

A Power-Managed Protocol Processor for Wireless Sensor Networks

M. Sheets, F. Burghardt, T. Karalar, J. Ammer, Y.H. Chee, J. Rabaey

Berkeley Wireless Research Center
University of California, Berkeley
msheets@eecs.berkeley.edu

Abstract

Wireless sensor network applications, such as environmental control in smart building and ecological monitoring, require low-power nodes that operate their entire lifetime without changing batteries. This paper describes the power management architecture for a digital protocol processor for a sensor network node. Eight subsystems implement the baseband through application protocol layers and are controlled by a centralized power manager. The prototype chip, implemented in 130nm CMOS, operates at 1.0V with an average power consumption of 150 μ W during normal operation.

Keywords: wireless sensor networks, power management and MT-CMOS

Introduction

A. Functionality

The presented chip (Charm) implements a protocol stack that is tailored for wireless sensor network applications. Subsystems generally follow the OSI reference model [1] and include the application, network, data link, and digital baseband portion of the physical layers. The protocol stack is augmented to include a location subsystem which triangulates the nodes position using information from fixed anchor nodes in the network [6].

As shown in Fig. 1, the Charm chip contains a synthesized 8051-compatible microcontroller with 64k of program/data RAM, two 1kB packet queues, a custom data-link layer (DLL), a neighborhood management subsystem, the digital portion of a custom baseband, a location computation subsystem, and several external interfaces. The chip is designed to interface to an external On-Off Keyed (OOK) two-channel radio [4] to form a sensor node [5].




Fig. 1 Block diagram of Charm protocol processor.

B. Power Management

Power management is critical to maximize battery life and enable operation on scavenged energy. Sensor network nodes typically have low average duty cycles, so deactivating idle subsystems can yield a significant power savings. Unfortunately, as process dimensions shrink, the effect of

CMOS leakage power is increasing and degrades the benefits of deactivating subsystems. Numerous circuit techniques have been proposed to reduce the effects of leakage power [2], with supply rail gating (MT-CMOS with sleep transistors) as the most promising. Unfortunately, supply rail gating is complicated by the need to restore the state of the logic upon wake-up. The approach in this paper reduces the power rail voltage to a data retention voltage (DRV) that maintains the state in the logic, while still reducing leakage current. Some overhead is required to control the power rails, such as [3], and this paper describes the implementation of a centralized power manager (PM).

Power Architecture

The activity profile varies widely between layers, so it is advantageous to form power domains (PD) that roughly correspond to the protocol stack layers. The supply of each PD can be separately switched between the nominal supply (vddhi) and the DRV (vddlo) using the power control circuit in Fig. 2a. The switch cell is implemented as a standard cell and is repeated every 30 μ m, to match the stride of the global power grid. The active mode PMOS is sized to reduce leakage by 80%, while still supplying enough active current for the cells surrounding it. Each switch contains a buffer powered by vddhi that forms a tree for distribution of the control signal. Switch cells are placed using a script that easily integrates into an industry standard P&R design flow, and the virtual vdd is connected to the standard cell supply, so no changes are required to the library provided by the foundry.

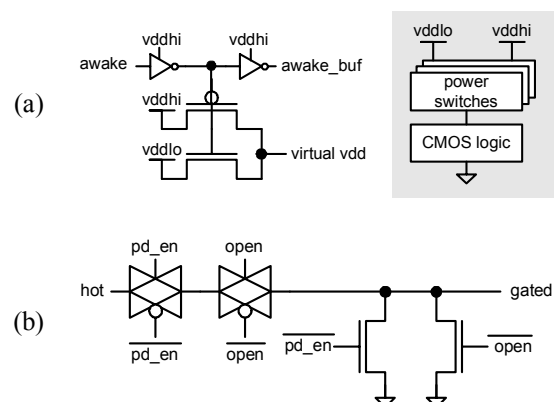


Fig. 2 Power domain circuits: a) power switch, b) signal wall.

Communication between protocol layers are typically infrequent bursts, such as a packet. Thus, a session-based approach is used to enable signaling between PDs. Within each PD, related I/O signals are bundled into ports that can be either open or closed. A PD can influence the port state by issuing power control messages through a standard Power

InterFace (PIF) to the power manager (PM). The PIF in each PD enables a plug-and-play approach for PD and PM. To prevent spurious signaling from corrupting a sleeping PD, the signal wall shown in Fig. 2b is used to ground PD I/O. Ground is chosen because, unlike the PD power supply, it is never gated. Signal walls are preferentially placed near the PD boundaries, since their control signals must use the ungated supply.

The PM implements a scheduling policy to ensure that PDs are active when necessary and put to sleep at all other times. Because it takes only one clock cycle to activate or deactivate a PD, a simple reactive policy is used. The PD is deactivated when 1) all its I/O ports are closed, and 2) the PDs *can_sleep* bit is set in the PM. The latter ensures that PDs can stay awake for independent processing until completion. The PM is programmed with the interconnections between PD ports and ensures that the appropriate target PD is awake whenever another PD opens a connected port.

Architectural decisions are made to allow block deactivation as much as possible, since this yields the lowest chip power. First, counters are common in protocol processors to set rates and detect timeouts, and these would ordinarily prevent a PD deactivation. Instead, a centralized system time wheel is included in the PM, and PDs can schedule a self-reactivation by setting a virtual alarm timer through the PIF. The PM sorts the alarms, selects only the most urgent one, and performs a low-power comparison to a single 24-bit counter that implements the system time-wheel. The system time-wheel operates on a derived, slower 80kHz timer clock to reduce switching activity for these always active circuits.

A second architectural power reduction method is to aggregate shared data in another PD. For example, there are TX/RX packet queues between the network layer (dw8051 microcontroller PD) and the dll PD. Both domains are sizeable and their time scales are quite different, so locating the packet queues in either would waste power. Instead, a separate, small queues PD allows independent access from both layers, enabling each to sleep more often. A similar approach is used to aggregate neighborhood information for use by the DLL, network, and locationing layers.




Fig. 3 Die photo and floorplan of Charm protocol processor.

Technology	130nm CMOS
Dimensions	2.7x2.7mm ²
Transistors	3.2M
Clocks freqs	8MHz (main), 80kHz (timer)
Memory	68kB
Core supply	1.0V
DRV supply	0.3-1.0V
I/O supply	1.8V
Leakage	53.6μW
Ave. power	150μW

Results

The Charm chip, shown in Fig. 3, is implemented in a 0.13μm triple-well, bulk digital CMOS process with six metal layers. Low-leakage transistors are used whenever possible, and the chip area is 2.7x2.7mm². The core logic is supplied by 1.0V and I/O voltage is 1.8V for easy connection with off-the-shelf components. The vddlo supply can be externally driven or internally generated by an on-chip voltage

converter. Since commercial clock oscillators consume several mW's of power, the chip includes a custom low-power clock generator that requires only 10μW to sustain an 8MHz oscillation with an external crystal.




Fig. 4 PD activity (active low) and measured power during TX.

With the power control logic disabled, the leakage power of the chip core is 250.1μW. When power control is active, there is 10μW of leakage (at 0.3V) inside the PDs, with an additional 43.6μW at the top level for buffers on long wires between PDs and always active logic. Thus, there is roughly a 5X reduction in leakage power during the sleep state, which is the most common. Fig. 4 shows the control signals for the power switches and the associated power consumption when transmitting a broadcast packet. The largest contributors to active power are the dll and baseband PDs (632μW and 71.4μW, respectively), since these periodically listen for data on the channel. The sampling rate is programmable and for the shown 100ms RX sampling period, the average power consumption of the chip is 132μW. Actual power consumption depends highly on the packet frequency and the amount of processing required for each packet, and for a typical low data rate network is about 150μW.

Acknowledgements

Authors want to thank ST Microelectronics for chip fabrication. Funding was provided by DARPA, GSRC, and the BWRC member companies.

References

- [1] Information technology – open systems interconnection – basic reference model. Technical Report ISO/IEC 7498-1:1994, International Organization for Standardization (ISO), 1994.
- [2] J. Kao, S. Narendra, and A. Chandrakasan. Subthreshold leakage modeling and reduction techniques. In IEEE/ACM ICCAD, pp. 141-8, San Jose, CA, USA, 2002.
- [3] H. Mizuno and T. Kawahara. ChipOS: open power management platform to overcome the power crisis in future LSIs. In IEEE ISSCC, pp. 344-5,463, San Francisco, CA, USA, 2001.
- [4] B. Otis, Y. H. Chee, and J. Rabaey. A 400μW-RX 1.6mW-TX super-regenerative transceiver for wireless sensor networks. In IEEE ISSCC, pp. 396-7,606, San Francisco, CA, 2005.
- [5] M. Sheets, B. Otis, F. Burghardt, J. Ammer, T. Karalar, P. Monat, and J. Rabaey. A (6x3)cm2 self-contained energy-scavenging wireless sensor network node. In Wireless Personal Multimedia Communications, Abano Terme, Italy, 2004.
- [6] T. Karalar, S. Yamashita, M. Sheets, J. Rabaey. A low power localization architecture and system for wireless sensor networks. In IEEE Signal Processing Systems Design and Implementation, pp. 89-94, Austin, TX, 2004.