ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

Overcoming Common Planar Phased Array Circuit Design Challenges

Joel Dobler, Product Applications Engineer, and
Sam Ringwood, System Platform Applications Engineer

Abstract

This article details the design and layout challenges associated with the electronic design of planar phased arrays and focuses on an RF front end containing power amplifiers (PAs), low noise amplifiers (LNAs), and beamformers. It discusses the signaling and timing involved with control of PA and LNA biasing, transmit and receive switching, memory loading, and beam advance. The PCB layout discussion highlights a single cell, consisting of a beamformer surrounded by four transmit/receive (TR) modules. Thermal management challenges are discussed, including component side heat sinking and heat sink cavity design with RF absorbers to avoid oscillations. The power management design for the RF front end is explained, including the power tree and the specific bias sequencing required to avoid PA damage.

Introduction

Two-dimensional planar phased array systems, where the RF circuitry and the antennal elements are on the opposite sides of the same PCB, offer a significant size advantage over three-dimensional blade-style structures. But that size advantage brings layout, power management, and thermal challenges. This is in addition to the traditional timing design challenges associated with radar systems. This article will explore how some of these challenges can be dealt with using well-planned device functionality and interfacing, careful PCB layout that maximizes usage of the limited available space, novel heat-sinking techniques, and well-planned power management that ensures safe turn-on and turn-off with correct supply sequencing.

Figure 1. A planar phased array front-end prototyping board (component side).

Phased array radar systems consist of many sections, which include the software and FPGA, ADCs and DACs, upconverters/downconverters, RF and beamforming circuitry, and the phased array antenna elements. This article concentrates on the RF front end and beamforming.

Radar System Design and Timing

In modern radars, the variable phase and amplitude blocks are consolidated into a beamformer integrated circuit (BFIC), with each BFIC containing several variable amplitude and phase blocks. Also, the PA, LNA, and transmit/receive switch can be integrated into a single TR module and be optimally designed to work directly with the BFIC.

A monostatic pulse radar has two critical timing characteristics: the transmit pulse duration (τ) and the pulse repetition time (PRT), which is the time between transmit pulses. Modern radars have transmit pulse durations that range between 10 µs and 100 µs. However, some applications feature pulse durations as short as 500 ns. Typical rise and fall times are between 500 ns and 1 µs and transmit duty cycles are between 1% and 20%. Typical pulsed radar timing is illustrated in Figure 2.

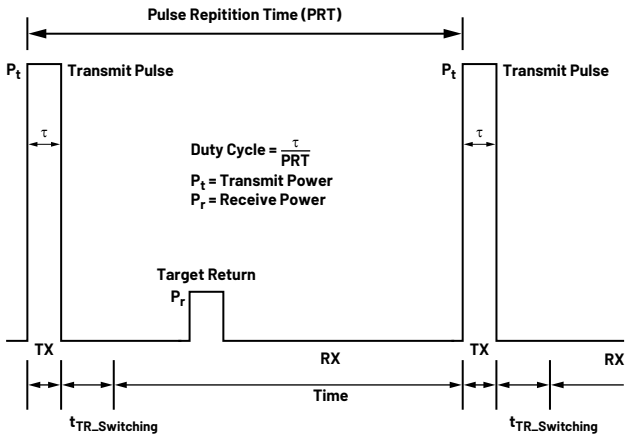

Figure 2. Simplified radar timing.

Radar metrics such as the minimum measurement range (R_{MIN}) or blind range, range resolution, and hits-per-scan are directly proportionally to τ and PRT. The minimum measurement range is also directly proportional to transmit-to-receive switching time:

$$R_{MIN} = \frac{c \times (\tau + t_{TR\_Switching})}{2}$$

Minimizing transmit-to-receive switching time is critical in search radar where targets can be relatively close to the transmitter.

Monostatic phased array radar oftentimes is a half-duplex system. Transmit-to-receive switching is accomplished with several control signals that must be sequenced together. These control signals perform the following functions:

- Toggle the transmit/receive switch
- Enable and disable the PA
- Enable and disable the LNA
- Enable and disable the BFIC's transmit and receive paths

Also, when the beam direction changes, new beam weights must be loaded into the variable amplitude and phase blocks. If local on-chip memory is available to store multiple sets of beam weights, the beam advance can be initiated with a single pulse of a control pin and with no SPI writes. An example of the sequencing and voltage levels of these controls is shown in Figure 3.

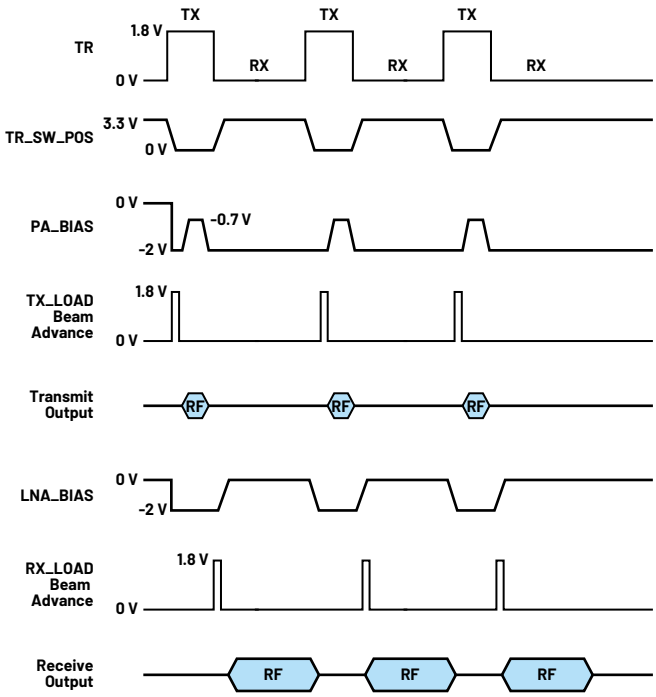

Figure 3. Transmit and receive timing with RF signaling, switch control, and amplifier bias.

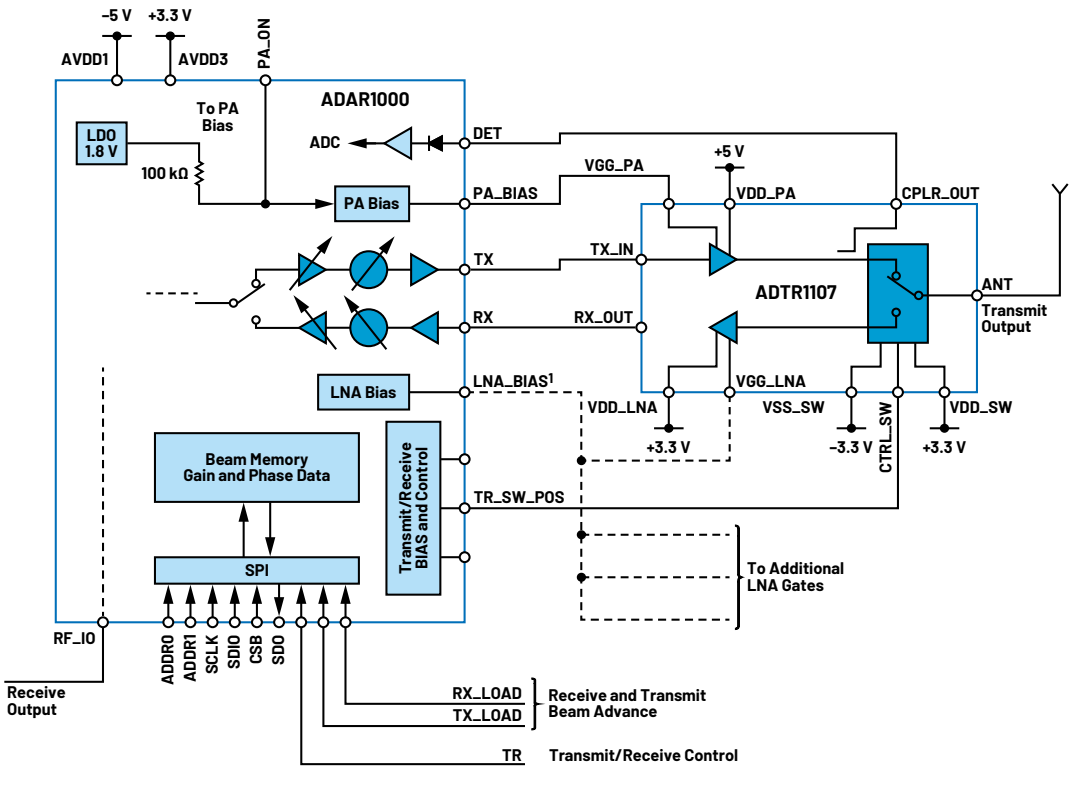

Figure 4. The ADAR1000 and ADTR1107 featuring a glueless interface.

An example of real-world control signaling and interfacing is shown in Figure 4. A single channel of the ADAR1000 BFIC is shown connected to a single ADTR1107 TR module. The two chips feature a single transmit/receive control signal going to the ADAR1000 via the TR pin. This pin in turn controls the PA bias, LNA bias, and the transmit/receive switch on the ADTR1107 TR module and controls the enabling/disabling of the ADAR1000 internal transmit and receive beamforming paths. Integrating the transmit/receive control not only reduces the board layout complexity by eliminating numerous control lines, but it also enables fast transmit-to-receive switching times so that metrics like the minimum measurement range are as short as possible.

Beam Memory

A discrete beamforming solution has to apply beam weights from a memory chip to the various variable amplitude and phase shifter blocks. The increased routing complexity and time required to do a beam advance are two major drawbacks of discrete system design. Modern phased array radars using integrated BFICs like the ADAR1000 have on-chip memory requiring a single load pulse and a few additional clock cycles to advance to a new beam. Newer BFICs like the ADAR3000 and ADAR3001 have simplified this to a single memory pulse.

In the example shown in Figure 4, the ADAR1000 BFIC features two on-chip locations that can store beam data: registers and random-access memory (RAM). There are advantages and disadvantages to using either type of memory for the storage of beam data. Both types are limited by the maximum SPI clock rate of 25 MHz when writing and reading beam data.

Sourcing beam data from the registers has the advantage of having only enough data for one beam position and that data is directly stored in the registers; thus, there is no need to operate the sequencer or perform any memory fetching. This mode of operation can be important in applications where users do not know what their next beam steering angle is a priori, and thus an on-the-fly scheme of writing new beam data needs to be implemented for each new beam steering angle. The downside of using the registers is the time it takes to clock in each new beam position, which is limited by the 25 MHz maximum clock rate. At this clock rate, it takes 4.16 µs to load either a receive or transmit beam position. Additionally, a quick load command is required to instruct the chip to use the new beam position data in the registers, since the ADAR1000 uses dual rank registers.

Sourcing data from the RAM is advantageous if the system knows the needed beam steering angles a priori and thus can clock in all the needed beam position data prior to the active operation of the system. Applications that employ a raster scan like weather and search radars can take advantage of storing beam positions in RAM. The time it takes to clock in each beam position into RAM is the same as the registers, but since this is normally done before active operation, all that is needed to apply a new beam position with the ADAR1000 is 6 clock cycles plus a load command. This operation takes as little as 320 ns.

RAM Partitioning

If RAM is used to source the beam position data, RAM partitioning can be used to update one section of the RAM while the other section is being accessed by the radar controller. Only one section is actively used to source beam position data, while new beam position data is written into the other section(s). This is possible due to the ADAR1000's use of start and stop RAM pointers, which enables the sequencer to start and stop at any arbitrary beam position in RAM.

A simple example is dividing the RAM into two sections: Section 1 contains beam positions 0 to 60, and Section 2 contains beam positions 61 to 120. Initially, start and stop pointer values are set to 0 and 60, respectively. Beam data is also clocked into beam positions 0 to 60 and beam position 0 is loaded. The beam is advanced from positions 0 to 60 several times (the sequencer wraps around and loads the start beam position after the stop beam position). In between beam advances of Section 1, Section 2 beam data is loaded.

When ready to switch from Section 1 data to Section 2 data, a new stop pointer value must be set. To ensure a smooth transition between beam positions 60 and 61, the new stop pointer must be written when the sequencer is at beam position 58 or before. On the first sequence through beam positions 61 to 120, a new start pointer value must be written. The new start pointer value must be set when the sequencer is at beam position 118 or before, or a smooth transition to beam position 61 will not occur. In general, any new beam data or new pointer values should be clocked in two beam positions ahead of the active beam position.

Writing new data to and then switching back to the Section 1 memory can be achieved using the same technique as outlined in the previous example.

PCB Layout

Phased array radar systems can vary in element count and size, from a 2 × 2 prototyping subarray to arrays that have 256, 512, or even 1024 elements. The array layout can be simplified down to a unit cell, typically connected to four elements. The cell size primary depends on the lattice spacing, usually $\lambda/2$ so that no grating lobes appear for beam steering all the way to the aperture horizon. Sometimes lattice spacing slightly greater than $\lambda/2$ is chosen for larger antenna gain and thinner beamwidths; however, this reduces the beam steering range that has no grating lobes. The BFIC and TR modules' location are usually constrained within the lattice spacing.

Cell Layout

An example cell layout on an Analog Devices planar phased array system board is shown in Figure 5. It features a four channel BFIC with four TR modules surrounding the BFIC.

The layout goals of the cell were to make the lattice spacing 15 mm ($\lambda/2$ at 10 GHz) and length-match the RF transmit and receive lines going from the BFIC to the four TR modules. The frequency of 10 GHz was chosen for $\lambda/2$ as it is at the center of X band, where many radar systems are operated, along with some satellite communications systems. Length matching of the transmit and receive interconnect lines between the BFIC and the TR module reduces the burden on calibration.

The glueless interface between the BFIC and the four TR modules results in almost no external components on the transmit, receive, or the coupler lines going to the power detectors. This makes the layout routing efficient, which is shown in Figure 5. The glueless interface along with the TR module chips orientated at 45° off axis relative the BFIC easily enables the routing of the transmit and receive lines to be length-matched, all while maintaining the 15 mm lattice spacing. Figure 6 shows the other side of the PCB, which contains RF connectors at 15 mm lattice spacing. In a real-world planar system, these connectors would be replaced with patch antennas.

Supply decoupling capacitors are kept at a minimum on this cell layout. The majority of the supply decoupling is for the BFIC. The TR modules also have supply decoupling on the board, but most of these capacitors are technically not needed due to internal supply decoupling. Additional capacitors on the board were a conservative design decision and the 15 mm lattice spacing provided enough board space for these capacitors.

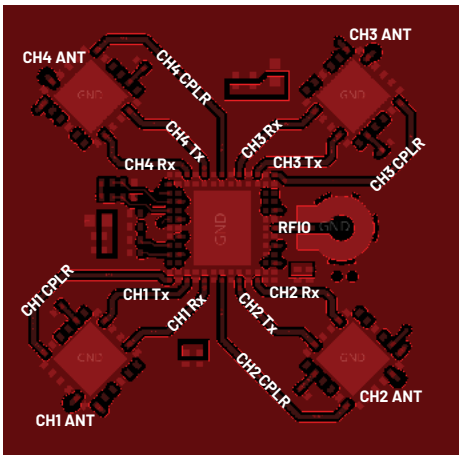

Figure 5. A layout of one X-band cell (top layer only).

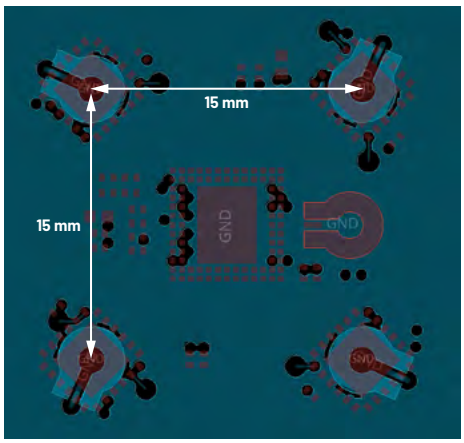

Figure 6. A layout of top and bottom layer showing locations of SMPM connectors (Layer 1 pins superimposed for reference).

Another important goal of the cell layout was to keep the switching transients on the bias lines to a minimum so that switching times could also be minimized. This was accomplished by keeping the length of these lines as short as possible to reduce line parasitics. With the RF lines and supply decoupling on the top layer, the PA bias and LNA bias lines from the BFIC to the TR module had to be routed on the PCB inner layers 4 and 5. The various digital control lines to the BFIC from the microcontroller were routed on Layer 5, which is shown in Figure 7 and Figure 8. On larger arrays with necessarily longer trace runs between the controller and BFICs and/or higher clock speeds, signal integrity simulations must be done to account for all trace propagation delays and ensure all digital timing is correctly synchronized.

On the ADI planar phased array system board, the RFIO trace is short and immediately goes to a connector; thus, its isolation to other ports is large. However, on a real-world planar phased array application board, care must be taken when routing the TR module antenna ports and the RFIO port, particularly the Channel 2 antenna and Channel 3 antenna ports shown in Figure 5 as they are on the same side of the cell layout as the RFIO port. While in transmit mode, the paths have larger gain relative to receive mode; thus, the isolation requirement between the paths is larger to prevent instability and oscillations.

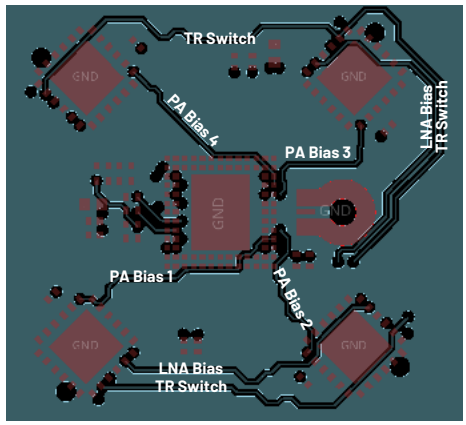

Figure 7. Layer 4: PA bias, LNA bias, and TR switch control routing (Layer 1 pins superimposed for reference).

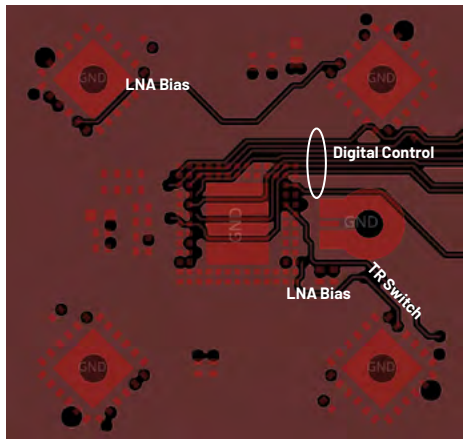

Figure 8. Layer 5: LNA bias, TR switch control, and digital routing (Layer 1 pins superimposed for reference).

Thermal Management

In a planar phased array system, with the antenna array on one side of the board and components on the other side, the heat sink must be located on the component side of the board. This presents a challenge for generating a thermal management solution that will effectively remove a sufficient amount of heat from the various components, particularly the PAs, so that none of the components exceeds its maximum junction temperature.

Options for Heat Management in a Planar Phased Array Antenna

Due to the antenna array's location on the back side of the board, the thermal management solution cannot be a scheme where heat is pulled from a component's ground paddles through thermal vias to a back-side mounted heat sink. Instead, the heat must flow directly from the top side of the component or must travel indirectly through the base the component, into the PCB moving laterally before coming back out to a component-side heat sink. This is shown in Figure 9. There are two possible options for heat conduction away from the components:

- The heat sink contacts large amounts of area of the top (component) side ground layer
- The heat sink contacts the top side of the components

The heat sink design largely depends on the power dissipation and thermal resistance values of the components. Most components have a low junction-to-case bottom thermal resistance ($\theta_{JC\text{-}BOTTOM}$) and relatively high junction-to-case-top thermal resistance. So, the scheme shown in Figure 9a is generally more effective.

The question becomes whether to additionally contact the top side of the components or not. If the junction to case-top thermal resistance ($\theta_{JC\text{-}TOP}$) is of the same order of magnitude as $\theta_{JC\text{-}BOTTOM}$ (for example, $\theta_{JC\text{-}TOP} < 5 \times \theta_{JC\text{-}BOTTOM}$) having the heat sink also contact the top side of the components' package would provide an additional useful parallel path for heat conduction. The metal of the heat sink should not contact the top of the package directly as this may cause mechanical stress. Instead, a piece of thermally conductive tape or thermal transfer pad should be used as shown in Figure 9b.

A real-world example of component side heat sinking of a planar phased array is shown in Figure 10. This is the same planar phased array front-end board shown in Figure 1, only with the heat sink attached. Thermal conducting compound can be seen at the interface between the heat sink and exposed ground metal on the board (white material along the edge of the heat sink). Also note the absence of fins in strategic areas to allow access to the RF input/output ports of the ADAR1000 BFICs.

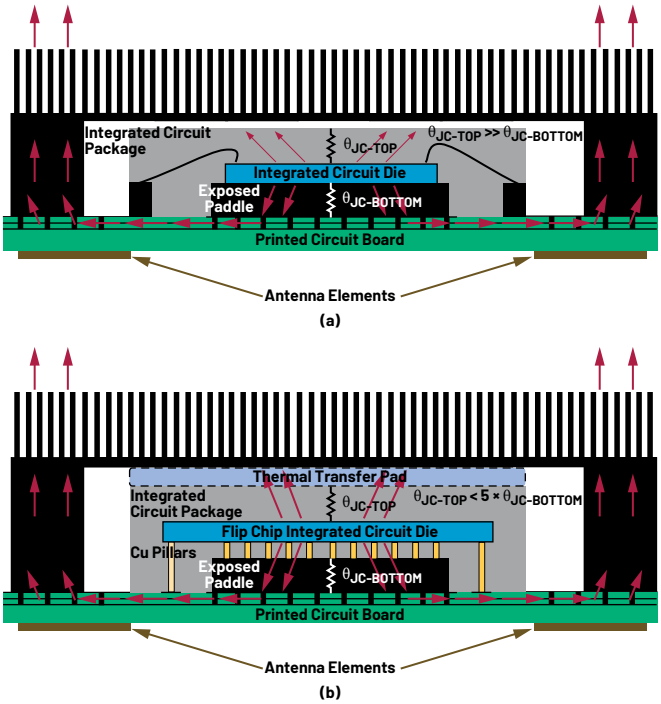

Figure 9. Component side heat sinking with (a) an LFCSP package with heat flowing primarily through the bottom of the package and (b) a copper-pillar flip-chip LGA package with heat flowing through the top and bottom of the package.

Figure 10. A planar phased array front end with component side heat sink.

Heat Sink Cavity Design

Component side heat sinking forces each cell of four TR modules and one beamformer to be inside a metal cavity. Care should be taken in the sizing and design of the cavity. An electromagnetic simulation should be performed to ensure that the cavity does not interact with the circuitry and result in instability or oscillations.

Analysis of metallic cavities whether for shielding purposes or thermal management cannot be overlooked at high frequencies. Resonant modes are generally supported when the largest cavity dimension is greater than or equal to one half wavelength ($\lambda/2$) of the operating frequency in free space. Energy radiated from circuitry and PCB traces within the cavity have nowhere to propagate beyond the confines of the cavity possibly resulting in instability and oscillatory behavior of active circuitry.

Various techniques can be employed to mitigate the undesired effects of cavity mode resonances. These techniques can be as complex as custom metallic structures within the cavity to only support resonances at frequencies outside the range of operation. A much simpler technique is the placement of an RF absorbing material within the cavity to attenuate the energy of the resonant modes. The absorber is composed of materials that create a high permittivity and high permeability of the electric and magnetic fields, respectively, for a range of operational frequencies. This is analogous to an electrical band-stop filter.

The ADI planar phased array system board heat sink mechanical dimensions of the cavities do support resonances at frequencies within its operational range. To mitigate the effects of the resonant modes, a die cut RF absorber is installed within each cavity to dampen resonances without impeding on the performance of the board. The RF absorber attenuation ranges from approximately 20 dB/cm to 50 dB/cm over the frequency band of operation, effectively lowering the Q factor of the cavity resonator.

An electromagnetic simulation of the heat sink cavity and RF absorber was performed using Keysight's EMPro simulation tool. A simple model was constructed to analyze a single cavity consisting of the mechanical design of the heat sink cavity,

the PCB material, and a bulk material to emulate the ICs attached to the PCB. Two simulations were performed using the finite element method simulation engine to calculate the eigenmode resonances within the cavity. The cavity material for the first simulation was defined as air and resulted in eigenmode resonances with a high Q factor. In the second simulation, RF absorber was used instead of air as the cavity material and resulted in no eigenmodes. The second simulation case concluded the selected absorber material lowered the Q factor of the cavity suppressing the resonant energy within the cavity. To ensure accuracy, the two simulation cases were verified with measured data of the hardware. The frequencies in which resonant mode are supported as well as the lowering of the Q factor by the RF absorber can be observed in the two gain measurements shown in Figure 11. These before and after measurements confirm the simulated predictions and highlight the importance of the electromagnetic analysis.

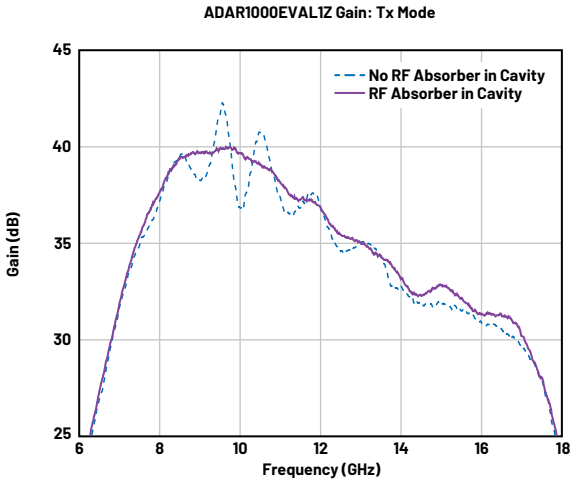

Figure 11. Gain measurements before and after the RF absorber was installed.

Power Management

Planar phased arrays that feature BFICs and TR modules usually require several different supply voltage domains. Additional supply domains may be required for support circuitry for the BFIC and TR module like a low voltage digital supply and/or an intermediate step-down voltage may be required to feed required supply domains.

For the power tree design, a common voltage such as 12 V is usually chosen to feed all the power management blocks. Common voltage rails of +5 V and +3.3 V can easily and efficiently be generated, as well as the common negative counterpart voltage rails of –5 V and –3.3 V.

PAs and LNAs that use depletion mode FET devices often require their gates to be biased to a negative voltage prior to bringing up the drain voltage in order to prevent damage to the device. Proper sequencing of the drain and gate voltages is thus crucial and must be considered when designing the power management solution.

Design of Power Management Tree

Figure 12 shows the power tree that was used to provide supply voltages for the BFICs and TR modules that are on the previously mentioned planar phased array board. The 12 V input drives a hot swap circuit featuring the ADM1172 hot swap controller that provides a protected 12 V with an 8.33 mA current limit to the branch of the power tree that powers the ADAR1000 and ADTR1107 ICs.

From the protected 12 V, the 3.3 V and 5 V supplies powering the ADAR1000 and ADTR1107 are generated by a pair of synchronous step-down Silent Switcher® regulators, the LT8642 and the LT8652, respectively. These regulators ensure low noise and low spurious tones on the critical RF supplies. The LT8642 provides 3.3 V for the ADAR1000's main supply and the V_{DD} for the switch and LNA of the ADTR1107, while the LT8652 provides the 5 V V_{DD} for the ADTR1107 PA.

To generate the negative supply voltages needed by the ADAR1000 and ADTR1107, an intermediate –6 V negative supply is generated from the protected 12 V rail, using the ADP5074 DC-to-DC inverting regulator. The ADP5074 in turn drives two negative linear regulators in parallel: the LT3093 and LT3094, generating the –3.3 V for the ADTR1107 switch V_{SS} and the –5 V for the ADAR1000, respectively.

A secondary branch off of the 12 V input powers the LT8606 step-down switching regulator and the ADP150 ultralow noise linear regulator, which provides a separate miscellaneous 3.3 V supply and a 1.8 V digital supply, respectively.

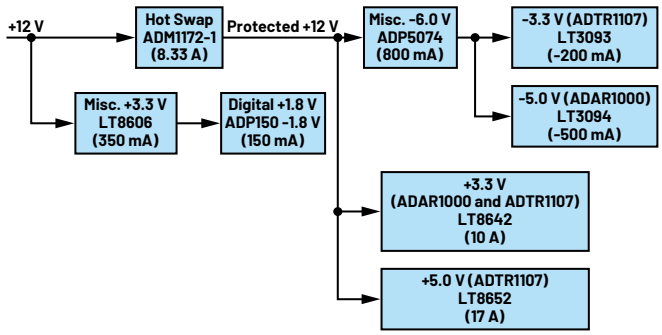

Figure 12. Power tree for a planar phase array system with BFICs and TR modules.

Power-Up Supply Sequencing

Powering up the tree shown in Figure 12 is complex. It requires the user to input specific control signals at specific times during the overall power-up sequence. The critical control signal sequencing is enumerated in Figure 13.

To start the power-up sequence, first apply the 12 V power to the board, after which the following are powered up:

- The 3.3 V miscellaneous and 1.8 V digital supplies from LT8606 and ADP150 turn on immediately.
- Once the ADM1172 hot swap gives its power good signal (HOT_SWAP_PG, Control Signal 1 in Figure 13), the ADP5074 turns on and provides its –6 V negative rail.

The power up of the LT8606, ADP150, and ADP5074 is done automatically and requires no input from the user.

Second, the remaining regulators power up in a specific sequence so that ADTR1107 PA is not damaged. The sequence is initiated by providing a rising edge power up signal (POWER_UP_DOWN_IN, Control Signal 2 in Figure 13) to a D-flipflop, which holds a high state at its output. The D-flipflop output is AND'ed with the hot swap power good signal. If both signals are high, then the sequencer enable (SEQ_EN) signal is asserted high on the ADM1186-2. The ADM1186-2 allows each LDO regulator 47 ms for power up (otherwise a fault condition is raised). An additional 2.2 ms of delay is inserted between each LDO regulator power up. The exact LDO regulator enable sequence is:

- The ADM1186-2's OUT1 pin (Control Signal 3 in Figure 13) asserts high and drives the LT8642 3.3 V LDO regulator enable pin; the LT8642 has 47 ms to assert its power good signal to ADM1186-2's VIN1 pin; a delay of 2.2 ms occurs.
- The ADM1186-2's OUT2 pin asserts high and drives the LT3093 –3.3 V LDO regulator enable pin; the LT3093 has 47 ms to assert its power good signal to ADM1186-2's VIN2 pin; a delay of 2.2 ms occurs.

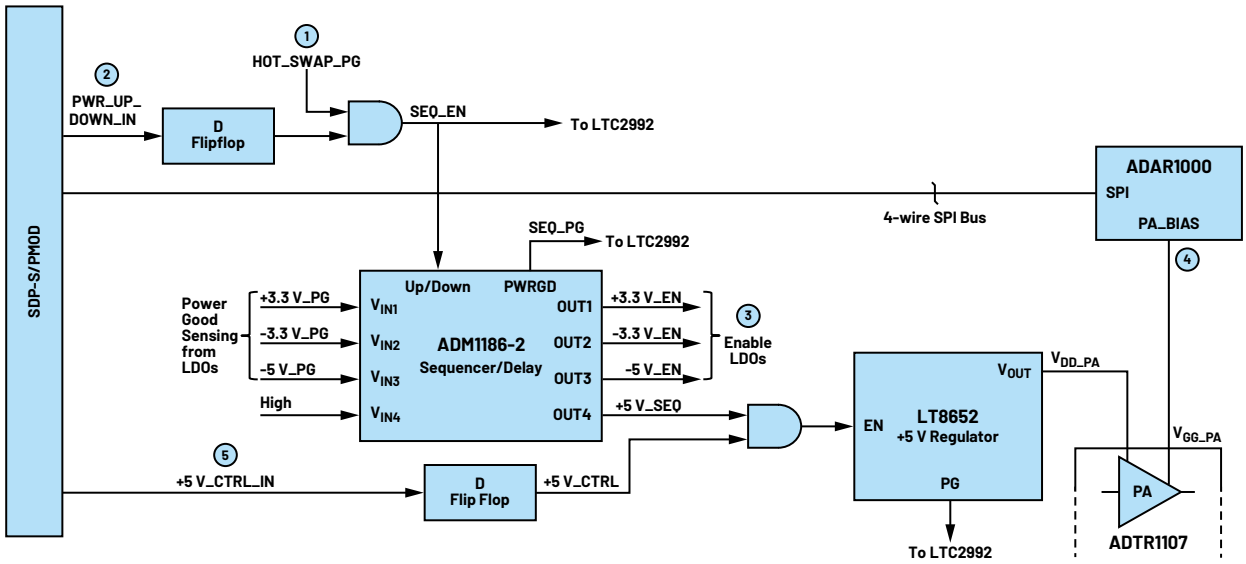

Figure 13. Power-up sequencing for ADTR1107 PA protection.

- The ADM1186-2's OUT3 pin asserts high and drives the LT3094 –5 V LDO regulator enable pin; the LT3094 has 47 ms to assert its power good signal to ADM1186-2's VIN3 pin; a delay of 2.2 ms occurs.
- The ADM1186-2's OUT4 pin asserts high and drives an AND gate whose other input is a 5 V control signal. The ADM8611-2's PWDGD pin will then assert high.

The ADAR1000 is now completely powered up along with the ADTR1107 switch and LNA.

Third, the ADAR1000 is programmed so that its PA bias outputs –2 V (Control Signal 4 in Figure 13), which is safe value for the ADTR1107 PA during its +5 V power up. Once all the PA bias outputs are programmed, a rising edge 5 V control signal (5 V_CTRL_IN, Control Signal 5 in Figure 13) drives a D-flipflop, which holds a high state at its output. The flipflop drives an AND gate whose other input is the ADM1186-2's OUT4 pin. The AND gate drives the LT8652 5 V regulator enable pin high and starts up the last supply rail. The overall power up sequencing solution is shown in Figure 13.

Power Monitoring

The LTC2992 monitors the power delivered by the LT8642 and LT8652 regulators. The power is monitored by measuring the voltage across a sense resistor that is on each regulator output. The voltage is sampled with an on-chip ADC and can be readback via the I²C port.

The LTC2992 also has four GPIO pins, which can be used to sense the state of digital signals. The four digital signals sensed are:

- Power good signal from the LT8652
- The sequencer enable signal that drives the ADM1186-2 UP/DOWN pin
- The sequencer power good signal from the ADM1186-2
- The 5 V control signal (5 V_CTRL) that enables the LT8652 regulator

The state of the GPIO pins can also be readback over the I²C port. The power monitoring solution can be seen in Figure 14.

The power monitoring and digital signal sensing that the LTC2992 provides help to ensure that the main supplies to the ADAR1000 and ADTR1107 are functioning properly. The power monitoring can also serve as a debug tool and built-in-test (BIT) for the ADAR1000 and ADTR1107 chips themselves. Individual chips in the array can be isolated and powered up to make sure that they are consuming the correct amount of power.

The GPIO sensing of the LTC2992 allows the user to know if the LT8652 has powered up properly by sensing its power good signal. With this, users can ensure that the 5 V control signal is low until they are ready for 5 V power up on the ADTR1107 PAs. The GPIO sensing also gives the sequencing state of the ADM1186-2—that is, it senses whether the LDO regulators are powering up or powering down, and if the power up/down sequence is complete.

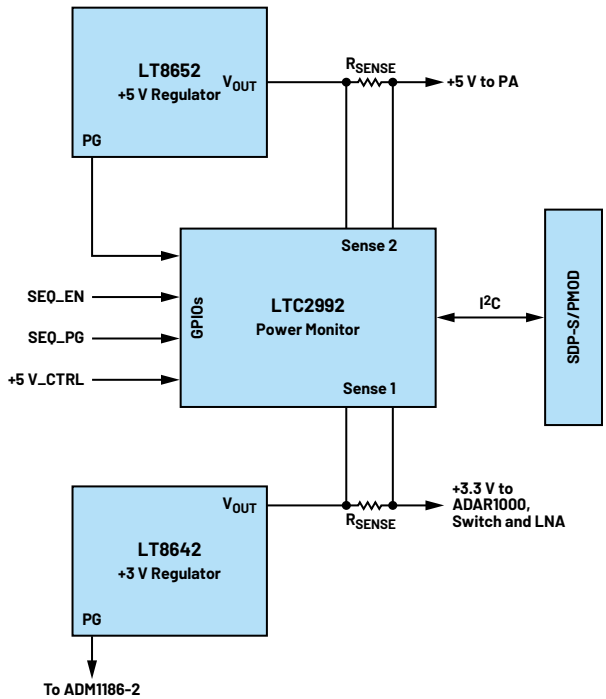

Figure 14. Power monitoring of 5 V and 3 V regulators and digital signal sensing.

Conclusion

In this article, we have looked at the challenges associated with designing the RF front end of a planar phased array system. The size restrictions that stem from the required lattice spacing at high frequencies demand novel approaches to circuit design. These challenges can be significantly reduced through the use of BFICs and TR modules, which easily interface together. This glueless interface has the knock-on benefit of shorter RF traces and less control lines, resulting in systems with fast responses that are easier to calibrate. Simple single-pin controls and on-chip memory of the BFIC enable quick TR transitions and beam advance. Because one side of the PCB is reserved for the patch antenna array, compromises must be made to implement an effective heat sink. Proper heat sink cavity analysis and design is vital to avoid oscillations. A power management solution with proper supply sequencing for the RF front end is crucial to the overall system design to keep noise and spurs to a minimum and avoid damage to the amplifiers.

The circuitry described in this article comes from the ADAR1000EVAL1Z X/ku band analog beamforming prototyping board, which can be purchased from Analog Devices.

About the Authors

Joel Dobler is a staff product applications engineer in the Aerospace and Defense Group, focusing on beamformer products, but also supporting vector modulators and programmable low-pass filters. He has worked for Analog Devices since 2006, supporting a wide range of RF products including logarithmic and rms detectors, digital and analog variable gain amplifiers, mixers, and IQ demodulators. He received his B.S.E.E. from Washington State University in 2005 and his M.E.E.E from Portland State University in 2007.

Sam Ringwood is a system platforms applications engineer for the Aerospace and Defense Business Unit located in Greensboro, North Carolina. Prior to joining ADI, Sam worked in RF test and RF design roles within the U.S. Nuclear Weapons Complex. He received his B.S.E.C.E degree in 2015 and M.S.E.E degree in 2016 from University of Missouri—Kansas City. Sam is currently focusing on creating full system solutions for aerospace and defense applications such as phased array radar.