



Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich



Cyrill Bannwart

Comparative Analysis of Dozer Network Performance

Semester Thesis SA-2011-13
March 2011 to June 2011

Advisor: Matthias Keller
Supervisor: Prof. Dr. L. Thiele

Abstract

Wireless Sensor Networks used in monitoring and surveillance scenarios are often deployed in areas that are difficult to access. As such they are required to work without any further human interaction and to recover from sensor node failures.

The goal of this semester thesis was to get a better understanding of the Dozer network protocol used to communicate in the WSNs maintained by the PermaSense project. Using Spearman's rank correlation the association between various parameters (e.g. failed transmissions or duplicate packets) collected in the WSNs since last summer have been analysed and compared between three deployments. The pre-analysis of the deployments showed that they have a different behaviour that can be attributed to their positioning, size, distance between nodes, and wireless link quality.

During the project multiple previously made assumptions on the Dozer protocol could be verified and we showed that RSSI (received signal strength indicator) measurements can be related to the behaviour of the deployments.

Finally a simple data reduction algorithm using a classification approach that is based on the cumulative distribution functions of the collected parameters has shown to be feasible to reduce the required bits to store and process the data.

Contents

1	Introduction	11
2	Problem Description	13
3	Deployments and Methods	15
3.1	Deployments	15
3.1.1	Jungfrauoch	15
3.1.2	Matterhorn	16
3.1.3	Thur	16
3.2	Time Range & Preprocessing	16
3.3	Analysed Parameters	16
3.3.1	RSSI (Received Signal Strength Indicator)	16
3.3.2	Transmission Failures	17
3.3.3	Events	17
3.3.4	Off-site Generated Parameters (Lost & Duplicate Packets)	18
3.4	Pre-Analysis	18
3.5	Analysis	18
3.5.1	Spearman's Rank Correlation Coefficient	19
3.5.2	Significance	19
3.6	Data Reduction	20
4	Results	21
4.1	Pre-Analysis	21
4.1.1	Hop Distribution	21
4.1.2	Child Distribution	22
4.1.3	RSSI (Received Signal Strength Indicator)	23
4.2	Analysis	23
4.2.1	Spearman's Rank Correlation	24
4.2.2	Deployment Comparison	26
4.2.3	Data Reduction	26

5 Related Work	29
5.1 Not All Wireless Sensor Networks Are Created Equal	29
5.2 Motes in the Jungle	30
5.3 Channel-Specific Wireless Sensor Network Path Analysis	31
6 Conclusion and Future Work	33
A MATLAB Files	35
B Additional Plots	37
References	41

List of Figures

4.1	Hop distributions for individual months	21
4.2	Child distributions for individual months	22
4.3	RSSI distributions	23
4.4	Correlation coefficients (Spearman's ρ) - Matterhorn	25
4.5	Correlation coefficients (Spearman's ρ)	26
4.6	Comparison of the RSSI related correlation coefficients	27
4.7	Correlation coefficients - with data reduction	27
B.1	Correlation coefficients (Spearman's ρ) (extended)	38
B.2	Correlation coefficients (Spearman's ρ) - with data reduction (extended) . . .	39

List of Tables

3.1	Assignment of ranks	19
3.2	Interpretation of the correlation	20
4.1	Results where the correlation is significant for > 40% nodes	25
A.1	Description of used MATLAB files	35

Chapter 1

Introduction

Wireless Sensor Networks (WSNs) are envisioned for many different monitoring and surveillance scenarios where the sensors are often deployed in areas that are difficult to access. As such they are required to work without any further human interaction and to recover from sensor node failures. To increase the lifetime of the WSNs their power consumption is the most important design factor [1].

In recent years new communication protocols that have a highly optimised power usage have been developed, such as the Dozer network protocol that according to its authors has shown to reduce the average radio duty cycle significantly while still guaranteeing a reliable data transfer [3].

For this semester thesis three deployments of the PermaSense project have been analysed. The PermaSense project is a real world application that investigates permafrost in the Swiss alps with WSNs currently using the mentioned Dozer protocol [12]. The WSNs collect environmental information that include temperatures, humidity or crack motions. To analyse and improve the networks additional data related to the communication and health of the network is collected.

This semester thesis report is partitioned in the following parts:

In chapter 2 we give a reduced problem statement for this semester thesis, that can also be found in original terms attached at the end of this report. Chapter 3 will first give an overview over the analyzed deployments and the collected data before presenting the used statistical measures. Subsequently chapter 4 presents the results of the analysis and in chapter 5 a summary of some of the related work that has been studied for this thesis will be given. The last chapter 6 then summarises and concludes the semester thesis report.

In the appendix A the list of MATLAB scripts and functions developed during the project can be found accompanied by a short description for the individual files. The second part of the appendix B contains further plots that have been included for reference.

Chapter 2

Problem Description

The PermaSense project maintains multiple WSN deployments that implement the Dozer network protocol for communication. The goal of this thesis is to compare the performance of the Dozer network protocols over different deployments. Since the above WSNs have been in use for multiple years, sufficiently large historical data sets exist for all considered deployments. To reduce the influence of already resolved issues the scope of the thesis is however limited to the data that has been retrieved from networks that are running the latest components only. The hardware components namely are the Tinynode 184 hardware [6] with a Semtech SX 1211 radio chip for communication. On the nodes a PermaDozer software image that is based on TinyOS 2 [7] is being used.

The available status information in particular consists of the following parameters that depending on their type of information are either sampled periodically, i. e. , every 2 minutes, or are created on the occurrence of an event:

- Snapshots of the current network topology
- RSSI (received signal strength indicator) measurements
- Information about the connection state of a node
- Number of failed packet transmissions
- Information on packet duplicates
- Event log revealing internal states of the Dozer protocol
- Radio/MCU (micro-controller) duty-cycle measurements
- Packet delays

Based on the status information we are interested in the quality and stability of the network links. A substantial interest also lies in the similarities and/or differences between the considered deployments. Further we are looking for methods to improve the processing and storage of the collected status information.

Chapter 3

Deployments and Methods

This section will take a look at the analyzed deployments to give an overview over their location and arrangements as well as the influences from their surroundings. The second part will present a description of the data that is collected during the operation of the Wireless Sensor Networks.

In a final section the statistical measure that has been employed in this semester project to analyse the collected data will be explained.

3.1 Deployments

The PermaSense project maintains multiple Wireless Sensor Networks which are mostly located in the mountains. After an early experiment at the Jungfrauoch in 2006, additional deployments have been built during the following years. At a later stage the experimental sensor network at the Jungfrauoch was replaced and only the actual sensors were kept.

For this thesis three deployments have been chosen. The Jungfrauoch and Matterhorn deployment are located in the mountains. Here, they are exposed to very harsh environmental conditions found at 3.500 metres above sea level. In contrast, the Thur deployment is located at the riverside of the river Thur.

3.1.1 Jungfrauoch

In the Swiss Alps on Jungfrauoch the first WSN of the PermaSense project has been installed in fall 2006. The WSN has the aim to quantify the spatial variability of thawing processes and heat transport in the near-surface layer [10]. After a few tests the data collection started in early spring 2007. However, due to software and hardware issues it soon became evident that the WSN had to be redesigned. Although the next generation WSN was first deployed at the Matterhorn, the Jungfrauoch deployment has later been upgraded and is consistent to the Matterhorn deployment since 2009.

The deployment at the Jungfrauoch has 17 nodes that are divided into two clusters to cover the north and the south facing slope of the mountain ridge. The distance between the nodes are of medium range up to 70 metres.

Near the Jungfrauoch multiple mobile phone as well as other antennas are located that may result in interferences with the analysed WSN. In this context it is noteworthy to know that due to the used Dozer protocol the WSNs are restricted to a single communication frequency.

3.1.2 Matterhorn

The deployment at the Matterhorn is installed at the north-east ridge called Hörnligrat. This particular site aims to investigate the temperature evolution in clefts as well as their dilatation [11].

In mid 2008 it was the first 2nd-generation PermaSense deployment and used the Dozer protocol for communication between the nodes. It is the largest analysed deployment and consist in total of 25 nodes that include the actual sensor nodes, additional relays as well as the base station used as the data sink. With short distances up to 40 metres between the nodes it has the highest density of the analysed deployments.

3.1.3 Thur

In contrast to the WSNs in the mountains, the deployment at the riverside of the river Thur is positioned at a level surface such that for most nodes the base station is in line of sight. It has only been operational since early 2010 and consists of 7 nodes only. The distance between the nodes is large and may be 100 metres and more. Just like at the Matterhorn deployment there are no known interferers, i. e. , mobile phone network stations, that could interfere with the communication between the nodes.

3.2 Time Range & Preprocessing

The analysed data that has been used for this semester thesis was collected during the period from 23. June 2010 to 1. March 2011. The start date was chosen such that only nodes running the latest hard- and software components, namely a Tinynode 184 with a Semtech SX1211 radio module and the software based on TinyOS, were used. This data range has been chosen to limit the influences of already resolved issues and to enable the comparison of the different deployments, that were previously running different software or even consisted of different hardware.

Using a fixed time range ensured that during the evaluation of different analysis methods the results remained consistent. This also allowed the local storage of intermediate results and the caching of queries at the database to speed up the data processing.

Being still under development the WSNs collected in addition to the desired environmental information (e. g. temperature) also data related to the communication behaviour and the health of the network. The different parameters will be described during the next few sections.

Depending on the type of information the parameters were either sampled periodically, i. e. , every 2 minutes, or on the occurrence of an event. To have consistent data all parameters have subsequently been aggregated to one value per day.

3.3 Analysed Parameters

3.3.1 RSSI (Received Signal Strength Indicator)

The Received Signal Strength Indicator is a physical layer information parameter that is constantly updated by the radio module whenever a radio signal is received. The measured parameter is the power in the received radio signal. When a node is connected it fetches the value from the register on the radio module every two minutes. During the analysis the stored

values are converted from the measured value to dBm using the formula in the Semtech SX1211 datasheet.

As the parameters from the physical layer are directly related to the wireless channel quality during the reception of a packet, there usually is a close correlation between the parameter and the link quality.

3.3.2 Transmission Failures

Each node generates state information packets, that include the period over which the state information was collected, the sending and receiving duration of the radio module, the activity of the micro controller and also the number of failed transmissions.

The number of failed transmissions are collected in the lowermost layer of the radio stack and represent for how many data packets no acknowledgement was received.

3.3.3 Events

Our implementation of Dozer stores certain events that occur during the execution of the Dozer protocol in a central database. Those events consist of an event id to identify the reason for the event being generated along with additional parameters related to the event.

The following list gives an overview over the most important events that may occur during the Dozer protocol execution.

- **CHILD NO DATAUPLOAD:**
The node could not send the data packet to its parent or more precisely did not receive the corresponding acknowledgement packet. For each failed attempt an event is generated, storing the number of attempts that failed in succession in the event parameter. The parameter gets reset to zero as soon as the transmission succeeds or when a new connection is set up. When the maximum number of tries is reached without being able to send the data, the child node will disconnect from its parent.
- **CHILD DISCONNECT:**
A child node has marked the connection to his parent as disconnected. Either because no beacons have been received from the parent for a certain amount of time or because the data upload to the parent failed multiple times.
- **CHILD CONNECT:**
The child node has found a new parent and has connected to it. The parameter of the event is the new parent ID.
- **PARENT DISCONNECT:**
The parent node has not received anything from its child node for a long time and thus marks it as not being connected. The event parameter contains the child node ID.
- **PARENT CONNECT:**
A connection to a new child node has been established. The event parameter contains the child node ID.
- **CHILD DISCONNECT ROUTINGLOOP:**
Every 30 seconds a beacon that contains the number of hops from the parent to sink is sent to the children. Whenever the number of hops is invalid or too large the children will disconnect.

- **RESET OVERHEARTIME:**
A node that is newly inserted into the network performs an initial neighbourhood discovery. The node listens for traffic on the channel and when traffic is received the node tries to gather additional knowledge about the neighbouring nodes by turning the radio on for a complete 30 seconds interval. The event occurs when the node turns its radio on for the 30 seconds interval.
- **CHILD FOUND PARENTS:**
The radio was on for 30 seconds and nodes that offer themselves as parents have been found. The parameter contains the number of possible parents.

The number of times a packet is not successfully sent to the parent node is counted using the «CHILD NO DATAUPLOAD» event. To calculate how often a child node changes its parent node the «CHILD CONNECT» and «CHILD DISCONNECT» events are used.

3.3.4 Off-site Generated Parameters (Lost & Duplicate Packets)

Using the sequence numbers of the data packets the number of lost packets between two successful communications can be detected off-site. Furthermore the sequence numbers is used to detect the packets that arrived multiple times at the base station.

3.4 Pre-Analysis

A short pre-analysis was performed to get a closer look at the behaviour of the three used WSN deployments. The analysis included a comparison of the deployments in terms of the number of hops that are required to deliver a data packet to the base station. In accordance with the Dozer protocol the number of hops is limited to 13 and the special value 14 is used to indicate that a node is not connected, which was however ignored during the analysis. To compare the deployments the distribution of the hops was used and to capture the development over time (dynamic) the distributions have been calculated for the individual months during the analysed time range.

In addition since the Dozer network protocol is organised in a tree-based manner the nodes may have multiple children. Thus the distribution that captures the number of children that the nodes have has been analysed as well, also using the individual month distributions to capture some of their dynamic.

Furthermore the wireless channel qualities have been compared for the deployments to see if there are similarities or differences between the deployments.

3.5 Analysis

The remaining analysis was looking at the question if the collected parameters were connected to each other or if there even was some redundancy between them. In addition since multiple deployments could be compared the next step was to determine if any found relations would exist in all three deployments and if this was not the case we were interested in the reasons.

But before answering these questions a statistical method that could be used to find dependencies between the collected parameters in a dataset with ~100 million entries had to be found. The statistical measure that was finally used in this thesis is described in the next section.

3.5.1 Spearman's Rank Correlation Coefficient

The Spearman's rank correlation coefficient, also called Spearman's rho, measures the association between two variables using monotonic functions. If there is an association between two variables, the knowledge of one value provides information about the likely value of the other. In contrast to the Pearson's correlation it will not require that there is a linear dependency between the two variables and since it is a distribution-free measure it does not rely on the assumption that the data comes from a given probability distribution (e. g. Pearson's r relies on a normal distribution) [8].

In fact the Spearman's rank correlation coefficient is the Pearson's correlation coefficient between the ranked variables. Using ranks produces variables that fulfill the condition of correlation as their relationship is now linear. In practice often the algorithmic formula 3.1 is being used to calculate the correlation [15].

$$\rho = 1 - 6 \sum_i \frac{d_i^2}{N(N^2 - 1)}, \text{ where } d \text{ is the difference in statistical rank of} \quad (3.1)$$

corresponding variables and N number of pairs of values.

To compute the ranks of a variable one has to assign positions to every value of the variable in declining order. Whenever a value exists multiple times, thus there are ties, the average of the positions is used to calculate the ranks as the example in Table 3.1 shows.

Table 3.1: Assignment of ranks

Variable X_i	Position	Rank
0.8	5	5
1.2	3	$\frac{4+3}{2} = 3.5$
18	1	1
1.2	4	$\frac{4+3}{2} = 3.5$
2.3	2	2

The general formula 3.2 also accounts for the existence of multiple ties in the dataset as the algorithmic formula gives just an approximation in these cases. The usage of the general formula is also the default behaviour when MATLAB's built in function `«[rho,pval] = corr(X,'type','Spearman')»` is being used [14].

$$\rho = \frac{\sum_i (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_i (x_i - \bar{x})^2 \sum_i (y_i - \bar{y})^2}} = \frac{Cov(x, y)}{s_x s_y}, \quad (3.2)$$

where x_i, y_i are the ranks, \bar{x}, \bar{y} the average thereof and s_x, s_y standard deviations.

A perfect Spearman correlation, that is a score of either -1 or $+1$, occurs when one variable is a perfect monotonic function of the other variable. The Spearman correlation coefficient is positive if Y increases when X increases and negative if Y decreases. As a rule of thumb the correlation coefficients can be interpreted according to Table 3.2.

3.5.2 Significance

Having calculated the correlation coefficients we also have to determine if the result is indeed significant that means it has to be unlikely that the value could have occurred by chance. The

Table 3.2: Interpretation of the correlation

Size of correlation	Interpretation
1	Perfect correlation
0.90 - 0.99	Very high correlation
0.70 - 0.90	High correlation
0.50 - 0.70	Moderate correlation
0.30 - 0.50	Low correlation
0.10 - 0.30	Very low correlation
0.01 - 0.10	Markedly low and negligible correlation
0	No correlation

Source: Bansal et al. [2]

significance level that has been used during this thesis was $\alpha = 0.05$ (5%). This is the most commonly used significance level.

When a correlation is to be computed without any statistical software, one will often use a precomputed table with the critical values that the resulting correlation has to exceed to be considered different from zero. Those tables with the critical values will also depend on the sample size as well as the chosen significance level.

Using a statistical software or also MATLAB's statistics toolbox we can however also compute the p-value that denotes the probability of obtaining a correlation coefficient that is at least as extreme as the one received under the assumption that the null hypothesis holds [13]. With the null hypothesis being that there is no association between the variables in the underlying population.

The computation of the significance was necessary since the dataset contained nodes that only existed temporarily and thus had only a few samples. Applying the correlation to such a node may result in a correlation coefficient that without calculating the significance could misleadingly have been interpreted as there being a correlation.

3.6 Data Reduction

During the semester thesis the amount of processed data became more and more evident such that the idea to reduce the required bits to process and store the data in the WSN arose. Finally a quantisation approach has been used to reduce the values of the different parameters to two bits with states «good», «normal» or «bad». As critical values to determine between the different states the 30% and 80% thresholds of the cumulative distribution functions of the parameters have been used.

Chapter 4

Results

First the results of the pre-analysis that explored the differences between the three analysed deployments will be presented before we turn our attention to the correlation results.

4.1 Pre-Analysis

4.1.1 Hop Distribution

Figure 4.1 shows the hop distributions of the analysed deployments. To capture the dynamic over the entire period the plots contain multiple distributions, each for an individual month.




Figure 4.1: Hop distributions for individual months

Matterhorn

In Figure 4.1a one can see that over time the Matterhorn deployment is quite stable. In contrast to the other two deployments we notice that there are nearly equally many nodes with two hops as there are with only one hop to the base station. Both hop counts combined make up for around 90% of all nodes, with the number of nodes that are connected over more than two hops declining rapidly.

Jungfrauojoch

Clearly visible in Figure 4.1b the Jungfrauojoch is the most unstable deployment during the analysed period. In fact this has previously been noticed to a certain extent by the members of the research group involved in the PermaSense project.

In contrast to the other two deployments there is a significant number of nodes that are connected via two hops, being the most often observed count. The variation over time is large reaching from under 0.5 to 0.7 relative frequency and in one month up to nearly 1.

Thur

The last of the three deployments that is depicted in Figure 4.1c is the Thur deployment. Also taking into account its small size it is very stable and in general exhibits a very low number of hops. In fact this was the expected behaviour since the deployment contains the fewest nodes and features a rather simple topology as most nodes can directly connect to the base station such that there are only rare cases where a node experiences more than two hops.

4.1.2 Child Distribution

The following Figures represent the number of children that a node has been observed to possess. A node may at most have 10 children, however this was rarely the case and even not possible for the smallest deployment.

Consistent with the hop distribution plots above, the dynamics have been visualised by plotting distributions for individual months.




Figure 4.2: Child distributions for individual months

Matterhorn

When examining the stability in Figure 4.2a we again notice that the deployment is probably as stable as the Thur deployment. However what can also be noted is that it is the deployment where the fewest nodes are so called leaves that do not have any children. This is most likely related to the large size and compactness of the deployment.

Jungfrauoch

It was seen for the hop distribution that there were lots of changes in the Jungfrauoch deployment in Figure 4.2b and as such we can also observe that instability in regard to the child count of the nodes. The Jungfrauoch deployment is the only one that constantly possesses nodes with 3 and more children and it is also the deployment where the nodes with most children have been observed.

Thur

The Thur deployment in Figure 4.2c is again quite stable. As expected, once again owing to the topology, we can find the most leaves as the nodes can communicate with the base station directly. The small number of deployed nodes further limits the possibility of having more than a few children and as such the deployment has very rarely nodes with three children, which is also the maximum value that has been observed for the deployment.

4.1.3 RSSI (Received Signal Strength Indicator)

In the second part of the pre-analysis the available physical layer information, that is the RSSI measurements of the deployments, have been compared. Figure 4.3 displays the cumulative distributions such that on the y axis is the RSSI values in dBm and the x-axis the percentage of the samples that are above that value.

The plots contain a red and a blue distribution curve for each deployment since the connections between the nodes are bi-directional and both endpoints of the connection make their independent signal strength measurements.




Figure 4.3: RSSI distributions

Comparing the different deployments one notices that the RSSI values for the Jungfrauoch are the lowest ones and as such that links of that deployment are inferior to the links of the other deployments. Secondly the differences between the red and blue curve show that there are asymmetric links that might be good in one direction but worse in the other.

4.2 Analysis

In this section the results involving the Spearman's rank correlation will be presented. First we will look at the correlation plots and explain the found relations between the collected

parameters and also set those findings in relation to the analysed deployments. In the second part we will evaluate the previously presented data reduction algorithm.

4.2.1 Spearman's Rank Correlation

The Spearman's rank correlation was calculated between the following parameter sets:

1. Number of «CHILD NO DATAUPLOAD» events
2. Number of failed transmissions
3. RSSI measurement (parent \Rightarrow child)
4. RSSI measurement (child \Rightarrow parent)
5. Number of duplicate packets
6. Number of parent changes
7. Number of packet losses

How to read the correlation plots

Although we have tried to only include the required information in the correlation plots they may not be self-explanatory. As such this short paragraph should help you to understand them:

The correlation has been calculated between any two previously shown parameters such that each bar represent a different set. Since the correlation of a parameter with itself is not of interest (always perfect) and the input order of the parameters does not matter there are

$$\frac{(\#\text{param.})^2 - \#\text{param.}}{2}$$

parameter sets / bars in the Figure. The ID of the input parameters is written on top of each bar with the name of the parameters being above this graph.

The black crosses within the bars are the correlation coefficients since they have been calculated for the individual nodes such that the whole bar represents the range of the received coefficients. Additionally the colour coding that has been applied to the bars shows for how many nodes the correlation was found to be significant as only those have been included in the plot.

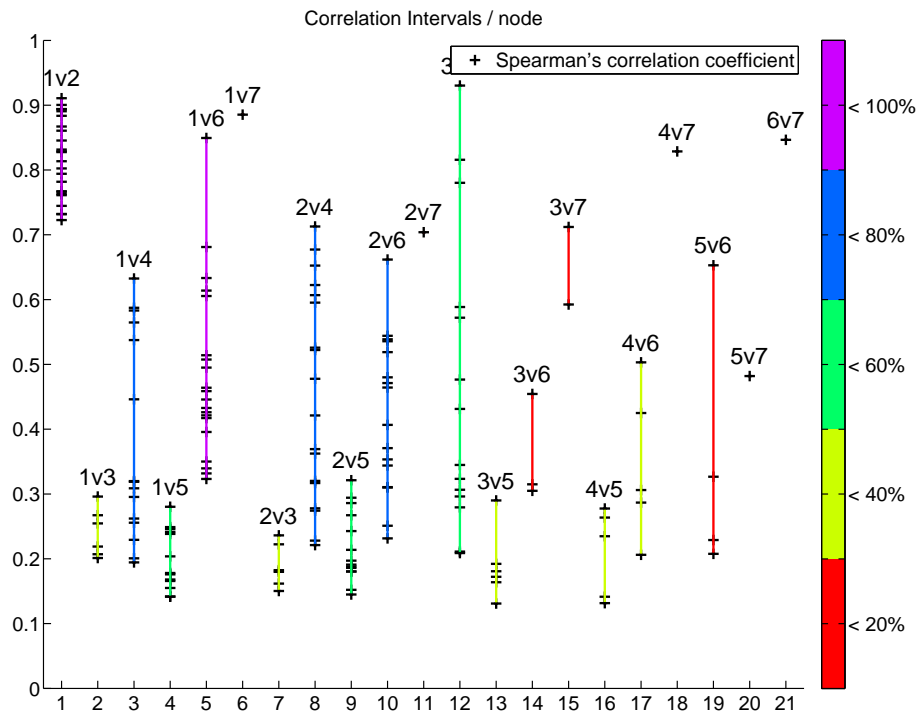
The results will be presented with the help of the Matterhorn deployment in Figure 4.4, however similar results were also found for the other two deployments. In Table 4.1 the results in which at least 40% of the nodes have a significant correlation are presented. The most important results and their explanations are also presented here.

The first bar (**1v2**) in the Figure represents the correlation between the «CHILD NO DATAUPLOAD» Dozer event and the number of failed transmissions.

There exists a very high correlation which has been highly expected to appear since the two parameters are directly connected to each other in software, where the only difference is that the transmission failures are counted in the radio stack itself and the events are individually sent and later counted during the analysis.

When the same event is compared to the number of parent changes as is the case in (**1v6**), a moderate to high correlation can be observed with the explanation that a node that repeatedly cannot send its data packet will eventually disconnect from its parent and try to find a new one. The value of the correlation is lower as before since the node will only disconnect if the data packet could not be sent for multiple times.

The correlation between the failed transmissions versus the duplicate packet counter, that is the case (**2v5**), is only low. The explanation for the low correlation is that there are two

Figure 4.4: Correlation coefficients (Spearman's ρ) - Matterhorn

reasons for a failed transmission: A lost data packet or a lost acknowledgement packet. However, only the case with the lost acknowledgement packet is associated with a duplicate packet at the base station.

Table 4.1: Results where the correlation is significant for > 40% nodes

Bar	Param.	max. Correlation	Sign.	Explanation
1	1 vs 2	Very High	100 %	The parameters are directly connected to each other in the software.
3	1 vs 4	Moderate	80 %	A deteriorating wireless link quality from child to parent results in failed transmissions. In contrast a deteriorating link from parent to child results in no beacons arriving and thus no data transmission at all, thus cannot be captured using the correlation (1 vs 3).
4	1 vs 5	Very Low	60 %	(see 2 vs 5)
5	1 vs 6	High	100 %	A node that repeatedly cannot send a data packet will disconnect from the parent and connect to a new one.
8	2 vs 4	Moderate	80 %	(see 1 vs 4)
9	2 vs 5	Low	60 %	When no ACK is received, the sender will retransmit the data resulting in a duplicate packet at the receiver if only the ACK was lost. A lost data packet however does not result in a duplicate.

10	2 vs 6	Moderate	80 %	(see 1 vs 6)
12	3 vs 4	High	60 %	For symmetrical links the correlation between up and down link is high and none to moderate for asymmetrical ones.

4.2.2 Deployment Comparison

When we compare the deployments in Figure 4.5, the most important differences can be found for the RSSI measurements. Thus in Figure 4.6 the correlations between the RSSI measurements and number of duplicate packets, parent changes and lost packets are shown next to each other for all deployments.




Figure 4.5: Correlation coefficients (Spearman's ρ)

While the Thur deployment does not show any useful correlations they exist for the Matterhorn and to a greater extent for the Jungfrauoch. We estimate that this is due to the fact that a majority of the links at the Jungfrauoch are usually already saturated such that a degradation of the link quality will result in lost acknowledgements or data packets and we see an increased rate of duplicates or disconnects.

In the pre-analysis we have shown that the wireless link quality of the Thur deployment was much better such that a variation in the quality will not immediately result in failed transmissions or topology changes. The quality of the links at the Matterhorn was in between the other two deployments which is consistent with results we see in this comparison.

4.2.3 Data Reduction

As has been described in the «Deployments and Methods» chapter, a data reduction algorithm to quantise the parameters to the values «good», «normal» and «bad» has been used. To verify that the results found in the previous sections still hold, the Spearman's rank correlation was computed using the quantised values. The resulting plots are presented in Figure 4.7.

When the correlation results with the data reduction applied (Figure 4.7) are compared to the results without the reduction (Figure 4.5), one notices that the values of the correlations and significances have slightly decreased for a few parameters. The previously made findings in the case without reduction are also existent in the reduced data set such that in theory the parameters could be reduced to a two bit value.

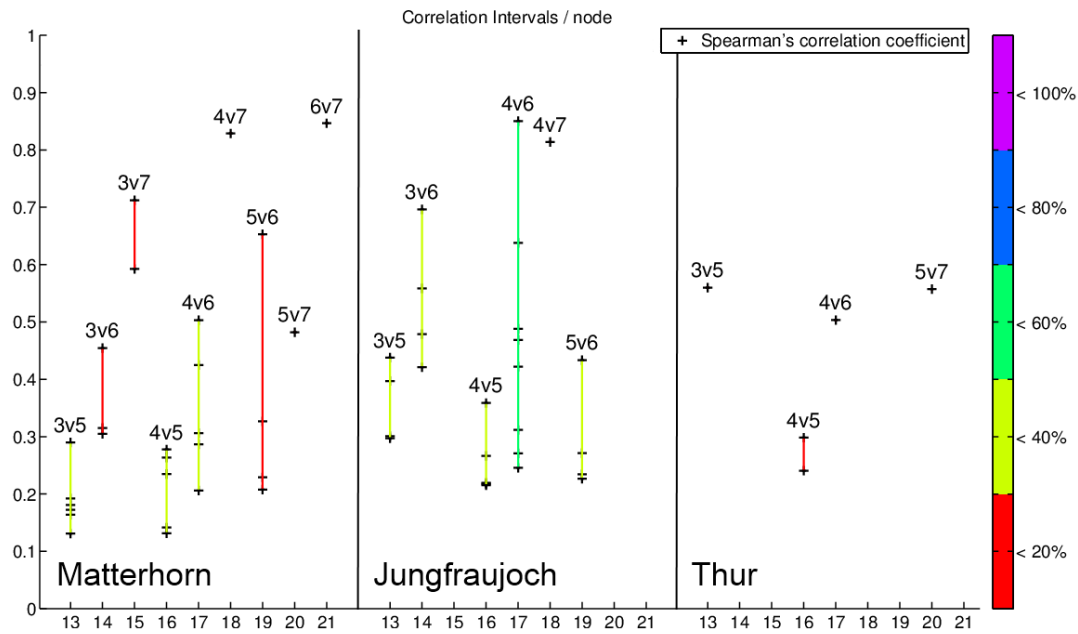


Figure 4.6: Comparison of the RSSI related correlation coefficients




Figure 4.7: Correlation coefficients - with data reduction

Since the critical values to classify the parameters were their cumulative distribution level values at 30% and 80% that have been chosen arbitrarily we believe that the results could be improved by selecting more sophisticated levels. In addition the quantisation using two bits only may be to vague for certain applications such that additional levels could enhance the results even further.

Chapter 5

Related Work

In this section we will present related work that has been studied for this semester thesis.

5.1 Not All Wireless Sensor Networks Are Created Equal: A Comparative Study on Tunnels

In Mottola et al. [9] multiple tunnel deployments that included an operational freeway tunnel nearby Trento, a non-operational tunnel and a vineyard that was used to represent an outdoor environment have been compared.

All three networks consisted of 20 WSN nodes that have been arranged in two opposing rows, as would be the case in the tunnel environment. Additionally it was payed attention that the nodes are never directly opposite to each other and at the ends of the rows the nodes had to be placed more densely, since those were the requirements by the TRITon project where WSNs are supposed to support adaptive light control in traffic tunnels. In the vineyard deployment the positions of the nodes were limited by the existence of poles to mount them, thus those constraints determined the node distance for all deployments.

The experiments of Mottola et al. targeted the physical, MAC and routing layers. For the physical layer the packet delivery ratio, the average RSSI as well as LQI (link quality indicator) and some observed environment parameters (e. g. average temperature or observed relative humidity) were taken into account. For the physical layer experiments the nodes had to be synchronized before the experiment during which each node broadcasted a packet every $N \times \delta$ time instants. To avoid collisions every node had a different offset. The results were then collected using a network-wide flooding. The MAC layer experiments required the additional control of the probability of concurrent transmissions. As output the same parameters as in the physical layer experiments were collected. The routing layer experiments collected information about the total message delivery ratio at the gateway node, the number of duplicates dropped inside the network, the parents in the tree along with the number of parent changes, the failed ACK transmissions at every node and the number of beacons at each node.

In the paper of Mottola et al. tried to compare the results of the physical layer to the higher layers. According to their results the temperature had no influence on the physical layer while the humidity only had an impact above a certain threshold. The communication range in the tunnels was much larger than in the vineyard and the link performance was quite stable over time. They assume that the particular shape of the tunnels act as a waveguide resulting in multi-paths effects and constructive interferences. Where the wireless links were either very

good or very bad the packet losses were independent of the fact that the previous k packets have been received. Not surprisingly they experienced that concurrent transmissions are very likely to generate collisions and packet losses.

According to the results of the paper the RSSI measurements did not reflect the link quality except for very good links. However, in the tunnels the LQI measurements could be used to identify links of intermediate quality. When asymmetric links existed they were mostly permanent and in the tunnel deployments the link symmetries were related to the node positions. Additionally the impact of vehicular traffic was analysed.

5.2 Motes in the Jungle: Lessons Learned from a Short-term WSN Deployment in the Ecuador Cloud Forest

The paper of Ceriotti et al. [4] describes an experimental WSN deployment in the Ecuadorian Andes which are between 1200 and 2800m above sea level. They used 18 TMote Sky sensor nodes of which 8 were stationary and arranged in a cross where the centre node served as a coordinator that configured and started the experiments. The experiments took place at the end of March and lasted for a week, since this was during the rainy season it was raining all days and the sensor nodes had to be packaged in a watertight box.

The goal was to characterise the physical connectivity and as such the software was based on TinyOS without any MAC protocol and to ensure that no packet collisions would occur the coordinator node assigned a predetermined time-schedule to every node. During the experiments the nodes collected information for each link that included the packet delivery ratio, the RSSI as well as the LQI.

Initially they determined the range of communication using different power levels and storing the aggregated average RSSI, LQI and PDR values. In a second step the connectivity among the nodes at different distances was analysed over a long time interval. However, due to the available terrain that was highly irregular and the undergrowth that interfered they later decided that the collected data was unusable.

In addition to the stationary experiments there were mobile ones where not only an aggregated value per link but test statistics for each individual message were recorded. Carrying a mobile sensor node the biologist moved around in the test area being filmed by a second team member.

Looking at the results that they have gained showed that when the nodes were in direct line of sight the links were symmetric in respect to the recorded PDR, but when there was a tree near in between this would result in a weaker link for the communication path originating near the tree. The RSSI and LQI however did not show any link asymmetry. Mid-range links had highly-variable quality and low RSSI and the variability was unpredictable such that they conclude that mid-range links cannot guarantee connectivity. Long-range links however were basically unusable but they still caused long-range interference and as such should be avoided. The RSSI measurements also showed that the influence of trees and bodies is more intense on shorter links while longer links were not affected.

Finally the authors of the paper also noted that the LQI varied significantly throughout the day, in fact they conclude that this may have influenced their short time measurements as some of their results showed that the quality of the links in the mobile scenario was superior to the quality measured in the stationary scenario.

5.3 Channel-Specific Wireless Sensor Network Path Analysis

In Doherty et al. [5] a monitoring sensor network was deployed and running for 25 days in a printing factory in Berkeley, California. The factory was divided into three areas with different propagation obstacles reaching from small job printing cells to a printing machine that takes a 5-ton paper roll. In total there were 46 deployed nodes. Each of those nodes reports its neighbours in range to a network manager that constructs the healthiest possible network. The nodes had between 4 and 26 possible neighbours and the farthest node was at least 3 hops away from the base station. As a constraint each node was only allowed to have up to 8 neighbours to communicate with. All sensor nodes report during 15 minutes the following data:

- channel number
- ID of neighbour
- number of transmissions
- number of transmit CCA (clear channel assessment) fails
- number of «No ACK» events
- number of ACK CCA fails
- number of receptions
- mean RSSI for all receptions
- mean LQI for all receptions

The stability of a channel-path was then defined as measured at the transmit end of the path to be:

$$\text{stability} = 1 - \frac{\#\text{NoACK}}{\#\text{Transmissions}} \quad (5.1)$$

The results when averaged over all paths and time showed that no channel was significantly better or worse than any other. The automated construction does not explicitly choose paths with high stability, however those with low stability fail more often and as such there is a self-selection of higher stability paths.

When Doherty et al. compared the stability of the path with the highest traffic they noticed that the stability was varying between the different channels. Also the four best channels according to the RSSI value matched the four most stable channels and the general shape of the LQI plot was similar to the one of the stability plot. However they also note that the RSSI and LQI values are only measured on successful packets, so that there is a bias.

Also the path-channel stability is dependent on the physical nature between two nodes and measurement of noise in the environment of a certain node is not sufficient to predict the frequency behaviour of their paths. Additionally there is a temporal behaviour that varies across channels on a given path.

To determine the path symmetry Doherty et al. used the measured RSSI values at both endpoints. When they measured a consistent difference between the two readings they attributed it to output power differences.

Chapter 6

Conclusion and Future Work

The goal of this semester thesis was to perform a comparative analysis on the Dozer network protocol using the collected data of three PermaSense deployments.

For the comparison of the large datasets, the Spearman's rank correlation has been used. The calculation, analysis and plotting of the results have been implemented using MATLAB.

The MATLAB scripts and functions have been programmed in a way such that additional parameters and deployments can be added easily. Also the duration of the included dataset and the aggregation periods are fully customisable. Additionally where it showed to be helpful, the data preprocessing parts have been modified to work well with the available CONDOR environment at ETH.

The results of the analysis showed that the collected data contained redundancies since some parameters were collected at multiple locations in the Dozer implementation. Also for all the additional associations that have been found we had logical explanations.

As a major driver we identified the wireless link quality that has been measured using the Received Signal Strength Indicator. The measurements showed that the three deployments have a different behaviour.

Of the seven parameters that have been taken into account the number of packet losses showed to be the only parameter that was not associated to any other parameter.

Furthermore, using a simple classification algorithm we could show that the found relations are still visible after a data reduction of the parameters to only two bits per value.

We think that taking into account our findings the communication between nodes in existing WSN deployments can be improved by inserting additional nodes that act as relays.

Also, currently the rating function of the Dozer protocol to chose the parents only uses the distance to the sink and the number of children. A future version of the protocol could be modified to include physical layer information, e. g. , RSSI or LQI that determine the link quality.

Appendix A

MATLAB Files

Table A.1: Description of used MATLAB files

Scripts:

correlation_<deployment><n>.m	<ol style="list-style-type: none">1. Define the deployment2. Load the locally stored data or fetch it from the database and store it locally.3. Prepare and aggregate the parameters.4. Calculate the correlations between the parameters and plot them.5. Determine the 30 and 80% threshold of the parameters CDF to be used in the classification.6. Calculate the correlation after applying the classification algorithm.
calc_average_<child/hop>_count.m	Calculate the average child/hop count and plot the child/hop distributions for all nodes.
calc_temperature_difference.m	Calculate and plot the difference between the max and min temperatures of all nodes.
classification_test.m	Calculates the percentage of matching classes for all parameters and nodes. <i>Requires the data from running correlation_<deployment>_n.m</i>
deployment_nodes.m	Plot the timing intervals during which the nodes were active.
doz_rssi.m	Fetch and prepare RSSI from database. (Use CONDOR)
eventstats.m	Plot time series of events.
eventstats_child.m	Plot classified «child no dataupload» events for the Matterhorn deployment.
plot_<hop/child>_distributions.m	Plot <hop/child> distributions for multiple months.
plot_rssi.m	Plot the RSSI distributions for separate up and down measurements. <i>Requires the data from running correlation_<deployment>_n.m</i>
validate_nodes.m	Testscript to determine nodes that do not behave as expected.

Functions:*Database access & Preprocessing:*

get_data.m	Fetch events and transmission failures.
fetchnodedata.m	Interface to facilitate the database access.
fetchevents.m	Events
fetchstatecounter.m	Statecounter including transmission failures.
get_duplicates.m	Fetch «packet duplicates» from database.
get_losses.m	Fetch «packet losses» from database.
fun_eventstats.m	Fetch Events. <i>Required by eventstats.m</i>
fetchrssi.m	Fetch RSSI values from database and convert returned layout. <i>Required by doz_rssi.m</i>

Aggregation:

get_cndu_data.m	Aggregate «child no dataupload» data.
get_fail_data.m	Aggregate «failed transmissions» data.
load_rssi_data.m	Load and aggregate RSSI data from files.
get_duplicates_data.m	Aggregate «packet duplicates» data.
get_parent_changes_data.m	Calculate and aggregate «parent changes» data.
get_losses_data.m	Aggregate «packet losses» data.

Calculations & Plotting:

do_correlation.m	Compute Spearman's rank correlation.
normalize_data.m	Normalize input data based on number of sent packages.
plot_correlation_interval<_n>.m	Plot the correlation between the input variables.
plot_distributions.m	Plot the distributions of the input variable. Calculate 30 and 80% thresholds.

Misc:

figpref.m	Set figure preferences.
print_figure.m	Print figure to a file.
mat2unixtime.m	Convert MATLAB date to Unixtime.
unixtime2mat.m	Convert Unixtime to MATLAB date.
location2nodes.m	Convert node position to node IDs.

Appendix B

Additional Plots

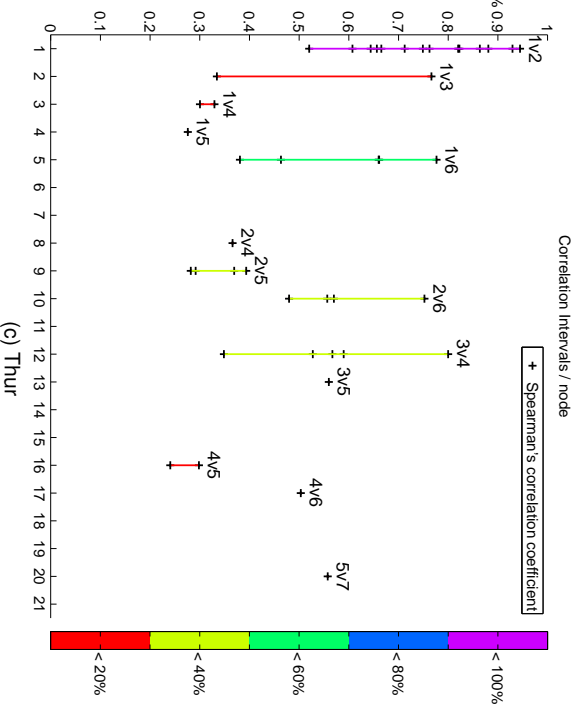

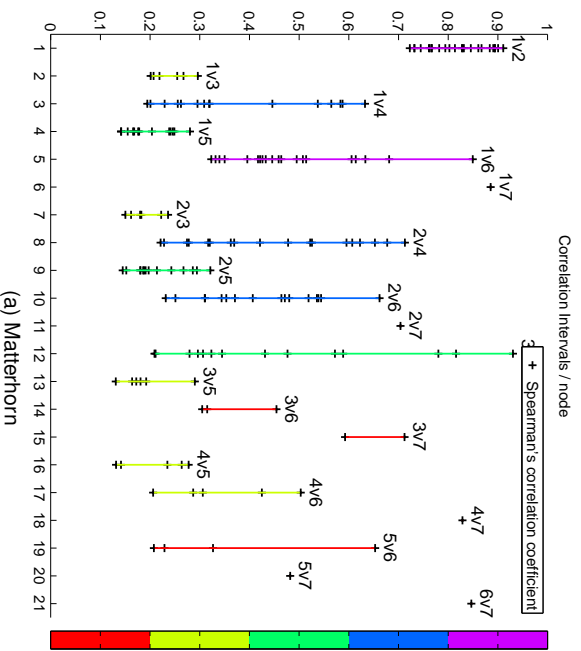


Figure B.1: Correlation coefficients (Spearman's ρ) (extended)




Figure B.2: Correlation coefficients (Spearman's ρ) - with data reduction (extended)

Bibliography

- [1] Mokhtar Aboelaze and Fadi Aloul. Current and Future Trends in Sensor Networks: A Survey. 2005.
- [2] I.K. Bansal, K.K. Vasishtha, T.C. Gyanani, M. C. Sharma, Snehlata Shukla, and Sunaina Kumar. *Correlation - It's Interpretation and Importance*. Indira Gandhi National Open University (IGNOU), 2008.
- [3] Nicolas Burri, Pascal von Rickenbach, and Roger Wattenhofer. Dozer: ultra-low power data gathering in sensor networks. In *Proceedings of the 6th international conference on Information processing in sensor networks*, IPSN '07, pages 450–459, New York, NY, USA, 2007. ACM.
- [4] Matteo Ceriotti, Matteo Chini, Amy L Murphy, Gian Pietro Picco, Francesca Cagnacci, and Bryony Tolhurst. Motes in the Jungle: Lessons Learned from a Short-Term WSN Deployment in the Ecuador Cloud Forest. In *REALWSN*, pages 25–36, 2010.
- [5] Lance Doherty, William Lindsay, Jonathan Simon, and Kristofer S. J. Pister. Channel-Specific Wireless Sensor Network Path Analysis. Technical report, Dust Networks, 2007. URL <http://www.dustnetworks.com/about/library/white-papers/channel-specific-wireless-sensor-network-path-analysis>.
- [6] Henri Dubois-Ferrière, Roger Meier, Laurent Fabre, and Pierre Metrailler. TinyNode: A Comprehensive Platform for Wireless Sensor Network Applications. In *Proceedings of Fifth International Conference on Information Processing in Sensor Networks*, pages 358–365, 2006.
- [7] P Levis, S Madden, J Polastre, R Szewczyk, K Whitehouse, A Woo, D Gay, J Hill, M Welsh, E Brewer, and D Culler. TinyOS: An Operating System for Sensor Networks. In Werner Weber, Jan M Rabaey, and Emile Aarts, editors, *Ambient Intelligence*, chapter 7, pages 115–148. Springer-Verlag, Berlin/Heidelberg, 2005.
- [8] Mathematics in Education and Industry. Spearman's rank correlation. Technical report, Mathematics in Education and Industry, 2011. URL www.mei.org.uk/files/pdf/Spearmanrcc.pdf.
- [9] Luca Mottola, Gian Pietro Picco, Matteo Ceriotti, Stefan Guna, and Amy L. Murphy. Not all wireless sensor networks are created equal. *ACM Transactions on Sensor Networks*, 7(2):1–33, August 2010.
- [10] PermaSense. Jungfrauojoch, . URL <http://www.permasense.ch/research/field-sites/jungfrauojoch.html>.
- [11] PermaSense. Matterhorn, . URL <http://www.permasense.ch/research/field-sites/matterhorn.html>.
- [12] Igor Talzi, Andreas Hasler, Stephan Gruber, and Christian Tschudin. PermaSense: Investigating Permafrost with a WSN in the Swiss Alps. In *EmNets '07 Proceedings of the 4th workshop on Embedded networked sensors*, 2007.

- [13] Wikipedia The Free Encyclopedia. p-value, . URL http://en.wikipedia.org/wiki/P_value.
- [14] Wikipedia The Free Encyclopedia. Spearman's rank correlation coefficient, . URL http://en.wikipedia.org/wiki/Spearman%27s_rank_correlation_coefficient.
- [15] Wolfram Research, Inc. Spearman Rank Correlation Coefficient. URL <http://mathworld.wolfram.com/SpearmanRankCorrelationCoefficient.html>.

SEMESTER THESIS

for
Cyrill BannwartAdvisor: Matthias Keller
Co-Advisor: Dr. Jan Beutel

Start date: March 7, 2011
End date: June 17, 2011

Comparative Analysis of Dozer Network Performance

Introduction

The PermaSense [1] project deploys and operates wireless sensor networks (WSNs) for recording geophysical measurements related to permafrost. There exist deployments at multiple locations (*i.e.*, Matterhorn, Jungfrauoch, Thur river), each deployment consists of a base station and a number of sensor nodes. A sensor node integrates a Tinynode [2] mote and a purpose-built sensor interface [3] (SIB) within a protective enclosure. The used software is based on the TinyOS operating system [4] and the Dozer ultra low-power network protocol [5].

Since the first deployment of a PermaSense WSN at the Matterhorn in Summer 2008, there have been several iterations in which software and hardware components were upgraded or even replaced. For instance, earlier versions of the network suffered from high amounts of lost packets and very frequent node reboots [6].

Problem Definition

In this thesis, we want to compare the performance of the Dozer network protocol over different deployments. Needed status information for running this comparison is (periodically) generated at all nodes, there exist sufficiently large historical data sets for all considered deployments. For reducing the influences of already resolved issues, we limit the scope of this thesis to data that has been retrieved from networks running latest components, namely Tinynode 184 hardware and a PermaDozer software image that is based on TinyOS 2. All nodes use the same radio chip for communication, namely a Semtech SX1211 radio.

In particular, we have access to the following information:

- Snapshots of the current network topology
- RSSI (received signal strength indicator) measurements
- Information about the connection state of a node
- Number of failed packet transmissions
- Information on packet duplicates
- Event log revealing internal states of the Dozer protocol
- Radio/MCU (micro-controller) duty-cycle measurements
- Packet delays

Depending on the type of information, the information is either sampled periodically, *i.e.*, every 2 minutes, or on the occurrence of an event.

Based on this information, we are interested in the quality and stability of networks links. Further interest lies on similarities found between snapshots of the network topology, *i.e.*, if many nodes are connected to the same parent node for a significant time amount. (Optional task) If time permits, further topics of interest include energy consumption and packet delays.

In order to reach these goals, the project should proceed according to the following steps:

1. The student should write a project plan and identify its milestones (thematically and concerning time). In particular there should be enough room for the final presentation and the report.
2. The student should find and study related work, *i.e.*, previously performed studies by other researchers [7, 8]. The results of this literature research should be written down as a chapter of the report.
3. The student should implement a set of MATLAB scripts that 1) fetch raw data from a database server, 2) generate needed intermediate, reduced results, 3) implement needed application logic for evaluating network stability and performance, and 4) export evaluation results into plots. Existing scripts used for performing similar tasks are provided to the student at the beginning of the project.
4. The student should interpret gathered evaluation results and document findings in the report.

Organization

- Duration of the Work:
This Semester Thesis starts March 7, 2011 and has to be finished no later than June 17, 2011.
- Project Plan:
A project plan with its milestones is held and updated continuously. Unforeseen difficulties that change the project plan have to be documented and should be discussed with the advisors in a timely manner.
- Subversion Repository:
All relevant files, *i.e.*, presentations, scripts, and text documents, should be stored in the Subversion repository that has been assigned to the project. Timely and frequent data synchronization prevents data loss and allows to collaborate more easily.

- Weekly Meetings/Reports:
In regular (weekly) meetings with the advisors, the current state of the work, potential difficulties as well as future directions are discussed. The day before the weekly meeting a brief status report should be sent to the advisors commenting on these issues, in order to allow an adequate meeting preparation for the student and the advisors.
- Initial Presentation:
Approximately two to three weeks after the start the student presents the objectives of the work as well as some background on the topic. The presentation should not exceed 5 minutes and consist of no more than three slides.
- Final Presentation:
By the end of the project, the student will present the achieved result. The presentation should not exceed 15 minutes.
- Documentation:
At the end of the project, no later than *June 17, 2011*, the student will have to hand in a written report. Together with the system implementation/software this report is the main outcome of the project. Code has to be commented extensively, allowing a follow-up project.
- Evaluation of the work:
The criteria for grading the work are described in [9].
- Finishing up:
The required resources (e.g., laptop, keys) should be cleaned up and handed back in.

References

- [1] A. Hasler, I. Talzi, J. Beutel, C. Tschudin, and S. Gruber, "Wireless sensor networks in permafrost research - concept, requirements, implementation and challenges," in *Proc. 9th Int'l Conf. on Permafrost (NICOP 2008)*, vol. 1, pp. 669–674, June 2008.
- [2] H. Dubois-Ferrière, L. Fabre, R. Meier, and P. Metrailler, "Tinynode: a comprehensive platform for wireless sensor network applications," in *Proc. 5th Int'l Conf. Information Processing Sensor Networks (IPSN '06)*, pp. 358–365, ACM Press, New York, 2006.
- [3] J. Beutel, S. Gruber, A. Hasler, R. Lim, A. Meier, C. Plessl, I. Talzi, L. Thiele, C. Tschudin, M. Woehrle, and M. Yuecel, "PermaDAQ: A scientific instrument for precision sensing and data recovery in environmental extremes," in *Proc. 7th Int'l Conf. Information Processing Sensor Networks (IPSN '09)*, pp. 265–276, ACM Press, New York, Apr. 2009.
- [4] P. Levis, S. Madden, J. Polastre, R. Szewczyk, K. Whitehouse, A. Woo, D. Gay, J. Hill, M. Welsh, E. Brewer, and D. Culler, *Ambient Intelligence*, ch. TinyOS: An Operating System for Sensor Networks, pp. 115–148. Springer, Berlin, 2005.
- [5] N. Burri, P. von Rickenbach, and R. Wattenhofer, "Dozer: ultra-low power data gathering in sensor networks," in *Proc. 6th Int'l Conf. Information Processing Sensor Networks (IPSN '07)*, pp. 450–459, ACM Press, New York, Apr. 2007.
- [6] M. Keller, J. Beutel, A. Meier, R. Lim, and L. Thiele, "Learning from sensor network data," in *Proc. 7th ACM Conf. Embedded Networked Sensor Systems (SenSys 2009)*, (New York, NY, USA), pp. 383–384, ACM, 2009.
- [7] L. Mottola, G. P. Picco, M. Ceriotti, c. Gună, and A. L. Murphy, "Not all wireless sensor networks are created equal: A comparative study on tunnels," *ACM Trans. Sen. Netw.*, vol. 7, pp. 15:1–15:33, September 2010.
- [8] M. Ceriotti, M. Chini, A. Murphy, G. Picco, F. Cagnacci, and B. Tolhurst, "Motes in the Jungle: Lessons Learned from a Short-Term WSN Deployment in the Ecuador Cloud Forest," *Real-World Wireless Sensor Networks*, pp. 25–36, 2010.

- [9] E. Zitzler, "Studien- und Diplomarbeiten, Merkblatt für Studenten und Betreuer." Computer Engineering and Networks Lab, ETH Zürich, Switzerland, Mar. 1998.