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

Holistic Design in mm-Wave Silicon ICs

Ali HAJIMIRI[†a)], ***Nonmember***

SUMMARY Millimeter-waves integrated circuits offer a unique opportunity for a holistic design approach encompassing RF, analog, and digital, as well as radiation and electromagnetics. The ability to deal with the complete system covering a broad range from the digital circuitry to on-chip antennas and everything in between offers unparalleled opportunities for completely new architectures and topologies, which were previously impossible due the traditional partitioning of various blocks in conventional design. This can open a plethora of new architectural and system level innovation within the integrated circuit platform. This paper reviews some of the challenges and opportunities for mm-wave ICs and presents several solutions to them.

key words: *mm-wave, silicon integration, on-chip antenna, phased array*

1. Introduction

Integrated circuit have had a phenomenal growth since their initial introduction [1]–[4]. New advancements in silicon integrated circuits has made it possible to integrate billions of transistors capable of operating at mm-wave frequency on the same die, thus creating an unprecedented opportunity to devise systems that leverage a very large number of transistors at mm-waves. Although conventional system architectures and their associated partitioning (e.g., antenna, filter, front-end, synthesizer, and digital back-end) can be directly ported to an integrated setting, the true potential of such integrated platforms can only be fully realized through a holistic co-design at multiple levels of abstraction (e.g., EM, RF, and digital) that could result in superior overall system architectures.

There are many existing and emerging applications that can greatly benefit from full integration of mm-wave systems in silicon. On the one hand, high-speed wireless communication can take advantage of the larger bandwidths available at mm-wave frequencies and leverage techniques, such as phased-array multiple antenna systems that are more conducive to mm-wave integrated circuits to improve the signal-to-noise ratio and hence increase the bit rates [7]–[9], [12], [13]. On the other hand, ranging and sensing applications, such as automotive radar at 24 GHz and 77 GHz, which can provide features such as autonomous cruise control, collision avoidance, low-visibility driving aid, early warning and brake priming, self-parking, and global traffic control resulting in a qualitative change in personal transportation [5], [6], [9]. The vehicular platform can also be used for applications such as "data mule" which enables transportation of large volumes of data through the storage space on each car. It uses the very-high speed mm-wave link in the car to create a very large scale high capacity (albeit high latency) data network.

Another class of applications involves non-commercial large arrays that can be used for a variety of systems including radar and tracking. The unmatched size, cost, and repeatability of silicon integrated circuits makes it possible to make very large scale phased arrays ("million element arrays") that cannot be done in a cost effective manner using the conventional module based approaches.

In spite of the numerous opportunities for mm-wave systems, we also need to deal with many challenges to successful implementation of such systems. In the subsequent sections we will discuss these opportunities and impediments and demonstrate some examples of the proposed holistic approach, encompassing EM, RF, and digital circuits.

2. Propagation

Perhaps the least discussed issues in IC design circles are the properties of wave propagation which are often problematic but could provide new opportunities depending on the function of the system in question. In general, the excess propagation loss (in addition to the standard inverse square law dependence) can result in substantial attenuation of the signal. A typical plot of the excess path loss versus frequency is shown in Fig. 1 [14]–[17]. Quite remarkable is the vertical axis, where the attenuation in dB is shown on a logarithmic scale itself! The need for "log of a log" to be able to show the attenuation in a meaningful fashion un-

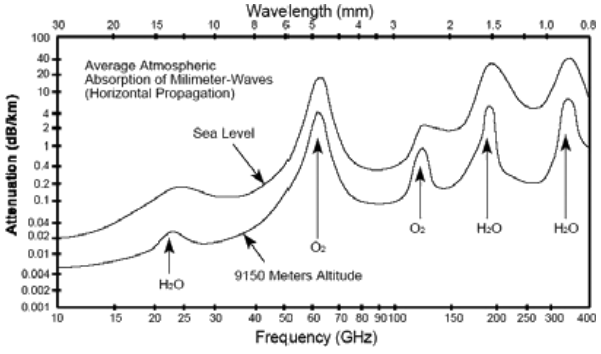

Fig. 1 Atmospheric absorption vs. frequency [14].

Manuscript received January 22, 2008.

[†]The author is with California Institute of Technology, Pasadena, CA 91125, USA.

a) E-mail: hajimiri@caltech.edu

DOI: 10.1093/ietele/e91–c.6.817

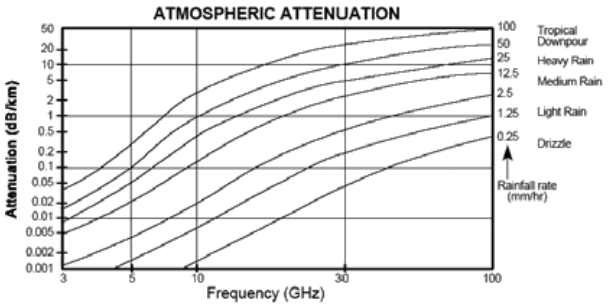

Fig. 2 Added absorption due to precipitation vs. frequency [17].

derscores the wide range of change in the excess path loss at different frequencies. The path loss is particularly large at some frequencies due to the absorption lines of certain elements (e.g., the oxygen absorption at 60 GHz results in an excess path loss of ∼20 dB/km compared to ∼0.4 dB/km at 94 GHz). This loss is exacerbated further by humidity and precipitation with an additional attenuation factor determined by the level of precipitation, as shown in Fig. 2 [14]–[17]. Note that the attenuation levels given in Fig. 2 are in excess of those in Fig. 1. This additional loss is of particular interest in outdoor applications, such as vehicular radar and point-to-point communication.

Propagation attenuation is not purely problematic, however. The added loss at frequencies such as 24 GHz[†] and 60 GHz can result in a faster drop in the signal power resulting in potentially higher frequency reuse ratio for a cellular implementation, where the same frequency can be used nearby by a different base station.

The propagation properties of mm-waves are closer to their yet higher frequency electromagnetic siblings, namely, visible light than are the more traditionally considered lower frequency RF signals. The behavior is more line-of-sight and there is less discernible diffraction compared to lower RF frequencies. This makes them particularly suitable for compact beam forming applications where a collimated beam of the signal can be transmitted in a desirable direction, for instance by using a phased array as discussed later. This is also particularly useful in wireless communications where the same frequency band can be used by two users in close proximity of each other as long as they are not on the same line of sight, which can be interpreted as an effective channel reuse ratio greater than unity.

The mm-waves also experience a larger loss upon reflection off of typical surfaces. This could pose a challenge to obtain reliable wireless coverage in the absence of line-of-sight view of the transmitter, but at the same time it produces a smaller delay spread in the channel impulse response [18]. This corresponds to a channel with smaller multi-path dispersion, which can in turn lead to a higher effective channel capacity.

3. Antennas

The most commonly used antennas are passive and as such do not generate power. The directivity of an antenna is

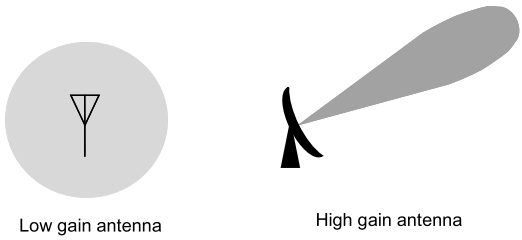

Fig. 3 A low gain isotropic antenna contrasted with a directional one.

a measure of how much it can focus the energy in a particular direction at the expense of lower signal transmission (and hence reception due to reciprocity) in other directions. The reference value of comparison is an omnidirectional (isotropic) antenna that transmits (and receives) power equally in all directions. Thus the directivity of an antenna is measured in dBi (dB with respect to isotropic). (Fig. 3). Of course, directivity comes at the cost of spatial coverage since higher gain antennas must have higher directivity due to conservation of energy. If the antenna has internal energy loss, (e.g., due to ohmic resistance of the conductor used) part of the energy entering the antenna terminal would turn into heat and the rest of it will turn into electromagnetic (EM) radiated energy. The antenna gain is the product of this energy efficiency and its directivity. Since in a lossless antenna the total radiated energy over all solid angles should be the same as that of an isotropic antenna over a unit sphere, we have:

$$\oint G d\Omega = 4\pi \tag{1}$$

where $G = G(\theta, \phi)$ is the gain as a function of direction. One can define an effective area for an antenna as a function of direction (θ and ϕ) as:

$$A(\theta, \phi) = \frac{P_r}{S(\theta, \phi)} \tag{2}$$

where P_r is the total received power and $S(\theta, \phi)$ is the incident power density in that direction. Obviously, directional antennas will have different areas (effective cross-section) when viewed from different directions. The antenna appears large in the direction of maximum gain and small when viewed from other directions. The antenna theorem [19] states:

$$\oint A d\Omega = \lambda^2 \tag{3}$$

where λ is the wavelength in the direction of interest. This indicates that, for example, the effective area of an isotropic antenna drops as the square of frequency. This has an important corollary, namely, for a given directionality, a given

[†]The mm-wave technically referred to frequencies in the 30–300 GHz range whose free space wavelength is between 1 mm and 1 cm. However, due to the nature of the applications at 24 GHz and the close proximity to the edge of the mm-wave frequencies, we will treat it in a similar fashion.

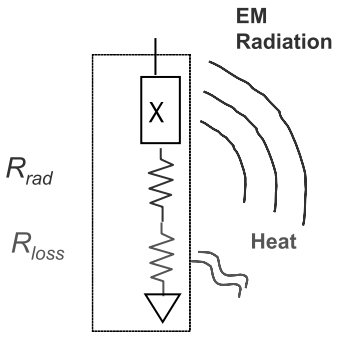

Fig. 4 Electrical equivalent model of an antenna.

antenna looks smaller at higher frequencies and hence it collects less power. Equivalently, we could say that the antenna has to become more directional at higher frequencies, if we wish to maintain the same collection area. This poses a challenge that could be turned into an opportunity by using phased arrays.

If we were to use a single relatively isotropic antenna for wireless communications, the antenna will collect less power at higher frequencies, making it more difficult to maintain a desired signal-to-noise ratio. Now, one could try to use a more directional antenna to recover the lost gain, but this solution will be limited to cases where the relative locations and orientations of neither transmitter nor receiver change rapidly (e.g. fixed point-to-point communications). One could try to solve this problem by using multiple antennas to increase the effective collection area, which again results in an increase in the directivity. However, by using phased arrays we could change the direction of the beam electronically and would hence effectively be able to cover a larger area. We will perform a more quantitative analysis of the impact of the antenna area reduction with frequency on channel capacity in a later section.

Another important parameter of an antenna is its radiation resistance. An antenna will present an impedance to the driving port which typically has reactive (imaginary) and resistive (real) parts, as in Fig. 4. The resistive part can be divided into two components, the radiation and the loss resistances. The loss resistance is due to the energy loss mechanism that converts energy to heat, such as ohmic loss in the antenna conductor. It is a parasitic element that should be minimized, as will be seen shortly. On the other hand, the radiation resistance is the desired part which models the energy converted from the electrical domain to electromagnetic radiation power leaving the system. Obviously, the reason that this energy conversion is modeled as a resistor is that the energy has been "lost" in the electrical domain[†]. It should be clear that the radiation resistance has nothing to do with the ohmic resistance of the antenna conductor, and even if it were made out of a superconductor, the radiation resistance would have remained the same.

The power efficiency of an antenna is the ratio of the electromagnetic energy radiated by it to the electrical energy entering it. This ratio is related to the radiation resistance, R_{rad}, and the loss resistance, R_{loss} in the following fashion:

$$\eta \equiv \frac{P_{rad}}{P_{elec}} = \frac{R_{rad}}{R_{rad} + R_{loss}} \tag{4}$$

We will see in the on-chip antennas section that maintaining a high R_{rad}/R_{loss} ratio is a challenge in silicon integrated antennas.

4. Silicon Integration

There are several challenges to integration of mm-wave systems on a silicon substrate in spite of its numerous advantages. We will discuss some of these challenges in this section.

Conventional silicon integrated circuit processes have a conductive substrate, which can cause energy loss due to magnetically induced eddy currents. Antennas can also couple energy into the substrate resulting in some additional loss. Additionally, the finite conductivity of the metal structures used in passive devices results in further energy loss in the system. The ohmic loss can be significantly more important at mm-waves because of the small skin depth (e.g., the skin depth of copper at 60 GHz is approximately 300 nm). This is particularly problematic in CMP processes which results in rougher surfaces, where most of the current follows. The additional "fill" and "cheese" rules for wide metal strips also degrade the current handling capabilities of the metal lines [11].

The other equally important effect is the high dielectric constant of most semiconductor substrates, such as silicon (e.g., $\epsilon_r = 11.7$ for silicon). The rectangular silicon substrate forms a dielectric waveguide that can sustain propagating modes depending on its physical dimensions. Unfortunately, many of the natural modes for a typical silicon die size fall right in the middle of the mm-wave bands of interest. This in turn creates an alternative mechanism for energy loss for on-chip components, where energy could easily leak into the substrate modes and be dissipated by ohmic loss within the substrate, or even worse, be radiated in undesirable directions [5], [11]. Also, energy can be easily coupled back into other on-chip components creating parasitic, yet potentially strong, coupling between seemingly unrelated elements on the same chip. Special care must be paid to the overall design to avoid such coupling. Also generation and pick-up of mm-wave signals must be minimized using techniques, such as fully-differential signaling.

For an on-chip antenna, the high dielectric constant of the substrate causes a considerable portion of the electromagnetic energy to be sucked into the substrate as opposed to being radiated into the air from the top side [5], [20]. In the case of silicon, the impact of this effect is further exacerbated by the conductivity of the substrate. This is a very important issue when trying to integrate the antennas with the rest of the chip, as will be discussed further in the section dealing with propagation and antennas.

[†]An analogous situation arises in a loudspeaker where the acoustic impedance is determined by the energy conversion from the electrical domain to acoustic domain.

Power generation in small feature size silicon technologies also presents a very serious challenge, mostly due to the lower breakdown voltages resulting from the scaling process and the shrinking of the depletion regions in the transistors. This necessitates the use of a lower supply voltage. Unfortunately, this is in direct conflict with the maximum power that can be generated using conventional power amplifier techniques, necessitating the use of parallel structures and novel power combing approaches as described in [21], [22].

Obviously, the modeling of the transistors (as well as the passive devices) becomes more challenging at these higher frequencies. The smaller parasitic components within the models are more prone to error, and hence special attention must be paid to guarantee reliable results, especially in the presence of process variations and environmental changes.

5. Interface at mm-Wave Frequencies

Even in the absence of the aforementioned challenges of making mm-wave integrated circuits in silicon, we still need to be concerned about the mm-wave interface to the outside world. External connections to the chip need to be rethought completely. As an example, consider the case where we connect the mm-wave input (or output) of the chip to a perfectly matched transmission line on a adjacent printed circuit board using a wirebond. Approximating a 1mm long wirebond as a 1 nH inductor with a Q of 30 we find that our connection has an impedance of approximately $12\,\Omega + j360\,\Omega$ at 60 GHz, making it very difficult to achieve a proper impedance match. Furthermore, such an element can easily turn into an unintended radiative element at mm-wave frequencies, creating very serious problems at the system level. Even a solder ball in flip-chip technology presents a non-trivial impedance to match to the outside world. Reported results based on direct wafer probing of the die entirely ignore this challenge. If one tries to make the interface with wirebond, every effort should be made to make it into a structure emulating a transmission line, as in [9]. On the other hand, these interface challenges make the totally-integrated solutions with on-chip antenna much more attractive by offering a solution where the high frequency electrical interface to the outside world is completely eliminated and replaced by an electromagnetic radiating interface. An example of such an approach can be found in [5], [6].

All of the above considerations demonstrate the important challenges in the area of mm-wave integrated circuits. It is clear that the direct application of microwave approaches to design of silicon-based integrated circuits, is neither efficient, nor optimum, since such techniques are fundamentally based on the underlying assumption of using a small number of well-characterized transistors. At the same time, direct application of lower frequency analog techniques to this problem would not be particularly successful either, as in the absence of proper treatment of the complex electromagnetics issues involved in such systems, the designers will not be able to come up with an implementation that closely matches the expected results. Thus, it is clear that to take full advantage of this significant opportunity, one needs to start at the system architecture level and devise new holistic methodologies and architectures along the way.

6. Capacity Limits

Shannon's theorem predicts that the maximum error-free bit rate of a channel with a bandwidth of BW and a signal-to-noise ratio of S/N is given by

$$R_B = BW \log_2\left(1 + \frac{S}{N}\right) \tag{5}$$

Although at first glance it may appear that going to higher frequencies automatically results in higher bit rates due to greater available bandwidths, the signal-to-noise term's degradation at very high frequencies diminishes this improvement.

To obtain a more realistic picture, we need to take several effects into account. Assuming a low directivity single-antenna to be able to communicate in different directions, we expect the effective antenna area to be similar in different directions and proportional to λ^2. Also, let us assume a BW that grows linearly with the center frequency. In our example, let us assume that the BW is about 1% of the center frequency, i.e., $BW \approx f_0/100$. Assuming a distance of 10 m and that the noise figure degrades linearly with frequency[†], we can calculate the Shannon capacity as a function of frequency, as plotted in Fig. 5 for two different initial NF values. It is noteworthy that the Shannon capacity approaches a maximum and goes back down, mainly due to lower collected signal power due to the lower antenna collection area and worse receiver noise performance.

Figure 5 is a clear indication that to be able to take full advantage of mm-wave frequencies, we must consider multiple antennas with electronic steering capability to increase the collection area, while maintaining good spatial coverage by steering the beam in the desired direction without any mechanical movement.

7. On Chip Antennas

On-chip antennas make it possible for the input and output signals to be radiated directly into and out of the chip with no need for an electrical interface at mm-wave frequencies. However, successful implementation of an on-chip antenna with reasonable gain and efficiency is a nontrivial task primarily due to the high dielectric constant and conductivity of the substrate.

To demonstrate this point more clearly, let us consider

[†]In this calculation it is assumed that we can generate power at higher frequencies with the same facility as lower frequencies. This assumption is clearly not accurate. Nonetheless inclusion of any degradation of power vs. frequency will only make the peak in capacity occur at a lower frequency, while maintaining the general behavior of the curve.

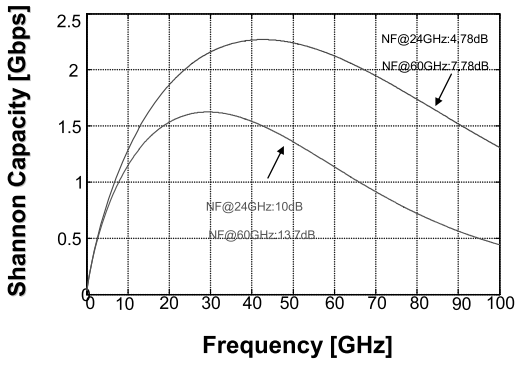

Fig. 5 Shannon capacity vs. frequency.

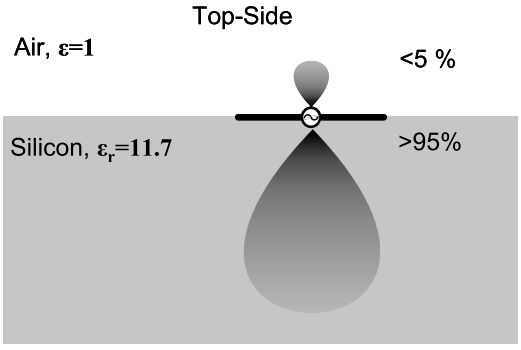

Fig. 6 Dipole antenna at the air-silicon boundary.

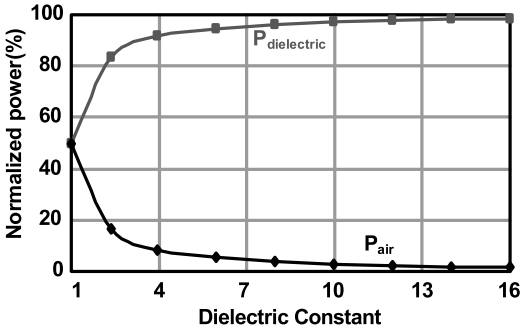

Fig. 7 Percentage of the radiated power into the air and dielectric vs. relative permittivity of the dielectric for a dipole antenna at the boundary.

a dipole antenna placed at the interface of air and a semi-infinite region of a dielectric material with a relative permittivity of ϵ_r, as shown in Fig. 6. The dielectric presents a lower electromagnetic impedance than the air and as a result more of the radiated electromagnetic power will end up in the dielectric. Figure 7 shows the percentage of the power going in each direction as a function of the ϵ_r [5], [20]. As can be seen, for silicon with an ϵ_r of approximately 11.7, more than 95% of the power will end up going into the silicon, as opposed to the air.

To avoid this problem, one may decide to use a ground plane in the lowest metal layer (e.g., M1) while placing the antenna in the top-most layer. While this may sound appealing initially, is not very practical due to the typical spacing

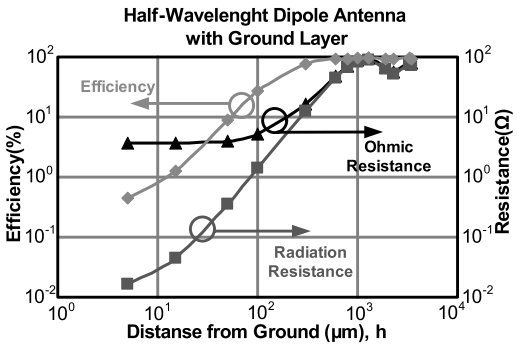

Fig. 8 The efficiency, radiation and loss resistance of a dipole vs. spacing from the ground plane.

between the top and bottom metal layers (e.g., in the 10–15 μm range in today's technologies). If one were to attempt this approach, the resulting antenna would have such a small radiation resistance due to the close proximity of the ground plane that its efficiency would be extremely low (e.g., <1%). The trade-off between these different parameters is shown in Fig. 8 [5].

Placing the ground plane directly underneath the substrate does not improve upon this situation. Although the spacing may be high enough for the radiation resistance to become acceptably large, the entire substrate will act as a waveguide with multiple propagation modes, some of which will typically fall very close the frequency of interest. The energy does get coupled into these modes where part of it is lost to heat in the conductive substrate and the remainder is radiated in undesirable directions.

One approach to solve this problem is to take advantage of the energy coupling into the substrate and radiating it from the backside of the chip. To do this, we must suppress the substrate modes and direct the energy into the desired radiative mode. This can be achieved by changing the "shape" of the substrate electromagnetically.

One way to do this is to use a dielectric lens [20] on the backside of the chip which couples the EM energy into a dominant radiative mode [23]. To reduce the antenna metal loss, three bottom layers are connected in parallel with vias to form the dipole antennas as close to the substrate as possible. To reduce the substrate loss, the fabricated chip is thinned down to 100 μm. Due to the layout constraints, antennas are placed at the chip edge and to maintain a uniform dielectric constant substrate underneath the antennas, an undoped silicon slab with the same thickness of silicon chip is abutted to the chip. For mechanical stability, a 500 μm silicon wafer is used to hold the chip and the silicon lens is attached to the backside of the undoped wafer. A 2-axis spherical far field measurement technique is utilized to measure the radiation pattern while a W-band horn antenna is used to irradiate the integrated dipoles. The 3-D measured patterns of two middle antennas are shown in Fig. 9. A maximum peak gain of about +8 dBi is achieved in this measurement. As seen in Fig. 9, the peaks of two antennas occur at the two different directions. The chip requires no high frequency

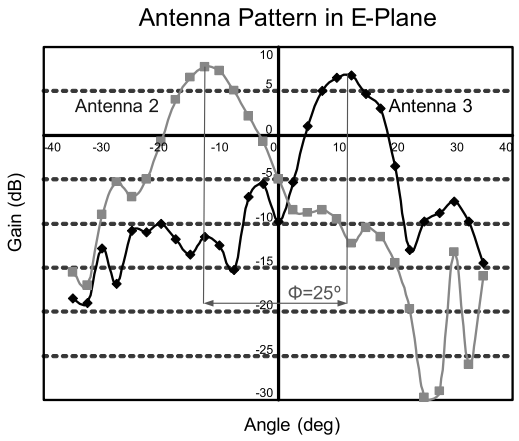

Fig. 9 Measured patterns of two adjacent antennas on a silicon chip with a backside lens.

electrical connection to the outside world.

8. On-Chip Transmission Lines

At a single element level, careful effort must be put into using the full capability of on-chip transmission lines designed for minimum loss to obtain the highest efficiency. In normal coplanar waveguide (CPW) structures made in CMOS processes with relatively high substrate conductivity, capacitive coupling to the substrate is often the dominant source of high-frequency loss. On the other hand, in an on-chip microstrip structure, where substrate losses are minimal, the close proximity of the ground plane to the signal line results in a narrow signal line for practical impedance levels. This constraint increases ohmic losses in the signal line that are further increased because of skin effect at high frequencies. Figure 10 shows the substrate shielded coplanar structure that is a combination of the two structures [32], [33]. As in a microstrip, the electric field does not penetrate into the substrate while the magnetic field and current distributions resemble that of a CPW structure. As seen in Fig. 10, a second shield layer is placed beneath the first shield layer, with metal stripes covering slots of the first layer, thereby completely isolating the coplanar structure from the substrate [34].

Slotting the bottom plate, as shown in Fig. 10, forces the return current to be mostly concentrated in the coplanar ground lines. In a coplanar structure, the separation between signal and return currents is greater than in a microstrip line. Therefore, the coplanar structure stores more magnetic energy, resulting in a larger distributed inductance per unit length. However, the proximity of the slotted ground line to the signal line results in high capacitance per unit length. Simultaneously high values of L_u and C_u, lead to slower wave velocity, while maintaining a given impedance level. This results in a shorted effective wavelength and hence a smaller loss for a given transmission line length in terms of wavelength.

Another way to look at this structure is to view it as a CPW structure with periodic capacitive loading. This is

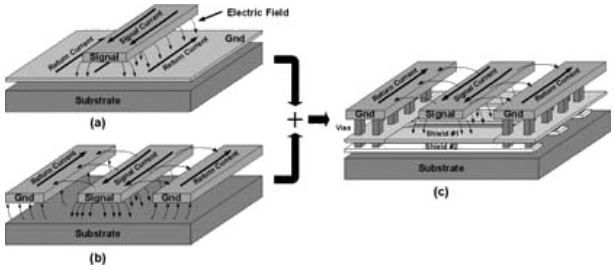

Fig. 10 Formation of slow-wave transmission line via combining coplanar and strip-line structures and slotting the bottom-plane to force the return current through the adjacent ground lines.

similar to the inductive loading concept [32], however in this case extra capacitance is added by placing the patterned ground beneath the coplanar structure. Therefore, the capacitance per unit length of the structure is increased, thereby slowing down the wave.

An example of this structure is given in [34], the velocity can be reduced by more than a factor of two, and as a result the wavelength at 24 GHz in this structure is 3 mm. Furthermore, reasonable impedance levels (~26 Ω) are achieved for signal line widths of 60 μm. These wide signal lines provide low loss per unit length (~0.9 dB/mm@24 GHz) and eliminate the need for tapering at the RF pad and transistor connections. The combination of lower loss-per-unit-length and the shorter length of the transmission lines lead to a much lower passive loss in the matching networks.

9. Power Generation

The breakdown voltage of silicon based transistors drops with scaling because of the necessary reduction in the device dimensions and the width of the depletion regions. If the supply voltages are not scaled the same voltage will drop across a smaller distance and hence result in a larger electric field. However, an unfortunate side effect of the reduced voltage swing is lower maximum power that can be delivered to a given load impedance, limiting the maximum generated mm-wave power using conventional techniques. This makes it necessary to consider alternative techniques of on-chip power generation in millimeter waves. These limitations on the supply voltage force us to reevaluate the design process and investigate methods to extract the most out of the individual transistors and minimize the passive loss.

An example of this was given in [34] where slow-wave transmission lines discussed earlier were used to design a 180 nm CMOS PA operating at 24 GHz with peak output power of +14.5 dBm fully-integrated power amplifier with on-chip 50 Ω input and output matching. The use of shielded-substrate coplanar waveguide structures for matching networks results in low passive loss and small die size. The PA achieves a power gain of 7 dB and a maximum single-ended output power of +14.5 dBm with a 3 dB bandwidth of 3.1 GHz, while drawing 100 mA from a 2.8 V supply. The chip area is 1.26 mm^2.

As another example, a 77 GHz, +17.5 dBm fully-

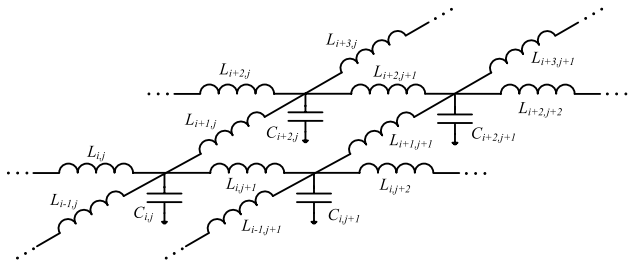

Fig. 11 An LC lattice forming a 2D transmission medium.

integrated PA with integrated 50 Ω input and output matching was demonstrated in 130 nm SiGe BiCMOS process [35]. The PA achieved a peak power gain of 17 dB and a maximum single-ended output power of +17.5 dBm with the power-added efficiency (PAE) of 12.8%. It offered a 3 dB bandwidth of 15 GHz and draws 165 mA from a 1.8 V supply. Microstrip tubs on-chip transmission lines are used resulting in large isolation between adjacent lines, enabling integration of the PA in a small area of 0.6 mm^2.

While the single-chain power amplifier techniques of the last two examples can be used to produce power at mm-waves, the big prize can only be achieved by utilizing the large number of fast transistors available in silicon. This can be done by using advance on-chip power combing techniques at mm-wave frequencies. Novel on-chip power combining techniques, such as distributed active transformers (DAT) [21] have been used in the microwave region to implement watt-level fully-integrated PAs for cellular applications [38]. While such techniques could be modified for mm-wave applications, it is possible to adopt power combing technique that utilize the comparable wavelengths of mm-wave signals to the dimensions of the chip. An example of such techniques are the so-called "optotronics" techniques proposed in [36], [37].

It is possible to create a two-dimensional (2D) propagation medium by creating a lattice of inductors and capacitors, as shown in Fig. 11. The lattice could be non-uniform in the sense that the values of the inductors and the capacitors can be different at different locations. As long as there is no abrupt change in the values between adjacent elements, it is possible to define a local characteristic impedance assuming isotropicity as

$$Z(\vec{r}) = \sqrt{L(\vec{r})/C(\vec{r})} \tag{6}$$

where $L(\vec{r})$ and $C(\vec{r})$ are the local inductance and capacitance values at the location defined by vector r. Also the propagation velocity can be defined as a function of the location as

$$v(\vec{r}) = 1/\sqrt{L(\vec{r})C(\vec{r})} \tag{7}$$

The local characteristic impedance and propagation velocity allow us to create surfaces where we can manipulate the characteristic impedance and propagation velocities independently, in effect manipulating the electrical signal in a similar fashion to optical waves, resulting in the concept of optotronics [36], [37].

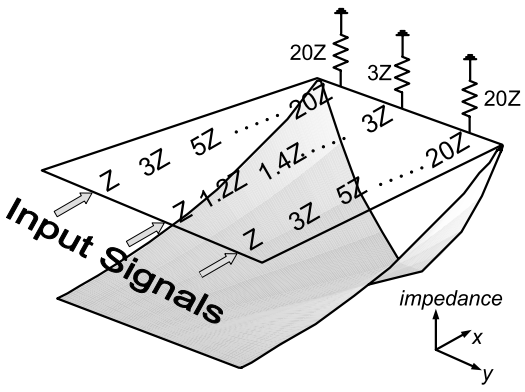

Fig. 12 Power combining through impedance transformation in an electrical funnel.

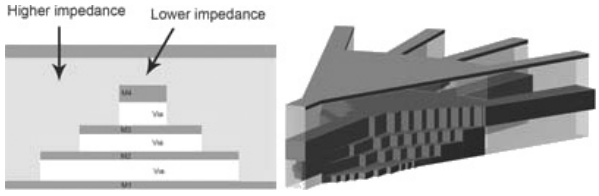

Fig. 13 An integrated circuit implementation of the 2D broadband combiner.

A very simple example of this is the electrical funnel combiner demonstrated in [37]. A rectangular surface can be modified to keep the propagation velocity constant along the y axis by keeping the *LC* product constant along any vertical line, while increasing the characteristic impedance at the top and bottom of the lattice at a faster rate as we move along the x axis to the right, as depicted in Fig. 12. This arrangement allows us to maintain the same voltage along the y axis, while concentrating more current (and hence more power) in the middle as we move to the right exploiting the graceful impedance transformation. Using multiple signal sources running inphase to drive the low-impedance left-hand side of the funnel, we can generate a planar wavefront moving to the right. The output node is at the center of the right boundary, while the other right boundary nodes are terminated to the higher local impedances at those points.

In a practical implementation, we can implement this combiner using a trapezoid since the characteristic impedance at the edges of the rectangular implementation keeps increasing making it is possible to discard the higher impedance parts of the mesh, as we move to the right. In an integrated setting with multiple metal layers, the gradual change in the capacitance can be achieved by using different bottom metal layers as the ground plane, as shown in Fig. 13. This approach has been used to combine the outputs of four distributed amplifiers in an 84 GHz power amplifier in 130 nm SiGe BiCMOS process [37]. The implemented PA has a peak power of 125 mW at 84 GHz with a −3 dB bandwidth of 24 GHz.

While the optotronics approach can be use to make

a high frequency broadband power combiner, it can have much broader applications in high-speed on-chip signal processing with structures such as lens and prism that can be used to perform functions such as Fourier transform at extremely fast rates.

10. Integrated Phased Arrays

Integration of a complete phased array system in silicon results in substantial improvements in cost, size, and reliability. It also offers various opportunities to perform in situ signal processing and conditioning without having to go off-chip, leading to additional savings in cost and power. The multiple signal paths, operating in harmony, provide benefits at the system and circuit level.

Multiple antenna phased-arrays imitate a directional antenna whose bearing can be controlled electronically [24]–[30]. This electronic steering makes it possible to obtain gain and directionality, while eliminating the need for continuous mechanical reorientation of the actual antennas. Additionally, the parallel nature of a phased array antenna transceiver alleviates the power handling and noise requirements for individual active devices used in the array. This makes the system more robust to the failure of individual components. In the past, such systems have been implemented using a large number of microwave modules, adding to their cost and manufacturing complexity [28], [29].

A phased-array transmitter or receiver consists of several signal paths each connected to a separate antenna. The antenna elements of the array can be arranged in different spatial configurations [26]. The array can be formed in one, two, or even three dimensions, with one or a two dimensional array being more common.

The principle of operation of a phased-array is similar for both receivers and transmitters. In a phased-array receiver, the radiated signal arrives at different times at each of the spatially separated antennas. The difference in the time of arrival of the signal at different antennas depends upon the angle of incidence and the spacing between the antennas. As shown in Fig. 14, an ideal phased-array receiver compensates for the time delay difference between the signals from different antennas and combines the signals coherently to enhance the reception from the desired direction(s), while rejecting emissions from other directions. Similarly, in a phased-array transmitter, the signals in different elements are delayed by different amounts so that the outgoing signals add up coherently only in the desired direction(s). Incoherent addition of the signals in other directions results in lower radiated power in those directions. Thus in a phased-array based system, the transmitter generates less interference in receivers that are not targeted. Furthermore, the receiver is also capable of nulling out some interferers as long as they do not originate from the same direction as the signal. Additionally, for a given power level at the receiver, the power that has to be generated is lower in a phased-array transmitter than in an isotropic transmitter.

In a transmitter with n elements (Fig. 15), if each element radiates P Watts, the total power that will be seen at the receiver in the desired direction is that of an omnidirectional transmitter with a power of n^2P Watts. The n^2 increase comes from the coherent addition of the signals in amplitude in the desired direction. For example, in a four element transmitter, the total power radiated in the beam direction is 12 dB higher than the power radiated by each element.

In receivers, the advantages of a phased array include better sensitivity and higher interference rejection capabilities. For a given receiver sensitivity, the output *SNR* sets an upper limit on the noise figure of the receiver. The noise figure, *NF*, is defined as the ratio of the total output noise power to the output noise power caused only by the source [31]. In an actual phased array receiver implementation, compared to the output *SNR* of a single-path receiver, the output *SNR* of the array is improved by a factor as high as n, depending on the noise and gain contribution of different stages [8], [9]. For instance, if the noise from the antennas is uncorrelated, an 8-path phased-array can improve the receiver sensitivity by 9 dB.

Thus, in a system based on phased arrays at the transmitter and receiver, the higher *SNR* and lower interference increases channel capacity. Furthermore, the directivity of

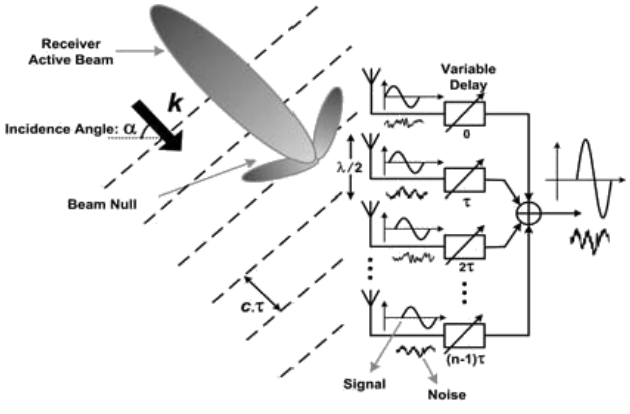

Fig. 14 The simplified architecture of a phased-array receiver.

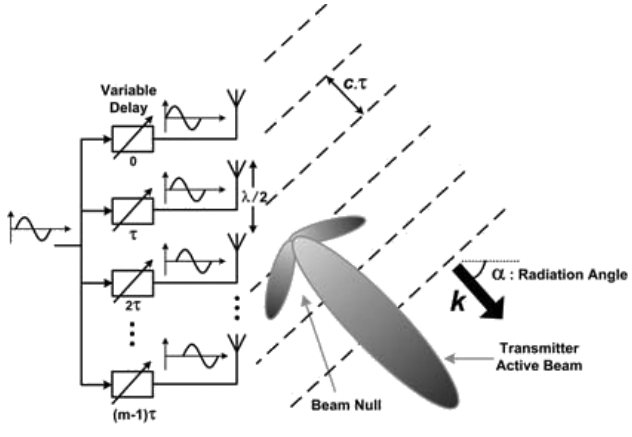

Fig. 15 The simplified architecture of a phased-array transmitter.

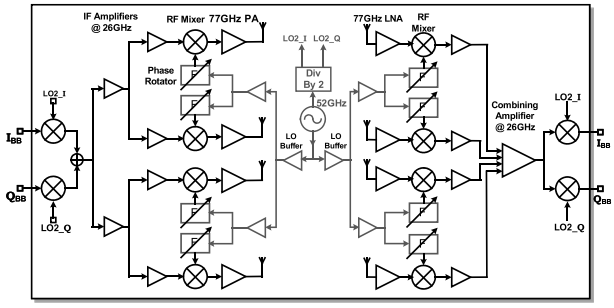

Fig. 16 The simplified architecture of the 77 GHz phased-array transceiver with on-chip antennas.

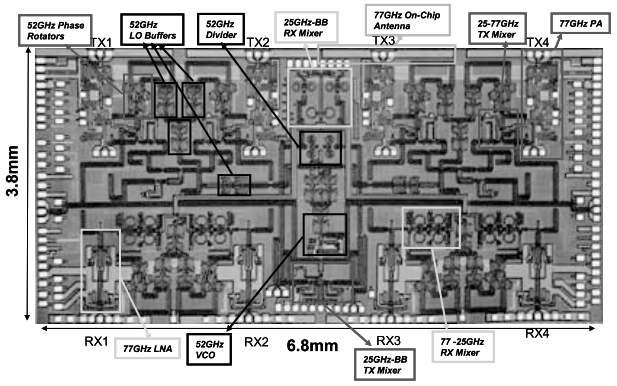

Fig. 17 The die micrograph of the 77 GHz phased-array transceiver with on-chip antennas.

the transmit-receive pairs can result in higher frequency reuse ratios, leading to higher network capacity.

For narrowband systems, the true-time delay necessary in each element of a phased-array can be approximated by a phase-shift. This approximation leads to some signal dispersion, due to the non-constant group delay, which increases as the bandwidth of the signal increases. This dispersion translates to a higher BER in communication systems and lower resolution in radar systems [8].

The phase shift necessary in each element of a phased-array can be achieved at RF, at baseband/IF or in the LO path. In integrated implementations, there are several advantages to using LO path phase shifting, as the gain in each element of the transmitter or receiver is less sensitive to the amplitude variations at the LO ports of the mixers [8], [9].

The various phases of the LO necessary in the LO-phase shifting approach can be generated in a central fashion by generating all the necessary phases at one place (e.g., a multi-phase oscillator), as in [8], [9], or in a decentralized fashion by distributing only one phase of the LO signal and generating the remaining phases locally using a phase rotator for each LO path, similar to [5], [6].

A fully integrated transceiver 4-element phased array transceiver with on-chip antennas has been designed and fabricated in a 0.13 μm SiGe BiCMOS process [5], [6]. The receiver consists of the complete down-conversion path with low-noise amplifier (LNA), frequency synthesizer, phase rotators, combining amplifiers, and on-chip dipole antennas, as shown in Fig. 16. A distributed active combining amplifier at an IF of 26 GHz is used to perform the signal combining. A 52-GHz first LO is generated on chip and is routed to different elements, where it can be phase shifted locally by the phase rotators. The local LO-path phase-shifting scheme enables a robust distribution network that scales well with increasing frequency and number of elements while providing high-resolution phase shifts. Measurements indicate a single-element LNA gain of 23 dB and a noise figure of 6.0 dB. Each of the four receive paths has a gain of 37 dB and a single-path overall noise figure of 8.0 dB. Each on-chip antenna has a gain of +8 dBi. Each element of the 2-step upconversion transmitter generates +12.5 dBm of output power at 77 GHz with a bandwidth of 2.5 GHz leading to a 4-element effective isotropic radiated power (EIRP) of 24.5 dBm. Each on-chip PA has a maximum saturated power of +17.5 dBm at 77 GHz. The entire phased array transceiver occupies an area of 3.8 mm × 6.8 mm, as shown in the die photo of Fig. 17.

11. Near-Field Direct Antenna Modulation

Conventional radio transmitters modulate the desired information at the baseband and then upconvert and amplify that information by modulating the desired signal to create the desired phase and amplitude for the symbols transmitted. This requires upconversion mixers and perhaps more importantly a linear (or linearized) PA for non-constant-envelope modulations. The modulated signal is fed to an antenna that radiates the modulated signal in all directions, albeit with different gain factors and time delays. Even in the case of a directional antenna, although a stronger signal is sent in the desired direction, essentially the same data constellation and information is sent in other directions that can be intercepted by an unintended receiver with good sensitivity.

The near-field direct antenna modulation technique is a fundamentally different approach to modulating and transmitting data that combines digital circuitry and electromagnetics to modulate the signal directly at the antenna through digital manipulation of the electromagnetic boundary conditions [10]. The direct antenna modulation makes it possible to simplify the rest of the transmitter by eliminating the need for upconversion mixers and linear power amplifiers. It can be used to create a secure communication channel by sending the desired data only in the intended direction, while a scrambled data stream is transmitted in other directions. Also, it can be used to *concurrently* transmit two (or multiple) streams of completely independent data in different directions both at the maximum rate, increasing the transmitted data rate significantly. As we will see next its very simple building block allows it to be used for transmitting data at very high rates not limited by the bandwidth of the upconversion chain.

Figure 18 shows the basis principle behind the direct antenna modulation technique, where an a dipole antenna in this case is driven by a CW signal of constant amplitude and phase. There is also a conductive metal line with compa-

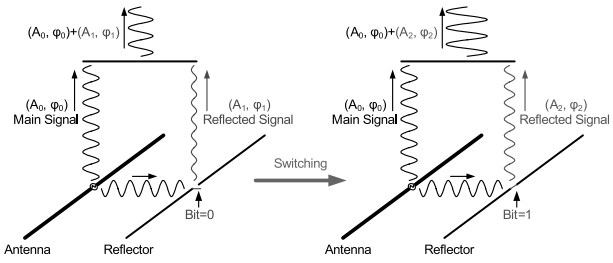

Fig. 18 Single bit modulation of the transmitted signal by digital switching the reflector in the near field.

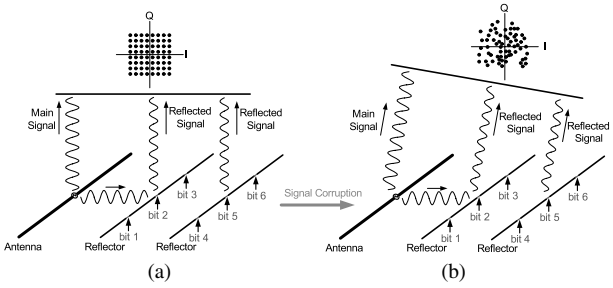

Fig. 19 Multiple reflectors with multiple switches to create different constellation points in (a) the intended direction and (b) in the unintended one [10].

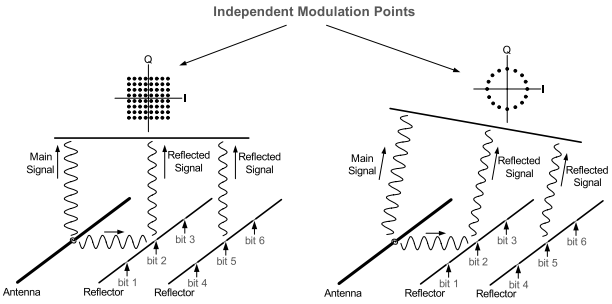

Fig. 20 Multiple reflectors with multiple switches creating different constellation points.

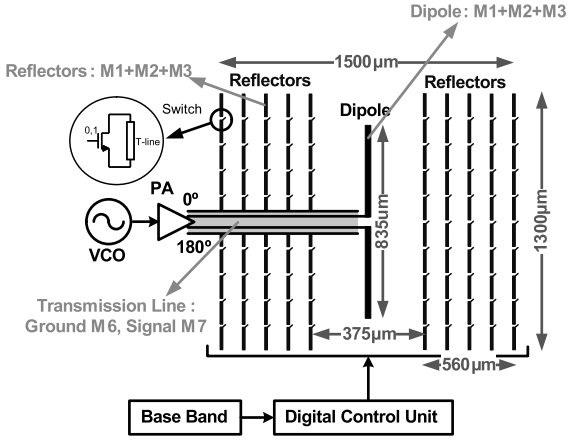

Fig. 21 The practical implementation of the digital antenna modulation scheme [10].

rable dimensions to the wavelength (a reflector) next to the antenna which can be shorted or opened at some point along its length using a switch. The reflected signal interferes with the main signal radiated by the antenna in a given direction. The amplitude and phase of the reflected signal depend on the boundary conditions that the reflector imposes and can be varied by turning the digital switch on or off. The two states of the switch result in two different phases and amplitudes in the direction of interest hence generating two distinct points in the I-Q plane. This provides a simple and crude one-bit digital modulation without changing the output power or phase of the PA driving the antenna hence allowing it to operate at its highest efficiency.

Although the single-reflector single-switch configuration of Fig. 18 provides a basic binary modulation, it offers limited control over where the two modulation points fall. More control over the constellation point (i.e., transmitted phase and amplitude) can be obtained by introducing multiple reflectors each with multiple switches, as shown in Fig. 19(a). A large enough number of reflectors and switches in close proximity of the antenna makes it possible to create more constellation points and eventually achieve full coverage across the constellation.

It is noteworthy that a set of switch combinations that generates constellation points of the given modulation in the intended direction generates different constellation points in a different direction, as shown in Fig. 19(b). This can be used to scramble the signal further in the unintended directions to implement a secure communication link by preventing an undesired eavesdropper from demodulating and recovering the signal.

For a large number of switches and reflectors, there are numerous different switch combinations that generate almost the same constellation point in a given direction. However, these different switch combinations can be used to generate distinct points in other directions simultaneously, allowing concurrent transmission of two independent streams of data at full rate in two different directions, as depicted in Fig. 20.

The viability of this technique is demonstrated via an integrated circuit implementation [10] with an on-chip dipole antenna with 10 reflectors (5 on each side), each with 9 tuned MOS switches along its length, resulting in a total number of 90 switches and a extremely large 2^{90} ($\sim 10^{27}$) switching combinations, as illustrated in Fig. 21. The very large number of switching combinations provides numerous ways to generate a desired point on the constellation (phase and amplitude) in a given direction, providing many additional degrees of freedom that can be used for concurrency or security. The ability to simultaneously send independent information to several directions at full rate using a single transmitter is not achievable using the conventional transmitter architectures.

It is self-evident at this point that the ability to switch the antenna reflectors in the near field using a large number of very fast switches is a direct result of our ability to integrate, antennas, RF, and digital circuits on the same substrate. This is an example of the holistic co-design of elec-

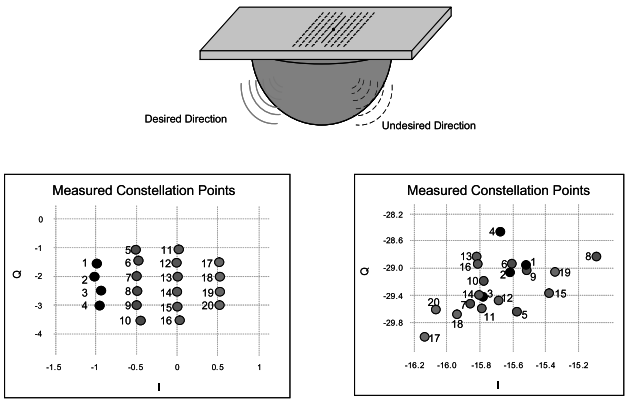

Fig. 22 The direct antenna modulation chip with a dielectric lens; a sample of the measured constellation points in the intended and unintended directions [10].

tromagnetic structures with analog and digital circuitry. In practice, designing such system pose challenges not only at the architecture and circuit levels, but also with respect to the simulation tools and methodology.

Making a high quality and fast integrated switch for use in mm-wave frequencies is challenging. Achieving a small on-impedance requires a relatively large transistor size, which itself limits the maximum achievable off-impedance due to its large parasitic capacitance. This can be alleviated by resonating the transistor parasitic capacitance using a transmission line connected between its drain and its source (Fig. 21).

A transmitter at 60 GHz using the direct antenna modulation scheme with an on-chip dipole antenna and reflectors is implemented in silicon. The antenna, the ten reflectors, the ninety switches, and the digital control circuitry for the switches occupy a die area of roughly 1.5 mm × 1.5 mm. The power is radiated from the backside of the chip with the aid of a hemispherical dielectric lens [5], as shown in Fig. 22.

An experimental demonstration of this technique is also shown in Fig. 22, where we show 20 points generated purely by switching the reflectors using 20 different switch combinations on the left hand side of Fig. 22. It should be noted that the input power and phase of the signal driving the antenna has not been changed to achieve the modulation. In addition, on the right side of Fig. 22, the received signal constellation at a different angle (about 90° off) for the same set of switch combinations is plotted. It is obvious that these points are completely scrambled, making it practically impossible for a receiver at the undesirable angle to recover this signal, no matter how sensitive it may be. This is in contrast with a traditional transmitter that sends essentially the same modulated signal, albeit with different amplitude scales and phase epochs, in all directions. This measurement demonstrates the feasibility of the idea for security of communication links. The antenna is driven by a three stage fully-differential PA with an output power of at least +7 dBm. The die photo is shown in Fig. 23.

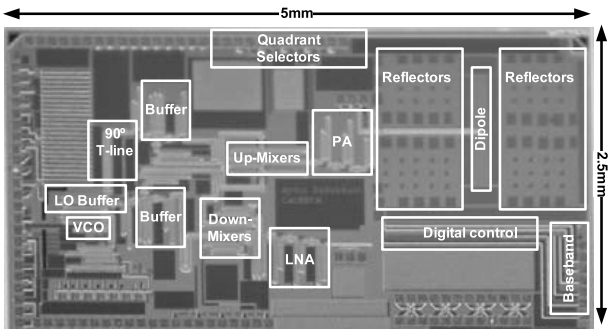

Fig. 23 The die micrograph of the direct antenna modulation transmitter.

12. Conclusions

For the first time, it is possible to have a very large number of extremely fast and reliable transistors capable of operating at mm-wave frequencies in close proximity of structures comparable to the wavelengths of interest on the same silicon substrate, where electromagnetics, analog, and digital circuits can be intertwined deeply. This unique combination makes is possible to consider novel holistic approaches to communication and radar systems operating at mm-wave frequencies only limited by our imagination and creativity.

Acknowledgments

The author would like to thank A. Babakhani, Y. Wang, and H. Wang, Prof. D.B. Rutledge, Prof. S. Weinreb of Caltech and Dr. A. Natarajan, Dr. I. Aoki, Dr. S. Kee, Dr. A. Komijani, Dr. X. Guan, Prof. H. Hashemi, Prof. J. Buckwalter, and Prof. E. Afshari formerly of Caltech for their numerous contribution to Caltech's mm-wave activities. We have benefitted from the support of Caltech's Lee Center for Advance Networking, Raytheon Company, National Science Foundation, and DARPA Trusted Foundry Program.

References

[1] J.S. Kilby, United States Patent Number 3,138,743, Filed Feb. 6, 1959, Granted June 23, 1964.
[2] J.S. Kilby, United States Patent Number 3,261,081, Filed Feb. 6, 1959, Granted July 19, 1966.
[3] R.F. Stewart, United States Patent Number 3,138,747, Filed Feb. 12, 1959, Granted June 23, 1964.
[4] G.E. Moore, "Cramming more components onto integrated circuits," Electronics, vol.38, no.8, pp.114–117, April 1965.
[5] 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.
[6] 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.
[7] X. Guan and A. Hajimiri, "A 24-GHz CMOS front-end," IEEE J. Solid-State Circuits, vol.39, no.2, pp.368–373, Feb. 2003.
[8] A. Hajimiri, A. Komijani, A. Natarajan, X. Guan, and H. Hashemi,

"Phased array systems in silicon," IEEE Commun. Mag., vol.42, no.8, pp.122–130, Aug. 2004.
[9] H. Hashemi, X. Guan, A. Komijani, and A. Hajimiri, "A 24-GHz SiGe phased-array receiver — LO phase shifting approach," IEEE Trans. Microw. Theory Tech., vol.53, no.2, pp.614–626, Feb. 2005.
[10] A. Babakhani, D.B. Rutledge, and A. Hajimiri, "A near-field modulation technique using antenna reflector switching," ISSCC Digest of Technical Papers, vol.51, pp.188–189, Feb. 2008.
[11] A. Hajimiri, "mm-wave silicon ICs: Challenges and opportunities," IEEE Custom Integrated Circuits Conference, Sept. 2007.
[12] C. Doan, S. Emami, A. Niknejad, and R. Brodersen, "Millimeter wave CMOS design," IEEE J. Solid-State Circuits, vol.40, no.1, pp.144–155, Jan. 2005.
[13] B. Floyd, S. Reynolds, U. Pfeiffer, T. Zwick, T. Beukema, and B. Gaucher, "SiGe bipolar transceiver circuits operating at 60 GHz," IEEE J. Solid-State Circuits, vol.40, no.1, pp.156–167, Jan. 2005.
[14] CCIR Doc. Rep. 719-3, "Attenuation by atmospheric gases," ITU 1990.
[15] L.J. Ippolito, "Propagation effects handbook for satellite systems design," NASA Doc., vol.1082, no.4, Feb. 1989.
[16] E.K. Smith, "Centimeter and millimeter wave attenuation and brightness temperature due to atmospheric oxygen and water vapor," Radio Science, vol.17, pp.1455–1464, Nov.-Dec. 1982.
[17] Federal Communications Commission, "Millimeter wave propagation: Spectrum management implications," Bulletin Number 70, July 1997.
[18] D. Lu and D. Rutledge, "Investigation of indoor radio channel from 2.4 GHz to 24 GHz," IEEE AP-S Int. Symp. Dig. Papers, pp.134–137, June 2003.
[19] J.D. Kraus and R.J. Marhefka, Antennas, McGraw Hill, 2001.
[20] D.B. Rutledge, D.P. Neikirk, and D.P. Kasilingam, Integrated circuit antennas, Infrared and Millimeter Waves, Millimeter Components and Techniques, Academic Press, New York, 1983.
[21] 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, March 2002.
[22] E. Afshari, H. Bhat, X. Li, and A. Hajimiri, "Electrical funnel: A broadband signal combining method," Dig. International Solid-State Circuits Conference, Feb. 2006.
[23] A. Babakhani and A. Hajimiri, "mm-wave phased arrays in silicon with integrated antennas," Proc. Antenna and Propagation Symposium, June 2007.
[24] V. Aulock and W.H. von Aulock, "Properties of phased arrays," Proc. IRE, vol.48, no.10, pp.1715–1728, Oct. 1960.
[25] R.C. Hansen, ed., Significant Phased Array Papers, Artech House, Norwood, MA, 1973.
[26] R.S. Elliott, Antenna Theory and Design, Prentice-Hall, Englewood Cliffs, NJ, 1981.
[27] M. Golio, ed., The RF and Microwave Handbook, Session 6.9, CRC Press LLC, 2000.
[28] D. Parker and D.C. Zimmermann, "Phased arrays — Part I: Theory and architectures," IEEE Trans. Microw. Theory Tech., vol.50, no.3, pp.678–687, March 2002.
[29] D. Parker and D.C. Zimmermann, "Phased-arrays — Part II: Implementations, applications, and future trends," IEEE Trans. Microw. Theory Tech., vol.50, no.3, pp.688–698, March 2002.
[30] B. Kane, L. Geis, M. Wyatt, D. Copeland, and J. Mogensen, "Smart phased array SoCs: A novel application for advanced SiGe HBT BiCMOS technology," Proc. IEEE, vol.93, no.9, pp.1656–1668, Sept. 2005.
[31] "IRE standards on electron tubes: Definition of terms," Proc. IRE, vol.45, pp.983–1010, July 1957. 57 IRE 7. S2.
[32] L. Zhu, "Guided-wave characteristics of periodic coplanar waveguides with inductive loading: Unit-length transmission parameters," IEEE Trans. Microw. Theory Tech., vol.51, no.10, pp.2133–2138, Oct. 2003.
[33] T.S.D. Cheung, J.R. Long, K. Vaed, R. Volant, A. Chinthakindi, C.M. Schnabel, J. Florkey, and K. Stein, "On-chip interconnect for mm-wave applications using an all-copper technology and wavelength reduction," ISSCC Digest of Technical Papers, vol.46, pp.396–397, Feb. 2003.
[34] A. Komijani, A. Natarajan, and A. Hajimiri, "A 24 GHz, +14.5 dBm fully-integrated power amplifier in 0.18 μm CMOS," IEEE J. Solid-State Circuits, vol.40, no.9, pp.1901–1908, Sept. 2005.
[35] A. Komijani and A. Hajimiri, "A wideband 77 GHz, 17.5 dBm power amplifier in silicon," IEEE J. Solid-State Circuits, vol.41, no.8, pp.1749–1756, Aug. 2006.
[36] E. Afshari, H.S. Bhat, A. Hajimiri, and J.E. Marsden, "Extremely wideband signal shaping using oneand two-dimensional nonuniform nonlinear transmission lines," J. Appl. Phys., vol.99, pp.54901–16, 2006.
[37] E. Afshari, H. Bhat, X. Li, and A. Hajimiri, "Electrical funnel: A novel broadband signal combining method," ISSCC Digest of Technical Papers, vol.49, pp.24–25, Feb. 2006.
[38] I. Aoki, S. Kee, R. Magoon, R. Aparicio, F. Bohn, J. Zachan, G. Hatcher, D. McClymont, and A. Hajimiri, "A fully-integrated quad-band GSM/GPRS CMOS power amplifier," ISSCC Digest of Technical Papers, vol.51, pp.570–571, Feb. 2008.

Ali Hajimiri received the B.S. degree in Electronics Engineering from the Sharif University of Technology, and the M.S. and Ph.D. degrees in electrical engineering from the Stanford University in 1996 and 1998, respectively. He was a Design Engineer with Philips Semiconductors, where he worked on a BiCMOS chipset for GSM and cellular units from 1993 to 1994. In 1995, he was with Sun Microsystems, where he worked on the UltraSPARC microprocessor's cache RAM design methodology. In 1997, he was with Lucent Technologies (Bell Labs), Murray Hill, NJ on a short-term assignment, where he investigated low-phase-noise integrated oscillators. In 1998, he joined the Faculty of the California Institute of Technology, Pasadena, where he is a Professor of Electrical Engineering and the director of Microelectronics Laboratory. His research interests are high-speed and RF integrated circuits. Dr. Hajimiri is the author of The Design of Low Noise Oscillators (Boston, MA: Springer, 1999) and has authored and coauthored more than one hundred refereed journal and conference technical articles. He holds more than two dozens U.S. and European patents. He is a member of the Technical Program Committee of the International Solid-State Circuits Conference (ISSCC). He has also served as an Associate Editor of the IEEE Journal of Solid-State Circuits (JSSC), an Associate Editor of IEEE Transactions on Circuits and Systems (TCAS): Part-II, a member of the Technical Program Committees of the International Conference on Computer Aided Design (ICCAD), Guest Editor of the IEEE Transactions on Microwave Theory and Techniques, and the Guest Editorial Board of Transactions of Institute of Electronics, Information and Communication Engineers of Japan (IEICE). Dr. Hajimiri was selected to the top 100 innovators (TR100) list in 2004 and is a Fellow of Okawa Foundation. He is a Distinguished Lecturer of the IEEE Solid-State and Microwave Societies. He is the recipient of Caltech's Graduate Students Council Teaching and Mentoring award as well as Associated Students of Caltech Undergraduate Excellence in Teaching Award. He was the Gold medal winner of the National Physics Competition and the Bronze Medal winner of the 21st International Physics Olympiad, Groningen, Netherlands. He was a co-recipient of the IEEE Journal of Solid-State circuits Best Paper Award of 2004, the International Solid-State Circuits Conference (ISSCC) Jack Kilby Outstanding Paper Award, two times co-recipient of CICC's best paper awards, and a three times winner of the IBM faculty partnership award as well as National Science Foundation CAREER award. He is a cofounder of Axiom Microdevices Inc.