InfoPad - Past, Present and Future

Robert W. Brodersen
rb@eecs.berkeley.edu
Dept of EECS, University of California, Berkeley

The InfoPad project was a large interdisciplinary effort, which had a goal of implementing a complete system for wireless access of multimedia data from a high bandwidth wired infrastructure. The basic approach was to minimize the cost and energy consumption of the portable user device by using network based resources. A system approach was taken in this investigation, which attempted to optimize all aspects of the design from the user interface to the low power design of the terminal hardware, which uncovered a need to deal with the issues of project management and intergroup communications. From what was learned in this project a new effort is now underway investigating the present and future capabilities of implementing single chip wireless systems in advanced CMOS technology.

I. Introduction

I would like to take this opportunity to thank SIGMOBILE for their recognition of the InfoPad project research activities. This was a large interdisciplinary effort that at its peak involved over 50 students, 7 faculty and 6 staff and 5 member companies. Particularly important to the success of the project was that it integrated activities in both computer science and electrical engineering as well as a number of disciplines within each of these areas. As is true in any good research project the technical topics in which we had our most interesting new results were in areas far different from those we initially anticipated. However, in this project we found that for success that contributions needed to extend into non- technical areas as well. In particular, the problems of project management, communications, goal setting and methods of defining progress were all issues which we had to address. What we learned about these aspects was certainly a valuable part of this large project experience not only for the students, but the faculty and staff as well.

The InfoPad project was declared successful and completed about 2 years ago and from that background we have moved into another large coordinated effort. After reviewing what was learned from InfoPad, we will discuss how this relates to our new initiative. I will then finish with some of the exciting new research opportunities that are opening up in the area of wireless communications and mobile computing made possible by the advent of "system on a chip" integration.

II. Past

The InfoPad project began to form around 1990 under DARPA funding which continued throughout the project. The first papers describing an overall system architecture with emphasis on the design of a "Portable Multimedia Terminal"[1,2]. The primary issue which was initially being addressed was to demonstrate a complete system solution for accessing and manipulating multimedia data (text/graphics, video and audio) from backbone network based information and compute servers as seen in Figure 1.

The user device, the InfoPad, was designed to be as light weight, low power and low cost as possible, by using network based compute and storage servers. In order for the user to *appear* to have large amounts of local storage and a high performance computing platform, it was necessary to optimize the wired and wireless communications links between the pad and the servers. This all needed to be done with the constraint of minimum power consumption in the portable units.

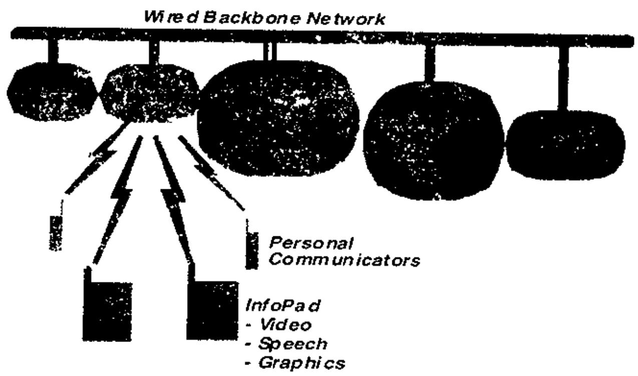

Fig. 1: The InfoPad system architecture

A number of important lessons and technologies were developed during the course of the InfoPad project research. It was hypothesized that a true system solution was necessary, because of the interrelationships of the various levels. The project would eventually include a tightly coordinated effort involving research in the user interface, middleware, backbone network, wireless link to the terminal, as well as the terminal's embedded software and low power hardware. Clearly, the problem of linking together so many research activities with a common focus resulted in tension due to the dependencies of one effort on another. While this clearly a challenge of such an approach, the benefits far outweighed the overhead in managing this process. A primary positive aspect was that students (and faculty) were exposed to the constraints, capabilities and research approaches of a number of disparate disciplines. This made it possible for individual researchers to develop a viewpoint of how their work fits into a larger picture, to develop more realistic individual constraints and for some to be able to contribute to new research areas.

Another aspect was the training that was provided in project management. It was found that because of the number of people involved in the project that common meetings of the entire effort were not useful. Splitting up into smaller efforts, however introduced bottlenecks in the communication, resulting in an intergroup dependency problem in which groups would be unaware or would misunderstand issues such as schedules, interface specifications, performance or functionality of something jointly being used or required. This was addressed by having sub-groups with rotating student leaders who would meet weekly to discuss these issues. These group leaders also had the job of allocating resources (e.g. new computers), defining the demonstrations and most importantly managing their peers on the development (*not their research - see below*) aspects of the project. Every 6 months we would rotate these leaders and the maturity gained by a student leader after they performed this function was impressive.

We found that the large project nature of InfoPad improved our interaction with industry. The goals of the project were clear, and because of the resources of the project, mechanisms for interaction were more easily set up and carried out. For example InfoPad had involvement from Intel, IBM, Hewlett Packard, Ericcson, Motorola, National and Texas Instruments and we were able to have an off-site retreat every 6 months for 2 to 3 days, which had on the order of 100 attendees. This had the dual use of bringing the entire project together to internally review the overall progress and plan for the future. The feedback and assistance that the industrial members were able to give, after they had sustained involvement in these retreats, became very beneficial.

In spite of all these very positive aspects, there was a dark cloud as well. After all, it is felt by many that the University research mission is to train individual students to perform research and not to be mired down in the overhead of a development project, no matter how advanced its goals. This challenge would show up in a number of ways. The demonstrations that would define our level of progress would consume enormous resources and a hard deadline for a demonstration would result in much work being performed at the last minute, that would either be poorly done or not useful to the long term goals. Another problem is that some students would be involved too much development work and not have sufficient time for their individual research efforts.

The solution to these problems involved in clearly defining two goals of the project: one was the support of the overall *development* goal which was to make and demonstrate a complete working system; and the second more important goal was to carry out and complete individual research projects (MS and Ph.D.). The student group leaders, staff and some faculty, would meet and plan the development aspects of the project (e.g. demonstrations, infrastructure hardware and software). It was then necessary to negotiate with the individual students and their faculty research advisor to be sure that the student had sufficient time for these activities. Clearly, most desirable was to align the development efforts with the individual research projects but this was not always possible, so that it was just necessary to be sure the burden was shared fairly. Said another way, no student received a Ph.D. degree for having merely worked on the InfoPad project, a clearly defined individual contribution had to be completed.

The question arises why have a large development project if individual research is the goal? Beyond the advantages described above, the fundamental reason is that the presence of the overall project provides test vehicles, infrastructure or realistic specifications for the individual research. In Figure 2 is shown a number of the Ph.D. research projects and how they cover the range of the InfoPad system.

To conclude the description of the overall project, the 5 groups that were defined are listed below along with a sampling of the references which were published about project research.

Portable Pad Design [3-9]- Design, fabrication and testing of the pad, its case and internal electronics and software.

Low Power Design Techniques and Tools [10-17] - CAD tools and methodology to support system level exploration and estimation of low power implementations and electrical-mechanical design

Multimedia Wireless Network Architecture [18-31]- Software to support InfoPad mobility (InfoNet) and support for Quality of Service on heterogeneous backbone networks (Medley)

Middleware Architecture and User Interface [32-36]- Provides interfaces through proxies between the conventional network and the InfoNet as well as "InfoPad aware" demonstration applications

Wireless link [37-41]- Implementation using commercial radios and research into customized monolithic CMOS implementations of TDMA and spread spectrum radios.

In order to keep the focus on real problems in system integration and to keep motivation high, it was decided to actually construct a number of terminals (10) and basestations and to attempt to use the resulting system design in our local environment. Though the reliability never was achieved to allow us to be actually dependent on the Pad, it is the platform we now use for demonstrating mobile network and radio link protocols and new radio transceivers.

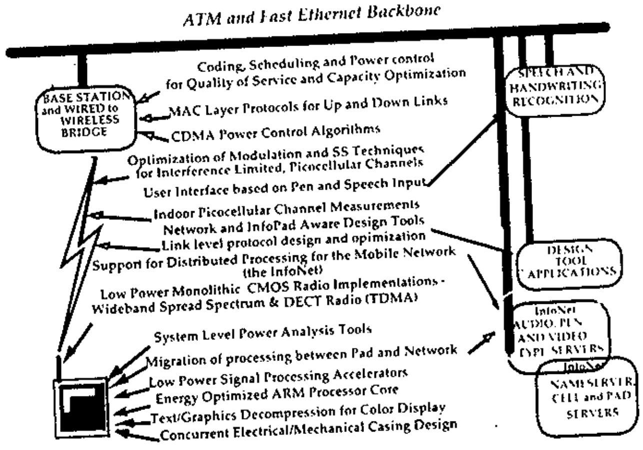

Figure 2: InfoPad research projects

III. Present

The InfoPad project was considered to be finished around the beginning of 1997 after about 6 years, and though we could have (and did) define various follow on projects, it was apparent that the original vision, though not realized in all respects, had been demonstrated and the next step was for industrial product development (something which is just now beginning to happen). It was also clear, however, that some portions of the InfoPad system implementation were further along than others. We had substantially met (or actually surpassed) what we thought was possible in low energy design, and there were spin-off projects in the area of user interface, middleware and mobile networking, so the area that seem to be the most in need of new research was the support for an inexpensive, high performance, wireless link.

The design of these links for this and other applications in state of the art CMOS technology is our new focus. We are driven by the realization that the improvements in CMOS technology has surpassed a threshold that now makes it possible to integrate the complete wireless system, including the RF, communication processing and protocol and interface functions on a single chip. This "system on a chip" design problem also is of broad importance, and the design of a single chip radio provides an excellent driver for developing this methodology. To insure generality of our results we are addressing a broad range of wireless link specifications ranging from the high end InfoPad requirements to ultra low power applications, as well as to address more fundamental issues in the general application of wireless connectivity.

In developing our new effort, we felt that again the advantages of a large integrated project was worth the investment and to properly design and deploy a wireless link, an integrated activity was needed from the connection into the wired infrastructure along the associated protocols to deal with mobility and the applications, through to the RF modeling of the CMOS technology which allows single chip integration of a complete radio. Though the focus is narrower than InfoPad, the goal of single chip integration of all these functions will require vastly more coordination of all the various components.

We also wanted to take some of the other lessons from InfoPad and apply them to our new endeavor. In particular, close involvement of industry is viewed critical because of the large investments being made in wireless technologies. We thus started the Berkeley Wireless Research Center (BWRC), with companies in each of the critical areas of our research focus, which include Ericcson, Lucent, Texas Instruments, SGS Thompson, Hewlett Packard and Cadence. This is designed to allow us to build on the industry base to do truly advanced research, instead of duplicating what was already being attempted. In order to substantially enhance our previous approach of sporadic meetings (retreats, occasional visits), we felt we needed personal contact. Though sending students to the companies and company involvement by sending junior employees to the University does help in this regard, it was felt that only by having senior (mentor level) researchers at the University from the companies that a new model of interaction could develop. This brought up the all important intellectual property issue, which we have addressed by the dual strategy of declaring all center research in the public domain and striving to be sure that our research projects are sufficiently advanced that they are pre- competitive.

III.A. Collaboration Support

As mentioned the communication problem between the disparate groups in the InfoPad project was a major issue, which we are now addressing by co-locating the students, staff and faculty. This coupled with the desire to bring in a substantial number of visitors and to support the collaboration that we feel is so necessary, required us to design a physical environment for that task. We wanted a space that was highly flexible, which could adapt to the changing needs of the project, as well as places for informal collaboration and meetings that are convenient to the work area. Finally we wanted to bring all our laboratory and testing areas together, from low level device testing up to full base stations for cellular experiments.

We did this by going off campus and leasing 11,000 square feet and constructing an interior which meets our needs for about 50 researchers. The centerpiece is a large "forum" area that is meant to foster collaboration, and around this, acoustically isolated through glass walls, is an open area with work cubicles which will have a "non-territorial" assignment policy. Each work area has a moveable file drawer which can either be at the workstation or in a "file barn" along one wall. A researcher can claim a space as long as they are actively at the center (which may be for years for the students and faculty), but if they are absent for a period and space is needed, the file drawer returns to the barn allowing reuse of the workstation space. There is a library to store books normally at a researchers desk and a large laboratory acoustically isolated off of the main area. The workstation environment and telecommunications have been designed so that all locations in the center are equivalent.

This interest in supporting the social and non-technical aspects of the research agenda, is clearly one of the lessons learned from InfoPad. The enhanced communication which we are striving to achieve in the Center, between our industrial mentors, involved faculty and the research students should provide a unique experience and hopefully highly efficient one for producing high quality research results.

IV. Future

Clearly the critical element in our present endeavor is the research agenda that we plan to pursue and to this end we have developed 3 research drivers. We believe that we are now on a threshold of wireless connectivity which will become the dominant method of interconnecting the myriad of future devices which will consume and produce data. These communications will range from the high end involving computers, high quality video, cameras and displays, to simple internet access devices and appliances for reading and listening; for communications between people and the control of utilities, to the lowest end in which very low bandwidth communications will be used to extract information from large numbers of distributed sensors. A fundamental belief we have is that many new applications will evolve when inexpensive, wireless communications become available and it is in these new areas that the most important new developments will occur.

IV.A. Universal Radio Systems

It is this latter belief, that the most important future uses of wireless communications are unknown, that drives our most aggressive future research agenda. The approach to developing new wireless systems has in the past been controlled by the standards setting process. This worked well in an environment when there was a single dominant wireless application, so it was clear what the optimizations should be in a new system design. However, now that application monopoly is due to be transformed into an anarchy of new uses, which will have an array of diverse requirements. The standards process therefore now has the enormously difficult task of standardizing the radio link for applications which possibly do not even exist yet.

Another transforming issue is the application of mainstream integrated circuit technology to radio implementation. In the past radio technology has been relatively static, with the designs involving hundreds of discrete components. This is about to change, since once single chip radio transceivers become a commercial reality, then radio design will ride the exponential curves of IC technology embodied in Moore's Law. The capability of the algorithms, protocols and interface circuitry which can go onto a single chip will then increase at an exponential rate, thus obsoleting the standards before they are even released. This technology obsolescence is already occurring in the 3rd generation UMTS standard in which the estimates of what can be done with implementation technology are already behind the capability of even present day circuitry.

The future is likely to continue to see an ever-increasing number of standards for technical, political and economic reasons. The strategy which has been pursued is to seek new frequency allocations to provide capabilities that were not previously supported (and thus making old standards obsolete). Detailed air interface specifications eventually result in gross inefficiencies since they are not adaptable to new requirements and implementation technologies. Therefore there is a need for a much better solution to the use of scarce frequency spectra.

Part of the answer may be found in observing how the unlicensed ISM bands are being used. These bands have made it possible for a variety of suppliers to try out new techniques and wireless link applications without the inertia of having to first work through the standards bodies. Particularly exciting are the Bluetooth and SWAP proposals, which will be able to exploit the latest in technology, demonstrate a new way to deploy important new wireless applications in home networking and cable replacement.

However, totally unregulated use or ad hoc requirements without technical basis, such as in the ISM bands, is also not the answer, because of the inter-system interference which would then be present. What is needed are a set of rules of etiquette or meta-level air interface requirements that allow multiple operators to co-exist without coordination, while being flexible enough to continuously upgrade to support new services as well as advances in communication engineering and implementation technologies; all the while still yielding the maximum possible channel capacity and robustness. In this way the increasingly scarce radio spectrum can be optimally utilized. It is not that every radio should necessarily support multiple standards, though this is desirable as well, but it is an approach to wireless connectivity which over time and in different implementations will be useful for widely different applications, with vastly different specifications and implemented in fundamentally different ways. If this approach is successful, there will be many different "Universal radio systems" operating in the same band with widely different characteristics.

We found in the InfoPad project that by having access to the entire range of the system design, that an understanding of trade-offs could be made, which would not have been possible otherwise. We feel that this will be a necessary component of our quest for the Universal radio system in that issues from all protocol levels must be re-evaluated along with the underlying communication engineering, and all of this in the context of a CMOS "system on a chip" implementation. The goal of this research focus is to redefine the way that frequency spectrum is allocated and used, so that maximum utilization of spectra will be achieved in spite of dramatic changes in use and implementation technology.

IV.B. Picoradios

It is believed in the future that while the Universal radio systems will provide maximum utilization, particularly for those applications that need it, there is another set of requirements which can benefit from the flexibility of wireless connectivity. This is the local, low bandwidth interconnection between small information sources and sinks, which would include wireless interconnection of home automation devices, such as security and monitoring, as well as control of temperature, lights and heat. In the office and lab, it may include the link to active badges, the control and data acquisition from test equipment and information from remote sensors and control of actuators. In the factory these might include sensors for failure detection or in robots or in automotive applications such as smart cars and freeways.

One characteristic of many of these applications is that they need to be extremely small and unobtrusive, on the size of a 10's to 100's of cubic millimeters, rather than 10's to 100's of cubic centimeters of typical cellular radios and we call these PicoRadios. This means that the limited amount of available energy will require ultra low power design, which if low enough may even make self-powering possible. The Picoradios will therefore need to operate at power levels of 10's to 100's of microWatts, which will be multiple orders of magnitude lower than the Universal radio designs. Fortunately, the bandwidth requirements of the Picoradio applications will also be 3-4 orders of magnitude below the high end requirements, and with implementations which can be optimized for a single task.

Meeting all of these requirements will require the most highly optimized, dedicated single chip solutions. This will require a design approach and architecture which is quite different from the highly flexible and adaptive Universal radio. Since cost is a critical component of these devices it will be necessary to have a fully-supported design methodology which will allow rapid implementation of these single chip radio systems. Again, however, a complete system approach will be needed since the networks will need to be self configuring, as well as support multihop and peer-to-peer capability and have sophisticated energy conserving protocols.

IV.C. Ultra-High Bandwidth Millimeter Radios

The most critical resource in wireless communications is the availability of spectra and though new spectra can not be fabricated there are large amounts of essentially unused spectra at higher frequencies. Though the frequency spectra in the region up to 5 GHz is essentially allocated, at frequencies above 20 GHz there are large amounts of unlicensed spectra with an unlicensed band of over a GHz available around 60GHz. We believe that the exponential improvements in CMOS technology means that the frequency bands in the 60 GHz range should be accessible by CMOS solutions in about 10 years and more importantly we can do test circuits to investigate this conjecture now. Critical issues are the determination of the architectures and system implementations that could exploit large bandwidths as well as modeling of CMOS in these high frequency regimes to facilitate predictions of how these devices will operate. The goal of this driver application is to explore the potential of future technologies in the wireless arena as well as the time scale when new capabilities will become available.

V. Conclusions

Large project University research has both advantages and disadvantages over the more typical individual investigator approach. These advantages include the developing and supporting of infrastructure, enhanced collaboration, improved ability to obtain support, and finally to providing an experience to students which will more closely approach their future work environment. The non-technical issues which arise from such large projects are important and must be addressed if the advantages are not to be lost in a mire of purely development activity. Particularly important is to recognize the two clearly distinct functions of a large project; one of which is to implement the development demonstrations and the other to provide resources and time for individual research accomplishment. Well planned project management is needed to resolve the tensions between these two goals.

The times are exciting in that a major technology break has occurred in which CMOS is now or will soon be the technology of choice to implement wireless links. The exponentially improving performance of CMOS at dramatically lower costs, provides an opportunity for providing fundamentally new ways of implementing wireless connectivity. The possibility exists for developing new paradigms for frequency allocation, the implementation of ultra low power, cost and size radios and the inexpensive exploitation of un- utilized frequency bands.

VI. Acknowledgments

I would like to acknowledge all the faculty, staff and students and the industry researchers who were involved in the InfoPad project. I also would like to thank the support by DARPA and the InfoPad member companies including Intel, Texas Instruments, IBM, Hewlett Packard, National and Motorola.

References

InfoPad Architecture

[1] Brodersen, R.W.; Chandrakasan, A.; Sheng, S. "Technologies for personal communications", 1991 Symposium on VLSI Circuits. Digest of Technical Papers. pp. 5-9.

[2] Sheng, S.; Chandrakasan, A.; Brodersen, R.W. "A portable multimedia terminal", IEEE Communications Magazine, vol.30, (no.12), Dec. 1992. pp. 64-75.

[3] Chandrakasan, A.P.; Burstein, A.; Brodersen, R.W., "A low-power chipset for a portable multimedia I/O terminal", IEEE Journal of Solid-State Circuits, vol.29, (no.12), Dec. 1994. p.1415-28

[4] Brewer, E.; Burd, T.; Burghardt, F.; Burstein, A.; Doering, R.; Lutz, K.; Narayansaramy, S.; Pering, T.; Richards, B.; Truman, T.; Katz, R.; Rabaey, J.; Brodersen, R., "Design of wireless portable systems", Digest of Papers", COMPCON '95, Digest of Papers, IEEE Comput. Soc. Press, 1995, p.169-76.

[5] Brodersen, R.W., "The network computer and its future", 1997 IEEE International Solid-State Circuits Conference. Digest of Technical Papers, Vol.40, 6-8 Feb. 1997, p.32- 6

[7] Truman, T.E.; Pering, T.; Doering, R.; Brodersen, R.W. "The InfoPad multimedia terminal: a portable device for wireless information access", IEEE Transactions on Computers, vol.47, (no.10), IEEE, Oct. 1998. p.1073-87.

[8] Wang, F.-C., Wright, P.K., and Richards, B., "A Multi-disciplinary Concurrent Design Environment for Consumer Electronic Product Design",Journal of Concurrent Engineering: Research and Application, Vol. 4, No.4, December 1996.

[9] Wang, F.-C., and Wright, P.K., "Concurrent Product Design and Manufacture: A Case Study on Infopad, a Wireless Network Computer", 1998 ASME International Mechanical Engineer Congress and Exposition, Anaheim, CA, November 15-20, 1998.

Low Power Design Techniques and Tools

[10] Chandrakasan, A.P.; Sheng, S.; Brodersen, R.W. "Low-power CMOS digital design", IEEE Journal of Solid- State Circuits, vol.27, (no.4), April 1992. p.473-84.

[11] Brodersen, R.W.; Chandrakasan, A.; Sheng, S. "Design techniques for portable systems" 1993 IEEE International Solid-State Circuits Conference. Digest of Technical Papers. p.168-9

[12] Chandrakasan, A.P.; Brodersen, R.W. Minimizing power consumption in digital CMOS circuits. Proceedings of the IEEE, vol.83, (no.4), April 1995. p.498-523

[13] Bentz, O.; Lidsky, D.B.; Rabaey, J.M. "Information based design environment", VLSI Signal Processing, VIII Sakai, Japan, 16-18 Sept. 1995, IEEE, 1995. p.237-46.

[14] Chandrakasan, A.P.; Brodersen, R.W., <u>Low Power Digital CMOS Design</u>, Kluwer Academic Publishers, 1995

[15] Landman, P.; Mehra, R.; Rabaey, J.M. "An integrated CAD environment for low-power design", IEEE Design & Test of Computers, vol.13, (no.2), IEEE, Summer 1996. p.72-82

[16] Landman, P.; Mehra, R.; Rabaey, J.M. "An integrated CAD environment for low-power design", IEEE Design & Test of Computers, vol.13, (no.2), Summer 1996. p.72-82

[17] Burd, T.D.; Brodersen, R.W. "Processor design for portable systems", Journal of VLSI Signal Processing, vol.13, (no.2-3), Kluwer Academic Publishers, Aug.-Sept. 1996. p.203-21

Multimedia wireless network architecture

[18] K. Keeton, B. A. Mah, S. Seshan, R. H. Katz and D. Ferrari, "Providing Connection Oriented Services to Mobile Hosts, 1993 USENIX Conf. on Mobile and Location Independent Computing, August, 1993

[19] Haskell, P.; Messerschmitt, D.G.; Nichols, K.; Pol, E.-J.; Morley, R.; Jardetzky, P. "Multimedia consumer applications on a heterogeneous network", Proceedings of the 1st International Workshop on Community Networking, San Francisco, CA, 13-14 July 1994, p.229-37

[20] Yun, L.C.; Messerschmitt, D.G. "Power control for variable QOS on a CDMA channel", Proceedings of MILCOM '94, Fort Monmouth, NJ, 2-5 Oct. 1994, IEEE, 1994, p.178-82

[21] S. Seshan, M.T. Le, F. Burghardt, J. M. Rabaey "Software Architecture of the InfoPad System", Mobidata Workshop, Nov. 1994

[22] H. Balakrishnan, S. Seshan, E. Amir, R. H. Katz, "Improving TCP/IP Performance over Wireless Networks," ACM Conference on Mobile Computing and Networks, Oakland, CA, (November 1995)..

[23] E. Amir, S. McCanne, R. H. Katz, "Receiver Driven Bandwidth Allocation for Light Weight Session," ACM Multimedia '97 Conference, Seattle, WA (November 1997).

[24] Yun, L.C.; Messerschmitt, D.G. "Variable quality of service in CDMA systems by statistical power control", IEEE International Conference on Communications , Seattle, WA, 18-22 June 1995, p.713-19

[25] Le, M.T.; Rabaey, J. "A global QoS management for wireless network", Proceedings of 1996 World Conference on Mobile Communications, Canberra, Australia, 2-6 Sept. 1996., p.205-18.

[26] T. Hodes, R. H. Katz, E. Servan-Schreiber, L. A. Rowe, "Composable Ad-Hoc Mobile Services for Universal Interaction," Third ACM Mobicom Conference, Budapest, Hungary, (September 1997).

[27] Lu, Y.; Brodersen, R. "Unified power control, error correction coding and scheduling for a CDMA downlink system", Wireless Networks, vol.3, (no.1), Baltzer, 1997, p.83- 90.

[28] Han, R.Y.; Messerschmitt, D.G. "Asymptotically reliable transport of multimedia/graphics over wireless channels", Proceedings of the SPIE , vol.2667, 29-31 Jan. 1996, p.99-110

[29] Truman, T.E.; Brodersen, R.W. "A measurement-based characterization of the time variation of an indoor wireless channel", 1997 IEEE 6th International Conference on Universal Personal Communications, Conference Record , San Diego, CA, USA, 12-16 Oct. 1997, p.25-32

[30] P.Haskell, D.Messershmitt, L.Yun, "Architecture of Integrated-Service Wireless Networks", chapter in Wireless Communications: Signal Processing Perspectives, V. Poor, G. Wornell Ed., Prentice Hall, 1998.

[31] L.C. Yun and D.G. Messerschmitt, ``Digital Video in a Fading Interference Wireless Environment'', IEEE Int. Conf. on Acoustics, Speech, and Signal Processing, Atlanta, GA., May 1996.

Middleware Software Architecture and User Interface

[32] Burstein, A.; Stolzle, A.; Brodersen, R.W. "Using speech recognition in a personal communications system",IEEE SUPERCOMM/ICC '92, 1992. p.1717-21

[33] Narayanaswamy, S.; Seshan, S.; Amir, E.; Brewer, E.; Brodersen, R.W.; Burghardt, F.; Burstein, A.; Yuan-Chi Chang; Fox, A.; Gilbert, J.M.; Han, R.; Katz, R.H.; Long, A.C.; Messerschmitt, D.G.; Rabaey, J.M. "Application and Network Support for InfoPad,". IEEE Personal Communications, vol.3, (no.2), April 1996, p.4-17

[34] Le, M.T.; Burghardt, F.; Seshan, S.; Rabaey, J., "InfoNet: The networking infrastructure of InfoPad", Digest of Papers. COMPCON '95, 5-9 March 1995, p.163-8

[35] Burstein, A.; Long, A.C., Jr.; Narayanaswamy, S.;

Han, R.; Brodersen, R.W. "The InfoPad user interface" Digest of Papers. COMPCON '95, 5-9 March 1995, p.159-62

[36] E. Amir, S. McCanne, R. H. Katz, "The Media Gateway Architecture: A Prototype for Active Services," ACM SIGCOMM'98 Conference, Vancouver, Canada, (September 1998).

CMOS Radio Implementations

[37] Gray, P.R.; Meyer, R.G. "Future directions in silicon ICs for RF personal communications", Proceedings of the IEEE 1995 Custom Integrated Circuits Conference 1-4 May 1995, p.83-90

[38] Sheng, S.; Lynn, L.; Peroulas, J.; Stone, K.; O'Donnell, I.; Brodersen, R. , "A low-power CMOS chipset for spread spectrum communications", IEEE International Solid- State Circuits Conference, Digest of Technical Papers, ISSCC, San Francisco, CA, 8-10 Feb. 1996, p346-347.

[39] Rudell, J.C.; Ou, J.-J.; Cho, T.B.; Chien, G.; Brianti, F.; Weldon, J.A.; Gray, P.R. "A 1.9-GHz wide-band IF double conversion CMOS receiver for cordless telephone applications", IEEE Journal of Solid-State Circuits, vol.32, (no.12), IEEE, Dec. 1997. p.2071-88

[40] Onodera, K.K.; Gray, P.R. "A 75-mW 128-MHz DS-CDMA baseband demodulator for high-speed wireless applications" IEEE Journal of Solid-State Circuits, vol.33, (no.5), IEEE, May 1998, p.753-61

[41] Sheng, S.; Brodersen, R.W., Low Power CMOS Wireless Communications, Kluwer Academic Publishers, 1998

Biography

Robert W. Brodersen received his PhD from MIT in 1972 and then was with the Central Research Laboratory at Texas Instruments for 3 years. Following that, he joined the Electrical Engineering and Computer Science faculty of the University of California at Berkeley, where he is now a professor. In addition to teaching, he is involved in research involving new applications of integrated circuits, which is focused in the areas of low power design and wireless communications; and the CAD tools necessary to support these activities.

He has won best paper awards for a number of journal and conference papers in the areas of integrated circuit design, CAD and communications, including in 1979 the W.G. Baker award for the best IEEE publication in all areas. In 1982 he became a Fellow of the IEEE and in 1983, he was co-recipient of the IEEE Morris Liebmann award. In 1986 he received the Technical Achievement awards in the IEEE Circuits and Systems Society and in 1991 from the Signal Processing Society.

In 1988 he was elected to be member of the National Academy of Engineering. In 1994, he was made the first holder of the John Whinnery Chair in Electrical Engineering and Computer Science. In 1996 he was the winner of the "Contributions to the IEEE Solid State Circuits award for "Contributions to the design of integrated circuits for signal processing" and most recently he was the recipient of the 1998 ACM SIGMOBILE Award for Outstanding Contributions to Research on Mobility of Systems, Users, Data and Computing.