IEEE STANDARDS ASSOCIATION

IEEE Standard for Terminology and Test Methods for Analog-to-Digital Converters

IEEE Instrumentation & Measurement Society

Sponsored by the
Waveform Generation Measurement and Analysis Technical Committee

IEEE
3 Park Avenue
New York, NY 10016-5997
USA

14 January 2011

IEEE Std 1241™-2010
(Revision of
IEEE Std 1241-2000)

IEEE Std 1241™-2010
(Revision of
IEEE Std 1241-2000)

IEEE Standard for Terminology and Test Methods for Analog-to-Digital Converters

Sponsor

Waveform Generation Measurement and Analysis Technical Committee
of the
IEEE Instrumentation & Measurement Society

Approved 17 June 2010

IEEE-SA Standards Board

Portions of Annex A, including text and Figure A.1 through Figure A.8, have been reprinted with permission from "ADC Parameters and Characteristics," which was written by Sergio Rapuano, Pasquale Daponte, Eulalia Balestrieri, Luca De Vito, Steven J. Tilden, Solomon Max, and Jerome Blair. The article can be found in *IEEE Instrumentation and Measurement Magazine,* vol. 8, no. 5, pp. 44–54, Dec. 2005 [B46]. © 2005 IEEE.

Abstract: The material presented in this standard is intended to provide common terminology and test methods for the testing and evaluation of analog-to-digital converters (ADCs). This standard considers only those ADCs whose output values have discrete values at discrete times, i.e., they are quantized and sampled. In general, this quantization is assumed to be nominally uniform (the input-output transfer curve is approximately a straight line) as discussed further in 1.3, Analog-to-digital converter background, and the sampling is assumed to be at a nominally uniform rate. Some but not all of the test methods in this standard can be used for ADCs that are designed for non-uniform quantization.
Keywords: ADC, analog-to-digital converter, code transition level, coherent sampling, DNL, ENOB, histogram, INL, LSB, missing codes, noise power ratio, noncoherent sampling, quantization error, quantization noise, SAR, SFDR, sine fitting

The Institute of Electrical and Electronics Engineers, Inc.
3 Park Avenue, New York, NY 10016-5997, USA

PDF: ISBN 978-0-7381-6239-3 STD96055
Print: ISBN 978-0-7381-6240-9 STDPD96055

IEEE Standards documents are developed within the IEEE Societies and the Standards Coordinating Committees of the IEEE Standards Association (IEEE-SA) Standards Board. The IEEE develops its standards through a consensus development process, approved by the American National Standards Institute, which brings together volunteers representing varied viewpoints and interests to achieve the final product. Volunteers are not necessarily members of the Institute and serve without compensation. While the IEEE administers the process and establishes rules to promote fairness in the consensus development process, the IEEE does not independently evaluate, test, or verify the accuracy of any of the information or the soundness of any judgments contained in its standards.

Use of an IEEE Standard is wholly voluntary. The IEEE disclaims liability for any personal injury, property or other damage, of any nature whatsoever, whether special, indirect, consequential, or compensatory, directly or indirectly resulting from the publication, use of, or reliance upon this, or any other IEEE Standard document.

The IEEE does not warrant or represent the accuracy or content of the material contained herein, and expressly disclaims any express or implied warranty, including any implied warranty of merchantability or fitness for a specific purpose, or that the use of the material contained herein is free from patent infringement. IEEE Standards documents are supplied "**AS IS**."

The existence of an IEEE Standard does not imply that there are no other ways to produce, test, measure, purchase, market, or provide other goods and services related to the scope of the IEEE Standard. Furthermore, the viewpoint expressed at the time a standard is approved and issued is subject to change brought about through developments in the state of the art and comments received from users of the standard. Every IEEE Standard is subjected to review at least every five years for revision or reaffirmation, or every ten years for stabilization. When a document is more than five years old and has not been reaffirmed, or more than ten years old and has not been stabilized, it is reasonable to conclude that its contents, although still of some value, do not wholly reflect the present state of the art. Users are cautioned to check to determine that they have the latest edition of any IEEE Standard.

In publishing and making this document available, the IEEE is not suggesting or rendering professional or other services for, or on behalf of, any person or entity. Nor is the IEEE undertaking to perform any duty owed by any other person or entity to another. Any person utilizing this, and any other IEEE Standards document, should rely upon his or her independent judgment in the exercise of reasonable care in any given circumstances or, as appropriate, seek the advice of a competent professional in determining the appropriateness of a given IEEE standard.

Interpretations: Occasionally questions may arise regarding the meaning of portions of standards as they relate to specific applications. When the need for interpretations is brought to the attention of IEEE, the Institute will initiate action to prepare appropriate responses. Since IEEE Standards represent a consensus of concerned interests, it is important to ensure that any interpretation has also received the concurrence of a balance of interests. For this reason, IEEE and the members of its societies and Standards Coordinating Committees are not able to provide an instant response to interpretation requests except in those cases where the matter has previously received formal consideration. A statement, written or oral, that is not processed in accordance with the IEEE-SA Standards Board Operations Manual shall not be considered the official position of IEEE or any of its committees and shall not be considered to be, nor be relied upon as, a formal interpretation of the IEEE. At lectures, symposia, seminars, or educational courses, an individual presenting information on IEEE standards shall make it clear that his or her views should be considered the personal views of that individual rather than the formal position, explanation, or interpretation of the IEEE.

Comments for revision of IEEE Standards are welcome from any interested party, regardless of membership affiliation with IEEE. Suggestions for changes in documents should be in the form of a proposed change of text, together with appropriate supporting comments. Recommendations to change the status of a stabilized standard should include a rationale as to why a revision or withdrawal is required. Comments and recommendations on standards, and requests for interpretations should be addressed to:

Secretary, IEEE-SA Standards Board
445 Hoes Lane
Piscataway, NJ 08854
USA

Introduction

This introduction is not part of IEEE Std 1241-2010, IEEE Standard for Terminology and Test Methods for Analog-to-Digital Converters.

This standard defines the terms, definitions, and test methods used to specify, characterize, and test analog-to-digital converters (ADCs). It is intended for the following:

- Individuals and organizations who specify ADCs to be purchased
- Individuals and organizations who purchase ADCs to be applied in their products
- Individuals and organizations whose responsibility is to characterize and write reports on ADCs available for use in specific applications
- Suppliers interested in providing high-quality and high-performance ADCs to acquirers

This standard is designed to help organizations and individuals

- Incorporate quality considerations during the definition, evaluation, selection, and acceptance of supplier ADCs for operational use in their equipment
- Determine how supplier ADCs should be evaluated, tested, and accepted for delivery to end users

This standard is intended to satisfy the following objectives:

- Promote consistency within organizations in acquiring third-party ADCs from component suppliers
- Provide useful practices on including quality considerations during acquisition planning
- Provide useful practices on evaluating and qualifying supplier capabilities to meet user requirements
- Provide useful practices on evaluating and qualifying supplier ADCs
- Assist individuals and organizations judging the quality and suitability of supplier ADCs for referral to end users

Several standards have previously been written that address the testing of analog-to-digital converters either directly or indirectly. These include

- IEEE Std 1057-2007, which describes the testing of waveform recorders. This standard has been used as a guide for many of the techniques described in this standard.
- IEEE Std 746-1984, which addresses the testing of analog-to-digital and digital-to-analog converters used for PCM television video signal processing.
- JESD99-1, which deals with the terms and definitions used to describe analog-to-digital and digital-to-analog converters. This standard does not include test methods.

IEEE Std 1241-2009 for analog-to-digital converters is intended to focus specifically on terms and definitions as well as test methods for ADCs for a wide range of applications.

This standard is a revision of IEEE Std 1241-2000. This version has added additional test methods, improved guidance for selecting tests, and additional terms. Some terminology pertaining to signal-to-noise ratio and related terms has been changed to be consistent with other standards.

Notice to users

Laws and regulations

Users of these documents should consult all applicable laws and regulations. Compliance with the provisions of this standard does not imply compliance to any applicable regulatory requirements. Implementers of the standard are responsible for observing or referring to the applicable regulatory requirements. IEEE does not, by the publication of its standards, intend to urge action that is not in compliance with applicable laws, and these documents may not be construed as doing so.

Copyrights

Updating of IEEE documents

Users of IEEE standards should be aware that these documents may be superseded at any time by the issuance of new editions or may be amended from time to time through the issuance of amendments, corrigenda, or errata. An official IEEE document at any point in time consists of the current edition of the document together with any amendments, corrigenda, or errata then in effect. In order to determine whether a given document is the current edition and whether it has been amended through the issuance of amendments, corrigenda, or errata, visit the IEEE Standards Association web site at http://ieeexplore.ieee.org/xpl/standards.jsp, or contact the IEEE at the address listed previously.

For more information about the IEEE Standards Association or the IEEE standards development process, visit the IEEE-SA web site at http://standards.ieee.org.

Errata

Errata, if any, for this and all other standards can be accessed at the following URL: http://standards.ieee.org/reading/ieee/updates/errata/index.html. Users are encouraged to check this URL for errata periodically.

Interpretations

Current interpretations can be accessed at http://standards.ieee.org/reading/ieee/interp/index.html.

Patents

Attention is called to the possibility that implementation of this standard may require use of subject matter covered by patent rights. By publication of this standard, no position is taken with respect to the existence or validity of any patent rights in connection therewith. The IEEE is not responsible for identifying Essential Patent Claims for which a license may be required, for conducting inquiries into the legal validity or scope of Patents Claims or determining whether any licensing terms or conditions provided in connection with submission of a Letter of Assurance, if any, or in any licensing agreements are reasonable or non-discriminatory. Users of this standard are expressly advised that determination of the validity of any patent rights, and the risk of infringement of such rights, is entirely their own responsibility. Further information may be obtained from the IEEE Standards Association.

Participants

At the time this standard was submitted to the IEEE-SA Standards Board for approval, the Analog-to-Digital Converters Working Group had the following membership:

Steve Tilden, *Chair*
Solomon Max, *Secretary*
Jerry Blair, *Editor*

Francisco Alegria
Eulalia Balestrieri
Niclas Björsell
John Calvin
Dominique Dallet
Pasquale Daponte
Luca De Vito
Alexander Goncharenko
Donald Greer
Richard Liggiero
Tom Linnenbrink*
Sergio Rapuano
Fang Xu

*Chair, TC-10

Contributions were also made in prior years by:

Pasquale Arpaia
B. N. Suresh Babu
Eulalia Balestrieri
Allan Belcher
David Bergman
Niclas Björsell
Jerry Blair
Eric Blom
William Boyer
Steve Broadstone
Paulo Carbone
Giovanni Chiorboli
Dominique Dallet
Pasquale Daponte
Luca De Vito
John Deyst
Robert Graham
Philip Green
David Hansen
Fred Irons
Dan Kien
Dan Knierim
Richard Kromer
Yves Langard
Richard Liggiero
Solomon Max
W. Thomas Meyer
Carlo Morandi
Norris Nahman
Bill Peterson
Sergio Rapuano
Pierre-Yves Roy
Otis M. Solomon
T. Michael Souders
Steve Tilden

The following members of the individual balloting committee voted on this standard. Balloters may have voted for approval, disapproval, or abstention.

Ali Al Awazi
Martin J. Bishop
Niclas Björsell
Jerome Blair
William Boyer
Paolo Carbone
Frans G. De Jong
Alexander Goncharenko
Ron Greenthaler
Donald Greer
Randall Groves
Timothy Harrington
Werner Hoelzl
Piotr Karocki
Jim Kulchisky
Donald Larson
Richard Liggiero
Thomas Linnenbrink
William Lumpkins
S. Max
Scott Misha
Michael S. Newman
John Noonan
Ulrich Pohl
Sergio Rapuano
Robert Robinson
Bartien Sayogo
Gil Shultz
James Smith
Joseph Stanco
Walter Struppler
Steven Tilden
Stephen Webb
Fang Xu

When the IEEE-SA Standards Board approved this standard on 17 June 2010, it had the following membership:

Robert M. Grow, *Chair*
Richard H. Hulett, *Vice Chair*
Steve M. Mills, *Past Chair*
Judith Gorman, *Secretary*

Karen Bartleson	Young Kyun Kim	Ronald C. Petersen
Victor Berman	Joseph L. Koepfinger*	Thomas Prevost
Ted Burse	John Kulick	Jon Walter Rosdahl
Clint Chaplin	David J. Law	Sam Sciacca
Andy Drozd	Hung Ling	Mike Seavey
Alexander Gelman	Oleg Logvinov	Curtis Siller
Jim Hughes	Ted Olsen	Don Wright

*Member Emeritus

Also included are the following nonvoting IEEE-SA Standards Board liaisons:

Satish Aggarwal, *NRC Representative*
Richard DeBlasio, *DOE Representative*
Michael Janezic, *NIST Representative*

Lisa Perry
IEEE Standards Program Manager, Document Development

Kathryn Bennett
IEEE Standards Program Manager, Technical Program Development

Contents

IEEE Standard for Terminology and Test Methods for Analog-to-Digital Converters

1. Overview

1.1 Scope

The material presented in this standard is intended to provide common terminology and test methods for the testing and evaluation of analog-to-digital converters (ADCs). This standard considers only those ADCs whose output values have discrete values at discrete times, i.e., they are quantized and sampled. In general, this quantization is assumed to be nominally uniform (the input-output transfer curve is approximately a straight line) as discussed further in 1.3, and the sampling is assumed to be at a nominally uniform rate. Some but not all of the test methods in this standard can be used for ADCs that are designed for non-uniform quantization.

1.2 Purpose

This standard identifies ADC error sources and provides test methods with which to perform the required error measurements. The information in this standard is useful both to manufacturers and to users of ADCs in that it provides a basis for evaluating and comparing existing devices, as well as providing a template for writing specifications for the procurement of new ones. In some applications, the information provided by the tests described in this standard can be used to correct ADC errors, e.g., correction for gain and offset errors. The reader should note that this standard has many similarities to IEEE Std 1057. Many of the tests and terms are nearly the same, since ADCs are a necessary part of digitizing waveform recorders.

1.3 Document organization

This standard is divided into fourteen clauses plus annexes. Clause 1 is a basic orientation. For further investigation, users of this standard can consult Clause 2, which contains references to other IEEE standards on waveform measurement and relevant International Standardization Organization (ISO) documents. The definitions of technical terms and symbols used in this standard are presented in Clause 3. Clauses 4 through 14 present a wide range of tests that measure the performance of an analog-to-digital converter. Annexes, containing the bibliography and informative comments on the tests presented in Clauses 4 through 14, augment the standard.

1.4 Analog-to-digital converter background

This standard considers only those ADCs whose output values have discrete values at discrete times, i.e., they are quantized and sampled. Although different methods exist for representing a continuous analog signal as a discrete sequence of binary words, an underlying model implicit in many of the tests in this standard assumes that the relationship between the input signal and the output values approximates the ideal staircase transfer curve depicted in Figure 1(a). Applying this model to a voltage-input ADC, the full-scale input range (FSR) of the ADC is divided into uniform intervals, known as code bins, with nominal width Q. The number of code transition levels in the discrete transfer function is equal to $2^N - 1$, where N is the number of digitized bits of the ADC. Note that there are ADCs that are designed such that N is not an integer, i.e., the number of code transition levels is not an integral power of two. Inputs below the first transition or above the last transition are represented by the most negative and positive output codes, respectively. Note, however, that two conventions exist for relating V_{min} (often called –FS, or negative full scale) and V_{max} (often called +FS, or positive full scale) to the nominal transition points between code levels, *mid-tread* and *mid-riser*.

The dotted lines at V_{min}, V_{max}, and $(V_{min} + V_{max})/2$ indicate what is often called the *mid-tread* convention, where the first transition is $Q/2$ above V_{min} and the last transition is $3Q/2$, below V_{max}. This convention gets its name from the fact that the midpoint of the range, $(V_{min} + V_{max})/2$, occurs in the middle of a code, i.e., on the *tread* of the staircase transfer function. The second convention, called the *mid-riser* convention, is indicated in the figure by dashed lines at V_{min}, V_{max}, and $(V_{min} + V_{max})/2$. In this convention, V_{min} is $-Q$ from the first transition, V_{max} is $+Q$ from the last transition, and the midpoint, $(V_{min} + V_{max})/2$, occurs on a staircase riser. The difference between the two conventions is a displacement along the voltage axis by an amount $Q/2$.

For all tests in this standard, this displacement has no effect on the results and either convention may be used. The one place where it does matter is when a device provides or expects user-provided reference signals. In this case, the manufacturer must provide the necessary information relating the reference levels to the code transitions. In both conventions the number of code transitions is $2^N - 1$ and the full-scale range, FSR, is from V_{min} to V_{max}. Even in an ideal ADC, the quantization process produces errors. These errors contribute to the difference between the actual transfer curve and the ideal straight-line transfer curve, which is plotted as a function of the input signal in Figure 1(b).

To use this standard, the user must understand how the transfer function maps its input values to output codewords, and how these output codewords are converted to the code bin numbering convention used in this standard. As shown in Figure 2(a), the lowest code bin is numbered 0, the next is 1, and so on up to the highest code bin, numbered $(2^N - 1)$. In addition to unsigned binary [Figure 1(a)], ADCs may use 2's complement, sign-magnitude, Gray, binary-coded decimal (BCD), or other output coding schemes. In these cases, a simple mapping of the ADC's consecutive output codes to the unsigned binary codes can be used in applying various tests in this standard. Note that in the case of an ADC whose number of distinct output codes is not an integral power of 2 (e.g., a BCD-coded ADC), the number of digitized bits N is still defined, but will not be an integer.

Real ADCs have other errors in addition to the nominal quantization error shown in Figure 1(b). All errors can be divided into the categories of static and dynamic, depending on the rate of change of the input signal at the time of digitization. A slowly varying input can be considered a static signal if its effects are equivalent to those of a constant signal. Static errors, which include the quantization error, usually result from non-ideal spacing of the code transition levels. Dynamic errors occur because of additional sources of error induced by the time variation of the analog signal being sampled. Sources include harmonic distortion from the analog input stages, signal-dependent variations in the time of samples, dynamic effects in internal amplifier and comparator stages, and frequency-dependent variation in the spacing of the quantization levels.

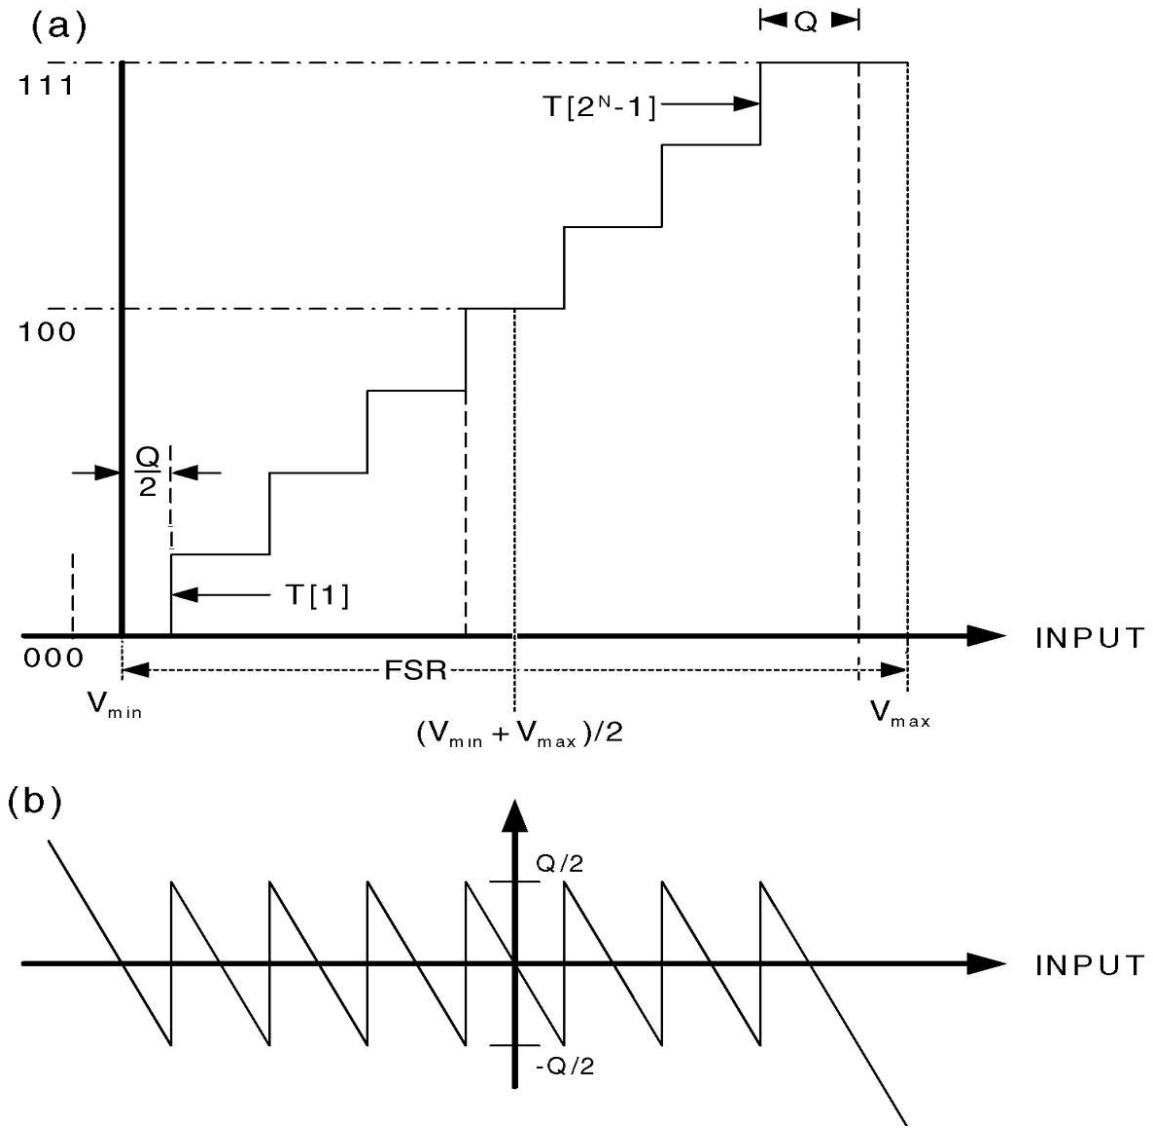

Figure 1—Staircase ADC transfer function, having full-scale range (FSR) and $2^N - 1$ levels, corresponding to *N*-bit quantization

There are two standard methods for characterizing the error of an ADC:

a) The transition levels are evaluated as $T[1]$ through $T[2^N - 1]$.

   — For $i = 1$ to $2^N - 2$, the quantity $T'[i] = (T[i] + T[i + 1])/2$ is evaluated.

   — The deviations from nominal of $T'[i]$ are used to define the errors of the ADC.

b) The transition levels are evaluated as $T[1]$ through $T[2^N - 1]$.

   — The deviations from nominal of $T[i]$ are used to define the errors of the ADC.

The two methods are shown graphically in Figure 2. The ADC characteristics shown in Figure 2 help illustrate the two methods of describing ADC errors. The following should be noted:

1) The location of the code edges $T[1]$, $T[2]$, $T[3]$, $T[4]$, $T[5]$, and $T[7]$ are the same as the ideal code edges illustrated in Figure 1.

2) The location of $T[6]$ is shifted from the nominal position by $+Q/2$.

3) Figure 2(b) shows the error in the ADC output for any input. Note that the maximum error ranges from $-Q/2$ to $+Q/2$ in the area where the code transitions are ideal. In the region around $T[6]$ the error peaks at a value of $+Q$ before returning to a value of 0.

4) The ADC is assumed to be noise-free. The ADC is assumed to be end-point calibrated.

5) Figure 2(c) is a plot of the ADC error based on code edges. The error is 0 for all codes except for code 1102 where the error is $+Q/2$.

6) Figure 2(d) is a plot of the ADC errors based on code centers.

The data for Figure 2 is tabulated below.

Code	Ideal T[code]	Actual T[code]	Figure 2(c)	Mid-step ideal	Mid-step actual	Figure 2(d)
0	NA	NA	NA	NA	NA	NA
1	0.5Q	0.5Q	0.0Q	1.0Q	1.0Q	0.0Q
2	1.5Q	1.5Q	0.0Q	2.0Q	2.0Q	0.0Q
3	2.5Q	2.5Q	0.0Q	3.0Q	3.0Q	0.0Q
4	3.5Q	3.5Q	0.0Q	4.0Q	4.0Q	0.0Q
5	4.5Q	4.5Q	0.0Q	5.0Q	5.25Q	0.25Q
6	5.5Q	6.0Q	0.5Q	6.0Q	6.25Q	0.25Q
7	6.5Q	6.5Q	0.0Q	NA	NA	NA

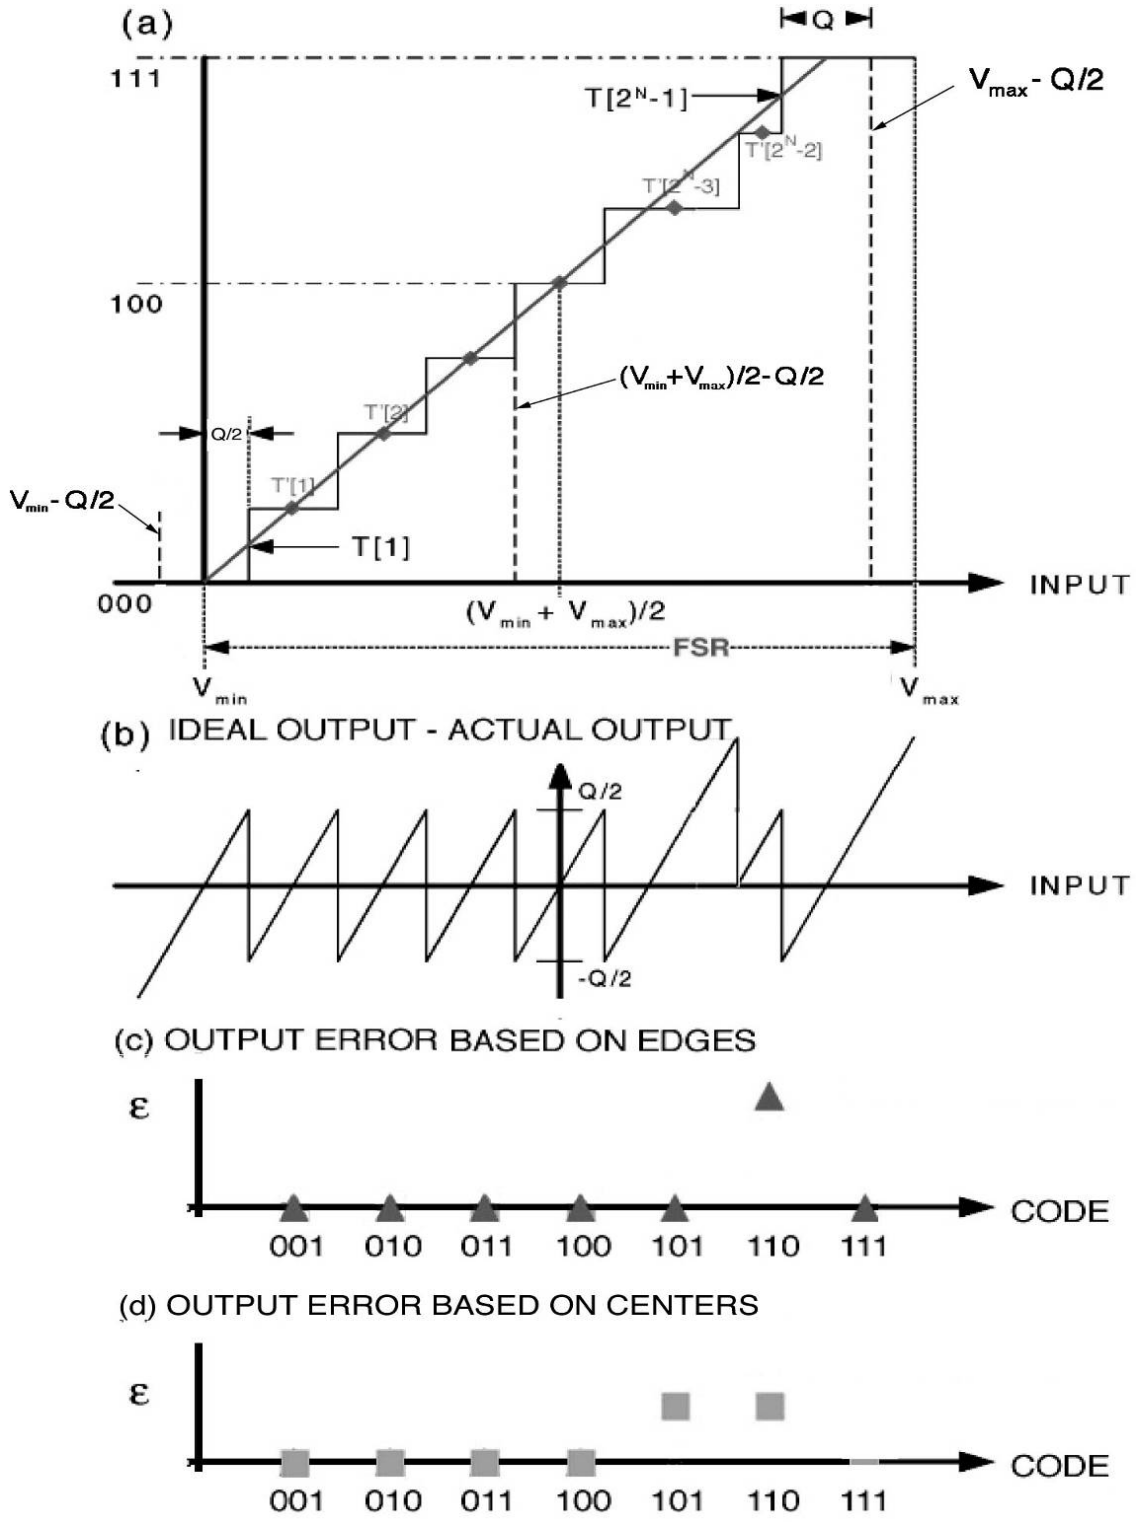

Figure 2—Methods for characterizing the error of an ADC

There are advantages and disadvantages to both methods of analyzing ADC errors

a) Code center error analysis (mid-tread):
   1) AC analyses such as THD, SINAD, and SNR are better defined by code center error analyses.
   2) For a given converter, INL errors are smaller when based on code centers.
   3) Histogram error analyses are more closely related to code center error analysis.

b) Code edge error analysis (mid-riser):
   1) ADCs that are used for instrumentation are better defined by code edge analysis.
   2) Servo measurement techniques of an ADC transfer function identify code edges.

1.5 Guidance to the user

Conditions described in this section are a general overview of the test environment. More detailed conditions about the environment will be included in relevant sections of the standard.

1.5.1 Interfacing

ADCs present unique interfacing challenges, and without careful attention users can experience substandard results. As with all mixed-signal devices, ADCs perform as expected only when the analog and digital domains are brought together in a well-controlled fashion. The user should fully understand the manufacturer's recommendations with regard to proper signal buffering and loading, input signal connections, transmission line matching, circuit layout patterns, power supply decoupling, and operating conditions. Edge characteristics for start-convert pulse(s) and clock(s) must be carefully chosen to maintain input signal purity with sufficient margin up to the analog input pin(s). Most manufacturers now provide excellent ADC evaluation boards, which demonstrate recommended layout techniques, signal conditioning, and interfacing for their ADCs. If the characteristics of a new ADC are not well understood, then these boards should be analyzed or used before starting a new layout.

1.5.2 Test conditions

ADC test specifications can be split into two groups: test conditions and test results. Typical examples of the former are: temperature, power supply voltages, clock frequency, and reference voltages. Examples of the latter are: power dissipation, effective number of bits, spurious free dynamic range (SFDR), and integral non-linearity (INL). The test methods defined in this standard describe the measurement of test results for given test conditions.

ADC specification sheets will often give allowed ranges for some test condition (e.g., power supply ranges). This implies that the ADC will function properly and that the test results will fall within their specified ranges for all test conditions within their specified ranges.

Since the test condition ranges are generally specified in continuous intervals, they describe an infinite number of possible test conditions, which obviously cannot be exhaustively tested. It is up to the manufacturer or tester of an ADC to determine, from design knowledge and/or testing, the effect of the test conditions on the test result, and from there to determine the appropriate set of test conditions needed to accurately characterize the range of test results. For example, knowledge of the design may be sufficient to know that the highest power dissipation (test result) will occur at the highest power supply voltage (test condition), so the power dissipation test need be run only at the high end of the supply voltage range to check that the dissipation is within the maximum of its specified range. It is very important that relevant test conditions be stated when presenting test results.

1.5.3 Electrical environment

When designing a test setup, one should consider the electrical environment. Prevent external disturbances from affecting the test results. Examples of external disturbances include common-mode noise and poor grounding in the set-up. Computers and high-power radio frequency (RF) signals like TV and radio transmitters can be sources of interference.

1.5.4 Test equipment

One must verify that the performance of the test equipment used for these tests significantly exceeds the desired performance of the ADC under evaluation. Users will likely need to include additional signal conditioning in the form of filters and pulse shapers. Accessories such as terminators, attenuators, delay lines, and other such devices are usually needed to match signal levels and to provide signal isolation to avoid corrupting the input stimuli.

Quality testing requires following established procedures, most notably those specified in IEEE Std ISO 9001: 2000 [B26]. In particular, traceability of instrumental calibration to a known standard is important. Commonly used test setups are described in 4.2.

1.5.5 Test selection

When choosing which parameters to measure, one should follow the outline and hints in this clause to develop a procedure that logically and efficiently performs all needed tests on each unique setup. The standard has been designed to facilitate the development of these test procedures. In this standard the discrete Fourier transform (DFT) is used extensively for the extraction of frequency domain parameters because it provides numerous evaluation parameters from a single data record. DFT testing is the most prevalent technique used in the ADC manufacturing community, although the sine-fit test, also described in the standard, provides meaningful data. Nearly every user requires that the ADC should meet or exceed a minimum signal-to-noise-and-distortion ratio (SINAD) limit for the application and that the nonlinearity of the ADC be well understood. Certainly, the extent to which this standard is applied will depend upon the application; hence, the procedure should be tailored for each unique characterization plan.

1.6 Manufacturer-supplied information

1.6.1 General information

Manufacturers shall supply the following general information:

a) ADC part identification

b) Physical characteristics: dimensions, packaging, pinouts

c) Power requirements

d) Environmental conditions: Reliable operating, non-operating, and specified performance temperature range; altitude limitations; humidity limits, operating and storage; vibration tolerance; and compliance with applicable electromagnetic interference specifications

e) Any special or peculiar characteristics

f) Compliance with other specifications

g) Control signal characteristics

h) Output signal characteristics

i) Pipeline delay (if any)

j) Exceptions to the above parameters where applicable

1.6.2 Minimum specifications

The manufacturer shall provide the following specifications (see Clause 3 for definitions):

a) Absolute accuracy (total unadjusted error)
b) Analog bandwidth (minimum, maximum, and/or typical)
c) Clock and digital control signals and level specifications
d) Input impedance (as applicable)
e) Input signal full-scale range with nominal reference signal levels (minimum, maximum, and/or typical)
f) Number of digitized bits (resolution)
g) Output coding format (binary, two's complement, etc.)
h) Output logic levels
i) Power dissipation (minimum, maximum, and/or typical)
j) Range of allowable sample rates (minimum, maximum)
k) Reference signal levels to be applied as required (minimum, maximum, and/or typical)
l) Supply currents (minimum, maximum, and/or typical)
m) Supply voltages (minimum, maximum, and/or typical)
n) Timing requirements for inputs and outputs

1.6.3 Additional specifications

a) Aperture delay time
b) Aperture uncertainty (aperture short-term time-base instability/jitter) (maximum, typical)
c) Common-mode rejection ratio (minimum, maximum, and/or typical)
d) Crosstalk (minimum, maximum, and/or typical)
e) Differential gain and differential phase (minimum, maximum, and/or typical)
f) Differential input impedance (as required)
g) Differential nonlinearity (minimum, maximum, and/or typical)
h) Effective number of bits (minimum, maximum, and/or typical)
i) Frequency response (minimum, maximum, and/or typical)
j) Gain error (minimum, maximum, and/or typical)
k) Harmonic distortion (minimum, maximum, and/or typical)
l) Hysteresis (maximum, typical)
m) Integral nonlinearity (minimum, maximum, and/or typical)
n) Intermodulation distortion (minimum, maximum, and/or typical)
o) Maximum common-mode signal levels (minimum, maximum, and/or typical)
p) Maximum static error (minimum, maximum, and/or typical)
q) Monotonicity
r) No missing codes resolution (minimum, maximum, and/or typical)

s) Random noise (maximum, typical)

t) Noise power ratio (NPR) (minimum, maximum, and/or typical)

u) Offset error (minimum, maximum, and/or typical)

v) Out-of-range/Overdrive recovery time (voltage or current)

w) Overshoot and precursors

x) Settling time (minimum, maximum, and/or typical)

y) Signal-to-noise ratio (minimum, maximum, and/or typical)

z) Slew rate limit (minimum, maximum, and/or typical)

aa) Spurious-free dynamic range (minimum, maximum, and/or typical)

bb) Transition duration of step response (rise time) (minimum, maximum, and/or typical)

cc) Word error rate

1.6.4 Pertinent ADC parameters

Table 1 is presented as a guide for many of the most common ADC applications. The wide range of ADC applications makes a comprehensive listing impossible. This table is intended to be a helpful starting point for users to apply this standard to their particular applications.

Table 1—Critical ADC parameters

Typical applications	Critical ADC parameters	Performance issues
Audio	SINAD, THD, noise	— Power consumption — Crosstalk and gain matching
Automatic control	Monotonicity Short-term settling, long-term stability, noise	— Transfer function — Crosstalk and gain matching — Temperature stability
Data acquisition	DNL, INL, gain, offset, noise, out-of-range recovery, settling time, full-scale step response, channel-to-channel crosstalk	— Channel-to-channel interaction — Accuracy, traceability (Sol Max)
Digital oscilloscope/waveform recorder	SINAD, ENOB, noise Bandwidth Out-of-range recovery Word error rate	— SINAD for wide bandwidth amplitude resolution — Low thermal noise for repeatability — Bit error rate
Geophysical	THD, SINAD, long-term stability, noise	— Millihertz response
Imaging	DNL, INL, SINAD, ENOB, noise Out-of-range recovery Full-scale step response	— DNL for sharp-edge detection — High-resolution at switching rate — Recovery from blooming
Radar and sonar	SINAD, IMD, ENOB SFDR Out-of-range recovery, noise	— SINAD and IMD for clutter cancellation and Doppler processing
Spectrum analysis	SINAD, ENOB SFDR, noise	— SINAD and SFDR for high linear dynamic range measurements
Spread spectrum communication	SINAD, IMD, ENOB SFDR, NPR Noise-to-distortion ratio, noise	— IMD for quantization of small signals in a strong interference environment — SFDR for spatial filtering — NPR for interchannel crosstalk
Telecommunication personal communications	SINAD, NPR, SFDR, IMD Bit error rate Word error rate, noise	— Wide input bandwidth channel bank — Interchannel crosstalk — Compression — Power consumption
Video	DNL, SINAD, SFDR, DG, DP, noise	— Differential gain and phase errors — Frequency response
Wideband digital receivers SIGINT, ELINT, COMINT	SFDR, IMD SINAD, noise	— Linear dynamic range for detection of low-level signals in a strong interference environment — Sampling frequency

COMINT = communications intelligence
DG = differential gain error
DNL = differential nonlinearity
DP = differential phase error
ELINT = electronic intelligence
ENOB = effective number of bits
IMD = intermodulation distortion
INL = integral nonlinearity
NPR = noise power ratio
SFDR = spurious free dynamic range
SIGINT = signal intelligence
SINAD = signal-to-noise-and-distortion ratio
THD = total harmonic distortion

2. Normative references

The following referenced documents are indispensable for the application of this document (i.e., they must be understood and used, so each referenced document is cited in text and its relationship to this document is explained). For dated references, only the edition cited applies. For undated referenced, the latest edition of the referenced document (including any amendments or corrigenda) applies.

IEEE Std 181™-2003, IEEE Standard on Transitions, Pulses, and Related Waveforms.[1, 2]

3. Definitions and symbols

For the purposes of this standard, the following terms and definitions apply. *The IEEE Standards Dictionary: Glossary of Terms & Definitions* should be referenced for terms not defined in this clause.[3]

3.1 Definitions

ac-coupled analog-to-digital converter: An analog-to-digital converter utilizing a network that passes only the varying ac portion, not the static dc portion, of the analog input signal to the quantizer.

alternation band: The range of input levels which causes the converter output to alternate between two adjacent codes. A property of some analog-to-digital converters, it is the complement of the hysteresis property.

analog-to-digital converter (ADC): A device that converts a continuous time signal into a discrete-time discrete-amplitude signal.

aperture: The interval during which the input to the ADC affects the output or the weighting function that determines the sampled output from the input signal.

aperture delay: The delay from a threshold crossing of the analog-to-digital converter clock which causes a sample of the analog input to be taken to the center of the aperture for that sample.

aperture duration (*p*%): The [50 − (*p*/2)]% to [50 + (*p*/2)]% transition duration of the step response of the ADC. If ringing of the step response causes multiple crossings of either of the levels, the *p*% aperture duration is the time from the first crossing of the first level to the last crossing of the second level. Note: The significance is that the output of the ADC is determined, with an error of (100 − *p*)% or less, by the input signal in an interval of this duration. Common values of *p* are 50, 80, and 99.9. For *p* = 80, this is the 10% to 90% transition duration of the step response.

aperture jitter: *See:* **aperture uncertainty**.

aperture uncertainty: The standard deviation of the apparent sampling time. *Syn:* **aperture jitter; timing jitter; timing phase noise**.

asynchronous sampling: Refers to sampling an input signal that is not phase locked to the analog-to-digital converter sampling frequency.

[1] The IEEE standards or products referred to in this clause are trademarks of the Institute of Electrical and Electronics Engineers, Inc.
[2] IEEE publications are available from the Institute of Electrical and Electronics Engineers, 445 Hoes Lane, Piscataway, NJ 08854, USA (http://standards.ieee.org/).
[3] *The IEEE Standards Dictionary: Glossary of Terms & Definitions* is available at http://shop.ieee.org/.

clock signal duty cycle: The fraction of the time the clock signal spends in excess of the logic threshold.

clock signal slew rate: The time derivative of the clock signal at the point where it crosses the logic threshold.

code bin *k*: A digital output that corresponds to a particular set of input values.

code bin width (*W*[*k*]): The difference of the code transition levels, T[k + 1] and T[k], that delimit the k^{th} bin:

$$W[k] = T[k+1] - T[k] \tag{1}$$

code transition level: The boundary between two adjacent code bins.

code transition level (*T*[*k*]): The value of the converter-input parameter at the transition point between two given adjacent code bins. The transition point is defined as the input value that causes 50% of the output codes to be greater than or equal to the upper code of the transition, and 50% to be less than the upper code of the transition. The transition level $T[k]$ lies between code bin $k - 1$ and code bin k.

coherent sampling: Sampling of a periodic waveform such that there is an integer number of waveform cycles in the data record. Coherent sampling occurs when the following relationship exists:

$$Mf_i = Jf_s \tag{2}$$

where

f_s is the sampling frequency
J is the integer number of cycles of the waveform in the data record
f_i is the frequency of the input waveform
M is the number of samples in the data record

common-mode out-of-range: A signal level whose magnitude is less than the specified maximum common-mode signal but greater than the maximum operating common-mode signal.

common-mode out-of-range recovery time: The time required for the analog-to-digital converter under test to return to its specified characteristics after the end of a common-mode out-of-range input signal.

common-mode range: The range of analog input signal swing at each differential input over which the common-mode rejection is specified. Common-mode range is also the sum of the largest simultaneously applied common-mode signal and differential signal.

common-mode rejection ratio (CMRR): The ratio of the input common-mode signal to the effect produced at the output of the analog-to-digital converter under test, in units of the input signal.

common-mode signal: The average value of the signals at the positive input and the negative input of a differential-input analog-to-digital converter. If the signal at the positive input is designated V_+, and the signal at the negative input is designated V_-, then the common-mode signal V_{cm} is:

$$V_{cm} = \frac{V_+ + V_-}{2} \tag{3}$$

conversion (clock) rate (f_s): The frequency at which digital output words are provided by the analog-to-digital converter on its output.

crosstalk: Undesired energy appearing in a signal as a result of coupling from other signals.

data valid time: A measure of the time, in analog-to-digital converter clock cycles, between the first clock transition after the data becomes valid at the digital outputs and the last clock transition before it becomes invalid.

differential-input impedance to ground: For a differential-input analog-to-digital converter, the impedance between the positive input and the negative input.

differential nonlinearity (DNL): The difference between a specified code bin width and the average code bin width, divided by the average code bin width.

differential signal: The difference between the signal at the positive and negative inputs of a differential-input analog-to-digital converter. If the signal at the positive input is designated V_+, and the signal at the negative input is designated V_-, then the differential signal (V_{dm}) is

$$V_{\mathrm{dm}} = V_+ - V_- \tag{4}$$

effective number of bits (ENOB): A measure of the signal-to-noise-and-distortion ratio used to compare actual analog-to-digital converter (ADC) performance to an ideal ADC.

epoch: The duration of time corresponding to a data record. For instance, for an M-sample record acquired at the uniform sampling period T_s, the epoch is MT_s.

equivalent-time sampling: A process by which consecutive samples of a repetitive waveform are acquired and assembled from multiple repetitions of the waveform, to produce a record of samples representing a single repetition of the waveform.

fall time (t_f): The time for the desired signal to go from 90% to 10% of the transition range.

full-scale range (FSR): The difference between the most positive and most negative analog inputs of a converter's operating range. For an N-bit converter, FSR is given by:

$$\mathrm{FSR} = (2^N)(\text{ideal code width}) \tag{5}$$

in analog input units.

full-scale signal: A full-scale signal is one whose peak-to-peak amplitude spans the entire range of input values recordable by the analog-to-digital converter under test.

full width at half maximum (FWHM): The width of a distribution measured at an amplitude of one half of the maximum amplitude.

gain and offset: (A) (independently based) Gain and offset are the values by which the input values are multiplied and then to which the input values are added, respectively, to minimize the mean squared deviation from the output values. **(B) (terminal based)** Gain and offset are the values by which the input values are multiplied and then to which the input values are added, respectively, to cause the deviations from the output values to be zero at the terminal points, that is, at the first and last codes.

harmonic distortion: For a pure sine-wave input, output components at frequencies that are an integer multiple of the applied sine-wave frequency which are induced by the input sine wave.

hysteresis: The maximum difference in values of a code transition level, when the transition level is approached by a changing input signal from either side of the transition.

ideal code bin width (*Q*): The ideal full-scale input range divided by the total number of code bins.

input impedance: The impedance between the signal input of the analog-to-digital converter under test and ground.

integral nonlinearity (INL): The maximum difference between the ideal and actual code transition levels after correcting for gain and offset.

k^{th} code transition level ($T[k]$): The input value corresponding to the transition between codes $k - 1$ and k.

NOTE—See Figure 3.

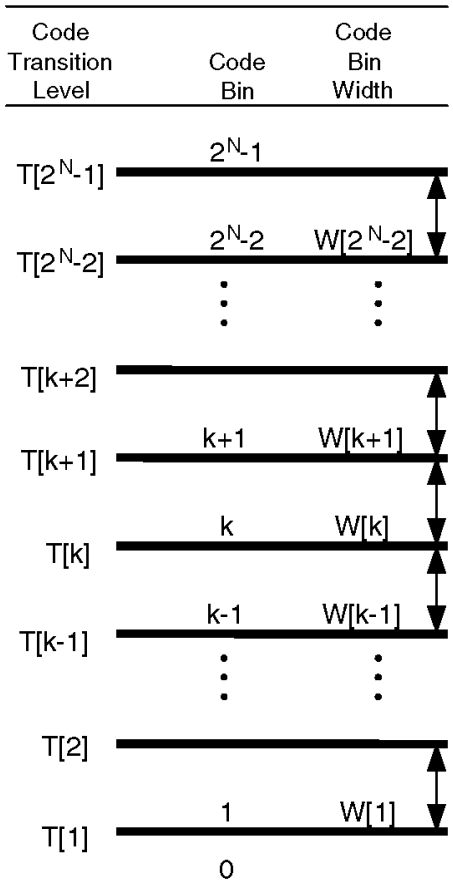

Figure 3—Definitions pertaining to input quantization

large signal: One whose peak-to-peak amplitude is as large as practical but is recorded by the instrument within, but not including, the maximum and minimum amplitude data codes. As a minimum, the signal must span at least 90% of the full-scale range of the analog-to-digital converter under test.

least significant bit (LSB): With reference to analog-to-digital converter input signal amplitude, an LSB is synonymous with one ideal code bin width.

logic level: Any level within one of two (or more) non-overlapping ranges of values, of a physical quantity, used to represent the logic.

long-term settling error: The absolute difference between the final value specified for short-term settling time and the value 1 s after the beginning of the step, expressed as a percentage of the step amplitude.

maximum common-mode signal level: The maximum level of the common-mode signal at which the common-mode rejection ratio is still valid.

maximum operating common-mode signal: The largest common-mode signal for which the analog-to-digital converter will meet its specifications when recording a simultaneously applied, normal-mode signal.

maximum input signal level: The input level beyond which damage to the device may occur.

monotonic analog-to-digital converter: An analog-to-digital converter that has output codes that do not decrease (increase) for a uniformly increasing (decreasing) input signal, disregarding random noise.

noise power ratio (NPR): The ratio of the average out-of-notch to the average in-notch power spectral density magnitudes for the discrete Fourier transform spectrum of the analog-to-digital converter output sample set.

noise (total): Any deviation between the output signal (converted to input units) and the input signal except deviations caused by linear time-invariant system response (gain and phase shift), or a dc level shift. For example, noise includes the effects of random errors (random noise), fixed pattern errors, nonlinearities (e.g., harmonic or intermodulation distortion), and aperture uncertainty. *See also:* **random noise**.

noncoherent sampling: Sampling of a waveform such that the relationship between the input frequency, sampling frequency, number of cycles in the data record, and the number of samples in the data record does not meet the definition of coherent sampling.

normal mode signal: The difference between the signal at the positive input and the negative input of a differential input analog-to-digital converter. *Syn:* differential signal.

offset: *See:* **gain and offset**.

out-of-range input: Any input whose magnitude is less than the maximum input signal of the analog-to-digital converter but greater than the full-scale range.

overshoot: The maximum amount by which the step response exceeds the high state, specified as a percentage of (recorded) pulse amplitude.

passband: The band of input signal frequencies that the analog-to-digital converter is intended to digitize with nominally constant gain.

phase nonlinearity: The deviation in phase response from a perfectly linear-phase response as a function of frequency.

pipeline delay: A measure of the latency in terms of analog-to-digital converter clock cycles between the clock transition that initiates sampling of the input and the presentation of the digitized value of that sample at the digital output.

precursor: In a step or pulse waveform, any deviation from the base state prior to the pulse transition.

probability density function (PDF): For a random variable, x, a positive real function, $f_x(x)$, which has the interpretation that $f_x(x)\,\mathrm{d}x$ is the probability that the random variable, X, lies in the interval $(x, x + \mathrm{d}x)$.

quantization: A process in which the continuous range of values of an input signal is divided into non-overlapping sub-ranges, and to each sub-range a discrete value of the output is uniquely assigned. Whenever the signal value falls within a given sub-range, the output has the corresponding discrete value. (*IEEE Standards Dictionary* [B21])

quantization error/quantization noise: The error caused by conversion of a variable having a continuous range of values to a quantized form having only discrete values, as in analog-to-digital conversion. The error is the difference between the original (analog) value and its quantized (digital) representation. (*IEEE Standards Dictionary* [B21])

random noise: A non-deterministic fluctuation in the output of an analog-to-digital converter, described by its frequency spectrum and its amplitude statistical properties. *See also:* **noise**.

record of data: A sequential collection of samples acquired by the analog-to-digital converter.

relatively prime: Describes integers whose greatest common divisor is 1.

residuals: In curve fitting, the differences between the recorded data and the fitted function.

rise time (t_r): The time for the signal to go from 10% to 90% of the transition range.

root-mean-square (rms): For a given set of data, the square root of the arithmetic mean of the squared values of each of the data.

root-sum-square (rss): For a given set of data, the square root of the sum of the squared values of each of the data.

sampling: The process of assigning discrete time values to a continuous time signal.

settling time: The time at which the step response enters and subsequently remains within a specified error band around the final value, measured from the mesial point (50%) of the response. The final value is defined to occur 1 s after the beginning of the step.

short-term settling time: Measured from the mesial point (50%) of the output, the time at which the **step response** enters and subsequently remains within a specified error band around the final value. The final value is defined to occur at a specified time less than 1 s after the beginning of the step.

signal-to-noise-and-distortion ratio (SINAD): For a pure sine-wave input of specified amplitude and frequency, the ratio of the root-mean-square (rms) amplitude of the analog-to-digital converter output signal to the rms amplitude of the output noise, where noise is defined as above to include not only random errors but also nonlinear distortion and the effects of sampling time errors.

single-ended analog-to-digital converter: A non-differential analog-to-digital converter, i.e., one that does not subtract the signals at two input terminals. Such a converter may add multiple inputs.

slew limit: The value of output transition rate of change for which an increased amplitude input step signal causes no change.

small signal: A signal whose peak-to-peak amplitude spans no more than 10% of the full range of the analog-to-digital converter under test.

signal-to-noise-ratio (SNR): For a pure sine-wave input of specified amplitude and frequency, the ratio of the root-mean-square (rms) amplitude of the analog-to-digital converter output signal to the rms amplitude of the output noise, this does not include the harmonic distortion components that are used for the estimate of THD. Note: This was called signal-to-non-harmonic ratio (SNHR) in the previous version of this standard.

spurious components: Persistent sine waves at frequencies other than the harmonic frequencies. *See:* **harmonic distortion**.

spurious-free dynamic range (SFDR): For a pure sine-wave input of specified amplitude and frequency, the ratio of the amplitude of the analog-to-digital converter's output averaged spectral component at the input frequency, f_i, to the amplitude of the largest harmonic or spurious spectral component observed over the full Nyquist band, max $\{|X(f_h)|$ or $|X(f_s)|\}$:

$$\mathrm{SFDR(dB)} = 20\log_{10}\left(\frac{\left|X_{\mathrm{avm}}(f_i)\right|}{\max\limits_{f_s,f_h}\left\{\left|X_{\mathrm{avm}}(f_h)\right| \text{ or } \left|X_{\mathrm{avm}}(f_s)\right|\right\}}\right) \tag{6}$$

where

X_{avm} is the averaged spectrum of the ADC output
f_i is the input signal frequency
f_h and f_s are the frequencies of the set of harmonic and spurious spectral components

step (or pulse) base state: The magnitude reference line at the base magnitude [IEEE Std 181-2003].

step (or pulse) high state: The magnitude reference line at the top magnitude [IEEE Std 181-2003].

step response: The recorded output response for an ideal input step with designated base state and high state.

synchronous sampling: Refers to sampling an input signal that has been phase locked to the analog-to-digital converter sampling frequency.

timing jitter: *See:* **aperture uncertainty**.

timing phase noise: *See:* **aperture uncertainty**.

total harmonic distortion (THD): For a pure sine-wave input of specified amplitude and frequency, the root-sum-of-squares (rss) of all the harmonic distortion components including their aliases in the spectral output of the analog-to-digital converter. Unless otherwise specified, THD is estimated by the rss of the second through the tenth harmonics, inclusive. THD is often expressed as a decibel ratio with respect to the root-mean-square amplitude of the output component at the input frequency.

total spurious distortion (TSD): For a pure sine-wave input of specified amplitude and frequency, the root-sum-square of the spurious components in the spectral output of the analog-to-digital converter. TSD is often expressed as a decibel ratio with respect to the root-mean-square amplitude of the output component at the input frequency.

transfer curve: The representation of the average digital output code of an analog-to-digital converter as a function of the input signal value.

transition duration of a step response: The duration between the 10% point and the 90% point on the recorded step response transition, for an ideal input step with designated base state and high state.

useful power bandwidth: The large signal analog input frequency at which a record of the analog-to-digital converter's output data will be degraded by less than a specified amount.

voltage standing wave ratio (VSWR): The ratio of the mismatch between the actual impedance and the desired or nominal impedance.

window: A set of coefficients with which corresponding samples in a data record are multiplied so as to more accurately estimate certain properties of the signal, particularly frequency domain properties. Generally the coefficient values increase smoothly toward the center of the record.

word error rate: The probability of receiving an erroneous code for an input after correction is made for gain, offset, and linearity errors, and a specified allowance is made for random noise. Typical causes of word errors are metastability and timing jitter of comparators within the analog-to-digital converter.

3.2 Symbols and acronyms

ε	error, used for total error and error band
ε_{rms}	root-mean-square value of ε
$\varepsilon[k]$	difference between $T[k]$ and ideal value of $T[k]$ computed from G and V_{os}
θ	phase, expressed as radians
π	ratio of the circumference of a circle to the diameter (constant)
ρ	reflection coefficient
σ	standard deviation; sometimes used as noise amplitude, which is the standard deviation of the random component of a signal
σ_σ	standard deviation of the standard deviation (e.g., standard deviation of the noise amplitude)
σ_t	aperture uncertainty
σ^2	variance; sometimes used to describe random noise power
τ	sampling period, the inverse of f_s
ω	$2\pi f$ = angular frequency, expressed in radians per second
ω_i	angular input frequency expressed in radians per second
a	general purpose real number
A	sinusoidal amplitude
B	test tolerance in fractions of a least significant bit (Q); also used as an amplitude
C	offset
CMR	common-mode range
D	general purpose integer
dBc	decibels referenced to the output fundamental frequency level
dBFS	decibels referenced to a full-scale sine wave
DNL	maximum differential nonlinearity over all k
$d[n]$	dither component of output sample $y[n]$
DNL[k]	differential nonlinearity of code k
$d_{est}[n]$	estimate of the dither component $d[n]$
$E_G(f)$	gain flatness error of frequency f
ENBW	equivalent noise bandwidth
ENOB	effective number of bits
$e_m[f]$	aliasing and first differencing magnitude errors
$e_p[f]$	aliasing and first differencing phase errors
FSR	full-scale range
f	frequency

$f(n)$	sine-wave component of output sample $y(n)$
f_{co}	upper frequency for which the amplitude response is −3 dB
f_d	sampling frequency of a record after decimation by an integer
f_{eq}	equivalent sampling rate
f_h	frequency of a harmonic of the input frequency
f_i	actual input frequency or approximate desired input frequency
f_{imf}	frequency of intermodulation distortion products
f_m	frequency of the mth component of a magnitude spectrum produced by a DFT
f_{opt}	optimum input frequency for testing
f_r	input signal reference frequency or input signal repetition rate
f_s	sampling frequency
f_{sp}	frequency of a persistent spurious tone
G	static gain of the ADC under test
$G(f)$	dynamic gain of the ADC under test, as a function of frequency f
H	average number of histogram samples received in two code bins sharing the same transition level
$H(f)$	frequency response of the ADC under test, as a function of frequency f
$H(f_k)$	DFT of $h(n)$
$H(i)$	number of histogram samples in bin i
$H_c(j)$	number in the jth bin of the cumulative histogram of samples
$h(n)$	discrete time impulse response of a system
INL	integral nonlinearity
INL(k)	integral nonlinearity at output code k
i	general purpose index
J	number of cycles in a record
k	code bin
L	general purpose integer
l	general purpose factor
M	number of sequential samples in a record
$M_+(x)$, $M_-(x)$	number of measurements of the output value at the input value x for increasing and decreasing inputs respectively
M_D	number of samples in one period of pseudorandom dither
M_d	number of samples in a record after decimation
mod	the standard function for which mod(I,J) is the remainder when I is divided by J (e.g., mod(12,10) = 2).
mse	mean square error
N	number of digitized bits. (Note that for certain ADCs, N may not be an integer value.)
NPR	noise power ratio
n	sample index within a record
p	probability
Q	ideal code bin width, expressed in input units
R	error parameter; also used as minimum number of records required
r	general purpose integer
S	set of samples collected over more than one record, also used as an error parameter or as total number of samples used in a histogram
$T[k]$	code transition level between code $k-1$ and code k
THD	estimate of total harmonic distortion
t_{eq}	average equivalent-time sampling period

t_f	top to base transition time; fall time
t_n	discrete sample times
t_r	base to top transition time; rise time
t_{wc}	the center point of the aperture time associated with an output sample
u	confidence level expressed as a fraction
V_{cm}	common-mode signal
V_{dm}	differential mode signal
V_o	input signal overdrive; the amount by which an input signal exceeds the ADC full-scale range
V_{os}	input offset of ADC, ideally = 0
VSWR	voltage standing wave ratio
$W[k]$	code bin width of code bin k
w	estimated word error rate
w'	worst-case word error rate
$w(n)$	window function coefficient (for a DFT)
X	number of standard deviations of a Gaussian distribution
$X_{avm}(f_m)$	the averaged magnitude spectral component at discrete frequency f_m after a DFT
x	ADC input signal value; or number of errors detected
$Y[k]$	the M-point discrete Fourier transform of the M-sample record $y[n]$
$y[n] = y_n$	the nth output data sample within a record
$y[n]$	average of y_n over M samples
Z_0	transmission line impedance
Z_t	ADC input impedance
$Z_{u/2}$	number of standard deviations that encompass $100(1 - u)$% of a Gaussian distribution about the center

4. General test methods

4.1 Introductory information on test methods

It is assumed that the user/manufacturer has defined operating limits for the device under test. These limits are categorized as *absolute* (the limit beyond which the device will be destroyed) and *operating* (the limit beyond which the device will not operate properly). These limits will vary from device to device, depending on the design. It is not the intention of this document to describe the method of setting these limits, only to verify the operation within them. All test procedures described herein apply only to parameters of a device that is operated within its specified limits.

The type of circuitry used to capture the digital data samples produced by the ADC is determined largely by the data rate. Slower ADCs may be interfaced directly to the computer. ADCs often require a buffer memory to acquire data at the ADC sample rate and then download accumulated samples to the computer at a slower rate. Even faster ADCs may require latches and/or de-multiplexers between the ADC and the buffer memory, and perhaps data decimation, as described in 4.2.2. A logic analyzer might be used as a buffer memory to capture data for some tests.

4.2 Test setup

A few general test setups can be used to perform most of the ADC tests presented in this standard. Test setups that use sine waves, arbitrary waveforms, and pulses are described in the following subclauses.

Some tests, such as those for VSWR and out-of-range signals, require setups other than those discussed in the following subclauses.

4.2.1 Sine-wave test setup

Figure 4 shows the sine-wave test setup. Sine waves are commonly used in ADC testing because appropriate sine-wave sources are readily available and because it is relatively easy to establish the quality of the sine wave (e.g., with a spectrum analyzer). A sine-wave generator provides the test signal while a clock generator provides the clock (or conversion) signal. Also, combining the output of two (or three) sine-wave generators can produce two-tone (or three-tone) test signals for intermodulation distortion testing. Additionally, a noise generator's output can be combined with a signal to provide low-level dither (Gray and Stockham [B17]).

If frequency synthesizers are used to generate the test and clock signals, the synthesizers can often be phase-locked to maintain precise phase relationships between the signal and the sampling clock. Phase-locking of synthesizers facilitates testing and simplifies subsequent digital signal processing, by preventing clock/signal walkthrough (beat patterns) that may artificially increase or reduce measured spurious output.

Both the clock and the test signals must be suitable for the test being performed. Filters may be required in either the clock or signal paths to reduce noise or harmonic distortion. For example, sub-harmonics in the clock path will degrade ADC performance, so the clock signal may require filtering to smooth edges which might otherwise feed through to the signal path. Also, low-pass or band-pass filters may be required in the signal path to eliminate noise or other undesirable signals (e.g., harmonics). The relative jitter between the clock and test signals must be low enough to prevent jitter artifacts from affecting the measured noise floor as described in 9.4.4.

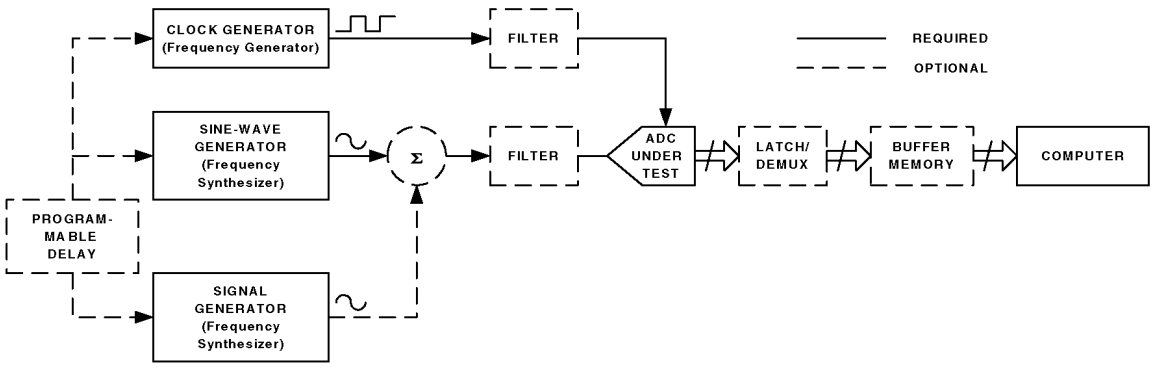

Figure 4—Setup for sine-wave testing

4.2.2 Arbitrary signal test setup

The arbitrary waveform test setup of Figure 5 can be used for arbitrary test signals, such as ramps, chirps, and steps. In this setup, the test signal is generated digitally and then converted to analog. Care must be taken to quantify the performance of the digital-to-analog converter (DAC) and filter in order to assess (or remove) its impact on the measured performance of the ADC under test. See 4.2.1 for comments on filters and data capture. It may be suitable to use an arbitrary waveform generator for the DAC and filter functions for some ADCs.

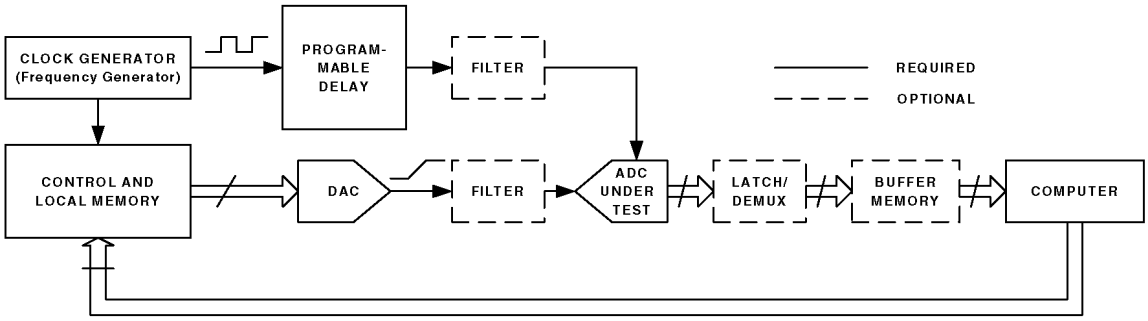

Figure 5—Setup for arbitrary signal testing

4.2.3 Step signal setup

Figure 6 shows a step waveform test setup to be used for testing with precision step signals that are not digitally generated. Precision pulses and step signals can be used to measure both time domain parameters (such as impulse response, transition duration, overshoot, and settling time) and frequency domain parameters (such as frequency response amplitude and phase, bandwidth, and gain flatness). Equivalent-time sampling can be employed, and certain data analysis tasks can be simplified, if the optional step repetition generator is phase locked to the sampling clock. See 4.2.1 for comments on data capture and on filters. Careful attention must be paid to the phase linearity of any filters placed in the pulse/step signal path.

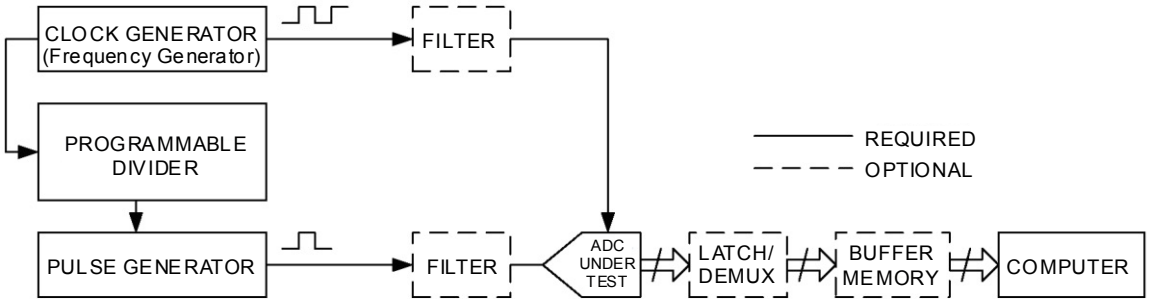

Figure 6—Setup for pulse and step signal testing

4.3 Taking a record of data

A record of data is a sequential series of samples acquired by test equipment interfaced to the ADC. The action of taking a record of data is defined as accumulating a set of samples from the ADC using the interfaced test equipment and transferring the sample set to a computer for analysis.

4.3.1 Use of output decimation in taking a record of data

For ADCs with very high sample rate, the limits of interfaced memory may make it impractical and/or uneconomical to store sequential samples from the ADC. In this case, output decimation may be used to take a record of data at a slower effective sample rate while still clocking the ADC at the high sample rate. A decimated record is a collection of every *D*th sequential sample acquired by the ADC. To do this, a divide-by-*D* counter is driven by the (high-frequency) ADC clock, and the output of the divider in turn

drives a bank of data latches (flip-flops), or triggers a logic analyzer to sample the data outputs of the ADC and pass every *D*th sample along to the memory system.

Output decimation may also be applied on tests that do not directly record the output data, for instance when recording histograms of output code occurrences or feeding back the output code to control the input signal level (see 6.2 and 6.3). The user should note that decimation involves a loss of information, which in special cases, such as hysteresis testing (see 8.7), may affect test results. When output decimation is used, the decimated sample rate, f_s/D, shall be used for any equations relating sample rate to input frequency (e.g., for equivalent-time sampling) or time measurements (e.g., step response parameters), but the actual ADC sample rate f_s shall be quoted as the sample rate in the test results.

For decimation of the output of an ADC system that has a time-interleaved architecture (i.e., L ADCs each sampling at f_s/L), the output decimation ratio D should be made relatively prime to the interleave ratio L; otherwise, the decimated output data stream will not contain data from all of the L interleaved converters.

4.4 Equivalent-time sampling and undersampling

The maximum sampling rate of a converter limits its useful measurement bandwidth. Furthermore, if the sampling rate is not at least twice the frequency of the highest frequency component of the input signal, then aliasing errors can result. If the input signal is repetitive, equivalent-time sampling can reduce these limitations. Equivalent-time sampling is a process by which consecutive samples of a repetitive waveform are acquired and assembled from multiple repetitions of the waveform, to produce a record of samples representing a single repetition of the waveform.

Several methods of equivalent-time sampling exist:

a) Apply a known delay between the input signal and the converter's time base

b) Independently measure the relative delay between the signal and the (time base) sample commands

c) Extract equivalent time samples from a single record using the converter's internal time base, provided that the input signal's repetition rate is selected appropriately

The method in 4.4.1 implements the extraction method.

4.4.1 Extraction method

The method below shows how to use equivalent-time sampling to increase the effective sampling rate by an integer factor D. By appropriate choice of the repetition rate of the input signal (f_i), D periods of the input waveform are recorded in a single record; then, upon rearranging the samples with a simple algorithm, a single period of the input signal is obtained which is effectively sampled at D times the real-time sampling rate. This is illustrated in Figure 7 for $D = 4$. To implement this method, choose integer D based on the required equivalent sample rate, f_{eq}, such that $f_{eq} = Df_s$, where f_s is the ADC sampling frequency. Next, L, the number of real-time samples taken during each repetition of the input waveform, is given by $L = \text{INT}(M/D)$, i.e., the integer value of M/D, where M is the number of samples in a record. Finally, the input signals repetition rate, f_i, is set [as shown in Equation (7)] such that

$$f_i = f_S \frac{D}{LD-1} \text{ with } LD \le M \tag{7}$$

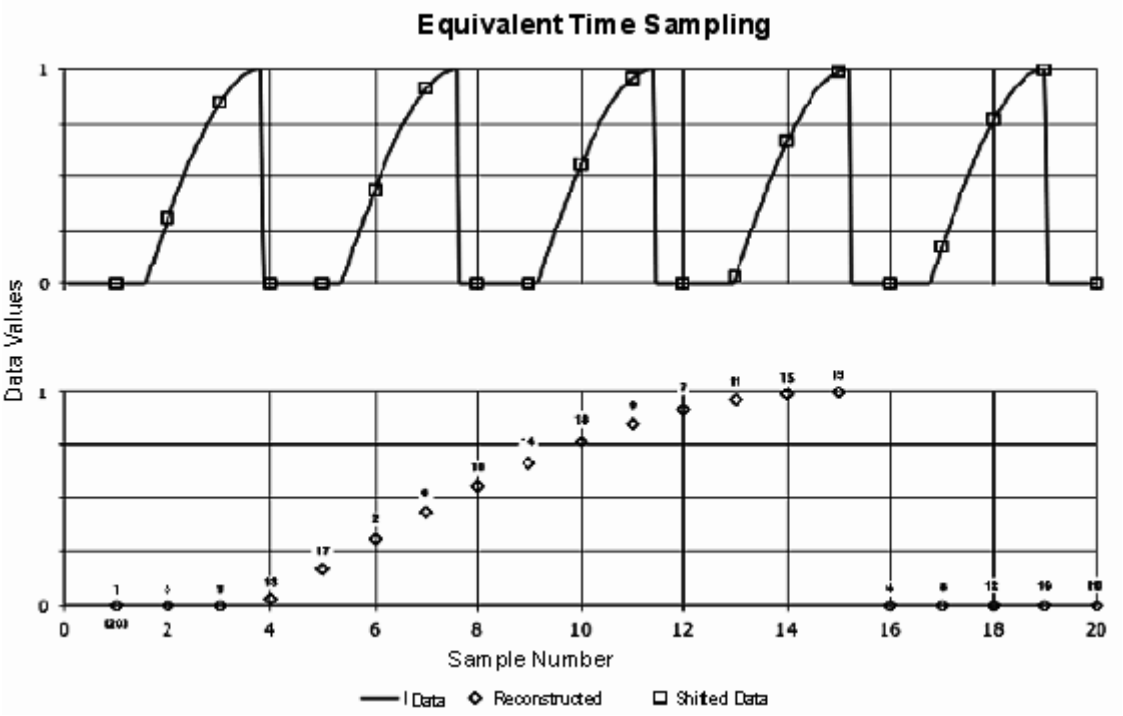

Figure 7—Equivalent-time sampling extraction

For the example of Figure 7, $D = 4$, $L = 5$, M could be 20, 21, 22, or 23, and $f_r = 4f_s/19$. (Note that data points between LD and M are not useable). The following pseudo-code implements the algorithm to rearrange the samples in equivalent time:

```
LET F = 0
FOR I = 1 TO L
        FOR J = 1 TO D
                F = F + 1
                I2 = I + (J − 1)*L
                E(F) = R(I2)
        NEXT J
NEXT I
```

where

D is the sample rate multiplier
M is the record length
L is the INT(M/D) where INT(*) designates the integer part of *
E(*) is the array containing LD equivalent-time samples
F is the equivalent-time sampling index
R(*) is the array containing real-time samples
I2 is the real-time sampling index

4.4.2 Comments on the extraction method for equivalent-time sampling

This method of achieving higher equivalent sampling rates requires that the repetition rate, f_i, of the input signal be precisely controlled. While the average equivalent-time sample rate is just Df_s, independent of f_i, the relative spacing of the equivalent-time samples becomes non-uniform when f_i deviates from the value

given by Equation (7). If f_i is too great, $D - 1$ out of D successive samples will occur too late while one sample will be correctly placed; if f_i is too small, $D - 1$ samples will occur too soon. In either case, the maximum sampling time error (Δt_{eq}) is given, to a good approximation, by Equation (8).

$$\Delta t_{eq} \approx \frac{M(D-1)}{Df_s} \frac{\Delta t_i}{t_i} \tag{8}$$

where

t_{eq} is the average equivalent-time sampling period, i.e., $1/(D \bullet f_s)$
Δt_{eq} is the maximum sampling time offset
$\Delta t_i/t_i$ is the proportional error in the repetition period (or repetition rate)

Note, however, that the errors are not cumulative; the average equivalent-time sampling period is still given by $1/Df_s$.

Of course, the assumption is made in Equation (7) that f_s is exactly known; if it is not exactly known, then the additional error given by an expression similar to Equation (8) will accrue. As an example, if $D = 4$, $M = 1024$, and the equivalent sampling period is known to 5%, then the repetition rate must be set, and the sampling rate must be known, each with an accuracy of $0.05/(1024 \times 3) = 16$ parts in 10^6.

To achieve such accuracy it is usually necessary to use a frequency-synthesized source. It may sometimes be necessary to measure the frequency of the input signal as well as the frequency of the ADC's clock generator with an accurate frequency counter to ensure that they are set with sufficient accuracy. If sufficient accuracy cannot be guaranteed for a specific record length, the accuracy might be improved by decreasing the record length. However, since the lowest frequency component that is represented in a record of length M is given by f_{eq}/M, this limits the range of frequencies that can be represented.

4.4.3 Alternate extraction method

This alternate method of extracting an equivalent time record is based upon a simple keep-one-of-each-D-samples decimation of an input record. While the decimation operation could be performed in either hardware or software, the method as outlined here implicitly assumes that the decimation is done in software.

At a sampling frequency, f_s, it is possible to adjust the input frequency, f_i, of a repetitive waveform to obtain exactly one cycle in a record of length M, as shown in Equation (9).

$$f_i = nf_s + \frac{f_s}{M} \qquad n = 0,1,2,3,... \tag{9}$$

Intentional aliasing of the input occurs when n is greater than or equal to one. The usual concern about not being able to tell which alias is observed does not apply in the practical case, as the input signal is known a priori.

When an ADC output data record is decimated by the factor D, the decimated output is sampled at a sampling frequency f_d; when a record of length M is decimated by keeping only one of every D samples it will be of length M_d, as shown in Equation (10) and Equation (11).

$$f_d = \frac{f_s}{D} \tag{10}$$

$$M_d = \frac{M}{D} \tag{11}$$

When f_d is substituted for f_s and M_d for M in Equation (9), and using the relation $f_d/M_d = f_s/M$, then one cycle in a record of length M_d results from the input frequency chosen using Equation (12).

$$f_i = nf_d + \frac{f_s}{M} \qquad n = 0,1,2,3,... \tag{12}$$

When a decimated record of length M_d contains a single cycle of f_i, the equivalent sampling rate, f_{eq}, is the number of points in the decimated record divided by the period of the single cycle. Thus f_{eq} is equivalent to the frequency f_i times the length M_d, as shown in Equation (13).

$$f_{eq} = f_i M_d \tag{13}$$

This method uses input frequencies determined by Equation (12) with n chosen to be an appropriate integer in order to yield an equivalent time record of M_d samples at various frequencies of interest. The resulting records contain a single cycle in a record of length M_d. Each increment of n yields a proportional increase in f_{eq}. If fast Fourier transform (FFT) testing is to be performed on the decimated record, then some integer number of cycles greater than one is usually desired. The input frequency, modified to provide exactly J cycles per record, is given by Equation (14):

$$f_i = nf_d + \frac{Jf_s}{M} \qquad n = 0,1,2,3,... \tag{14}$$

where J is the number of cycles in record of length M_d

and the equivalent sampling rate f_{eq} is found by Equation (15):

$$f_{eq} = \frac{f_i M_d}{J} \tag{15}$$

Thus increasing J simultaneously increases the number of cycles in the decimated record and decreases the effective sampling rate. The parameter J is not necessarily an integer; it could be adjusted to give a non-integer number of cycles and exact integer equivalent sampling rate increases if that were desired.

4.4.4 Comments on alternate extraction method

This method uses simple keep-one-of-D samples decimation of the ADC output data stream to produce the equivalent-time data samples. Figure 8 depicts this operation. For D = 2, one ADC output sample is discarded, while the next is retained in the decimated record. If the decimation is done in software, then this method requires acquiring D times more ADC output samples than will be used. However, this method does not require complicated data reshuffling. Also, small frequency errors result in small errors in the equivalent sampling rate increase, but do not result in apparent systematic jitter of the equivalent-time output samples. This method is useful when an ADC is to be characterized for operation near a given frequency. It does not yield exact integer times increase in effective sampling frequency.

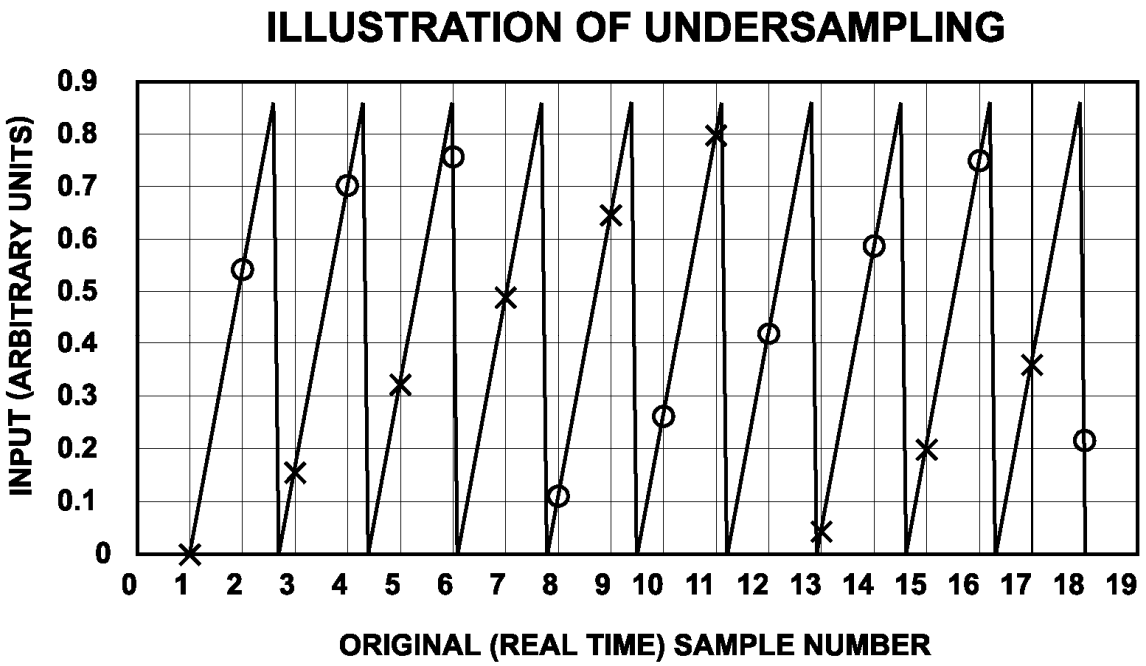

Figure 8—Continuous time waveform showing samples to be kept after decimation-by-2; samples marked "×" are kept; those marked "o" are discarded

This alternate extraction is a simple method for generating equivalent time records of arbitrary numbers of cycles per record which can utilize the same set of analog input frequencies required for coherent FFT-based tests. See 3.1 and 12.2 for coherent FFT input frequency discussions.

Specifically, when the integer M is a power of 2 and the integer D is also chosen to be a power of two, then the input frequencies chosen by Equation (14) will also be optimum input frequencies as described by 5.4.1, when the parameter J is chosen to be an odd integer. It is thus possible to generate a coherent DFT of both the real-time signal and the equivalent-time signal from the same length M data record if this alternate extraction method is used with these restrictions upon M, D, and J.

As in any under-sampled application, small errors due to inclusion of undesired signals from the unwanted base band, or other aliases, may pollute the signal. For narrowband or sine signals, the use of a band-pass filter at the ADC input minimizes any aliased artifacts.

One application of equivalent-time sampling is reconstructing an equivalent-time record to readily observe ADC anomalies, especially dynamic anomalies that become apparent only at input frequencies approaching or above the Nyquist frequency of the ADC. The equivalent-time record can also be used to measure settling times when a non-sinusoidal periodic waveform is used as the input signal. When the equivalent-sampling frequency is chosen to be much greater than the original sampling frequency the equivalent-time record can be used to measure settling times with much greater resolution. This is especially useful when the expected settling time is actually a fraction of one real-time sampling period.

FFTs of the equivalent-time signal are normalized to the equivalent-sampling rate. This can be used to *unscramble* aliased harmonics. This eases the differentiation of harmonics from any spurious tones.

5. Sine-wave testing and fitting

5.1 Introductory information on sine-wave testing and fitting

There are a number of tests that use sine waves as input signals, and there are often questions as to why sine waves are used to test instruments that are often intended to record transients. The primary reason is that very accurate sine-wave signals can be produced, and the accuracy of the signals can be readily tested with a spectrum analyzer. Another reason for the usefulness of sine waves as test signals is that they are eigenfunctions for linear time invariant (LTI) systems. This means that when a pure sine wave is supplied to the input of an LTI system, the output is a pure sine wave of the same frequency but with altered amplitude and phase. Since the analysis of the test results usually assumes that both the amplitude and phase are unknown, the LTI system distortion is ignored. The test results give the nonlinear and time-varying errors.

5.2 Curve fitting test method

Apply a sine-wave signal with appropriate amplitude, frequency, and spectral purity to the input of the ADC. Trigger the ADC to collect a record of data. The trigger does not have to be synchronized to the signal. For the four-parameter method, calculate the values of A_0, B_0, C_0 and f_0 that give the best fit, in the least squares sense, to the recorded signal to a function of the form

$$x[n] = A_0 \cos(2\pi f_0 t_n) + B_0 \sin(2\pi f_0 t_n) + C_0 \tag{16}$$

where

t_n is the nominal time associated with the nth data value

For the three-parameter fit, the known frequency is substituted for the parameter f_0. Subtract the fitted signal from the recorded signal to obtain the residuals.

There are many algorithms for performing the least squares curve fits above. This standard describes methods for both the three- and four-parameter fit.

The sine-wave curve fit is used in several specific test methods described later in this document. Analyze the residuals and fit parameters using methods described for the specific test being performed.

The fitted signal may be converted to the amplitude and phase representation of the function using

$$A_0 \cos(2\pi f t) + B_0 \sin(2\pi f t) = A \cos(2\pi f_0 t + \phi) \tag{17}$$

where

$$A = \sqrt{{A_0}^2 + {B_0}^2}$$
$$\phi = \arctan(B_0, A_0)$$

where

arctan is the standard inverse tangent of two arguments, called ATAN2 in many programming languages, and returns a value in the range of 0 to 2π or in the range from $-\pi$ to π

5.3 Comment on three-parameter versus four-parameter sine fit

The recommended method is the four-parameter method with a record containing at least five cycles of the signal. Even if the input frequency is accurately known a priori, the four-parameter method usually determines it to even better accuracy. The computations are more complicated with the four-parameter method. For records containing less than five cycles, the four-parameter method can underestimate harmonic distortion by compensating for it with a change in frequency. In this case, the three-parameter method may be more appropriate. When frequencies are chosen precisely as recommended for coherent sampling, with an exact integer number of cycles in a record, a DFT will give the same results as a three-parameter sine-wave fit.

5.4 Choice of frequencies and record length

Typically ADCs are tested at several different frequencies. There are several factors that enter into the selection of frequencies. These factors affect frequency selection on three different scales: fine, medium, and coarse. Although the frequency selection decisions are made in the order: coarse, medium, then fine, the criteria are described in the reverse order.

5.4.1 Fine-scale frequency selection

On the fine scale, one wants to select a frequency for which the sampled values will all be different. For example, if a frequency of 200 kHz was used to test an ADC with a sampling rate of 1 MHz, the same five different phases of the sine wave would be sampled many times. The recommended approach is to use a record length, M, and a frequency, f_i such that M uniformly distributed phases will be sampled. This is easily accomplished by choosing f_i as follows:

$$f_i = \left(\frac{J}{M} \right) f_s \tag{18}$$

where

J is an integer which is relatively prime to M
f_s is the sampling frequency
M is the record length

The condition of being relatively prime means that M and J have no common factors, i.e., their greatest common divisor is one. For the recommended frequency there are exactly J cycles in a record (called coherent sampling in this standard). If M is a power of two, then any odd value for J meets the relatively prime condition.

If the signal frequency meets the above conditions exactly, the maximum phase difference between successive sampled phases will be $2\pi/M$. The accuracy required of the signal frequency depends strongly on the frequency and on whether the frequency deviation is in the positive or negative direction from the nominal value.

For any value of J, relatively prime to M, there is a unique value, I, between 0 and $M - 1$, which satisfies $\text{mod}(IJ, M) = 1$. The number I is the multiplicative inverse (mod M) of J, and its value determines the frequency accuracy required.

For an exact input frequency, the maximum difference between successive sampled phases is $2\pi/M$. *If the frequency does not have the exact value specified in Equation (18), then the maximum magnitude of the phase difference will be larger.* If the larger value is written in the form $(1 + \rho) \times (2\pi/M)$, then the error, ε_f, in the frequency must satisfy:

$$\left|\frac{\varepsilon_f}{f_i}\right| \le \begin{cases} \dfrac{\rho f_s}{IMf_i} & \text{for } \varepsilon_f > 0 \\ \dfrac{\rho f_s}{(M-1)IMf_i} & \text{for } \varepsilon_f < 0 \end{cases} \tag{19}$$

where

ε_f is the error in input frequency
ρ $(1 + \rho)$ is the error limit factor between successively sample phases as defined above
f_i is the input frequency
f_s is the sampling frequency
I is chosen such that mod(IJ, M) = 1
M is the record length

For positive frequency deviations, I should be as small as possible. For negative frequency deviations, I should be as large as possible (the maximum value is $M - 1$). Tables are available, Blair [B9], that give values of J corresponding to small values of I and small values of $M - I$ for all power-of-two record lengths between 2^8 and 2^{20}. One can choose a frequency from these tables that is close to the desired frequency and determine the accuracy requirement and required direction of deviation from information in the table. If the tables do not contain values of J with acceptable frequencies, one can compute the value for I for an interval of acceptable values of J and select the value with the smallest values of I or $M - I$. The actual deviation of the frequency from the ideal value can be determined from the sine-wave fitting result.

For an ADC with N bits of resolution and an ideal transfer characteristic, the minimum record size that will produce a representative sample in every code bin (in the absence of random noise) is $2^N \times \pi$, when the input frequency is chosen as above. The smallest power-of-two record length is 4×2^N. To achieve one sample in each code bin with this slightly longer record length one can let $1 + \rho = 2\rho/4$, or $\rho = 0.57$.

At low input frequencies, the accuracy requirement for the input frequency becomes much less stringent. However, if a four-parameter fit is being used, the frequency should be chosen large enough that there are at least five cycles in the record. Otherwise, errors may be underestimated.

5.4.2 Medium-scale frequency selection

The rules for fine-scale frequency selection give many frequencies that optimize the spacing between the samples recorded. On the medium scale, one selects frequencies to cause errors from different sources to occur at different frequencies. For example, for an ADC with a sampling rate of 2 GSa/s if a frequency of 400 MHz were selected, third harmonic distortion would be at a frequency of 1200 MHz. Since this is above the Nyquist frequency of 1000 MHz it would be aliased down to 800 MHz. This is the same frequency as second harmonic distortion, so the two would be indistinguishable. With a frequency of 420 MHz, the second harmonic is at 840 MHz while the third harmonic aliases down to 760 MHz, allowing one to distinguish between the two.

In the coarse-scale frequency selection, one normally selects nice round numbers (e.g., 250 MHz, 500 MHz). These round numbers then have to be modified to separate the errors from different harmonics and, perhaps, from interleaving. It is important to take aliasing into account in this step. The resulting frequencies should then be modified a second time to meet the criteria for fine-scale frequency selection.

5.4.3 Coarse-scale frequency selection

Select several test frequencies spanning the range of major expected frequency components in the final-use input signal. The highest frequency signal should have at least as large of a maximum slew rate (derivative with respect to time) as the maximum slew rate of final-use input signal.

The test frequencies can be categorized as low, medium, and high. Low frequencies are low enough to not cause significant dynamic errors (e.g., frequency-dependant distortion and time jitter) in the ADC. Frequencies less than a few percent of the analog bandwidth may generally be considered low. Medium frequencies are those high enough to cause some dynamic effects, but still well below the analog bandwidth. These will generally be in the range of 10% to 30% of the analog bandwidth. High frequencies are near enough to the analog bandwidth that the amplitude roll-off is a significant factor. The test frequencies should include at least one frequency in each category.

5.4.4 Special considerations with very long record lengths

It is sometimes necessary to test an ADC with a record length much longer than that for which the frequency accuracy condition in Equation (19) can be met. This is the case when one wants to quantify errors that can only be detected with very long record lengths, e.g., drift and clock phase noise. In this case, the frequency selection should be based on a shorter base record length. The long record length should then be selected as an integer multiple of the base record length.

5.5 Selecting signal amplitudes

Signal amplitude of between 90% of full scale and 100% of full scale should be used for each frequency. This is referred to in the standard as a large amplitude signal. This amplitude is the amplitude of the signal at the input of the ADC. The frequency response of the ADC may substantially change the measured amplitude. For example, if a 90%-of-full-scale signal is supplied at the −3dB bandwidth of the ADC, the measured signal will have amplitude of only 63% of full scale. This is acceptable. If the ADC has a gain of greater than 110% at a test frequency, a 90%-of-full-scale signal will saturate the ADC. This may require reducing the amplitude of the test signal. This shall be reported in the test results.

It is useful to also measure a lower amplitude signal at each frequency. If the lower amplitude signal is obtained by leaving the oscillator amplitude unchanged and adding an attenuator, the test results can distinguish between distortion and noise in the oscillator and distortion and noise in the ADC. Typical values for the attenuator are between 6 dB and 12 dB.

5.6 Presenting sine-wave data

There are three common ways of presenting the results of sine-wave test data for visual analysis. Particular calculations on the test data are given at various other places in the standard. This subclause describes some general methods of presentation. The three presentations are as follows:

- Power spectrum of the residuals (Blair [B8])
- Modulo time (2π) plot of the residuals versus phase angle (Irons et al. [B24])
- Plot of the residuals versus. time

See Annex D for details. The illustrations in Annex D are from waveform recorders but are equally meaningful for ADCs.

5.7 Impurities of sine-wave sources

A number of tests in this standard use sine-wave sources, and the analyses of the test results assume that the signal is a pure sine wave. This subclause describes how the impurities of a sine wave are described quantitatively, how they are measured, and how to control them.

The impurities in a sine wave are identified as follows:

— Harmonic distortion

— Spurious components

— Wideband noise

— Amplitude modulation

— Phase modulation

Harmonic distortion is the presence of sinusoidal signals at frequencies that are integer multiples of the signal frequency. If the signal is periodic at the nominal frequency but its shape is not pure sinusoidal, it will have harmonic distortion. Harmonic distortion is typically specified in dBc, decibels relative to the carrier (the signal at the desired frequency). So, if the harmonic distortion is –30 dBc, it is about 3% of the signal (in voltage, not power.)

Spurious components are sinusoidal signals at frequencies that are not integer multiples of the signal frequency. They can result from extraneous signals coupling to the output of the signal source or from artifacts occurring in the ADC. They are also specified in dBc.

Wideband noise is a random signal that is spread over a large frequency range. It is measured in dBc/Hz, the power in a 1 Hz bandwidth relative to the power of the signal. This is often specified in units such as $nV/\sqrt{Hz}$, because the noise has the property that if the measuring bandwidth is multiplied by some number, R, the observed noise signal has an rms value that is multiplied by $\sqrt{R}$. Such sources of noise often have a "(1/*f*) corner." Such a signal will exhibit increasing power at lower frequencies that can cause instability in measurements. Such behavior is beyond the scope of the current standard.

When the signal is a pure sinusoid with amplitude that varies with time, it is said to have amplitude modulation. Amplitude modulation is typically expressed as *x*%, where *x* specifies the fluctuations in amplitude of the signal. Amplitude modulation adds to the spectrum of the signal a spectrum that is spread over a frequency range of $\pm BW_a$ around the signal frequency, where BW_a is the bandwidth of the fluctuations in the amplitude of the signal. Analysis of amplitude modulation usually assumes that the amplitude is changing slowly relative to the signal, in other words, that BW_a is significantly smaller than the signal frequency.

When a signal is of the form $a(t)\sin(2\pi ft + \varphi(t))$, where $\varphi(t)$ varies with time, it is said to have phase modulation. When $\varphi(t)$ varies randomly the phenomenon is called phase noise. Phase modulation has the same effect on the spectrum of the signal as amplitude modulation when the phase modulation is small, as it will likely be if it arises as an unwanted impurity. The added spectrum has a bandwidth equal to the bandwidth of the variations in φ. Phase noise is usually specified in dBc/Hz at specified offset frequencies from the carrier.

5.8 Estimating impurity problems from sine-fitting results

One approach to dealing with the problem of potential sine-wave impurities is to assume that they are negligible and proceed with sine-fitting tests. The results of the sine-fitting tests can then be used to determine potential problems with the signal source. This subclause gives guidance on how to accomplish this.

This approach requires performing the sine-fit test with two different amplitudes of the same frequency. The first is a large amplitude signal, and the second is reduced in amplitude by a factor of R from the large signal. The reduced amplitude signal must be obtained from the large amplitude signal by applying an attenuator, rather than reducing the amplitude at the signal source. Values of R between 3 and 10 are

appropriate. It is a good practice to test ADCs at two amplitudes in this way even if the signal source is known to be perfect.

One must also determine the frequency spectrum of the residuals using one of the methods given elsewhere in the standard. The accuracy with which the various distortion components can be determined improves with record length, so the longest reasonable record length should be used.

Harmonic distortion: Observe the harmonic distortion in the residuals. If it is negligible, it is reasonable to assume that the harmonic distortion in the signal source is negligible. Usually, harmonic distortion within the ADC will be a factor of R^{n-1} (for the n^{th} harmonic) or smaller relative to the signal for the low amplitude signal than for the large amplitude signal. Harmonic distortion in the signal source will remain the same relative to the signal. If the reduction factor in the harmonic distortion is less than R^{n-1}, the signal source should be tested for harmonic distortion as discussed in 5.9.

Other spurious components: This is handled in much the same way as harmonic distortion. Other spurious components in the signal source will be the same, relative to the signal, for the attenuated signal. If there are significant components that remain the same relative to the signal, the signal source should be tested. There can be components within the ADC, such as difference signals between an internal clock and the input signal, that also are proportional to the input signal.

Use of the frequency domain with long record lengths has the same effect as averaging. The noise (quantization included) is reduced by a factor of $\sqrt{M}$ relative to the harmonics and spurs, where M is the record length.

Wideband noise: Observe the wideband noise component of the residuals. If it is significantly smaller for the attenuated signal, the signal source should be tested for wideband noise.

Amplitude and phase modulation: The spectrum of the residuals will have a peak at the signal frequency if either of these is significant. The two types of modulation can be distinguished by looking at the modulo time plot (see Irons, et al. [B24]) of the residuals. Amplitude modulation will appear as random noise multiplied by a sinusoidal envelope at the frequency of, and in phase with, the input signal. Phase modulation is the same except that the envelope is 90 degrees out of phase with the signal. The observed phase modulation will be the difference between that of the signal and that of the clock of ADC. One way to discriminate between the two is to simultaneously test two ADCs that have independent clocks with the same signal. By correlating the residuals from the two ADCs, one can determine how much of the phase modulation is due to the signal and how much is due to the ADC clocks.

5.9 Measuring and controlling sine-wave impurities

The primary means of measuring input sine-wave impurities is with a spectrum analyzer. The primary means of controlling them is with filters.

Spectrum analyzer basics: A spectrum analyzer shows power in dBm on the vertical axis as a function of frequency on the horizontal axis. It is calibrated so that placing a sine wave at the input causes it to read the power of the sine wave on the vertical axis. The user sets the start and stop frequencies for the horizontal axis. There are a number of other controls. The two most important are the attenuation and the resolution bandwidth. Other controls that affect the accuracy are the vertical bandwidth and the sweep speed. The spectrum analyzer should be put in the mode in which the vertical bandwidth and the sweep speed are automatically determined from the start and stop frequencies and the resolution bandwidth. In this discussion it is assumed that the power shown on the vertical axis takes into account the attenuation setting. Since spectrum analyzers have a quite limited set of available attenuation values, it may be desirable to use an external attenuator. In this case, the user must do the calculations to take the attenuator into account. When measuring noise and harmonic distortion with a spectrum analyzer, the noise and distortion of the spectrum analyzer itself must be considered.

Harmonic distortion: If the signal is larger than the recommended input signal to the mixer of the spectrum analyzer, attenuate the signal with an external attenuator to get it within range. Set the start and stop frequencies to display the harmonics of interest. Reduce the resolution bandwidth until the displayed noise floor is at least 16 dB below the amplitude of the smallest harmonic to be measured. Determine the size (in dBm) of the desired harmonics. Now increase the internal attenuation of the spectrum analyzer by 10 dB (usually the smallest step.) If there is an accurate measurement of the signal, the size of the harmonic will stay the same and the noise floor will rise 10 dB. In this case, the measurement is good. If the amplitude of the harmonic changes, then the spectrum analyzer's harmonic distortion is contributing to the observed error. In this case, the process of decreasing the resolution bandwidth and increasing the attenuation must be repeated until the size of the harmonic does not change. If this cannot be accomplished, a better spectrum analyzer is needed, with either lower resolution bandwidth or lower noise.

Filtering the signal with band-pass and/or low-pass filters can reduce harmonic distortion.

Other spurious components: The measurement here is the same as for harmonic distortion, but the spectrum analyzer is not likely to contribute. As for harmonic distortion, filtering the signal with band-pass and/or low-pass filters can reduce spurious components.

Wideband noise: For this it is best to have the signal as large as possible without damaging the spectrum analyzer. Use the minimum of attenuation and, perhaps, a low noise amplifier. The setting of the resolution bandwidth is not critical. Fluctuations in the noise floor can be reduced by averaging or by reducing the video bandwidth. Measure the height of the noise floor both with the signal connected and with the signal source replaced by a terminator. If these are the same, then the spectrum analyzer does not have low enough noise for this measurement. Convert each measurement from dBm to mW. Subtract the value of the measurement without the signal from the measurement with the signal. Convert this difference back to dBm. Subtract 2 dB from this to correct for the fact that the spectrum analyzer is calibrated for sine waves rather than noise. Divide the result by the resolution bandwidth to get the noise in dBm/Hz. The noise can be reduced with filters.

The −2 dB correction mentioned above is the combination of corrections for two different sources of error that occur in spectrum analyzers when measuring noise. One is a –2.5 dB correction due to the fact the spectrum analyzer uses an envelope detector and a log amplifier rather than obtaining a true rms value. The other is a +0.5 dB correction due to the fact that the resolution bandwidth is given as a –3 dB bandwidth rather than as a noise bandwidth. This subject is discussed in detail in Application note 1303 by Agilent Technologies [B1].

Amplitude and phase modulation: These cannot be measured with common spectrum analyzers. They require specialized and expensive equipment. They cannot be easily reduced. The effect of both generally decreases with record length. The best control is to use short enough record lengths that these phenomena are not significant. Their effect at a given record length can be determined using the sine-fit method of measuring them given earlier.

6. Locating code transitions

6.1 Introductory information on locating code transitions

In most cases, determining the code transition levels as discussed in 1.4 can represent the transfer characteristic of an ADC. Quantitatively, a code transition level is the value of the converter input parameter that causes half of the digital output codes to be greater than or equal to, and half less than, a given output code.

Note that it is not always possible to define a unique value for a particular code transition level. For instance, feedback from the output to the input of an ADC can cause either no, or a range of, input

parameter values to cause an equal distribution of output codes on either side of a transition. See 6.5 for an alternate approach in these cases.

Once code transition levels have been measured, then all static parameters, including integral and differential nonlinearity, missing codes, gain, and offset can be computed. There are three test methods in wide use: the feedback loop, the ramp histogram, and the sine-wave histogram; they are described in the following subclauses.

6.2 Locating code transitions using a feedback loop

A widely used test method for determining transition levels is based on a feedback loop. In this method an input is applied to the ADC, the converter is triggered, and the results of the conversion are compared to a desired value. If the ADC output is below the desired value, the input is raised by a fixed amount. If the ADC output is equal to or above the desired value, the input is lowered by a fixed amount. This process is repeated until the ADC input has settled to a stable, average value.

After the loop has settled, the input value can be either measured or, if the input source is well calibrated, computed from its transfer function.

6.2.1 Test method

A block diagram is given in Figure 9a. In this diagram, an N_{DAC}-bit digital-to-analog converter (DAC) generates the feedback signal, but other implementations are possible, including the classic analog one shown in Figure 9b. For clarity, this method will be discussed in terms of a DAC-generated input. In this test, N_1 and N_2 of Figure 9a are equal and assigned the value N_0. The DAC's value is decremented or incremented by N_0 after each conversion cycle according to the result of the comparison between the ADC's output code, k, and a designated reference code, k_{in}. Once the code transition level $T[k_{in}]$ has been reached, the feedback loop causes the input signal to oscillate across this transition in steps that can be chosen to be as small as desired down to the DAC resolution. The ADC input level is calculated from the known transfer function of the DAC, or is measured by an optional voltmeter.

In an ideal noiseless ADC, the asymptotic state of the test is an alternation between the values k_{in} and $k_{in} - 1$, and the transition level is known only to an accuracy of N_0. Repeated tests with smaller values of N_0 can determine the transition level as precisely as desired.

In a real-world ADC, one with internal noise, the situation becomes more complex because the noise affects the properties of the asymptotic state of the test. Instead of a simple alternation about the desired code value, there will be a random walk about the transition level. The properties of this random walk depend on the relative values of the noise, the step size, and the code width. Choosing the optimum step size, N_0, is a tradeoff between speed of convergence and the desired accuracy. If N_0 is well chosen, this test is faster than either the histogram or the ramp techniques discussed in following clauses. The remainder of this subclause offers guidance in choosing the step size and the number of samples to be taken to achieve a desired accuracy.

In an ADC with no pipeline delay, N_0 would typically be set to the rms value of the converter noise. Converters with pipeline delay, P, would typically have N_0 set to the rms value of the noise divided by (P + 1). Because this is a statistical process, the desired accuracy and the step size being used in the feedback loop determine the number of samples in the average. In general, there will be a setup period of M_1 samples, followed by an averaging period of M_2 samples. Optimizing the procedure requires some care in choosing M_1 and M_2. Papers by Max [B39], [B40] give detailed guidance for doing this; this subclause gives rules of thumb for estimating values that work in most cases.

For the case where the step size, N_0, is greater than or equal to the noise, M_1 can be set to be eight. The initial DAC setting is assumed to differ from the true code edge by less than three times the rms value of

the input noise of the ADC. This initial setting is usually evaluated by an input adjusting routine, which initially sets a large value for N_0, and eventually reduces the size of N_0 in binary steps to the point where the appropriate step size is reached for the final settling. If the step size is less than the noise level, the setup requires additional time. The number of setup samples is inversely proportional to the ratio of step size to noise value. For example, if the step size is one-half the noise value, M_1 would be 16; for one-quarter it would be 32.

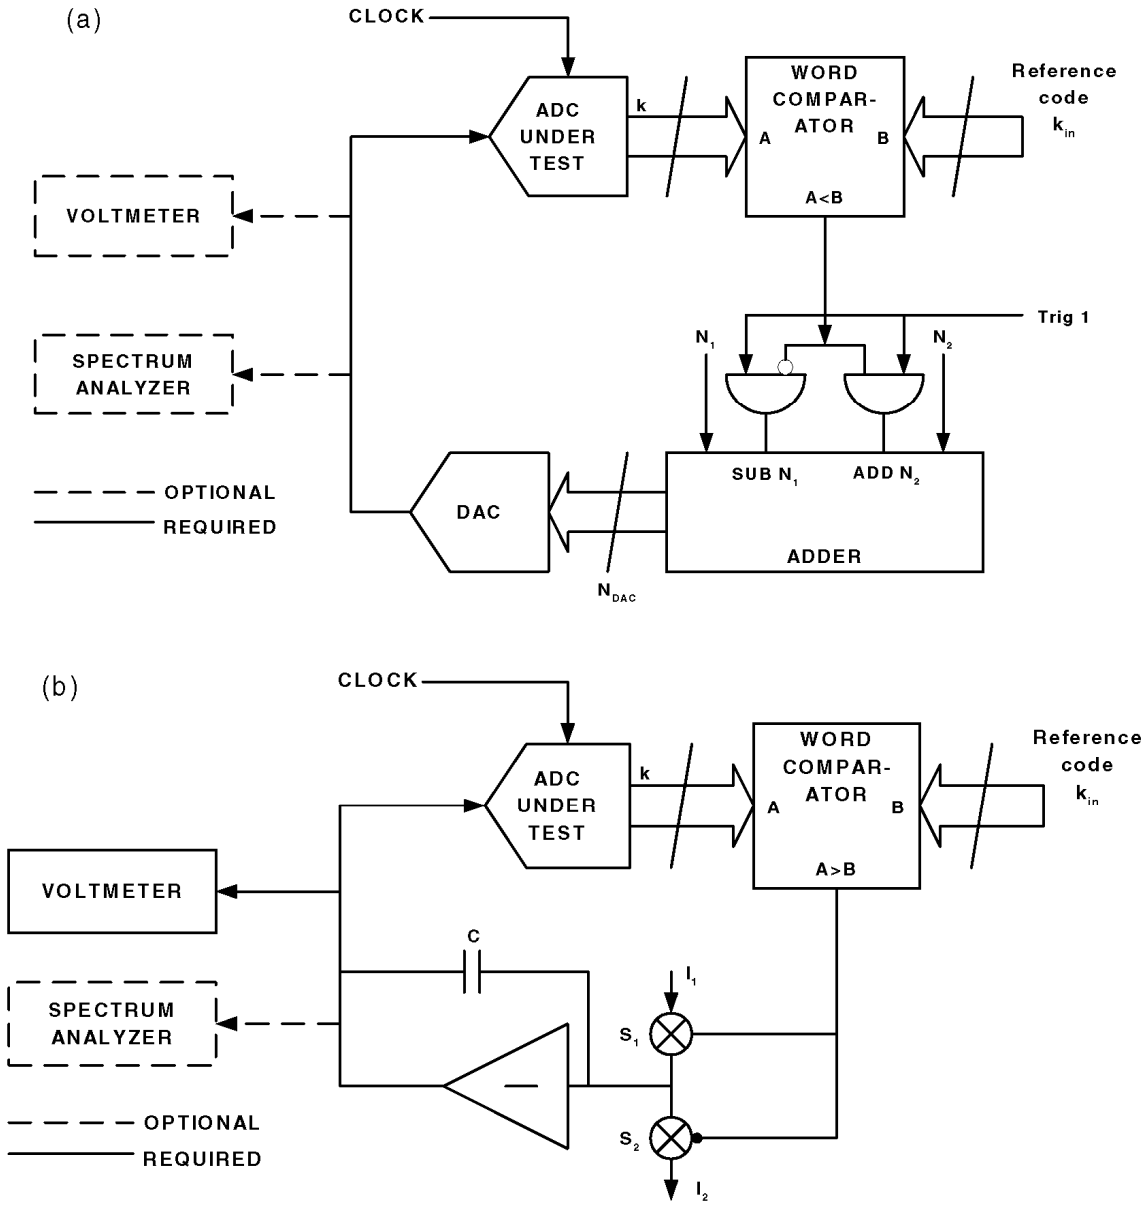

Figure 9—(a) Block diagram of feedback loop (digital method); (b) Block diagram of feedback loop (analog method)

Together, the desired accuracy of the noise measurement and the step size being used determine the choice of M_2. Figure 10 shows a plot that can be used for this purpose in most cases. For example, using this plot, one sees that at a step size corresponding to 0.5 output code widths and for a desired accuracy of 0.3 code

widths, 16 samples must be averaged. Users needing a more careful determination of M_2 are referred to Max [B39], [B40].

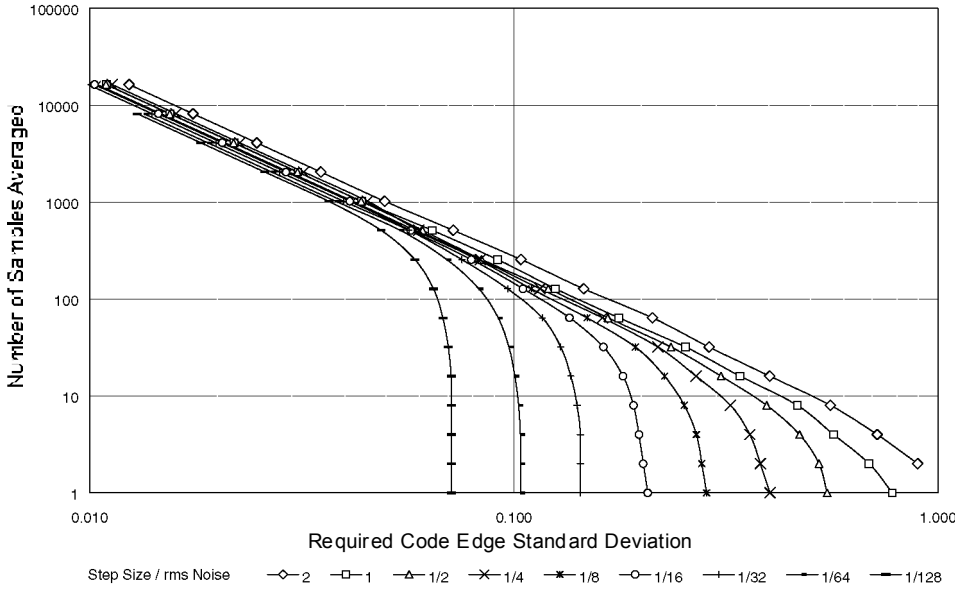

Figure 10—Number of points averaged versus required standard deviation

In principle, this technique can determine the shape of the noise at each code transition. This is beyond the scope of this standard and readers are referred to the papers by Max [B39], [B40].

6.3 Alternate code transition location method based on ramp histogram

In this approach, a histogram of code occurrences is generated in response to an input signal level which ramps linearly between the extremes of the full-scale range of the ADC. After a sufficiently large number of samples [determined from Equation (25)], the histogram of the output provides an accurate measure of the differential nonlinearity of the ADC. Integral nonlinearity can be directly computed by numerically integrating the differential nonlinearity data.

The input ramp shall be generated synchronously with the sampling clock, by a high-resolution DAC or arbitrary waveform generator with suitable linearity. Absolute signal level measurements can be made at the terminal codes to compute offset and gain errors. The statistics of this process, as noted in the comments below, can be used to calculate how many hits per bin are required to achieve a given confidence level based on the equivalent input noise level.

The location of the code transitions, $T[k]$, can be extracted by manipulating the data that is collected in a histogram test with a ramp input. The code transition levels are given by Equation (20)

$$T[k] = C + A \times H_c[k-1] \qquad \text{for } k = 1, 2, \ldots, (2^N - 1) \tag{20}$$

where

A is a gain factor
C is an offset factor

$H_c[j]$ is equal to $\sum_{i=0}^{j} H[i]$

$H[i]$ is the number of histogram samples received in code bin i

S is equal to $\sum_{i=0}^{2^N-1} H[i]$ = the total number of histogram samples

The values of C and A can be computed directly from the collected data and the direct measurement of $T[1]$ and $T[2^N - 1]$. The expressions for A and C are given by Equation (21) and Equation (22).

$$A = \frac{(T[2^N - 1] - T[1])}{(S - H[2^N - 1] - H[0])} \tag{21}$$

$$C = T[1] - \left(\frac{H[0] \times (T[2^N - 1] - T[1])}{(S - H[2^N - 1] - H[0])} \right) \tag{22}$$

It should be noted that if code bins 0 and $2^N - 1$ are excluded (defined as having zero width) then the expressions reduce to Equation (23) and Equation (24).

$$A = \frac{\left(T[2^N - 1] - T[1]\right)}{S} \tag{23}$$

$$C = T[1] \tag{24}$$

6.3.1 Comments on number of samples to be averaged per transition level for a given confidence level

The precision of the measured values of the code transition levels depends on the total number of histogram samples measured. Increasing the number of samples decreases the uncertainty while ramping the input. A larger total number of samples reduce the uncertainty. Nonlinearity of the ramp input signal would produce errors in the code transition levels. Noise on the ramp signal or the ADC under test will cause uncertainty in the measured code transition levels. Specifically, the uncertainty in LSBs due to noise in the estimate of a transition level is approximated by Equation (25).

$$\varepsilon \approx \frac{\sigma}{\sqrt{H}} \tag{25}$$

where

σ is the standard deviation of the noise, in units of ideal code bin widths (LSBs)

H is the average number of histogram samples received in each of the code bins that share the given transition level

6.3.2 Comments on ramp characteristics

The ramp method is generally used when static characteristics of the device under test are being measured. The sine-wave histogram is generally used for dynamic testing. The ramp method is more efficient in measuring the device characteristics.

6.3.3 Comments on histogram testing

Histogram methods can produce erroneous results if the device being tested has output codes that are swapped with other codes or exhibits other types of non-monotonic behavior. Such converters can produce seemingly good results, yet have large errors in the actual code transitions. To avoid these issues, converters should also be tested for SINAD performance to confirm that non-monotonic behavior is not significant.

6.4 Alternate code transition location method, based on sine-wave histogram

This method of locating code transitions is often easier to implement than the prior one, especially if one is interested only in determining nonlinearities. A pure sine wave of amplitude sufficient to slightly overdrive the ADC is input to the ADC under test. The frequency of the sine wave and the ADC sampling frequency shall be specified. Multiple records of ADC output data are taken and a histogram constructed. Selection of the sine-wave frequency, the number of samples per record, and the number of data records taken are discussed in 6.4.1 and 6.4.2. If the input range of the ADC is not symmetrical around the middle of the full-scale range, then a constant must be added to the sine wave so that the peaks of the combined signal are equidistant from the center of the full-scale range. The amount of overdrive required depends on the combined noise of the input signal and the ADC additive noise. If the amplitude and offset of the sine wave are precisely known, this method gives the transition levels to the same precision. If the amplitude of the sine wave and the offset are unknown, this method gives the transition levels to within a gain and offset error; i.e., the calculated transition levels, $T'[k]$, will be related to the true transition levels, $T[k]$, by the relation in Equation (26)

$$T'[k] = a \times T[k] + b \tag{26}$$

where a and b are constants.

This test assumes that the ADC is monotonic and has no hysteresis. See 8.6 and 8.7 for definitions of these terms.

The sine-wave frequency must be chosen as described below in 6.4.1. Note that different frequencies may produce different results, due to dynamic errors.

Take many records of data (the required amount is covered in 6.4.2) and keep track of the total number of samples received in each code bin. The transition levels are then given by Equation (27)

$$T[k] = C - A\cos\left[\frac{\pi \times H_c[k\text{-}1]}{S}\right] \quad \text{for } k = 1, 2, ..., \left(2^N - 1\right) \tag{27}$$

where

A is the amplitude of the sine wave
C is the offset (dc level) of the applied signal
$H_c[j]$ is equal to $\sum_{i=0}^{j} H[i]$
$H[i]$ is the total number of samples received in code bin i
S is the total number of samples

If A and C are unknown, they can be determined from the data and an independent estimate of any two of the transition levels $T[k]$. Errors in the values of A and/or C do not induce any errors in the determination of differential or integral nonlinearity from the calculated transition levels because they only induce gain and

offset errors in the transition levels, as shown in Equation (26). These results are derived by Vanden Bossche et al. [B58].

This test assumes that the transfer function is monotonic (see 8.6).

6.4.1 Comment on the selection of the sine-wave frequency and data record length

The frequency of the sine wave and the record length of the data collected must be carefully selected in order for the error estimates of the preceding clause to apply. There must be an exact integer number, *J*, of cycles in a record, and the number of cycles in a record must be relatively prime to the number of samples in the record. This guarantees that the samples in each record are uniformly distributed in phase from 0 to 2π.

If the test frequency is low enough that dynamic errors do not arise, this method will give the same results as the static test method. If the frequency is chosen large enough that the dynamic errors are significant, the user shall be warned that some dynamic errors will appear in the results while others will be averaged out by the histogram calculations.

A frequency that meets the above requirements can be selected as follows. Choose the number of cycles per record, *J*, and a record length, *M*, such that *J* and *M* have no common factors. Choose the frequency using Equation (28).

$$f_i = \frac{J}{M} f_s \tag{28}$$

where

f_i is the input signal frequency
f_s is the sampling frequency

In order for the test tolerances in the derivations in 6.4.2 to be valid, the accuracy required of the ratio of the input signal frequency to the sampling frequency is given by Equation (29).

$$\frac{\Delta a}{a} \leq \frac{1}{2JM} \tag{29}$$

where

$$a \equiv \frac{f_i}{f_s} \tag{30}$$

With larger values of *M*, fewer total samples (the number of records times the number of samples per record) will be required to obtain any given accuracy, but greater accuracy will be required of the signal frequency. The best approach is to use the largest value of *M* compatible with the frequency accuracy obtainable. The frequency accuracy specified by Equation (29) and Equation (30) guarantee that the phase separation between samples is within ±25% of the ideal. This tolerance is assumed in the derivation of Equation (31).

6.4.2 Comments on sine-wave histogram testing

The same testing inaccuracies can occur in sine-wave histogram testing that are indicated in 6.3.3. Again, to avoid these issues, converters should also be tested for SINAD performance to confirm that non-monotonic behavior is not significant.

6.4.3 Comment on the amount of sine-wave overdrive and the number of records required

The minimum amount of overdrive required in the method in 6.4 depends on the combined noise level of the signal source and the ADC. In the absence of noise, the overdrive need only be sufficient to receive at least one count in each of the first and last code bins. If noise is present, it will modify the probabilities of samples falling in various code bins, and the effect will be largest near the peaks where the curvature of the probability density is greatest. This effect can be made as small as desired by making the overdrive large enough. The amount of overdrive required to obtain a specified accuracy also depends on whether the specified accuracy is for the code bin widths (i.e., differential nonlinearity) or for the transition levels (i.e., integral nonlinearity).

Input overdrive

The overdrive required to obtain a specified tolerance in code bin widths is given by Equation (31).

$$V_o \geq \text{Maximum of } (3\sigma) \text{ or } \left(\sigma \times \sqrt{\frac{3}{2B}} \right) \tag{31}$$

where

σ is the rms value of the random noise in input units
B is the desired tolerance as a fraction of the code bin width
V_o is the input overdrive: the difference between the positive (negative) peaks of the sine wave and the most positive (negative) transition level of the ADC

This amount of overdrive guarantees that the error caused by the noise is ≤1/3 of the desired tolerance.

The overdrive required to obtain a specified tolerance in transition levels is given by Equation (32).

$$V_o \geq \text{Maximum of } (2\sigma) \text{ or } \frac{\sigma^2 2^N}{VB} \tag{32}$$

where

V is the full-scale range of the instrument in input units
N is the number of bits of the ADC

The values of V_o in Equation (31) and Equation (32) are adequate to keep the errors due to noise equal to or less than $B/3$ code bin widths so that these errors are negligible when added to the statistical errors due to taking a finite number of samples.

The number of records required depends on several factors. It depends on the combined noise level of the ADC and the signal source, on the desired test tolerance and confidence level, on whether the tolerance and confidence level is for INL (transition levels) or DNL (code bin widths) and on whether one wants to obtain the desired confidence for a particular width or transition level or for the worst case for all widths or transition levels. The number of records required for a given test tolerance and a given confidence in code bin widths is given by Equation (33).

$$R = D\left[\frac{2^{N-1} K_u}{B} \right] \left[\frac{c\pi}{M} \right] \left\{ 1.13 \left[\frac{\sigma^*}{V} + \frac{c}{2} \sigma_\varphi \right] + 0.25 \left[\frac{c\pi}{M} \right] \right\} \tag{33}$$

where

R is the minimum required number of records

D is equal to 1 for INL and $D = 2$ for DNL
M is the number of samples per record
c is equal to $1 + 2(V_o/V)$
V is the full-scale range of the ADC under test
V_o is the input overdrive
K_u is equal to $Z_{u/2}$ for obtaining the specified confidence in an individual transition level or code bin width
K_u is equal to $Z_{N,u/2}$ for obtaining the specified confidence in the worst-case transition level or code bin width
u is equal to $1 - v$, with v the desired confidence level expressed as a fraction
σ^* is for INL, σ, the rms random noise effects including additive noise and jitter
σ^* is for DNL, the minimum of σ or $Q/2.26$
$\sigma\pi$ is the rms random phase error of the input signal relative to the sampling time, in radians
N is the number of bits of the ADC
B is the desired test tolerance as a fraction of the code bin width

The values for $Z_{u/2}$ and $Z_{N,u/2}$ can be obtained from Table 2. For values of N between those in the table, use linear interpolation. $Z_{u/2}$ is defined such that the probability is $(1 - u)$ that the absolute value of a Gaussian distributed random variable, having a mean of zero and a standard deviation of one, is less than or equal to $Z_{u/2}$. $Z_{N,u/2}$ is defined such that the probability is $(1 - u)$ that the maximum of the absolute values of 2^N Gaussian distributed random variables, having means of zero and standard deviations of one, will be less than or equal to $Z_{N,u/2}$.

Table 2—Values of $Z_{u/2}$ and $Z_{N,u/2}$

u	$Z_{u/2}$	$Z_{4,u/2}$	$Z_{8,u/2}$	$Z_{12,u/2}$	$Z_{16,u/2}$	$Z_{20,u/2}$	$Z_{24,u/2}$
0.2	1.28	2.46	3.33	4.04	4.64	5.19	5.68
0.1	1.64	2.72	3.53	4.21	4.80	5.33	5.81
0.05	1.96	2.95	3.72	4.37	4.94	5.46	5.93
0.02	2.33	3.22	3.95	4.57	5.12	5.62	6.08
0.01	2.58	3.42	4.11	4.71	5.25	5.74	6.19
0.005	2.81	3.60	4.27	4.85	5.38	5.85	6.30
0.002	3.09	3.84	4.47	5.03	5.54	6.01	6.44
0.001	3.29	4.00	4.62	5.16	5.66	6.12	6.54

Equation (28), Equation (29), Equation (30), Equation (31), and the values in Table 2, are derived by Blair [B7]. For further information, see Papoulis [B45].

6.5 Determining the static transfer curve

The transfer curve of an ADC is the average output code, $\bar{y}$, as a function of a particular input signal level, x. The transfer curve, $\bar{y}(x)$, is used as a basis for alternate definitions of many of the static parameters of an ADC, e.g., gain and offset, INL, and monotonicity. The transfer curve (and the parameter definitions based on it) is especially useful in specifying ADCs where it may be impractical or impossible to measure the code transition levels. Examples of such ADCs are those with a high number of digitized bits (it may be impractical to search for all 2^{20} code transition levels for a 20-bit ADC), those with non-monotonic behavior and/or output-to-input crosstalk (which can result in either undefined or multiply-defined code transition levels), and those ADCs that are actually composed of multiple time-interleaved sample-holds and/or ADCs.

Another useful measurement that can be made at the same time as the transfer curve is the deviation of the output codes about the average, again as a function of the input signal value, x.

To estimate the transfer curve, $\bar{y}(x)$, and the standard deviation of the output as a function of the input, $\sigma_y(x)$, a dc input source is required whose output signal range spans slightly more than the full-scale range of the ADC, and that has an accuracy, resolution, and noise better than the desired accuracy of the measurement.

a) Set the level of the input signal, x, slightly below the bottom of the ADC input range (such that further lowering of the input level would not change the ADC output).

b) Acquire M samples from the ADC: $y_0, y_1, y_2, \ldots, y_{M-1}$. M shall be chosen large enough that the standard deviation of the sample mean (the standard deviation of the samples divided by the square root of M) is small compared to the desired accuracy of the measurement.

c) Record the sample mean (estimated average), $\bar{y}(x)$, and standard deviation, $\sigma_y(x)$, of these M samples in arrays indexed by the current input level x.

d) Increment the input level x by a specified amount. Preferably, the increment is roughly equal to the deviation of the additive random noise present within the ADC at the analog input, or $Q/8$, whichever is larger.

e) Repeat steps b), c), and d) until the input level x is set slightly above the top of the ADC input range (such that increasing the input level further would not change the ADC output).

6.5.1 Alternate method

The order of the loops in the above method (the inner loop being the acquisition of M samples in Step b) may be reversed. In this case, the input signal becomes a repetitive ramp, increasing by the specified amount per sample interval. M records of data are acquired, each record triggered so as to contain points digitized from when the ramp was set at a specific point slightly below the bottom of the ADC input range to when the ramp was slightly above the top of the range. The M records are then compared on a point-by-point basis to find the mean and deviation.

Both techniques can also be used with a decreasing rather than increasing input level. If the results differ significantly (e.g., due to hysteresis), the transfer curve is the average of the results using an increasing and a decreasing input level.

7. Analog input

7.1 Input characteristics

The input impedance is the impedance between the signal input and ground. The input impedance may be specified at various frequencies. The minimum specification for input impedance is the static input resistance and, if significant, the input capacitance, inductance, and leakage current.

7.1.1 Static input resistance

The static input resistance is the ratio of the change of an applied static input signal to the resulting (static) change of input current. If the best model of the ADC includes a significant current source, it shall be specified.

7.2 Static input impedance versus input signal level

Some devices may vary in input impedance over the specified operating range of input signal levels. Perform the test in 7.1 with the input signal level set at a minimum of three equally spaced signal levels over the range of the device. Measure the impedance at each of the input settings as defined in 7.1.

7.2.1 Test method

Connect a vector impedance meter of desired accuracy and appropriate output level to the device. Vary the frequency over the desired range and record results. Note that if a test fixture is used which can affect the measurement the construction details of the fixture shall be stated as well as the results obtained in the fixture. Any readings obtained in the fixture shall be recomputed to remove the effect of the fixture.

7.3 Static input current

Apply a series of three equally spaced signal voltage levels across the maximum useable input range of the device. Measure the current entering the device, using an appropriate current meter for each of the signal levels.

7.4 Static gain and offset

Static gain and offset are the values by which the input values are multiplied and then to which the input values are added, respectively, to minimize the mean squared deviation from the output values. For static measurement of gain, see 7.4.1 and 7.4.2. Unless otherwise specified in this standard, static gain and offset will be taken to mean independently based gain and offset measured as in 7.4.1. For dynamic measurement of gain, see 11.3.

7.4.1 Static gain and offset (independently based)

Independently based static gain and offset are the values by which static input values are multiplied and then to which the input values are added, respectively, to minimize the mean squared deviation from the output values. Unless otherwise specified, static gain and offset will be taken to mean independently based static gain and offset.

Test method: Locate the code transition levels as per Clause 6. The transfer characteristic can then be represented by Equation (34).

$$G \times T[k] + V_{os} + \varepsilon[k] = Q(k-1) + T_1 \tag{34}$$

where

$T[k]$ is the input value corresponding to the transition between codes k and $k - 1$
T_1 is the ideal value corresponding to $T[1]$
V_{os} is the output offset in units of the input quantity, nominally equal to zero
G is the gain, nominally equal to unity; G is greater than 1 if the actual thresholds occur at smaller input voltages than nominal
Q is the ideal width of a code bin, that is, the full-scale range divided by the total number of codes
$\varepsilon[k]$ is the residual error corresponding to the k^{th} code transition

The expression on the right side of Equation (34) gives the ideal code transition level, in input units, as a function of k. The variable, k, is assumed to be the value of the binary coded output. Using conventional linear least-squares estimation techniques, independently based static offset and gain are the values of V_{os}

and G that minimize the mean squared value of $\varepsilon[k]$ over all k. The value of G that minimizes the mean squared value of $\varepsilon[k]$ is given by Equation (35).

$$G = \frac{Q(2^N - 1)\left(\sum_{k=1}^{2^N-1} kT[k] - 2^{(N-1)}\sum_{k=1}^{2^N-1} T[k]\right)}{(2^N - 1)\sum_{k=1}^{2^N-1} T^2[k] - \left(\sum_{k=1}^{2^N-1} T[k]\right)^2} \tag{35}$$

and the value of V_{os} that minimizes the mean squared value of $\varepsilon[k]$ is shown in Equation (36)

$$V_{os} = T[1] + Q(2^{(N-1)} - 1) - \frac{G}{(2^N - 1)}\sum_{k=1}^{2^N-1} T[k] \tag{36}$$

Given these values for G and V_{os}, $\varepsilon[k]$ is the independently based integral nonlinearity (see 8.2).

7.4.1.1 Alternate method for determining gain and offset

The independently based static gain and offset may alternatively be found by using a least-squares fit of a straight line to the transfer curve (see 7.4.1). In order to avoid having the ends of the transfer curve, where the ADC is overdriven, affect the fit, the straight line is fitted just to that portion of the transfer curve where the average output code is between its minimum value plus twice its deviation and its maximum value minus twice its deviation. This method may give slightly different results to those of the fit of straight line to the code transition levels, but the differences are insignificant in practical cases.

7.4.2 Static gain and offset (terminal based)

Static terminal-based gain and offset are the values by which static input values are multiplied, and then to which the input values are added, respectively, to cause the deviations from the output values to be zero at the terminal points, that is, the first and last codes.

Test method: Locate the code transitions as per Clause 6. The transfer characteristics can be represented by Equation (34). Terminal-based static gain and offset are the values of G and V_{OS} that cause $\varepsilon[1] = 0$ and $\varepsilon[2^N - 1] = 0$, where N is the number of digitized bits and $2^N - 1$ is the highest code transition defined. The terminal-based gain and offset are given explicitly by:

$$G = \frac{Q(2^N - 2)}{T[2^N - 1] - T[1]} \tag{37}$$

and

$$V_{os} = T_1 - GT[1] \tag{38}$$

The variables are the same as those in Equation (34).

8. Linearity

8.1 General comments on linearity

For a perfectly linear ADC, all of the code bin widths would be exactly equal. This section describes a number of common measurements of nonlinearities and related effects in both the time and frequency domains.

8.2 Integral nonlinearity

The integral nonlinearity is the difference between the ideal and measured code transition levels after correcting for static gain and offset. Integral nonlinearity is usually expressed as a percentage of full scale or in units of LSBs. It will be independently based or terminal based depending on how static gain and offset are defined. When the integral nonlinearity is given as one number without code bin specification, it is the maximum absolute value integral nonlinearity of the entire range.

8.2.1 Integral nonlinearity test method

Find the static gain and offset per the method described in 7.4.1 or 7.4.2, as appropriate for independently based or terminal-based static gain and offset. The static integral nonlinearity as a function of k is given in percent by Equation (39):

$$\mathrm{INL}[k] = 100\% \times \frac{\varepsilon[k]}{2^N \times Q} = 100\% \times \frac{\varepsilon[k]}{V_{FS}} \tag{39}$$

where

INL[k] is the integral nonlinearity at output code k

$$\varepsilon[k] = \frac{GT[k] + V_{os} - T_{\mathrm{nom}}[k]}{Q} \tag{40}$$

Q is the ideal code bin width, expressed in input units (the full-scale input range divided by the total number of code states)
V_{FS} is the full-scale range of the ADC in input units
N is the number of digitized bits per sample for the ADC

The maximum INL is the maximum value of |INL[k]| for all k.

Note that if code transitions are determined by a histogram method, and the test signal parameters are inaccurately known, then the gain and offset determined here may be in error. However, the error in gain and offset will not materially affect the calculated INL.

Note that code transition levels are undefined at any codes where the ADC is not monotonic or where the codes are missing.

8.2.2 Alternate test method for determining INL

An alternative method can determine INL from the transfer curve. This is useful in cases where the code transition levels are difficult or impossible to determine, e.g., when the ADC is actually a set of interleaved ADCs. To use this method, first determine the transfer curve according to 6.5. Then smooth it with a boxcar function of width Q, centered on zero, and with an area of one. This creates a running average of

length Q. The INL is the maximum absolute deviation of the smoothed transfer curve from the best-fit straight line as found in determining gain and offset in 7.4.1.1. The search for the maximum is taken over the same range of input levels as used to fit the straight line in 7.4.1.1. Because of this, this test method may not find the worst-case INL if it occurs within twice the deviation of the first or last transition level.

8.3 Absolute accuracy error

The maximum difference between any measured code transition level and its ideal value. It is often expressed as a percentage of full-scale or in LSB.

8.3.1 Test method

Locate the code transition levels per Clause 6. The absolute accuracy error (AAE) is given, in percent, by

$$AAE = 100 \times \frac{\max\{T[k] - Q \times (k-1) - T[1]\}}{Q \times 2^N} \quad \text{(in units of \%)} \tag{41}$$

$$AAE = \frac{\max\{T[k] - Q \times (k-1) - T[1]\}}{Q} \quad \text{(in units of LSB)} \tag{42}$$

where

$T[k]$ is the code transition level for the k^{th} transition (between codes $k - 1$ and k)
Q is the ideal width of a code bin
N is the number of digitized bits

This measurement includes linearity, offset, and gain errors.

8.4 Differential nonlinearity and missing codes

Differential nonlinearity (DNL) is the difference, after correcting for static gain, between a specified code bin width and the ideal code bin width, divided by the ideal code bin width. When given as one number without code bin specification, it is the maximum value differential nonlinearity of the entire range.

8.4.1 Differential nonlinearity and missing codes test method

Locate the code transition levels by any of the methods of Clause 6. Differential nonlinearity is calculated using Equation (43).

$$\text{DNL}[k] = \frac{(W[k] - Q)}{Q} \tag{43}$$

where

$W[k]$ is $G \times (T[k+1] - T[k])$
Q is the ideal code bin width
G is the static gain

Neither the width of the top bin, $W[2^N - 1]$, nor that of the bottom bin, $W[0]$, is defined. A code k is defined to be a missing code if Equation (44) is true.

$$DNL[k] \leq -0.9 \tag{44}$$

Perfect differential nonlinearity coincides with DNL = 0.

The maximum differential nonlinearity is the maximum value of |DNL[k]| for all k. In addition, the RMS value of the DNL can be given as shown in Equation (45).

$$\mathrm{DNL}_{\mathrm{RMS}} = \sqrt{\frac{1}{2^N - 2} \sum_{k=1}^{2^N - 2} \{\mathrm{DNL}[k]\}^2} \tag{45}$$

8.5 Example INL and DNL data

Figure 11 and Figure 12 show sample plots based on the measurement of transition levels per Clause 6 for a 12-bit ADC, from an end-point calibration. Figure 11 shows the INL error and Figure 12 shows DNL error.

Note that a DNL of 0 for a particular code means that code width was equal to the average width. A DNL of −1 means the code was missing. And a DNL of +1 means the code width was twice as wide as the average.

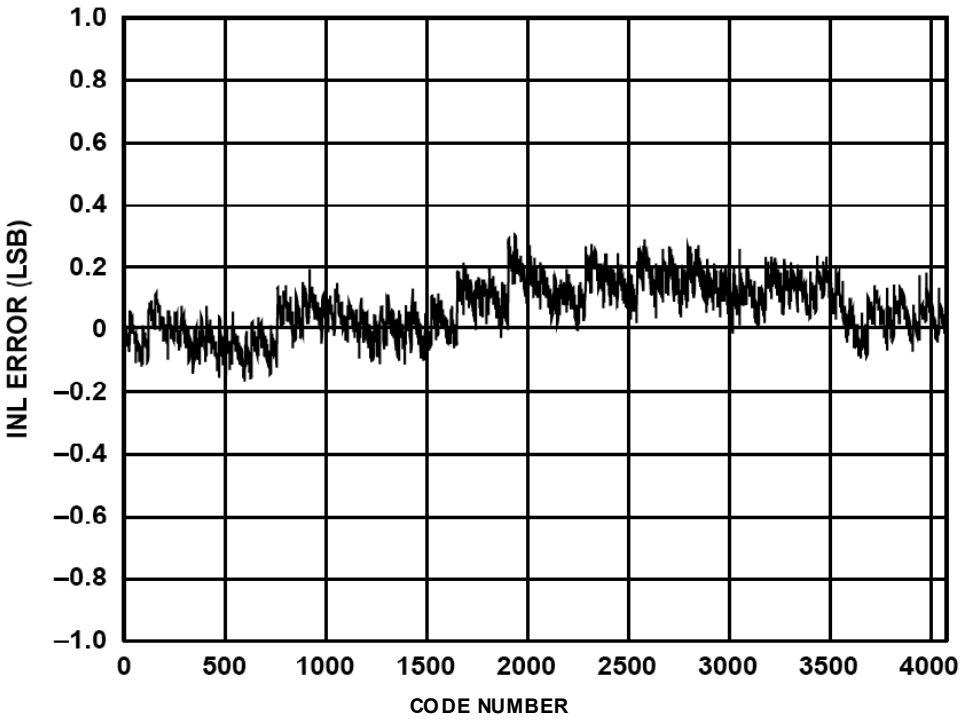

Figure 11—Example of calculated INL from a terminal-based calibrated 12-bit device

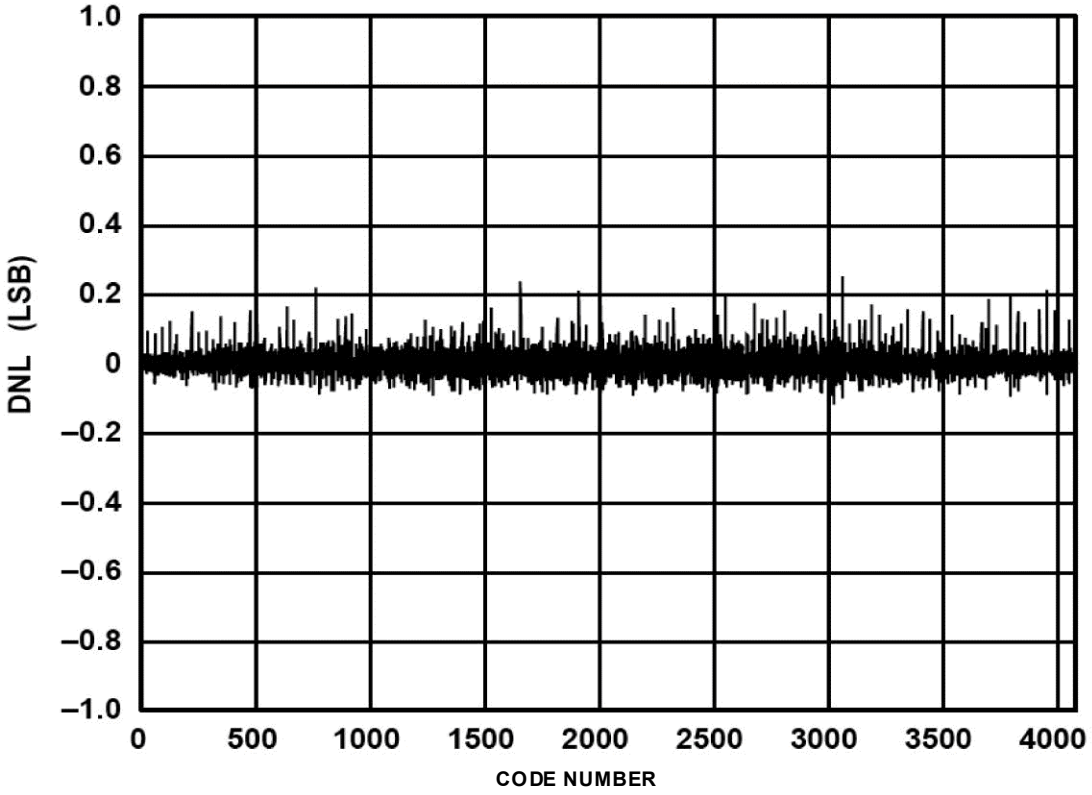

Figure 12—DNL from a terminal-based calibrated 12-bit device

8.6 Monotonicity

A monotonic non-inverting ADC produces output codes that are consistently increasing with increasing input stimulus and consistently decreasing with decreasing input stimulus, changing in the same direction relative to the change in input stimulus. If the input stimulus and output codes change consistently in opposite directions, e.g., a higher input produces a lower output code; the ADC is monotonic and inverting.

8.6.1 Test method

Determine the transfer curve of the ADC using both increasing and decreasing input levels, according to 6.5. Then the ADC is considered non-monotonic if, for any pair of input levels x_1 and x_2, with $x_1 < x_2$ as shown in Equation (46) and Equation (47).

$$\bar{y}(x_1) - \frac{3\times\sigma_y(x_1)}{\sqrt{M}} > \bar{y}(x_2) + \frac{3\times\sigma_y(x_2)}{\sqrt{M}}, \quad \text{(non-inverting)} \tag{46}$$

$$\bar{y}(x_2) - \frac{3\times\sigma_y(x_2)}{\sqrt{M}} > \bar{y}(x_1) + \frac{3\times\sigma_y(x_1)}{\sqrt{M}}, \quad \text{(inverting)} \tag{47}$$

using the notation of 6.5, where $\bar{y}$ (x) is the mean output code value and $\sigma_y(x)$ is the standard deviation of the output code value, for a given static input signal level x, and M is the number of samples taken at each x value. Note that it is best to keep $M > 20$ so that the standard deviation estimates $\sigma_y(x_1)$ and $\sigma_y(x_2)$ are statistically valid.

8.7 Hysteresis

The measured value of the ADC transfer curve may depend on the direction by which the transfer curve is traversed (i.e., increasing or decreasing signal). The reported hysteresis of the ADC, if any, is the maximum of such differences.

8.7.1 Hysteresis test method

Determine the transfer curve of the ADC using both increasing and decreasing input levels (see 6.5). Let $\bar{y}_+(x)$ and $\bar{y}_-(x)$ denote the mean output values measured at input level x for increasing and decreasing input levels, respectively, and let $\sigma_{y+}(x)$ and $\sigma_{y-}(x)$ be the calculated standard deviation of those measured values, for increasing and decreasing input levels, respectively. Let $M_+(x)$ and $M_-(x)$ be the number of measurements of the output value at input value x, for increasing and decreasing input values, respectively.

If, for all levels x, the difference between the measured mean ADC output values for increasing and decreasing input levels is within the random measurement uncertainty, i.e., if Equation (48) is true.

$$\left|\bar{y}_+(x)-\bar{y}_-(x)\right| \le \left|\frac{3\times\sigma_{y+}(x)}{\sqrt{M_+(x)}}-\frac{3\times\sigma_{y-}(x)}{\sqrt{M_-(x)}}\right| \tag{48}$$

then the ADC is said to have no hysteresis. Note that it is best to keep M_+ and M_- greater than about 20, so that the standard deviation estimates are statistically valid. If the ADC does have hysteresis, the amount is given by the magnitude of the largest observed difference, converted to the units of the input signals. Thus in Equation (49)

$$\text{hysteresis} = \frac{\left|\bar{y}_+(x)-\bar{y}_-(x)\right|_{\max}}{G} \tag{49}$$

8.7.2 Alternate hysteresis test method

Another method of measuring ADC hysteresis, using a feedback loop, is described in 9.5.3. There, the magnitude of the hysteresis at the code transition level of interest is converted into the position of a peak in a frequency spectrum, which is measurable with a spectrum analyzer.

8.7.3 Comment on hysteresis and alternation

It is important to distinguish between real hysteresis, and the effect of transition noise. Hysteresis is present only when the rising and falling transitions differ significantly from the transition noise.

Hysteresis may be present only when there is a significant rate of change of the input. If the effect of the slope of the input is a significant parameter, then the specification shall specify a rate of change of the input that resulted in the measured hysteresis.

A phenomenon called “alternation” may also be observed. This occurs when a significant range of input voltages results in adjacent ADC output codes being alternately expressed. Alternation is a complementary function to hysteresis since both hysteresis and alternation are unwanted feedbacks to the input from previous output codes. If the feedback is positive then hysteresis is observed. If the feedback is negative then alternation is observed.

To illustrate the phenomenon of hysteresis and alternation, consider the 1-bit ADC shown in Figure 13.

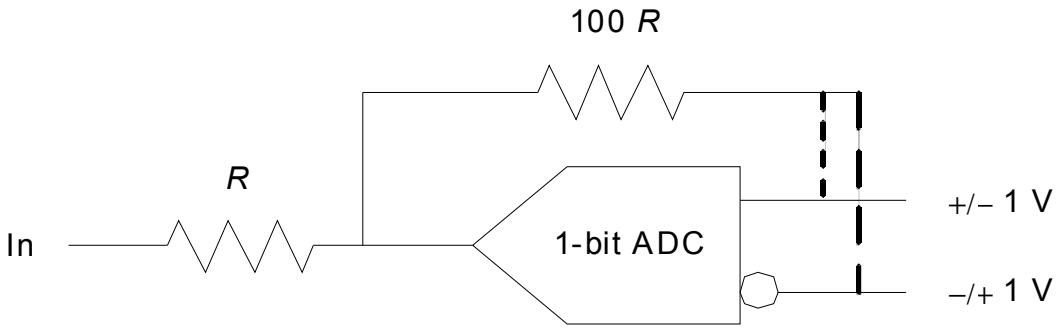

Figure 13—Block diagram to demonstrate hysteresis and alternation

Consider the case when the "True" output is connected to the divider. The input is gradually changed from a negative voltage, which has caused the "True" output to be at the −1 V level. The output of the ADC will go to its high state only after the input voltage has exceeded 10 mV. If the initial voltage had started from a positive value, and then decreased, it would have to go lower than −10 mV to force the output to the low level. This is the phenomenon of hysteresis.

Consider the case when the "False" output is connected to the divider. The input is gradually changed from a negative voltage that has caused the "False" output to be at the +1 V level. The "True" output of the ADC will go to its high state after the input voltage has exceeded −10 mV. From that point onward the output will alternate between the high and low states until the input voltage exceeds +10 mV. If the initial voltage had started from a positive value, and then decreased, the alternation would resume when the input dropped below +10 mV. This is the phenomenon of alternation.

8.8 Harmonic and spurious distortion

Dynamic errors and integral nonlinearities contribute to harmonic distortion whenever an ADC is sampling a periodic signal. This subclause describes different measures that are used to quantify such behavior.

For a pure sine-wave input, the output harmonic distortion components are found at spectral values whose non-aliased frequencies are integer multiples of the applied sine-wave frequency. Their amplitudes, which depend upon the amplitude and frequency of the applied sine wave, are generally given as a decibel (dB) ratio with respect to the amplitude of the applied sine-wave input. Their frequencies are usually expressed as a multiple of the frequency of the applied sine wave. For pure sine-wave input, spurious or non-harmonic components are persistent sine waves at frequencies other than the fundamental (input) or those described above as harmonic components. Usually, their amplitudes are expressed as a decibel ratio with respect to a full-scale signal (dBFS).

The measures described in the following subclauses derive from the use of a spectrally pure, large amplitude sine-wave input as the test signal. The test signal should be properly filtered to lower all harmonics to a level that is well below the desired measurement resolution for the device under test. The amplitude and frequency of the test signal shall be specified in the test results.

8.8.1 Total harmonic distortion

The scaled square root of the sum-of-squares of a specified set of harmonic distortion components including their aliases for an input of a pure sine wave of specified frequency and amplitude. THD amplitude is always expressed relative to the amplitude of the applied signal, either as a percent or in decibels.

The total harmonic distortion is given by the ratio:

$$\mathrm{THD} = \frac{\sqrt{\frac{1}{M^2}\sum_h |X[f_h]|^2}}{A_{\mathrm{rms}}} \tag{50}$$

where

$X[f_h]$ is the complex value of the spectral component at frequency f_h
f_h is the h^{th} harmonic frequency of the DFT of the ADC output data record computed using Equation (C.1)
M is the number of samples in the data record
h is the set of harmonics over which the sum is taken
A_{rms} is the rms value of the input sine wave

$$A_{\mathrm{rms}} = \frac{1}{M}\sqrt{(X_{\mathrm{avm}}(f_i))^2 + (X_{\mathrm{avm}}(f_s - f_i))^2} \tag{51}$$

The summation is taken over all of the harmonics used that are described in the various test methods given below. The quantity under the square root is the square of the rms value of the signal consisting of the specified harmonics.

The members of the set of harmonics, f_h, used in Equation (50), must be specified. The choice of harmonic components included in the set is a tradeoff between the desire to include all harmonics with a significant portion of the harmonic distortion energy, but not to include DFT bins whose energy content is dominated by random noise. Unless otherwise specified, to estimate THD, the set is normally composed of the lowest nine harmonics, second through tenth inclusive, of the input sine wave.

8.8.1.1 Coherent sampling THD test method

To estimate THD, apply a test signal consisting of a pure, large amplitude sine wave at frequency f_i chosen to meet the criteria for coherent sampling. See C.1.2 of Annex C for a discussion of coherent sampling and the DFT. The sine-wave frequency shall not be Nyquist or any multiples thereof.

Acquire K data records of M points each from the ADC under test at sample frequency f_s. Let $x_k[n]$ represent the k^{th} record of sine-wave data for k = 1, 2, …, K. For each $x_k[n]$ record, compute the DFT, $X_k[m]$, where m is an integer from 0 to M − 1. The K sets of data are used to compute an averaged magnitude spectrum of the DFT at each basis frequency f_m:

$$X_{\mathrm{avm}}[m] = \frac{1}{K}\sum_{k=1}^{K} |X_k[m]|, \quad m = 0, 1, 2, \ldots, M-1 \tag{52}$$

The averaged spectral magnitude, X_{avm}, is used because it has a smaller variance than the non-averaged spectral magnitude, X. The standard deviation of the random errors in X_{avm} is less than $|X|$ by a factor approximately equal to the square root of K (Jenkins and Watts [B28]).

Identify the set of bin numbers, n_h, which corresponds to the chosen set of harmonics of the input test frequency. For a test tone at frequency f_i, the harmonics are aliased so that f_h lies between 0 and the sampling frequency f_s. Aliasing is accounted for by means of Equation (53). The default value for N_H is 10:

$$n_h = \mathrm{mod}(hn_i, M) \quad h = \pm[2, 3, \ldots, N_H] \tag{53}$$

where

n_i is the bin number of the input frequency, $n_i = Mf_i/f_s$
N_H is the number of the highest order harmonic used

This procedure locates both Euler components of the aliased harmonics. It is important that n_i be chosen so that all of the n_h are different and that none of them is equal to either n_i or 0. THD is then given by:

$$\mathrm{THD} = \frac{\sqrt{\frac{1}{M^2}\left(\sum_{h=2}^{N_h} X_{\mathrm{avm}}[n_h]^2 + \sum_{h=-N_h}^{-2} X_{\mathrm{avm}}[n_h]^2\right)}}{\frac{1}{M}\sqrt{X_{\mathrm{avm}}[n_i]^2 + X_{\mathrm{avm}}[M-n_i]^2}} \tag{54}$$

8.8.1.2 Noncoherent sampling test method 1 (windowed DFT)

This method uses the windowed DFT to reduce the problems caused by spectral leakage. It is of value when the input frequency does not satisfy the condition for coherent sampling with sufficient accuracy. Since the windowed DFT is used, each spectral line splits into several lines, the number depending on the window. Therefore, it is necessary to use the values from several DFT bins to calculate the rms value of the input signal and each harmonic. Choose an L^{th} order cosine window, for some small integer L, following the guidance in Annex C.

To estimate THD, apply a test signal consisting of a pure, large amplitude sine wave. Acquire K data records of M points each from the ADC under test at sample frequency, f_s. Let $x_k[n]$ represent the k^{th} record of sine-wave data for $k = 1, 2, ..., K$. For each $x_k[n]$ record, compute the windowed DFT, $X_{w,k}[m]$, where m is an integer between 0 and $M - 1$. The K sets of data are used to compute an averaged magnitude spectrum of the windowed DFT at each basis frequency f_m:

$$X_{\mathrm{wavg}}[m] = \frac{1}{K}\sum_{k=1}^{K}\left|X_{w,k}[m]\right| \qquad \text{for } m = 0,1,2,...,M-1 \tag{55}$$

where

$X_{\mathrm{wavg}}[m]$ is the spectrum averaged over K records of the DFTs of the windowed data records
$|X_{w,k}[m]|$ is the magnitude of the DFT of each windowed record computed using Equation (C.5)

Identify the set of bin numbers, n_h, which corresponds to the chosen set of harmonics of the input test frequency. For a test tone at frequency, f_i, the harmonics are aliased so that f_h lies between zero and the sampling frequency, f_s. Aliasing is accounted for by means of Equation (53). The default value for N_H is 10.

The resulting values will not be integers if the input frequency does not exactly fulfill the requirement for coherent sampling. This procedure locates both Euler components of the aliased harmonics. It is important that n_i be chosen so that all of the n_h are sufficiently different that the range of bin numbers used in their calculations (given below) do not overlap and that none of them contains either 0 or the range of bin numbers used for calculating $X_{\mathrm{wavg}}[n_i]$. Let $X_{\mathrm{avm}}[n]$ be:

$$X_{\mathrm{avm}}[n]^2 = \frac{1}{\mathrm{NNPG}}\sum_{k=-(l+1)}^{L+1} X_{\mathrm{wavg}}[n+k]^2 \text{ for } \begin{cases} n = n_i \text{ for } h = \pm 1 \\ n = n_h \text{ for } h = \pm\{2,2,...,N_h\} \end{cases} \tag{56}$$

where

NNPG is the normalized noise power gain of the window, which is defined in Equation (C.6)

n are the values for n_i and n_h rounded to the nearest integer
X_{wavg} is given by Equation (55)
L is the order of the cosine window function used

The THD is then given by Equation (54).

8.8.1.3 Noncoherent sampling test method 2 (sine fitting)

This method uses sine fitting to determine the input signal and harmonic amplitudes rather than the DFT. It is somewhat more computationally intensive than using fast algorithms for the DFT. Its advantage over the method in the previous clause is that it is less sensitive to noise and more thoroughly eliminates spectral leakage.

The data used are the same as for the other methods. To maximize accuracy each data record shall be truncated so that it has approximately an integer number of cycles of the input signal. It is assumed that this truncated record is being used throughout this clause. The sine fits are performed as described in Clause 5.

Perform either a three-parameter or four-parameter sine fit to the data to determine the input amplitude A_1 and, if using a four-parameter fit, the input frequency, f_i. Calculate the residuals as described in 5.2.

For each harmonic number, h, between 2 and N_H, perform a three-parameter sine fit to the residuals with a frequency of hf_i to determine the harmonic amplitude A_h.

THD is given by

$$\text{THD} = \frac{\sqrt{\sum_{h=2}^{N_H} {A_h}^2}}{A_1} \tag{57}$$

where

A_h is the amplitude of the h^{th} harmonic

The reason for truncating the records to an approximate integer number of cycles is to allow each harmonic to be accurately determined separately.

If multiple records are used, the value for the A_i used in Equation (57) shall be the average of the values from the individual records.

8.8.1.4 Comments on record lengths, sample rate, and input frequency for noncoherent sampling using curve fitting

The uncertainty in the calculated harmonic distortion due to noise is proportional to the square root of the reciprocal of the record length. Thus, longer record lengths reduce the effects of noise. However, longer record lengths are more susceptible to errors due to frequency instability, which could cause the frequency to not be constant throughout the record. The record length can be truncated to minimize leakage between components by including in the record only the number of points that are close to an integer number of cycles of the fundamental. Truncate the record length to be approximately an integer number of cycles. If the record length is not truncated correctly, there will be some “leakage” between the calculated harmonic values.

To demonstrate this effect, if the harmonics are at random phases with respect to the fundamental, and the harmonics are all 1% of the fundamental, the following distribution of deviation from the nominal of the

THD and the fundamental are observed when 6.5 cycles of the signal are present in the record. The observed leakage can theoretically be reduced by solving many simultaneous equations, but the simplified algorithm has well-bounded errors as shown in the plot in Figure 14. The errors introduced by noncoherent sampling can be reduced by taking multiple records, and averaging the results.

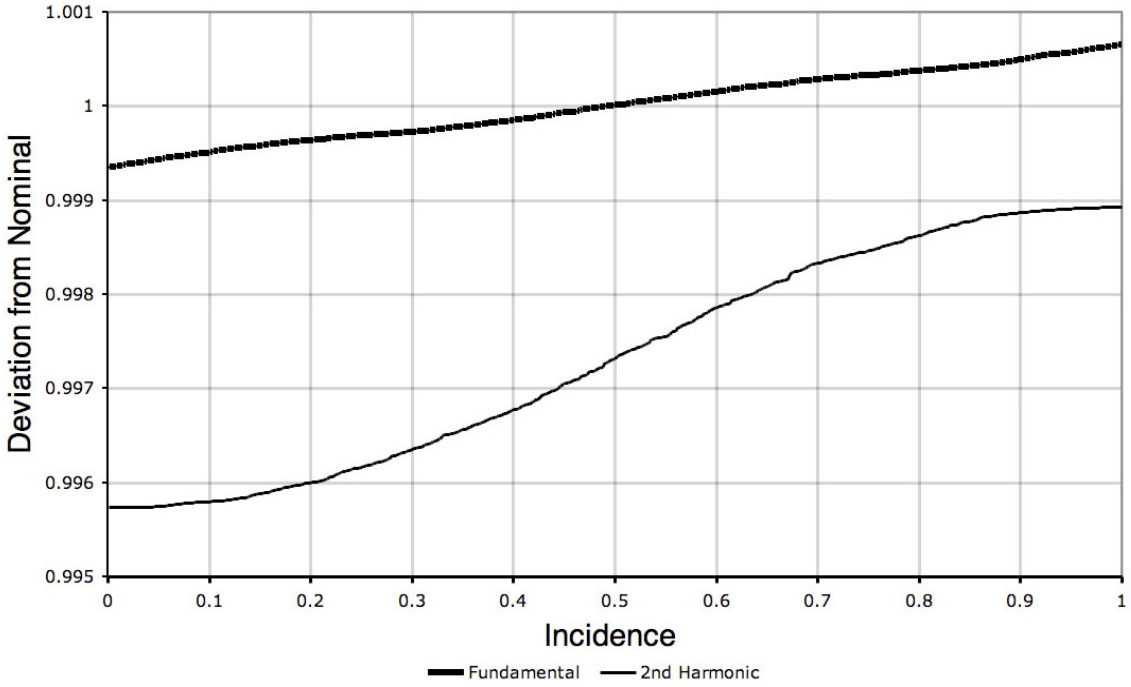

Figure 14—Plot of distribution of calculated values of fundamental and harmonic amplitudes for noncoherent sampling with random phase between the second harmonic and the fundamental

The error in the calculation of THD is a strong function of the number of cycles of the waveform that are included in the record. The worst deviations occur when $M + 0.5$ cycles of the fundamental are included in the record. The plot in Figure 15 illustrates the variation in the THD as M is increased (the phases of the harmonics are fixed in this plot at $\pi/4$, and the record length is fixed at 16 384 samples). The error in the THD is interpreted as a fraction of the measured THD. From the graph one observes that with a ten-cycle record a measured THD of 1% could actually be between 1.014% and 0.986%.

It can be observed that the THD deviation varies inversely with the number of cycles over a wide range. The largest deviations are concurrent with a smaller number of cycles. When a small number of cycles are present in the record, it is fairly easily to correct the THD deviations by truncating the record length.

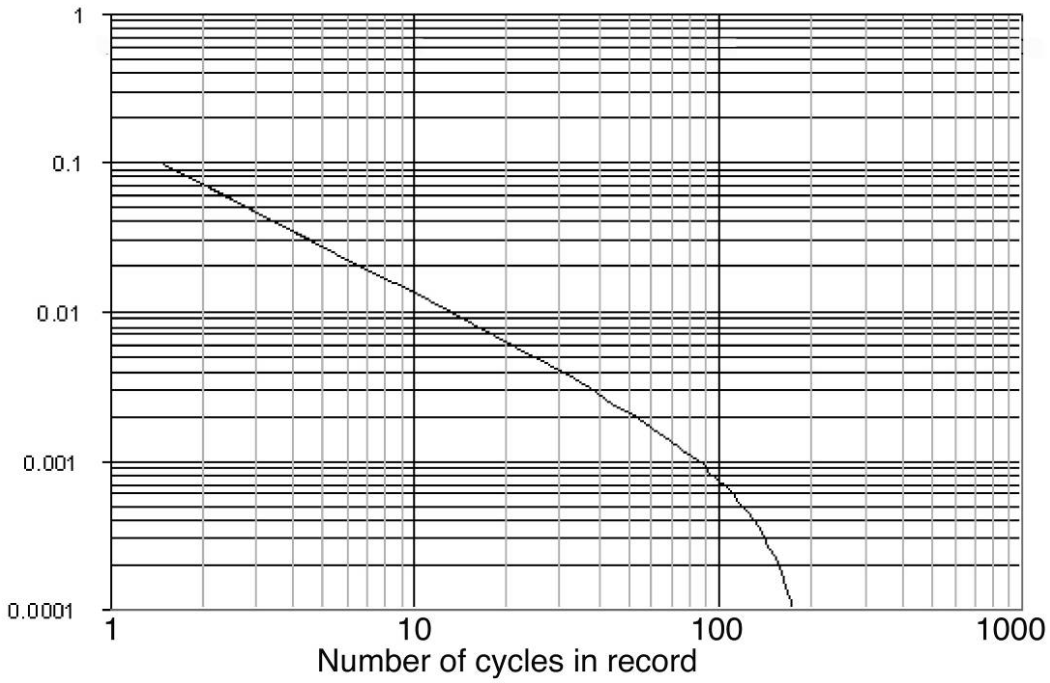

Figure 15—THD deviation from true value as a function of number of cycles in the record

8.8.2 Spurious free dynamic range

The spurious free dynamic range (SFDR) is the frequency domain difference in decibels between the input signal level and the level of the largest spurious or harmonic component for a large, pure sine-wave signal input. This parameter is used to indicate the ADC's usable dynamic range beyond which problems occur in spectral analysis. SFDR is a function of both the amplitude and the frequency of the input sine wave, and possibly of the ADC sample frequency as well as input noise or dither. Thus, the amplitude and frequency of the input, and the sample frequency for which SFDR measurement(s) are made shall be specified.

8.8.2.1 Coherent sampling SFDR test method

The test procedure for estimating SFDR using coherent sampling is given as follows:

a) Apply a test signal consisting of a pure, large amplitude sine wave at frequency f_i chosen to meet the criteria for coherent sampling. See C.1.2 of Annex C for a discussion of coherent sampling and the DFT.

b) Collect K records of data, $x_k[n]$, each containing M samples.

c) Compute the magnitude of the DFT of each record, $|X_k[n]|$, using Equation (C.1) in Annex C..

d) Compute the average over all of the K DFT records, $X_{\text{avm}}[n]$, at each spectral component using Equation (52).

e) Compute SFDR using Equation (58) as follows:

$$\text{SFDR} = 20\log\left(\frac{X_{\text{avm}}[n_1]}{X_{\text{avm}}[n_{nf}]}\right) \tag{58}$$

where

$X_{avm}[n]$ is the averaged magnitude of the spectral component at frequency index n
n_{nf} is the set of frequency indices that are not the fundamental or dc
n_1 is the index for the fundamental frequency

8.8.2.2 Noncoherent SFDR sampling test method

The test procedure for estimating SFDR using noncoherent sampling is given as follows:

a) Apply a test signal consisting of a pure, large amplitude sine wave at frequency f_i.

b) Collect K records of data, $x_k[n]$, each containing M samples.

c) Apply an appropriate window function to each record per Annex C.

d) Compute the magnitude of the DFT of each windowed record, $|X_k[n]|$, using Equation (C.5).

e) Compute the average over all of the K DFT records, $X_{avm}[n]$, at each spectral component using Equation (52).

f) Compute SFDR using Equation (58).

8.9 Intermodulation distortion

Intermodulation distortion may occur due to ADC nonlinearities when sampling a signal composed of two or more sine waves or narrowband signal groups. This subclause describes different measures that are used to quantify such behavior. Intermodulation distortion spectral components may occur at sum and difference frequencies for all possible integer multiples of the input frequency tones or signal group frequencies. The power series model of the ADC transfer function can be used to predict the intermodulation distortion phenomenon.

The measure described in the following subclause is based upon the use of an input composed of two independent pure sine waves. The cautionary comments given in 8.9.1.2 also apply to this test.

8.9.1 Intermodulation distortion test method using two tones

Apply a test signal consisting of the sum of two independent, pure sine waves with frequencies, f_{r1} and f_{r2}, at values that are an odd number of DFT bins away from $f_s/2$, with $f_{r2} > f_{r1}$. The difference, Δf, between f_{r2} and f_{r1} is then always an even number of DFT bins.

Take K records of data. Compute the averaged magnitude spectrum, $X_{avm}[f_m]$, as specified in Equation (52) in 8.8.1.1 as described for the THD test.

Intermodulation distortion magnitudes for a two-tone input signal are found at specified sum and difference frequencies, f_{imf}, noted below in Equation (59) and Equation (60).

The difference frequencies are:

$$f_{imf} = \left|(i)f_{r2} - (j)f_{r1}\right| \tag{59}$$

And the sum frequencies are:

$$f_{imf} = (i)f_{r2} + (j)f_{r1} \tag{60}$$

where $i, j = 0, 1, 2, 3, \ldots$ are integers, such that $|i| + |j| > 1$.

8.9.1.1 Comments on test procedure

There are no specific guidelines to specify what frequencies and signals shall be used to perform intermodulation tests since the test parameters are influenced by each individual application. The size of Δf depends upon the application and the information desired.

The range for the integers i and j only need span nonnegative values; however, conjugate Euler frequencies can be determined using negative integers if desired. Range limits of three or four are appropriate for an ADC whose harmonic distortion test (see 8.8) shows that second and third harmonic distortion is dominant. Note that for small Δf, the intermodulation frequencies are clustered around harmonics of f_{r1} and f_{r2}. The location of the aliased intermodulation frequencies, within the sampling band, follows the modulo f_s procedure specified in 8.8.1.1 through Equation (53) and Equation (54).

Two-tone intermodulation distortion is generally a function of the amplitudes, $X_{avm}(f_{r1})$ and $X_{avm}(f_{r2})$, and the frequencies, $f_{r1}$and f_{r2}, of the input components. Thus, the amplitudes and frequencies of the input components for which intermodulation distortion measurement(s) are made shall be specified.

8.9.1.2 Additional comments

Note that the term "m^{th}-order" is commonly used to describe specific nonlinear system behavior such as "third-order" intercept points, etc. The "m^{th}-order intermodulation products" are found for those values of i and j that satisfy $m = |i| + |j|$, for the sum and difference frequencies defined by Equation (59) and Equation (60). For example, for $m = 3$, $(i, j) = (3,0)$, $(2,1)$, $(1,2)$, and $(0,3)$. The frequencies found for $i = 0$ or $j = 0$ correspond to harmonic distortion. However, the measured distortion may be different than that measured for single sine-wave input due to the presence of the other input sine wave.

A typical set of intermodulation distortion tests might involve three pairs of frequencies f_{r1} and f_{r2}, e.g., pairs of frequencies close to 0, $f_s/4$, and $f_s/2$, for a conventional Nyquist-band-limited ADC application. The three pairs of frequencies would be exercised at different input amplitude combinations, e.g., each at −7 dB, −20 dB, and −40 dB below full scale (dBFS); or one tone could be held at −7 dBFS while the other is incremented in equal steps from the noise floor to −7 dBFS. Other ADC applications, such as intermediate-frequency (IF) sampling, may require intermodulation distortion tests with input frequencies spanning from $f_s/2$ to f_s; etc.

One caution about this test is that the intermodulation distortion of the test-input signal must be significantly smaller than the specification of the ADC to be tested. Intermodulation distortion can easily occur between two signal generators that have output-leveling circuitry and are coupled to one another through balanced, or so-called isolated, ports of a hybrid, and other coupling circuits. In addition, the hybrids, or passive filters, used to combine two tones should be operated well within their linear range limits in order to avoid the generation of intermodulation distortion in the resulting test signal input to the ADC.

8.9.2 Intermodulation distortion test methods using more than two tones

Multi-tone intermodulation distortion tests are often used to evaluate ADC overall linearity performance in systems such as broadband data communication, in which, harmonics of a single-tone or intermodulation products of dual-tones would be outside the band of interest and would not be relevant to overall performance. When the frequency spacing between adjacent tones is constant, frequencies of intermodulation products will be at the same frequencies as test tones. This makes intermodulation products measurement impossible. In order to overcome this problem, some tones are intentionally removed in order to measure the intermodulation products that originated from the other tones.

A typical test procedure uses an arbitrary waveform generator, or by default, a computer-controlled DAC to generate a signal composed of a set of sine waves having frequencies that are set at DFT bin center frequencies. Gaps between the tones are used as observation points to measure intermodulation distortion in the spectrum of X_{avm} as the amplitudes of the tones are uniformly increased from the noise floor to a level where the signal starts to be clipped as it exceeds the ADC full-scale range. Such a test provides results similar to the NPR test, but allows for better simulation of expected signal group waveforms.

8.9.2.1 Multi-tone power ratio

Multi-tone power ratio (MTPR) is defined as a figure of merit for applications such as asymmetric digital subscriber line (ADSL), where groups of frequency domain impulses (tones) uniformly spaced over the bandwidth of interest, with the characteristic that periodically a tone is "missing." The output waveforms are analyzed to determine how much power has bled into the "missing" tone. The ratio of the power in the spectral notches to the power of the frequency components is the MTPR. Figure 16 illustrates the system input and system output of an MTPR test.

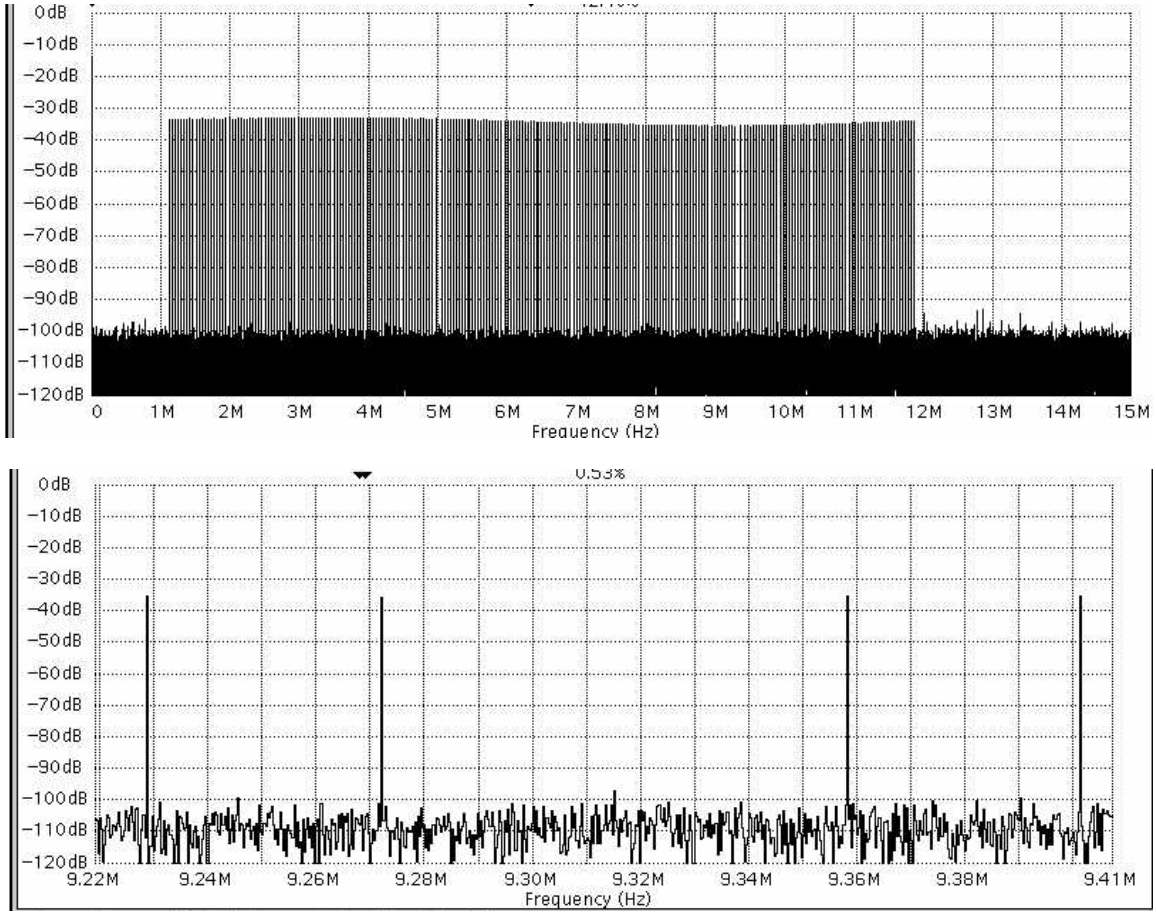

Figure 16—Full power spectrum of a multi-tone for MTPR test (upper) and close-up around a missing tone (lower)

8.9.2.2 MTPR test method

An MTPR test is commonly performed to determine the non-linearity of a discrete multi-tone (DMT) system, also known as orthogonal frequency division multiplexing (OFDM). A DMT waveform is a signal consisting of multiple discrete frequency components. To execute an MTPR test, a DMT waveform is created with frequency components that span the region to be tested. Additionally, this waveform contains missing frequency components, or spectral notches. The frequency range, number of frequency component, and missing tones are specified by the corresponding standard.

Once the DMT waveform has been created, it is loaded onto the AWG, using the test setup described in 4.2.2. Take a record of data, and measure the true rms voltage at one of the carriers and at the frequency of the missing tone. The MTPR is given by

$$\mathrm{MTPR} = 20\log\left(\frac{\mathrm{rms\,carrier}}{\mathrm{rms\,missingtone}}\right) \tag{61}$$

8.9.2.3 Comments on MTPR

The drawback to DMT in general is the bothersome peak-to-average (PAR) ratio associated with the channel waveform due to possible subcarrier instantaneous summation. Each carrier is given a starting phase to constrain the PAR. Specifically, each tone's starting phase is adjusted to establish a desired PAR.

8.10 Noise power ratio

The NPR is the ratio of the average out-of-notch to the average in-notch power spectral density magnitudes for the DFT spectrum of the ADC output sample set which is collected from a notch-filtered broadband white-noise generator as the input.

The dynamic performance of an ADC with broad bandwidth input is sometimes characterized by measuring NPR. In ADC applications where the input signal contains a large number of noncoherent tones or narrow bandwidth signals, it is generally desired that distortion, due to combinations of strong signal components, should not interfere with detection of weaker signal components. An example of such an input signal is one which contains a large number of frequency-division multiplexed (FDM) voice channels. Since it is impossible to design a test that embodies the specific features of all possible applications, NPR has been adopted as a figure of merit for characterizing ADC performance in response to broad bandwidth signals. As explained below, the test leads to a number, the maximum NPR, by comparison of measured data to an ideal curve.

Analog-to-digital converters possessing measured NPRs that closely match theoretical NPR, for an ideal *N*-bit device, are desirable candidates for broadband signal applications, e.g., a signal containing many FDM channels.

Using a notch-filtered broadband white-noise generator as the input to the ADC under test, the NPR is defined as the ratio of the average out-of-notch to the average in-notch power spectral density magnitudes for the DFT spectrum of the ADC output sample set.

8.10.1 Test method for noise power ratio

Use an arbitrary waveform generator (AWG) or a noise generator as the input test signal. Use the test setup shown in Figure 5. The normal procedure is to create a curve in a series of steps which proceeds as follows. A random noise process is generated such that it possesses an approximately uniform spectrum up to a chosen cutoff frequency, f_{co}, which is less than half the sampling frequency. A narrow band of frequencies is then removed from the noise using a notch filter, or preferably, the AWG pattern is tailored to remove the signals in the notch. To obtain a meaningful measurement, the depth of the notch must be at least 10 dB greater than the NPR value being measured. In addition, the width of the notch shall be narrow compared to the overall noise bandwidth. With the notched noise applied to the ADC input, the frequency spectrum of a captured code sequence is computed. See Figure 17 for an example spectrum and Figure 18 for the time domain signal equivalent. The NPR is then calculated, in decibels, from:

$$\mathrm{NPR} = 10\log_{10}\left(\frac{P_{No}}{P_{Ni}}\right) \quad \mathrm{dB} \tag{62}$$

where

P_{No} is the average power spectral density outside the notched frequency band
P_{Ni} is the average power spectral density inside the notched band

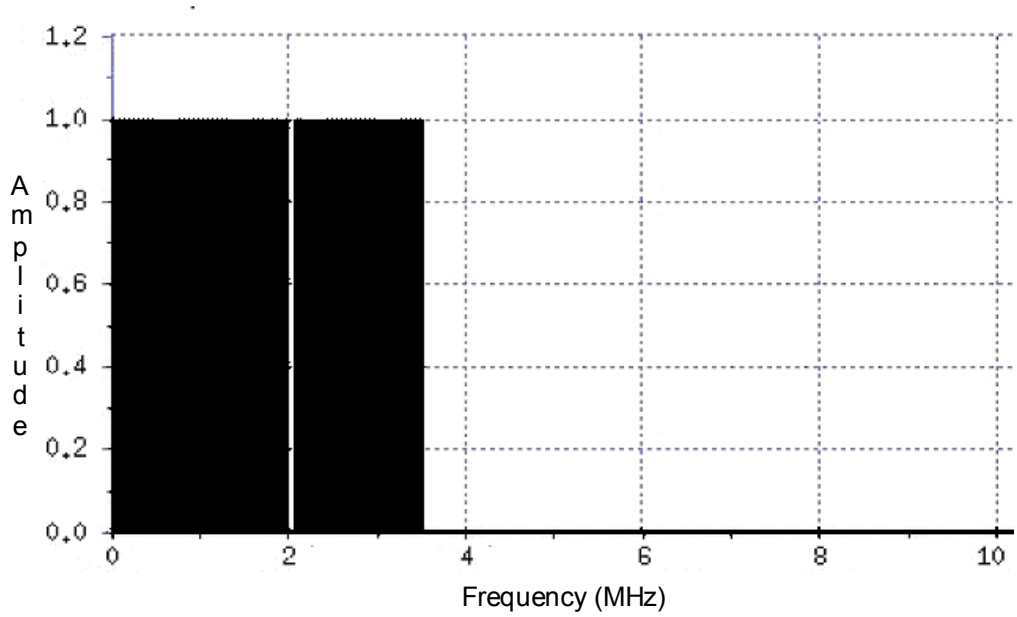

Figure 17—Spectrum of NPR test signal

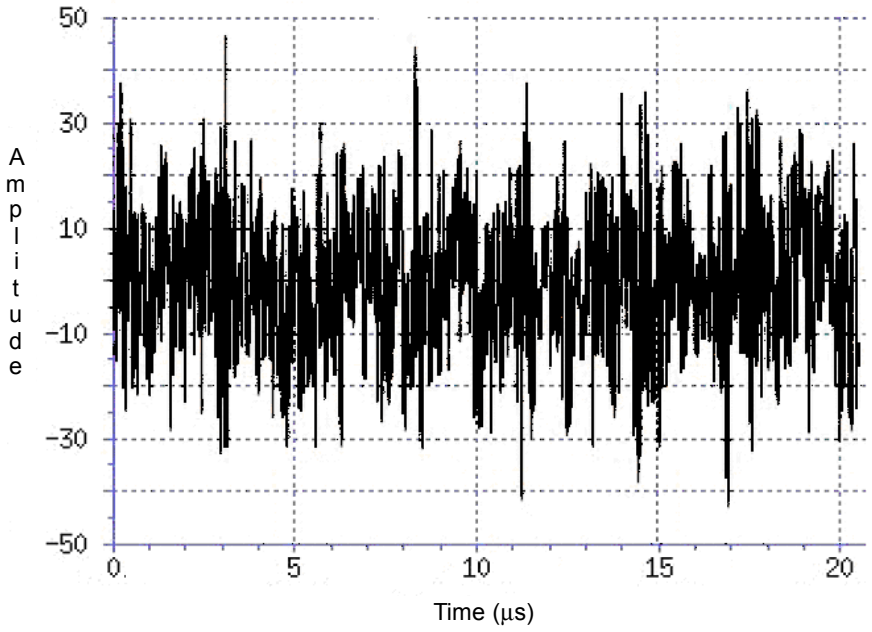

Figure 18—Time domain representation of NPR test signal

Table 3 lists theoretical maximum NPR, NPR_{max}, values for ideal *N*-bit quantizers for both Gaussian and uniform random noise sources. These values were obtained from Equation (63) and Equation (64) given in 8.10.2.4. The parameter α is the signal input level relative to the FSR in decibels.

Table 3—Maximum NPR for Gaussian and uniform noise sources

Source	Uniform		Gaussian	
Number of Bits	α (dB)	NPR_{max} (dB)	α (dB)	NPR_{max} (dB)
6	−4.65	36.20	−10.31	29.94
8	−4.77	48.16	−11.87	40.60
10	−4.77	60.21	−13.04	51.56
12	−4.77	72.25	−14.02	62.71
14	−4.77	84.29	−14.80	74.01
16	−4.77	96.33	−15.45	85.40
18	−4.77	108.37	−16.04	96.88
20	−4.77	120.41	−16.62	108.41

8.10.2 NPR testing issues

Some experimentation with the following test procedures may be necessary to obtain reliable measures of the NPR.

8.10.2.1 Input signal filtering

In practice, it is usually necessary to low-pass filter the input noise signal to prevent aliasing and to obtain a uniform noise power across the input signal spectrum. When the noise bandwidth is low-pass filtered to obtain a bandwidth less than the Nyquist frequency, the peak data listed in Table 3 are valid for the Gaussian input signal since they are dependent upon the average input power. However, for the uniformly distributed input, the data deteriorate toward the Gaussian values since, as the signal is low-pass filtered, the convolution of signal with filter response converges toward a Gaussian process as the bandwidth is lowered from the full Nyquist band. If one plots NPR measurements based upon the input prior to low-pass filtering, then the input at maximum NPR will shift according to the bandwidth ratio, in decibels, of the filter cutoff frequency, f_{co}, to the Nyquist frequency, $f_s/2$.

8.10.2.2 Notch filter width

Another factor that affects measured NPR is the width of the notch filter. Assuming that the measured NPR is obtained from DFT spectral estimates, widening of the notch filter and averaging the noise power contained in the DFT bins inside the notch improves the estimated NPR when compared to using a single bin for the estimate of average noise power in the notch. Making the notch too wide, however, degrades the estimated NPR since the assumption for a uniform noise spectrum could be jeopardized.

8.10.2.3 Windowing

For some cases, the depth of the filtered notch may be degraded due to spectral leakage. The use of windowing eliminates this effect at the expense of a small change in the noise floor. See Annex C for additional details on the effects of windowing.

8.10.2.4 Measured and theoretical NPR

It is customary to plot measured and theoretical NPR versus the mean noise power of the input noise process as described by Daboczi [B15] and Irons et al. [B24]. For a non-ideal *N*-bit ADC, measured NPR curves follow theoretical response at small input power levels given that the ADC does not have excess internal

noise. Measured NPR curves normally depart from theory prior to reaching theoretical maximum NPR due to ADC-generated harmonic and intermodulation distortions. It is also true that the measured curve will depart from theory for very small power levels where peak-to-peak signals are less than one LSB of the ADC.

The maximum measured NPR value is used to specify ADC response to broadband signals with a single number, but it is also necessary to specify the type of noise source used for the test. The theoretical NPR equations for Gaussian and uniform distribution signals are given in Equation (63) and Equation (64) respectively. See Irons et al. [B24] for more information.

$$\mathrm{NPR}_G = \frac{\alpha^2}{\left[\frac{2^{-2N}}{3} - \frac{2\alpha}{\sqrt{2\pi}} e^{-\frac{1}{2\alpha^2}} + \left(1+\alpha^2\right)\frac{2}{\sqrt{\pi}} \int_{\frac{1}{\sqrt{2}\alpha}}^{\infty} e^{-v^2} dv\right]} \tag{63}$$

$$\mathrm{NPR}_U = \frac{\left(\sqrt{3}\alpha\right)^3}{\left[\left(\sqrt{3}\alpha\right)2^{-2N} + \left(\sqrt{3}\alpha - 1\right)^3 u\left(\sqrt{3}\alpha - 1\right)\right]} \tag{64}$$

where

α is equal to 2 $V_{\text{in rms}}$/FSR
$V_{\text{in rms}}$ is the rms of the actual input to the ADC
N is the number of ADC bits
u is the unit step function
v is the dummy integration variable

Example plots of Equation (63) and Equation (64) compared with simulated measured data are shown in Figure 19(a) and Figure 19(b), respectively, for an 8-bit quantizer. An example simulated broadband signal for a 12-bit AWG-generated notch-filtered uniform distribution spectrum is shown in Figure 20(a). The plot has several features that should be noted. It is symmetric about its center, $f_s/2$; the middle notch is an anti-aliasing filter with cutoff f_{co}; and the left notch is the NPR measuring filter. The corresponding histogram for a 16K-sample set is shown in Figure 20(b). Note that this histogram looks nearly uniform, but the rounding is due to spectral convolution with the anti-aliasing and notch filters.

8.10.2.5 Comments on NPR

Insight into the NPR response can be obtained by considering an ideal N-bit ADC. The mean-squared quantization error for such a device is $Q^2/12$. Assuming that the input signal does not saturate the ADC, the quantization error is independent of the input power level. This ideal quantization noise exhibits a uniform spectrum with the noise power evenly distributed over the full Nyquist band. When the input noise power is greater than the quantization power, an increase of 1dB in input power yields a 1 dB increase in NPR, since the quantization power spectral density in the notch remains constant. The linear slope of the NPR curve is thus maintained as long as the ADC operates within its unsaturated signal range and other nonlinearities are not present in the ADC's response to the broadband test signal.

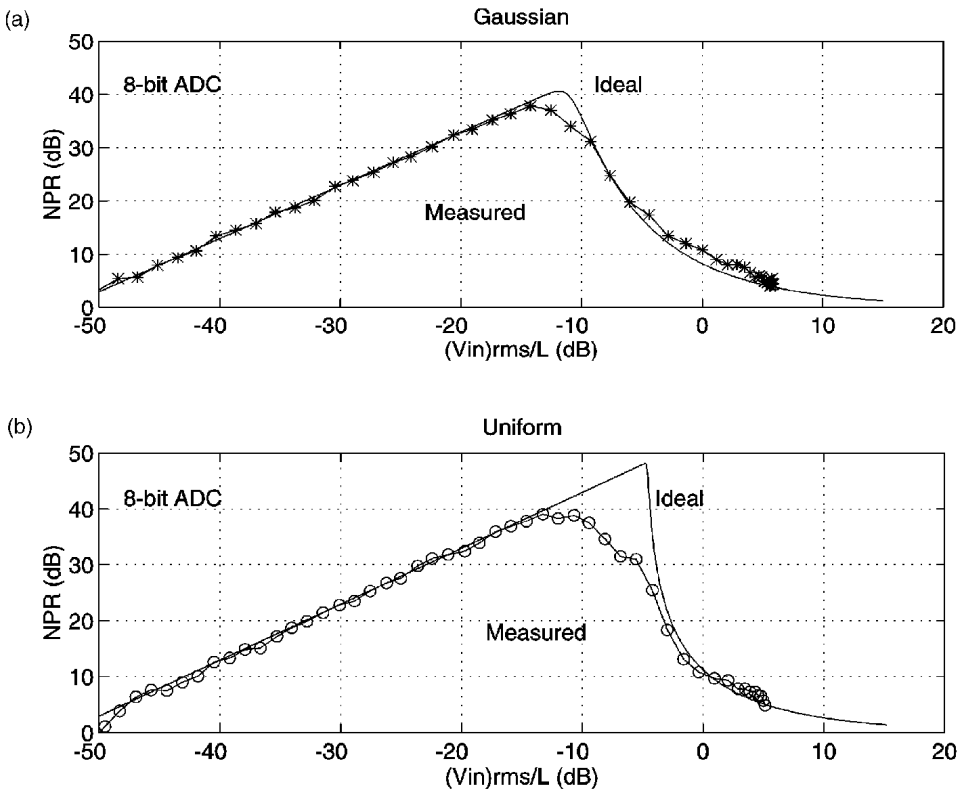

Figure 19—(a) Plot of ideal and measured NPR for a Gaussian noise input from Equation (63); (b) Plot of ideal and measured NPR for a uniform noise input from Equation (64)

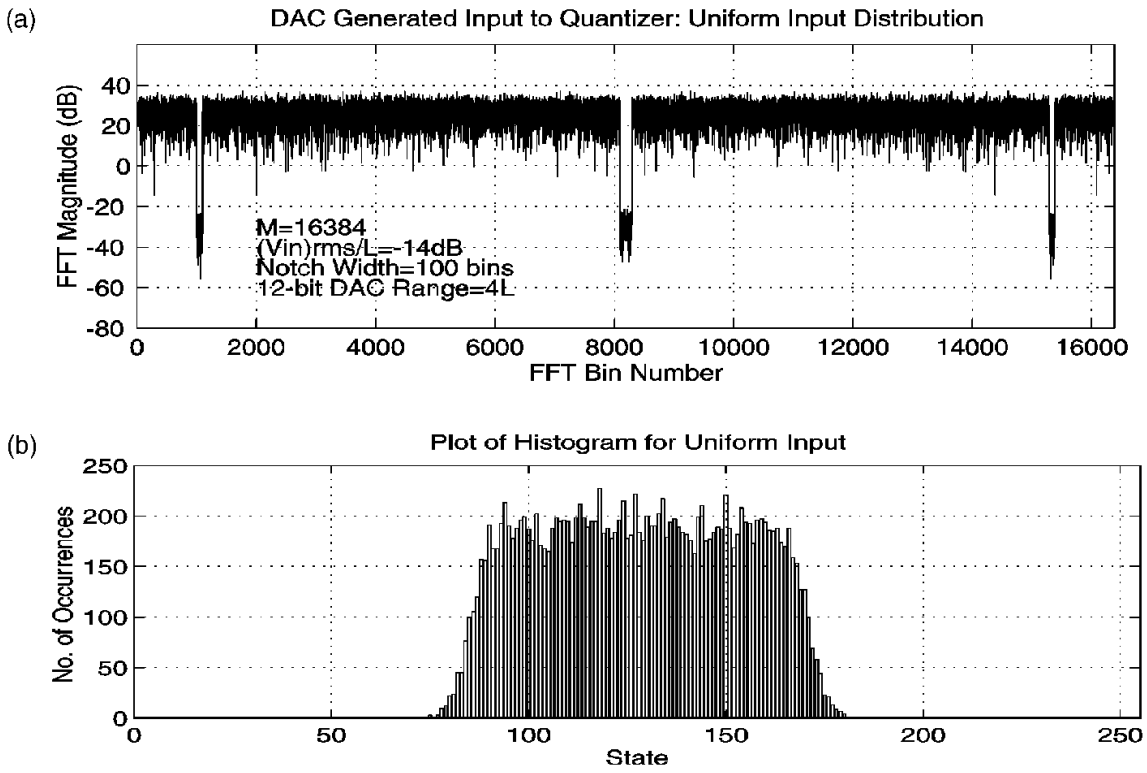

Figure 20—(a) DAC-generated input to quantizer; (b) Histogram from 16K-sample set

9. Noise (total)

9.1 General comments concering noise

Noise has historically been an ambiguous term. In this standard, noise is any deviation between the output signal (converted to input units) and the input signal except deviations caused by linear time invariant system response (gain and phase shift), a dc level shift, or total harmonic distortion. For example, noise includes the effects of random errors, nonlinearities producing harmonics at frequencies greater than those used in measuring total harmonic distortion, quantization errors, spurious signals, and aperture uncertainty.

9.2 Signal-to-noise-and-distortion ratio (SINAD)

Signal-to-noise-and-distortion ratio (SINAD) is the ratio of root-mean-square (rms) signal to rms noise and distortion (NAD). Unless otherwise specified, SINAD is measured using sine-wave input signals. SINAD depends on the amplitude and frequency of the applied sine wave. The amplitude and frequency at which the measurement is made shall be specified.

9.2.1 SINAD test method

To estimate SINAD, apply a sine wave of specified frequency and amplitude to the ADC input. A large signal is preferred. The frequency of the input sine wave is called the fundamental frequency. Almost any error source in the sine-wave input other than gain accuracy and dc offset can affect the test result. It is recommended that a sine-wave source of good short-term stability be used and that the sine-wave input be highly filtered to remove distortion and random noise from the input signal.

Take a record of data: To find NAD, fit a sine wave to the record at the fundamental frequency as per 5.2. Compute the estimate of the rms value of the noise and distortion as shown in Equation (65).

$$\mathrm{NAD} = \sqrt{\frac{1}{M}\sum_{n=1}^{M}(x[n]-x'[n])^2} \tag{65}$$

where

$x[n]$ is the sample data set
$x'[n]$ is the data set of the best-fit sine wave
M is the number of samples in the record

The signal-to-noise-and-distortion ratio, SINAD, is given by Equation (66)

$$\mathrm{SINAD} = \frac{A_{\mathrm{rms}}}{\mathrm{NAD}} \tag{66}$$

where

$$\mathrm{A_{rms}} = \frac{\text{Sine Wave Peak Amplitude}}{\sqrt{2}}$$

The values of $x[n]$, $x'[n]$, and $\mathrm{A_{rms}}$ must all be in the same units, typically either input units or LSB.

9.2.2 Coherent sampling test method for SINAD in the frequency domain

SINAD can be determined equivalently from the frequency domain as a consequence of Parseval's Theorem. Apply an appropriate sine wave as described in the test procedures above in 4.2.1. Compute the DFT of the measured waveform. SINAD is the ratio of the rms input signal to the rms noise and distortion (NAD). Both quantities can be determined from the DFT of data records as was done for THD in 8.8.1.1. The rms signal, A_{rms}, is obtained from Equation (51). NAD is found from the sum of all the remaining Fourier components after the bins at dc and at the test frequencies have been deleted from the spectrum.

$$\text{NAD} = \frac{1}{\sqrt{M(M-3)}} \sqrt{\sum_{m \in S_0} X_{\text{avm}}[f_m]^2} \tag{67}$$

where

S_0 is the set of all integers between 1 and $M - 1$, excluding the two values that correspond to the fundamental frequency and the zero-frequency term
X_{avm} is the averaged spectral magnitude, as defined in Equation (52)

SINAD is then given by substitution of NAD and A_{rms} into Equation (66).

9.2.3 Comments on SINAD and SNR

These ratios are both proportional to the test signal, A_{rms}, and it is customary to use a near full-scale signal for these measures. However, if clipping should occur, the measures will be severely degraded. In addition, these measures are generally a function of amplitude, X_{avm} (f_i), and the frequency, f_i, of the input sine wave. Thus, the amplitude and frequency of the input, for which the SINAD and/or SNR measurements are made, shall be specified.

9.3 Signal-to-noise ratio (SNR)

The SNR is the ratio of the rms signal to the rms noise for a sine-wave input signal of a specified frequency and amplitude. The rms noise is determined by determining the rms noise and distortion as described in 9.2.1 or 9.2.2 and then determining the distortion as described in 8.8.1.1, 8.8.1.2, or 8.8.1.3. The SNR is given by

$$\text{SNR} = \frac{A_{\text{rms}}}{\eta}, \tag{68}$$

where

A_{rms} is the rms signal as determined in one of the following clauses and
η is the rms noise as determined in one of the following clauses

9.3.1 Coherent sampling test method for SNR

Determine the NAD as described in 9.2.2. Determine A_{rms} and THD as described in 8.8.1.1. Let

$$\eta = \sqrt{\text{NAD}^2 - A_{\text{rms}}^2 \text{THD}^2} \tag{69}$$

The SNR is defined by Equation (68).

9.3.2 Noncoherent sampling test method 1 (windowed DFT)

Determine the NAD as described in 9.2.1. Determine A_{rms} and THD as described in 8.8.1.2. Use Equation (69) and Equation (68) to determine SNR.

9.3.3 Noncoherent sampling test method 2 (sine fitting)

Determine the NAD as described in 9.2.1. Determine A_{rms} and THD as described in 8.8.1.3. Use Equation (69) and Equation (68) to determine SNR.

9.4 Effective number of bits (ENOB)

For an input sine wave of specified frequency and amplitude, after correction for gain and offset, the effective number of bits (ENOB) is the number of bits of an ideal ADC for which the rms quantization error is equal to the rms noise and distortion of the ADC under test. ENOB is given by:

$$\text{ENOB} = \log_2\left(\frac{\text{FSR}/G}{\text{NAD}\sqrt{12}}\right) \approx N - \log_2\left(\frac{\text{NAD}}{\varepsilon_Q}\right) \tag{70}$$

where

N is the specified number of bits in the ADC
FSR is the specified full-scale range of the ADC
G is the measured gain (nominally = 1) and is defined in Equation (35) or Equation (37)
NAD is rms noise and distortion and is defined in Equation (67)
ε_Q is the rms ideal quantization error

The quantity ENOB depends on the amplitude and frequency of the applied sine wave. The amplitude and frequency at which the measurement was made shall be specified.

9.4.1 Comment on ideal quantization error

The input signal value corresponding to an ADC output code is best assumed to be the center of that code's bin. An input signal falling into a code bin not at the center generates quantization error amounting to the difference of the signal value from the center of the bin. To evaluate the rms of this error over many samples, the probability distribution of the signal over a code bin must be known.

9.4.2 Comment on the relationship of SINAD and ENOB

SINAD and ENOB are related by Equation (71):

$$\begin{aligned}\text{ENOB} &= \log_2(\text{SINAD}) - \frac{1}{2}\log_2(1.5) - \log_2\left(\frac{2G\times A}{\text{FSR}}\right)\\ &= \log_2(\text{SINAD}) - \log_2\left(\frac{2G\times A}{\text{FSR}}\right) - 0.292\end{aligned} \tag{71}$$

or equivalently by Equation (72)

$$\text{SINAD} = \left(\sqrt{1.5}\right)\times\left(\frac{2G\times A}{\text{FSR}}\right)\times 2^{\text{ENOB}} \tag{72}$$

where

G is the measured gain (nominally = 1) and is defined in Equation (35) or Equation (37)
A is the amplitude of the fitted sine wave during the test in the same units as FSR
FSR is the full-scale range of the ADC input

9.4.3 Comment on significance of record length

The effect of random noise on the results of sine-wave tests is decreased when the record size is increased. Thus, longer record sizes lead to more reproducible results. The record size should not be made so large that frequency drift or close-in phase noise in either the sine-wave source or the ADC clock signal affects the results, or the test time becomes inordinately long. See 5.4 for further discussion.

9.4.4 Comment on effects of jitter or phase noise on sine-wave tests

Time jitter (also called phase noise in the frequency domain) in the sine-wave signal source produces both random and systematic errors for sine-wave tests. A consequence of jitter (see Souders et al. [B53]) is that it spreads the energy of the original sine wave over a broad spectrum of frequencies, reducing the amplitude of the fundamental component (for $\sigma_t << 1/f$) approximately by the factor shown in Equation (73).

$$\text{Amplitude multiplicative factor} = 1 - \frac{(2\pi f \sigma_t)^2}{2} \tag{73}$$

where

f is the signal frequency fundamental component, in hertz
σ_t is the standard deviation of the jitter, in seconds

The energy lost in the fundamental component shows up as broadband noise that has an rms value (computed over a complete period of the input sine wave) given by Equation (74)

$$V_{\text{noise}} \approx V_p \frac{2\pi f \sigma_t}{\sqrt{2}} \tag{74}$$

where V_p is the sine-wave peak amplitude.

The jitter-induced noise is distributed according to the time-derivative of the signal, approaching zero at the sine-wave peaks, and reaching a maximum at the zero crossings given approximately by Equation (75).

$$\sigma_{\max} = 2\pi f V_p \sigma_t \tag{75}$$

where $\sigma_{\max}$ is the maximum value of the standard deviation of the amplitude noise, occurring at the zero crossings.

If a sine fit is performed on the sampled sine wave with jitter, the amplitude of the estimated sine wave will be reduced by the factor given in Equation (73). Furthermore, if repeated acquisitions of the waveform are averaged, the result will be a sine wave with amplitude reduced by the same factor. Note that if the original sine wave were measured using a “true rms” responding instrument, e.g., an instrument that uses thermal transfer techniques, the measured value will NOT be reduced by this factor; this is because the total energy is not changed by jitter, it is only redistributed.

Jitter in the sine-wave source limits the signal-to-noise ratio and the number of effective bits that can be measured. Substituting Equation (74) for rms noise in the effective bits formula [Equation (76)] gives

$$\mathrm{ENOB} = \log_2\left[\frac{\mathrm{FSR}}{V_{\mathrm{noise}} \times \sqrt{12}}\right] \tag{76}$$

and, for a large signal sine wave (i.e., $V_p \times$ FSR/2), this gives [Equation (77)]

$$\mathrm{ENOB} = -0.7925 - \log_2\left[\frac{2\pi f\, \sigma_t}{\sqrt{2}}\right] = -0.2925 - \log_2\left[2\pi f \sigma_t\right] \tag{77}$$

Table 4 gives the maximum effective bits that can be measured as a function of the jitter standard deviation, expressed as a fraction of the sine-wave period.

Table 4—Maximum effective bits versus normalized jitter (fraction of sine-wave period)

$\sigma_t f$	10^{-1}	10^{-2}	10^{-3}	10^{-4}	10^{-5}	10^{-6}
$\mathrm{ENOB}_{\mathrm{max}}$	0.4	3.7	7.0	10.3	13.7	17.0

It is unusual to find jitter specified for sine-wave sources; instead, phase noise is more commonly specified. Unfortunately, it is not a simple task to compute jitter from typical phase noise specifications. An alternative approach is to measure the jitter directly. Wideband sampling oscilloscopes, e.g., often provide simple procedures for measuring signal jitter, with resolution that is usually adequate for most ADC applications.

9.4.5 Effects of harmonic distortion on sine-wave tests

Harmonic distortion in the sine-wave signal source can cause direct errors in measurements of SINAD, effective bits, and total harmonic distortion. In the worst case, the distortion of the ADC under test is dominated by a single harmonic component (frequency) of amplitude A_H, and the same harmonic component (and same phase) dominates the distortion of the signal source, with smaller amplitude, B_H. The true total harmonic distortion of the ADC is given by 20 log10(A_H/A), where A is the amplitude of the fundamental component. On the other hand, the measured total harmonic distortion, assuming $B_H << A_H$, is given approximately by 20 $\log_{10}$(A_H(1+B_H/A_H)/A). The difference between measured and true THD is given in Table 5 for several values of B_H/A_H. Similar results can be computed for SINAD and effective bits measurements, and example errors for these parameters are also included in Table 5.

Table 5—Worst-case change in measured performance

B_H/A_H	0.25	0.125	0.1	0.0625	0.05	Units
SINAD	−1.94	−1.02	−0.83	−0.53	−0.42	dB
THD	1.94	1.02	0.83	0.53	0.42	dB
ENOB	−0.32	−0.17	−0.14	−0.09	−0.07	bits

More common but less-serious cases occur when the distortion is spread out over many frequencies, and especially when the distortion in the signal source occurs at frequencies that are different from the major distortion frequencies of the ADC under test. In those cases the components combine orthogonally (i.e., root-sum-squares). In this case, if B_H is the amplitude of the major distortion component of the sine-wave generator, and it is a different frequency from that of the major distortion component (with amplitude A_H)

of the ADC, then the measured THD is given approximately by $20\log10(A_H(1 + B_H^2/2A_H^2)/A)$. Table 6 reports errors in SINAD, THD, and ENOB that occur under these less-serious conditions.

Table 6—Change in measured performance, assuming orthogonal components

B_H/A_H	0.25	0.125	0.1	0.0625	0.05	units
SINAD	−0.27	−0.07	−0.04	−0.02	−0.01	dB
THD	0.27	0.07	0.04	0.02	0.01	dB
ENOB	−0.04	−0.01	−0.01	−0.00	−0.00	bits

As a rule of thumb, a value of B_H/A_H of 0.1 (−20 dB) should be adequate for both cases, and a value of 0.25 (−12 dB) should be adequate if the components are orthogonal. It is relatively easy to measure the distortion of a sine-wave source using a spectrum analyzer, provided that the distortion in not lower than −80 dB. If lower levels of distortion are required, then carefully designed notch filters should be used to reject the fundamental component, passing only the distortion components on to the spectrum analyzer. Care must be taken to account for any attenuation of the harmonic components caused by the filter. If the signal source does not have adequate spectral purity by itself, it can be improved with low-pass or band-pass filters. THD values of −80 dB can usually be achieved with relatively inexpensive commercial filters. Achieving significantly lower THD values usually requires specially designed filters (including the notch filters used for verification) constructed with linear, passive components. Amplifiers and iron core inductors, for example, often generate additional distortion that is difficult to remove.

9.5 Random noise

Random noise is a nondeterministic fluctuation in the output of an ADC, typically described by its frequency spectrum and its amplitude statistical properties. For the measurements in this clause, the following noise characteristics are assumed: The amplitude probability density function is stationary and has zero mean. (A nonzero mean is the same as offset error.) Also the noise is assumed to be additive. Random noise test results shall include a description of the impedance at the input to the ADC.

9.5.1 Test method

Set up the ADC under test with a static input signal, of specified output impedance, having noise level at least four times less than the level of accuracy required for the ADC random noise measurement. The value(s) of the static input signal(s) shall be reported. Take two records of data M-samples long, and subtract one from the other; the subtraction eliminates fixed-pattern errors that occur in the same location in successive records. In Equation (78) the noise variance may be estimated from

$$\delta = \sigma^2 = \frac{1}{2M}\sum_{n=1}^{M}\left(x_a[n] - x_b[n]\right)^2 \tag{78}$$

where

δ is the mean squared difference between the two test records
σ^2 is the mean squared noise
$x_a[n]$, $x_b[n]$ are the samples from the two noise records
M is the number of samples in each record

When the noise is $Q/2$ or less, the method above can produce either an underestimate or an overestimate of the noise. If the signal is near the center of a code bin, the noise will not affect the result and an underestimate will be obtained. If the signal is near the boundary of two code bins, the recorded value will

randomly toggle between two adjacent values and give an overestimate. To overcome these difficulties use the method in 9.5.2.

9.5.2 Alternative test method for low noise ADCs

Connect the output of a triangle wave generator to the signal input of the ADC. Adjust the output amplitude to about ten code transition levels peak-to-peak (see 9.5.3). Trigger the ADC on the beginning of the positive-going portion of the triangle. Adjust the frequency of the triangle wave generator such that one period of the triangle wave subtends one record length. The record length shall be commensurate with the desired measurement accuracy (see 9.5.3). Capture two records, and find their difference as given above in 9.5.1. Use Equation (79) and Equation (80) to get the noise variance σ^2:

$$\text{mse} = \frac{1}{M}\sum_{n=1}^{M}(y_{an} - y_{bn})^2 \tag{79}$$

$$\sigma^2 = \frac{1}{\sqrt{\left(\frac{2}{\delta}\right)^2 + \left(\frac{Q}{0.886\delta}\right)^4}} \tag{80}$$

where

δ	is the mean square difference between the samples in the two test data records and is given by Equation (78)
σ^2	is the mean squared noise
$x_a[n]$ and $x_b[n]$	are the samples from the two noise records
M	is the number of samples

As random noise increases, these equations converge to that used in 9.5.1. For a derivation of Equation (80), see IEEE Std 1057-2007, Annex D.

Information about the precision of the estimates of random noise can be found in Alegria and Cruz Serra [B2].

The analysis above is for noise whose amplitude can be described by a Gaussian probability density function (PDF). Equation (80) can be modified for other PDFs. For example, for a uniform PDF the factor 0.866 changes to 0.886, and for a bimodal PDF the factor is 1.000.

9.5.2.1 Note on amplitude of triangle wave used for test

The triangle wave provides a means of slowly slewing the ADC over multiple code bin thresholds at a relatively constant rate. The subtraction process removes the contribution of the triangle wave to the result to the extent that the two repetitions are identical. Any differences due to noise, jitter, etc. will contribute to the apparent result. Consequently, unless the output of the generator can be independently judged to have a sufficiently low noise level, it is best to keep the amplitude low. This means that only a part of the full-scale range of the ADC can be explored with each measurement.

9.5.2.2 Note on desired accuracy

The standard deviation of an estimate of random noise standard deviation is given by

$$\sigma_\sigma = \frac{\sigma}{\sqrt{M\text{ -}1}} \tag{81}$$

where

σ_σ is the standard deviation of the estimate of random noise amplitude
σ is the random noise standard deviation
M is the number of independent random noise samples

9.5.3 Alternative random noise and hysteresis test method based on a feedback loop

An alternative method for measuring random noise, hysteresis, and alternation bands employs the feedback loop method shown in Figure 9(a) and Figure 9(b). The method often becomes impractical when conversion times are significantly longer than 10 μs. The feedback loop technique is shown in the figures for converters with voltage inputs. The method is easily extended to converters with other forms of input. The parameters N_1 and N_2 are used to define the magnitude of the change of the DAC input if the ADC output, k, is less than or greater than or equal to the code "k_{in}". The clock signals "Trig" and "Trig1" can generally be identical to each other.

This test method can determine the parameters of the random noise at a given code transition level of the ADC under test. The feedback loop works to find the transition level $T[k_{in}]$ at the lower edge of the code bin whose code is the reference code value "k_{in}". The stable dc source is set to a value near the value which corresponds to that code. The stable source must have a noise level that is significantly less than that expected from the ADC under test. If $N_1 = N_2 = N$, and the change in the DAC output generated by a change in input code of N, is less than the standard deviation of the random noise being measured, then the voltage at the ADC input will adjust itself to a value which will cause the ADC output code to be greater than k_{in} 50% of the time.

Changing the ratio of N_1 to N_2 will force a change in the duty cycle of the codes being produced by the ADC under test. Table 7 illustrates how the average ADC input voltage is affected by a change in the ratio N_1/N_2. The tabulated values assume that the equivalent noise at the input to the ADC under test is Gaussian, with an rms value of σ.

A typical measurement sequence would involve calculating the average ADC input twice. The first time, the ratio N_1/N_2 would be set to 2.0. The second time the ratio would be set to 0.5. Table 7 shows that the difference of the two measurements will be $2 \times 0.431\sigma$. The value of σ can be evaluated by dividing the difference in the voltage readings by 0.862.

Table 7—ADC input versus N1/N2 for alternate noise measurement method

N_1/N_2	ADC_{INPUT}	N_1/N_2	ADC_{INPUT}	N_1/N_2	ADC_{INPUT}
0.2	$x + 0.967\sigma$	1.0	x	5.0	$x - 0.967\sigma$
0.25	$x + 0.841\sigma$	2.0	$x - 0.431\sigma$	10.0	$x - 1.335\sigma$
0.3333	$x + 0.674\sigma$	3.0	$x - 0.674\sigma$	15.0	$x - 1.534\sigma$
0.5	$x + 0.431\sigma$	4.0	$x - 0.841\sigma$	20.0	$x - 1.668\sigma$

The circuits whose block diagrams are shown in Figure 9 can be used for evaluating the size of the hysteresis or the alternation band. The spectral content of the ADC input is measured while $N_1 = N_2 = N$. The spectrum can be measured by performing a windowed Fourier analysis of the record recorded in the test sequence described in 4.3. A spectrum analyzer can be used to monitor the ADC input. If the ADC is ideal, there will be no obvious frequency lines noted at less than f_s. If an alternation band is present there will be a significant component at $f_s/2$. If hysteresis is present there will be a lower frequency component generated by a ramp spanning the hysteresis band. Hysteresis or alternation will become apparent only if the noise of the converter is smaller than the hysteresis or the alternation band.

10. Step response parameters

10.1 Step response definition

The step response is the recorded response of the ADC under test to a perfect step input signal with designated base state and high state.

10.2 Test method for acquiring an estimate of the step response

Use the step signal test setup (Figure 6). Use a suitable input step signal generator: in lieu of a perfect step, a suitable input step signal is one that has transition duration, overshoot, and settling time no greater than one-fourth of those expected from the ADC under test. Using this 4:1 performance ratio as a guideline, one would expect approximately a 3% error in the transition duration estimate. For overshoot and settling time, the uncertainty in the input step parameters will add equal uncertainty in the corresponding estimated ADC quantities. If smaller uncertainties are required, then deconvolution techniques may be applied to remove the input step signal imperfection (see Daboczi [B15]).

If possible, before and after acquiring the actual step response data, determine the (static) initial and final values of the step, which are respectively the base state and high state or the high state and base state. If the pre-acquisition measurements of the initial and final values differ significantly from the post-acquisition measurements of the initial and final values (due to system drift during the data acquisition), estimate the initial and final values of the step as the average of their respective pre-acquisition and post-acquisition measurements. Acquire a record or records of samples of the ADC step response that are sufficiently long to include all features of interest, e.g., precursors, and electrical and thermal settling. In order to produce high-resolution sampling of step response features and reduce errors due to aliasing, it may be necessary to use the method of equivalent-time sampling, as described in 4.4. To reduce the effects of random noise and aperture uncertainty, it is recommended that multiple records of the step response be ensemble averaged.

Significant error sources in acquiring the step response estimate include aliasing, jitter, random noise, ADC nonlinearities (such as slew rate limiting), input step imperfections, and inaccuracies or imprecision of the frequencies used to control the ADC and the step generator during equivalent-time sampling (see Souders and Flach [B50] for more information).

The following subclauses describe how to determine ADC parameters from measured step response data. The initial and final values of the applied input step signal shall be specified in the test results.

10.2.1 Comment on test results

On some ADCs, typically those that do not employ sharp cutoff anti-aliasing filters, the step response can be nonlinear when the slew rate approaches the slew rate limit (10.3). The value of this slew rate limit is dependent upon the ADC under test. The degree of resulting nonlinearity of the step response increases with the steepness of the applied step. Because of this nonlinearity, the measured step response of such an ADC can be misleading. To eliminate the gross nonlinearities, the slew rate of the applied input step must be sufficiently below the slew rate limit. The statement of results shall include the transition duration and amplitude of the applied step.

10.3 Slew rate limit

The slew rate limit is the value of output transition rate of change for which increased amplitude input step causes no change.

10.3.1 Test method

Record the step response (see 10.2) for an input step having amplitude 10% of full scale. Determine and store the maximum rate of change of the output transition. Repeat this process, increasing the amplitude of the input step each time. When the maximum rate of change ceases to increase with increasing step amplitude, slew limiting is taking place and the slew rate limit is the largest recorded value for the maximum rate of change.

10.4 Settling time parameters

10.4.1 Settling time

Measured from the 50% reference level instant of the output transition, the settling time is the time at which the step response enters and subsequently remains within a specified error band around the final value. The final value is defined to occur 1 s after the beginning of the step unless otherwise specified.

10.4.2 Short-term settling time

Measured from the 50% reference level instant of the output transition, the short-term settling time is the time at which the step response enters and subsequently remains within a specified error band around the final value. The final value is defined to occur at a specified time less than 1 s after the beginning of the step.

10.4.3 Long-term settling error

The long-term settling error is the maximum absolute difference between the final value specified for short-term settling time and the value 1 s after the beginning of the step, expressed as a percentage of the step amplitude.

10.4.4 Test method for settling time and short-term settling time

Record the step response (as per 10.2) to an input step using a record length sufficient to represent the step over the duration specified, or for at least 1 s when the duration is not specified. Two or more overlapping records with different sample rates may be required to achieve the necessary time resolution and the required duration. To reduce noise or quantization errors, it may be desirable to digitally filter the step response data before computing settling time parameters. For example, apply a moving average filter of the form as shown in Equation (82).

$$y_n = \frac{1}{(2r+1)} \sum_{s=-r}^{r} x_{n-s} \tag{82}$$

where

x_{ns} is the value of the $(ns)^{th}$ data point of the unfiltered step response
y_n is the value of the nth data point of the filtered step response
r is an integer defining the width of the moving average window

If such a filter is used, the width of its window, $(2r + 1)$, shall be specified in the test results.

Determine the time of occurrence of the first 50% point on the transition of the recorded waveform. Counting from that time, the settling time (or the short-term settling time) is the time at which the output

waveform last enters the bound given by $V(t) \pm e$, where $V(t)$ is the value at the end of the specified duration and e is the specified error. When the duration is not specified, $V(t)$ is the value 1 s after the beginning of the step.

To measure the long-term settling error, record the same step used to determine the short-term settling time, with a record that spans at least a 1 s interval from the beginning of the step. The long-term settling error is the absolute difference between the value 1 s after the beginning of the step and the value at the end of the specified duration following the step, expressed as a percentage of the step amplitude.

10.4.5 Comment on settling time

The term settling time refers to the time required to settle to the steady state, dc value, to within the given tolerance. The dc value is assumed to be the value after a constant input has been applied for at least 1 s. Changes that occur after 1 s are considered drift, and may be due to room temperature fluctuations, component aging, and similar effects.

The term *short-term settling time* refers to the time required to settle to a relative value (perhaps different from the steady-state value), defined as the value at the end of a specified duration, for record lengths less than 1 s. If static offset, gain, and linearity corrections are used to assign true values to short-term settling data, the results will have an uncertainty given by the long-term settling error. The uncertainty results because of longer-term settling phenomena, such as thermal imbalances that may occur after the short-term duration is complete, but which affect a steady-state measurement.

Note that only short-term settling time can be specified for ac-coupled ADCs.

10.5 Transition duration of step response

The transition duration of the step response is the duration between the 10% point and the 90% point on the recorded output response transition, for an ideal input step with designated base state and high state. The algorithm used to determine the base state and high state of the output step must be defined. The methods of IEEE Std 181-2003 are preferred.

10.5.1 Test method

Record the step response (see 10.2) and determine the 10% and 90% points of the output transition using methods in IEEE Std 181-2003. Linear interpolation is used to determine the 10% and 90% points when insufficient data points are available on the transition. The transition duration of the step response is the time between the first 10% point and the last 90% point on the transition.

10.6 Overshoot and precursors

Overshoot is the maximum amount by which the step response exceeds the high state, and is specified as a percent of the (recorded) pulse amplitude. Precursors are any deviations from the base state prior to the pulse transition. They are specified in terms of their maximum amplitude as a percent of the pulse amplitude.

10.6.1 Test method

Record the step response (see 10.2). Using IEEE Std 181-2003, determine the waveform's post-transition overshoot and the pre-transition overshoot and undershoot. The first value is the ADC step response overshoot, and the other two values are the other two values are the ADC step response precursors.

11. Frequency response parameters

11.1 Bandwidth (BW)

The ADC's bandwidth is the width of the passband of its frequency response. Specifically, it is the difference between the upper and lower −3 dB frequencies, which are the frequencies at which the gain of the ADC is −3 dB of the gain at a specified reference frequency within the passband. The reference frequency is typically chosen as a frequency where the gain is at or near its peak value in the passband. Many ADCs have no lower −3 dB frequency because their passbands extend down to zero frequency (dc); in such cases, their bandwidths are simply the values of their upper −3 dB frequency. In such cases, the reference frequency is often chosen to be zero, so that the reference gain is the dc gain. If a lower −3 dB frequency does exist, the upper and lower −3 dB frequencies shall be specified in the test results along with the bandwidth, instead of the bandwidth alone, which is their difference.

The determination of the large signal upper −3 dB frequency can be nontrivial if an ADC starts to exhibit slew rate-induced nonlinearities below or near this frequency. It is recommended that an alternate figure of merit, useful power bandwidth, be used to describe such ADCs, as determined in 11.1.3. Linear system theory no longer applies to devices at these slew rates, and an upper −3 dB frequency, if it could be uniquely identified, would not be useful in the usual manner.

Below are two methods of determining bandwidth and methods for determining the useful power bandwidth. The first method uses sine-wave inputs, and can be done very quickly if the reference frequency and the approximate limit frequencies are known. The disadvantage of this method is the typically low accuracy of estimates of the analog input amplitudes, which reduces the accuracy of the bandwidth result as well. The second method uses the frequency response as determined by the DFT of the derivative of the step response, from 11.3. Its disadvantages are high noise at higher frequencies, and aliasing and first-differencing errors resulting from the frequency response estimation. The results of using the step response method are invalid in the presence of slew rate induced errors; this method is generally more useful for ADCs which contain analog-bandwidth-limiting circuitry before the quantizer(s). Converters prone to slew rate induced errors shall specify useful power bandwidth and use the third method below.

Bandwidth may be measured at any stated signal amplitude and sampling rate. When the sampling rate is not specified, bandwidth is measured at the maximum sampling rate.

11.1.1 Bandwidth test method

This method uses sine waves of known frequencies and amplitudes to determine bandwidth. A large signal (defined in 3.1) sine wave is used, unless the small-signal bandwidth is to be determined. When small-signal bandwidth is to be determined, the peak-to-peak input amplitude used is less than 1/10 of full scale. Use the sine-wave test setup (see 4.2.1). The input sine-wave source shall produce sinusoids of high spectral purity, harmonic distortion lower than that of the ADC under test, and shall have stable output during the measurement time. The tested input frequencies shall not be sub-harmonics of the ADC sampling rate; such frequencies can produce incorrect answers in this test.

Select an input reference frequency at which the ADC dynamic gain is at or near its peak value in the passband. The reference frequency shall be stated with the test results. Connect the sine generator to the ADC input, set its frequency to the reference frequency, and acquire a sufficient number of data records from the ADC output to determine the maximum peak-to-peak range of the signal, using a three-parameter (B.1) or four-parameter (B.2) sine fit. Using an ac voltmeter or other means, measure the amplitude of the applied input sinusoid at the same reference plane as that represented by the ADC input port. This input amplitude measurement must be done with care if high accuracy is required (see Kinard and Ti-Xiong [B32] and Laug et al. [B34]). If the measured input amplitude parameter is the rms amplitude, A_{rms}, it must be converted to peak-to-peak amplitude, $A_{\text{pk–pk}}$, using Equation (83).

$$A_{\mathrm{pk-pk}} = A_{\mathrm{rms}}\sqrt{2} \tag{83}$$

Divide the peak-to-peak ADC output amplitude by the measured peak-to-peak input amplitude to determine the reference gain.

If the chosen reference frequency is zero, the reference gain is the static gain, as determined in 7.4. Alternatively, to using the static gain as determined in 7.4, use a precision dc signal source to provide a constant input signal. Approximate the dc gain by the constant output signal level minus the measured static dc offset (see 7.4.1), divided by the input dc level.

Once the reference gain is determined, change the input frequency to another value that is not a sampling-rate sub-harmonic. Measure the maximum peak-to-peak range of the recorded data, and divide it by the measured input amplitude to find the gain at this frequency. Repeat this as necessary to find the upper (and, if it exists, lower) frequency closest to the reference frequency, at which the gain is 3 dB below the reference gain. If no lower −3 dB frequency exists, the upper −3 dB frequency is the bandwidth. If a lower −3 dB frequency exists, the difference between the upper and lower −3 dB frequencies is the bandwidth of the ADC.

11.1.2 Alternative bandwidth test method using time domain techniques

This method is used to determine the ADC bandwidth via the ADC frequency response determined in 11.3.1. It is desirable to have as many samples in the record as possible, to increase the resolution with which the bandwidth can be resolved from the DFT of the derivative of the step response. The sampling rate, or equivalent-time sampling rate, shall be high enough to make aliasing errors negligible (see 11.3.2).

Choose the reference frequency f_{ref} from the DFT bins $f_k = k/(MT_s)$: choose the one within the passband such that the dynamic gain is at or near the peak gain of the passband. The reference frequency shall be stated. Next, search the DFT bins to find the upper and, if applicable, the lower frequency samples closest to the reference frequency, at which the gain is −3 dB below the reference gain. The bandwidth is the difference between these upper and lower −3 dB sample frequencies (or, if a lower −3 dB frequency does not exist, the bandwidth is simply the upper −3 dB frequency value). To improve the bandwidth estimate, interpolate between the frequency samples above and below −3 dB in amplitude, to better estimate the actual −3 dB frequency.

11.1.3 Useful power bandwidth test method

The useful power bandwidth of a device is the large signal analog input frequency at which a record of the ADC's output data will be degraded by less than a specified amount. The type of degradation used to denote the useful bandwidth is dependent upon the architecture of the ADC and should be chosen on the basis of the type of slew rate degradation which the ADC typically exhibits. ADCs whose architecture begins to show spurious sparkle codes at high slew rates shall use the frequency at which these sparkle codes begin to appear, to a specified confidence level. ADCs that begin to start missing codes shall specify either a no-missing-codes power bandwidth or an equivalent metric such as SNR greater than a specified number of decibels. A compromise measure of degradation, SINAD greater than a specified number of decibels, is most useful with ADCs that simultaneously show small-amplitude spurious sparkle codes and missing codes.

To test for useful power bandwidth at a stated frequency when a converter exhibits spurious sparkle codes, apply a large signal (90% of full-scale range or greater) sine wave at the frequency using the test setup of 4.2.1. Perform repeated SINAD tests, using the methods of 9.2.1 or 9.2.2 and storing the results of each test. The number of tests to perform is governed by both the number of samples used in each test and the desired confidence level of the results. The set of test results is then subjected to a statistical test for uniformity: the variance of the test results is computed and compared to a limit. ADCs exhibiting spurious sparkle codes in one to a few of the SINAD tests will result in a variance which is larger than the limit,

while ADCs exhibiting no sparkle codes will show test result variance below the limit. While ADCs always exhibiting sparkle codes will also meet the variance limit comparison, a simultaneous limit on the value of the test result will enable detection of such devices. Optionally, ADCs failing the variance comparison or the test value limit can be retested at a lower applied input frequency to see whether they meet a new, lower frequency, useful power bandwidth test.

To test for useful power bandwidth at a stated frequency when a converter exhibits missing codes, apply a large signal sine wave at the frequency to be tested, using the test setup of 4.2.1. Perform an SNR test using the methods of 9.2.2, and compare the test result to the stated limit.

To test for useful power bandwidth at a stated frequency when a converter exhibits missing codes and/or small amplitude spurious sparkle codes, apply a large signal sine wave at the frequency to be tested using the test setup of 4.2.1. Perform a SINAD test using the methods of 9.2.1 or 9.2.2 and compare the test result to the stated limit.

While it is possible in each of the above tests to decrease/increase the input test frequency until the test passes/fails, and to thus determine the actual useful power bandwidth of each device, it is often more economical to pick a single, conservative test frequency and to test that all ADCs exhibit a useable power bandwidth greater than this minimum. ADCs tested in such a manner shall state the input frequency used in the test and specify this frequency as the minimum useful power bandwidth.

11.2 Gain error (gain flatness)

Gain error, also known as gain flatness, is the difference between the dynamic gain, $G(f)$, of the ADC at a given frequency and its gain at a specified reference frequency, divided by its gain at the reference frequency. The dynamic gain of the ADC under test at a frequency f is the magnitude of the frequency response at that frequency. The reference frequency is chosen to be a frequency whose gain is at or near the peak gain of the ADC passband; typically it is the same frequency as the one used in the bandwidth test (see 11.1). For dc-coupled ADCs, the reference frequency is typically dc ($f = 0$). To determine gain error, first determine the dynamic gain. This may be done by using the sine-wave-based methods of 11.1 or from the differentiated step response method of 11.3. The gain error at frequency f is shown in Equation (84).

$$E_G(f) = \frac{G(f) - G(f_{\text{ref}})}{G(f_{\text{ref}})} \times 100\% \tag{84}$$

where

f_{ref} is the chosen reference frequency

11.3 Frequency response and gain from step response

The frequency response of an ADC is its complex response (magnitude and phase) versus frequency. It is also the Fourier transform of its impulse response. The preferred method of presentation is in the form of plots of magnitude (gain) and phase versus frequency.

11.3.1 Frequency response and dynamic gain test method

Record the step response of the ADC under test (see 10.2), using the step signal test setup in Figure 6, an appropriate step signal, $s(t)$, and equivalent-time sampling if necessary (see 4.4). Determine to sufficient accuracy the step signal's input amplitude, s_0 (the magnitude of the difference between the steps input base state and input high state). Select the ADC's (equivalent-time) sampling rate, T_s, high enough to give negligible aliasing errors based on the ADC bandwidth (see 11.3.2); if the bandwidth is unknown prior to this test, the test may have to be repeated, once the bandwidth is known, at a sufficient sample rate to make

the aliasing errors negligible. Acquire a record of M samples of the step signal, with an epoch (MT_s) long enough to enable the high state of the step to settle to within the desired accuracy. Estimate the ADC's discrete-time impulse response, $h[n]$, by taking the discrete derivative of the step response samples, $s[n]$, in units of the output quantity, and dividing it by the step's input amplitude, s_0, in units of the input quantity. The discrete derivative is often estimated by the first difference of the samples in the record (see Souders and Flach [B50]) as shown in Equation (85).

$$h[n] = \frac{1}{s_0}\frac{d[s(nT_s)]}{dt} \cong \begin{cases} \dfrac{s[n+1]-s[n]}{s_0 T_s} & \text{for } n = 0,1,2,\ldots,M-2 \\ \dfrac{s[n]-s[n-1]}{s_0 T_s} & \text{for } n = M-1 \end{cases} \tag{85}$$

Calculate the DFT of the impulse response using a non-weighted (rectangular) window. Multiply the result by the value of the sampling period, T_s. The result is an estimate, $H(f_k)$, of the frequency response of the converter, at the frequencies $f_k = k/(MT_s)$ given in Equation (86).

$$H(f_k) = T_s \sum_{n=0}^{M-1} h[n] \exp\left(\frac{-j2\pi nk}{M}\right) \quad \text{for } k = 0.1.1,\ldots,\frac{M}{2} \tag{86}$$

Note that the frequency response is estimated only at discrete frequencies f_k. To estimate the frequency response at other frequencies, linearly interpolate between the closest discrete frequencies.

For most Fourier transform calculations the phase spectrum typically is *wrapped*, that is, its values are modulo (2π); in other words, only the remainder after dividing by 2π is given. The wrapping is partly due to the delay between the start of the record and the position in the record of the step transition. This delay introduces a phase term that is linearly related to frequency. The delay and the *linear phase* term that it induces are usually arbitrary quantities because the actual delay between the recorded signal and the time of the input step's transition is usually indeterminate. However, the portion of the phase spectrum that is not linearly related to frequency is often of interest, since this indicates effects on the phase due to the ADC under test. The *nonlinear phase* portion of the phase response can be made more apparent by unwrapping the phase (Souders et al. [B52]). A simple method to do this is to create a simple program to subtract 2π following each 2π discontinuity. Noise will usually impose a limit on how high in frequency such an approach can be effective. The result is a plot of the nonlinear phase contribution.

This test method makes use of the natural roll-off of the ADC under test as an anti-aliasing filter, attenuating the frequency components of the step that are beyond the Nyquist limit. Bounds on the magnitude and phase errors from aliasing and first differencing are given in 11.3.2.

Note that the digital differentiation operation accentuates high-frequency noise components, such as that due to quantization, and the equivalent noise increases as the square root of record length.

Ideally, $H(0)$ shall equal the static gain as calculated in 7.4. This may not be the case, due to nonlinearities in the ADC, incomplete settling of the step signal, and other non-ideal behavior associated with the signal used for the test. Other non-ideal behavior could include: errors due to period-to-period jitter in the test square wave, i.e., short- versus long-term jitter effect on equivalent time measures; any hysteresis error introduced as the ADC cycles periodically through its saturated and cutoff states; distortion due to bandwidth reduction architectures that translate harmonics by means of decimation filters; etc. These types of errors are all architecture dependent and so it is not possible to write general procedures to account for such effects.

11.3.2 Aliasing and first differencing error bounds

Bounds can be calculated for the errors in the dynamic gain or frequency response estimated above in 11.3.1. Assume that the frequency magnitude response of the ADC under test rolls off at least −20 dB per decade for frequencies higher than the −3 dB frequency (see 11.1), corresponding to the roll-off for a single pole. Then the aliasing and first differencing errors, $e_m[f]$, in the magnitude response, as measured above, will be no greater in magnitude (positive or negative) than (see Souders et al. [B51]) as shown in Equation (87).

$$e_m(f) = \frac{400 \times f_{co} \times f}{f_s^{\,2}} \%\ \text{of the step's amplitude at the ADC output,} \tag{87}$$

and valid for $f < f_s/2$ and $f_s \geq 2f_{co}$

where

f is the frequency of interest
f_{co} is the cutoff frequency (bandwidth) of the ADC under test
f_s is the sampling rate.

For the phase response, the aliasing and first differencing errors, $e_p[f]$, will be no greater in magnitude than (see Souders et al. [B51]) as shown in Equation (88).

$$e_p[f] = 270 \frac{f}{f_s} \text{ degrees} \tag{88}$$

valid for $f \leq f_s/4$ and $f_s \geq 2f_{co}$

Example: If the expected cutoff frequency of the ADC under test is f_{co} = 10 MHz and an equivalent-time sampling rate of f_s = 100 MHz is chosen, what is the maximum aliasing and first differencing error that can be expected at half the cutoff frequency (5 MHz)? [See Equation (89) and Equation (90).]

$$e_m = \frac{4 \times 5 \times 10^6 \times 10^7}{10^{16}} = 2\% \tag{89}$$

$$e_p = 270 \frac{5 \times 10^6}{10^8} = 13.5° \tag{90}$$

As stated above, these error bounds assume that a single pole dominates the ADC's frequency response roll-off. Tighter error bounds may be applicable, in the case when the ADC's magnitude response rolls off at −40 dB per decade of frequency or faster, and when a correction is applied for the first-differencing error (see Blair [B8]).

Note that these expressions give the upper and lower bounds only, and cannot be used as corrections.

11.3.3 Comment on frequency response tests

Significant amounts of nonlinearity and signal distortion in the ADC under test may result in inaccurate or even pathological values for frequency response, bandwidth, and gain flatness. Specifically, the errors in the step response caused by nonlinearity, as described in 10.2, can become errors in the measured bandwidth, gain flatness, and frequency response, as determined in 11.1, 11.2, and 11.3, respectively. As noted in 10.2, the slew rate of the step signal used to determine the step response of the ADC has to be significantly below the slew rate limit, if any, in order to avoid nonlinearities.

12. Differential gain and phase

12.1 Introductory information on differential gain and phase

Differential gain and differential phase are parameters that quantify the suitability of circuits primarily for use with color composite video signals. These parameters are unrelated to any other gain or phase measurements outlined in this standard. They can and should only be measured on circuits specified to operate at bandwidths high enough to support the digitization of video test signals without aliasing.

In a National Television Standards Committee (NTSC) coded color video signal, the chrominance information is contained in amplitude-modulated 3.579 545 MHz sub-carriers that are displaced in phase by $\pi/2$ radians. The luminance information, which can be used without color demodulation to decode the black and white portion of the signal, is broadcast as amplitude-modulated signal.

If the level of the luminance signal were to affect either the amplitude or the phase of the chrominance signal, then the resulting color displayed would not be the same as intended. Differential gain is present when the gain of the chrominance signal is affected by changes in the luminance level. Differential phase is present when the phase of the chrominance signal is affected by changes in the luminance level. Differential gain distortion causes incorrect color saturation; differential phase causes incorrect hues to be reproduced.

To test for differential gain and phase, a small amplitude sine of 3.58 MHz (or close to that) is measured for changes in sine-wave amplitude and phase as the dc offset is varied. Each dc offset represents a different value of amplitude-modulated luminance signal. As NTSC standards define differential gain as the largest amplitude deviation between any two levels, the number which shall be reported is the worst-case peak-to-peak deviation of sine-wave amplitude over all luminance levels expressed as percent of the sine amplitude. Similarly, the number reported for differential phase shall be the worst-case peak-to-peak deviation of sine-wave phase over all luminance levels, expressed in degrees. As a consequence of this definition, it is not necessary to specify differential gain or differential phase as signed numbers; the absolute value of the peak-to-peak error unambiguously meets the definition.

12.2 Method for testing a general purpose ADC

When testing a general purpose ADC (or any ADC that could be used with the sync tips at either end of the input range) it is important to scale and offset the input test signal so that almost all of the input full-scale range gets tested. In most applications this will involve setting a starting dc offset just large enough so that the 3.58 MHz sine-wave output code does not clip, and then increasing the dc offset until just before the sine-wave output code would begin to clip at the opposite peak.

The suggested test for a general purpose ADC uses a six-level stepped waveform, generalized from National Television System Committee standard test signals of 140 IRE units peak to peak. While IRE units assume standard impedances and levels in volts, it is assumed here that an adjustable gain is placed ahead of the ADC under test in order to map the input full-scale range of the ADC to 140 IRE units. When this gain is adjusted such that the sine-wave amplitude measures 40/140 = 28.6% of the full-scale range peak to peak, then the gain from the sine source to the ADC is correctly mapped. When the dc steps change the output by 20/140 = 14.3%, then the gain from the stepped dc source is mapped correctly. It is further assumed that an adjustable offset has been summed with the input in order to keep the test signal from clipping at either end of the input range. The test setup is shown in Figure 21(a) and the waveform diagram is shown in Figure 21(b).

The waveform of Figure 21(b) uses six stepped dc levels of 20/140, 40/140, 60/140, 80/140, 100/140, and 120/140 of full scale at the ADC input. The test signal has been generalized from the NTSC standard test signal in that the length of time spent at each step has been made arbitrary, and the sync and timing information has been deleted.

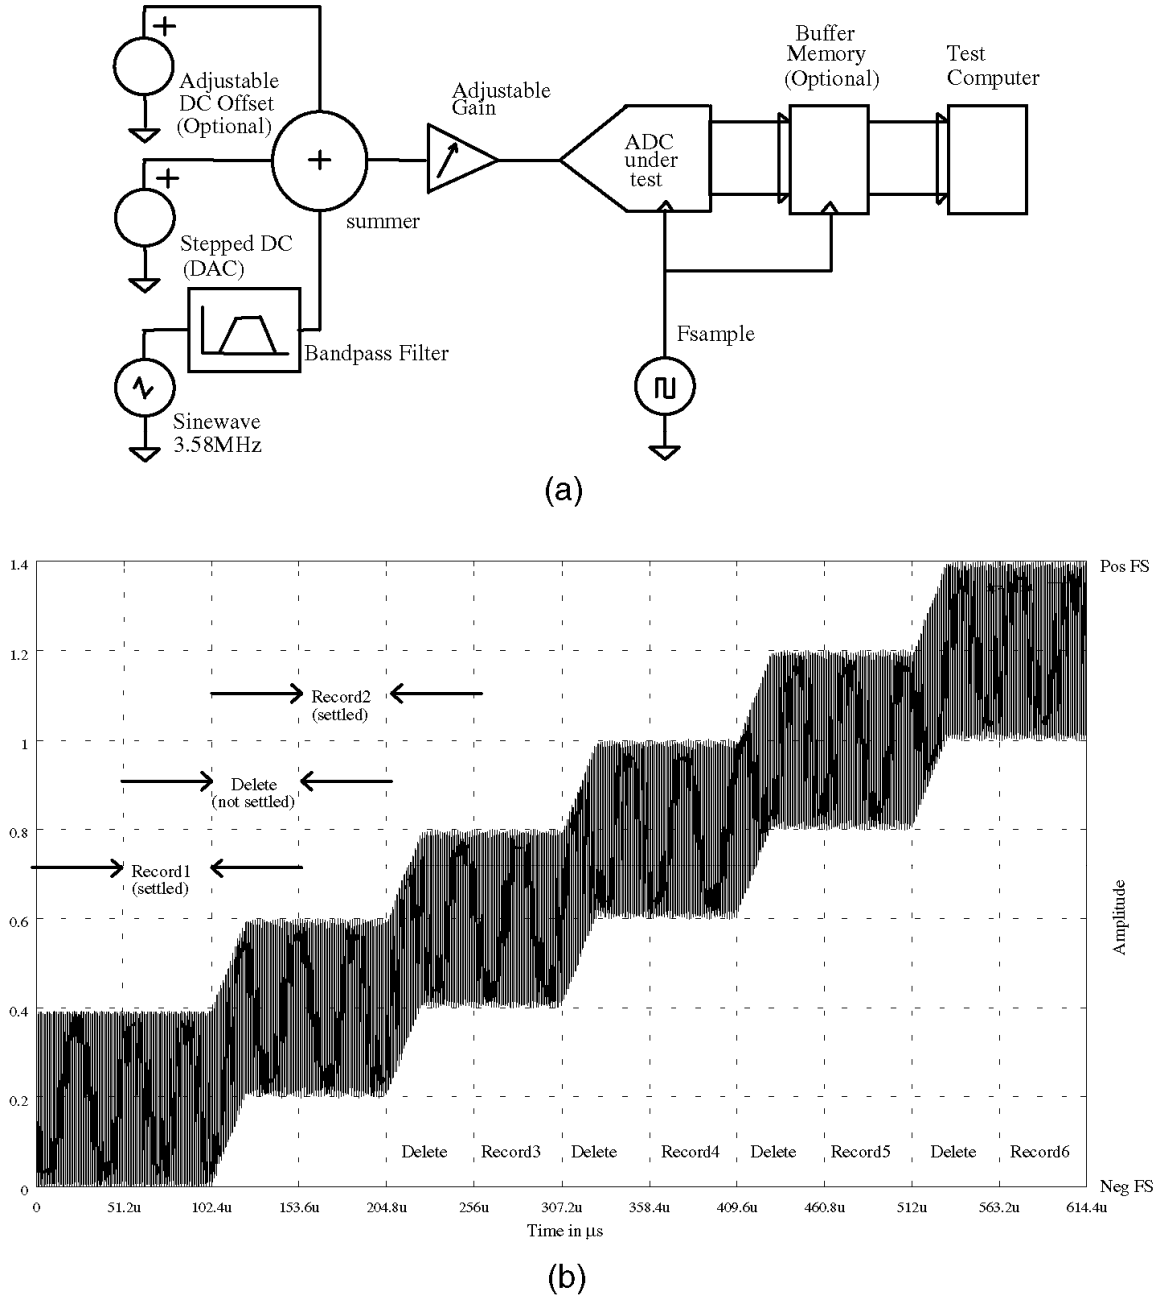

Figure 21—(a) Setup for differential phase/gain testing; (b) Example of stepped sinusoidal waveform used in differential phase and gain test

In order to maximize the number of unique output codes obtained in each record, the input frequency and the sampling frequency shall be picked such that the aperture point *walks through* the waveform. This requires that the sampling frequency NOT be equal to an integer multiple of the input sine frequency. Typically the sampling frequency is fixed at some value and the sine-wave input frequency is offset by a delta frequency sufficient to force at least one walk-through within *M* samples, where *M* is the length of each record sent to the FFT/DFT routine. If the input sine generator is only capable of generating exactly 3.579 545 MHz, then the ADC sampling frequency shall be offset by a small delta of sufficient value to force the aperture point to walk through a complete cycle of the sine at least once in each record.

If a coherent FFT or DFT based test is assumed, then the usual non-integer restrictions on the ratio of sampling rate to input frequency that achieve aperture point walkthrough will apply. The use of a non-windowed FFT imposes an additional requirement that an integer number of cycles of the sine wave be present in the record length chosen. When combined, these conditions are equivalent to requiring the aperture point to walk through a complete integer number of cycles within each record. These considerations result in the following restrictions shown in Equation (91) (also see 5.4.1):

a) J and M are mutually prime

b) $$f_i = \frac{f_s J}{M} \quad \text{for } j = 1,2,3,... \tag{91}$$

where

J is some integer to be chosen, usually less than $M/2$
M is the number of samples in each record input to any FFT/DFT
f_s is the ADC sampling frequency
f_i is the frequency of the input sine

If M is chosen to be some power of 2, such as 2048, then any J that is odd is mutually prime to M. In this case, the odd number J yielding the frequency closest to 3.58 MHz could be chosen to get the input frequency. If the source to be used is constrained to 3.579 545 MHz, then the relation above can be solved for the appropriate f_s.

The suggested test of a general purpose ADC acquires one long, contiguous record of ADC digital output data. Records 1 through 6, small pieces of the contiguous record, are each sent to an FFT, which yields complex data. The complex numbers corresponding to the frequency of the input sine wave are converted to magnitude/phase polar coordinates. One (usually the first) pair of magnitude/phase numbers is chosen to be the reference pair. The subsequent magnitude numbers from the FFT of records 2 through 6 are normalized to the reference magnitude and converted to percent change in magnitude.

The phase numbers are converted to degrees and the reference phase is subtracted, yielding change in phase in degrees. In general, if the first phase is used as the reference, then the phase from subsequent FFTs must also be adjusted by subtracting the equivalent phase of the time delay between the beginning of the reference record and the record to be adjusted. This additional adjustment factor is computed as

t_d = delay in seconds between beginning of reference record and the beginning of the record being adjusted

where

phase adjustment (degrees) is equal to $\text{mod}_{360}(360 f_i t_d)$

When the delay between records is set to exactly an integer number of record lengths, the above rules used to pick the input frequency will provide a phase adjustment factor is zero (see 12.2.1). The phase adjustment factor may be useful in cases where the dc step settles in much less time than one record length, such as tests utilizing very long records. In this case, the long continuous record will fit into a smaller memory size and some time which would have been spent waiting for the full record-length delay will be saved.

Five pairs of change in gain and change in phase numbers are obtained from the six-level stepped waveform. If the amplitude change from the reference is positive for one change and negative for another, then the worst-case sum of the positive change plus the absolute value of the negative change is reported as the differential gain. Similarly, if the change in phase is positive for one record and negative for another, then the worst-case sum of the positive change plus the absolute value of the negative change is reported. By convention, the reported numbers are always positive.

12.2.1 Method for neglecting phase adjustment factor

ADC output samples, which were acquired while the dc offset step was settling, are deleted. A smaller record of data beginning at some number-of-samples (delay = d) offset into the time record is kept. For FFTs of an ideal ADC to show identically the same phase for each sub-record, it is necessary that records to be kept each begin with a sample of the sine wave at the same phase. If the frequency restrictions above have been followed, it is known that the input frequency will be at the same phase at each integer multiple of M samples. Thus the next record of data to be kept would begin at sample number = $d + lM$. The integer l shall be chosen for each dc step so that the small record to be kept occurs after the dc offset step has settled. In all, six small records of length M samples shall be *kept* if a six-level stepped modulated waveform is used as the test input. The waveform of Figure 13(b) shows the case where record length M is 1024 with l = 2, 4, 6, …, and d = 1024 samples. It was assumed that f_s was 20 MHz and as shown in Equation (92).

$$f_i = \frac{183 f_s}{1024} = 3.574\,218\,75 \quad \text{MHz} \tag{92}$$

12.3 Method for testing a special purpose ADC

There may be systems which contain clamping, or dc restoring circuitry, or which require sync stripping for functionality, or which for other reasons set the ADC input range such that the video information is never able to be found in portions at either end of the ADC input range. If the special purpose ADC or system is designed this way, then a conventional NTSC test waveform should be used. The sine wave should start at a dc offset equal to the blanking level in its application (which could be at a dc offset of up to 23% of full scale). The dc offset should similarly be increased in steps toward the reference white level of the intended application, which will be at a level well before the sine-wave output code begins to clip. If the input full-scale range of the special purpose ADC is designed to allow user-adjusted gain before the ADC, then the full-scale input span should be mapped to 180 IRE units (as in a potential application which includes sync tips and a small amount of headroom within the input full-scale range). In this case, a sine wave of 40/180 = 22.2% of the full-scale range peak to peak corresponds to 40 IRE units.

Note that the gain and offset suggested above will never test sections of ADC output codes at the lower and upper end of its input range for their contribution to differential phase or gain errors. It is expected, due to the nature of the special purpose circuitry and its intended mode of operation, that these output codes will never be used in digitizing the active video portion of an NTSC signal.

If an NTSC standard signal generator is used, where the input frequency is constrained to be exactly 3.579 545 MHz, the sampling frequency at the ADC should be adjusted as in 12.2 to produce an integral number of cycles in the record length and thus produce spectral lines in the FFTs aligned with the center of their bin.

In the case where the sampling frequency is locked to a color burst which is derived from the input signal, a nonstandard input signal with an exact 3.579 545 MHz color burst could be used, while the frequency digitized at each dc offset and used in the tests could be offset from this frequency. Systems where both the input sine frequency and the sampling frequency are constrained to an integer ratio may have to resort to extraordinary measures, such as repeated tests over multiple sine-wave phases, or repeated tests using a ramped dc offset, in order to obtain the accuracy and repeatability available when testing over a wide alphabet of output codes.

Note that only two dc offsets are required to be tested during each horizontal interval; one sine at the reference dc offset and a second sine at the new dc offset. Thus, systems dependent upon sync information within an NTSC signal need not squeeze all dc offsets into a single horizontal interval as shown in standard waveforms. If dc offset settling times or FFT record length considerations warrant, the test can be spread

over many sync intervals by acquiring records for only one pair of dc offset sine waves in each horizontal interval.

12.4 Comments on differential phase and differential gain testing

The choice of generator(s) used for the tests can limit the achievable accuracy and resolution of the tests. While an arbitrary waveform generator could conceivably be used to provide the test stimulus of a combined stepped dc offset and sine wave, the DAC used in the generator and any filtering or de-glitching output circuitry present could color the measured results with its own differential gain and phase characteristics. A better solution uses a continuous low-phase-noise sine-wave source summed with a stepped dc offset. These signals may come from separate generators, each optimized in bandwidth for minimum noise.

The measurement of differential phase and gain for a Phase Alternation Line (PAL) encoded video application can generally be regarded as an equivalent problem to the NTSC test except that the color sub-carrier frequency is defined to be 4.4332 MHz. This frequency should be used for the sine-wave test signal.

The method for testing general purpose ADCs utilizes an input waveform featuring user-chosen duration for each dc stepped sine segment. By making the length of each segment long, a record of 2048 samples or more may be used as the input to the FFT. The influence of quantization noise upon the results diminishes in proportion to the square root of the number of samples of each record. Using longer records enhances the repeatability of results.

The NTSC 40 IRE standard waveform contains offset sine waves from −20 to +120 IRE units, but also includes sync tips extending to −40 IRE. An NTSC waveform generator can be used to perform the general purpose ADC test by offsetting and scaling either the input signal or the ADC input range such that negative full-scale corresponds to −20 IRE and positive full-scale corresponds to +120 IRE.

When an NTSC standard test generator is used as the input source, the dc-stepped segments are of fixed duration. Typically only a short record length (≤64 samples) is possible for use in each FFT, and the repeatability of the results will be worse than if a longer record length were used.

Note that differential gain and phase are not the only error sources which affect the actual color of ideally reconstructed digitized NTSC video. Quantization noise is another major error source in the digitized video. When video is digitized on a well-designed converter with a low number of bits, the quantization noise can exceed the color pollution effects of differential phase and differential gain by many times. The total instantaneous color pollution of a system will result from the sum of quantization noise degradations and the numbers measured here.

Other, earlier methods of testing utilized a reconstruction DAC and an instantaneous error display. The results displayed required visual interpretation to distinguish that portion of the instantaneous error due to quantization noise from that due to systematic differential phase or gain distortion. The tests of this standard, unlike earlier methods, yield measurements of differential phase and differential gain that are inherently distinct from the effects induced by quantization noise. Quantization noise will affect the repeatability of the numbers measured here, but this random variation can be made arbitrarily small by increasing the record length M.

A similar test method was used in Carbone [B12].

13. Aperture effects

13.1 Introductory information on aperture effects

Sampling, in the real world, does not occur instantaneously. As a consequence, an output value produced by an ADC is a weighted average of the analog input signal over the period of time during which the actual sampling occurs. The term aperture can refer to this period of time in general, or more specifically to the weighting as a function of time during this period. By the latter definition, the aperture, or aperture weighting function, is the time reversal of the impulse response of the ADC. An ADC generally requires a sample-and-hold function of some kind in order to reduce the aperture effects. In many cases the aperture effects are related to the sample-and-hold step response.

13.2 Aperture duration

Aperture duration may be closely related to the transition duration of the step response. Unless otherwise indicated, aperture duration is the full width at half maximum (FWHM) of the aperture weighting function. For a Gaussian aperture, this is roughly 0.92 times the 10% to 90% transition duration of the step response. Aperture duration can also be described as the length of time necessary to encompass a specified percentage of the area under the aperture weighting function, see 3.1, Definitions. The length of time necessary to encompass the center 80% of the area under the aperture is identically the 10% to 90% transition duration of the step response. Some ADCs, such as successive approximation converters without sample-and-holds or track-and-holds, will have aperture duration equal to the total conversion time, and the output will represent some value that has occurred during the conversion time. In this case, the ADC output value is a highly non-linear function of the input signal. Since the concept of an aperture weighting function is based on a linear model for the ADC, the term *aperture weighting function* is not meaningful in this and similar situations. Other converter architectures will have other aperture effects.

The assignment of a definite aperture weighting function to an ADC is only valid if the sampling process is linear. In some cases this is a good approximation, while in other cases it is not. Considering the linear case the output of a sample at sampling time, t_0, has the form

$$v = \int_{-\infty}^{0} w(t) v_{\text{in}}(t + t_0)\, dt, \tag{93}$$

where

$w(t)$ is the aperture weighting function, satisfying $w(t) = 0$ for $t > 0$
$v_{\text{in}}(t)$ is the input signal

This expression is valid, because it is the most general expression for a value that depends linearly on $v_{\text{in}}(t)$, and that is causal (i.e., depends only on the past of $v_{\text{in}}(t)$). It is assumed that the units are adjusted so that the integral of w is one, i.e., so that the output equals the input if the input is constant. Let $g(t) = w(-t)$, then

$$v = \int_{-\infty}^{0} g(-t) v_{\text{in}}(t + t_0)\, dt = \int_{-\infty}^{t_0} g(t_0 - t') v_{\text{in}}(t')\, dt' = \int_{-\infty}^{\infty} g(t_0 - t') v_{\text{in}}(t')\, dt' \tag{94}$$

where the substitution $t' = t + t_0$ was made. The change in the upper limit of integration in the last step is valid, because $g(t_0 - t') = 0$ for $t' > t_0$.

The last expression shows that the output is equal to the result of passing the input signal through a filter whose impulse response is the time reversal of the aperture weighting function. This has significant implications for the testing and analysis of ADCs. It means that any method for measuring the impulse response or the step response yields a method for measuring the aperture weighting function. It also means

that the combined effect of two or more components, such as an input buffer amplifier and a sample-and-hold circuit, can be obtained by convolving their individual impulse responses.

The following figures illustrate the meaning of the *p*% aperture duration defined in 3.1. Figure 22 shows the impulse response of a sample-and-hold circuit while it is in the sample mode.

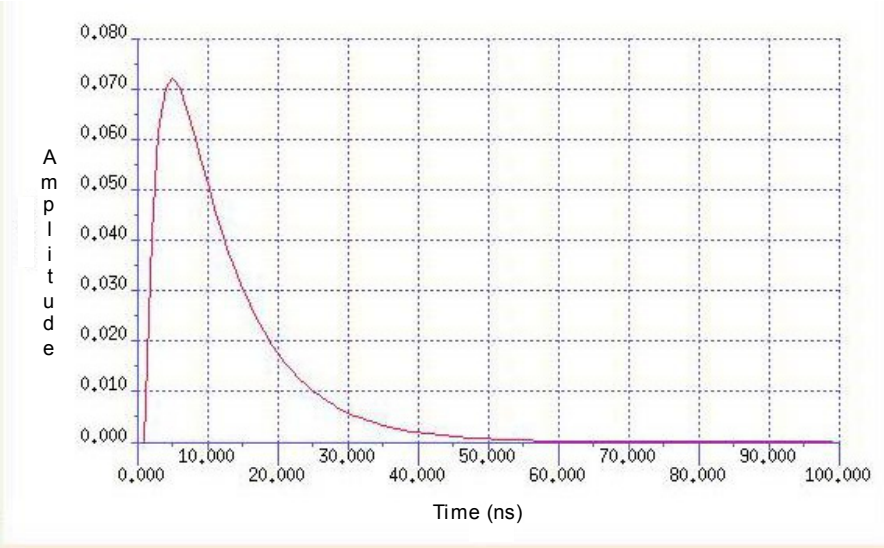

Figure 22—Impulse response of a sample-and-hold amplifier while in the sample mode

The step response of the sample-and-hold amplifier whose impulse response is shown in Figure 22 is shown in Figure 23.

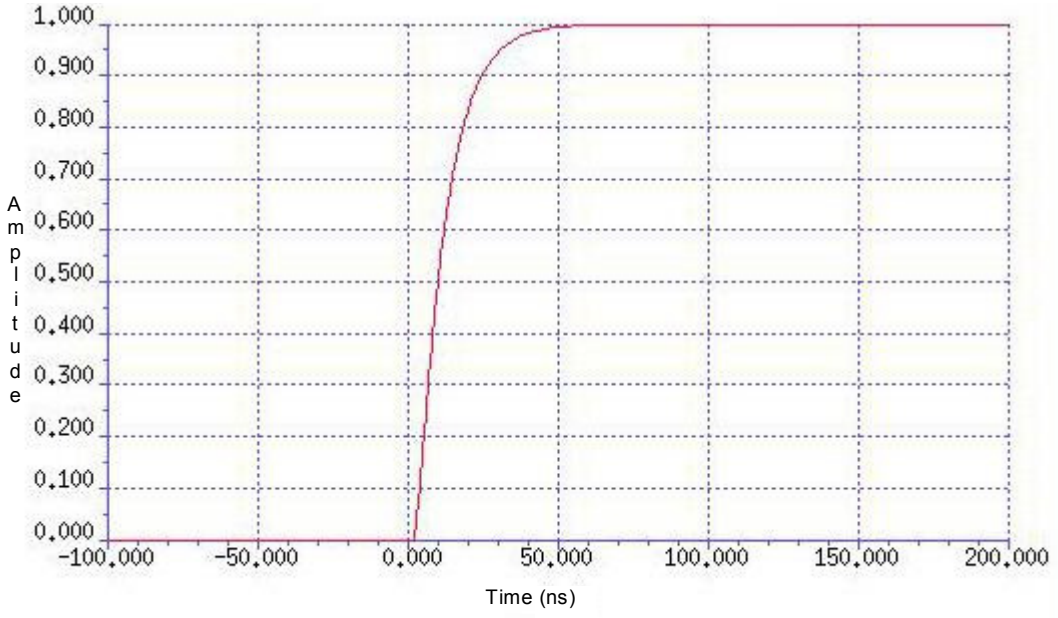

Figure 23—The step response corresponding to the impulse response of the previous figure

The aperture-duration is illustrated for the cases of p = 50, 80, and 99.9 in the following figures.

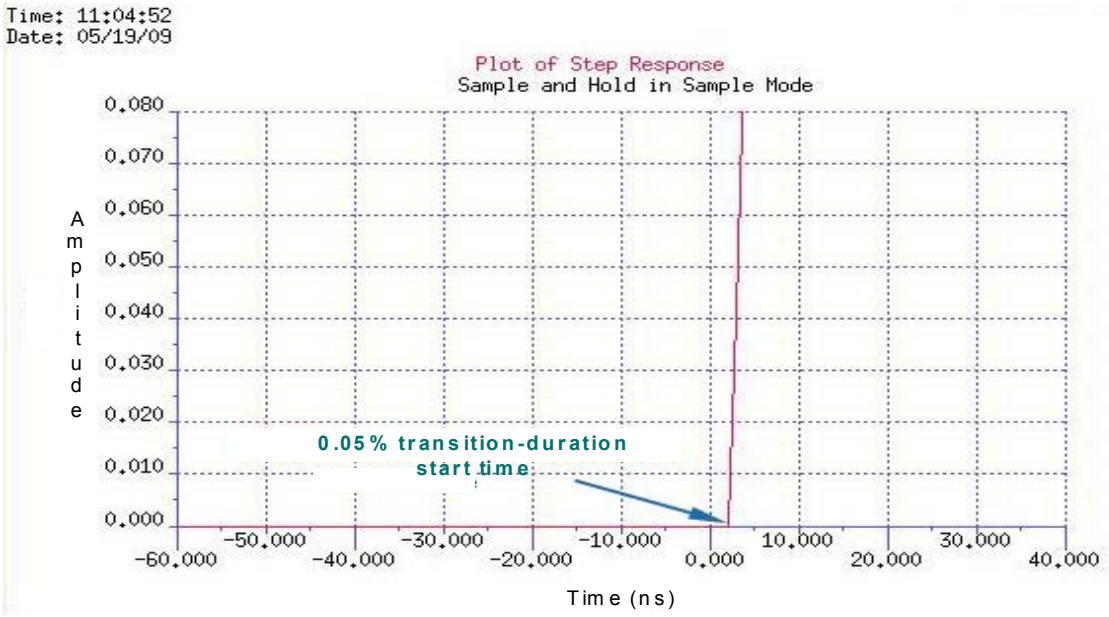

Figure 24—View of start of transition-duration for p = 99.9%

The transition-duration start time is at t = 2 ns for the step response shown in Figure 23.

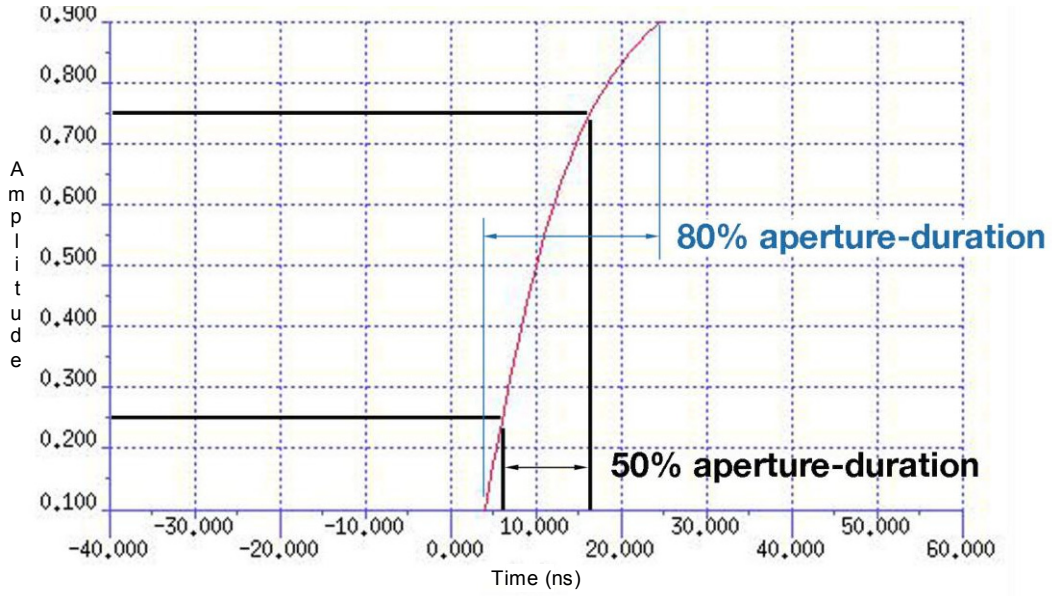

Figure 25—Aperture duration for p = 50%, and p = 80%

The aperture-duration shown in Figure 25 for p = 50% and p = 80% are 12 ns and 20 ns respectively. The 99.9% aperture-duration is computed from Figure 26 and Figure 24.

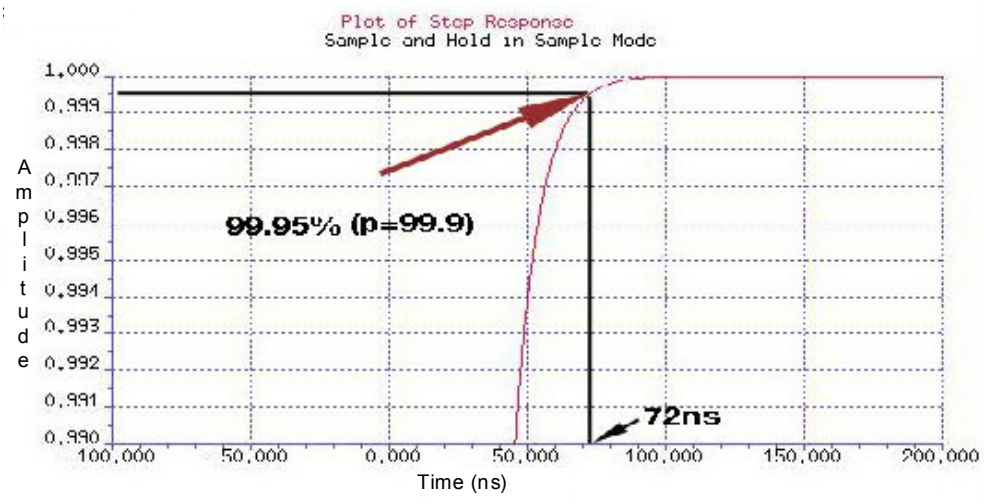

Figure 26—Evaluation of aperture-duration stop time for *p* = 99.9%

The impulse response is frequently a ringing response as is shown in Figure 27.

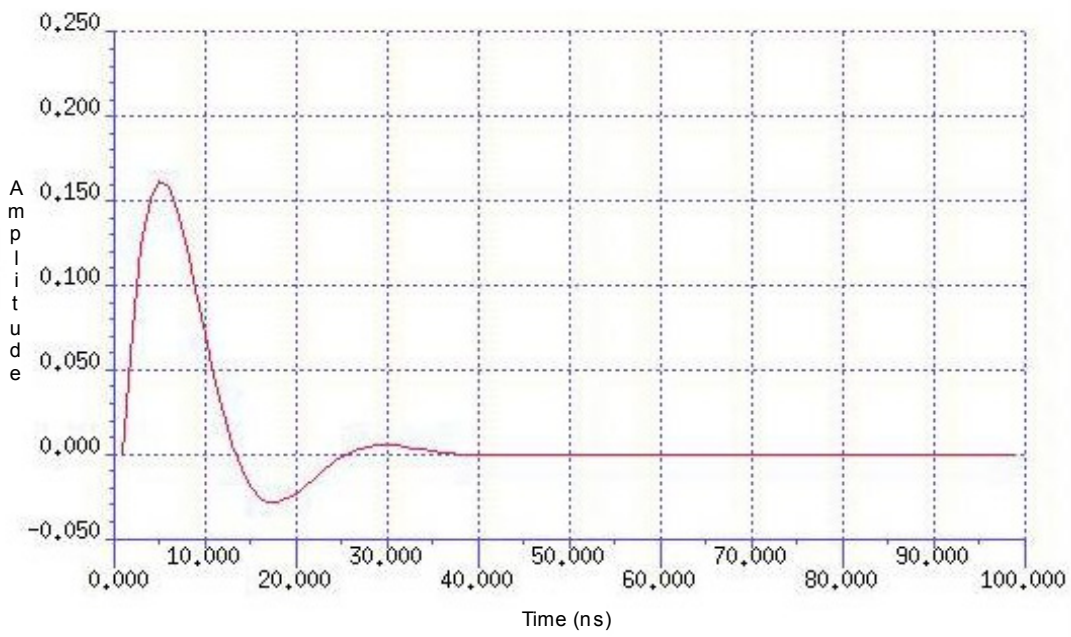

Figure 27—Ringing impulse response

This results in a ringing step response as shown in Figure 28.

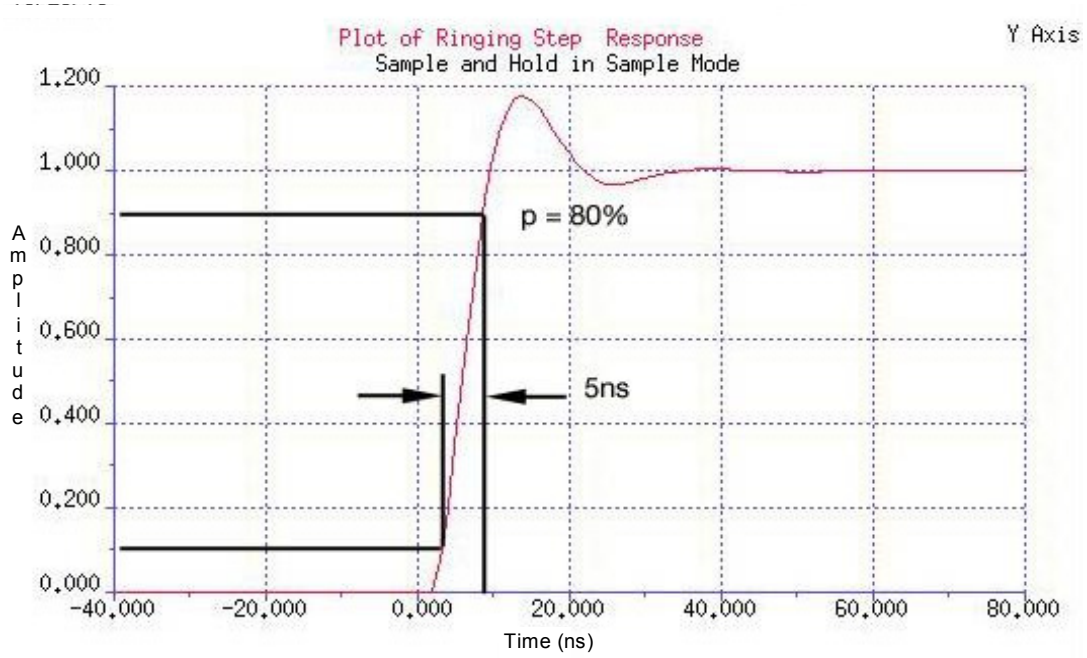

Figure 28—Ringing step response

The transition-duration ends at the last crossing of the reference level as illustrated by Figure 29.

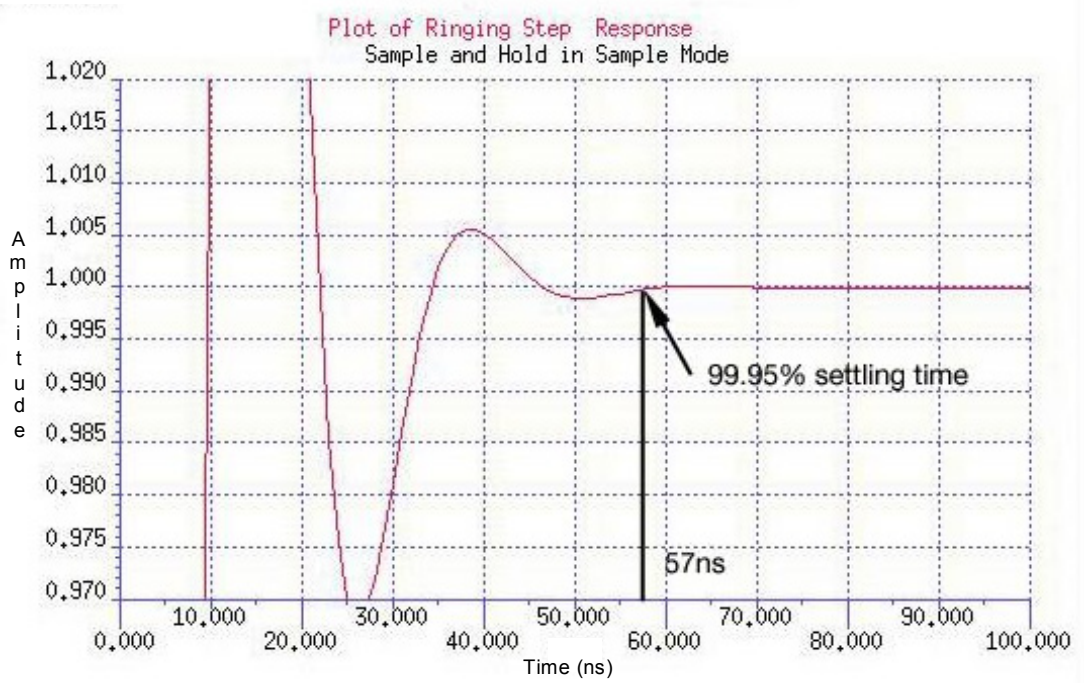

Figure 29—Zoomed ringing transition

The mathematical description is valid for a linear system, which is often a good approximation. However, it is seldom exact. For example, a sample-and-hold circuit usually charges a capacitor through a forward-biased diode or transistor junction, which is not linear. A more extreme case is a successive approximation

ADC without a sample-and-hold. The output value depends in a very nonlinear (but calculable) manner on the input signal during the entire conversion time. For a nonlinear system, the step response, and its p% duration, will depend on the initial and final values of the step used for its measurement, which must be specified along with any specification of the aperture duration.

13.2.1 Test method

Record the step response of the ADC under test as described in Clause 10. The p% aperture-duration is the length of time for the step response to go from $50 - (p/2)$% to $50 + (p/2)$% of its final value.

13.2.2 Comment on selecting the value of *p*

Low values of p (50 to 90) are typically used in communication applications and larger values of p (99 to 99.9) in data acquisition applications.

13.3 Aperture delay

Aperture delay is the delay from a threshold crossing of the ADC clock, which causes a sample of the analog input to be taken, to the center of the aperture for that sample. The center of the aperture is defined as in Equation (95).

$$t_{wc} = \frac{\int_{-\infty}^{\infty} tw(t)dt}{\int_{-\infty}^{\infty} w(t)dt} \tag{95}$$

where

t is the time from the threshold crossing
$w(t)$ is the aperture weighting function

The aperture delay can be either positive or negative, depending on whether there is greater delay in the clock or the analog input path in the ADC.

13.3.1 Test method

Apply a ramp to the analog input and a clock to the clock input of the ADC. Instead of a ramp signal at the analog input, a portion of another waveform (e.g., a sine wave) can be used, provided that the slew rate of the waveform does not vary substantially over the aperture duration of the ADC (e.g., a sine wave of frequency less than half the analog bandwidth of the ADC). The ramp signal slew rate should be as high as possible without exceeding the slew rate limit of the ADC's input or causing excessive dynamic errors. Synchronize the ramp and the clock such that the ADC samples the ramp near the center of the ADC's full-scale range.

Using a time-interval meter or oscilloscope of sufficient resolution and accuracy, measure the time delay from when the clock input crosses its threshold to when the analog input crosses the dc value corresponding to the center of the ADC's full-scale range. For ADCs with very high sample rate, where the aperture delay may be similar in magnitude to the clock period, extra care must be taken to measure from the correct clock edge. This can be done by repeating the measurement at various clock frequencies; the aperture delay as a function of clock periods should be nearly constant.

13.3.2 Comment on aperture delay

This test is difficult to make, and generally unnecessary, since the absolute delay from trigger to data acquisition is generally absorbed by other delays. The figure that actually determines aperture effects are the

differences between a fixed delay and the actual aperture time. This is called aperture uncertainty. Note that parallel converters can have an aperture time which varies with the signal level at the input to the converter.

13.4 Aperture jitter

Aperture jitter, sometimes called *aperture uncertainty*, is the standard deviation of the aperture delay.

NOTE—The concept of aperture jitter can apply to ADCs that contain digital filtering of the output data (e.g., sigma-delta ADCs), even though the aperture for a given output value cannot be associated with a single clock edge, if a sufficiently stable clock source is applied to the ADC. An arbitrary decision can be made about which clock edge to associate with the sampling of a given output value, so long as the association is consistent across all samples taken during the measurement of the aperture uncertainty. Aperture jitter in a given setup may depend on the transition duration and noise on the clock signal (see 11.1).

13.4.1 Test method

Couple the output of a stable signal generator to both the analog input and the clock input of the ADC (see Figure 30), using appropriate signal splitters, attenuators, dc blocks, frequency multipliers/dividers, etc., to produce a signal amplitude, offset, and frequency at each port are appropriate for that port. If at all possible, do not use any active components in this coupling, as any jitter in those components would contribute to the overall measured aperture uncertainty. The slew rate of the signal at the analog input port should be as high as possible without exceeding the slew rate limit of the ADC's input or incurring significant attenuation due to the bandwidth limitations of the ADC's input, as measured using the test methods of 11.1 or 11.2. In particular, the slew rate must be large enough that $\sigma_A > 2\sigma_B$ in Equation (96). An inadequate slew rate signal at the analog input port would result in inadequate aperture uncertainty measurement sensitivity, whereas an excessive slew rate signal would result in an estimate for the aperture uncertainty that is lower than the actual value.

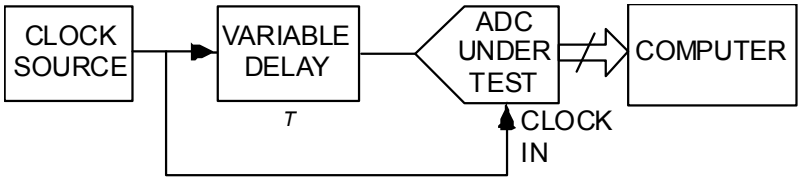

Figure 30—Test setup for aperture jitter test method; *T* is the value of the delay

Adjust the delay of the path from the signal generator to the analog input port to be longer than the delay to the clock input port by the amount of the aperture delay (see 13.3), such that each active clock edge is sampling itself at its midpoint (Figure 31). If the frequency of the signal at the analog input is divided down from the clock frequency, decimate the output record by the same ratio and adjust the relative phase of the analog input and decimation dividers such that each recorded output value is a sample of the input signal at its midpoint.

Measure the apparent random noise of the ADC in this configuration according to the test methods of 9.5. It may be necessary to add a summing node at the analog input port for low-noise ADCs in order to apply the slow triangle wave as well as the fast clock signal to the analog input. The apparent noise will include the effect of aperture uncertainty multiplied by the slew rate of the analog input signal.

Break the connection between the signal generator and the analog's input port, and terminate both ends of the broken connection appropriately to prevent reflections. Measure the random noise of the ADC according to the test methods in 9.5. The aperture uncertainty is then given by Equation (96).

$$\sigma_T = \frac{\sqrt{\sigma_A^2 - \sigma_B^2}}{S_{\text{eff}}} \tag{96}$$

where σ_A^2 is the measured noise variance with a clock signal applied to the analog input port, σ_B^2 is the measured noise variance without the clock signal applied to the analog input port, and S_{eff} is the effective slew rate (magnitude of the slope, see the following) of the clock signal at the analog input port at the sampling instant.

The effective slew rate must be measured with the ADC under test. If it is measured with an instrument with smaller (larger) transition duration than the ADC under test, the calculated aperture jitter will be smaller (larger) than the true value. To measure the S_{eff}, determine the average value of the ADC output with two values of the variable delay, t_1 and t_2, near the value used for the jitter measurement (one of the values can be the same as the delay used for the jitter measurement), and calculate S_{eff} as follows

$$S_{\text{eff}} = \frac{|m_2 - m_1|}{|t_2 - t_1|}, \tag{97}$$

where m_1 and m_2 are the mean values of the ADC output with the variable delay set to t_1 and t_2.

Figure 31 illustrates the method. The upper graph shows the clock signal as a function of time with the trigger point and the aperture delay shown. The lower graph shows the ADC output as a function of the variable delay. The time scale is expanded by a factor of two in the lower graph. With zero delay the clock is sampled the point that is one aperture delay after the trigger point. The shape of the lower graph is the time reversal of the upper. Delays close to the aperture delay give sample points on the rising edge (in this example) of the clock.

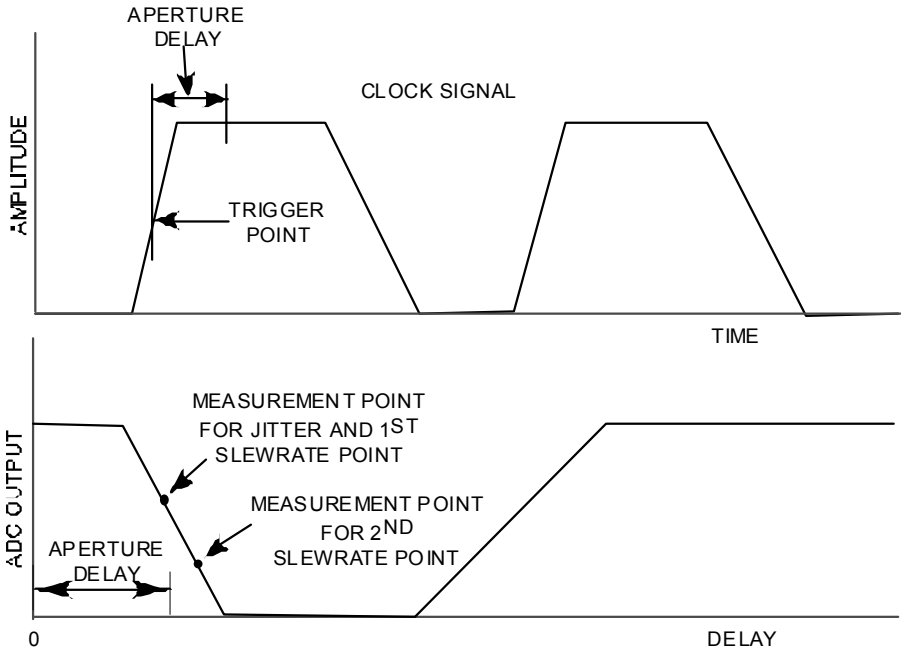

Figure 31—The upper graph shows the clock signal as a function of time; the lower graph shows the ADC output as a function of the variable delay

14. Additional tests and specification

14.1 Digital logic signals

Digital logic signals may be of many different types, depending on the ADC. The logic family name (CMOS, TTL, ECL-100k, ECL-10k, LVDS, etc.) should be used if the digital signals adhere to the standards for the family. If there are deviations from the standards (e.g., LVDS-like outputs that follow the specified LVDS high and low levels but have shorter transition durations than that specified for LVDS), then the differences should be clearly stated.

Users of ADCs will not usually need to measure the parameters of the logic signals, but manufacturers should do so to verify compliance with accepted logic standards.

This section is worded in terms of a positive logic binary system, where the higher signal level (high state) is associated with the binary value 1 and the lower (base state) with the value 0. The changes to adapt to a negative logic system are straightforward. The determination of base state and high state can be accomplished using the histogram method, peak method, or user defined limits in accordance with IEEE Std 181-2003, but the method used must be defined. In general, the logic parameters need to be determined under a suitable range of operating conditions. These operating conditions include power supply voltages, input voltages, load impedances, and temperature.

14.2 Pipeline delay

Apply to the analog input a known steady-state signal different from that used to obtain the values currently on the output. Initiate a series of conversions. Use an oscilloscope and/or logic analyzer to observe the input, clock, and output signals. The pipeline delay is the number of clock cycles between the clock transition that initiates the conversion and the clock transition that causes the corresponding data to appear as valid data at the output. To prevent confusion between pipeline delay and aperture and/or propagation delays, the converter shall be clocked at a low enough speed that the sampling of the analog input signal and the appearance of valid output data can each be unambiguously associated with a particular clock transition. The pipeline delay of an ADC is independent of the clock frequency. If the polarity of the clock transition at the start of conversion is the same as its polarity at the start of data validity, the delay will be an integral number of clock cycles. If the polarities are opposite, it will be a half-integral number of clock cycles. Note that even if the clock is asymmetric, each phase is considered to be one-half of a cycle.

14.3 Out-of-range recovery

An out-of-range input is any input whose magnitude is less than the maximum input value of the ADC but is greater than the full-scale value for the selected range. An out-of-range input may produce changes in the characteristics of the input channel, such as saturation of an amplifier or temporary changes in component values caused by thermal effects. The out-of-range recovery time is the time from the end of out-of-range to when the input channel returns to its specified characteristics. Out-of-range recovery occurs according to two different criteria. Relative recovery is achieved when the ADC's normal transfer characteristic is restored in all respects, except for signal propagation time through the ADC. Absolute recovery is achieved when the ADC's normal transfer characteristic is completely regained. Relative recovery is adequate when data before and after the out-of-range need not be related in time. When the data before and after the pulse must be related in time, then the ADC must recover absolutely.

14.3.1 Test method for absolute out-of-range recovery

Arrange a network capable of simultaneously applying both a high-purity sine wave and a specified out-of-range pulse (e.g., twice full scale) with a flat base state. Apply a high-purity, large signal sine wave of a convenient, non-harmonically related frequency (e.g., 1/20th the sampling frequency). Take a record of data with the out-of-range pulse occurring near the center of the record. Fit a sine wave to the data prior to the out-of-range pulse. Extrapolate the fitted sine wave to the end of the record. The measure of out-of-range recovery is the deviation of recorded data from the fitted sine wave. Out-of-range recovery time is measured from the last full-scale point associated with the pulse to the first point that deviates less than, and stays within, the desired tolerance of the fitted sine wave.

As a test of the method, record only the sine wave. Fit a sine wave to the portion of the record occurring prior to the point at which the out-of-range pulse will be introduced. Extend the fitted sine wave in the portion of the record where the out-of-range recovery is expected to occur. The observed deviation indicates the resolution obtainable when the pulse is applied.

14.3.2 Test method for relative out-of-range recovery

When the occurrence of events before the out-of-range pulse is not relevant to data acquired after the pulse, relative recovery is an appropriate criterion. Relative recovery may also be used when record length precludes the above method. To measure relative recovery, record several records of the sine wave. Fit each record of data with a sine wave. Find the average amplitude, frequency, and dc offset of the fitted sine waves. Take a record of data in which the out-of-range pulse is removed very early in the record. Synchronize the previously fitted, average sine wave to the latter portion of the record (e.g., the last F record) by varying the phase only. Extend the synchronized sine wave across the entire record. Observe deviations as before.

14.3.3 Comments on test methods

In a high-frequency 50 Ω system, the sine wave and the pulse must be added using a resistive adder. An isolating reactive adder generally does not work because the top of the test pulse droops due to the adder's limited low-frequency response. This droop causes undershoot when the pulse returns to its initial level. The resistive adder feeds some of the pulse back to the sine-wave generator, which may degrade the quality of the sine wave. This effect can be checked by performing the test on a sine wave/pulse combination that does not go beyond the full-scale of the ADC. Placing as large an attenuator as possible at the sine wave input to the resistive adder can reduce the degradation.

The out-of-range test pulse must return cleanly to its initial level. Any aberrations degrade the sine-fit results.

14.4 Differential input specifications

An ADC with differential inputs produces output codes that are a function of the difference between two input signal levels. The two input signals are typically called positive and negative. Such devices have a number of performance features in addition to those found in single-ended ADCs. These include the impedance of each input (positive and negative) to ground and to each other, maximum common-mode signal, maximum operating common-mode signal, common-mode rejection ratio, and common-mode out-of-range recovery time.

14.4.1 Input impedance to ground (for differential input ADCs)

This is the impedance between the positive input and ground or the negative input and ground. This impedance may be specified at several different frequencies. When the frequency is not specified, the

impedance given is the static value. Alternatively, the input impedance can be represented as the parallel combination of passive resistance and capacitance elements.

14.4.1.1 Test method

Perform the measurement described in 7.1 or 7.2 for each of the inputs. When determining the impedance of the positive (negative) input, the negative (positive) input shall be appropriately terminated and this termination shall be specified.

14.4.2 Common-mode rejection ratio (CMRR) and maximum common-mode signal level

CMRR is the ratio of the input common-mode signal to the effect produced at the output of the ADC in units of the input, T_k. The output codes can be converted to input units by using Equation (98).

$$V_{\text{out}} = Q[k-1] + T_1 \tag{98}$$

CMRR is normally specified as a minimum value in decibels. CMRR may be specified at various frequencies. The maximum common-mode signal level is the maximum level of the common-mode signal at which the CMRR is still valid. The maximum common-mode signal level must also be specified.

14.4.2.1 Test method

Arrange a network capable of simultaneously applying identical amplitude sine-wave signals to both differential inputs. The two common-mode signal levels must be identical to within the desired accuracy of the measurement. This is accomplished by connecting both inputs together to a single source with equal-length cables as shown in Figure 32. The common-mode signal level (V_{in}) must be large enough to discern an effect in the output data, and it must be equal to or below the specified maximum common-mode signal level. Take a record of data. Perform a DFT on the output record, and identify the frequency component corresponding to the common-mode frequency, V_{out}, converted into input units using Equation (98) at the common-mode sine-wave frequency. Compute CMRR in decibels from Equation (99).

$$\text{CMRR} = 20\log_{10}\left(\frac{V_{\text{in}}}{V_{\text{out}}}\right) \tag{99}$$

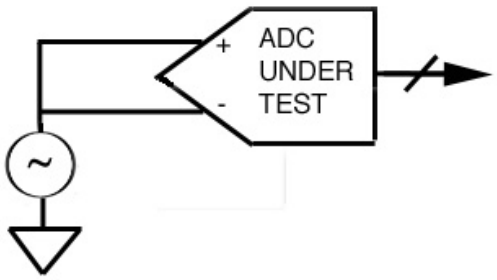

Figure 32—Test setup for measuring common-mode rejection

Low-noise converters, with low-frequency common-mode inputs, may not generate identifiable signals at the common-mode input frequency. When this is the case, the normal mode input to the converter shall be biased so as to generate at least two codes at the output, when no common-mode signal is present. If the output code cannot be easily adjusted, i.e., is a fixed value with no variation, assign V_{out} the value of $Q/2$.

14.4.3 Maximum operating common-mode signal

The maximum operating common-mode signal is the largest common-mode signal for which the ADC will meet the effective number of bits specifications in recording a simultaneously applied normal mode signal. Many ADCs have an absolute limit on the signal applied to either input. This limit should not be exceeded during the test process.

14.4.3.1 Test method

Arrange a network capable of simultaneously applying a common-mode sine wave at a specified frequency or a dc level signal to the inputs and a normal mode large signal sine-wave test signal at a different frequency to the inputs (see Figure 33). Adjust the initial common-mode signal level to the specified maximum common-mode signal level. Take a record of data. Compute effective number of bits. Raise or lower the common-mode signal amplitude to determine the largest amplitude for which the effective number of bits specification is met. Repeat the measurement at common-mode and normal mode sine-wave frequencies of interest. The maximum operating common-mode signal is the one which allows the device under test to still meet its effective number of bits specification.

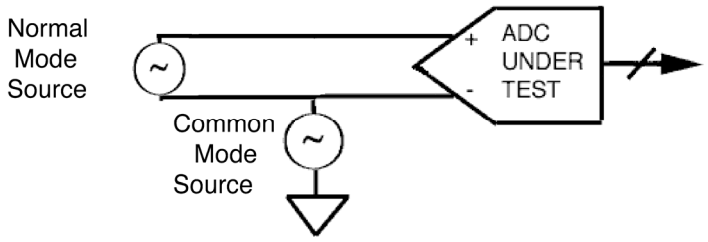

Figure 33—Block diagram of maximum operating common-mode signal

14.4.4 Common-mode out-of-range recovery time

The common-mode out-of-range recovery time is time required for the ADC to return to its specified characteristics after the end of a common-mode out-of-range pulse. A common-mode out-of-range input is a signal level whose magnitude is less than the specified maximum common-mode signal but greater than the maximum operating common-mode signal.

Differential amplifiers often have poor CMRR at high frequencies and performance will be degraded following a high-level common-mode pulse. The output may be driven off scale by a common-mode pulse. Comments concerning absolute and relative recovery times for normal mode out-of-range inputs in 14.3 will generally apply to common-mode out-of-range inputs.

14.4.4.1 Test method for common-mode out-of-range recovery time

Arrange a network capable of simultaneously applying both a high-purity sine wave and a common-mode out-of-range pulse of specified amplitude, transition duration, and width. Measure absolute and relative recovery times as described in 14.3.1 and 14.3.2.

14.5 Comments on reference signals

Many ADCs provide for one or more reference signals, which can be either inputs to control the ADC gain, or outputs that can be used to monitor the ADC operating characteristics. A common example is voltage

references to set V_{max}, the full-scale signal levels for the converter and, occasionally, V_{min}, the negative full-scale voltage of the ADC.

When such reference signals are control inputs, it is necessary to take into account the electrical properties such as input impedance, and to avoid applying inputs that exceed the ADC specifications. It is also of importance to note that the ADC's reference bandwidth may dramatically affect the ADC linearity, SNR, SINAD, and THD performance.

Methods to determine the properties of reference signals are beyond the scope of this standard, in part because they may not be the same for all ADC architectures.

14.6 Power supply parameters

14.6.1 Power consumption

Power consumption refers to the average power dissipated by the device under test from the main power supplies to the device for a specific set of operating conditions. As the power consumption may vary for different operating conditions, it is important that the setup used for determining the power consumption be specified. The setup configuration should mimic actual operating conditions and include such things as clock signals and required output loads for proper operation.

Note that since power consumption may vary depending upon several parameters (such as analog input voltage, clock polarity, etc.), it is mandatory to state the operating conditions for all inputs and outputs during this test. Manufacturers should try as many combinations as possible to arrive at the maximum consumption configuration for the device.

14.6.1.1 Power consumption test method

Connect the ADC under test to the appropriate power supplies set at specified, worst-case, or maximum, operating values for the device. Connect all appropriate signals necessary to operate the device as well as any necessary loads (i.e., signal termination resistors for inputs and outputs). Measure the current and voltage supplied to the device from each power supply. Compute the average power for each supply independently and sum these measures to determine the total power dissipation.

14.6.1.2 Power supply voltage effects

Changes in the power supply voltage can affect the ADC output. Different effects may be observed if the power supply voltages are changed, or if an ac waveform is superimposed on the power supply voltages. Variations of the ADC output due to power supply changes are usually a function of the analog input signal.

Although information about performance of the ADC with different power supply voltages can be achieved by performing gain, offset, and linearity measurements at the minimum, nominal, and maximum supply voltages given in the ADC specifications, most of the information about the effects of power supply changes can be measured by observing the effect on the output codes when the power supply voltages are changed.

The output of the ADC may be modified if ac signals are superimposed on the power supply voltages. Most effects of high frequency injection are mitigated and complicated by properly bypassing and filtering the supplies, however, it is fairly straightforward to measure the effects of low-frequency ripple superimposed on the supplies, and these effects are usually described by evaluating the power supply rejection ratio ($PSRR_{ac}$).

14.6.1.3 Power supply voltage effects test method

This test requires adjustable power supplies powering the ADC under test, and a method of collecting a data record. Static PSRR is measured by implementing the following steps:

1) Set the power supply voltages to their nominal values. The power supply voltage is defined as V_{PS}.
2) Apply an input signal to the ADC input that is approximately 95% of full scale.
3) Collect a record of the output codes. Compute the quantity $M1$ − the average of the output codes.
4) Change one of the power supply voltages to a new value (typically the specified tolerance on the power supply voltage). The change in the power supply voltage is defined as ΔV_{PS}.
5) Measure the new ADC average output code as $M2$.
6) Calculate the PSRR using the equation:

$$\text{PSRR} = \frac{(M2 - M1)}{\Delta V_{PS} / V_{PS}} \tag{100}$$

7) Restore the power supply to the nominal voltage.
8) Repeat steps 3 through 6 for all the other power supplies. If desired, repeat at other input signal levels.
9) The PSRR shall be less than the limits set by the data sheet.

Note that PSRR is a dimensionless ratio. It can be expressed in units of %/% or in decibels.

If it is desired to measure the effects of a ripple voltage superimposed on the supply voltages, then PSRR_{ac} can be measured. PSRR_{ac} at low frequencies is usually tested by adding a low frequency (typically 1 kHz) ripple directly through the regulator of the DAC power supply. The component of the ripple at the output is noted by collecting a record of data. The PSRR_{ac} is then calculated using Equation (101). High frequency ripple can be injected either with transformers or with resistive ladders. Because high frequency ripple injection often conflicts with bypass capacitor requirements, there is no general way to describe how to superimpose high frequency noise on the power supply.

The ADC ac PSRR is defined as follows:

$$\text{PSRR}_{ac} = \left(\frac{\text{ac}_{out} / 2^N}{V_{ripple} / V_{PS}} \right) \tag{101}$$

where

V_{ripple} is the rms ripple voltage added to power supply (typically around 50 mV_{rms} to 100 mV_{rms})
ac_{out} is the corresponding change in the rms output code as measured by computing the rms of the ripple frequency component present in the collected data record

This test requires an adjustable power supply that can accept a low-frequency ac ripple source as an input.

14.6.1.4 Comments on Power Supply Rejection Ratio (PSRR)

The power supply rejection ratio (PSRR) is a measure of how immune the ADC transfer function is to changes in the applied power voltages. It is specified in terms of decibels or percent of full-scale output change versus percent change in the applied power supply voltage. PSRR is specified for dc or low-

frequency changes in power supply voltages. PSRR is generally measured with the ADC generating close to a full-scale output (positive and/or negative).

PSRR is defined as follows:

$$\mathrm{PSRR} = \left(\frac{\Delta_{\mathrm{out}} / 2^N}{\Delta V_{\mathrm{PS}} / V_{\mathrm{PS}}} \right) \tag{102}$$

where:

Δ_{out} is the measured change in the output
ΔV_{PS} is the change in the power supply voltage
V_{PS} is the nominal value of the power supply voltage

The units of PSRR can be noted as "%FSR/%Power Supply Change" or it can be noted in decibels by the relationship:

$$\mathrm{PSRR}_{\mathrm{dB}} = 20 \log(\mathrm{PSRR}) \tag{103}$$

Power supply sensitivity (PSS) and power supply rejection (PSR) are terms that are often used to describe the effect of power supply changes on the output. This standard considers them to be equivalent to PSRR.

Annex A

(informative)

ADC architectures

Portions of this annex have been reprinted with permission from Rapuano, S., Daponte, P., Balestrieri, E., De Vito, L., Tilden, S. J., Max, S., and Blair, J., "ADC Parameters and Characteristics," *IEEE Instrumentation and Measurement Magazine,* vol. 8, no. 5, pp. 44–54, Dec. 2005 [B46]. © 2005 IEEE.

A.1 Integrating ADCs

Integrating ADCs provide high resolution and can reject both line frequency and noise. The integrating architecture provides an approach to converting a quasi-static analog signal into its digital representation. Integrating ADCs have usually low speed, low cost, and high resolution.

An integrating converter integrates the input signal and correlates the integration time with a digital counter. The output of the counter is proportional to the amplitude of the sample. In a dual-slope converter (Figure A.1), the input sample is integrated for a fixed time dictated by the digital counter. When the counter overflows, a switch is thrown, a positive reference voltage is connected to the integrator, and the counter is simultaneously reset. When the integrator output reaches zero (fully discharged) a comparator switches state, thereby latching the counter output.

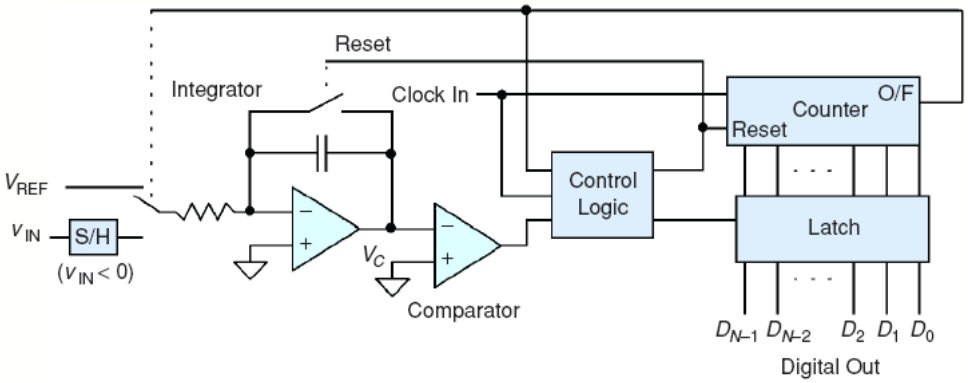

Reprinted from Rapuano et al., *IEEE Instrumentation and Measurement Magazine,* 2005 [B46]. © 2005 IEEE.

Figure A.1—Dual slope integrating ADC architecture

The main disadvantage of such architecture is that it can require 2^{N+1} clock pulses to perform a full-scale conversion in an N-bit converter. The advantage of the dual-slope architecture is that the precision and nonlinearity issues that reduced the accuracy of the single slope version are cancelled out because the same circuitry is used for both the reference voltage and the sample voltage integration (see Rauth and Randal [B47]).

There are additional integrating ADC architectures that are appropriate for different applications.

A.2 Flash ADCs

Flash ADCs are the fastest, operating from many megasamples per second (MS/s) to tens of gigasamples per second (GS/s). They perform their multibit conversion directly, but they require stringent analog design to manage the large number of comparators and reference voltages required. Figure A.2 shows a converter with N-bit resolution that has $2^N - 1$ comparators connected in parallel, with reference voltages set by a resistor network and spaced $V_{FS}/2^N$ apart (see Rapuano et al. [B46]). A change of input voltage usually causes the output state transition in more than one comparator. These output changes are combined in a decoder-logic unit that produces a parallel N-bit output from the converter. Although flash converters are the fastest types available, their resolution is constrained by the available die size and by excessive input capacitance and power consumption from the large number of comparators used. Their repetitive structure demands precise matching between the parallel comparator sections, because any mismatch can cause static errors.

Flash ADCs are also prone to sporadic and erratic outputs known as sparkle codes. Sparkle codes have two major sources: meta-stability in the $2^N - 1$ comparators and the thermometer-code bubbles. Mismatched comparator delays can turn a logical 1 into 0 (or vice versa), causing the appearance of bubbles in an otherwise normal thermometer code. Because the ADCs encoder unit cannot detect this error, it generates an out-of-sequence code that also appears as an output spark (see Rapuano et al. [B46]).

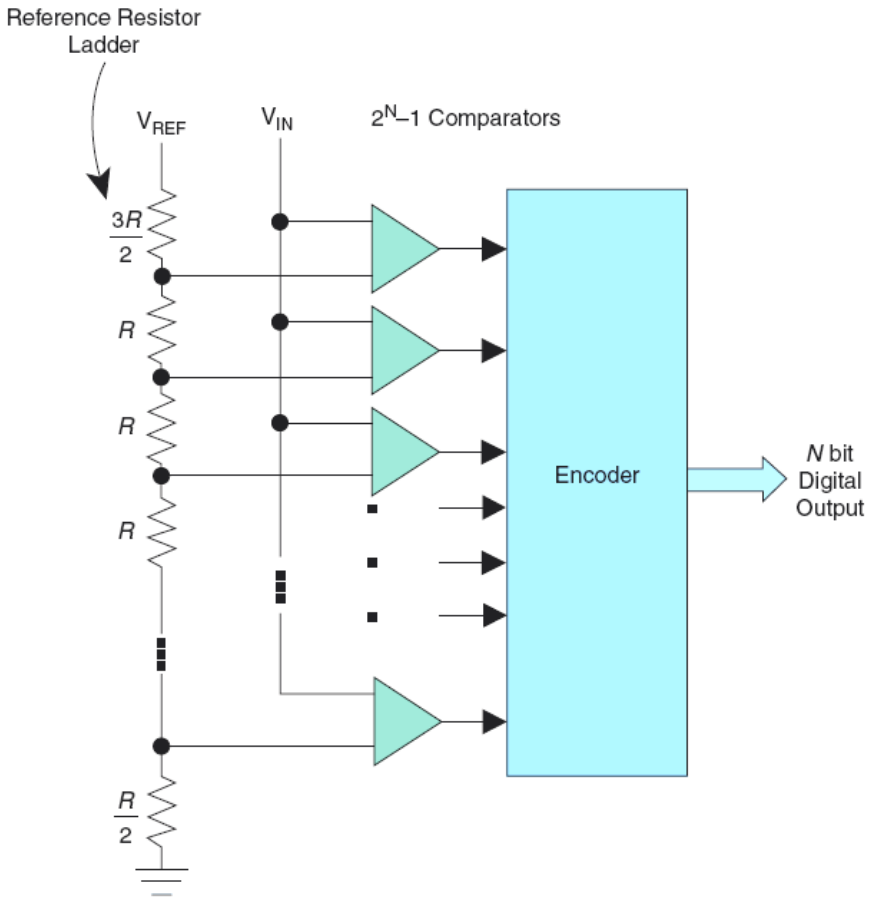

Reprinted from Rapuano et al., *IEEE Instrumentation and Measurement Magazine,* 2005 [B46]. © 2005 IEEE.

Figure A.2—Block scheme of a flash ADC

A.3 Pipelined and Subranging ADCs

The pipelined ADC has become the most popular ADC architecture for sampling rates from a few MS/s to >500 MS/s, with typical resolutions of 8 bits to 16 bits. It has its origins in the subranging architecture. Figure A.3 shows a block diagram of a simple 6-bit, two-stage subranging ADC. The output of the sample-and-hold amplifier (SHA) is digitized by the first stage 3-bit sub-ADC (SADC), usually a flash converter. The coarse 3-bit most-significant bit conversion is converted back to an analog signal using a 3-bit sub-digital to analog converter (SDAC). Then the SDAC output is subtracted from the SHA output, the difference is amplified, and this residue signal is digitized by a second-stage, 3-bit SADC to generate the three LSBs of the total 6-bit output word. This architecture is useful for resolutions up to about 8 bit; however, maintaining better than 8-bit alignment between the two stages (over temperature variations, in particular) can be difficult. There is no particular requirement for an equal number of bits per stage in the subranging architecture. In addition, there can be more than two stages (see Rapuano et al. [B46]).

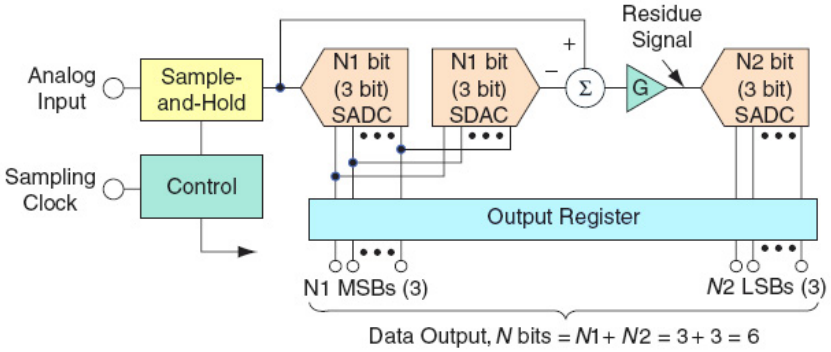

Reprinted from Rapuano et al., *IEEE Instrumentation and Measurement Magazine,* 2005 [B46]. © 2005 IEEE.

Figure A.3—Subranging ADC

To increase the speed of the basic subranging ADC, the pipelined architecture shown in Figure A.4 has become very popular. This pipelined ADC has a digitally corrected subranging architecture in which each of the two stages operates on the data for one half of the conversion cycle and then passes its residue output to the next stage in the pipeline prior to the next phase of the sampling clock. The interstage track-and-hold (T/H) serves as an analog delay line; it is timed to enter the hold mode when the first-stage conversion is complete. This allows more settling time for the internal SADCs, SDACs, and amplifiers and allows the pipelined converter to operate at a much higher overall sampling rate than a non-pipelined version.

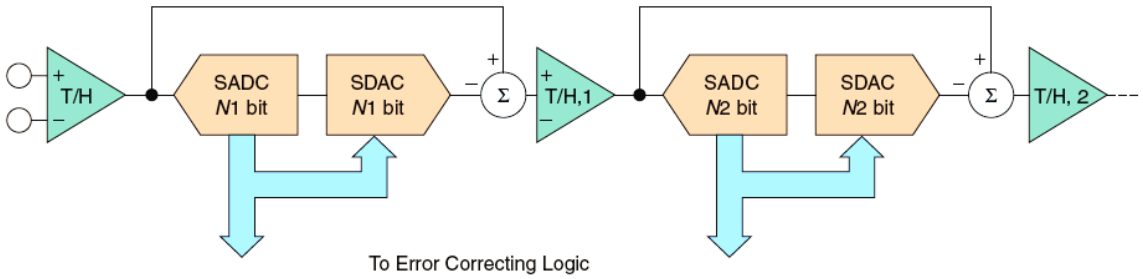

Reprinted from Rapuano et al., *IEEE Instrumentation and Measurement Magazine,* 2005 [B46]. © 2005 IEEE.

Figure A.4—Pipelined ADC architecture

There are many design tradeoffs that can be made in the design of a pipelined ADC, such as the number of stages, the number of bits/stage, number of correction bits, and the timing. An *N*-bit converter can require only *N* comparators, and the output of each comparator is a bit in the final converter output.

To cause the digital data from the individual stages corresponding to a particular sample arrives at the error correction logic simultaneously; the appropriate number of shift registers must be added to each of the outputs of the pipelined stages (see Rapuano et al. [B46]).

The disadvantages of the pipeline architecture are that the comparators must be very precise to prevent error compounding down the pipeline, and there is an *N* clock latency (for an *N*-bit converter) as the pipeline is filled.

However, after the initial latency, a new output is produced for each subsequent clock cycle (see Rauth and Randal [B47]).

It is important to clarify the distinction between subranging and pipelined ADCs. Although pipelined ADCs are generally subranging (with error correction), subranging ADCs are not necessarily pipelined. As a matter of fact, the pipelined subranging architecture is predominant because of the demands for high sampling rates, where internal settling time is of utmost importance (see Rapuano et al. [B46]).

A.4 SAR ADCs

The conversion technique based on an SAR employs a comparator to weigh the applied input voltage against the output of an *N*-bit DAC. Using the DAC output as a reference, this process approaches the final result as a sum of *N* weighting steps, in which each step is a single-bit conversion (see Rapuano et al. [B46]).

The successive approximation converter illustrated in Figure A.5, includes a shift register, a S/H circuit, a comparator, an output register (successive approximation register or SAR), and a DAC.

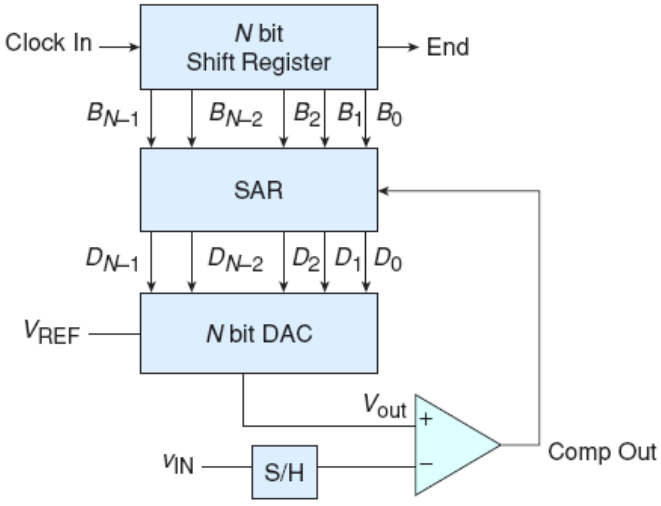

Reprinted from Rapuano et al., *IEEE Instrumentation and Measurement Magazine,* 2005 [B46]. © 2005 IEEE.

Figure A.5—SAR ADC architecture

The easiest way of understanding the successive approximation converter is algorithmically (see Rauth and Randal [B47]).

1) A 1 is applied to the shift register input. For each bit converted, the 1 is shifted to the right 1-bit position. $B_{N-1} = 1$ and B_{N-2} through $B_0 = 0$.
2) The MSB of the SAR, D_{N-1}, is initially set to 1, while the remaining bits D_{N-2} through D_0 are set to 0.
3) Since the SAR output controls the DAC and the SAR output is 100 ...0, the DAC output will be set to $V_{REF}/2$.
4) Next, V_{IN} is compared to V_{REF} /2. If V_{REF} /2 is greater than V_{IN}, then the comparator output is a 1 and the comparator resets D_{N-1} to 0. If V_{REF} /2 is less than VIN, the comparator output is a 0 and the D_{N-1} remains a 1. D_{N-1} is the actual MSB final digital output code.
5) The 1 applied to the shift register is then shifted by one position so that $B_{N-2} = 1$, while the remaining bits are all 0.
6) D_{N-2} is set to a 1, D_{N-3} through D_0 remain 0, while D_{N-1} remains the value from the MSB conversion.
7) The output of the DAC will now either equal V_{REF} /4 (if $D_{N-1} = 0$) or $3V_{REF}/4$ (if $D_{N-1} = 1$) and the comparison repeats.

SAR ADCs are frequently the architecture of choice for inexpensive applications that need medium to high resolution; typically with sample rates less than 5 MS/s. SAR ADCs most commonly range in resolution from 8 bits to 18 bits and provide low power consumption as well as a small form factor (see Rapuano et al. [B46]).

A.5 Σ-Δ ADCs

Σ-Δ ADCs have relatively simple structures. Also called over-sampling converters, they consist of a Σ-Δ modulator followed by a digital decimation filter. The modulator, whose architecture is similar to that of a dual-slope ADC, includes an integrator and a comparator with a feedback loop that contains a 1-bit DAC. This internal DAC is simply a switch that connects the comparator input to a positive or negative reference voltage. The Σ-Δ ADC also includes a clock unit that provides proper timing for the modulator and digital filter (Figure A.6). Through a series of iterations, the integrator, comparator, DAC, and summing junction produce a serial bitstream that represents the oversampled input voltage. Once digitized, the oversampled signal goes through a digital filter to remove frequency components at or above the Nyquist frequency. A decimator then removes the oversampled data (see Rauth and Randal [B47]).

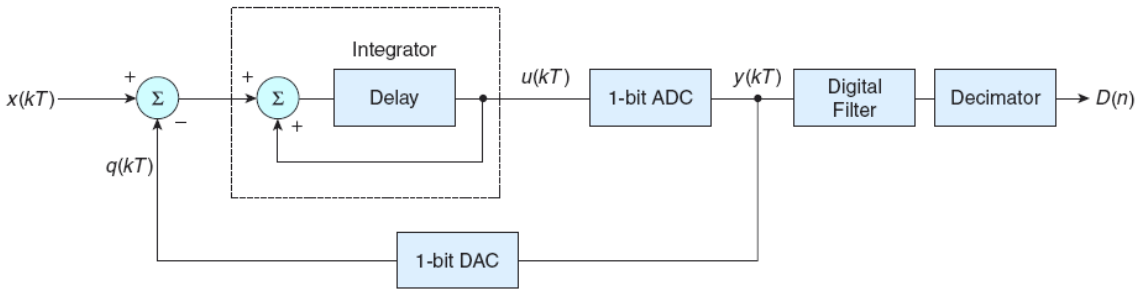

Reprinted from Rapuano et al., *IEEE Instrumentation and Measurement Magazine,* 2005 [B46]. © 2005 IEEE.

Figure A.6—Σ-Δ converter architecture

Low-bandwidth signals applied to the input of a Σ-Δ ADC are quantized with very low (1 bit) resolution, but with an exceedingly high, internal sampling frequency (100 MS/s is typical). Combined with digital post-filtering, this oversampling reduces the sampling rate to kilosamples/second (kS/s) or hundreds of kS/s and increases the ADC resolution (i.e., dynamic range) to 16 bits or more. Although slower than pipeline ADCs and limited to lower input bandwidths, the Σ-Δ principle has developed a strong position in the data-converter market. It offers four major advantages: low cost, low power, high-resolution conversion, and DSP compatibility for system integration because of the digital filter included with the conversion circuitry. Σ-Δ ADCs were originally used predominately in lower speed applications requiring a trade-off of speed for resolution (Rapuano et al. [B46]), but newer designs can also used in many higher-speed applications.

A.6 Time-Interleaved ADCs

Interleaving multiple ADCs is usually performed to increase the effective sampling rate, especially if there are no off-the-shelf ADCs available that fulfil the desired sample rate, linearity, and ac requirements of such applications. However, time-interleaving data converters is not an easy task, because even with perfectly linear components, gain/offset mismatches and timing errors can cause undesired spurs in the output spectrum. The simplified block diagram in Figure A.7 sketches a single-channel, time-interleaved data acquisition system in which two ADCs double the system's sampling rate. This rate ($f_{\text{SYSTEM CLK}}$) is a clock signal at twice the rate of $f_{\text{CLK1}} = f_{\text{CLK2}}$. Because f_{CLK1} is delayed with respect to f_{CLK2} by the period of $f_{\text{SYSTEM CLK}}$, the two ADCs sample the analog input signal alternately, producing an overall sample rate equal to $f_{\text{SYSTEM CLK}}$. Each converter operates at half the sampling frequency (Rapuano et al. [B46]).

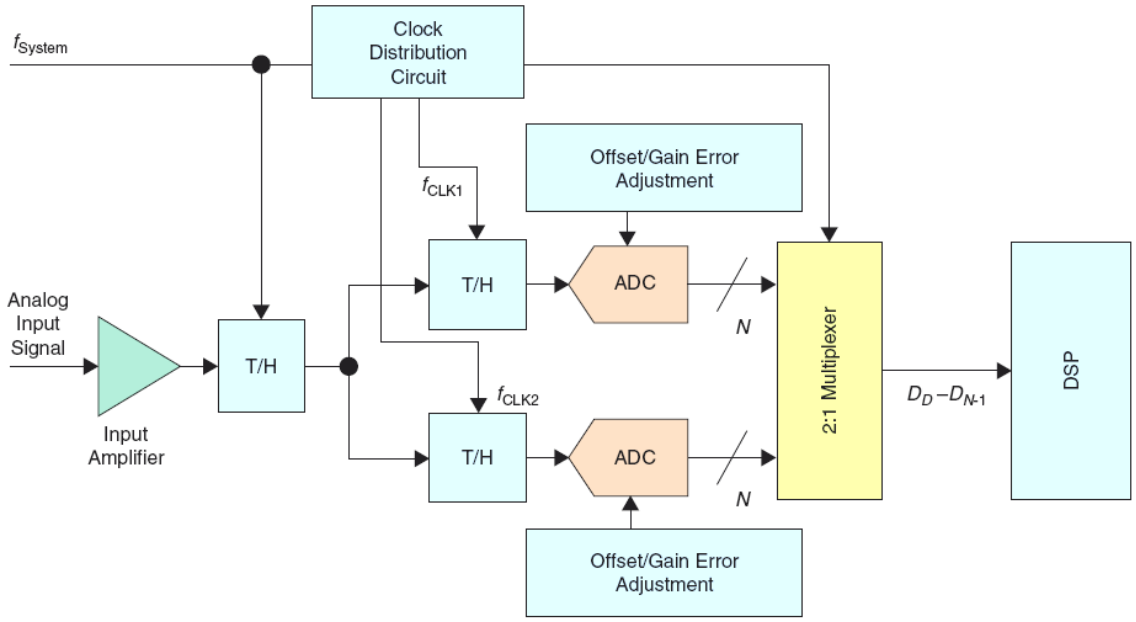

Reprinted from Rapuano et al., *IEEE Instrumentation and Measurement Magazine,* 2005 [B46]. © 2005 IEEE.

Figure A.7—Time-interleaved data acquisition system

The channel-to-channel matching of offset and gain in separate ADCs are parameters of concern in a time-interleaved system. If one ADC shows an offset and the other a gain error, the digitized signal represents not only the original input signal but also an undesired error in the output. An offset discrepancy and gain mismatches show up as anomalies in the output. For interleaving designs, it is therefore necessary to choose ADCs with integrated gain and offset correction or include external circuitry that corrects these mismatches.

Integral nonlinearity (INL) of ±1 LSB is quite common for individual ADCs, but in an interleaving system such errors can easily double. The appearance of nonlinearity introduces distortion into the system, which degrades dynamic parameters such as SINAD and ENOB. Most of the errors discussed above can be overcome using calibration procedures in the time domain, careful circuit design and layout, a suitable selection of data converters, and digital post-processing. Unfortunately, this approach is complex and entails extra cost, lengthy calibration, and mathematical analysis (see Rapuano et al. [B46]).

A.7 Folding and Interpolating ADCs

One way to increase the resolution of flash converters is to employ analog pre-processing. There are two ways to perform this: interpolation and folding. By using interpolation, certain signals or reference levels are not generated. Instead, the existing signals are used to interpolate these signals (using resistor strings and preamps) before applying them to the comparators. The main advantage of this is the reduction of input capacitance.

The folding principle aims at reducing the number of comparators by performing continuous-time subranging conversion (see Rapuano et al. [B46]). By connecting the outputs of two or more differential pairs, a linear input signal is folded repeatedly into different sectors. The folded signal is then quantized to provide the least significant bits (LSBs). The loss of most significant bits (MSBs) information caused by folding is recovered using a separate coarse quantizer. Compared to flash architecture, the number of comparators required for the folding architecture is cut to $2^N/M$, where M is the number of analog folding. This will reduce the power consumption, chip area, and device count. The number of comparators can be further reduced using interpolation technique, where resistor networks across different comparator outputs are used to generate extra zero crossings.

The complete system diagram of a folding and interpolation 8-bit ADC is shown in Figure A.8. Interpolation technique reduces the number of preamplifiers at the input of the ADC. The error sources for these kinds of ADCs include, in addition to the standard error sources, the reference inaccuracy, the folding amplifier input offset, the tail-current mismatch (gain error of segments), and the interpolation error (see Rapuano et al. [B46]).

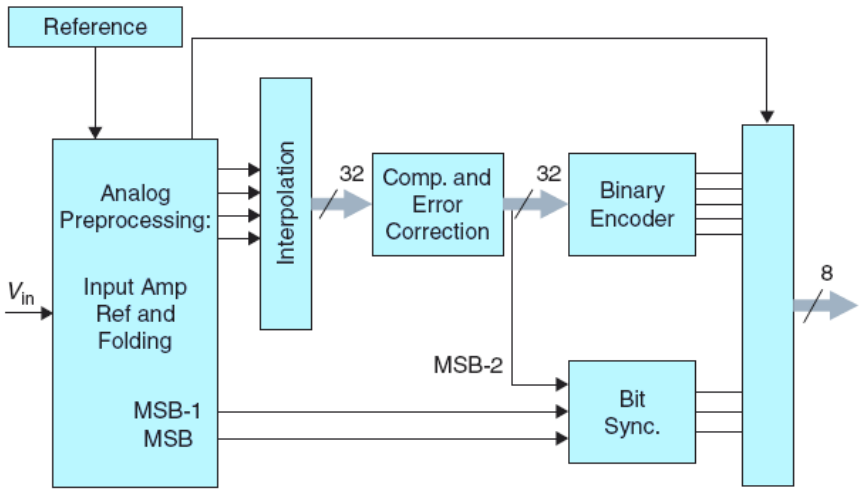

Reprinted from Rapuano et al., *IEEE Instrumentation and Measurement Magazine,* 2005 [B46]. © 2005 IEEE.

Figure A.8—Folding and interpolation conversion scheme

Annex B

(informative)

Sine-wave fitting algorithms

B.1 An algorithm for three-parameter (known frequency) least-squares fit to sine-wave data

This algorithm provides a least-squares method for fitting digitized waveform data to a sine wave in the case where the frequency of the sine wave is known. The algorithm is presented using matrix notation.

The equation for the sine wave is given in Equation (16).

Assuming the data record contains the sequence of M samples $x[1]$, $x[2]$, …,$x[M]$ taken at times t_1, t_2, …, t_M, this algorithm finds the values of A_0, B_0, and C_0 that minimize the following sum of squared differences [see Equation (B.1)]:

$$\sum_{n=1}^{M}\left[x[n]-A_0\cos(2\pi f_0 t_n)-B_0\sin(2\pi f_0 t_n)-C_0\right]^2 \tag{B.1}$$

where

f_0 is the known frequency of the sine wave applied to the ADC input.

Since the equation is linear, a closed form (noniterative) solution can be computed. To find the values for A_0, B_0, and C_0, first create Matrix (B.2), Matrix (B.3), and Matrix (B.4).

$$D_0=\begin{bmatrix}\cos(2\pi f_0 t_1) & \sin(2\pi f_0 t_1) & 1\\ \cos(2\pi f_0 t_2) & \sin(2\pi f_0 t_2) & 1\\ \cdot & \cdot & \cdot\\ \cdot & \cdot & \cdot\\ \cos(2\pi f_0 t_M) & \sin(2\pi f_0 t_M) & 1\end{bmatrix} \tag{B.2}$$

$$x=\begin{bmatrix}x[1]\\ x[2]\\ \cdot\\ \cdot\\ x[M]\end{bmatrix} \tag{B.3}$$

$$s_0=\begin{bmatrix}\mathrm{A}_0\\ \mathrm{B}_0\\ \mathrm{C}_0\end{bmatrix} \tag{B.4}$$

In matrix notation, the sum of squared differences in Equation (B.1) is given by that shown in Equation (B.5).

$$\left(x-D_0 s_0\right)^T\left(x-D_0 s_0\right) \tag{B.5}$$

where $(*)^T$ designates the transpose of the vector or matrix (*).

The least-squares solution, $\hat{s}_0$, that minimizes Expression (B.5) is given by:

$$\hat{s}_0 = \left(D_0{}^T D_0\right)^{-1}\left(D_0{}^T x\right) \tag{B.6}$$

The components of $\hat{s}_0$ can then be used in Equation (16) to compute the fitted function. Note that although the value of the least squares fit is given by Equation (B.6), one might compute the value by a more numerically stable method, such as the Q-R decomposition (see Stewart [B56], pp. 208–249).

To convert to the amplitude and phase (polar) form:

$$y_n{}' = A_0 \cos(\omega_0 t_n) + B_0 \sin(\omega_0 t_n) + C_0 \tag{B.7}$$

To convert to the amplitude and phase (polar) form:

$$x_n = A\cos(2\pi f_0 t_n + \phi) + C \tag{B.8}$$

use Equation (B.9), Equation (B.10), and Equation (B.11).

$$A = \sqrt{A_0^2 + B_0^2} \tag{B.9}$$

$$\varphi = \tan^{-1}\left[\frac{B_0}{A_0}\right] \quad \text{if } A_0 > 0 \tag{B.10}$$

$$\varphi = \tan^{-1}\left[\frac{B_0}{A_0}\right] \pm \pi \quad \text{if } A_0 < 0 \tag{B.11}$$

The residuals, r_n, of the fit are given by Expression (B.12).

$$r[n] = x[n] - A_0 \cos(2\pi f_0 t_n) - B_0 \sin(2\pi f t_n) - C_0 \tag{B.12}$$

and the root-mean-square (rms) error is given by Equation (B.13).

$$e_{\text{rms}} = \sqrt{\frac{1}{M}\sum_{n=1}^{M} r[n]^2} \tag{B.13}$$

B.2 An algorithm for four-parameter least-squares fit to sine-wave data

This algorithm provides a least-squares method for fitting digitized data to a sine wave in the case where the frequency of the sine wave is not known. The equation for the sine wave is the same as that given in Equation (16) except that the input sine-wave frequency, f_0, is not known. Since the equation is no longer linear, a closed form solution cannot be computed, and iterative techniques must be used.

Assuming the data record contains the sequence of M samples, $x[1]$, $x[2]$,…, $x[M]$, taken at times t_1, t_2, …, t_M, this algorithm uses an iterative process to estimate the parameters A_i, B_i, C_i, and f_i, which minimize the following sum of squared differences:

$$\sum_{n=1}^{M}\left[x[n]-A_i\cos(2\pi f_i t_n)-B_i\sin(2\pi f_i t_n)-C_i\right]^2 \tag{B.14}$$

where

i is the iteration number

The algorithm is given below.

a) Set iteration index $i = 0$.

b) Make an initial estimate of the frequency f_0 of the recorded data. The frequency may be estimated by using a DFT (either on the full record or a portion of it), or by taking the inverse of the average time between zero crossings, or simply by using the best measurement of the applied input frequency. A very effective general method is to use the interpolated Fast Fourier Transform described in Schoukens et. al. [B49] and Bilau et al. [B6]. This uses an interpolation formula on the DFT coefficients of the data record and is used in the software available at Markus [B37]

c) Perform a pre-fit using the three-parameter matrix algorithm to estimate A_0, B_0, and C_0.

d) Set $i = i + 1$ for the next iteration.

e) $f_i = f_{i-1} + \Delta f_{i-1} \qquad (\Delta f_{i-1} = 0 \quad \text{for } i = 1)$

Create the following matrices:

$$x = \begin{bmatrix} x[1] \\ x[2] \\ . \\ . \\ x[M] \end{bmatrix} \tag{B.15}$$

$$D = \begin{bmatrix} \cos(2\pi f t_1) & \sin(2\pi f t_1) & 1 & -At_1\sin(2\pi f t_1)+Bt_1\cos(2\pi f t_1) \\ \cos(2\pi f t_2) & \sin(2\pi f t_2) & 1 & -At_2\sin(2\pi f t_2)+Bt_2\cos(2\pi f t_2) \\ . & . & . & . \\ . & . & . & . \\ \cos(2\pi f t_M) & \sin(2\pi f t_M) & 1 & -At_M\sin(2\pi f t_M)+Bt_M\cos(2\pi f t_M) \end{bmatrix} \tag{B.16}$$

$$s_i = \begin{bmatrix} A_i \\ B_i \\ C_i \\ \Delta f_i \end{bmatrix} \tag{B.17}$$

f) Compute the least-squares solution, $\hat{s}_i$,

$$\hat{s}_i = \left(D_i^T D_i\right)^{-1}\left(D_i^T x\right) \tag{B.18}$$

Note that although the value of the least-squares fit is given by Equation (B.18), one might compute the value by a more numerically stable method, such as the Q-R decomposition (see Stewart [B56]).

g) Update the frequency estimate using:

$$f_i = f_{i-1} + \Delta f_{i-1} \qquad (\Delta f_{i-1} = 0 \quad \text{for } i = 1) \tag{B.19}$$

h) Compute the amplitude, A, and phase, φ, for the polar form.

$$x[n] = A\cos(2\pi f_i t_n + \phi) + C \tag{B.20}$$

using

$$A = \sqrt{A_i^2 + B_i^2} \tag{B.21}$$

and

$$\phi = \text{atan2}[B_i, A_i] \tag{B.22}$$

i) Repeat Steps d) through h), re-computing the model based on the new values of A_i, B_i, and f_i, calculated from the previous iteration. It appears to be best to iterate a fixed number of times. Based on experience, six iterations have proven more than adequate. This method doubles the number of significant digits in f at each iteration and converges very rapidly.

The residuals, $r[n]$, of the fit are given by:

$$r[n] = x[n] - A_i\cos(2\pi f_i t_n) - B_i\sin(2\pi f_i t_n) - C_i \tag{B.23}$$

and the rms error is given by

$$e_{\text{rms}} = \sqrt{\frac{1}{M}\sum_{n=1}^{M} r[n]^2} \tag{B.24}$$

Annex C

(normative)

Discrete Fourier transforms and windowing

The discrete Fourier transform (DFT) is a mathematical operation that converts sampled data from the time domain to the frequency domain; see Oppenheim and Willsky [B44] or Stearns and Hush [B55]. (See IDFT below for the inverse operation.) In this standard, the discrete Fourier transform (DFT) is defined in Equation (C.1) for a record of data $x[n]$ that is M samples long.

$$X[k] = \sum_{n=0}^{M-1} x[n] \times \exp(-j2\pi kn / M) \qquad \text{for } 0 \le k \le M-1 \tag{C.1}$$

where

M is the number of sequential samples in the data record

The inverse discrete Fourier transform (IDFT) is the mathematical operation that converts sampled data from the frequency domain to the time domain. The formula is:

$$x[n] = \frac{1}{M} \sum X[k] \times \exp(+j2\pi kn / M)$$

$$x[n] = \frac{1}{M} \sum_{n=0}^{M-1} X[k] \times \exp(+j2\pi kn / M) \tag{C.2}$$

Note that in this standard, the IDFT includes the requisite $(1/M)$ normalization. Other definitions of a DFT may use other normalization factors.

The frequency corresponding to $X[k]$, f_k, is

$$f_k = \begin{cases} \dfrac{k}{M} f_s & \text{for } 0 \le k \le M/2 \\ -\left(1 - \dfrac{k}{M}\right) f_s & \text{for } M/2 + 1 \le k \le M-1 \end{cases} \tag{C.3}$$

where

f_k is the frequency corresponding to the k^{th} DFT component
f_s is the sampling frequency

Note that the second half of the frequency components can be interpreted as negative frequencies, and for real signals their transform values are the complex conjugates of the values for the corresponding positive frequencies. Sometimes, to simplify mathematical expressions (such as modulo), the upper formula will be used for all k, mapping the negative frequencies to frequencies between the Nyquist frequency and twice the Nyquist frequency. The frequencies are separated by an interval of $\Delta f = f_s/M$. The interval of frequencies from $f_k - \Delta f/2$ to $f_k + \Delta f/2$ is called the k^{th} frequency bin. Associated with any frequency, f, is the

bin number, $f/\Delta f$. If f is at the center of the k^{th} frequency bin, the bin number will be the integer, k, Non-integer values of the bin number correspond to frequencies that are not at the center of a bin.

$X[f_k]$ may be written in place of $X[k]$ to make the frequency dependence more obvious.

Example DFT plots are shown in Figure C.1 and Figure C.2 below. Both show the DFT of a signal sampled at 100 samples per second. The Nyquist frequency is 50 samples per second. The signal is a 21 Hz sine wave with amplitude 1 with second harmonic distortion with amplitude 0.5 and third harmonic distortion with amplitude 0.25. Figure C.1 shows the DFT spectrum using positive and negative frequencies. All signals are symmetrical having the same magnitude for positive and negative frequencies. The third harmonic, which is at 63 Hz is aliased down to 37 Hz. Figure C.2 shows the negative frequencies mapped up to the frequencies between the Nyquist frequency and twice the Nyquist frequency. They are moved to the right 100 frequency units (the sampling frequency). Note that one of the third harmonic terms is at the correct frequency of 63 Hz. However the other frequencies that are mapped for the negative side are at locations that are not very intuitive. The Figure C.1 gives a better visual picture, while the Figure C.2 is often more convenient in mathematical calculations, specifically in the tracking of frequencies that are above the Nyquist frequency.

The "energies" in the time and frequency representations of the data are related by the Parseval relation

$$\sum_{n=0}^{M-1}|x[n]|^2 = \frac{1}{M}\sum_{k=0}^{M-1}|X[k]|^2 \tag{C.4}$$

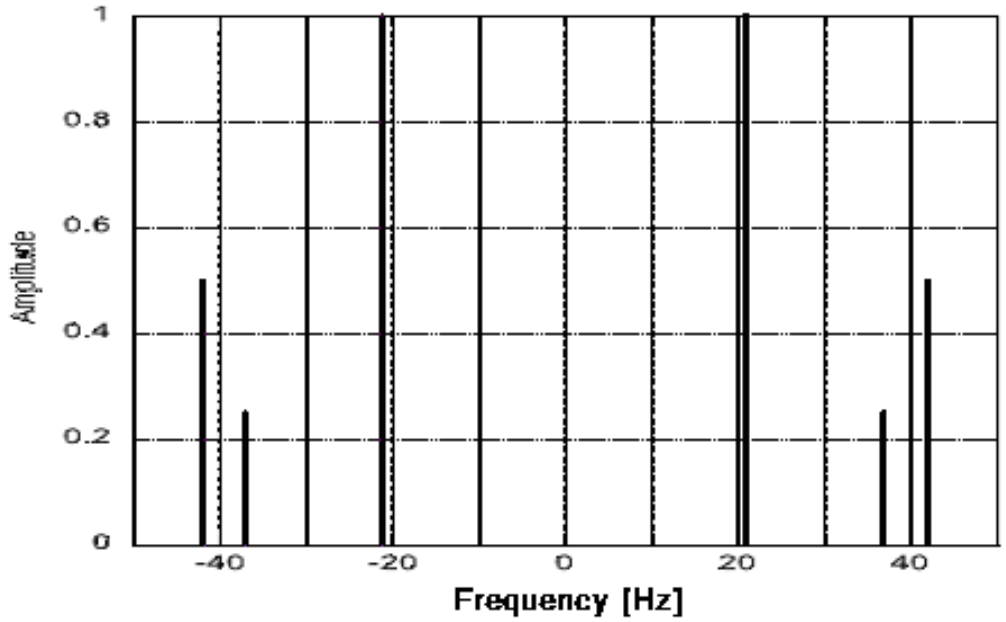

Figure C.1—Example DFT spectrum using positive and negative frequencies

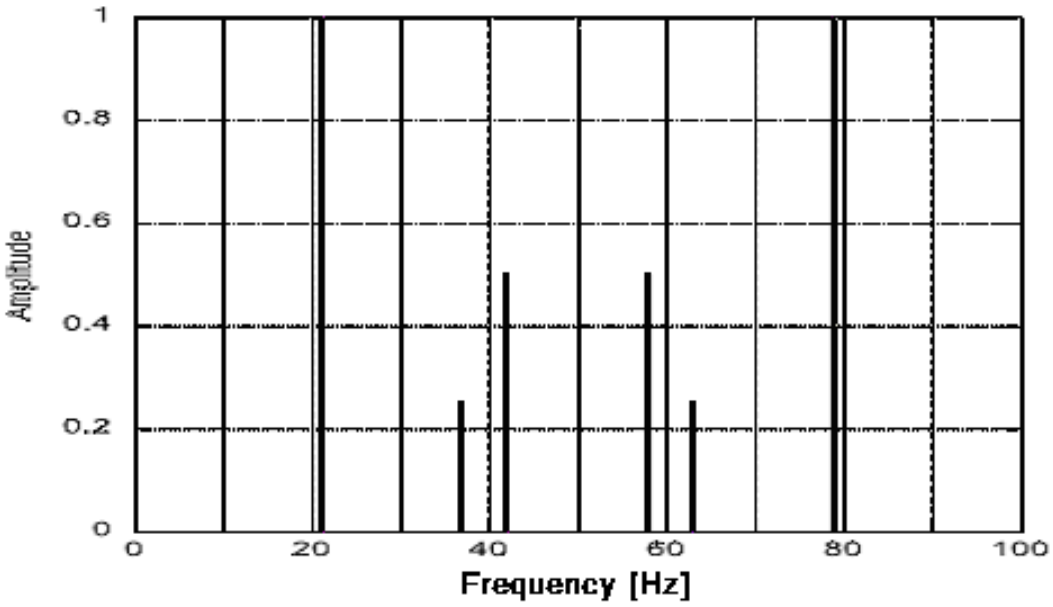

Figure C.2—Example DFT showing the negative frequencies mapped up to the frequencies between the Nyquist frequency and twice the Nyquist frequency

C.1 The windowed DFT and spectral leakage

A windowed DFT is defined in Equation (C.5). The window factors, $w[n]$, are near one at the center of the data interval and approach zero at the ends of the interval.

$$X_w[k] = \sum_{n=0}^{M-1} w[n]x[n]\exp(-j2\pi nk/M) \tag{C.5}$$

Windowing is multiplication in the time domain and corresponds to convolution in the frequency domain. Some authors refer to the function $w[n]$ as the temporal weighting function and its Fourier transform, $W[k]$, as the window function; whereas some refer to $w[n]$ as the window function. When all the window factors are one, i.e., no windowing, the window is called the rectangular window. This rectangular window is also called non-weighting. Window functions other than the rectangular are used to reduce the effects of spectral leakage, which is described in the next subclause.

C.1.1 Spectral leakage

The rectangular-windowed DFT of a sinusoid with a frequency at the center of a DFT bin produces a single spectral line. If the frequency of the sinusoid is not at the center of a bin, the DFT will produce lines in all frequency bins (see Harris [B19]). The leakage is described by the continuous-time Fourier transform, $W[f]$, of the window function. Two examples are shown in Figure C.3 below. The solid line is the Hann window (defined in C.2), and the dashed line is the rectangular window. The horizontal axis is the frequency converted to DFT bin number, and the vertical axis is $W[f]$ expressed in decibels. Since $W[-f] = W[f]$, the data are not shown for negative values. If the signal being analyzed is a sinusoid with amplitude A and frequency with a bin number of x (not necessarily an integer), the k^{th} spectral line will have an amplitude of $A \times W[k-x]$.

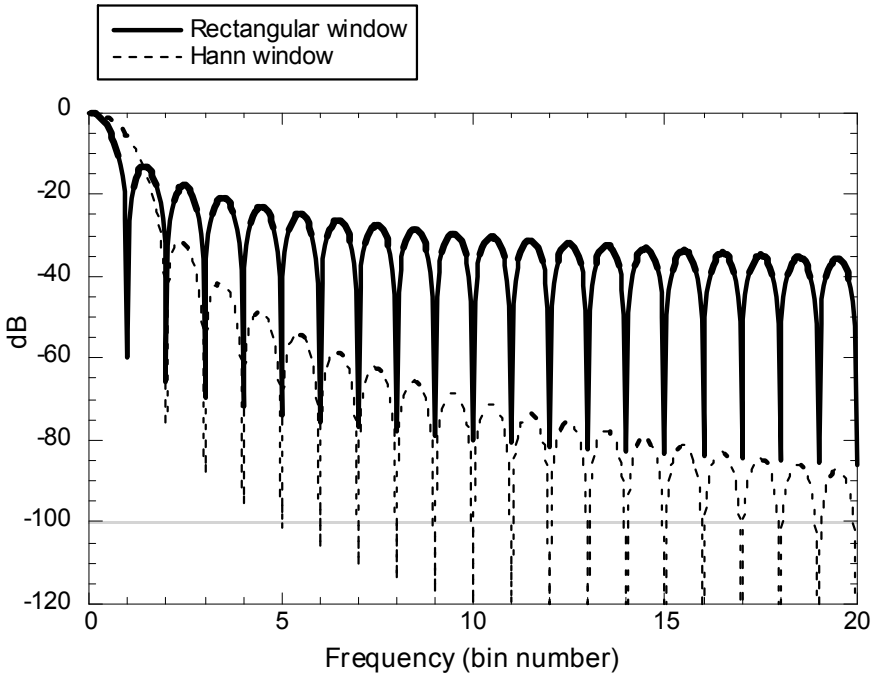

Figure C.3—Fourier transform of the Hann and the rectangular windows

The negative spikes in the plot are where the function is zero. For the rectangular window the function is zero at all integers except zero, while the function for the Hann window is zero at all integers except zero and ±1. Thus, for a coherently sampled sinusoid (one with an integer bin number) the rectangular window will produce a single line, while the Hann window will produce three lines. However, for incoherent sampling both windows produce lines at all frequencies. The amplitudes of the lines fall off much more slowly for the rectangular window than for the Hann window.

Figure C.4 shows the same information for the Hann window (solid) and the Blackman window (dashed, also defined in C.2). Note that the difference is much smaller than the difference between the rectangular and Hann windows. This illustrates that once the user chooses to use a good window other than the rectangular, it is often not critical which one is chosen.

C.1.2 Coherent sampling and sine-fitting methods of reducing spectral leakage

There are several methods of dealing with spectral leakage other than the judicious choice of windows. The basic problem is that the energy of some large sinusoidal component of the signal leaks into spectral bins that have smaller components that need to be identified. This is often solved simply by choosing a sufficiently long record length.

As seen from the figures in the previous section, the simplest windows (Hann and Blackman) are both down by more than 87 dB at an offset of 20 frequency bins. As the record becomes longer, the frequency deviation corresponding to 20 frequency bins becomes smaller. For example, with a record length of 4096, 20 frequency bins is less than 1% of the Nyquist frequency.

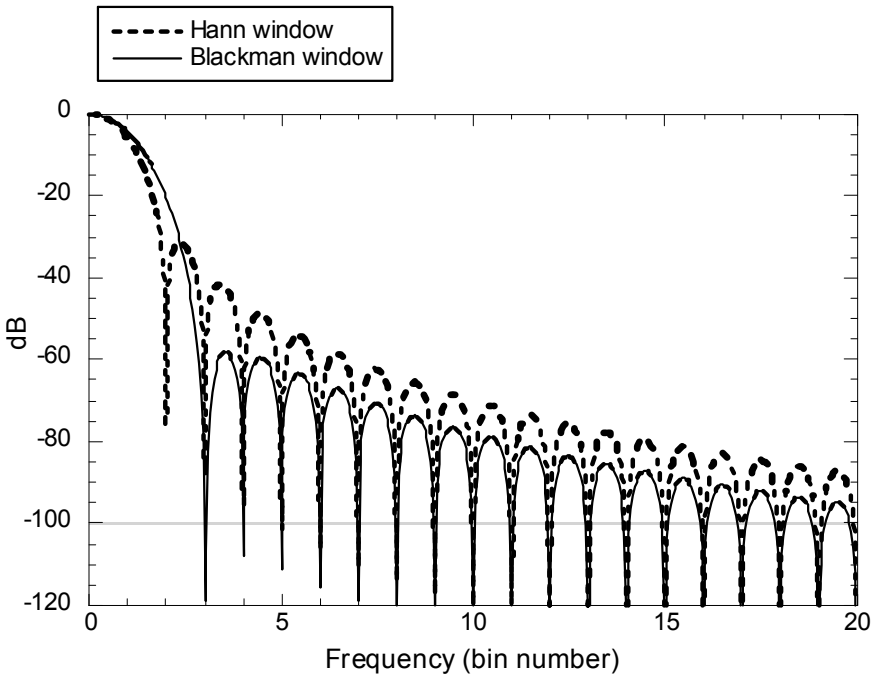

Figure C.4—Fourier transforms of the Hann and the Blackman windows

Coherent sampling, in which the bin number of the applied signal is an exact integer, is the usually the best approach for dealing with spectral leakage. It eliminates the leakage problem completely if exact coherence is obtained. The recommended frequencies given elsewhere in this standard, for best achieving other ends, typically result in coherent sampling. However, sometimes the accuracy and/or precision of the frequency setting of available oscillators does not allow for coherent sampling.

Sine fitting is another approach to reduce leakage problems. Use the four-parameter sine-fitting algorithm described in B.2 to determine the frequency, amplitude and phase of large noncoherent sinusoidal components of the data record. Then subtract these components from the data record and perform the DFT on the result. This typically reduces leakage problems by 30 dB to 50 dB. The user can also first truncate the record to a length containing approximately an integer number of periods.

C.2 Some useful windows and their characteristics

The only window parameter used in this standard is the normalized noise power gain (NNPG). The normalized noise power gain is given by

$$\text{NNPG} = \frac{1}{M}\sum_{n=0}^{M-1} w[n]^2 \tag{C.6}$$

The NNPG is a conversion factor that gives a Parseval-like relationship for the windowed DFT,

$$\frac{1}{M}\sum_{n=0}^{M-1} x[n]^2 = \frac{1}{M \times NNPG}\sum_{k=0}^{M-1} X_w[k]^2 \tag{C.7}$$

where

$X_w[i]$ is the Fourier transform of a windowed input sequence $w[i]\,x[i]$

The value of NNPG is one for the rectangular window and is less than one for all other windows.

The windows described here and their characteristics are covered in detail in Harris [B19]] and Nuttall [B41]. These windows are of the form

$$w[n] = \sum_{l=0}^{L} a_l (-1)^l \cos(2\pi n l / M) \tag{C.8}$$

where

a_l is the lth coefficient used to define the window function. The sum of the coefficients equals one
L is the order, or number of coefficients in the window

These windows are commonly used and are discussed at length in Nutall [B41]. Any window of this form is referred to as a cosine window of order L.

For these windows the coefficients of the DFT of the window function are proportional to the coefficients multiplying the cosine terms. There is one nonzero coefficient at zero frequency, L nonzero terms at positive frequency and L at negative frequency. Therefore, a coherently sampled sine wave produces L lines on each side of the central bin.

Equation (C.6), the Parseval relation, and the fact the a_i are the Fourier coefficients of the window function yields

$$\text{NNPG} \cong a_0^2 + \frac{1}{2}\sum_{l=1}^{L} a_i^2 \tag{C.9}$$

The windows shown earlier are described in Table C.1.

Table C.1—Properties of various data windows

Name	L	NNPG	a_0	a_1	a_2
Rectangular	0	1	1	0	0
Hann	1	0.625	0.5	0.5	0
Blackman	2	0.305	0.42	0.5	0.08

C.3 Window selection

The spectral leakage from a large sinusoidal component can interfere with some other feature one wishes to measure. The magnitude of the problem is determined by inspection of the Fourier transform of the window sequence, as described earlier, and depends on the size of the interfering component(s) and their distance from the desired smaller components. If the selected window causes an interference that is too large one must either select a window that is sufficiently small at the location that is important for the particular problem or use one of the other methods for ameliorating spectral leakage.

Annex D

(informative)

Presentation of sine-wave data

This annex describes common ways of presenting and displaying the results of sine-wave tests.

D.1 ENOB presentation

The most common specification that results for sine-wave tests is the effective number of bits (ENOB). This gives an overall measure of several different sources of error. Its value will depend on the frequency and the amplitude of the test signal. Figure D.1 shows an example of the presentation of the ENOB results.

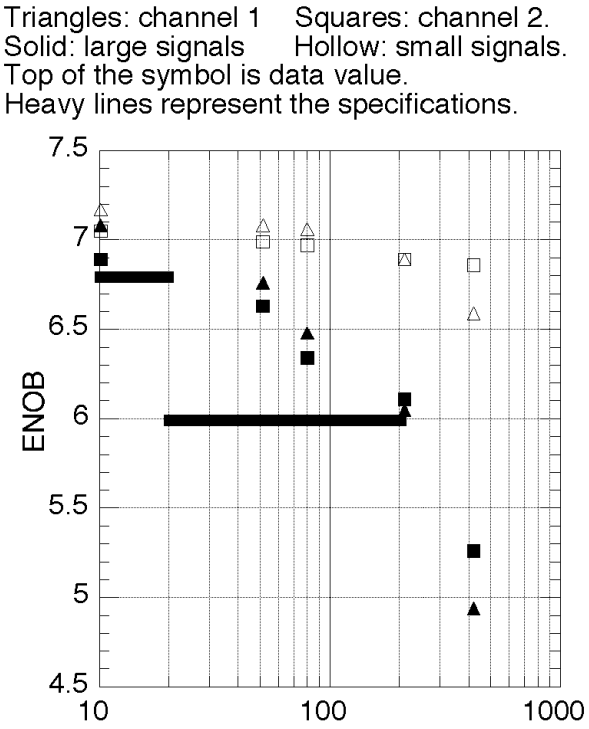

Figure D.1—Example plot of ENOB vs. frequency for a two-channel ADC

The value of ENOB is plotted as a function of frequency, with the frequency on a logarithmic scale. Different test conditions, in this case different signal amplitudes and different channels, can be represented with different symbols. In this example the specifications of the ADC are shown as a curve (a solid line). Points above the curve meet the specifications, while those below do not.

The decrease in ENOB at high frequencies is typical and is usually caused by errors in the time values of the data points. A simple method for estimating the time-base errors from the ENOB plot is given by Blair in [B10] and [B11]. The method of 12.2.2 of IEEE Std 1057-2007 [B23] can also be used to estimate time-base errors from the reduction of ENOB at high frequency. This requires first converting from ENOB back to NAD using the defining formula for ENOB in 9.4.

The ENOB is based on NAD, and the NAD can be decomposed into its separate noise and distortion components. The noise and distortion can be plotted individually as above with the vertical scale in decibels.

D.2 Presentation of residuals

Much can be learned by viewing the residuals of a particular sine-wave test. In particular, one can determine the sources of error that are responsible for an ENOB that is less than expected. The residuals can be viewed in either the time domain or the frequency domain, and it is most informative to look at both. First, the time and frequency domain displays that are used in software available from the ADC Test Software web site at the Budapest University of Technology and Economics (see Markus [B37]) are covered. The time-domain presentation is called the modulo-time plot and is described in detail by Irons and Hummel [B25]. The frequency domain presentation is called the power spectral distribution (PSD) and is described by Blair [B8]. Typical displays of each are shown in Figure D.2.

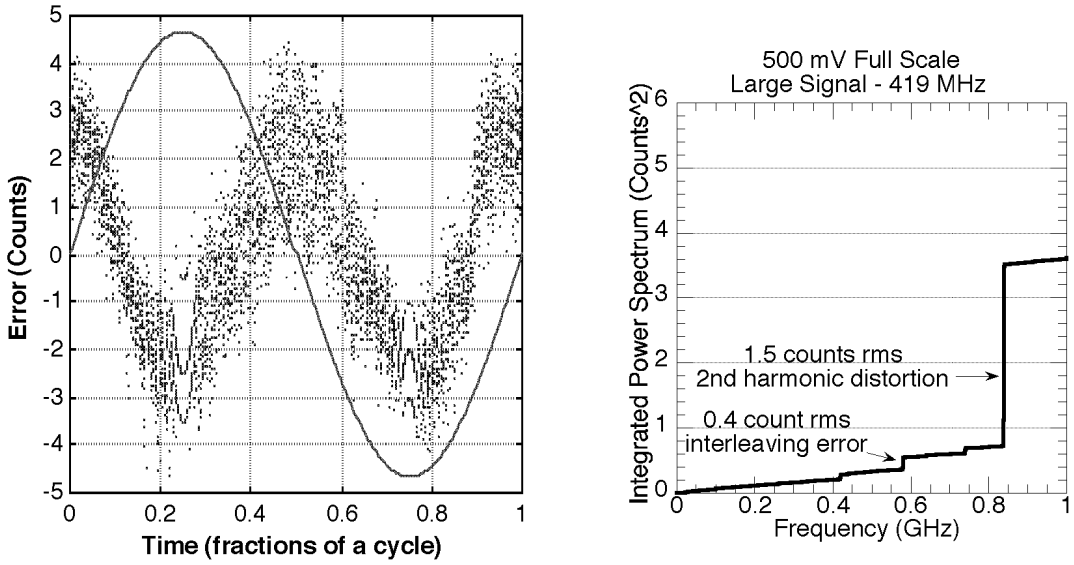

Figure D.2—Example plots of sine-fit residuals. The left side shows a modulo time plot, and the right side shows a power spectral distribution. Both represent the same data.

The construction of the modulo time plot is described first. From the time coordinate of each data point in the record the phase value, between 0 and 2π, relative to the input sine wave, is calculated as determined by the fit. This phase angle, divided by 2π, is shown on the horizontal axis of the plot. The residual value is shown on the vertical axis. On this particular plot the units are LSB, but any other units, such as volts, might be used. A scaled replica of the fitted input signal is displayed as the continuous curve. On this plot the second harmonic distortion of approximately ±2.5 LSB superimposed on random noise of about ±2 LSB can be clearly seen. The plot clearly displays the phase relationship of the harmonic to the input signal, having its negative peaks at the peaks of the input signal and having its positive peaks at the zero crossings.

The PSD is the integral of the power spectral density. Its units are the square of units in which signal amplitude is measured, e.g., volts, some other physical quantity or, as in the case of Figure D.2, LSBs. The value of the PSD at any frequency, f, could be obtained by filtering the residuals with a low-pass filter with cutoff frequency *f* and calculating the mean-square value of the result. In practice it is calculated by a more

efficient manner (see Blair [B8]). A jump in the PSD represents energy concentrated at a single frequency. A straight line represents white noise with density equal to the slope of the line.

The interpretation of a PSD will be illustrated using Figure D.2. At 840 MHz, the second harmonic of the input signal, the PSD jumps from 0.8 LSB^2 to 3.6 LSB^2, an increase of 2.8 LSB^2, or (taking the square root) 1.7 LSB. This is the rms value of the sine-wave component at 840 MHz. The peak value of 2.4 is obtained by multiplying this by $\sqrt{2}$. This agrees with the second harmonic observed in the modulo time plot.

This data was taken from a waveform recorder that consists of two interleaved ADCs with sampling rates of 1 GSa/s, giving a combined sampling rate of 2 GSa/s. An interleaving error would occur at 581 MHz which is the difference between the applied frequency of 419 MHz and the sampling frequency of 1 GHz. There is a jump of 0.2 LSB^2 at this frequency, which corresponds to an rms error of $\sqrt{0.2} \cong 0.4$LSB. This is masked by other errors in the modulo time plot. The primary sources of interleaving error are differences in gain, offset and delay between the interleaved channels.

The PSD is well approximated by a straight line plus jumps at a few discrete frequencies. The straight-line portion corresponds to white noise. It can be approximated by extrapolating the straight line segment of the PSD between 0 and 400 MHz. This line has amplitude of 0.5 LSB^2 or 0.7 LSB rms at the Nyquist frequency of 1 GHz. The peak-to-peak (95% probability) amplitude is expected to be four times the rms amplitude, or 2.8 LSB. This agrees with the noise level observed in the modulo time plot.

Figure D.3 shows the modulo time and PSD plots for another channel of the same recorder. The primary difference between the PSD here and that in Figure D.2 is the much larger interleaving error component. In this case, the interleaving error is large enough to be visible in the modulo time plot. It appears as two distinct error curves, one for each of the interleaved channels.

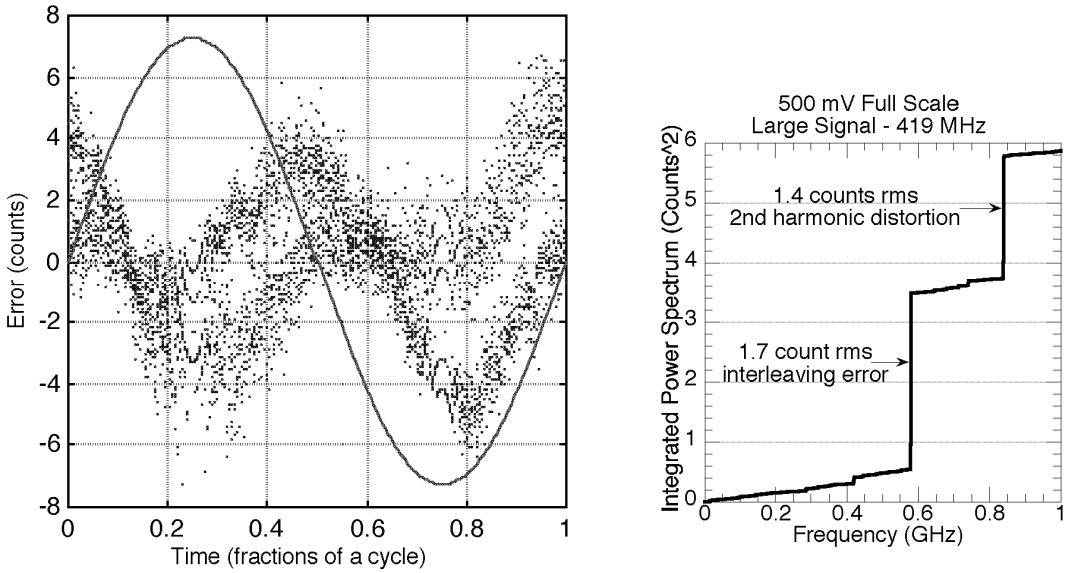

Figure D.3—Same as Figure D.2, but for channel 1 of the recorder, showing the interleaving error

D.3 Other examples of presentations of sine-wave test results

There are other commonly used methods for presenting the residuals, one in the time domain and one in the frequency domain. The time domain method is simply to plot the residuals as a function of time. This allows one to see phenomena that change slowly over the record length. The alternate frequency domain approach is to calculate the DFT of the recorded signal (not the residuals) and plot the amplitudes of the coefficients on a decibel scale. This approach is only meaningful when using coherent sampling, because otherwise the calculated magnitude of a frequency component can be as much as 6 dB below the true value. This approach allows one to easily see the SFDR, which will be the difference between the signal component and the largest spurious component. However, if the spurious component does not have an integer number of cycles in the record it can be shown as much as 6 dB below its true value.

A method of aiding the visualization of a data record, when coherent sampling is used, can be implemented by unwrapping the data record. When an integral number of cycles of the fundamental signal are included in the record, the unwrapping is straightforward. Figure D.4 shows an example of a 1024 point data record of a 10-bit data recorder containing 7 cycles of a fundamental signal. The method used is described by Blair [B9].

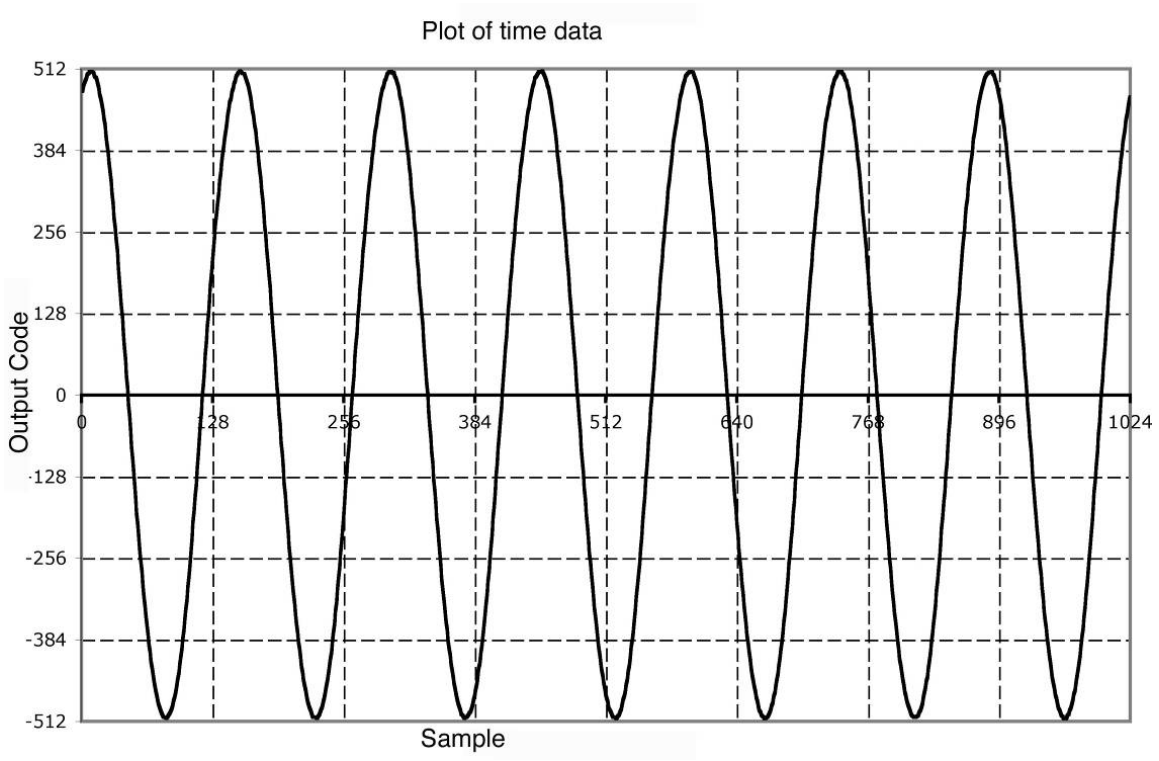

Figure D.4—Plot of 1024 point record containing 7 cycles of the fundamental from a 10-bit data recorder

A DFT of the record identifies that the record has fundamental amplitude of 510.47, at a phase angle of 1.234 radians. Second and third harmonic signals are also present.

The DFT computed magnitude spectrum of the signal is shown plotted in Figure D.5. The horizontal axis is in units of the inverse of the record duration, i.e., if the record duration is 10 μs, then the first frequency on the axis would be 100 kHz, and the last frequency would be half the sample rate, or 51.2 MHz. The average of the record, the dc component, is not included in the graph. The vertical axis is given in units of decibels,

where 0 dB corresponds to the amplitude of the fundamental. The plot is similar to a graph that one would observe on a spectrum analyzer viewing the signal.

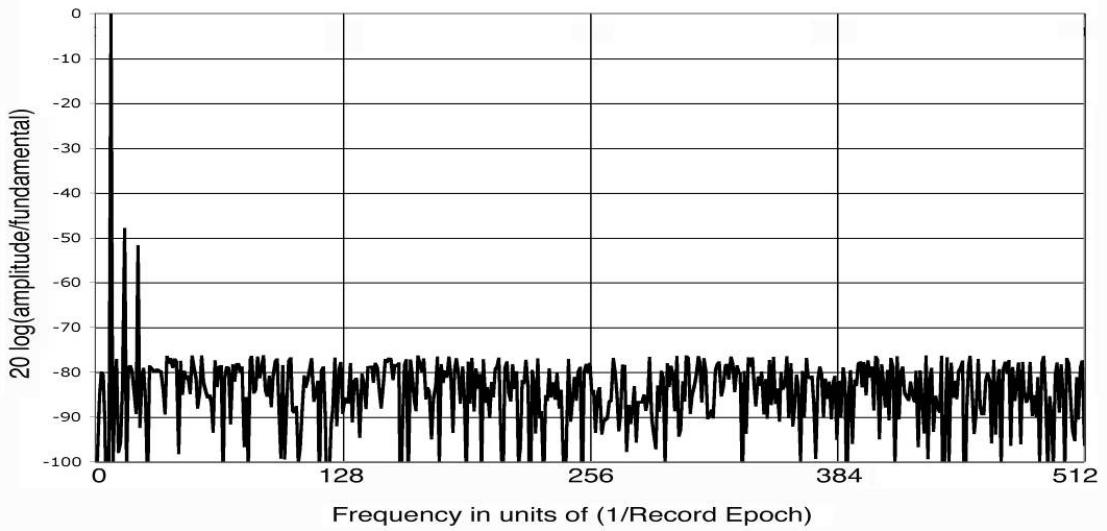

Figure D.5—Plot of the magnitude of the spectrum of the signal shown in Figure D.4

To visualize the nature of the measurement, the calculated fundamental is subtracted, from the data and a residual signal is calculated. A plot of the residual signal is shown in Figure D.6.

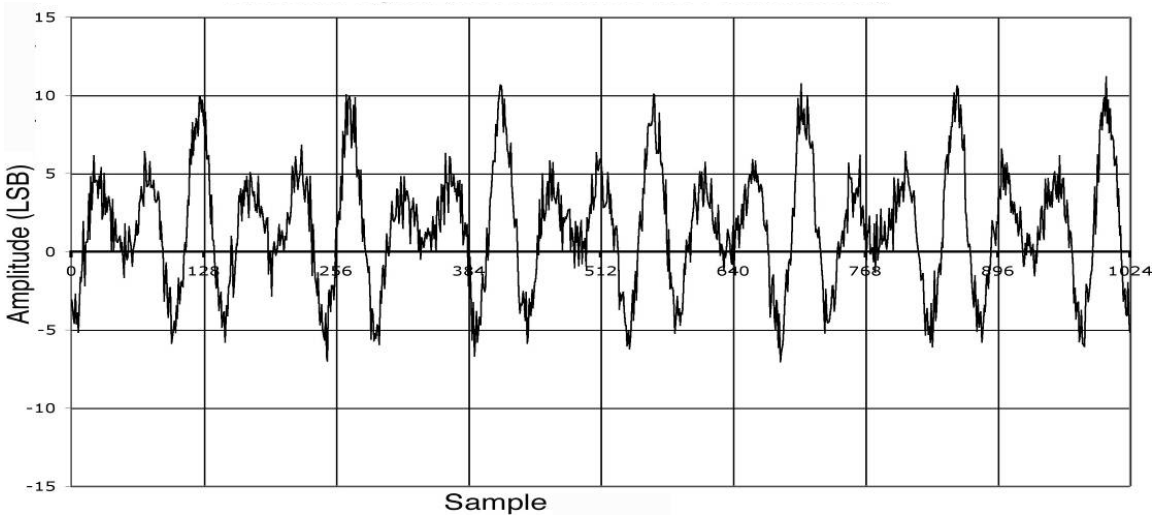

Figure D.6—Plot of the residual signal (record minus the fundamental)

Figure D.6 can be understood better if the data are reordered. This is a feature that can be used when an integral number of cycles of the fundamental are measured. In our example the data contain exactly seven cycles of the fundamental. The parameter *C* is an integer computed from the relationship;

$$C = \frac{MK + 1}{m} \tag{D.1}$$

where

M is the size of the data record being reordered. $M = 1024$ in our example
K is an integer which is determined by trying successive integers until $(MK + 1)$ is a multiple of m
m is the number of cycles of the fundamental present in the record; m is assumed to be relatively prime to M, so that the data are not redundant

The data are then reordered by applying the algorithm described by the pseudo code shown below:

```
m = 7
j = 1
for i = 1 to 1024
     x'[i] = x[j]    (x'[i] contains the unfolded record)
     j = mod(j + C, 1024)
end for
```

In the example the value of C computes to be 439. It results when K is 3. If the data in the example are reordered, the data plot shown in Figure D.7 is generated. The value of C is computed by selecting successive values of K until $MK + 1$ can be factored by m.

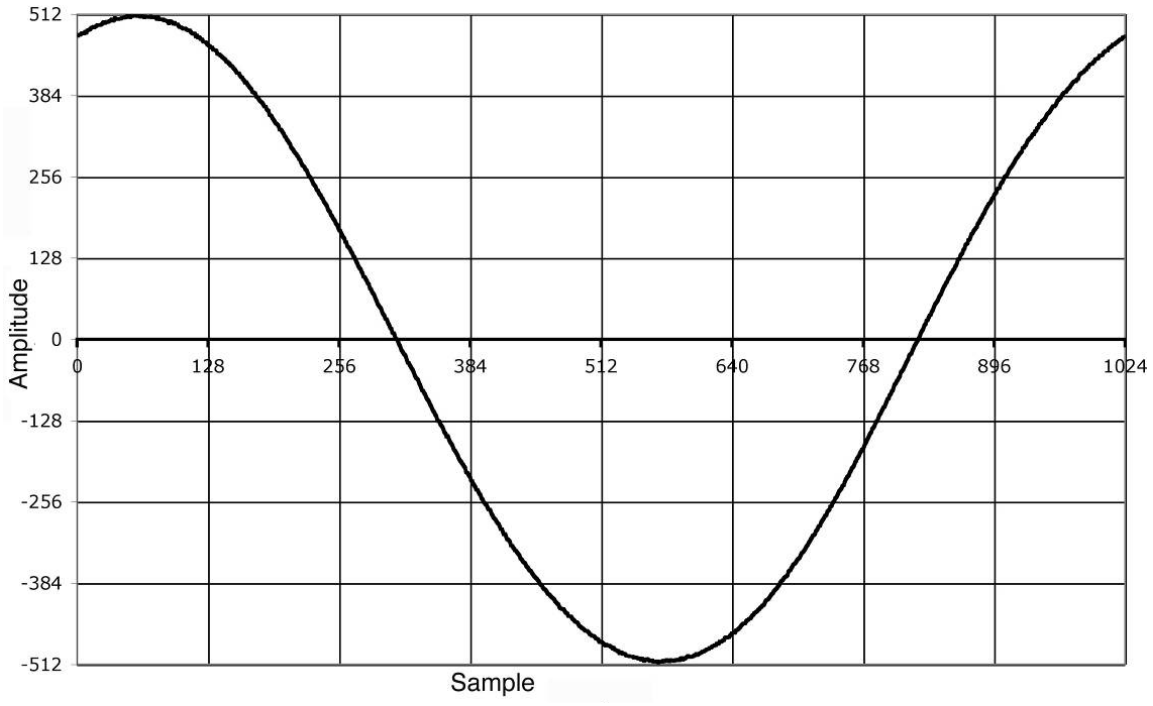

Figure D.7—Plot of unwound data; original record is reordered

The residual can be similarly reordered. The resulting plot is shown in Figure D.8.

It can be noted from Figure D.8 that the distortion components are more easily identified from the unfolded residual.

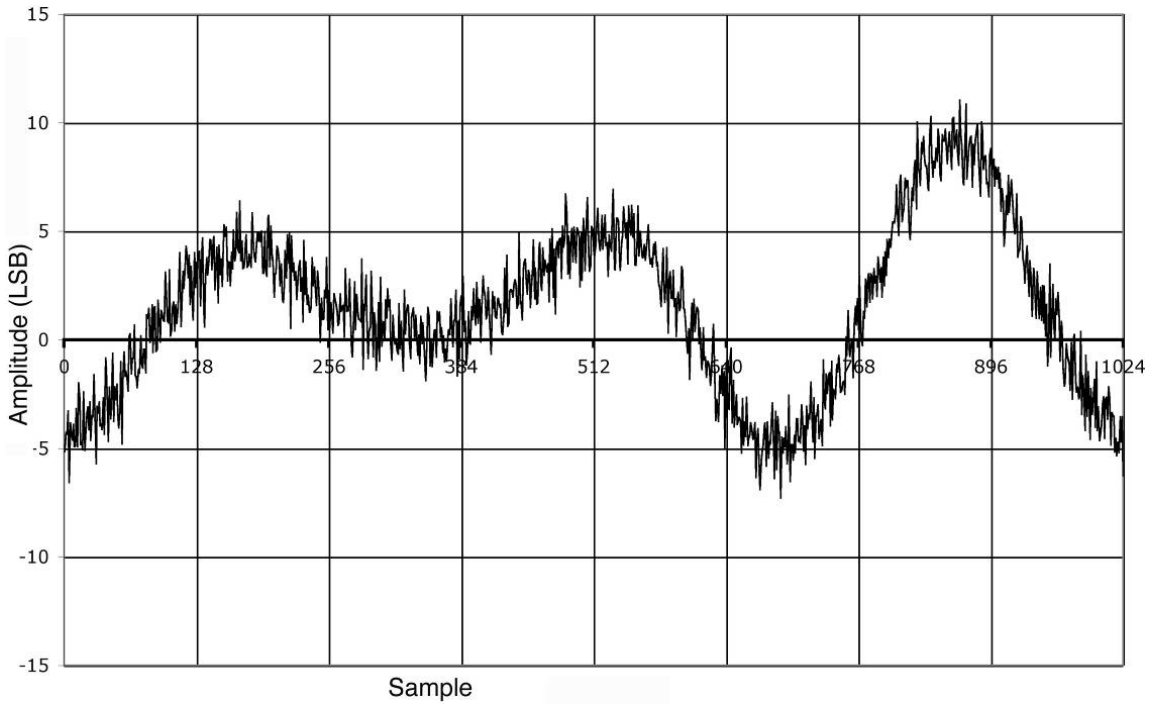

Figure D.8—Plot of unfolded residual

Annex E

(informative)

Bibliography

[B1] Agilent Technologies, "Spectrum Analyzer Measurements and Noise," Application Note 1303, 2003, available from www.agilent.com.

[B2] Alegria, F., Cruz Serra, A., "Uncertainty of ADC random noise estimates obtained with the IEEE 1057 Standard test," *IEEE Transactions on Instrumentation and Measurements,* vol. 54, no. 1, pp. 110–116, Feb. 2005.

[B3] Babu, B. N., and Wollman, H. B., "Testing an ADC linearized with pseudorandom dither," *IEEE Transactions on Instrumentation and Measurement,* vol. 47, no. 4, Aug. 1998.

[B4] Babu, B. N., and Wollman H. B., "Dynamic converter characterization for radar," *Proceedings of the IEEE Second International Conference on Analog to Digital Conversion,* Cambridge, UK, July 1994.

[B5] Benetazzo, A. L., Narduzzi, C., Offelli, C., and Petri, D., "ADC performance analysis by a frequency domain approach," *Conference Record of Instrumentation and Measurement Technology Conference,* New York, NY, May 1992.

[B6] Bilau, T. Z., Megyeri, T., Sarhegyi, A., Markus, J., and Kollar, I., "Four parameter fitting of sinewave testing result: Iteration and convergence," *Computer Standards and Interfaces,* vol. 26, pp. 51–56, 2004.

[B7] Blair, J., "Histogram measurement of ADC nonlinearities using sinewaves," *IEEE Transactions on Instrumentation and Measurement,* vol. 43, no. 3, pp. 373–383, June 1994.

[B8] Blair, J., "A method for characterizing waveform recorder errors using the power spectral distribution," *IEEE Transactions on Instrumentation and Measurement,* vol. IM-41, no. 5, Oct. 1992.

[B9] Blair, J., "Selecting test frequencies for sinewave tests of ADCs," *IEEE Transactions on Instrumentation and Measurement,* vol. 54, no. 1, February 2005.

[B10] Blair, J., "Recorders," in *Wiley Survey of Instrumentation and Measurement,* edited by S. Dyer, pp. 1035–1051, John Wiley and Sons, 2001.

[B11] Blair, J., "Recorders," in *Wiley Encyclopedia of Electrical and Electronics Engineering*, vol. 18, pp.274–290, John Wiley and Sons, 1999.

[B12] Carbone, J., "Testing flash ADCs, theory and practice," Teradyne A500 Family Test Technique Note MS51, Teradyne Inc., Boston, MA, Nov. 1995.

[B13] Carrier, P., "A microprocessor based method for testing transition noise in analog to digital converters," *Proceedings 1983 IEEE International Test Conference,* Philadelphia, PA, Oct. 1983.

[B14] Coleman, B., Meehan, P., Reidy, J., and Weeks, P., "Coherent sampling helps when specifying DSP ADCs," *EDN*, pp. 145–152, Oct. 15, 1987.

[B15] Daboczi, T., "Uncertainty of signal reconstruction in the case of jittery and noisy measurements," *IEEE Transactions on Instrumentation and Measurement,* vol. 47, no. 5, p. 1062, Oct. 1998.

[B16] Gray, G. A., and Zeoli, G. W., "Quantization and saturation due to analog-to-digital conversion," *IEEE Transactions on Aerospace and Electronic Systems,* pp. 222–223, Jan. 1971.

[B17] Gray, R. M., and Stockham, J. G., Jr., "Dithered quantizers," *IEEE Transactions on Information Theory,* vol. 39, pp. 805–812, May 1993.

[B18] Hamming, R. W., *Digital Filters.* 2nd ed. Englewood Cliffs, NJ: Prentice-Hall, 1983.

[B19] Harris, F. J., "On the use of windows for harmonic analysis with the discrete Fourier transform," *Proceedings of the IEEE,* vol. 66, no. 1, pp. 51–83, Jan. 1978.

[B20] Honda, M., *The Impedance Measurement Handbook, A Guide to Measurement Technology and Techniques.* Yokogawa-Hewlett-Packard Ltd, Hachioji, Japan, 1989.

[B21] *IEEE Standards Dictionary: Glossary of Terms & Definitions.*[4] [5]

[B22] IEEE Std 746™-1984, IEEE Standard for Performance Measurements of A/D and D/A Converters for PMC Television Video Circuits.

[B23] IEEE Std 1057™-2007, IEEE Standard for Digitizing Waveform Recorders.

[B24] Irons, F. H., Riley, K. J., and Hummels, D. M., "The noise power ratio—Theory and ADC testing," *Proceedings of the 16th IEEE Instrumentation and Measurement Technology Conference*, vol. 3, pp. 1415–1420, Venice, Italy, 1999.

[B25] Irons, F., and Hummel D., "The modulo time plot, a useful data acquisition diagnostic tool," *IEEE Transactions on Instrumentation and Measurements,* vol. 45, no. 3, pp. 734–738, June 1996.

[B26] ISO 9001:2000, Quality Management Systems Requirements.[6]

[B27] ISO 10012-2:1997, Quality Assurance Requirements for Measuring Equipment—Part 2: Guidelines for Control of Measurement Processes.

[B28] Jenkins, G. M., Watts D. G. *Spectral Analysis and Its Applications,* San Francisco, Holden-Day, 1968.

[B29] Jenq, Y. C., "High-precision sinusoidal frequency estimator based on weighted least square method," *IEEE Transactions on Instrumentation and Measurement,* pp. 124–127, Mar. 1987.

[B30] JESD 99-1:1989, Terms, Definitions, and Letter Symbols for Analog-to-Digital and Digital-to-Analog Converters.[7]

[B31] Kester, W., "High Speed Design Seminar," *Analog Devices,* Norwood, MA, 1990.

[B32] Kinard, J. R., and Ti-Xiong, C., "Determination of ac–dc difference in the 0.1–100 MHz frequency range," *IEEE Transactions on Instrumentation and Measurement,* vol. 380, no. 2, pp. 360–367, Apr. 1989.

[B33] Kuffel, J., McComb, T., and Malewski, R., "Comparative evaluation of computer methods for calculating the best fit sinusoid to the high-purity sinewave," *IEEE Transactions on Instrumentation and Measurement,* vol. IM-36, no. 2, June 1987.

[B34] Laug, O. B., Souders, T. M., and Flach, D. R., "A custom integrated circuit comparator for high-performance sampling applications," IEEE Transactions on Instrumentation and Measurement, vol. 41, no. 6, pp. 850–855, Dec. 1992.

[B35] Lindquist, C. S., Adaptive & Digital Signal Processing, International Series in Signal Processing and Filtering, vol. 2. Miami, FL, Steward & Sons, 1989.

[B36] Linnenbrink, T., "Effective bits: is that all there is?" *IEEE Transactions on Instrumentation and Measurement,* vol. IM-33, no. 3. Sept. 1984.

[B37] Markus, J., "ADC Test Data Evaluation Program for MATLAB."[8]

[B38] Max, S. M., "Fast accurate and complete ADC testing," *Proceedings of the 1989 International Test Conference,* Washington, DC, Paper 5.1, pp. 111–117.

[4] IEEE publications are available from the Institute of Electrical and Electronics Engineers, 445 Hoes Lane, Piscataway, NJ 08854, USA (http://standards.ieee.org/).

[5] The IEEE standards or products referred to in Annex E are trademarks owned by the Institute of Electrical and Electronics Engineers, Incorporated.

[6] ISO publications are available from the ISO Central Secretariat, Case Postale 56, 1 rue de Varembé, CH-1211, Genève 20, Switzerland/Suisse (http://www.iso.ch/). ISO publications are also available in the United States from the Sales Department, American National Standards Institute, 25 West 43rd Street, 4th Floor, New York, NY 10036, USA (http://www.ansi.org/).

[7] JEDEC publications are available from JEDEC, 2001 I Street NW, Washington, DC 20006, USA.

[8] Available for download at http://www.mit.bme.hu/projects/adctest/.

[B39] Max, S. M., "Optimum measurement of ADC code transitions using a feedback loop," *Proceedings of the 16th IEEE Instrumentation and Measurement Technology Conference,* vol. 3, pp. 1415–1420.

[B40] Max, S. M., "Testing high speed high accuracy analog to digital converters embedded in systems on a chip," *Proceedings of the 1999 International Test Conference,* Washington, DC, Paper 29.3, pp. 763–771.

[B41] Nuttall, A. H., "Some windows with very good side lobe behavior," *IEEE Transactions on Acoustics, Speech, and Signal Processing,* vol. ASSP-29, no. 1, pp. 84–91, Feb. 1981.

[B42] Oppenheim, A. V., and Schafer, R. W., *Discrete-Time Signal Processing*. Englewood Cliffs, NJ: Prentice-Hall, 1989.

[B43] Oppenheim, A. V., and Schafer, R. W., *Digital Signal Processing.* Englewood Cliffs, NJ: Prentice-Hall, 1975.

[B44] Oppenheim, A. V., and Willsky, A. S., *Signals and Systems.* Englewood Cliffs, NJ: Prentice-Hall, 1983.

[B45] Papoulis, A., *Probability and Statistics.* Englewood Cliffs, NJ: Prentice-Hall, 1990.

[B46] Rapuano, S., Daponte, P., Balestrieri, E., De Vito, L., Tilden, S. J., Max, S., and Blair, J., "ADC Parameters and Characteristics," *IEEE Instrumentation & Measurement Magazine,* vol. 8, no. 5, pp.44–54, Dec. 2005.

[B47] Rauth, D. A., and Randal, V. T., "Analog-to-digital conversion", IEEE *Instrumentation & Measurement Magazine,* vol. 8, no. 4, pp.44–54, Oct. 2005.

[B48] Roberts, L. G., "PCM Television Bandwidth Reduction Using Pseudorandom Noise." M.S. Thesis, Massachusetts Institute of Technology, Cambridge, MA, Feb. 1961.

[B49] Schoukens, J., Pintelon, H., and Van Hamme, H., "The interpolated fast Fourier transform: A comparative study," *IEEE Transactions on Instrumentation and Measurement,* vol. 41, no. 2, pp 2226–2232, Feb. 1992.

[B50] Souders, T. M., and Flach, D. R., "Accurate frequency response determinations from discrete step response data," *IEEE Transactions on Instrumentation and Measurement,* vol. IM-36, no. 2, June 1987.

[B51] Souders, T. M., Flach, D. R., and Blair, J. J., "Step and frequency response testing of waveform recorders," *Seventh IEEE Instrumentation and Measurement Technology Conference (IMTC-90), Conference Record.,* pp. 214–220, 13–15 Feb. 1990..

[B52] Souders, T. M., Flach, D. R., and Blair, J. J., "Step and frequency response testing of waveform recorders," A Guide to Waveform Recorder Testing, prepared by the Waveform Measurements and Analysis Committee of the IEEE Instrumentation and Measurement Society, pp. 14–20, Apr. 1990.

[B53] Souders, T. M., Flach, D. R., Hagwood, C., and Yang, G. L., "The effects of timing jitter in sampling systems," *IEEE Transactions on Instrumentation and Measurement,* vol. 39, Feb. 1990.

[B54] Stearns, S. D., and David, R. A., *Signal Processing Algorithms.* Englewood Cliffs, NJ: Prentice-Hall, 1988.

[B55] Stearns, S. D., and Hush, D. R., *Digital Signal Analysis,* 2nd ed. Englewood Cliffs, NJ: Prentice-Hall, 1990.

[B56] Stewart, G. W., *An Introduction to Matrix Computations,* New York, Academic Press, 1973.

[B57] Strickland, J. A., "Time-domain reflectometry measurements," Tektronix Inc., Beaverton, OR, 1970.

[B58] Vanden Bossche, M., Schoukens, J., and Renneboog, J., "Dynamic testing and diagnostics of ADCs," *IEEE Transactions on Circuits and Systems,* vol. CAS-33, no. 8, Aug. 1986.