ISSUE FOUR • 2008

Agilent Measurement Journal

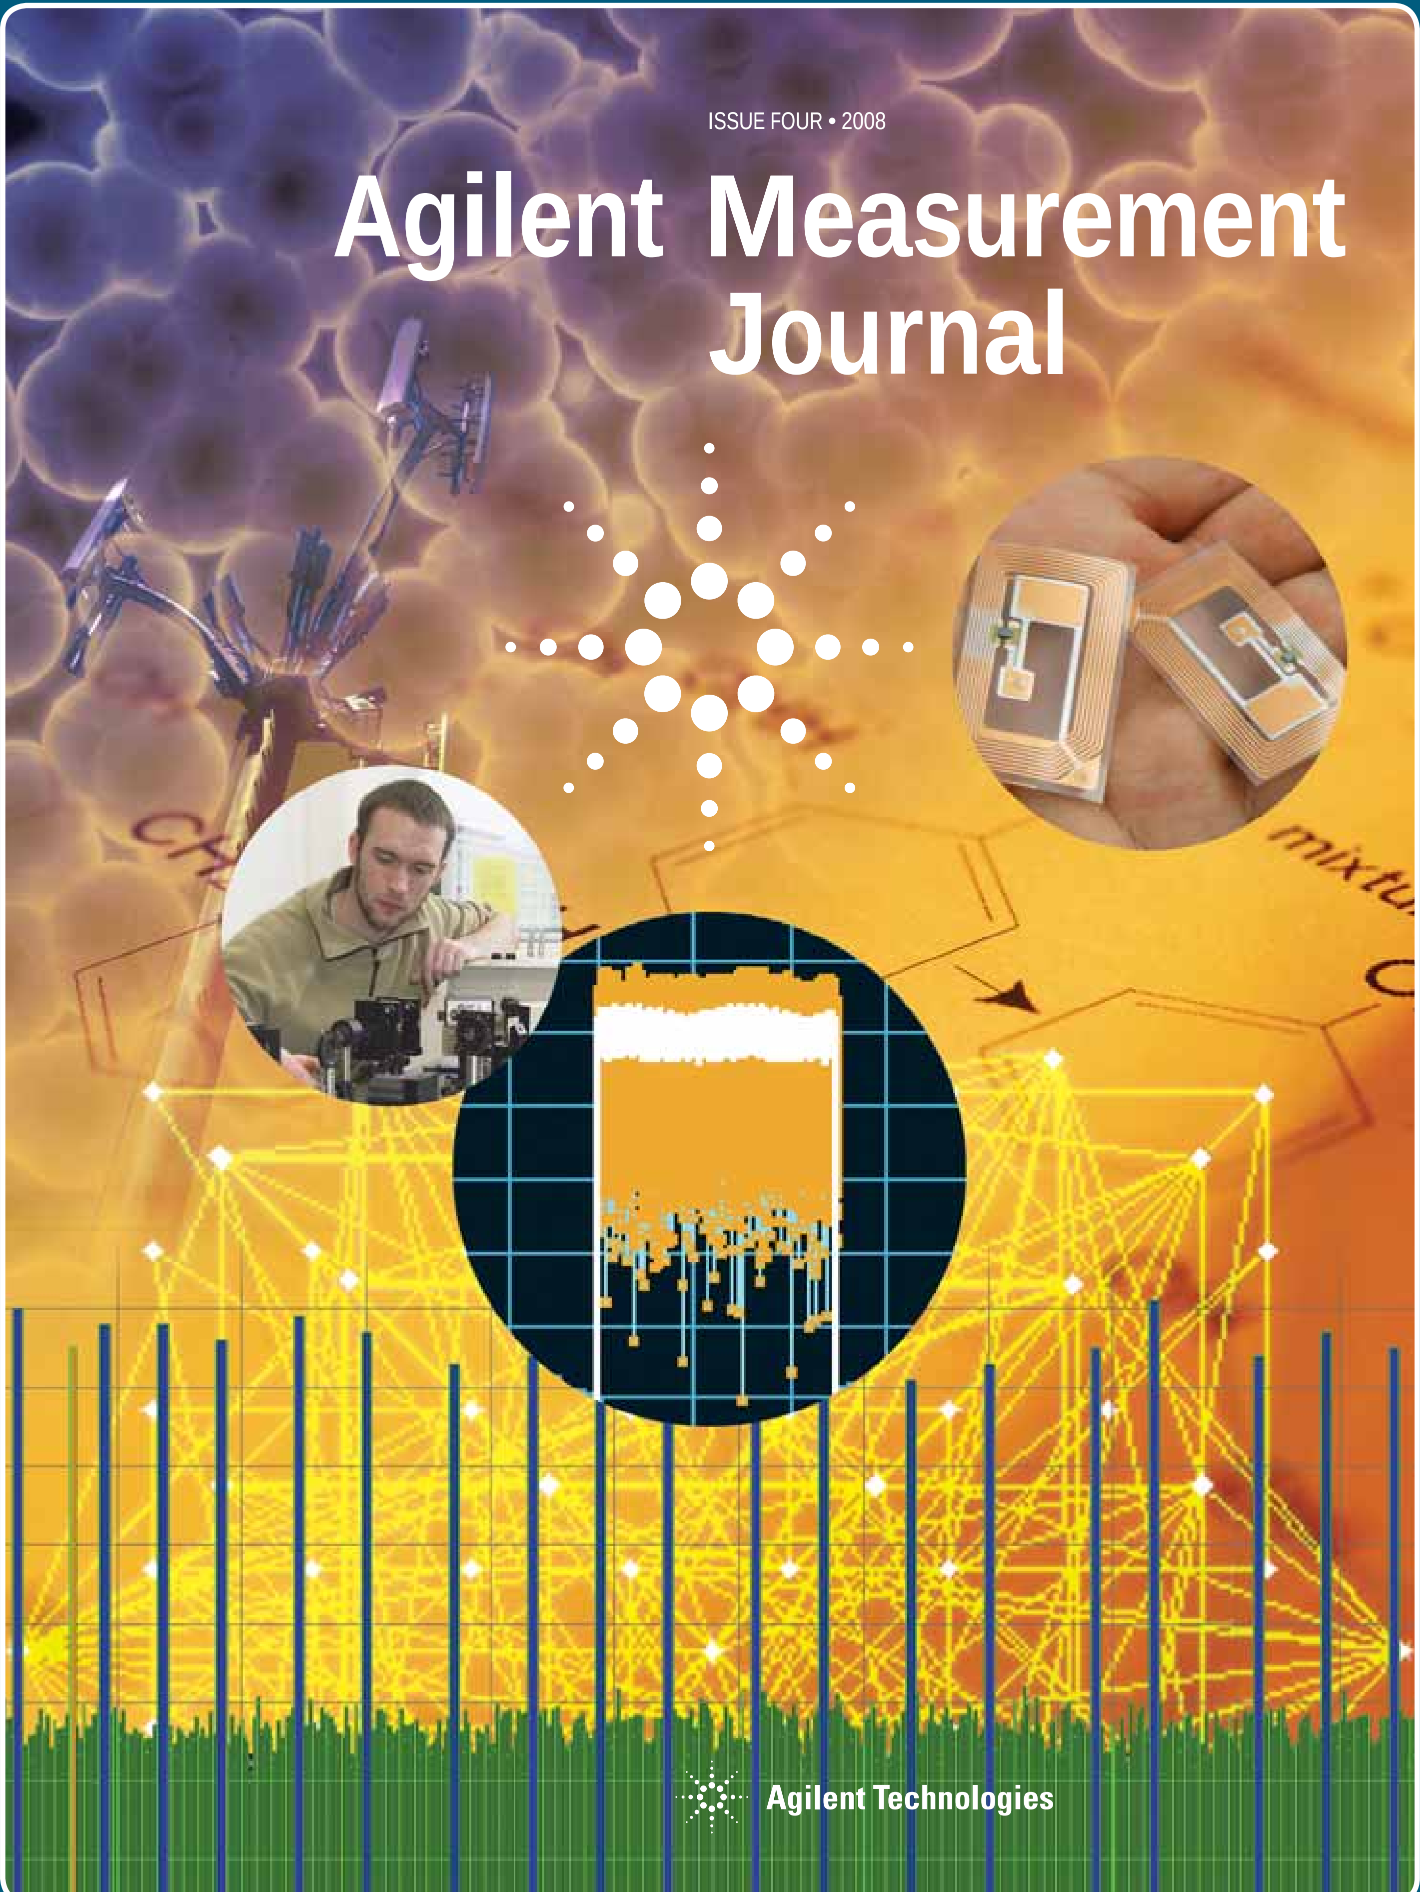

Agilent Technologies



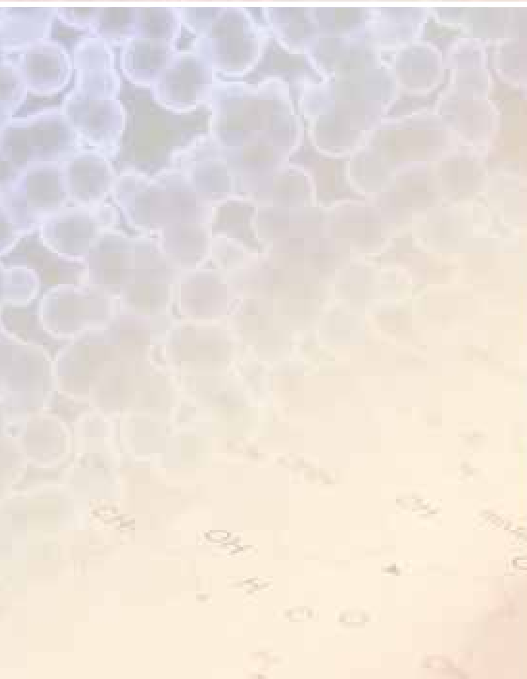



Delivering Confidence through Compliance with Standards

Darlene J.S. Solomon
Chief Technology Officer, Agilent Technologies
Vice President, Agilent Laboratories

Confidence is the overarching benefit that comes from metrology standards, industry standards and regulatory standards. Metrology standards give manufacturers confidence in measurement results and in the quality of their products. Industry standards offer end users confidence that devices from different manufacturers will interoperate. Regulatory standards give a nation's population confidence that its well-being is protected.

If we work backwards, the roots of many standards come from two main sources: manufacturers and their customers. In most cases, those roots take hold when a manufacturer's researchers talk face-to-face with customers. That's when researchers see the body language and hear the tone of voice that underlines the obstacles in the customer's day-to-day work.

After an accumulation of conversations with numerous customers around the world, ideas begin to form. When a promising idea is transformed into a useful solution to a real problem, it often takes the form of a closed, proprietary product. In rare cases, the idea takes on a life of its own, driven by a universal need for a consistent, predictable solution. Some of these ideas become the *de facto* choice in a particular field. The best of these ideas are so beneficial, they are formally ratified as the standard for an entire industry.

Several Agilent innovations have followed this latter path. In the 1970s, customers needed a better way to connect computers to automated test equipment. An idea for a parallel interface became the IEEE-488 standard — better known as HP-IB, the Hewlett-Packard Interface Bus, or GPIB, the General Purpose Interface Bus.

More recently, Agilent led the development of a new measurement interface that is more in tune with current needs for smaller, modular instruments: LAN eXtensions for Instrumentation (LXI). Leveraging the pervasiveness and cost advantages of LAN (Ethernet) technology, LXI will work alongside GPIB while providing additional capabilities that accrue from the use of a distributed architecture enabled by the LAN environment. The benefits include greatly increased visibility into test-system operation and a much richer set of synchronization and triggering capabilities. These new capabilities provide not only greater flexibility for system designers but in many cases more precise system timing. A key component of the LXI capability is another standard that has roots in Agilent Labs, the IEEE 1588 precision time protocol (PTP).

Agilent is a strong believer in the benefits of standards. As you'll see in this issue of *Agilent Measurement Journal*, we're an active participant on many fronts: the adoption and promotion of industry standards, the development of metrology standards, the creation of traceable measurements and more. The end result for our customers is greater confidence in their ability to get good measurement results, deliver a quality product and ensure the satisfaction of their customers.

Visit the *Journal's* online version to listen to Darlene's podcast interview at **www.agilent.com/go/journal**
Interviewer: Frank Elsesser, *eMarketing Manager, Agilent Technologies*

Agilent Measurement Journal

Contents

Insight Department

Visit the *Journal's* online version to listen to Darlene's podcast interview at **www.agilent.com / go / journal**

Interviewer: Frank Elsesser
eMarketing Manager, Agilent Technologies

Emerging Innovations Department

www.agilent.com / go / journal

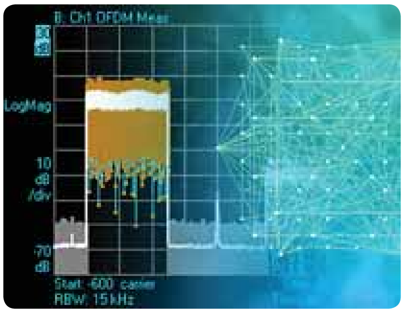

Issue Four | 2008

Campus Connection *Department*

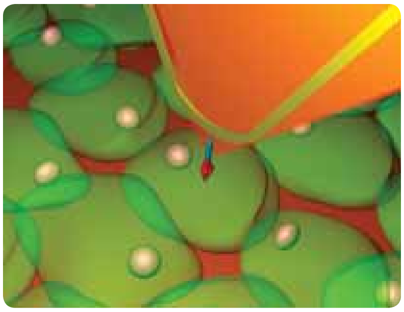

Subscribe to *Agilent Measurement Journal*

Give yourself an edge in today's dynamic world of technology: Subscribe to the *Journal* and receive your own copy of the print edition in the mail.

To activate your free subscription, go to **www.agilent.com / go / journal** and look for the link "Manage your *Agilent Measurement Journal* subscription".

AGILENT MEASUREMENT JOURNAL

William P. Sullivan | President and Chief Executive Officer

Darlene J.S. Solomon | Agilent Chief Technology Officer and Vice President, Agilent Laboratories

Heidi Garcia | Editor-in-Chief

EDITORIAL

Please e-mail inquiries and requests to journal@agilent.com

Agilent **Measurement Journal**

Emerging Innovations

Researchers find key to greater DNA sequencing throughput

Adapting Agilent's *in situ* oligonucleotide synthesis technology to oligonucleotide library synthesis, a team of researchers from Agilent, Harvard Medical School, Virginia Commonwealth University, Stanford University and Codon Devices believes it has unlocked the door to significantly faster and more cost-effective DNA sequencing. The researchers synthesized complex oligonucleotide probe mixtures, applying the mixtures to simultaneously capture and amplify approximately 10,000 human exons in a single reaction. They also demonstrated how this technique can be integrated with ultra-high-throughput sequencing for economical, high-quality targeted variation discovery. The team published the results of its findings in the online version of *Nature Methods*, describing how previous-generation DNA sequencing has been constrained by a lack of parallel "front end" methods to replace traditional polymerase chain reactions.

Agilent introduces first of its kind, single-platform Ethernet OAM test software

New capabilities for the Agilent N2X multiservice test solution provide end-to-end Ethernet operations, administration and maintenance (E-OAM), eliminating one of the final barriers to wide-scale adoption of Carrier Ethernet. This single-platform solution for protocol emulation and conformance testing will enable network equipment manufacturers and service providers to ensure their devices meet performance and scalability requirements, conform to new international standards and are capable of interoperating in large, multivendor networks. The solution also enables comprehensive testing of emerging Ethernet infrastructure technologies as well as the application services, such as Internet protocol television (IPTV), that operate over them.

Agilent takes oscilloscope probing to the extreme

Extreme temperatures, both hot and cold, have always posed a challenge to accurate oscilloscope measurements. Agilent's new N5450A InfiniiMax extension cables for the InfiniiMax Series probing system and Agilent Infiniium oscilloscopes now provide the industry's widest temperature range coverage: –55º C to +150º C. The solution is ideal for cellular, automotive electronics, storage device and consumer electronics equipment designers who must validate their designs' functionality under extreme temperatures. The extension cables also provide testers with an extra probing distance of 92 cm and two different operational temperature ranges based on probe-head configurations.

Streamlining test system creation and deployment in aerospace and defense

Aerospace and defense companies have long been hindered by inefficient test system architectures that require every system connection and protocol conversion to be uniquely programmed. The result is a time-consuming change to every connected system element when only one element requires modification. To solve this problem, the Agilent Virtual Rack platform decouples all system elements within a matrix-based architecture, resulting in an integrated system with components that are independent of specific interfaces or programming languages. System developers only need to define the endpoints — Virtual Rack handles all the required integration through its "storeroom" of thousands of hardware, software and firmware elements from numerous vendors.

Many services, one Agilent next-generation network test solution

Agilent's N2X multiservice test solution claimed the title as the first IPTV service quality test solution for IPv6 (next-generation network) triple-play architectures. The system address a wide range of patterns such as channel zapping and triple-play traffic, providing network equipment manufacturers and service providers with the ability to characterize individual IPv6 network elements or entire networks, ensuring they meet IPTV quality of experience (QoE) expectations prior to deployment. N2X tests the scalability limits of each network device by measuring media delivery index (MDI) delay factor and loss metrics on a per-subscriber basis.

Agilent's LXI-based oscilloscope earns praise

French instrumentation and industrial automation magazine *Mesures* recognized Agilent's LAN eXtensions for Instrumentation (LXI) standard 6000L Series low-profile oscilloscope as a leading, innovative product for test systems in its annual top-products list. The 6000L Series digital storage oscilloscopes (DSOs) provide four measurement channels in a 1U-high form factor. The equipment simplifies test setup and execution time by allowing engineers to control the instrument through a LAN connection and a Web-based interface.

Agilent and BioNanomatrix collaborate on new genome imaging technology

Genetic imaging and analysis equipment company BioNanomatrix will leverage Agilent's expertise in measurement instrumentation to develop a technology that will provide nanoscale single-genome molecule identification and analysis. The solution is intended to provide scientists with rapid, comprehensive and cost-effective ultra-high-resolution DNA analyses. BioNanomatrix's unique nanoscale whole-genome imaging and analysis technology, with sensitivity at the level of the single molecule, has the potential to enable a number of important new applications.

Network operators introduced to lightning-quick data capture and warehousing solution

Mobile broadband network operators now have a quicker and easier way to manage the massive amounts of data associated with their transaction detail records (TDRs) with a new data capture, warehousing and reporting solution from Agilent. Operating on the Teradata Corporation data warehouse platform, the assureME intelligence solution is capable of processing more than 2 billion TDRs daily, far outpacing other monitoring systems that only can handle 200 to 300 million TDRs. Agilent's probe-based network monitoring solution integrates both operational support system (OSS) and business support system (BSS) data and generates intuitive reports based on the mobile network operator's specific needs.

Integrating ICP-MS with enterprise content management

Agilent has integrated the 7500 Series inductively coupled plasma mass spectrometer (ICP-MS) with the Agilent OpenLAB enterprise content management (ECM) system, enabling pharmaceutical laboratory information to be shared and managed on a wider scale. This new solution lets lab managers, bench chemists/operators and IT professionals exchange data across the room or around the world in a secure environment without impeding workflows. To help ensure data security, access to ICP-MS ChemStation/OpenLAB is controlled through a unique username and password, and all captures and changes to data are recorded in a time-stamped, automatically generated audit trail.

Trying Early Device Implementations at the IEEE 1588 PlugFest

A report on the 2007 International IEEE Symposium on Precision Clock Synchronization for Measurement, Control and Communication

Martin Curran-Gray
Research Engineer, Agilent Technologies
martin_curran-gray@agilent.com

INDUSTRY BRIEF

Held October 1-3, 2007, in Vienna, Austria, the 2007 International IEEE Symposium on Precision Clock Synchronization (ISPCS) for Measurement, Control and Communication was the year's main conference event for the IEEE 1588 precision time protocol (PTP).[1] The hosting organization was the Austrian Academy of Sciences Research Unit for Integrated Sensor Systems.[2, 3]

As at previous annual conferences, a "plugfest" was held to allow informal implementation testing of various devices running the protocol. During the three-day conference, the PlugFest occurred during the entire first day and part of the second.

Bugs were found and bugs were fixed, but that's what Plug-Fests are about. Engineers from a wide variety of companies, organizations, agencies and education institutions convene to try out early implementations of their devices in what is usually a friendly, intelligent and helpful environment. Infused with a spirit of mutual cooperation, seemingly everyone was eager to help one another find and resolve problems — and thereby move the standard closer to commercial realization.

Overview: The 1588 PTP

The PTP is commonly known as just "1588," which is the specification number originally assigned by the IEEE during its development. Its tagline is "precision clock synchronization protocol for networked measurement and control systems," a description that conveys its twofold purpose: to enable different pieces of equipment to collaborate and to create a means of providing control or measurement in the operation or characterization of a system.

The protocol describes a set of messages and behaviors that allow devices with inherently different levels of precision to operate concurrently by orienting them — within the defined accuracy and precision — to a common clock. In operation, the PTP provides a means for devices connected by some form of network (usually a LAN) to establish and distribute a shared notion of time. This allows creation of spatially dispersed or "distributed" systems that use a more decentralized operating paradigm than traditional, collocated, controller-centric test systems. While approaches such as Network Time Protocol (NTP) enable millisecond synchronization across a LAN, 1588 was created to achieve synchronization of a microsecond or better.

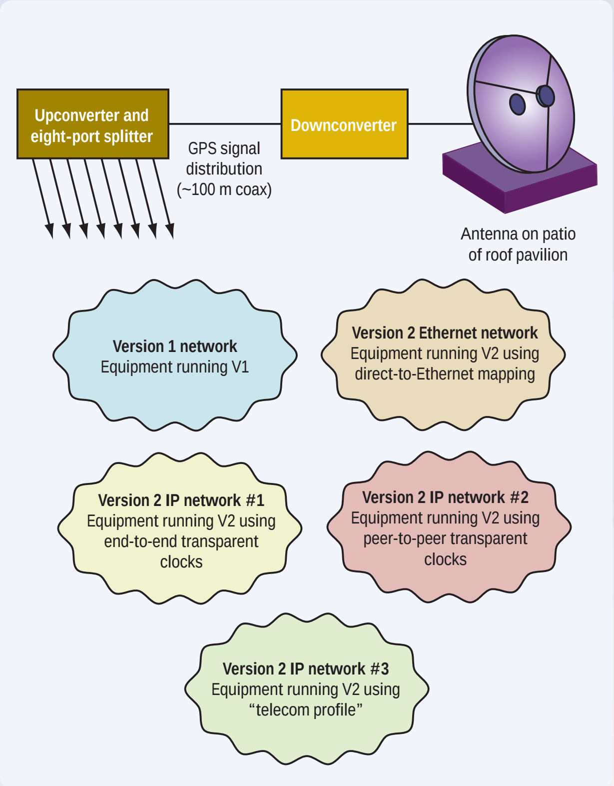

Figure 1. A conceptual diagram of the PlugFest test network

The current international standard, IEC 61588, is based on IEEE-1588 2002, which is generally called “version 1” or “V1” of the PTP. V1 was published in 2002 and version 2 (V2) is due to be completed by early 2008. Both versions were represented at the PlugFest by various types of PTP devices:

- Ordinary clocks, which are end devices providing or using the “time”
- Boundary and transparent clocks, which are specialized switches that allow higher performance than normal networking gear
- A management node

The test networks

Because interest in 1588 has expanded beyond the aims of the initial version, V2 operation is much more varied. As a result, a total of five separate Ethernet networks were set up to support the expected amount of testing (Figure 1). Ethernet was the dominant LAN technology at the conference; however, V2 will describe ways to carry PTP on other network types.

Testing as much as possible

Compared with 2006, PlugFest participation doubled in terms of attendees and the number of organizations they represented. Nearly three times as many devices were brought to this year’s event, representing industries such as industrial automation, test and measurement, telecom components, networking systems and components, ASIC/FPGA cores and microprocessors.

Participants received a draft test plan in advance, and the majority seemed willing to follow the plan. While there were similarities among the test plans for the five networks, some differences were necessary to account for the diverse requirements of applications that used the wider range of options allowed in V2 compared to V1.

Tests included basic synchronization and resolution of the “Best Master Clock” algorithm, which determines which clock will act as the network’s time source — the “Grandmaster” in PTP jargon. For the PlugFest, Grandmaster properties were changed to force the algorithm to re-run, pick a new master and verify that the “slave” devices moved over to the new master. In some networks, message generation frequency was altered, enabling participants to observe changes in the resultant synchronization performance between the master and its slave devices.

There were plans to run precise synchronization measurements; however, with the large number of participants, time ran short so activities focused on functional and interoperability testing.

Highlighting the results

The V1 network proceeded through most of the test plan fairly smoothly, although some of the tests were skipped because no one had implemented every feature of the specification. The end-to-end V2 IP network also completed its tests, as did the Layer 2 direct-Ethernet-mapping network. The peer-to-peer V2 network successfully completed a portion of its suggested plan.

The telecom profile group had the hardest job. It was slower than the other groups, but this was due to greater complexity, variations in the implementations and the newness of the profile’s definition. Ultimately, though, the special negotiation between master and slaves was made to work, and correct unicast operation and synchronization was achieved with the much higher message rates that this profile requires.

The V1 and V2 end-to-end IP networks also were connected together in a test that resulted in a hiccup from one V1 ordinary clock but it recovered and everything else kept running. This was significant because there is no specified compatibility between V1 and V2; however, they can coexist if the implementations are designed with this in mind (as apparently occurred in this case).

Conclusion

From testing of system-wide clock synchronization or device and system interoperability to resolution of compatibility issues, the 2007 PTP PlugFest made significant progress. It will be interesting to see how many people and devices participate in ISPCS 2008 in Ann Arbor, Michigan, and how much progress is made in the implementation of V2.

References

1. For more information, please visit www.ispcs.org.

2. The academy’s web site is available at www.oeaw.ac.at/english/home.html.

3. Details are available at www.oeaw.ac.at/fiss/.

INDUSTRY BRIEF

Achieving Greater Confidence in Measurement Accuracy through Consistency in Calibration Services

Jean-Claude Krynicki
European Quality and Metrology Manager, Agilent Technologies
jean-claude_krynicki@agilent.com

One of the best ways to maintain confidence in the measurements provided by an electronic instrument is to follow its recommended calibration interval. Maintaining calibration — and up-to-date certification — may also be required by internal and external quality auditors.

These days, many companies no longer maintain internal calibration labs and instead are subcontracting or outsourcing their calibration work. Unfortunately, there is wide variation in the depth and quality of the calibrations performed by service providers. One way to gauge calibration services is by the quality and metrology standards the provider follows.

Since 1987, the common denominator for quality has been the ISO 9001:2000 standard, which is part of the umbrella ISO 9000 standard and is administered by the International Standards Organization (ISO). ISO 9001 applies to products, services and software, and has been extended with several industry-specific standards.

To use an analogy, ISO 9001:2000 certification can be viewed as a "passport" required by all businesses. This passport can be endorsed with "technical visas" for specific topics in different countries or regions. For example, ISO 17025:2005 covers the "competence of test and calibration laboratories" globally while ANSI/NCSL Z540.3 is used for Department of Defense (DoD) activities in North America.

The role of standards and guides

The general objective of quality and metrology standards is to provide confidence in measurements and to facilitate product acceptance; they also play a role in reducing technical barriers to trade. As applied to instrument service operations, the end user can expect to see standards requirements reflected within service deliverables, helping them satisfy the needs of internal and external auditors. This is especially important when all or part of the metrology process is subcontracted to a service provider: Essential documents such as calibration certificates, quality system certification, accreditation information or traceability must fully address all relevant audit requirements.

Taking a broader view, international acceptance of service deliverables will provide access to the best worldwide metrology resources and facilitate the movement of equipment from country to country. For multinational companies such as Agilent, this type of standardization allows consistent delivery of services across a global network of service centers.

Within the worldwide deployment of quality and metrology standards, it is expected that requirements are understandable without the support of a specialist, do not conflict with one another, and are not susceptible to multiple interpretations. From a business perspective, the objective is to also see a positive return on investment from benefits such as compliance with requirements through reductions in audits, increases in productivity and gains in competitive advantage. Convergence of standards is expected to also help instrument users such as contract manufacturers who must deal with a wide variety of standards that apply to the countries and industries they supply.

Standardizing calibration certificates

As a specific example, a calibration certificate is the main tangible proof of calibration usually scrutinized by auditors. As a result, service providers must apply special care to the design and ergonomics of the certificate itself. It is also necessary to find a good compromise between worldwide standardization and specific requirements while ensuring consistency between new and in-service instruments.

ISO 17025:2005 is one of a few standards that provide a detailed list of requirements for calibration certificates. Because Agilent participated in the design of ISO 17025, it was possible for the company to create a standard document in 1999. Currently, up to 30 pieces of information are included in the administrative part of an Agilent calibration certificate. One section presents information useful to the end user (e.g., type of service delivered, equipment reference, as-received status) while another section addresses the needs of auditors (e.g., traceability and quality data).

From the very beginning, Agilent's Calibration Certificate Working Group has been monitoring customer inputs on a weekly basis. Along the way, one issue arose concerning the inclusion of a "next calibration due" statement on the document. The Z540 standard seemed to require this usage while ISO 17025 forbade this kind of information because it must be defined by the user's quality system, not by the laboratory performing the calibration service. To provide appropriate flexibility, Agilent lets the customer choose to include or exclude this sentence.

Over time, an Agilent-internal software application was created to issue the document in a service center or onsite. This was followed by the development of Infoline, an online service that gives our customers access to the archived certificate, which is stored as a PDF file.

More recently, Agilent has revisited the standardization of measurement reports, addressing the need to update those produced for older or discontinued products. In such cases, the original calibration software from the production period is still in use: our service organization can calibrate a total of 3,000 Agilent models. Customer feedback regarding this important document is constantly monitored using the Agilent Customer Satisfaction Survey: To date, over 5,000 customer responses have been recorded and the average rating is 9.1 on a scale of 10.

Assessing acceptance criteria

The objective of calibration is to validate the previous usage of the instrument and to provide proof of fitness for use until the next calibration cycle. This is documented with a statement of conformity on the calibration certificate.

When a conformity statement is required, the presence of uncertainty adds a new element to the decision-making process. This issue came to light with the adoption of Military Standard (MIL-STD) 45662A and has resulted in the definition

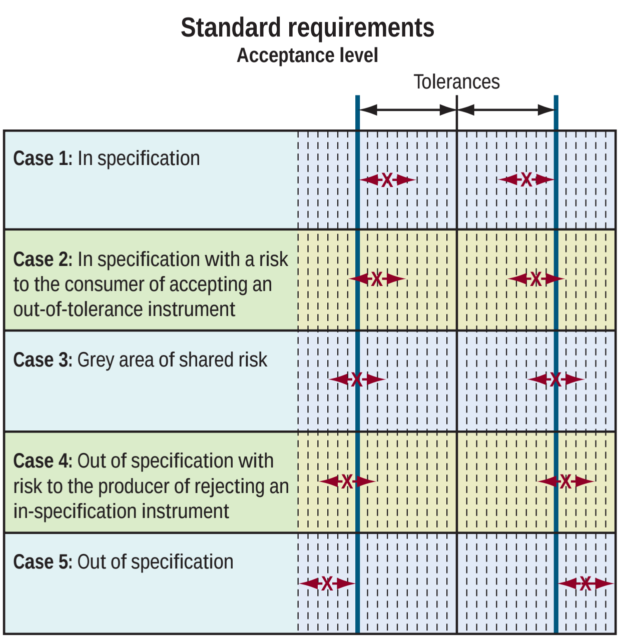

Figure 1. Each of these five possible measurement outcomes has consequences for either the customer or the manufacturer.

of two metrics, the Test Accuracy Ratio (TAR) and its successor, the Test Uncertainty Ratio (TUR). These are also called the guardband management process and producer-consumer risk.

The graph in Figure 1 shows five possible outcomes for a measurement, with the arrows showing the uncertainty around the specifications. Case 1 clearly corresponds to an acceptance situation while Case 5 would be a rejection. Cases 2 and 3 introduce a risk to the consumer (accepting an instrument out of specification) and Case 4 presents a risk for the producer (rejecting an instrument within tolerances).

During the design of ISO 17025, the relationship between acceptance criteria and uncertainty resulted in multiple versions of the draft standard, ranging from the most stringent ("measurement value extended by uncertainty shall fall within the appropriate specification limit") to a more practical approach ("a statement of compliance should be made only if the ratio of the uncertainty of measurement to the specified tolerance is reasonably small, e.g., 1:3") or to the acceptance of an indeterminate status (neither compliance nor noncompliance can be proved).

Ultimately, the standard incorporated a realistic conclusion: The uncertainty must be taken into account. Because ISO 17025 is mainly applied to testing, the expectation was to find the acceptance criteria explained within the technical standard itself.

In the dimensional area, where acceptance criteria determine the acceptance or rejection of a batch of parts without any possible adjustment, the situation is quite critical. This drove the creation of ISO-14253-1, the only standard covering this topic. For calibration, some accreditation bodies have proposed the use of this standard as the default choice if there is no other agreement with a customer.

ANSI/NCSL Z540.3 was the first standard to introduce a formal maximum allowance for consumer risk. However, the stated two-percent maximum risk is for a population of parameters or instruments, not for a single measurement.

Apart from purely textual content analysis and purely theoretical approaches, there is a new trend in quality and metrology in dealing with the tradeoff between risk and cost. Referring back to Figure 1, the expectation behind acceptance criteria is that Cases 2, 3 and 4 could be moved to Case 1 by a simple adjustment. Figure 2 assumes that a calibration alone will satisfy 90 percent of the cases while the other 10 percent will require an adjustment and a second full calibration. The function can be written as

$$P = (1 - X) \times T_c + X(T_a + T_c)$$

where P is the price of calibration, X the percentage of equipment initially found in tolerance, T_c the time to calibrate and T_a the time to adjust. Any instrument function with only one point occurring in the grey zone of Cases 2 or 3 will result in an additional adjustment and calibration process. For complex instruments having hundreds of calibration points, this situation has a high probability of occurring up to the point at which the calibration price will more than double. This is why some service providers now invoice for adjustment fees separately.

Considering technical calibration schedules

Another issue associated with the conformity statement relates to the number of points and functions controlled during the calibration. This helps determine how representative the calibration is compared to the overall instrument specifications. For Agilent instruments, the recommended calibration schedule is included in the product service manual or other similar documents. A "full calibration according to manufacturer specifications" is the rule for our service operations. Among independent calibration providers, however, the trend is toward reducing the price of calibration services — and this is often accomplished by skipping several measurement points and functions. This may be due to technical infeasibility (e.g., the lack of essential reference equipment) or insufficient accreditation to cover all test parameters.

In theory, the end user should know exactly which measurements are performed with the equipment and the technical calibration schedule should cover those traceability points. In practice, instruments do not stay in a dedicated configuration. Instead, the trend within many industries is toward placing equipment in a shared pool, which can translate into limited knowledge of previous usage. As a result, the manufacturer's specifications become the reference.

When talking with our customers, we sometimes face questions such as, "How does an accredited calibration that covers 30 percent of the instrument specification compare to a commercial calibration covering 100 percent of the parameters?" To address such questions, standards bodies have tried to define ethical limits such as the minimum number of points to declare a unit in specification per family of instrument. Much work remains, however, because this approach must account for over 60,000 instruments on the market classified in 200 families. This process was expected to publish its first technical recommendation — AFNOR FD X 07-25 — in late 2007.

Conclusion

The panorama of region- and industry-specific quality and metrology standards is showing signs of convergence, making it possible to address the global market with consistent, transportable certification, accreditation and services. In the quality domain, it is possible for multinational companies to achieve global ISO 9000 certification, allowing better alignment of internal processes, sharing of best practices, and creation of consistent sets of deliverables for customers worldwide.

Greater consistency should ultimately provide greater — and universal — confidence in the precision and accuracy of calibrated test equipment. This will happen only when all service providers adhere to a clear, consistent set of standards — and can achieve and maintain the relevant quality and metrology accreditations and certifications.

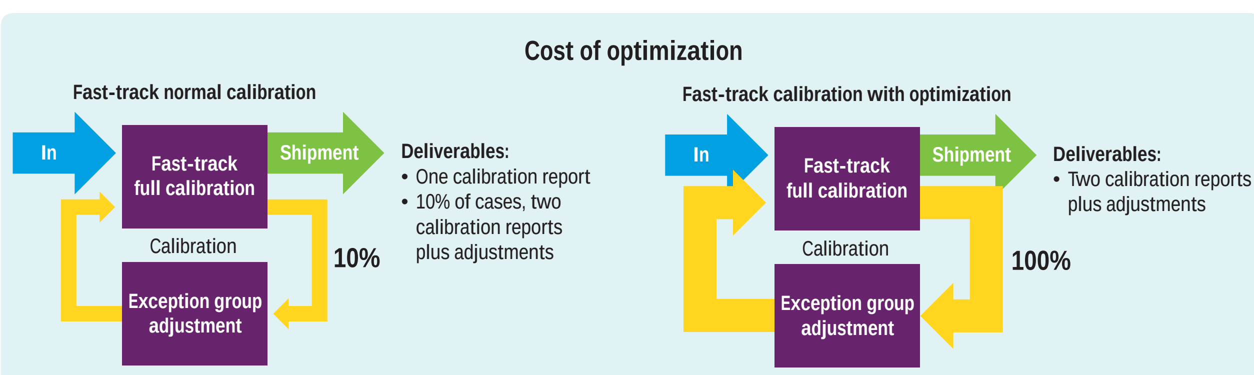

Figure 2. It may be possible to move an instance of Case 2 or 3 into Case 1 with a calibration alone (left) or an adjustment and an additional full calibration (right).

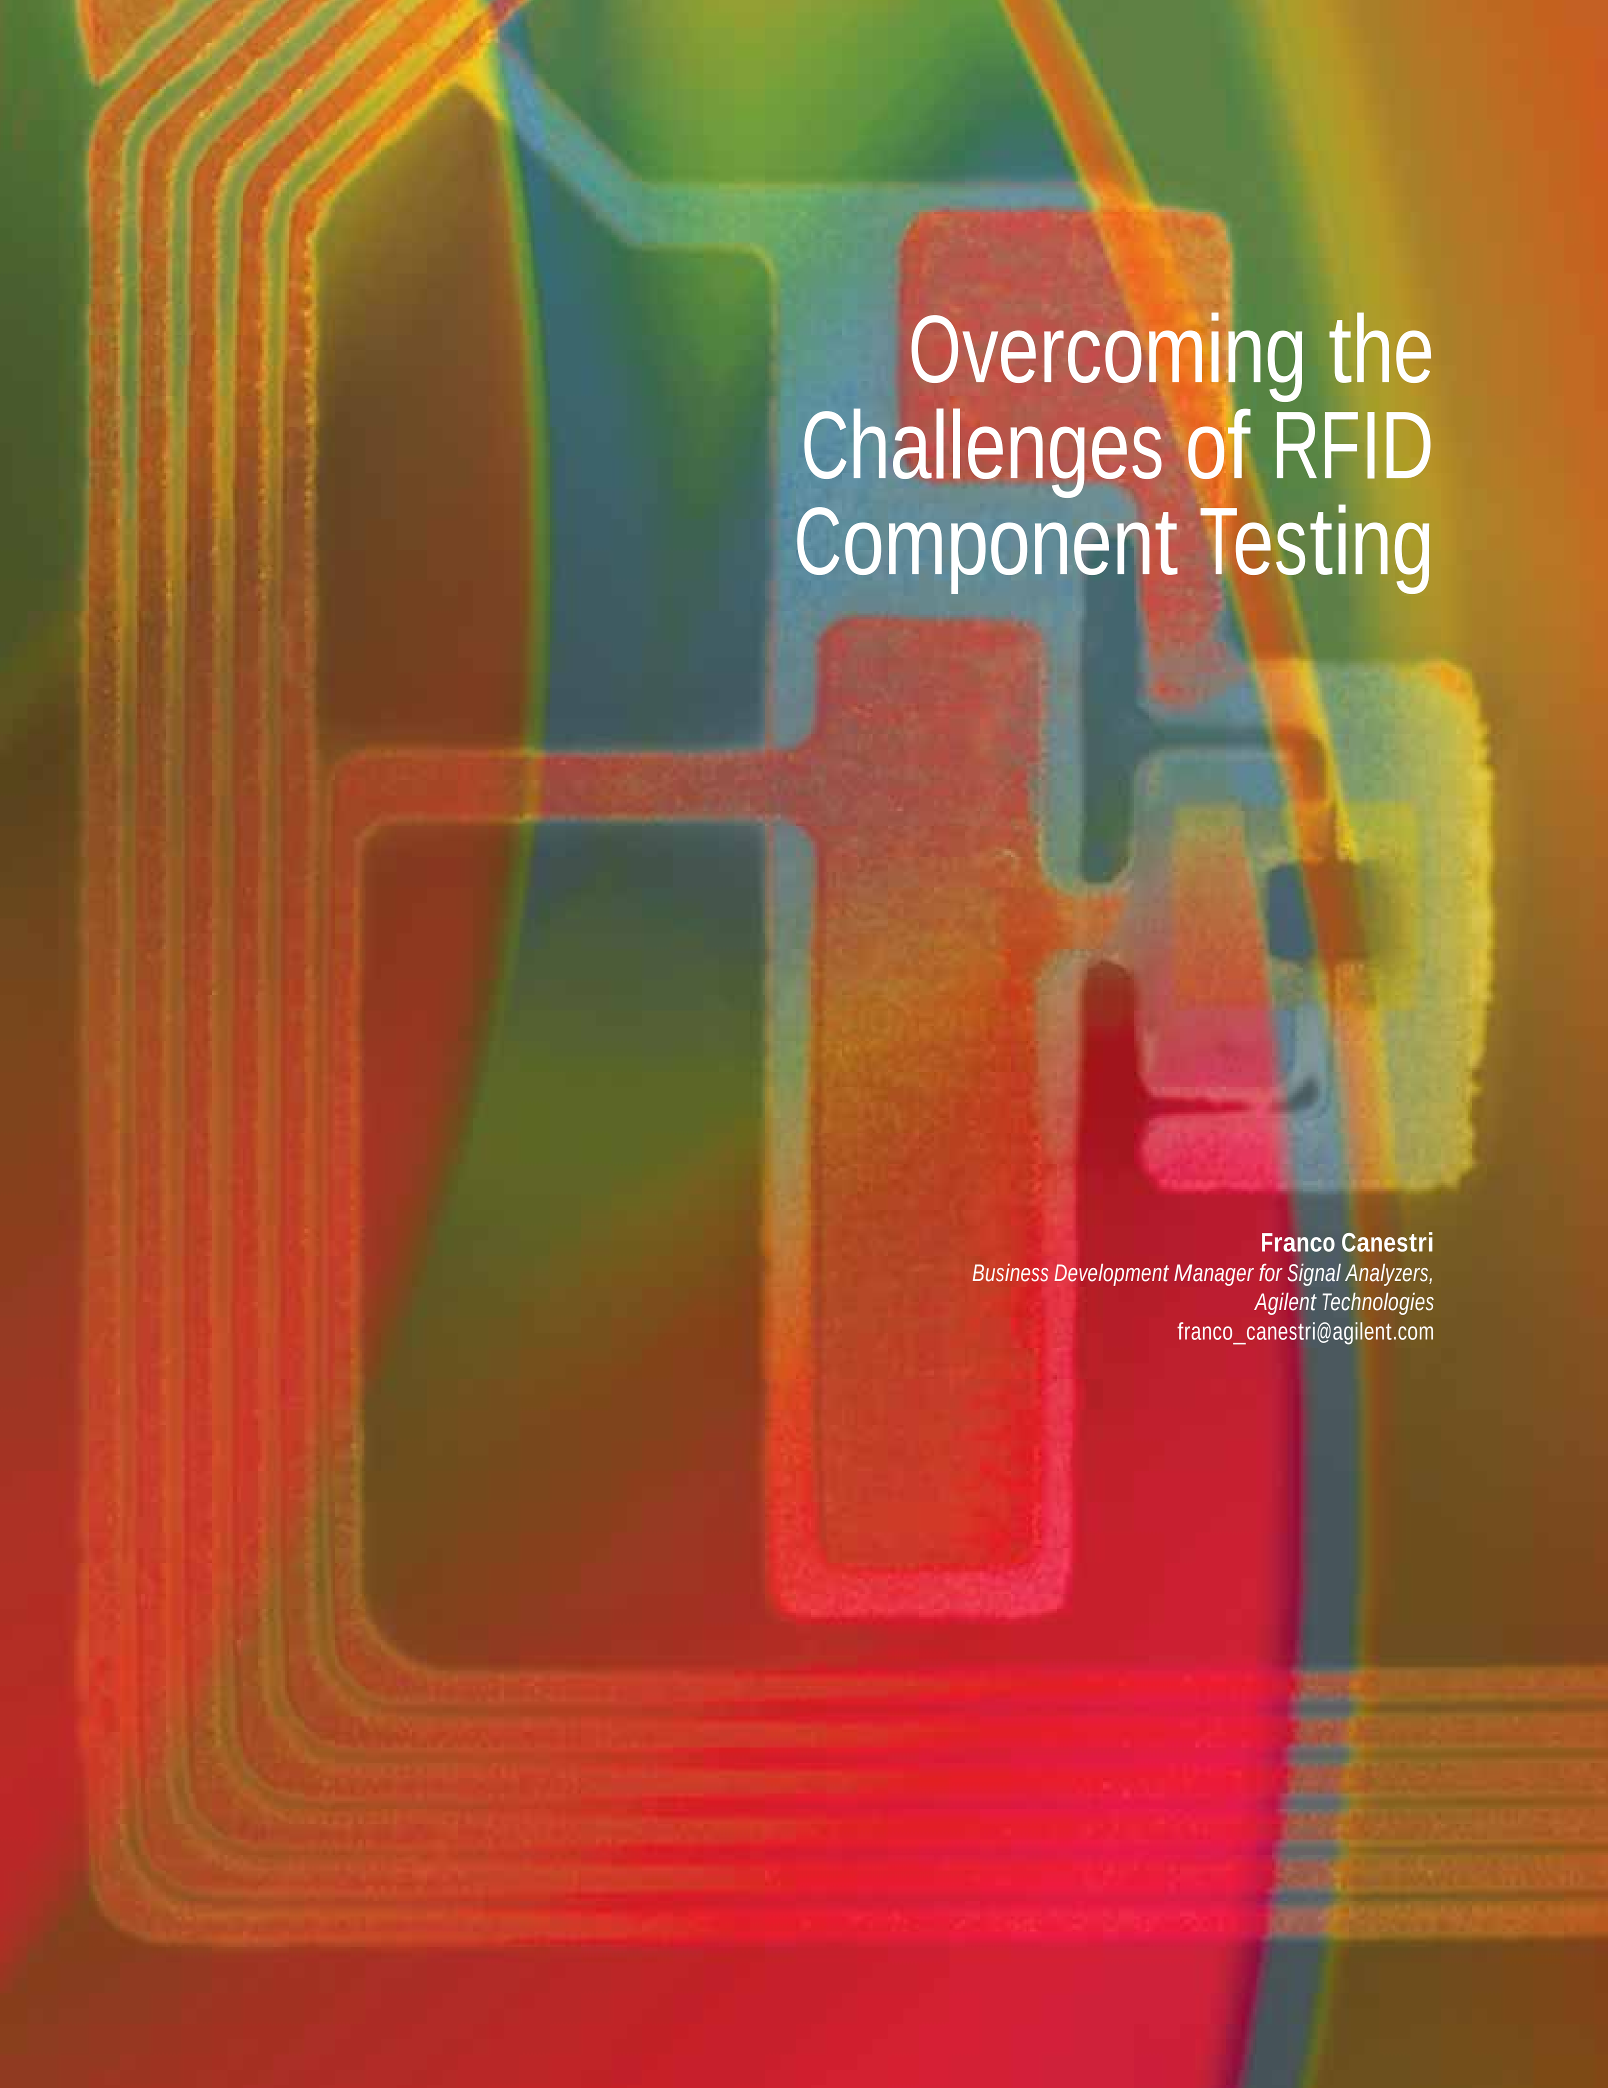

Overcoming the Challenges of RFID Component Testing

Franco Canestri
Business Development Manager for Signal Analyzers,
Agilent Technologies
franco_canestri@agilent.com

Radio frequency identification (RFID) technology is gaining favor as a cost-effective solution for automatic data capture. One key advantage is the use of wireless radio frequency (RF) signals, which enable contactless identification of objects over a greater distance than is possible with alternatives such as barcode readers. Current applications of RFID range from inventory tracking, industrial automation and access control to electronic toll systems, e-passports, medical applications and animal identification.

In development, manufacturing and installation, every RFID application presents difficult measurement challenges, particularly in the analysis of transmit signals, modulation schemes and data signals. Solving those challenges depends on advanced signal-analysis equipment that can measure and characterize the performance of RFID devices and systems. This article focuses on the technical aspects of RFID data transmission and acquisition, typical test requirements and the desirable attributes of suitable test equipment.

Looking at RFID, past and present

Various types of RFID technology have been in development for decades. Basic forms, such as military "identify friend or foe" (IFF) systems, were developed and used during World War II. RFID development continued through the years and many foundational elements of today's technology were created during the 1970s and '80s. Until recently, however, high costs and low levels of standardization hindered widespread implementation of RFID technology. With today's technology advances, the availability of small, cheap, disposable "tags" is driving greater demand for RFID solutions. As a result, widespread use is now an obtainable goal in a variety of applications including inventory control, secure building access, toll-road payments and animal tracking.

Product tracking and inventory control: Passive RFID tags are used to track the movement of trucks, pallets or individual items. Compared to barcodes, RFID tags work across greater distances because the reader does not have to "see" the tag and scan it with laser light. This capability greatly decreases the time required to store and warehouse inventory. What's more, readers can often detect the responses of several tags at once, further reducing the time required for inventory counting. This also reduces the chance of human error that often occurs during the physical inventorying of individual items.

ID badges and access control: Proximity readers can grant access to users carrying RFID-based badges encoded with a specific data response that matches the level of security required for different locations or situations. These badges are more durable than those that use a magnetic strip or barcode because they require only proximity to the card reader rather than physical contact.

Transportation payments: In toll-plaza applications, active RFID tags use their onboard battery power only when they receive a wake-up signal from a reader. With active tags, readings can be obtained at distances of two hundred feet or more while vehicles are traveling at highway speeds.

Animal identification: Many animal owners are implanting their small pets with passive RFID tags. Every tag has a unique identification number that can be accessed if a pet is brought to a veterinary clinic or humane society facility equipped with a compatible reader. After reading the ID number, a database search can provide the owner's contact information.

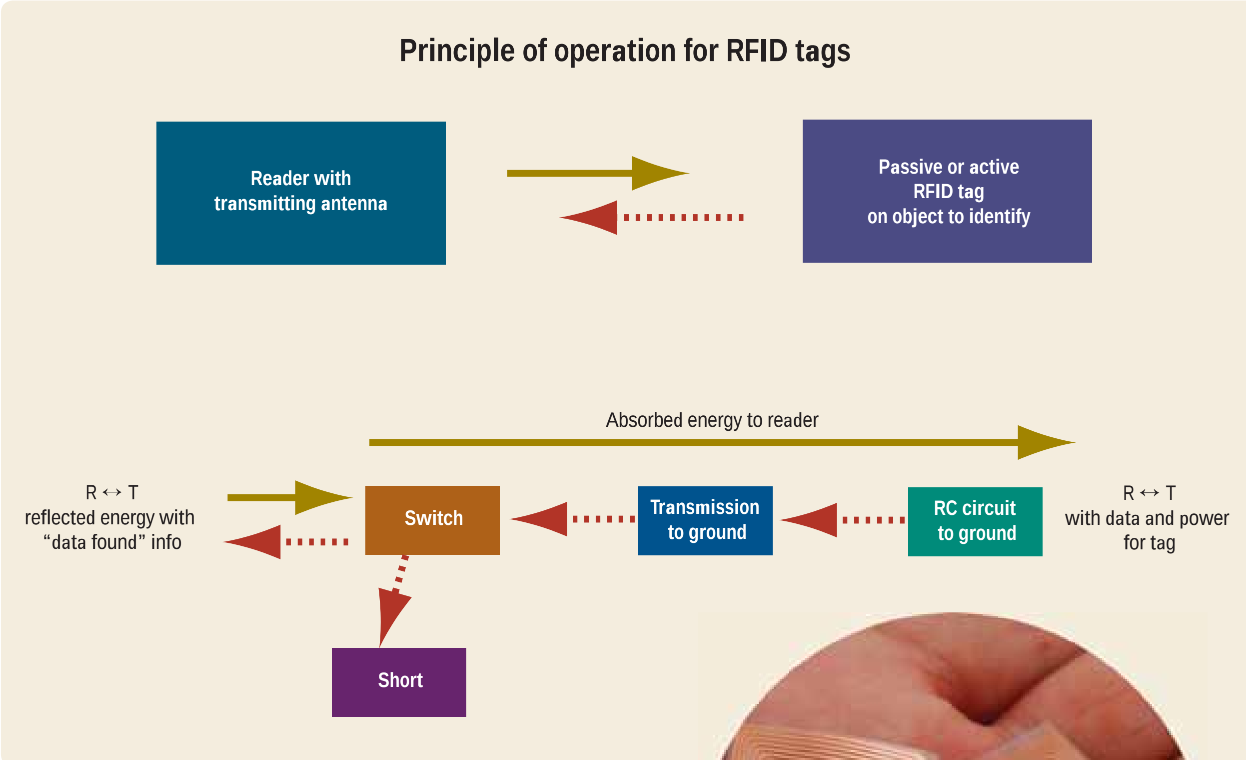

Figure 1. Diagram of RFID operation

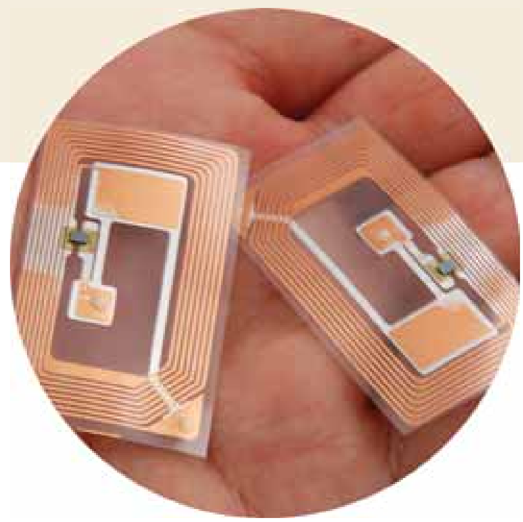

Examining the challenges of RFID testing

The three main elements of an RFID system are an antenna, a transponder (the tag) and a transceiver (the reader). When a passive transponder receives an RF read signal, a small portion of that signal activates the tag. The tag then reacts according to the absorption parameters of its design, reflecting its data to the reader through backscattering (Figure 1).

These systems use simple modulation and coding/decoding algorithms that are often spectrally inefficient. As a result, any given transmission rate needs a wide enough RF bandwidth to allow the transmitted data to be delivered within a serial information stream. This further complicates the coding and decoding processes.

Tag design affects the efficiency of the data transfer. This is due mainly to low-precision timing sources onboard the tag. Another factor is the need for transmit power strong enough to activate the tag under anticollision protocols and allow reading of all tags within range of the reader.

This combination of factors presents many measurement challenges, particularly in the analysis of transient signals, bandwidth-inefficient modulation schemes and backscattered data. The backscattered data, for example, is usually at the same frequency as the absorbed energy from the reader but its amplitude is much smaller and therefore difficult to measure accurately. Performance considerations include tag-reading speed, a tag's ability to operate in a dense reader environment, and the distance between tag and reader. To meet these needs, a variety of instruments are needed to address these challenges, especially in R&D and manufacturing environments for the reader: oscilloscopes, logic analyzers, power meters and advanced spectrum or signal analyzers.

Defining the optimum test solution

A swept spectrum analyzer with fast auto-tuning as well as vector signal analysis (VSA) software with spectrogram displays can easily analyze the power characteristics of complex RFID transmissions. Such an analyzer can recognize the modulation of a transient RFID signal and obtain the requested measurements of frequency, bandwidth and power (typically power versus time). One feature of particular benefit, particularly in the R&D phase, is the ability to make triggered, gap-free recordings of the signals for later analysis.

This combination of capabilities provides the best solution for measurements of spectrally inefficient RFID modulations and their unique decoding needs. Built-in demodulation and decoding capabilities also enable measurements of transient signals by triggering on specific spectral events in a timely manner.

With three-dimensional color spectrogram displays, the user can monitor the evolution of a transient signal in real time and as a trend. The automatic setting of thresholds and markers on a display trace also allows the numerical analysis of rapidly changing signals. For detailed analysis, a multi-trace display configuration with “average,” “maximum hold” and “minimum hold” detectors and markers enables the identification of significant transient changes within specific frequency segments.

Surveying RFID test requirements

Around the world, various government agencies regulate the testing of RFID signals in terms of power, bandwidth and frequency. These regulations are intended to protect users and other devices from harmful interference, ensuring that transmitters do not cross-talk or compete with neighboring RF channels. As a result, numerous standard parameters must be tested during R&D and manufacturing processes. The three most important parameters are reader/tag analysis, time-domain analysis and standards compliance.

Reader / tag analysis

When integrating an RFID communications system, digital, baseband, IF and RF signals are present. The close proximity of the components often leads to crosstalk and the presence of unwanted signals in the output.

By calculating the fast Fourier transform (FFT) of an error-vector magnitude (EVM)-versus-time trace, any deterministic components in the error trace will show up as spectra in the error-vector spectrum. The result, for example, might show a spur below the center frequency. Examining the spur’s absolute frequency and its frequency offset from the carrier will often reveal the interference source.

Time-domain analysis

Analysis of burst and CW signals in the time domain provides additional insights into RFID performance. Any RFID interrogator signal can be recorded in the time domain and displayed in a log-magnitude format. This is an easy way to see the power envelope of the signal. If multiple display markers are available, they can be used to measure power or voltage at an instant in time. The time axis can be configured as relative to the beginning of the acquisition record or at the trigger point.

Standards compliance

Examples include EPCglobal Class1-Gen2 and other ISO 18000 standards for ultra high frequency (UHF) operation, plus others such as ISO 18062 and ISO 14443 that operate in the lower high frequency (HF) range.[1] Measurements typically allow manual setting of demodulation format, line coding and bit-rate control. In all cases, a ten-step process can help ensure successful digital demodulation of RFID signals. This process is based on the capabilities of the Agilent 89600 VSA software.

1. EPC is Electronic Product Code and Class1-Gen2 refers to class 1, generation 2 of the EPCglobal standard.

Step 1: Set center frequency and span

The correct values depend on the signal of interest. For some signals, the process is simple: Use the carrier frequency as the center frequency and select a span wide enough to include any desired out-of-band signals. If the standard allows for frequency-hopped signals (e.g., EPCglobal Class1-Gen2), the measurements are more difficult but a similar approach can be used, starting by setting the center frequency to the midpoint of the channel-hopping bands. The span should be set to include the entire frequency range of all active channels and allow for extra spectral content on the edges.

Measuring hopped signals is comparatively easy if the signal analyzer has a gapless recording capability and the ability to change center frequency and span during post-capture analysis. By changing the analyzer's center frequency and span to the values of the desired hop and selecting the appropriate portion of the recording for playback, any hop can be thoroughly analyzed with no need for prior knowledge of the hop timing or sequence.

Step 2: Set input range

The key is setting the input range as low as possible without allowing overloads. If the input range is set too high, noise will increase and cause greater errors. If the input range is set too low, it will overload the ADC. Selecting the optimal input range will provide the best possible data.

Step 3: Set up triggering (if required)

Rather than the typical magnitude increase, many RFID signals transition from a base magnitude to a lower amplitude. Solutions such as the 89600 VSA software can capture such signals with a "below level" mode that triggers when the amplitude falls below a user-defined value.

Step 4: Select a digital demodulation mode

Many VSAs support various demodulation formats and standards. The 89600 VSA software has an RFID-specific demodulation mode that can be selected, along with the capability to select default settings for many of the current standards.

Step 5: Specify the modulation format

The type of forward (interrogator) modulation scheme depends on the applicable standard. With EPCglobal Class1-Gen2, for example, DSB-ASK, SSB-ASK, PR-ASK, FSK 2 or OOK may be used.[2] Because the interrogator and tag may not have the same modulation scheme, the types of return modulation typically include DSB-ASK, FSK 2 and OOK.

Similar to the forward and reverse modulation schemes, interrogators and tags have different line encodings (dictated by the relevant standard). Line coding helps guard against data corruption caused by noise and interference from other systems.

Step 6: Specify the symbol rate

In EPCglobal RFID demodulation, two important parameters must be set to match the analyzed signal: *Tari* and *bit rate*. (Note that other formats may use only bit rate.) Tari is defined in the standards as the length of a data zero for pulse interval encoding (PIE). In the 89600 VSA software, the markers tool can be used to measure this value. Similarly, the bit rate can be measured by zooming in on a data burst and measuring the time between the start and end points of a high/low/high interval.

Step 7: Select result length and points/symbol

The acquisition length can be set to capture different time widths in the signal. With the "burst search" feature of the 89600 VSA software, the acquisition length and sync search length are the same. Because the software does not require symbol-clock timing, the clock can be manually adjusted to account for signal nonlinearities.

Step 8: Select filter shape and alpha
(measured and reference)

In signal analysis using digital demodulation, adding a filter to the measurement can help reduce intersymbol interference or splatter into adjacent channels. Depending on the signal analyzer or software, several filter types can be applied, each offering tradeoffs between resolution in the time or frequency domain. In many cases, the user can specify the alpha or bandwidth/time (BT) product (for Gaussian filters) to define the filter's shape and width, or create a user-defined filter. Note that filtering is not included in any of the RFID standards but may be a useful design and measurement tool, especially for reader/interrogators, which are essentially fully functional transmitters.

2. DSB-ASK is double-sideband amplitude-shift keying; SSB-ASK is single-sideband ASK; PR-ASK is phase-reversal ASK; FSK 2 is two-level frequency-shift keying; OOK is on-and-off keying.

Step 9: Select the burst to analyze

In the RFID demodulation mode of the 89600 VSA software, it is possible to perform modulation analysis, CW analysis or both simultaneously. The CW mode provides a summary table that includes values such as rise, fall and settling times, as well as other values that are helpful in assessing a signal versus standards requirements. Individual bursts of data can also be examined using the "burst index" feature (Figure 2).

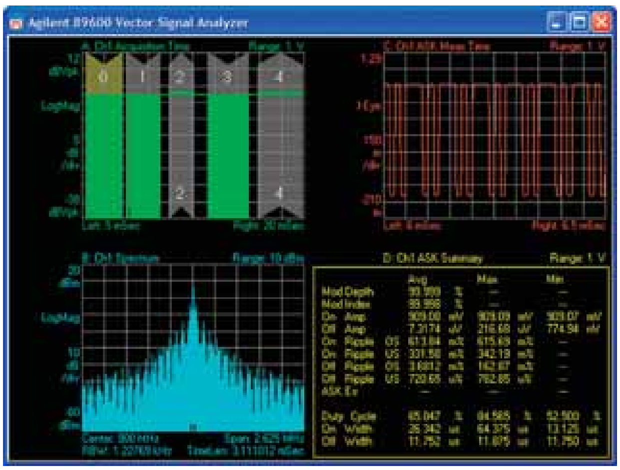

Figure 2. In this analysis of a demodulated RFID signal, bursts are identified with down arrows representing interrogator bursts and up arrows indicating tag return transmissions. Note that tag amplitude levels are well below those of the reader.

It is also possible to automatically identify bursts from either a tag or the reader. What's more, synchronous search parameters can be set to enable selective analysis of all or part of a burst. This type of analysis could be used, for example, to verify preamble settings or data sent N symbols after a specific burst pattern.

Step 10: Check diagram shape and examine error table

Eye diagrams provide an indication of a signal's noise level and, through the width of the eye's opening, also show the amount of distortion. EVM results are automatically calculated and reported in tabular format by the 89600 VSA software. This makes it easy to observe both numerical values and visual indications of signal errors. In addition, verification of proper encoding can be performed by viewing raw demodulated bits and comparing them to the encoded bit stream.

Conclusion

The growing availability of affordable, cost-effective technology is making RFID a more attractive solution for automatic data capture in a wide range of applications. Whatever the application, testing of RFID systems during research, manufacturing and installation presents many challenges. Fortunately, current-generation auto-tuned spectrum analyzers — and the demodulation capabilities of solutions such as the Agilent 89600 VSA software — can meet those challenges.

Additional reading

- Landt, J. 2006. *Shrouds of Time: The history of RFID.* AIM, Inc.

- *Understanding RFID* educational video. 2006. Intermec Education Services.

- Wadham, R. 2006. Radio Frequency Identification. *Library Mosaics*, Vol. 14, No. 5: 22.

- Singh, J., et al. 2006. The State of RFID Applications in Libraries. *Information Technology & Libraries,* Vol. 1: 24–32.

- Hara, Y. 2007. Hitachi advances paper-thin RFID chip. *EETimes.*

- Roberti, M. 2007. A 5-Cent Breakthrough. *RFID Journal.*

Acknowledgement

Portions of this article were previously published in the September 2007 "Emerging Wireless Technology" supplement to *RF Design* magazine.

The author would like to thank Elizabeth Aguilar and Ben Zarlingo of Agilent's Signal Analysis Division and Sarah Ditlevson, Agilent engineering intern, Signal Analysis Division, for their contributions to this version of the article.

3GPP LTE: Introducing Single-Carrier FDMA

Moray Rumney BSc, C. Eng, MIET
Lead Technologist, Agilent Technologies
moray_rumney@agilent.com

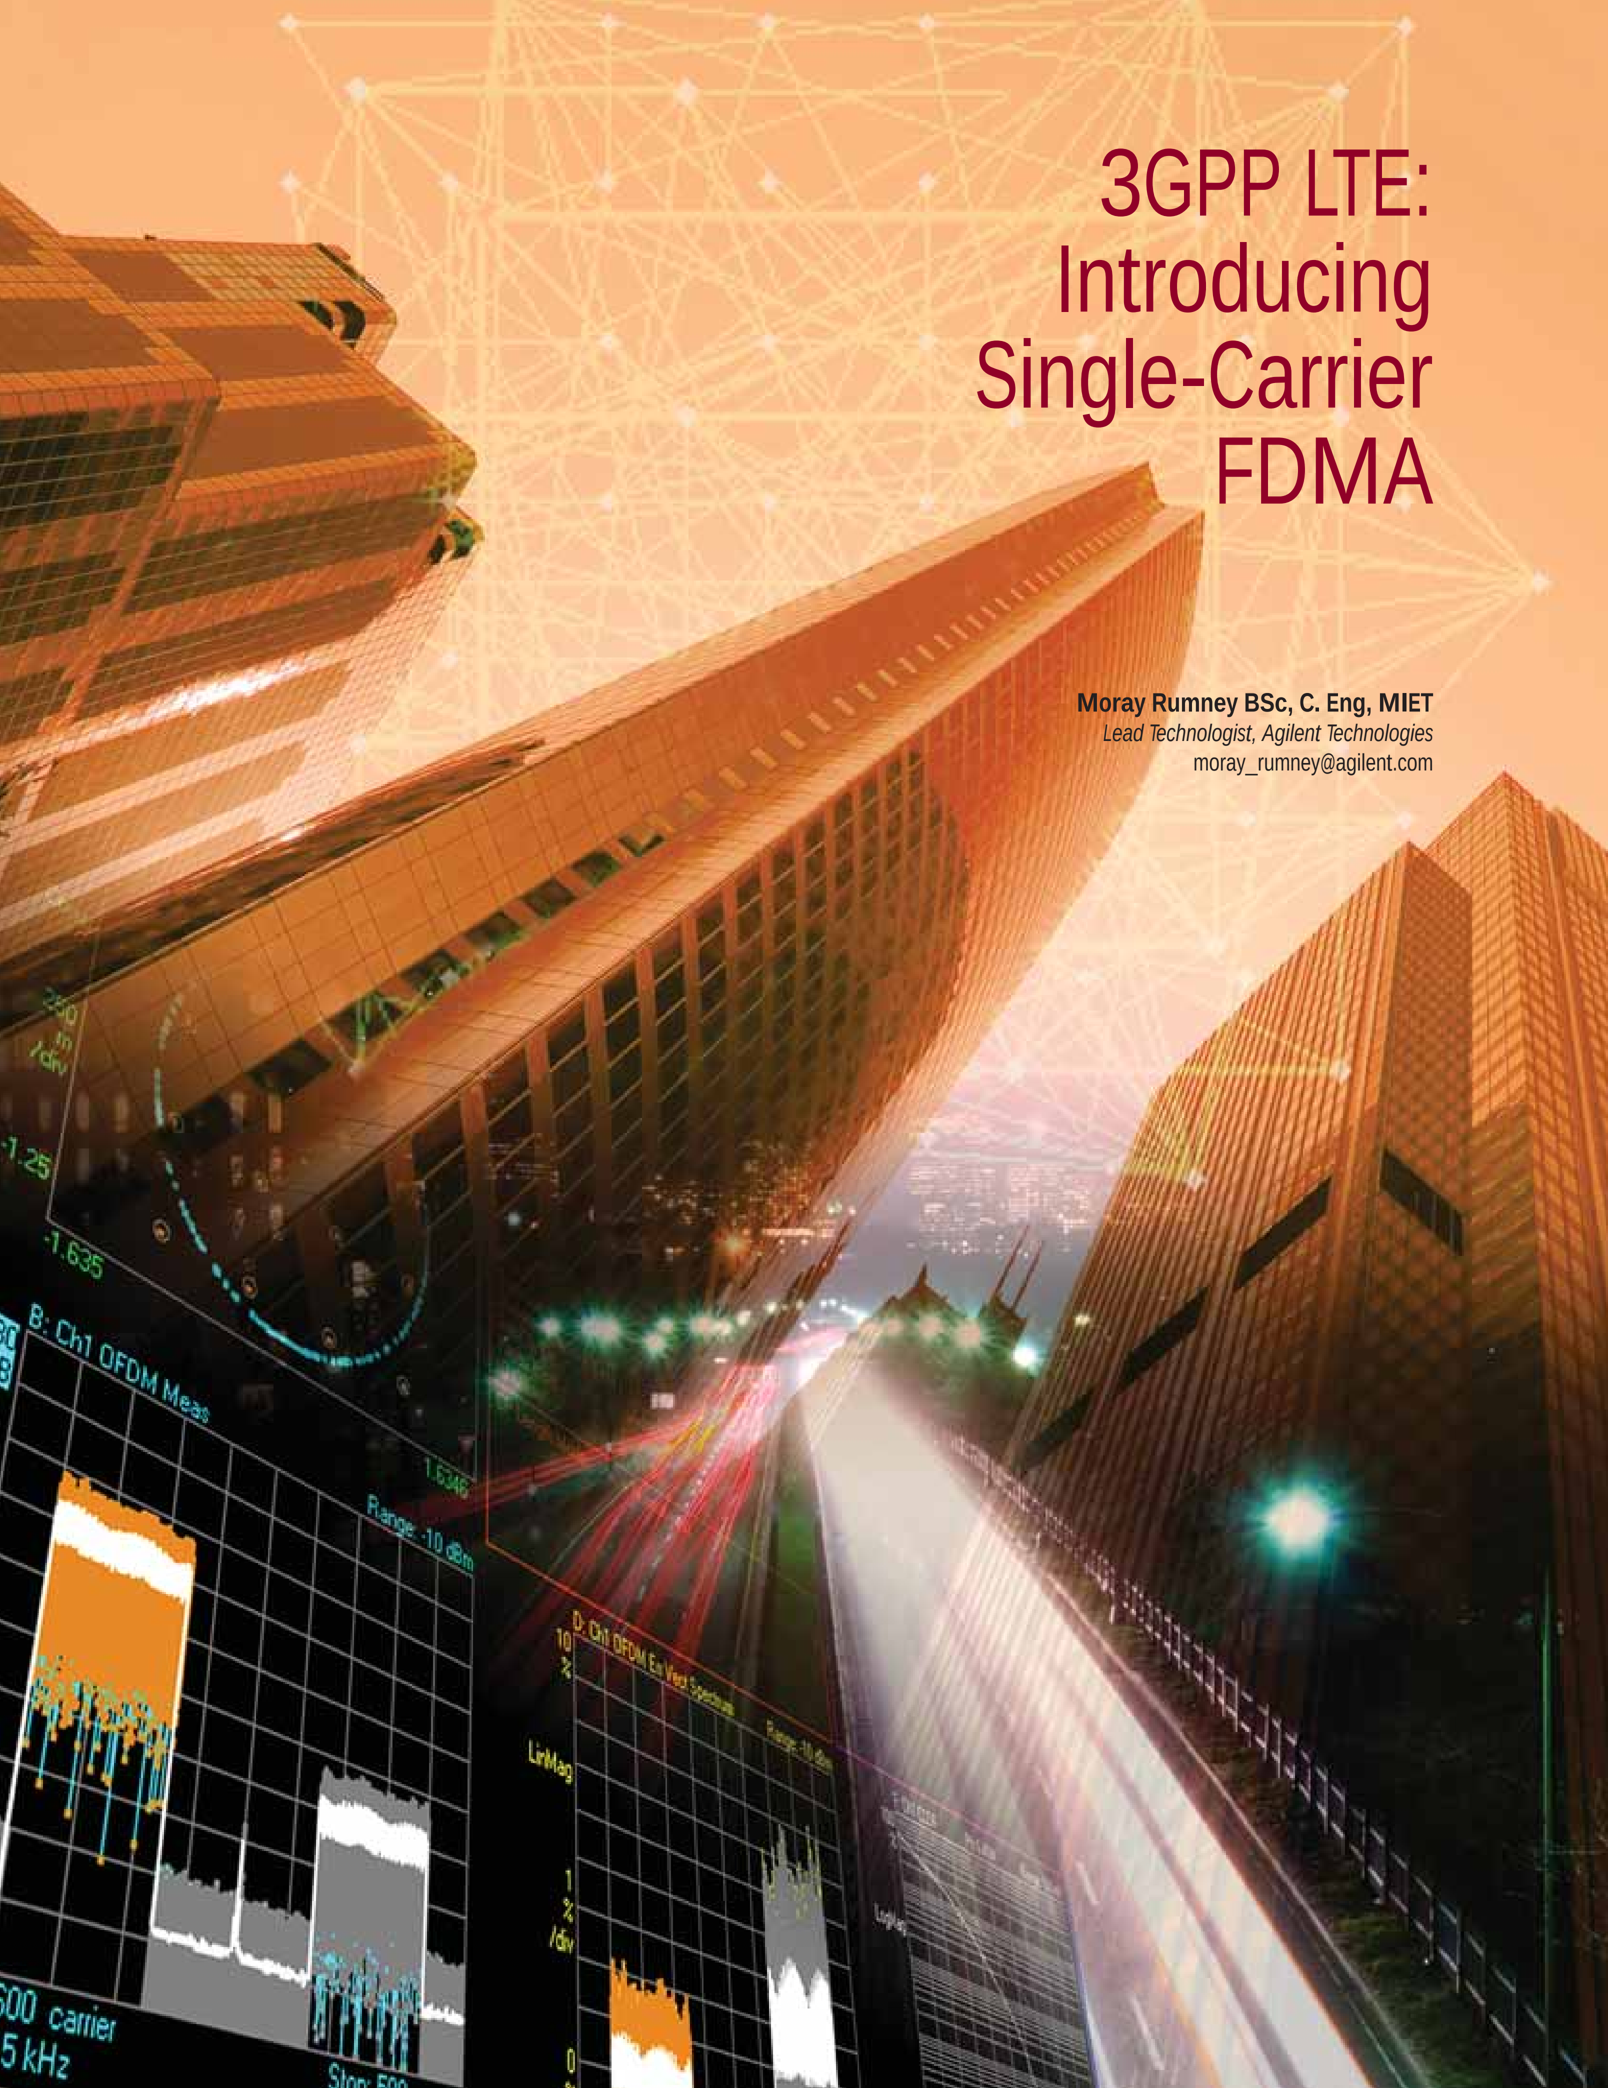

Close on the heels of IEEE's new 802.16e standard—better known as Mobile WiMAX™ — follows the response from the Third-Generation Partnership Project (3GPP) in the form of its Long-Term Evolution (LTE) project. We featured WiMAX™ in Issue Three of *Agilent Measurement Journal* and in this article we explore what LTE aims to bring to the wireless ecosystem. After considering the broader aspects of LTE, we take a closer look at the uplink, which uses a new modulation format called single-carrier frequency-division multiple access (SC-FDMA). These are interesting times because it is rare that the communications industry rolls out a new modulation format.

From both a technical and practical point of view, there is much to understand, examine and evaluate in the capabilities and benefits that SC-FDMA brings to LTE. SC-FDMA is a hybrid modulation scheme that combines the low peak-to-average ratio (PAR) of traditional single-carrier formats such as GSM with the multipath resistance and in-channel frequency scheduling flexibility of orthogonal frequency-division multiplexing (OFDM).

Acronyms galore: LTE history and context

LTE's study phase began in late 2004. The overall goal was to select technology that would keep 3GPP's Universal Mobile Telecommunications System (UMTS) at the forefront of mobile wireless well into the next decade. Key project objectives were set in the following areas: peak data throughput; spectral efficiency; flexible channel bandwidths; latency; device complexity; and overall system cost. The main decision was whether to pursue the objectives by continuing to evolve the existing W-CDMA air interface (which incorporates HSPA*) or adopt a new air interface based on OFDM. At the conclusion of the study phase, 3GPP decided that the project objectives could not be entirely met by evolving HSPA. As a result, the LTE evolved radio access network (RAN) is based on a completely new OFDM air interface.

This does not mean the end of 3GPP's interest in GSM and W-CDMA. Rather, the investment in these technologies — and their remaining untapped potential — means that LTE is not the only format being developed in 3GPP Release 8. For example, the EDGE Evolution project will be pushing GSM to newer levels and the HSPA+ project — the runner-up to OFDM for LTE — will continue to evolve the underlying W-CDMA, HSDPA and HSUPA technologies. For an overview of how these formats inter-relate, please see "What Next for Mobile Telephony?" in Issue Three of *Agilent Measurement Journal*.

By using OFDM, LTE is aligning with similar decisions made by 3GPP2 for Ultra-Mobile Broadband (UMB) and by IEEE 802.16 for WiMAX. For an overview of OFDM technology, please see "Understanding the Use of OFDM in IEEE 802.16 (WiMAX)" in Issue Two of *Agilent Measurement Journal*. Although the article explains the basics of OFDM with reference to WiMAX, the general principles apply to LTE and UMB as well.

Within the formal 3GPP specifications, the LTE evolved RAN is split into two parts: the Evolved UMTS Terrestrial Radio Access (E-UTRA) describing the mobile part; and the Evolved UMTS Terrestrial Radio Access Network (E-UTRAN) for the base station. For simplicity, this article refers to the new air interface by its project name, LTE. This is becoming common usage just as happened with another project name, UMTS, which has been synonymous with W-CDMA since 1999. In addition to developing LTE, 3GPP is also working on a complementary project known as System Architecture Evolution (SAE), which defines the split between LTE and a new Evolved Packet Core (EPC), which is a flatter packet-only core network that aims to deliver the higher throughput, lower cost and lower latency promised by LTE. The EPC is also designed to provide seamless interworking with existing 3GPP and non-3GPP access technologies.

*HSPA (high-speed packet access) refers collectively to high-speed downlink packet access (HSDPA) and high-speed uplink packet access (HSUPA), the latter being formally known as the Enhanced Dedicated Channel (E-DCH).

LTE objectives and timing

The sidebar *LTE at a glance* (page 25) describes the major objectives of the LTE project and some of the key system attributes. Figure 1 shows an overall timeline for the LTE project. Compared to UMTS, the overall timescale is shorter, due largely to a much smoother standardization process. The development of LTE will avoid the 8000-plus change requests ultimately applied over a four-year period to the "frozen" UMTS Release 99 specifications. The instability and subsequent delays in the UMTS standard led to commercial deployment of a proprietary system in Japan before the worldwide standard was available. It is expected that the surprises and delays of UMTS will be averted with LTE, meaning its introduction should be more predictable and better able to avoid a proprietary launch. The dates in Figure 1 are acknowledged as aggressive and may slip; however, progress is solid and, as UMTS proved, trying to rush the process can be counterproductive.

OFDM: The choice of next-generation wireless

With LTE joining UMB and WiMAX in choosing OFDM as the underlying modulation technology, it could be argued that there is now little to choose between these cellular systems. Of the five major new cellular systems, only HSPA+ and EDGE Evolution do not use OFDM, a difference clearly driven by the practical need for backwards compatibility with their respective installed bases.

OFDM has been around since the mid 1960s and is now used in a number of non-cellular wireless systems such as Digital Video Broadcast (DVB), Digital Audio Broadcast (DAB), Asymmetric Digital Subscriber Line (ADSL) and some of the 802.11 family of Wi-Fi standards. In contrast, it has taken longer for OFDM to be adopted into cellular standards: It was briefly evaluated in the late 1980s during the early stages of GSM and again a decade later as a candidate technology for UMTS but was not adopted in either case. The primary issue was the processing power needed to perform the fast Fourier transform (FFT) operations at the heart of OFDM. In the '80s and '90s, suitable processors were too expensive and power-hungry for mobile applications. Since then, Moore's Law has come to the rescue for first WiMAX then UMB and now LTE.

Assessing the advantages of OFDM

The primary advantage of OFDM is its resistance to the damaging effects of multipath delay spread (fading) in the radio channel. Without multipath protection, the symbols in the received signal can overlap in time, leading to inter-symbol interference (ISI). In OFDM systems designed for use in multipath environments, ISI can be avoided by inserting a guard period, known as the cyclic prefix (CP), between each transmitted data symbol. The CP is a copy of the end of the symbol inserted at the beginning. By sampling the received signal at the optimum time, the receiver can avoid all ISI caused by delay spread up to the length of the CP.

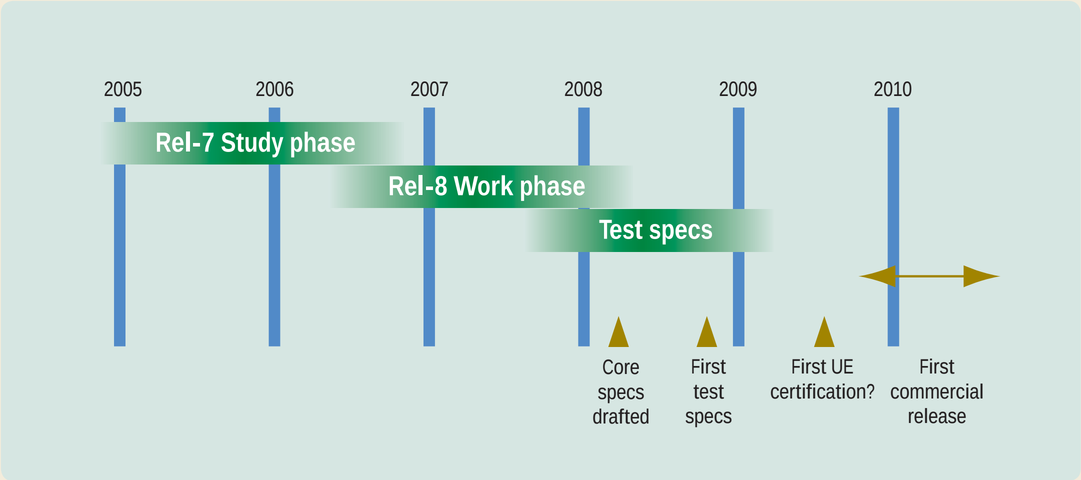

Figure 1. LTE timing

The CP is chosen to be slightly longer than the longest expected delay spread in the radio channel. For the cellular LTE system, the standard CP length has been set at 4.69 µs, enabling the system to cope with path delay variations up to about 1.4 km. Note that this figure represents the difference in path length due to reflections, not the size of the cell.* Inserting a CP between every symbol reduces the data handling capacity of the system by the ratio of the CP to the symbol length. For LTE, the symbol length is 66.7 µs, which gives a small but significant seven percent loss of capacity when using the standard CP.

The ideal symbol length in OFDM systems is defined by the reciprocal of the subcarrier spacing and is chosen to be long compared to the expected delay spread. LTE has chosen 15 kHz subcarrier spacing, giving 66.7 µs for the symbol length. In a single-carrier system, the symbol length is closely related to the occupied bandwidth. For example, GSM has 200 kHz channel spacing and a 270.833 ksps symbol rate, giving a 3.69 µs symbol length that is 18 times shorter than that of LTE. In contrast, W-CDMA has 5 MHz channel spacing and a 3.84 Msps symbol rate, producing a 0.26 µs symbol length — 256 times shorter than LTE. It would be impractical to insert a 4.69 µs CP between such short symbols because capacity would drop by more than half with GSM and by a factor of 20 with W-CDMA. Systems that use short symbol lengths compared to the delay spread must rely on receiver-side channel equalizers to recover the original signal.

The link between channel bandwidth and symbol length puts single-carrier systems at a disadvantage versus OFDM when the channel bandwidths get wider. Consider a radio channel with 1 µs of delay spread: A 5 MHz single-carrier signal would experience approximately five symbols of ISI and a 20 MHz signal would experience approximately 20 symbols of ISI. The amount of ISI determines how hard the equalizer has to work and there exists a practical upper limit of about 5 MHz beyond which equalizer costs rise and performance drops off.

Each 15 kHz subcarrier in LTE is capable of transmitting 15 ksps, giving LTE a raw symbol rate of 18 Msps at its 20 MHz system bandwidth (1200 subcarriers, 18 MHz). Using 64QAM — the most complex of the LTE modulation formats — in which one symbol represents six bits, the raw capacity is 108 Mbps. Note that actual peak rates as described in the LTE sidebar are derived by subtracting coding and control overheads and adding gains from features such as spatial multiplexing.

OFDM’s other main advantage over single-carrier systems is the ease with which it can adapt to frequency and phase distortions in the received signal, whether caused by transmitter impairments or radio-channel imperfections. Transmitted and received signals are represented in the frequency domain by subcarrier phase and amplitude. By seeding the transmitted signal across the frequency domain with many reference signals (RS, known in other systems as pilots) of predetermined amplitude and phase, the receiver can easily correct for frequency-dependent signal distortions prior to demodulation. This correction is particularly necessary when using higher-order modulation formats (e.g., 16QAM, 64QAM) that are susceptible to erroneous symbol demodulation caused by even small errors in phase and amplitude.

This ability to easily manipulate phase and frequency also lends itself to the processing required for multiple-input/multiple-output (MIMO) antenna techniques such as spatial multiplexing and beamforming. The required manipulations of signal phase and amplitude are much easier to implement in OFDM systems than in single-carrier systems, which represent signals in the time domain.

To summarize the advantages, OFDM systems transmit multiple low-rate subcarriers — resistant to multipath — that combine by the hundreds and thousands to provide a truly scalable system bandwidth and associated data rates. In addition, the frequency-domain representation of signals simplifies the correction of signal errors in the receiver and reduces the complexity of MIMO implementation. By contrast, single-carrier systems do not scale well with bandwidth and are impractical at much above 5 MHz if path delay differences are long.

*Longer CP lengths are available for use in larger cells and for specialist multi-cell broadcast applications. This provides protection for up to 10 km delay spread but with a proportional reduction in the achievable data rates.

Examining the disadvantages of OFDM

OFDM has two big disadvantages when compared to single-carrier systems. First, as the number of subcarriers increases, the composite time-domain signal starts to look like Gaussian noise, which has a high peak-to-average ratio (PAR) that can cause problems for amplifiers. Allowing the peaks to distort is unacceptable because this causes spectral regrowth in the adjacent channels. Modifying an amplifier to avoid distortion often requires increases in cost, size and power consumption. There exist techniques to limit the peaks (e.g., clipping and tone reservation*) but all have limits and can consume significant processing power while degrading in-channel signal quality.

The other main disadvantage of OFDM systems is caused by tight spacing of subcarriers. To minimize the lost efficiency caused by inserting the CP, it is desirable to have very long symbols, which means closely spaced subcarriers; however, apart from increasing the required processing, close subcarriers start to lose their orthogonality (independence from each other) due to frequency errors. Three key problems associated with close subcarriers cause lost performance. First, any frequency error in the receiver will cause energy from one subcarrier's symbol to interfere with the next. Second, phase noise in the received signal causes similar ISI in the subcarriers but on both sides. Third, Doppler shift can cause havoc. It is easy to remove a fixed Doppler shift but consider the case when multipath is involved and signals are arriving at the receiver from both front and back: The received signals are shifted both higher and lower in frequency and it takes considerable processing power to recover the original signal. To balance the desire for long symbols with the problems caused by close subcarrier spacing, LTE has adopted 15 kHz spacing, with a narrower 7.5 kHz chosen for use with LTE's solution for mobile TV, the evolved Multimedia Broadcast Multicast Service (eMBMS).[1]

*Tone reservation is an advanced form of clipping in which the time-domain signal is shaped such that the error energy falls on specific, reserved in-channel frequencies, ensuring less distortion in the wanted part of the signal.

Introducing SC-FDMA

The undesirable high PAR of OFDM led 3GPP to choose a different modulation format for the LTE uplink. This difference contributed to the inability of TTA, the Korean standards body, to persuade 3GPP (in 2006) to merge LTE and WiMAX. Pure OFDM is used in the WiMAX uplink but LTE continued to use SC-FDMA, a new hybrid modulation scheme that cleverly combines the low PAR of single-carrier systems with the multipath resistance and flexible subcarrier frequency allocation offered by OFDM.

When a new concept in modulation comes along (e.g., OFDM or CDMA), it can take a long time before the literature starts to make sense. Yet, after everyone "gets it," we often look back at what previously seemed to be impenetrable explanations and wonder what the fuss was about! So it may be with SC-FDMA. The Release 8 3GPP specifications do little to explain the concept. For a formal definition of SC-FDMA, a student of signal processing need look no further than TS 36.211, which gives the mathematical description of the time-domain representation of an SC-FDMA symbol.[2] For the majority of us who find the formal mathematical approach hard to follow, we will present here a graphical comparison of the differences between OFDM and SC-FDMA.

Comparing OFDM and SC-FDMA

Figure 2 shows how a series of QPSK symbols are mapped into time and frequency by the two different modulation schemes. Rather than using OFDM, we will now shift to the term OFDMA, which stands for orthogonal frequency-division multiple access. OFDMA is simply an elaboration of OFDM used by LTE and other systems that increases system flexibility by multiplexing multiple users onto the same subcarriers. This can benefit the efficient trunking of many low-rate users onto a shared channel as well as enable per-user frequency hopping to mitigate the effects of narrowband fading. For clarity, the example here uses only four (N) subcarriers over two symbol periods with the payload data represented by QPSK modulation. Real LTE signals are allocated in units of 12 adjacent subcarriers (180 kHz) called resource blocks that last for 0.5 ms and usually contain seven symbols whose modulation can be QPSK, 16QAM or 64QAM.

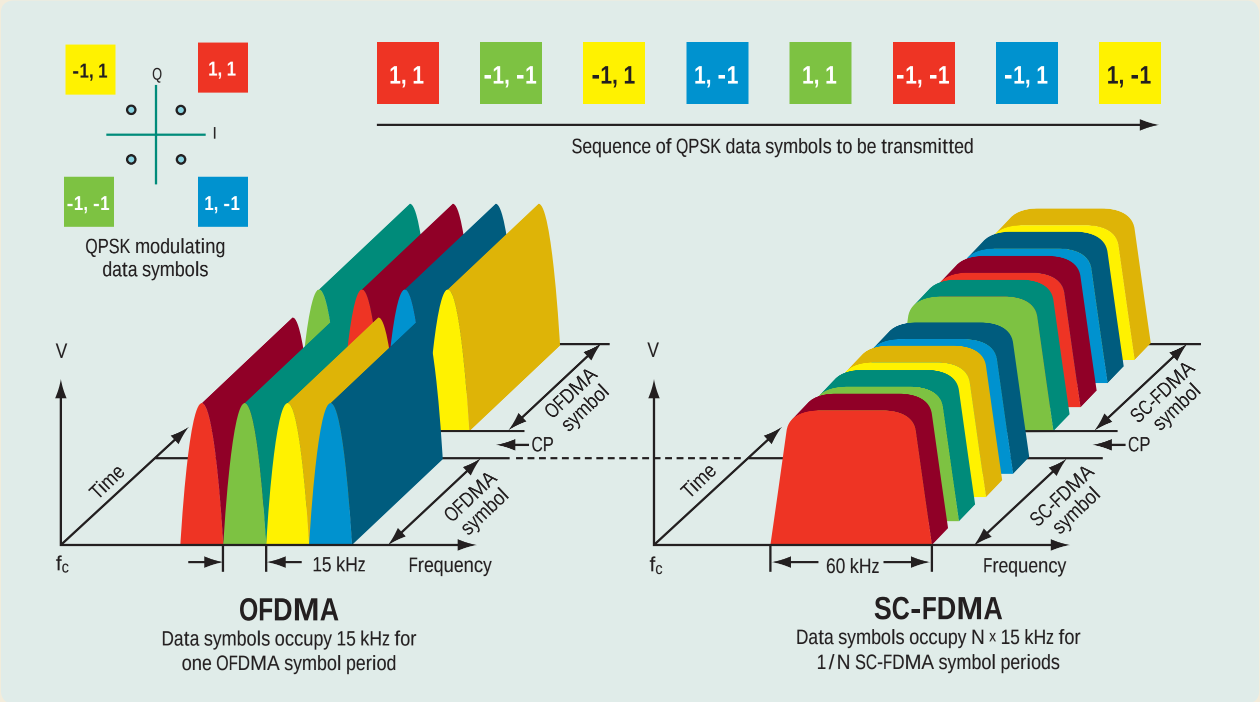

Figure 2. Comparison of how OFDMA and SC-FDMA transmit a sequence of QPSK data symbols

The LTE downlink uses traditional OFDMA methods and differs from other systems such as UMB and WiMAX only in details of the OFDM numerology (that is subcarrier spacing, symbol length, bandwidth, etc.). On the left side of Figure 2, N adjacent 15 kHz subcarriers — already positioned at the desired place in the channel bandwidth — are each modulated for the OFDMA symbol period of 66.7 µs by one QPSK data symbol. In this simple four-subcarrier example, four symbols are taken in parallel. These are QPSK data symbols so only the phase of each subcarrier is modulated and the subcarrier power remains constant between symbols. After one OFDMA symbol period has elapsed, the CP is inserted and the next four symbols are transmitted in parallel. For visual clarity, the CP is shown as a gap; however, it is actually filled with a copy of the end of the next symbol, meaning the transmission power is continuous but has a phase discontinuity at the symbol boundary. To create the transmitted signal, an inverse FFT is performed on each subcarrier to create N time-domain signals that are vector summed to create the final time-domain waveform used for transmission.

SC-FDMA signal generation begins with a special precoding process but then continues as with OFDMA. Before outlining the generation process it is helpful to first describe the end result as shown on the right side of Figure 2. The most obvious difference between the two schemes is that OFDMA transmits the four QPSK data symbols in parallel, one per subcarrier, while SC-FDMA transmits the four QPSK data symbols in series at four times the rate, with each data symbol occupying N x 15 kHz bandwidth. Visually, the OFDMA signal is clearly multi-carrier and the SC-FDMA signal looks more like single-carrier, which explains the “SC” in its name. Note that OFDMA and SC-FDMA symbol lengths are the same at 66.7 µs; however, the SC-FDMA symbol contains N “sub-symbols” that represent the modulating data.

It is the parallel transmission of multiple symbols that creates the undesirable high PAR of OFDMA. By transmitting the N data symbols in series at N times the rate, the SC-FDMA occupied bandwidth is the same as multi-carrier OFDMA but — crucially — the PAR is the same as that used for the original data symbols. This should make heuristic sense without delving into the mathematics: Adding together many narrowband QPSK waveforms in OFDMA will always create higher peaks than would be seen in the wider-bandwidth single-carrier QPSK waveform of SC-FDMA. As the number of subcarriers N increases, the PAR of OFDMA with random modulating data approaches Gaussian noise statistics but, regardless of the value of N, the SC-FDMA PAR remains the same as that used for the original data symbols.

Having seen what SC-FDMA looks like, we will now explain the precoding process that brings it about. Figure 3 shows the first steps, which create a time-domain waveform of the QPSK data sub-symbols. Using the four color-coded QPSK data symbols from Figure 2, the process creates one SC-FDMA symbol in the time domain by computing the trajectory traced by moving from one QPSK data symbol to the next. This is done at N times the rate of the SC-FDMA symbol such that one SC-FDMA symbol contains N consecutive QPSK data symbols. For simplicity, we will not discuss time-domain filtering of the data symbol transitions even though such filtering will be present in any real implementation.

Having created an IQ representation in the time domain of one SC-FDMA symbol, the next stage is to represent it in the frequency domain using a discrete Fourier transform (DFT; Figure 4). The DFT sampling frequency is chosen such that the time-domain waveform of one SC-FDMA symbol is fully represented by N DFT bins spaced 15 kHz apart, with each bin representing one subcarrier with amplitude and phase held constant for 66.7 µs. There is always a one-to-one correlation between the number of data symbols to be transmitted during one SC-FDMA symbol period and the number of DFT bins created — and this in turn becomes the number of occupied subcarriers. This should make intuitive sense: When an increasing number of data symbols is transmitted during one SC-FDMA period, the time-domain waveform changes faster, generating a higher bandwidth and hence requiring more DFT bins to fully represent the signal in the frequency domain.

Note from Figure 4 that there is no longer a direct relationship between the amplitude and phase of the individual DFT bins and the original QPSK data symbols. This is quite different from the OFDMA example in which data symbols directly modulate the subcarriers.

The next stage is to shift the baseband DFT representation of the time-domain SC-FDMA symbol to the desired part of the overall channel bandwidth. Because the signal is now represented as a DFT, frequency shifting is a very simple process achieved by copying the N bins into a larger DFT space that can be up to the size of the system channel bandwidth — of which there are six to choose from in LTE, spanning 1.4 MHz to 20 MHz. The elegance of the DFT lets us position the signal anywhere in the channel bandwidth, thus executing the frequency-division multiple access (FDMA) essential for efficiently sharing the uplink between multiple users. *This explains the origin of the latter portion of "SC-FDMA."

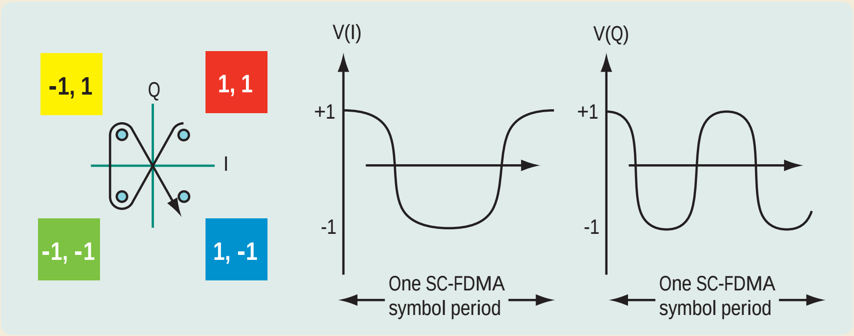

Figure 3. Creating the time-domain waveform of an SC-FDMA symbol

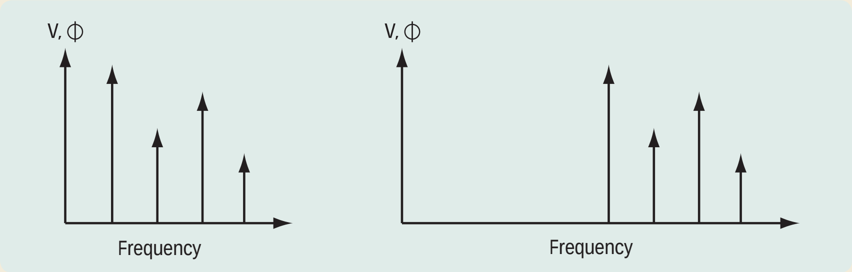

Figure 4. Baseband and shifted frequency domain representations of an SC-FDMA symbol

*Although 3GPP did consider a distributed form of subcarrier allocation for the uplink that would have alleviated susceptibility to narrowband fading, it instead opted for the adjacent allocation described here combined with the possibility of frequency hopping at the slot (0.5 ms) level.

To conclude SC-FDMA signal generation, the process follows the same steps as for OFDMA. Performing an inverse FFT converts the frequency-shifted signal to the time domain and inserting the CP provides OFDMA's fundamental robustness against multipath.

If we now return to Figure 2's representation of OFDMA and SC-FDMA, we can consider how each signal would look depending on the analysis bandwidth. Table 1 summarizes the differences between the modulation formats.

When analyzed one subcarrier at a time, OFDMA resembles the original data symbols. At full bandwidth, however, the signal looks like Gaussian noise in terms of its PAR statistics and the constellation. The opposite is true for SC-FDMA. Its relationship to the original data symbols is evident when analyzing the entire signal bandwidth whereupon the constellation (and hence low PAR) of the original data symbols can be observed rotating at N times the SC-FDMA symbol rate (ignoring the seven percent rate reduction due to adding the CP). When analyzed at the subcarrier bandwidth, the SC-FDMA PAR and constellation are meaningless because these are N times narrower than the information bandwidth of the data symbols.

Table 1. Analysis of OFDMA and SC-FDMA at different bandwidths

Modulation format	OFDMA		SC-FDMA	
Analysis bandwidth	15 kHz	Signal BW (N x 15 kHz)	15 kHz	Signal BW (N x 15 kHz)
Peak-to-average power ratio	Same as data symbol	High PAR (Gaussian)	Not meaningful (< data symbol)	Same as data symbol
Observable IQ constellation	Same as data symbol at 66.7 µs rate	Not meaningful (Gaussian)	Not meaningful (< data symbol)	Same as data symbol at N x 66.7 µs rate

LTE at a glance

November 2004 LTE / SAE High-level requirements

- Reduced cost per bit
- More lower-cost services with better user experience
- Flexible use of new and existing frequency bands
- Simplified lower-cost network with open interfaces
- Reduced terminal complexity and reasonable power consumption

Speed

Downlink peak data rates (64QAM)

Antenna configuration	SISO	2x2 MIMO	4x4 MIMO
Peak data rate (Mbps)	100	172.8	326.4

Uplink peak data rates (single antenna)

Modulation depth	QPSK	16QAM	64QAM
Peak data rate (Mbps)	50	57.6	86.4

Services

Packet-switched voice and data. No circuit-switched services supported.

Flexible channel bandwidths

Bandwidth MHz	Access mode
1.4	FDD and TDD
3	FDD and TDD
5	FDD and TDD
10	FDD and TDD
15	FDD and TDD
20	FDD and TDD

The 1.6 MHz and 3.2 MHz TDD bandwidths have recently been deleted, and the six remaining bandwidths apply to both FDD and TDD.

Mobility

Optimized: 0 to 15 km/h

High performance: 15 to 120 km/h

Functional: 120 to 350 km/h

Under consideration: 350 to 500 km/h

Spectral Efficiency

3-4x Rel-6 HSDPA (downlink)

2-3x Rel-6 HSUPA (uplink)

Latency

Idle to active < 100 ms

Small packets < 5 ms

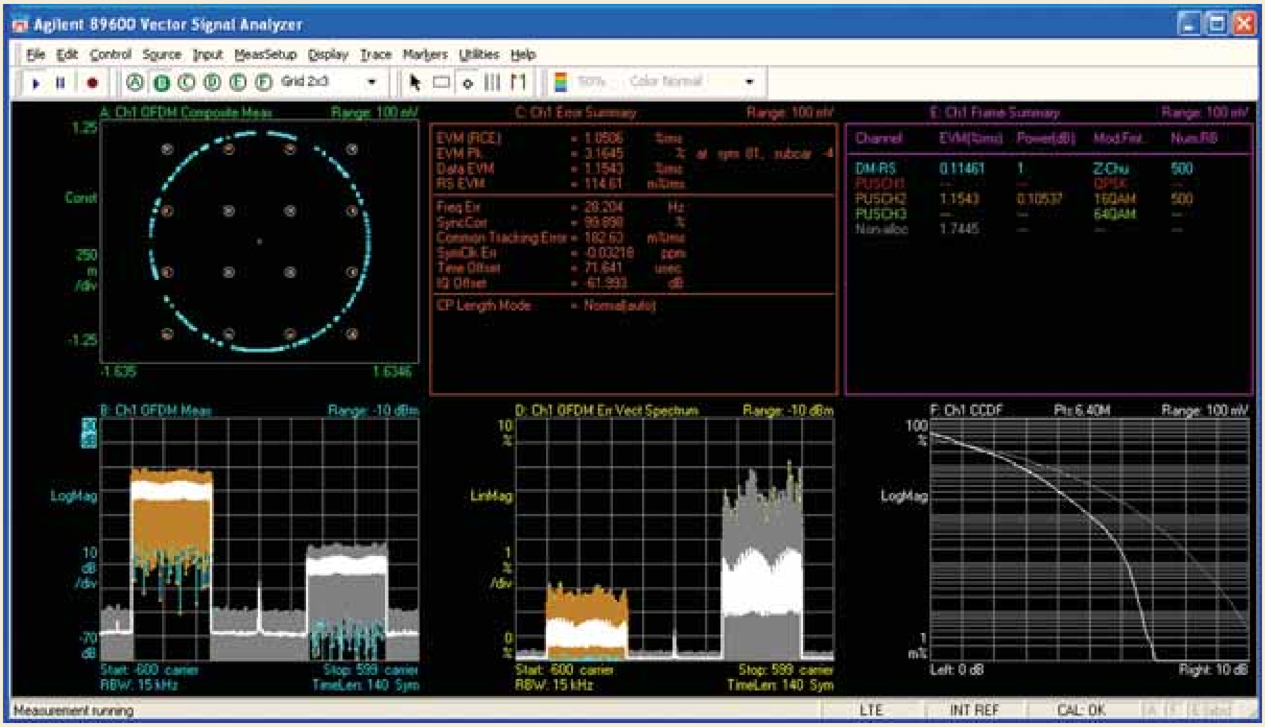

Figure 5. Analysis of a 16QAM SC-FDMA signal

Multipath resistance with short data symbols?

At this point it is reasonable to ask, “How can SC-FDMA still be resistant to multipath when the data symbols are still short?” In OFDMA, the modulating data symbols are constant over the 66.7 µs OFDMA symbol period but an SC-FDMA symbol is not constant over time since it contains N sub-symbols of much shorter duration. The multipath resistance of the OFDMA demodulation process seems to rely on the long data symbols that map directly onto the subcarriers. Fortunately, it is the constant nature of each subcarrier — not the data symbols — that provides the resistance to delay spread. As shown earlier, the DFT of the time-varying SC-FDMA symbol generated a set of DFT bins constant in time during the SC-FDMA symbol period even though the modulating data symbols varied over the same period. It is inherent to the DFT process that the time-varying SC-FDMA symbol — made of N serial data symbols — is represented in the frequency domain by N time-invariant subcarriers. Thus, even SC-FDMA with its short data symbols benefits from multipath protection.

It may seem counterintuitive that N time-invariant DFT bins can fully represent a time-varying signal. However, the DFT principle is simply illustrated by considering the sum of two fixed sine waves at different frequencies: The result is a non-sinusoidal time-varying signal — fully represented by two fixed sine waves.

Examining a real SC-FDMA signal

Figure 5 shows some of the measurements that can be made on a typical SC-FDMA signal. The constellation in trace A (top left) shows this is a 16QAM signal. The unity circle represents the RS (every seventh symbol), which do not use SC-FDMA but are phase modulated using an orthogonal Zadoff-Chu sequence.[3] Trace B (lower left) shows the signal power versus frequency. The frequency scale is in 15 kHz subcarriers numbered from –600 to 599, which represents a bandwidth of 18 MHz. From this we can conclude this must be a 20 MHz channel and the allocated signal bandwidth is 5 MHz towards the lower end. The brown dots represent the instantaneous subcarrier amplitude and the white dots the average over 10 ms. In the center of the trace, the spike represents the LO leakage (IQ offset) of the signal; the large image to the right is an OFDM artifact deliberately created using 0.5 dB IQ gain imbalance in the signal. Both the LO leakage and the power in non-allocated subcarriers will be limited by the 3GPP specifications.

Trace C (top middle) shows a summary of the measured impairments including the error vector magnitude (EVM), frequency error and IQ offset. Note the data EVM at 1.15 percent is much higher than the RS EVM at 0.114 percent. This is due to a +0.1 dB boost in the data power as reported in trace E, which was ignored (for illustration) by the receiver to create data-specific EVM. Also note the RS power boost is reported as +1 dB, which can also be observed in the IQ constellation because the unity circle does not pass through eight of the 16QAM points. Trace D (lower middle) shows the distribution of EVM by subcarrier. The average and peak of the allocated signal EVM is in line with the numbers in trace C. The EVM for the non-allocated subcarriers reads much higher although the way this in-channel impairment will be specified will be as a power ratio between the wanted signal and the unwanted signal. The ratio for this signal is around 30 dB as can be seen in trace B. The blue dots in trace D also show the EVM of the RS, which is very low.

Trace E (top right) shows the ability to measure EVM by modulation type from one capture. This signal uses only the RS phase modulation and 16QAM so the QPSK and 64QAM results are blank. Finally, trace F (lower right) shows the PAR — the whole point of SC-FDMA — in the form of a complementary cumulative distribution function (CCDF) measurement. It is not possible to come up with a single figure of merit for the PAR advantage of SC-FDMA over OFDMA because it depends on the data rate. The PAR of OFDMA is always higher than SC-FDMA even for narrow frequency allocations; however, when data rates rise and the frequency allocation gets wider, the SC-FDMA PAR remains constant but OFDMA gets worse and approaches Gaussian noise. A 5 MHz OFDMA 16QAM signal would look very much like Gaussian noise. From the white trace it can be seen at 0.01 percent probability the SC-FDMA signal is 3 dB better than the Gaussian reference trace, and as every amplifier designer knows, even a tenth of a decibel shaved from the peak power budget is money in the bank.

Conclusion

In essence, SC-FDMA means “create a single-carrier waveform and shift it to the desired part of the frequency domain.” After a careful consideration of the characteristics of OFDMA and the new SC-FDMA, we can conclude that SC-FDMA provides the advantages of OFDMA — especially robust resistance to multipath — without the problem of high PAR. The use of SC-FDMA in LTE, however, is restricted to the uplink because the increased time-domain processing would be a considerable burden on the base station, which has to manage the dynamics of multi-user transmission.

It will be interesting to see if LTE — the latest of the three new OFDMA cellular standards — has indeed identified a superior solution for the uplink or whether the pure OFDMA used in WiMAX or the OFDMA/CDMA combination used in UMB prove to be just as successful when all the factors are taken into account. Today, the experts disagree so we will have to wait on the ultimate arbiter, time, before we find out for sure.

References

1. 3GPP TS 36.201 v8.0.0 section 4.2.1
2. 3GPP TS 36.211 v8.0.0 subclause 5.6
3. 3GPP TS 36.211 v8.0.0 subclause 5.5

Correction

In Issue Three of *Agilent Measurement Journal*, the article “What Next for Mobile Telephony?” presents the “G factor” equation (page 35) in its general form as $G = \hat{I}_{or} / (\hat{I}_{or} + I_{oc})$. The article should have used the 3GPP form of the equation, which is $G = \hat{I}_{or} / I_{oc}$.

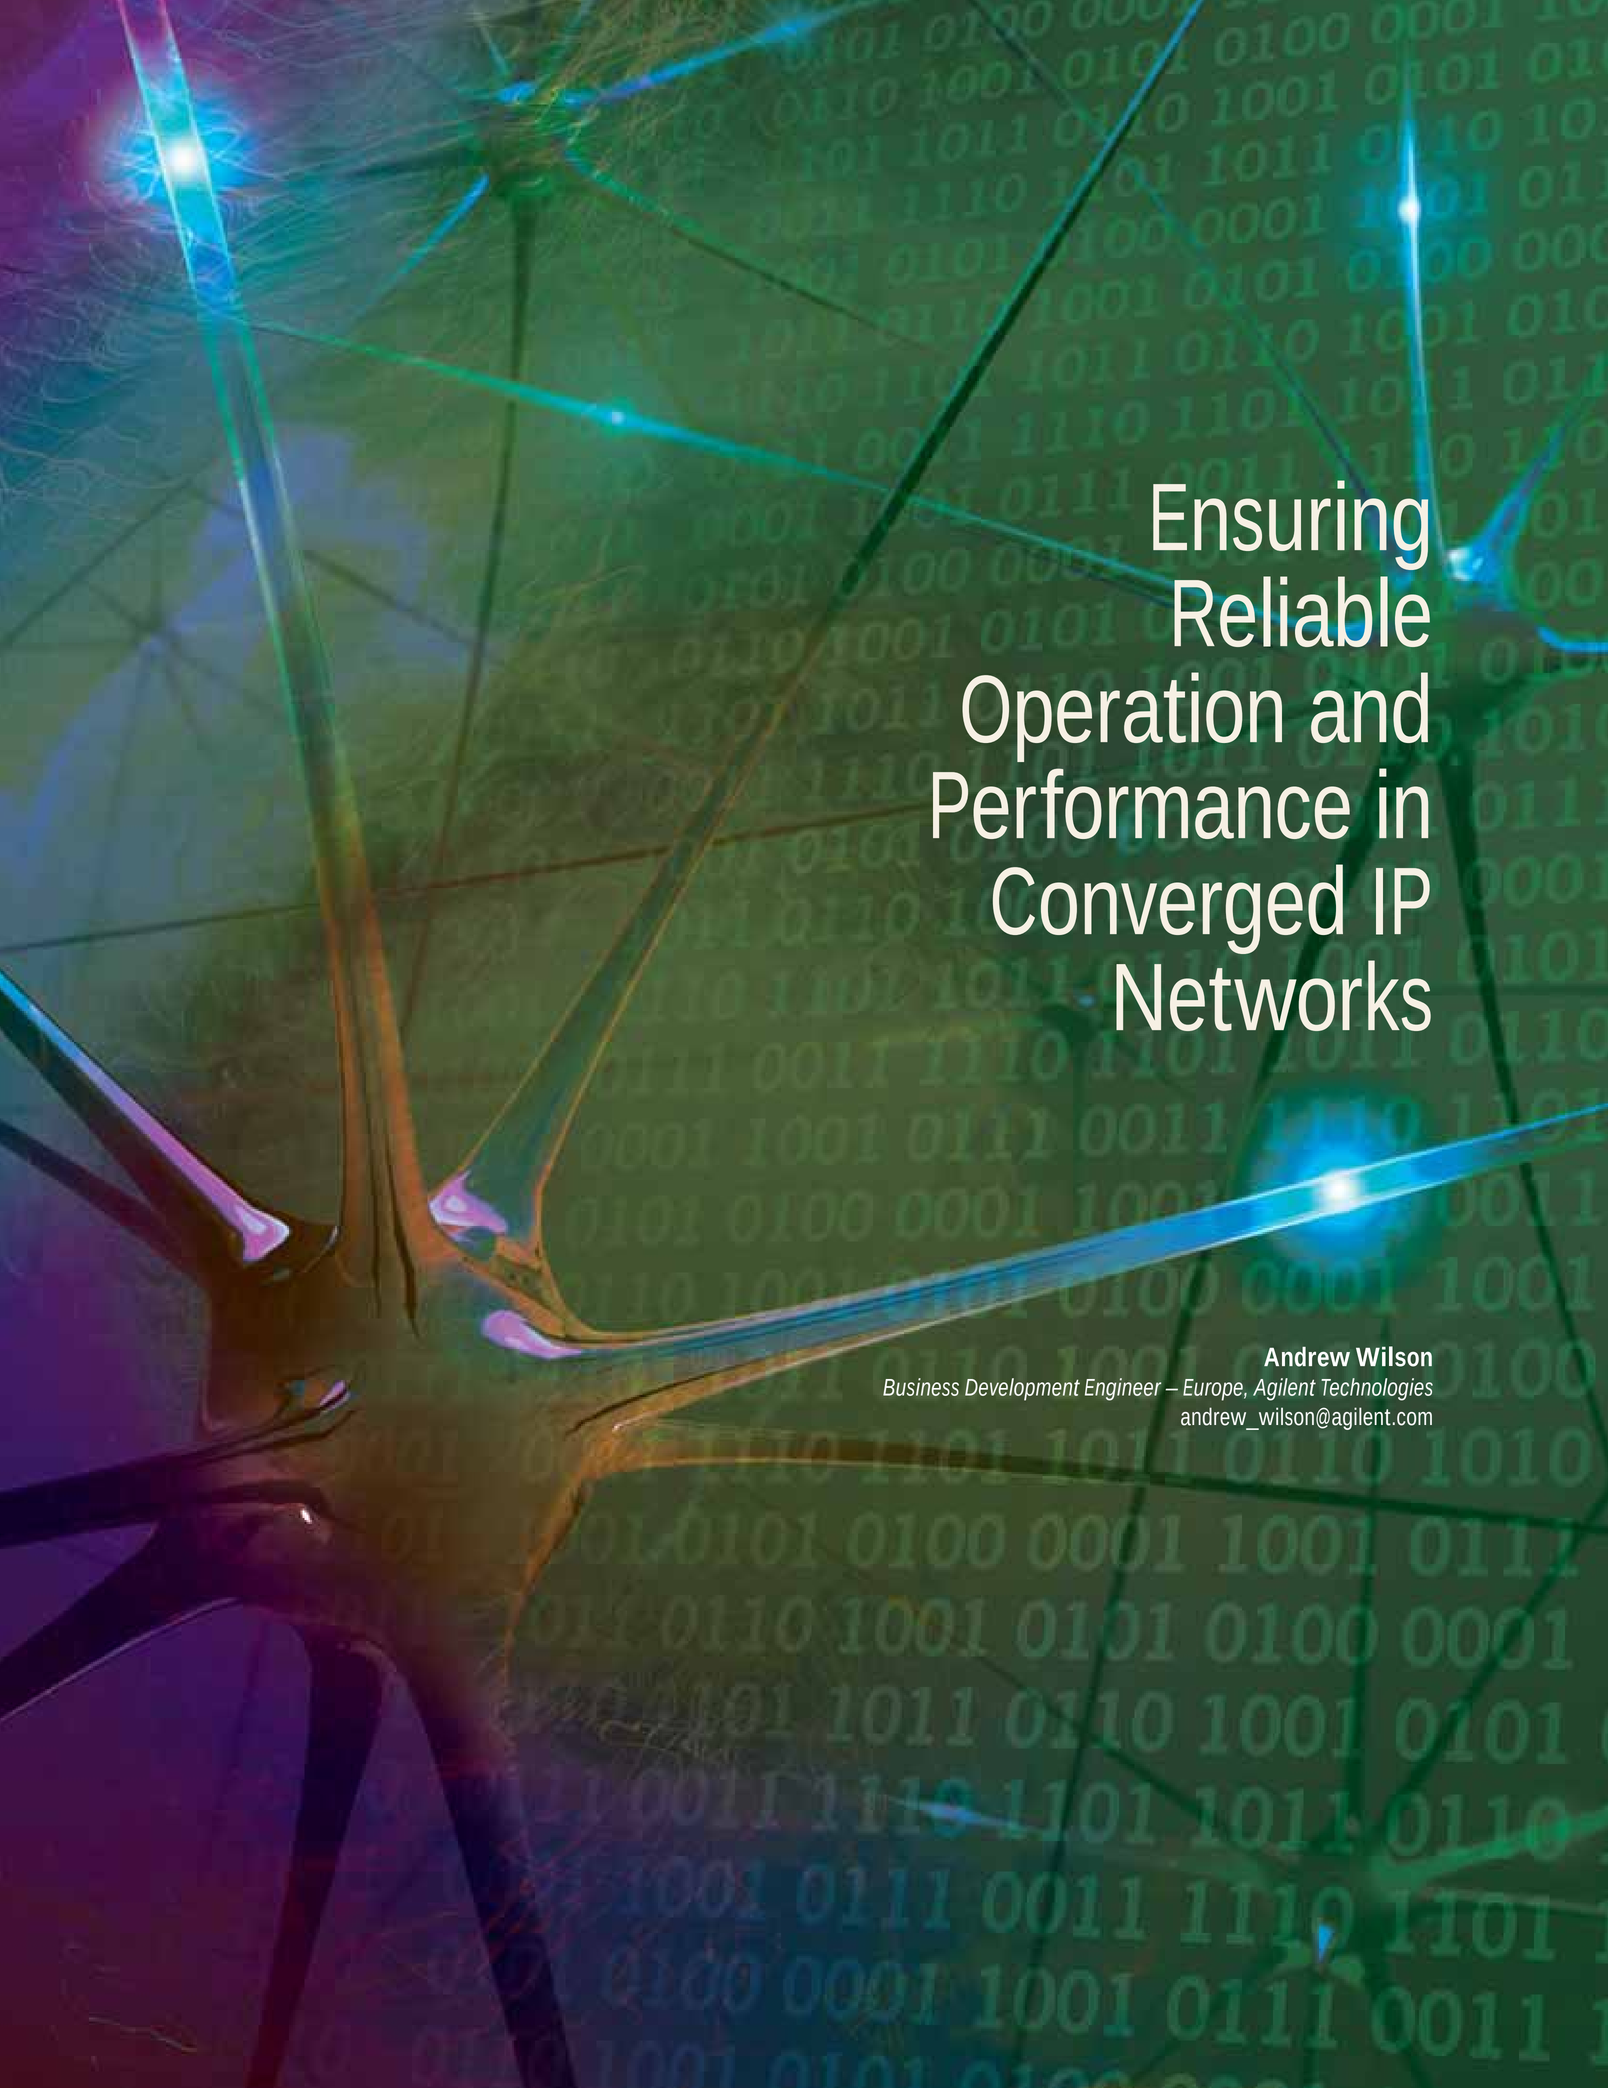

Ensuring Reliable Operation and Performance in Converged IP Networks

Andrew Wilson
Business Development Engineer – Europe, Agilent Technologies
andrew_wilson@agilent.com

The integration or convergence of voice, data and video on a single network infrastructure is a major goal for the telecom industry. The combined benefits to service providers (economy of scale) and service users (convenience) are tantalizing. Add the possibilities of new, interactive services such as video-on-demand, camera angle control or targeted content, and the move to convergence becomes seemingly inevitable.

The inevitable, however, may be delayed by one crucial factor: Existing technologies for phone and TV services are currently so good — and any deficiencies so easily identified — that end-user tolerance of problems is extremely low. In that context, the success of converged services delivered via Internet Protocol (IP) networks depends on two key factors:

- Replicating existing phone and TV services well enough to provide a positive quality of experience (QoE).
- Maintaining quality of service (QoS) for each user independent of total traffic loading and type.

To help satisfy those needs, service providers and equipment makers need tools that can perform rigorous testing of next-generation routers, switches, gateways and so on during design and development. This article explores a few ways to perform thorough, robust testing of devices and trial networks to ensure positive QoE in the field.

Looking back: A very brief history of telecom

In its first 150 years, the circuit-switched world of telecom never achieved convergence, despite heroic attempts. Voice-band modems — still alive and well in the standard fax machine —came closest in that they provide an end-to-end data connection compatible with the voice channel and the signaling mechanism of the existing public switched telephone network (PSTN). Unfortunately, modems are too slow in both transmission speed and call setup time to support delivery of real-time video or transfer of large files.

Later, the industry devised integrated services digital network (ISDN) to extend the digital core of the circuit-switched network out to end users. While this allowed fast data-circuit setup, ISDN access transmission rates were still insufficient for the services people wanted to use. As a result, ISDN is rapidly disappearing from general use, remaining strong only in niche applications such as broadcast-quality sound connections over the public network.

Switching to packets and routers

Enter the Internet in the 1980s. Through a combination of suitability and chance, it has become pervasive. Catalyzed by the content of the World Wide Web and mass access to high-bandwidth data connections via cable or digital subscriber line (DSL), huge investments have been made in networks based on routed IP datagrams or packets. These IP networks — and the use of IP — offers a presumption of convergence driven by the apparent ease of using a common packet infrastructure to carry all service types. It's also straightforward to demonstrate that, in and of itself, the adaptation of "tricky" real-time services such as voice and video into and out of IP packets causes no perceptible quality problems. In some ways, the business models that propelled some recent infrastructure investments seem to assume that the shift to a converged-services IP network has a smooth evolution path.

The shift can be smooth, but not without some work. The snag is that IP networks were never intended for real-time services such as phone calls and broadcast or on-demand TV, nor to isolate service quality for individual users. Rather, they were designed to be used as a base for file transfer in conjunction with a number of other significant protocols — in particular Transmission Control Protocol (TCP), which provides the necessary reliability. Unfortunately, TCP is unsuitable for use with voice or real-time (streaming) video because it would introduce large and indeterminate delays. To allow for convergence with a mixture of TCP and real-time traffic, almost every aspect of IP networks has had to be hardened and extended to make them work with these diverse services at a huge scale. Today, however, the reality of the technology still lags content and service providers' commercial aspirations.

Reviewing IP networking

The principal element in IP networks is the IP router, a device that forwards IP packets from incoming physical ports onto outgoing ports based on simple numerical addresses (e.g., the URL www.agilent.com is just the IPv4 address 157.238.197.58) and other service tags, all carried within the administrative area of the packet called the *header*. Packets pass from one router to the next in their journey through the network.

One fundamental characteristic of IP networks is that there is no inherent determinism about how long a packet should take to traverse the network. These times can vary depending on what types and quantities of traffic are present. In extreme conditions (or during faults), packets can be discarded altogether.

Until recently, routers based prioritization on only header information, with no regard to the data content of the packet. Such an egalitarian approach essentially leaves IP packets to fight it out for priority. To resolve this, routers increasingly are designed to look inside packets and make forwarding-priority decisions based on content as well as header information. While the actual policies may vary and evolve, they will always include elements such as service type and user identity. The networks of the future must intervene actively in the journey of every packet from end-to-end in order to provide the required QoS to each service and user.

Testing with extreme realism

To accurately test the performance of these devices, it's necessary to construct test traffic that represents huge numbers of users and all plausible service types. Each type has its own challenges: Phone and TV services have significant user-signaling aspects that impact QoE as much as content quality. What's more, integration of these signaling (or control-plane) aspects into the common IP structure means that forwarding devices must be tested using a traffic mix that also contains these packet types.

TV service is particularly demanding on the control-plane side because specialized routers must terminate and interpret control-plane packets in order to select channels. Reaction time between receiving a *channel join* request and commencing service-packet delivery to a user is a vital measurement, especially at high scale (Figure 1). By its very nature, TV can cause coordinated subscriber behavior (e.g., commercial breaks stimulate channel surfing) that can cause large peaks in the signaling load.

Looking at data content, TV also is the most difficult service to accommodate in the converged IP network. To conserve network bandwidth, TV service needs to use a network mechanism called *multicast* in which the traffic streams associated with each available channel are sent to a manageable number of replication points distributed throughout the network. From a test point of view, the intermediate routing devices need to be tested to ensure they can properly forward multicast packets. To facilitate that, the test equipment must be capable of emulating the control protocols used to relate video sources to multicast replication points.

Even with multicast, it's necessary to minimize bandwidth per TV channel because subscriber connections are usually limited to a few megabits per second. This requires extreme compression techniques (e.g., MPEG) that are particularly sensitive to packet loss and delay variations. Test traffic needs to simulate the structure of such compressed video.

Simulating TCP traffic in particular poses a huge challenge. For routers to recognize and hence prioritize a TCP connection associated with a particular customer, the connection needs to look "real." To achieve this, test traffic must simulate the handshaking and retransmission mechanisms that define TCP over long sequences of packets. It also must do so at a realistic scale of tens of thousands of active TCP connections on a single router port.

The requirements for realism don't stop at the TCP level because the purpose of TCP connections is to carry web pages (HTTP), files (FTP) or peer-to-peer traffic. What's most demanding is simulating large numbers of fully realistic transactions in a predictable and consistent manner. As with any test stimulus, a lack of precision or determinism put into the system-under-test adds uncertainty to every measurement of what comes out. Therefore, accurate simulations are exceptionally difficult to achieve at high scale without resorting to processor-based schemes. Unfortunately, such approaches lack determinism and also struggle to provide real-time variation in traffic load, which makes it difficult to quickly

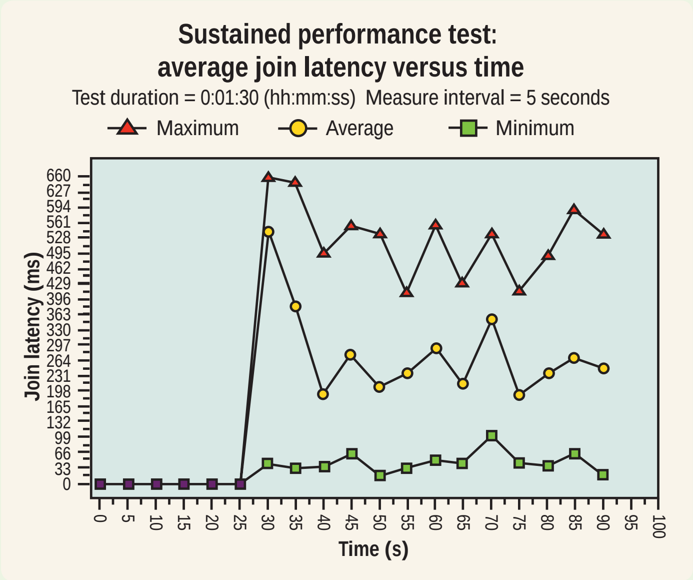

Figure 1. Measurements of channel-join latency versus time help gauge QoS in IPTV systems.

find the point at which devices reach their limits. To provide a deterministic stimulus, Agilent recommends — and provides — hardware-based high-scale TCP load generators (Figure 2).

Making meaningful measurements

Testing IP performance requires more than just the blasting of devices with a realistic mix and scale of traffic types. It is equally important to make measurements that are per-user and sensitive to all aspects of network performance that affect user QoE. Measures of simple packet loss and delay are certainly central but they aren't enough. In the case of video QoS, which goes beyond the channel-change measurements described earlier, it is vital to measure factors such as media delivery index (MDI) at large scale. Though based on packet loss and delay, MDI is optimized to relate directly to the effect of network-induced impairments on real TVs and set-top boxes. It's also very compelling to provide video and sound decode to allow selective subjective assessment of quality from either real or simulated video channels.

All network devices introduce delay, and its effect varies according to the type of service. For phone service, absolute delay must be significantly less than one second or else conversation is difficult. For broadcast TV, absolute delay is relatively unimportant (other than its effect on channel-change times) — but delay variation should be tightly controlled because the TV receiver must ultimately remove such variations by using buffer memory with a finite — and cost-sensitive — capacity. It's worth noting that delay measurements averaged over many user traffic streams are of little use. What's needed are detailed per-user measurements of packet delay characteristics (Figure 3).

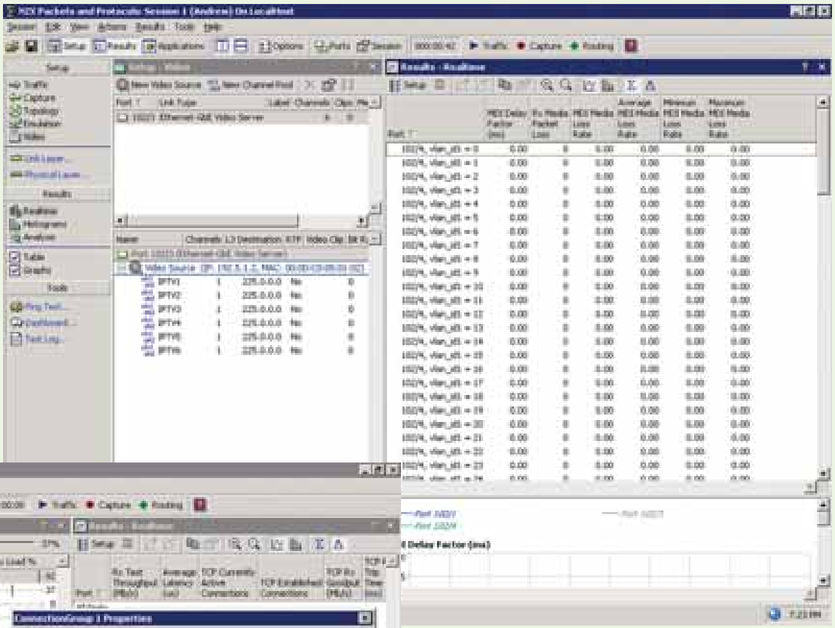

Figure 3. Simultaneous per-user measurement of delay variation and packet loss provides a true picture of IPTV delivery quality.

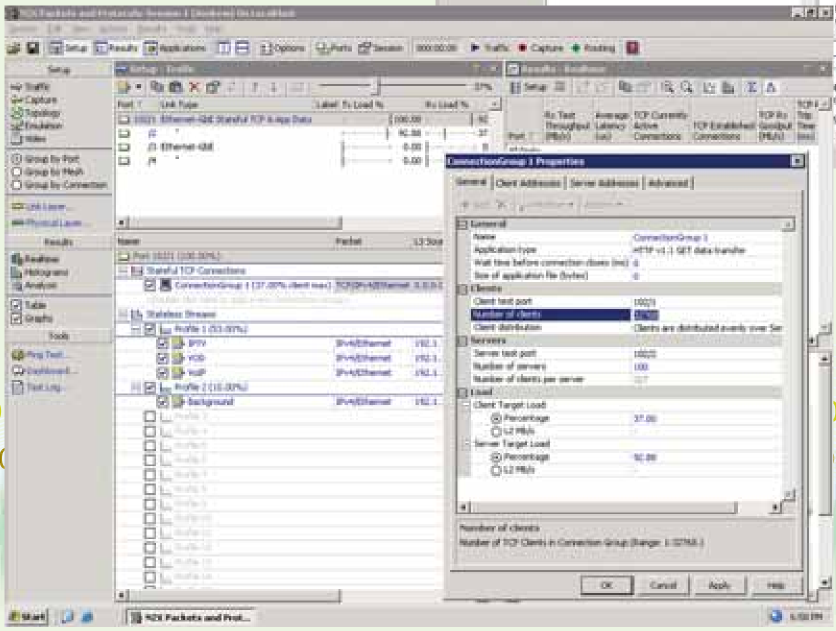

Figure 2. Varying the load dynamically over tens of thousands of active TCP connections allows accurate determination of device limits.

For TCP traffic, raw packet rates are an inconclusive metric. It's far more useful to make throughput measurements that account for the retransmission characteristics of TCP, again at the huge scale of real networks. As an example, a forwarding device may be capable of passing IP packets with a low rate of loss. However, if the lost packets happen to be spread out, then the number of TCP retransmissions can be large and therefore cause much slower overall throughput of application traffic.

With numerous network-user results accumulated, the test instrument must store, sort and present them in a way that makes it easy to spot problems and verify performance. With typical scale numbers of hundreds of thousands of users, the ability to extract useful insights depends on the filtering and grouping of results according to traffic characteristics and result values. Finding significant results would otherwise be akin to looking for a needle in a haystack (Figure 4).

Conclusion

The tantalizing possibility of converged voice, data and video is at last within our grasp. Such convergence is hugely attractive to service providers, but users will embrace an offered service only if it performs at least as well as existing mechanisms for phone and TV. To achieve convergence, the IP networks of today must evolve more in the next few years than telecom has in all the 150 years of its existence. This will be made possible by rigorous testing that accurately simulates real traffic types at real-world scale in the design, development and network-trial phases.

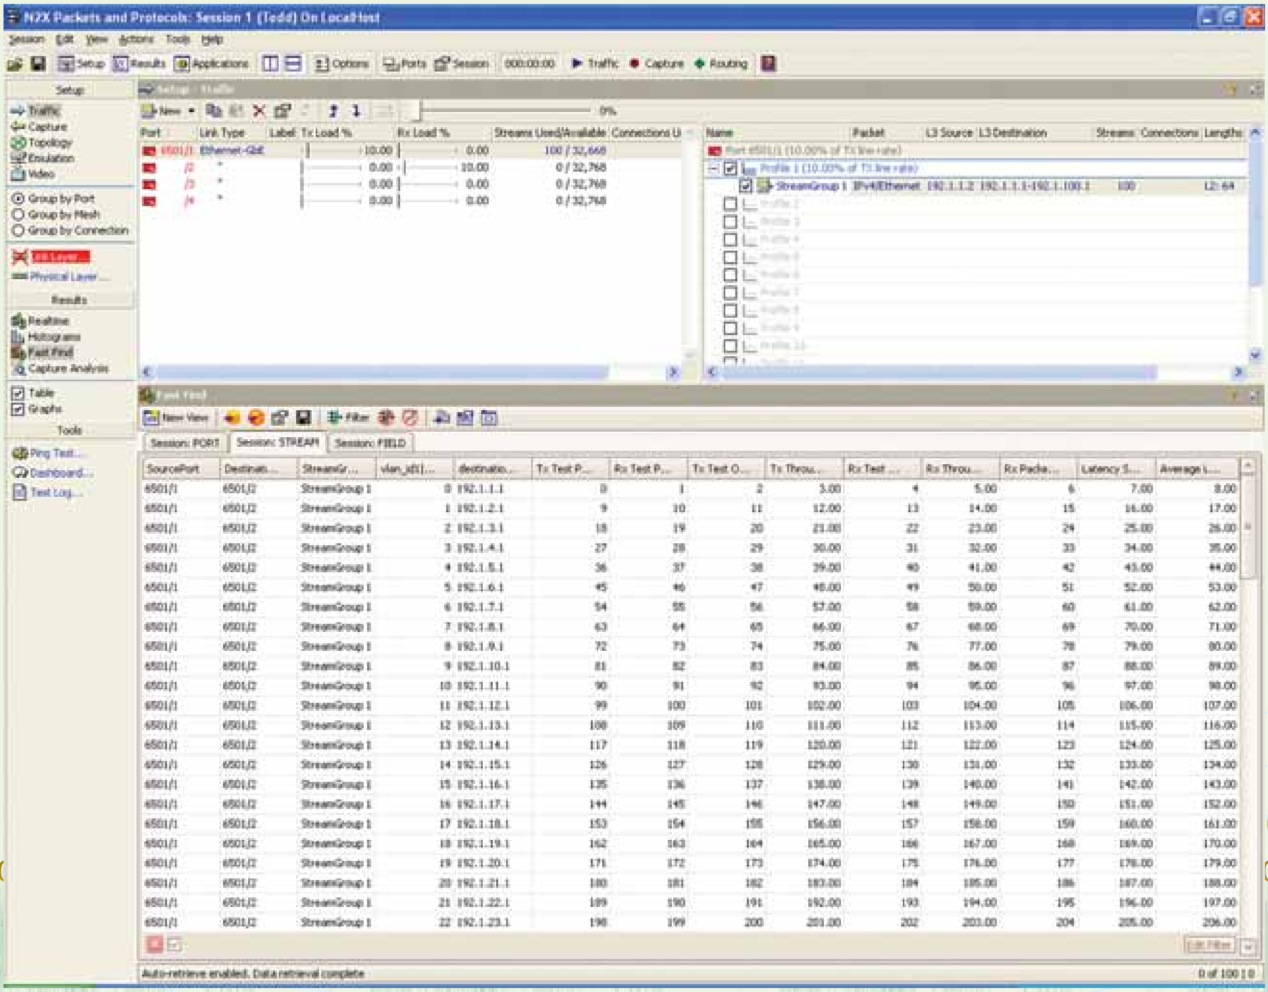

Figure 4. Tools for results analysis must scale up to match the vast quantities of collected test data.

Testing Storage Area Networks and Devices at 8.5-Gb/s Fibre Channel Rates

Yenyi Fu
Product Manager, Agilent Technologies
yenyi_fu@agilent.com

Consumer desire for online access to on-demand services and content is affecting seemingly every industry — from entertainment to banking, medicine to education and beyond. Examples include movies, video games and TV programming, all of which are rapidly moving to on-demand and online delivery.

Storage area networks (SANs) are used for these services, storing and delivering the large amounts of data that support the associated business models. However, growth in these areas places tremendous pressure on SAN developers to meet ever-increasing performance and scalability requirements. New application growth also is driving the addition of new features and capabilities into SANs, further increasing their complexity.

These factors are driving a test strategy reassessment for current SANs and SAN devices. Existing test strategies and environments, which are appropriate for small-scale SAN infrastructures, must be redefined to keep pace with the increasing complexity of SANs. The new strategies also must be compatible with the current reality of test budgets that are becoming tighter every year.

Examining traditional test environments

The early days of SAN resembled the early days of LAN and WAN: Lacking off-the-shelf test solutions, test teams developed proprietary test environments based on real computer servers and storage devices. Real applications and proprietary test software were used to load the SAN infrastructure with traffic representative of multiple simultaneous applications. While this approach provides obvious benefits for interoperability testing, the limitations of these hand-built solutions became more visible under the rigorous demands placed on SANs.

These limitations occur because typical servers and storage systems are not designed for testing. Inside each device there are usually multiple software layers (e.g., applications, file systems, operating systems, drivers) that cause indeterminism in the traffic injected into storage networks. The key consequence is poor measurement repeatability.

To complicate matters, servers and storage devices are designed to conform to protocol specifications. This makes it difficult to inject errors, create protocol violations and, in general, test device or system reliability under such conditions. When errors do occur in real devices, they provide only high-level logs and lack the capabilities needed to drill down and find the root cause. Worse still, it is almost impossible to recreate any of these error conditions on demand.

Device performance is another factor. As storage networks move to higher speeds, such as 4.25 and 8.5 Gb/s, devices may not be able to load the network at 100 percent capacity, making maximum-performance network device measurements even more difficult.

To address the limitations of server and storage equipment as test devices, several innovations have been added to recent test solutions. Some of these address the three key challenges highlighted earlier: measurement reliability, system scalability and cost effectiveness.

Addressing the measurement challenge

To fully test a SAN device or network, it is important to understand its behavior under a variety of conditions. This knowledge gives the designer added confidence in the device and its performance, whatever conditions may occur in the field. To perform such tests, designers must have a quick way to set up a wide range of possible conditions.

This need influenced the design of the Agilent SAN Tester. While real devices are hindered by operating systems and drivers, dedicated test hardware and software on the SAN tester makes it possible to configure various behaviors. What's more, the SAN tester's dedicated hardware is capable of testing at the full line rate of 8.5 Gb/s, enabling validation of SAN network devices at maximum performance.

Real-time statistics included in the SAN tester also enhance understanding and simplify interpretation of measurement results (Figure 1). Unlike previous approaches that required multiple devices and manual correlation, the SAN tester collects measured results in one place and presents them in a clear, tabular format. For example, error statistics collected during an overnight test run can highlight data corruption issues, or a latency graph can point to types of traffic that cause major lags in system performance.

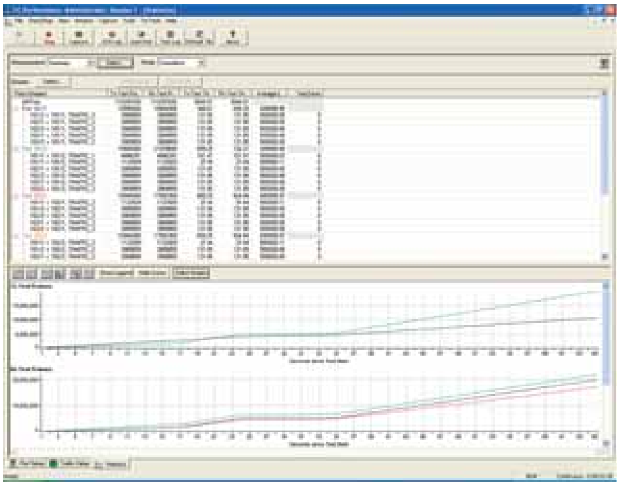

Figure 1. The ability to view real-time statistics in tabular and chart formats enables rapid understanding of device behavior.

In most cases, test systems easily can create operating scenarios that are expected by the device (positive testing). Because the SAN tester has been purpose designed, however, it also can create unexpected operating scenarios (negative testing). These negative tests are essential to understanding the behavior of a device under stressful or unexpected conditions. For example, a switch in a network might be connected to a faulty server that is rebooting every few minutes. A test engineer would want to understand how the switch behaves in such unexpected conditions. By testing for this scenario, the engineer can determine if the switch might cause a catastrophic failure across an entire network when deployed in the field. This is one of several types of negative testing the SAN tester is designed to address.

Protocol analyzers are another important tool in SAN testing. Detecting failure symptoms and retrieving the root cause of functional problems becomes more challenging as the number of ports in a SAN increases. Having access to traffic history before a failure is essential; a protocol analyzer can provide better visibility into the problem and a clear picture of ongoing communication and protocol exchanges. The protocol analyzer should provide tools and features that enable clear visualization of the system being debugged, and different analyzers use different approaches. One example is the Agilent Fibre Channel protocol analyzer, which includes advanced graphical user interface (GUI) capabilities that highlight the information flows moving through a multi-node system (Figure 2). To help the user quickly isolate problems, each type of frame is color coded, eliminating the need to stop and read the text in each one.

Addressing the scalability challenge

When using hand-built test systems based on real servers and storage devices, scaling up to validate a configuration that has several hundred ports might require dozens of large racks filled with equipment. It also would require a lot time to maintain all of the equipment and ensure proper operation, as well as a lot of electricity to power everything.

Dedicated SAN-test platforms replace those racks of real hardware — and reduce the physical footprint — by emulating the behavior of servers and storage devices (Figure 3). Each port of the test device can be configured to mimic multiple types of active devices. As one example, each physical port of the Agilent 173x Fibre Channel SAN test system can mimic 126 devices. Creating a test environment of 1000 devices requires only eight tester ports, which would occupy just 2U of vertical rack space.

In terms of control and coordination, a dedicated test system offers additional advantages. To control real devices, the test engineer typically needs to log in to each device individually (or can write scripts to automate the process). In contrast, a dedicated test system typically includes a GUI that provides a single point of control over hundreds or thousands of emulated devices.

Addressing the cost challenge

In a recent survey conducted by Agilent, a majority of respondents identified “budget” as the most important factor when assessing testing and test needs. This was not a surprise: while SANs — and their test needs — are growing, budgets are being held constant. For test teams and data-center staff, this means finding creative ways to do more with less.

One solution is dedicated test systems. On a per-port basis, test tools might be more expensive than real devices; however, because the test device can emulate hundreds of devices on a single physical port, the total cost of test will be lower.

Multifunction test equipment provides another way to reduce the cost of testing. Historically, SAN testers and protocol analyzers have been separate devices manufactured by different vendors. As a result, organizations bought at least two instruments and two applications to cover their test and debug needs. This model was costly because it needed to account for peak usage of individual instruments and the ports

Figure 2. Advanced GUI capabilities highlight information flows in multi-node systems and helps the user quickly isolate problems.

Figure 3. A large-scale system has a much smaller physical footprint than the numerous real-world server and storage devices it can emulate.

they provided. For example, early in development or functional validation, multiple SAN testers may be needed to generate sufficient traffic to understand the behavior of SAN equipment during both positive and negative testing. Later, during the testing of more-mature devices, not as many SAN tester ports are required for stimuli; instead, more analyzers are required for monitoring and troubleshooting of errors.

Multifunction test modules can address these issues. For example, the Agilent Fibre Channel SAN test system lets the user switch between two different applications, SAN tester and protocol analyzer (Figure 4). When the organization needs more stimuli, most of the hardware can be used in the SAN-tester mode. When the organization needs more protocol analyzers, most of the hardware can be used in the analysis mode. Having this level of flexibility in one tester helps reduce the total cost of test.

Conclusion

Storage area networks are growing in size, complexity, features and capabilities. These changes are driving the adoption of dedicated test solutions, which offer many advantages for testing today's SANs and SAN devices: a single point of control over thousands of emulated devices, repeatable measurements and a lower cost of test. Dedicated testers that include multiple applications in one instrument — a SAN tester and protocol analyzer, for example — give the user greater flexibility and more control across the complete development cycle of new SAN products.

Additional reading

For more information about test solutions for Fibre Channel, please visit www.agilent.com/find/8G.

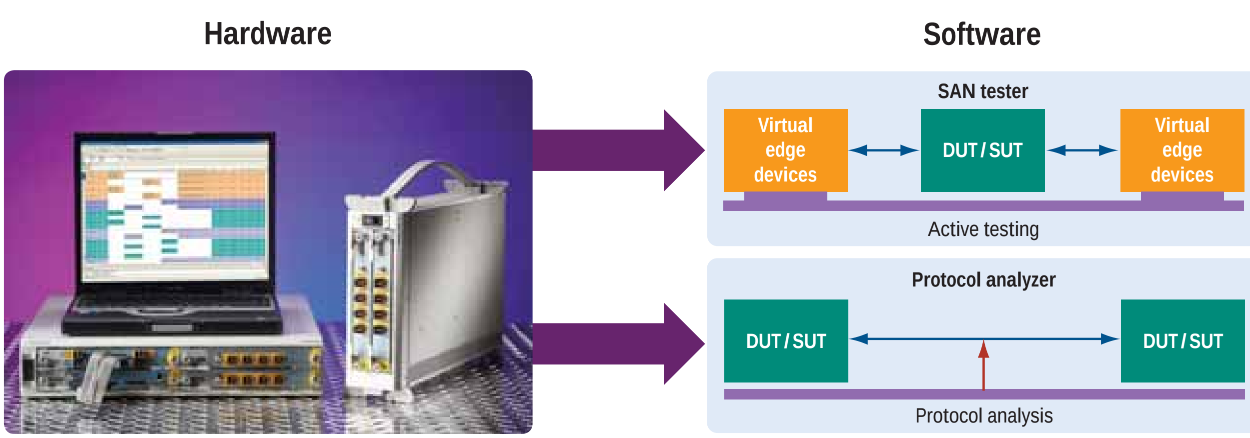

Figure 4. The Agilent 173x Fibre Channel SAN test system provides the capabilities of a SAN tester and protocol analyzer in one hardware platform.

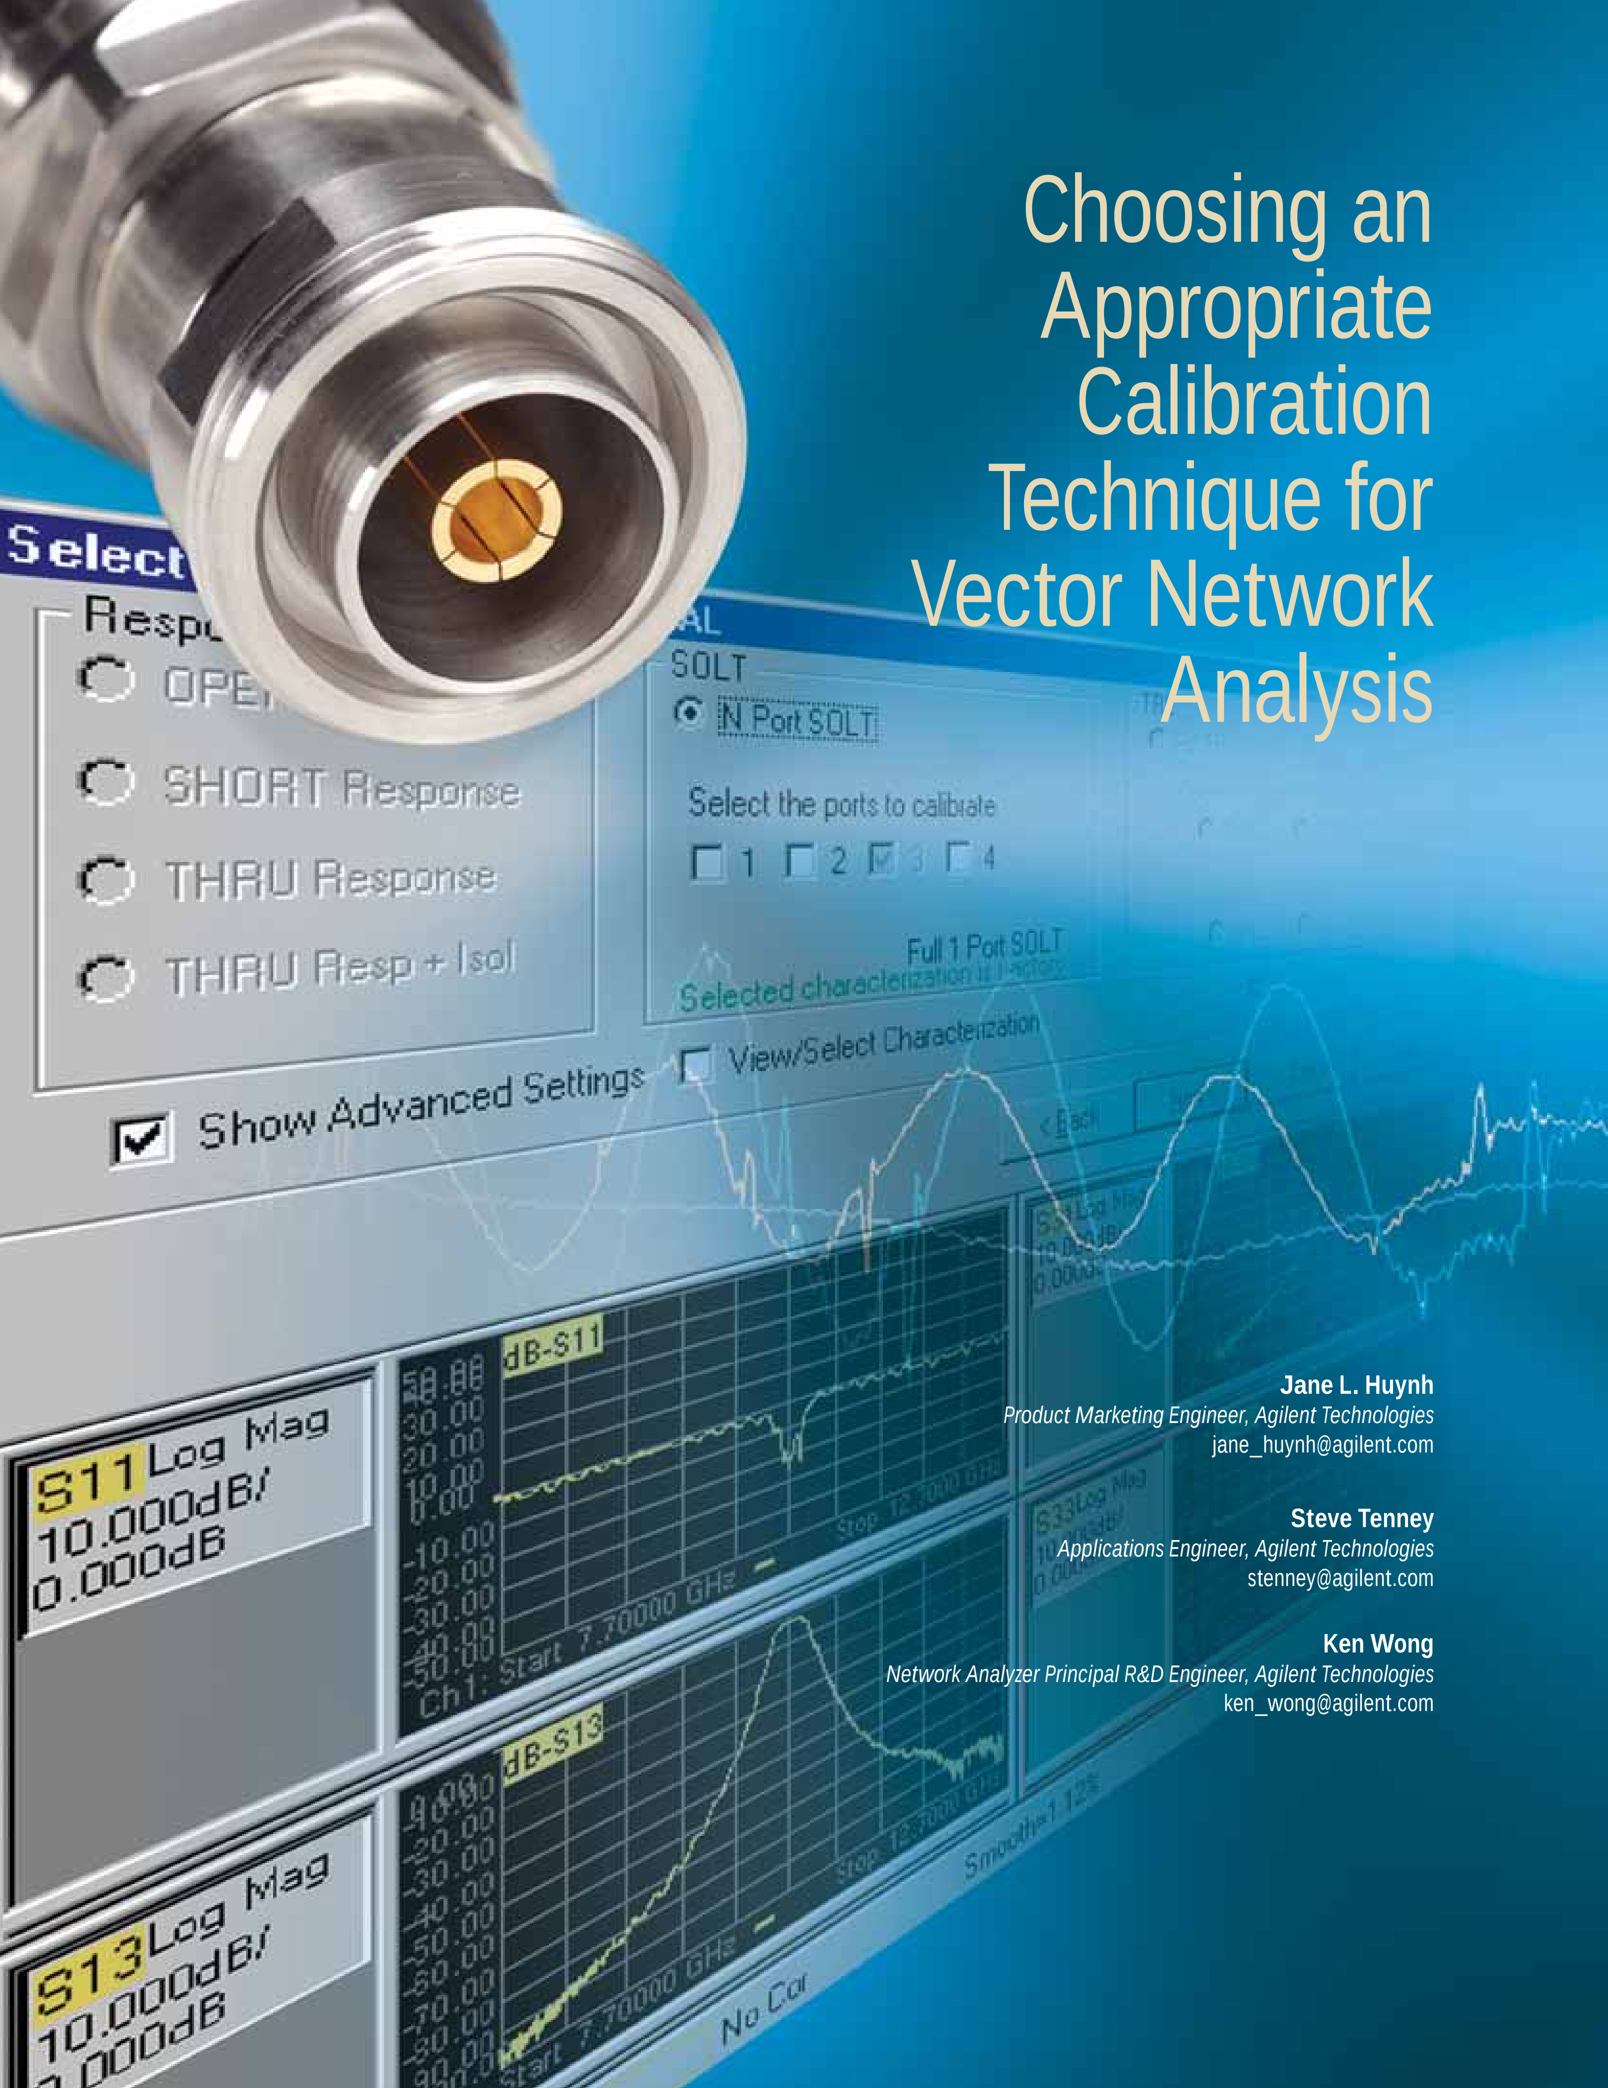

Choosing an Appropriate Calibration Technique for Vector Network Analysis

Jane L. Huynh
Product Marketing Engineer, Agilent Technologies
jane_huynh@agilent.com

Steve Tenney
Applications Engineer, Agilent Technologies
stenney@agilent.com

Ken Wong
Network Analyzer Principal R&D Engineer, Agilent Technologies
ken_wong@agilent.com

In network analysis, “It depends” is often the accurate — but sometimes unsatisfying — answer to the question, “Which calibration technique is most accurate?” The “best” answer depends on factors such as the application, the required accuracy and the care taken during the calibration process itself.

Compared to past-generation vector network analyzers (VNAs), today’s instruments offer more choices in calibration methods. While greater choice is beneficial, it also can lead to greater confusion. Fortunately, a few key points of comparison can quickly narrow the field and identify the most appropriate calibration technique.

This article discusses common network analyzer calibration techniques and their relative accuracies. The focus is on sound measurement practices and other factors that can enhance accuracy.

Types of calibration

With the significant processing power and flexibility of today’s network analyzers, the availability of targeted application-specific calibration methods is growing. Examples of application-specific calibration types include mixer/converter calibration (for frequency-offset devices), noise figure calibration and in-fixture measurements.

Full one- and two-port vector calibrations are our focus. As such, vector calibration methods targeted at all sources of error in the network analyzer will be reviewed. These methods provide greater accuracy than approaches such as response calibration that do not consider all error terms.

For a discussion of calibration accuracy, we can limit the discussion to a few common calibration types from which most modern methods are derived. In general, there are three commonly used families of calibration techniques: SOLT (short-open-load-thru), TRL (thru-reflect-line) and ECal (electronic calibration) modules. Within each, there are implementations usually targeted at specific measurement requirements such as broadband frequencies or on-wafer probing. These commonly used calibration techniques are summarized in Table 1 along with the key benefits of each.

Understanding systematic errors in network analyzers

Figure 1 summarizes the sources of systematic errors within a typical network analyzer. The ability to measure phase enables VNAs to accurately account for all sources of error. Directivity error affects the accuracy of reflection measurements. Isolation error affects the accuracy of transmission measurements. Source and load errors relate to the mismatch between the device-under-test (DUT) and the impedance of the analyzer’s measurement ports. Reflection and transmission tracking relate to the difference in frequency response of the analyzer’s reference and measurement receivers.

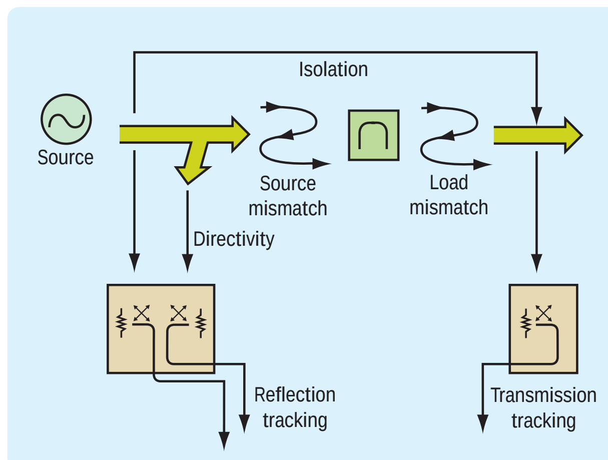

Figure 1. Six systematic errors are present in the forward direction of a network analyzer.

Exploring SOLT calibration

The first calibration method most network analyzer users learn is SOLT. When properly implemented, SOLT calibration can provide excellent accuracy and repeatability.

This calibration method requires short, open and load calibration standards. If sexed connectors (e.g., male or female) are on the DUT, then one of each standard is required for male and female connections. Connecting the two measurement planes together makes the thru connection.

Table 1. The benefits of each calibration technique can help determine which one is best for a particular application.

Calibration techniques	Description	Key benefits
SOLT family		
SOLT fixed load (also known as two one-port + thru)	Short, open, load, thru using a fixed impedance (usually 50 or 75 Ω) broadband load.	Widely used and understood. Applicable in nearly all cases. Requires well-defined standards.
SOLT sliding load	Short, open, load, thru using a fixed impedance (usually 50 or 75 Ω) sliding broadband load.	Widely used and understood. Applicable in nearly all cases. Requires well-defined standards.Higher accuracy due to better characterization of coupler directivity. More accurate at higher frequencies than a fixed load.
SOLT offset[1]	Offset short, open, load, with thru	Ideal for specific applications such as wave-guide calibration.
SOLR	Short, open, load, reciprocal-thru (also called “unknown thru”)	Ideal for setups such as: • noninsertable • multiport • physically long devices where cables should not be moved after calibration
QSOLT	Quick SOLT	Ideal for multiport setups in which a minimum set of connections and disconnections are selected to optimize calibration. The more ports, the greater the time savings.
TRL family[2]		
TRL	Thru, reflect, line	Good choices for calibrating inside fixtures or on-wafer probing. Steps vary depending on available standards (e.g., TRL requires one line with finite length; LRL requires two lines, each with a different finite length).
TRM	Thru, reflect, match	
LRL	Line, reflect, line	
LRM	Line, reflect, match	
LRM+	Line, reflect, offset match (called “LRM-plus”)	
ECal Modules		
ECal: Available as two- or four-port modules in various frequency ranges	Uses predefined impedance states compared to measured values. USB controlled.	Fast and repeatable. Ideal choice for eliminating human errors during calibration. Can be as accurate as SOLT sliding load manual calibration.

1. Different types of offset calibration standards are available and may or may not be covered here. For more information, refer to Reference 2.

2. TRL is not applicable for one-port calibration.

The SOLT calibration method utilizes a 12-term error-correction model. The 12 terms refer to six error terms each in the forward and reverse direction through the DUT. Figure 2 shows the forward-direction error terms: E_D (directivity), E_S (source match), E_L (load match), E_{RF} (reflection tracking), E_{TF} (transmission tracking) and E_X (crosstalk). When performed properly, this method makes it possible to measure hundredths of a decibel of magnitude and millidegrees of phase.

SOLT calibration standards are contained in widely available calibration kits that include a variety of connector types. They are relatively inexpensive and can last for years with proper care.

Some SOLT calibration kits contain sliding loads, which make it possible to vary the line length of the path while maintaining a constant load impedance (usually 50 or 75 Ω). This is especially important at higher frequencies where good fixed loads are difficult to implement. The change in line length is directly proportional to a change in electrical length, which causes a phase shift in the measurement path. By using several different lengths with corresponding phase shifts during calibration, directivity of the network analyzer can more accurately measured (Figure 3).

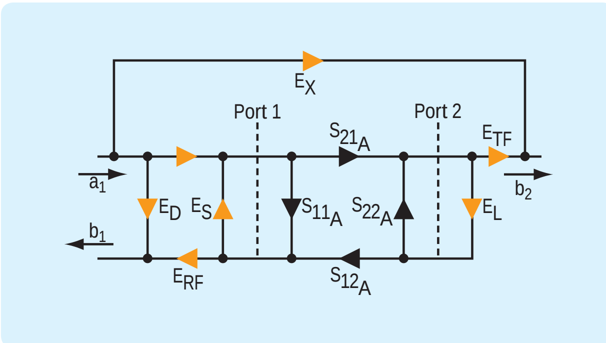

Figure 2. The systematic errors shown in Figure 1 are described here by the 12-term error model flow graph. There is one flow graph for the forward direction (six terms shown here) and another for the reverse direction (six more terms).

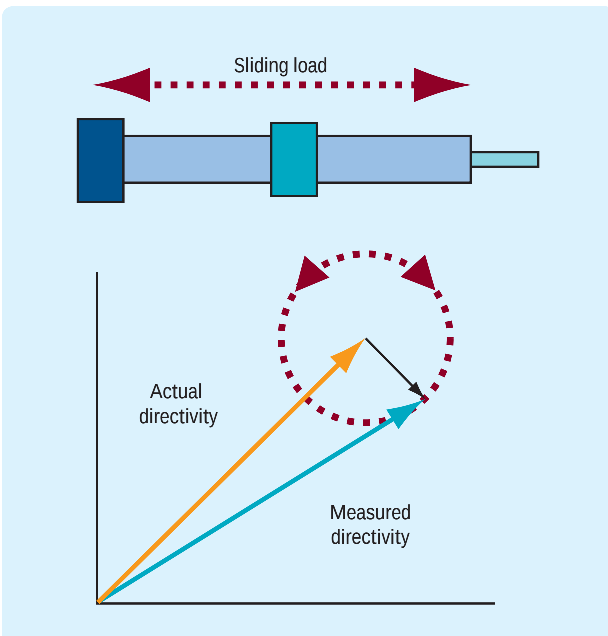

Figure 3. Using a sliding load can improve calibration accuracy of directivity measurements at high frequencies.

SOL with reciprocal thru is better known as "unknown thru." This method allows the use of devices such as cables, circuit board traces or ECal modules as the thru during calibration, as long as a few basic guidelines are followed.[1] This method is especially useful when dealing with noninsertable devices — those with same-sex or incompatible connectors that require an adapter to complete the thru connection during calibration. This adapter represents an error in the calibration. Unknown thru is useful for eliminating the need for a precision or calibrated adapter and minimizing cable movement during calibration. It is generally more convenient and accurate than other adapter-removal techniques.

Other calibration techniques based on SOLT include offsetting one of the calibration standards. This "offset SOLT" method is effective for waveguide and other high frequency applications. For instance, an offset load can be considered a compound standard consisting of two known offsets (transmission lines) of different lengths and a load element.[2]

"Quick SOLT" or QSOLT is used in multiport applications where the number of measurement ports on the network analyzer is greater than two. This is sometimes referred to as an N-port solution, where N represents the number of ports. The number of calibration steps is proportional to the number of ports. QSOLT uses mathematical algorithms requiring fewer calibration steps to solve for the full N-port error model with a minimum set of connections.

Most SOLT calibrations are performed manually and are relatively easy to implement. Agilent network analyzers provide guided (step-by-step) calibration, which reduces human error and improves repeatability. As with most measurements, however, proper calibration techniques must be practiced to ensure the maximum performance SOLT calibration is capable of providing.

Understanding TRL calibration

TRL calibration is extremely accurate — in most cases more accurate than SOLT calibration. However, few calibration kits contain TRL standards. TRL calibration is most often performed in situations requiring a high level of accuracy when the available calibration standards do not share the same connection type as the DUT. This usually is the case when using test fixtures or making on-wafer measurements with probes. Therefore, in some cases it becomes necessary to construct and characterize standards in the same media type as the DUT configuration. It is easier to manufacture and characterize three TRL standards than the four SOLT standards.

TRL calibration has another important advantage: The standards need not be defined as completely or accurately as SOLT standards. While SOLT standards are completely characterized and stored as the standard definition, TRL standards are modeled rather than being completely characterized. However, TRL calibration accuracy is directly proportional to the quality and repeatability of the TRL standards. Physical discontinuities such as bends in the transmission lines and beads in coaxial structures will degrade TRL calibration. The interface must be clean and allow repeatable connections.

Utilizing ECal modules

To ensure temperature stability, ECal modules use a set of solid-state impedance standards on a heated plate. During network analyzer calibration, these solid-state impedance standards are measured with correction off. These raw (uncorrected) measurements are then compared with the expected values of the impedance standards stored on flash memory inside the ECal module. The network analyzer reads the impedance state values and compares them to the measured values. The difference is used to calculate the calibration factors (or error terms).

ECal is an excellent option for minimizing human errors during calibration. To use ECal, the module is connected to the network analyzer via USB. Once the ECal module warms up, it is connected to the network analyzer test ports and ECal is selected in the calibration menu. The module automatically senses the port connections and begins its calibration process. The process, which typically takes less than 30 seconds, is highly repeatable and, when properly performed, provides accuracy that rivals many manual calibration techniques.[3]

Unlike other methods, ECal modules are flexible and can be recharacterized with different connectors by performing the user characterization function available on some VNAs. ECal modules are available with coaxial connectors. After connecting the coaxial-to-waveguide adapters, a user characterization can be performed, enabling the ECal module to behave as a waveguide module (Figure 4).[4]

Figure 4. After adding waveguide adapters, an ECal module can be recharacterized and used as a waveguide module.

Calibration techniques and tradeoffs

When deciding which calibration technique to apply, users often consider accuracy versus ease-of-use. Ideally, it is most useful to apply the method that requires the highest accuracy and the lowest possible skill level. Unfortunately, there is a tradeoff between these two factors.

Table 2 leverages the results of Agilent research and summarizes the calibration techniques discussed here, comparing required skill level, repeatability, cost and accuracy. The different categories are rated as low (L), medium (M) and high (H). The derivation of these accuracies is beyond the scope of this article; however, details are available from the references listed at the end.

In general, the SOLT values provided in Table 2 are for two-port calibrations because this is the most common application. In contrast, one-port SOL calibrations do not require the thru so will typically be slightly more accurate. Subsequently, any uncertainty due to the thru standard would not be included as part of the calibration. Only the data for two-port TRL and its derivatives are provided because TRL calibrations are not applicable to one-port measurements.

Examining accuracy considerations

Accuracy is the ability of an instrument to measure the actual value within a stated error specification. A typical network analyzer is capable of measuring magnitude within hundredths of a decibel and phase within a few millidegrees.

This level of accuracy is obtainable only if proper calibration and measurement practices are followed. In other words, the question, "Which calibration technique is the most accurate?" is meaningless if poor calibration or measurement techniques are employed. Indeed, proper connector care, torque, instrument operation and other best practices can have a greater impact on the measurement accuracy of a network analyzer than the selected calibration method. Poor calibration practices can lead to inaccuracies much worse than those shown in Table 2.

Network analyzers, broadly speaking, have no inherent guaranteed accuracy. These instruments rely on the measurement of known calibration standards during the calibration process

Table 2. This comparison of calibration types, tradeoffs and accuracies can help identify which one is best for a particular application. The ratings are low (L), medium (M) and high (H).

Calibration techniques	Tradeoffs							
	One-port calibration				Two-port calibration			
	Cost	Repeatability	Skill level	Accuracy	Cost	Repeatability	Skill level	Accuracy
SOLT family[1]								
SOLT fixed load	L	H	L	L-M (10-1%)	L-H	L-H	L-H	L-H (10-0.5%)
SOLT sliding load	M-H	M	M	M-H (2-0.3%)	M-H	M	M	M-H (5-0.3%)
SOLT offset[2]	M-H	L-H	L-H	M-H (1-0.2%)	M-H	L-H	L-H	M-H (3-0.05%)
SOLR (unknown thru)	*Not applicable*				L-M	M-H	L-M	L-H[3] (10-0.5%)
QSOLT					L-H	M-H	M-H	L-H[4]
TRL family[5]								
TRL	*Not applicable*				H	M	H	H+ (<1%)
TRM					L-M	M-H	L-M	L-M (10-2%)
LRL					H+	L-M	H+	L-M (<3%)
LRM					H	L-M	H	L-M (12-2%)
ECal modules								
ECal: Available as two- or four-port modules	M	H	L	L-H (10-0.2%)	M	H	L	L-H (5-0.3%)

1. The thru (T) standard is applicable only with two-port calibration.
2. Different types of offset calibration standards are available and may or may not be covered here. For more information, refer to Reference 2.
3. Because SOLR eliminates the need for adapters during calibration (along with their associated uncertainties), it may exceed the accuracy of SOLT.
4. Accuracy of QSOLT depends on the quality of the calibration and the accuracy of the reduction algorithm applied to the calibrated data.
5. TRL is not applicable to one-port calibration.

to establish a baseline. A network analyzer's deviation from a standard during calibration defines its accuracy. Put another way, any error introduced during calibration will directly affect the measurement accuracy of the network analyzer. It is essential to perform a proper calibration to ensure optimal measurement performance.

The following key topics must be considered when using a network analyzer to ensure the best overall accuracy. To begin, recognize that the calibration of a network analyzer is a measurement: It requires the care and precision of any RF/microwave measurement.

Considering mismatch measurement uncertainties

Consider a circuit trace used to connect a network analyzer to a DUT with a 10 percent error in its characteristic impedance (e.g., 55 Ω with a 50-Ω network analyzer). A simple calculation yields a reflection coefficient (ρ) for the source and load of 0.05, equal to a VSWR of 1:1.10.[5] This results in a mismatch loss of 0.01 dB and mismatch uncertainty of ±0.02 dB.

Now consider a circuit trace between the network analyzer and DUT with a reflection coefficient of 0.13, equal to a VSWR of 1:1.30.[5] This is well below a VSWR of 1:2.0, which is often specified in some commercial applications. In this case, the mismatch loss jumps to 0.075 dB, and mismatch uncertainty of ± 0.15 dB. Figure 5 illustrates this relationship.

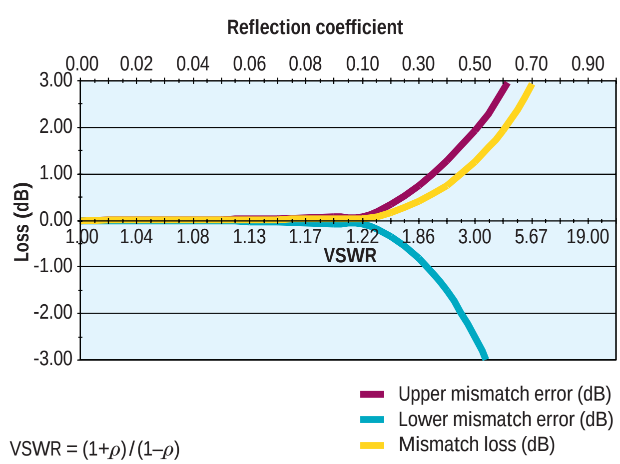

VSWR = $(1+\rho)/(1-\rho)$

Maximum mismatch error limits (dB) = $20\log_{10}(1\pm|\rho_{source}\rho_{load}|)$

Mismatch loss (dB) = $-10\log_{10}(1-\rho^2)$

where ρ = reflection coeficient and $\rho_{source} = \rho_{load}$

Figure 5. This graph shows the relationship between reflection coefficient, mismatch loss and uncertainty.

A connection between the network analyzer and the DUT may include cables, adapters, circuit traces and fixtures. The variances in the characteristic impedance of each component can be caused by poor quality components and cables, dirty or damaged connectors and improper torque, to mention a few. These have a cumulative effect that will produce a significant level of measurement uncertainty.

These examples demonstrate the importance of using good measurement practices, including proper connector care, torque, cleaning, and minimizing adapters, as well as using quality cables and components during calibration and device measurements.

Minimizing noise during calibration

In addition to using good measurement practices during calibration, the impact of the environment and noise should be minimized. Doing so will improve the repeatability, accuracy and stability of the measurement results.

When radiated or conducted electrical noise is present, filtering and shielding should be employed to reduce the noise. Certain instrument settings can further reduce the effect of noise on the measurement:

- Set the stimulus power of the network analyzer as high as practical to maximize the signal-to-noise ratio of the DUT. Balance this with the input requirements of the DUT and other measurement considerations.
- Reduce the IF bandwidth. The tradeoff is slower sweep speed. If required, use averaging to further improve signal-to-noise ratio by reducing the effect of random white noise. Averaging can be turned off after calibration but the measurement noise will increase.
- To reduce the risk of procedural errors, a test run of the instrument settings and DUT connections is strongly recommended prior to performing a calibration. Proper calibration practices will ensure minimal errors.

Outlining best measurement practices

When using network analyzers, optimize measurement accuracy and repeatability by following the steps listed below. It is important to test run the setup before calibrating. Ensuring proper connection to the DUT and instrument settings will improve accuracy. Any changes made to the measurement or setup after calibration can decrease measurement accuracy.

1. Warm up the network analyzer
2. Clean, inspect and gauge all connectors and cables
3. Connect cables and adapters to the analyzer
4. Preset the analyzer to a known condition
5. Set up stimulus frequencies, number of points and power
6. Set velocity factor, IF bandwidth and averaging (if required)
7. Connect the DUT to verify setup, cables, adapters and operation
8. Select S-parameter(s) to be measured
9. If using special functions, select the proper setup as required
10. Select the display format
11. Scale the display
12. Remove the DUT
13. Select the proper calibration kit or calibration standards definitions
14. Calibrate
15. Confirm the quality of the calibration
16. Reconnect the DUT
17. Save the instrument state
18. If required, turn off averaging, select time domain, port extensions or other modes

The most important steps happen *before* calibration. For example, it is essential to ensure proper instrument settings, DUT connections and measurement setup prior to performing a calibration. Equally important is understanding the correct usage of the calibration kit and standards for a given application. Some network analyzers will turn off correction if parameters are changed after calibration. All network analyzers will have reduced measurement accuracy if the calibration is performed improperly.

Conclusion

"Which calibration technique is the most accurate?" is not easy to answer. First, one must consider the application as well as the required measurement accuracy and other factors.

Most network analyzer calibrations fall into three broad categories: SOLT, TRL and ECal. Each has its own strengths and associated accuracies. Proper training and experience often will result in developing a familiarity over time with a particular method; however, engineers should also exercise caution to prevent forming bad habits during calibration.

Ultimately, precision and accuracy can be (and often are) laid to waste by poor calibration and measurement practices. The average user may not have the opportunity to use a network analyzer frequently. When using an RF instrument capable of just one or two decibels of accuracy, a dirty connector that causes 0.1 dB of error will seem insignificant. When using a network analyzer, however, 0.1 dB is well within the measurement range of the instrument. As a result, it is useful to develop and review good measurement practices to avoid adversely impacting calibration accuracy.

References

1. Please see www.agilent.com/find/na for an overview of these guidelines.

2. Agilent Application Note 1287-11: *Specifying Calibration Standards and Kits for Agilent Vector Network Analyzers*. Literature number 5989-4840EN available from www.agilent.com.

3. Agilent White Paper: *Calibration – Electronic vs. Mechanical Calibration Kits: Calibration Methods and Accuracy*. Literature number 5988-9477EN available from www.agilent.com.

4. Agilent White Paper: *User Characterization: Electronic Calibration Feature Allows Users to Customize to Specific Needs.* Literature number 5988-9478EN available from www.agilent.com.

5. Agilent Application Note 339-10: *Input-output Impedance and Reflection Coefficient Measurements*. Literature number 5950-2925 available from www.agilent.com.

Additional Reading

- Agilent Application Note 1287-3: *Applying Error Correction to Network Analyzer Measurements*. Literature number 5965-7709E available from www.agilent.com.

- Agilent Application Note: *On-Wafer SOLT Calibration Using 4-port PNA-L Network Analyzers (N5230A Options x4x)*. Literature number 5989-2287EN available from www.agilent.com.

Making Traceable EVM Measurements with Digital Oscilloscopes

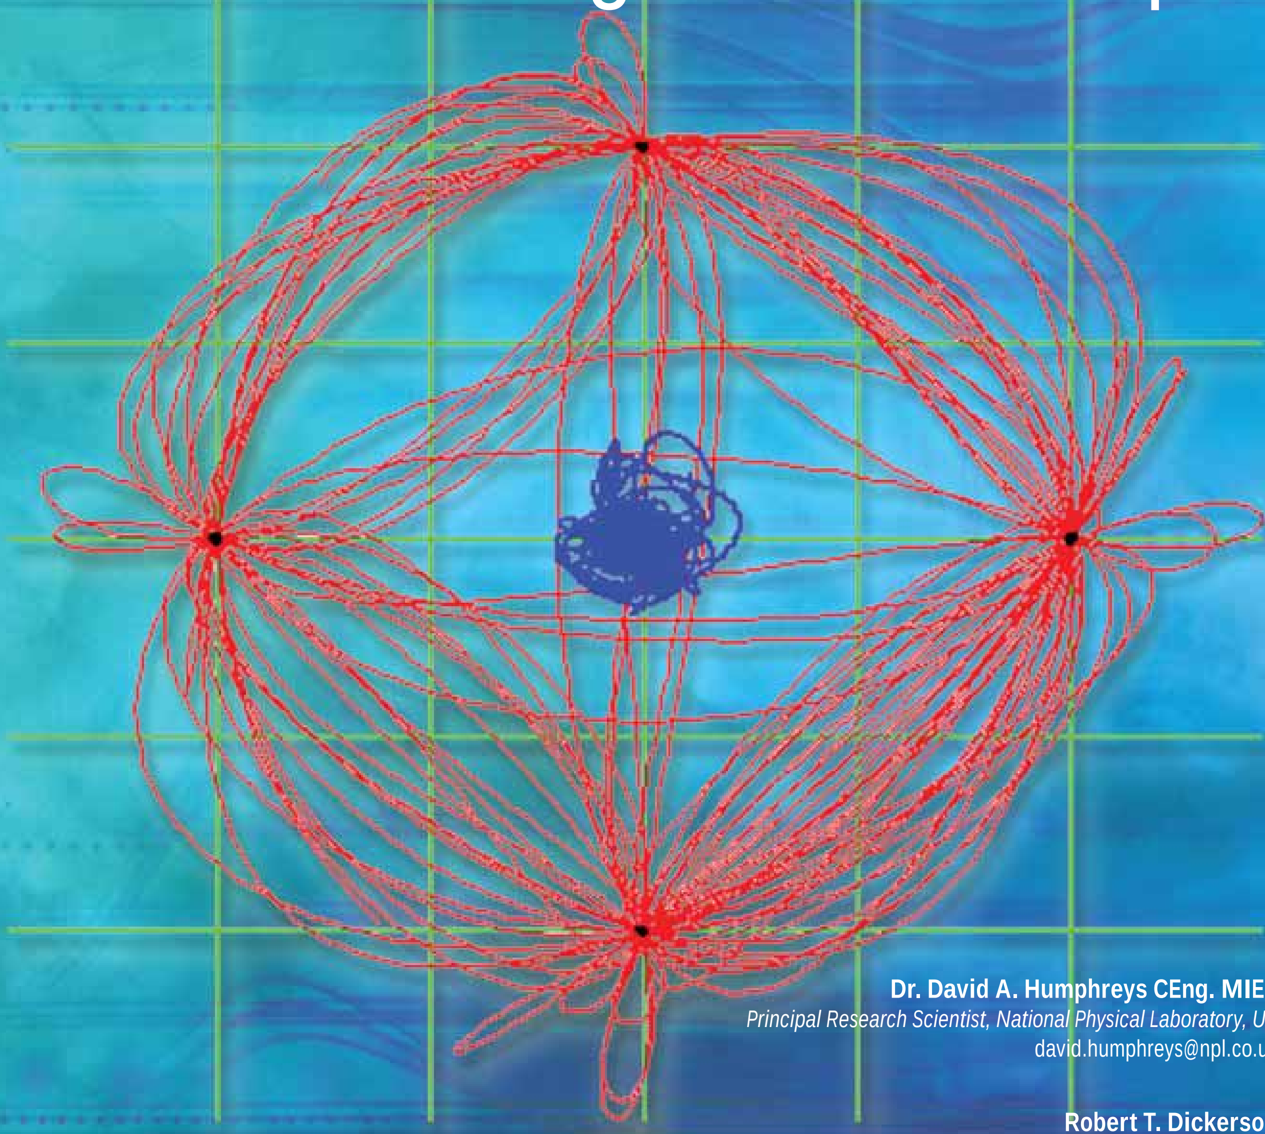

Dr. David A. Humphreys CEng. MIET
Principal Research Scientist, National Physical Laboratory, UK
david.humphreys@npl.co.uk

Robert T. Dickerson
Measurement Engineer, Agilent Technologies
rob.dickerson@agilent.com

Ann Yanecek
New Product Test System Engineer, Agilent Technologies
ann_yanecek@agilent.com

Wireless communication systems use scarce radio spectrum efficiently by applying sophisticated mathematical coding algorithms, agreed by specification standards bodies such as the Third-Generation Partnership Project (3GPP) and implemented in silicon by device and equipment manufacturers. Component and system performance is specified using parameters such as error vector magnitude (EVM), which quantifies the relative root mean square (RMS) error in the signal.

To ensure acceptable operation of wireless systems, EVM is measured with dedicated test instrumentation during design, manufacturing, installation and maintenance. Parametric measures are simple and useful, and they can answer yes/no questions — but can't explain why a problem exists. What's more, finding the root problem is made more difficult if two measurement instruments disagree.

RF waveform metrology is important in aerospace/defense, instrumentation and telecommunications applications because it helps answer why. Modeling and simulation provide accurate predictions of what should happen but eventually real measurements of actual devices are required. As such, waveform metrology is valuable for developing new designs, diagnosing problems and verifying implementations. Traceability of waveform measurements provides a common reference and supports the ISO 17025 standard for testing and calibration.

This article provides an overview of joint research conducted by Agilent and the National Physical Laboratory (NPL) of the United Kingdom. The NPL's industry-focused "waveform metrology for wireless communications" project is scheduled to run until 2009.[1] Its aim is to provide traceability for W-CDMA and other formats such as WiMAX™ and ultra-wideband (UWB).

Overview of W-CDMA

3GPP W-CDMA was leveraged from the W-CDMA system developed in Japan and Europe. The 3GPP version is viewed as the next-generation replacement for PDC (Japan) and GSM (worldwide) and is designed to have a higher data rate than 2G systems — up to 2 Mbps for W-CDMA versus 14.4 kbps for GSM. W-CDMA allows several users to efficiently share the same RF carrier by assigning unique codes to each user and dynamically adjusting data rates and link budgets to balance the cumulative demand from all active users.

Figure 1 shows the mapping of W-CDMA logical channels (dedicated to control and traffic) to the W-CDMA physical channel (dedicated to physical data/control) of an uplink signal. 3GPP defines a logical channel as an information stream dedicated to the transfer of a specific type of information over the radio interface. It also defines a physical channel having specific carrier frequency, scrambling code, channelization code and relative phase.[2]

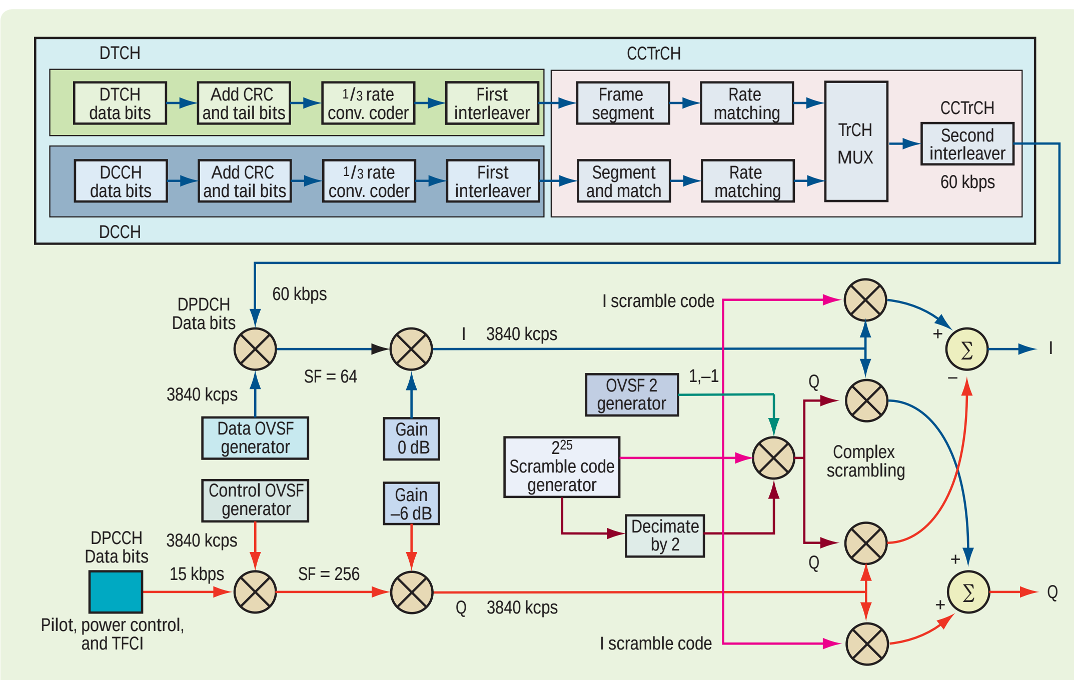

Figure 1. Uplink data-channel air interface

In Figure 1, voice data is carried on the logical dedicated traffic channel (DTCH; green box), and signaling data is carried on the dedicated control channel (DCCH; blue box). Each logical channel is channel coded and interleaved, then segmented to conform to the physical layer's 10 ms frame structure and rate-adjusted to match the physical layer data-block size. The traffic and control channels are then multiplexed together to form the coded composite transport channel (CCTrCH; pink box). After a second interleaving, this transport channel is mapped onto the physical data channel (DPDCH; blue path), which is then spread using orthogonal variable-spreading-factor (OVSF) codes to attain the desired 3.84 Mbps rate. The pilot, power control and other control data are mapped onto the physical control channel (DPCCH; red path), which is also spread to 3.84 Mbps and scaled to be –6 dB relative to the DPDCH.

The composite spread signal, containing data on the in-phase (I) path (blue) and control information on the quadrature (Q) path (red), is scrambled using a complex function called hybrid phase-shift keying (HPSK). HPSK scrambling offers an important benefit: It reduces the zero-crossing transitions in the IQ plane. This reduces the peak-to-average power ratio of the signal and ultimately simplifies mobile design.

Not shown in Figure 1 is the final physical channel in which the data is passed through a root-raised-cosine (RRC) filter and an IQ modulator. The filtered modulated signal is then applied to the RF carrier.

The downlink is created in a similar manner with two exceptions: the DPDCH and DPCCH are time-multiplexed and a different channelization process is used. Please see Reference 3 for a detailed look at W-CDMA downlink signals. For more information on W-CDMA uplink signals and HPSK, please see Reference 2.

Overview of EVM

It takes a wide array of measurements to characterize a transmitter — and each measurement provides a different insight into transmitter performance. As an example, a signal's constellation can be used to evaluate the accuracy of an RF waveform modulated by a complex waveform. A symbol is the smallest data unit being transmitted and each one has a known location within the constellation of a particular data pattern. From this, a reference signal can be created to compare the expected and actual constellations. This comparison — the EVM — is a critical measure of W-CDMA link performance.

Paraphrasing, the 3GPP specifications for terminal conformance define EVM as follows: The error vector magnitude measures the difference — the error vector — between reference and measured waveforms. Both waveforms pass through a matched RRC filter with 3.84-MHz bandwidth and $\alpha = 0.22$ rolloff. Both waveforms are then modified by selecting frequency, absolute phase, absolute amplitude and chip-clock timing values that minimize the error vector. The final EVM result is the square root of the ratio of the mean error-vector power to the mean reference power expressed as a percentage (Figure 2).[4]

The error vector shown in the figure represents the quantity measured for each chip. The individual chip-error vectors are then combined as defined in the EVM equation to yield the EVM measurement result.

$$\mathrm{EVM} = \sqrt{\frac{\sum_{\gamma=0}^{N-1} \left| Z'(\gamma) - R'(\gamma) \right|^2}{\sum_{\gamma=0}^{N-1} \left| R'(\gamma) \right|^2}} \times 100\%$$

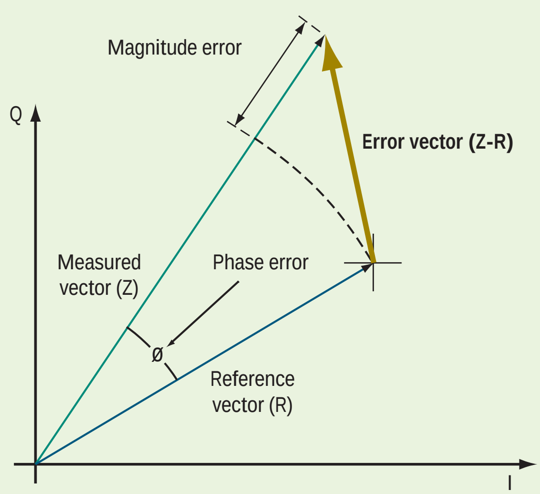

Figure 2. Definition of EVM

For single-channel EVM (QPSK EVM), the baseband I and Q signals are recovered and passed through an RRC filter. The obtained samples are then passed through a QPSK decoder at the chip-timing rate to determine the correct chip location for each sample. An assumption is made about the error of the true chip that allows the sampled chip to be placed in the correct quadrant. Once all of the chips have been decoded, they are QPSK encoded and passed through a raised-cosine filter to create the reference signal. This reference signal is then compared to the measured signal to produce an EVM result.

Waveform metrology and choice of instruments

A W-CDMA signal can be treated as a carrier signal with a complex modulation term, $V(t) = \mathrm{Re}(M(t) \times \exp(j\omega t))$, where $V(t)$ is the time-varying RF voltage at an angular frequency ω and M(t) is the complex W-CDMA modulation term. The measured waveform will be imperfect, containing the modulation waveform and an error waveform $M_{meas}(t) = M(t) + M_{error}(t)$. Combining averaging with a knowledge of the measuring instrument characteristics reduces the instrument noise and makes it possible to recover the error waveform.

Three types of general-purpose instruments seem to be logical choices for waveform metrology: sampling oscilloscopes, digital oscilloscopes and spectrum analyzers. Sampling oscilloscopes can demodulate repetitive RF signals and provide the closest link to the primary standards; however, the low frame-update rate of W-CDMA signals makes this approach impractical.[5] Instruments such as the Agilent PSA E4440 Series performance spectrum analyzers can demodulate complex waveforms and are in widespread industrial use. NPL plans to use these instruments in later stages of the wireless project.

Today, NPL recovers the W-CDMA modulation waveform mathematically from single-acquisition digital oscilloscope traces. These versatile instruments can be used to measure both repetitive and single-shot waveforms. Their weakness, though, when compared with sampling oscilloscopes or spectrum analyzers, is that the voltage-scale resolution is typically eight bits. Digital oscilloscopes also have complex noise and nonlinearity behavior, which restricts the choice of measurement frequency. To compensate for this shortcoming, digital oscilloscopes provide much higher levels of oversampling, typically > 500 points per chip.

Digital-oscilloscope timebase accuracy is typically 1×10^{-6} or better. NPL uses an additional synthesizer, phase-locked to the W-CDMA signal, to provide a pilot tone that compensates for residual timing fluctuation, improves the timebase performance by a factor of 10,000, and keeps additional EVM contributions below 0.1 percent at 2.5 GHz. To avoid the introduction of correlated errors into the result, the signal and pilot frequencies are chosen using a predictive model.

Diagnosing errors with RF waveform metrology

An NPL scientist worked with researchers at the Centre for Communications Research (CCR) at Bristol University, UK in January 2007 (Figure 3).[6, 7] They used waveform metrology to locate a software fault that gave EVM values of eight to 10 percent when measured with commercial test equipment. The cause was a software error that shifted the symbol timing by 87.2 ns (0.335 chips) every 20 symbols (2560 chips; Figure 4). When calculated with the correct timing reference, the source EVM is less than 0.6 percent.

Figure 3. NPL and CCR staff performing RF metrology waveform measurements at Bristol University

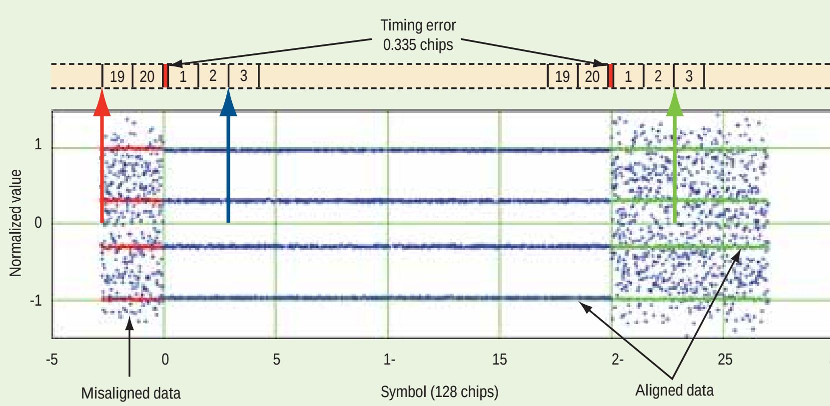

Figure 4. Constellation diagrams showing symbol-timing errors and alignment

The traceability chain

Communications waveforms are dynamic, linking signals at different frequencies. They derive traceability from primary systems that measure time-resolved voltage. The NPL electrical-risetime calibration system has excellent uncertainties for 70 GHz sampling oscilloscopes and pulse generators with transition durations as short as 5.1 ps.[8] NPL maintains confidence in the results by intercomparing measurements with other national measurement institutes such as the National Institute of Standards and Technology (NIST) in the United States and Physikalisch-Technische Bundesanstalt (PTB) in Germany.[9, 10]

In the NPL electro-optic sampling (EOS) system, the electric field on a coplanar line induces measurable rotation $\Delta\theta$ of the polarization of light passing through a 20 µm-thick lithium tantalate probe positioned just above the line (Figure 5).[11] The electrical impulse and optical probe signal are generated using the same 200 fs optical pulse from a modelocked titanium-sapphire (Ti:sapphire) laser. NPL uses a coplanar photoconductor or photodiode to generate pulses down to below 1 ps duration (Figure 6). The relative delay of the two optical pulses is varied to measure V(t).

Communications test equipment requires specific modulated waveforms and cannot interpret the simple impulse waveform used with the primary standard. The strategy for traceability is to use a multi-step process to link the simple waveforms of the primary EOS system to the complex waveforms used by communication systems using generalized instrumentation such as digital oscilloscopes (Figure 7).

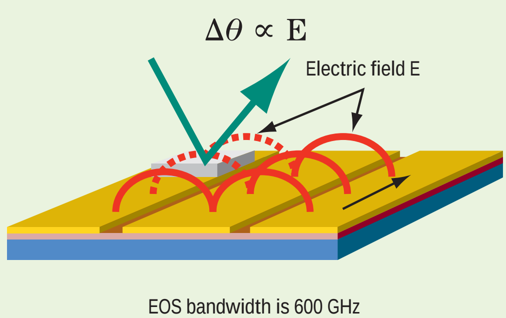

Figure 5. Electro-optic sampling: The electric field on the line rotates the polarization of the light

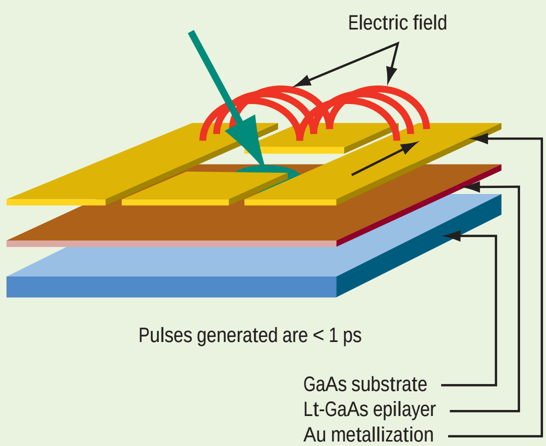

Figure 6. Photoconductive pulse generator

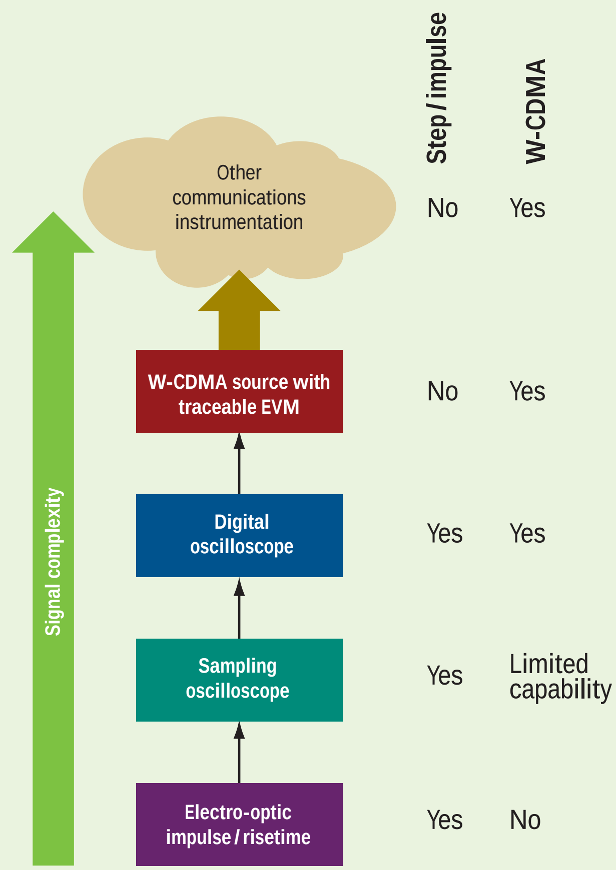

Figure 7. Linking W-CDMA to primary standards

Combining several measurements gives statistical uncertainty values for the modulation waveform. The waveforms must be aligned, correcting the phase rotation and residual frequency offset before they can be averaged. Typical measured results (three pilot tones x 10 acquisitions) for an Agilent N5182A MXG vector signal generator at 900 MHz show RMS deviations from the mean of < 0.25 percent for an average of 10 aligned traces (Figure 8). The EVM and the uncertainty of the EVM, calculated from the modulation waveform and its uncertainty, was measured as 0.41 ± 0.02 percent, which is considerably less than the typical value (0.8 percent).

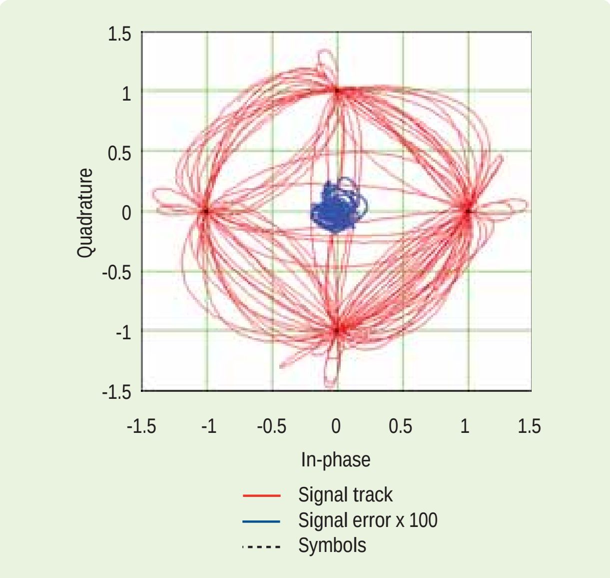

Figure 8. Typical W-CDMA modulation waveform results for 30 traces (red). The error between different results sets is < 0.25 percent (blue).

Conclusion

As test equipment and RF modulation formats become increasingly complex, it becomes more difficult for instrument manufacturers to prove equipment specifications. Agilent and NPL have collaborated to develop RF waveform metrology into a powerful tool that traceably links industrially important W-CDMA test equipment to primary standards. As one example, the measured EVM performance of the Agilent N5182A MXG vector signal generator is considerably better than the typical specified value. NPL has also used these techniques to help Bristol University diagnose subtle faults in signal coding.

Communications is a constantly evolving field and the supporting metrology must evolve at a similar rate. With the future rollouts of WiMAX and UWB, RF spectrum availability will move to higher frequencies and data rates. As the volume of data becomes excessive, this will provide new challenges for RF waveform metrology based on oscilloscopes.

References

1. www.npl.co.uk/wireless_communications/digital_comms_parameters

2. Agilent Application Note 1356, *Designing and Testing W-CDMA User Equipment.* Publication number 5980-1238EN available from www.agilent.com.

3. Agilent Application Note 1355, *Designing and Testing 3GPP W-CDMA Base Transceiver Stations*. Publication number 5980-1239EN available from www.agilent.com.

4. 3GPP TS 34.121 V5.6.0: Terminal conformance specification; Radio transmission and reception (FDD).

5. Humphreys, D.A., Dickerson, R.T. 2007. *Traceable Measurement of Error Vector Magnitude (EVM) in WCDMA Signals.* 2007 International Waveform Diversity & Design Conference, Pisa, Italy, June 4-8, 2007: 270-274.

6. www.npl.co.uk/measurement_for_innovators/

7. www.bristol.ac.uk/ccr/

8. Harper, M.R., Smith, A.J.A., Basu, A., and Humphreys, D.A. 2004. Calibration of a 70 GHz Oscilloscope. *CPEM 2004*. London, June 27 to July 2, 2004: 530-531.

9. Paulter, N.G., Smith, A.J.A., Larson, D.R., Sounders, T.M. and Roddie, A.G. 2003. NIST-NPL interlaboratory pulse measurement intercomparison. *EEE Trans. Instrum. & Meas.,* Vol. 52, No. 6, Dec 2003: 1825-1833.

10. Bieler, M., Seitz, S., Spitzer, M., Hein, G., Pierz, K., Seigner, U., Basu, M.A., Smith, A.J.A., and Harper, M. 2007. Rise-time calibration of 50-GHz sampling oscilloscopes: intercomparison between PTB and NPL. *IEEE Trans. Instrum. & Meas.*, Vol. 56, No. 2, April 2007: 266-270.

11. Smith, A.J.A., Roddie, A.G., and Henderson, D. 1996. Electro-optic sampling of low temperature GaAs pulse generators for oscilloscope calibration. *Optical and Quantum Electronics*, 28 (1996): 933-943.

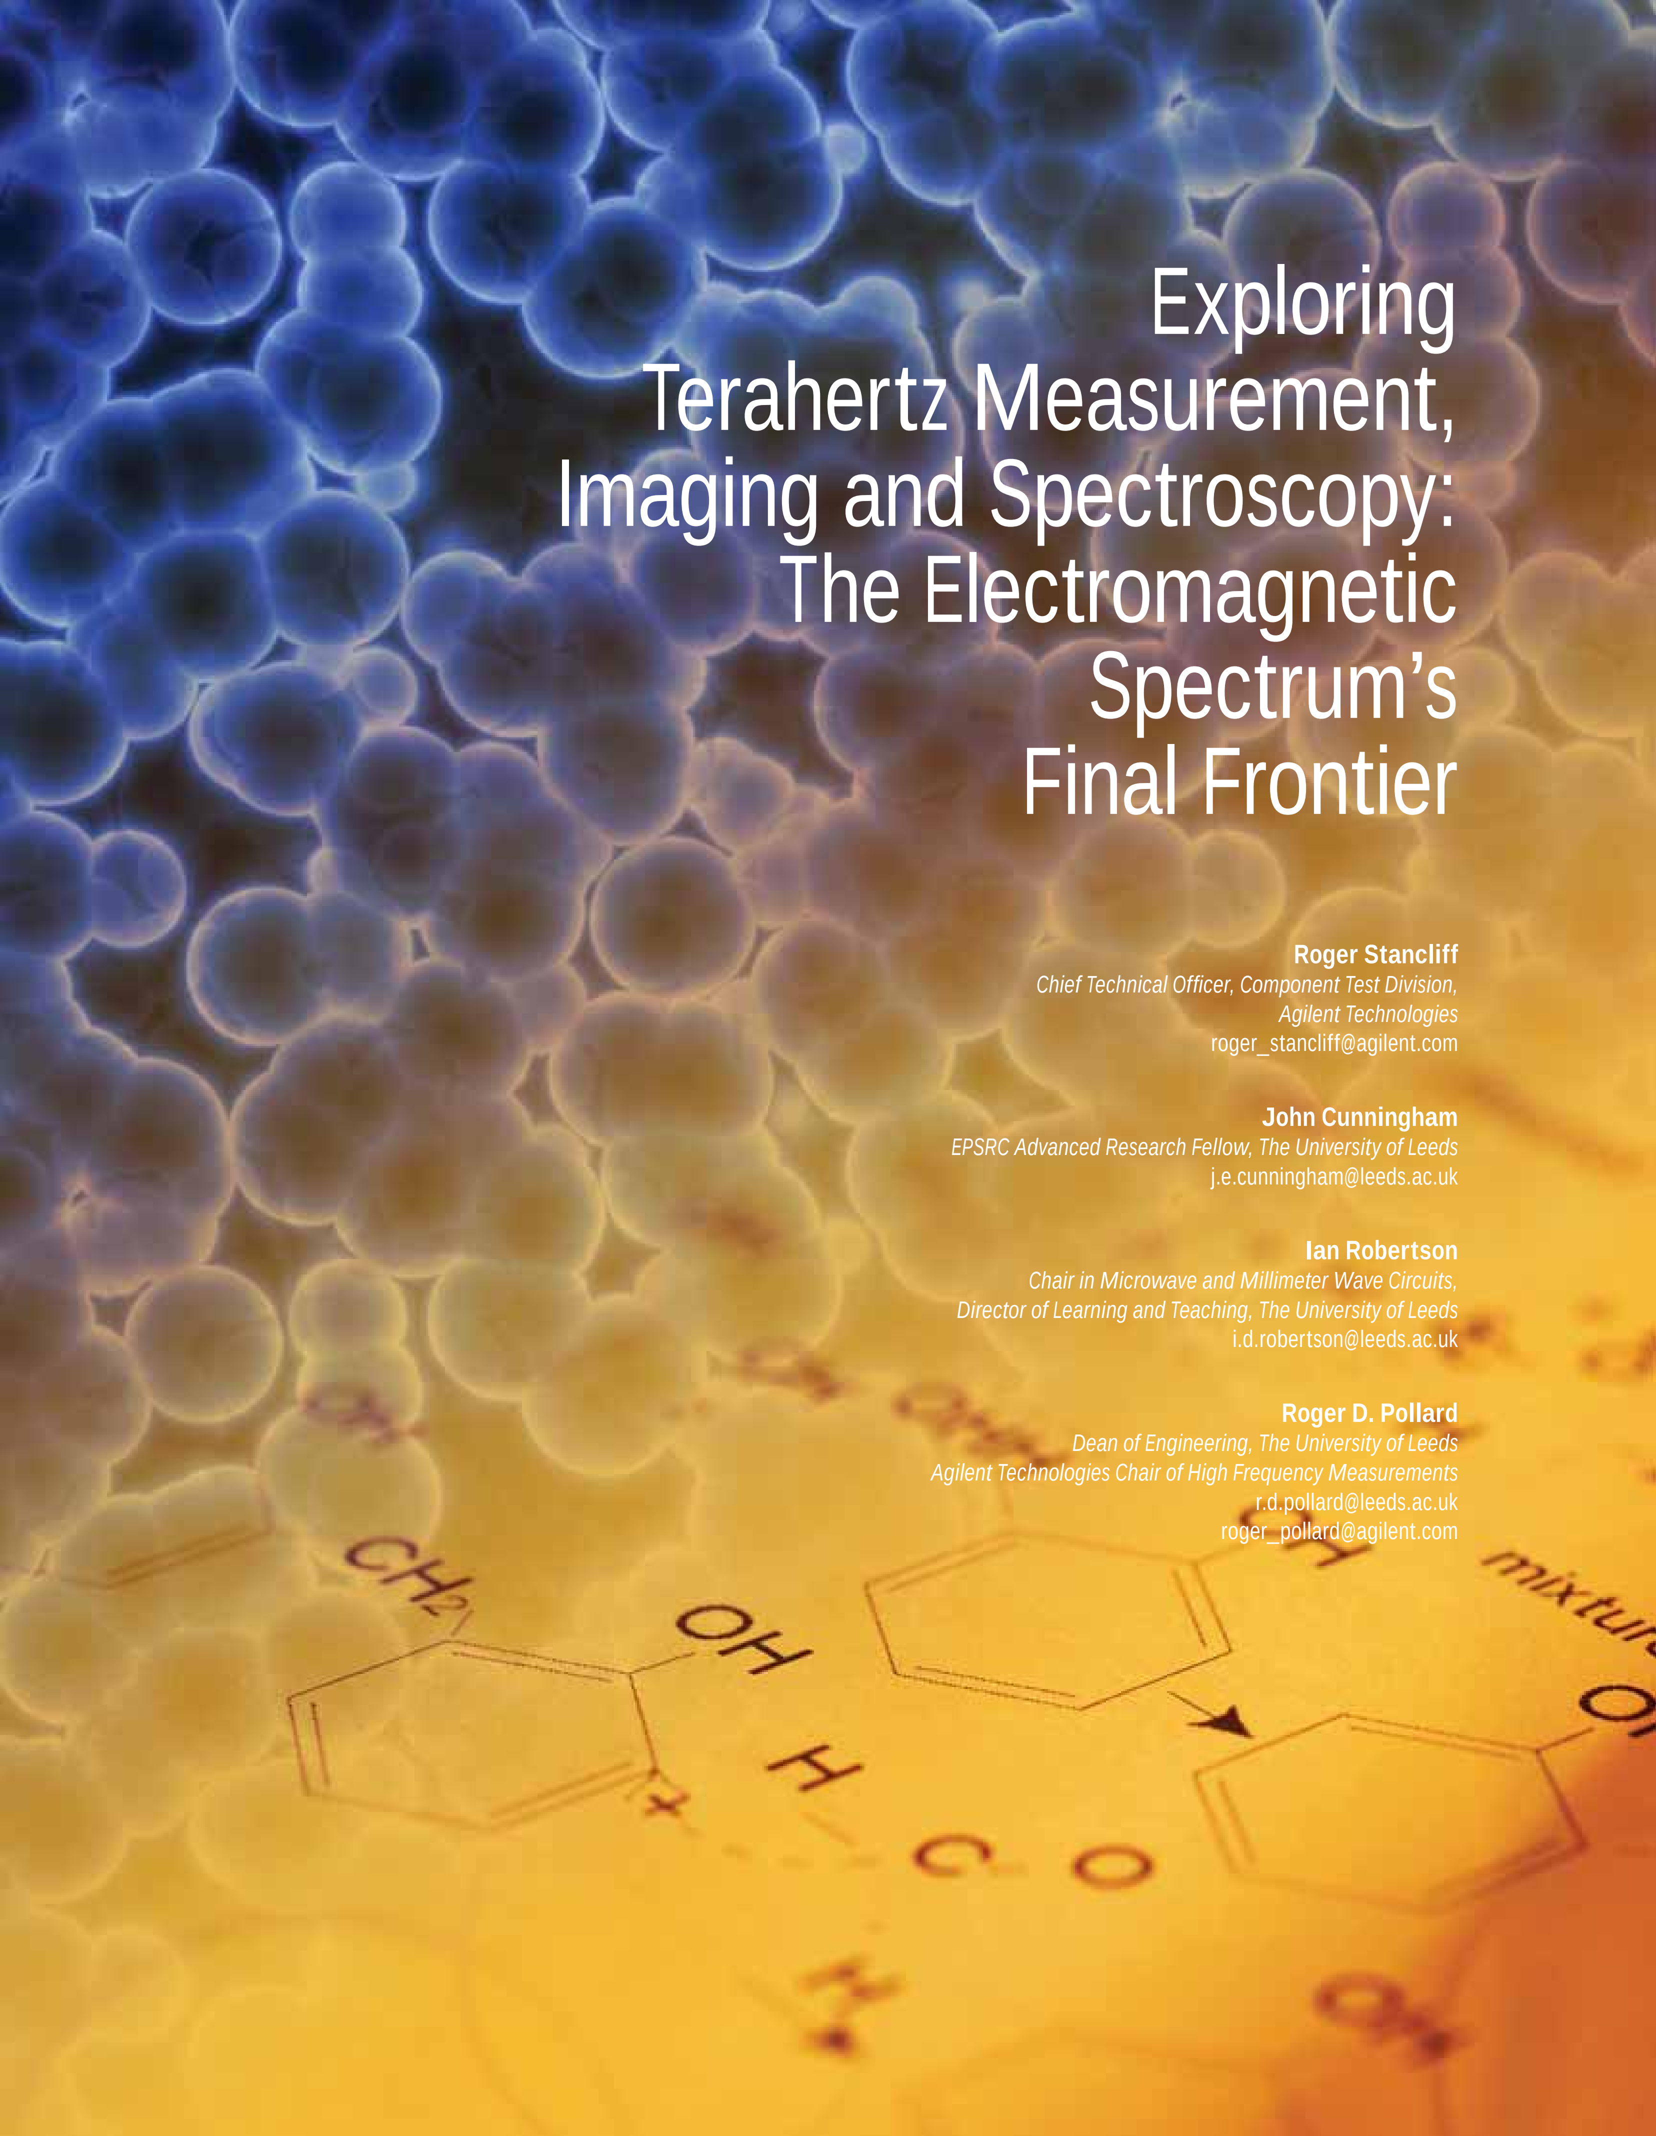

Exploring Terahertz Measurement, Imaging and Spectroscopy: The Electromagnetic Spectrum's Final Frontier

Roger Stancliff
Chief Technical Officer, Component Test Division, Agilent Technologies
roger_stancliff@agilent.com

John Cunningham
EPSRC Advanced Research Fellow, The University of Leeds
j.e.cunningham@leeds.ac.uk

Ian Robertson
Chair in Microwave and Millimeter Wave Circuits, Director of Learning and Teaching, The University of Leeds
i.d.robertson@leeds.ac.uk

Roger D. Pollard
Dean of Engineering, The University of Leeds
Agilent Technologies Chair of High Frequency Measurements
r.d.pollard@leeds.ac.uk
roger_pollard@agilent.com

The portion of the electromagnetic spectrum between the microwave and millimeter region (10^9 to 10^{11} hertz) and the infrared (IR) and optical region (10^{13} to 10^{15} hertz) is characterized by a lack of emitters and detectors capable of enabling useful measurements. The reason is that the electronic devices used at microwave frequencies are very difficult to extend up in frequency and the optical ones used at IR and optical frequencies are hard to extend down in frequency.

The popular name for this part of the spectrum is "terahertz" (0.1 to 10 THz). Interest in terahertz is accelerating since many materials exhibit unique terahertz frequency-range properties that provide high contrast for imaging and spectroscopic materials identification. There also is a need for measurement equipment to be expanded into the terahertz region not only to support these applications but also to measure devices that, due to Moore's law, are rapidly pushing up toward 1 THz and beyond.

For many years, Leeds University has performed some of the world's best research in terahertz. In the past five years its program has expanded to include involvement with most aspects of the terahertz research going on around the world. As a leading provider of microwave, millimeter wave and IR/optical measurement equipment, Agilent Technologies is supporting some of this research with an eye toward expanding our measurement coverage into this area — and exploring new possibilities in measurement, imaging and spectroscopy. There also are some other aspects of collaboration and partnership with Leeds that have grown up over the last 30 years and we will discuss these in the context of good models for university/industrial relationships.

Introduction: Why terahertz?

The terahertz part of the electromagnetic spectrum has four important characteristics that make it particularly interesting. First, wave propagation at these frequencies is very strongly affected by the presence of water. This could make terahertz useful for very short-haul secure communications. It also enables a large contrast between materials with different moisture content. A great example of the usefulness of this property is in the detection of cancerous tissues. Cancer tumors grab onto blood vessels (vascularize) to feed their rapid growth. As a result, they have much higher moisture content than normal surrounding tissues and easily can be detected.

Second, many solid materials have unique phonon resonances in the terahertz range. This leads to the ability to use spectroscopy to identify materials with high confidence. Examples include the detection of a range of plastic explosives and of different crystal forms (polymorphs) of pharmaceutical compounds. The latter is important for detecting counterfeit drugs and for protecting a company's intellectual property.

Third, many gas molecules have rotational moments at terahertz that lead to rapid identification with spectroscopy.

Fourth, the penetration of terahertz radiation into materials makes it suitable for many imaging applications. For example, new applications are emerging in nondestructive inspection of semiconductors, pharmaceutical compounds and medical diagnostics. In each case, the unique spectral response of materials to terahertz radiation reveals features that are difficult or impossible to see in any other frequency range. Resolution can be improved by using near-field techniques such as those discussed below. This approach will open even more exciting applications.

Figure 1. Research conducted in the University of Leeds Terahertz Photonics Laboratory focuses on developing new technology and applications for terahertz spectroscopy.

The main reason that terahertz application solutions have not become widespread is that terahertz waves are hard to generate, guide and detect. When significant progress is made in these areas, broad commercialization should occur, driven by the many unique capabilities of terahertz.

Terahertz research at Leeds: An overview

The University of Leeds Terahertz Photonics Laboratory, in the School of Electronic and Electrical Engineering, is a dedicated facility for investigating and utilizing terahertz frequency components and systems. Significantly refurbished in April 2004, it is probably the largest university facility for terahertz research in Europe, and one of the largest in the world.

The facility contains eight optical benches that include an impressive array of equipment: five pumped Ti: sapphire laser systems, a Bruker 66V FTIR spectroscopy system, three continuous flow helium cryostats and one 1.2 K Oxford Instruments pumped helium bath cryostat (including an 8 T superconducting magnet). All instruments feature optical access, purged and evacuated terahertz time-domain spectroscopy systems — including an ultra-broad bandwidth capability — and on-chip, guided-wave terahertz apparatus. Associated with this facility is a new III-V semiconductor molecular beam epitaxy (MBE) laboratory and a semiconductor nanotechnology cleanroom.

The laboratory's work is presently coordinated by four academic staff who lead a team of approximately 10 post-doctoral workers and 20 Ph.D. students. The main focus is on developing new technology and applications for terahertz spectroscopy. Three major programs are currently underway: spectroscopy of materials with security significance (such as drugs-of-abuse and explosives); fundamental investigations of condensed matter systems; and medical imaging and spectroscopy. The recent installation of a state-of-the-art MBE growth facility has enabled successful demonstration of terahertz quantum cascade lasers, which are now being supplied to several groups worldwide, including Harvard University.

Agilent-funded research: Terahertz evanescent-field microscope

John Cunningham
University of Leeds

In recent years, there has been great progress in the use of terahertz frequency radiation for medical imaging applications. In particular, the potential of pulsed terahertz techniques has been proven in the laboratory diagnosis of basal cell carcinomas, melanomas and dental caries.

However, a number of technical challenges must be addressed and overcome before this frequency range can be routinely applied to biological imaging. The two key technical issues to address are the complexity and cost of the terahertz time-domain apparatus and the spatial resolution available. Furthermore, three additional issues are limiting the rate of progress: the lack of a fundamental understanding of the image contrast mechanisms obtained, the nature of the interaction of terahertz radiation with tissue materials and the interplay between radiation of absorption and scattering phenomena in inhomogeneous materials. Nevertheless, there is clear and demonstrated potential for the future industrial development of terahertz imaging systems.

Existing terahertz imaging and spectroscopy systems typically utilize the technique of free-space time-domain spectroscopy. This offers a limited spatial resolution determined by the diffraction limit of the radiation (of order 1 mm). A number of prototype near-field terahertz microscopy systems have been demonstrated based on modifications of time-domain spectroscopy; however, these have severe operational limitations since they either require powerful optical radiation focused on the sample area — causing damage — or they waste a large portion of the valuable terahertz power by constricting the terahertz beam using an aperture. The objectives of this project are to develop a new microscopy technique with high spatial resolution that aims to overcome these difficulties and to perform more systematic studies of the interaction of terahertz radiation with sensitive materials such as biological samples (Figure 2).

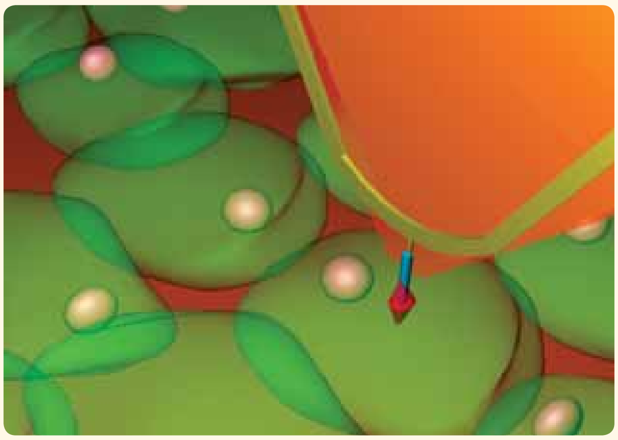

Figure 2. Artist's impression of new terahertz imaging technique being used to scan a biological substrate. An on-chip waveguide (yellow) directs terahertz radiation into a resonant filter (blue), whose electrical properties are determined by the surrounding evanescent field (red). Scanning the sample produces a contrast map of terahertz effective permittivity.

The technology behind the technique can confine terahertz radiation to an active region (a scanning "tip") that is several tens of square microns in area. The tip, which takes the form of a filter in a planar on-chip terahertz waveguide, is scanned a small and controllable distance (<1 µm) away from the sample under study. The tip region is highly sensitive to changes in the dielectric permittivity of its surroundings, which alter the tip's high-frequency electrical properties (its resonant frequency and peak attenuation).

These properties are read out using an optoelectronic photoconductive measurement scheme that provides a bandwidth extending from a few tens of gigahertz to more than 1 THz. The material property of interest is the high-frequency dielectric constant (refractive index) of the sample being scanned. This property is directly accessed from the size of the shift in the electrical resonance of the band-stop filter induced by the samples as they are scanned in close proximity to the band-stop filter (tip). Work over the last year has successfully demonstrated the technique's proof of principle in scanning semiconductor substrates, allowing it to be optimized for resolution and sensitivity.

Agilent Ph.D.s at Leeds

Roger Pollard

Throughout the course of the relationship with the University of Leeds, there has been Agilent (and previously Hewlett-Packard) support for Ph.D. students at Leeds. Here are four examples of topics that have been the subject of successful Ph.D. theses:

- *The Design of Microwave and Millimetre-Wave Power Combining Arrays* (Andrew Adams)
- *On-Wafer Microwave Noise Characterization* (Caroline Collins)
- *Network Analyzer Techniques for the Characterization of Lightwave Components* (Bala Elamaran)
- *Optimum Design of Microwave Filters with Finite Dissipation* (Andrew Guyette)

Much of the work has been published in internationally recognized peer-reviewed journals; some of the technology has found its way into Agilent products; and all of the students are now employed in senior positions in the microwave industry (including some at Agilent).

Agilent-funded research: Terahertz connectors and cables

Yun Hua Zhang, Ian Robertson and Roger Pollard
University of Leeds

There is a growing demand for precision measurements above 200 GHz. Advanced devices such as SiGe, GaAs and InP transistors can operate to 500 GHz, and even regular CMOS technology is now reporting potential operating frequencies well beyond 100 GHz as device dimensions scale down to tens of nanometers. At such high frequencies, laboratory measurements are currently conducted with either rectangular waveguides or in some cases, free-space beams. The former are expensive and narrowband; the latter requires setups that are too fragile to be employed in most applications outside the specialist terahertz research lab. The electronics and physics communities require terahertz connectors and cables not only to extend the capabilities of traditional measurement equipment but also to enable a wide range of new medical sensing and imaging systems.

My relationship with Agilent

Roger Pollard

My Agilent odyssey started in 1981 as a young faculty member at the University of Leeds looking to take a brief sabbatical — one different from the usual route of going to another university in another country to teach the same subjects to someone else's students. Some research background in network measurements made the Hewlett-Packard (now Agilent) division in Santa Rosa, California seem the ideal place, so I penned a request to a conference contact. Three months later, a three-line message asked, "Can you start next week?" I spent the next seven months as part of the team that developed the HP 8510A, little realizing the impact this landmark product would have on the microwave industry. I returned to Leeds and settled back into academic life, but the following summer received a surprise call: "So, when can we expect you?" Now, nearly 27 years later, I still spend part of my summers at Agilent in Santa Rosa. I'm grateful to have had the opportunity to collaborate with some of the finest engineers in the world.

This project has compared myriad potential solutions to this terahertz cable requirement. Coaxial cables and connectors are beyond the limits of conventional fabrication. Rectangular waveguides are restricted to bandwidths of less than one octave. Furthermore, metallic guides of any form have high conductor loss in the terahertz region and there is a need for ultra-low surface roughness. As a result, the project is concentrating on investigating the application of dielectric waveguide techniques. Two materials have emerged as strong candidates for this application: PTFE (Teflon) and polypropylene have reported loss tangents of approximately 0.006 at 1 THz. PTFE has proven itself extensively in microwave connector applications; polypropylene is very easy to fabricate.

Standard circular dielectric waveguide structures — equivalent to conventional optical fibers — have significant dispersion and other limitations. This project has demonstrated that the photonic crystal fiber (PCF) technique (sometimes referred to as the "holey fiber") from optical applications can be applied to the design of a terahertz dielectric waveguide. Modeling work has been conducted using the full vectorial Effective Index Method (EIM).[1] The use of a microstructured guide is intended to achieve endless single-mode behavior (important for ultra-broadband applications in measurement and imaging) with a flattened dispersion characteristic, controllable mode area (for ease of transition design) and low loss.

Figure 3 shows the cross section of a conventional hexagonal PCF. To demonstrate the broadband single-mode behavior of PCF in the terahertz band, Figure 4 compares the mode behavior of a step-index fiber and a PCF. For the PCF, the single-mode range is from 202 GHz to 2376 GHz; in contrast, the single-mode range of step-index fiber is from 200 GHz to 540 GHz.

To illustrate the mechanism of single-mode behavior of PCF, it is possible to introduce normalized V_{eff} by making an analogy to normalized V in step-index fiber:

$$V_{eff} = \frac{2\pi}{\lambda} a_{eff} \sqrt{n_{core}^2 - n_{clad}^2}$$

PCF cross section

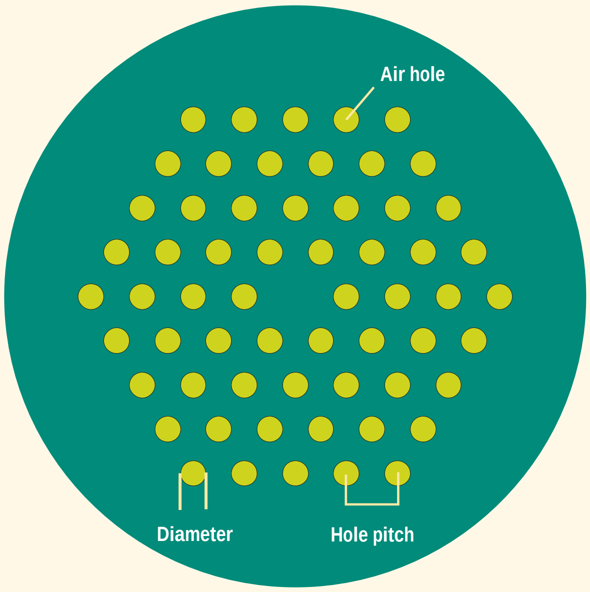

Figure 3. Conventional hexagonal photonic crystal fiber (PCF), shown in cross section, can be applied to the design of terahertz dielectric waveguide.[2]

The effective refractive index of the cladding region n_{clad} is frequency dependent. When the frequency is relatively high, the field becomes more concentrated in the silica region and away from the air holes, which leads to high n_{clad}. Consequently, the gap between n_{core} and n_{clad} becomes small at high frequency and therefore the frequency range for single-mode operation can be extended.

The project is developing full-wave modeling techniques for terahertz PCF dielectric waveguide structures. The modeling will then lead to a study of the optimum design of a dielectric guide cross-section for single-mode behavior and low dispersion. The design of transitions from the dielectric guide to conventional sub-millimeterwave rectangular waveguide will be investigated in order to demonstrate practical measurements in the 140-GHz to 325-GHz range using a sub-millimeterwave Agilent 8510XF system.

After experimental determination of the waveguiding properties (phase propagation constant and loss) of the flexible dielectric guide, work will concentrate on whether the guide can be engineered into a practical ruggedized cable with connectors. For connectors, such high frequencies pose severe challenges on fabrication tolerances and planarity of interface. What's more, it is essential — but difficult — to design a metrology-grade connector with 60-dB repeatability levels.

Applications around future products such as security and medical scanners require cables and connectors that can be replaced by nonspecialist field engineers. Connector design issues will be studied in terms of connectability (ease of connect/disconnect, standardization of coupling mechanism, threads/engagement/latching), repeatability and lifetime, and effect of mechanical tolerances. The longer term goal of the research is to work towards an open standard for terahertz connectors and guiding structures.

Stepping backward, then forward

Joel Dunsmore, Ph.D.
Senior R&D Engineer/Scientist, Agilent Technologies
joel_dunsmore@agilent.com

It started with an observation of unusual time-domain behavior when tuning a multicavity filter: The time-response changes corresponded to changes in the tuning screws. From this, I found that these filters could be tuned exactly using only the time-domain response — something that had never before been observed. Taking this beyond mere observation, however, would require a detailed look at the mathematics of filter responses and Fourier transforms. As it happened, Roger Pollard was visiting for the summer and we discussed a Ph.D. topic focused on a theoretical explanation of the observed behavior. I quickly discovered that this was akin to jumping in at the midpoint of a Ph.D. program. I spent the next two years seemingly going backward until a broad study of the area enabled me to confidently describe the research problem. Outside interests fell by the wayside as some aspect of research, experiment or writing absorbed any free time. Three years later I had answered the problem, and also came to understand that learning the process of research is just as important the research itself.

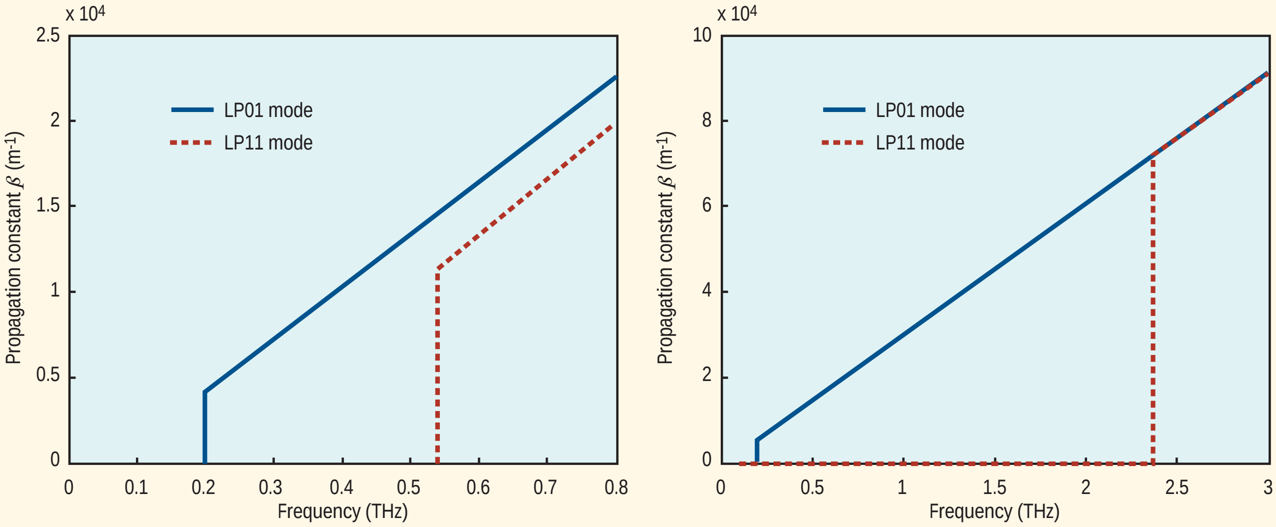

(a) Propagation modes of step index fiber

(b) Propagation modes behavior of PCF

Figure 4. These diagrams enable a comparison of mode behaviors: (a) Propagation modes of step index fiber, n_{core} = 1.45, n_{clad} = 1, r = 0.21 mm; (b) Propagation modes of PCF, n_{core} = 1.45, air hole diameter d = 0.350 mm, air hole pitch Λ = 0.725 mm.

Creating synergy across work and research

Dave Blackham, Ph.D.
R&D Engineer, Agilent Technologies
dave_blackham@agilent.com

After joining the company, I used the continuing education program to earn MSEM and MSEE degrees. One afternoon, I crossed paths with Roger Pollard and he congratulated me on completing my MSEE. I replied with a wish of a way I could pursue a Ph.D. without having to take a leave of absence from work. He offered to help me out. At the time, I was looking into microwave characterization of materials using vector network analyzers as a way to expand our business. Roger helped me set up a Ph.D. program in that area — and let me do it remotely from Santa Rosa. There was a definite synergy in pursing a Ph.D. simultaneously with developing a product. The Ph.D. requirements not only demanded much more effort but also yielded a detailed understanding of my work and how it compared to published work. I am grateful to Roger for mentoring me through the whole process.

Looking to the future of terahertz

Aside from the work on terahertz guiding and terahertz microscopy, researchers at Leeds and other universities are exploring many different ways to generate and detect terahertz. Electronic approaches to signal generation have been pushed up to 2 THz and solid-state lasers have operated below 2 THz. Still, none of these are currently available from commercial vendors. For terahertz to become commercially attractive, further breakthroughs in signal generation and detection will be required. For example, the existence of a signal generator with >1 mw output power, tunability over an octave of frequency, operable at room temperature and priced less than $10,000 would begin to enable a range of small and reasonably priced equipment.

The closest technology to achieve this currently would be backward-wave oscillator (BWO) tubes, but they tune only 20 to 30 percent of bandwidth and are large because of the required magnets. Miniaturizing this type of e-beam device is one possible approach to solving the signal generation problem; however, the continuing progress in semiconductors (following Moore's law) may provide an optimum long-term solution. Continuing development of solid-state lasers that achieve reasonable power at room temperature — even if they can't be tuned — would enable inexpensive solutions for spot problems in which a narrow spectral range is enough. Today's femto-second laser terahertz generation techniques will continue to be refined but, just as coherent oscillators replaced arc generators in RF systems in the 1920s and 1930s, it is reasonable to expect a similar transformation in terahertz.

Continued good engineering progress will gradually increase usage of current solutions and continue to build market demand for more. It's always impossible to schedule invention and breakthroughs, but it seems likely that within the next five years, a significant breakthrough in terahertz generation will occur. This will enable rapid progress towards common availability of terahertz solutions.

Working with the leading university research groups is probably the best way to get early warning of when such breakthroughs will occur. From that inflection point, rapid time-to-market with the new technology will determine commercial success.

References

1. Midrio, M., Singh, M.P. and Someda, C.G. 2000. The Space-Filling Mode of the Holey Fibers: An Analytical Vectorial Solution. *IEEE J. Lightw. Technol.*, Vol. 18, No. 7, July 2000: 1031-1037.

2. Saitoh, K., and Koshiba, M. 2005. Numerical Modeling of Photonic Crystal Fibers. *IEEE J. Lightw. Technol.*, Vol. 23, No.11, November 2005: 3580-3590.

University connections are about people

Jack Wenstrand
Director of University Relations, Agilent Technologies
jack_wenstrand@agilent.com

A successful academic collaboration has many facets, providing rich benefits to the participating individuals and institutions. These benefits go far beyond the progress made on the initial research topic. One mark of a successful collaboration is that it continues long after the initial interaction is completed, and every such continued relationship has at its center a committed individual. I can think of no longer-running, deeper, more valuable single academic collaboration at Agilent than the one with Professor Roger Pollard, dean of the Faculty of Engineering at the University of Leeds.

Roger's long association with Agilent and regular sabbatical visits have benefited the company in many ways: direct technical contribution to several of our world-leading network analyzer products; regular doses of intellectual rigor for our staff; and three Leeds Ph.D.s superbly trained by Roger now in leadership positions at Agilent with a fourth in the making. What's more, Roger's detailed knowledge of the unique capabilities of our equipment and its utility for teaching and research has helped expand the market for our products.

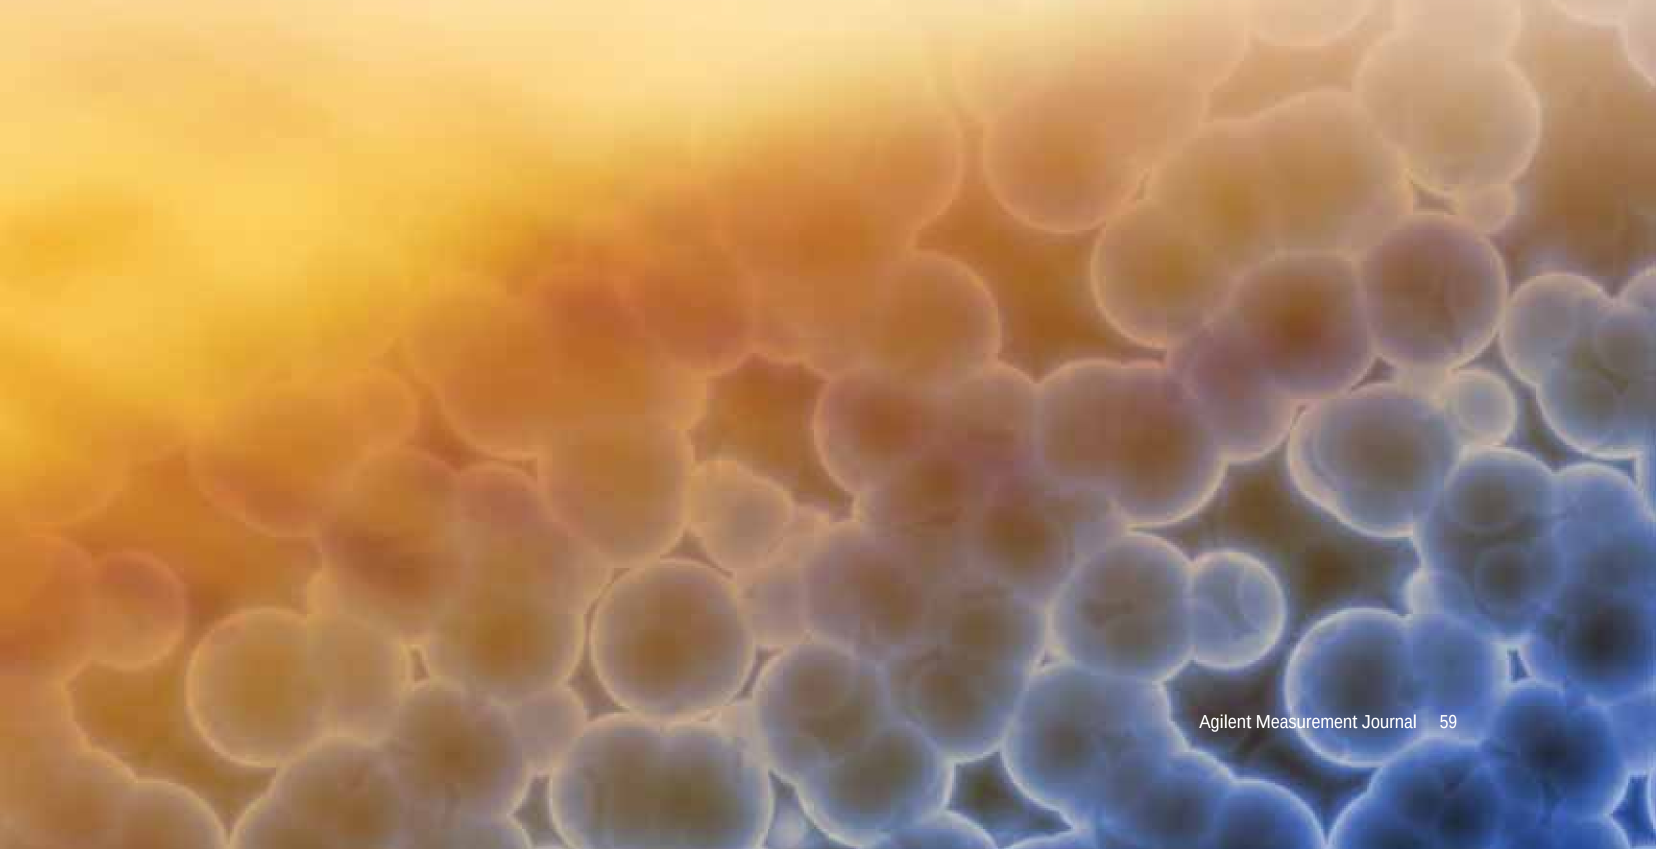

Interpreting Quoted Specifications when Selecting Digitizers

Richard Soden
Product Marketing Manager, Agilent Technologies
richard_soden@agilent.com

Selecting the best digitizer solution for any application is a complex task. At a minimum, it involves an in-depth evaluation of what must be measured, how it can be measured and the required degree of accuracy. These factors must then be used in the comparison of a multitude of available acquisition schemes and the associated individual device specifications.

Digitizer manufacturers typically present a device's "banner specifications" such as bandwidth, resolution and sampling rate. Although these are offered as an indication of instrument quality, they often have little or no impact on the ultimate measurement fidelity in many applications.

There is no simple answer to the question, "Which digitizer should I use?" Instead, each application and device must be reviewed on a case-by-case basis. No system is ideal: It will add noise to the analog signal being sampled, create signal distortion (harmonics) and suffer from some amount of clock jitter. These elements and more contribute to system performance and should be reviewed within each set of device specifications.

Comparing devices

Great care should be taken when comparing devices between manufacturers because they often use different test procedures and limits. For example, one manufacturer may quote a harmonic distortion value that includes the first six harmonics while another manufacturer includes only the first five, thereby producing a lower distortion value. In such cases it is unclear as to which device presents the lowest distortion properties for any given input signal.

As a rule, device specifications should be compared like to like. This is not always possible, however, so specifications and the definitions of those specifications must be reviewed for each device (Figure 1).

Figure 1. Example of high-speed PCI digitizer that carries detailed performance specifications.

Understanding accuracy and resolution

Resolution is intended to define the fineness of detail that can be distinguished. Usually quoted as a number of bits, it indicates the number of discrete levels used in the encoding of the input signal. As such, a resolution of 8, 10, 12 or 14 bits should represent the ability to distinguish one part in 256, 1,024, 4,096 or 16,384, respectively (Figure 2).

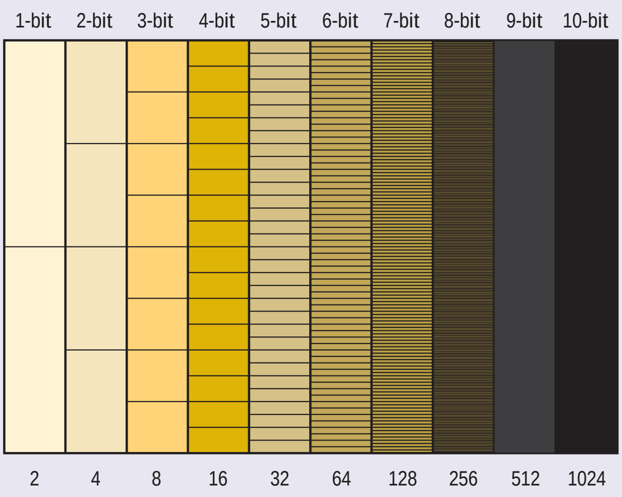

Figure 2. Resolution quoted as bits indicates the number of discrete measurement levels used.

Tip: Resolution is often taken as an indication of measurement accuracy; however, overall accuracy can be one of the most difficult specifications to determine from the specifications presented on a data sheet. In fact, a resolution value quoted as "bits" does not include any indication of the noise or distortion levels that effectively diminish a device's ability to distinguish between discrete levels. A better evaluation of system performance with respect to measurement accuracy is provided by the "effective bits" or "effective number of bits" (ENOB).

Determining the effective number of bits

When assessing a system in terms of effective bits, all error sources are included. Evaluation of system performance can be made without considering the individual error sources.

Theoretically, a 16-bit digitizer will be able to measure one part in 65,536; however, operating under real-world conditions, a good 16-bit board has at most 13.5 bits of accuracy. What's more, a poorly designed 14-bit board may perform less accurately than a well-designed 12-bit system. This is especially true when reviewing effective bits as a function of input signal frequency.

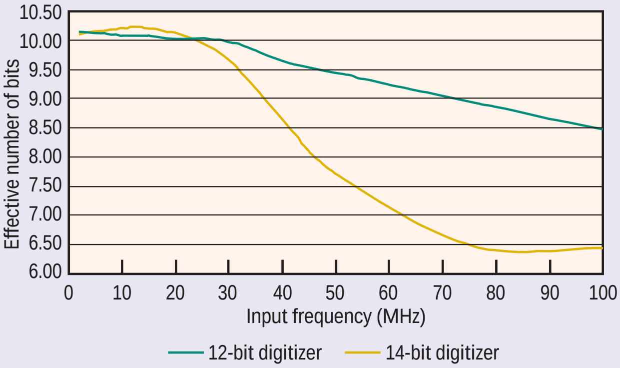

Figure 3. This diagram uses ENOB to compare commercial 14-bit and 12-bit digitizers that have 100 MSa / s sampling rates, 2 to 100 MHz input frequency range and input signal amplitude at 97.5 percent of 1V full scale.

The ENOB value can be calculated from the signal-to-noise-and-distortion (SINAD) ratio:

$$ENOB = \frac{SINAD - 1.76}{6.02}$$

Tip: ENOB specifications should be compared carefully because the value can be highly dependent on amplitude and frequency. An ENOB value should be quoted with the signal input level used, sampling rate and bandwidth over which the specification is measured.

From the example shown in Figure 3, if quoted only at 10 MHz, the 14-bit digitizer shows more effective bits than the 12-bit digitizer quoted at 50 MHz. This is clearly misleading because the 12-bit system shows better performance from about 25 MHz and beyond.

Applying SINAD

The signal-to-noise ratio (SNR) is the ratio, expressed in decibels, of the root mean square (RMS) value of the input signal at the output to the RMS value of the sum of all other spectral components. In contrast, the SINAD ratio, also expressed in decibels, is the ratio of the RMS value of the input signal at the output to the RMS value of all of the other spectral components, including harmonics. Because SINAD includes both noise and distortion, it can be used to directly calculate the effective number of bits.

Tip: Compare SINAD specifications carefully because the value can be highly dependent on amplitude and frequency. As with ENOB, the value should be quoted with the signal input level used, sampling rate and bandwidth over which the specification is measured.

Factoring in sources of sampling error

Analogous to bit resolution, sampling rate, quoted in samples per second (Sa/s), is often regarded as a measure of temporal accuracy. As with resolution, however, the given sampling rate of a digitizer does not indicate the levels of temporal noise added to the signal.

Sampling rate and all device-timing functions of a digitizer are driven by clock circuitry, either internal or external, that is assumed to be stable. Commonly based around a crystal oscillator, timing circuits are susceptible to frequency drift and spurious signals.

Clock accuracy, quoted as parts per million (ppm), indicates frequency drift. In the process of digitizing a signal, it is assumed that the samples are equally spaced in time so any clock inaccuracy will introduce a false frequency shift in the measurement of an input signal.

Sampling jitter indicates the random inaccuracies in the sampling of a signal. Normally quoted in picoseconds (ps), jitter introduces random noise into the signal by dispersion of the actual signal in time.

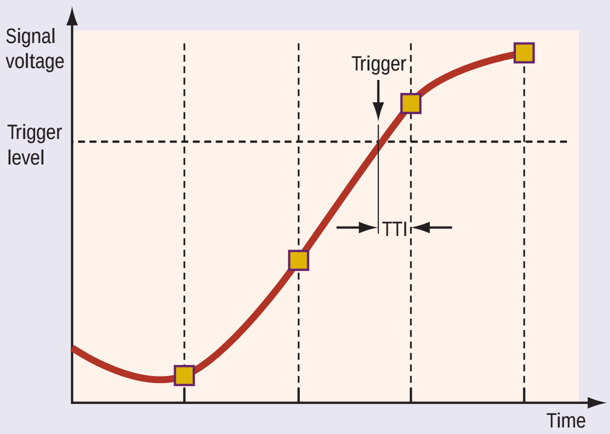

Figure 4. Trigger time interpolation allows the precise placement of an external trigger between sample points at a given resolution, normally quoted in picoseconds.

Some digitizers include a capability called time-to-digital conversion (TDC) or trigger time interpolation (TTI). Implemented in circuitry within the device architecture, this function allows the precise placement of an external trigger with respect to the internal clock of the digitizer, even when that trigger falls between sample points (Figure 4). TDC performance is quoted in picoseconds and is particularly important when measuring repetitive signals. If no interpolation is used, phase errors may be introduced in the measurement.

Tip: A poorly designed digitizer system may need an external reference clock to achieve timing performance comparable to that of a well-designed device. It should also be noted that in some cases, particularly measurement of repetitive signals, the sampling-rate value can be ignored completely if large acquisition memories are present in the system.

Characterizing bandwidth with the Bode plot

The analog bandwidth (BW) of a device defines the difference between the lowest- and highest-frequency signal components that can be measured before the input signal is attenuated to –3 dB of its original value. This is often quoted as a single value in hertz that indicates the device's upper frequency limit. However, a more useful measure of a digitizer's frequency response is provided by a Bode plot (or Bode diagram), which indicates gain as a function of frequency (Figure 5).

A Bode plot shows the levels of attenuation over all frequencies. Ideally, the curve should be flat (following an asymptotic limit) and then roll off steeply to the bandwidth limit. This is very difficult to obtain and often the Bode plot of a device can show various inconsistencies in gain as a function of input frequency, even for devices that possess the same bandwidth values.

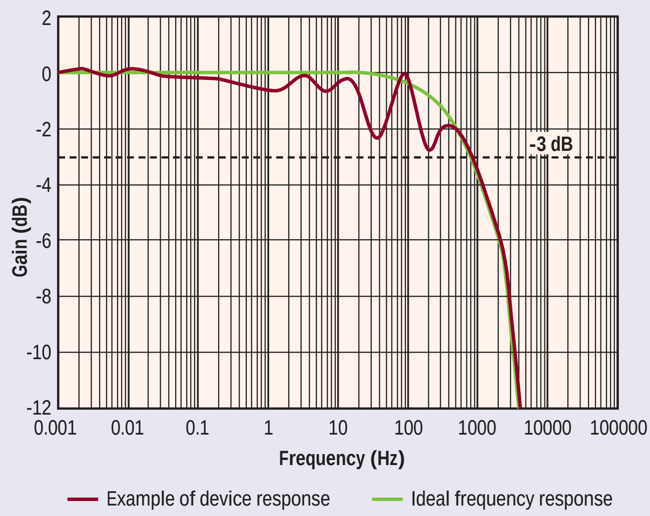

Figure 5. A Bode diagram shows signal gain or attenuation as a function of frequency.

Tip: Note that when measuring the relative amplitude of fixed-frequency signals, digitizer bandwidth and the Bode plot will be of little relevance because input signals of the same frequency will all experience equal attenuation. Bandwidth and the Bode plot have the greatest influence when measuring signals over a broad frequency range.

Considering total harmonic distortion

Total harmonic distortion (THD) is the ratio in dB of the total level of the first few harmonics to the level of the input signal at the output. THD is usually calculated from an FFT of a test signal where f is the fundamental (input) frequency and f_1 through f_n are the first n harmonic frequencies:

$$THD = \sqrt{\frac{f_1^2 + f_2^2 + f_3^2 + \ldots + f_n^2}{f^2}}$$

Tip: Harmonic distortion values are most meaningful when quoted with the number of harmonics used in the calculation; the input frequencies over which the measurement is valid; the sampling rate of the measurement quoted; the full-scale voltage range; and the input voltage level as a percentage of that full-scale range.

Examining spurious-free dynamic range

Spurious-free dynamic range (SFDR) is the difference, expressed in dB, between the RMS values of the input signal at the output and the peak spurious signal where a spurious signal is any output signal that was not present at the input (Figure 6).

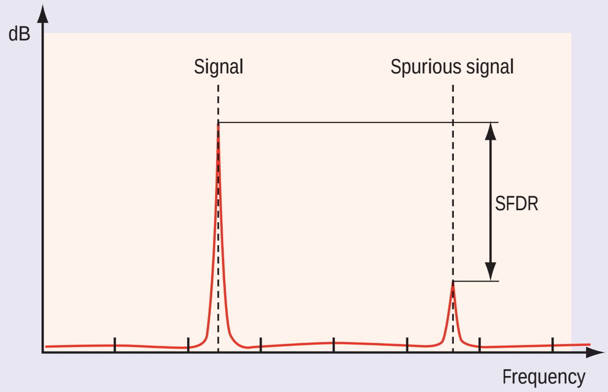

Figure 6. Spurious-free dynamic range affects the usable dynamic range in a measurement.

Tip: SFDR values should be quoted with the input frequencies over which the measurement is valid; the sampling rate of the measurement quoted; the full-scale voltage range; and the input voltage as a percentage of that full-scale range.

Conclusion

The handful of key specifications described here only begins to illustrate the non-trivial task of selecting a product based on published specifications. For example, ENOB (and SINAD) provide some indication of amplitude measurement accuracy. Clock accuracy and sampling jitter provide an indication of the frequency measurement accuracy and the temporal noise introduced in the measurement.

There are additional useful specifications, each with individual implications for specific applications. For example, linearity addresses the fact that the discrete measurement levels described by the bit resolution are not necessarily equally spaced, leading to amplitude measurement error. Voltage standing wave ratio (VSWR) describes the reflection of the input signal at the input terminals, an effect that diminishes the signal amplitude before it is measured and also introduces signal echoes into the system. VSWR also is frequency dependent, introducing more complexity to the selection process. Noise ratios have been addressed here, but other specifications such as the Sparkle Code Rate indicate the probability of a sampled point exceeding a specified deviation threshold. In general, spurious noise could have an effect on any peak-detection routine based on the digitization of a signal.

In most cases, obtaining the best measurement system for a specific application can be ensured by consulting not only the specification sheets of the various devices but also the manufacturers themselves.

Contact Agilent

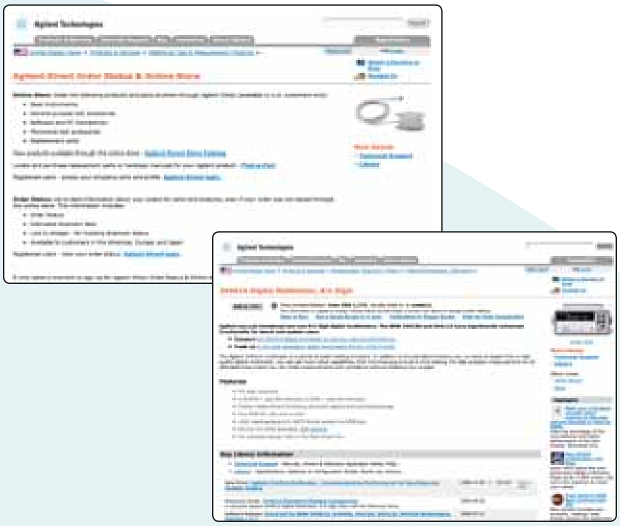

Agilent Direct Store

- Quickly build and place your order, 24x7
- Easily select and configure your products
- Share carts within your company (e.g., technical team and purchasing)
- Flexible payment options
- Secure transactions
- Order basic instruments, general purpose test accessories, microwave test accessories and replacement parts online today (available to U.S. customers only)

Agilent Direct

www.agilent.com / find / onlinestore

Email Update Service

Get the latest information and newest test skills
Subscribe to our free Email Update Service for the latest Agilent product, support and application information — customized for your interests and preferences.

Only the information you want, when you want it
Easily identify your specific interests from a selection of more than 100 product and application areas and specify the types of information you would like to receive:

- **Support information** such as new-product information and specifications; application notes; tradeshows, events and live or virtual seminars.
- **What's New**, including education/training and consulting services; firmware and software upgrades; and warranty, calibration and repair.

Total commitment to privacy
Agilent is proud to participate in the Better Business Bureau On-Line Privacy Program, and we adhere to all of the program's standards.

Agilent Email Updates

www.agilent.com / find / emailupdates

Agilent Test and Measurement RSS Feeds

Subscribing to Agilent´s RSS feeds lets you receive Agilent news as soon as it is published. With over 50 topics — ranging from "AC Power Sources/Power Analyzers" to "Wireless Test Systems" — you may find several feeds of interest.

www.agilent.com / find / rss

www.agilent.com / go / journal

Printed in U.S.A. January 1, 2008

5989-7680EN

Agilent Technologies