System Design Issues for Low-Power, Low-Cost Short Range Wireless Networking

Pravin Bhagwat, Chatshick Bisdikian, Ibrahim Korpeoglu, Arvind Krishna, Mahmoud Naghshineh
IBM Thomas J. Watson Research Center,
New York, USA
e-mail: {pravinb,bisdik,ibrahimk,arvind,mahmoud}@us.ibm.com

ABSTRACT

The emergence of battery powered handheld devices as popular computing devices is presenting new challenges. Among the most important challenges is the need to provide a low cost, low power, indoor wireless networking access to handheld devices. The constraints posed by battery power and cost require a careful re-evaluation of system design issues at all layer of the protocol stack. In this paper, we present the design of a short range wireless networking system called BlueSky which is being developed at IBM Research to address these challenges. We show that the optimization objectives for short range indoor wireless systems are quite different from those for traditional cellular wireless systems. We also argue that in the next millennium the primary optimization criteria for the design of (short range) wireless systems will shift from the traditional spectral efficiency towards battery lifetime and cost.

I. INTRODUCTION

Handheld computing devices of today are as powerful as the desktop PCs of five years ago, and their cost is as little as one tenth that of traditional Personal Computers. This has been made possible by advances in CPU, memory and integration technologies. Unfortunately, battery technology has not kept pace with this development. Energy density of batteries has only doubled in the last 35 years whereas CPU speed keeps doubling every 18 months [1]. The slow improvement in battery lifetime shows that energy consumption will be one of the most important factors in designing systems and wireless networking support for portable devices in the future. Without motivating the need for wireless networking, it suffices to say that access to information is a primary driver for handheld computing and online access is a primary component of information access [2].

The major focus of past research on power control has been maximizing spectral efficiency, i.e., maximizing throughput or system capacity for cellular telephone systems. However, for battery powered PDAs connected via licence free wireless links, the optimization objective is shifting from link efficiency to battery efficiency. Cost and battery lifetime constraints are presenting new challeges for the design of wireless network architectures, protocols and applications.

In this paper we identify some of the important issues and principles in designing *short range* wireless networking support for handheld, portable devices with the aim of extending the battery lifetime and reducing the cost of the communication attachment and network support infrastructure. We focus only on providing *indoor* wireless connectivity for environments such as office buildings, airports, hotels, homes, etc.

This paper is organized as follows. In section II we discuss important issues that need to be considered in designing low-power, low-cost wireless networks and identify principles that need to be followed. In section III we present our initial findings from the system BlueSky which we are building to address these issues. Finally, in section IV we present our conclusions.

II. OBSERVATIONS, ISSUES, and DESIGN PRINCIPLES

Traditionally, power control in cellular telephone networking systems has focused on balancing received power levels, reducing the interference from other channels, or balancing the SNR ratios for the purpose of increasing system capacity [3], [4]. Besides these capacity-centric approaches, there have been also some user-centric approaches which try to optimize user-related objectives such as bit error rate (BER), throughput and delay [5], [6]. Research in the area of multi user detection, adaptive equalization, coding and modulation does not consider increasing the battery lifetime and reducing the cost of network components as the primary objective. Most designs trade complexity for efficiency resulting in handsets which are costly and power-hungry, and networks which require complex switching and control equipments [7].

With the proliferation of low-cost, low-power handheld devices, it is important to design wireless networks with the aim of increasing the handsets' battery life and reducing the cost of user communication equipment. The issues span many dimensions: from initial design and planning of the overall network architecture to the selection and design of individual network components, and to the design of networking algorithms and protocols.

A. *Planning the Network*

The first issue in planning an indoor wireless network from power and cost perspective is the selection of a *network topology* and architecture. Although cellular topology is the most commonly used network topology for wide area wireless networks, there are other alternatives such as ad-hoc topologies. From power consumption point of view, however, cellular topology is more favorable (see [8] for a discussion of these issues). In a cellular system, base stations are deployed to cover a well-defined fixed geographical area (a building, a city, a region) and they serve as the access points for mobile devices to the wired network infrastructure. They are just one-hop away from the mobile devices. In an ad-hoc system, there may or may not be base stations that provide access to the wired network infrastructure. Even if they exist, they may be multiple hops away from the mobile devices. In access point based approach, complex functionality can be shifted to the access points which are not power limited. Design of asymmetric network protocols [9] which put more processing and intelligence to the access points rather than mobile devices are possible.

The second issue is the *range of communication*. Due to power-distance relationship, short range radio link is more favorable from battery lifetime point of view, although it requires more access points to be installed to cover a fixed area. This is due to the fact that the received power (P_r) is inversely proportional to the distance (d) from the transmitter: $P_r \propto P_t d^{-\beta}$, where $2 \leq \beta \leq 4$ [10]. The transmitter power of a mobile device should be designed for the maximum possible distance to the corresponding access point. Even with adaptive power control, as long as the mobile devices are somewhat evenly distributed over the area covered by the access point, the lifetime of the batteries is inversely proportional to the range covered by the base station (more precisely, the battery lifetime is proportional to $d_{AP}^{-\beta}$ where d_{AP} is the average range of an access point). The range of the wireless link affects the number of access points that needs to be deployed to cover a fixed area. Simple physical area arguments under a fixed transmitted power assumption imply that the number of access points needed is inversely proportional to the second degree with respect to the range of the wireless link: $N_A \propto d_{AP}^{-2}$. Thus, the implication is that the battery lifetime is proportional to $N_A^{0.5\beta}$. For example, for a typical office building a range of 30 m usually implies that $\beta = 4$ at carrier frequencies of 2.4 GHz and this implies that if we double the number of access points, the average battery life time will be extended four times which is a considerable usability enhancement. The reader should realize that the cost paid for this extended lifetime is in infrastructure cost, namely the access points, and it is hence important to keep their cost low. Making the range shorter than a minimum amount also does not make sense given the vagaries of radio range, e.g., the range should not be less than an average sized room since it might be hard to constrain the radio range to be smaller. Keeping the range too short also increases the mobile handoff rate which may be unacceptable beyond a threshold.

System considerations for supporting data communication also favor the use of short range link. Information access implies that the average device uses the network in a bursty mode, as opposed to voice communication which makes relatively constant rate use of the network. Further, while using the network, the information access rate could be between 100 Kbps - 1 Mbps. The peak aggregate rate requirements are also mitigated with small cells since they have a smaller number of actively transmitting/receiving users. Ideally, we would design cells in which we expect just one or two active users, on average. Smaller peak aggregate rates imply that the system could afford to use lower power requirements and cheaper components due to simpler designs.

The third issue is the selection of the wireless *link capacity*. Handheld devices have small screens (usually 2 bit grayscale), very little memory (1-2 MB) and one tenth the power of a PC (16-33 Mhz). Therefore, the network bandwidth demand of the applications that will run on these devices are expected to be quite low. In fact, most handheld devices are connected only through a serial port with a speed of 19200 bps or at most 57600 bps. Since the range of the wireless link is very small, not many handheld devices are expected to be in the same cell at the same time. Hence the total bandwidth need of these devices in a cell is not very high. Therefore providing a high-bandwidth (more than 1 Mb) link does not make much sense in this environment. High bandwidth requires complex radio transceiver and high-speed DSP processing which increases the power consumption and cost.

B. Choosing and Designing the Components

There is a trade-off between cost, power, quality and utilization in selecting and designing network components. In providing low-cost, low-power short range wireless networks, lower cost components which do not utilize the network resources very well may be preferred to expensive ones. Simple designs with suboptimal performance which consume less power may be favorable to complex designs.

Very sophisticated radio transceivers with advanced modulations and DSP circuits that utilize the channel bandwidth very efficiently and provide good channel quality have been designed and are in the market. But these are far from satisfying the requirements of being low-power and low-cost. Providing efficient and high-quality radio link requires complex DSP processing which consumes high power. Building ASICs is costly unless they are produced in large quantities. For short range links the number of users per cell is small and the bandwidth is free. Therefore capacity and bandwidth utilization are secondary issues compared to cost and power. Using off-the-shelf radio components are preferable over designing custom components, resulting in higher cost savings at marginal bandwidth wastage. Reduction of complexity also translates into power saving. For example, Smith [11] shows the implications of selecting the radio device on power consumption: an adaptive direct sequence spread spectrum radio which consumes 8 W of power is replaced with a simpler frequency hopping radio which only consumes 0.8 W of power without sacrificing much of the functionality. The second radio device consumes 10 times less energy.

A networking adapter can be connected to the host device in several ways; via serial, parallel, or USB port, PCMCIA, PCI, or AT interface. The choice of interface has implications on power consumption and cost. For example, attaching a device to the serial port may not be the best approach in terms of performance, but it may be the least expensive approach in terms of cost.

Continuous coverage using short range radio links is only feasible if the cost of access points is very low. This is only possible by reducing the hardware and software complexity of access point function and pushing the intelligence in the backend systems and servers.

C. Designing the Network Protocols and Algorithms

Power-aware, intelligent network protocols, and software and algorithms are the essential parts of the overall design of low-power, low-cost wireless networks. All layers of networking stack can play a role in reducing the energy consumption of mobile devices. At the physical layer, new low power RF circuitry and DSP designs and coding and modulation techniques are useful [12], [13]. At the data-link level, MAC layer controls the access to the medium and decides which stations transmit and when they transmit. It can employ intelligent medium access control and packet scheduling techniques to put mobile devices into sleep modes and conserve energy without sacrificing the network performance. At the logical data-link layer, FEC and ARQ schemes for error correction should be revisited from the power efficiency point of view. The channel error characteristics are important factors for designing low power link layer protocols, since there is no point in transmitting packets and wasting energy when the link is in the error state [14], [15]. Compression of data payloads and headers [16], [17] are crucial in coping with very low bandwidth links and decreasing transmitter power consumption by reducing the packet sizes. At the network layer, routing algorithms may use the total transmission power necessary to route packets from source to destination in selecting the routes in ad-hoc networks [18]. At the transport and application layers, the traffic models of the streams are important in designing when and how long to put the respective network cards into sleep mode [19]. Stemm [20] points out that network interface cards consume as much power as the PDAs, hence strategies that will put them into sleep mode when there is no network activity may be very effective in reducing the power consumption.

It is important to revisit the design and implementation of compression and encryption algorithms from power consumption point of view. Encryption protocols such as public-key encryption are asymmetric protocols where the computation required in two communicating peers is not the same. The design of the encryption algorithms should be so that minimum computation is done at the mobile peer. The power consumed by encryption and compression protocols are due to CPU operations and memory accesses. Each operation in CPU consumes different amount of power (multiplication consumes more energy than addition) and memory reads and writes are also not equal in terms of energy consumption [13]. Implementations of these algorithms should take into account power consumption besides time and space efficiency.

Often, the interaction between several components needs to be carefully studied to design optimal power saving strategies. For example, a compression algorithm should weigh the cost of CPU power consumption against the cost of transmitting an uncompressed packet. Depending on the relative power consumption of CPU and radio, compression may or may not always be beneficial. This trade-off is new since CPU power consumption is roughly the same as the radio power consumption. Similar argument is valid for any transformation technique that is not mandatory for operational correctness but is useful to increase the system performance, such as FEC coding.

The amount of power consumed per unit job in each of the system components may be different in each system. For wide area networks, the transmitter power is in the order of watts [1] (ARDIS requires 4 W) and may be the primary contributor for power consumed per packet transmission in which case compression may be very effective in reducing the total power consumption. For short range links, the required transmission power may be low (cordless telephone handsets consumes 10 mW) so that CPU and memory may be the primary contributors of power consumption in which case compression is not desirable from energy efficiency point of view.

Strategies should also be adaptive to the changing network conditions. Adaptive power control techniques where a device emits just enough power to reach the receiver are useful for mobile networks where the distance between transmitters and receivers is not constant. We also need to consider the wireless channel conditions when applying transformation techniques such as compression and FEC coding on the network data. For example, under light traffic loads, compression may be wasteful if CPU is the primary consumer of power. On the other hand, compression may be necessary when available bandwidth per user drops under high load conditions. Similarly, FEC should adapted as a function of changing link conditions. When the channel is in good state, there is no need to do FEC and waste both computation and transmission power since FEC adds redundant bits.

III. BLUESKY SYSTEM

We took the above discussed issues into consideration while designing a low-cost, low-power wireless network called BlueSky [21]. The system consists of wireless *BlueSky attachments* that are connected to the mobile PDAs and provide short range wireless access to the wired network backbone through the *BlueSky access points*. *BlueSky backend network servers* located deeper in the network are used to offload the protocol processing and other services as much as possible from the mobile devices and access points.

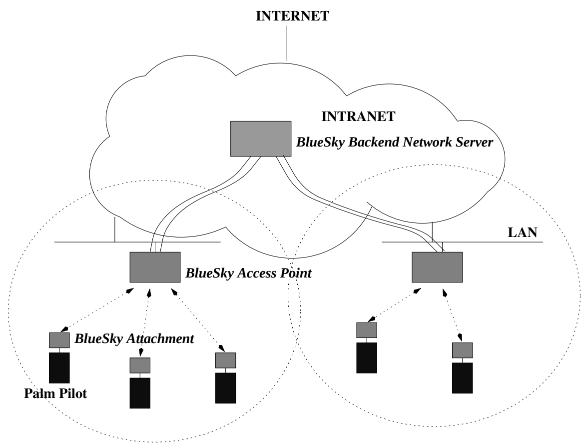

Fig. 1. BlueSky System Architecture

A. Architecture and Components

Figure 1 shows the system architecture. We designed BlueSky as an external serial port attachment so that it can be plugged into any PDA, laptop, or serial port enabled device. A BlueSky MAC protocol, that is running on the attachment, packetizes the byte stream received from the handheld device and sends it to the access point over the wireless link.

Since the cost of the radio is an important factor in the overall cost of the BlueSky attachment and access point, we chose a cost effective radio which runs in ISM 915 Mhz band and uses simple FSK modulation similar to what is used in cordless phones. The radio and access point can select one of the ten available frequency channels. The link bandwidth is 150 Kbps and the the effective range is between 20-30 m depending on the interference, multipath effects, etc. The bandwidth can be shared among 20 mobile devices.

Cost has been a very important factor in the design of our access point. Besides offloading some of the complex functionality from mobile devices to the access points, such as coordination and scheduling components of the MAC protocol, we also pushed some of the network layer and data link layer functionality to the backend servers. BlueSky access points are not network layer routers like in [22] but are *bridges* that forward the link layer traffic between mobile devices and the backend servers.

B. MAC Protocol

A power-conscious, polling based, asymmetric MAC protocol is designed for the shared wireless link. The portion of the MAC protocol that is running on the access point, called AP-MAC, is more intelligent and coordinates the access to the medium. It includes a scheduler component which schedules MAC-level events such as transport-payload, poll-and-receive-payload, invite-stations, etc. The scheduler and MAC layer are separate components which interface to each other with well-defined functions. Various scheduling algorithms can be implemented based on round-robin, fair-queueing, power-aware schemes. The corresponding MAC protocol that is running on a mobile device, called RS-MAC (remote station MAC), responds just to the commands sent from the AP-MAC and is much simpler, requires much less memory and processing power, and has smaller code size than the AP-MAC.

MAC protocol supports sleep modes. The attachments can go into sleep mode after informing the access point or the access point can initiate the sleep requests to the attachments. While an attachment is in sleep mode, the access point buffers the packets destined to it. The attachment also buffers the packets destined to the access point. When the attachment wakes up, it can learn from the access point about the queued packets and can indicate to the access point the amount of data it has buffered so that the access point can re-schedule the events.

Choice of a power-conscious MAC protocol is an important design decision. The alternatives for MAC layer protocols [23] include a CSMA-like or a centralized polling-based MAC protocol. Centralized, polling-based schemes are more easily adapted for low energy consumption. In a polling-based system, the possibility of collisions is much smaller than a CSMA system. Collisions cause waste of energy. Additionally, in a polling-based system, the central station knows the state of the queues in the downlink buffers and can use advanced power-conscious scheduling techniques to reduce the energy consumption in mobile stations. For example, sending the packets in bursts and putting the mobile device into sleep mode between burst periods is more energy efficient than sending the packets continuously and smoothly. But, batching the packets and sending them together at once increases the average packet delay and may require more buffers. Therefore, there is a tradeoff between battery lifetime, cost, and delay objectives. The packet size of the MAC layer and whether it uses variable or fixed-size packets may also have implications on power consumption. Allowing variable-length packets is a desirable property from power consumption point of view, since a station does not have to fragment a big, variable-size packet into smaller fixed-size packets and wait until all pieces are transmitted. It may instead transmit a whole variable-size packet at once and save more energy. However, using variable-length packets complicates the scheduler component.

The choice of the ARQ protocol for reliable transmission over the wireless link is affected by the cost/performance tradeoff. We chose a simple stop-and-wait ARQ protocol. Although schemes such as selective-repeat ARQ achieves higher throughput and link utilization, they need more buffers in the mobile devices, which means higher cost.

C. Other Logical Link Layer Issues

Besides sequencing and reliability, other logical link layer issues involve transformation on the data packets, such as compressing the payloads, compressing the headers, encrypting the data, coding the data for error resilience, all of which are very important for wireless links and affects the power consumption in the wireless network adaptors and in the mobile end-systems. It may not be always desirable to do these transformations from the energy efficiency point of view. There may be tradeoffs between CPU power consumption, memory power consumption, and network card power consumption while transmitting and receiving the data.

We carried out some measurements on a 3COM Palm Pilot to validate our reasoning. The power consumption of the Pilot varies between 50 - 150 mW depending on the intensity of task performed (device on but not doing anything, CPU in computation loop, the serial port being exercised). The Bluetooth Special Interest Group [24] is defining radios that appear to draw current between 0.3 mA and 30 mA which corresponds to the power consumption in the same orders as the power consumption of a Palm Pilot. With these numbers, there is indeed a great deal of tradeoffs which can be achieved, unlike current systems where either the wireless link or the computing device dominate the power consumption.

D. Light-weight Mobility Support

Short range wireless links imply frequent handoffs. With each handoff, mobility support processing occurs at the mobile devices and access points, and mobility related messages are exchanged over the wireless link. Redundant mobility related protocol processing and message exchange cause extra energy to be spent. Therefore, it is highly desirable to use a very light-weight mobility support which incurs least amount of processing in the mobile devices and causes very little extra message overhead to the wireless link. To achieve this, our approach is to support mobility at the lowest layer possible in the networking stack. All layers above it will be transparent to the mobility. The lower the layer that mobility is supported, the less the number of layers are involved in mobility processing and message exchanges. We provided mobility support at the data-link layer, similar in concept to MDL [25], as opposed to the traditional approaches which support mobility at the network layer, such as Mobile IP [26].

IV. CONCLUSION

Battery lifetime will dominate system and network design issues in the next millennium. Battery efficiency is increasing at a very slow speed compared to the improvements in silicon technology. With the proliferation of battery-powered handheld devices, low-power and low-cost system and network design has become a more important issue. In this paper, we identified important systems and network design issues and gave some design principles that will help the wireless network designers in designing the next generation wireless systems for portable devices. We gave examples of how we addressed some of the issues in our low-cost, low-power wireless networking platform, called BlueSky.

REFERENCES

[1] J. R. Lorch and A . J. Smith, "Software Strategies for Portable Computer Energy Management," *IEEE Personal Communications*, June 1998.

[2] A. Joshi and A. Krishna (Guest Editors), "Mobile Access to Web Resources," *IEEE Personal Communications*, October 1998.

[3] S. V. Hanly, "An Algorithm for Combined Cell-Site Selection and Power Control to Maximize Cellular Spread Spectrum Capacity," *IEEE Journal on Selected Areas in Communication*, September 1995.

[4] Roy D. Yates, "A Framework for Uplink Power Control in Cellular Radio Systems," *IEEE Journal on Selected Areas in Communication*, September 1995.

[5] J. Wu and R. Kohno, "Performance Evaluation of Wireless Multimedia CDMA Networks Using Adaptive Transmission Control," *IEEE Journal on Selected Areas in Communications*, December 1996.

[6] Y. Lu and R. Broderson, "Unified Power Control, Error Correction Coding and Scheduling for a CDMA Downlink System," in *INFOCOM Proceedings*, 1996.

[7] D. C. Cox, "Wireless Personal Communications: What Is It?," in *IEEE Personal Communications*, April 1995.

[8] H. Ahmadi, A. Krishna, and R.O. LaMaire, "Design Issues in Wireless LANs," *Journal of High-Speed Networks (JHSN)*, vol. 5, no. 1, pp. 87–104, 1996.

[9] S. Paul, E. Ayanoglu, T. F. La Porta, K-W. H. Chen, K. K. Sabnani, and R. D. Gitlin, "An Asymmetric Protocol for Digital Cellular Communications," in *INFOCOM Proceedings*, 1995, pp. 1053–1062.

[10] K. Pahlavan and A. H. Levesque, *Wireless Information Networks*, John Wiley & Sons, Inc., 1995.

[11] W. Mangione-Smith, P. S. Ghang, S. Nazareth, P. Lattieri, W. Boring, and R. Jain, "A Low-Power Architecture for Wireless Multimedia Systems: Lessons Learned from Building a Power Hog," in *International Symposium on Low Power Electronics and Design Digest of Technical Papers*, August 1996.

[12] L. Larson, "Radio Frequency Integrated Circuit Technology for Low-Power Wireless Communications," *IEEE Personal Communications*, June 1998.

[13] T. H. Meng, A. C. Hung, E. K. Tsern, and B. M. Gordon, "Low-Power Signal Processing System Design for Wireless Applications," *IEEE Personal Communications*, June 1998.

[14] P. Bhagwat, P. Bhattacharya, A. Krishna, and S. Tripathi, "Enhancing Throughput Over Wireless LANs Using Channel State Dependent Packet Scheduling," *ACM/Baltzer Wireless Networks*, vol. 3, no. 1, 1997.

[15] M. Zorzi and R. R. Rao, "Error Control and Energy Consumption in Communications for Nomadic Computing," *IEEE Transactions on Computers*, March 1997.

[16] V. Jacobson, "Compressing TCP/IP Headers for Low-Speed Serial Links," February 1990, Internet RFC 1144.

[17] M. Degermark, M. Engan, B. Nordgren, and S. Pink, "Low-loss TCP/IP Header Compression for Wireless Networks," in *MOBICOM Conference Proceedings*, 1996.

[18] S. Sing, M. Woo, and C. S. Raghavendra, "Power-Aware Routing in Mobile Ad-Hoc Networks," in *MOBICOM Conference Proceedings*, 1998.

[19] R. Kravets and P. Krishnan, "Power Management Techniques for Mobile Communication," in *MOBICOM Conference Proceedings*, 1998.

[20] M. Stemm and R. H. Katz, "Measuring and Reducing Energy Consumption in Network Interfaces in Handheld Devices," in *Proc. of 3rd International Workshop on Mobile Multimedia Communications (MoMuc-3)*, September 1996.

[21] C. Bisdikian, P. Bhagwat, B. P.Gaucher, F. J. Janniello, I. Korpeoglu, M. Naghshineh, and P. Pandoh, "WiSAP: A Wireless Personal Access Network for Handheld Computing Devices," *IEEE Personal Communications*, December 1998.

[22] R. Caceres and V.N. Padmanabhan, "Fast and Scalable Handoffs for Wireless Internetworks," in *MOBICOM Conference Proceedings*, November 1996.

[23] R.O. LaMaire, A. Krishna, J. Panian, and P. Bhagwat, "Wireless LANs and Mobile Networking: Standards and Future Directions," *IEEE Communications Magazine*, vol. 34, no. 8, pp. 86–94, August 1996.

[24] J. Haartsen, M. Naghshineh, J. Inouye, O. J. Joeressen, and W. Alle, "Bluetooth: Vision, Goals, and Architecture," *Mobile Computing and Communications Review*, vol. 2, no. 4, October 1998.

[25] B. Patel, P. Bhattacharya, Y. Rekhter, and A. Krishna, "An Architecture and Implementation Towards Multiprotocol Mobility," *IEEE Personal Communications*, vol. 2, no. 5, pp. 32–43, June 1995.

[26] C. E. Perkins, "IP Mobility Support," 1996, RFC 2002.