A Simulation Study of TCP over the IEEE 802.15.3 MAC

Kwan-Wu Chin and Darryn Lowe
University of Wollongong
Australia, 2522
{kwanwu, darrynl}@uow.edu.au

Abstract

This paper presents the impact of IEEE 802.15.3 MAC's channel time allocation methods on a TCP flow's performance. We show the importance of having super rate and appropriately sized channel time allocations (CTAs).

1. Introduction

The IEEE 802.15.3 medium access control (MAC) is designed to support bandwidth intensive multimedia applications in wireless personal area networks. The main channel access period provided by the IEEE 802.15.3 MAC is called the channel time allocation period (CTAP), which offers a TDMA approach wherein the sending device reserves a channel time allocation (CTA) from the piconet controller (PNC) in advance of its actual transmission [2].

This paper investigates how to allocate CTAs to a TCP flow given the caveat that each CTA is inherently uni-directional. This means TCP will need to be allocated at least two complimentary CTAs (dubbed as Dual-UniCTA) where one is for data packets and the other is for acknowledgement packets (hereafter referred to as sender and receiver CTAs respectively).

Figure 1 shows a TCP flow with both sender and receiver CTAs. Also shown is the growth of the TCP congestion window (signified by the number of packets transmitted) at each superframe cycle. We can see that both CTAs are grossly under utilized because, after sending a *cwnd* of data packets, the sender has to wait until the receiver's CTA before it can obtain the acknowledgment packets it needs before it can queue more data. Then, it will then need to wait again, this time until its own CTA reoccurs in the next superframe before it is able to transmit. This example clearly illustrates the important of the duration and position of CTAs since both these parameters have a significant effect on a flow's congestion window (*cwnd*) growth rate.

2. Simulation and Results

We investigate CTA methods using our *ns-2* IEEE 802.15.3 MAC implementation described in [1] which uses an Ultra Wideband (UWB) physical layer modeling the DS-UWB proposal currently before the IEEE. Our topology consisted of a PNC and two other nodes where we set one of them as the FTP sender and the other the receiver. Other parameters of the simulation are shown in Table 1.

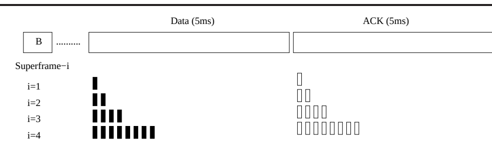

Figure 1. **Example of a TCP flow growing its congestion window in the CTAP.**

Parameters	Values
Superframe length	20ms
Duration of the CAP	1ms
Desired CTA time units	5ms.
TCP flavour	Newreno
Bandwidth	1 Gbps
TCP Packet Size	1460 bytes
Slow-Start Threshold	43 packets
Simulation runtime	10s

Table 1. **Parameters used in simulation studies.**

2.1. Block Sizes and Congestion Window Growth

In this experiment, we are interested in determining the impact of different CTA allocation methods on a flow's *cwnd* growth. Figure 2 shows the speed at which the *cwnd* grows and identifies the benefit of having multiple CTAs within a superframe since each pair of sender and receiver CTAs provides an opportunity for a flow to grow its *cwnd*. Note that, in this experiment, TCP's slow-start threshold is set to 43 and will therefore enter the congestion avoidance phase early in the connection.

Figure 3 shows the critical role played by TCP's maximum congestion window with regard to CTA utilization. Utilization is defined as the number of packets that are actually transmitted over the number of packets that could possibly be transmitted in a given CTA duration. As we increase the window size, utilization increases up to the point where the flow is limited by the overheads associated with transmitting the acknowledgment packets. This is

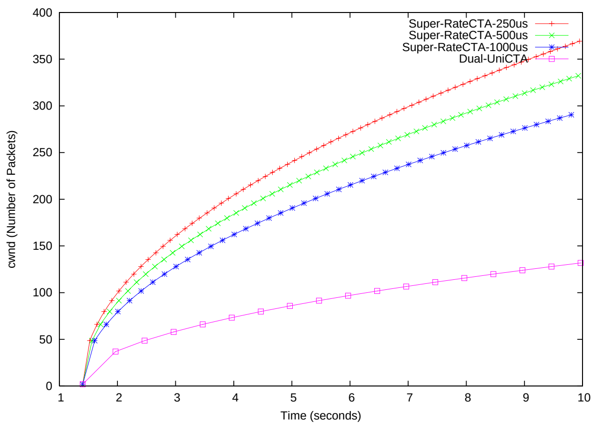

Figure 2. This shows the effect of a 5ms time block being broken up into blocks of 250μs, 500μs, 1000μs and 5000μs. The Imme-ACK policy was used in this experiment.

because each discrete acknowledgment packet/group incurs two SIFS (10μs each), plus the transmission time of the acknowledgment itself, which limits the growth of flows.

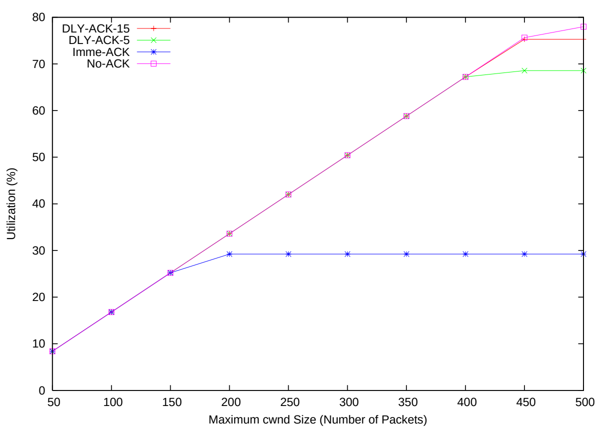

Figure 3. This shows the utilization of CTAs when we increase TCP's maximum congestion window size.

Figure 4 shows the maximum throughput achieved for all the CTA allocation methods considered in our simulation for maximum congestion windows that range from 100 to 1000 packets. We find that when a flow uses the Dly-ACK policy, and a CTA duration of 250μs, the flow's throughput is not affected by TCP's maximum congestion window even when it is set a value as low as 100 packets. In fact, the throughput of this scenario approximates that of one that is free of any acknowledgment overheads. The reason is as follows. If we set the duration of a single CTA to coincide with the time required to transmit a maximum congestion window of packets, say 200μs, then there will be 25 super-rate CTAs, each of a duration of 200μs within a single superframe. Note that the aggregate channel time is still the 5ms that was originally requested. All 25 super-rate CTAs will be fully utilized as each one only requires a small *cwnd*. Therefore, we conclude that having the schedular break a single channel time request into a set of smaller super-rate CTAs avoids both the slow *cwnd* growth problem and the low-utilization problem previously identified in Figures 3 and 4 respectively.

In sum, Figure 4 shows that the throughput of a TCP flow is very sensitive to the configuration of the IEEE 802.15.3 MAC. In other words, while keeping the total air-time constant, TCP throughput can vary from less than 10Mbps to more than 200Mbps. It is therefore critical that designers choose appropriate CTA allocation and MAC acknowledgement strategies.

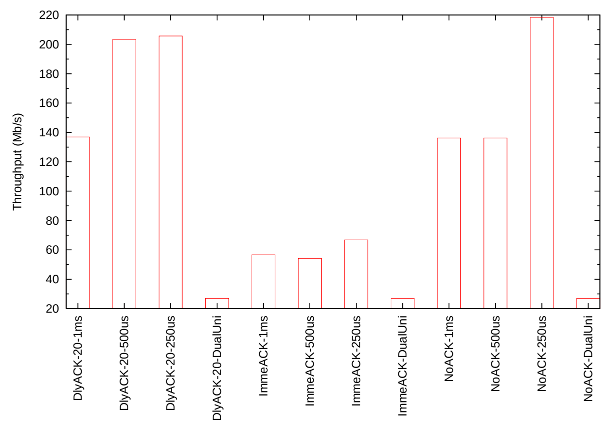

Figure 4. This shows the maximum throughput achieved for different CTA allocation strategies and acknowledgment policies.

3. Conclusion

In this paper, we have looked at channel allocation methods and investigated their suitability for supporting TCP flows. Our results show that super-rate allocation with appropriately sized CTAs enables a TCP flow to quickly achieve high throughput by utilizing CTAs efficiently.

References

[1] K.-W. Chin and D. Lowe. A simulation study of the IEEE 802.15.3 MAC protocol. In *Australian Telecommunications and Network Applications Conference (ATNAC)*, Sydney, Australia, 2002.

[2] IEEE 802.15.3 Working Group. Part 15.3: Wireless medium access control (MAC) and physical layer (PHY) specifications for high rate wireless personal area networks (WPAN). IEEE Draft Standard, Draft P802.15.3/D16, Feb. 2003.