AIDA: Adaptive Application Independent Data Aggregation in Wireless Sensor Networks

Tian He Brian M. Blum John A Stankovic Tarek Abdelzaher
Department of Computer Science
University of Virginia

Abstract

Sensor networks, a novel paradigm in distributed wireless communication technology, have been proposed for use in various applications including military surveillance and environmental monitoring. These systems could deploy heterogeneous collections of sensors capable of observing and reporting on various dynamic properties of their surroundings in a time sensitive manner. Such systems suffer bandwidth, energy, and throughput constraints that limit the quantity of information transferred from end to end. These factors coupled with unpredictable traffic patterns and dynamic network topologies make the task of designing optimal protocols for such networks difficult. Mechanisms to perform data centric aggregation utilizing application specific knowledge provide a means to augmenting throughput, but have limitations due to their lack of adaptation and reliance on application specific decisions. We therefore propose a novel aggregation scheme that adaptively performs application independent data aggregation in a time sensitive manner. Our work isolates aggregation decisions into a module that resides between the network and the data link layer and does not require any modifications to the currently existing MAC and network layer protocols. We take advantage of queuing delay and the broadcast nature of wireless communication to concatenate network units into an aggregate using a novel adaptive feedback scheme to schedule the delivery of this aggregate to the MAC layer for transmission. In our evaluation we show that end-to-end transmission delay is reduced by as much as 80% under heavy traffic loads. Additionally, we show as much as a 50% reduction in transmission energy consumption with the addition of only 2 bytes of header overhead per network unit. Theoretical analysis, simulation, and a test-bed implementation on Berkeley's MICA motes are provided to validate our claims.

1. Introduction

Wireless Sensor Networks have emerged as a new information-gathering paradigm based on the collaborative effort of a large number of sensing nodes. In such networks, nodes deployed in a remote environment must self-configure without any *a priori* information about the network topology or global view. Nodes will act in response to environmental events and relay collected and possibly aggregated information through the formed multi-hop wireless network in accordance with desired system functionality. The inherently dynamic and distributed behavior of these networks, coupled with inherent physical limitations such as small instruction and data memory, constrained energy resources, short communication radii, and a low bandwidth medium in which to communicate, make developing communication protocols difficult.

Research on hardware for such devices has taken place at Berkeley [14][30][32] and various other research institutions [24] throughout the world. Using such hardware as a basis for development, the software architecture and communication stack residing on these devices has prompted prolific research in the areas of ad-hoc networking [10][15][17][20], data aggregation [16][21][26], cluster formation [25], distributed services [22], group formation [6], channel contention [3][5][7][19], power conservation [4][12], and much more. Work on the utility of such an innovative technology has unearthed potential applications including environmental monitoring, event tracking [1], disaster relief, and search and rescue.

In this work we address the problems of low bandwidth and energy limitations inherent to sensor devices. These networks' ever-changing and unpredictable state demands a self-configuring, adaptive elixir. Our solution is a novel adaptive application independent data aggregation (AIDA) component that fits seamlessly into the current sensor network communication stack. Our goal is to maximize utilization of the communication channel (single frequency) with energy savings coming as an ancillary benefit. With significant costs incurred from channel contention, packet header overhead, and data padding for fixed sized packets, this work abates such costs by employing varying degrees of data aggregation at forwarding nodes in accordance with current local traffic patterns.

Data aggregation techniques have been extensively investigated in recent literature. Our work, as a novel data aggregation approach, distinguishes itself from current state of the art solutions in three respects. First, all prior Application Dependent Data Aggregation (ADDA shown in Figure 1b) relies on application layer information and must have a bi-directional interface, and therefore dependence with, the data centric routing protocol implemented. AIDA isolates aggregation decisions from application specifics by performing adaptive aggregation in an intermediate layer that resides between the traditional data-link and network layer protocols (Figure 1.a). AIDA is totally transparent to other layers and can be swapped into the network stack without modifying any existing interface. Second, no prior work in data aggregation adapts itself to the traffic situation in a time sensitive manner. AIDA takes the timely delivery of messages as well as protocol overhead into account to adaptively adjust aggregation strategies in accordance with assessed traffic conditions and expected sensor network requirements. Simulation results show that AIDA can adapt to varying traffic situations and dramatically reduce network congestion and transmission energy consumption. Third, previous data aggregation schemes (e.g., data centric routing [16]) perform in-network processing to reduce the amount of application data transmitted. These in-network processes (e.g. averaging) can achieve higher degrees of aggregation; however data are less available to the application (e.g. standard deviation of the data set can not obtained from the average). In contrast, AIDA performs loss-less aggregation allowing the upper layer to decide whether information compression is appropriate at the time. Very important, our design enables AIDA to remain complementary to other data aggregation strategies (Figure 1.c) while providing significant timesaving benefits in the lower layers of the communication stack.

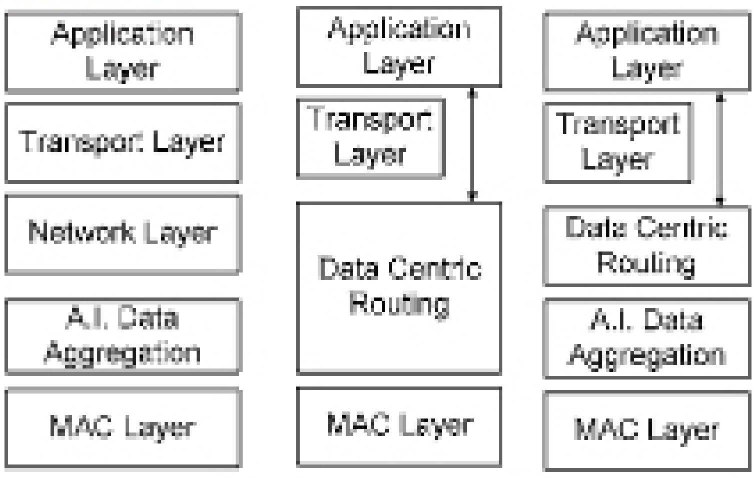

a. AIDA b. ADDA c. Both

Figure 1: Architectural Designs

This paper attempts to address the aforementioned problems through a novel adaptive time sensitive data aggregation component. As an introduction to sensor networks, and to provide a more in depth discussion of the type of research taking place within this field, we begin section 2 with a discussion of related and ongoing work. Section 3 addresses the need for adaptation, data aggregation, and real-time data delivery. Section 4 then presents specific details about our protocol. Sections 5 and 6 describe our simulation environment, the type of experiments run, and a discussion of the results we obtain in both simulation and in the Berkeley MICA test-bed. Finally, we conclude in Section 7.

2. Leveraging Previous Work

Efforts to maximize channel utilization have been spread across various layers of the sensor network communication stack. Starting at the MAC layer, these include attempts to minimize collisions through contention-based mechanisms designed for a lossy wireless medium. Such work includes 802.11 [3], MACA [19], MACAW [5], FAMA[7], S-MAC[29], and Multi-Hop Scheduling [18], to name a few. All of these solutions reside within the data-link layer of the communication stack and, therefore, can coexist with the higher layer aggregation component we provide.

Similar to the data link layer the network layer, and more specifically the routing component, has brought about significant efforts to avoid congestion and maximize use of the communication medium. Such schemes include distributing the traffic load to route around congestion [10] and using a minimal hop path to reduce the total number of transmissions [27]. Beyond the routing layer the communication stack in sensor networks becomes more amorphous. Clustering [25], group formation [6], and other higher layer hierarchical components serve to combine node responsibilities and come to consensus on what data to send. Often such information is application specific and must rely on a general understanding of exactly what the network is tasked to do. Additionally, the hierarchical and grouping components often utilize various forms of data aggregation through consensus algorithms or other forms of local processing.

Basic schemes [16] for the aggregation of data include the Center at the Nearest Source (CNS), where data is aggregated at the source nearest to the destination; Shortest Path Trees (SPT), where data is sent along the shortest path from source to sink and aggregated at common intermediate hops along the way; and Greedy Incremental Trees (GIT), which builds an aggregation tree sequentially to merge paths and provide more aggregation opportunities.

An extremely popular data aggregation scheme for sensor networks, Directed Diffusion [15][11], is a data-centric architecture where named (application specific) data gets propagated along paths back to the requestor. Effective paths are reinforced as they are used to optimize communication from point to point. Specifically designed for sensor networks, Directed Diffusion aggregates data along these reinforced paths to reduce the quantity of data transmitted across the network. Similarly Data Placement [26] is designed for applications where multiple sinks coexist and use in-network caching to update and distribute data to leaf nodes at the minimally requested rate. LEACH [12] is a high layer protocol that provides clustering and local processing to aggregate sensor data and reduce global communication. Many other data aggregation schemes exist that also provide network, transport, and application level mechanisms taking advantage of application specific knowledge about the data in question. All of these schemes reside either at or above the network layer and are orthogonal and can coexist with our work.

Aggregation scheme comparison studies have demonstrated the effect of network parameters and the utility of aggregation mechanisms in a wide variety of applications [16][21]. These studies discuss potential savings that aggregation can provide and are noted to explicate the potential for such work to improve network throughput.

To date, very few sensor network papers have addressed the need for incorporating adaptive behavior into their protocols. Sensor networks exhibit complex distributed behavior rendering static pre-configuration utterly useless as network traffic, often initiated by environmental events of interest, transitions from one extreme to another. Several protocols have taken a first stab at addressing the need for adaptive behavior in such dynamic networks. RAP [23] and SPEED [10] utilize locally available information to adjust priority levels or make more informed routing decisions in response to network congestion and changing traffic patterns. SPIN [13] makes adaptive decisions to participate in data dissemination based on current energy levels and the cost of communication. [30] uses adaptive rate control at the data-link layer to fine tune contention parameters in response to local traffic conditions. GAF [28] monitors network connectivity and turns nodes on/off to adapt network density for energy-conservation. While many more examples of online adaptation exist, these solutions provide relevant examples of how adaptation is beneficial in dynamic and unpredictable sensor networks and serve as a starting point to introduce adaptive behavior into these complex systems.

In addition to maximizing channel utilization and adapting to dynamic network conditions, energy conservation has become a central focus in sensor network research. Similar to data aggregation, work in energy conservation for sensor networks has been considered at various levels of the communication stack. Aside from minimizing power consumption at the hardware level [24], MAC layer protocols developed for energy savings mostly take advantage of overhearing and scheduling to allow nodes to sleep while they are not transmitting or receiving messages [8][12][27]. At the network and routing layer, schemes work to minimize power along the transmission path [26], set routes according to the energy remaining at nodes along that path [31], and use mechanisms to save power through the distribution of messages among various paths from source to destination [10]. Finally, higher layer protocols that often incorporate routing semantics exist to form groups and rotate leadership responsibilities allowing non-leader nodes to sleep and conserve their energy [4]. Again all of these protocols involve layered decisions that should adhere to strict modular programming interfaces allowing our work to coexist with them.