# Analysis of User Behavior from Billing Records of a CDPD Wireless Network

Luc A. Andriantiatsaholiniaina and Ljiljana Trajkovic*
School of Engineering Science
Simon Fraser University
Burnaby, BC, Canada V5A 1S6
(604) 291-3998
{luca, ljilja}@cs.sfu.ca

## ABSTRACT

Collection of user statistics and network traffic is crucial for understanding user behavior and for creating network workload models. It is also valuable for the management of commercial wireless networks. In this paper, we report on the analysis of billing records collected from the Telus Mobility Cellular Digital Packet Data (CDPD) network. The longest continuous billing record that we examined covered approximately twenty one days, spanning the Christmas and New Year holiday seasons. We used various tools to graphically illustrate the billing data. We observed that network activities exhibit daily and weekly cycles. Furthermore, the clustering analysis revealed four distinct behavioral classes of users. Analysis of billing data provided useful information about the usage of an operational wireless network.

## Keywords

Wireless networks, CDPD networks, mobile networks, billing data, clustering analysis, traffic analysis.

## 1. INTRODUCTION

A number of recent research reports have been devoted to collection and analysis of traces from operational wireless networks. Analysis of these traces provided useful information about the behavior of network users. Collected traces also enabled traffic-based performance evaluations of wireless networks.

In this paper, we present an analysis of a billing record obtained from Telus Mobility CDPD network. The record lasts approximately twenty one days (from December 22, 2000 to January 11, 2001). Our objective was to identify patterns of user behavior and network activities.

In order to protect users' privacy, the mobile Network Entity Identifiers (NEIs) (also called User IDs) have been sanitized (scrambled). Billing data consist of three types of events: registration, deregistration, and IP data [2]. Registration events occur when a user (an IP-enabled device) attempts to identify itself to the network in order to gain access to network services. Deregistration events occur when a user leaves a cell. These deregistration events are optional. Data events describe the actual network traffic generated by users.

We performed data extraction in a variety of ways. Java[tm] [6] programs were used to parse the billing records and to write billing data as a series of Structured Query Language (SQL) statements that load the billing data into a MySQL database [8].

The main difficulty in analyzing billing records has been dealing with its sheer volume. Hence, we used data mining and a machine learning technique called clustering. Clustering is used for discovering hidden patterns and trends in a given data set. It groups data into categories with similar behavior. In our analysis, we used *k-means* clustering algorithm from the S-PLUS statistical package [12].

We describe the aggregate network traffic, user behavior, and cell activities. Our analysis may

* This research was funded by the Grant-in-Aid from Telus Mobility and the BC Advanced Systems Institute Fellowship.

provide useful information about the access to the CDPD network even if the billing record may be short and not representative because of holiday seasons. We were also able to provide a classification of network users and its cells. These results, although preliminary, may help in modeling network traffic and in creating workload models for commercial wireless networks.

The remainder of this paper is organized as follows. In Section 2, we present an overview of CDPD protocol and the CDPD billing data. We present data analysis in Section 3. It includes aggregate network characteristics, a detailed analysis of user behavior, and additional results based on the analysis of cell activities. Related work is described in Section 4. We conclude with Section 5.

## 2. BACKGROUND

In this section we present a brief description of CDPD protocol and the CDPD billing data.

### 2.1 CDPD protocol

CDPD [1, 11] is a standard protocol developed for commercial public mobile data communication networks. The CDPD communication architecture is based on the Open Systems Interconnection (OSI) Reference Model [15]. It deals only with the lower three OSI layers. CDPD network's function is to enable data transmission between Mobile End Systems (M-ESs) and Fixed End Systems (F-ESs). The topology of a simple CDPD network is shown in Figure 1 [7]. M-ESs are connected to the backbone network through the Mobile Data Base Station (MDBS), the Mobile Data Intermediate System (MD-IS) (also called a mobile router), and an Intermediate system. The network also includes several Fixed End Systems (F-ESs) connected to the wired backbone network.

CDPD is a multiple access protocol. Stations that want to transmit data must compete for access in the shared communications medium. CDPD shares some characteristics with multiple access protocols, such as Ethernet (IEEE 802.3), while still having significant differences. CDPD differs from other multiple access protocols mainly in two aspects: the wireless transmission medium and the mechanism for collision detection.

Mobile Data Base Station (MDBS) communicates with M-ESs over the airlink interface, providing data link layer and physical layer functions to paired radio channels (forward and reverse) within its cell. Forward channel is contentionless, and, therefore, always available for the MDBS to broadcast data to M-ESs. Multiple M-ESs have to compete with each other to access the reverse channel for data transmission to the MDBS.

Although the CDPD network supports multiple protocols at the network layer, all CDPD network support services are based on International Standards Organization (ISO)/OSI protocol suites. The two categories of CDPD support services are CDPD network support and CDPD network application support services. We are in particular interested in CDPD network application support services, which include network management, message handling, and accounting [2].

#### *2.1.1 Network management services*

The CDPD network provides comprehensive mobile data communication services to subscribers. To ensure high-level network availability, the CDPD network is designed to incorporate network management services that allow CDPD service provider to operate the network. The network management services provide timely information to the network operator to detect network faults, to exercise controls, to correct faults, and to configure the network for optimal operation.

#### *2.1.2 Message handling services*

Message handling services provide a generic message, store, and forward services that are useful to other CDPD network support services.

#### *2.1.3 Accounting services*

Accounting services provide information on how and by whom the CDPD network resources are used. The CDPD accounting services collect network usage data for every subscriber to enable compilation for billing purposes. The collected data may include packet count, packet

size, source and destination addresses, cell ID used by M-ESs, and an approximate time of transmission.

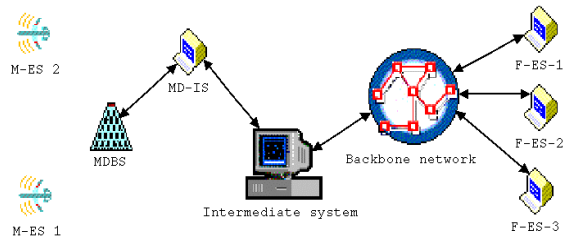


**Figure 1. Topology of a CDPD network.**

## 2.2 Description of the CDPD billing data

Analysis of network records often focuses on the overall network behavior (network elements, network events, and network activity characterization) and on traffic characteristics from a user point of view (when, for how long, and in which manner customers use the network). Records could be session data, transport layer data, application layer data, and movement data [3, 4, 5, 13, 14].

The billing records from a CDPD network that we analyzed summarize the network activity. They are used by the service provider to bill its customers for network usage. The structure of the billing record is a set of directories and files. Each directory is named after a sequence number (e.g., 528/, 529/) and contains up to 100 data files also named sequentially (e.g., 00074-52800, 00074-52801, 00074-52802). Each file contains a record of approximately fifteen minutes of network activities. A file consists of a header that contains timestamp and sequencing information, and a series of TMS rows. The smallest number of recorded TMS rows in a file is 300 and the largest is 2,694. Each row represents a single event of either registration, deregistration, or IP data type. TMS rows include registration and deregistration event timestamps, the number of packets and octets transferred, as well as the number of discarded packets. The billing records we analyzed consisted of 1,888 files.

The billing record starts at 11:30 AM on December 22, 2000 and lasts until 6:30 AM on January 11, 2001 (duration of twenty one days). There is a discontinuity in the billing record on January 2, 2001 from 7:30 PM until 10:30 PM, and the directory 536/ contains only 88 files.

# 3. DATA ANALYSIS

The purpose of our analysis is to characterize the CDPD network and the behavior of its users. This characterization deals with network elements, network characteristics over time, user behavior, and cell activities.

## 3.1 Aggregate network characteristics

### *3.1.1 Network elements*

A total of 60 unique cell IDs were located in the billing record. There were 2,096 unique users (user IDs). Of these 2,096 users, only 1,730 actually generated IP data events. We note that it is not known whether the billing record actually covers the entire CDPD network, nor whether all network elements were active during the collection period. Hence, above numbers may represent minimum values.

Table 1 shows the breakdown by event type in the billing record. The number of deregistration events is an order of magnitude smaller than either the numbers of registration or IP data events. This is most likely because deregistration events are optional and because registration events may fail. Most users who generated large number of registration events failed to access the network and never succeeded in making a connection.

**Table 1. Breakdown of events by type.**

| Event type | Number of events | Percentage of total events |
|---|---|---|
| Registration | 619,268 | 39.19 |
| IP data | 889,227 | 56.27 |
| Deregistration | 71,741 | 4.54 |
| Total | 1,580,236 | 100.00 |

Table 2 shows the registration statistics and the possible reasons for rejected registrations. We notice a large proportion of suspicious users who may be trying to fraudulently use the network.

**Table 2. Distribution of registrations.**

| Number of registrations | | | | |
|---|---|---|---|---|
| Total | Accepted | Rejected | | |
| | | Suspicious users | No credentials | No reasons |
| 619,268 (100%) | 166,525 (26.8%) | 152,200 (24.5%) | 152,772 (24.6%) | 147,771 (23.8%) |

Another aspect of network behavior of interest is the percentage of aggregate discarded packets (Table 3). This percentage is fairly low, and it would be even lower if the control packets were included in the calculation.

**Table 3. Discarded packets.**

| Data packets | Discarded | Percentage |
|---|---|---|
| 21,046,695 | 237,712 | 1.13% |

*3.1.2 Number of network users over time*

The first aspect of the network's dynamic behavior that we investigated was the number of network users over time. Figures 2 and 3 show the time series. In Figure 2 (and in all subsequent figures) the discontinuity in the billing record is represented by zero number of users and it is marked (a gray star).

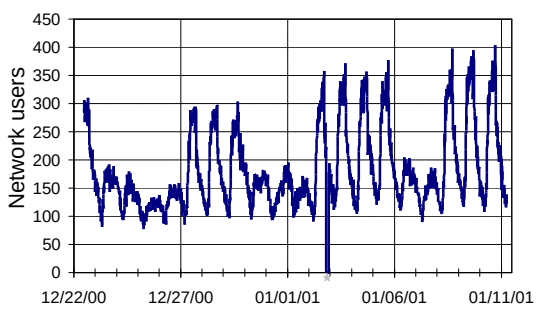


**Figure 2. Number of network users over time (December 22, 2000 to January 11, 2001).**

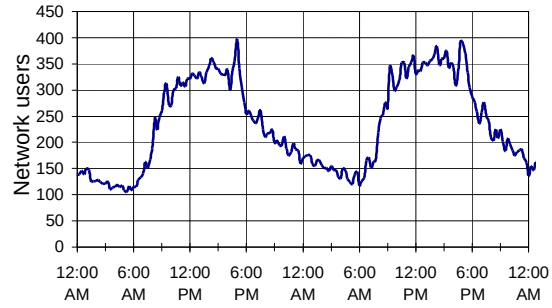


**Figure 3. Daily cycle of network users (January 8, 2000: 12 AM to January 10, 2001: 12 AM).**

Two cycles may be observed in Figure 2. The first is a daily cycle between peaks and valleys within a day. In general, the peaks occur around 5:00 PM and the valleys around 5:30 AM (Figure 3). The second cycle (weekly) is slightly irregular, but clearly visible: there are two ranges of visible peaks. The higher peaks tend to fall on weekdays and the smaller ones on weekends and holidays. We observed numerous low peaks in the first half of the collection period. Those additional low peaks occur between December 23 and December 26, and again between December 29 and January 1. The next pair of low peaks appears between January 6 and 7, which is the first "regular workweek" weekend of the year. It is possible that these irregularities are due to holiday seasons. Unfortunately, the billing records were not long enough to enable us to predict the usage pattern. Nevertheless, this reduced activity on weekends matches the behavior reported in past studies [3, 4, 5, 13, 14].

*3.1.3 Growth of users over time*

Figure 4 shows daily and cumulative number of network users during the collection period (twenty one days). The average number of unique users in the network is approximately 272 per day, with standard deviation 91.2. The number of the daily network users grew approximately by 7. We notice a decrease of the growth in number of users on weekends and on holidays. This is illustrated in Figure 4 with data plots falling under the regression lines.

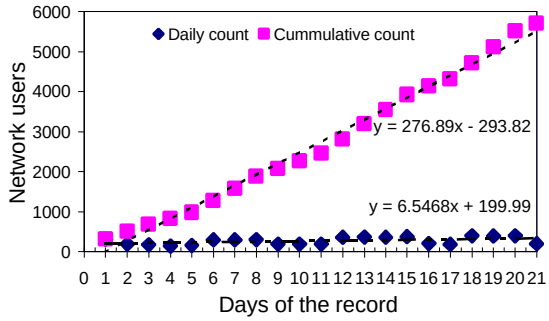


**Figure 4. Cumulative and daily growth of users (December 22, 2000 to January 11, 2001). Dashed lines represent regression lines.**

*3.1.4 Data traffic over time*

An alternative way of measuring the network utilization is to observe the number of data packets, data octets, and discarded packets over time. Figure 5 shows the number of data packets over time. The network seems to have more activity (peaks) throughout certain days. Note the increased activity at the beginning of the trace. Again, the number of data packets over time exhibits the same periodicities (daily and weekly) as the number of network users over time. The first cycle, shown in Figure 6 (depicting two days), is a cycling between peaks and valleys. Generally, the peaks occur around 5:00 PM and the valleys occur around 6:00 AM.

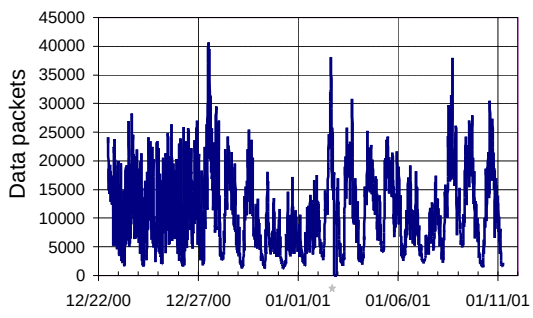


**Figure 5. Number of data packets over time (December 22, 2000 to January 11, 2001).**

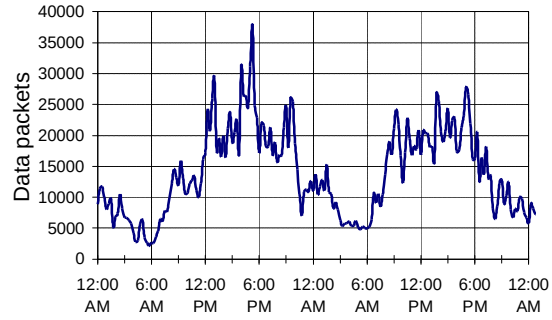


**Figure 6. Daily cycle of data packets (January 8, 2000: 12 AM to January 10, 2001: 12 AM).**

When data octets (as opposed to packets) are plotted over time (Figure 7), the increased activity at the beginning of the trace is absent. However, the two cycles appear again. When discarded packets are plotted over time (Figure 8), the two cycles that appear in Figures 5 and 7, disappear. Nevertheless, closer examination of the number of discarded packets over time shows that cycles are present and most peaks occur around 12:00 AM (Figure 9).

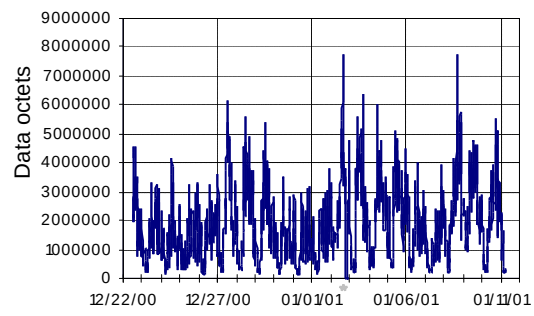


**Figure 7. Number of data octets over time (December 22, 2000 to January 11, 2001).**

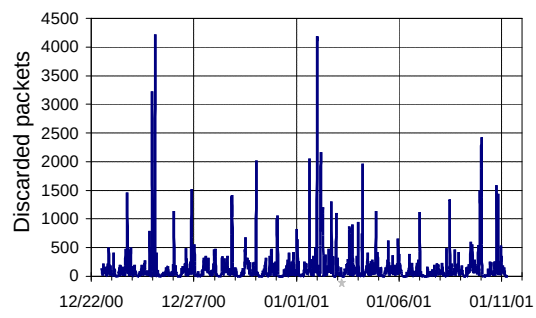


**Figure 8. Number of discarded packets over time (December 22, 2000 to January 11, 2001).**

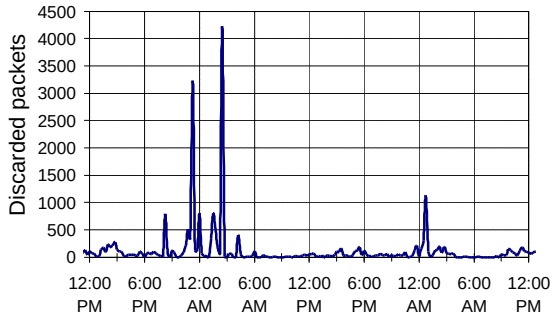


**Figure 9. Daily cycle of discarded packets (December 24, 2000: 12 PM to December 26, 2000: 12 PM).**

### *3.1.5 Snapshot of network characteristics over time*

Figure 10 illustrates the summary of network characteristics over time: number of unique users, data packets, data octets, control packets, control octets, and discarded packets. Figure 10 is a symmetric matrix of graphs. Their x and y variables are shown on the diagonal. It describes the inter-relationships between these variables. For example, the first graph in the first column shows the number of unique users vs. time. The first graph in the first row is its rotated mirror image. Among other observations, we notice the cyclic pattern of the number of users, data packets, and data octets over time. The proportionality between the number of control packets and control octets is also clearly observable.

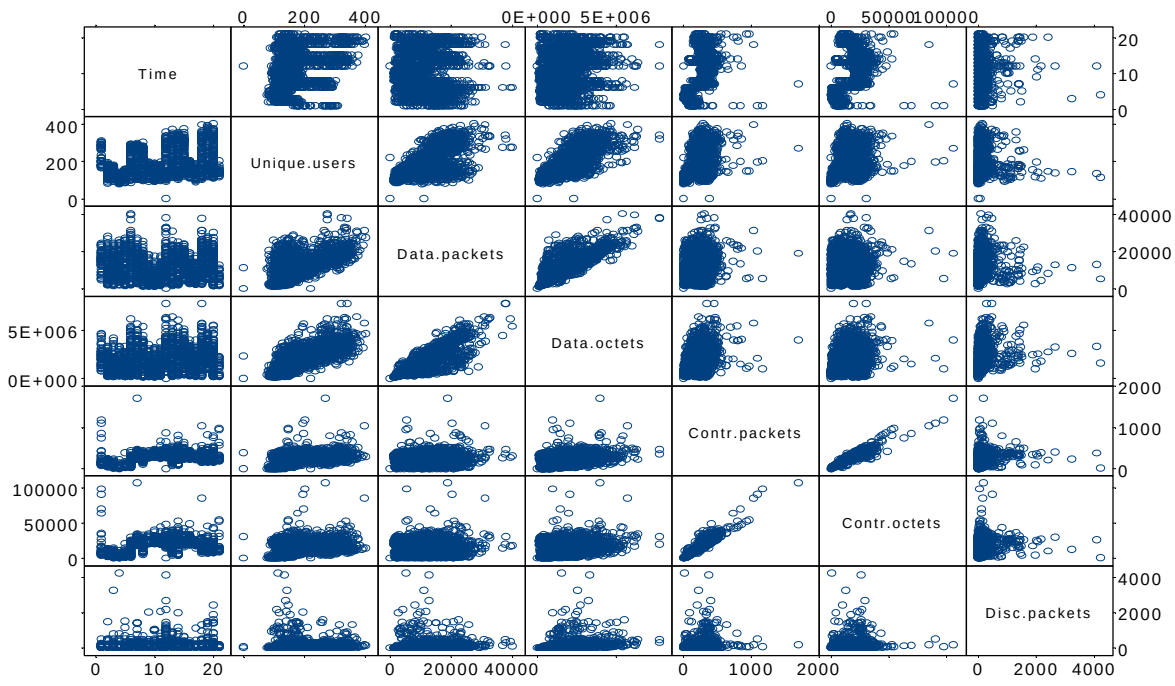


**Figure 10. Matrix of graphs summarizing network characteristics over time (number of network users, data packets, data octets, control packets, control octets, and discarded packets) and their inter-relationships.**

## 3.2 User behavior

### *3.2.1 Events per user*

In this section, we consider the events in the billing record by user ID (mobile NEI). Figure 11 shows the number of total events per user. In Figure 11, (and in all subsequent figures) the user IDs are not shown due to lack of space. It

can be seen that a few users (approximately ten) account for the majority of events. Figures 12, 13, and 14 show the breakdown of events per user by event type. Only a small number of users account for the majority of registration (Figure 12), deregistration (Figure 13), or IP data events (Figure 14). The disproportionate number of total events is due to the large number of registration events for a few users. There are users whose registrations are always accepted. Nevertheless, there are also 366 users who have registered with the network never gaining access to the system.

Several cases of unusual user behavior are shown in Table 4. These users account for 60.75% of all the registration events. User 61.131.154.9 alone accounts for 6.06% of the total number of events in the entire record. It might be that some users are trying to fraudulently gain access to the network because their user IDs differ from the IDs of authorized users. User IDs of subscribed users commence with 37, 59, or 91. For example, user 59.206.117.22 with the most deregistration events (Figure 13) has only one less registration event than the number of deregistration events (2,136 vs. 2,137) with all registrations accepted.

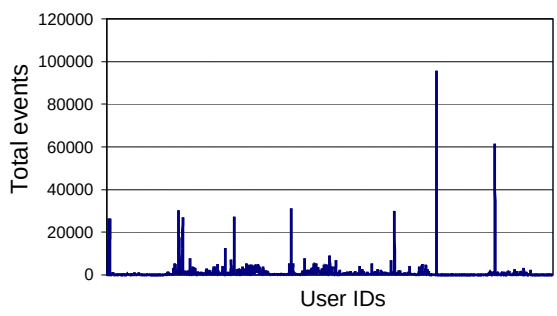


**Figure 11. Number of total events (sum of registration, deregistration, and IP data events) per user.**

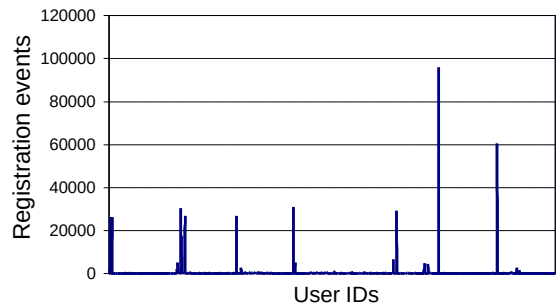


**Figure 12. Number of registration events per user. Approximately ten users account for the majority of events.**

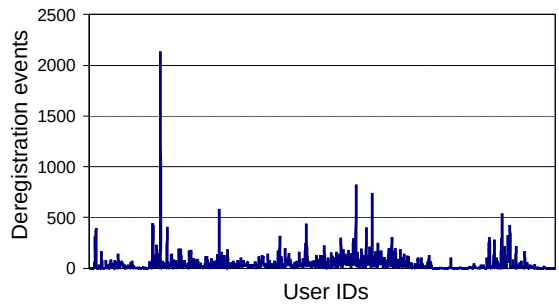


**Figure 13. Number of deregistration events per user. The user 59.206.117.22 has the most deregistration events (2,137).**

IP data events are shown in Figure 14. While there are still a few users that appear as outliers in the number of IP data events that they generate, events are far more evenly distributed.

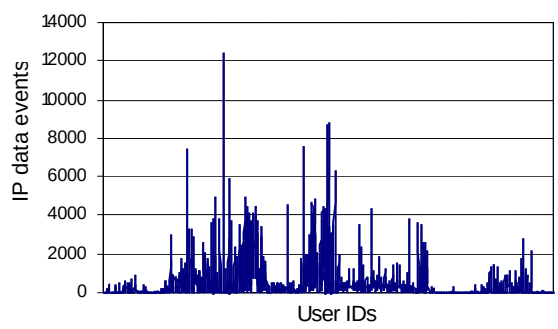


**Figure 14. Number of IP data events per user. Data events are far more evenly distributed.**

**Table 4. Users with unusual behavior.**

| User IDs | Regis. events | Deregis. events | Regis. accepted | Regis. accepted % |
|---|---|---|---|---|
| 59.206.117.22 | 2,136 | 2,137 | 2,136 | 100.00 |
| 102.133.238.133 | 2,190 | 0 | 0 | 0.00 |
| 91.171.227.151 | 2,611 | 0 | 0 | 0.00 |
| 59.207.111.139 | 2,895 | 0 | 0 | 0.00 |
| 59.207.215.20 | 4,623 | 56 | 56 | 1.21 |
| 59.206.117.28 | 5,024 | 7 | 380 | 7.56 |
| 59.207.215.117 | 4,899 | 1 | 8 | 0.16 |
| 59.207.206.47 | 5,225 | 0 | 0 | 0.00 |
| 59.207.211.210 | 6,720 | 9 | 9 | 0.13 |
| 59.206.118.178 | 12,328 | 0 | 0 | 0.00 |
| 247.33.244.20 | 26,421 | 0 | 0 | 0.00 |
| 59.206.118.180 | 26,901 | 0 | 0 | 0.00 |
| 59.206.127.247 | 27,019 | 37 | 37 | 0.14 |
| 59.207.211.23 | 29,373 | 50 | 74 | 0.25 |
| 59.206.118.147 | 30,299 | 0 | 0 | 0.00 |
| 59.207.206.36 | 31,042 | 0 | 0 | 0.00 |
| 91.171.159.54 | 60,598 | 101 | 102 | 0.17 |
| 61.131.154.9 | 95,838 | 0 | 0 | 0.00 |

### *3.2.2 Data traffic per user*

Data packets, octets, and discarded packets were plotted over time in Figures 5, 7, and 8. Figure 15 shows the number of discarded packets per user. The motivation is to locate users whose packets were unfairly discarded. Specific examples of inequity in terms of discarded packets are shown in Table 5. We notice that some users have over 10% of discarded packets, while other users have very low packet loss. These listed users represent the extremes in terms of packet loss. Note that user 59.207.209.71 sends almost five times more packets (with no loss) than user 59.206.125.155 (with more than 30% loss). Further examination of these two users' personal records may explain this discrepancy.

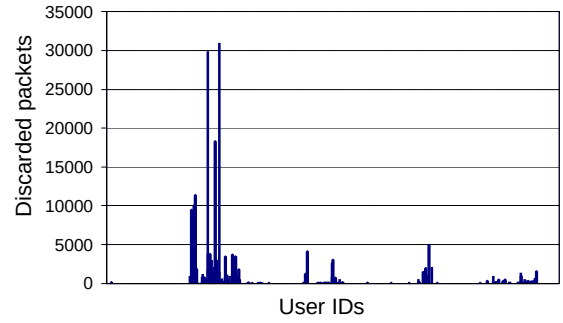


**Figure 15. Number of discarded packets per user. Discarded packets are unfairly distributed among users.**

**Table 5. Inequity of discarded packets per user.**

| User IDs | Data packets | Control packets | Disc. packets | % Disc. packets |
|---|---|---|---|---|
| Users with high data and low percentage of loss | | | | |
| 59.207.209.71 | 285,223 | 98 | 0 | 0.00 |
| 59.206.125.156 | 395,176 | 4,039 | 3,806 | 0.96 |
| 59.206.120.20 | 447,523 | 5,237 | 9,474 | 2.12 |
| 59.207.209.151 | 1,926,522 | 194 | 308 | 0.02 |
| 59.207.209.164 | 1,984,159 | 42 | 660 | 0.03 |
| Users with low data and high percentage of loss | | | | |
| 59.206.120.6 | 100,553 | 1,696 | 11,358 | 11.30 |
| 59.206.125.71 | 4,238 | 204 | 504 | 11.89 |
| 59.206.125.213 | 115,911 | 1244 | 18,282 | 15.77 |
| 59.206.125.208 | 8,244 | 251 | 2,158 | 26.18 |
| 59.206.125.245 | 74,767 | 2,181 | 30,966 | 41.42 |
| 59.206.125.155 | 60,342 | 2,586 | 29,871 | 49.50 |

### *3.2.3 Mobility of users*

Figure 16 shows the number of unique cells visited by each user. It shows that most users move around (at least a bit). The average number of cells visited is 9.41, with standard deviation 9.05, minimum 1, and maximum 47. Figure 17 shows the number of network users grouped by the number of visited cells. From the 2,096 recorded users, 78.58 % move between multiple cells. There were 449 stationary users (21.42%). This observation regarding the mobility of users matches the behavior found in similar studies of wireless network access by Tang and Baker [9] who found that 58% of the wireless users moved

between multiple locations during the data collection. Hutchins and Zegura [5] also observed 64% wireless users who moved between multiple access points.

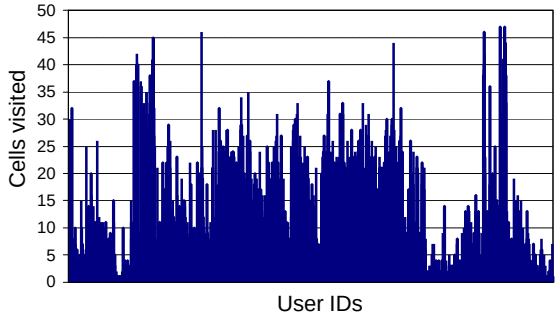


**Figure 16. Number of cells visited by users.**

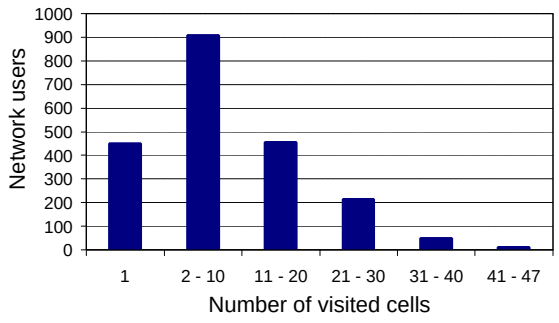


**Figure 17. Number of network users grouped by their range of mobility.**

*3.2.4 Clustering of users*

We used *k-means* clustering algorithm to find possible classification of users. *K-means* is an iterative algorithm, where the number of clusters *k* is supplied in advance [4, 9, 14]. We used eight variables when clustering network users: number of cell visited, data packets, control packets, discarded packets, IP data events, deregistration events, registration events, and total events. The total population is 2,096 users. Figure 18 shows a matrix of graphs that represents these variables and their inter-relationships. The best clustering results (Table 6) are those providing four distinct classes of users: *class 1* (96.7% of users) consists of customers with relatively low network usage, *class 2* (2.6% of users) consists of customers having medium network usage, *class 3* (0.5% of users) consists of customers having relatively high network usage, and *class 4* (only users 59.207.209.164 and 59.207.209.151) consists customers having exceptional behavior.

We also experimented with other clustering tools available in S-PLUS. *K-means* algorithm proved to be more efficient and produced results that best matched graphical analysis of the CDPD data.

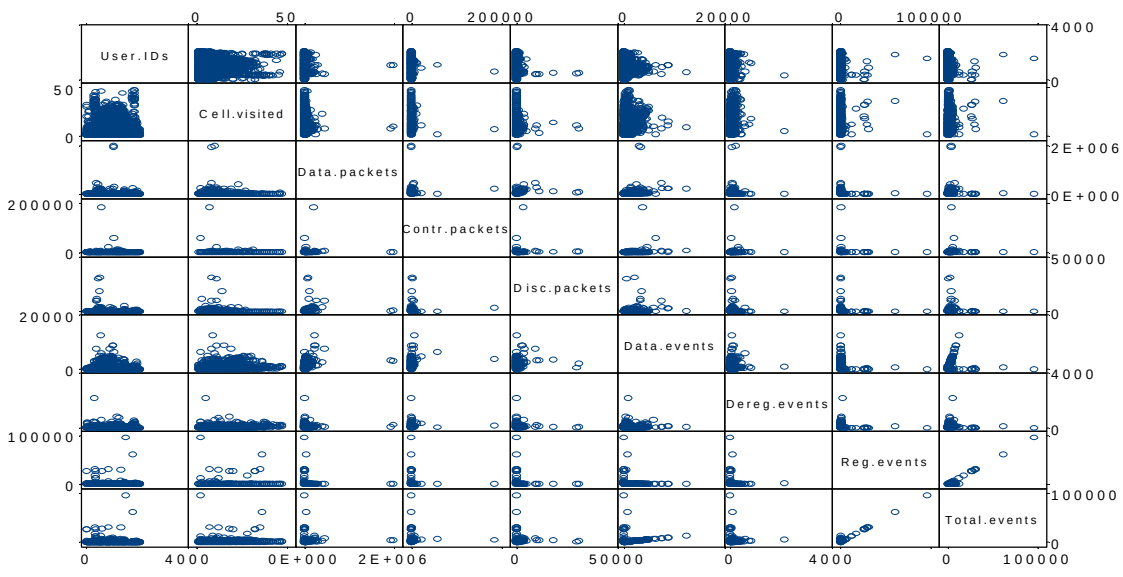


**Figure 18. Matrix of graphs showing the variables used in the clustering analysis of users. Data plots may give an insight for the possible number of classes to be used in *k-means* clustering algorithm. One class counting for only two users is easily observable.**

**Table 6. *K-means* clustering results for the classification of users (center values).**

| Cluster ID | Unique users | Cells visited | Data packets | Control packets | Disc. packets | IP Data events | Deregis. events | Regis. events | Total events | Network usage |
|---|---|---|---|---|---|---|---|---|---|---|
| 1 | 2,028 | 9.3 | 4,541.0 | 116.2 | 23.7 | 356.6 | 32.7 | 300.1 | 689.6 | Low |
| 2 | 55 | 11.4 | 90,768.0 | 391.3 | 2,734.7 | 1,775.3 | 72.7 | 153.7 | 2,001.8 | Medium |
| 3 | 11 | 10.7 | 267,608.0 | 19,138.0 | 3,469.4 | 5,678.7 | 101.0 | 167.4 | 5,947.2 | High |
| 4 | 2 | 9.0 | 1,955,340.5 | 118.0 | 484.0 | 3,226.0 | 143.5 | 379.0 | 3,748.5 | Very high |

## 3.3 Cell activities

### *3.3.1 Events per cell*

There are 60 unique cell IDs that appear in the billing record. In all subsequent figures, we show only 30 cell IDs due to lack of space. Figure 19 shows the total number of events per cell (cell IDs are ordered by their identification numbers). Figure 20 shows the breakdown by event type. A few cells are significantly busier than others. A geographical network diagram may have provided the reason for this behavior. The lack of correlation between different types of events is fairly visible.

In general, most events occurring in most cells are IP data and registration events. This indicates less than 50% overhead (in terms of events, not the actual bits).

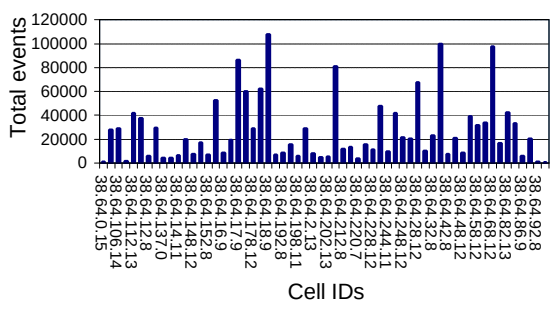


**Figure 19. Total number of events (sum of registration, deregistration, and IP data events) per cell.**

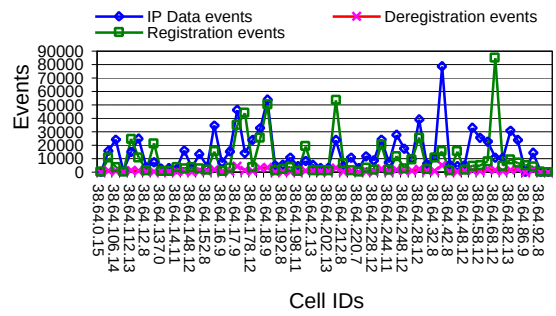


**Figure 20. Breakdown of events by type: registration, deregistration, and IP data events per cell.**

### *3.3.2 Cell user density*

Figure 21 shows the number of network users per cell. The distribution of users among cells varies. The average number of users in a cell is 329.05, with standard deviation 183.03, minimum 28, and maximum 860.

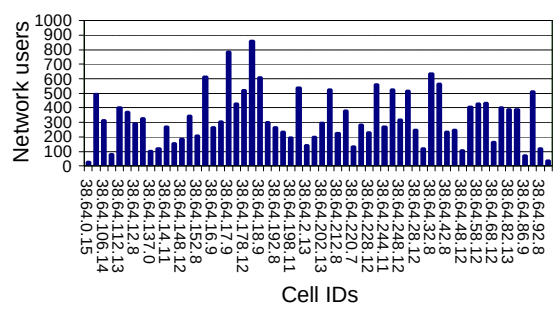


**Figure 21. Number of network users per cell.**

### *3.3.3 Clustering of cells*

We employed *k-means* clustering algorithm to classify cell activities using five variables: number of total events, IP data events, deregistration events, registration events, and

network users. Figure 22 shows these variables. The best clustering results (Table 7) produced three distinct classes consisting of 37 (62%), 15 (25%), and 8 (13%) cells.

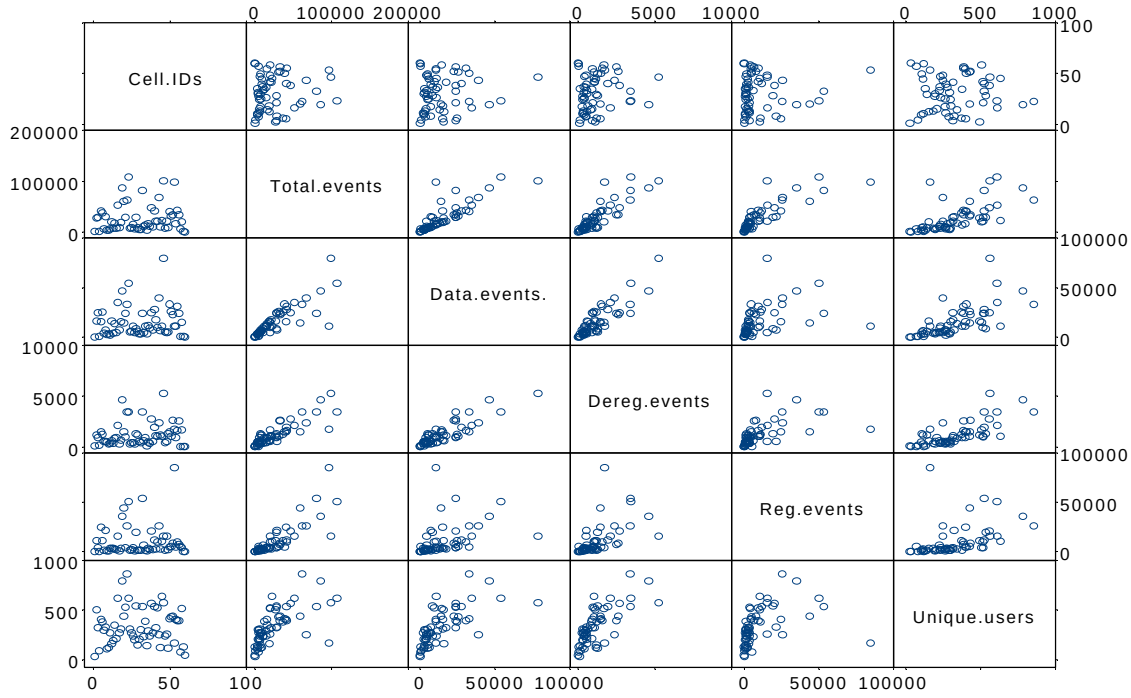


**Figure 22. Matrix of graphs showing the variables used in cell clustering analysis.**

**Table 7. *K-means* clustering results for the classification of cells (center values).**

| Cluster ID | Number of cells | Network users | Total events | IP Data events | Deregis. events | Regis. events | Cell activities |
|---|---|---|---|---|---|---|---|
| 1 | 37 | 238.67 | 10,115.57 | 6,624.486 | 596.567 | 2,894.514 | Low |
| 2 | 15 | 448.00 | 36,276.73 | 22,877.867 | 1,589.20 | 11,809.667 | Medium |
| 3 | 8 | 524.00 | 82,726.13 | 37,619.125 | 3,228.75 | 41,878.25 | High |

## 4. RELATED WORK

Most closely related to our work are studies of access wireless networks and user behavior. Hutchins and Zegura [5] described an analysis of data collected from a campus area network providing wireless 802.11b access capabilities. They presented an analysis of 138 days of session data, 54 days of transport layer flow data, and movement data taken from 109 wireless access points spread across 18 buildings. They observed linear growth in number of users and noticed daily behavioral cycles, with a peak usage time in early evening and minimum usage in early morning. Hutchings et al., [3, 4] also analyzed data sets of several months of RADIUS authentication data [10] taken from a large national dial-up Internet Service Provider. They noticed differences in number of network users on weekends and weekdays. Our findings in a CDPD network show similar trends.

Tang and Baker [14] described an analysis of a 12-week study of the local-area wireless network located at Stanford University. Their study

consisted of 74 users utilizing 13 wireless access points (within the single campus building) to campus and Internet resources. They found that more users are active on fewer days, while fewer users are active on many days. Tang and Baker [13] also presented an analysis of seven weeks of data from a metropolitan-area wireless network, Metricom's Ricochet Network service. Among other findings, they reported that more than half of the users in the study move between multiple locations during the data collection. Our observations are comparable.

## 5. CONCLUSIONS

Analysis of the CDPD billing records yielded some interesting data, despite the fact that the billing records contain only a high-level description of the network usage. Our analysis shows that CDPD network users have cyclic behavior. Two cycles of periodicity are observable. The first is a daily period between peaks and valleys. In general, the peaks occur around 5:00 PM and the valleys around 5:30 AM. The second cycle is a weekly cycle between two ranges of peaks: the higher peaks appear on weekdays and the smaller on weekends and holidays. A small subset of users accounts for a large number of rejected registration events, packet loss, and control packets. High number of rejected registrations may originate from users who attempt to use the network fraudulently. Clustering analysis with *k-means* algorithm revealed that it was possible to classify network users and cell activities. Analysis of longer billing records and using more powerful visualization and clustering tools may provide additional insights.

Although the billing record we examined may not be representative, our analysis provided useful information about the access to an operational wireless network. It may also be useful for modeling network traffic and for creating workload models for commercial wireless networks access. This type of analysis may also prove useful for better management of wireless networks.

## 6. ACKNOWLEDGMENTS

We would like to thank Telus Mobility for providing the billing records. Special thanks are due to S. Petrovic, K. McNair, and P. Chan for initial contributions to this research project.

## 7. REFERENCES

[1] Agosta, J., and Russell, T., *CDPD: Cellular Digital Packet Data Standards and Technology*. Reading, MA: McGraw-Hill, 1996.

[2] CDPD System Specification, Release 1.1, CDPD Forum, pp. 600-1, 600-2 and pp. 630-1 – 630-44, January 1995.

[3] Hutchins, R., Zegura, E. W., Liashenko, A., and Enslow, P. H., "Internet user access via dial-up networks – traffic characterization and statistics," in *Proc. Int. Conference on Network Protocols 2001,* Riverside, CA, USA, November 2001, pp. 314-322.

[4] Hutchins, R., Zegura, E. W., Liashenko, A., and Enslow, P. H., "Usage characteristics of dial-in Internet users: a national study," 2001: http://citeseer.nj.nec.com/458891.html.

[5] Hutchins, R., and Zegura, E. W., "Measurements from a wireless campus network," in *Proc. Int. Conference on Communications 2002,* New York City, NY, USA, May 2002.

[6] Java[tm] programming language: http://java.sun.com.

[7] Jiang, M., Nikolic, M., Hardy, S., and Trajkovic, Lj., "Impact of self-similarity on wireless data network performance," in *Proc. Int. Conference on Communications 2001,* Helsinki, Finland, June 2001, pp. 477-481.

[8] MySQL: http://www.mysql.com.

[9] Russell, S., and Norvig, P., *Artificial Intelligence: A Modern Approach*. Upper Saddle River, NJ: Prentice-Hall, 1995.

[10] Rigney, C., Rubens, A., Simpson, W., and Willens, S., "Remote authentication dial in user service (RADIUS)," Internet RFC 2138, April 1997.

[11] Sreetharan, M., and Kumar, R., *Cellular Digital Packet Data.* Norwood, MA: Artech House, 1996.

[12] S-PLUS: http://www.insightful.com.

[13] Tang, D., and Baker, M., "Analysis of a metropolitan-area wireless network," in *Proc. MOBICOM 1999,* Seattle, WA, USA, August 1999, pp. 13-23.

[14] Tang, D., and Baker, M., "Analysis of a local-area wireless network," in *Proc. MOBICOM 2000,* Boston, MA, USA, August 2000, pp. 1-10.

[15] Walrand, J., and Varaiya, P., *High-Performance Communication Networks*, 2nd edition, San Francisco, CA: Morgan-Kaufmann, 2000.