A Digital Engine Indicator

(a 21st century approach to engine performance analysis)

Background

Over the last 7-8years we (Louise and I) have built a steam yacht called “Befur” *(“B for boat” you understand!)* which we are still attempting to turn into a reliable, retirement project, mostly providing a distraction in retirement and an attempt to stay fit and have fun. You can read about it (and see higher resolutions for the pictures in this article) in our blog – just say “Befur” to Google, and it should take you there, or visit http://sy-befur.co.uk

Befur is a Selway Fisher, strip-plank Golden Bay 26 hull, a John King 30sqft Yarrow boiler, a 18.5sqM junk rig sail with a Leak Compound engine – the performance measurement of which is the subject of this article. *I have included costs for most of the components, so you can see how staggeringly cheap this technology is today*.

The History of Engine Indicators

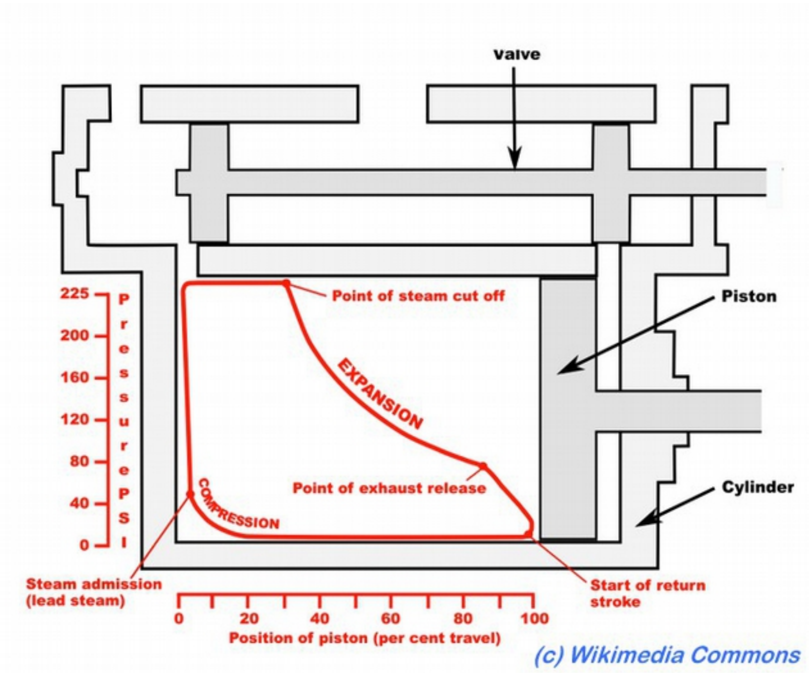

Traditionally the measurement of steam engines (and very large marine diesel engines) is undertaken with an Engine Indicator, which produces “PV diagrams”, which are plots of cylinder pressure by cylinder volume. These “indicate” the power being produced by the engine, allow analysis of valve events, and various types of fault finding. (Search Wikipedia for “Indicator Diagram” or “PV Diagram” for good descriptions).

These engine indicators were really Heath-Robinson contraptions of springs, bellows, string and pointers.

On railway locomotives all this was attached to the front buffer beam inside an “indicator shelter”, in which very brave engineers rode (often at high speed) to measure the locomotive.

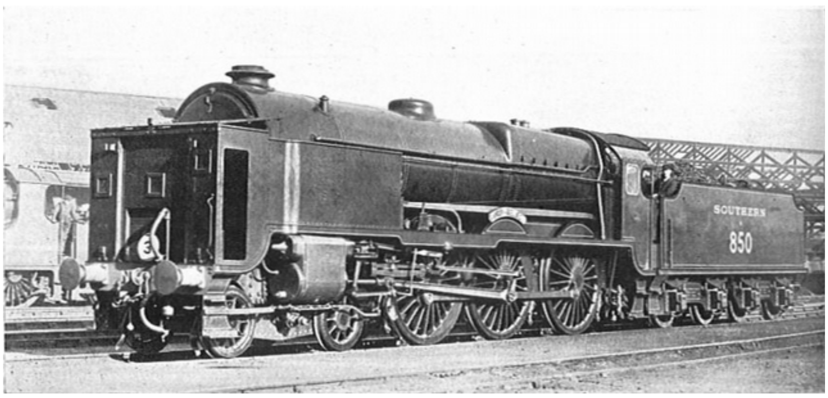

Visit Preston Services' site and find “indicator-sets” listed under “parts & fittings”, or visit this page

http://www.traintesting.com/bulletin_8.htm for more engine testing talk!

(It would appear that in the 1890's an Indicator would have cost much more than a months wages!)

The History of the Author

By background I am an engineer in instrumentation and computing (for Shell Research), and latterly in software, telecommunications and "big data". But for fun I dabble in mechanical engineering and things computing.

This background left me thinking that it would be good to instrument and analyse the performance of Befur's Leak Compound, and that modern technology should allow us to drop the string and sealing wax solutions, and build something a bit more 21st Century.

Overview Of System

There are a number of components needed to make a working system:

- **Sensors**: The very first requirement are pressure sensors capable of measuring the pressure quickly enough and accurately enough to get some meaningful data. Searching eBay will find the ones we chose, costing £15.00 or so. These claim a response time of <1ms (so cira 1,000 measurements a second or a frequency response of 500hz) and available in ranges from 30 to 500psi. Remember NPT is not BSP (I didn't!!!!)

- **Data Acquisition**: The next part of the puzzle are Analog to Digital converters (A2Ds) and some clever timing system. As it happens suitable systems have been developed to manage our car engines, and toasters! The PIC range of microcontrollers by Microchip produce a number of suitable chips.

  FIGURE 1-1: CURIOSITY DEVELOPMENT BOARD KIT

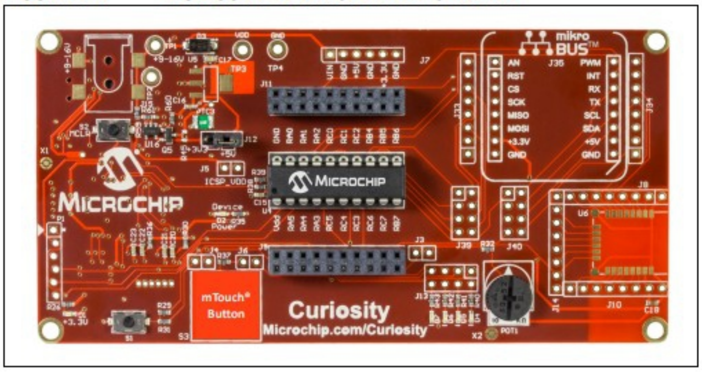

  These combine a basic 8-bit microcontroller (computer), a number of A2Ds and "Angular Timers", which are designed to take a single crankshaft pulse at TDC and divide each rotation into almost any number of equal divisions, which can then trigger the A2Ds to read the sensors.

  We opted for the PIC16F1619 which contains a 12-channel A2D, Angular Timers, I2C, DACs, Enhanced USART,SPI, I2C, RS-232, RS-485, & LIN compatible communications, 8K of flash, 1k of RAM and a host of other features in a chip that costs the princely sum of £1:32!!!!!!

  Testing shows this can manage to take 72 samples per rev (every 5-degrees of rotation) from 4 sensor channels at over 1,700rpm. This is faster than the engine will run!!!!, and exceeds the frequency response of the sensors (2,125 samples per second), but shows that the system is capable of meeting it's workload.

  (*We bought a development board (pic) for this, to save having to make one ~$45.)*

- **Data Processing**: Once the data is captured, one needs a system to draw the PV diagrams and save the data for export. For this we chose the ubiquitous Raspberry Pi.

We chose a Raspberry Pi 3 Model A+ (£23:40) – this boasts a 64-bit quad core processor running at 1.4 GHz, dual-band 2.4 GHz and 5 GHz wireless LAN, and Bluetooth 4.2/BLE – (*that's a description of a machine that would have cost a few hundreds of thousand pounds when I started with computers in the '70s!)*

We added the standard Pi 7-inch touchscreen and "SmartiPi" case that cost an additional £90 all in. This system fetches the collected data from the PIC, deals with things like angularity corrections, RPM computation, graphing of 4 channels (HP top, bottom and LP top, bottom), and providing operator data entry for test meta data.

- **Data Analysis**: We have found it useful to provide a mechanism to further process and graph the data, and this is done via a Windows machine, using OpenOffice Calc (price £0:00!!!) spreadsheets, with some Office Basic macros...

Other Stuff

Sensor Mounting

The sensors do not like to get too hot, so we have chosen to mount them via a length of 6mm PTFE hose, which should stand about 250degC. We soon discovered that the push-in hose connectors do NOT stand that kind of temperature! So they have been replaced at the engine end with brass compression fittings (don't forget to fit the brass tubes that push into the end of the pipe to withstand the forces from the compression olive). Hose failure is still a risk, so some care needs to be exercised when in action.

Crank Position Sensing

The crank position needs to be indicated to the PIC microcontroller via a 5v pulse. To do this we have employed a small Hall Effect sensor (more eBay purchases). This produces the required pulse when a magnet is passed near by. The sensor comes with a little circular penny-sized magnet that can be stuck to the engine flywheel. The sensor is held with a dial-test-indicator stand, so it can be positioned to "fire" as the magnet passes it at top-dead-centre. (There is some scope for positioning error here, especially when you change from ahead to astern). We have included an "adjustment" feature in the Data Analysis subsystem to allow for this. Perhaps a better approach would be to have an optical sensor and notch in the flywheel, which would provide a more accurate pulse. But the magnet approach means we can quickly attach the system to almost any engine.

Power & A Lack Of Time

The whole system fits inside the Pi case, and an (ebay!!!) power pack (battery) is hot glued to the back. This is very capable of providing the ~3amps of 5volt power for several hours.

We have discovered that the Pi does not have a real-time clock chip, it fetches its time from the Internet... which is OK until there is no Internet (e.g. on the boat!).

This means that when it starts on the boat, it just continues counting from the last time it was switched on, so your data files might have strange timestamps – our's did!!!

The system on the kitchen table.

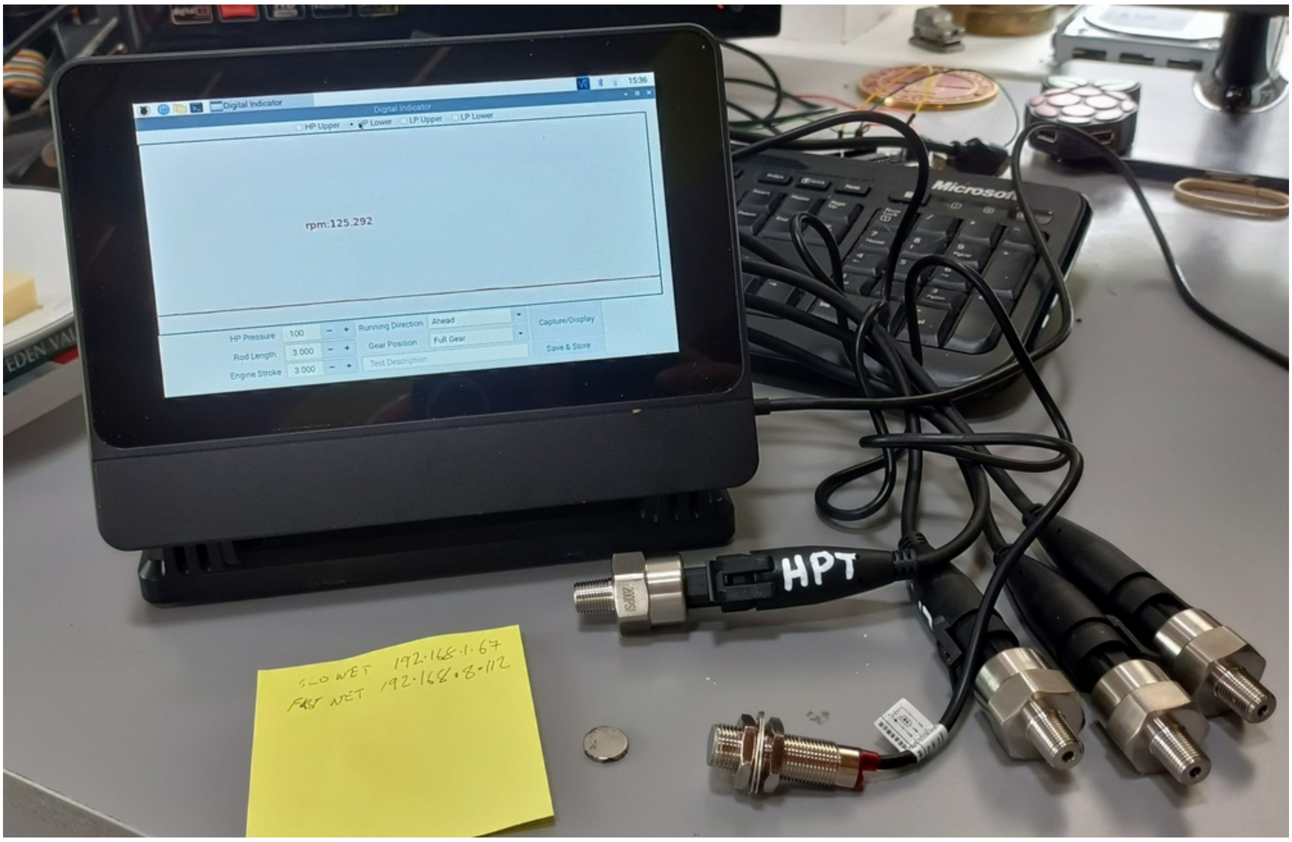

Operation

When the system is powered up and the crank sensor starts sending pulses, the PIC microcontroller immediately and continuously starts sampling the sensors 72 times per rev under interrupt control, storing the data in its memory, and computing RPM data.

When the user clicks “Capture/Display” the Pi asks the PIC to pass it one set of data. This approach means we don't load up the PIC while it's sampling. (in a future revision it might make sense to get the PIC to average the data from 4 or 5 revs, which would help reduce noise (ringing etc.) in the data capture).

The PI then does the angularity adjustment, draws a graph of the data, and lets the user inspect each sensor's results. They can also enter test data (gear position, HP pressure, ahead/astern, rod lengths and stroke (these are needed for a bit of Pythagoras and trigonometry to do that angularity correction), and any other test data like (“flat out!” or “lots of knocking”). Then clicking “Save/Store” will save the data onto the Pi's SD card.

The Results

We have had a couple of data capture sessions. These done with the boat moored up, and while we are a long way from being able to put a figure on the IHP of the engine, it is giving us a feel for the issues we should be addressing in terms of engine performance.... and the issues we still need to address in the development of the Digital Engine Indicator.

As the pictures that follow show, the data recorded from HP cylinder suggests things are not too bad, and we can see the effect of notching up the gear...

Full Gear - Ahead - HP inlet 150psi – 214.2rpm

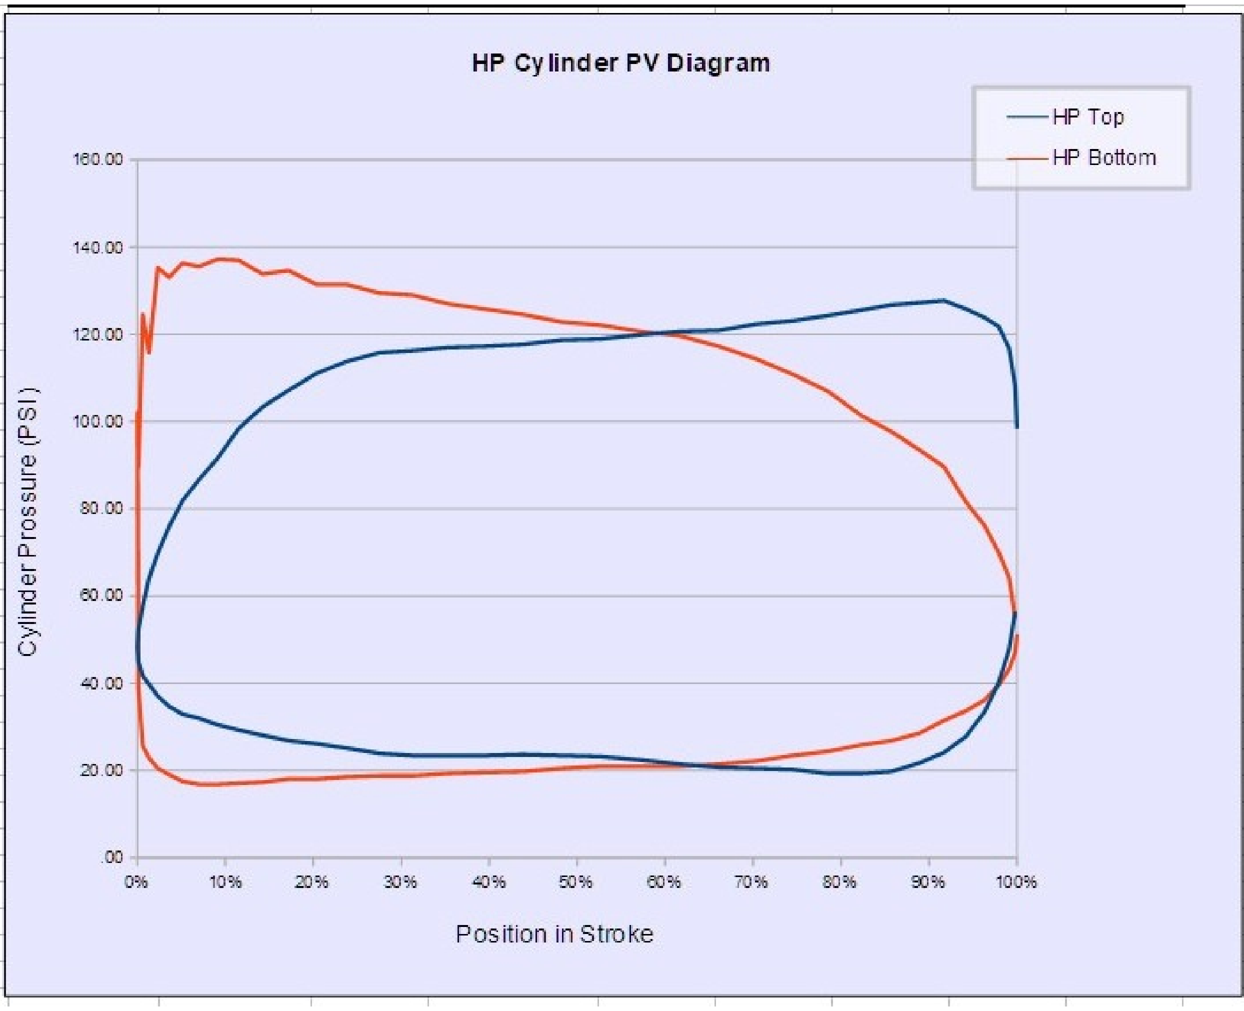

Commentary

This looks plausible, with peak cylinder pressures of around 140psi, and minimums of about 20psi, which match the LP data we are seeing. There is some drop off in pressure during the stroke, probably due to gas-flow issues in the ports, and throttling of the throttle valve!

The top and bottom of the cylinder appear similar. We can see in the top left of the diagram a limited example of "ringing" (which any electronics engineer would recognise) This, I think, is caused by oscillations/resonance of the steam and water in the tubes to the sensors. As we will see in a minute, this is much worse on the LP, at slower speeds, and on the lower sensors, because the tubes to the LP sensors are longer (the computer sits on the hot-well!)

There appears to be some apparent compression at the ends of the stroke, but this might be more gas-flow issues.

But we are getting apparently meaningful data at over 200rpm.

...although this is *<u>all</u>* dangerous ground with the designer interpreting the data – no peer review here!

Notched-Up – Ahead – HP Inlet 150psi – 182rpm

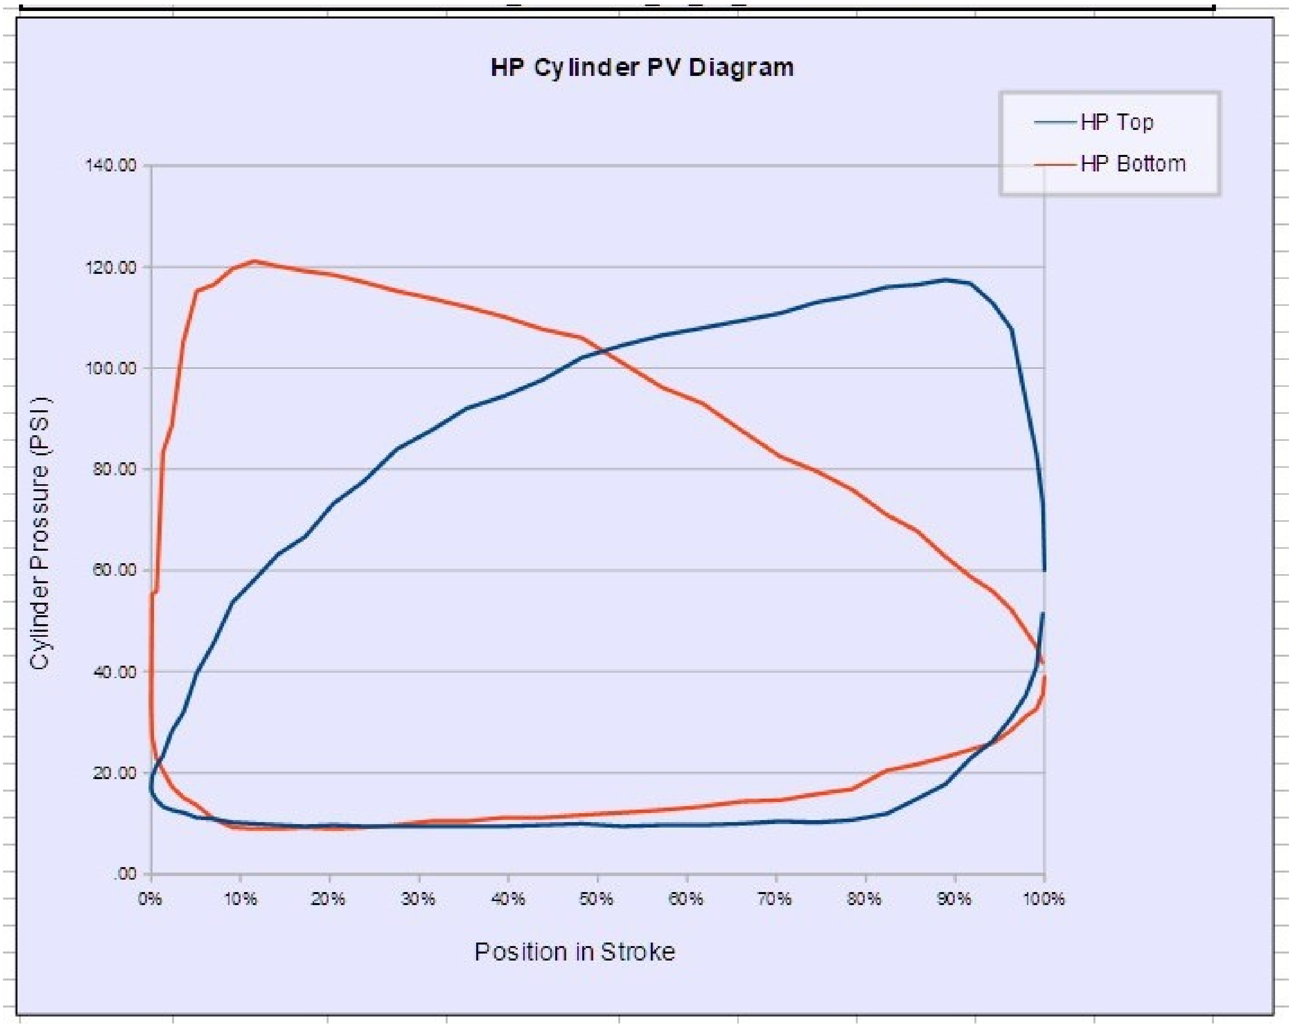

Commentary

Again this looks plausible, with cylinder pressures reducing much more during the stroke, suggesting we are getting a degree of expansive working of the steam.

I suspect that the more “triangular” nature of the enclosed areas indicates we are getting less power from the engine, and the fact that the RPM has dropped from 214 to 182 probably bears this out.

It looks like the minimum cylinder pressures are a bit lower, not really sure why this is...

It seems that there is a visible curve towards maximum cylinder pressures, and perhaps this is a sign of the (piston) valve ports closing.

In both cases (full and notched up) peak pressures in the top of the cylinder are a bit lower than those in the bottom. I guess this might be the effect of lower cylinder volume being smaller (because of the piston rod) or maybe it's a sign of poor valve events, or my imagination or instrumentation error.

In both diagrams the “compression” sections of the curves are very different, again suggesting some timing errors in the valve gear.... and remembering just how hard it was to balance these when building the engine I am surprised it loos as good as it does!

Full Gear – Astern – HP Inlet 50psi – 101.9rpm

Commentary

Here we start to see the problems we need to address on the engine, and some of the instrumentation errors.

Firstly, I had to phase shift the data by 10-degrees to allow for the fact we are running astern (magnet approaches Hall Effect sensor from the opposite direction). This adjustment can only be done by “inspection/guess”, one could move the magnet before switching from ahead to astern, but this is simpler, if a little more error-prone.

Secondly, the curves look a little more “wobbly”. I think this is because we are using a 200psi sensor to measure pressures in the range 7 to 45psi. Also the A2Ds are 10-bit (1 in 1024), but the sensors output from 0.5 to 4.5volts – so actually the sensors are only going to be able to resolve 0.03psi, so we are measuring down in the noise of the system, and we are using the battery voltage as a reference, where we really should be using some kind of regulator to give a more reliable reference.

On a mechanical front, we can see there is now a more pronounced difference between top and bottom data, and again I suspect this is as I “moved” the timing errors to Astern when I built the engine....as we generally try to steam forward, perhaps this is OK.

The LP Engine

If we now turn our attention to the LP end of the engine things look MUCH less pretty, and point to work we need to do this winter.

All of the error factors we looked at above are still present (oscillations in tubes, lower pressures being measured etc.) but now we add another one – vacuum!

Vacuum

I can't find any differential pressure sensors that are affordable, and as capable as the ones we are using. So the ones we are using say they will measure from 0psi to their maximum pressure with output of 0.5v to 4.5v, and that this should be a linear range.

However, when we actually connect them to a cylinder where the pressure dips below 0psi, we see the output voltage drop below 0.5volts too. So clearly the sensors *do* measure some amount of vacuum, but we must view the results with some scepticism.

Obviously we should attempt to calibrate the sensors, and if anyone knows to generate a given and stable level of vacuum to allow this calibration please let me know! (At the moment I think it's all down to long lengths of water-filled hose thrown over a cliff!)

So, the data below must be viewed with caution, although we can make some assessment of the probable levels of error.

LP cylinder – Full Gear – Ahead – 150psi – 214rpm

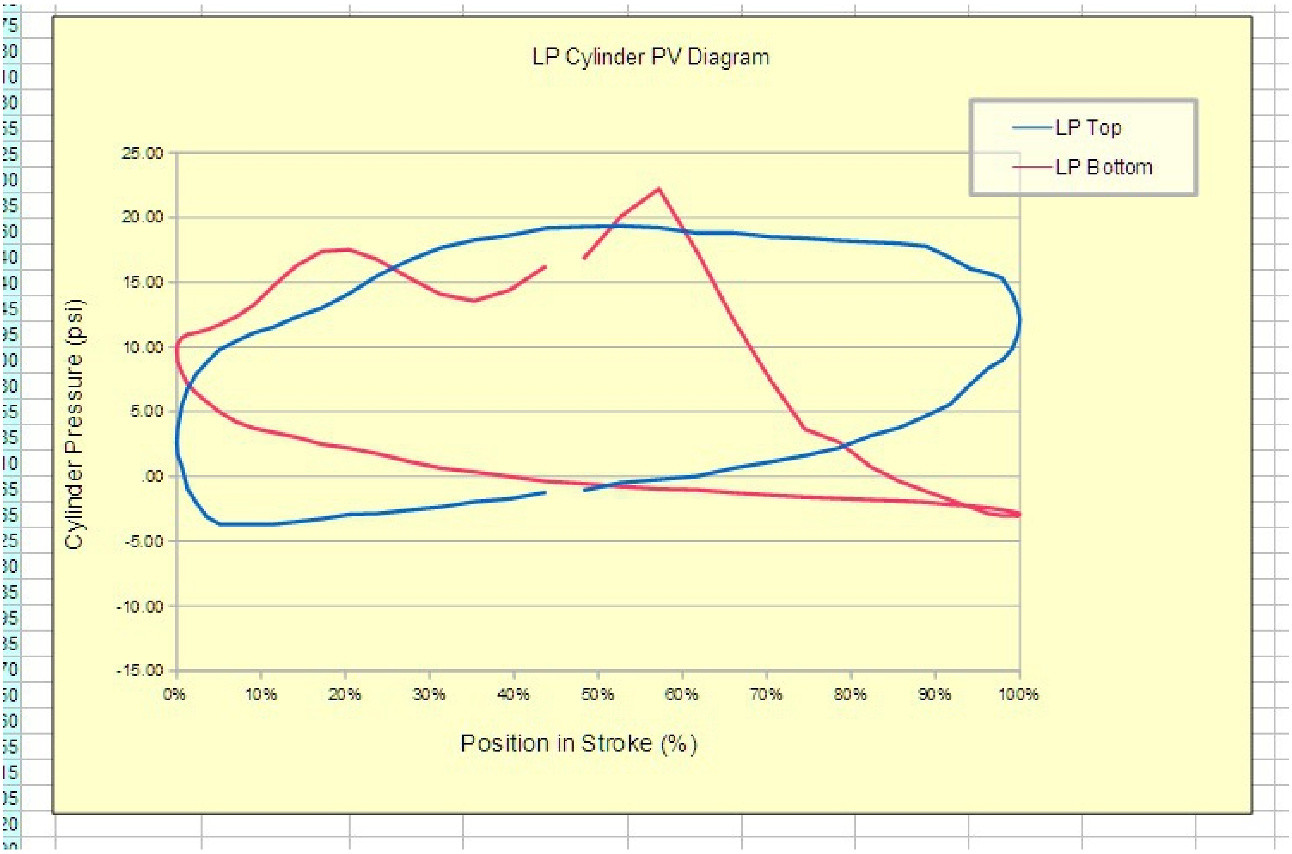

Commentary

What can one say! The max pressure matches the HP minimums, so that's OK. The minimum pressures are about -5psi, and the vacuum gauge was reading 10-12” of Hg (-4.9 to -5.8psi) so that looks believable. The top end data is OK'ish, but the bottom looks like a disaster to me! (thoughts from the readers?)

LP cylinder – Full Gear – Astern – 50psi – 88rpm

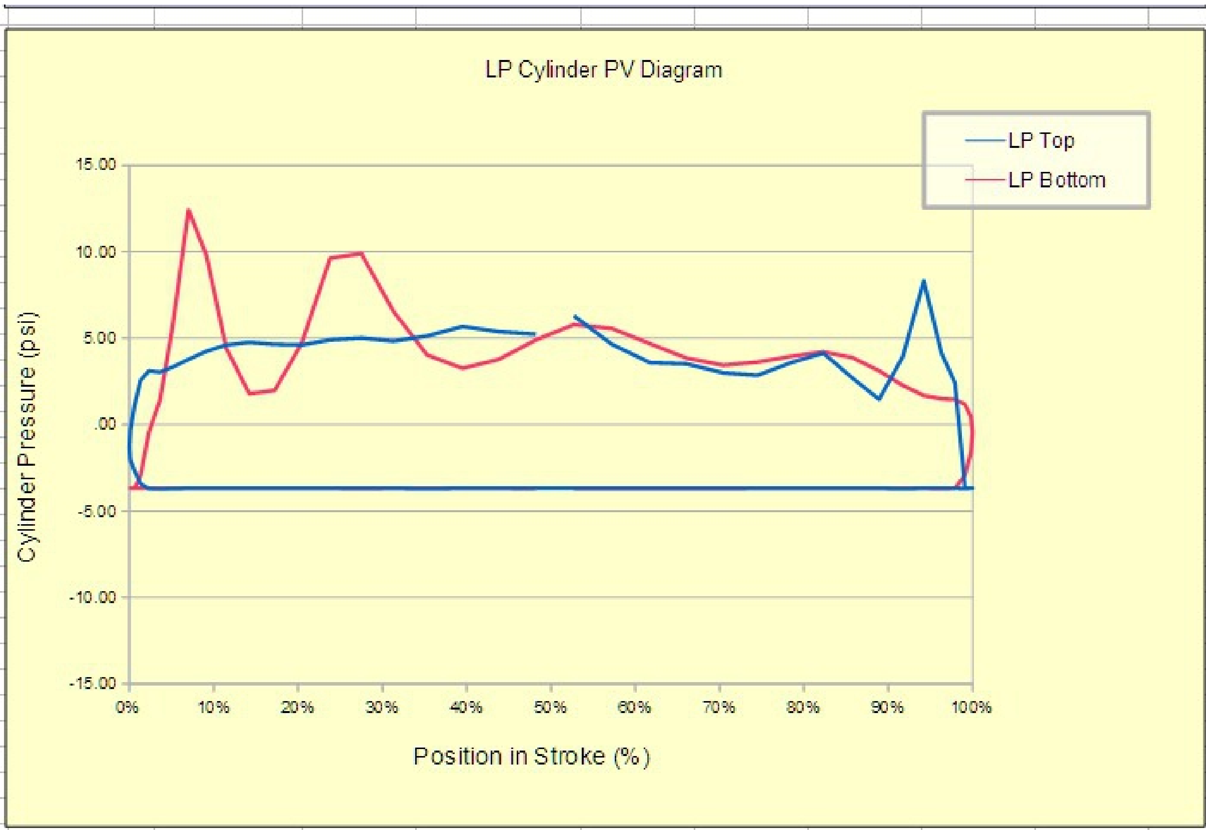

Commentary

Well, this looks like a time to stop.

The most obvious thing here is the apparent flat-bottom to the curve, this looks like the negative pressure limit of the sensors... I guess it could be that the condenser “clamps” the cylinder pressure when the valves open, but judging from the oscillations of the vacuum gauge that seems improbable. Maybe the slow engine speed is aggravating things.

I think the way forward is to indicate the engine with the condenser disconnected (well opening a large hole in the piping) to kill the vacuum and see what the sensor says then. That should remove the doubt/speculation, but prevent meaningful power computations.

Also we can see the “ringing” which seems much worse at slower engine speeds, and here (at 88rpm) it's the worst we've seen. You can see slugs of water bouncing around in the tubes to the sensors.... I am wondering if the problem can be reduced by pre-filling the tubes with water... but I think that's research for 2022....

Onwards!