

Integrated Self-Healing for mm-Wave Power Amplifiers

Steven M. Bowers, *Student Member, IEEE*, Kaushik Sengupta, *Member, IEEE*,
Kaushik Dasgupta, *Student Member, IEEE*, Benjamin D. Parker, and Ali Hajimiri, *Fellow, IEEE*

Abstract—Self-healing as a technique for improving performance and yield of millimeter-wave power amplifiers (PAs) against process variation and transistor mismatch, load impedance mismatch, and partial and total transistor failure is described and investigated. A 28-GHz PA is presented with three types of sensors, two types of actuators, data converters, and a digital algorithm block that are all integrated on a single chip to show the validity of the technique. Two algorithms are implemented to either maximize output power or to minimize dc power for a desired output power. Measurements from 20 chips show increased RF output power up to 3 dB or reduced dc power by 50% in backoff with a 50- Ω load. Self-healing with up to 4-1 voltage standing-wave ratio load impedance mismatch is verified and linear operation under nonconstant envelope modulation is shown to improve with healing. Self-healing after laser cutter induced transistor failure is verified and increases RF output power by up to 5.4 dB. The aggregate yield of the PA across several representative specifications is increased from 0% to 80% with self-healing.

Index Terms—Actuators, CMOS integrated circuits (ICs), digital control, power amplifiers (PAs), power generation, sensors, thermal sensors.

I. INTRODUCTION

CONTINUAL advances in integrated circuit (IC) fabrication have opened up numerous new applications and design possibilities for millimeter-wave (mm-wave) systems that previously were not possible and/or not economically feasible [1]–[4]. In addition, improvements in power generation in silicon processes have made silicon power amplifiers (PAs) viable [5]–[10]. Making the PAs in a silicon process allows for greater integration with the rest of the transceiver and reduces the cost. These advances also come with new challenges because with every reduction in minimum feature size, as the industry moves to smaller and smaller process nodes, variation

Manuscript received October 23, 2012; revised January 11, 2013; accepted January 15, 2013. Date of publication February 13, 2013; date of current version March 07, 2013. This work was supported by the Air Force Research Laboratory. This paper is an expanded paper from the IEEE RFIC Symposium, Montreal, QC, Canada, June 17–19, 2012.

S. M. Bowers, K. Dasgupta, and A. Hajimiri are with the Department of Electrical Engineering, California Institute of Technology, Pasadena, CA 91125 USA (e-mail: sbowers@caltech.edu).

K. Sengupta was with the Department of Electrical Engineering, California Institute of Technology, Pasadena, CA 91125 USA. He is now with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08540 USA.

B. D. Parker is with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10598 USA.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TMTT.2013.2243750




Fig. 1. Block diagram of a generic mm-wave self-healing system.

between ICs, as well as between transistors within a single IC, continues to increase [11]–[13]. This is compounded by the fact that the digital processing market is the primary driving force behind this scaling, leading to foundries optimizing their device models mainly for digital use. Thus, models that are reliable at mm-wave frequencies are often not available early in the node's development stage. These transistor variations and model inaccuracies, as well as other sources of performance degradation, such as environmental variations, critically impact high-power mm-wave designs in nanometer-scale CMOS technologies, and reduce the yields of such designs. Environmental variations can be caused by many things, but one of the important issues for PAs is antenna load impedance mismatch. This occurs when the environment interacts in the near field of the antenna and changes the load impedance looking into the antenna. The ability of PAs to function properly under load impedance mismatch conditions becomes an especially critical issue when they are being used in a phased array [1], [2], [14]–[17]. In a phased array, the load impedance that needs to be driven by a PA changes when signal from other nearby elements of the array is coupled back through the antenna. Due to the fact that a mm-wave PA is tuned to provide the optimal impedance to the driving stage, any change in load impedance will be away from that optimal and will degrade the PA's performance [18].

Self-healing is a method that can mitigate these issues by identifying any degradation and modifying the circuit to improve its performance post fabrication with minimum overhead by incorporating a feedback loop into the system that takes advantage of the very low cost of digital processing available in these advanced processes, as shown in Fig. 1 [19]–[22].

In order to adjust the performance of the circuit post fabrication, actuators must be implemented that contain a sufficient actuation space to keep the circuit operating as close to optimum as possible under the expected variations and sources of performance degradation. In order to know how to set those actuators, the important performance metrics for the mm-wave circuit must be identified. Also, high-performance low-power reliable

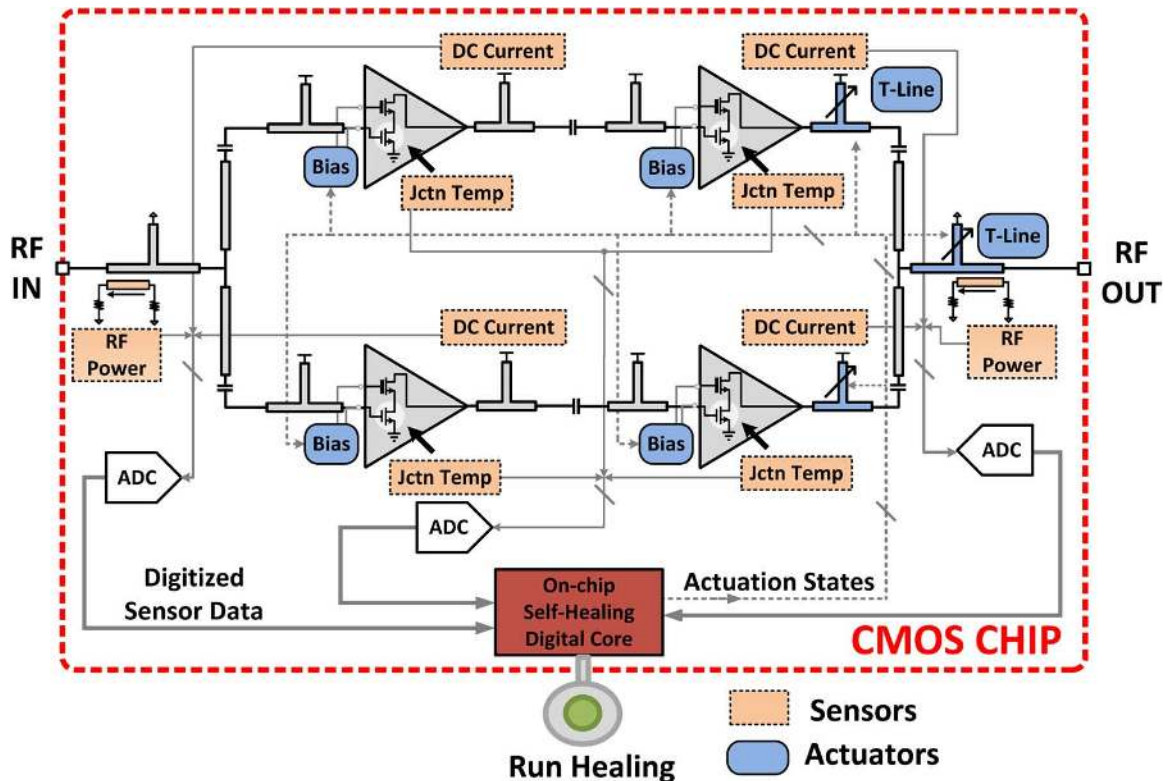


Fig. 2. Self-healing PA system showing amplifier and self-healing loop including sensors to detect performance, ADCs, an on-chip digital algorithm core, and actuators to bring amplifier to the optimum state after healing.

sensors must be implemented that can extract that information. Finally, a digital algorithm needs to be developed that takes in the sensor data and sets the actuation state, converging on the optimal state for the specific chip and current environment. An important aspect of the self-healing method is to integrate the entire self-healing loop, including data converters and digital processing, on a single chip because it is essential that the additional self-healing circuitry introduces minimal degradation of the RF performance with a low overhead in terms of area, power consumption, and connectivity to the outside world. Additionally, it is vital that the sensors are designed to be robust against the very same variations they are trying to detect since any feedback system is limited by the accuracy of its sensors, which is compounded by the fact that all calibration must be done automatically and internally within the chip to maintain viability in a mass production setting.

This paper describes the design of self-healing for mm-wave PAs, where on-chip self-healing improves output power, dc power consumption, and linearity under a broad range of non-idealities, such as process variations, modeling inaccuracies, load mismatch, and even partial and total transistor failure. We demonstrate the practicality of the technique in a 28-GHz self-healing PA system (Fig. 2). Two types of actuators are implemented: bias actuators to control the dc operating points of the amplifying stages, and tunable transmission line stubs to actuate the output power-combining network. Also incorporated into the design are three types of sensors: RF power sensors to measure the input and output RF power, and dc current and junction temperature sensors to measure the dc power

of the amplifier. On-chip data converters interface the actuators and sensors with the integrated self-healing digital algorithm core to complete the self-healing loop. The design of the PA itself is presented in Section II. The design and block-level measurements of the actuators, sensors, and data converters are discussed in Sections III–V, respectively, while Section VI discusses the integrated digital algorithm block. Finally, the system measurements are shown in Section VII.

II. PA DESIGN

The PA itself is a two-stage 2–1 power-combining class-AB amplifier, as shown in Fig. 2 [23], [24]. The input dividing network is matched to $50\ \Omega$, and the interstage and combining network are designed to present the optimal impedance to each stage for maximum power transfer at saturation, when the interstage networks are loaded with the gates of the output stages, and the output combiner is loaded with $50\ \Omega$. The transistors of the first stage are designed at half the width as those in the second stage to ensure that the first stage is capable of saturating the second stage, even when the loss of the interstage match is considered. Each amplifying stage is a cascode (CG) stage, and uses 56-nm analog transistors for the common-source (CS) transistors, and 112-nm-thick gate-oxide transistors for the CG. This is to increase the breakdown voltage, which enables larger voltage swings at the output of each stage. For the same output power, having a larger voltage swing will raise the optimal impedance looking into the power combiner, and thus lower the impedance transformation from the $50\text{-}\Omega$ output load, which reduces the loss of the combining network. A schematic




Fig. 6. Tunable transmission line stub capable of shorting out the signal to the ground at various points along the line.

changes to the reactive parasitics as well as load impedance mismatch requires a change in the matching networks themselves. This is a very difficult task, as the power levels going through the output matching network are very high, and any incurred loss can have significant impacts on the RF performance of the circuit. Also, the drain voltage swing on the transistors in the CG output stage is designed to be on the edge of breakdown for two stacked transistors, and the impedance match to $50\ \Omega$ only increases the voltage swing on the line near the load. On a $50\text{-}\Omega$ load, a peak power delivery of 16.5 dBm creates a peak-to-peak voltage swing of more than 6 V so any actuation transistors must be carefully placed in such a way that they do not experience breakdown when the PA is supplying maximum power. Also, because the power levels within the output combiner are at their highest levels for the entire circuit, any additional attenuation to the signal that is caused by the actuators will have the greatest impact on output power. Finally, under the above two constraints, the actuators still need to cover a large enough actuation space to cover the entire region of expected variation in the circuit, as well as externally.

Tunable transmission line stubs were implemented to provide the necessary actuation. To tune the stubs, transistor switches were used to short out the stub at various points to control the effective length of the stub (Fig. 6). The reason for tuning the stubs only is that by placing the actuation points on the stub near its ground, the voltage fluctuation on the signal lines when the transistors are off is kept within breakdown limits. One concern of using transistors as switches at these frequencies is the tradeoff between the on resistance and off capacitance. However, because the transmission line itself is a distributed reactive element, it can be modeled as a series of inductors and capacitors, the capacitance of transistors in their off state can be incorporated into the transmission line model, allowing for much larger transistors to be used, which reduces the on resistance.




Fig. 7. Measurement of a single tunable transmission line stub and its effective inductance at each state.

The on resistance of the switches is the dominant source of added loss due to the actuators, which means that keeping it very low keeps the performance cost of the actuators low. Turning all of the switches off increases the effective length of the transmission line to a maximum while turning them all on moves to the minimum effective length. Since the voltage swing on the line when all of the switches are off increases the closer it gets to the stub's junction with the rest of the matching network, actuation switches are placed from the shorted side of the stub to about 60% of the way to the junction.

There is another tradeoff between the achievable resolution of the effective length of the stub and the performance cost of the actuators, and for the described design, a series of seven actuation transistors along the line, and thus, eight different effective stub lengths produces adequate resolution to cover the actuation space while having minimal impact on the amplifier output power. An ideal shorted stub transmission line that is less than $1/4$ of a wavelength can be modeled as a parallel inductor to ground, and thus the actuation of each stub can be looked at as creating a variable inductor. Measurements were taken on a breakout cell of one of the tunable stubs, and show a wide range of tuning from an effective inductance of 25 pH up to 71 pH, as shown in Fig. 7.

When all three tunable stubs are incorporated into the entire matching network, they enable two dimensional control of the network, allowing matching from $50\ \Omega$ to an actuation space around the typical optimal impedance, matching from a variety of output loads in the case of load impedance mismatch back to the typical optimal impedance, or more practically, some combination of both. Fig. 8 shows the simulated actuation space of impedances looking into the combiner due to the tunable transmission line stubs with a $50\text{-}\Omega$ output load. Thus, if there are transistor variations that cause the optimal impedance for the amplification stages to move within that space, the actuators will be able to provide that impedance to deliver maximum power. Looking at the inverse case, Fig. 9 is the simulated actuation space of output load impedances that can still be transformed




Fig. 8. Actuation space of impedances that can be matched to a load impedance of 50Ω .




Fig. 9. Actuation space of load impedances that can be transformed to the typical optimal impedance of $(7 + 7j) \Omega$.

back to the optimal impedance of $(7 + 7j) \Omega$ of the output stage with typical transistors.

IV. SENSOR DESIGN AND MEASUREMENTS

In this section, we will discuss the design, implementation, and measurement results of individual sensors, which are responsible for on-chip characterization of the PA [25]. Fig. 2 delineates the placement of the various sensor blocks in order to measure the critical metrics such as output power, gain, and power-added efficiency. The RF power sensors, at the input and output ports, can measure true input and output power in the presence of load mismatch and have low insertion loss. This measurement enables us to self-heal for a significant variation of load mismatch events across the Smith chart, as we will discuss in detail in the measurement section. In order to measure drain and power-added efficiency, we propose direct measurement of the dc current drawn by the PA during operation, while simultaneously performing supply voltage regulation with ultra-low (≈ 30 mV) headroom. This ensures that the efficiency of the PA is not sacrificed during sensing and regulation. We will also discuss how parameters such as efficiency and dc power drawn

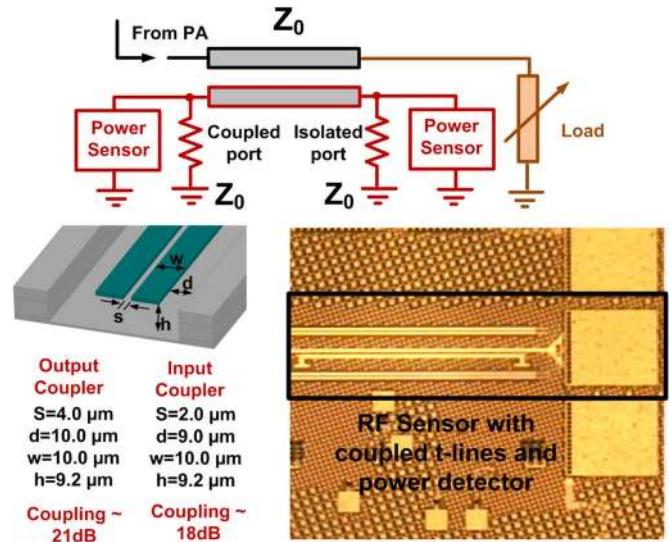


Fig. 10. RF sensor showing coupled transmission lines with power detectors at the coupled and isolated ports.

can be measured in the nonelectrical domain such as local temperature change around the PA core during operation. Since temperature change is directly correlated with dc power dissipated, this allows a low overhead measurement of PA operation. In totality, there are 12 sensor outputs, which are multiplexed and digitized by three 8-bit successive-approximation-register (SAR)-based analog-to-digital converters (ADCs) that are discussed in Section VI.

The critical performance metrics of the sensors are their dynamic range, responsivity, response time, and power consumption. The responsivity along with the resolution of the ADCs determine the minimum change detectable. The response time of the sensor is determined by our desired maximum healing time, which entails an iterative run through all possible actuation states and reading sensor data corresponding to all actuation states. The response time of the sensors was required to be less than $1 \mu\text{s}$ to ensure accurate measurements even with the digital running at 50 MHz. The power consumption of all the sensor blocks is less than 6% of the nominal power consumption of the PA.

A. Sensing Input and Output RF Power

The RF sensors, as shown in Fig. 2, are responsible for measurement of true RF power at the input and output. Ordinarily, a voltage rectifier that measures the RF swing at the output can serve as a power detector, as long as the load remains constant. However, in case of load mismatch events, the voltage swing is not necessarily proportional to the RF power delivered to the load. In order to facilitate measurement of both forward and reflected waves that can occur in such cases, two $50\text{-}\Omega$ coupled transmission lines are implemented at the input and output ports, as illustrated in Fig. 10. The coupled and isolated ports are matched to 50Ω at the input of two RF power detectors, as shown in Fig. 10. In case of a load mismatch event, the true power delivered to the load can be calculated from the power detected at both the ports. If P_C and P_I are the power detected at the coupled and isolated ports, then power delivered

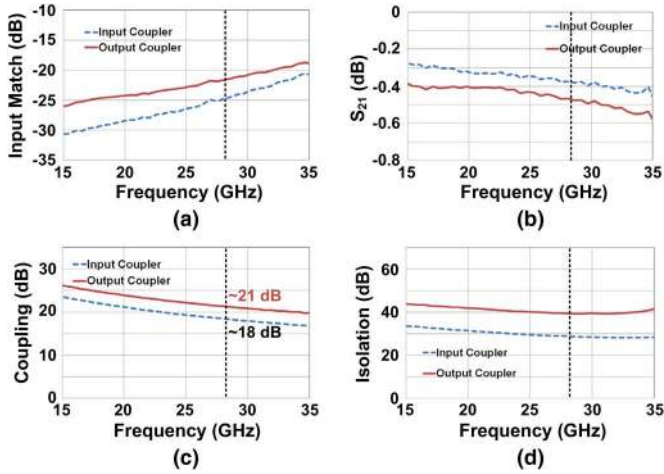


Fig. 11. Measured S -parameters of the input and output RF sensors showing: (a) input match, (b) through, (c) coupling, and (d) isolation.




Fig. 12. (a) Power detector schematic showing rectification, low-pass filtering, and current amplification. (b) Simulated predicted and through power of the RF power sensor.

to the load, neglecting loss in the transmission line, is given by $P_L = (P_C - P_I)/k$, where k is the coupling coefficient. The sensors were designed to be compact with low-insertion loss by keeping the length of transmission line couplers only $220 \mu\text{m}$ ($\approx \lambda/20$ at 28 GHz). The coupling coefficients of 18 and 21 dB for the input and output couplers, respectively, ensured only a very small fraction of the RF power was sensed at the input and output ports, so as to not affect output power and gain. The transmission lines were implemented in the top $2.1\text{-}\mu\text{m}$ -thick aluminum layer in a ground tub structure, as illustrated in Fig. 10. Fig. 11 shows the measured S -parameters of the input and output sensors. The input match stays below 18 dB and the isolation is better than 29 dB between 15–35 GHz. The insertion loss was measured between 0.4–0.5 dB at 28 GHz.

The power detector circuit at the coupled and isolated ports is demonstrated in Fig. 12. A $4\text{-}\mu\text{m}$ -wide input transistor $M1$ biased at cutoff with a $50\text{-}\Omega$ resistor for input match, rectifies the incoming RF power. The rectified current is amplified in the current domain, after low-pass filtering and the analog sensor output is measured across a resistor and subsequently digitized by the on-chip ADC, as shown in Fig. 12. The RF power sensors are characterized on the full chip by varying the input power and measuring the response of all the digitized sensor data corresponding to the detectors at the input and output couplers.




Fig. 13. Measured analog: (a) output and (b) input sensor response at 28 GHz over six chips, enumerated from the on-chip ADC data.

The results measured over six dies are shown in Fig. 13. The input and output sensors show respective mean responsivities of 54 and 8.3 mV/mW , implying that the LSB of the sensor outputs correspond to an RF power change of 55 and $360 \mu\text{W}$, respectively. Since the analog responsivity is enumerated back directly from the measured ADC characteristics, the measurement spread include both the RF sensor and ADC variations. It should be noted that the self-healing PA is not calibrated against a known RF power source at any point of time. Therefore, in order to ensure that the healing algorithm converges to at least the desired absolute output power, the level is set to 1.5σ above the mean sensor data corresponding to that desired output power. This results in the PA correctly self-healing to at least the desired output power level with 90% confidence. As can be seen from Fig. 13, the sensor response is monotonic over its entire range and the 3σ spread measured over six chips, which includes the variation of the data-converter, is 2.2 and 1.1 dB for the input and output sensors, respectively. The dc power drawn by the sensor depends on the output power level and can consume a maximum of 1.2 mW.

B. Sensing DC Current Drawn by the PA

One of the operable modes of the self-healing algorithm seeks to find the optimum actuation state, which reaches the desired RF output power with maximum possible efficiency. Measurement of the dc component of the current drawn by the mm-wave PA during operation is challenging [26]. Mirroring the PA current through a sense transistor will require minimizing mismatches through careful layout techniques, which, in turn, would increase interconnect lengths and parasitics, degrading performance at high frequencies. Additionally, since the PA may be operating in the large-signal region, the drain voltages of the mirror transistor will also need to follow that of the PA at mm-wave frequencies to accurately mirror the current.

Therefore, two functions of supply regulation and current sensing are accomplished simultaneously through the dc current sensing circuit, as shown in Fig. 14. The current drawn by the PA passes through the 1-mm -wide transistor M_1 , which is scaled down by a factor of 100 and mirrored through matched transistor M_2 , as shown in Fig. 14. Since any increase in overhead voltage leads to lower efficiency, M_1 and M_2 are kept in a deep triode region with a headroom of 20–30 mV over a 2.12-V supply. In order to ensure accurate scaling, the source nodes of




Fig. 14. DC sensor which accomplishes dynamic current sensing of the PA and supply regulation simultaneously.




Fig. 15. Measured supply regulation over the PA bias actuation range showing the overhead varying from 10 to 40 mV over a 2.12-V supply.

M_1 and M_2 are held at equal potential with the operational amplifier A_1 . The voltage at the PA drain is set to a controllable reference voltage V_{ref} . The current through M_2 is converted into the dc sensor voltage by passing through a resistor, as shown in Fig. 14.

Fig. 15 shows the measured supply regulation over the entire PA bias current actuation range. The drain voltage of the PA is held within a 30-mV range, implying a overhead varying from 10 to 40 mV over a 2.12-V supply. The sensor performance is measured by changing the bias actuation settings of the PA (input or output stage) and measuring the digitized data. Fig. 16 shows the measured sensor response across the bias actuation range of the output stage over five chips. As before, the spread includes variations from both the on-chip sensor as well as the ADC. The sensor response is monotonic over the entire range with a mean responsivity of 4.2 mV/mA. Since the self-healing PA is never calibrated with a known current source, the spread of the actual bias current for a given sensor reading is important in on-chip evaluation of the absolute value of efficiency. As can be seen from Fig. 16, the 3σ spread of the sensor with the ADC is less than 14% over all measured chips. In order to monitor individual currents drawn by the stages, separate sensors accompany each of the input and output stage. Each sensor consumes less than 3 mW of dc power.

C. Local Temperature Sensing to Measure PA Efficiency

The dc and RF power sensors, as described above, can adequately measure on-chip gain, output power, and efficiency




Fig. 16. Measured analog response of the dc sensor across the PA actuation range over five chips, enumerated from the on-chip ADC data.


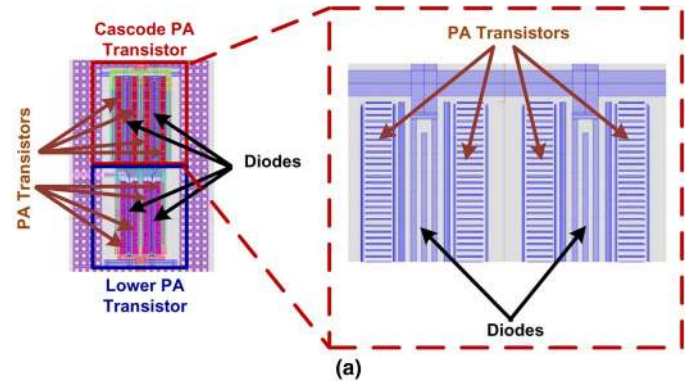


Fig. 17. (a) Layout of the PA transistor with the interspersed diodes in between and (b) the simulated thermal profile for 80-mW power dissipation in the process.

of the PA. However, the effects of finite efficiency of the amplifier can manifest itself in nonelectrical domains such as thermal. PA efficiency is directly related to the fraction of the dc power drawn by the PA that gets dissipated across the transistor channel. As a result, during operation, local temperature around the PA core is likely to rise in proportion to the dissipated power. Therefore, instead of directly measuring the PA current, a low overhead method of measuring PA efficiency is used to sense the local temperature rise during the PA operation. Fig. 17 illustrates the concept. Both the CS and CG PA transistors consist of four parallel active cores with p-n junction thermal sensor diodes laid in between. In order to predict accurately the temperature change during PA operation, thermal simulations were carried out using Ansoft ePhysics 2.0. The silicon substrate is an excellent thermal pathway; however, metal connections and vias also affect the temperature change, and therefore need to be incorporated in the simulation, as shown in Fig. 17 [27]. Simulations indicate that during PA operation, a thermally active region of 20–30

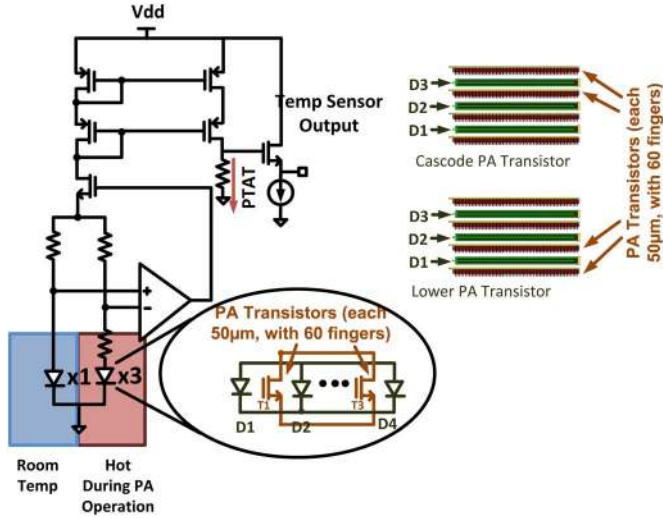


Fig. 18. Schematic of temperature sensor showing “hot” sensor diodes interspersed within PA transistor.




Fig. 19. Thermal sensor test structure response to: (a) global temperature change and (b) power dissipation through local temperature change due to transistor heating. The same expected change in temperature changes the output by almost $10\times$.

μm surrounding the amplifier core, which houses the sensor diodes, can experience a rise in temperature of nearly 10°C .

The placement of the diodes within the PA and the sensor schematic is shown in Fig. 18. As can be seen, the reference diode $x1$ is maintained at a relatively constant temperature during PA operation by being placed outside the local thermally sensitive region around the PA core ($\sim 40\ \mu\text{m}$). This leads to a larger difference between the diode voltages for a given PA temperature rise, leading to almost a factor of a $10\times$ increase in responsivity. This is illustrated in Fig. 19, which shows that the sensor output is almost $10\times$ more sensitive to local temperature change during PA operation (reference diode colder than sensor diode) as compared to global temperature variation (reference diode same temperature as sensor diode). The measured results are in close agreement with predicted results from thermal and circuit simulations. As can be seen in Fig. 20, the thermal sensor exhibits a similar response to the dc sensor demonstrating a measured responsivity of $2\ \text{mV/mW}$ of dissipated power under a 2.1-V supply. The response of the thermal sensor is, however, limited by the thermal time constant, which can be of the order of tens of microseconds. Therefore, during the process of self-healing, the dc power drawn by the PA is measured using the dc current sensor for a faster response. However, the thermal sensor gives a low power method of constantly monitoring excess power dissipation.




Fig. 20. Temperature sensor and dc sensor outputs after zero current offset calibration show a similar response with increasing dc power dissipation in a single stage of the PA.

Each thermal sensor draws less than $1\ \text{mW}$ of dc power. The sensor performances are summarized in Table I.

V. DATA CONVERTERS

Data converters are necessary to interface between the analog sensor data and the digital core and back again to the dc operating point actuators (the inputs of the tunable transmission line stub actuators are switches, and thus do not need digital-to-analog conversion), and thus both DACs and ADCs were designed and implemented on-chip.

A. DAC

In the present design, bias control was implemented both for the CS, as well as the CG transistors through DACs. These DACs need to be designed to be extremely low power so as not to add to the self-healing overhead. In addition, these DACs need to be able to drive the transistor loads at high speeds to keep the healing time to a minimum. The DACs are implemented as switchable transistor-based current sources through a load resistor. Transistors were sized in a binary fashion and laid out in a common-centroid fashion to ensure good matching, as well as monotonicity. Both kinds of DACs are implemented using long channel-length transistors. The DAC for the CS transistor has a range of actuation from $450\ \text{mV}$ to $1.05\ \text{V}$, whereas the DAC for the CG transistor has a range from 1.1 to $1.95\ \text{V}$. Each DAC has a 6-bit control and the CS DAC has one extra control bit, which sets the voltage to 0, thereby turning off the PA. This is essential when the digital core calibrates all the sensors. Fig. 21 shows the schematics of the implemented DACs.

The DACs were measured by programming them through the digital core and probing the dc voltage at their output. Both types of DACs were verified to operate until $25\ \text{MHz}$, which was limited by the test setup. Fig. 22 shows measurement results for the CS and the CG DACs.

B. ADC

On-chip ADCs have been implemented to digitize the analog sensor outputs for use with the digital healing algorithm. These ADCs need to be low power designs to reduce the overhead due to self-healing; at the same time they need to operate at high enough speeds so as not to slow down the digital healing algorithm, which, in turn, affects the healing time.

TABLE I
PA SENSOR PERFORMANCE SUMMARY

Sensors	Measured Entities	Responsivity	Range	Sensor 1-b resolution
RF Power	True In/Op Power	54 mV/mW (in) 8.3 mV/mW (op)	0-10 mW (in) 0-100 mW (op)	55 μ W (in) 360 μ W (op)
DC Sensor	DC drawn by PA (in, op stages)	8.5 mV/mA (in) 4.2 mV/mA (op)	0-60 mA (in) 0-120 mA (op)	350 μ A (in) 715 μ A (op)
Thermal Sensor	Power dissipated by PA	4.0 mV/mA (in) 2.0 mV/mA (op)	0-130 mW (in) 0-260 mW (op)	0.75 mW (in) 1.5 mW (op)




Fig. 21. Binary-weighted current-mode DAC schematic.




Fig. 22. Measured response of: (a) CS and (b) CG DACs.

Flash ADCs [28], [29] have traditionally been the fastest, but the power consumption overhead, as well as usable resolution, limits use of such ADCs in integrated self-healing systems. Pipelined ADCs [30] offer high sample rates while operating under relatively low power; however, the area constraint can limit their use in these applications. The SAR-based ADCs [31] are ideally suited for medium data rate applications while consuming low power. In addition to that, due to the inherent serial data output feature of the SAR ADCs, it makes it ideal for self-healing systems where the overhead due to routing multiple high-speed digital signals across the chip, specifically from sensors to the digital core, also needs to be minimized.

To accurately reflect the behavior of the sensors, a resolution of 8 bits was chosen for the SAR ADCs. Data from four dc sensors, four thermal sensors, and four RF sensors were multiplexed and fed to three ADCs placed throughout the chip. The digital core controls the sensor read order, as well as initialization of the digitization sequence. Fig. 23(a) shows the block diagram of the implemented SAR ADC. The external digital in-

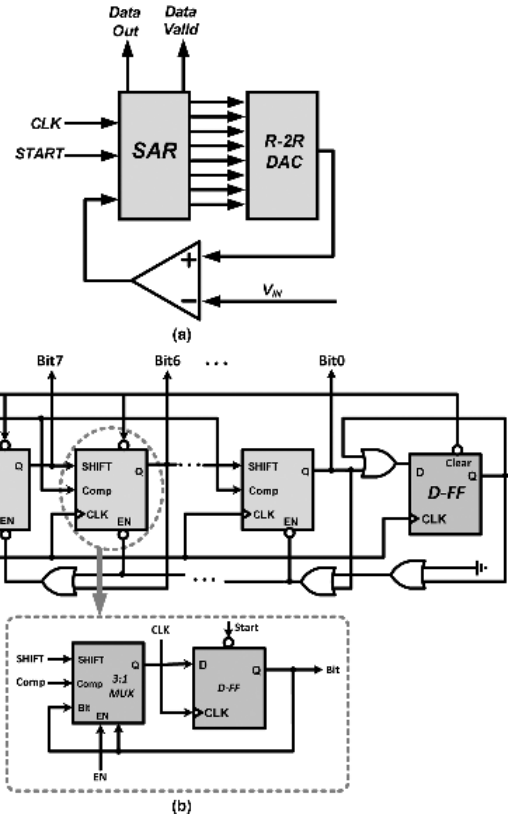


Fig. 23. (a) Block diagram of SAR ADC. (b) Synchronous SAR block.

puts to the SAR block are the global clock and the START signal from the core. Output of the ADC is a serial data output line and a data valid line, which informs the core that the digitization sequence has been completed. An R-2R based DAC is used to enable the successive approximation. Layout of the R-2R DAC was done in a common centroid fashion to improve matching accuracy. The input clock gets divided by 2 to generate quadrature clocks for the data output chain and the clocked comparator. This is to make sure that the input to the comparator does not change over the comparison window itself. The external START signal also gets retimed using this divided clock.

A typical successive approximation block is composed of a sequencer and a code register. The sequencer sets the code register to a state that gets set or reset on the next clock cycle depending on the comparison result. The problem with such a design is that although the sequencer is operating synchronously, the code register is asynchronous, leading to potential issues commonly found in asynchronous circuits, such as glitches, bit swallowing, etc. Fig. 23(b) shows the detailed implementation




Fig. 24. Simulated voltages for an input voltage of 400 mV. Data output is decimal 73 for an 8-bit output.




Fig. 25. (a) ADC measured and simulated characteristics ($V_{refn} = 350$ mV, and $V_{refp} = 950$ mV). (b) Measured DNL.

of the fully synchronous SAR block [32] used in the design. Eight MUX-based flip flops, one for each output bit and one extra flip-flop, for signal completion are used.

Fig. 24 shows the simulated voltages for an input voltage of 400 mV corresponding to higher and lower reference voltages of 900 and 200 mV, respectively. Note that only the first 8 bits after the DATA valid goes high are read by the digital core.

The ADC was measured as part of the full system where the sensor output voltages were probed and the corresponding data read out through the digital core. Due to test setup limitations, the operation of the ADC was verified up to 25 MHz, which translates to 2.5 Ms/s for an 8-bit SAR with start and data ready bits. Fig. 25(a) shows the characteristics for the ADC. The results show good agreement between measurement and simulation. Fig. 25(b) shows the measured dynamic nonlinearity (DNL) of the ADC. Average DNL of -0.04 LSB and a worst case DNL of -0.605 LSB were measured which ensures that the ADC reads are monotonic. The ADC draws 1.6 mW from a 1-V supply.

VI. DIGITAL ALGORITHM

An integral part of any self-healing system is a global digital controller, which can be implemented either as a microcontroller or a synthesized digital block. The algorithm for this design was implemented as a custom digital block that was coded with VHDL and synthesized on chip, which communicates with the actuators and the ADCs. The on-chip healing algorithm is made up of a global state machine that controls component blocks for actuation load, sensor read, and optimization.

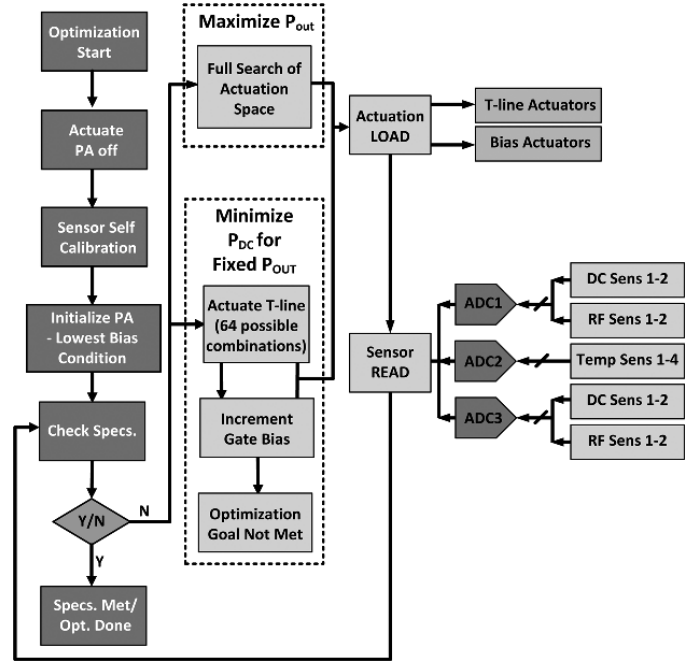


Fig. 26. Flowchart showing details of optimization component.

Due to this modular code setup, many different types of complex optimization algorithms can be incorporated into this general fully integrated self-healing framework.

Fig. 26 shows a basic flowchart of the optimization component synthesized for this design along with the connections to the actuators, sensors, and ADCs. Two optimization algorithms have been synthesized in the current implementation, which find the optimum solution within the search space of 262 144 possible states. The first finds the state that produces the maximum RF output power, and the other finds the state with the lowest dc power for a given output power. Both algorithms start off with an offset calibration step where the PA is turned off digitally and all the sensor readings are recorded via the ADCs. This information is used as a self calibration step that removes any offsets in the sensor/ADC chain. The algorithms then continue as they read data from the sensors, compute the RF output power and dc power consumption, and change the actuation states to achieve the desired performance.

The digital core uses a test setup limited clock of 25 MHz and takes $3 \mu\text{s}$ per optimization iteration (set actuators, read all sensors, decide on next actuation state), which leads to a maximum healing time of 0.8 s in the case when the algorithm goes through every single actuation state.

VII. SYSTEM MEASUREMENTS

To measure the entire PA self-healing system, the chip was mounted on a printed circuit board (PCB) and probed in the setup depicted in Fig. 27. In order to evaluate the benefits of the self-healing, a nominal actuation state was selected as the default state without self-healing. That state was chosen to be the one that had best performance at saturation in simulation with typical transistors, and represents the optimal state with the information the designer has at the time of fabrication. The dc power consumption of all of the self-healing blocks is omitted




Fig. 27. Measurement setup: PA chip is mounted on a PCB and probed. Calibrated mm-wave load tuner can set load impedances up to the 4-1 VSWR circle at the tips of the probe.




Fig. 28. (a) Self-healing for maximum output power at small signal and near 1-dB compression point. Histograms of 20 chips in: (b) back-off and (c) near 1-dB compression show improvement in output power, as well as decrease in variation between chips with healing.

when considering the performance of the default state, to obtain the best estimate of the performance of the amplifier without self-healing. Since the entire self-healing loop is integrated on the CMOS chip, to run the self-healing algorithm, a mode of operation, along with the desired output power (if necessary), and the “run” command are the only things sent to the chip. The entire rest of the process is locally automated within the chip.

A. Self-Healing Measurements With 50- Ω Load

First, we will examine the measurement results when the algorithm to maximize output power was run when the output load impedance was 50 Ω . Self-healing is run twice, once at small signal, and again near the 1-dB compression point of the amplifier, and the input power is swept for the resulting state in each case, as shown in Fig. 28. The optimal states that the algorithm found at different power levels is different, showing that there is not a single optimal state that is valid over all power levels. This gives the self-healing PA a unique advantage because we do not need to design it specifically for a single power level, but as long as the optimum is within the actuation space, we can let the algorithm find the optimum state for whatever power

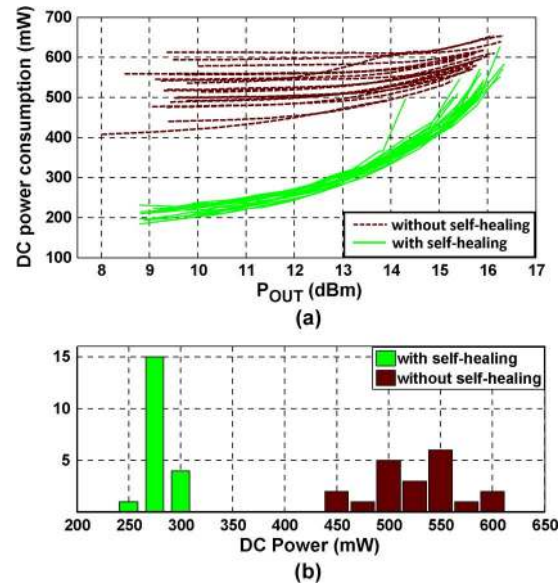


Fig. 29. (a) Self-healing to minimize dc power while maintaining specified output power for 20 chips with: (b) a histogram when $P_{OUT} = 12.5$ dBm.

level it is currently operating at. After healing the PA achieves a small-signal gain of 21.5 dB, with output powers of 16 dBm at saturation and 12.5 dBm at the 1-dB compression point while drawing 520 mW of dc power at 28 GHz.

Since the default state without self-healing is designed for an amplifier near saturation, the default state is closer to optimal near saturation, and thus there is not as much room for improvement. The measurements show that in the back-off region, where the default state is further from optimal, the ability of the system to self-heal the PA increases. To evaluate the viability of the system over process variation, 20 chips were measured, and histograms, showing operation at small signal and near the 1-dB compression point, are plotted in Fig. 28.

Next, the algorithm to minimize dc power while maintaining a specified output power is shown for 20 measured chips in Fig. 29. For each individual chip, we see as expected that the dc power consumption before self-healing does not change much until the amplifier begins to saturate. Variation between each chip, however, is significant even at low power levels. When the self-healing algorithm is turned on, significant dc power consumption savings are observed. When trying to achieve very high levels of output power, dc power is reduced compared to the same chip before self-healing, but large power levels are still required. However, further into the back-off region, the power savings increase, and the reduction in variation between the chips is dramatic. Near the 1-dB compression point with an output power of 12.5 dBm, there is a 47% reduction in average power level, and a 78% decrease in the standard deviation between chips. At small-signal levels below 9 dBm output power levels, reductions of greater than 50% for every measured chip are observed.

B. Self-Healing Under Load Mismatch

To imitate the effects of antenna load impedance mismatch, a focus microwaves mm-wave load tuner [33] capable of




Fig. 30. Contours before and after self-healing for maximum output power under load impedance mismatch.




Fig. 31. Histograms of ten chips verify self-healing across multiple devices. One is near the: (a) optimum load and the other is on the (b) edge of the 4–1 VSWR circle.

producing loads within the 12–1 voltage standing-wave ratio (VSWR) circle was used. When the measurement setup is calibrated with the loss of the connecting cable and probe, it is capable of producing impedances at the probe tips within the 4–1 VSWR circle. Since the RF power sensors are capable of detecting the power delivered to the load regardless of its impedance, the algorithm does not need to determine what the load impedance is, as long as it can optimize the desired performance metric. Measurements were taken by sweeping through the achievable load impedances, and at each impedance, running self-healing for each algorithm. The results for the algorithm to maximize output power are presented as a contour plot on Smith charts before and after self-healing in Fig. 30. There is improvement across the entire 4–1 VSWR circle, e.g., the 15.5-dBm contour just barely shows up before healing, but occupies a large area of the Smith chart when self-healing is turned on.

To quantify self-healing for load impedance mismatch over process variation, ten chips were measured, and two histograms are plotted in Fig. 31. The first is taken with a load impedance of 33Ω , where the output power is at a maximum, and the second is on the edge of the 4–1 VSWR circle at $(23 - 44j) \Omega$.

The measurement results of the second algorithm to minimize dc power while maintaining a specified power are shown on the two Smith charts in Fig. 32. The specified output power was set to 12.5 dBm, and the outermost contour of each plot shows the boundary where that power specification is met in each case. Since there is only one actuation state without self-healing, the dc power level does not change with the load impedance, but when self-healing is applied, substantial savings of up to 35% are achieved, especially near load impedances around 50Ω , the impedance the match was designed for.

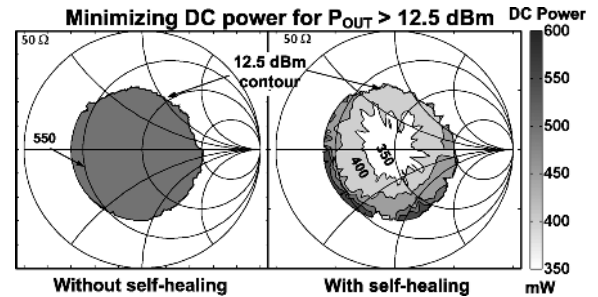


Fig. 32. Contours before and after self-healing for minimum dc power while maintaining a specified 12.5-dBm output power.

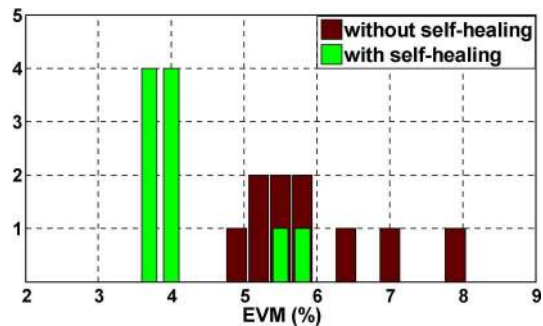


Fig. 33. Histogram showing linearity improvement measuring the EVM for ten chips at 12.5-dBm output power and 100-ks/s 16 QAM modulation before and after self-healing.

C. PA Linearity Measurement

One reason that a linear PA was chosen for this design was to enable amplification of signals with nonconstant envelope modulation. To verify the linearity of the PA and its operation with nonconstant envelope modulation, the error vector magnitude (EVM) of a 100-ks/s 16 quadrature amplitude modulation (QAM) signal was measured. A histogram of the ten chips measured shows EVM when the PAs are outputting 12.5 dBm, showing a reduction in the average value of the EVM from 5.9% down to 4.2% (Fig. 33). The improvement in the linearity after self-healing is attributed to being able to provide the same output power without forcing the PA as far into saturation.

D. Self-Healing Laser-Induced Transistor Failure

Finally, self-healing was verified in the case of partial and total transistor failure. To achieve this, various parts of one of the two output stages of the amplifier were gradually cut out. First, half of the CS transistor was cut, then half of the CG transistor was cut, and finally the entire stage was cut. At each stage of transistor failure, self-healing to maximize output power was run, and the resulting output power plotted against input power is shown in Fig. 34, with close-up photographs of the output stage before and after laser cutting shown in Fig. 35.

The first trend to note is that as more and more of the transistor is cut out, the performance continues to degrade; however at the same time, the ability to heal that degradation is increased across all input power levels. This is because the designed bias points and matching network are no longer near the optimal state for the transistors after being cut. When the design is further from optimal, the ability to heal that nonoptimality is increased. For

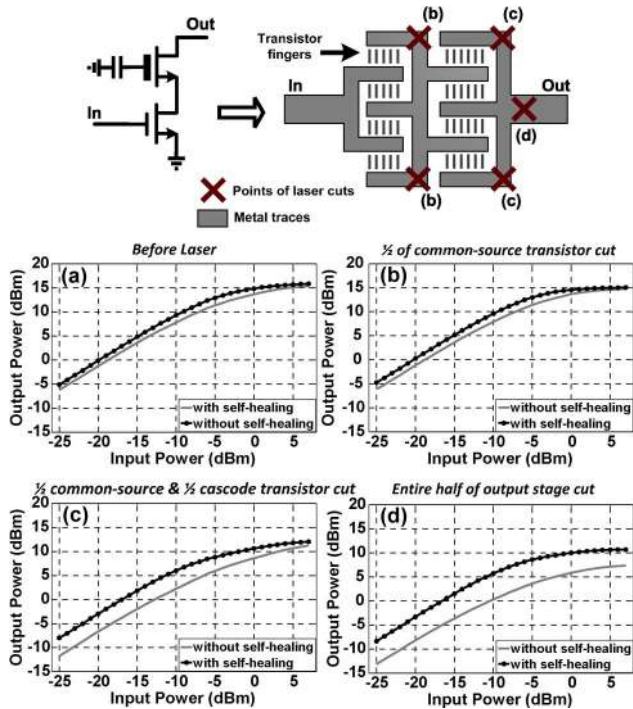


Fig. 34. PA output stage showing the various points of the transistors that were laser cut to verify self-healing for transistor failure, and the measured output power improvement with self-healing at each level of transistor failure.




Fig. 35. Chip photograph showing functional blocks, and close-up of the output stage before and after laser cutting.

example, at small signal, the PA without healing loses 7.2 dB in output power from before the lasers were used to after the entire stage was cut. 3 dB of that loss is because only one of the two output stages is proving power, but an additional 4.2-dB loss can be attributed to mismatch in the output combing network once the stage is cut. The loss in output power in both cases once self-healing has been run is only 3.3 dB. Taking into account the 3-dB loss in available power with only one output stage, there is only a 0.3-dB loss due to mismatch, even in this extreme case where one half of the network that is designed to be driven is just floating as an open stub. Also, because the impedance match is nonoptimal both at small signal and at saturation, the improvement due to self-healing is seen at all power levels, with an improvement in saturated power of 3.3 dB.

E. Yield Improvement and Performance Summary

In order to look at yield improvement, metrics and yield specifications must be defined. For this PA, the metrics are saturated output power, small-signal gain, power-added efficiency, and 4–1 VSWR tolerance, which is defined in decibels as the maximum output power falloff for any load impedance within the 4–1 VSWR circle compared to the 50- Ω nominal load. Demonstrative yield specifications are set to show healing improvement for saturated output power, small-signal gain, and power-added efficiency at >15.5 dBm, >20 dB, and >6% respectively, with self-healing resulting in a yield improvement from 20% to 90%, 20% to 100%, and 5% to 100% for each metric over 20 chips tested, and best achieved results of 16.5 dBm and 23.7 dB, and 7.2% respectively. For 4–1 VSWR tolerance, the yield specification was set to <3 dB, resulting in a yield improvement from 0% to 80% with self-healing and a best achieved result of 2.28 dB over ten chips measured. This results in an aggregate yield improvement due to self-healing from 0% to 80%.

To the best of our knowledge, there is no standalone integrated self-healing PA to compare with this work. To do a fair comparison of this self-healing PA to other PAs without self-healing, additional data, such as variation between multiple chips, tolerance to VSWR events, efficiency in the back-off region, and amplifier operation in the case of transistor failure would be necessary and has not been reported in those publications.

VIII. CONCLUSION

In this paper, we have presented a mm-wave self-healing PA system that is capable of healing against process variation and mismatch, load impedance mismatch, and even partial and total transistor failure that requires no external calibration of any kind. Robust sensors, actuators, and ADCs enable an on-chip digital algorithm block to heal for either maximum output power or minimum dc power while maintaining a specified output power, and aggregate yield over several performance specifications is improved from 0% to 80%.

ACKNOWLEDGMENT

The authors would like to thank Prof. A. Babakhani, Rice University, Houston, TX, USA (formerly with the California Institute of Technology) and A. Chang, Massachusetts Institute of Technology (MIT), Cambridge, MA, USA (formerly with the California Institute of Technology), for their valuable technical discussions and contributions, as well as Dr. S. Raman, Defense Advanced Research Projects Agency (DARPA), Arlington, VA, USA, T. Quach, Air Force Research Laboratory (AFRL), Dayton, OH, USA, and C. Maxey, Booz Allen Hamilton (BAH), McLean, VA, USA, for support.

The views expressed are those of the authors and do not reflect the official policy of the Department of Defense (DoD) or the U.S. Government.

REFERENCES

- [1] A. Natarajan, A. Komijani, X. Guan, A. Babakhani, and A. Hajimiri, "A 77-GHz phased-array transceiver with on-chip antennas in silicon: Transmitter and local LO-path phase shifting," *IEEE J. Solid-State Circuits*, vol. 41, no. 12, pp. 2807–2819, Dec. 2006.

- [2] A. Babakhani, X. Guan, A. Komijani, A. Natarajan, and A. Hajimiri, "A 77-GHz phased-array transceiver with on-chip antennas in silicon: Receiver and antennas," *IEEE J. Solid-State Circuits*, vol. 41, no. 12, pp. 2795–2806, Dec. 2006.
- [3] B. Razavi, "A millimeter-wave CMOS heterodyne receiver with on-chip LO and divider," *IEEE J. Solid-State Circuits*, vol. 43, no. 2, pp. 477–485, Feb. 2008.
- [4] M. Varonen, M. Karkkainen, M. Kantanen, and K. Halonen, "Millimeter-wave integrated circuits in 65-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 43, no. 9, pp. 1991–2002, Sep. 2008.
- [5] I. Aoki, S. D. Kee, D. B. Rutledge, and A. Hajimiri, "Distributed active transformer—A new power-combining and impedance-transformation technique," *IEEE Trans. Microw. Theory Techn.*, vol. 50, no. 1, pp. 316–331, Jan. 2002.
- [6] I. Aoki, S. D. Kee, D. B. Rutledge, and A. Hajimiri, "Fully integrated CMOS power amplifier design using the distributed active-transformer architecture," *IEEE J. Solid-State Circuits*, vol. 37, no. 3, pp. 371–383, Mar. 2002.
- [7] A. Komijani, A. Natarajan, and A. Hajimiri, "A 24-GHz, +14.5-dBm fully integrated power amplifier in 0.18- μ m CMOS," *IEEE J. Solid-State Circuits*, vol. 40, no. 9, pp. 1901–1908, Sep. 2005.
- [8] K. Sengupta and A. Hajimiri, "A compact self-similar power combining topology," in *IEEE MTT-S Int. Microw. Symp. Dig.*, May 2010, pp. 244–247.
- [9] R. Gupta, B. Ballweber, and D. Allstot, "Design and optimization of CMOS RF power amplifiers," *IEEE J. Solid-State Circuits*, vol. 36, no. 2, pp. 166–175, Feb. 2001.
- [10] G. Liu, T.-J. K. Liu, and A. Niknejad, "A 1.2 V, 2.4 GHz fully integrated linear CMOS power amplifier with efficiency enhancement," in *IEEE Custom Integr. Circuits Conf.*, Sep. 2006, pp. 141–144.
- [11] K. Bernstein, D. J. Frank, A. E. Gattiker, W. Haensch, B. L. Ji, S. R. Nassif, E. J. Nowak, D. J. Pearson, and N. J. Rohrer, "High-performance CMOS variability in the 65-nm regime and beyond," *IBM J. Res. Develop.*, vol. 50, no. 4.5, pp. 433–449, Jul. 2006.
- [12] T. Mizuno, J. Okumura, and A. Toriumi, "Experimental study of threshold voltage fluctuation due to statistical variation of channel dopant number in MOSFET's," *IEEE Trans. Electron Devices*, vol. 41, no. 11, pp. 2216–2221, Nov. 1994.
- [13] P. Stolk and D. Klaassen, "The effect of statistical dopant fluctuations on MOS device performance," in *Int. Electron Devices Meeting*, Dec. 1996, pp. 627–630.
- [14] H. Hashemi, X. Guan, A. Komijani, and A. Hajimiri, "A 24-GHz SiGe phased-array receiver-LO phase-shifting approach," *IEEE Trans. Microw. Theory Techn.*, vol. 53, no. 2, pp. 614–626, Feb. 2005.
- [15] A. Natarajan, A. Komijani, and A. Hajimiri, "A 24 GHz phased-array transmitter in 0.18 μ m CMOS," in *IEEE Int. Solid-State Circuits Conf. Tech. Dig.*, Feb. 2005, vol. 1, pp. 212–594.
- [16] S. Alalusi and R. Brodersen, "A 60 GHz phased array in CMOS," in *IEEE Custom Integr. Circuits Conf.*, Sep. 2006, pp. 393–396.
- [17] A. Valdes-Garcia, S. T. Nicolson, J.-W. Lai, A. Natarajan, P.-Y. Chen, S. Reynolds, J.-H. C. Zhan, D. Kam, D. Liu, and B. Floyd, "A fully integrated 16-element phased-array transmitter in SiGe BiCMOS for 60-GHz communications," *IEEE J. Solid-State Circuits*, vol. 45, no. 12, pp. 2757–2773, Dec. 2010.
- [18] A. Keerti and A.-V. H. Pham, "RF characterization of SiGe HBT power amplifiers under load mismatch," *IEEE Trans. Microw. Theory Techn.*, vol. 55, no. 2, pp. 207–214, Feb. 2007.
- [19] F. Bohn, K. Dasgupta, and A. Hajimiri, "Closed-loop spurious tone reduction for self-healing frequency synthesizers," in *IEEE Radio Freq. Integr. Circuits Symp.*, Jun. 2011, pp. 1–4.
- [20] H. Wang, K. Dasgupta, and A. Hajimiri, "A broadband self-healing phase synthesis scheme," in *IEEE Radio Freq. Integr. Circuits Symp.*, Jun. 2011, pp. 1–4.
- [21] A. Tang, F. Hsiao, D. Murphy, I.-N. Ku, J. Liu, S. D'Souza, N.-Y. Wang, H. Wu, Y.-H. Wang, M. Tang, G. Virbila, M. Pham, D. Yang, Q. Gu, Y.-C. Wu, Y.-C. Kuan, C. Chien, and M. Chang, "A low-overhead self-healing embedded system for ensuring high yield and long-term sustainability of 60 GHz 4 Gb/s radio-on-a-chip," in *IEEE Int. Solid-State Circuits Conf. Tech. Dig.*, Feb. 2012, pp. 316–318.
- [22] J. Liu, R. Berenguer, and M. Chang, "Millimeter-wave self-healing power amplifier with adaptive amplitude and phase linearization in 65-nm CMOS," *IEEE Trans. Microw. Theory Techn.*, vol. 60, no. 5, pp. 1342–1352, May 2012.
- [23] S. M. Bowers, K. Sengupta, K. Dasgupta, and A. Hajimiri, "A fully-integrated self-healing power amplifier," in *IEEE Radio Freq. Integr. Circuits Symp.*, Jun. 2012, pp. 221–224.
- [24] C. Law and A.-V. Pham, "A high-gain 60 GHz power amplifier with 20 dBm output power in 90 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. Tech. Dig.*, Feb. 2010, pp. 426–427.
- [25] K. Sengupta, K. Dasgupta, S. M. Bowers, and A. Hajimiri, "On-chip sensing and actuation methods for integrated self-healing mm-wave CMOS power amplifier," in *IEEE MTT-S Int. Microw. Symp. Dig.*, Jun. 2012, pp. 1–3.
- [26] C. F. Lee and P. Mok, "A monolithic current-mode CMOS DC–DC converter with on-chip current-sensing technique," *IEEE J. Solid-State Circuits*, vol. 39, no. 1, pp. 3–14, Jan. 2004.
- [27] S.-S. Lee and D. Allstot, "Electrothermal simulation of integrated circuits," *IEEE J. Solid-State Circuits*, vol. 28, no. 12, pp. 1283–1293, Dec. 1993.
- [28] S. Park, Y. Palaskas, A. Ravi, R. Bishop, and M. Flynn, "A 3.5 GS/s 5-b flash ADC in 90 nm CMOS," in *IEEE Custom Integr. Circuits Conf.*, Sep. 2006, pp. 489–492.
- [29] Y.-Z. Lin, Y.-T. Liu, and S.-J. Chang, "A 5-bit 4.2-GS/s flash ADC in 0.13- μ m CMOS," in *IEEE Custom Integr. Circuits Conf.*, Sep. 2007, pp. 213–216.
- [30] J. Li and U.-K. Moon, "A 1.8-v 67-mW 10-bit 100-MS/s pipelined ADC using time-shifted CDS technique," *IEEE J. Solid-State Circuits*, vol. 39, no. 9, pp. 1468–1476, Sep. 2004.
- [31] S. Mortezaipoor and E. Lee, "A 1-V, 8-bit successive approximation ADC in standard CMOS process," *IEEE J. Solid-State Circuits*, vol. 35, no. 4, pp. 642–646, Apr. 2000.
- [32] J. Russell and H. , "An improved successive-approximation register design for use in A/D converters," *IEEE Trans. Circuits Syst.*, vol. CAS-25, no. 7, pp. 550–554, Jul. 1978.
- [33] "Basics on load pull and noise measurements," Focus Microw., Montreal, QC, Canada, Appl. Note 8, Jun. 1994.



Steven M. Bowers (S'08) received the B.S. degree in electrical engineering from the University of California at San Diego, La Jolla, CA, USA, in 2007, the M.S. degree in electrical engineering from the California Institute of Technology, Pasadena, CA, USA, in 2009, and is currently working toward the Ph.D. degree at the California Institute of Technology.

His research interests include holistic integration of high-frequency analog circuits, advanced digital circuits, and novel electromagnetic structures to enable the next generation of mm-wave applications, specifically in adaptive and self-healing mm-wave circuits and mm-wave power generation and radiation.

Mr. Bowers was the recipient of the Caltech Institute Fellowship (2007), the Analog Devices Outstanding Student Designer Award (2009), and the IEEE RFIC Symposium Best Student Paper award (2012).



Kaushik Sengupta (M'12) received the B.Tech. and M.Tech. degrees in electronics and electrical communication engineering from the Indian Institute of Technology (IIT), Kharagpur, India, both in 2007, and the M.S. and Ph.D. degrees in electrical engineering from the California Institute of Technology, Pasadena, CA, USA, in 2008 and 2012, respectively.

In February 2013, he joined the faculty of the Department of Electrical Engineering, Princeton University, Princeton, NJ. During his undergraduate studies, he performed research at the University of Southern California, and the Massachusetts Institute of Technology, in the summers of 2005 and 2006, where he was involved with nonlinear integrated systems for high-purity signal generation and low-power RF identification (RFID) tags, respectively. His research interests are in the areas of high-frequency ICs, electromagnetics, optics for various applications in sensing, imaging and high-speed communication.

Dr. Sengupta was the recipient of the IBM Ph.D. fellowship (2011–2012), the IEEE Solid State Circuits Society Predoctoral Achievement Award, the IEEE Microwave Theory and Techniques Graduate Fellowship, and the Analog Devices Outstanding Student Designer Award (2011). He was also the recipient of the Prime Minister Gold Medal Award of IIT (2007), the Caltech Institute Fellowship, the Most Innovative Student Project Award of the Indian National Academy of Engineering (2007), and the IEEE Microwave Theory and Techniques Undergraduate Fellowship (2006). He was the corecipient of IEEE RFIC Symposium Best Student Paper Award in 2012.



Kaushik Dasgupta (S'09) received the B.Tech. and M.Tech. degrees in electronics and electrical communication engineering from the Indian Institute of Technology (IIT), Kharagpur, India, both in 2008, the M.S. degree in electrical engineering from the California Institute of Technology, Pasadena, CA, USA, in 2010, and is currently working toward the Ph.D. degree in electrical engineering at the California Institute of Technology.

During the summers of 2006 and 2007, he was a Research Intern with the System on Chip laboratory, University of Washington, Seattle, WA, USA. His research interests are mm-wave integrated transmitters and receivers and high-frequency reconfigurable systems, as well as ICs for biomedical applications.

Mr. Dasgupta was the recipient of the Caltech Atwood Fellowship (2008), the Analog Devices Outstanding Student Designer Award (2009), the Jagadish Bose National Science Talent Scholarship (2003), and the Best Undergraduate Project of the Electronics and Electrical Communication Engineering Department, IIT. He was the corecipient of the IEEE RFIC Symposium Best Student Paper award (2012).



Benjamin D. Parker received the B.S. degree in physics from Bowdoin College, Brunswick, ME, USA, in 1979, and the M.S. degree in physics from Brown University, Providence, RI, USA, in 1981. His graduate research concerned the optical properties of adsorbed layers on metal surfaces.

In 1986, he joined the GaAs Group, IBM T. J. Watson Research Center, Yorktown Heights, NY, USA, where he was involved with the characterization of III-V semiconductors. In 1991, he joined the Mixed-Signal Communications Integrated

Circuits Design Group, IBM T. J. Watson Research Center, where he has been involved with the design and verification of digital circuits in high-speed serial communications.



Ali Hajimiri (F'10) received the B.S. degree in electronics engineering from the Sharif University of Technology, Tehran, Iran, in 1994, and the M.S. and Ph.D. degrees in electrical engineering from the Stanford University, Stanford, CA, USA, in 1996 and 1998, respectively.

From 1993 to 1994, he was with Philips Semiconductors, where he was involved with a BiCMOS chipset for GSM and cellular units. In 1995, he was with Sun Microsystems, where he was involved with the UltraSPARC microprocessor's cache RAM design methodology. During the summer of 1997, he was with Lucent Technologies (Bell Laboratories), Murray Hill, NJ, USA, where he investigated low phase-noise integrated oscillators. In 1998, he joined the faculty of the California Institute of Technology, Pasadena, CA, USA, where he is currently the Thomas G. Myers Professor of Electrical Engineering and Director of the Microelectronics Laboratory. In 2002, he cofounded Axiom Microdevices Inc., whose fully integrated CMOS PA has shipped close to 200 million units, and was acquired by Skyworks Inc. in 2009. He has authored or coauthored over 150 refereed journal and conference technical papers. He authored *The Design of Low Noise Oscillators* (Springer, 1999). He was a Guest Editorial Board member of the *Transactions of Institute of Electronics, Information and Communication Engineers of Japan*. He holds over 60 U.S. and European patents. His research interests are high-speed and high-frequency ICs for applications in sensors, biomedical devices, photonics, and communication systems.

Dr. Hajimiri has served on the Technical Program Committee, International Solid-State Circuits Conference (ISSCC). He is an associate editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS and an associate editor of the IEEE TRANSACTIONS ON CIRCUITS SYSTEMS—II: EXPRESS BRIEFS. He is a member of the Technical Program Committee, International Conference on Computer Aided Design (ICCAD). He was guest editor for the IEEE TRANSACTIONS ON MICROWAVE THEORY AND TECHNIQUES. He was selected to the TR35 top innovator's list (formerly TR100) in 2004. He has served as a Distinguished Lecturer of the IEEE Solid-State Circuits and IEEE Microwave Theory and Techniques Society. He was the recipient of the California Institute of Technology Graduate Students Council Teaching and Mentoring Award, as well as the Associated Students of the California Institute of Technology Undergraduate Excellence in Teaching Award. He was the Gold Medal recipient of the National Physics Competition and the Bronze Medal recipient of the 21st International Physics Olympiad, Groningen, The Netherlands. He was a corecipient of the IEEE JOURNAL OF SOLID-STATE CIRCUITS Best Paper Award of 2004 and the International Solid-State Circuits Conference (ISSCC) Jack Kilby Outstanding Paper Award. He was a two-time corecipient of the CICC Best Paper Award, and a three-time winner of the IBM Faculty Partnership Award, as well as the National Science Foundation CAREER Award and Okawa Foundation Award.