



Modem pool traffic survey continued

Anni Matinlauri
amatinla@cc.hut.fi



Preface

This is the third analysis of the modem pool situation of HUT. Jani Lakkakorpi made the two previous ones in 1997 and 1998. When the first survey was made the capacity of the HUT staff modem pool was 120 simultaneous connections and the capacity of the student pool was 60 simultaneous connections. It was discovered that the staff pool had excess capacity and thus the staff capacity was reduced to 60 while the student capacity was increased to 180.

The data I analyzed is from 2001 and 2002. The staff modem pool capacity was still 60 simultaneous connections and the student pool had 180 simultaneous connections. The purpose of this survey is to monitor the situation of the modem pools, to see whether there is enough capacity or not. Also as using a modem is not so popular anymore it is interesting to see the development in the number of user. The pricing policy of Elisa (ex-Helsinki Telephone Company) can also be seen from the data. In the 1998 survey the evening and weekend rates were no longer fixed, you had to pay extra for every minute that exceeded 30 minutes. Today there is yet again a fixed rate for evenings and weekends, this rate is for Internet usage and it's called Elisa Nettitaksa. Elisa also has an Internet connection service where you are charged per minute.

There was a vast number of pictures produced from the two years of data. I have not included all of them, the general tendencies are visible from the chosen graphs. I have also used some parts of Lakkakorpi's survey here because the Matlab-functions I used are the same.



Contents

Preface	2
1. Log files	4
2. Processing data with Matlab	6
2.1. System occupancy averages	6
2.2. Holding times	19
2.3 Arrivals	29
3. Conclusions	33
4. References	34
Appendix	
I. Matlab -functions	35



1. Log files

The modem pool log files are text files (created in the HUT computing center by Kimmo Laaksonen) with the following structure:

- Each line contains the information of a single connection
- Comma is the field separator
- Information fields starting from left are:

1. Starting time of a connection YYMMDDHHMMSS
2. Weekday (0=sunday, 1=monday etc.)
3. Encrypted user ID
4. Session duration (sec.)
5. Bytes in
6. Bytes out
7. Caller phone number (3 last digits removed) or 99999 if unknown
8. Port type code:
 - Async = 0 (=modem)
 - ISDN = 2
 - ISDN-V120 = 3
 - ISDN-V110 = 4
9. Termination cause code:
 - User request = 1
 - Lost carrier = 2
 - Lost service = 3
 - Idle time-out = 4
 - Session time-out = 5
 - Admin reset = 6
 - Admin reboot = 7
 - Port error = 8
 - NAS error = 9
 - NAS request = 10
 - NAS reboot = 11
 - Port unneeded = 12
 - Port preempted = 13
 - Port suspended = 14
 - Service unavailable = 15
 - Callback = 16
 - User error = 17
 - Host request = 18
10. Terminating connection speed (if known) e.g 33600 LAPM/V42BIS
11. Initial connection speed (if different)



Usually, a log file contains the user information of one month. This seems to be an appropriate amount of information to be processed at one time. Here is a small piece from the beginning of the student modem pool log file (September 2001):

```
2001,09,09,01,48,55,0,"oaAEVJ1auKc",478,103524,617782,"016642",0,1,"33600 LAPM/V42BIS", ""
```

```
2001,09,09,01,56,29,0,"dA8gGnLAOm.",316,31809,151834,"096820",0,1,"40000 LAPM/V42BIS", "48000 LAPM/V42BIS"
```

```
2001,09,09,01,57,34,0,"VTPhtUvSUqY",3502,448459,2592980,"99999",2,1,"64000", ""
```

```
2001,09,09,02,00,43,0,"rD91BHf1MkM",3986,152149,1041171,"99999",0,1,"38666 LAPM/V42BIS", "48000 LAPM/V42BIS"
```

```
2001,09,09,02,21,09,0,"SgVVeG2/oCw",2105,255265,1563373,"098747",2,1,"64000", ""
```

```
2001,09,09,02,21,15,0,"SgVVeG2/oCw",2099,265395,1557592,"098747",2,1,"64000", ""
```

This not a good format because not all this information is needed. So I wrote a C-program to extract the essential data. I needed two kinds of data files. The first has the time in seconds from the beginning of the first day, the date and the information if it is an arrival or a departure (1/-1).

```
128284 4 -1
128297 4 -1
128299 4 1
128315 4 1
128359 4 -1
128372 4 1
128440 4 -1
```

The second type of a file is just a bit different. Instead of the event is has the holding time in seconds.

```
405249 1 745
405255 1 1694
405430 1 576
405448 1 1186
405540 1 79
405577 1 117
405681 1 3487
```



2. Processing data with Matlab

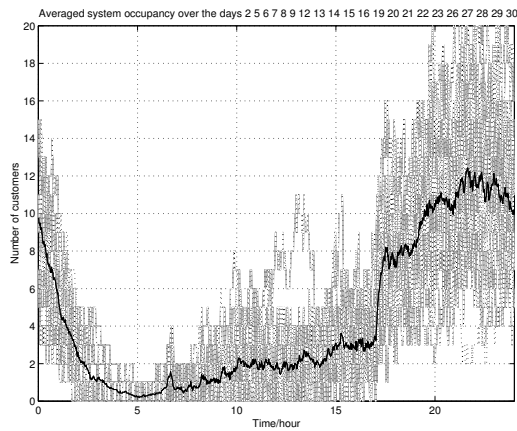
2.1 System occupancy averages*

System occupancy is the number of customers at the pool in a particular moment of time. Matlab-function *events.m* counts and prints the system occupancy during the measurement period. It takes two arguments: the moments of time and the information whether the event was an arrival or a departure. Because the behavior of the system can be considered periodical, we can compute and print out the system occupancy average over selected days. The traffic in these days should be quite similar - for example we can compute the system occupancy average for weekdays and weekends (Sat. & Sun.). We use a Matlab-function *average.m* to compute and print the system occupancy average over the selected days. Functions *events.m* and *average.m* are written by Esa Hyytiä [2].


By using the following macro *sep_stu_ave_w.m* we can easily (just by typing *sep_stu_ave_w* in Matlab) compute and print out the weekdays' system occupancy average in September in the student modem pool. To see the variation within different days, all days are first plotted with broken lines, and finally the average line is plotted with a continuous line.

```
load sepstudents.times;
[ x, y ] = events( sepstudents(:,1), sepstudents(:,2) );
for i=10:36,
    if ( i ~= [ 12 13 19 20 26 27 33 34 ] )
        [ xi, yi ] = average( x, y, [i] );
        plot( xi, yi, 'w' );
        plot( xi, yi, 'c:' );
        hold on;
    end;
end;
[ x0, y0 ] = average( x, y, [ 10 11 14 15 16 17 18 21 22 23 24 25 28
29 30 31 32 35 36 ] );
axis([0, 24, 0, 120]);
plot( x0, y0, 'r' );
hold off;
```


*This is all a direct quotation from J.Lakkakorpi's Traffic in modem pools of HUT.




Graph 1
Average number of customers in the staff modem pool on weekdays of January 1998.




Graph 2
Average number of customers in the staff modem pool on weekdays of May 1998.



Graph 3
Average number of customers in the staff modem pool on weekdays of June 1998.




Graph 4
Average number of customers in the staff modem pool on weekdays of September 1998.




Graph 5
Average number of customers in the staff modem pool on weekdays of October 1997.

Graphs 1 through 5 describe the traffic in the staff modem pool on weekdays. There are no significant changes (from October 1997) here. (Graph 5 is reprinted here for comparison. [1]) The number of pool users is slowly increasing, but the use of this pool is still quite weak.


This page is from J.Lakkakorpi's work




Graph 6
Average number of customers in the staff modem pool on weekdays in January 2001.




Graph 7
Average number of customers in the staff modem pool on weekdays in March 2001.



Graph 8
Average number of customers in the staff modem pool on weekdays in June 2001.




Graph 9
Average number of customers in the staff modem pool on weekdays in September 2001.




Graph 10
Average number of customers in the staff modem pool on weekdays in October 2001.

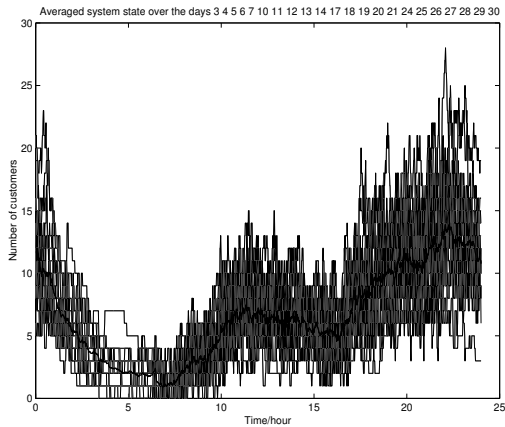
Graphs 6 through 10 describe the traffic in the staff modem pool on weekday in 2001. The number of users has increased from 1998 and 1997. The busiest time is still evenings after 5 pm, though in the student pool the sharp increase in the number of connections after 5 pm is more obvious.



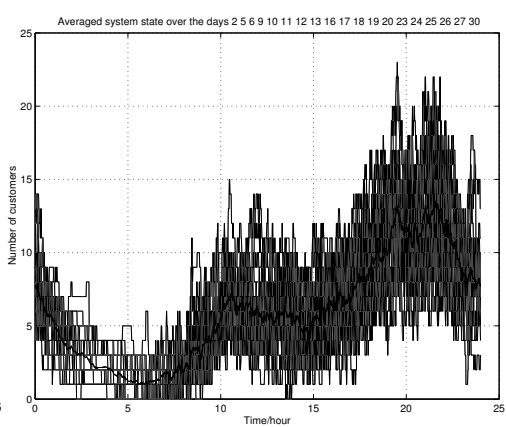
Graph 11
Average number of customers in the staff modem pool on weekdays in January 2002.




Graph 12
Average number of customers in the staff modem pool on weekdays in March 2002.



Graph 13
Average number of customers in the staff modem pool on weekdays in June 2002.




Graph 14
Average number of customers in the staff modem pool on weekdays in September 2002.




Graph 15
Average number of customers in the staff modem pool on weekdays in October 2002.

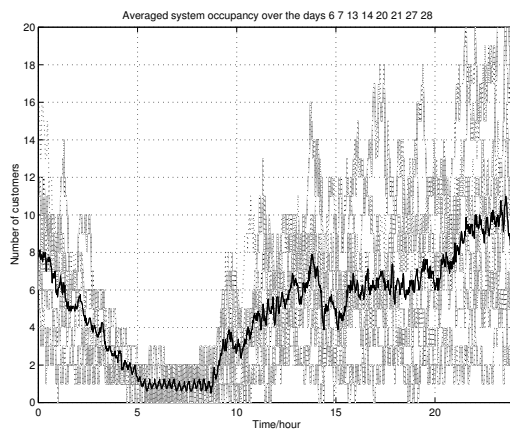
The graphs 11 through 15 describe the weekdays in the staff modem pool in 2002. The number of users has decreased from the previous year and the sharp increase after 5 pm is no longer there, the number of connections rise steadily trough the day.




Graph 16
Average number of customers in the staff modem pool on weekends of January 1998.



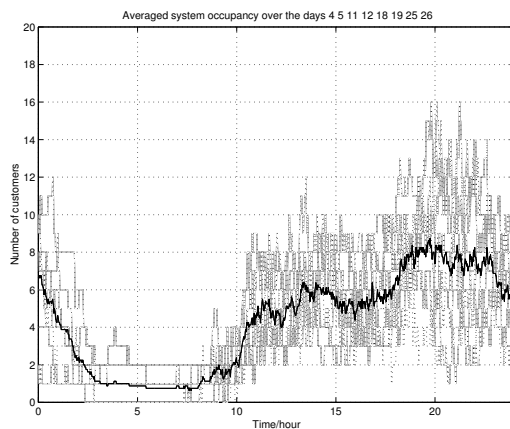
Graph 17
Average number of customers in the staff modem pool on weekends of May 1998.



Graph 18
Average number of customers in the staff modem pool on weekends of June 1998.




Graph 19
Average number of customers in the staff modem pool on weekends of September 1998.



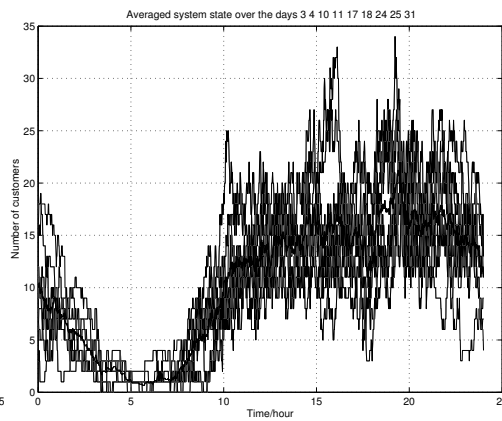
Graph 20
Average number of customers in the staff modem pool on weekends of October 1997.

Graphs 16 through 20 describe the traffic in the staff modem pool on weekends. (Graph 15 is reprinted here for comparison. [1]) The daytime use of this pool seems to be a bit more intense during weekends (which is quite natural, since people are at home).


This page is from J.Lakkakorpi's work



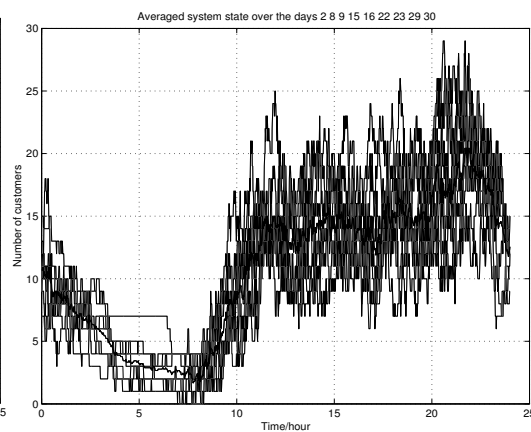
Graph 21
Average number of customers in the staff modem pool on weekends in January 2001.




Graph 22
Average number of customers in the staff modem pool on weekends in March 2001.



Graph 23
Average number of customers in the staff modem pool on weekends in June 2001.




Graph 24
Average number of customers in the staff modem pool on weekends in September 2001.




Graph 25
Average number of customers in the staff modem pool on weekends in October 2001.


The graphs 21 through 25 describe the staff modem pool on weekends in 2001. The number of users has increased a bit since 1998 and the number of connections/hour increases slightly all day.




Graph 26
Average number of customers in the staff modem pool on weekends in January 2002.




Graph 27
Average number of customers in the staff modem pool on weekends in March 2002.



Graph 28
Average number of customers in the staff modem pool on weekends in June 2002.

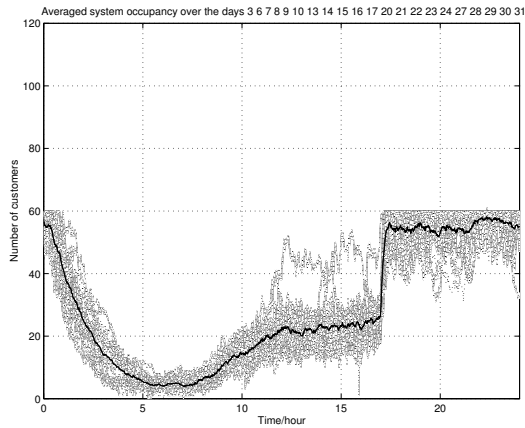


Graph 29
Average number of customers in the staff modem pool on weekends in September 2002.

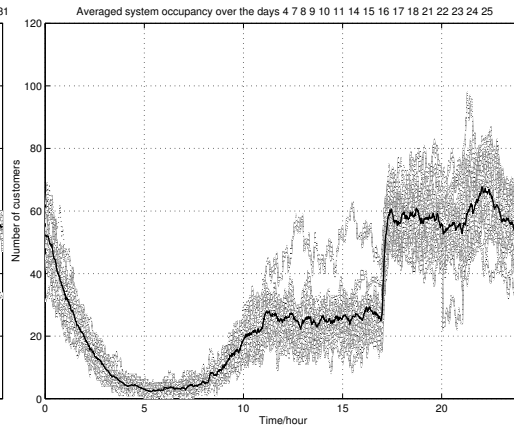


Graph 30
Average number of customers in the staff modem pool on weekends in October 2002.


The graphs 26 through 30 describe the staff modem pool on weekends in 2002. The number of connections has decreased a bit from the previous year.




Graph 31
Average number of customers in the student modem pool on weekdays of January 1998.




Graph 32
Average number of customers in the student modem pool on weekdays of May 1998.



Graph 33
Average number of customers in the student modem pool on weekdays of June 1998.




Graph 34
Average number of customers in the student modem pool on weekdays of September 1998.




Graph 35
Average number of customers in the student modem pool on weekdays of October 1997.

Graphs 31 through 35 describe the traffic in the student modem pool on weekdays. Traffic in January is still very similar to traffic in October 1997 [1] (graph 30 is reprinted here for comparison), but the next three graphs differ from these. Now there seems to be enough capacity in this modem pool (for 180 users), too. The peak at 17.00 (cheaper rates until 07.00 next morning) is very clear in every graph.


This page is from J.Lakkakorpi's work




Graph 36
Average number of customers in the student modem pool on weekdays in January 2001.




Graph 37
Average number of customers in the student modem pool on weekdays in March 2001.



Graph 38
Average number of customers in the student modem pool on weekdays in June 2001.




Graph 39
Average number of customers in the student modem pool on weekdays in September 2001.




Graph 40
Average number of customers in the student modem pool on weekdays in October 2001.


The graphs 36 through 40 describe the student modem pool on weekdays in 2001. The number of connections has increased from 1997 and there still is an obvious sharp increase in the number of connections after 5 pm.




Graph 41
Average number of customers in the student modem pool on weekdays in January 2002.




Graph 42
Average number of customers in the student modem pool on weekdays in March 2002.



Graph 43
Average number of customers in the student modem pool on weekdays in June 2002.




Graph 44
Average number of customers in the student modem pool on weekdays in September 2002.




Graph 45
Average number of customers in the student modem pool on weekdays in October 2002.


The graphs 41 through 45 describe the student modem pool on weekdays in 2002. The number of connections has noticeably decreased from the previous year and the 5 pm increase lessens a bit in the course of the year.




Graph 46
Average number of customers in the student modem pool on weekends of January 1998.




Graph 47
Average number of customers in the student modem pool on weekends of May 1998.



Graph 48
Average number of customers in the student modem pool on weekends of June 1998.




Graph 49
Average number of customers in the student modem pool on weekends of September 1998.




Graph 50
Average number of customers in the student modem pool on weekends of October 1997.

Graphs 46 through 50 describe the traffic in the student modem pool on weekends. Traffic in January is still very similar to traffic in October 1997 [1] (graph 45 is reprinted here for comparison), but the next three graphs differ from these. Now there seems to be enough capacity in this modem pool (for 180 users), too.


This page is from J.Lakkakorpi's work




Graph 51
Average number of customers in the student modem pool on weekends in January 2001.




Graph 52
Average number of customers in the student modem pool on weekends in March 2001.



Graph 53
Average number of customers in the student modem pool on weekends in June 2001.




Graph 54
Average number of customers in the student modem pool on weekends in September 2001.




Graph 55
Average number of customers in the student modem pool on weekends in October 2001.


The graphs 51 through 55 describe the student modem pool on weekends in 2001. The number of connections first increases and then decreases during the year. The number of connections/hour increases steadily all day.




Graph 56
Average number of customers in the student modem pool on weekends in January 2002.




Graph 57
Average number of customers in the student modem pool on weekends in March 2002.



Graph 58
Average number of customers in the student modem pool on weekends in June 2002.



Graph 59
Average number of customers in the student modem pool on weekends in September 2002.



Graph 60
Average number of customers in the student modem pool on weekends in October 2002.

The graphs 56 through 60 describe the student modem pool on weekends in 2002. The number of connections keeps decreasing all through the year. The number of connections/hour slightly increases all day.



2.2 Holding times*

Holding time is the duration of a single connection. Matlab-function *tail_holding_times.m* computes and prints out the tail distribution of the holding times of a particular measurement period. The step for these holding times is set to 100 seconds. The following macro *hold_times_sep_stu_w_d.m* computes the tail distribution for the weekday (07.00 - 17.00) holding times of the student modem pool. It uses the function mentioned earlier.

```
load sepstudents.data;

len = length(sepstudents(:,3));
newlen = 0;

for i=1:len,
    if ( ( sepstudents(i,3) ~= [6 0] ) &
        ( ( mod(sepstudents(i,1),86400) < 61200 ) &
          ( mod(sepstudents(i,1),86400) > 25200 ) ) ),
        newlen = (newlen + 1);
    end;
end;


x = zeros(newlen,1);
y = zeros(newlen,1);

k=1;

for j=1:len,
    if ( ( sepstudents(j,3) ~= [6 0] ) &
        ( ( mod(sepstudents(j,1),86400) < 61200 ) &
          ( mod(sepstudents(j,1),86400) > 25200 ) ) ),
        x(k) = sepstudents(j,3);
        y(k) = sepstudents(j,4);
        k = ( k + 1 );
    end;
end;


tail_holding_times( x, y );
axis([0,120,0,15500]);
text(40, 7750, sprintf('The total number of daytime
(0700 - 1700)\nstudent calls on workdays: %d', newlen));
```

*Also quoted from J.Lakkakorpi's work.




Graph 61

The tail distribution of the holding times in the staff modem pool on weekdays (07.00-17.00) of January 1998.




Graph 62

The tail distribution of the holding times in the staff modem pool on weekdays (07.00-17.00) of May 1998.




Graph 63

The tail distribution of the holding times in the staff modem pool on weekdays (07.00-17.00) of June 1998.



Graph 64


The tail distribution of the holding times in the staff modem pool on weekdays (07.00-17.00) of September 1998.



Graph 65


The tail distribution of the holding times in the staff modem pool on weekdays (07.00-17.00) of October 1997.

Graphs 61 through 65 describe the holding times in the staff modem pool on weekdays (07.00 - 17.00). No major changes here. Graph 60 is reprinted here for comparison [1]. Note the different scales in these graphs.




Graph 66

The distribution of the holding times in the staff modem pool on weekdays (07.00-17.00) in September 2001.



Graph 67

The tail distribution of the holding times in the staff modem pool on weekdays (07.00-17.00) in September 2001.




Graph 68

The logarithmic tail distribution of the holding times in the staff modem pool on weekdays (07.00-17.00) in September 2001.


On the previous page there is an example of J.Lakkakorpi's work. I have not included any more in this section because the amount of holding times graphs there are.

The graphs 66 through 68 describe the distributions of holding times in September 2001. On the next page graphs 69 through 74 also describe the holding times in 2001.


The holding times seem to be quite long and there's a bit more calls than 1998. The most calls come in the evenings. The weekends and the daytimes are quite even but which is more popular depends on the season.




Graph 69
The distribution of the holding times in the staff modem pool on weekday evenings (17.00-07.00) in September 2001.




Graph 72
The distribution of the holding times in the staff modem pool on weekends in September 2001.




Graph 70
The tail distribution of the holding times in the staff modem pool on weekday evenings (17.00-07.00) in September 2001.




Graph 73
The tail distribution of the holding times in the staff modem pool on weekends in September 2001.



Graph 71
The logarithmic tail distribution of the holding times in the staff modem pool on weekday evenings (17.00-07.00) in September 2001.




Graph 74
The logarithmic tail distribution of the holding times in the staff modem pool on weekends in September 2001.




Graph 75

The distribution of the holding times in the staff modem pool on weekdays (07.00-17.00) in September 2002.



Graph 77


The distribution of the logarithmic tail holding times in the staff modem pool on weekdays (07.00-17.00) in September 2002.




Graph 76

The distribution of the tail holding times in the staff modem pool on weekdays (07.00-17.00) in September 2002.


The graphs 75 through 83 describe the staff modem pool in 2002. There's no major change from 2001. There is only a slight decrease in the number of users in the evenings and weekends. The length of the average holding time doesn't really change.




Graph 78
The distribution of the holding times in the staff modem pool on weekday evenings (17.00-07.00) in September 2002.



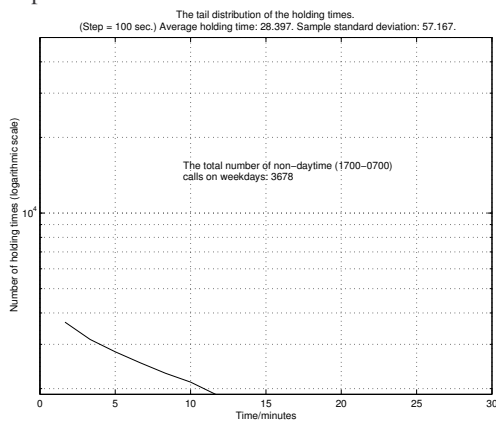
Graph 81
The distribution of the holding times in the staff modem pool on weekends in September 2002.




Graph 79
The distribution of the tail holding times in the staff modem pool on weekday evenings (17.00-07.00) in September 2002.




Graph 82
The tail distribution of the holding times in the staff modem pool on weekends in September 2002.




Graph 80
The distribution of the holding times in the staff modem pool on weekday evenings (17.00-07.00) in September 2002.




Graph 83
The tail distribution of the holding times in the staff modem pool on weekends in September 2002.



Graph 84
The distribution of the holding times in the student modem pool on weekdays (07.00-17.00) in September 2001.

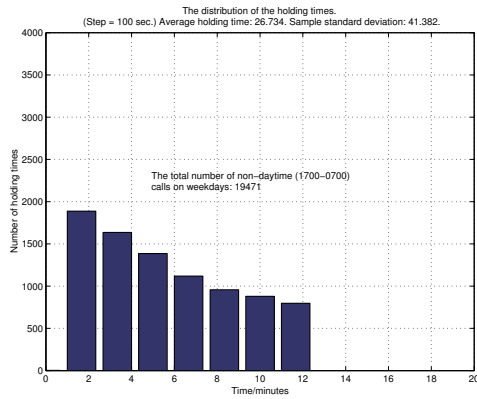


Graph 85
The tail distribution of the holding times in the student modem pool on weekdays (07.00-17.00) in September 2001.




Graph 86
The distribution of the logarithmic tail holding times in the student modem pool on weekdays (07.00-17.00) in September 2001.


The graphs 84 through 92 describe the student modem pool in 2001. The number of users decreases slightly in the course of the year and the evenings are the most popular time. The average holding time stays pretty much the same with evenings and weekends having longer holding times than the days.




Graph 87
The distribution of the holding times in the student modem pool on weekday evenings (17.00-07.00) in September 2001.




Graph 90
The distribution of the holding times in the student modem pool on weekends in September 2001.




Graph 88
The distribution of the tail holding times in the student modem pool on weekday evenings (17.00-07.00) in September 2001.




Graph 91
The tail distribution of the holding times in the student modem pool on weekends in September 2001.




Graph 89
The distribution of the holding times in the student modem pool on weekday evenings (17.00-07.00) in September 2001.




Graph 92
The tail distribution of the holding times in the student modem pool on weekends in September 2001.



Graph 93
The distribution of the holding times in the student modem pool on weekdays (07.00-17.00) in September 2002.




Graph 94
The tail distribution of the holding times in the student modem pool on weekdays (07.00-17.00) in September 2002.




Graph 95
The distribution of the logarithmic tail holding times in the student modem pool on weekdays (07.00-17.00) in September 2002.


The graphs 93 through 101 describe the student modem pool in 2002. The number of calls keeps on decreasing but the average holding time remains about the same than previous year.




Graph 96
The distribution of the holding times in the student modem pool on weekday evenings (17.00-07.00) in September 2002.




Graph 99
The distribution of the holding times in the student modem pool on weekends in September 2002.




Graph 97
The distribution of the tail holding times in the student modem pool on weekday evenings (17.00-07.00) in September 2002.



Graph 100
The tail distribution of the holding times in the student modem pool on weekends in September 2002.



Graph 98
The distribution of the holding times in the student modem pool on weekday evenings (17.00-07.00) in September 2002.




Graph 101
The tail distribution of the holding times in the student modem pool on weekends in September 2002.



2.3 Arrivals


An arrival happens when someone starts a modem session. A Matlab function `arrivals.m` calculates the arrivals that happen within a half an hour. A function `arr_average` is used to calculate the average arrivals over chosen days. The macro `arr_hlo_vko.m` prints out a graph of the average arrivals in the staff pool on weekdays.

```
load hlo9710.ev;
[ x, y ] = arrivals( hlo9710(:,1), hlo9710(:,3) );
for i=2:31,
    if ( i~= [ 4 5 11 12 18 19 25 26 ] )
        [ xi, yi ] = arr_average( x, y, [i] );
        plot( xi, yi, 'w' );
        plot( xi, yi, 'c:' );
        hold on;
    end;
end;
[ x0, y0 ] = arr_average( x, y, [ 2 3 6 7 8 9 10 13 14 15 16 17 20 21
    22 23 24 27 28 29 30 31 ] );
plot( x0, y0, 'r' );
hold off;
```



Graph 102


The arrivals in the student modem pool on weekdays in January 2001.




Graph 103

The arrivals in the student modem pool on weekends in January 2001.


The graphs 102 through 105 describe the arrivals in the student pool in 2001. During that year there seems to be a slight decrease in the number of arrivals. Weekday evening after 5 pm is the most popular time to arrive. On weekends students seem to be waking up after 10 am, or at least that's when they connect to the pool. The pool also has enough capacity unlike in 1997.




Graph 104
The arrivals in the student modem pool on weekdays in September 2001.




Graph 105
The arrivals in the student modem pool on weekends in September 2001.




Graph 106
The arrivals in the student modem pool on weekdays in January 2002.



Graph 107
The arrivals in the student modem pool on weekends in January 2002.




Graph 108
The arrivals in the student modem pool on weekdays in October 2002.




Graph 109
The arrivals in the student modem pool on weekends in October 2002.


The graphs 106 through 109 describe the student modem pool in 2002. No drastic changes here compared to 2001 only a slight decrease in the number of users.




Graph 110
The arrivals in the staff modem pool on weekdays in January 2001.



Graph 111
The arrivals in the staff modem pool on weekends in January 2001.




Graph 112
The arrivals in the staff modem pool on weekdays in September 2001.




Graph 113
The arrivals in the staff modem pool on weekends in September 2001.


The graphs 110 through 113 describe the staff modem pool in 2001. The number of arrivals stays roughly the same through the year. The pool has enough capacity for every arrival.




Graph 114
The arrivals in the staff modem pool on weekdays in January 2002.



Graph 115
The arrivals in the staff modem pool on weekends in January 2002.



Graph 116
The arrivals in the staff modem pool on weekdays in October 2002.



Graph 117
The arrivals in the staff modem pool on weekends in October 2002.

The graphs 114 through 117 describe the staff modem pool in 2002. There's really no change to 2001, unlike the student pool the number of users is not declining.



3. Conclusions

There was a vast amount of data for this survey and all kinds of graphs could have been produced from it. However for comparability's sake the same Matlab-functions were used as in Lakkakorpi's work. Quite a lot of graphs were produced to see the development in the modem pool situation clearly.

The first survey by J.Lakkakorpi found that the staff pool had too much capacity and the student pool had too little. This had been fixed by the second survey in 1998. In 2001 and 2002 there definitely was no congestion at all, in fact the user rates started to drop. From 2001 to the end of 2002 there was a significant decrease in the user rates. It seems that the peak of modem usage was in the end of the 90's. Nowadays better and faster ways to connect such as ADSL are fairly cheap and very popular.

Another thing that was looked into in the previous surveys was the effect of the pricing policy. It seems that in the student pool the busiest time is the weekday evenings after 5 pm. Weekends are also fairly popular but on those days the user rates simply climb steadily all day. With the staff pool the weekday 5 pm peak is a lot smaller. The length of the holding times has increased a bit from 1998 but during 2001 and 2002 they don't undergo any drastic change. Naturally on weekday evenings and weekends the holding times tend to be longer. This is especially because of the pricing policy of Elisa. They offer a fixed rate for evenings after 5 pm and weekends. It's obvious from the graphs that those are the most popular times to use the student pool. With staff the money doesn't seem to matter that much since there is a more steady increase in the user rates all through the day.



4. References

- [1] Jani Lakkakorpi: TKK:n modeemipoolien lokitiedostojen analysointi (erikoistyö)
- [2] Esa Hyytiä: Modem pools of Helsinki University of Technology (memo)
- [3] Jani Lakkakorpi: Traffic in modem pools of Helsinki University of Technology



Appendix I. Matlab-functions

events.m:

```
function [x,y] = events(xe,ye);
%
% This function plots events. It also returns
% data plotted. X-axis is scaled to days.
%
len = length(xe);
yy = cumsum(ye);
y = zeros(2*len-1,1);
x = zeros(2*len-1,1);
for i=1:(len-1)
    y(2*i-1) = yy(i);
    y(2*i) = yy(i);
    x(2*i-1) = xe(i);
    x(2*i) = xe(i+1);
end;
x(2*len-1) = xe(len);
y(2*len-1) = ye(len);
x = x / 86400;           % Scaled to days
plot(x,y);
grid;
xlabel('Time/days');
ylabel('Number of customers in the system');
```

average.m:

```
function [ax,ay] = average(x, y, days )
%
% function [ax,ay] = average(x, y, days )
%
% This function calculates average number of customers
% in system on given days. Time ax is in hours.
%
% First we make a vector of system occupancy with discrete
% one minute steps

const = 24*60;
x2    = x * const;
xmax  = floor(max( x2 )) + 1;
y2    = zeros( xmax, 1 );

t1 = floor( x(1)*const );
for i=2:length(x)           % first event of that day
    t2 = floor( x(i)*const );
    curr_y = y(i-1);
    while( t1 < t2 )
        y2(t1) = curr_y;
        t1 = t1 + 1;
    end;
end;
y2(t1) = curr_y;

%
% Ok, now collect given days from this "minute" vector
%
ay = zeros( const, 1 );
ax = (1:const)/60;
```



```
for d=1:length( days )
    t1 = (days(d)-1) * const;
    t2 = t1 + const;
    if ( t1 < 1 )
        t1 = 1;
    end;
    if ( t2 > xmax )
        t2 = xmax;
    end;
    while( t1 < t2 )
        tmod = mod(t1,const) + 1;
        ay(tmod) = ay(tmod) + y2(t1);
        t1 = t1 + 1;
    end;
end;
ay = ay / length( days );

plot( ax, ay );
grid;
daystr = sprintf( ' %d', days );
title( sprintf( 'Averaged system occupancy over the days%s', daystr
) );
xlabel( 'Time/hour' );
ylabel( 'Number of customers' );
```

tail_holding_times.m:

```
function [ x, y ] = holding_times( xe, ye );
%
% This function plots the tail distribution of
% the holding times of the connections.
% Vector xe contains the days of the week. Vector ye contains
% the corresponding holding times.
%
ys = sort(ye); % First we sort the holding times
len = length(ye); % The number of all
connections
points = floor( max(ye) / 100 ); % The number of
observation points
step = 100; % The step of holding
times (100sec)

y = zeros(points+2, 1);
x = zeros(points+2, 1);

x(1) = 0; % Unsuccessful logins?
i=1;

while ( ys(i) == 0 )
    y(1) = ( y(1) + 1 );
    i = ( i + 1 );
end;

start = i; % First non-zero holding
time

left = (len - y(1));

for j=2:(points+2),
    x(j) = ( (j-1)*step );
    y(j) = left;

    while ( (ys(i) > (j-2)*step) & (ys(i) <= ( (j-1)*step )) & (i <
len) ),
        i = ( i + 1 );
        left = (left - 1);
    end;
```



```
end;

dev2 = 0;
ave = sum(ys)/(len-y(1));      % Average does not contain zero
holding times!
for k=start:len,              % Neither does sample standard
deviation...
    dev2 = ( dev2 + 1 / ( (len-y(1)) - 1 ) * ( ys(k) - ave )^2 );
end;
dev = sqrt(dev2);
x = x/60;                      % Scaled to minutes

bar(x, y);
grid;
results = sprintf('Average holding time: %0.5g. Sample standard
deviation: %0.5g.', ave/60,dev/60 );
title( sprintf( 'The tail distribution of the holding times. \n(Step
= 100 sec.) %s', results) );
xlabel('Time/minutes');
ylabel('Number of holding times');
```