# An Approach to Capacity Planning of a Local Service Provider as an Element of Internet Infrastructure

Prof. Victor N. Vasiliev[†], Prof. Natalia S. Ruzanova[‡],
Dr. Timo Alanko[§], Dr. Iouri A. Bogoiavlenski[¶]

[†,‡]Administration of University of Petrozavodsk
[§]Department of Computer Science, University of Helsinki
[¶]Department of Computer Science, University of Petrozavodsk

[†,‡,¶]Lenin St., 33, Petrozavodsk, Republic of Karelia, 185640, Russia
[§]P.O.Box 26 (Teollisuuskatu 23) FIN-00014 University of Helsinki, Finland

E-mail: [†,‡,¶]{rector, ruzanova, ybgv}@mainpgu.karelia.ru,
[§]alanko@cs.Helsinki.FI

**Abstract**

Growth of the Internet infrastructure in the Republic of Karelia is presented and discussed. The importance of problems related with the growth planning is pointed out. The focus of the paper is the capacity planning problem of an Internet Local Service Provider system. A decomposition of the system into four main subsystems is proposed: the customer–access servers, the internal Internet servers, the internal LAN, and the external channels to higher level service providers. The decomposition is based on an analysis of the evolution of the Federal Petrozavodsk RUNNet Node. The Local Service Provider system capacity planning problem is treated in some detail. A customer–access server model and Ethernet segment modeling considerations get some attention as components of an LSP composite model.

# 1 Introduction

Despite of serious crisis phenomena in Russian economics, some of its branches develop with fast rate. During the last years an intensive growth of the Internet infrastructure can be observed, both with regard to the number of customers and to the services provided. Indeed, the information–technology industry providing Internet–related services can be included into the major growing branches in the Russian economy.

As it is typical worldwide, Russian Universities play a role of conductors and catalysts of the Internet technologies. Since 1994, the University of Petrozavodsk provides Internet services in the Republic of Karelia; the main users of these services are pupils of Karelian schools, and students and staff of educational, scientific and cultural organizations.

In this paper we shortly describe the main characteristics of the growing Internet infrastructure in the Republic of Karelia. We consider the planning problem due to the fast growth of the infrastructure and proclaim that, in this respect, the most important part of the infrastructure are the Local Service Providers (LSP), i.e. the Internet service providers closest to the customers. Then we give an overview on the evolution of the Federal Petrozavodsk RUNNet Node (FPRN), which is the oldest and most powerful LSP in the Republic. We conclude that an LSP has a stable functional structure, and describe its main subsystems. This makes it possible to decompose the LSP capacity planning problem into more tractable parts. Basing on LSP modeling through a hierarchical composition of subsystem models we formulate an approach to efficient solving methods of the capacity planning problem. Results of some first steps in this direction, presented among the papers of these Proceedings, are shortly discussed.

This paper presents continuation of research, started in [1, 2].

# 2 Growth of Internet Infrastructure in Republic of Karelia

Since 1995 the growth of the Russian Internet segment and its Karelian part has had an explosive character, quite in the same lines that have been typical for the global Internet. The growth has exceeded all expectations.

There were in 1997 twelve national Internet providers in Russia: Demos/Internet, EUNet/Relcom, FREENet, GlasNet, Global One Russia, MSUNet RELARN–IP, RSSI, RUHEP/Radio–MSU, RUNNet, RosNet, and Sovam Teleport; in addition to these, there were approximately two hundred relatively small local service providers.

The National Russian Backbone is under construction (RBNet project) [13], and there were in 1997 year two Internet Exchange points to transmit traffic between different provider networks (in Moscow and in Saint–Petersburg).

In January 1997 there were more than 50.000 hosts and near 3.000 second–level domains within the domain "ru", and there were close to 19.000 hosts and 600 second–level domains within the domain "su" [14].

The process of the Internet infrastructure growth has been—and continues to be—very fast: the development is still at a level which is far from any kind of saturation.

The Internet infrastructure of the Republic of Karelia, shown in Figure 1, reflects the state of the network at Spring 1998. There are two main commercial operators in the Republic. The "Rostelecom", a branch of the federal level enterprise with the same name, provides various kinds of communication services in the Northern Region. The main area of activity of "Electrosviaz", the other operator, is providing of phone services for Petrozavodsk population. Both operators are now developing modern Internet services and compete for customers.

Such parts of the infrastructure as communications channels exist, and have been developed, in the Republic for a long time. But Internet access points, shown in Figure 1, started to emerge not earlier than three years ago, after the creation of the Federal Petrozavodsk RUNNet Node (FPRN) of the University of Petrozavodsk. Its first hardware configuration was rather simple, and the range of offered services was small. However, the operational FPRN has demonstrated the advantages, profits and pleasures of Internet to the people of the Republic. More important is that FPRN has been—and still is—a kind of incubator of staff and users for the Internet access points.

Today, even if the Internet infrastructure already covers an essential part of the Republic, it still has big growth potentials both in territorial expansion and in improving the range and quality of services, provided by the access points.

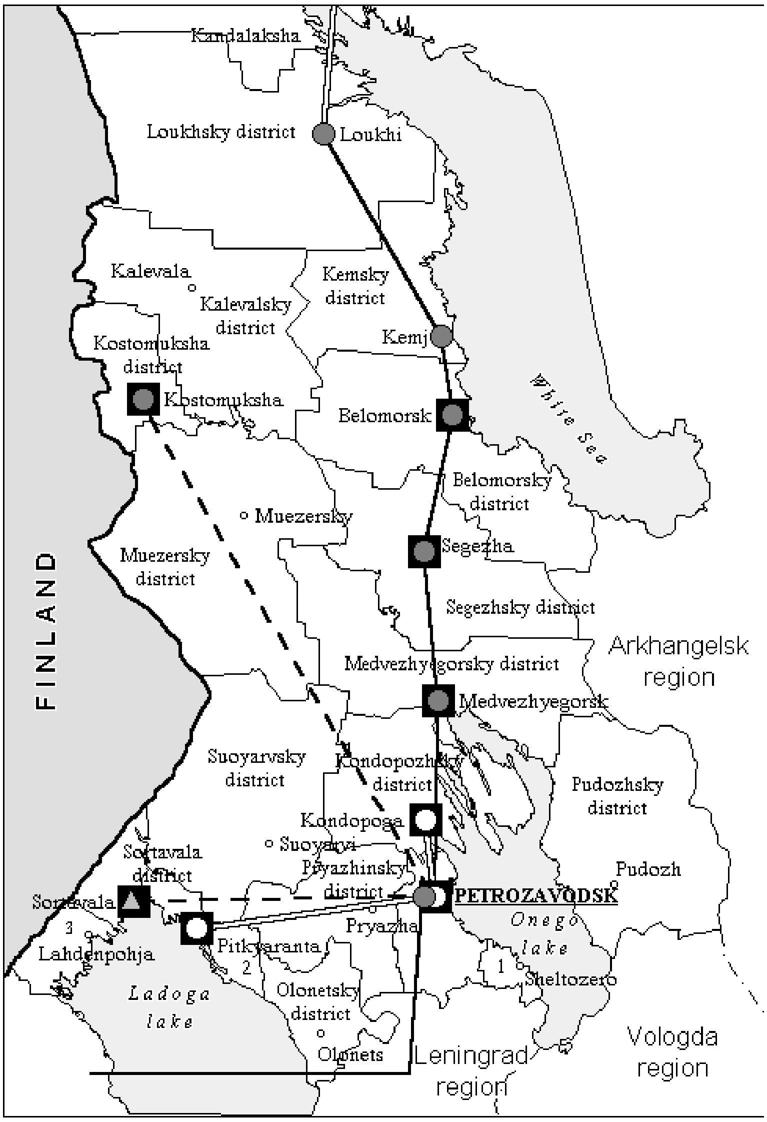


Figure 1. The Internet infrastructure of the Republic of Karelia, Spring 1998.

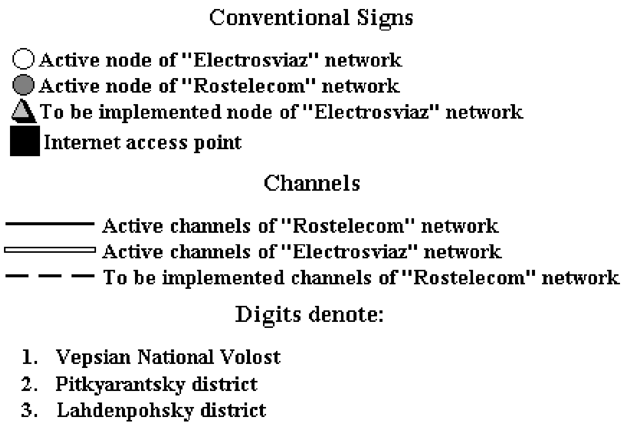


Figure 2. Conventional signs for Figure 1

## 3 Internet Infrastructure Planning and Local Service Providers

The expansion of Internet infrastructure, described in the previous section, can be carried out in many different ways, but in any case it involves many organizations and enterprises, capital investments, customers and other related interests. What is needed are approaches to the infrastructure planning problem, with a goal to keep the process as economic and as manageable as possible.

Nowadays an ever increasing number of research efforts are concerned with Internet behavior, treating the Internet as a more or less individual object with its own characteristics. A good review with further references of the efforts in this area can be found in [3]. The report points out, that "As the nation and world become increasingly dependent on the Internet, it is critical that we develop mechanisms to enable infrastructure–wide planning and analysis and to promote continued efficient scaling of the Internet". These efforts are mainly concerned with the first stages of performance analysis: definitions of metrics [4], [5] and monitoring [6]. Creation of the new IETF working groups IPPM (IP Performance Metrics) [7] and RTFM (Real–time Traffic Flow Measurement) [8] reflects the importance of this research direction.

It is evident, that the problem of Internet infrastructure planning as a whole is huge, and, hence, if we want to solve the problem, we have to decompose it into more manageable subproblems. Internet can be considered as a complex system with a virtual hierarchical structure, as shown in Figure 3. For us the most interesting components of the structure are the Local Service Providers (LSP)—especially as the FPRN, which is supported by the University of Petrozavodsk, can be considered as a typical LSP. The notion of LSP covers both enterprises selling Internet services, and organizations providing services for their own personnel. The LSP's are the service providers which are closest to the customers.

Our research concentrates on problems connected with the capacity planning of LSP's, and we use the FPRN as a modeling and experimental environment. The problem area is of uttermost importance as the LSP's carry most of the burden due to the growth of Internet and their capacity and functionality are crucial with respect to the performance of the whole infrastructure.

In Summer 1997 there were 31 national–level providers and 4009 LSP's in the United States [10]. The relationship between the number of the national–level providers and of the number of LSP's is close to that observed in Russia (12 national providers and about 200 LSP's in spring 1997 [12]); the same holds for other countries of the Internet "second wave" (SW), which follow the worldwide Internet tendencies with a certain delay.

The importance of the problem is strengthened by the fact, that LSP's have a pronounced influence on several interesting Internet service characteristics, such as the end-to-end performance of a transmission flow. Notice, as well, that in each Internet access point in Figure 3 there exists at least one LSP.

The term "infrastructure–wide planning" [3] can be understood in a very wide sense, including even psychological, sociological, and ecological aspects. We will consider only aspects related with the data communications technology remaining inside well–known and well–developed performance analysis and capacity planning methodology (see, for example [11]).

The capacity planning problem is rather difficult due to the complexity

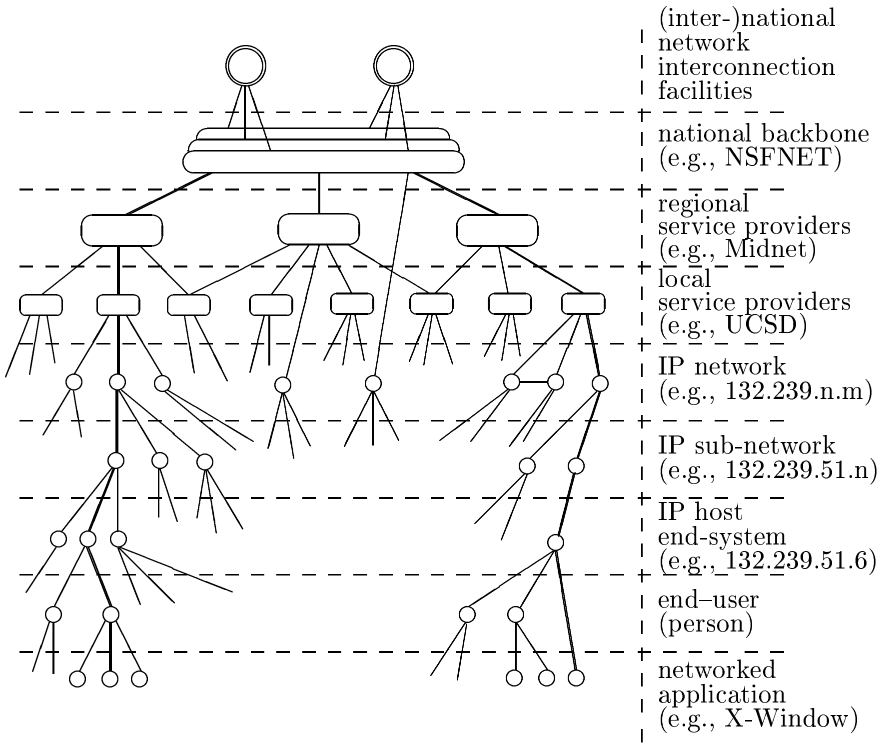


Figure 3. Virtual hierarchical structure of Internet

of an LSP itself and due to the rapid rate its features change, including the technologies used. As we have already pointed out, practical solution techniques should be based on a decomposition of the problem area. As we want to develop planning methods which remain practical for a longer period of time we have to produce the decomposition carefully. Hence, the first step should be an analysis of the LSP growth process with the aim to distinguish essential, time–invariant features of the system and, first of all, of its structure. For this purpose, let us review the FPRN growth process itself.

## 4 Growth of Federal Petrozavodsk RUNNet Node

The Federal Petrozavodsk RUNNet Node was started at the University of Petrozavodsk within the frame of Branch V of the Scientific–Technical Program "Universities of Russia". The FPRN acts since 1994 as a Local Service Provider for staff and students of the University, and for customers from different governmental, educational, scientific, and cultural organizations of the city of Petrozavodsk and of the Republic of Karelia.

Since the starting a fast growth has been characteristic of FPRN, and this with respect to all the main features, such as range of Internet services, hardware capacity, internal structure complexity and so on. FPRN is a typical Russian LSP. This growth can be characterized by three main factors.

First, hardware, and more slowly, software of FPRN permanently change to provide modern services at a sufficient level. As an example we give some growth characteristics of FPRN for a one–year period (Table 1).

Second, requests for all kinds of telecommunication services in Russia exceed by far the possibilities of the existing technology. This means that there exists potential for extensive growth.

Take the University users as an example. In 1997, about 100 students and 550 members of staff had access to the e–mail service. On the other hand, there were 5500 students and 1500 members of staff still needing it. Service requirements are especially large at the big faculties of Medicine, Economics and Law. These faculties are very popular among the Uni-

Table 1.

| **Characteristic** | **May 1996** | **May 1997** |
|---|---|---|
| Mailexchanger hardware | P75/32M/1G | P200/32M/2G |
| FUNet router | i486–33/2M/20M | i486–75/16M/170M |
| New link to RUNNet/SPb | not exists | 256 kbps |
| RUNNET router | i486–66/16M/540M | CISCO 2509 |
| External Nets | not exists | Karelian Research Center of Russian Academy of Science, National Library |
| Modems in Pool | 6 | 10 |
| Modems connection type | IRQ, DigiBoard | CISCO, Ethernet |
| Mail external users | 120 | 250 |
| On–line users | 53 | 112 |
| Segments in LAN | 1 | 2 |
| Workstations in LAN | 115 | 190 |
| Mail internal users | 315 | 655 |

versity entrants and among adults needing retraining, and for them the international contacts supported by Internet are of vital interest.

In 1997, the number of regional customers were about 250, but a cautious estimate of the near–future potential is close to 2500. Currently, growth of the customer base is restrained by weaknesses of the communication infrastructure, but its modernization, based on digital technologies, is in progress.

Third, there are efforts of international organizations to enlarge the coverage of Internet within Russia. As universities are considered to be conductors and catalysts of the adoption process into information technologies, remarkable investments are made in their information infrastructures. An example: funded by the Soros Foundation, special Internet computer classes have been established in 32 Universities of Russia. The equipment in Petrozavodsk consists of some 50 computers for the end users, new LSP servers, and hardware for a WWW design laboratory. Different kinds of application development, based on the equipment, are coordinated by the branch of the “Open Society Institute”, created at the University at the Soros foundation’s expense.

Today FPRN has a rather complex structure; it is a multi–server multilink system providing a full range of Internet services. Its structural development is roughly presented in Figures 4, 5 and 6.

Figure 4 shows the general scheme of FPRN before March 1997. The scheme is rather simple, both workstations and servers are connected to one Ethernet segment, and the capacity of the external channels is low—9.6 kbps to FUNet via Joensuu and 64 kbps to RUNNet via Saint–Petersburg (SPb).

Figure 5 shows FPRN scheme after the 1st of March, 1997, when the "University LAN segment" was divided into two segments, connected by bridges. The administrative segment "CoreLAN" includes the main FPRN servers, and the users segment "University LAN" includes two NetWare servers and close to 190 workstations. The capacity of the external channel to SPb is now already 256 kbps.

Figure 6 presents FPRN scheme at the 1st of November 1998. FPRN LAN already consists of several segments, connected by switches inside the main building. The LAN segments of remote buildings are connected to the main building via routers. The FPRN servers, the external channels routers, and the access server are interconnected via fast routers and switches. The capacity of a new external channel to Moscow "Rostelecom" is 2 Mbps. The development became possible due to funds of Soros Foundation and cooperation with "Rostelecom".

As a last notice we want to point out that for historical reasons network platforms and protocols at the University have developed to form a rich variety of different species. For example, currently there are in use at FPRN: NetWare and Unix servers, TCP/IP protocol suite and an applied subset of it, IPX/SPX/NCP, NetBios, AppleTalk, UUCP, UUCP over TCP/IP, and some other protocols of a more local interest. This state of affairs contributes to the complexity of any performance analysis, capacity planning, and workload predictions.

Thus, FPRN is a highly complex and rapidly growing LSP system. The factors, listed above, lead to a prolonged growth and to significant structural changes in the FPRN traffic, and these again make methods for LSP capacity planning a most interesting area of concern.

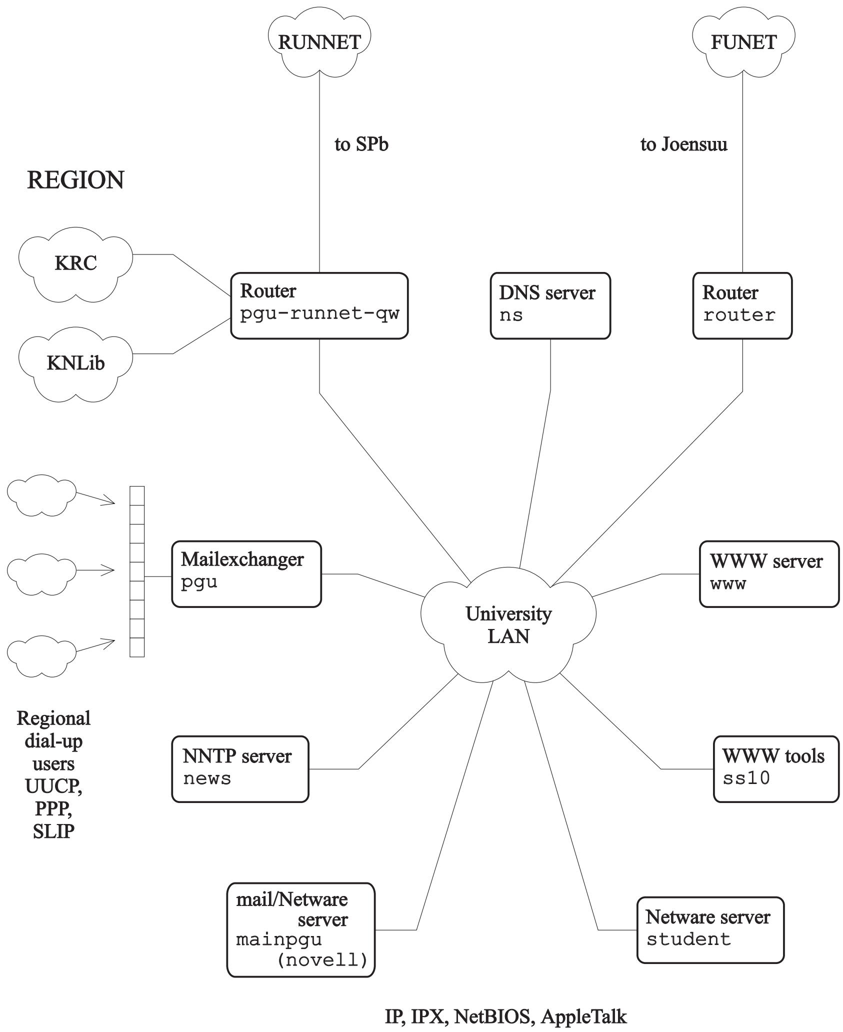


Figure 4. Federal Petrozavodsk RUNNet Node scheme before 1st March 1997

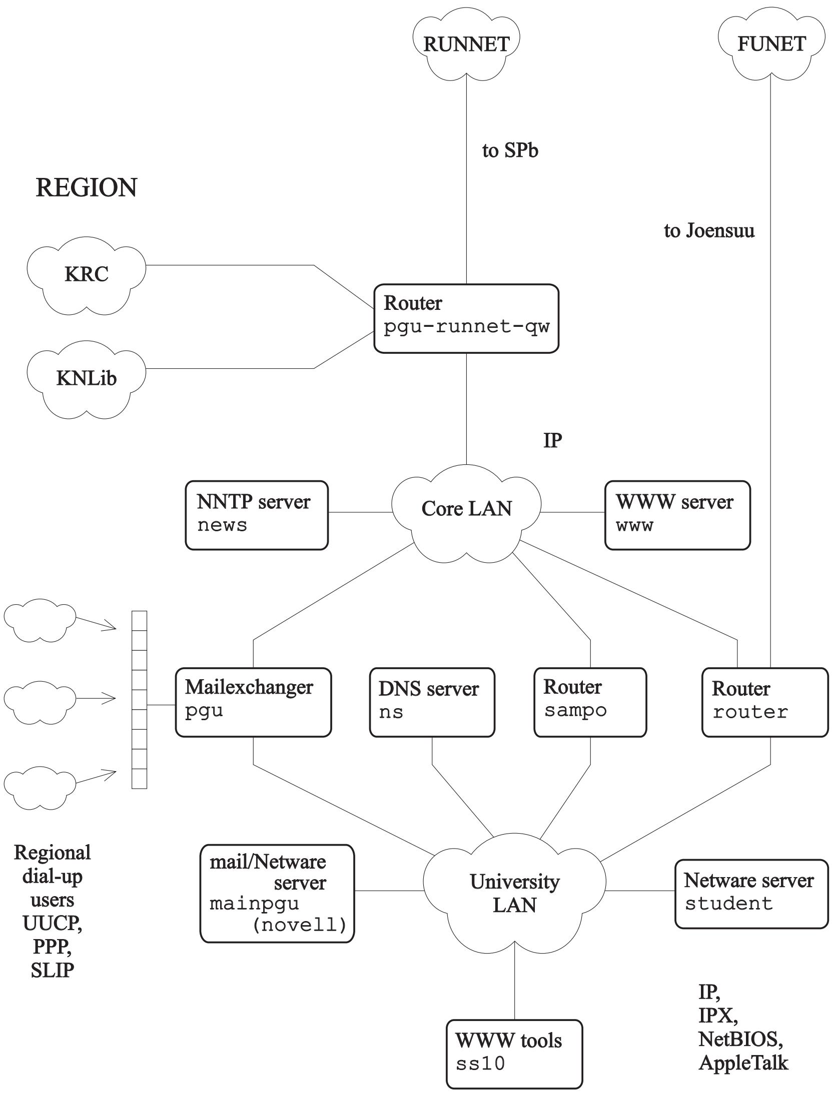


Figure 5. Federal Petrozavodsk RUNNet Node scheme after 1st March 1997

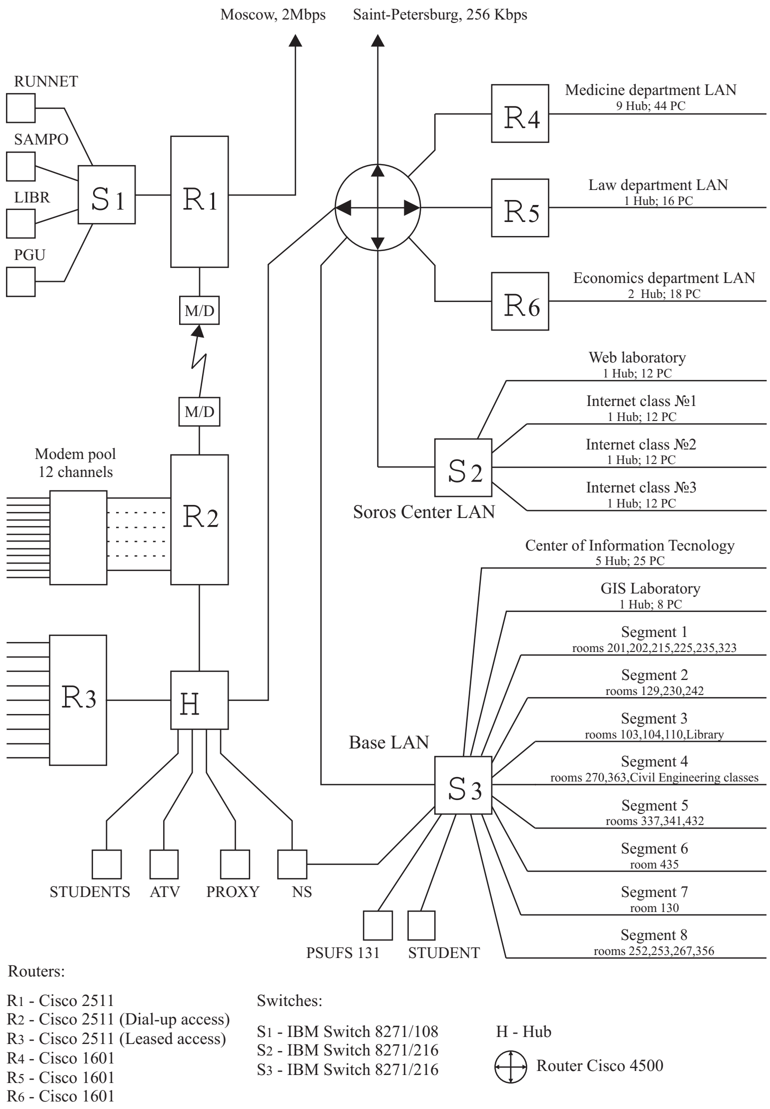


Figure 6. Federal Petrozavodsk RUNNet Node scheme at 1st November 1998

# 5 Local Service Provider Subsystems and Capacity Planning Problem

It can be seen from the Figures 4, 5, and 6, that through all the changes the structure of FPRN remains essentially same; it consists of four main subsystems: customer–access servers, internal Internet servers, the internal LAN, and external channels to higher–level service providers. The stability of this structure is quite natural as the structure elements correspond to the main functions of LSP. This has been noted by other researchers as well, see for example [13].

Thus, it is quite reasonable to use a composite LSP model, which consists of four subsystem models, all connected together. An advantage of the approach is that the different upgrades—hardware, software or network technology—can be treated separately: the basic work in the analysis concentrates on one individual submodel, and after the upgrade one single model has to be revalidated.

A typical capacity planning problem in an LSP environment is how to provide stable quality of service for an increasing number of customers when—at the same time—all the service needs grow: the diversity of Internet services grows and their demands for bandwidth grows.

Practice shows that the natural desire of LSP administrators to survive on the market forces them into a never–ending spiral of fast upgrades. The starting phase of each subsequent round of the spiral is characterized by the wish of the LSP administrator to satisfy the customer's needs for more bandwidth and, accordingly, to balance this with increasing costs of the "LSP–ware".

From a capacity planning point of view there is a subtle feedback in the spiral process. Typically, in a small environment, increasing capacity improves the quality of service, improved quality increases usage of the system, which again worsens the quality of service; worsening the quality decreases the usage. Hence, if the system is close to the saturation point, its state becomes—at least to some extent—unstable.

Due to the fast growth of the service requirements the capacity planning of an LSP must be based on reliable predictions of growth and on reasonable decision making. It requires adequate knowledge of the current performance of the system, of the capacity limits of the LSP subsystems, and of a reasonable forecast about the future workload.

As the first step we try to investigate problems of customer–access server modeling [15] and problems related with the current performance and workload characterization of an Ethernet segment [16].

The external channels of the LSP and the internal servers are topics for further research.

# 6 Performance Analysis and Capacity Planning of a Customer–Access Server

Customer–access servers (ACS) have some special features which create certain challenges to mathematical modeling. The ACS's have a compound structure of traffic, they are often heterogeneous, and they are based on multichannel techniques.

The compound structure of ACS traffic is specific to a situation where remote users generate independent service processes. Typically a service request—a customer in the queuing model—consists of a sequence of independent units. Hence, one can consider the service process to consist of a sequence of consecutive steps. The problem is to construct and analyze distributions which reflect this structure of the customer service process. Further, such general characteristics of ACS as the distribution of the number of busy channels or the loss probability are also of vital interest. Thus, the analyst has to use mathematical schemes which allow to obtain information about the system's behavior at different levels of detail, and, if necessary, make it possible to consider general and detailed characteristics all together.

A queuing system with service behavior specified to model the ACS characteristics is investigated in paper [15]. The corresponding random process is constructed, and theorems on its properties are proved; in addition, some steady–state characteristics of the queuing system are obtained. The applicability of the results has been tried in a real ACS environment using the FPRN modem pool as the experimental basis. Results of the pool modeling are presented, including a fitness test for the distributions obtained from modeling and the measured data of the real pool traffic.

The example of the real ACS modeling presents a customer call loss probability, which depends on three parameters. These are the arrival rate of remote calls, the message–group size, and the service rate, which depends on the hardware and software characteristics of the server and of the modem pool. The obtained dependence allows to plan when and how to update the modem pool to keep its performance at the required level.

# 7 Performance and Workload Characterization of an Ethernet Segment

Internal LAN is one of the most important LSP subsystem that, in most cases, is implemented by Ethernet technology and consists of several segments. The LAN traffic reflects the behavior of the local users, their interaction with an LSP servers and with external channels.

The performance of an Ethernet segment depends on many parameters: specifications of the standard and choices made for the specific configuration, characteristics of the hosts, implementation of interfaces, and the nature of the workload. The situation is getting worse by the unpredictable qualitative changes in the workload characteristics, to be expected during the life cycle of an Ethernet segment.

First stages of a segment capacity planning problem solution are [11] current performance estimation and workload characterization through traffic measurements. The paper [16] describes a practical software methodology for solution of this two problems.

There exists a number of software packages for measuring, collection, processing, and visualization of Ethernet segment traffic data. However, the packages are rather expensive and they are not flexible enough because their source code is not available. Therefore the research presented in [16] has had two concurrent goals: to obtain current performance and workload measured characteristics of Ethernet segment, which provides data communication service for local users of FPRN, and to develop software tools, which are flexible enough to be used as a base for obtaining analogous characteristics of arbitrary Ethernet segment.

The tools operate under OS Linux environment and are implemented as set of filters, based on free software packages. The set of filters developed uses special two layer packets capturing scheme. The set is easy to use and flexible enough. It may be used for an Ethernet segment performance monitoring and for partitioning of workload on various classes according to requirements, which may appear at a segment capacity planning.

The set of filters was used to obtain measured performance and workload characteristics of local users segment of FPRN. Packets delivery were considered as main service of the segment. Measurements were specified to obtain bytes and packets throughput, utilization, and main fractions of traffic appertaining to different low levels protocols and to servers connected to the segment. There were measured additionally packets length empirical distributions for each fraction of traffic separately, and number of active hosts (sending at least one packet during a minute).

The results obtained allowed to develop a simple technique for Ethernet segment capacity planning. Further, it gave reason to formulate a practical approach for a case, where the Ethernet segment, as an element of LSP, transfers both internal and external traffic, for example, from an access server to a router.

In this case the two components of traffic can be measured independently, and the influence of the external traffic will be taken into account using current averages, only. Hence, the technique can by applied for composite modeling of the LSP traffic. Unfortunately, due to space limitations the details of the approach will be postponed to a follow–up article.

## 8 Conclusion

Considering the growth process of the Internet infrastructure in the Republic of Karelia we emphasize the importance of conducting the planning process in some reasonable way. We suggest a decomposition of the planning problem and concentration on such important elements of the infrastructure as the local service providers.

Proceeding to the capacity planning of the providers we notice, again, that solution techniques in this area can conveniently be based on decomposition techniques. The analysis of the FPRN growth–process shows that an LSP can be decomposed into four main subsystems: the customer–access servers, the internal Internet servers, the internal LAN, and the external channels to higher–level service providers.

Approaches to modeling of a customer–access server and to modeling of an Ethernet segment are briefly described with references to more detailed papers of our colleagues, also presented in these Proceedings.

Further research will be devoted to composite models of LSP workload and LSP performance.

## Acknowledgments

We want to thank the technical administrators of FPRN Sergei V. Matcko and Alexei M. Minin, who spent a lot of time in consultations on FPRN configurations and on working procedures. Discussions on the research with Kimmo Raatikainen, Markku Kojo and Tiina Niklander, taken place during FDPW'97 in Petrozavodsk, were very useful.

The students Vadim Ponomarev and Oleg Shtchoukin spent in 1996 their last term with the Department of Computer Science at the University of Helsinki. They have carried out a remarkable amount of work with the measurements and with the statistical processing of the monitored data, using knowledge received in networking courses of Markku Kojo, given during FDPW'95 in Petrozavodsk and then, with assistance of Jarkko Sevanto, in Helsinki.

And we express our gratitude to Eugen Borzenin, Ian Kovru, Maria Shreders, and Sergei Zavialov from the Department of Geographical Information Systems for the carefully prepared Figure 3 of the paper with map of Internet infrastructure of the Republic of Karelia.

## References

[1] *Timo Alanko, Iouri A. Bogoiavlenski* Performance analysis and capacity planning for a message server in a regional data–communication network. Plan of Joint Research. Universities of Helsinki & Petrozavodsk, Departments of Computer Science, 1995

[2] *Victor N. Vasiliev, Natalia S. Ruzanova, Iouri A. Bogoiavlenski* Performance Analysis of Petrozavodsk RUNNet Node in conditions of its configuration fast change, in Abstracts of Conference Telematica'97, Saint–Petersburg, 1997, pp.52–54 (in Russian)

[3] *Tracie Monk and K. Claffy* End-to-End Dynamics, Internet Data Acquisition and Analysis: Status and Next Steps, National Laboratory for Applied Network Research, USA

http://ftp.uni-mannheim.de/ftp/info/inet_97/f1/f1_3.htm

[4] The Internet Statistics and Metrics Analysis (ISMA) workshop (February 1996 and May 1997):

http://www.nlanr.net/ISMA

[5] *Vern Paxson* Towards a Framework for Defining Internet Performance Metrics

http://www.advanced.org/Almes/INet96/Paxson/

[6] *Les Cottrell and Connie Logg* Network Monitoring for the LAN and WAN

http://www.slac.stanford.edu/~cottrell/tcom/ornl.htm, 1996

[7] The IETF's IP Performance Metrics working group

http://www.advanced.org/IPPM

[8] The IETF's Real–time Traffic Flow Measurement working group

http://www.ietf.org/html.charters/rtfm-charter.html

[9] *Kimberly C. Claffy et al.* Tracking Long–Term Growth on the NSFNET, CACM, August 1994, V. 37, Number 8, pp.34–45

[10] *Larry Press* Tracking the Global Diffusion of the Internet, CACM, November 1997, V. 40, N 11, pp.11–17

[11] *D. A. Menasce, V. A. F. Almeida, and L. W. Dowdy* Capacity Planning and Performance Modeling. Prentice Hall, 1994, p.412

[12] *A. V. Sigalov* Internet. Russian Resources, Saint–Petersburg, Piter Press, 1997 (in Russian)

[13] *Valeri A. Vasenin* Russian Academic Networks and Internet, Moscow State University, REFIA, Moscow, 1997 (in Russian)

[14] http://www.nw.com

[15] *Olga I. Bogoiavlenskaia* Access Systems Modeling by Queues with Compound Customers, in this Proceedings

[16] *Vadim A. Ponomarev, Iouri A. Bogoiavlenski, Timo Alanko* An Ethernet Segment Performance and Workload Characterization Using Set of Filters Based on Free Software Tools, in this Proceedings

[17] *Van Jacobson, Creig Leres, Steven McCanne* Package tcpdump, ftp://ftp.ee.lbl.gov/tcpdump.tar.Z