UDC 681.5

D. FOTY

Gilgamesh Associates / Sarissa Radio, Inc.
Fletcher VT / Palo Alto CA / Atlanta GA USA

NEXT-GENERATION WIRELESS NETWORKS: NEW BUILDING BLOCKS AND NEW NETWORK TOPOLOGIES

As data content mushrooms and becomes increasingly dispersed, the demand for (and need for) bandwidth – both wired and wireless – continues to increase; however, technology has demonstrably not been up to this task. A major intellectual gap has been the attempt to address this demand via "digital thinking," in which transistors and power can be used gratuitously to overwhelm the problem; this is simply not working. Despite a number of challenges, millimeter-wave radio is slowly emerging as a unique way to transcend the "digital mindset" and address a wide variety of bandwidth challenges – by making bandwidth available without outlandish demands in complexity and power consumption. In addition to the "usual" applications, the basic nature of millimeter-wave radio offers the prospect of revolutionizing network topologies, allowing for the development of new network infrastructures – such as mesh networks.

network topologies, bandwidth

Introduction

At this time, the demand for (and need for) increased data bandwidth continues relentlessly. Of particular note, not only is the "scale" of the situation – more content and thus more data – the only factor; more of our data is dispersed in a wider fashion, and more of it is riding aboard the various mobile devices which have become staples of everyday life. As data becomes more mobile – and as the amount of storage available in mobile devices increases drastically – this "data diaspora" requires that capabilities for moving, sharing, and synchronizing data across various devices be drastically improved from what is now available.

However, present-day approaches to technology are failing to keep up with this progress and are thus failing to address those needs. This is largely due to a misunderstanding of the constraints on the situation, and thus in attempts to address this opportunity via the long-used "digital" approach of deploying large numbers of transistors consuming large amounts of power. Clearly, a fundamentally different approach to this challenge is required.

This paper will examine a different route to address this challenge – one that is simple, direct, and more robust. New technologies based on millimeter-wave radio offer the opportunity to transcend the "transistors and power" mindset, offering very large amounts of wireless bandwidth without the now-unacceptable levels of complexity and power consumption. In addition, a successful and cost-effective millimeter-wave radio technology can go beyond just this particular set of applications – it can also provide the long-sought building blocks for new data network infrastructures, such as dense-mesh and hopping networks. These sorts of networks can be used to provide ambient intelligence and ubiquitous data connectivity.

Reliability/Dependability - "Impose" vs. "Build-In"

Next-generation wireless networks must break new ground on a variety of fronts. In addition to the usual factors of data rate, power consumption, and range, other factors include form factors, dependability, and reliability. As a prequel to this discussion, we will first consider several interesting guideposts for the examination of this problem.

First, we can note a trenchant observation by the noted technology analyst George Gilder [1]:

Every economic era is based on a key abundance and a key scarcity.

This is a very useful observation. Right now, as Gilder noted (back in 1996!), we have an abundance transistors (and thus, of bits and MIPs) and of watts; problems are "solved" by throwing large numbers of transistors at them, and by allowing for the unlimited consumption of (large amounts of) power. However, as Gilder also notes, we face a shortage of bandwidth – at all levels. Both the demand and the *need* for bandwidth continue to increase – but those needs cannot be met by the expedient and profligate use of transistors and watts. In the wireless realm in particular, we face a shortage of natural wireless bandwidth (due to the limited nature of the presently-used carrier frequencies). This paucity of wireless bandwidth is a key opportunity – provided other factors are not ruined in the attempt to reach that prize.

We can also note another trenchant observation, this one by the late Czech-American engineer/philosopher Petr Beckmann [2]:

In a healthy society, engineering design gets smarter and smarter; in an [unhealthy society], it gets bigger and bigger.

This is also a useful consideration. It is very easy to try to address a problem with brute force – by just making things bigger and more complicated, rather than by thinking about things and coming up with simple yet clever engineering. In the integrated circuit realm, this "bigger is better" approach often manifests itself as "solving" a problem by the simple expedient of consuming more power.

A third consideration is to understand the difference between a problem of "scale" and a problem of "knowledge." Many technology problems are simply ones of "scale" – the basic principles are clear and settled, while the difficulties are centered on the sheer amount of something which is in play. In contrast, problems of "knowledge" require clear thinking about the underlying circumstances, and how present methods of handling them are inadequate. In technology, tremendous disasters often develop when a "problem of knowledge" is approached as being a "problem of scale." A corollary is that there is often confusion about what is a "digital problem" and what is really an "analog problem."

A fourth and final consideration is that of basic engineering trade-offs. Reflecting the previous paragraph, with all things remaining equal, good engineering involves managing the trade-offs among a number of countervailing factors – a "gain" in one factor can only be had by paying some price in one or more of the other interlocking factors. It is a separate task to find ways to expand the scope of the trade-off matrix, and to thus expand the capabilities of the trade-off space. Mixing these two separate things together frequently leads to comical outcomes.

This in essence presages the situation today with the development of next-generation high-speed wireless data networks. The "old" approaches can only be made to work by providing gains (e.g., data rate) at a terrible price in other factors (e.g., power consumption and range). In a similar way, many of the problems related to network reliability and dependability are inherently related to the basic coarseness of the network; rather than trying to "impose" these attributes *ex post facto* onto the network, a more sensible approach is to change the basic nature of the network itself.

We will also employ the principle once articulated by the great American philosopher Yogi Berra: "You can observe a lot just by watching."

Demand for Wireless Bandwidth

Despite various ups and downs, the overall demand for data bandwidth (both wired and wireless) always increases. Analogies are possible, but a good one is a comparison to the situation with equities and equity prices; while there are various ups and downs, the overall long-term trend is for equity values to appreciate.

What makes the present situation in data bandwidth historically unique is the convergence of previously disparate content – audio, video, voice, and data – into the same collection. What this also represents is a simple reality – all of these items (and more) are actually "data" *in toto*, even though this was not previously the case. This convergence provides a general template, but one can also note that an implicit factor is the need for increased mobility for that "data."

Of particular note, video is rapidly emerging as the high-content driver of demand for bandwidth. Last summer, it was noted that the popular youtube.com video-sharing web site had reached a daily level of traffic of 250 *terrabytes* per day; this translates into an *average* data throughput of 23Gbps! More recently, John Chambers, CEO of Cisco Systems, explicitly noted that video is emerging as the "killer app" for growth in data networking.

An additional factor is an interesting problem with compression technology, which is supposed to reduce the amount of "data" which must be stored and transmitted in video applications. For example, the "raw" bandwidth requirement of the new High Definition Television (HDTV) technology is on the order of 1.5 – 2.2 Gbps. The original intention was that this data rate would not be necessary, because compression methods would serve to reduce the *de facto* required data rate. However, the image quality in the new screens has outrun the abilities of compression technology; "compressed" video produces a discernibly poorer image.

Taken together, the overall demand for bandwidth tells us that the concept of "anywhere, anytime" communications is moving closer to becoming a reality. This concept goes back at least a decade [3], and it was noted back then that a high-quality, high-speed wireless data communications capability is required for this vision to become a reality [4]. This is illustrated by a 1994 postulation (fig. 1), showing a variety of layers and capabilities for access to data.

As this 1994 vision becomes a reality, some parts are slightly different than expected. What is now happening is a "data diaspora" – that is, there are more and more devices with data access, and more and more of them are mobile. (For example, mobile/portable devices now vastly outnumber desktop PCs). However, there is no longer a need to trade off data storage and computational abilities in return for enhanced mobility; portable devices, taking advantage of stunning advances in flash storage and micro-disk-drives, now have huge amounts of on-board storage ability. Moving data around the "diaspora" is demanding, and there is a growing need to *synchronize* all of this disparate data across disparate devices.

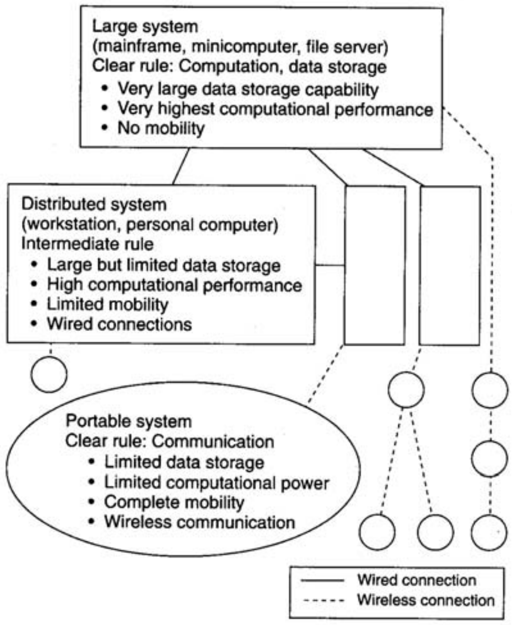

Fig. 1. A 1994 vision of an integrated computation/communications infrastructure; this concept is still a work-in-progress, requiring further advancements in high-speed wireless data technology.

However, wired data rates are way ahead of wireless data rates, and wireless data rates simply have not been keeping up; this is a major headache for future mixed networks, as a "seamless infrastructure" is increasingly required. For example, contemporary wired data rates include Ultra3 SCSI at 1.28Gbps and FC-AL Fiber Channel at 800Mbps – 3.2Gbps; in contrast, WLAN ("Wi-Fi") operates in geographically-restricted spots at no better than 54Mbps, more mobile 3G networks can do no better than 2 – 2.5 Mbps, and the highly mobile "EDGE" types of networks can only manage 130 – 250 *kbps*. In addition, as this wireless data list notes, the wireless infrastructure is fragmented into several specialized "tiers" that do not work well with each other. Both data rates *and* coherence must be improved if wireless data networking is to integrate smoothly into overall data networking.

For portable devices, there are four interlocking issues that will dominate all technology questions – storage, computation, mobility, and communication. Connected with those constraints, there are five major factors driving wireless bandwidth:

- rapid growth of multimedia content;
- dispersal of data (the "data diaspora");
- the need for ubiquitous connectivity (and thus equivalent wired/wireless data rates);
- growth in local storage due to high-capacity flash-drives and micro-drives;
- the emergence of new network topologies, such as mesh and sensor networks.

To date, there have been several attempts to address these needs; unfortunately, these efforts have been undertaken without proper consideration of the engineering principles mentioned earlier, leading to less-than-successful results.

What's Not Working

With these obvious needs in wireless data networking, why haven't they been met? At the philosophical level, the various "popular" approaches have neglected to take into account the principles elucidated earlier – things like Beckmann's postulate ("bigger" vs. "more clever"), confusing "scale" and "knowledge" problems, and trying to use "digital methods" to solve what is essentially an "analog problem." The landscape is now littered with failed (and failing) technologies.

The essence of all of these failed/failing methods has been an attempt to use a "trick" approach of trying to "fake" bandwidth by using a very large number of data channels over a relatively large frequency range; for example, the "ultra-wideband" (UWB) approach (and its cousins) involves trying to use a large number of parallel channels over the immense frequency range of 3.1 – 10.6 GHz. These approaches are the radio equivalent of parallel processing – but there are inherent problems which should have been obvious at the outset. First, breaking up the data at one end for multi-channel transmission and then reconstituting it at the other end requires an immense amount of baseband processing; in addition to the obvious complexity and integrated-circuit size required to put a virtual supercomputer at each end to do the baseband processing, this kind of baseband-centric approach is doomed to require immense amounts of power to function. This is a classic example of misdirected engineering – here, trying to treat an analog problem as a digital problem, and trying to solve a radio (analog) problem via an immense amount of baseband (digital) processing. A further problem is the nearly-insurmountable difficulty presented by the immense relative width of the frequency range – one must either use a large number of separate antennae, or try to design a complicated and finicky antenna; in this environment, minor factors become important, and equipment shapes/sizes become critical factors in system performance. A final indignity is that the immense relative frequency spread leads to nearly incomprehensible and basically unworkable in-room propagation characteristics.

The blunt message in these difficulties is that the limitations inherent in today's "popular" carrier frequencies are a natural barrier limiting their further exploitation. At the simplest level, lower frequencies provide less "natural" bandwidth. In addition, by trying to use a wide frequency range, the propagation and absorption behavior across the frequency spread varies widely; for example, building-material attenuation is very low at the popular 2.4GHz, but is much higher at 5GHz. In any UWB-type of approach, the propagation characteristics of the different channels will vary to extreme degrees.

What has been missing from thinking regarding this approach is a proper understanding of the "Iron Triangle" of wireless data communications. Three factors – data rate, range, and power – are intertwined, and they also trade off *directly against each other*. This is the "trade-off matrix" which limits room-for-maneuver, as one factor can only be improved at the expense of the others; in UWB types of approaches, the data rate can only be improved by paying a severe penalty in both the power consumption required and in the data transmission range. These factors trade off directly – and one must either work with the trade-off matrix at hand, or find some new way to expand that trade-off matrix.

Thus, UWB methods have choked on the severity of the baseband demands and the extreme relative range of the operating frequencies. Operating range has constantly been scaled back, portable applications have been discarded, use scenarios have contracted dramatically, data rates have been greatly reduced, and the propagation characteristics have proven to be troubling at best; similar problems have infected related technologies such as WUSB and 802.11n.

This clumsy approach was destined not to work, and that has proven to be the case. A fundamentally different approach to the challenge is required.

Moving to Higher Carrier Frequencies

An alternative strategy is to move to higher carrier frequencies for communications. The use of higher carrier frequencies is the most "natural" way to produce higher bandwidth – and to do so in a way that provides the greatest simplicity along with low power consumption. Essentially, this is a direct method of expanding the trade-off matrix – that is, expanding the aforementioned "Iron Triangle," making the overall trade-off interactions much more favorable and addressing the "key scarcity" problem mentioned earlier.

Of particular note, the "millimeter wave" region of the electromagnetic spectrum (30 – 300 GHz) is a good candidate for addressing our present problems, challenges, and opportunities. Millimeter-wave radio offers two key benefits. First, as noted above, the higher carrier frequencies provide a "natural" abundance of bandwidth. Second, the shorter wavelengths allow for a very small size ("form factor") in system equipment and the system antenna; this second point has a number of interesting implications.

To rephrase this slightly, millimeter-wave radio offers the promise of very high data bandwidth – but with low power consumption and reasonable transmission ranges. With a simple transmission scheme thus enabled, the propagation characteristics of a radio system are greatly simplified – and it is this simplicity that leads to both low system costs and high reliability/dependability. Furthermore, the short wavelengths allow for very small form factors; in addition to the engineering benefits of small size and simplicity, such tiny systems offer aesthetic benefits – tiny systems are less obtrusive and thus are more acceptable for general use in the everyday world.

Of course, developing millimeter-wave radio technology requires that two factors be present. First, as will be discussed below, basic technology must be capable of operation at sufficiently-high frequencies – and with sufficiently low cost – to allow systems to be constructed. Second, national and international regulatory bodies must make appropriate spectrum allocations. In the United States, the Federal Communications Commission (FCC) has already allocated *licensed* bands at 70, 80, and 90 GHz. However, of larger interest is the band now available at 60GHz; this is an *unlicensed* allocation of some 5GHz of bandwidth – an allocation that has already been made available internationally.

Underlying Technology Issues

The development of successful, useful, *low-cost* millimeter-wave radio systems requires the coming-together of a number of technological factors. What is critical to grasp here is that the key questions are not ones of "scientific feasibility" – as millimeter-wave radio systems have existed for years, albeit it in very expensive and specialized applications. The key questions involve the abilities of "mainstream" technologies to realize the goals in a cost-effective manner. This requires careful evaluation and intelligent choices among the possible options and roadblocks.

Millimeter-wave radio is not a new technology; it has been used in a few specialized applications for many years, but has been a difficult and expensive technology. Integrated circuits for millimeter-wave applications have used very expensive heterostructure bipolar transistors (HBTs) based on gallium arsenide and other (relatively) exotic III-V semiconductor materials; the supply of these technologies is very limited, and their overall cost renders them impractical for "mainstream" applications. At the other end of the spectrum, CMOS is a very ubiquitous and very low cost technology that is in wide use; however, the success of CMOS has been based mostly in digital applications, and RF-CMOS has remained problematic for many reasons. In recent years, there has been considerable hype surrounding the possible use of CMOS for millimeter-wave radio; however, these reports have been in very artificial situations where the constraints were far below those required for real products – instead being carefully set up to produce conference papers. The reality of the situation is that in the product marketplace today, there are no CMOS products available at frequencies above about 5GHz; the notion that there is some magic way to "leap" CMOS to performance levels more than ten times higher simply lacks credibility. The most interesting integrated circuit technology for millimeter-wave radio at this time is modern silicon-germanium BiCMOS technology – which now offers high bipolar speeds along with reasonable costs and trustworthy behavior.

At the design level, the most damaging problem in today's RF-IC world – yet one that it is considered virtually uncouth to mention in polite company – is the problem of excessive design iterations. In the "digital world," the design process tends to be straightforward, but it is too-often forgotten that this ease is largely due to the inherent coarseness of the design task. It is assumed that the pre-fabrication design will proceed simply to an end goal, that the silicon results will fall into place rather nicely, and that very few silicon iterations will be necessary. That is very often the case in digital design, but when that thinking is transferred to the analog/RF world, very different things happen. At the design level, designs often will not close up – simply because the problem is a more precise one and our ways of dealing with that problem are presently inadequate to the task. Even worse, the silicon results are also unable to close up with the design – and the time lag for another silicon run is on the order of three months. With time and money wasting, the final madness of the situation is the very expensive panic exercise of "massive silicon guessing" – that is, running a very large number of silicon lots in parallel with minor variations between them, hoping to bump into the right answer. This approach is frightfully expensive, and also guarantees neither the attainment nor the retention of success – since again it is trying to apply "scale" to what is really a problem of insufficient "knowledge."

Another item to note is that in the RF-IC world in particular (and in the integrated circuit world in general), the era of "independent" IC design – where a "bare-die" IC could be built and verified, victory declared, and the IC sold along to the system integrator – is drawing to a close. The interactions of the IC with packaging and the final system are becoming too strong too neglect; sharply rising test costs are one symptom of this emerging situation. Essentially, everything is becoming system design; parts of that effort are no longer able to operate independently, as the interactions of *all* pieces are critical to what ultimately matters – which is final *system* performance.

One final problem of note is not technical, but organizational. In the world of RF-IC design, the integrated circuit world bumps up against the worlds of data communications and telecommunications. These latter spheres are accustomed to much higher levels of regulation and "standardization" than is the integrated circuit world. Problems occur when the datacomm/telecomm worlds try to impose their levels of "standardization" onto the IC world – which is actually their supplier. This over-standardization eliminates nearly all room for IC innovation, as the "standardization" dictates far too many details – essentially, over-regulating the IC developers and "regulating away" the early and profitable parts of the normal product lifecycle. This must not be allowed to happen, as it destroys innovation in the IC realm – thus destroying suppliers, cutting off the *raison-d'être* for future investment, and eliminating the path to even better products. This problem manifested itself in the "bloodbath" which swept through the 802.11 chipset supplier market – where there were plenty of "successful" (in terms of working products) companies that expired due to inability to turn a profit. This over-standardization also spills over into "standards" committees – where the slow pace of the data-comm/telecomm world simply doesn't mesh well with the fast-paced IC world. Committees move too slowly, reduce technology practice to lowest common denominator thinking, and turn into valueless turf wars – while also being demonstrably susceptible to corruption.

These challenges are not technical, but intellectual. They are also critical, since misunderstanding these challenges is a serious impediment to progress. The insights of the various observers cited earlier provide a valuable resource here.

Extended Implications

The regulatory existence of millimeter-wave spectrum allocations would seem to imply that a new generation of wireless data technology could be developed that would basically mirror existing (lower-frequency) technology – essentially, that technology repeated but this time on steroids. However, there are a number of extended mitigating factors that need to be considered – factors that make the millimeter-wave radio fundamentally different, and which provide some interesting implications for future development.

As noted above, millimeter-wave radio offers two key benefits – a "natural" abundance of bandwidth (without the need for "tricks" and games), and the natural ability to make very small systems and antennae (which has both practical and aesthetic benefits).

However, the millimeter-wave region faces a fundamental limitation. As noted earlier, lower frequencies (such as the popular 2.4GHz) have insufficient natural bandwidth, but propagate well through both air and common building materials. In contrast, millimeter-wave frequencies offer very large natural bandwidth, but propagate modestly through air and poorly through building materials. This is a serious nature-based conundrum faced by attempts to extend wireless data technology. For example, a notion beloved in many sectors of that of an "information furnace" for a house – a single, simple "box" sits in the cellar and is able to blast out multi-gigabit data channels (e.g., for streaming video) with ease. This is a pleasant vision – but the basic laws of physics cited above prohibit its realization.

As noted earlier, many of the extant wireless network reliability/dependability problems are largely related to the relative coarseness of the network. This strongly hints that rather than trying to attempt an *ex post facto* "fix" of imposing reliability/dependability on an older network structure, perhaps it's time to think about changing the nature of the networks themselves.

Considering the combination of much higher "natural" bandwidth available with millimeter-wave radio and the concomitant long-propagation difficulties, the natural suggestion from the propagation physics is that the simplest way to address the problem is to densify the network grid – that is, to reform networks to a basis of a larger number of nodes that are much closer together.

These sorts of networks are already a topic of considerable discussion, and there are already some early practical deployments. These are mesh networks and hopping networks, comprised of a large number of transceiver nodes in a dense coverage environment. Mesh networks are self-organizing and easily scalable; they also are more robust and reliable by default, since the large number of nodes provides built-in redundancy in case of node failure or traffic overload. These mesh networks can be further "smartened up" by including *ad hoc* (very free range but still self-organizing) and sensor capabilities.

While there is a great deal of higher-level thinking on mesh networks (at the software and "organization" levels), there is very little thinking about what building blocks would best form the basis of such networks. For example, Microsoft Research has already put considerable effort into pondering the software needs and implications for these sorts of networks – while simply assuming that the building blocks will magically appear.

Millimeter-wave radio systems provide ideal, tiny (thumb-sized) nodes for advanced mesh networks. The data rates are naturally high, and the systems are unobtrusively small. The antenna (even a small dish!) can be built *inside the package* for a truly complete network node. These complete system nodes are also very inexpensive, so that mesh and hopping networks can be constructed at reasonable cost – and can thus be deployed widely.

The small antenna size also makes possible a further intriguing possibility. Phased-array antennae were originally developed for radar systems; since the beam was formed by carefully-phased driving of numerous simple antennae, the beam can be directed without the need for any mechanical moving parts (as is the case with a dish antenna). In addition to the obvious reliability advantages (few or no moving parts), this kind of antenna system can quickly and easily track multiple targets with ease.

The use of millimeter-wave frequencies allows for the construction of phased array antenna systems that are only a few square centimeters in size; these can also be placed directly inside the very small transceiver system package. Mesh networks based on this sort of building block are very simple yet very powerful – since they require neither a fixed directed antenna nor an isotropic transmission at all times (which wastes power). The various mesh nodes can link and interact completely at the software level, where network management is able to quickly and efficiently direct traffic and connect interacting nodes as necessary (and, when necessary, only briefly). This sort of network presents enormous software challenges; however, much of that thinking is already being done – and, harkening to an earlier comment, this particular situation is indeed more a problem of scale than one of knowledge.

A final interesting opportunity is the ability of these sorts of networks to build out data capability to what are today known as "underserved areas." In most of these situations, electrification has been provided, but data connectivity remains problematic – both logistically and financially. The electrification aspect is key, since electric utilities bring two critical items into play. The first and most obvious is that power is available at any utility pole, by definition. The second is subtle but equally critical – electric utilities also possess the legal rights-of-way that are required for the utility poles and cables. Thus, millimeter-wave nodes can easily (both technically and legally) be deployed onto existing utility poles – providing the backbone for mesh and hopping networks that will easily provide data access to the aforementioned "underserved areas."

A final note about this kind of network is that we must learn to think about our data networks a little bit differently – as the analysis of the "information furnace" indicated. Our thinking about wireless networks has tended to be centered on the concept of coarse networks with a relatively small number of powerful nodes. The "new thinking" required by mesh networks demands that we think of wireless data access in the same way that we now thinking of lighting – a large number of "lower power" nodes are placed ubiquitously and where they provide the most convenient usage. By making wireless networks ubiquitous, like lighting they essentially become a normal and expected part of the ambient background.

Taken together, this is an ambitious and demanding agenda; however the various benefits are obvious, as noted in several instances. In particular, this generalized approach unifies the network topologies and transcends the present fragmentation into "tiers" of differing mobility and capability. The use of millimeter-wave radio provides a ubiquitous, ambient, high-speed wireless data environment; this is a move up from "smart networks" to "brilliant networks." This presents a huge software challenge, but the present (and growing) nature of the "data diaspora" demands that next-generation networks move in this direction.

Summary/Conclusions

This paper has examined the present situation in high-speed wireless data communications. Despite the steadily-growing demand for wireless bandwidth, technology development in this area has become somewhat stultified – as it has focused on attempts to extend present-day technology with a variety of short-term "tricks." The "tricks" are only made possible by paying a very serious price in other engineering aspects of the technology, such as transmission distance, power consumption, and equipment form factor.

A better alternative is to take a route that uses higher carrier frequencies and demands an extensive amount of innovation; the challenges that are inherent in such an approach are manifold, demanding both tangible and intellectual innovation. In particular, millimeter-wave radio offers the opportunity to achieve high data rates in the most "cost effective" (both literally and engineering-wise) manner… while also opening the door to completely new network topologies, such as dense mesh networks.

To enumerate the most specifically-important points:

- The decentralization of data is accelerating far beyond earlier expectations; a major factor now is the availability of large-capacity micro-storage units (both flash drives and micro-drives) for portable devices; multimedia content is expanding rapidly to fill up this storage space.
- New network topologies are required, particularly as mesh and sensor networks become common; wireless data access must become like room lighting ("illumination), and will expand data access into new realms.
- Wireless data rates must improve rapidly, so as to catch up with wired data rates; this is required if wired and wireless technologies are to co-exist seamlessly.
- The telecomm/datacomm level of "acceptable" standardization is unacceptable in the IC/system world, as it stifles innovation and fatally distorts the business constraints; its imposition must be resisted.
- Attempts to increase bandwidth by using massive digital baseband processing are now counterproductive – in terms of cost, complexity, power consumption, and form factor; the use of higher carrier frequencies restores the proper balance between "analog" and "digital" tasking.
- Due to all the interactions, isolated integrated circuit design is insufficient, and must be replaced with
- system design; fundamentally new approaches to IC design and system design are required.

The key interlocking factors to watch are:

- The "data diaspora" is a combination of both more data and more devices.
- Millimeter-wave radio will enable the useful realization of new network topologies – and each of these will synergistically drive the other.
- Increased wireless data bandwidth will finally allow for the development of seamless, mixed (wireless and wired) networks.
- Relaxed standardization is critical if suitable "building blocks" are to be successfully developed and deployed.
- Attempts to increase wireless data bandwidth via more and more digital baseband processing have not worked out – and will not work out.
- All factors must be treated together as part of "system design" – rather than being over-segmented into tiny, unrelated component parts which can merely be tossed together at the end.
- Finally – and most importantly – *none of these factors can be isolated from the others*.

References

1. Gilder G. Wired (magazine), December 1996.

2. Beckmann P. A History of π. – New York: St. Martin's Press, 1971.

3. Borel J. General Introduction, Motivation, and Requirements for LV-LP ICs // Proceedings of the 1993 European Solid State Device Research Conference (ESSDERC). – P. 911-918.

4. Foty D., Nowak E. MOSFET Technology for Low-Voltage/Low-Power Appliations // IEEE Micro, June 1994. – P. 68-77.

Поступила в редакцию 7.03.2007

Рецензент: д-р техн. наук, проф. И.А. Фурман, Харьковский национальный технический университет сельского хозяйства им. Петра Василенко, Харьков.