Using Network Activity Data to Model the Utilization of a Trunked Radio System

Nikola Cackov, Božidar Vujičić, Svetlana Vujičić, and Ljiljana Trajković
Simon Fraser University
Vancouver, Canada
{ncackov, bvujicic, svujicic, ljilja}@cs.sfu.ca

Keywords: emergency communications, trunked radio system, network utilization, traffic analysis.

Abstract
In this paper, we analyze the system utilization of a deployed radio communication network operated by a public safety agency. The network consists of a central site and a number of cells, each cell having a different capacity (number of available radio channels). The network is circuit-switched. Hence, the system utilization is a distribution of the number of concurrent calls in each cell vs. time. To determine the traffic load variations, we analyzed activity data from a sample week in 2002 and in 2003. We created an OPNET model based on the collected activity data and used the model to evaluate the utilization of system resources and to locate network bottlenecks. Our analysis may be used to address existing and future network congestion problems.

1. INTRODUCTION

The analysis of deployed networks is used to determine their performance and to identify and locate possible network congestion. The scope of the analysis and the parameters of interest depend on the network and its characteristics, such as technology, topology, and user behavior [1]–[3]. Furthermore, such analysis may be used to improve network reliability, which is essential for networks used by public safety agencies (police, fire department, and ambulance).

E-Comm, Emergency Communications for South West British Columbia Incorporated [4], provides radio communications to mobile users that belong to several public safety agencies. The wireless section of E-Comm's network has a cellular architecture. Each cell covers a certain geographical area within the Greater Vancouver Regional District (GVRD). Individual radio users access the network using trunking [5], [6], which implies sharing a set of frequencies (radio channels) among all agencies, rather than dedicating subsets of frequencies to individual agencies. The number of frequencies available in each cell is predefined and it determines the cell's capacity.

Prior work in the area of traffic analysis in trunked radio systems involved statistical analysis of the traffic [7], or creating a mathematical model based on queuing theory [8]. Here, we adopted a simulation approach to modeling the network utilization.

We used activity data collected by E-Comm to examine network utilization over a sample week in February 2002 and in March 2003. E-Comm activity data is a table that contains records of the network events. The relevant data were extracted into a format suitable for trace-driven network simulations. A network model was created using the OPNET network simulator [9]. We examined the instantaneous utilization of radio channels (the number of occupied radio channels) in each cell in order to determine how the amount of traffic changed between the two sample weeks and whether it would be possible to increase the number of network users.

We describe the E-Comm network in Section 2. The data model is given in Section 3, while the OPNET network model is presented in Section 4. Simulation results are reported in Section 5. We conclude with Section 6.

2. E-COMM NETWORK

E-Comm employs the Enhanced Digital Access Communication System (EDACS) [10]. EDACS has a complex architecture consisting of various interconnected network elements: data and PBX gateways, radio transceivers and repeaters, dispatch and management consoles, and network switches. A simplified schematic diagram of the system is shown in Fig. 1.

We consider only the circuit-switched network segment that carries user information among mobile users. It consists of 11 cells connected to a central digital switch. Each cell has a number of radio repeaters capable of transferring data using a set of frequencies. Repeaters belonging to one cell use an identical set of frequencies. This method of radio transmission is known as simulcast. The number of frequencies in each cell is predefined and it determines the number of available radio channels. Each radio channel uses one frequency. Thus, we can define the capacity of a cell as a number of radio channels in the cell. This number also determines the maximum number of simultaneous radio transmissions in a given cell. Every radio transmission is treated as one call. Each call in a given cell occupies one radio channel.

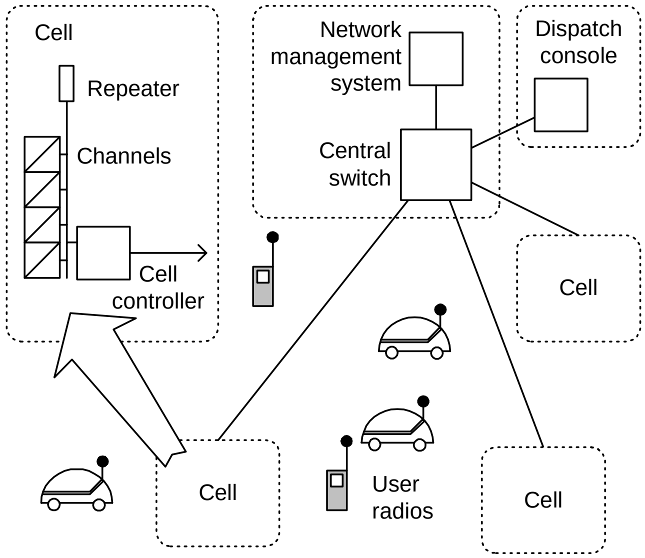

Fig. 1. General structure of a trunked radio system.

In each cell, one frequency is dedicated to the exchange of control information before, during, and after the call. Thus, the capacity of a cell (number of user channels) is one less than the number of available frequencies. No protocol and traffic data for the control channel were available. Hence, we only analyzed the utilization of user channels.

Network users are organized into talk groups. A call is established by using a push-to-talk mechanism. A user (member of a talk group) talks to other members in the talk group by pressing the push-to-talk button on the mobile radio transceiver. The system then determines the locations of the talk group members and checks for availability of radio channels in every cell where the members are located. If there is at least one free channel in every cell, the call is established. The one-way communication (call) lasts as long as the initiator holds the push-to-talk button. If there are no available channels in at least one cell with members of the talk group, the call is queued. The call is discarded if it cannot be established after a certain period of time. We found that the number of queued calls is negligible (< 0.5 %) compared to the number of established calls.

3. DATA MODEL

Activity data from the deployed network recorded at E-Comm consist of records of network events: established, queued, and dropped calls, as well as talk group dynamics. In order to create a data model, we extracted two weeks of data, called "sample data", from E-Comm's records. The 2002 sample data spans the week between 8:00 on February 1, 2002 and 8:00 on February 8, 2002. The 2003 sample data spans the period starting at 0:00 on March 20, 2003 and ending at 24:00 on March 26, 2003. The accuracy of the data is 10 ms for call durations. The resolution of the timestamps is 1 s and 10 ms for the 2002 and 2003 sample data, respectively. In order to analyze the utilization of radio channels in individual cells, the sample data were filtered. The data models, formatted as trace files, were used as input for the trace-driven OPNET simulations. Fig. 2 illustrates the stages of the data transformation.

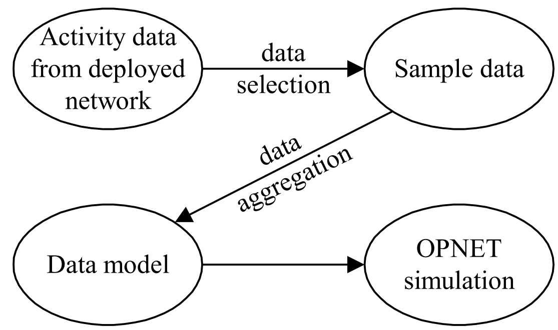

Fig. 2. Creating the data model. Resolution of time stamps in the activity data is 10 ms. Resolution for sample data is 1 s and 10 ms for 2002 and 2003, respectively.

From the sample data, we extracted only the records relevant to established voice calls. They indicate the caller, the called talkgroup, the time of a call, and how long a channel in a given cell was occupied. An excerpt from the 2003 sample data showing only the relevant fields is given in Table 1. Each user has a unique user ID and each talk group has its unique identification number. In this example, all four rows correspond to one call. The call began at approximately 0:00:10 on March 20, 2003. It lasted 4,870 ms, and involved cells 4, 8, 9, and 10. (To maintain confidentiality of the data, the ID's of the caller and the callee were relabeled A and B, respectively.) All timestamps end in either 0 or 9. Hence, the actual resolution of the timestamps is 10 ms. One call can be represented by one or several rows in the sample data. The number of rows depends on the number of cells where the call terminates. A call should be uniquely identified by four fields in the table: timestamp, duration, caller (always the ID of the user who initiated the call), and callee (usually ID of a talkgroup that receives the call). Nevertheless, as shown in Table 1, there are discrepancies in the records. In many cases of the sample data, timestamps that correspond to a single call differ by up to 1 s, while the call durations differ by up to several tens of milliseconds. For the data model, we arbitrarily chose the smallest timestamp. The largest call duration was chosen in order to simulate the worst-case scenario. We also modified

Table 1. Excerpt from the 2003 activity table.

Timestamp	Duration (ms)	Caller	Callee	Cell
2003-03-20 0:00:10.639	4,870	A	B	4
2003-03-20 0:00:10.599	4,830	A	B	8
2003-03-20 0:00:10.529	4,860	A	B	9
2003-03-20 0:00:10.510	4,870	A	B	10

the format of the timestamp. The original timestamp represented the date and time of the beginning of a call. For simulation purposes, it was convenient to express the timestamp as a difference between the original timestamp and an arbitrary reference time. The reference times were chosen to be 8:00 on February 1, 2002 and 0:00 on March 20 in the 2002 and 2003 data models, respectively. In order to create trace files for OPNET simulations, we modified the sample data so that one row corresponds to one call. As a result, one record in a trace file (our data model) that corresponds to the four rows of data shown in Table 1 is: {10,510, 4,870, 4, 8, 9, 10}, where 10,510 is the timestamp in milliseconds, 4,870 is the duration of the call in milliseconds, and the remaining numbers are the cell IDs where the call terminated.

4. NETWORK MODEL

We used the OPNET [9] network simulator to analyze the data models. OPNET network models have a hierarchical architecture with three layers: network, node, and process. Network topology represents the top layer of the OPNET network model and consists of network nodes. Nodes consist of interconnected modules that perform defined tasks and exchange information using packet streams and statistical wires. The functionality of each module is defined using a process model (finite-state machine). The process model (written in C and C++ and using various OPNET specific functions) describes the behavior of the finite-state machine and defines the transitions between states.

4.1 OPNET Network Model

The OPNET network layer model, shown in Fig. 3, consists of a central switch and eleven cells located in various regions of GVRD. The cells are connected to the switch via point-to-point simplex links. In the deployed system, after call establishment, voice information flows from the originating cell to the central site and then to the destination cell(s). The call establishment procedure is performed by exchanging information in the control channel. Our data model does not contain information about the originating cell of a call, and, therefore, all traffic in the OPNET model is generated in the central site and is then sent to the corresponding cells.

Each link has a number of channels equal to the number of frequencies available in a cell. Table 2 shows the cell capacities in terms of number of user channels. One occupied frequency in a cell corresponds to one busy channel in the link that connects the cell with the central site. Therefore, in our simulations, we observed the number of used frequencies (the instantaneous utilization of radio resources) by monitoring the utilization of each point-to-point link (number of occupied channels). All channels in a link have an identical bit rate (1,000 bits per second in our simulations). When a call is forwarded to a cell, the central site generates a packet of 1,000 × *CD* bits, where *CD* is the duration of the call in seconds. The packet is sent to an idle channel in the link that connects the central site and the cell. The channel in the corresponding link will be occupied exactly *CD* seconds, starting from the moment the call is established.

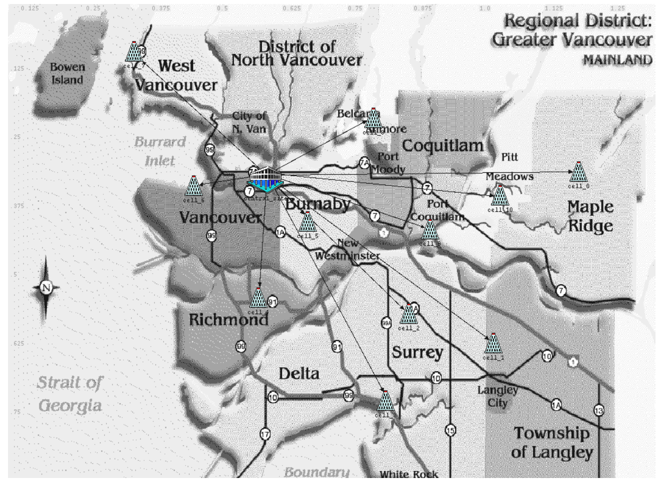

Fig. 3. Topology of E-Comm's network modeled in OPNET [11].

Table 2. Number of user channels per cell.

Cell	1	2	3	4	5	6	7	8	9	10	11
Channels	12	7	4	5	3	7	6	4	6	6	3

4.2 OPNET Node and Process Models

We created OPNET node and process models for several elements of E-Comm's network. We used standard OPNET process models for point-to-point transmitters, point-to-point receivers, and packet sinks.

The OPNET network model consists of two types of nodes: central site and cell. The node model of the central site (network switch) is shown in Fig. 4. Its functions are: reading the trace file, generating packets that correspond to calls, sending the packets to appropriate cells, and collecting statistics. These functions are implemented in the modules that constitute the central site node model: *source*, *dispatcher*, *channel_selector*, and *tx* (point-to-point transmitter), shown in Fig. 4. There is one *source* module, one *dispatcher* module, and there are eleven pairs of *channel_selector* and *tx* modules (one pair for each cell). The *source* module reads the trace file and forwards information (call duration and destination cells) about the calls that need to be established to the *dispatcher* module. The central module in the node model is the *dispatcher*. Its process model is shown in Fig. 5. It consists of four states. In the *init* state, initialization of the statistics that are collected during the simulation is performed. After the *init* state, the process proceeds to the *idle* state. When the *dispatcher* receives a notification from the *source* that a call needs to be established, it proceeds to *call* state. In this state, it checks for availability of free channels in the cells and it decides whether or not the call can be established. If the call

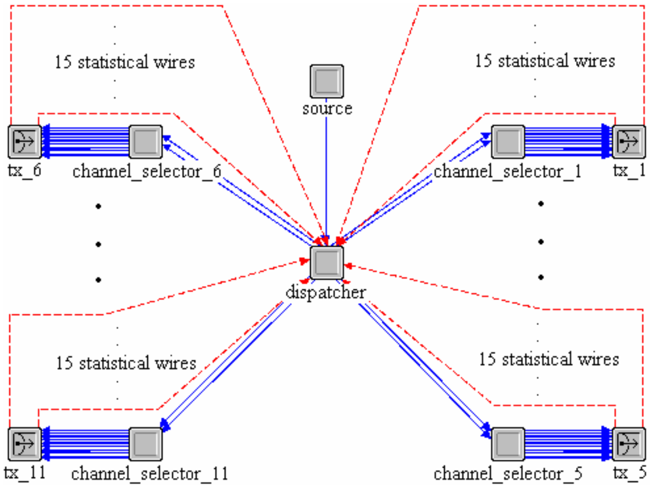

Fig. 4. Node model of the central site.

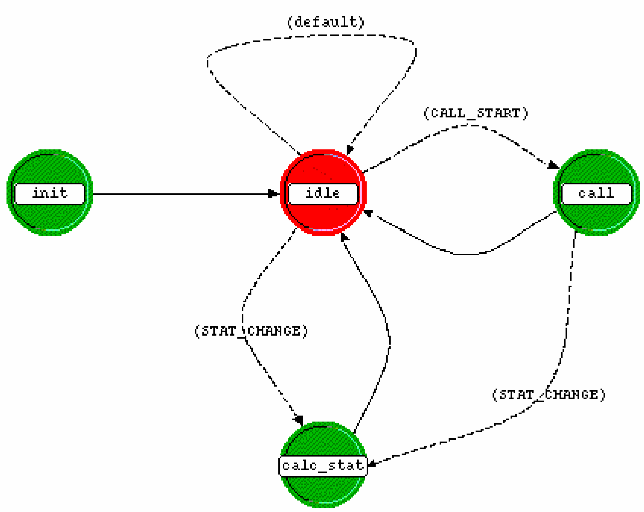

Fig. 5. Process model of the dispatcher module in the central site node model.

can be established, the process creates a packet of a length proportional to the duration of the call and forwards it to the corresponding *channel_selector* module(s). If the call cannot be established, the number of discarded calls is updated and no further actions are performed. After this state, the process returns to *idle* state. The *dispatcher* is connected to each transmitter *tx* by statistical wires that monitor the channel occupancy in each link. One statistical wire monitors a single channel. When the *dispatcher* receives a notification via a statistical wire that the status of a channel has changed, it proceeds to the *calc_stat* state. There, the values of the collected statistics are updated and the process returns to the *idle* state. Each *channel_selector* module registers free and occupied channels in its connected link. When a packet from the *dispatcher* arrives, the *channel_selector* sends the packet via one of the free channels and marks the channel as busy. When a cell receives a packet, which is equivalent to a call being completed, the *channel_selector* marks the corresponding channel as free.

The node model of a cell is shown in Fig. 6. It consists of a point-to-point receiver *rx*, a *receiver* module, and a *sink*. When a packet arrives, the *receiver* module notifies the corresponding *channel_selector* in the central site of the free channel in the link and sends the packet to the *sink*.

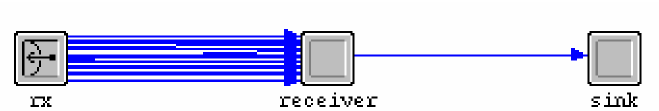

Fig. 6. Node model of the cell.

5. OPNET SIMULATION RESULTS

We simulated two weeks of network activity. Simulation results for the 2002 and 2003 data are shown in Figs. 7 and 8, respectively. The horizontal axis in both figures represents time and it is common for all 13 graphs. In Fig. 7, the first tick (marked 0d) corresponds to 8:00 on February 1, 2002, while the last tick (marked 7d) corresponds to 8:00 on February 8, 2002. In Fig. 8, the first tick (marked 0d) corresponds to 0:00 on March 20, 2003, while the last tick (marked 7d) corresponds to 24:00 on March 26, 2003.

The first eleven graphs in Figs. 7 and 8 show the number of occupied radio channels. They are named "Occupied channels [*x*]", where *x* corresponds to the cell ID (1 to 11). For example, "Occupied channels [1]" shows the number of busy radio channels in cell 1. As shown in Fig. 7, there were no busy channels in cell 1 during the first three days because there were no user activity data for cell 1 in the 2002 sample data for that period. The second graphs from the bottom in both Figs. 7 and 8 correspond to the statistic called "Call drop". It is equal to 1 when a call is discarded. The bottom graphs in Figs. 7 and 8, called "Discarded calls", show the total number of discarded calls over time.

The "Occupied channels [*x*]" graphs reflect the presence of diurnal cycles in the activity data. The minimum number of used channels is observed at approximately 2 PM every day, while maximum utilization is reached between 9 PM and 3 AM. As expected, call drops are experienced during periods of high utilization. The simulation results from the 2002 data, shown in Fig. 7, indicate that most cells seldom reach their capacity. Cell 5, with a capacity of 3 channels, is an exception and often has all available channels occupied. Cells 4 and 6 have all channels occupied only once during the 2002 sample week. Two cells (cells 1 and 10) never reach their capacity. Fig. 8 shows simulation results from the 2003 data. During the busy hours several cells are operating at their full capacity. (Cell capacities are given in Table 2.) For example, cells 2, 4, 7, and in particular cell 9 have all their available channels occupied during the periods of high utilization. However, cell 1 has at least one free channel at all times. The average number of used channels in every cell in both 2002 and 2003 sample

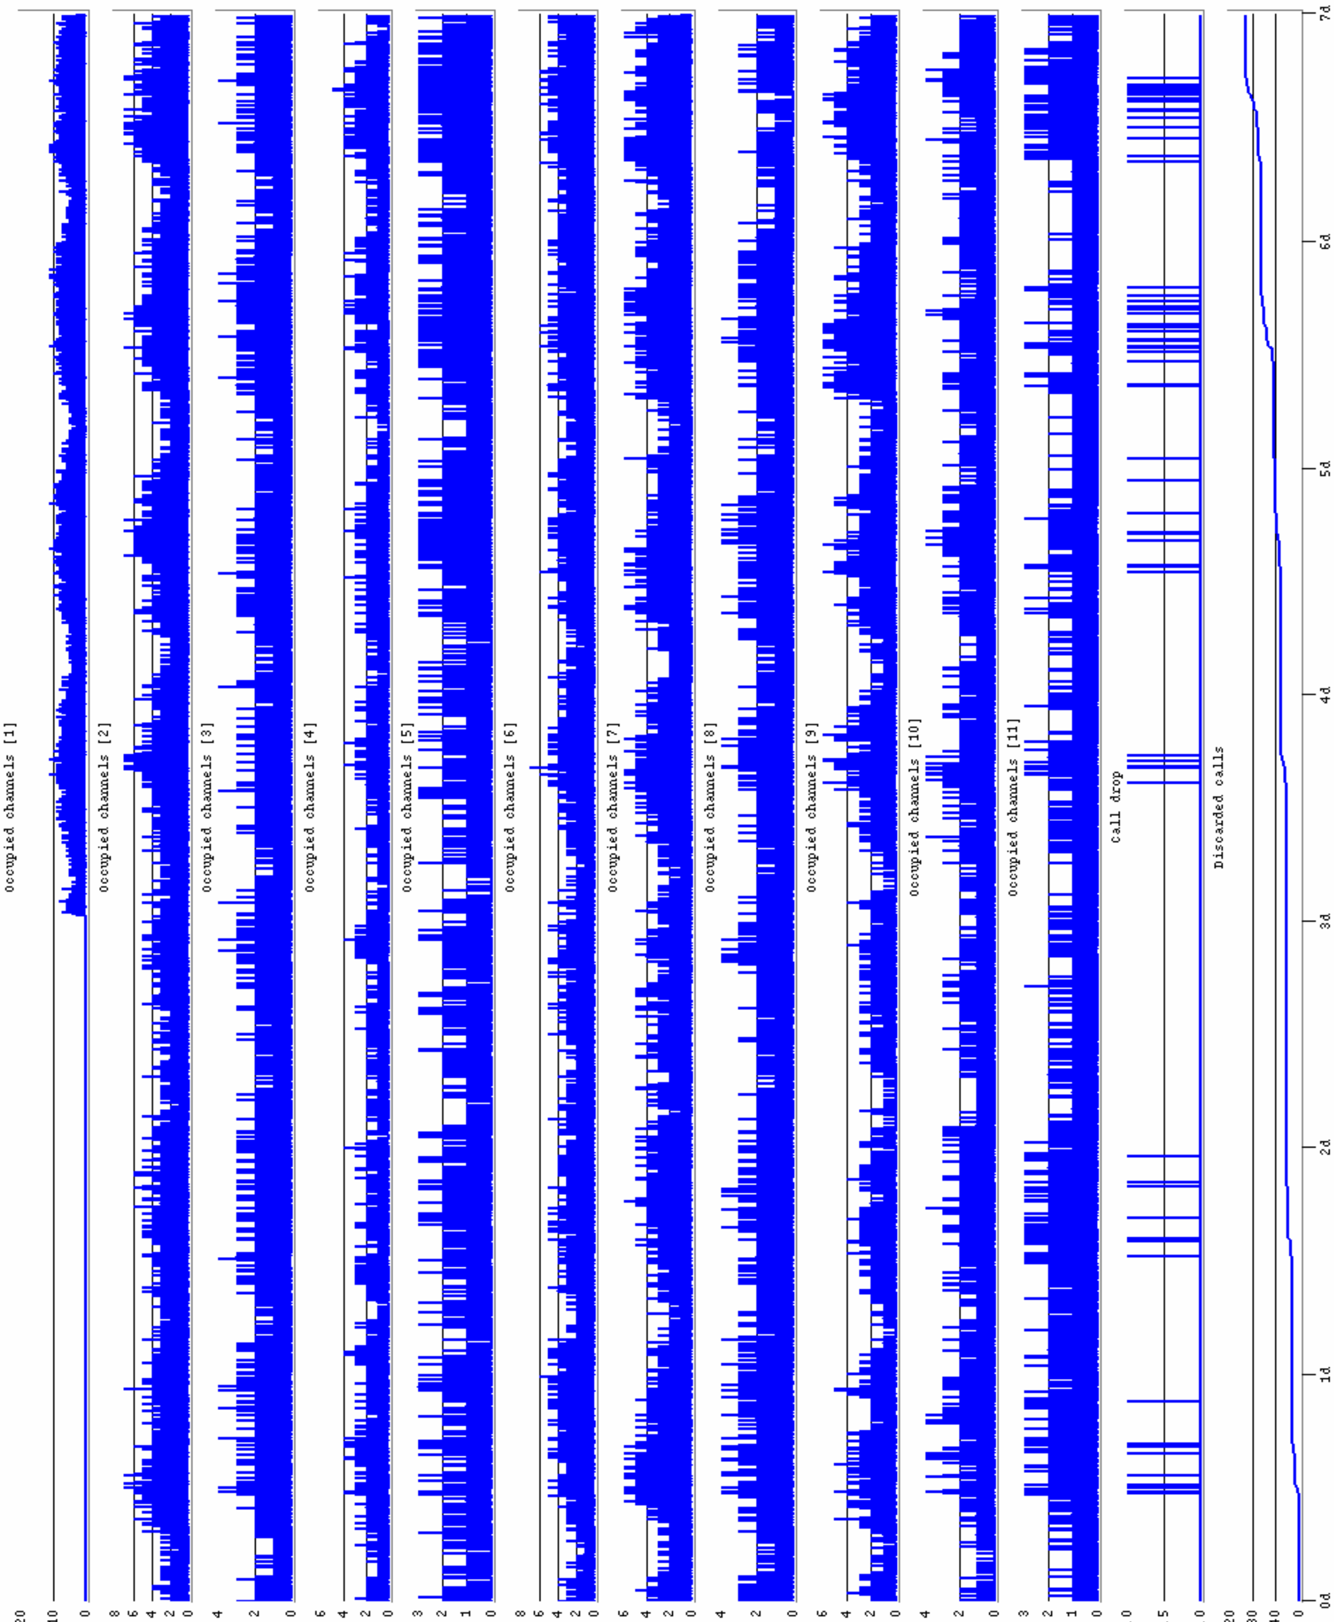

Fig. 7. Statistics collected during the simulation of the week of network activity starting at 8:00 on February 1, 2002 and ending at 8:00 on February 8, 2002. “Occupied channels” ([1], [2], …, [11]) show the utilization of each cell (number of occupied radio channels). “Call drop” indicates the moments when calls are discarded. “Discarded calls” statistic is the cumulative number of discarded calls.

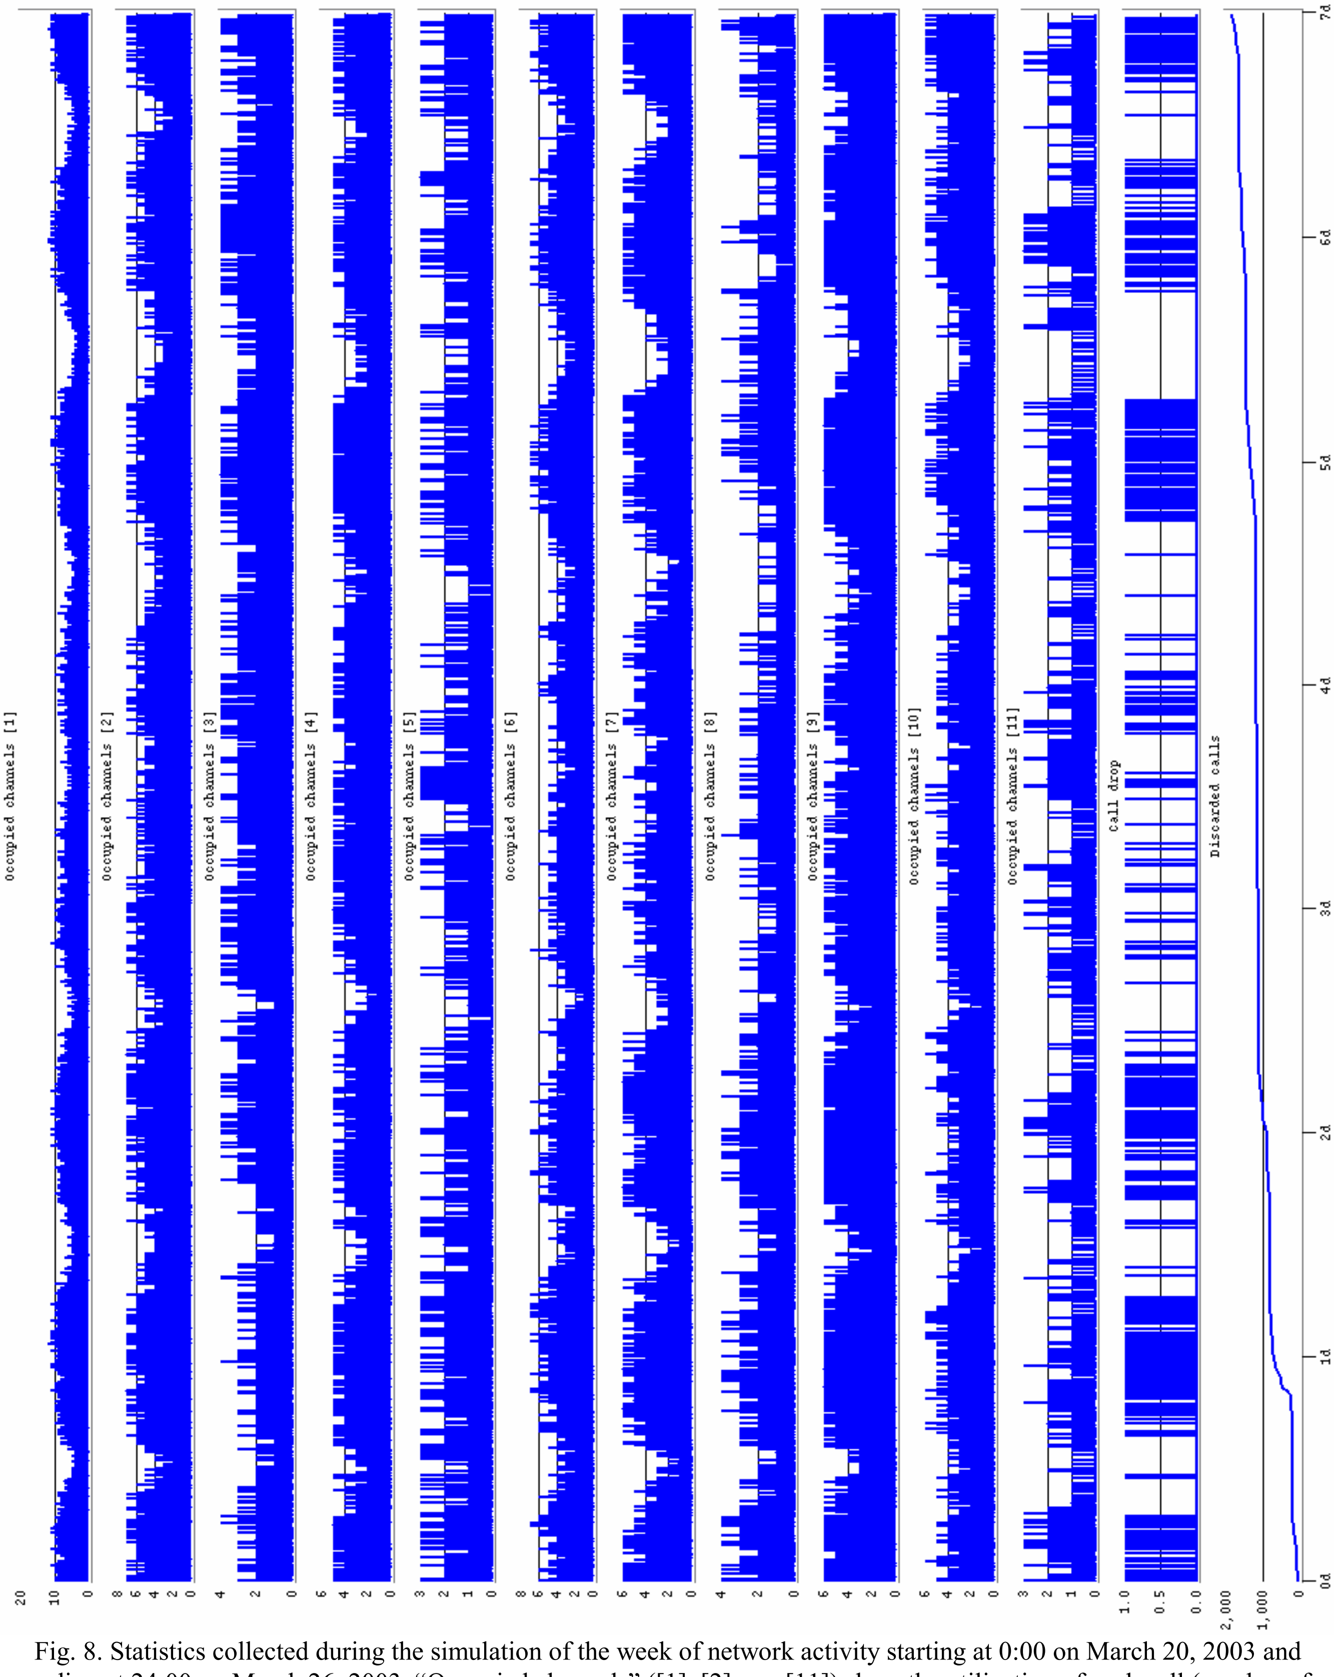

Fig. 8. Statistics collected during the simulation of the week of network activity starting at 0:00 on March 20, 2003 and ending at 24:00 on March 26, 2003. “Occupied channels” ([1], [2], …, [11]) show the utilization of each cell (number of occupied radio channels). “Call drop” indicates the moments when calls are discarded. “Discarded calls” statistic is the cumulative number of discarded calls.

data is small compared to its capacity. This is to be expected because of the time variability of the traffic and the design requirement that the system meets certain quality of service criteria during busy hours. The maximum and average numbers of used channels for each cell for both sample weeks are given in Table 3. The average number of used channels in the cells increased from 2002 to 2003, except for cell 11. The average utilization of cell 10 increased by a factor of five, while the average utilization of cell 11 remained unchanged.

Table 3. Maximum and average number of used channels.

	2002		2003	
Cell	Max	Average	Max	Average
1	11	2.5	11	2.6
2	7	0.8	7	1.6
3	4	0.3	4	0.5
4	5	0.3	5	1.1
5	3	0.2	3	0.3
6	7	0.7	7	1.2
7	6	0.7	6	1.1
8	4	0.3	4	0.4
9	6	0.4	6	1.6
10	4	0.2	6	1.0
11	3	0.2	3	0.2

Our model does not match the behavior of the deployed system with respect to discarded calls. The trace file (data model) was created from the sample data by considering only the established calls. Hence, we expected no discarded calls during simulations. However, the simulation results showed 91 and 1,812 discarded calls in the 2002 and 2003 results, respectively. The total number of calls in the 2002 sample week was 403,590, while in 2003 it was 645,167. The number of discarded calls in OPNET simulations is small compared to the total number of calls and, hence, it does not affect the network utilization.

Discarded calls in the simulation results dealing with sample data from 2002 are due to the coarse resolution of the timestamps in the sample model. The original timestamp resolution in the activity data from the deployed network is 10 ms. Both the sample data and the data model have a timestamp resolution of 1 s. Hence, simulation of the network's behavior did not precisely match the behavior of the deployed system. Due to rounding errors, our model may have larger utilization (number of used channels) than the deployed system at the same time instance, as illustrated in Fig. 9.

An example of a channel utilization pattern in an arbitrary cell is shown in Fig. 9 (dashed lines). One call starts at 0 s and lasts 300 ms. Another call starts at 0.5 s and lasts 400 ms. There are no other calls in the time interval between 0 s and 1 s. Hence, one channel is busy from 0 s to 0.3 s, all channels are free from 0.3 s to 0.5 s, and again one channel is busy from 0.5 s to 0.9 s. However, the resolution of the timestamp in the data model is 1 s. In case the timestamp is rounded down to the nearest integer number of seconds, the channel utilization pattern is shown in Fig. 9 (solid line). In this case, both calls begin at 0 s and, hence, there are two busy channels. The number of busy channels becomes one when the call that lasts 300 ms ends. Finally, at 0.4 ms all channels are free. If the cell had only one available channel, in the real system both calls could be established. Nevertheless, in the model, one of the calls would be discarded.

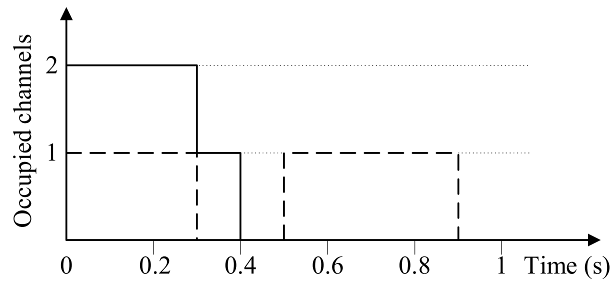

Fig. 9. Example of two calls that could be established in the deployed system using one channel (dashed lines) but require two channels in the model (solid line).

Discarded calls in the simulation results dealing with sample data from 2003 are due to certain discrepancies in the sample data. One discrepancy, described in Section 3, is that records corresponding to one call may have multiple values for the timestamp and the duration. The trace file (data model) was created by taking the largest value for the call duration. Hence, it is possible for the model to exhibit larger utilization than the deployed network. Another discrepancy is the existence of records showing overlapping usage of radio channels in the cells. One channel in a given cell can be occupied by only one call at a time. As indicated in Table 4, channel 4 in cell 10 is busy with a call lasting from 0:00:33.370 until 0:00:42.790 (= 0:00:33.370 + 9.420). However, there is a record corresponding to a call that originates at 0:00:42.769 occupying the very same channel 4 in cell 10. These records of overlapping usage of channels are due to errors in the recording of the activity data in the deployed network. In our model, the calls in cell 10 shown in Table 4 are established using two channels instead of one.

Table 4. Overlapping usage of channels.

Timestamp	Duration (ms)	Cell	Channel
3/20/03 0:00:33.370	9420	10	4
…	…	…	…
3/20/03 0:00:42.769	4290	10	4

In general, the collected simulation statistics cannot indicate which cell(s) caused a call to be discarded. If all available channels in two or more cells are occupied at the time when a call is discarded, the call may have been discarded from any of the maximally utilized cells. The busiest cells are most likely to cause call drops. From graphs shown in Figs. 7 and 8, it is possible to identify cells that are operating near the full capacities given in Table 2. For example, Fig. 7 shows that cell 5 has all its channels occupied during busy hours, when most of the call drops occur. Similarly, from the simulation results for the 2003 sample week shown in Fig. 8, cells 4 and 9 may be identified as the busiest cells. When the capacities of cell 5 (for the 2002 data) and cells 4 and 9 (for the 2003 data) were increased by one channel, while maintaining other cells' capacities the same as in Table 2, the number of discarded calls decreased significantly, as shown in Table 5. This implies that the originally identified busiest cells indeed contributed to most of the call drops.

Although our data model consists only of established calls, the discarded calls can be used to locate network bottlenecks. They are indicators of network congestion and occur at moments of high network utilization. They also may indicate future network congestion problems if the traffic load of busy cells increases.

Table 5. Simulation results for five simulation runs and various cell capacities.

Sample data	Cell no.	Capacity	No. of discarded calls
2002		Table 2	91
2002	5	3 + 1	62
2003		Table 2	1,812
2003	9	6 + 1	679
2003	4	5 + 1	521
	9	6 + 1	

Table 3 and Figs. 7 and 8 illustrate the difference in network traffic loads in February 2002 and March 2003. The number of calls increased by ~ 60 %. This increase was not uniform among the cells. Several cells (4, 9, and 10) experienced a larger increase in the traffic load than others (1, 7, 8, and 11). If the same trend of non-uniform increase of traffic load continues, cells 4, 9, and 10 may cause queued and discarded calls. As shown, these queued and discarded calls can be avoided by increasing the capacities of those cells.

6. CONCLUSIONS

In this paper, we used collected activity data to analyze a deployed, operating network. We simulated two weeks of network activity using the OPNET simulator. Simulation results show diurnal cycles in the network utilization. We also observed that the average number of busy channels in most cells is small compared to their capacities, with busy periods of high utilization. The data model used for the trace driven simulations included only established calls. Because of the coarse granularity of the data model and certain discrepancies in the sample data, OPNET simulation results indicated a small number of discarded calls that were not present in the deployed system. The discarded calls did not affect the network utilization. Between February 2002 and March 2003, the number of calls increased by nearly 60 %. This increase caused several cells that were underutilized in 2002 to be used near their full capacity in 2003. Hence, our analysis may be used to address existing and future network congestion problems.

ACKNOWLEDGMENT

The authors thank D. Sharp from Planetworks and the management and technical staff at E-Comm for providing access to the activity data and technical support for data analysis. This research was supported by the NSERC Grant 216844-03 and Canada Foundation for Innovation.

REFERENCES

[1] L. A. Andriantiatsaholiniaina and L. Trajković, “Analysis of user behavior from billing records of a CDPD wireless network,” in *Proc. Workshop on Wireless Local Networks 2002,* pp. 781-790, Tampa, FL, Nov. 2002.

[2] D. Tang and M. Baker, “Analysis of a local-area wireless network,” in *Proc. of ACM Mobicom 2000,* pp. 1-10, Boston, MA, Aug. 2000.

[3] D. Tang and M. Baker, “Analysis of a metropolitan-area wireless network,” in *Proc. of ACM Mobicom ’99,* pp. 13-23, Seattle, WA, Sept. 1999.

[4] E-Comm, Emergency Communications for SW British Columbia Incorporated. (2004, May). [Online]. Available: http://www.ecomm.bc.ca.

[5] R. J. Orsulak, R. R. Seach, J. P. Camacho, and R. J. Matheson. (2004, May). “Land mobile spectrum planning options,” National Telecommunications and Information Administration, Washington, DC, Spectrum Engineering Reports, Oct. 1995. [Online]. Available: http://www.ntia.doc.gov/osmhome/reports/slye_rpt/cover.html.

[6] K. Zdunek, “Design considerations for trunked radio systems,” in *Proc. of IEEE Vehicular Technology Conference*, pp. 194-201, Dallas, TX, May 1986.

[7] P. Cohen, D. Haccoun, and H. H. Hoang, “Traffic analysis for different classes of land mobile communication systems,” in *Proc. of IEEE Vehicular Technology Conference*, pp. 283-285, Toronto, ON, May 1983.

[8] H. H. Hoang, R. Malhame, and C. Rosenberg, “Communication load and delay in multichannel land mobile systems for dispatch traffic: A queuing analysis,” in *Proc. of IEEE Vehicular Technology Conference*, pp. 773-777, Denver, CO, May 1992.

[9] *OPNET documentation V.9.0.A*, OPNET Technologies, Inc., Bethesda, MD, 2001.

[10] EDACS Trunking Information. (2004, May). [Online]. Available: http://www.radioreference.com.

[11] In Vancouver! Vancouver Travel Guide. (2004, May). [Online]. Available: http://www.vancouver-bc.com/maps.html.