Nano-Resolution

Oceanic, Atmospheric, and Seismic Sensors With Parts-Per-Billion Resolution

Dr. Theo P. Schaad
Paroscientific, Inc.
Redmond, WA

10 March 2009

Summary

New practical techniques have been developed to measure the frequency outputs of inherently digital quartz resonator sensors to parts-per-billion, or, to nano-precision. Applications include measuring nano-bar barometric pressure fluctuations for infrasound detection, measuring acceleration or gravity to nano-g's, and measuring water level fluctuations to microns (µm) in deep-sea depth sensors. This paper discusses new developments in nano-counting techniques with regression (FIR) filters and digital multi-stage IIR low-pass filters.

Background

Inherently digital quartz resonator sensors (designed by Paroscientific, Quartz Sensors Inc., and Quartz Seismic Sensors) utilize resonators that change frequency under load. The measurand-induced load is generated by pressure, acceleration, gravity, weight, etc., depending on the sensor type. The changes in frequency under full-scale measurand excursion are typically scaled to 10 % frequency variations from a nominal base frequency of 35 kHz. The resolution as a fraction of measurand full scale depends on the precision of the frequency counter and it is therefore desirable to develop practical frequency counters with higher resolution.

Three frequency counter methods have been developed:

- Start-stop counting, also known as reciprocal counting; with a refinement called interpolating counting
- Regression (FIR) counting
- Nano-counting with multi-stage digital IIR low-pass filters

Start-Stop Counting

In our traditional, start-stop counting method, the measurand frequency output signal (at nominal 35 kHz) gates a high frequency clock between a start transition and a stop transition N periods later. The measurand period is calculated by dividing the time interval by N. The period or frequency resolution of the start-stop counter consists mainly of a digital count error plus a trigger error dependent on the frequency jitter of the signal source. At each transition, the clock is measured as the nearest integer. The error is therefore between 0 and 0.5 counts at both start and stop transition. The total error is less than 1 count. Statistically, one can use the concept of an rms clock error, which is 0.29 counts per rounding, or 0.41 counts for two independent transitions. For instance, if the counter high frequency clock is at 14.7 MHz, one count period is 68 ns (nano-seconds) and the rms time uncertainty of the measurement is (0.41)*(68 ns) = 28 ns. The fractional period or frequency resolution is simply the time uncertainty divided by the time interval. For instance, at 1 Hz sampling, it is 2.8E-8, or -151 dB. (The resolution of the measurand is ten times less, or -131dB in this example.)
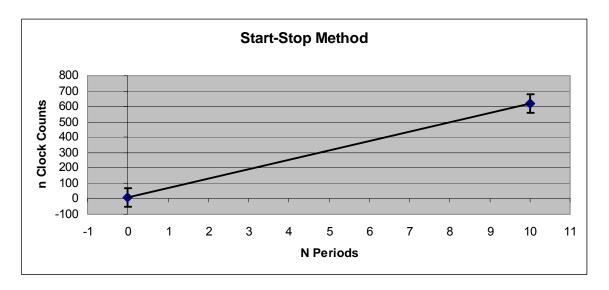

Figure 1: The principle of the start-stop method is to measure time by counting clock pulses between a "start" and a "stop" transition of the signal period. The measurand signal period is then simply the total time divided by N. The method can also be viewed as finding the slope between the end-points. Because of clock rounding and trigger errors, the end-points are not precisely known, leading to a finite resolution error in the frequency measurement. (The numbers in the plot are only for illustration; the actual count errors are a fraction of a clock count.)

Regression Counting

Regression counting is based on a linear regression (least-squares) fit. The simple premise is that the start-stop method involves finding the slope between clock counts and number of signal periods, wherein the error of the slope depends on the measurement errors of the end-points ("start" and "stop"). It is well known that a slope estimate improves if many

measurements are taken between the end-points. The error of the slope, and therefore the frequency resolution, can be estimated by statistical error analysis if we treat the individual sub-samples as independent and Gaussian.
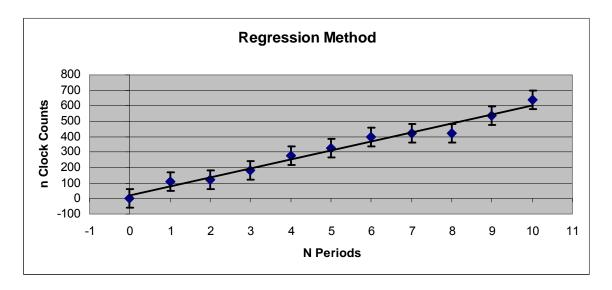

Figure 2: The regression method is based on measuring clock counts many times at a sub-sample frequency between end-points. The slope is estimated by regression analysis and is a better measure of the signal period compared to the start-stop method in Figure 1. (The numbers and errors in the plot are for illustration of the principle only.)

Statistically, the fractional error in frequency (or period τ) is

$$d\tau / \tau = (\sqrt{12})\, \sigma /(N^{3/2}\, h\, \tau)$$

Wherein N is the number of sub-samples measured, h is the number of signal periods per sub-sample (referred to as the signal "divider"), and σ is the combined error per sub-sample (rounding error and trigger error). In comparison with the "old" start-stop scheme, the improvement is the ratio ("old" error divided by "new" error), which is

$$\text{Improvement} = \sqrt{(N/6)}$$

At 1 Hz sampling, with a divider of h=4, nominal signal period τ=28.7 microseconds, and clock frequency fc=59 MHz, we measure about N=8710 transitions per sample and potentially improve the measurement by -32 dB (about 40 times) over the start-stop counting. The rms rounding error σ is 0.29/fc, or 5 ns (nano-seconds). Hence, the best possible frequency resolution (at 1 Hz) is (dτ / τ) = 1.86E-10 = -195 dB.

In addition, the electronic frequency jitter (trigger error) is typically 15 ns, which lowers the achievable resolution by 10 dB. With regression counting, we can therefore achieve a measurand resolution of a few parts per billion (in the time-domain) at a one Hz data rate.

Note that the resolution improves as the 3/2 power of the time window. Averages of independent samples added in quadrature only improve by one-half power, whereas the reciprocal counting method improves linearly with time.

The regression technique in this application can be viewed as digital signal processing on an over-sampled time series. What is unusual is that a digital filter is being applied to an over-sampled frequency count, whereas the more common case is filtering a digitized analog signal. The regression method belongs to a class of digital filters called FIR filters, or Finite-Impulse-Response filters. FIR filters are considered to be stable in the sense that a miscount will not be propagated into the next sample. It can be applied to a finite sub-sample time-series, for instance within the time interval of the data sampling period (we define the data sampling rate as the rate of reporting filtered data). Each sample is independent, and the sensor could even be turned off between samples (for instance, to save power).

The Digiquartz sensors accurately measure geophysical phenomena over time scales from a fraction of a second to spans of decades. An example is given in Figures 3 and 4 in which lunar gravitational effects are measured over a period of a week and an earthquake is measured with a sampling time of seconds.
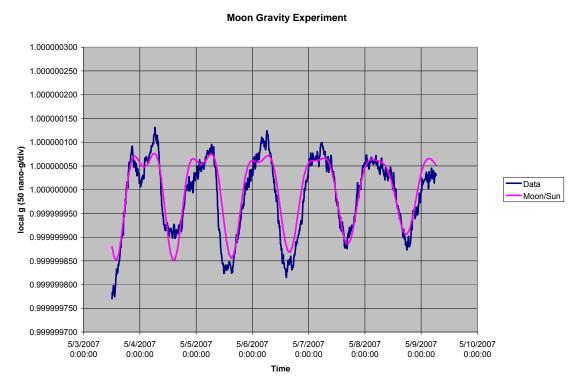

Figure 3: Local vertical gravity changes measured with an absolute 3-g accelerometer using the regression FIR counting method and compared to predictions from the positions of the sun and the moon. The residuals are about 40 nano-g's over 6 days of data.

4

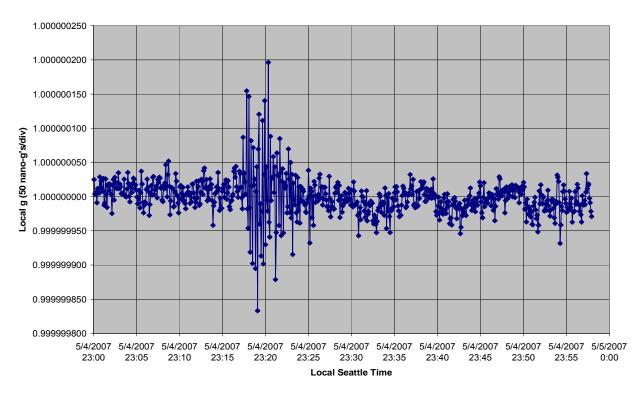**Earthquake 2007-05-04 M=4.8 NW of Vancouver Island, BC**

Figure 4: First earthquake measured with a 3-g seismic accelerometer using the regression FIR counting method. Ground acceleration from a distant earthquake was measured with a resolution of a few nano-g's.

Nano-Counting with IIR Filter

Recently Spahr Webb showed us another nano-counting method based on IIR low-pass filtering that is an improvement over the regression method. The IIR Filter reduces the high-frequency noise (typ. 10 to 100 Hz or higher) that could lead to aliased signals in the long-period region of interest. Each low-pass filter reduces the high-frequency noise above the filter cutoff by -20 dB/dec. A 5-stage IIR low-pass filter is effective at -100 dB/dec above the filter cutoff frequency. Infinite-Impulse-Response (IIR) filters are fundamentally different from Finite-Impulse-Response (FIR) filters, as they run continuously in the background.

As with regression counting, we chose to over-sample at every fourth transition of the signal (signal divider h = 4). At a nominal signal period of 28.7 microseconds, the sub-sample frequency is near 9 kHz it is essential to choose the fastest clock available, e.g. 59 MHz (command VP=4).

The sensor signal at the sub-sample frequency f_s is measured with the "traditional" start-stop method. The raw measurement of this sub-sample is the z_i counter. A high-speed clock of 59 MHz produces about 7000 integer clock counts (per sub-sample). The IIR algorithm then works on the sub-samples to produce a filtered output at a slower data output sampling rate.

Each IIR filter is of the form

$$y_n = \alpha\, x_n + (1 - \alpha)\, y_{n-1}$$

Each stage is a single-pole low-pass filter with a corner frequency f_c of

$$f_c = (\alpha\, f_s) / (2\pi)$$

The filter constant α is programmable by the IA=L command (L is an integer from 1 to 16), such that $\alpha = 2^{-L}$. The filter frequency cutoff, defined as the knee at -3 dB and controlled by α, essentially replaces the integration time of the previous counter method. The sampling period can be adjusted separately by the PI command, but if the sampling (data reporting) rate is higher than the cutoff, consecutive samples are filtered.

The frequency resolution of the IIR scheme can be predicted by numeric simulation for different values of α. With a filter cutoff near 1 Hz and a sub-sample frequency of 9 kHz, the improvement over the start-stop method is near -30 dB, similar to the regression technique. With a four times higher sub-sample frequency (measuring each transition of the signal period without divider), the predicted improvements are near -40 dB.

Empirical resolution measurements were made of the IIR filter sensor/interface in the time domain (Allan deviation) with an isolated crystal that is used in Digiquartz transducers, standard production oscillator electronics, and the new IIR interface, and firmware Rev. R5.0 (or later). In the IIR nano-counting scheme, the resolution in the time domain depends on the corner frequency of the filter. An effective time interval can be defined as the inverse of the sampling rate after setting the sampling rate to twice the filter cutoff frequency (such that the cutoff frequency is at the Nyquist limit).

The results are shown in Figure 5. At time intervals of 1 second or longer, the barometric resolution is better than 0.0001 Pa, or 0.1 mPa (one part per billion of full scale).
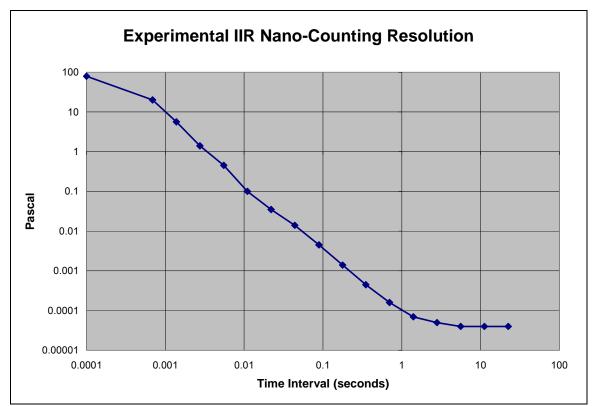
Figure 5: Experimental instrument/interface resolution scaled to a Digiquartz 216B barometric standard in the time domain (Allan deviation vs. effective sample time)

Infrasonic Measurements with Nano-Barometers

Two Paroscientific High-Precision Barometers were connected to the same pressure line leading to an external wind-insensitive pressure port. The barometers were kept indoors at normal room temperature to mitigate thermal effects on the sensors from extreme temperature changes and sunlight. The two sensor outputs were measured with a single interface counter board to ensure synchronicity of the measurement. The measurement mode was temperature-compensated pressure in units of psia. The setting was PI=1000 (1 Hz data reporting), IA=10 (alpha=0.001, frequency cutoff = 1.4 Hz), IS=5 (5 IIR filter stages), and VP=4 (quadrupled counter clock at 59 MHz). The purpose of the side-by-side installation of two independent barometers was to measure the relative incoherent noise floor between the two sensors.

On a normal calm and sunny summer day in the Pacific Northwest, there is a stationary high-pressure system over the ocean. The time-series of July 23, 2008, is shown in Figure 6.
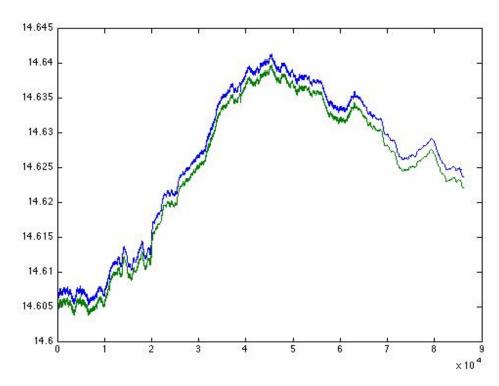
Figure 6: Barometric data recorded in psia on July 23, 2008 in Seattle, WA. The time axis is in seconds. The traces of two independent barometers are slightly offset.
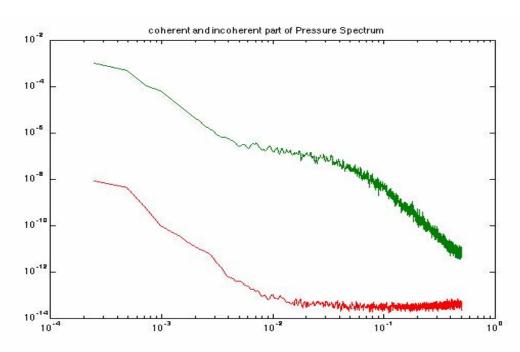

Figure 7: Spectral density plot in psi^2/Hz. The x-axis is frequency in Hz. The green curve is the coherent barometric data. The red curve is the incoherent noise floor of the instrument, flat to about 100 seconds. There is no evidence of micro-baroms in this data set. (Plot courtesy of Spahr Webb)

In pressure units of Pascal, the noise floor is at 7.2E-7 Pa^2/Hz. Integrating the spectral noise floor to the Nyquist limit gives a time-domain resolution of 0.0006 Pa or 6 nano-bars. This noise floor has been superimposed on the Infrasound Ambient Spectrum shown in the following figure.
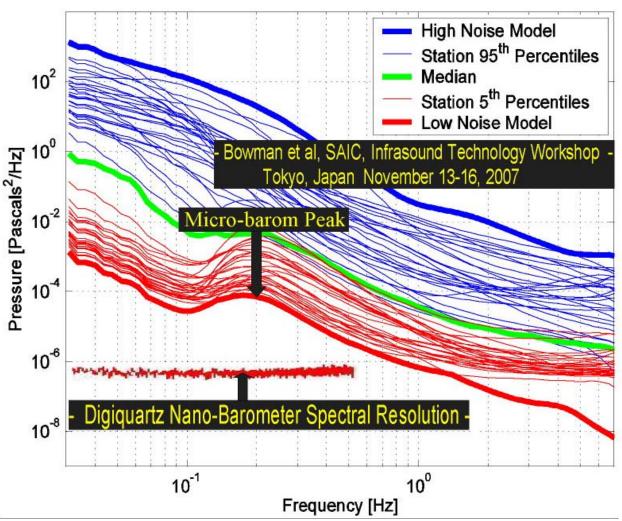

Figure 8: Digiquartz Nano-Barometer spectral resolution compared with ambient spectrum

Micro-Baroms

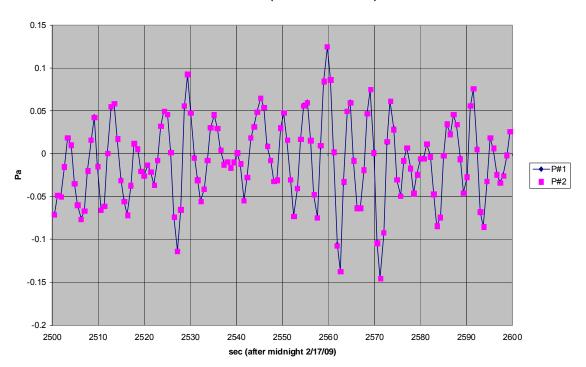**Micro-Baroms (filtered 0.05 to 0.7 Hz)**

Figure 9: Micro-baroms measured with two independent Nano-Barometers

Micro-baroms are natural infrasound signals that originate at the low-pressure center of storms over the Pacific Ocean (among other sources). In Seattle, WA these signals are a constant background in the winter months. The frequency is near 0.2 Hz and the amplitude is typically 1 micro-bar (hence the name), or 0.1 Pa (Pascal).

Recently, we tested two Nano-Barometers in a similar way as shown previously in Figure 6. This time the micro-baroms were prevalent over the 2 days of testing. Figure 9 shows a random 100-second time interval in the wee-hours of 2/17/09. As can be seen, there were about 20 waves in the interval (0.2 Hz). The blue curve is from the data of nano-barometer/infrasound sensor #1. The colored dots are from the data of nano-barometer/infrasound sensor #2. The data was filtered at 0.05 Hz with a single-stage high-pass filter to remove the static offset between the two units. Visually the two curves are identical; by simple computation, the residuals are 0.4 mPa per unit (0.55 mPa difference). This number is consistent with self-noise of 0.2 mPa (at 1.4 Hz sampling) plus potential errors from trigger, synchronization, scale factors, and thermal errors.

Nano-Precision Counting with Depth Sensors

In general, the parts-per-billion results obtained with a nano-barometer are transferable to other pressure and depth sensor ranges, properly scaled to the full scale of the device. To illustrate the point, the resolution of a 7000-meter depth sensor (10 kpsia pressure sensor) was measured by comparing how well it tracked ambient barometric pressure compared to a nano-barometer.

The 7000-meter depth sensor tracked the nano-barometer to about 0.25 Pa (Allan deviation) as shown in Figures 10, 11, and 12. Compared to its full scale of 70 Mpa, the time-series resolution is 4 parts per billion at 1 Hz. The time-series clearly shows that the short-time resolution (at 1 Hz) is relatively "noisy", but the longer-term tracking is excellent. At even longer times (hours), the traces tend to separate slightly depending on the thermal long-term stability of the depth sensor and the counter clock, with optimal results expected in a stable ocean environment. A spectral analysis was performed on this data set.
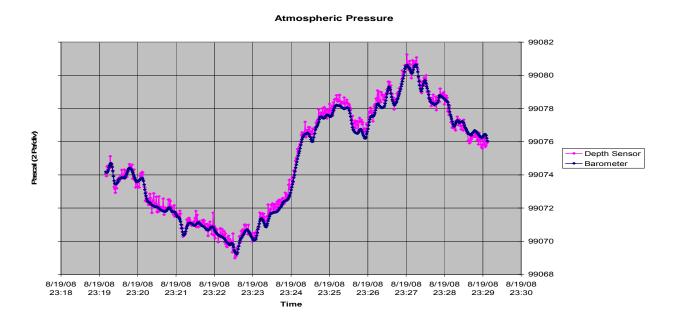

Figure 10: 10-minute time series of a 7000 meter depth sensor (red) and a nano-barometer (blue) in parallel. The barometer is 700 times more accurate and serves as a reference standard. The depth sensor resolution is 0.25 Pa, or 0.03 mm of water depth.
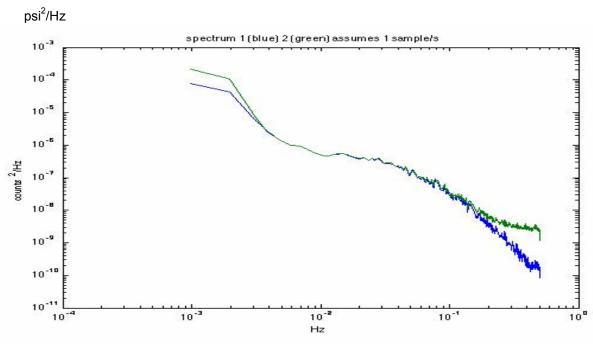
psi^2/Hz

Figure 11: Spectral analysis of a daylong time-series of barometric data. The green curve is a 7000-meter depth sensor and the blue curve is the reference pressure taken with a nano-barometer. Between about 0.003 and 0.1 Hz, the two signals are highly coherent and track each other. At a time scale of a few seconds, the depth sensor resolution is a few parts-per-billion. The vertical scale is in psi^2/Hz.
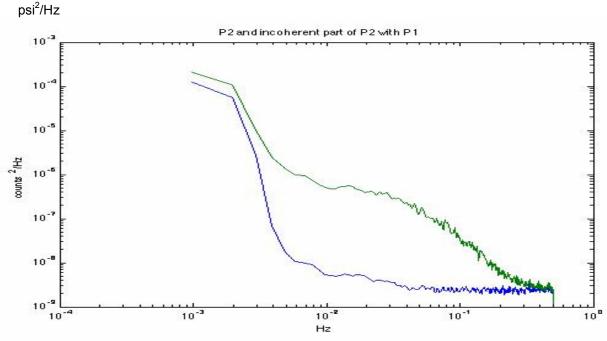
psi^2/Hz

Figure 12: Green curve as in Figure 11; blue curve is the instrument noise floor of the 7000 meter depth sensor. The flat portion is below 3E-9 psi^2/Hz or less than 0.14 Pa^2/Hz. The spectrum for a lower range depth sensor (e.g. 2000 meters) would have a noise floor of 0.01 Pa^2/Hz. (Plots courtesy of Spahr Webb)

Earthquakes have been measured with the IIR Filter algorithm using a 3-g triaxial accelerometer developed by Quartz Seismic Sensors.

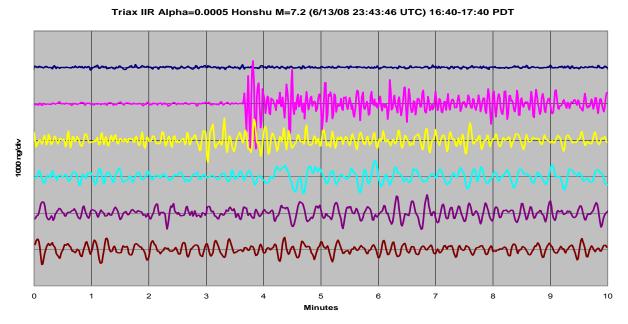**Triax IIR Alpha=0.0005 Honshu M=7.2 (6/13/08 23:43:46 UTC) 16:40-17:40 PDT**

Figure 13: IIR nano-counting applied to a 3-g seismic sensor. Honshu, Japan, earthquake of 13 June 2008 measured in Seattle, WA. The graph is shown in a seismograph format with 10 minutes of data per line.

How to Get Started with Nano-Counting

To get started with a nano-barometer (or any other nano-resolution sensor by Paroscientific, Quartz Sensors, or Quartz Seismic Sensors), the sensor must be attached to a Paroscientific Intelligent Board with the new filter algorithms. The experimental data described in this report was collected with the latest beta version of the firmware that runs the nano-counting IIR or FIR algorithm (Version VR=X5.0K or later). The production-version of the Operation Manual is at: http://www.paroscientific.com/manuals/8819-001.pdf.

A communication program (for instance DQI 2.0) can be used to check the settings of the user-selectable parameters. For single sensors, the sensor configuration should be at XX=1. The three counter algorithms are start-stop (XM=0), regression (XM=2), and IIR nano-counting (XM=1). Choose IIR nano-counting (*0100EZ*0100XM=1). Set the clock to quadruple speed (VP=4). Start with a filter setting of IA=11 (alpha=0.0005). Set the data rate to 10 Hz (PI=100). For outputs in bar, set the units to UN=3. For outputs in hPa (mbar), set the units to UN=2. For outputs in Pascal, set the units to UN=0 and the unit conversion to UF=6894.757 (Pa/psia).

Using software DQI 2.0, the data can be plotted with the monitoring program. For outdoors applications, we recommend placing the nano-barometer in a protected place like an earthquake vault and plumbing the pressure sensor inlet to a wind-insensitive high-

performance pressure port placed in an unobstructed location as close as possible to the sensor.

Conclusion

Two new methods have been described to measure frequency-output sensors at a resolution of parts-per-billion. The two methods are regression counting and IIR nano-counting. Generally, both methods provide up to 100 times higher resolution than our start-stop counting technique.

We would like to thank Spahr Webb at Lamont-Doherty Earth Observatory of Columbia University for many helpful suggestions, and Randy Quayle at Paroscientific, Inc., for his implementation of the new counting algorithms.