proves to be a useful diagnostic aid when used with learners.

When I first rang the test bell (before Abel kicked in) I commented on the rope noise – a familiar sound in towers where the rope rasps on a rope guide. The noise coincided with when I pulled, and I instinctively tried to tidy my handling to reduce it. I was told that it wasn't from the guide but from the rope tension sensor, and that one of the local ringers was reluctant to ring that bell because of it. I didn't notice any effect on the handling, and once the sound of bells started I didn't notice the rope noise anymore. Although I didn't notice it, there might have been a small effect on the rope for which I was automatically compensating, but which would be more significant for a less proficient ringer, so this is something to consider when moving towards a marketable system.

Overall, what I saw convinced me that the system had satisfied the requirements of the 'feasibility' challenge, but there is obviously some way to go to meet the 'marketability' challenge.

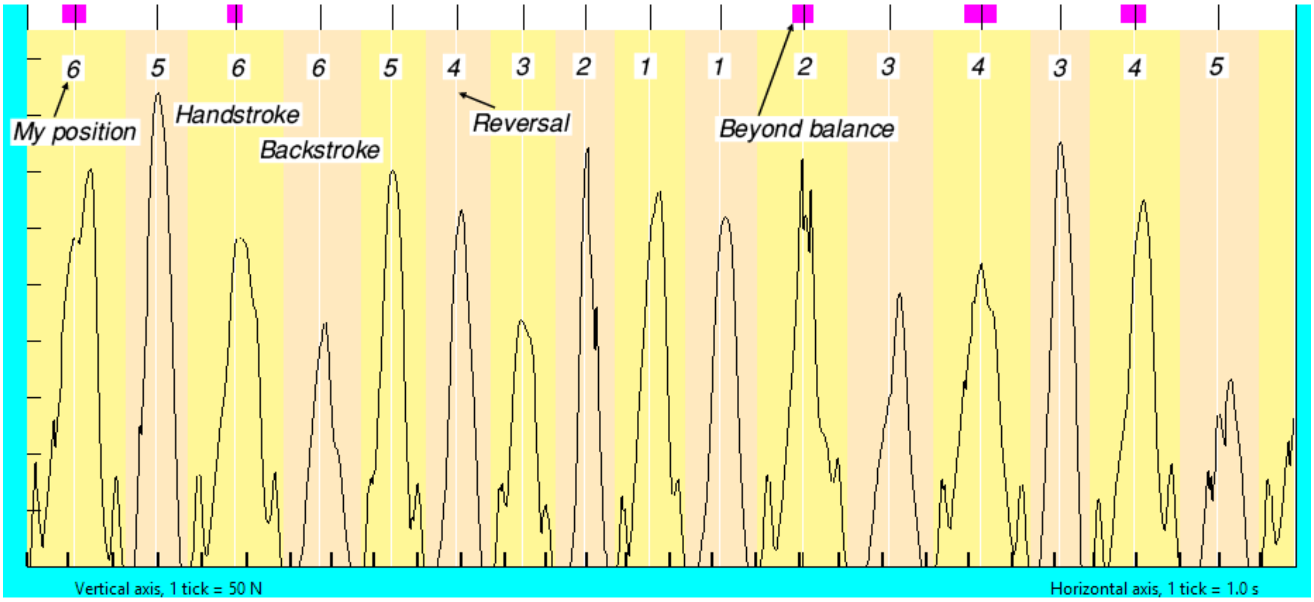

Figure 1: Extract from my ringing Plain Bob Minor – from the 5-6 up dodge to just after 3-4 up dodge

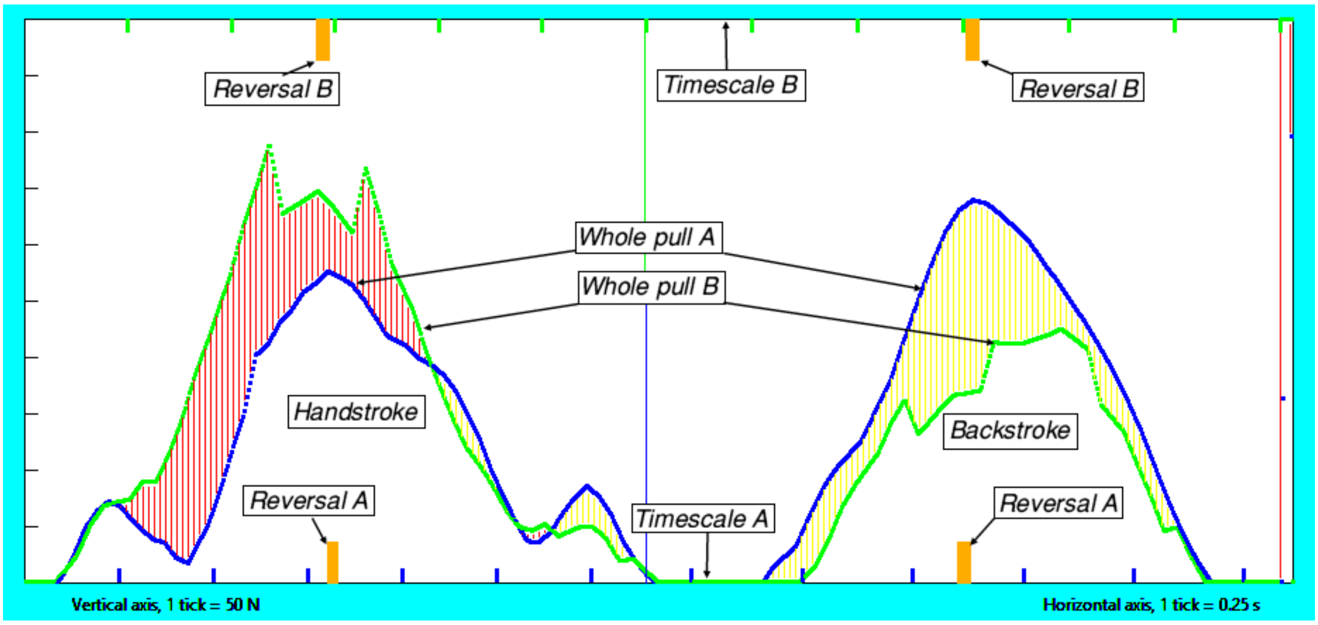

Figure 2: Comparison of two different whole pulls, one with the backstroke going well over the balance

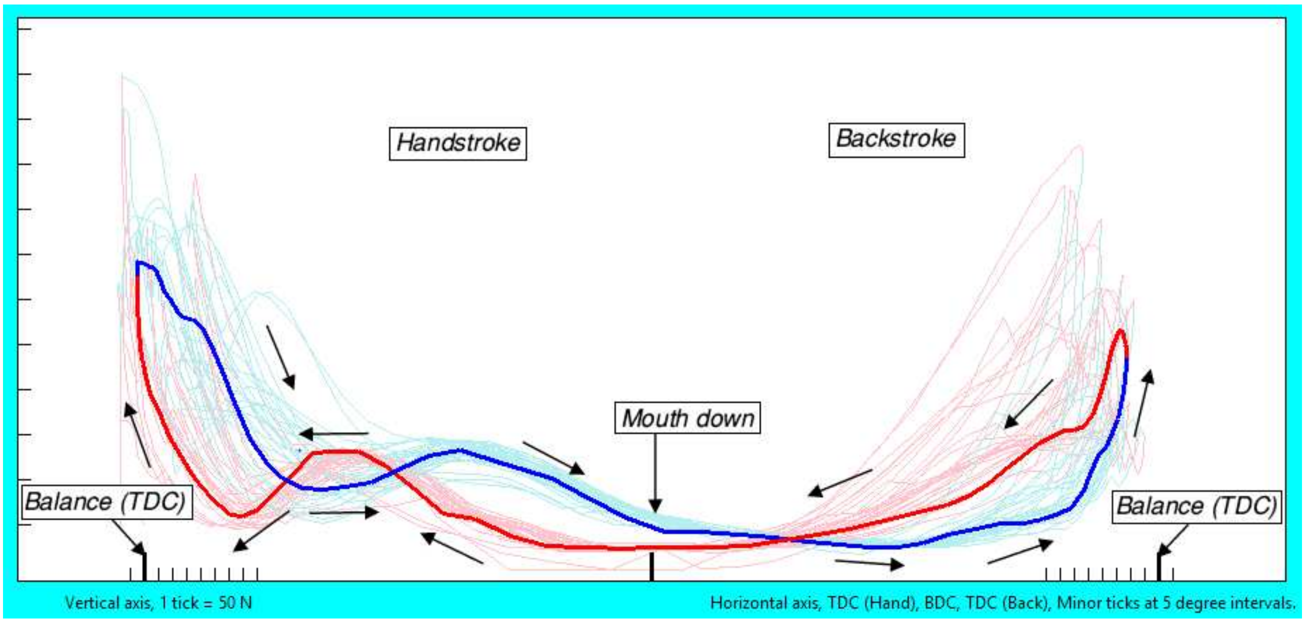

Figure 3: Force v bell position with one cycle highlighted
(Blue, L-R = handstroke to backstroke; Red, R-L = backstroke to handstroke)

Mk 1 force gauge

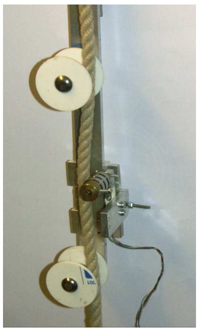

Mk 2 force gauge

Solid state gyro on wheel

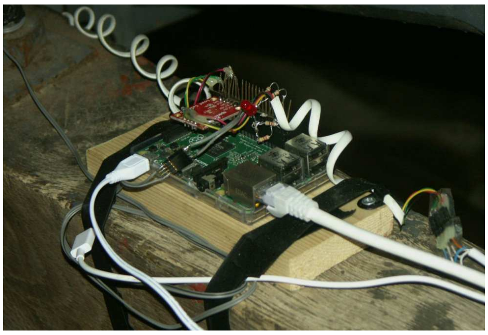

Processor board on frame

The Oswestry prototype equipment

The system uses two sensors: one device to measure rope tension and another to measure bell movement. Both communicate with a microprocessor on a circuit board fixed to the bell frame, which transmits data to a laptop in the ringing room.

The rope tension is detected with a force gauge just below the bell chamber floor. It has three rollers with the rope passing between them. The middle roller is offset slightly, which deflects the rope by a few millimetres so that tension in the rope applies a force to the roller, which is attached to a strain gauge to measure the force.

A solid state gyroscope is attached to the wheel. It measures angular velocity, which is integrated to give bell angle, using the known start and end point when the bell is on the stay at handstroke as a reference to correct any drift in the calculated value of the angle.

The developer's perspective

I was interested to know how Richard and his team got to this point.

What motivated you to get involved in the first place?

When I ran the ringing at Morpeth Clock Tower I could see several ringers pulled much too hard. Talking about it didn't get through to them and I felt a visual way to contrast their pulling with what a good ringer did would help.

I had been intrigued by Andrew Chin's idea of inserting a load cell in the rope. I realised it would need to be lightweight and easily fitted. I designed a device to clip on the rope just above the sally, small enough to pass through guides and ceiling bosses (made with a lot of help from Andy Hamer at Newcastle University) but I realised I had not allowed for the rope's high velocity through the guides, so we couldn't try it out.

In 2013 I moved to Oswestry, which has a high ringing room ceiling, and I tried the device. By March 2014 I was taking force measurements but realised that interpreting them would be difficult without knowing the position of the bell. I could not think of an economic way to do that so I shelved the project.

The introduction of the Raspberry Pi [a small, economic and simple computer designed to promote basic computer skills teaching in schools and developing countries], and cheap Micro-Electro-Mechanical Systems (MEMS), inspired me to explore their potential. I mounted a 3-axis accelerometer on the headstock to monitor bell position and the point of strike. I published the results (RW 2016 p453) which could have had several potential uses including accurate simulation of striking during lowering, and use as an odd struckness meter. I realised it could also be used for force measurement but didn't give that priority. John's challenge moved me forward with the pullometer.

What determined your approach?

When I started thinking about the Challenge the two approaches of measuring force and movement came together. The three-pulley approach had been mentioned in the background to the challenge and I intended to work on that as well as my previous approach with a rope mounted sensor, but I only had time to work on one, because what followed was a lot of hard work, including mastering new software packages and different types of hardware.

What did you learn along the way? Did you make any compromises?

Initial progress was rapid, using the new three-pulley approach to measure force, with the centre pulley mounted on a strain gauged shear beam connected to an HX711 Load Cell Amplifier.

That initial testing had the bell silenced but when I let it strike, the accelerometer picked up the vibration, which defeated the analysis I was using. Ringing with other bells was even worse since their vibrations were transmitted through the frame.

I had already added a tri-axis gyroscope, intending to use it to identify which stroke was being recorded, so with the accelerometer signals proving difficult I gave up on that and extended the use of the gyroscope to give position and acceleration as well.

Would you do anything differently with the benefit of hindsight?

We used banks of 19mm roller bearings because they were cheaper than pulleys, but they contribute to the rope noise. The Mk 2 design should reduce the noise because it uses a different sensor that needs less deviation of the rope path, but a source of cheap pulleys would be the ideal way forward.

What we have developed will need significant effort to install, and I can see the merit of something easily transportable, even if it is less accurate. Perhaps there is a need for more than one type of system – one with maximum accuracy for training centres and towers where the band has the capability to install it, and a simpler device that's portable and easier to set up for more widespread use.

What do you see as the strengths and limitations of the current design?

It gathers and presents a lot of information. I'm sure the presentation could be slicker, but I've been constrained by having to learn new software. I'm confident it can be improved, and I have made several improvements since the demonstration.

The challenge asked for a graph-like presentation, but I know some ringers 'turn off' when shown a graph. I think it would be worth exploring an alternative presentation with an animated bell, and perhaps a stick-figure ringer, with a superimposed arrow showing the force vector.

How would you like to go forward?

That's to be determined. As well as developing the tool I'm keen to help develop its use for the benefit of the whole ringing community. To this end, anything I offer will be made open source.

The next article will describe other developments together with ideas about the way ahead.