Proceedings of the 12th International IEEE Conference on Intelligent Transportation Systems, St. Louis, MO, USA, October 3-7, 2009

Integrated Software-Hardware Design for Ultra-Low Power Infrastructure Monitoring

Jingxian Wu and Scott C. Smith
Department of Electrical Engineering
University of Arkansas
Fayetteville, AR 72701, USA
{wuj, smithsco}@uark.edu

Abstract— **Ultra-low power communication and computing are of critical importance to the development of an infrastructure monitoring system, which is expected to have a very long life cycle under the constraints of extremely limited battery capacity or small energy scavenging devices. This paper discusses the enabling technologies for ultra-low power infrastructure monitoring, and includes a survey on existing ultra-low power infrastructure monitoring technologies. The survey covers both ultra-low power wireless communication technologies, which are usually developed in the form of software, such as protocols and algorithms, and ultra-low power digital hardware design techniques. Traditionally, the communication protocols and hardware are designed in isolation. The isolated design approach fails to capture the interaction between software and hardware or the fundamental tradeoff between communication and computing. We propose a new integrated software-hardware design approach to unify the development process of communication software and hardware. Several new integrated design methods are proposed, and they have the potential to lead to new ultra-low power technologies with power efficiency far beyond existing technologies.**

Keywords: infrastructure monitoring, ultra-low power, wireless sensor network, asynchronous digital design, integrated software-hardware design

I. INTRODUCTION

Wireless sensor networks (WSNs) developed for the automatic and remote monitoring of the health of critical military or civilian infrastructure possess many unique features that are not available in conventional wireless networks. Much infrastructure, such as bridges, tunnels, and buildings, has an extremely long life cycle on the order of years or decades, with a very slow rate of change, e.g., new data might only need to be collected once every few days or even months. As a result, infrastructure monitoring systems have extremely long delay tolerance with ultra-low data rate. In addition, data collected in the real world often contain redundancies due to the spatial correlation inherent in the monitored subject(s). The redundancy/correlation can be used to facilitate the design of infrastructure monitoring systems.

The long life cycle of the monitored infrastructure imposes formidable challenges to system design. The developed monitoring system is expected to have a very long life cycle (years) under the constraints of extremely limited battery capacity or small energy scavenging devices. Hence, an extremely stringent power budget, usually on the order of tens of micro-Watts, is required to power the operation of the entire system, including sensing, signal processing (DSP), and wireless communication with the outside world. Ultra-low power operation is therefore critically important to the development of infrastructure monitoring systems.

One of the objectives of this paper is to provide a brief survey of the enabling technologies for ultra-low power infrastructure monitoring. Existing ultra-low power infrastructure monitoring technologies can be classified into two categories, ultra-low power wireless communications, and ultra-low power digital hardware design, both of which are indispensible for the implementation of a practical infrastructure monitoring system. We review ultra-low power technologies in both of these two areas. It should be noted that the survey presented in this paper is by no means an exhaustive review of ultra-low low power communication and computing technologies. Instead, this is merely an account of some of the representative technologies available in the literature or recently developed by the authors. It aims at providing insights into the design of future infrastructure monitoring systems, and pointing out possible future research directions.

Ultra-low power wireless communication technologies are usually developed in the form of software, such as communication protocols and digital signal processing algorithms. The developed protocols and algorithms need to be implemented in digital hardware, which is another main source of power consumption. Traditionally, communication software and the underlying hardware are designed in isolation. There is minimal interaction between the communication protocol engineers and the digital hardware design engineers. This is also the case for almost all of the ultra-low power wireless sensing technologies. This isolated design methodology fails to capture the possible interactions between communication protocol and communication hardware, or the fundamental tradeoffs between communication and computing.

The second objective of this paper is to propose a new design paradigm, namely, integrated software-hardware design for ultra-low power communication, to fully exploit the potential benefits inherent in the rich interactions and tradeoffs between communication protocol and digital hardware. The integrated software-hardware design approach aims to unify the design process of communication protocol and underlying digital hardware to combine the unique features of the two processes and generate new ultra-low power technologies not possible with the traditional isolated design approach. We

978-1-4244-5521-8/09/$26.00 ©2009 IEEE

present in this paper several possible new integrated design approaches, such as complexity optimization, supply voltage scalability, power saving sleep mode, and mixed Application Specific Processor (ASP) and Application Specific Integrated Circuit (ASIC) design. It is expected that an integrated software-hardware design approach can lead to ultra-low power communication technologies with power efficiency far beyond that of existing software or hardware technologies.

II. Ultra-Low Power Wireless Communications

In this section, we give a brief survey on algorithms and protocols that have been developed specifically for low power or ultra-low power wireless communications.

A. Physical Layer

There are very limited works in the literature devoted to the development of low power communication technologies in the physical layer, which covers the operations of coding/decoding, modulation/demodulation, equalization, estimation, detection, digital filtering, and other DSP operations, etc.

Most of the existing physical layer low power communication technologies are designed by exploiting the tradeoff relationship between power efficiency, which is defined as the amount of transmission power, or signal to noise ratio (SNR), required to achieve certain bit error rate (BER), and bandwidth efficiency, which is defined as the maximum data rate supported in unit bandwidth. Power efficiency and bandwidth efficiency feature the most fundamental tradeoff relationship in the physical layer of a digital communication system. The tradeoff relationship can be explained in Fig. 1, where the BER performance of a binary phase shift keying (BPSK) modulated system with rate 1/2 convolutional code is compared with its uncoded counterpart. At the BER level of 10^{-3}, the SNR requirement of the coded system is 10 dB lower than that of the uncoded system, which means 10 dB less power is required by the coded system to achieve the same performance as the uncoded system. The superior performance is achieved at the cost of bandwidth efficiency and complexity. The bandwidth efficiency of the rate 1/2 convolutional code is only half of the uncoded system, due to the fact that each data bit is represented as 2 code bits after the encoding process.

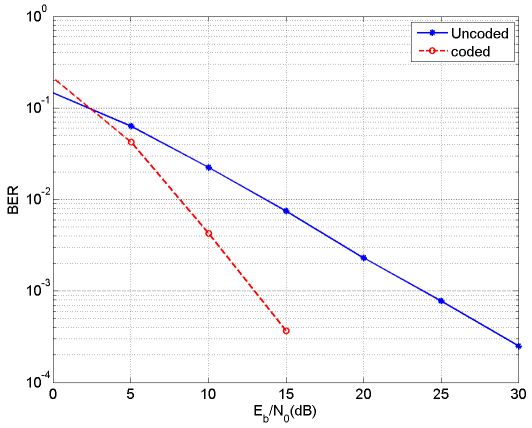

Fig. 1. BER of coded and uncoded systems.

Similar tradeoff relationships can also be achieved by employing modulation or other advanced communication or signal processing techniques, such as spread spectrum (SS) communication, multicarrier communication, and ultra-wide band (UWB) communication [1], [2].

The power saving factor, which is defined as the ratio between power consumption before and after the employment of power saving technology, for physical layer technologies is usually on the order of two to tens.

B. Media Access Control Layer

The power saving techniques in the media access control (MAC) layer can be classified into two categories.

The first category reduces power consumption by reducing or minimizing the duty cycle of the sensor nodes. The low data rate and long delay tolerance of infrastructure monitoring systems lead to extremely low duty cycle, and this property can be utilized to facilitate the MAC layer design.

To take advantage of the low duty cycle, a sensor node is in sleep mode most of the time to save power consumption. It should be noted that sleep mode is different from idle listening. The power consumption of idle listening is usually much higher compared to sleep mode, where power consumption can be reduced to as low as a few nano-Watts with advanced digital design techniques. However, while in sleep mode, the sensor node will not be able to receive any incoming signal. This poses new challenges for the design of efficient MAC algorithms that can cope with a network of nodes in sleep mode most of the time.

To solve the above problem, a preamble sampling scheme is presented in [3], [4] for an ad hoc wireless network. With preamble sampling, all sensor nodes in the network wake up periodically to sample the channel and detect if there is any information. All the sensor nodes share the same wake up period, T_W, and once woke-up, the sensor node listens to the channel for a short period of time, $T_P << T_W$. The relative sampling schedule offsets of different sensor nodes are independent and constant. When a node has information to transmit, it will always prefix a wake-up preamble with size equal to T_W to all data frames to ensure the message will be detected by the intended receiver. The long wake-up preamble incurs a large overhead in information delivery.

The second category of MAC layer power saving schemes reduce power consumption by carefully coordinating the transmissions among the sensor nodes to reduce collisions or interference among signals from different nodes, thus reducing the amount of wasted transmissions or retransmissions.

In recognition of the long delay tolerance and low data rate of infrastructure monitoring systems, a new exponentially-interval MAC (EI-MAC) scheme is proposed in [5]. In the EI-MAC, each node transmits at random, with the interval between two consecutive transmission attempts following an independent exponential distribution with mean $1/\lambda$. The transmission schedules of different nodes are independent. The message has a fixed length of $\tau << 1/\lambda$. If the signals from more than 2 nodes overlap with each other at the basestation or fusion center (FC), a collision is declared. In case of a collision,

the message will be discarded and no retransmission will be attempted. Since the monitored object changes very slowly in infrastructure monitoring systems, the loss of a very small percentage of packets will not significantly affect the integrity of the reconstructed information. The product of the parameters, $\lambda\tau$, defines the duty cycle of the system. It has been proved in [5] that if the duty cycle satisfies $\lambda\tau \approx 1/(n \log n)$, with n being the number of sensor nodes in the network, the probability of collision tends to 0 as $n \to \infty$. In other words, collision free MAC can be asymptotically achieved given a small enough duty cycle; and this condition is easily achieved for infrastructure monitoring systems.

C. Network Layer

There are a large number of various routing algorithms proposed for wireless sensor networks [6]. However, only a few of these schemes are designed with power consumption reduction as one of the objectives.

Low-energy adaptive clustering hierarchy (LEACH) [7] is one of the most popular energy-aware hierarchical routing algorithms. The key idea of LEACH can be summarized into three points. First, the nodes are grouped into clusters. Nodes within a cluster transmit to a cluster head, and only the cluster head can transmit to the basestation. Such localized control will save transmission power since the majority of the communication happens locally within a cluster. Second, the selection of cluster head changes randomly over time. The random rotation of cluster head balances the power consumption of the nodes. Third, the cluster head performs localized data fusion or data compression to reduce the amount of information that needs to be delivered to the basestation, thus further reducing overall power consumption of the entire network. Simulation shows that the LEACH scheme can achieve a power reduction factor of as much as 8. One of the limitations of LEACH is that it assumes that all cluster heads can directly transmit to the basestation, and such an assumption is not always true, especially for networks covering a large area.

In a spatially distributed wireless network, establishing multi-hop networks usually leads to better energy efficiency compared to single-hop communications, given the fact that the power of a signal decays on the order of $d^{-\alpha}$, where d is the transmission distance and α 7 2 is the pathloss exponent. A multi-hop energy-aware routing algorithm is proposed in [8] for energy constrained wireless ad hoc sensor networks. Unlike most other energy-aware routing algorithms that select the routing path by optimizing energy usage at a node, the scheme in [8] uses a new metric of network survivability, i.e., the connectivity of the entire network is maintained for as long as possible, and the energy health of the entire network is on the same order. Instead of finding a single optimum path and then consuming the energy of the nodes along that path, the energy-aware routing in [8] selects a set of good (optimum or sub-optimum) paths, and chooses one in a probabilistic fashion. In this way, the energy of the nodes in the network is consumed in a more balanced manner, and the survivability of the entire network is improved. Simulation results show that, compared to conventional routing algorithms, the proposed energy-aware routing reduces average energy consumption by 21.5%, and the lifetime of the network is increased by 44%.

D. Application Layer

The application layer achieves power reduction by exploiting the contents of the data transmitted in the network. Data collected in the real world often contains redundancies due to spatial correlation inherent in the monitored subject(s). This redundancy/correlation can be used to facilitate ultra-low power wireless sensing.

Motivated in part by recent results for compressive sensing [9]–[15], two compressive wireless sensing schemes, compressive transmission (CT) and compressive detection (CD), are proposed in [5]. Compressive sensing demonstrates that a relatively small number of random projections of a sparse signal, i.e., a signal that has many coefficients close to or equal to zero when represented in a certain domain, can retain most of the salient information in the original signal. In the CT scheme presented in [5], n spatially distributed sensing nodes first transmit the collected information to a compressing node, which projects the n-dimensional data to a transform domain with dimension $k << n$ by exploiting the spatial domain correlation of the original n-dimensional data. The compression node then transmits the k-dimensional information to FC. In the CD scheme, the n-dimensional data is directly transmitted to the FC, which then estimates the k transform domain coefficients instead of the original n data samples. Fig. 2 shows the comparison among the original data, which is a 1024×1024 pixel airport image, the data recovered from conventional direct sensing (DS) scheme, and the data recovered from the CT and CD schemes, respectively. It is obvious from the figure that the CT and CD schemes preserve a much higher fidelity compared to the conventional DS scheme. By exploiting the special structure of the application data, the CT and CD schemes can achieve a better fidelity with less transmission power. Numerical examples show that the CT and CD schemes can achieve a power saving factor as high as 25.

E. Cross-Layer

The interactions among protocol layers can be utilized to significantly reduce the power consumption of a wireless network. Recently, the cross-layer design approach has attracted great attention for its potential to achieve more effective optimization with more available knowledge in multiple layers.

Routing with compression can be considered as a cross-layer technique since it involves both routing in the network layer, and data aggregation in the application layer. Several cross-layer routing with compression algorithms are presented and analyzed in [16], where the spatial correlation of data in the application layer is used to reduce the total amount of information to be delivered to FC, and to facilitate the routing process. Specifically, two schemes are investigated and compared. The first scheme is routing-driven compression, where data is routed through shortest paths to FC, with compression taking place wherever possible along the paths; the second scheme is compression-driven routing, where the route is selected such that the data from all the nodes can be compressed sequentially. The interaction between compression

in the application layer, and routing in the network layer, allows flexible design of the routing algorithms and high efficiency delivery of information.

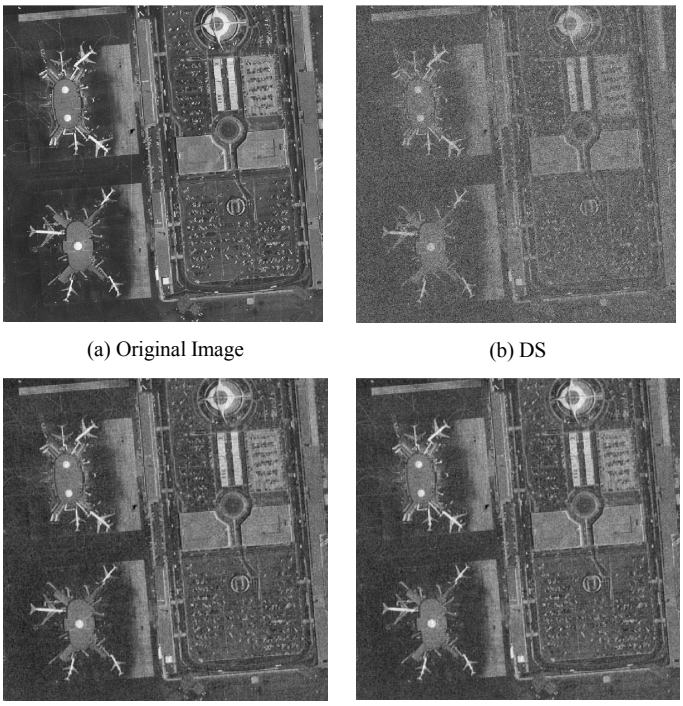

(a) Original Image (b) DS

(c) CD (d) CT

Fig. 2. Comparison of different compressive sensing schemes.

III. ULTRA-LOW POWER DIGITAL DESIGN

Power consumption has long been one of the top considerations for digital integrated circuit (IC) designers. The high density integration of today's IC led to substantially increased power dissipation in both active mode, when gates are switching due to the IC processing data, and idle mode, when gates are not switching while the IC is waiting for the next set of data to process. Power used during active mode, called dynamic power, is proportional to the square of the IC supply voltage, V_{dd}^2, such that utilizing a smaller technology, which requires a smaller supply voltage, decreases dynamic power for the same circuit operating at the same speed [17]. However, power dissipated during idle mode, called static power, which is due to leakage current flowing through transistors that are supposedly turned off, significantly increases as transistor size is reduced to today's deep submicron region [18]. Hence, methods such as asynchronous digital design [19] and Multi-Threshold CMOS (MTCMOS) [20] have been developed to substantially reduce dynamic and static power, respectively, to achieve significantly decreased energy consumption for low power applications.

A. Asynchronous Digital Design

For the last three decades the focus of digital design has been primarily on synchronous, clocked architectures. Asynchronous, clockless circuits require less power, generate less noise, and produce less electro-magnetic interference (EMI), compared to their synchronous counterparts, without degrading performance. Therefore, as demand increases for designs with higher performance, greater complexity, and decreased feature size, asynchronous paradigms will become more prevalent in the multi-billion dollar semiconductor industry, as predicted by the International Technology Roadmap for Semiconductors (ITRS) [21], [22]. ITRS shows that asynchronous circuits accounted for 11% of chip area in 2008, compared to 7% in 2007, and estimates they will account for 23% of chip area by 2014, and 35% of chip area by 2019 [23].

The ***delay-insensitive (DI)*** asynchronous NULL Convention Logic (NCL) paradigm [24], which assumes that delays in both logic elements and interconnects are unbounded, such that timing analysis is not required, offers the best opportunity for integrating asynchronous digital design into the predominantly synchronous semiconductor design industry for the following reasons:

1) The framework for NCL systems consists of DI combinational logic sandwiched between DI registers, as shown in Fig. 3, which is very similar to synchronous systems, such that the automated design of NCL circuits can follow the same fundamental steps as synchronous circuit design automation. This will enable the developed DI design flow to be more easily incorporated into the chip design industry, since the tools and design process will already be familiar to designers, such that the learning curve is relatively flat.
2) NCL systems are delay-insensitive, making the design process much easier to automate than other non-DI asynchronous paradigms, since minimal delay analysis is necessary to ensure correct circuit operation.
3) NCL systems have power, noise, and EMI advantages compared to synchronous circuits, performance and design reuse advantages compared to synchronous and non-DI asynchronous paradigms, area and performance advantages compared to other DI paradigms, and have a number of advantages for designing complex systems, like Systems-on-Chip (SoCs), including substantially reduced crosstalk between analog and digital circuits, ease of integrating multi-rate circuits, and facilitation of component reuse and technology migration..

As the trend towards higher clock frequency and smaller feature size continues, power consumption, noise, and EMI of synchronous designs increase significantly. With the absence of

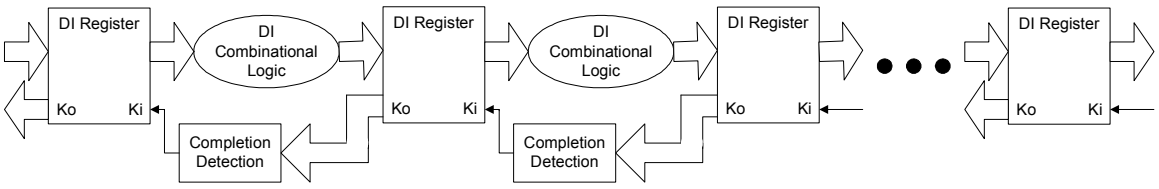

Fig. 3. NCL system framework: input wavefronts are controlled by local handshaking signals, K_i and K_o, and Completion Detection, instead of by a global clock signal.

a clock, NCL systems aim to reduce power consumption, noise, and EMI. NCL circuits designed using CMOS exhibit an inherent idle behavior since they only switch when useful work is being performed, unlike clocked Boolean circuits that switch every clock pulse, unless specifically disabled through specialized circuitry, which itself requires additional area and power. NCL circuits adhere to monotonic transitions between DATA and NULL, so there is no glitching, unlike clocked Boolean circuits that produce substantial glitch power. NCL systems better distribute switching over time and area, reducing the occurrence of hot spots, peak power demand, and system noise, unlike clocked Boolean circuits where much of the circuitry switches simultaneously at the clock edge [19].

Furthermore, NCL systems are very tolerant of power supply variations, allowing cheaper power supplies to be used and supply voltage to be dramatically reduced, even below the threshold voltages of the transistors, while maintaining correct circuit operation, which can yield ultra-low power consumption. The strong tolerance to power supply variation is critical to ultra-low power consumption. It is well known that the supply voltage of a battery gradually drops with the consumption of its internal energy. Traditional ICs will stop working when the supply voltage drops below a certain threshold (e.g., 1.2V), at which level there might still be a significant amount of energy remaining inside the battery. ICs designed with NCL, on the other hand, can operate with a supply voltage as low as 0.1 volt, when most of the energy inside the battery has been consumed. Therefore, NCL can use most of the energy inside a battery, and lasts much longer than a synchronous IC. In addition, many of the power scavenger or power storage devices, such as a super capacitor, can only provide a supply voltage below 1V. The power supply scalability of NCL makes it possible for the IC to work with these devices to obtain a very long life cycle.

B. Multi-Threshold CMOS (MTCMOS)

It is well-known that lowering V_{DD} is the most efficient approach to reducing power consumption. However, a lower V_{DD} causes insufficient gate overdrive, which in turn causes increased delay. In order to reduce power while maintaining performance, transistor threshold voltage, V_t, needs to be scaled with V_{DD}. Although current fabrication technology has significantly reduced threshold voltage, this decrease causes an exponential increase in leakage power, due to the increase in sub-threshold leakage. To reduce leakage, MTCMOS uses high-V_t transistors to gate power and ground of a low-V_t logic block [25]. When the high-V_t transistors are turned on, the low-V_t logic is connected to virtual ground and power, and switching is performed through fast, low-V_t devices. When the circuit enters sleep mode, the high-V_t gating transistors are turned off, resulting in a very low sub-threshold leakage current from V_{DD} to ground. MTCMOS is a very attractive technique for reducing sub-threshold leakage current during standby mode because existing designs can be easily modified into MTCMOS blocks by simply adding high-V_t power supply switches/transistors; and the semiconductor fabrication process only requires one additional implant processing step to provide the extra threshold voltage level [25].

One method of designing MTCMOS Boolean gates is shown in Fig. 4 [20]. A high-V_t PMOS transistor, *P1*, a standard-V_t PMOS transistor, *P0*, and a standard-V_t NMOS transistor, *N0*, form the sleep circuitry. In normal operating mode, *Sleep* is logic0; the circuit behaves exactly like a normal CMOS circuit. In sleep mode, *Sleep* is logic1, which causes *P0* and *N0* to be ON and *P1* to be OFF. Since *P0* is ON, the pull-up network, *PUN*, is now between two points at equal voltage potential, V_{DD}; hence, no leakage current should flow through the *PUN*. Similarly, *N0* is ON so no leakage current should flow through the pull-down network, *PDN*, since it is also between two points of equal potential, 0V. Since *Out* is connected to *X2*, during sleep mode the output value will always be logic0. The leakage power dissipation during sleep mode will only be through *P1*, which is a high-V_t device that is turned OFF; therefore, leakage will be minimal.

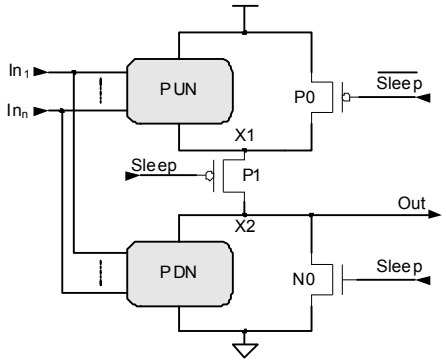

Fig. 4. Generic MTCMOS Boolean gate [20].

Unfortunately, three serious drawbacks hinder the widespread use of MTCMOS: 1) the generation of *Sleep* signals is timing critical, often requiring complex logic circuits; 2) synchronous circuits lose data when the power transistors are turned off; and 3) applying MTCMOS on a per gate basis for Boolean circuits requires excessive overhead, and applying the technique on a per stage basis requires proper sizing of the sleep transistors for correct circuit operation, which is a very difficult task. However, all three of these drawbacks can be eliminated by utilizing NCL in conjunction with the MTCMOS technique, as explained in the following section.

C. Delay-Insensitive Multi-threshold LOGic (DIMLOG)

Delay-Insensitive Multi-threshold LOGic (DIMLOG) [26] combines the MTCMOS technique with the naturally energy efficient delay-insensitive NCL paradigm. This combination yields an ultra-low power digital design methodology that solves the three drawbacks of utilizing MTCMOS with the synchronous paradigm: 1) the handshaking signal, K_o, between every two NCL pipeline stages naturally serves as the *Sleep* signal for the combinational logic block in the next stage, without any additional circuitry or delay analysis; 2) DIMLOG circuits do not lose data during sleep mode, since sleep mode is applied in lieu of the NULL cycle, which causes all

threshold gates in the combinational logic block to become logic0 anyway; 3) in synchronous MTCMOS design, sleep transistors are added to each stage in order to reduce area overhead, not to each gate, which makes proper sizing of the sleep transistors a very difficult task, critical for correct circuit operation. However, the DIMLOG methodology applies the MTCMOS concept at the gate level instead of at the stage level, since NCL gates are more complex and powerful than Boolean gates, requiring fewer gates to implement the same logic function; hence, implementing MTCMOS in NCL threshold gates will not result in a large area overhead.

As an example, Fig. 5 shows both the original TH23 gate and its MTCMOS version. In Fig. 5b, the circled transistors are high-V_t while the others are low-V_t. During regular operation, the *Sleep* signal is logic0 and its complement is logic1, such that the gate functions normally. During sleep mode, the *Sleep* signal becomes logic1, which turns on the output pull-down transistor, forcing the output to logic0 immediately. It also turns off the high-V_t transistor above the output inverter to disconnect its power. At the same time, the *Sleep* complement signal becomes logic0, disconnecting ground from the gate logic. Since the power-to-ground paths are broken and all PMOS transistors are high-V_t, the leakage power is minimized.

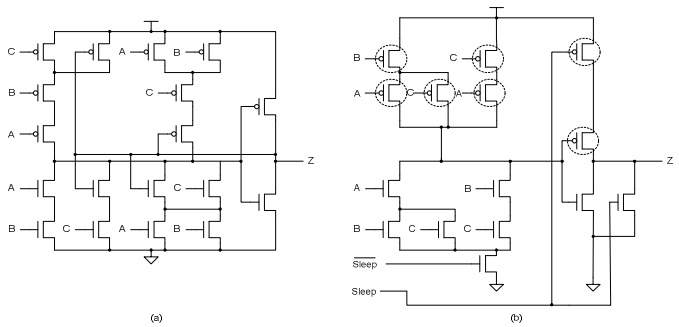

Fig. 5. (a) Original TH23 gate (b) MTCMOS TH23 gate.

8-bit×8-bit pipelined unsigned array multipliers were designed and compared using the 1.2V IBM 8RF-DM 130 nm CMOS process, which showed that an initial DIMLOG design reduced leakage power and active energy by 150× and 1.8×, respectively, compared to the original NCL version [26], while requiring negative area overhead. A subsequent DIMLOG design further reduced leakage power by 22.3%, active energy by 17.5%, and achieved 8.1% higher throughput [27]. Recent results show that an improved DIMLOG architecture requires less than 40% active energy and 62% leakage power, compared to a synchronous MTCMOS implementation, for an IEEE-standard single-precision floating-point co-processor. Additionally, supply voltage can be reduced from 1.2V to a sub-threshold voltage of 0.1V to reduce power, without affecting correct DIMLOG circuit operation; however, processing time increases as supply voltage is dropped. Therefore, the minimum energy point was found to be at a supply voltage of 0.13V. The complete DIMLOG methodology is detailed in [28].

IV. Integrated Software-Hardware Design

In almost all the existing ultra-low power communication systems, the communication software and the physical hardware are designed in isolation. There are usually minimal interactions between the communication protocol engineers and the hardware design engineers during system design and development. This isolated design methodology fails to capture possible interactions between the communication protocol and communication hardware. The software development only focuses on the communication aspect of the system, yet overlooks the computation requirement and the physical limitation imposed by the underlying hardware; the hardware development does not take advantage of the unique properties of an infrastructure monitoring system, such as low data rate and long latency tolerance.

We propose a new integrated communication protocol and digital hardware design methodology to fully exploit the interactions and mutual benefits between these two. The key idea is to achieve an optimum tradeoff between communication and computation. We will develop new design methodologies to integrate the development process of communication protocols and digital hardware, such that the total power savings achieved by the integrated design exceeds the sum of the power savings achieved by simply combining the individually developed components.

A. Complexity Optimization

Many of the communication protocols and algorithms achieve ultra-low power communication at the cost of increased complexity. For example, in the physical layer, the employment of channel coding provides protection against channel distortion, but this demands more sophisticated processing and computation in the underlying hardware. However, from the perspective of hardware design, increased complexity leads to increased power consumption. With a highly complicated channel coding scheme, it is possible that the increased power in hardware will dominate the power savings in communication protocol. Therefore, during the design of ultra-low power communication technologies, it is essential to identify the power consumption and power saving from all possible sources, in both communication protocol and hardware, and identify the optimum tradeoff point. From the perspective of protocol design, the power consumption decreases as complexity increases; from the perspective of hardware design, the power consumption increases as complexity increases. The optimum operation point is therefore the intersection between these two curves. In reality, there might be more than two tradeoff parameters, yet the optimum or sub-optimum operation point can be identified through a similar approach. Such a global complexity-power tradeoff analysis is impossible with the conventional isolated software and hardware design methodology.

B. Adaptive Supply Voltage Scaling

The asynchronous digital design with NCL allows for extreme supply voltage scaling to sub-threshold operation. Sub-threshold operation achieves ultra-low power consumption at the cost of increased delay. The property of supply voltage scalability of asynchronous digital hardware can be integrated

in the communication protocol design to further reduce system power consumption. For example, a high supply voltage can be used for circuits corresponding to critical communication protocols, and those non-critical operations can operate at a lower supply voltage.

We propose an adaptive modulation, coding, and voltage scaling (AMCV) technique to take advantage of the power savings in both communication protocol and asynchronous digital circuit. From the perspective of communication software, the transmitter can select different combinations of modulation and coding schemes from a list of candidates based on the channel condition and targeted quality of service (QoS) to achieve a tradeoff between power efficiency and bandwidth efficiency. With the supply voltage scalability of asynchronous digital circuits, a third dimension, supply voltage, can also be adaptively adjusted to reduce power consumption. Reducing the supply voltage leads to slower computation (from nano-second/machine cycle to micro-second/machine cycle), yet it significantly reduces the power consumption, since power is proportional to V_{DD}^2. Such a reduction in computation speed will not apparently affect an infrastructure monitoring system given its long delay tolerance. A system with AMCV will have a finite list of modulation schemes, coding schemes, and predefined voltage levels. The transmitter will choose a certain modulation, coding and voltage level based on the current operation condition. For example, the transmission of critical delay sensitive information, such as important network control information, can be performed with a more complex modulation and coding scheme using a higher supply voltage. A lower supply voltage with simpler modulation and coding schemes can be used for the transmission of regular sensing data. The scalability of the supply voltage adds great flexibility to the design of adaptive communication systems without requiring the design of different sets of hardware.

C. *Partial Sleep Mode*

To achieve low power consumption for infrastructure monitoring, the underlying hardware is required to, 1) have extremely low leakage current during idle mode, and, 2) the transition between idle mode and active mode should incur minimum overhead. The NCL asynchronous digital design technique using CMOS gates inherently meets both of these requirements. After the completion of each NCL stage, the output of the DI register goes back to the NULL state, which corresponds to idle mode (i.e., all gates are logic0). NCL can then be combined with MTCMOS, as explained in Section III.C, to convert the idle mode into sleep mode, where power and ground are disconnected from the NCL gates, to substantially decrease leakage power to achieve ultra-low power consumption. The handshaking signal, K_o, between every two NCL pipeline stages naturally serves as the *Sleep* signal for the combinational logic block in the next stage. Therefore, no additional circuitry or control signals are required to achieve the transition between sleep mode and active mode.

Joint software-hardware design can be performed by leveraging the inherent sleep mode and natural sleep to active transition of NCL hardware. The two properties of wireless sensor nodes, extremely low duty cycle and periodic channel listening, can be achieved simultaneously with a partial sleep mode design in the digital hardware. In partial sleep mode, the entire baseband processing unit is in sleep mode, except for the circuitry corresponding to idle listening and the RF chain. The idle listening circuitry can operate independent of the remaining circuit, and it can operate with sub-threshold supply voltage to further reduce power consumption. Note that NCL registers do not sleep, they just remain idle, such that they are ready to process the next input whenever it arrives, and still dissipate only minute leakage power during idle mode. During regular operation, the MAC protocol periodically switches between full sleep mode, where the entire circuit is turned off, and partial sleep mode, where only the idle listening circuitry is active. Once an incoming signal is detected during partial sleep mode, the circuit will convert to active mode. The communication protocol with the corresponding circuit can be designed in such a way that the circuit in sleep mode will not affect the proper operation of the remaining active circuit, since the combinational logic operates asynchronously in NCL.

D. *Mixed ASP and ASIC Design*

The ultra-low power digital signal processing algorithm and communication protocols, explained in Section II, can be implemented in hardware using the ultra-low power DIMLOG architecture explained in Section III.C. The two extreme methods to do this are: 1) an ASP could be developed, which consists of a general purpose processor optimized specifically for executing the instructions needed for the communications algorithms; and 2) the communications algorithms could be implemented in fully custom hardware, as an ASIC. Comparing the ASP to the ASIC, the ASP would be much less complex and require far less chip area to implement, but would have much slower performance. Therefore, the ASP would require less power over a longer period of time, whereas the ASIC would require more power over a shorter period of time, for an identical communiqué.

To achieve the optimum power tradeoff between ASP, ASIC, and various communication algorithms, it is expected that both ASIC and ASP should be used. According to the properties of different communication algorithms and protocols, some portion of the operation should be implemented as an ASIC while the rest is implemented on an ASP, in order to minimize energy usage. For example, the constantly repeated operations, such as data sensing and idle listening, can be implemented as an ASIC; the operations that require regular adaptation, such as AMCV, or the irregularly performed operations, such as network topology probing, can be implemented on an ASP to improve flexibility and reduce required hardware.

V. CONCLUSIONS

The stringent power budget of infrastructure monitoring systems necessitates the development of ultra-low power communication technologies and ultra-low power digital hardware design techniques that are tailored specifically for infrastructure monitoring systems. A large number of wireless communication technologies have been proposed in the literature to achieve ultra-low power communication and computing. These technologies can achieve tens to hundreds

order of magnitude power reduction factor . To achieve further power reduction, we proposed several integrated software-hardware design approaches, such as complexity optimization, adaptive supply voltage scaling, partial sleep mode, and mixed ASP and ASIC design, which can exploit the unique interactions and tradeoffs between communication protocols and the underlying hardware. These methods have the potential to achieve a power reduction factor that is unattainable in the traditional isolated software and hardware design approaches. Future work will focus on the analysis, implementation, and thorough verification of the various integrated design approaches through computer simulation, hardware implementation, and field measurements.

ACKNOWLEDGEMENT

This work was supported in part by the Arkansas Biosciences Institute under Grant 2410.

REFERENCES

[1] C. Schurgers, O. Aberthorne, and M. B. Srivastava, "Modulation scaling for energy aware communication systems," in *Proc. Intern. Symp. Low Power Electronics Design*, pp. 96–99, 2001.

[2] M. Z. Win and R. A. Scholtz, "Ultra-wide bandwidth time-hopping spread-spectrum impulse radio for wireless multiple-access communications," *IEEE Trans. Commun.*, vol. 48, pp. 679–691, 2000.

[3] A. El-Hoiydi, "Spatial TDMA and CSMA with preamble sampling for low power ad hoc wireless sensor network," *Intern. Symp. Computers Commun.*, pp. 685–692, July 2002.

[4] A. El-Hoiydi and J.-D. Decotignie, "WiseMAC: an ultra low power MAC protocol for the downlink of infrastructure wireless sensor networks," *Intern. Symp. Computers Commun.*, pp. 244–251, July 2004.

[5] J. Wu, "Low power compressive wireless sensing for ditstributed wireless networks," submitted to *IEEE Militray Commun. Conf.*, 2009.

[6] K. Akkava and M. Younis, "A survey on routing protocols for wireless sensor networks," *Ad Hoc Networks*, vol. 3, pp. 325–349, 2005.

[7] W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan, "Energy-efficient communication protocol for wireless microsensor networks," *33rd Hawaii Intern. Conf. System Sci.*, 2000.

[8] R. C. Shah and J. M. Rabaey, "Energy aware routing for low energy ad hoc sensor networks," *IEEE Wireless Commun. Networking Conf*, pp. 350–355, Mar. 2002.

[9] R. G. Baraniuk, "Compressive sensing," *IEEE Signal Processing Magazine*, vol. 24, pp. 118–121, July 2007.

[10] D. L. Donoho, "Compressed sensing," *IEEE Trans. Information Theory*, vol. 52, pp. 1289–1306, Apr. 2006.

[11] E. J. Candes, "Compressive sampling," *Intern. Congress Mathematicians*, pp. 1–20, 2006.

[12] V. K. Goyal, A. K. Fletcher, and S. Rangan, "Compressive sampling and lossy compression," *IEEE Signal Processing Magazine*, vol. 25, pp. 48–56, Mar. 2008.

[13] E. J. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Information Theory*, vol. 52, pp. 489–509, Feb. 2006.

[14] A. K. Fletcher, S. Rangan, and V. K. Goyal, "On the rate-distortion performance of compressed sensing," *IEEE Intern. Conf. Acoustics, Speech and Sig. Processing*, vol. 3, pp. 885–888, Apr. 2007.

[15] E. J. Candes and T. Tao, "Decoding by linear programming," *IEEE Trans. Information Theory*, vol. 51, pp. 4203–4215, Dec. 2005.

[16] S. Pattem, B. Krishnamachair, and R. Govindan, "The impact of spatial correlation on routing with compression in wireless sensor networks," *Intern. Conf. Information Processing Sensor Networks*, pp. 28–35, Apr. 2004.

[17] A. Wang and A. Chandrakasan, "Optimal Supply and Threshold Scaling for Subthreshold CMOS Circuits," *IEEE Computer Society Annual Symposium on VLSI*, 2002.

[18] S. Mutoh, T. Douseki, Y. Matsuya, T. Aoki, S. Shigematsu, J. Yamada, "1-V Power Supply High-Speed Digital Circuit Technology with Multithreshold-Voltage CMOS," *IEEE Journal of Solid-State Circuits*, Vol. 30/8, pp. 847-854, 1995.

[19] J. McCardle and D. Chester, "Measuring an Asynchronous Processor's Power and Noise," *Synopsys User Group Conference (SNUG)*, Boston, 2001.

[20] P. Lakshmikanthan, K. Sahni, and A. Nunez, "Design of Ultra-Low Power Combinational Standard Library Cells Using a Novel Leakage Reduction Methodology," *IEEE International SoC Conference*, 2006.

[21] http://www.itrs.net/Links/2003ITRS/Design2003.pdf (available August 2009).

[22] http://www.itrs.net/Links/2007ITRS/2007_Chapters/2007_Design.pdf (available August 2009).

[23] http://www.itrs.net/Links/2008ITRS/Home2008.htm (available August 2009).

[24] K. M. Fant and S. A. Brandt, "NULL Convention Logic: A Complete and Consistent Logic for Asynchronous Digital Circuit Synthesis," *International Conference on Application Specific Systems, Architectures, and Processors*, pp. 261-273, 1996.

[25] J. Kao and A. P. Chandrakasan, "Dual-Threshold Voltage Techniques for Low-Power Digital Circuits," *IEEE Journal of Solid-State Circuits*, Vol. 35/7, 2000.

[26] A. Bailey, A. Al Zahrani, G. Fu, J. Di, and S. C. Smith, "Multi-Threshold Asynchronous Circuit Design for Ultra-Low Power," *Journal of Low Power Electronics*, Vol. 4/3, pp. 337-348, 2008.

[27] A. Al Zahrani, A. Bailey, G. Fu, and J. Di, "Glitch-Free Design for Multi-Threshold CMOS NCL Circuits," *Great Lakes Symposium on VLSI*, 2009.

[28] Scott C. Smith and Jia Di, "Chapter 6.2: Multi-Threshold CMOS (MTCMOS) for NCL (MTNCL)" in *Designing Asynchronous Circuits using NULL Convention Logic (NCL)*, Synthesis Lectures on Digital Circuits and Systems, Vol. 4/1, July 2009, Morgan & Claypool Publishers (doi: 10.2200/S00202ED1V01Y200907DCS023).