Transport Services for Energy Constrained Environments

N. Riga[†], A. Medina[‡], I. Matta[†], C. Partridge[‡], J. Redi[‡], I. Castineyra[‡]
[‡]BBN Technologies, [†]Boston University

ABSTRACT

JAVeLEN (Joint Architecture Vision for Low Energy Networking) is a network architecture whose design targets the reduction of the energy-per-bit used for data delivery in tactical wireless mobile ad-hoc networks (MANETs). It comprises the physical, MAC, routing, and transport layers of the communication stack. In this extended abstract we briefly summarize our work in progress on the design of JTP, the JAVeLEN Transport Protocol. The central question of our JTP research is, given a network-wide energy efficiency objective, how should a transport protocol be designed so that such objective is achieved while taking into account application semantics. JTP achieves that goal by exploiting reliability semantics weaker than those offered by TCP when applications tolerate it. JTP incorporates as well additional QoS provisions for applications.

1. INTRODUCTION

In the context of MANETs, maximum energy gains would be achieved by an overall network architecture design targeting energy efficiency as a common optimization target at all layers. In this abstract we briefly describe our work in progress on the design and implementation of flexible transport services for energy-constrained environments.

In the context of energy efficiency, there has been work on satisfying QoS requirements of applications while considering energy usage. This work has targeted mostly routing protocols and strategies [1] as well as support for dynamic resource management (with respect to rate and delay) at the routing and lower layers of the protocol stack [2]. Work on transport protocols for MANETs has been centered mainly around devising enhancements to TCP to enable it to operate in the context of the intrinsic characteristics of MANET environments [3], and some new protocols have been proposed to provide TCP-like reliability semantics over MANETs [5]. To the best of our knowledge, there is no work on designing transport protocols that seek to achieve a higher network-wide energy efficiency by exploiting energy-gain opportunities available by the capability of applications to operate under varying reliability semantics.

In this abstract we describe briefly the JAVeLEN architecture, emphasizing the major design aspects included in each layer; and summarize our ongoing work on the design and development of JTP, the JAVeLEN Transport Protocol, highlighting the main goals set and decisions we have made for the protocol.

2. JAVELEN SUMMARY

MANETs impose substantial communication challenges, such as unpredictable RF attenuation, node mobility, wide range of traffic demands, and very limited energy availability. Most of these challenges have been addressed in isolation by protocol designers in both industry and academia. The problem with such a piece-meal approach is that when these mechanisms and protocols are put together to operate in a single environment, it is likely that contradicting design goals, as well as inefficiencies in some layers will offset the gains obtained at other layers.

BBN has developed JAVeLEN, a Joint Architecture Vision for Low Energy Networking. JAVeLEN encompasses a set of interoperating mechanisms and protocols from the modem up to the transport layer, that collectively target a $10^{2.5}$ factor in energy reduction under varying network size, traffic rates, mobility, and network density. Following is a summary of the JAVeLEN layers.

Physical Layer: Delivers bits between adjacent nodes. Encompasses the use of three physical waveforms, a *low-data rate, energy optimized* Hail waveform; and two *high-data rate time and frequency hopping* waveforms, one optimized for short packets and the other optimized for long packets. JAVeLEN's physical layer design achieves energy gains of $10 - 100$ times.

Point-to-Point (P2P): Delivers packets between adjacent nodes and manages idle times. Encompasses the exploitation of each physical waveform given their best operating points; it employs pseudo-random codes for implementing uncorrelated wakeup strategies; and it performs link and neighbor establishment based on the characteristics of the current RF environment. JAVeLEN's P2P layer design achieves energy gains of $100 - 200$ times.

Path Management: Determines routes in the network topology using energy-driven link cost metrics. Encompasses the use of *Hazy-sighted scoping* to control the dissemination of routing information; it employs *energy conserving multipoint relaying* using knowledge of transmission power at nodes to build a connected dominating set for route information dissemination; route generation is energy-sensitive, and battery-aware link biasing is employed to send packets along the paths that require the minimum amount of power for their delivery. JAVeLEN's routing layer design achieves energy gains of $10 - 100$ times.

Transport: This layer leverages the rest of the JAVeLEN architecture in support of applications. Its design is summarized in the next section.

3. JTP OVERVIEW

The main question we seek to answer in our research is: for a given optimization metric, how should transport services be designed so that every aspect of the corresponding protocol is influenced by the need of improved performance with respect to m. In the context of JAVeLEN, the metric m correspond to *energy efficiency*. Therefore, every design decision for the architecture has the underlying goal of improving that metric. In a network architecture, some layers lend themselves better for direct improvements of a given performance metric. For example, at the physical layer, the design of the radio components can be focused on aggressively reducing the energy-per-bit requirements. Similarly, at the routing layer, the design can be focused on the reduction in the number and frequency of routing information messages that are sent across the network. These design tradeoffs will translate into overall energy gains only if the upper layer protocols are capable of leveraging their benefits while satisfying the semantic requirements of applications. Difficulties arise, however, when the optimization metric we seek to optimize (e.g. energy efficiency) may conflict with the optimization requirements of applications (e.g. reliability). The challenge we face is then that of designing a transport protocol that carefully accounts for energy efficiency, while still satisfying application semantics.

The challenge above can be stated as an optimization problem to maximize a gain function such as gain $= \frac{1}{E} \times G$, where E is the network-wide energy consumed to transfer certain amount of data, and G is the amount of data that effectively reaches the destination. JTP seeks to strike a proper balance between the goodput requirements of an application and the energy well-being of the underlying target network. Following we briefly describe the principles guiding JTP's design. For each design principle, we list the corresponding design choices/mechanisms included in JTP.

Flexibility. In a MANET, energy savings should be enabled whenever they are possible. Furthermore, the reliability semantics of different applications are different. Therefore, transport services should support applications with different QoS requirements while leveraging the energy savings achieved by lower layers. Mechanisms: extensible packetization modules; interface between applications and JTP to express QoS requirements and mapping of application data units (ADU) to JTP data units (JDU).

Robustness. Transport services for a MANET environment should achieve energy savings by adapting efficiently to the often challenged conditions of MANETs. Mechanisms: point-to-point operations; local re-routing; mid-path node caching.

Self-tuning properties. Given the characteristics of the underlying network infrastructure (e.g. mobility, fading, etc.), JTP actions and decisions will likely be challenged continuously. Therefore, JTP design choices must be *self-tuning* to discrepancies and/or inefficiencies caused by sudden changes in network conditions. Mechanisms: epoch-based data transfers; QoS assurance without per-flow state.

Cross-layer optimizations. Since energy efficiency percolates all layers in a MANET architecture, the design of JTP should allow for cross-layer optimizations wherever and whenever they facilitate the achievement of energy gains. This aspect must be approached with enough care to prevent the extreme design choice of having monolithic network design. Mechanisms: explicit feedback notification from MAC to JTP; explicit retransmission control from JTP to MAC; gathering of routing information from the routing layer to aid in local re-routing.

Some of the mechanisms mentioned in the above paragraphs include, among others (the list is not complete), the following functionalities:

Adjustable reliability semantics. Since not all applications require the same level of reliability, JTP enables them to express their requirements so that decisions made during data transfers can be influenced by the actual benefit/cost tradeoffs. For example, if an application is able to tolerate a 20% loss rate, JTP will use such information to, among other things, reduce the energy consumed in retransmissions.

Receiver-controlled retransmissions. In addition to enabling applications to express their QoS requirements, JTP also approaches the reliability semantics of applications by allowing the receiver side of the application to decide which packets should be retransmitted by the sender at any given time.

Epoch-based data transfers. JTP conceptually divides a data transfer into packet epochs. Protocol parameters are adjusted after each epoch of packets has been transferred. In addition, packets carry dynamic state enabling them to continuously gather network information along the path they traverse and communicate them to the receiver side. Based on such information, the receiver can adjust transfer parameters at any time, effectively starting a new packet epoch. This mechanism enables JTP to be self-tuning.

Feedback minimization. The previous design choices are combined to enable JTP to significantly reduce the feedback required to satisfy a given level of reliability. Specifically, only negative selective acknowledgments are sent by the receiver side of a transfer.

Energy-expenditure control. While JTP must target the satisfaction of the reliability requirements of applications, it must adapt to the changing nature of a MANET environment. JTP will take into account the QoS requirements expressed by applications and at any moment, it will establish the proper tradeoffs between trying to deliver a given packet to the destination, given the packet's energy expenditure level, and dropping the packet and possibly inducing an end-to-end retransmission. The mechanisms involved in this aspect of JTP include the computation of a per-packet *energy budget* which will be carried in packet headers, and used by JTP at mid-path nodes to control energy expenditure when network conditions substantially degrade, or when the reliability semantics of applications provide flexibility for JTP to use only the energy resources required to satisfy the specific requirements imposed by such semantics.

Mid-path transport services for end-to-end services. While JTP implements end-to-end services, its design encompasses point-to-point interactions of JTP modules along the path of a data transfer. Such interactions make JTP robust to the changing nature of a MANET, and enable energy efficiency gains by means of local re-routing and caching mechanisms that are part of JTP. These mechanisms require interactions between JTP and the MAC layer, as well as between JTP and routing—interactions normally not available to the transport layer.

Certain QoS assurance. JTP's design encompasses mechanisms to ensure certain level of service assurance with respect to delay and bandwidth, within the constraints imposed by a MANET environment. Specifically, JTP design leverages and extends techniques designed for the Internet to provide QoS guarantees without per-flow state [4].

4. REFERENCES

[1] I. Jawhar and J. Wu. Quality of Service Routing in Mobile Ad Hoc Networks. *Resource Management in Wireless Networking (Springer)*, 16:365–400, 2005.

[2] S-B. Lee, A. Gahng-Seop, X. Zhang, and A.T. Campbell. INSIGNIA: An IP-Based Quality of Service Framework for Mobile Ad Hoc Networks. *Journal of Parallel and Distributed Computing*, 60(4):374–406, April 2000.

[3] J. Liu and S. Singh. ATCP: TCP for Mobile Ad Hoc Networks. *IEEE Journal on Selected Areas in Communications*, 19(7):1300–1315, 2001.

[4] Ion Stoica and Hui Zhang. Providing Guaranteed Services Without Per Flow management. In *Proceedings of ACM SIGCOMM*, pages 81–94, Cambridge, Massachusetts, 1999.

[5] K. Sundaresan, V. Anantharaman, H-Y. Hsieh, and R. Sivakumar. ATP: A Reliable Transport Protocol for Ad-hoc Networks. In *Proceedings of ACM MobiHoc*, Annapolis, ML, 2003.