Reliability in Distributed TCP Caching

Thiemo Voigt, Adam Dunkels, Juan Alonso
{thiemo, adam, alonso}@sics.se

Abstract: Distributed TCP Caching (DTC) substantially improves TCP performance in wireless sensor networks. We study DTC performance gains with two link layer error control mechanisms, namely link layer retransmissions and forward error control. Our results suggest that applying forward error correction to TCP acknowledgements is the best choice for our ongoing real implementation of DTC.

1 Introduction

Most wireless sensor networks (WSN) do not operate in isolation but are connected to an external network through which monitoring and controlling entities can reach the sensor network. By running TCP/IP in the sensor network it is possible to directly connect the sensor network with a wired network infrastructure, without proxies or middle-boxes. UDP could be used for transporting sensor data, TCP for administrative tasks that require reliability and compatibility with existing application protocols such as configuration and monitoring of individual sensor nodes, downloads of binary code and data aggregation descriptions to sensor nodes. In particular, downloading code to designated nodes, such as cluster heads in a certain geographical region requires a reliable unicast protocol.

We have previously presented solutions that make TCP/IP viable in sensor networks [DVA04]. One of them called Distributed TCP Caching (DTC) greatly enhances TCP performance in WSNs both in terms of energy efficiency and throughput [DAVR04]. In this paper, we study the performance gains for DTC of two link layer error control mechanisms, namely forward error correction (FEC) and local link layer retransmissions. Our results indicate that using FEC is more efficient than local link layer retransmissions. In particular, applying FEC to TCP acknowledgements seems a vital candidate for our ongoing real DTC implementation. While others have previously discussed the use of error control mechanisms for sensor networks [KLN03, JE02] for sensor data, we investigate the use of error control for an enhanced TCP suitable for WSNs.

2 Distributed TCP Caching

To make TCP/IP usable in WSNs, some problems must be tackled. Our solutions [DVA04] include *spatial IP address assignment* that assigns sensor nodes semi-unique IP addresses that reveal location information and e.g. allow to address sensor nodes in a geographical region. To minimize the header overhead we propose *shared context header compression*. We use a specific form of *application overlay networks* to implement data-centric routing and data aggregation for TCP/IP sensor networks. The end-to-end acknowledgment and retransmission scheme employed by TCP causes expensive retransmissions along every hop of the path between the sender and the receiver if a packet is lost. The aim of *Distributed TCP Caching (DTC)* is to improve TCP performance in WSNs both in term of energy efficiency and throughput.

DTC's key idea is to avoid end-to-end retransmissions by caching TCP segments in the network and by performing local retransmissions. Since sensor nodes are generally resource constrained, we assume that each node can only cache one segment. Each node caches one unacknowledged TCP segment and takes extra care to cache a segment that presumably has not been received successfully by the next node. We use feedback from a link layer that uses positive acknowledgments to infer packet drops on the next-hop. Our design works with either active or passive acknowledgements [JT87]. In our simulations we use active link layer acknowledgements. A TCP segment that is forwarded but for which no link layer acknowledgment is received may have been lost in transit. Therefore, the segment is *locked* in the cache indicating that it should not be overwritten by a TCP segment with a higher sequence number. A locked segment is removed from the cache only when an acknowledgment that acknowledges the cached segment is received, or when the segment times out. To avoid end-to-end retransmissions, DTC needs to respond faster to packet drops than regular TCP. DTC uses ordinary TCP mechanisms to detect packet loss: timeouts, duplicate and selective acknowledgments. Every node participating in DTC maintains a soft TCP state for connections that pass through the node

Previous simulations have shown that this mechanism greatly improves TCP performance in wireless sensor networks in several ways [DAVR04]:

- DTC substantially reduces the overall number of TCP segment transmissions.
- DTC decreases the number of end-to-end retransmissions.
- DTC shifts the burden of the energy consumption from nodes close to the base station into the network.

The last point is important since nodes close to the base station usually are the first to run out of energy [MRK+04].

3 Evaluation

In this section, we evaluate two different schemes for improved reliability, namely link layer retransmissions and FEC. Our simulations are done using the same environment as described previously [DAVR04]. Data transfers contain 500 segments. We use a chain topology where node n is in transmission range of node $n-1$ and $n+1$, but node $n-1$ is not in range of node $n+1$. We run the simulations until the sender receives the acknowledgement for the 500th packet. Since TCP data segments are larger than acknowledgements, we set the packet loss probability for data segments to twice the loss probability for TCP acknowledgements and to four times the loss probability of link level acknowledgements. Our simulation consists of 30 runs, the results denote the average. If not mentioned otherwise, our chain consists of 11 hops and the per-hop packet loss probability is 10% for TCP payload segments, thus 5 % for TCP and 2.5% for link layer acknowledgements.

Wireless communication is often the major power consumer during sensor operation [RSPS02]. Therefore, the number of packet transmissions is one indicator of the energy efficiency of a network protocol given a MAC layer that schedules packets accordingly. To compare the energy consumption of the different reliability schemes we compute the energy consumption for transmitting packets.

3.1 Reliability using Link Layer Retransmissions

In this section we evaluate if a more reliable link layer that allows for retransmissions of TCP segments can improve DTC performance. For our simulations we assume the buffer model shown in Figure 1.

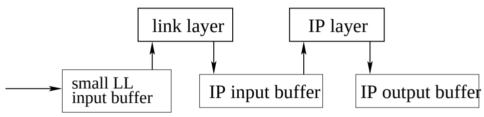

Figure 1: Buffer model for retransmissions

In this model, a transmitted TCP segment resides in the output buffer until it is acknowledged by the next hop or until the link level retransmission timer expires. In the latter case, the TCP segment is either retransmitted when another link layer retransmission is allowed or removed from the output buffer. TCP segments that arrive when the output buffer is occupied, are placed in the IP input buffer after a link level acknowledgement has been transmitted. If the input buffer is also occupied, an arriving TCP segment must be dropped. Link level acknowledgements can always be processed. We assume a maximum of one link layer retransmission per hop for each packet.

retransmission strategy	transmitted PL	transmitted Acks	energy consumption (mAs)
no link level retr.	7795	5596	3633
link level retr. Acks	7182	6917	3550
link level retr. payload	7945	5854	3720
link level retr. Acks and payload	6883	6416	3357

Table 1: Different retransmission strategies

Our simulation results for the different retransmission strategies are shown in Table 1. The average energy consumption Ec for data transfers is computed as $Ec = te(tPL(st + PLtime) + tAcks(st + Atime))$. We set the start-time st of the radio to 1 ms. $PLtime$ is the time to transmit one payload packet of size Bytes, $Atime$ the time to transmit a 40 Byte acknowledgement assuming 19.2 Kbit/s data rate. The transmit power consumption te of our ESB nodes [CS] is about 6.5 mA. tPL and $tAcks$ are the numbers of transmitted TCP payload and acknowledgements taken from the table above. Our results show that link layer retransmissions reduce the energy consumption for most retransmission strategies. However, performing link layer retransmissions of payload segments only increases power consumption. There are two major reasons for this: First, unnecessary retransmissions of payload packets when the link layer acknowledgement is lost. Second, some packets must be dropped due to full buffers at the next hop.

3.2 Reliability using FEC

In the next simulations we use FEC to increase reliability. We assume a FEC coding scheme that increases the number of bytes to be transmitted with 40%. For example, using Hamming code with data lengths of 8 Bits leads to 50% overhead, while the overhead for data lengths of 16 Bits is 31%.

TCP Ack loss probability	transmitted PL	transmitted Acks	Energy consumption (mAs)
5%	7795	5596	3873
3,5%	7089	5370	3567
2%	6637	5290	3381
1%	6382	5289	3283
0%	6153	5264	3191

Table 2: Improvement with FEC for acknowledgements

In the first experiment, we apply FEC to TCP acknowledgements. In all simulations, the packet loss rate for TCP payload segments is fixed to 10%. The results in Table 2 show that with a reduction of the acknowledgement loss rate from 5% to 2%, FEC for TCP acknowledgement leads to about the same energy consumption as link layer retransmissions for payload and acknowledgment. Interesting is that given all acknowledgments arrive, the number of transmitted payload packets is almost the theoretical optimum which is 6111 for 11 hops with 10% payload packet loss rate [DAVR04].

The benefit of applying FEC to TCP payload segments (see Table 3) is not as high. This is due to the increased packet length. We have also simulations applying FEC to both payload segments and acknowledgements. However, the benefit is not as high as for applying FEC to TCP acknowledgements. For example, assuming FEC reduces packet loss rate to 2.5% for payload segments and 1% for TCP acknowledgements leads to an energy consumption of 3404 mAs. Another argument for applying FEC to TCP acknowledgements is that we assume a routing scheme that finds stable routes with low loss rates, e.g. [WTC03]. However, since we assume symmetric routes we cannot take for granted that the path in the direction that acknowledgements travel is a good as the payload path.

Payload segments loss probability	transmitted PL	transmitted Acks	Energy consumption (mAs)
10%	7795	5596	4828
7.5%	7201	5312	4477
5%	6143	4637	3831
2.5%	5797	4473	3629
1%	5582	4363	3499
0%	5248	4146	3295

Table 3: Improvement with FEC for payload segments

4 Conclusions

In this paper we have investigated the performance implications of link layer retransmissions and FEC for Distributed TCP caching. Our results suggest that applying FEC to TCP acknowledgements is the best candidate for our ongoing real DTC implementation.

References

[CS] CST Group at FU Berlin. Embedded Sensor Board. Web page. Visited 2003-10-21.

[DAVR04] Dunkels, A., Alonso, J., Voigt, T., und Ritter, H.: Distributed TCP Caching for Wireless Sensor Networks. Technical Report 6. Swedish Institute of Computer Science. 2004.

[DVA04] Dunkels, A., Voigt, T., und Alonso, J.: Making TCP/IP Viable for Wireless Sensor Networks. In: *Work-in-Progress Session of the first European Workshop on Wireless Sensor Networks (EWSN 2004)*. Berlin, Germany. January 2004.

[JE02] Jeong, J. und Ee, C. T.: Forward Error Correction in Sensor Networks. Technical report. University of California, Berkeley. May 2002.

[JT87] Jubin, J. und Tornow, J.: The darpa packet radio network protocols. *Proceedings of the IEEE*. 75(1):21–32. January 1987.

[KLN03] Karl, H., Löbbers, M., und Nieberg, T.: A data aggregation framework for wireless sensor networks. In: *Dutch Technology Foundation ProRISC Workshop on Circuits, Systems and Signal Processing*. November 2003.

[MRK+04] Mhatre, V., Rosenberg, C., Kofman, D., Mazumdar, R., und Shroff, N.: Design of surveillance sensor grids with a lifetime constraint. In: *First European Workshop on Wireless Sensor Networks (EWSN 2004)*. Berlin, Germany. January 2004.

[RSPS02] Raghunathan, V., Schurgers, C., Park, S., und Srivastava, M.: Energy aware wireless microsensor networks. *IEEE Signal Processing Magazine*. 19(2):40–50. March 2002.

[WTC03] Woo, A., Tong, T., und Culler, D.: Taming the underlying challenges of reliable multihop routing in sensor networks. In: *The First ACM Conference on Embedded Networked Sensor Systems (SenSys 2003)*. Los Angeles, California. November 2003.