INVITED PAPER   *Special Section on Analog Circuits and Related SoC Integration Technologies*

Built-In Measurements in Low-Cost Digital-RF Transceivers

Oren ELIEZER[†a)] *and* Robert Bogdan STASZEWSKI[††], *Nonmembers*

SUMMARY   Digital RF solutions have been shown to be advantageous in various design aspects, such as accurate modeling, design reuse, and scaling when migrating to the next CMOS process node. Consequently, the majority of new low-cost and feature cell phones are now based on this approach. However, another equally important aspect of this approach to wireless transceiver SoC design, which is instrumental in allowing fast and low-cost productization, is in creating the inherent capability to assess performance and allow for low-cost built-in calibration and compensation, as well as characterization and final-testing. These internal capabilities can often rely solely on the SoCs existing processing resources, representing a zero cost adder, requiring only the development of the appropriate algorithms. This paper presents various examples of built-in measurements that have been demonstrated in wireless transceivers offered by Texas Instruments in recent years, based on the digital-RF processor (DRP[TM]) technology, and highlights the importance of the various types presented; built-in self-calibration and compensation, built-in self-characterization, and built-in self-testing (BiST). The accompanying statistical approach to the design and productization of such products is also discussed, and fundamental terms related with these, such as 'soft specifications', are defined.
key words: system-on-chip (SoC), digital RF processor (DRP), design for testability (DfT), design for manufacturability (DfM), built-in self-testing (BiST), soft specifications

1. Introduction

The past decade has seen an explosion in the consumer markets of many RF transceiver based products, such as mobile phones, wireless headsets, wireless local-area-networks (WLAN), global-positioning satellite (GPS) receivers, and various entertainment and computing devices. This growth has been both a driver and the outcome of the ever-increasing levels of integration, which have been enabled by the advent of CMOS technology. It has also created a persistent cost-reduction trend, which is supported not only by the increased level of integration and the resultant miniaturization, but also by the introduction of fundamentally new approaches to design and productization. These include built-in mechanisms that allow for self calibration and compensation, for production testing, and even for characterization. All this has become possible with the integration of the analog/RF functions of a transceiver with its digital baseband processor and memory on a single CMOS die, as demonstrated in [1]–[4].

The motivation for developing digital alternatives to functions that were traditionally designed in analog, as well

as a survey of a few examples, may also be found in [5], [6]. In addition to the pioneering effort by Texas Instruments, there have been a large number of other commercial realizations of all-digital PLL's, for example [23]–[25]. An extensive recent survey of transceiver solutions and related patents filed by several companies may be found in [25].

2. Integrating Analog/RF with Digital

For many years, digital circuitry has been benefiting from automatically-inserted testing capabilities and high production yield, dominated by the inevitable defect-density experienced in the chip fabrication. Contrarily, analog/RF functions have been designed and tested using entirely different methodologies, which are incompatible with a digital SoC environment in terms of the testing costs and the targeted yields. For example, digital processors, logic, and memory, would be tested using primarily built-in mechanisms, and would rely on automatic-test-equipment (ATE) that is fully digital and is of low cost. Furthermore, many such digital devices could be tested in parallel, and the production yields would be typically very high (well above 95%), resulting in overall low testing costs. But when these digital elements are at the core of a modern RF transceiver SoC, where the RF/mixed-signal functions typically constitute a small portion of the SoC area, as shown in the example GSM SoC of Fig. 1, the testing requirements and costs, as well as the resultant yields, can potentially be dominated by the analog circuitry. Since analog circuitry is significantly more prone to performance variability, particularly when implemented in the most advanced CMOS processes, its performance can more easily deviate from the targeted specifications for it. The consequences of this are: the perceived
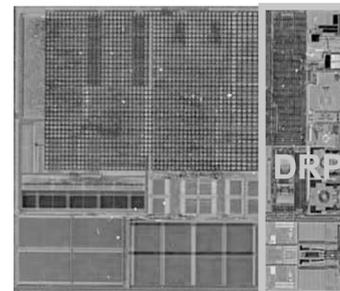
Fig. 1   Die micrograph of TI's GSM 90 nm CMOS SoC with the DRP quad-band transceiver part occupying only about 20% of the die.
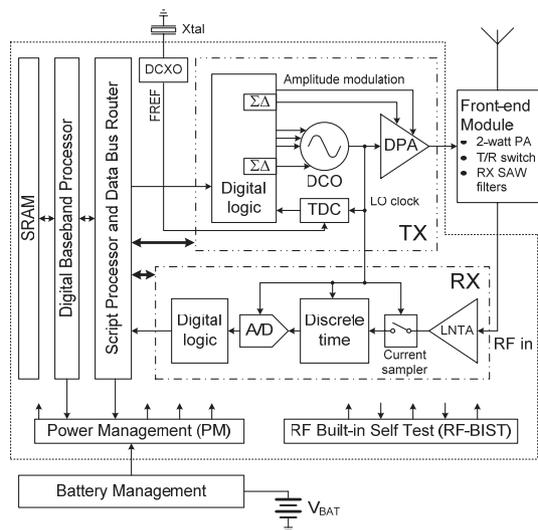
Fig. 2 Block diagram of the DRP-based cellular transceiver SoC.

Fig. 3 Processor access into ADPLL's internal signals allows testability.

need for accurate performance testing that is not supported by the low-cost digital ATE, yield losses, and, ultimately, increased production costs. The traditional approach to design and testing of RF/analog circuitry may have been acceptable when these functions were on a separate die, implemented in the most suitable fabrication process and allowed the assumption of mixed-signal ATE and production yields well below 90%. However, these are no longer valid assumptions in today's era of highly integrated, digitally-dominated, low-cost SoCs. Hence, in order to minimize production costs and ensure cost-competitiveness of the overall SoC, a new approach to design and productization must be adopted. The design aspects associated with this new approach include extensive digitization of functions that were traditionally designed in analog circuitry, reliance on built-in calibration/compensation that is typically based on digital processing, and a design-for-testability philosophy wherein a product's datasheet provides statistical distributions, rather than hard-limit based specifications, in order to minimize testing costs and maximize production yields. The architectural aspects of the digitally-intensive transceiver are not covered here, and may be found in [1], or in the Japanese translation of this book [2], as well as in [3]–[6] and in [25].

An example of a highly-integrated digitally extensive transceiver is shown in Fig. 2. In this example a Digital-RF Processor (DRP^{TM}) for GSM/EDGE is shown, which requires very few external functions to realize a GSM cellphone (e.g., battery management and a high-power front-end module). In addition to the baseband processor that is integrated with the analog/RF transceiver, the SoC includes a dedicated processor that is used to control the transceiver and to serve for various calibration and testing algorithms.

A particular transceiver function of interest demonstrating the principles discussed here, which is covered in detail in [1], [2], is the all-digital-PLL (ADPLL), shown in greater detail in Fig. 3. This sub-system, which realizes the functions of carrier generation and modulation for the trans-
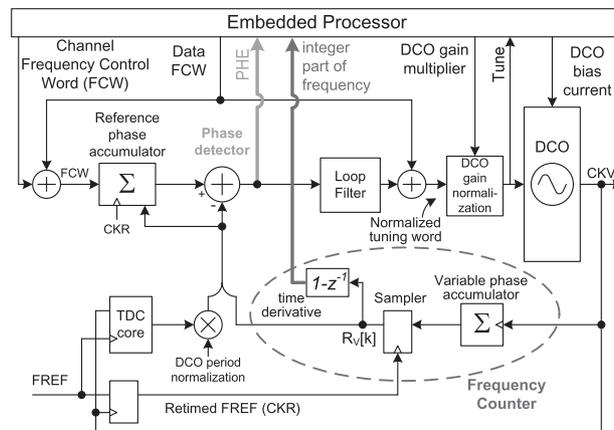
mitter, as well as local-oscillator generation for the receiver, serves as a classic example for the approach presented here. Due to its digital nature, various internal signals in the AD-PLL, such as its frequency measurements and phase error, are inherently digital, allowing simplified implementation of built-in compensation and production tests, some of which are discussed in Sect. 6.

Since the new approach greatly relies on the ability to internally measure and process signals of interest, typically corresponding to RF/analog performance parameters, it is important to distinguish between different types of internal measurements, which is the focus of the next section.

3. The Different Types of Built-In Measurements

Built-in measurements may serve various different needs, all of which enable the minimization of productization costs. Three of these are defined here in separate subsections, where the distinct purposes and requirements associated with each type of built-in measurement are outlined.

3.1 Production Testing (or 'Final Testing')

Although the term 'test' is often used in place of 'measurement', an internal measurement should be considered a built-in self-test (BiST) only when it is used to verify that the tested function satisfies some predefined criteria that is necessary to allow the device-under-test (DUT) to be delivered to a customer. Ideally, the integrity of the tested function, i.e., the absence of fabrication defects in it, would serve as sufficient criteria. Such reduced criteria would preferably result in the need for minimal and crude testing of that function and would serve in place of traditional accurate performance measurements, as demonstrated in [7]–[9].

3.2 Performance Adjustments (Calibration/Compensation)

Performance adjustments are required in order to compensate for intolerable performance variations that would be
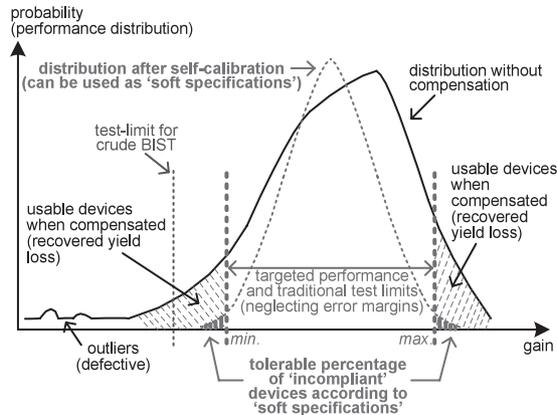
Fig. 4 Test strategies with and without self-compensation.

experienced as a result of variations in the fabrication process and in the operating conditions, including supply voltages and temperature. The built-in measurements associated with this need are required to be accurate enough to ensure, with sufficiently high probability, that the internally-measured performance can be brought to within the targeted range for it. This need is discussed in Sect. 7 and is demonstrated in [10]–[16].

In the absence of self-compensation capabilities, the final testing might be forced to include relatively accurate parametric measurements, requiring extended test time on costly mixed-signal ATE, and resulting in yield loss associated with the percentage of devices whose performance has shifted outside of the targeted window, as shown in Fig. 4. For this reason, a device's ability to self-sufficiently compensate for performance variations directly affects its productization costs and is therefore crucial for achieving cost competitiveness in a consumer market's mass-volume production environment.

3.3 Built-In Self-Characterization

The characterization of the performance of a particular function in a transceiver may pertain to one specific circuit in it or to a sub-system in the SoC that involves the proper functionality and coordination of multiple functions. Characterization activities may have different purposes depending on the productization phase in which they are carried out. For example, in an early post-silicon phase, accurate characterization, typically performed manually on a small number of fabricated samples, is carried out for the purpose of determining the extent by which the actual performance of the fabricated samples deviates from what was targeted. Certain deviations, resulting from design errors/oversights or modeling shortcomings that are established in this debugging phase, may require design modifications and therefore typically necessitate full and accurate understanding. This is usually obtained by performing measurements on different samples from different production lots, to experience process variations, as well as under different operating conditions. A later stage in which characterization is carried out

targets the identification of the statistical distribution of a parameter of interest over a large volume of devices. Although high accuracy may be desirable, it is also desirable to carry out such measurements in an ATE environment while relying as much as possible on built-in measurements, to allow for a reduced-time low-cost operation. An example of a built-in measurement capability that was used for the purpose of characterization is given in [16].

4. Built-in Self Testing (BiST) Fundamentals

Production test costs are proportional to the test time, and hence it is desirable to minimize the duration allocated to such measurements. This becomes more feasible as their accuracy requirements are relaxed and as more devices can be tested in parallel due to the minimized need for ATE resources to perform a given test. The production testing, also called 'final testing', is only meant to screen out defective and incompliant devices. Therefore, the measurement accuracy in final testing only needs to be sufficient to satisfy the targeted probabilities for making the two types of erroneous decisions in the screening of devices:

- failing a device that was suitable for delivery, thereby resulting in yield loss, or
- passing a device that should have been screened out, resulting in degraded delivery-quality, expressed in defective parts-per-million (DPPM).

Although the probability for both events would, ideally, be zero, non-zero probabilities are allowed in the low-cost mass-volume applications targeted here. This may be exploited as the test program is developed by allowing a budget for yield loss associated with the first type of error, as well as DPPM associated with the second. This means that the testing of transceiver functions does not need to be based on accurate measurements that are to be compared against the corresponding targeted design specifications, resulting in overly expensive testing, without offering a competitive reward in the consumer market.

The statistical distribution of an amplifier's gain, representing one function in a transceiver, is given as an example in Fig. 4, where it is assumed that the gain must be higher than *min* (e.g., 26 dB), in order to meet a receiver's targeted sensitivity, and lower than *max* (e.g., 30 dB), in order to meet the linearity requirements (i.e., ability to withstand strong blockers). This statistical distribution is assumed to have been established over high volume of devices (e.g., a million units), and the probability for the outliers on the left side, representing amplifiers with production defects, is visually exaggerated for clarity. This figure illustrates that when the distribution of performance can be narrowed through the use of self-calibration and compensation, there is no need for accurate measurement of it during final testing. Instead, a crude measurement may be used, while the function of interest (an amplifier, in this case) is undergoing self-compensation, to determine whether it is defective or not. This serves to reduce the test costs as well as the
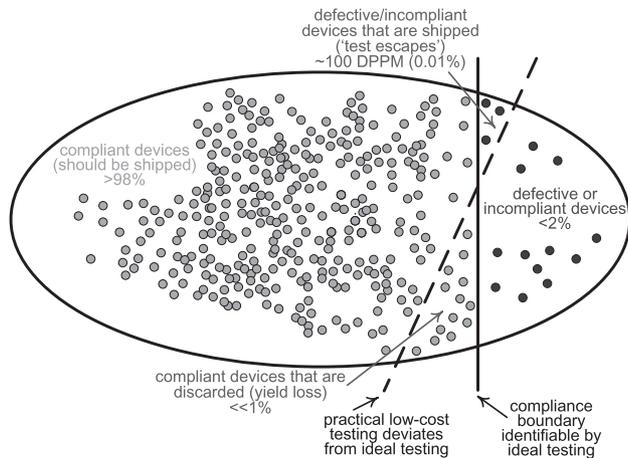
Fig. 5 Practical versus ideal production testing.

unnecessary yield loss shown in conjunction with the alternative parametric test, for which the limits *max* and *min* are applied (i.e., a strict specification-compliance test).

In essence, the goal of final testing should be limited to simply distinguishing between the populations of compliant devices that are to be shipped and those that are to be screened out, as shown in Fig. 5. The statistical targets suggested as examples in that figure may be realistic, but these generally depend on the cost/profit analysis for a given SoC. Such analysis is to be used to derive the budget for test-time and the targeted yield, and must also consider the quality (DPPM) committed to a customer.

Statistical analyses, demonstrating the validity of the crude-measurement approach for final-testing, are provided in [17] and summarized in the following subsections.

4.1 The Quality of a High-Volume Shipment (DPPM)

The number of defective or incompliant devices D in a shipment of large quantity S defines its quality and reflects on the reliability of the screening process performed at production by the SoC vendor. In the consumer market of low-cost SoCs, agreed-upon values for the fraction that D represents may be on the order of hundreds of defective parts-per-million (DPPM). This fraction, expressed as the ratio in (1), corresponds to the probability, denoted P_{sd}, for a defective device to pass the test and be shipped:

$$P_{sd} = D/S \qquad (1)$$

If the quality requirement is, for example, 100 DPPM, this probability must be limited to $P_{sd} \leq 10^{-4} = 0.01\%$. Possible reasons for the misidentification of a defective or incompliant device may be the test inaccuracy or the absence of a test that would detect the fault or incompliance in that device, which can occur whenever the test coverage is below 100%.

The behavior of a defective function, such as an amplifier, for example, is assumed to be substantially different from that of a normal one. If its typical gain is 20 dB, for

example, then a production defect is more likely to result in it exhibiting attenuation, rather than in a slight gain loss, thereby creating more than 20 dB of a difference and allowing a crude measurement to detect such fault. To meet the quality target P_{sd}, the probability for passing that device despite its defective block, denoted P_{pd}, must satisfy the inequality

$$P_{pd} = P_d \cdot (1 - P_{dd}) \ll P_{sd}, \qquad (2)$$

where P_d denotes the probability of that block being defective and P_{dd} is the probability of detecting such defect. If, for example, a defect in a low-noise amplifier (LNA) of a receiver's front-end occurs in one out of 10,000 devices (i.e., $P_d = 10^{-4}$) and the test procedure implemented for it will detect only $P_{dd} = 90\%$ of the defective LNAs, then $P_{pd} = 10^{-4} \times (1 - 0.9) = 10^{-5}$, which is an order of magnitude below the target of $P_{sd} = 0.01\%$. In this example, defective LNAs could account for up to 10% of the total population of defective devices D.

4.2 Yield Loss when Failing 'Good' Devices

This type of erroneous decision, often referred to as 'overkill', occurs when a device that is not faulty is failed, resulting in yield loss. This is typically the result of an inaccurate measurement or the consequence of applying an overly strict test limit in an attempt to account for measurement errors. Assuming low probabilities P_i ($i = 1 \ldots M$) for mistakenly failing each of the $M \gg 1$ individual tests for the various compliant blocks tested, the probability, P_{fc}, for failing a compliant device, may be approximated as the sum of the small probabilities in (3).

$$P_{fc} \cong \sum_{i=1}^{M} P_i \qquad (3)$$

Naturally, it is desirable for the production yield, or probability P_s that a device passes all the tests and is shipped, to be dominated by actual fabrication defects rather than by 'overkill'. This will be the case whenever inequality (4) is satisfied.

$$P_{fc} \ll 1 - P_s \qquad (4)$$

However, if the defect density is sufficiently low, allowing the targeted yield to be met with margin, the decision error probability P_{fc} may be allowed to dominate the yield, as expressed in the approximated equality (5).

$$P_s \cong 1 - P_{fc} \qquad (5)$$

If, for example, the yield target for a particular SoC is $P_s = 98\%$, but the fabrication defect density is very low, resulting in less than 0.5% faulty devices, then the probability of failing a compliant device may be allowed to slightly exceed $P_{fc}=1.5\%$. Hence, if an alternative, more accurate but costly test program could reduce this probability, it may not be an economically desirable solution and it may be more

worthwhile to allow the 1.5% 'overkill' in this example.

4.3 Defining the Pass/Fail Criteria for a Test

Traditionally, the pass/fail criteria for a particular RF/analog test has been linked to a corresponding parameter in the design specifications. Consequently, design targets often propagate into the test program, resulting in explicit measurements of performance parameters. Common examples of these are: receiver sensitivity, linearity (such as intermodulation effects and performance in the presence of interference), phase-trajectory-error (PTE) in GSM transmitters, and error-vector-magnitude (EVM) in complex transmitters. If large volumes reveal that a design target is not met at high enough probability, the committed performance in a product's data sheet may be relaxed, but would typically still maintain the traditional format of *min./typ./max.*, implying that devices that have passed the test must have exhibited performance within that committed range. The relaxation is designed to reduce the loss in production yield caused by incompliance with that targeted specification, while the reduction in committed performance is perceived as acceptable. However, with such approach, the test program is still forced to measure that performance parameter accurately, resulting in the need for accurate analog ATE, as well as the need to fail devices that do not meet the criteria, both of which impact profitability.

Contrarily, the approach described in [17] proposes the use of 'soft specifications', wherein the statistical distribution of a performance parameter, characterized over sufficiently large volumes, is communicated to customers. For example, a receiver's sensitivity may be within a certain range around the typical value with some 99% probability, and up to 2 dB inferior to that value with less than 1% probability. Although 1% is a much higher percentage than the agreed upon DPPM, these devices do not have to be considered defective and the representative quality of the customer's product, targeted at the consumer market, should still be satisfactory. This eliminates not only the waste associated with failing these slightly inferior devices, but mostly the high costs associated with identifying them in final-testing, a cost that may be considered 'money left on the table'.

5. Structural versus Parametric Testing

The costly parametric testing, still widely employed for RF/analog functions, is in great contrast to the mature structural approach adopted long ago in the testing of digital functions, wherein the circuitry is tested for production defects using scan-chains. The additional logic allowing these scan chains is automatically inserted as part of the digital design flow and is designed, with accompanying test-vectors, to allow for maximal coverage of the digital logic to which it is added. This means that almost all possible defects in the logic being tested, such as stuck-at-zero/one faulty nodes in it, could be detected. It is also common to run multiple scan chains in parallel in order to allow for faster testing of a large amount of logic. Thus, a digital function is deemed defective based on the integrity of the elements comprising it rather than through verification of its intended function (e.g., a multiplier is not tested by multiplying numbers in it).

There are various reasons why such structural approach has not yet become as prevalent for analog/RF circuitry. While a more limited set of building blocks exists in digital circuitry, and the definition and detection of faults in them may be relatively simple, the analog equivalent is considerably more complex. The higher complexity of this challenge in RF/analog circuitry, the relatively recent trend of integrating it into a digitally extensive SoC, and the common misconception as to what the goal of testing should be, all account for the lagging of a structural approach in analog/RF. It should be noted, however, that the development of fault models for analog circuitry has been of interest for many years [18], [19], and structural testing of analog/RF functions has been demonstrated. It is, therefore, anticipated that this trend will continue and as CAD tools for analog design progress, automatically-inserted structural built-in-self-testing may become part of the design flow, just as with digital circuitry.

An analog/RF function may be tested in a structural manner at a few different levels. These may be at the subsystem level, where a few different circuits are involved in a single test (e.g., an entire receiver front-end), at the circuit level (e.g., one amplifier or oscillator at a time), or at the component level (e.g., individual capacitors or transistors involved in each measurement). The choice of a particular level of structural testing should consider the required test resources, test time and test coverage. The next section provides examples of structural BiSTs, as well as BiST replacements for parametric tests that are traditionally performed externally.

6. The Implementation of Built-In Self-Testing (BiST)

An example for the benefits of a structural test performed at the component level may be seen in the testing of a digitally-controlled-oscillator (DCO) that has multiple banks of capacitors allowing it fine frequency resolution over a wide range, as described in [1]–[4].

A functional test of such oscillator, typically performed at very few frequencies, may not reveal a defect in one of the capacitors, which might manifest itself during normal operation. A component-level structural test, however, where each capacitor is tested individually, as described in [7], can serve to provide full coverage of the capacitor banks and would detect a defect that may occur in one of them without necessitating any parametric RF measurement.

As shown in [7], when a tested capacitance is toggled by the BiST software running on the internal processor, a digital periodic waveform is observed on the ADPLL's phase-error signal ('PHE' in Fig. 3). A crude examination of this waveform is sufficient to detect disconnected or shorted
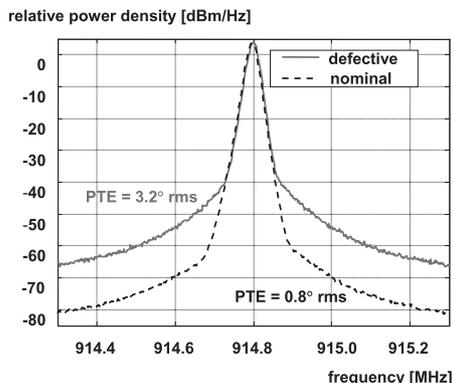
Fig. 6 Phase noise performance of a defective DCO.

capacitors, thus representing an efficient structural built-in test of 100% coverage.

An additional test described in [7] addresses the phase noise of the DCO. It is assumed that certain defects may result in degraded phase noise performance, which would be harder to detect than complete malfunctioning of the oscillator. For this test, the variance of the phase error signal, *PHE*, within the ADPLL, which is readily available in digital form, is computed by the transceiver's digital processor in order to determine the amount of noise exhibited by the RF oscillator. Thus, without requiring explicit frequency conversion and digitization functions, both of which are inherently provided by the ADPLL, the BiST algorithm is capable of quantifying the phase noise quite accurately and distinguishing between a defective DCO, shown in Fig. 6 to exhibit excessive noise, and a nominal one. This test, although based on a relatively accurate measurement of a performance parameter, may still be considered structural. This is because it targets specific circuitry that is responsible for assuring performance parameters that would otherwise be measured with external RF measurements. These parameters are the modulation accuracy of the transmitter, which is expressed in terms of the phase-error-trajectory (PTE), and the spectral mask compliance during transmission. Another performance parameter impacted by the oscillator's phase noise is the receiver's immunity to blockers, since the excessive phase noise of a defective DCO may result in reciprocal mixing that violates the receiver's targeted specifications.

In this example, rather than measuring the performance parameters of interest directly, which may be done using an expensive RF tester, the structural approach verifies the integrity of the circuitry responsible for assuring this performance. Once its integrity is verified, it is assumed that the sub-system would deliver the targeted performance without measuring it explicitly.

A few examples for BiSTs of parametric nature, which were designed to replace parametric tests that are performed externally, may be found in [9]. The common practice of using a loop-back test, where the locally available transmitter is used to test the receiver and vice versa, may be found there as well. It should be noted that when such test is used for testing the receiver's sensitivity, the injected signal

level may not be very accurately controllable, and it is also impractical to assume that a signal approaching sensitivity levels may be created in such manner. Nevertheless, such loopback test may effectively serve for crude testing, which should be sufficient in final testing, where the purpose of the test is merely to detect defects that significantly impact performance.

A 'design marginality' defines a situation in which the performance of a given circuit or system violates its targeted specifications at a probability higher than the afforded DPPM due to the natural distribution of performance resulting from the process variations. It is also assumed that the probability of failing the specifications is still low enough for the yield-loss to be affordable, e.g., 3% and not 60%, such that the motivation to redesign and refabricate is low (and, in practice, may not be affordable due to schedule constraints). In such scenario, if a test limit were to be designed that would screen out the incompliant devices, inaccuracies in it could result in significant overkill. This is because the slope of the distribution curve may be high when considering probabilities as high as a few percent, such that adding a margin that would prevent excessive test escapes could result in costly overkill. This scenario must be avoided even at the cost of considerable added area and design effort, as it defeats the purpose of low-cost testing and can significantly impact profitability in large volumes.

It should be stressed that design-for-manufacturability (DfM), which implies low-cost productization and high yield, is of paramount importance in consumer market high-volume SoCs, as profitability greatly depends on it.

At the earliest stages of the design of a potentially marginal circuit, built-in compensation mechanisms, allowing 'self-healing' of that circuit, must be incorporated, such that a parameter of interest could to be tuned into its targeted range over all process corners with a probability approaching 100%. Even the duplication of a circuit, with selection-switching built around it, as shown in Fig. 8, and discussed in Sect. 7, may be a cost effective solution. This is because the additional silicon area may have a negligible impact on the overall cost, whereas the previously discussed consequences of higher cost testing and yield losses may have a worse impact on profitability.

A post-fabrication approach to dealing with such scenario would be the redefinition of the performance specifications, which are anyway usually somewhat arbitrary numerical limits that presumably correspond to user experiences, as stated in [17]. As explained in Sect. 4, the newly defined soft specifications can correspond to the actual statistical distribution characterized in the fabricated device over a sufficiently large volume. This is less feasible when the performance parameter of interest is derived from a standard or a regulatory limit, although many of these are also set somewhat arbitrarily and can tolerate small violations of low probability.

7. Self-Calibration/Compensation

Several examples for built-in calibration/compensation may be found in [1]–[4], where parameters such as the DCO's modulation gain and the time-to-digital (TDC) step size are internally measured and calibration is realized digitally. Additional examples of similar nature may be found in [10]–[15], such as an oscillator's bias current that is adjusted to achieve optimal phase noise performance [10].

A different class of built-in compensation, targeting the minimization of the impact of self-interference, which is not a design parameter, is introduced in [13]–[15]. The concept of design-for-interference-mitigation (DfIM) is introduced in [13], suggesting that provisions of the type made for BiST and self-calibration, shown in Fig. 7, should be extended for the development of self-interference mitigation solutions (e.g., RF loop back, analog mux for receiver's analog-to-digital converters to allow for internal measurements of analog signals of interest, ability to control clock rates and phases, and ability to control various parameters in the ADPLL).

A different approach to self-compensation may be the selection of a circuit best suited for a particular process corner, as illustrated in Fig. 8. The system is assumed to be capable of measuring the performance parameter internally and selecting the best suited circuit. Alternatively, the system may be able to ascertain the process corner based on a measurement obtained in another circuit, and apply this information across all relevant circuits in the SoC. The area
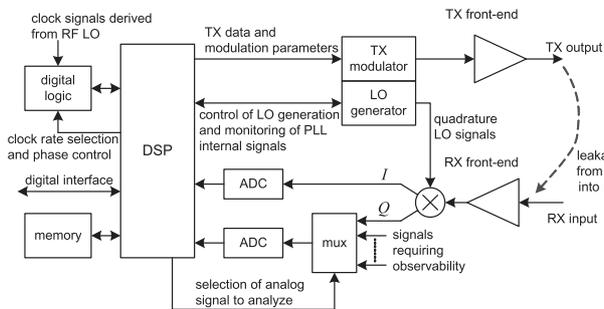

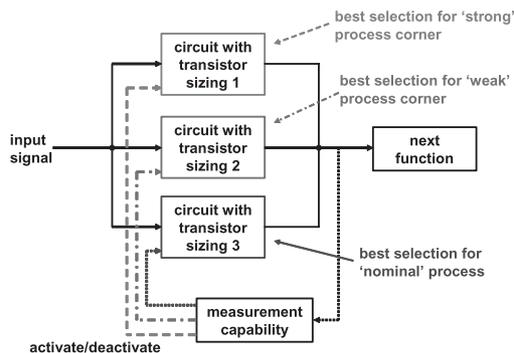**Fig. 7** Built-in testability and compensation provisions in a transceiver.

Fig. 8 'Self healing' based on automatic circuit selection.

penalty associated with this may be well worth the advantage of self-healing provided by such solution. There should be no power consumption penalty associated with the duplication of circuitry, as only one circuit would be activated, and no pin-count penalty, as this capability is based on fully internal circuitry.

8. Conclusion

The advances in transceiver SoC integration, fueled by the consumer market of high-volume devices, such as mobile phones, has allowed the introduction of revolutionary transceiver architectures, where many building blocks that were traditionally implemented in analog circuitry have been replaced with digital counterparts.

However, many of the functions in a typical transceiver continue to rely on the analog nature, for which performance variability and production testing may be challenging. Several examples were surveyed here, where the SoC's existing resources, as well as minimal dedicated circuitry that was added to allow testability, are used in order to accommodate different types of built-in measurements that address these challenges. Distinction was made between measurements targeting the self-compensation of a circuit, such as the self-adjustment of bias current in an oscillator [10], and those targeting the detection of defects in it as part of 'final testing' at production, such as the scanning of DCO capacitors [7].

In order to maintain reasonable profit margins and ensure cost competitiveness for a consumer market SoC, a design-for-manufacturability (DfM) approach must be employed, which would consider all aspects of self-compensation and testability for its analog circuitry. A design effort and associated silicon area must be allocated to this, as has been common practice in digital designs for many years.

Furthermore, it is advantageous to employ a soft-specifications strategy wherein the distribution of performance, determined through characterization over a large volume of devices, defines the performance specifications in a statistical fashion, thereby eliminating the need to accurately measure performance parameters against hard limits that are defined as per the traditional specification methodology.

References

[1] R.B. Staszewski and P.T. Balsara, All-digital frequency synthesizer in deep-submicron CMOS, John Wiley & Sons, New Jersey, ISBN: 978-0471772552, 2006.

[2] R.B. Staszewski and P.T. Balsara, All-Digital Frequency Synthesizer in Deep-Submicron CMOS, (Japanese translation), Japan: CQ Press (Semiconductor Series), ISBN: 978-4-7898-3122-2, 2010.

[3] R.B. Staszewski, K. Muhammad, D. Leipold, C.-M. Hung, Y.-C. Ho, J.L. Wallberg, C. Fernando, K. Maggio, R. Staszewski, T. Jung, J. Koh, S. John, I.Y. Deng, V. Sarda, O. Moreira-Tamayo, V. Mayega, R. Katz, O. Friedman, O.E. Eliezer, E. de-Obaldia, and P.T. Balsara, "All-digital TX frequency synthesizer and discrete-time receiver for Bluetooth radio in 130-nm CMOS," IEEE J. Solid-State Circuits,

vol.39, no.12, pp.2278–2291, Dec. 2004.

[4] R.B. Staszewski, J. Wallberg, S. Rezeq, C.-M. Hung, O. Eliezer, S. Vemulapalli, C. Fernando, K. Maggio, R. Staszewski, N. Barton, M.-C. Lee, P. Cruise, M. Entezari, K. Muhammad, and D. Leipold, "All-digital PLL and transmitter for mobile phones," IEEE J. Solid-State Circuits, vol.40, no.12, pp.2469–2480, Dec. 2005.

[5] A. Matsuzawa, "Digital-Centric RF CMOS Technologies," IEICE Trans. Electron., vol.E91-C, no.11, pp.1720–1725, Nov. 2008.

[6] S. Tanaka, "Evolutional trend of mixed analog and digital RF circuits," IEICE Trans. Electron., vol.E92-C, no.6, pp.757–768, June 2009.

[7] O. Eliezer, I. Bashir, R.B. Staszewski, and P.T. Balsara, "Built-in self testing of a DRP-based GSM transmitter," IEEE Radio Frequency Integrated Circuits Symposium, RFIC 2007, pp.339–342, June 2007.

[8] C. Mao, D. Mannath, V. Besong, O. Eliezer, and S. Larson, "The implementation of built-in self tests in a digital radio processor (DRP^{TM})," Proc. IEEE Workshop on RTL and High Level Testing, WRTLT'07, 2007.

[9] D. Webster, J. Cavazos, D. Guy, P. Patchen, and D.Y.C. Lie, "Structural verification of a WLAN system using simple BiSTs," IEEE Dallas Circuits and Systems Workshop (DCAS-10), 2010.

[10] I. Bashir, R.B. Staszewski, O. Eliezer, et al., "An SoC with automatic bias optimization of an RF oscillator," IEEE Radio Frequency Integrated Circuits Symposium, RFIC 2009, pp.259–262, June 2009.

[11] I. Elahi, K. Muhammad, and P. Balsara, "I/Q mismatch compensation using adaptive decorrelation in a low-IF receiver in 90-nm CMOS process," JSSCC, vol.41, pp.395–404, Feb. 2006.

[12] J. Mehta, B. Staszewski, O. Eliezer, et al., "An efficient linearization scheme for a digital polar EDGE transmitter," IEEE Trans. Circuits Syst. II: Express Briefs, vol.57, no.3, pp.193–197, March 2010.

[13] O. Eliezer, B. Staszewski, and P. Balsara, "A methodological approach for the minimization of self-interference effects in highly integrated transceiver SoCs," 2009 International IEEE Conference on Microwaves, Communications, Antennas and Electronic Systems, COMCAS 2009, Nov. 2009.

[14] O. Eliezer, B. Staszewski, I. Bashir, S. Bhatara, and P.T. Balsara, "A phase domain approach for mitigation of self-interference in wireless transceivers," IEEE J. Solid-State Circuits, vol.44, no.5, pp.1436–1453, May 2009

[15] I. Bashir, B. Staszewski, O. Eliezer, B. Banerjee, and P. Balsara, "A novel approach for mitigation of RF oscillator pulling in a polar transmitter," IEEE J. Solid-State Circuits, vol.46, no.2, pp.403–415, Feb. 2011.

[16] O. Eliezer, B. Staszewski, J. Mehta, F. Jabbar, and I. Bashir, "Accurate self-characterization of mismatches in a capacitor array of a digitally-controlled oscillator," IEEE Dallas Circuits and Systems Workshop (DCAS-10), 2010.

[17] O. Eliezer, R.B. Staszewski, and D. Mannath, "A statistical approach for design and testing of analog circuitry in low-cost SoCs" Proc. 53rd IEEE International Midwest Symposium on Circuits and Systems (MWSCAS 2010), 2010.

[18] M. Soma, "An experimental approach to analog fault models," IEEE 1991 Custom Integrated Circuits Conference, pp.13.6.1–13.6.4, 1991.

[19] C. Hoffmann, "A new design flow and testability measure for the generation of a structural test and BIST for analogue and mixed-signal circuits," Proc. 2002 IEEE Design, Automation and Test in Europe Conference and Exhibition (DATE '02), 2002.

[20] R. Staszewski, B. Staszewski, T. Jung, T. Murphy, I. Bashir, O. Eliezer, K. Muhammad, and M. Entezari, "Software assisted digital RF processor (DRP^{TM}) for single-chip GSM radio in 90 nm CMOS," IEEE J. Solid-State Circuits, vol.45, no.2, pp.276–288, Feb. 2010.

[21] R.B. Staszewski, I. Bashir, and O. Eliezer, "RF built-in self test of a wireless transmitter," IEEE Trans. Circuits Syst. II, vol.54, no.2, pp.186–190, Feb. 2007.

[22] R.B. Staszewski, J. Wallberg, C.-M. Hung, G. Feygin, M. Entezari, and D. Leipold, "LMS-based calibration of an RF digitally-controlled oscillator for mobile phones," IEEE Trans. Circuits Syst. II, vol.53, no.3, pp.225–229, March 2006.

[23] T. Tokairin, M. Okada, M. Kitsunezuka, T. Maeda, and M. Fukaishi, "A 2.1-to-2.8-GHz low-phase-noise all-digital frequency synthesizer with a time-windowed time-to-digital converter," IEEE J. Solid-State Circuits, vol.45, no.12, pp.2582–2590, Dec. 2010.

[24] K. Takinami, R. Strandberg, P.C.P. Liang, G. Le Grand de Mercey, T. Wong, and M. Hassibi, "A rotary-traveling-wave-oscillator-based all-digital PLL with a 32-phase embedded phase-to-digital converter in 65 nm CMOS," 2011 IEEE International Solid-State Circuits Conference, ISSCC 2011, pp.100–101, Feb. 2011.

[25] R.B. Staszewski, "Digital RF technology for expanding programmability of RF transceivers," SK Telecom Journal — Reconfigurable RF Systems, vol.20, no.5, pp.721–738, Oct. 2010.

Oren Eliezer received his BSEE and MSEE degrees in electrical engineering from Tel-Aviv University, Israel, in 1988 and 1996 respectively, and his Ph.D. from The University of Texas at Dallas (UTD), USA, in 2008. He served in The Israel Defense Forces from 1988 to 1994, where he specialized in wireless communications, and after his military service he co-founded Butterfly Communications in Israel and served as the company's chief engineer. Following Butterfly's acquisition by Texas Instruments (TI) in 1999, he was relocated to Dallas in 2002, where he took part in the development of TI's Digital RF Processor (DRP^{TM}) technology and was elected Senior Member of the Technical Staff. At TI he specialized in digital transceiver architecture, built-in testing, calibration and compensation, and mitigation of self-interference in RF SoCs. He has authored and coauthored over 40 conference and journal papers on these topics and over 30 issued/pending patents. He is currently the Chief Technology Officer at Xtendwave in Dallas, and participates in the research at the Texas Analog Center of Excellence (TxACE) at UTD.

Robert Bogdan Staszewski received his BSEE (*summa cum laude*), MSEE and Ph.D. degrees from The University of Texas at Dallas, USA, in 1991, 1992 and 2002, respectively. From 1991 to 1995 he was with Alcatel in Richardson, Texas. He joined Texas Instruments in Dallas, Texas in 1995. In 1999 he co-started a Digital RF Processor (DRP^{TM}) group with a mission to invent new digitally-intensive approaches to traditional RF functions. He has served as the CTO of the DRP group between 2007 and 2009. Since July 2009 he is Associate Professor at Delft University of Technology in The Netherlands. He has authored and co-authored one book, three book chapters, 110 journal and conference publications, and holds 70 issued and 40 pending US patents. His research interests include nanoscale CMOS architectures and circuits for frequency synthesizers, transmitters and receivers. He was elected IEEE Fellow in 2009 for his pioneering work on digital RF.