Proceedings of 20th International Congress on Acoustics, ICA 2010

23-27 August 2010, Sydney, Australia

Signals for the approximate characterization of acoustic systems

James S. Martin, Peter H. Rogers and Michael D. Gray

Georgia Institute of Technology, Woodruff School of Mechanical Engineering, Atlanta GA, USA 30332-0405

PACS: 43.60.EK, 43.60.WY, 43.25.FE, 43.40.PH

ABSTRACT

In attempting to experimentally characterize an acoustic system with a progressive, non-essential nonlinearity that is driven by a small amplitude excitation as a nominally equivalent linear time-invariant (LTI) system, the relevant metric of measurement quality is the signal-to-noise-and-distortion (SINAD) ratio associated with the data rather than merely the signal-to-noise ratio (SNR) of the measurement. This sort of measurement involves an inherent trade-off between noise and distortion. At small drive levels, noise is problematic; whereas at larger drive levels, nonlinear distortion rather than noise limits the quality of the approximation. Although the noise can be independently characterized in a signal-free measurement, the distortion is not similarly accessible and must be neglected, assumed, or approximated. It is proposed that a reasonable empirical approximation to the SINAD can be achieved by applying a linear pulse compression to measured data such that the pulse response of the system along with some noise may be temporally separated from a significant portion of the measured noise and distortion. The approximate SINAD ratio, which is referred to as the apparent signal-to-noise ratio (ASNR) can then be formed with the ratio of energies computed for each of these separated components. This scheme was tested on both numerical and experimental systems and found to yield a suitable surrogate for the SINAD, which could be used to evaluate the relative merits of various candidate probing signals.

BACKGROUND

Several authors have suggested that particular time-domain probing signals are advantageous for the characterization of LTI systems in comparison with logical alternatives [1, 2, 3, 4]. These comparisons have focused on either observations of the characteristics of the signals under consideration or on idealized assumptions of the system under study. A generalized scheme for evaluating the relative merits of a particular drive signal and an optimal signal level has not been proposed, and relevant experimental comparisons have not been made. This is largely due to the difficulty involved in directly measuring the appropriate metric of performance: the SINAD ratio.

Four features are appropriate for a candidate probing signal. First, the signal must be of sufficient duration to sample the reverberation of the system under study. In the case of transient or impulsive signals, this implies that the signal (or signal record) include sufficient quiet time. In the case of periodic signals, it implies that the period must be longer than the reverberation time. Here *reverberation time* is considered in a qualitative sense as indicating the time required for the response of the system to decay to a level of relative insignificance for the purpose of the system's characterization rather than the time required for a specific level of decay. Second, the bandwidth of the probing signal must be sufficient for the characterization task, which is in-turn dependent on a subjective notion of *characterization* as all analytical representations of physical systems are approximations or abstractions. Third, the peak level(s) and/or energy of the signal must be sufficiently low that the system's behaviour is linear to the extent that it is required to be so for the purposes of the characterization and sufficiently high that the response can be distinguished from ambient noise. Although this trade-off might be easily resolved by arbitrarily extending integration times, there is a fourth requirement that the signal or the period in which it is sampled must be sufficiently short in time that the system under study does not vary from the state in which it must be characterized or, in more pragmatic sense, that the information acquired is still relevant.

For an approximation to the SINAD to be useful, it need not give the same numerical result provided that it can be used to correctly identify the optimal level for a particular signal and correctly rank the merit of a given signal with regard to others.

INTRODUCTION

A linear pulse compression [5] can be generally represented in the time domain as:

$$r(t) = DFT^{-1}\left[DFT(v(t)) \cdot DFT(r'(t)) \Big/ DFT(v'(t)) \right] \qquad [1]$$

where *DFT* and DFT^{-1} represent the discrete Fourier transform and its inverse, $v'(t)$ and $r'(t)$ represent the measured drive signal and response, and $v(t)$ and $r(t)$ represent the compression pulse shape and compressed response respectively. Provided that appropriate characteristics are selected for the signals and transforms, $r(t)$ is the response that would have been observed from a linear system if it had been subjected to the drive $v(t)$. This is a generalized formulation that reduces to the commonly used case of cross-correlation of the system's input and output when $v(t)$ is chosen to be the autocorrelation of $v'(t)$.

When Eq. 1 is applied to a system with observable nonlinearities, processing artefacts will result. These will not occur entirely at times representative of the system's actual response to $v(t)$. A simple way of viewing this is that the system-generated harmonics of a particular spectral component of $v'(t)$ will be incorrectly phased in the compression based on the corresponding fundamental components of $v'(t)$. These are temporally separate in a spectrographic sense because $v'(t)$ is not impulsive. Thus the artefacts will also tend to be temporally separated from the actual response in $r(t)$. The separation can be very large if the period of $v'(t)$ is much longer than the pulse width of $v(t)$, and can be separated from any components of the system's response if the period is longer than the reverberation time of the system. This does not mean that one of these separate parts accurately represents the linear response of the system. It is merely coincident with this response rather than identical to it. In contrast, the artefacts are entirely related to the nonlinear response.

Figure 1 depicts the temporal nature of these artefacts when they are produced by the quadratic-spring system described in the next section. In each of the four cases shown the system and the compression signal are identical. The linear response is therefore also identical. The portion of the response that is coincident with the linear response has been clipped at the far left side of this figure in order to emphasize the much smaller artefacts. The artefacts shown in this figure are representative of those that were produced by the 11 probing signals that were considered in the two subsequent sections.

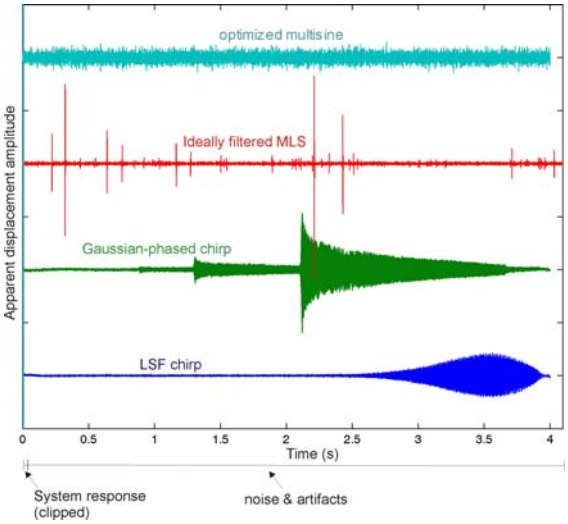

Figure 1. Typical nonlinear compression artefacts

MODELLED RESPONSES

In order to assess the viability of the proposed technique a simple system was considered: a parallel combination of a linear and a quadratic spring with force as the input variable and displacement as the output variable. The exact analytical solution for the system's response was therefore given by the quadratic formula and is represented in Eq. 2 and Eq.3.

$$F(t) = K_1 x(t) + K_2 x^2(t) \quad [2]$$

$$x(t) = \frac{-K_1 + \sqrt{K_1^2 + 4K_2 F(t)}}{2K_2} \quad [3]$$

Because this system has infinite bandwidth, the input was coupled to a 4-pole Butterworth filter with a corner frequency at ¼ of the Nyquist frequency to restrict this. The reason for including this additional complexity was that several of the candidate probing signals that were studied were discontinuous and therefore also had infinite bandwidth. This creates a difficulty in that, without bandwidth restrictions, the system's nonlinearity will exactly mimic the response of a different linear system at any specific drive level using the discontinuous drive.

Although, there was no relevant temporal scale for this system the same timescale was used as for the experimental results described in the following section. All of the drives and responses were sampled at 16 kHz. All of the probing signals covered a band of at least 50 Hz to 2 kHz. The same compression pulse was used for each of the simulations. This is given by:

$$v(t) = -A2\sqrt{2}\pi f_c (t - t_c) e^{-\left[(t-t_c)\sqrt{2}\pi f_c\right]^2} \quad [4]$$

where v is the source drive signal, A is an appropriate amplitude coefficient, f_c is the pulse centre frequency (450 Hz), and t_c is the time at which the waveform is centred (5 mS). The amplitude spectrum of this is given by a similar function:

$$|V(\omega)| = A\sqrt{\frac{\pi}{2}} f_c \omega e^{-\left[\omega/(2\sqrt{2}\pi f_c)\right]^2} \quad [5]$$

where ω is the radian frequency. The selection of compression pulse is indicative of the goals of system characterization both temporally and spectrally. In this case the choice seems somewhat arbitrary as it was intended only to match the characterization goals of the experimental system. The ASNR was formed from the ratio of energy in the first 12.5 ms to the energy at all later times.

Eleven probing signals were used. All of them were periodic (i.e. circular). These signals fall into two classes. The first of these are the seven analog signals which were generated from analytical expressions. These include (1) a linearly-swept-frequency (LSF) flat amplitude chirp, (2) a randomly phased spectrally flat multisine, (3) an LSF chirp with an amplitude envelope corresponding to the Gaussian compression pulse, (4) a flat amplitude chirp with a time-dependant phase that produced a Gaussian spectrum corresponding to the compression pulse, (5) a spectrally flat multisine with phasing numerically adjusted to minimize its crest factor (CF), which is the ratio of the signals peak value to its RMS value [6], (6) a multisine with a Gaussian spectrum and numerically optimized phasing and (7) a multisine with phases drawn from the discrete transform of a 2^{14}-1 bit maximum length sequence (MLS). The second class of signals are the four signals made up of binary sequences with zero-order-hold interpolation filters. These signals contained only values of ±1 by definition and were thus discontinuous regardless of the density of their sampling. The first of these signals was formed using a base band 2^{14}-1 bit MLS. This was sampled at 4 kHz resampled at 16 kHz such that each bit was represented as four identical values. All of the other signals were sufficiently sampled at 16 kHz by definition. The fundamental

period of the two MLS-based signals was 4.096 seconds. For all of the other signals it was 4.000 seconds. The other three binary signals were formed by applying a signum function to three of the analog signals: the LSF Chirp, the Gaussian-phased chirp, and the numerically optimized multisine.

Nonlinear compression artefacts for all of these signals were similar to one of the forms depicted in figure 1 with the source of probing signal phase determining this form. For linearly phased signals, artefacts were similar to those depicted in the first (blue) curve. Gaussian phasing produced artefacts like those in the second (green) curve. MLS-based phasing produced artefacts like those in the third (red) curve. Random phasing or numerically optimized phasing, which was seeded with a randomly phased starting point, produced artefacts like those in the fourth (cyan) curve.

Figure 2 shows a comparison for the exact SINAD, which was computed by subtracting the response of the system with the nonlinear spring and noise disabled to the compression pulse from the pulse-compressed response of the noisy nonlinear system, to the ASNR that was computed from the ratio of the time-separated components in the nonlinear responses. In this case and all of the modelled responses the value of K_2 was 1% of the value of K_1. Identical white noise was added to all of the responses at a level that was selected to place the optimal drive level within the dynamic range of the simulation. The probing signal represented in the figure is the LSF flat amplitude chirp. Its driving levels spanned a range of 30 dB in 1 dB steps with an upper limit of system collapse (i.e. the system does not recover from the peak compressive component of the probing signal). The two curves have been shifted vertically to match at the lowest drive level (-30 dB).

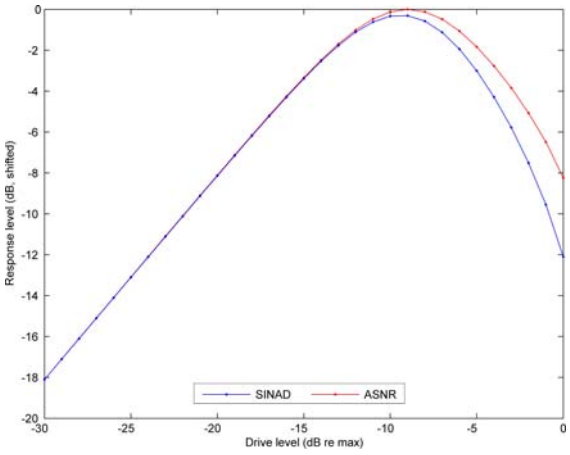

Figure 2. System output relative to noise and distortion

It is apparent from the figure that the ASNR and the actual SINAD are in good agreement with regard to both the extent of linear response and the optimal drive level (-9dB). This agreement is typical for all of the probing signals that were tested although the other sets of curves were shifted horizontally and vertically with respect to the set shown in the figure. The ASNR slightly overestimates the optimum drive level (by less than the 1-dB step size). This is also typical. The reason for this is that the nonlinear distortion and noise that are coincident with the response following a compression are neglected in the ASNR, which therefore consistently underestimates the total noise and distortion. The consistency and small size of this underestimate should make it easily correctable or negligible.

Figure 3 depicts the relative merits of the four binary probing signals for the characterization of the modelled system. Each bar indicates the peak Y value of a curve similar to those shown in figure 1 scaled to the highest peak for each of the eleven signals. Three values are depicted for each drive signal. The first of these, the uncompressed SINAD which is shown in red, represents the peak SINAD computed directly from the probing signal: the energy in the simulated response divided by the difference between the actual response and the noiseless linear response to the probing signal. The second bar, the peak SINAD which is shown in blue, represents a similar ratio where the pulse compression is included in both responses. The third bar, which is shown in green, represents the peak ASNR which was computed from the energy ratio of the time-separated components.

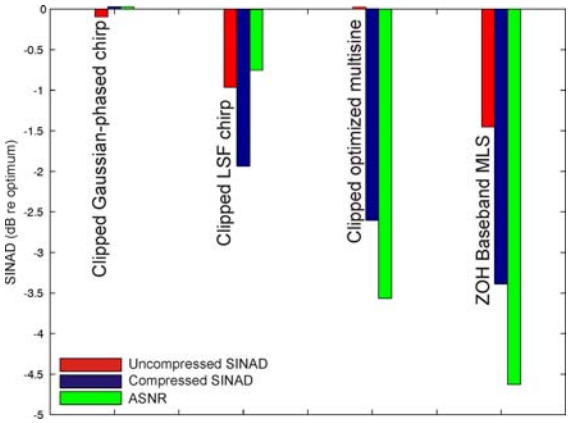

Figure 3. Performance of binary signals in modelled system

It is clear from the figure that the ASNR and SINAD are in agreement with regard to ordering these four signals based on their merit. There are quantitative disagreements between these metrics (particularly for the two best signals) with regard to the actual figures of merit. These are on the order of 1 dB. The uncompressed SINAD does not agree with the other two metrics with regard the ordering and merely sorts the four signals based on their crest factors evaluated after the low-pass filter (all the signals were binary and had CF≡1 before the filter) these range from CF=1.40 to CF=1.53. The disagreements are therefore indicative of the suitability of each signal for both the system and for the characterization task as it has been defined rather than for the system alone.

Figure 4 depicts the response of the system to the seven analog signals using the same convention and vertical scale as figure 3.

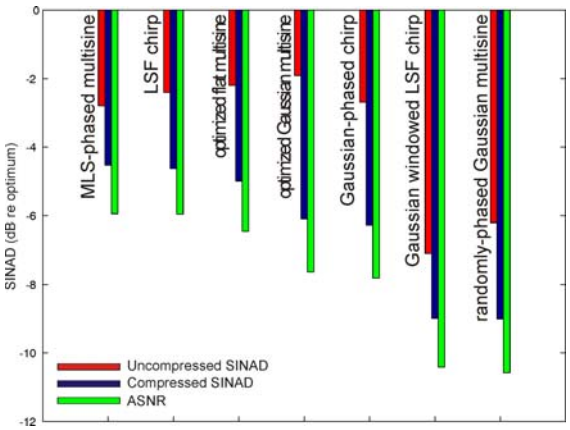

Figure 4. Performance of analog signals in modelled system

It is apparent that all of these signals, which had lower crest factors than the binary signals both in their archetypal form and after the system's low-pass filter, are inferior to the best binary signal by 4dB to 11dB. It is interesting to note that comparing the analog to the binary signals shows no clear relationship between the origin of the probing signal's phase and its merit. The best binary signal was based on a Gaussian phasing whereas the best analog signal was derived from an MLS. It is also interesting to note the extent to which the superiority of the analog signals is not well predicted by their crest factors. The best analog signal had CF=2.68 pre-filter and CF=2.45 post-filter compared to 1.41 and 1.45 for the optimized Gaussian multisine. As with the binary signals, post-filter crest factors correlate well with the uncompressed SINAD. In this regard, it is interesting that the best analog signal has no prescribed relationship with either the Gaussian spectrum of the compression pulse or the flat spectrum of the additive noise, but has a greater affinity for this compression task than the signals that were based on it.

MEASURED SEISMIC RESPONSES

An experiment was conducted in which the eleven probing signals previously described were used as the inputs to a seismic system that was constructed to image shallow-buried landmines [7]. These signals were used as voltage inputs for an audio amplifier that drove an electrodynamic shaker coupled to the surface of a large tank filled with damp compacted sand. The response of this system was the output of a vertically oriented accelerometer, which was buried at one of four fixed distances from the source. The closest accelerometer was directly below the foot coupling the source to the ground, and the furthest was 1.6 m away. The configuration of this experiment has been described in a previous paper [8]. In a linear sense, this is a theoretically complex system with multiple modes of propagation and multiple paths of propagation between the source and the receiver. Nonlinearity further complicates the system such that a predictive numerical model is not practical. Previous experiments have documented various manifestations of nonlinearity in this system including spall, nonlinear dispersion, and dynamic fluidization, in both the propagation path and the source-to-surface coupling.

The noise in the experiment was ambient and uncontrolled. The tank is built into the foundation of a building, and most of the observed noise is from building activities coupled through the structure. Although noise was sampled simultaneously on all four accelerometers, it undoubtedly varied between adjustments to probing signal levels and probing signal types. This may explain some variability in the data which was not apparent in the model, where identical noise was used in each simulation. Although similar in there overall shape, the experimental ASNR curves showed some fine-scale roughness that is unlike the curve in figure 1. This was probably due to the variability of the measured noise. There was no experimental way to directly measure the SINAD that the ASNR was used to approximate.

Table 1 indicates the computed peak ASNR values for each of the four transfer functions (accelerometer locations) and each of the 11 probing signals. This was determined by incrementing the drive signal amplitude over a 47-dB dynamic range in 3-dB steps. The table has been ordered by signal merit determined at a range of 1m from the source, although there is a clear reordering of signals from the other locations.

It is apparent from the table that there is very little resemblance to the ordering of probing signals for the modelled system. This is reinforces the idea that signal merit is system-dependent and is not intrinsic to particular signals regardless of their application. There is no clear demarcation between the analog and binary signals in the physical system, although the latter obviously improve at closer ranges. This suggests that the nonlinearity of the source-to-surface coupling, which is a force (driving current)-to-displacement response, is better modelled by lumped-element nonlinearity than are the nonlinearities associated with the wave propagation.

Table 1. Measured ASNR for a seismic system

Signal	CF	ASNR R=0 m (dB re max)	ASNR R=0.4 m (dB) (dB re max)	ASNR R=1 m (dB) (dB re max)	ASNR R=1.6 m (dB) (dB re max)
LSF flat chirp	1.47	0.00	0.00	0.00	0.00
Clipped Gaussian-phased chirp	1.00	-0.16	-1.97	-3.34	-3.05
Gaussian-phased chirp	1.67	-2.96	-5.09	-3.43	-2.13
ZOH baseband MLS	1.00	-1.62	-1.63	-3.84	-3.78
Optimized flat multisine	1.42	-2.15	-2.21	-4.29	-4.54
clipped optimized multisine	1.00	-1.78	-4.04	-4.78	-4.34
Optimized Gaussian multisine	1.42	-2.44	-5.40	-5.06	-4.24
Gaussian-windowed LSF Chirp	2.77	-5.16	-7.14	-5.58	-3.86
MLS-phased multisine	2.68	-2.66	-3.50	-5.60	-5.60
Clipped LSF chirp	1.00	-0.86	-2.91	-5.77	-5.78
Randomly-phased Gaussian Multisine	4.54	-5.91	-9.86	-10.03	-9.08

The LSF chirp yields the best performance at all ranges, although the magnitude of its advantage over the best of the binary signals is negligible at very short range. There is no obvious feature of this signal or precedent in the previously modelled responses to account for its superiority. Both its spectral content and crest factor are comparable to the optimized-flat-multisine probing signal which shows a deficit of more than 4dB at longer ranges. This, along with the modelled responses, points to the importance of probing signal phase (independent of its impact on a signal's crest factor) in determining merit.

TIME-VARYING SYSTEMS

It is difficult to assess the effectiveness of probing signals for the interrogation of slowly time varying systems because the source(s) of system variability on longer timescales are not often controllable. In the previous two examples, where time invariance was assumed, there were three time scales relevant to the characterization problem: the minimum timescale implied by the bandwidth of the probing signal, the reverberation time of the system, and the total interrogation time. To correctly assess the impact of nonlinearities on the approximate linear characterization of the system, it was necessary that the interrogation time be significantly longer than either of the other timescales. If the system is time varying, this variation introduces a fourth relevant time scale, which cannot be neglected unless it is much greater than the total interrogation time.

This problem was examined in the context of an ultrasonic vibrometry system that is the subject of a forthcoming paper. Here long timescale variations are produced by audio-frequency shear waves, which propagate very slowly (~5 m/s) in the medium of interest (a tissue-mimicking phantom). The waves are produced by a shaker that is mechanically coupled to the surface of the phantom. The characterization goal, in this case, is determining the instantaneous pulse response of the medium to incident ultrasound i.e. the location of each scattering site along the beam of an interrogating

ultrasonic transducer. This is equivalent to reproducing the pulse response of the system to the most temporally discrete pulse that the ultrasonic bandwidth would permit. In the case of the system that was under study, the ultrasonic bandwidth spanned $f_c = 2\pi\omega_c = 2.5$ MHz ± 400 kHz and the response to a drive pulse of the form

$$v(t) = \sin(\omega_c(t - t_c))e^{-\omega_c^2 \frac{1}{1200}(t-t_c)^2} \quad [6]$$

was sought. Low-frequency motion was generated in the range of 20 Hz to 250 Hz. The period of the ultrasonic interrogation signal was 500 µs, which was slightly longer than the reverberation time of the phantom. The total interrogation time was 0.5 s (i.e. 1,000 cycles). Post-processing involved the sequential compression of 500 µs segments of the backscattered ultrasonic signal. Thus, there was a significant change in the medium during each interval of its interrogation. Since the low-frequency source could be phase locked to the ultrasonic signal, it was possible to investigate the effects of this change with different probing signals.

In the experiment three probing signals were used. Drive levels for each were determined by reducing the drive in 6 dB increments until there was no perceptible change in the processed data other than noise, and then increasing the averaging time to make the effects of the noise negligible. This required averaging for up to 17 minutes (2,000 shear-wave cycles). Three different probing signals were tested. The first was the compression pulse itself, which was used as the baseline for "truth", the best achievable result if noise could be neglected. This "truth" included errors from sources unrelated to the selection of probing signal, and thus did not represent a perfect characterization of the medium. The other two signals were an LSF chirp and a randomly phased multisine. Their spectral content and crest factors were nominally the same, thus the only discrepancy from "truth' resulted from the effects of the time variance on the signal.

Figure 5 depicts the discrepancies in arrival times of a low-frequency shear pulse determined using the three probing signals, where these were computed by a cross-correlation of the computed displacements with the low-frequency drive signal. The blue dots were computed by subtraction the arrival times computed from the "truth" measurement from those computed from the chirp measurement. The red dots were computed by subtracting this from the times computed from the multisine measurements. The spatial region indicated on the horizontal axis (5 to 20 cm) was the region over which the effects of noise could be neglected when 17 minutes of averaging was used for the "truth" measurement.

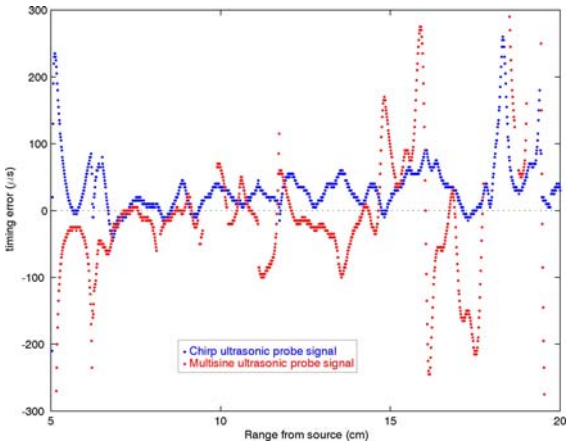

Figure 5. Timing errors for two different probing signals

It is clear from the figure that different errors result from the different probing signals and that the chirp signal is superior to the multisine in this regard. This is probably due to the fact that the chirp's spectral and temporal centres, in a spectrographic sense, are coincident. This example illustrates the importance of the phasing of a probing signal when the approximating a system as time invariant, regardless of nonlinearity.

CONCLUSIONS

Errors in LTI system characterization and approximation have been explored in three different contexts. In a very simple modelled system with nonessential nonlinearity, an easily accessible metric, the ASNR, was found to be a suitable surrogate for the normally inaccessible SINAD both in determining optimal drive level and the relative merit of different probing signal types. In a complicated experimental system, the ASNR was used to compare the merits of various probing signals. This comparison revealed the importance of a signal's phasing in determining its merit independently of the effect of this on the signal's crest factor. The comparison also suggested that signal merit is system-dependent and cannot be judged outside of this context. In a second experimental system the relevance of phasing in determining signal merit when confronting slow time variance was revealed.

This work is ongoing. More complex modelled systems need to be studied in order to determine the robustness of the ASNR as a surrogate for the SINAD. Similarly, different experimental systems, which can be used to directly validate these models, need to be studied. Time variance also needs to be addressed in a more systematic way with a greater variety of probing signals and a more accessible baseline truth.

REFERENCES

1 K.R. Godfrey, H.A. Barker, and A.J. Tucker, "Comparison of Perturbation Signals for Linear System Identification in the Frequency Domain," *IEE Proc. Control Theory Appl.*, Vol. **146**, no. 6, pp. 535-48, November 1999.

2 D. D. Rife, and J. Vanderkooy, "Transfer Function Measurements with Maximum-Length Sequences," *J. Audio Engin. Soc.*, Vol. **37**, no.6, pp. 419-44, June 1989.

3 J. Schoukens, R. Pintelon, E. Van Der Ouderaa, and J. Renneboog, "Survey of Excitation Signals for FFT Based Signal Analyzers," *IEEE Trans. on Instrum. Meas.*, Vol. **37**, no.3, pp. 342-52, September 1988.

4 G.E.J. Bold, M.D. Johns, and T.G. Birdsall, "Signals for Optimal Bandlimited System Interrogation with Application to Underwater Acoustics," *J. Acous. Soc. Am.*, Vol. **81**, no.4, pp. 991-99, April 1987.

5 P.E. Blankenship and E.M. Hofstetter, "Digital Pulse Compression via Fast Convolution", *IEEE Trans. on Acoustics, Speech, and Signal Processing*, Vol. **ASSP-23**, No. 2, April 1975.

6 E. Van Der Ouderaa, J. Schoukens, and J. Renneboog, "Peak Factor Minimization Using a Time-Frequency Swapping Algorithm," *IEEE Trans. Instrum. Meas.*, Vol. **37**, no.1, pp. 145-147, March 1988.

7 W.R. Scott, J.S. Martin, and G.D. Larson, "Experimental Model for a Seismic Landmine Detection System," *IEEE Trans. on Geoscience and Remote Sensing*, vol. **39**, no. 6, June 2001, pp. 1155-64.

8 J.S. Martin, W.R. Scott, G.D. Larson, P.H. Rogers and G.S. McCall, "Probing signal design for seismic landmine detection," in *Proc. SPIE: 2004 Annu. Int. Symp. Aerospace/Defense Sensing, Simulation, and Controls*, Orlando, FL, Vol. **5415**, pp. 133-144, April 2004.