READING THE WAVES

Steve Sabin and Randall Chitwood, SETPOINT Vibration, USA, show how modern commercial process historians can also serve as full-featured vibration monitoring software.

When the first operational data historians (i.e. process historians) emerged nearly four decades ago to begin replacing pen and paper chart recorders, they focused on relatively slow-changing process variable data from batch and continuous processes. Data from such processes did not typically need to be characterised in the historian with resolution better than every few minutes or seconds. Indeed, while the underlying control loops may have needed to function at sub-second speeds, summary data at 5 or 10 sec. intervals was often considered more than adequate and the concept of breaking the 'one second data' barrier became the process industry's equivalent of the 4 min. mile. Given the requirement for tens, or even hundreds, of thousands of process points at collection speeds measured in seconds or minutes, it is not surprising that process historians focused and evolved along a path of handling a large number of points at intervals measured in seconds rather than milliseconds, sampling speeds for individual points.

At the same time, online software for historising and analysing machinery vibration was evolving along a parallel, but separate, path. Unlike process historian software, its focus was on a much smaller number of total measurement points (or 'tags' in process historian parlance) – often fewer than 500 – but at much faster sampling rates. To put this in context, adequately capturing a waveform from a typical vibration sensor requires sampling rates of 40 kHz, or, as Nyquist dictates, 20 000 times faster than 1 sec. of process data without aliasing. An uncompressed, high-fidelity audio signal requires approximately the same sample rates as a typical vibration signal and is, therefore, a useful proxy when thinking of vibration data requirements. A conventional compact disc (600 MB) can hold approximately 60 min. of two-channel (stereo) audio data. Likewise, it can hold approximately 60 min. of uncompressed vibration data from two sensors.

Clearly, the sampling speeds of vibration data and its associated storage requirements meant that systems, which were originally designed for historising process data, were not able – until recently – to address high speed vibration data applications.

Therefore, within this historical context, two separate systems and infrastructures evolved: one for the specialised data generated by the vibration monitoring subsystems; and one for the process variable type data generated by virtually every other subsystem in the plant, whether process control, turbine control, motor control, or any other type of control or monitoring. This is depicted in Figure 1.

Industry tolerance

While the need for two infrastructures has admittedly been a necessity in the past, it is probably most accurate to say that users have 'tolerated' this situation rather than 'preferred' it. This is particularly true for IT departments tasked with managing the computing, software, network, and security infrastructures for a plurality of systems. Furthermore, it is also true for the machinery engineers in a typical plant, who must convince their IT departments that they need a separate infrastructure that often consists of different computers, operating systems, client applications, firewalls, remote access environments, security models, etc. Both stakeholders are inconvenienced in multiple ways, and both would like a solution that is less expensive to deploy and sustain by sharing, rather than duplicating, infrastructure.

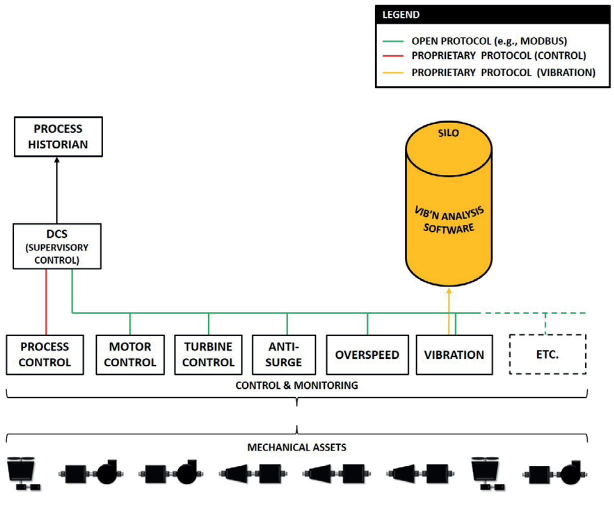

Figure 1. Most data in the plant flows from purpose-built controllers/monitors into the Distributed Control System (DCS), and then into the process historian. However, the special needs of vibration analysis software have historically required a stand-alone 'silo' that uses its own infrastructure, separate from the process historian.

Convergence

There has long been an appetite to converge both process and vibration data because the interaction between the machine and the process surrounding it is inevitable. Process conditions can adversely impact a machine, and failures are often due to external influences rather than normal wear and tear. For instance, in the case of pumps, it might be process conditions that lead to cavitation; in the case of compressors, it might be process conditions that lead to a surge; and, in the case of a turbine generator, it might be changing steam conditions that result in excessive differential expansion and a mechanical rub. This means that the ability to correlate process and vibration data is frequently necessary when diagnosing problems in rotating machinery – particularly critical machinery that is an integral part of the process flow. Various cumbersome, and usually expensive, methods of integrating machinery and process data have existed for years, but often meant that the data had to be replicated in both systems. This lead to the inevitable 'two versions of the truth' that never agreed precisely with one another. Until recently, the technology had simply not advanced sufficiently to enable convergence of these types of disparate data into a single repository. As a result, separate repositories could share data with one another, but were still just that – separate repositories.

Meanwhile, technology has progressed at a staggering rate along the familiar Moore's Law trajectory (Figure 2). Indeed, as process data historians have pushed the envelope of data collection speeds, they now easily surpass the 1 million tag/sec. mark and continue to climb even higher, with speed constraints primarily imposed by the ability of hard drives to write to their platters. The advent of affordable solid-state drives and raid arrays has circumvented even this constraint, and the speed limit is climbing ever higher.

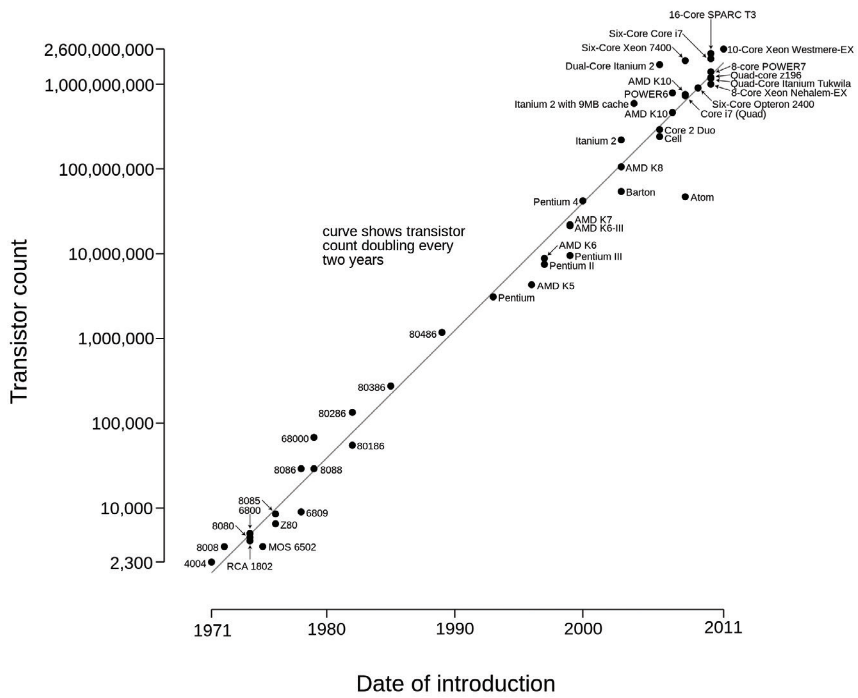

Figure 2. Microprocessor transistor counts 1971 – 2011 and Moore's Law.

100 000 x 10 = 1000 x 1000

Process historians evolved along a path where many thousands of points needed to be scanned, but at relatively slow rates. Update rates of 1 sec. for process data are often considered to be extremely fast. In a typical process plant, the Distributed Control System (DCS) can have tens of thousands of loops to control and monitor, and the process historian can likewise have hundreds of thousands of points. In contrast, vibration systems usually have fewer than 1000 total points across all of the critical monitoring machinery, but each point must be updated with millisecond scan rates.

Modern process historians are equally adept at addressing both of the following situations: a large number of points at update rates measured in seconds, or a smaller number of points at update rates measured in milliseconds or even microseconds. In other words, historians can address 100 000 points at 10 sec. scan rates as easily as 1000 points at 1 msec. scan rates. In either instance, the rate at which the historian must acquire and store data equates to 1 million value per sec.

A fork in the road

The reason that a process historian, such as the PI System, can be used for vibration data is two-fold. Firstly, a vibration waveform is nothing more than an exceedingly fast trend. Instead of data samples separated in time by seconds, data samples are separated in time by milliseconds. Both are time series data, and one is simply faster than the other. Process historians are very good at handling time series data. Secondly, the PI System had evolved not just in terms of speed, but in terms of how it organised measurement points. Measurement points in the vibration world are organised into ascending hierarchies starting with points, bearings, machine cases, machine trains, process units, plants, etc. Thus, a measurement point might be described as 'X probe, inboard bearing, compressor K-101, refrigeration train, plant 2'. In the case of the PI System, it had recently been enhanced to allow measurement points (or 'tags') to be organised in terms of such hierarchies. Therefore, it could handle both the data structures required and the data speeds required.

It is precisely these capabilities that made the PI System viable for high speed vibration data, allowing the process historian and vibration historian to converge into a single infrastructure, without the need of proprietary vibration servers and networks.

Deadbands for waveforms

One of the inherent challenges in the collection of any type of data, and particularly with vibration waveform data,

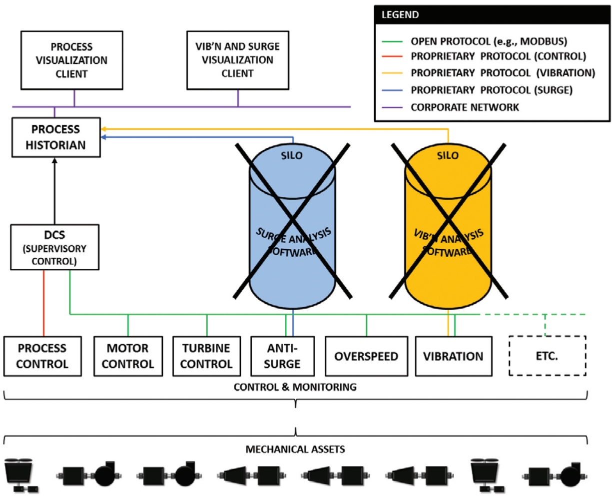

Figure 3. By using the process historian, the need for separate silos can be eliminated. This approach is available not just for vibration data, but also for compressor surge/thermodynamic performance data as well.

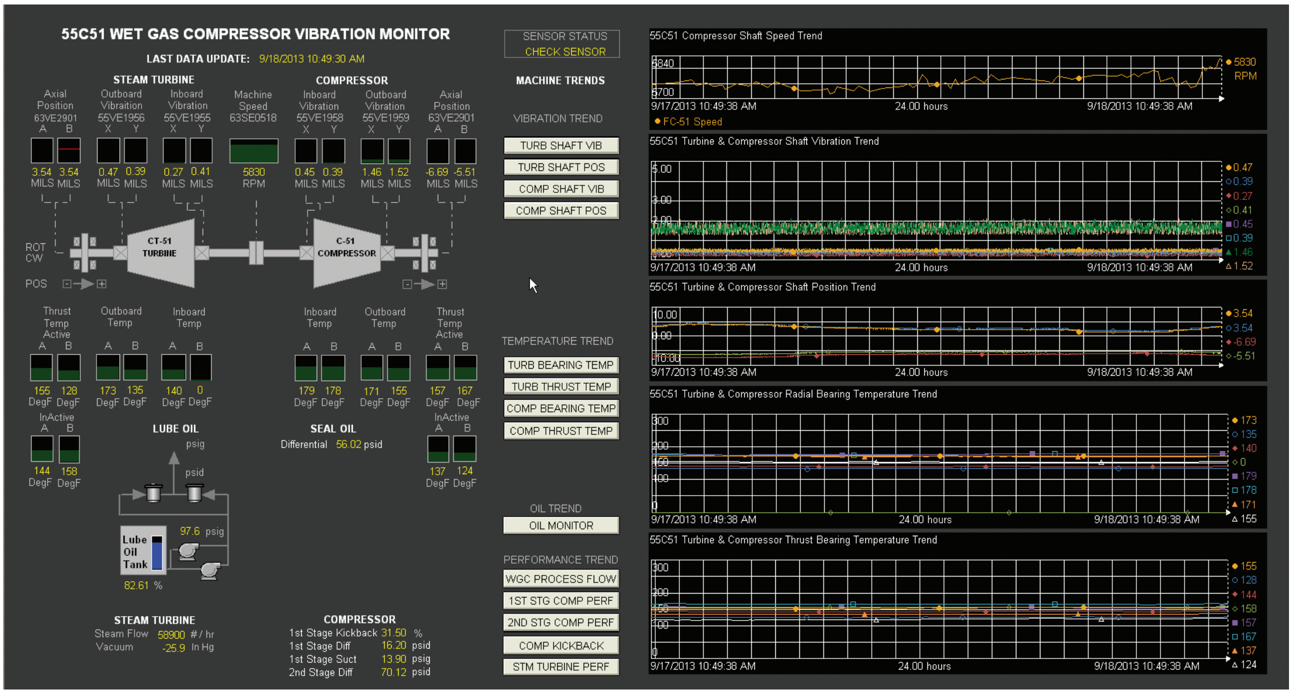

Figure 4. Typical screen for a turbine driven compressor train built by a customer using PI ProcessBook. When trend and status information beyond the capabilities of ProcessBook are required, the user can use the buttons on this screen to launch the specialised visualisation tools, such as the one shown in Figures 5 and 6.

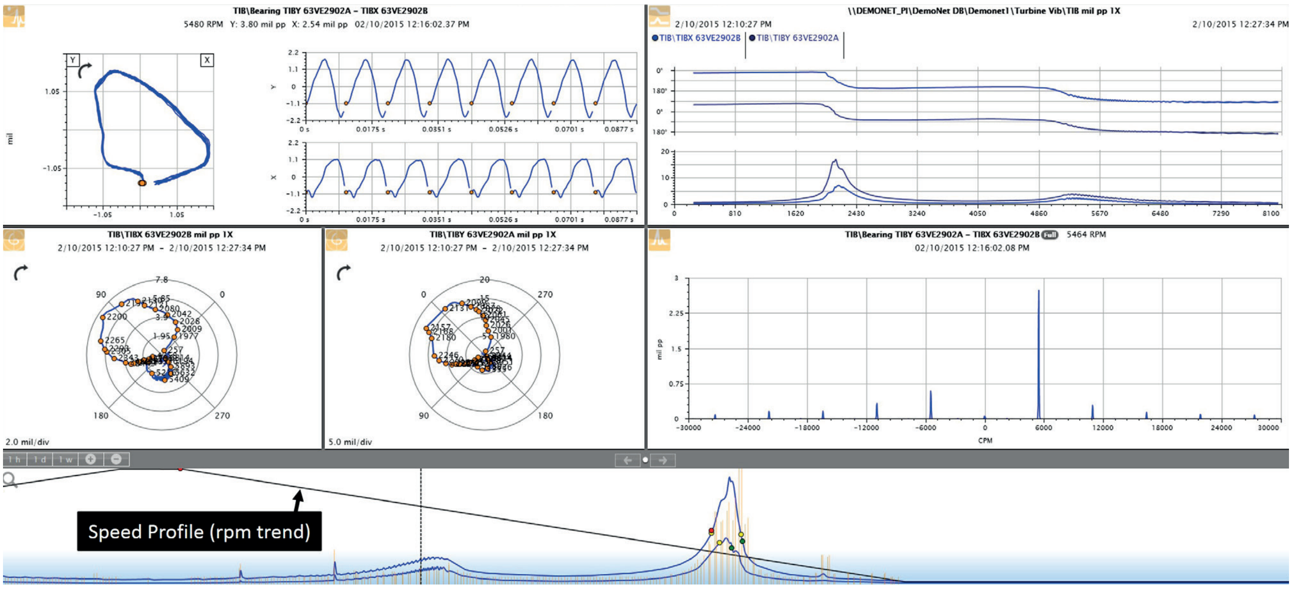

Figure 5. Waveform data collected during a machine trip.

is that it is undesirable to store everything. This is partly because of the storage space required to save everything and network bandwidths needed to move everything, but also because the vast majority of machinery data is uninteresting. One shaft orbit is very similar to the next, except when problems occur. Therefore, not all vibration data needs to be saved, but only vibration data that has changed from the last stored waveform.

This basic idea has existed since the inception of process historians. It is embodied in the concept of a deadband, whereby a trend line that is not changing does not need dozens or hundreds of points stored. It only needs two points stored, and a straight line drawn between them. The deadband simply defines how much change must occur between one point and the next to recognise it as a new data point, rather than just a linear extrapolation of the previous data point.

Initially, the primary imperative for deadbands around process historian points was the limited storage space of most computers. In the 1980s, when historians were beginning to debut, a large hard drive was 20 – 40 MB. Clearly, the ability to store data efficiently was paramount, and the sophistication of deadbands and other compression algorithms were likewise paramount. In a day when a 1 TB drive can be purchased for less than US$100, it is tempting to think that compression is no longer that important. There is some merit to this, but anyone who has ever had to move that 1 TB across a network will understand that although there is room to store all of the data, it can take a long time to move it.

Deadbands have historically been used with scalar data that can be characterised by amplitude and time tags. However, the concept of only saving data when something changes can apply to any type of data. It is now routinely being used by the SETPOINT system to decide which vibration waveforms to keep and which ones to ignore.

Consider a shaft turning at 6000 RPM (100 revolutions/sec.) on a typical compressor train in an LNG plant. Additionally, assume that the vibration is sampled 128 times per revolution for 16 shaft revolutions (waveform consists of 2048 samples with 78 microseconds between each sample). When the system first turns on (and assuming the machine is already running at steady-state speed), it stores the waveform corresponding to those first 16 shaft revolutions into the historian as a time series, just as with any other time series data, but with microseconds between the values rather than seconds, minutes or hours. It then examines (but does not necessarily store) the next 16 shaft revolutions (revolutions 17 – 32). It examines multiple attributes of the waveform, such as overall amplitude, gap voltage, frequency content, period (machine speed), etc., and compares these to similar attributes of the initial 16 revolutions already stored. If things have adequately changed, it also stores the new waveform. If things have not adequately changed, it only keeps the first waveform and discards the other. The system continues on in this way indefinitely, collecting and examining every waveform, but only saving those that represent sufficient change from the baseline. This, in essence, is how the deadbanding of waveforms, rather than just scalar data, is accomplished.

By utilising the concept of a deadband for not only conventional trend data, but also for high speed waveform data, the system is able to continuously collect and analyse every waveform from every rotation of the machine's shaft. However, it only stores this data when a waveform has sufficiently changed relative to its constantly updating baseline. The system can, therefore, go for long intervals without saving extraneous waveforms if conditions are not changing, yet store waveforms quickly when conditions are rapidly changing, such as during machine start-ups, process upsets, or incipient mechanical failures. Therefore, two years of data from a typical 300-channel system only requires a modest 1 TB of storage space, employing a combination of compression algorithms native to the process historian and new algorithms for high speed waveform data.

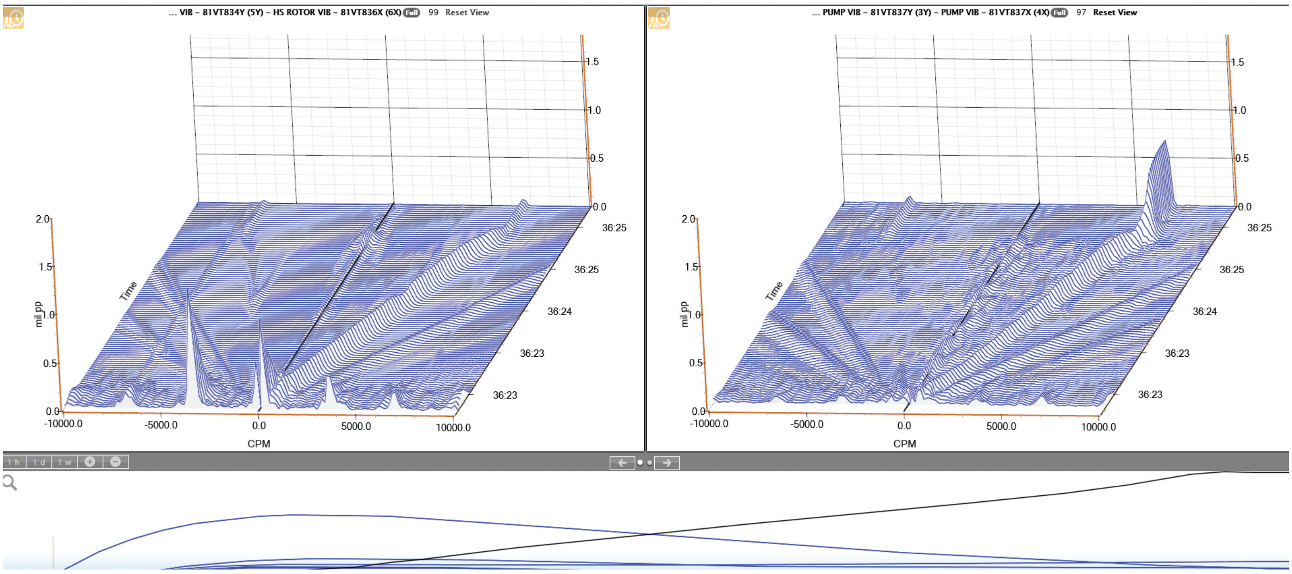

Figure 6. Electric motor driven start-up lasting only 3 sec.

Boost mode

Some machines can exhibit extremely fast starts and stops, such as motor driven pumps and smaller compressor trains and blowers driven by motors or steam turbines. On such machines, the transient conditions can be brief and the ability to store every waveform, as with an analogue tape recorder, can be important. Normally, the system will save no more than one waveform from each vibration sensor every 2 – 3 sec., and it simply chooses the most interesting waveform during that interval. If nothing is interesting, then nothing is saved. The system, therefore, can go for many minutes or hours without the need to save anything. However, during fast starts or stops, the system can be placed into a special mode in which every single waveform is saved for the duration of the start-up or shutdown.

Operating experience

Experience gained over the last two years has demonstrated that a commercial process historian, namely the OSIsoft PI System, is capable of handling vibration data without constraining the types of plots, data resolution, or features that are required by rotating machinery engineers and vibration analysts. By using the process historian as the single 'system of record' for not just process data, but also for vibration data and compressor surge/thermodynamic performance data, the ease of correlating process, vibration, and other relevant machinery information is greatly facilitated and IT infrastructure is eliminated (Figure 3). Indeed, the identical considerations discussed in this article for vibration data and its special needs can (and are) being extrapolated to compressor surge/thermodynamic performance data, so that it too can use the process historian rather than necessitating a separate 'silo' just for its own data.

Vibration data is stored in the process historian as time series and event records, just as any other process data, and can be visualised using standard process historian tools, such as PI ProcessBook (Figure 4). When users need to access more specialised plot types that are not native to the process historian, a separate visualisation tool is used. Its purpose is simply to display the vibration data resident in the historian that cannot be presented using standard trends, bar graphs, and other tools native to the historian. Examples of the vibration visualisation tool are shown in Figures 5 and 6.

The system described here is in use at more than two dozen sites globally, many of which did not require the purchase of any additional software as they were already users of the OSIsoft PI System and had the requisite number of tags and server capacity to add the desired vibration data.

Conclusion

The industries in which critical machinery is used, and in which continuous vibration monitoring is practised, have conditioned themselves to accept the necessity of separate databases and tools for each kind of data, be it process data, vibration data, thermodynamic performance data, or compressor surge data. The historical reasons for this have largely been technical in nature, due to the limitations of process historians. However, advances in the speed at which the historian can write data to its database and in the data structures it uses have removed the impediments to the inclusion of vibration data, including the extremely high speed waveform data necessary for machinery engineers and vibration analysts to perform their jobs. The benefits of using a single system extend not only to these individuals, but also to corporate IT departments that want to standardise and rationalise the number of software platforms used across the enterprise, while simultaneously providing high levels of reliability and security. As the process historian is almost always considered a 'mission critical' application with a broad base of users and IT support resources, the inclusion of other types of data, which has previously been impractical or impossible – such as vibration and surge data – is now a viable alternative. **LNG**