Mat-2.177 Seminar on Case Studies in Operations Research

Modeling Peer-To-Peer Networks

Project Report

06/21/2007

Client: Nokia Research Center

Contact: Jukka Nurminen

Nokia Project Group

Lauri Hyry, project manager

Emilia Etelä-Aho

Heidi Kettunen

Eero Nevalainen

Table of contents

I. Introduction

This study is part of the course "Mat-2.177 Seminar on Case Studies in Operations Research" in spring 2007. The client of this project is Nokia Research Center, which is a separate unit within Nokia.

The study concerns modeling peer-to-peer networks (P2P) and observing message diffusion through simulation. The aim is to observe how heterogeneity of nodes in a network of mobile phone users affects the message diffusion. The client did not provide the group any specific topic or data to be studied, and consequently narrowing down the scope of the project was done partly as the work progressed.

A pure peer-to-peer network is defined as a network of equal members, or *peer*s, connected to each other in some way. A P2P network structure has emerged especially in IT applications, but it may also be used to model for example social phenomena. For IT purposes, P2P differs from the traditional server-client relationship, as it requires no central members, i.e. nodes, and thus the responsibility for the functioning of the network is not assigned for a single node.

The client was interested in information propagation in an unstructured P2P network, as new applications such as advertisement linked to mobile phone use is emerging. Also, applications such as iPTV, in which users follow TV broadcast on the Internet with a file sharing applications, are emerging.

II. The progress and execution of the project

i. Initial Project objectives

The objective of this study was to examine in practice the message diffusion in the P2P context with 2 or 3 different cases. These cases were to be simulated with Matlab programming to see how messages would be circulated in the network.

The intent was to create the model in stages: (i) new-year's humorous SMS messages with all nodes treating the message homogeneously, (ii) a network with the message having some content characteristics and nodes preferring some content over another and (iii) mobility and dynamic behavior effect on message diffusion.

A literature review was planned in order to get to know the P2P field better, but also to see whether this sort of research has been done previously. The overall picture on P2P studies was also set to help in narrowing down the scope of the project as it progressed.

ii. Earlier progress of the project

The progress of the project through halfway and problems encountered were reported in detail in "Status Report" of the project in March 2007. Matlab and Java were used in programming the simulation model, and the implementation of the model was discovered to be more difficult than estimated. This was mainly due to incomplete interface documentation, but it was also acknowledged that the initial timetable underestimated the time needed for programming. As result, it was accepted that possibly no time would be left for implementing the more complex features to the model (such as mobility).

iii. The project scope at the end

The final extent of the project has been limited to the creation and exploration of the simple model as the heterogeneous simulation model was not successfully implemented in time.

III. P2P network overview

i. Nodes

Network nodes can have different characteristics; e.g. homogeneity, heterogeneity, credibility, and activity. Heterogeneity means that nodes in the network may differ in their preferences on delivered messages while homogeneous nodes send and receive all kinds of information in a similar way. Messages sent by high-credibility nodes are taken more seriously than those sent by others. Active nodes send messages more promptly than inactive ones. There might also be limits on how much information one node can handle at same time.

ii. Networks

Large, if not all, P2P networks include routing, and as such are structure networks. This means that in addition to the direct information flow between nodes, some information may be routed from one node to another through intermediate nodes. Reliability of the message delivery may e.g. refer to successful information distribution under transfer capacity constraints. Filtering, in P2P context, means blocking messages, whose content is considered inappropriate.

Scalability can be used to explore how the network can handle increases in network traffic, or the joining and leaving of nodes. If the network nodes know only one other node, the network may handle greater amounts of traffic as all information received by a node is sent further only to one node. When one node leaves such a network, it may result in reducing the formerly bigger network to two isolated smaller networks.

P2P networks are relatively new. Traditionally computer networks have had a client-server architecture: the network may consist of many client nodes e.g. personal computers (clients) and one central node e.g. web-server (server). In a television network, based on computer technology that is built on client-server architecture, the content is downloaded only from the central node – the television server - to the client nodes.

However, P2P networks do not necessarily need any kind of central node. Instead of communicating with central nodes, the client nodes communicate directly with each other. P2P television networks, without central television servers, have to provide the content to the client nodes in a completely different manner. In practice, client nodes of P2P network often disconnect from the network and then later re-join in. Because every node has its part in transferring the content, the disconnecting and connecting may affect some nodes in the network. On the other hand, with client-server architecture, the disconnection of a client node does not affect other clients, who mainly communicate with the central node. The continuous transformation of the P2P network, as nodes connect and disconnect, also makes the routing of information traffic more complex when compared to client-server architectures. In P2P television network, the goal is to make sure that all the clients receive their content despite the transformation of the network, and as such the disconnection of a node should not isolate others from the network.

Complex routing requirements of P2P networks may be hard to fulfill but P2P networks can also have many advantages considering network bandwidth, flexibility and even reliability. Communication failure between some nodes in network doesn't greatly affect the overall performance of P2P network. Communication problem that isolates the television server from the rest of client-server network stops the television content broadcast completely. The requirements for bandwidth of client-server television increase rapidly if the network grows (more clients watch the television broadcast). In P2P television networks the content received by one client comes from many different sources that send only a piece of the content. These small information streams are then combined into one single TV-broadcast on each node, and in this way bandwidth requirements are significantly lower.

iii. About mobile phones

Modern mobile phones are more than just phones. They are often capable of playing video, music, surfing the Internet and running games. The increased diversity of features in mobile phones, or mobile multimedia devices has taken them closer to traditional PC's. These devices are also capable of using wireless TCP/IP networks and communicating over Bluetooth just to name few. There already exist P2P applications for mobile devices. These applications are called Mobile Peer-To-Peer applications, (MP2P).

IV. Literature review

i. Overview on Peer-To-Peer network research

Literature overview on peer-to-peer papers shows that P2P networks have been studied especially regarding file sharing applications, since in that field the use of P2P architecture has become very popular in recent years. Studies on different specific algorithms or efficient network structures were abundant in this field, but papers are focused on narrow subjects within the field, such as specific software such as Gnutella or eDonkey. As a result, no fundamental theory or prevailing mathematical model was found regarding P2P networking. Overall, studies were usually conducted by formulating (programming) a peer-to peer model with certain assumptions, and then verifying the underlying behavior or improved performance by simulation. Usually, simulation results showed exponential characteristics; a result which is intuitively clear as well.

The papers more carefully studied within the project work were related to diffusion models and on the other hand heterogeneity in the P2P context. This area did not include as many articles as mentioned above, since network reliability and efficiency of file sharing networks seems to be a more useful research area at the moment. Research was conducted by reading through articles given by the supervising professor of the course, and by searching the library databases for related papers.

ii. Epidemic diffusion models

Existing mathematical models of epidemic models can be used in both medical research and in modeling information diffusion in peer-to-peer networks. Diffusion in P2P networks can be assumed to act similarly to epidemic spreading of diseases. Khelil et al (2002) used epidemic models to describe and to adapt the information dissemination in MANETs. MANET is a mobile ad hoc network with no predetermined structure. They developed an epidemic model for a simple information diffusion algorithm based on simulation results. The main focus was to analytically investigate the impact of node density on information diffusion.

Although these epidemic models have been used in different areas of research, the main pitfall lies in the following assumption: it is assumed that an object, which carries the message, will meet other objects in a random pattern. In this traditional model, it is assumed that only with random mobility, the critical mass of the nodes is reached for the information to spread fast. Thus, traditional diffusion models are not suitable for modeling real world mobility due to their random mobility patterns. (Kurhinen & Vuori)

Papadopouli & Schulzrinne (2000) developed a simple epidemic stochastic model to determine the delay until data spreads to all mobile devices. This model is a simple continuous-time Markov process i.e. pure birth process. Their network consisted initially of only few nodes. Later, in 2001, Papadopouli & Schulzrinne introduced a diffusion-controlled model which is used for data dissemination between mobile queries and static servers. In this type of static trapping model, mobile nodes perform diffusive motion in the space and are absorbed on the straps i.e. servers when they step on them. However, the static strapping model assumes that nodes are fixed, and they do not cooperate with each other.

iii. Information diffusion in a MP2P network

Information delivery in a rapidly changing network topology is a challenging task. In general, there are two different methods to deliver messages: proactive or reactive. This means that network nodes can communicate proactively being active parts in transmitting data, or the nodes can communicate reactively where the communication link is established on-demand when required. Both of these methods are applicable to mobile peer-to-peer networks. However, in the proactive method, the information is sent to all of the nodes whether they benefit form it or not. (Kurhinen & Vuori)

Both of these methods have a common feature: messages are sent via other network nodes. This is the only way to transmit data between nodes if they do not have a direct data link between them. In MP2P context, the network is changing constantly when the nodes are moving. In a single-hop solution, this is simplified by assuming that a source node transmits data to the target without any intermediate nodes. Thus, by this simplification, a complex routing problem is avoided. In a single-hop model, nodes have two possible ways to get information: (i) either nodes find their ways to a location, where they can request the information reactively, or (ii) they can receive that information from a node who is serving it proactively. (Kurhinen & Vuori)

The model of Kurhinen & Vuori is called an exchange pipe model and it can be thought to represent a busy promenade where the users of mobile devices are packed tight. The density of the mobile users is much lower outside of the promenade. People enter and leave this promenade from its ends. Kurhinen &

Vuori created a simulation model to observe thickenings of mobile nodes in which information is able to transmit from node to node. They concluded that the diffusion curves for the diffusion models were not S-shaped as in the traditional diffusion models but the most intense growth took place at the very beginning of the diffusion process. According to their results, the information diffusion is almost explosive in this type of pipe. However, this model is best suited for dense urban areas where mobile penetration is high.

iv. Heterogeneity in P2P context

Heterogeneity of nodes in the P2P networks refers here to the different approaches nodes use in handling information. The model developed within this project is an unstructured P2P network imitating an ad hoc real life network as opposed to structured networks used e.g. in most applications file sharing applications. Research on efficient topologies of file sharing P2P networks is rich, and these papers often connect heterogeneity to different characteristics of the nodes affecting network performance. Such characteristics may include broadband speed or total file size in the content sharing environment, and the network is structured so that nodes with larger capacities are used for routing to enhance the P2P performance and reliability.

These stronger nodes are often called "super nodes" to differ them from the "weak nodes". Information routing tactics through stronger nodes were introduced to file sharing applications early on, as a pure P2P network with limited connection capacities of the weaker nodes and random information movement becomes quickly clogged. Efficient topologies and related algorithms have been studied quite widely; for example, Mudhakar et al. (2006) have modeled and simulated heterogeneity-aware topology effect on P2P network performance. They have concluded that a hierarchical structure for information routing, with strong nodes on the top of the pyramid and no reciprocal interaction allowed on lower levels, enhances the performance and reliability of the overall network. The super node approach brings the P2P closer to the traditional server-client IT structure.

The host-client characteristics are studied also in the paper of Wei Yu et al. (2005) on computer virus propagation in a file sharing P2P network. In short, they show by simulation the intuitive result that an attack strategy of hitting server nodes results in more rapid network virus infection compared to random selection of individual target nodes or spreading the virus online to random connected peers.

V. Simulation models

i. The simple model for the P2P network

The model formulation and simulation somewhat overlap the project work of one group member for the course 'Mat-2.170 Simulation'. Some results presented below are also included in his project work report.

The framework for the model was programmed using Matlab and Java. Matlab creates a given number *n* of nodes to the network, after which a predetermined number of connections *m* is randomly distributed between nodes. Two nodes may have only one direct connection between them, and no nodes may be left outside the network altogether (Figure 1).

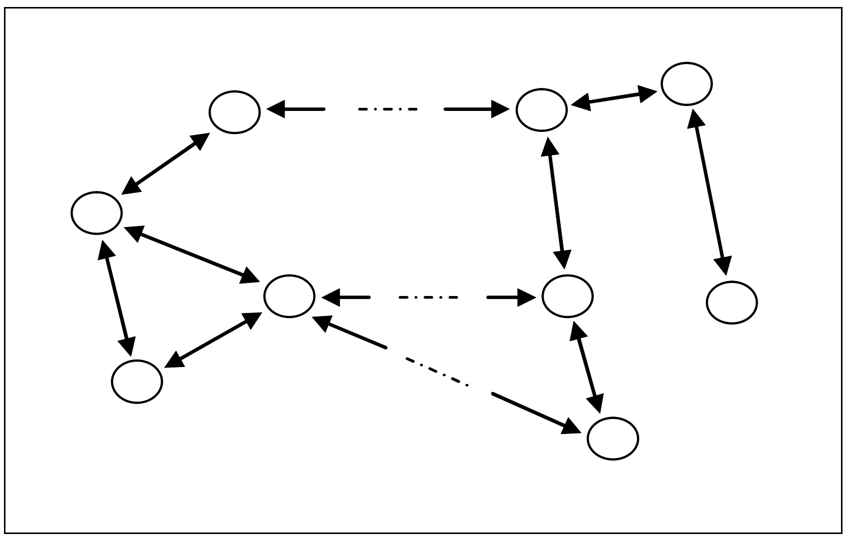

Figure 1 A P2P network with *n* nodes and *m* connections

There is an activity parameter *a* created for the nodes to describe their willingness to send the received SMS message on to friends. The binary decision to pass the message on is made separately for each of the friends, i.e. peers directly connected to the sender.

The values for node-specific activity follow the beta distribution with known parameters. Behavior is studied and compared with two distributions; with parameters $\alpha = 2$, $\beta = 2$ and $\alpha = 2$, $\beta = 5$ (Figure 2). The latter one describes a significantly lower activity of sending the SMS on, this could be used for example to illustrate the attitude towards advertisements on in a social context versus normal content messages.

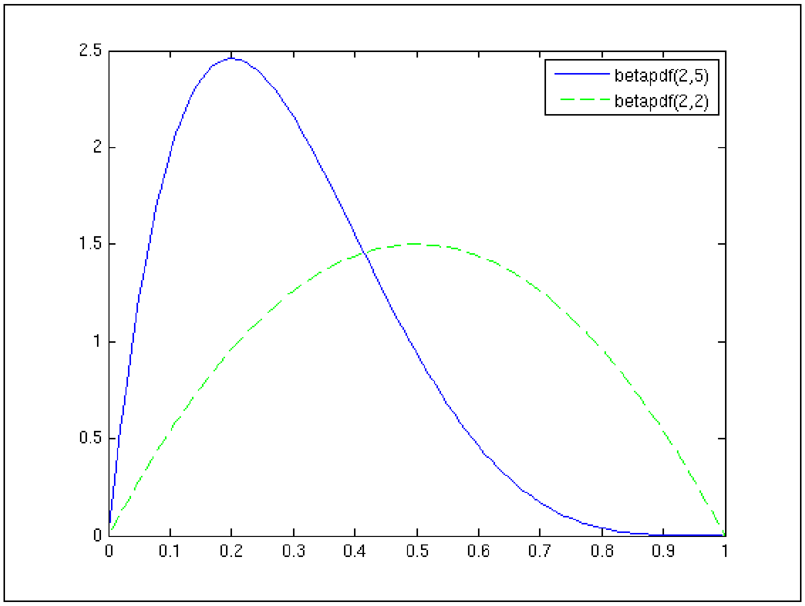

Figure 2 Probability density functions for the beta distributions used

After receiving the message, the node starts going through its "address book", i.e. the connected nodes, in a random order. A uniformly distributed random value [0,1] is compared to the node's activity level *a*, and if smaller, the message is sent to the friend in question. Thus, nodes with a low activity parameter value will send messages on less likely compared to highly active ones. If the node has received the message already once, it will not be sent on to anyone.

The sending process takes time, more specifically 2 seconds (illustrates also the network delays etc), if the message is sent, and 0.5 seconds, if the sending decision is negative. Furthermore, there is an initial delay of 3 seconds before the node starts checking its address book after receiving the message. Simulation runs in the network are done in discrete time. At this first stage, the message has no content parameter, so the message does not affect the activity of the node. The total time taken by the node to send and go through the address book may be limited; this is noted by total time limit *t* (Table 1).

Table 1 Variables in the network

	Range
n, number of nodes	
m, number of total connections	[n, n*(n-1)/2]
m/*n* = average number of connections per node	
a = activity level, [0,1]	a ~ Beta(α,β)
***t*, sending constraint for total time of one node**	

1. Model validation

The validation of model was done by simulating message propagation with quite small network consisting of a hundred of nodes. As there was no real-life data available from the message propagation of such a network, the results provided by the model were analyzed statistically. Validation included regression analysis and investigation of simulation results and parameters (Figures 3 and 4). In the process of validation, some distinct faulty model behavior was identified and fixed.

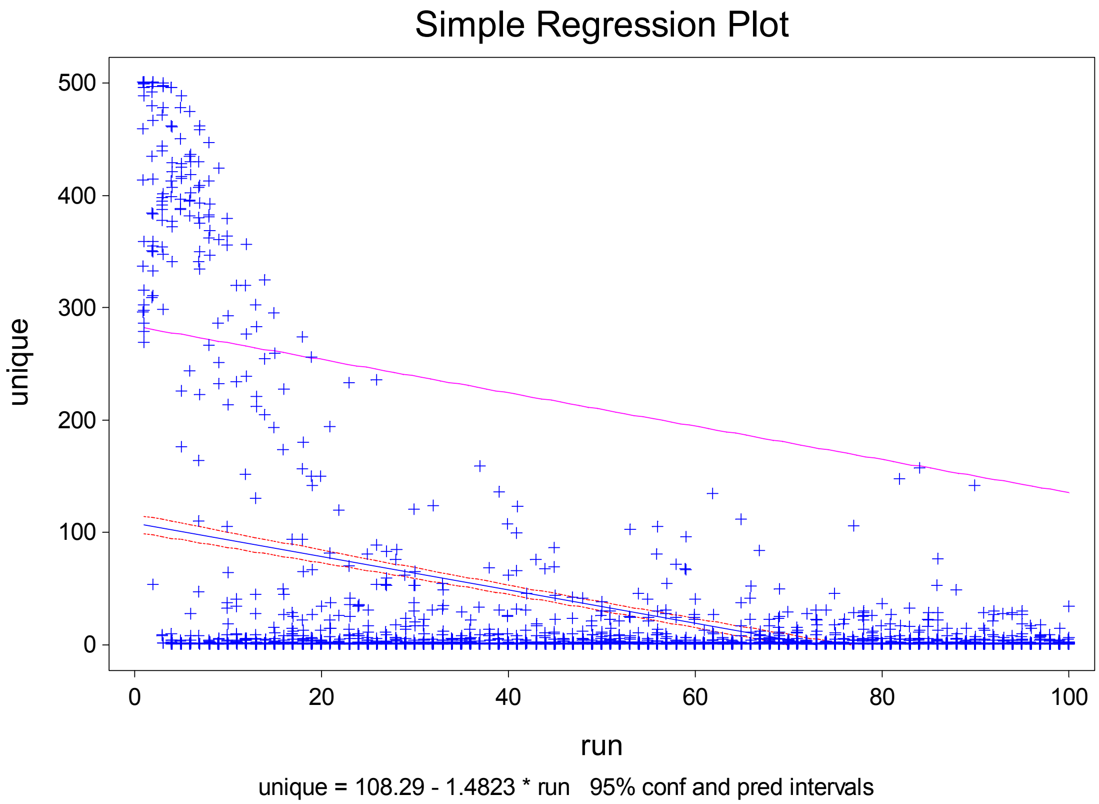

Figure 3 Regression plot from validation runs before fixes. Vertical axis describes the number of nodes that received the message at least once. The number of simulation is shown on horizontal axis. The propagation level of message depends strongly on simulation run of the model.

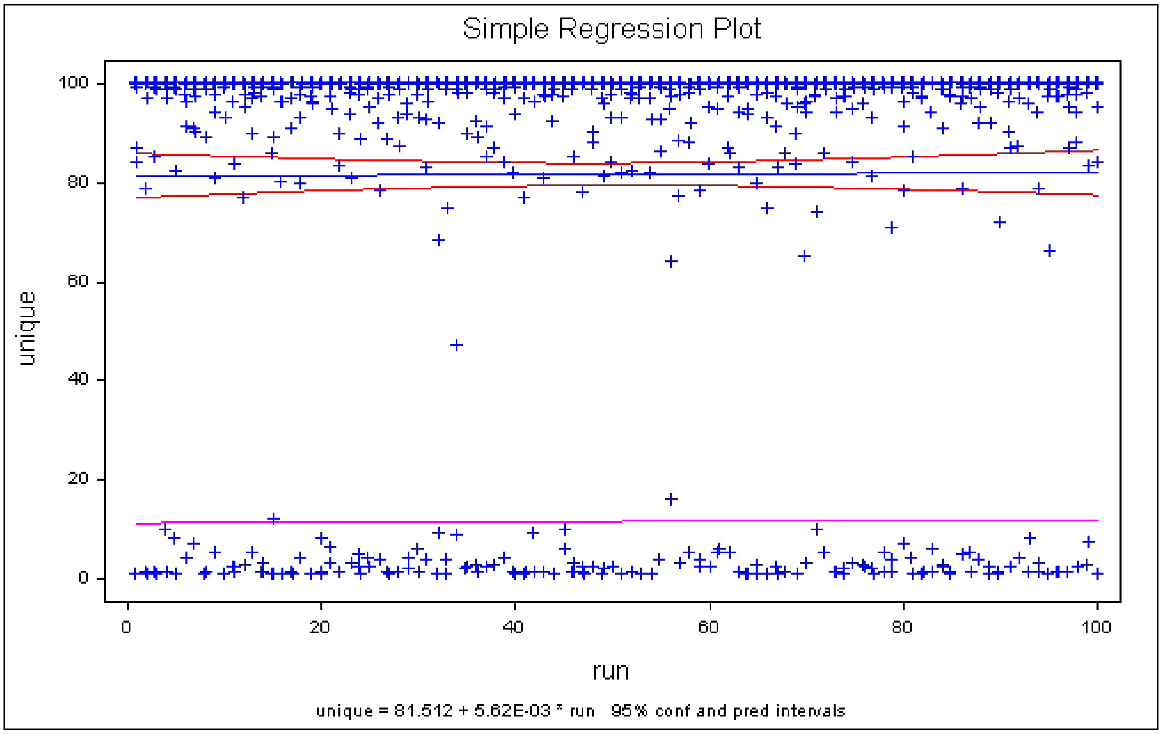

Figure 4 Regression plot from validation runs after fixes: the propagation of message does not depend on the simulation number.

The validation process was a challenging task. Even the simplest model runs used in quick validation checks took roughly ten minutes. It was also hard to identify whether the model failures were caused by Matlab that controlled the simulation, Java virtual machine that ran the simulation or both of them.

2. Varying the number of average connections between nodes

At the beginning of the simulation, one message is sent to one of the network peers, which will then decide on passing it on to his peers. The first simple simulation run is done with 100 nodes with varying the average number of connections per node in the network from 1 to 10. The time constraint for one node is set to 60 seconds (so it does not affect the process), and underlying activity distribution is Beta(2,2). Each m/n (marked *average connections* Figure 5) is simulated with 100 simulations.

Among others, the time T, at which the message diffusion ends, and the number of nodes reached r are recorded. Naturally, only one variable of the network is varied at time.

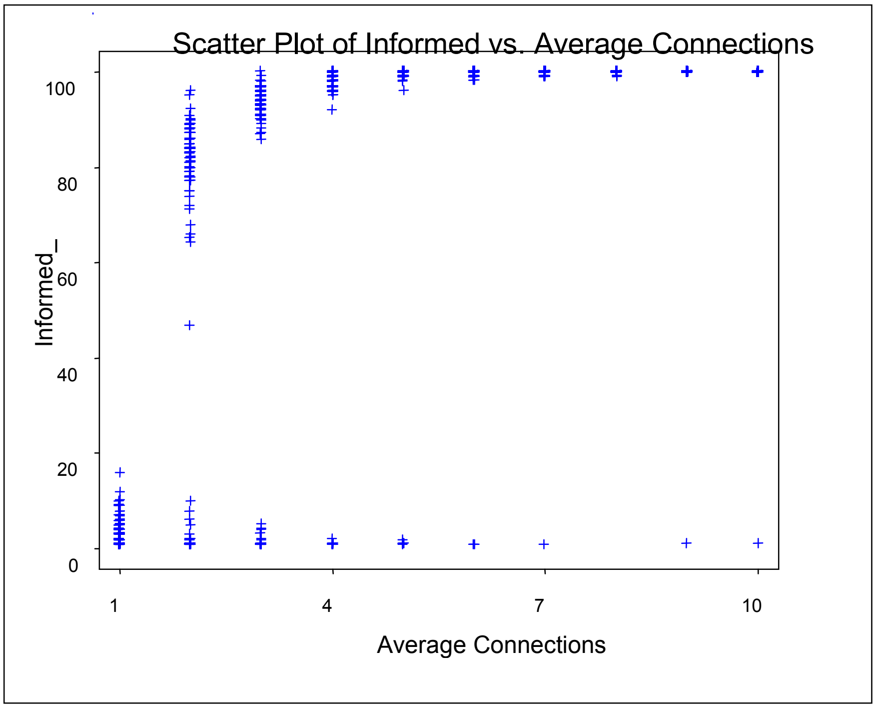

Figure 5 Network of 100 nodes, number of informed vs. average connections per peer, Beta(2,2)

Naturally, with only one connection allowed between peers, the message never reaches everyone in this context. This particular situation is a line, and not a network as such, and the overall probability of every one of the 100 nodes deciding to pass on the message is extremely low. This arises from the fact that the average probability of sending the message further to one node is 0.5 (average of Beta(2,2) distribution is 0.5). The simulation runs illustrate the exponential characteristics of the network well; the message diffusion stops at the initial steps, or reaches almost everyone very quickly as the number of connections is increased. Already with 4 connections per node we have a high saturation, and only with m/n of 2, there occurs one situation where the message has reached halfway, but the diffusion stops. This implies that either a very small or large number of nodes are reached.

The rest of the simulations are done with networks of 500 nodes. The connections per peer are varied between 1 and 18 (Figure 6). The number of simulation runs is lowered to 10 to save time, as the creation and simulation in Matlab of a 500-node network may take long time. Time constraint for one node is set to 90 seconds.

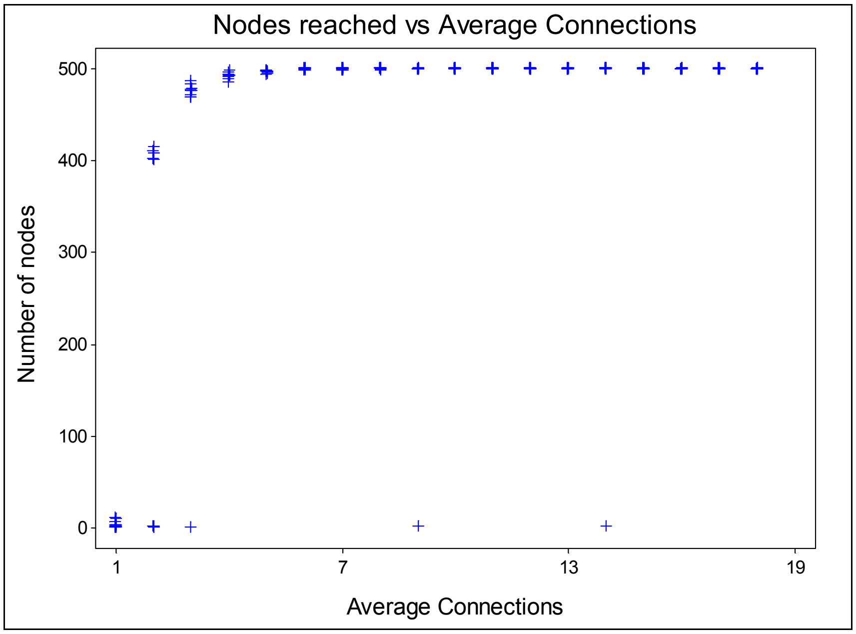

Figure 6 Network of 500 nodes, number of informed vs. average connections per peer, Beta(2,2)

The simulation runs illustrate the exponential characteristics of the network even better now. Naturally, also in the network of 500 nodes, if only one connection is allowed between peers, the message never reaches everyone. Figure 6 shows that if the average number of connections between peers is over six, all nodes in the network will almost every time receive the message. There occurs only four situations where the diffusion of the message delivery stops before it has reached any nodes.

The same simulation with 500 nodes, ten runs and activity distribution of Beta(2,5) provides the following results, Figure 7.

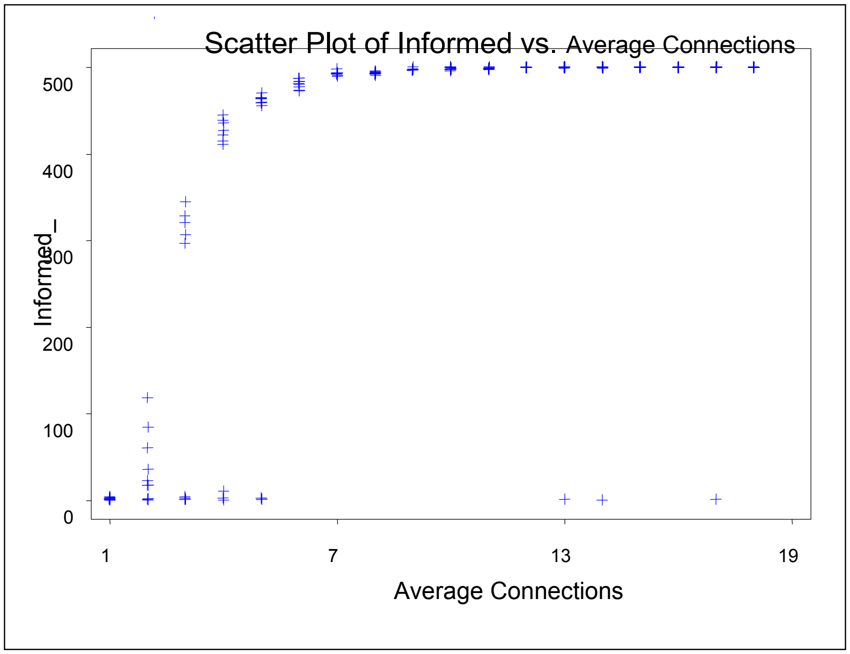

Figure 7 Network of 500 nodes, number of informed vs. average connections per peer, Beta(2,5)

The bias of the activity values to left by using Beta(2,5) can be seen by comparing Figures 6 and 7. The sending probability is much lower at 0.2857 (average of Beta(2,5) distribution), and thus reaching all nodes requires eight or more average connections per node. In simulations, where the message sent to the first random node gets immediately terminated, the node in question may have only few connections or really low activity level or both.

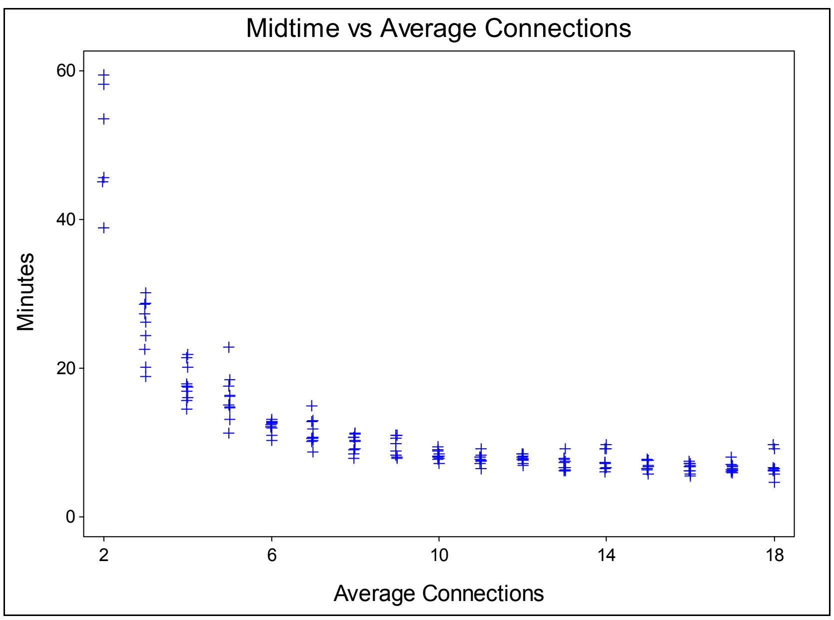

Figure 8 Network of 500 nodes, simulation time vs. average connections per peer, Beta(2,2)

Figure 8 illustrates the same simulation with simulation time on the vertical axis. On average, in the network of 500 nodes, it takes 6 min 44 s to reach all the nodes if the number of connections is between 13 and 18.

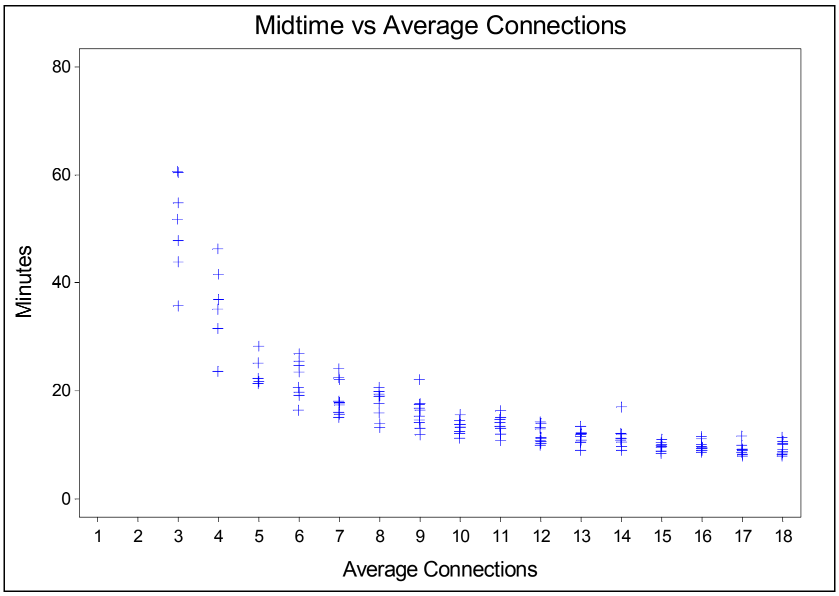

Figure 9 Network of 500 nodes, simulation time vs. average connections per peer, Beta(2,5)

Figure 10 illustrates the same simulation with simulation time on the vertical axis. On average, it takes 9 min 42 s to reach all the nodes if the number of connections is between 13 and 18.

3. Varying the network size

In the last simulation run with the simple model, the number of average connections is set constant to 18 and the number of the nodes varies between 100 and 2000. All other assumptions are kept the same.

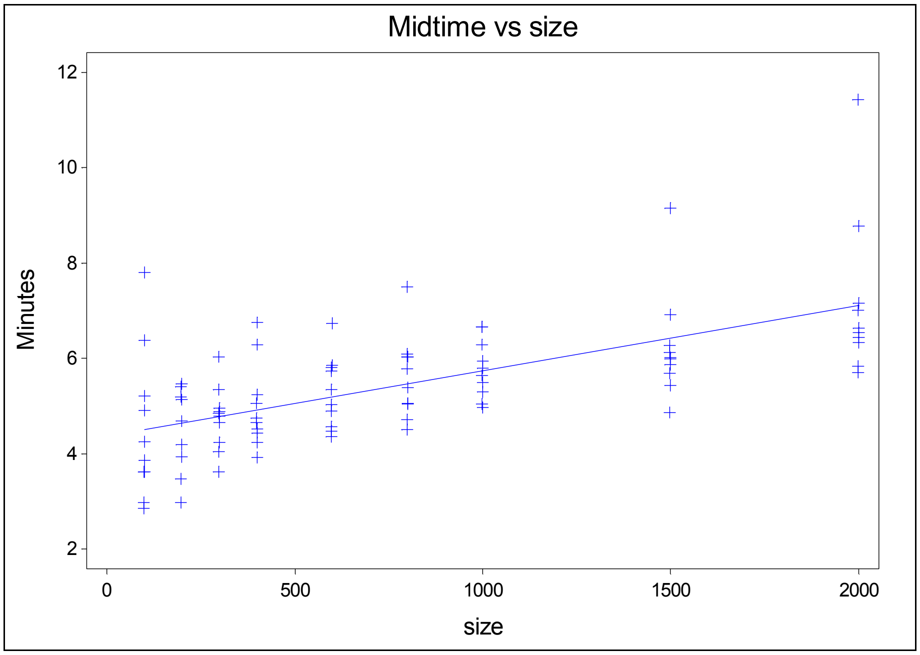

Figure 10 Simulation time vs. network size, Beta(2,5)

Figure 11 illustrates the results of the simulation run. Naturally, it takes longer time to reach all the nodes if the network is larger. Figure 12 illustrates the same model with percentage shares of the reached nodes. As the size of the network increases, so does the time to reach certain percentage of the nodes. More importantly, it may be noted that the time to reach one quartile of the nodes is drastically longer than the time taken to reach the next quartile, i.e. half of the population, and this shows the exponential nature of the message diffusion. For the next quartiles, the difference is not as significant anymore, which reflects the fact that nodes do not send on the message twice. The exponential spreading does not take place for too long as the network runs out of new nodes, and so the message diffusion rate in the network bears resemblance to the diffusion S-curve model.

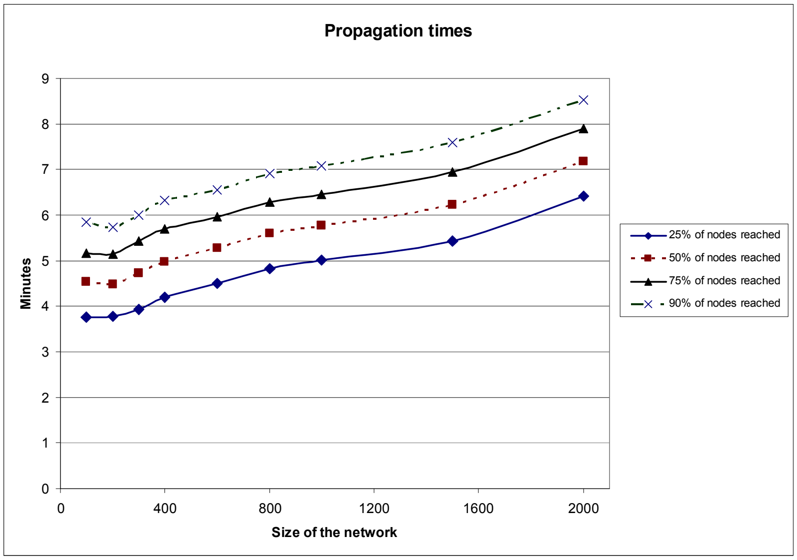

Figure 11 Propagation times versus size of the network

ii. The network model with preferences

Initially, the aim was to model a network of heterogeneous nodes with the following assumptions:

- The network has *n* different interest groups
- Each message in the network has a certain *interest value* for each interest group. The interest value thus describes the different aspects that people may have over the subject of message.
- Each node in the network has a certain *node activity* value in each interest group. The activity of group describes also the credibility of node when it sends a message. The most active nodes can be also seen having commonly recognized knowledge on subject and thus their information is taken more seriously.
- Figure 13 illustrates a small example of a network and its operation which has been parameterized as described

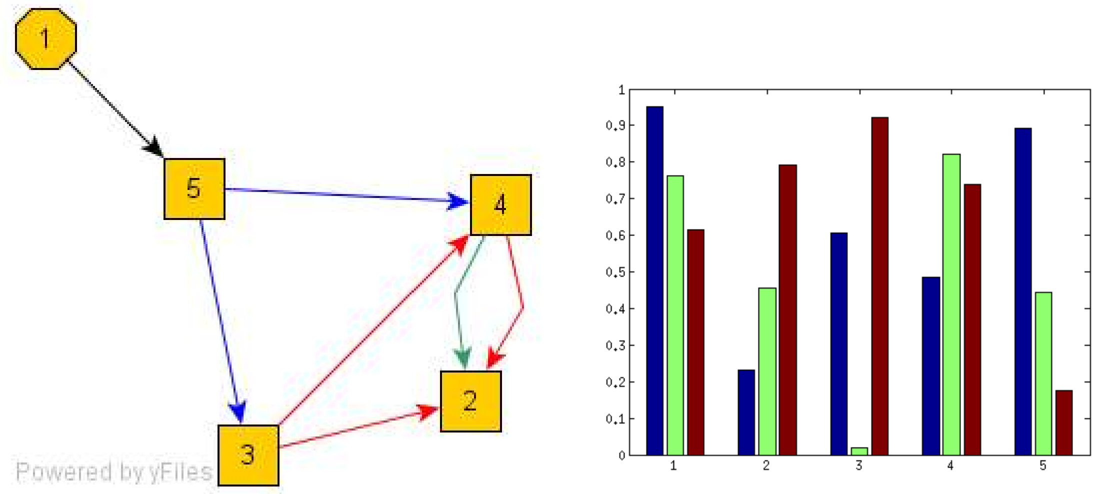

Figure 12 Overview of network with heterogeneous nodes

The object and values of index 1 represent the message that is arriving to the local network. Indexes 2-5 are nodes in the network. The message is first received by node 5.

After receiving the message, node 5 may decide to send it to one or more interest groups. The probability for sending the message for each group can be described e.g. as a function of nodes of interests, sender's credibility and messages interests.

Because node 5 is very active in the first group (darkest) and the message has also a high interest value for that group, the node will most likely send the message to the first ”communications channel” immediately.

Additionally, it may be assumed that the messages interest values decreases in time (for example exponentially). The probability for receiving and sending the message for each node could also be affected by the interest values.

In this example, likely receivers of the message through the first communications channels are the nodes 3 and 4. Depending on the receiving order, the message may spread on from the node 3 through the third channel or from the node 4 through either the third or the second channels. Finally, it could be assumed that the nodes also have the option to “re-broadcast” the message to increase its visibility value in the communications channel.

1. Status of the model and further development suggestions

Simulation model with interest groups is an extension of the New Year's SMS model. This extended model uses the same type of network as in the New Year's message propagation; the only difference is that now the nodes and messages belong to certain interest groups. Therefore, the message propagation depends on the node preferences.

The first version of the model has been implemented but it still requires validation and debugging in order to provide correct results. The main identified problem is that the simulation results depend strongly on the simulation run. Quite similar behavior was described in the chapter V.i.1.

As already mentioned, there are still problems that need to be solved. However, the model seems promising and once it is working correctly, it could be used to model several different scenarios of advanced message propagation. For example, the model could be used to investigate an optimal propagation strategy to reach as many people as possible in the network. In real-life, this information could be useful when e.g. gathering people to some event.

Although the model is suitable for depicting real-life message propagation, modeling accurately people's behavior can become very challenging. Therefore, without any research on the users' message sending behavior, the model provides only referring results.

VI. Conclusions

The first part of the study consisted of literature review on peer-to-peer networking. The review concentrated on the usage of diffusion model in the P2P context, and heterogeneity of the nodes. It can be concluded that based on the review, there seems to not exist any similar studies carried out previously.

The simulation results of the simple model with 500 nodes showed that when the number of connections between nodes increases, so does the likelihood to reach all the nodes within the network. Activity of the nodes following Beta(2,2) and Beta(2,5) distributions lead to slightly different results. Messages will reach the whole network faster if the node activity follows a Beta(2,2)-distribution. However, it was surprising that the entire network could be reached only with several connections; average connections below four resulted in small part of the nodes always remaining uninformed. This was due to the fact that the nodes do not attempt to send the message several times.

When the number of connection was kept constant at 18, quite naturally, it took more time to reach the nodes in the network if the number of nodes was increased. In all of the simulation runs, exponential behavior of the message diffusion can be clearly seen.

The interface between Java and Matlab has not been commercialized yet. The interface is complicated and many important features are still missing. There has been a great amount of work done over the simulation model but still the most important cases could not be finished.

More work is needed on the model to provide better results.

VII. Schedule, envisioned risks and their materialization

The most important risks related to the project success identified at the planning stage and follow-up report were concerning the programming work and related difficulties, the programming software and the actual results given by the model.

The programming work has caused some further delays for the project, and another unforeseen delay factor has been the long duration of the actual simulation runs on the model; the creation and simulation of a 500-node network may take several hours.

These issues have put pressure on the task of writing the final report, and some parts of the initial targets have not been implemented in the project. However, it may be noted that as the underlying model is well constructed, applying new features the model is relatively easy.

The coordination of the team by leader could have been better. The literature review departed to a different path from the simulation part as no similar cases were found from the literature. On the other hand, the implementation of even the simple models was problematic, and no new features and objectives were needed from the literature. Therefore the lack of coordination might not be the cause of failure for not reaching the more complicated objectives.

As stated in the status report, the risk having trivial results was increasingly present as delays occurred. Naturally, a simplified model of the real life produces intuitive results.

The project work required 450 hours of team work including seminars.

VIII. References

Khelil, A et al. 2002. *An epidemic model for information diffusion in MANETs*. ACM Press, New York, NY, USA. Proceedings of the 5th ACM international workshop on Modeling analysis and simulation of wireless and mobile systems 2002, Atlanta, Georgia, USA *September 28 - 28, 2002*

Kurhinen, J. and Vuori, J. *MP2P Network as an Information Diffusion Channel.*

Kurhinen, J. and Vuori, J. *Information Diffusion in a Single-Hop Mobile Peer-to-Peer Network*. Proceedings of the 10th IEEE Symposium on Computers and Communications, in press.

Mudhakar, S., Bugra, G., and Ling. *Large Scaling Unstructured Peer-to-Peer Networks.*

Papadopouli,M. and Schulzrinne, H. 2000. *Seven Degrees of Separation in Mobile Ad Hoc Networks*. IEEE GLOBECOM. Nov. 27th-Dec. 1st. San Francisco.

Papadopouli, M. and Schulzrinne, H. 2001. *Effects of power conservation, wireless coverage and cooperation on data dissemination among mobile devices.* ACM SIGMOBILE Symposium on Mobile Ad Hoc Networking & Computing (MobiHoc) October 4-5. Long Beach, California.

Wei, Y. et al. 2005. *Peer-to-Peer System-based Active Worm Attacks: Modeling and Analysis*. IEEE.