Analysis of Timekeeping Metrics and their Impact on Mobile Ad hoc Network Energy

Divya Ramachandran
August 2003

Abstract

Power efficiency is increasingly becoming an issue in wireless mobile networks. Routing protocols for efficient communication between nodes are complicated and require large amounts of energy, so research areas are now focusing on decreasing this energy consumption. Here, we applied timekeeping metrics to the design of a power model for mobile networks. These metrics had previously been proven to predict the behavior and performance of routing protocols. Using NS-2, multiple simulations were run to gather energy data for different mobility scenarios. Even after breaking the energy down into components, no direct correlation was found between the metrics and energy consumption. However, some trends suggest that indirectly, with the incorporation of a few more factors, the metrics can be applied to a model that accurately predicts the power consumption in a network.

Introduction & Background

As mobile devices become prevalent in our day-to-day lives, more researchers are focusing on efficient network communication. Wireless nodes in these networks are scattered across wide ranges, and must find ways to communicate with one another as well as a central host. In more recent studies, the nodes as well as the host are all mobile. Routing protocols such as Dynamic Source Routing (DSR) and Ad hoc On-Demand Distance Vector Routing (AODV) are just two of many algorithms designed for time-efficient dynamic route computation between nodes. But it is now becoming increasingly important to concentrate on power efficiency. With such complicated algorithms designed for higher success rates, the energy consumption also begins to rise. It is vital that the energy usage stays low to ensure that the nodes stay in operation for a long duration of time. In order to approach this problem, a thorough analysis is necessary to pinpoint different factors in an active network that might affect the power consumption. For the past two months, I have been working with Professor Li-Shiuan Peh at Princeton University on a project to analyze the power consumption in a mobile sensor network.

Previously, Professor Peh and her graduate student Yong Wang worked on a model of the successful packet delivery rate in mobile sensor networks. They identified some key timekeeping metrics that described the behavior of the routing protocols. With these metrics, they designed a model that served to analyze and predict the success rate of a protocol in a specified scenario. In approaching the power consumption problem, I looked for correlations between these same metrics and energy dissipation. There are three components of power in a mobile network: transmit, receive, and idle power. These values were set to .660, .395 and .035 W respectively, which are the default values in the *Network Simulator* software further detailed in the next section. After ascertaining some data of the overall energy consumption, I broke down the values into these three energy components. After more analysis, I examined the energy consumption differences between data packets and control packets. Further explanations of these simulations and their results are found in the next section.

Experiments & Analysis

All of the simulations were run using the *Network Simulator (ns-2)*. Among a variety of functions, *ns-2* provides support for simulating different routing protocols over wireless networks. All my simulations were run using DSR. The energy model in *ns-2* is a simple one that decrements each of the three energy components after corresponding events. For example, after a packet has been transmitted by a particular node, the transmit energy of that node is decremented by the energy dissipated. This energy value is the product of the specified power (.660 W) and the time taken to transmit the packet.

My initial simulations consisted of sixteen different scenarios. In the mobility model, the nodes move to a destination at a random speed between 0 and a specified maximum speed. The node then stays at that destination for a length of time, called the pause time. I used maximum speeds of 1, 5, 10 and 20, and pause times of 0, 10, 20 and 100; pairs of these two variables made up the sixteen scenarios. The simulation area covered 1500mX1500m containing 30 nodes, and the network lifetime was 900s. I first gathered data to see the correlation between these two variables and the total energy dissipated. These results can be seen in Figure 1. We analyzed the data as follows: for the lower pause times, there is a steady increase in energy consumption in accordance with the increasing maximum speed. As the speed increases, more links are breaking between nodes. This calls for more route computations; the nodes are also covering more distance, so as they move into the range of other nodes, they receive more broadcast packets. In the cases with higher pause times, there is a decrease in the energy consumption when the maximum speed reaches 10 and 20. At this point, the nodes are covering long distances, and may remain out of range for a longer period of time. In this case, there is less traffic between nodes.

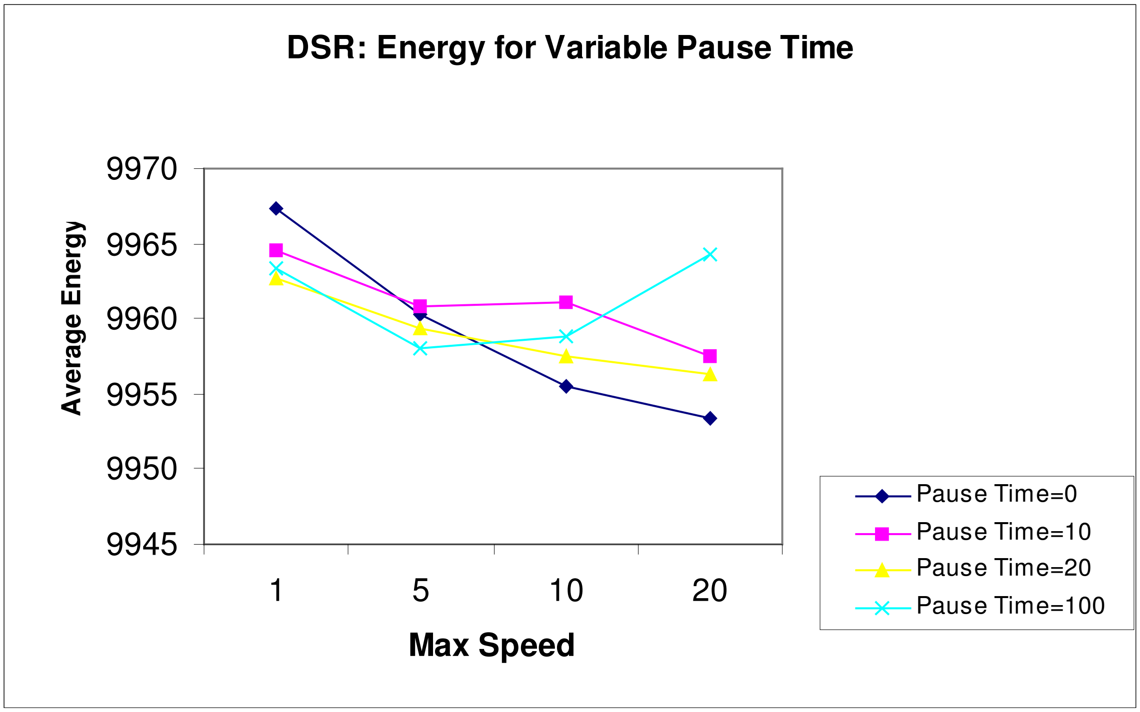

Figure 1

In order to further test these hypotheses, I added some variables in the *ns-2* code to track the separate transmit, receive and idle energy values. The results from these cases are seen in Figure 2. These graphs show the amount of energy dissipated. Surprisingly, the idle energy values remained nearly constant across all of the cases, and covered the bulk of the energy. It could be that the packet traffic was relatively low, and required little communication between nodes. I did not further investigate this issue; rather, I concentrated on the receive and transmit energy components. The receive energy was significantly higher than the transmit. This might be explained by two reasons: DSR uses broadcast messages for route computations. Often, a node receives the same message from multiple nodes, but only transmits it once. Furthermore, nodes listen for packets and information that might be useful to them. They forward only the control packets, and drop any data packets. In this case again, there may be a significantly larger number of received packets than sent packets.

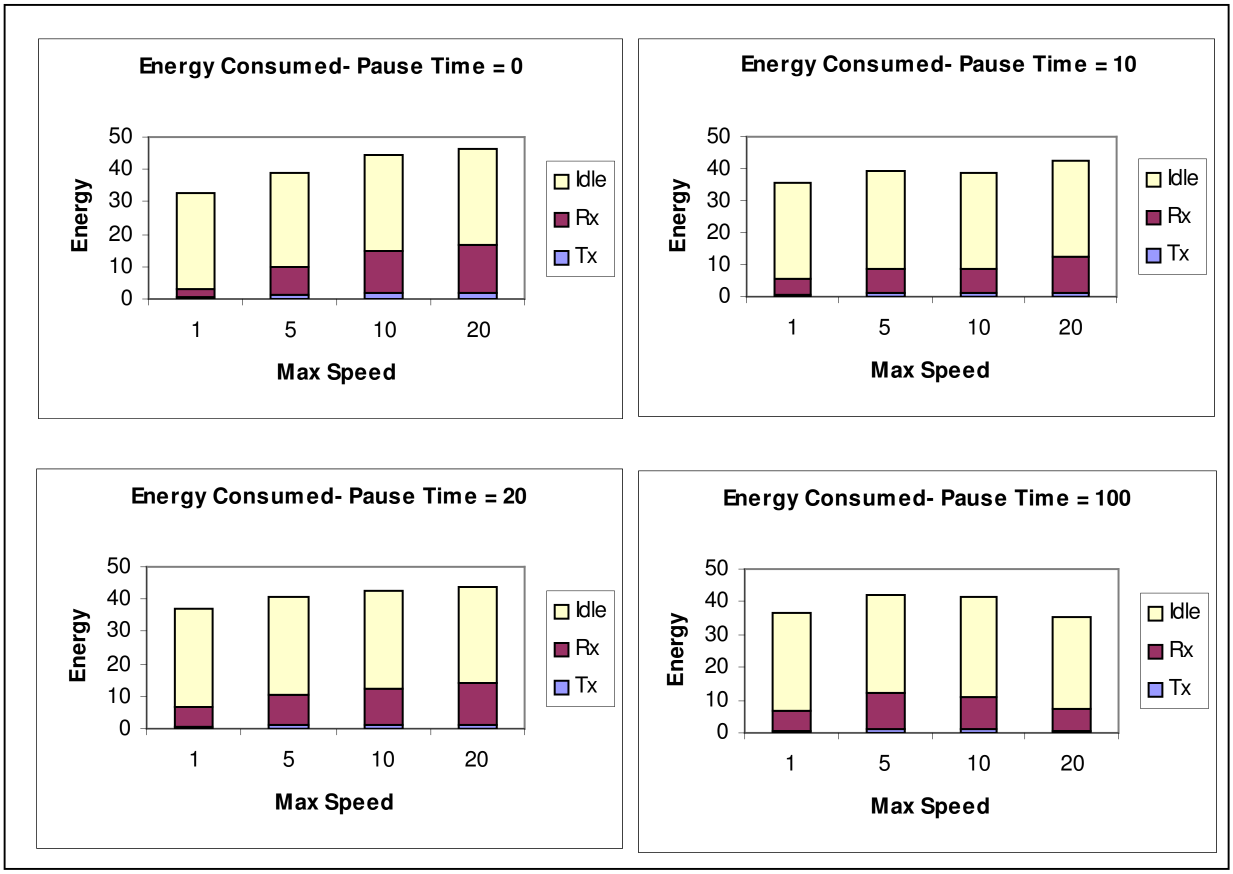

Figure 2

With this basic information available, I approached the problem of applying the timekeeping metrics that were previously correlated to success rate, to power consumption. The three metrics used were Single Path Duration (SPD), Multiple Path Duration (MPD) and Multiple Path Interval (MPI). SPD is defined as the minimum duration of a single link along a path. MPD is the period of time during which at least one path exists between any pair of nodes. MPI is the average period of time during which no path exists between pairs of nodes in the network. Yong provided me with values of SPD, MPD and MPI that correspond to the maximum speeds of 1, 10 and 20, and pause times of 0, 10, 20 and 100. Therefore, all further data are plotted against the twelve combinations of these values.

Hypothetically, these metrics should have some correlation with the power consumption in a network. A shorter SPD would mean that links break often and routes must be frequently recomputed. This would result in higher control packet traffic, and therefore higher power consumption. A similar correlation would be expected for MPD. MPI should probably have somewhat opposite results, as longer intervals between paths would mean more communication is necessary to deliver packets. The results, however, shown in Figure 3, did not exhibit any of the expected trends. Their lack of any sort of consistent pattern led us to believe we should further divide the data. A short SPD might require more control packets but less data packets, and a longer SPD might allow for more successful data packet deliveries with fewer control packets. Measuring these two conflicting events together might skew the data. The same would be the case for MPD, and with MPI, the longer intervals would correspond to more control packets and less data packets.

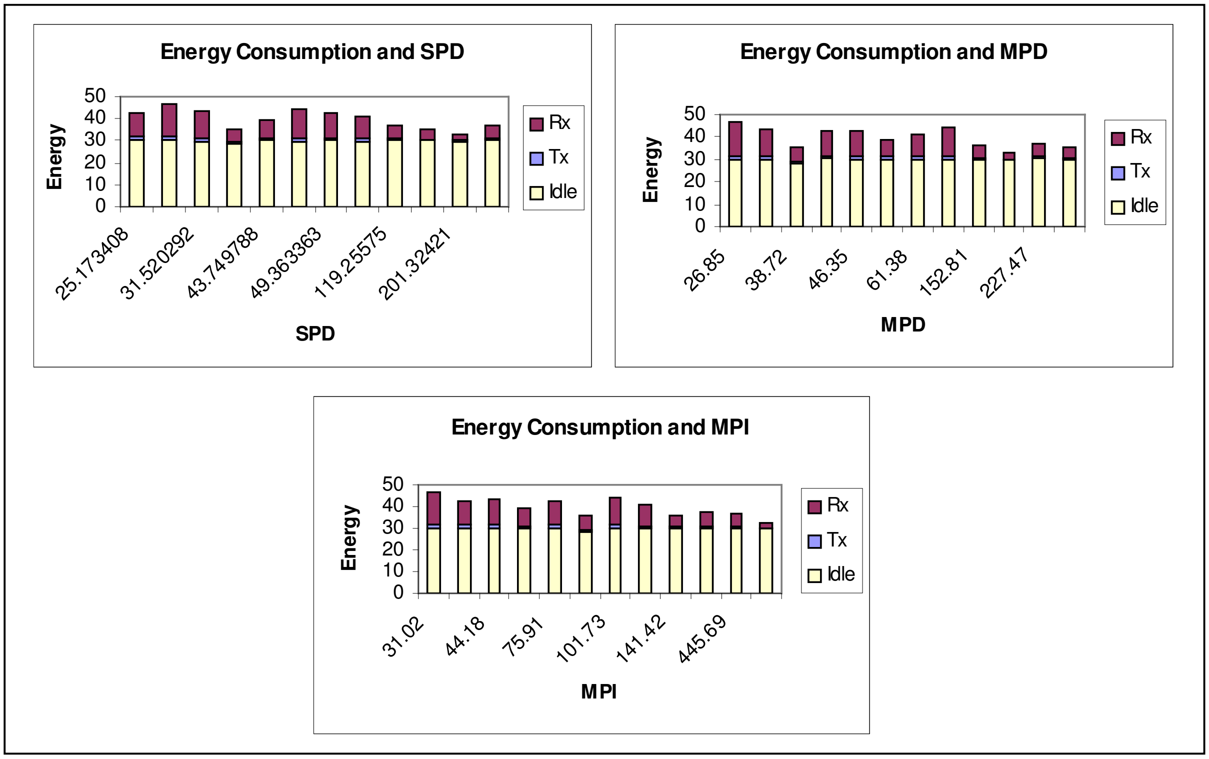

Figure 3

Figures 4a, 4b and 4c show the transmit and receive energy data divided into control and data packets for each of the three metrics. I derived these values by post-processing the trace files using a gawk script. We expected to see opposite trends in the energy consumed by data and control packets. But no direct correlation is apparent from this data.

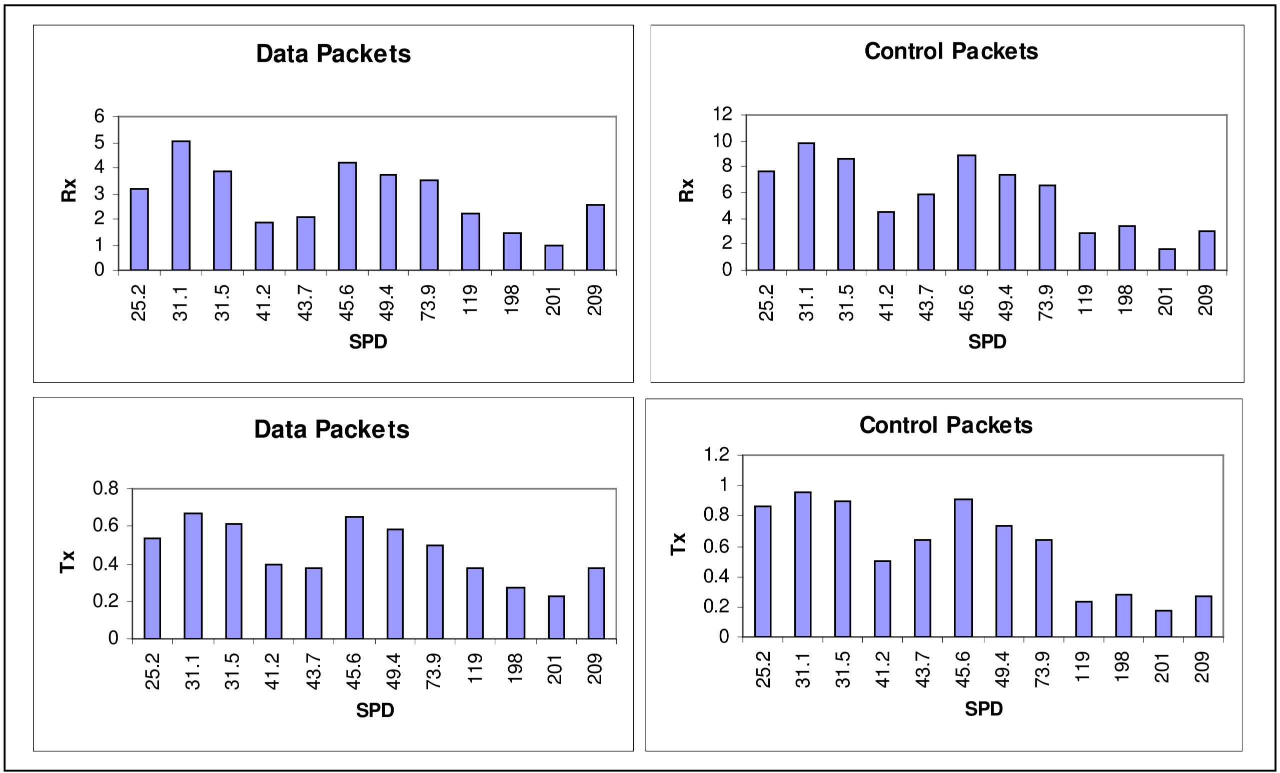

Figure 4a

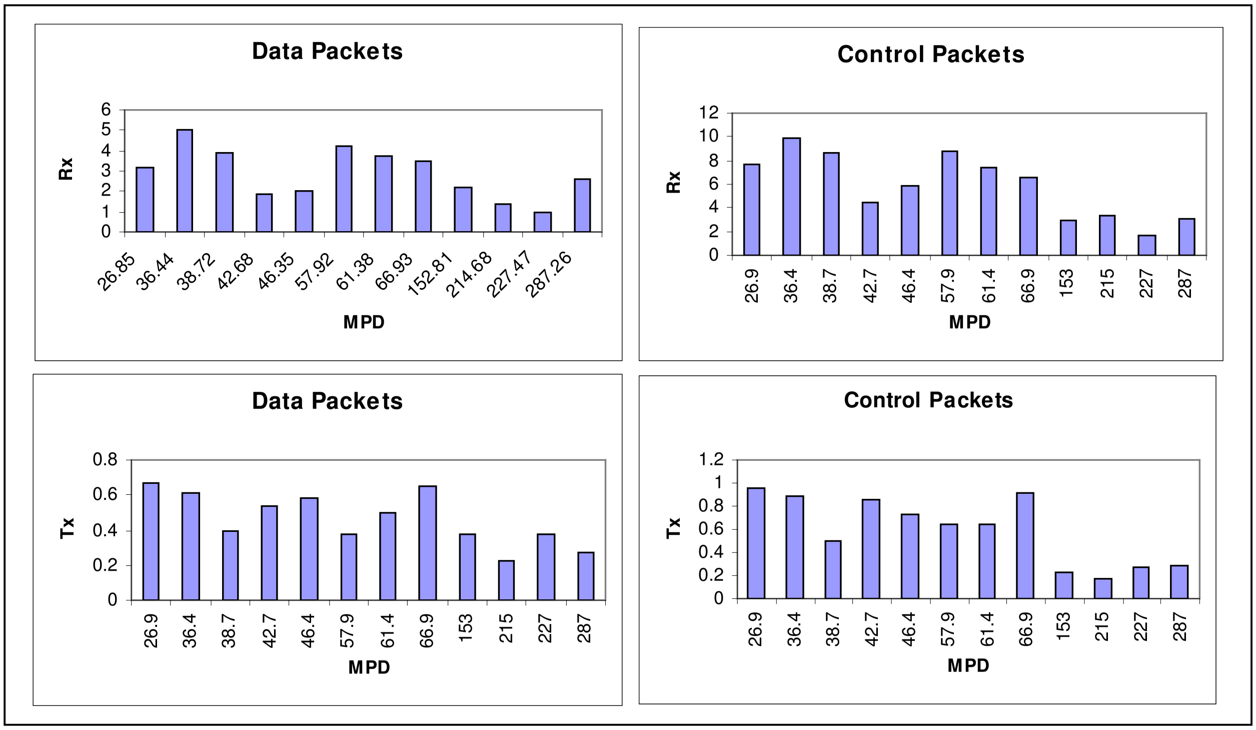

Figure 4b

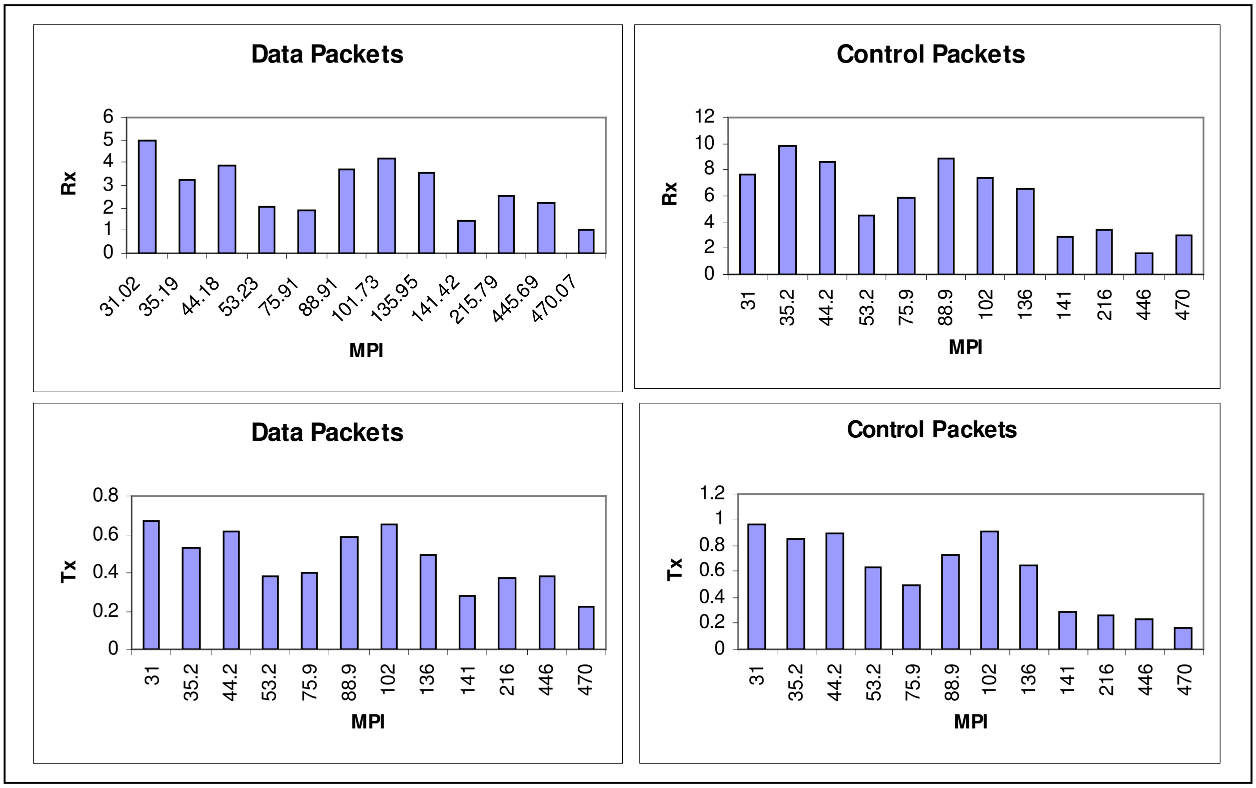

Figure 4c

I further analyzed the SPD metric to determine whether or not energy was entirely independent of these timekeeping metrics. According to the previous hypothesis, the energy trends expected were based on the assumption that at different values of SPD, communication between nodes would vary. This implied that the number of packets sent and received would increase and decrease according to the SPD of a given scenario. I gathered data for the corresponding number of packets being delivered for the energy consumed in Figure 4a. This time, the results did correspond logically with the hypothesis, as seen in Figure 5.

The number of data packets transmitted is specified in the script, and is therefore at a constant average of about 70 packets. The number of control packets sent exhibits the desired trend. As SPD values increase, the number of control packets transmitted steadily decreases. Clearly, the existence of a path reduces the need to compute new routes. The received packets show a more vague example of this trend. We can see a very slight decrease in the number of control packets received as the SPD increases; the last group of points is in a lower range than those at the beginning.

When comparing Figures 4a and 5, we would have expected to see similar trends in the corresponding graphs. But one striking inconsistency is that although the number of data packets sent remains constant, the energy consumed in this case is random. A possible explanation for this is that the varying distances traveled by the packets may directly affect the energy consumption. The more number of hops a packet makes, the more energy must be consumed. Therefore, there are likely other factors that, in addition to the timekeeping metrics, contribute to the energy consumption in a network.

Another factor that could also be affecting the number of packets transmitted and received is the number of packets that are dropped. A dropped packet must be retransmitted, and can therefore cause the network to consume more energy. Figure 6 shows the very apparent negative correlation between the number of packets dropped and SPD. As routes last longer, fewer packets are dropped, both for data and control packets. According to our initial hypothesis we would have expected a steady rise in the number of data packets received as routes endured for long periods of time. In Figure 5 we do not see this; however, Figure 6 shows an example of why this might be the case. As the number of packet drops decreases, there are fewer retransmits, and this might counteract the fact that more packets are being successfully sent.

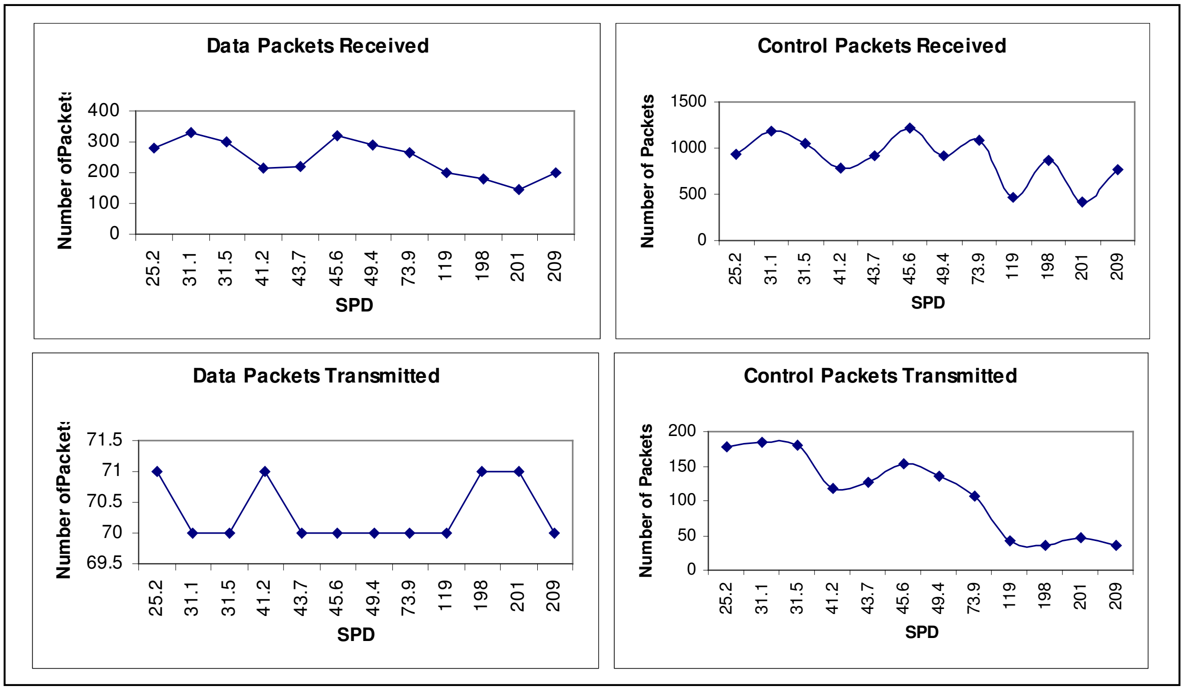

Figure 5

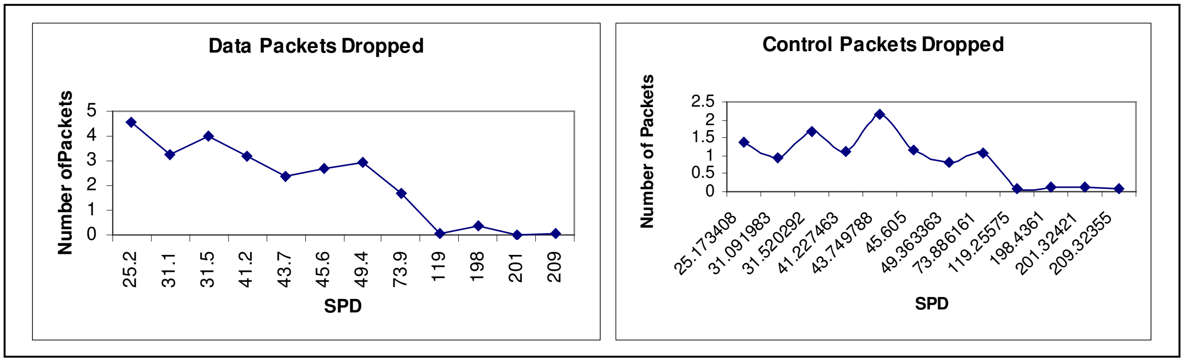

Figure 6

Future Work & Conclusions

From the experiments conducted, there is no way to conclude that the timekeeping metrics are directly applicable to a power consumption model for a network. However, there is no reason yet to disregard them; with more time, a number of remaining hypotheses can still be tested. There may be as few as one other factor that must be taken into account in order to have a proper model. For example, the power consumption must surely depend on the number of hops a packet must cover. When considering the SPD metric, a single path may exist, allowing for immediate transmission of data, but the energy consumed will depend on the length of the path. Also, it must be noted that the metrics were previously proven applicable to the successful delivery rate of the packets. However, failed attempts do consume energy until the packet is dropped. These cases are ignored by the metrics and probably have a great impact on the energy. Therefore, there are still multiple factors that perhaps also need to be incorporated into the equations in order to design an accurate model. We can conclude that modeling power solely on timekeeping metrics is not sufficient for the multifaceted protocols designed for efficient network communication.