

Milestone 8.1 Chapter 4

Citation for published version (APA):

Brouns, F., Kester, L., Sloep, P., Van Rosmalen, P., Berlanga, A., & Koné, M. (2007). *Milestone 8.1 Chapter 4*.

Document status and date:

Published: 29/01/2007

Document Version:

Peer reviewed version

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

<https://www.ou.nl/taverne-agreement>

Take down policy

If you believe that this document breaches copyright please contact us at:

pure-support@ou.nl

providing details and we will investigate your claim.

Downloaded from <https://research.ou.nl/> on date: 24 Jan. 2025

Open Universiteit
www.ou.nl



State of the art review for network management	1
Objectives	1
Models and methods	1
Domain model.....	2
Social space.....	4
Community characteristics	6
Community population	6
Community guidelines	7
Existing practices	9
Existing tools	12
Methods and policies for self-organisation in the network	14
References	17
Appendix 1. Activity diagram	19
Appendix 2. Use case diagram	27
Appendix 3. Class diagram	33

State of the art review for network management

Objectives

Successful communities are characterized by boundaries which protect their collective good, populated with a heterogeneous group of members to assure their liveliness and equipped with guidelines that foster ongoing interactions among its members. Although norms and rules may arise out of member interaction, policies and purpose of the community need to be communicated in order to set initial boundaries within which to act (Weber, 2004). These policies can then be renegotiated by community members as the network evolves. Task 8.4 will describe models that help explain and understand the functioning of networked communities as well as tools that help manage them, all the while preserving a maximum of user autonomy and control. It will also present policies and identify services that foster successful, self-organizing communities, including their ontological requirements. Such policies and tools are important because they allow for the emergence of network communities that are increasingly self-governed, self-organized and decentralized.

Models and methods

As far as the functioning of communities is concerned, two kinds of models may be discerned. There are models that aim at the dynamic behaviour of communities and mainly have an explanatory function. To the extent that they provide successful explanations, they may also be used as a basis for community design and management. And then there are also conceptual models or domain models that seek to make an inventory of pertinent terminology and describe the way terms are related to each other. Such models do not explain nor provide development or design guidelines; their function is to provide a useful vocabulary. In this section, we will first describe a domain model and then describe various dynamic models, that each focus on a particular aspect of communities.

Domain model

Community

A learning network is, among other things, a *community* of people (*members*) who share the intention to learn something about a particular *domain of knowledge*. Actually, calling a learning network a community presumes already too much, as its connotation is one of people who somehow interact and have a shared history. We do not assume this to be the case up front, although it may, as a matter of contingent fact, happen to be true for some of the members. Eventually, it will become true. Either way, we assume that *strengthening the social ties within a (learning) community will positively affect learning*. So, through active participation in the community the learning goals people have set for themselves will be attained more effectively, more efficiently, more attractively; or, put differently, reshaping a learning network as a community enhances the quality of the members' learning experience.

There are various ways to back up this assumption, such as social constructivist learning theories and, more generally, the notion that we as social animals – the term goes back as far as Aristotle – simply perform better in a social context. Justifying this assumption is not our present concern.

Activities and roles

The strengthening of social ties does not come about automatically. Mechanisms that allow or even stimulate the members to interact will have to be implemented. Typically, members engage in some sort of joint *activity*, i.e. they individually carry out *tasks* that fit into some overall activity. It is through their joint participation in an activity that mutual ties are strengthened. With respect to a particular activity, community members may be classified as *participants* and *non-participants*. The latter stay out off the activity entirely, may not even be aware of its taking place, the former carry out tasks. With respect to some activity, participants adopt *roles* that are specific to that activity (although similarly named roles are likely to occur in other activities).

For example, the activity of peer-tutoring in ad hoc, transient (sub)communities (Kester et al., submitted; Sloep et al., submitted) may take place in a learning network's community. Those outside the ad hoc community are non-participants, those inside are participants. In the example, there are two kinds of participants, a tutee, whose task it is to ask a question, and several peer-tutors, whose task it are to answer the tutee's question. Asking a question, answering a question, reacting to someone's answer, rephrasing the original question are all tasks in the overall activity of peer-tutoring.

Playing an initiation game in the course of joining a particular learning network is another example of an activity. Here the initiation game is the overall activity and the notion of a role is to be taken quite literally. Tasks are moves made in the game. A concrete example could be two opposing teams trying to find the way out of a virtual maze. The sole purpose is letting participants get to know each other. All participants may have the same role, although they are on different teams, or there may be role differentiation. Sharing bookmarks through a public site, such as Del.icio.us, is another, relatively simple example of an activity. Here every member might be a participant. There are two roles, active providers and lurkers. The entire activity consists of two tasks only, sharing a bookmark and finding someone else's bookmarks. Lurkers do only the latter, providers do both.

With respect to learning, joint activities may be grouped into two broad categories, depending on the goals of the participants and a few other characteristics (Strijbos, 2004). On the one hand there is *collaborative* work, such as found in project teams. The people on such a team have different responsibilities (division of labour); each participant contributes to a single common goal such as producing a joint report or software product. As a consequence their individual tasks are somewhat synchronized in time, a bit like in online project organization. In a learning situation, the goal to be pursued is set by a teacher in the form of an assignment; often also the allotment of the work is also done by the teacher, as is the setting of an overall time-frame, often also of detailed milestones. Typically, the common end result will be assessed and no distinction will be made between individual contributions. Often, individuals will be assessed too, but that then pertains to their behaviour as a group member.

On the other hand, there is *cooperative* work. Participants all pursue their own private goals, there is no division of labour, and all have their own schedules they adhere to; so there's no synchronization of individual tasks. Teachers are not involved in this, with the exception perhaps of suggesting that teaming up with someone else may be useful. Typically, only individual products are assessed, irrespective of whether they were the result of an individual or joint effort.

Participant characteristics

Participants in a particular activity need to be describable in terms that are relevant to their role in a given activity. Participants may be described in many different ways, but what matters here are the characteristics that are relevant to the activity they participate in; so this is an abstraction process. Characterizing them is a way of identifying them. Non-participants do not need to be characterized, at least with respect to the activity in question.

For example, tutees in the peer-tutoring activity are characterized as group by their asking questions. An individual tutee is characterized by the specific question he or she asks. Perhaps there are other relevant characteristics (number of questions asked, courses taken, current course), but these need not been taken into account immediately. Peer-tutors are characterized by their role as question answerers. Each tutor is more or less suitable as an answerer. Their suitability really is an aggregate of content competence, tutor competence, eligibility, and availability. Other characteristics could be taken into account, such as number of questions answered, recent history of questions answered. (Kester et al., submitted; Sloep et al., submitted). With respect to the initiation game – getting out off the maze – all participants adopt the same role within one team. The team is characterized by the progress it has made; this characterization is identical for all members on the same team. With respect to the bookmarking activity, a participant is either a provider or a consumer. A particular participant could be characterized by his or her historical record, by the total number of bookmarks contributed, by the diversity of the bookmarks, etc.

Proximate and ultimate goals

Every activity in a community has a particular purpose, i.e. the activity's participants engage in the activity because through their participation they pursue a particular goal. One should distinguish between *proximate* and *ultimate* goals. Being part of a learning network, we assume all participants to have a similar ultimate goal, which is to become more proficient with respect to a particular aspect of the domain of

knowledge that the learning network covers. The ultimate goals only differ with respect to their exact elaboration: although some learning network users may have identical ultimate goals, most will not. This variation really only means that in a learning network one may learn a variety of different things. Proximate goals are mere means to an end, to the ultimate goal. Proximate goals are therefore like instruments, they help to achieve something else. One's proximate goal is determined by the kind of activity one engages in, they are the goal a participant in an activity pursues. They are also role-specific. Ultimate goals are connected with activities in that activities channel the efforts members make to reach their ultimate goals. Some activities are better channels than others in that they require less effort for a similar effect. Proximate goals are the role-bound goals set by a particular activity. It is through the activity that the achievement of a proximate goal contributes to the achievement of an ultimate goal.

Pursuing a proximate goal in some activity may be a more or less effective means of moving towards the achievement of one's ultimate goals. How effective it is depends on the activity, on its kind but also on the way it has been designed. There are many degrees of freedom in the design of any activity, at least some will affect the activity's effectiveness.

Tutees in the peer-tutoring activity have as their proximate goal to have their question answered satisfactorily and as quickly as possible. Peer-tutors do not have that same proximate goal, theirs is to have their own questions answered when, at another moment of time they have one; or, getting a better grasp of the subject matter by explaining it to someone else. The peer-tutoring activity should be designed in such a way that questions will be answered quickly and efficiently. This brings us to questions about dynamic community behaviour and its theories and models.

Social space

A sound social space is characterized by affective work relationships, strong group cohesiveness, trust (i.e., perceived reliability of the word of other group members and genuine interest in the welfare of group members), respect, belonging (i.e., recognition of membership) and satisfaction (Kreijns, 2004; Nichani, 2001; Rovai, 2002). Social interaction enhances the emergence of social space. Interaction directed towards the completion of assigned tasks, however, could negatively influence aspects of this social space. When a task entails peer assessment, for example, fear of criticism or reluctance to criticize could interfere with feelings of trust (Rovai, 2002). An individual's expectations of the community could also negatively influence social interaction and hence the emergence of social space. According to Brown (2001), individuals who felt that people needed to join voluntarily or felt that face-to-face association was necessary, only developed a sense of belonging and trust if they joined a face-to-face community of their own volition. So social interaction and, as a consequence, the emergence of social space is facilitated only if socio-emotional-driven interaction is stimulated and not merely task-driven interaction; the same facilitation is observed when people's expectations about a community are met.

More generally still, three social prerequisites should be met in order for social interaction, in particular cooperation, to occur: (1) any two individuals must be likely to meet again in the future (*continuity*), (2) all individuals must be able to identify each other (*recognisability*) and (3) all individuals must be able to know how any

other person has behaved in the past (*history*). If individuals only meet once, they are very much tempted to behave selfishly, which negatively influences the cooperation process. In addition, if individuals are not identifiable and no history of a person's behaviour is available, group members are more likely to act selfishly because they cannot be held accountable for their actions (Kollock, 1998).

Relevant theories

The *social exchange theory* of Thibaut and Kelly (1959) applies Skinner's behaviourism to groups. Individuals strive to maximize their rewards and minimize their costs. Within groups, individuals no longer control their outcomes. Interdependences are created: actions (tasks) of each group member potentially influence outcomes of actions (tasks) of every other group member. Members negotiate throughout their interaction to secure greater personal rewards while minimizing costs.

Systems theory (Millen, 1978; McClure, 1998; Tubbs, 2001) regards groups as systems of interacting individuals. Groups can set goals, and work towards these goals through united action. The task of the group is to analyze inputs, provide feedback to members, and generate decisions regarding group actions. The analysis is focused on the information input that is fed into the group, the processes during group work and the products that are generated as output. Inputs include any factors that are present when the group work begins, such as characteristics of individual members (skill, experience, training, motivation) and group-level factors (group structure and cohesiveness). Processes include communication, planning, conflict and leadership. The outputs include aspects of the group's performance (products, decisions, and errors) and changes in the factors that serve as inputs. Larger groups may be built on a number of smaller groups. This organization is initiated by the system itself, and may undergo both gradual and rapid change.

Expectation-states theory (Berger, Wagner & Zelditch, 1992; Wagner & Berger, 1993) focuses on the cognitive processes that occur within each individual in the group. Newcomers form an impression of the group, and search for information about the other group members. Group members search their memories for stored information about the group and tasks it must face; they take note of the actions of others and try to understand what caused the other member to act in a particular way. Group members allocate status within the group by two types of cues: specific status characteristics (i.e. qualities attested to each individual's level of ability to perform the specific task at hand) and diffuse status characteristics (i.e. general qualities that group members think are relevant to ability and evaluation). Members with the most status-earning characteristics will rise to the top.

Level-of-aspiration theory (Lewin, Dembo, Festinger, & Sears, 1944; Zander, 1971) is a compromise between ideal goals that people set and more realistic expectations that they develop over time. Applied to groups, group members compare their performance to the group standards and eventually revise their strategies. A group's level of aspiration often slightly exceeds those of individual members. Also, members raise it more after success than they lower it after failure. Difficult goals challenge members to work harder; groups that fail consistently have low group morale and high turnover in membership.

Complexity theory (Waldrop, 1992; Kauffman, 1995) states that critically interacting components self-organize to form potentially evolving structures exhibiting a hierarchy of emergent system properties. This theory takes the view that systems

are best regarded as wholes, and studied as such, rejecting the traditional emphasis on simplification and reduction as inadequate techniques on which to base this sort of scientific work. Such techniques, whilst valuable in investigation and data collection, fail in their application at system level due to the inherent nonlinearity of strongly interconnected systems - the causes and effects are not separate and the whole is not the sum of the parts. The approaches used in complexity theory are based on a number of new mathematical techniques, originating from fields as diverse as physics, biology, artificial intelligence, politics and telecommunications, and this interdisciplinary viewpoint is the crucial aspect, reflecting the general applicability of the theory to systems in all areas.

Self-organization theory (Varela, Thompson, & Rosch, 1991; Maturana & Varela, 1992) contends that the behaviour of the system as a whole, and often of the individual parts, is a complex aggregation of the interactions of all the parts. No part controls the whole, or can even control another part outside the influence of the rest of the system. Such systems are said to be 'self-organizing' and the behaviour of aggregates of components is said to be 'emergent'. In these systems, which certainly include living organisms, ecosystems, and social or ecosocial systems, there are no isolated controlling agencies. There is no all-powerful father, boss, or king. There are no control hierarchies among components: no generals, captains, or soldiers. Self-organizing systems are inherently 'democratic'.

Community characteristics

Communities are characterized by (1) boundaries, (2) rules, (3) monitoring possibilities and (4) sanctioning mechanisms (Kollock & Smith, 1996; Koper, et al., 2004). Successful communities have clearly defined boundaries. These boundaries protect the collective good of the community to outsiders and encourage ongoing interaction because the group members are likely to meet again. In addition, communities have a set of rules that govern the use of common resources and that point out that is responsible for producing and maintaining the collective goods. Community members should be responsible for setting and modifying these rules themselves. Individual accountability facilitates cooperation. By monitoring each other's actions in a community, community members see whether their fellow members comply with the rules; if they do, this will make them more willing to comply themselves. A transparent community with clear boundaries and rules allows group members to sanction the behaviour of other group members. This happens mostly by informal social control mechanisms but sometimes more firm measures are necessary. These measures could be as severe as banishment from the group. So, monitoring and sanctioning, if used wisely, are important facilitators of cooperative relations (Kollock & Smidt, 1996).

Community population

The thriving of a community also depends on the characteristics of the people in it. First of all, people differ with regard to their experiences with communities. Often students are divided in *veterans* and *newbies*. Brown (2001) found that veterans showed good community behaviour. They were supporting and encouraging peers, sharing knowledge and experiences, reflecting on past learning, and sustaining friendships and/or acquaintances begun earlier. Newbies, however, depended much less on other group members and were wont to rapidly call for tutor help. They preferred a tight class structure with frequent interaction and helpful assessment from the tutor. It seems therefore wise to populate a community with both veterans and newbies. Because of their experience, veterans model good community

behaviour to the newbies. Newbies can turn to veterans for support and encouragement instead of to the tutor. Although this helps to create an online community, veterans need an incentive to continue to interact with newbies. Veterans are willing to do their 'duty' in the beginning but after a while tend to restrict their communication to veterans only, which hinders community building (Brown, 2001).

Second, most people are trend-followers, but it is the trendsetters that make the difference. Nichani (2001) describes three types of trendsetters, that each could have a big influence on the thriving of a community: connectors, mavens and salesmen. Connectors form the 'social glue' of a community; they are very sociable and attentive and have a talent for making friends. Mavens are the information experts that have a talent for collecting information and who are willing to tell others about it. Salesmen are persuaders, they have a tendency to reach out to the unconvinced and persuade them, in this case to join the community. The absence of these trendsetters in a community, which then consists of trend-followers only, will negatively influence elementary features such as belonging, trust and social interaction.

Finally and related to the issue of trendsetting, participants of online newsgroups differ in their inclination to either lurk or post in a community. A lurker, by definition, belongs to a community but never posts in it. The percentage of lurkers in communities is very variable (i.e., ranging from 0% to 99%; Preece, Nonneke & Andrews, 2004). For example, lurkers appear to make up 45.5% of health support communities while the lurker population in software support communities could be as high as 82% (for an overview, see Preece et al., 2004). Reasons for not posting range from 'didn't need to post', 'needed to find out about the group', 'couldn't make the software work', 'didn't like the group' to 'had nothing to offer' (Preece et al., 2004). Posters and lurkers are attracted to and join a community for the same reasons. However, posters feel their needs are better met, perceive more benefit and feel a greater sense of membership than lurkers. Partly because posters do not regard lurkers as inferior members, lurking is not necessarily a problem in active communities (see also Weber, 2004). Without a critical mass of posters, however, a community will never thrive (Preece, et al., 2004).

Community guidelines

Reward and incentive mechanisms need to be in place to encourage users to share, use and contribute knowledge. Additional policies, such as member participation, terms of use, quality standards and procedures, including their ontological requirements need to be considered.

Guidelines for fostering social space

The *recognisability* of users may be assured by forbidding the use of aliases such as screen names; this seems a reasonable demand to make in the context of a network devoted to learning. If one does not want to be this strict, users that go by a pseudonym should adopt one and only one persistent pseudonym, i.e. a single pseudonym they keep throughout their membership of the network and use in all interactions.

A *historical record* of user activities is maintained by logging all user-activities. The ones most significant for knowledge sharing - activities that reflect content competency and sharing competency - become part of the user's *profile*. Content

competency reflects the user's mastery of the content within the network. Hereto, the profile contains the products that resulted from the learning activities of a user (i.e., papers, reports, assessments). Sharing competency refers to the ability of a user to satisfactorily support peers during a process of knowledge sharing. This information could be acquired by letting users rate each other's performance. To enhance individual accountability (Slavin, 1995), both content and sharing competency of a user is made visible to the members of a particular ad hoc, transient community (there seems to be no reason to stigmatize a person at this stage by making it always available within the entire network). For the same reason, rating should not be anonymous, at most singularly and persistently pseudonymous.

Continuity of contact is guaranteed by demanding that all community members are accessible. But continuity of contact only makes sense if there is extra value that having access to others. Therefore, learning network users should be allowed and stimulated to maintain a rich online identity. This should preferably be done through a digital dossier or *portfolio*. It should contain information on a user's background, but it should also be updated regularly and automatically, almost as a track record of someone's presence in the learning network. The portfolio is part of a user's profile.

Guidelines for community characteristics

The (proximate and ultimate) goals learners have form the incentive for the process of knowledge sharing. Indirectly this goal strongly influences the amount of social interaction during knowledge sharing within the community. Clearly, a goal that can be reached by only one correct solution will elicit less social interaction than a goal that can be reached through various solutions.

Different interaction-structures can be implemented to mediate the effects of a goal on the social interaction. For example, if the goal can be reached by a limited number of solutions then a peer-tutoring structure could stimulate social interaction. King, Staffieri, and Adelgais (1998) advocate a three-step structure that consists of communication guidelines (i.e., listening, encouraging and giving feedback), an explanation procedure (i.e., the TEL WHY-procedure; telling in one's own words, explaining why and how, and linking of content), and questioning guidelines (e.g., asking comprehension questions or thinking questions). Other examples of structuring interaction within groups are "...*Group Investigation* (Sharan & Sharan, 1992), *Student Teams Achievement Division* (Slavin, 1995), 'Jigsaw' (Aronson & Thibodeau, 1992; Bielaczycs, 2001), *Structural Approach* by Kagan (1994)) (each structure is a scenario to teach specific skills and, although not likewise articulated, it is implicitly assumed that no situation is identical), *Progressive Inquiry* (Rahikainen, Lallimo, & Hakkarainen, 2001), *the use of scripts* (O'Donnell, 1999; Weinberger, Fischer, & Mandl, 2001), *scenarios that prescribe collaboration activity* (Wessner, Pfister, & Miao, 1999), *feedback rules or requirements of a minimum degree of contributions to a discussion* (Harasim et al., 1995)." (*vide* Strijbos, 2004; p.33). From our perspective, 'high-structuring' methods such as peer-tutoring or Jigsaw are most suitable for goals that can be reached by a limited number of solutions because they guarantee a minimum amount of social interaction. 'Low-structuring' methods such as Progressive Inquiry, however, are most suitable for goals that can be reached by various solutions because these methods support rather than elicit social interaction (e.g. negotiation, argumentation) which is believed to be necessary under these circumstances.

Guidelines for the community population

Specialization of roles has been associated with effective self-organizing systems. Roles help position and clarify the relationship between members. Roles also delineate the responsibility of each member for the production or maintenance of collective goods, and to stimulate the transition of lurkers and passive members to more active poster and co-developers of knowledge.

In other words a community should consist of a mix of members with complementary expertise, all related to the goal of the community. So if, for example, 'answering a content-related question' is the goal of the community, it should consist of members with different levels of expertise related to the content-question since heterogeneity in levels of expertise can have differential effects on learning. Although King et al (1998) found that peer-tutors do not necessarily have to be more competent or more knowledgeable than their tutee counterparts; a study of Hinds, Patterson, and Pfeffer (2001) indicates that tutors equal in competence convey qualitatively different knowledge than more distant tutors. The near tutors - those who are similar to their tutees in expertise level - use more concrete statements during their interactions with the tutee. In contrast, the distant tutors - those with a higher level of expertise - convey more abstract and advanced concepts. Heterogeneity in level of expertise between members thus leads to a wide spectrum of knowledge shared in the community.

Existing practices

Virtual learning communities

Scholieren.com www.scholieren.com

Community population: [Www.scholieren.com](http://www.scholieren.com) is a website maintained by Dutch pupils and meant for Dutch pupils. The community has over 63.000 subscribers, who can post contributions to the site. Most contributors are between 15 and 18 years old; a few are adults. Guests are allowed to view the contents of the site. Scholieren.com dates back to 1997. According to the editors, it is one of the most popular sites for pupils.

Community characteristics: At the website, all kinds of materials are exchanged that can be useful for pupils, for example extracts. Besides, the website contains various discussion forums, in which pupils can post their questions and problems. Their topics include anything that a pupil can come across, and include much more than education. The forums on homework are classified by subject area. Usually, contributors do not have to wait very long for reactions to their message. It is not unusual to receive five reactions within the first few hours after posting.

Social space: The discussions are moderated by one of the pupils from the website editorial staff. Moderators must be at least 15 years old. There is a distinction between a 'moderator' and a 'moderator+'. A moderator is responsible for one subforum, and can move, remove, adjust and close topics. A moderator+ keeps an eye on one whole forum. A moderator+ can interfere with a subforum when the moderator hasn't taken appropriate action. A moderator+ can also ban people from the community. Materials for exchanges are placed in a database. The main categories are book reports and papers. Pupils can search the database and they can post requests for book reports of specific books.

The success of Scholieren.com seems to be determined by several factors. One is the enormous number of subscribers. At any moment there will be a few hundred of the 63.000 subscribers online. The large number of subscribers is the result not only of

the huge size of the target group, but also of the urgency of the problems that are discussed. Pupils have a lot of homework, and often they get stuck and need help. Further, the community consists mainly of peers, which makes it easier for individuals to ask questions. This is strengthened by the possibility to use a nickname (pseudonym), which almost everyone does. Finally, the community is moderated, so that disorderly and undesirable behaviour are reduced.

Fifth dimension www.5d.org

Community population: The Fifth Dimension (5D) is described by its founder Michael Cole as 'a specially designed cultural medium for promoting the all-around intellectual and social development of 6- to 12-year-old children' (Cole 1999). Children and university members can enter the community throughout the year. As a result, at any one moment the community is a mix of newcomers and old timers, in which some children have more experience with the norms and computer aspects of 5D than some Wizard's assistants. This leads to a change in the power relations between children and adults.

Community characteristics: 5D has been developed in the United States in the eighties and it is designed to address certain long-standing problems in American education, in particular the distressingly low academic achievement of many American children, the widely perceived need for them to gain a qualitatively richer experience with new information technologies, and the failure of apparently successful educational innovations to survive beyond the period of innovation and external funding. 5D is a virtual learning environment. 5D runs parallel to the university year and the school year. Working occurs in periods of eight weeks, in which members are active in 5D between one and four days a week. Its heart is formed by activities, which are presented to the children in the form of a cardboard maze, divided into 20 rooms, each of which gives access to two activities. Three quarter of the activities is computer activities, including computer games and educational software. According to the rules of 5D (enshrined in a Constitution, which each child receives upon entering the activity system), children progress through the maze; the whole process takes from several months up to several years. In order to carry out a task, children must first consult a task card that defines progress on that task. Each task on one task card is described at three levels: beginner level, good level and expert level. Higher levels of achievement increase children's freedom of choice in moving within the maze. They also give the child the possibility to alter their avatar, which is very plain in the beginning. Children who complete all the rooms in the maze attain expert status and access to new activities.

Social space: A very important role is played by the Wizard, the ultimate authority. The Wizard is the creator of the Constitution, helps children who experience difficulties, reprimands them in case of antisocial behaviour or working below their abilities, and settles disputes between members of the community. Children report their achievements to the Wizard, and are thus forced to explicate what they are doing. The Wizard is supported by the Wizard's assistants. 5D is run from universities, and the Wizard and the Wizard's assistants are university teachers and undergraduate students. Working in 5D is done from schools, youth clubs, day care, libraries and churches. Success and sustainability of 5D is defined by the interaction between 5D and the environment from which children participate. A quiet environment such as a library shows better learning results, but less chance of sustainability, as 5D is seen as a disturbing element. A noisy environment such as a youth club shows the opposite. [adapted from Cole, 1995 and 1999].

Success factors in 5D include the following. There are several non-personal mechanisms for settling disputes, such as the constitution and the Wizard. Achieving

a higher level is rewarded in several ways, e.g. freedom of movement is increased, more activities come within reach, new duties and responsibilities are acquired, the avatar can be altered. There is a constant flow of newcomers at all moments, resulting in more equal power relations. Success and sustainability is also defined by the interaction between 5D and the environment from which children participate. Finally, a further interesting characteristic is the possibility of carrying out the same task at several levels of proficiency.

Notschool www.notschool.net

Community population: Users are selected by panels based on set selection criteria.

Community characteristics: Notschool started as an online research project aimed at young people of school age who have been out of the traditional education systems for personal or logistical reasons. Notschool looks at ways to get these people back into learning. It is aimed at those people for whom traditional alternatives such as home tutoring have not worked. It started out with a virtual community of 100 people, but is now being used in Education Authorities in the UK and overseas with over 1700 people. The participants were given the opportunity to develop their self-esteem and be reintroduced to learning, through the support of mentors, buddies, experts and the use of new technology. Four key factors distinguished the adults in Notschool.net from those in mainstream schools:

- Teams at local level were not teachers.
- They all had 24/7 access to up to date technology.
- They all had unlimited 24 hour access to broad bandwidth.
- They all had good levels of computer literacy.

Social space: Notschool consists of a highly structured community, with a central support team and several local teams. Each local team contains several mentors and researchers (i.e. the pupils). No titles were used so no distinction could be made between adults or those with authority. Everybody could see who was online. Every comment was attributable. Everyone could see who had read what at what time. All words and phrases relating to school were avoided or changed into more acceptable words.

ESP network www.esp.uva.nl

Community population: Teleprojects are collaborative distance learning projects designed by teachers from various countries around a part of curriculum that is thought to be mutually relevant. Leading idea of a teleproject is the combination of local research of pupils around a certain topic and exchange of, and conversation on research results with partner-schools, using a foreign language and electronic mail. Both domain specific teachers, foreign language teachers, and informatics teachers can help improve the activities of their pupils and make new educationally relevant activities possible when participating in teleprojects.

Community characteristics: The European Schools Project started in 1988. A central concept was introduced and refined to structure computer-supported collaborative learning between primary and secondary schools: the teleproject. The concept encompasses Internet-based collaborations between teachers and pupils around 'conversation' topics that are thought to be relevant for learning and teaching of all participants in the project. The topics demand active and authentic learning of the pupils, while for the collaboration a mutual foreign language, and electronic mail or Virtual Learning Environments are used.

Professional network communities

Cisco Netpro

Community population: Anonymous browsing of forums. Registration is required to add or reply to posts, and rate.

Community characteristics: There are forums for every possible topic related to networking professional. There is an expert section where experts present events on certain topics. TechTalks are live events featuring technical presentations and the opportunity for viewers to have their questions answered online. Previously broadcast TechTalks are available for viewing at your convenience.

Social space: Users can rate topics and indicate whether conversations contain suitable answers. Within each forum the top 5 rated users are displayed, indicating their points, average rating, and badge status. Badges can be earned by number of points. Over forums, the top experts can be listed, showing points, average rating, badge status and the number of posts with satisfactory answers.

Existing tools

There are several applications which provide some of the aspects or functionalities that are required for management of social networks or allow implementation and/or enforcement of policies. Some are listed below. Also mentioned are some techniques that are required for implementation of network management tools.

Relationships

The application should allow creating and managing expressions of personal relationships and build new ones. The FOAF (Friend of a Friend) standard can be used here.

History of learner's activity

General logging techniques and the use of e-portfolio systems as well as social network analysis can be used.

- Social Network Analysis

Social network analysis (SNA) determines and visualizes the patterns in interaction between people, groups, organisations, etc in social networks. People form the nodes in the network, while the links between the nodes shown relationships or flows. There are several commercial and some non-commercial applications available to perform these analyses. Examples are Analytic Technology, Inflow, Jung, and Visone.

Awareness and accountability

An interesting example is Babble/Loops. It is a socially translucent system, using awareness and accountability, which support computer-mediated communication, allowing threaded and persistent conversation. A social proxy is a graphical representation of users depicting their presence and activities in relation to the conversation (Erickson et al., 1999; Erickson et al., 2006).

Recognisability

- Applications like Orkut (<http://www.orkut.com>), LinkedIn® (<http://www.linkedin.com/>), Friendster® (<http://www.friendster.com/>), MySpace (<http://www.myspace.com>) and (<http://www.facebook.com>) Facebook are social network sites that allow people to create their profile and make that available to others. People can link to others and can comment on each other's profile.

- the ASA-system, under development at the Open Universiteit Nederland, that uses peer-tutoring in ad hoc, transient communities as a means to strengthen the social fabric of a learning community.

- LiveJournal (<http://www.livejournal.com>) offers functionality at several aspects. It is an open source content management system that lies behind successful online communities, such as LiveJournal.com. LiveJournal combines FOAF services with blogging services; based around journal it allows people to create communities. Users keep ownership and control over their space.

- ELGG, <http://www.elgg.org> is a learning application centred around user's profiles. When a user creates a profile it is automatically linked to others with the same interest, but also to resources. Weblogs are used for own reflections, but also communications with the community. Connections to other people and resources are used to build networks to enhance the learning experience. Users have control over how and what they present and can control who sees what.

- social bookmarking and tagging applications, such as Furl, Del.icio.us and non-commercial variants.

Methods and policies for self-organisation in the network

Successful communities are characterized by boundaries which protect their collective good, populated with a heterogeneous group of members to assure their liveliness and equipped with guidelines that foster ongoing interactions among its members. Although norms and rules may arise out of member interaction, policies and purpose of the community need to be communicated in order to set initial boundaries within which to act (Weber, 2004).

A learning network is, among other things, a community of people (members) who share the intention to learn something about a particular domain of knowledge. We assume that strengthening the social ties within a (learning) community will positively affect learning. So, through active participation in the community the learning goals people have set for themselves will be attained more effectively, more efficiently, more attractively; or, put differently, reshaping a learning network as a community enhances the quality of the members' learning experience. Mechanisms that allow or promote strengthening of social ties involve users engaging in joint activities in different roles. Role specific user characteristics and descriptors related to a particular activity are required. Users should be recognisable and identifiable.

Ad hoc transient communities are seen as the vehicle to organise this (Kester et al., submitted; Sloep et al., submitted). Ad hoc transient communities serve a specific goal, are limited in time (i.e. dissolve when the goal has been attained, and operate according to social exchange policies that enhance social embedding and knowledge exchange.

A generic use case diagram is presented in figure 1. In the remainder of the document a detailed requirement description is given for a first prototype to supply a specific form of ad hoc transient communities; that is ad hoc transient communities for peer tutoring.

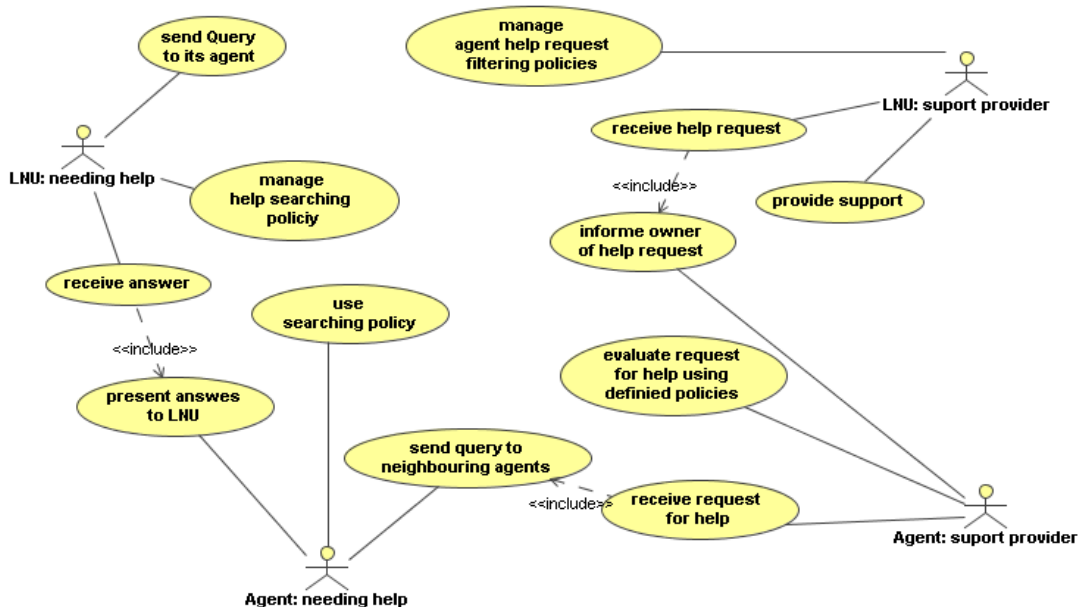


Figure 1: Generic use case

Ad hoc transient communities for peer tutoring

This section describes the requirements for ad hoc transient communities as well as the requirements for a system implementing ad hoc transient communities. Innovative educational technology and ICT are applied to create and populate ad hoc transient communities in which peer tutors instead of institutional tutors provide support to tutees. LSA technology is used to select suitable peers and possible (fragments of) answers in the learning network. The users are assisted by personal agents, and a central matchmaking agent provides the glue in the system. There are four criteria for selection of suitable tutors: content competence, tutor competence, eligibility and availability. Algorithms for these criteria are developed but need to be transformed into more flexible and dynamic rules. The policies for population of the ad hoc transient communities to conform to the boundary, heterogeneity and accountability conditions are present implicitly in the model. The system does not describe those external systems it relies on, which should be available in the learning network; such as portfolio data to determine content competence and tutor competence; dossier data for personal preference, availability, reputation, rating scores; logging data of network use, etc.

The system distinguishes three types of actors: the learning network user (LNU), who can take on the role of tutee (i.e. a LNU, usually a learner, who has a question) or the role of tutor (i.e. a LNU or peer learner who provides an answer to a question). Every LNU is assisted by a personal agent (not necessarily embodied in one agent) who represent the LNU and acts on behalf of the LNU in both roles of tutee and tutor. The third actor, the Matchmaker agent orchestrates the processes involved in the selection of the tutor(s) and suitable answer fragments, and the population of the ad hoc transient communities.

Abbreviations used:

LNU: learning network user

LN: learning network

AN: activity node

CQ: content question, i.e. a question relating to content

ID: the LNU's identifier that uniquely identifies the user in the learning network

The term dossier is used to refer to personal data, portfolio data and logging data that are required by the system to be able to determine content competence, tutor competence, eligibility and availability. These are not necessarily represented in one system.

The term community refers to an ad hoc transient community that is created within a learning network with the purpose to provide an answer to a content related question.

In the following paragraphs a short run-through of the system is provided in a narrative and is indicated how the main network policies are implemented in the system. A more detailed step by step approach and algorithms are provided in the description of the activity diagram. The corresponding UML use case, activity and class diagrams are provided in the appendices.

Narrative

Every LNU is assisted by a personal agent. The personal agent is the intermediate between the user and the learning network, sometimes acting on behalf of the user; at other times automating tasks for the user (e.g. maintaining availability records).

The learning network contains one or more Matchmaking agents who deal with users' request through their personal agents. A user in a learning network does not fully understand the topic he is studying. The answer can not be found in the learning materials and the user (tutee) decides to ask a question in the learning network community. The learning network either provides a mechanism for this, or the LNU agent ensures that the question can be posed and is relayed to the learning network. Via the LNU agent the validity of the question is determined, if necessary in interaction with the user. The matchmaker agent then takes on the question and will form an ad hoc transient community, populated with the user asking the question (tutee), a number of suitable peer tutors and (fragments) of documents found in the learning network that either contain (part of) the answer or can be used as input for the answer. This process involves several steps. First the activity node is determined from which the question seems to arise, or find the first activity nodes which are most closely related to the question. This is required to be able to select peer tutors who are competent on the question subject and find the most suitable text fragments. The matchmaker then, through the LSA engine, selects suitable text fragments. The matchmaker agent also selects suitable tutors, taking four criteria into consideration i.e. content competence, tutor competence, eligibility and availability. Content competence is related to the level of mastery of activity nodes (registered in the user's portfolio or dossier) and is measured relative to the mastery of the activity nodes by the tutee and the origin activity node. For tutor competence measures like quality of contribution and rating of tutees are relevant (also related to data in dossier and portfolio). The eligibility is a measure of preference. Availability is related to time constraints, but also takes into account work load and past performance (based on data in portfolio and dossier). The matchmaker agent then invites possible tutors to participate in the ad hoc community. The question forms part of the invitation. This could entail several invitation rounds or reselection of tutors when an insufficient number of tutors accept the invitation. Once sufficient tutors have accepted the invitation, the ad hoc community is created and populated with the question and the possible answer text fragments. Tutee and tutor are granted access to the community and engage in a discussion to arise at the answer. The tutee can rate contributions and tutors. The tutee decides when the question is sufficiently answered, or failing that, that the community can be closed. The answer and question are stored, as are the ratings. Answer and question are made available to the learning network.

Policies

The boundary policy is met because the goal of the ad hoc transient community is clear and tutee is responsible for closing the community; this is also communicated to the participants via the invitation and in the ad hoc community. The population of the ad hoc community with tutee and several tutors, selected on several criteria ensures heterogeneity of the community. Accountability is ensured because users need to log in, maintain a profile and portfolio, and the system logs required data. For a more extensive description of policies see also Berlanga et al. (submitted).

Evaluation plan

A prototypical implementation of the system has been created that will be used to evaluate the effect of the policies on the effectiveness of the ad hoc transient communities. The first pilot will focus on heterogeneity aspects as well as parameterisation of the LSA engine and general variables of the system. A new prototypical system will be build based on previous experience and outcomes of the first pilot.

References

- Berlanga, A.J., Kester, L., Sloep, P.B., Brouns, F., Van Rosmalen, P., Koné, M. & Koper, R. (submitted). Foster Knowledge Sharing in Ad Hoc Transient Communities
- Berger, J., Wagner, D. G. & Zelditch, M. (1992). A Working Strategy for Constructing Theories: State Organizing Processes (pp. 107-123). In G. Ritzer (Ed.) *Metatheorizing*, Newbury Park, CA: Sage.
- Brown, R. E. (2001). The process of community-building in distance learning classes. *Journal of Asynchronous Learning Networks*. 5(2), 18-35.
- Cole, M. (1999) Cultural Psychology: Some general principles and a concrete example. In Y. Engeström, R. Miettinen & R.L Punamäki (Eds.) *Perspectives on Activity Theory*. Cambridge: Cambridge University Press.
- Erickson, T., Kellogg, W. A., Laff, M., Sussman, J. Wolf, T. V., Halverson, C. A., Edwards, D. A. A Persistent Chat Space for Work Groups: The Design, Evaluation and Deployment of Loops. *The Proceedings of DIS 2006*. New York: ACM Press, 2006.
- Erickson, T. Smith, D. N., Kellogg, W. A., Laff, M. R., Richards, J. T., and Bradner, E. "Socially Translucent Systems: Social Proxies, Persistent Conversation, and the Design of 'Babble.'" In *Human Factors in Computing Systems: The Proceedings of CHI '99*. ACM Press, 1999.
- Harasim, L., Hiltz, R., Teles, L., & Turoff, M. (1995). *Learning Networks: A field guide to teaching and learning online*. Cambridge: MIT Press.
- Hinds, P. J., Patterson, M., & Pfeffer, J. (2001). Bothered by abstraction: The effect of expertise on knowledge transfer and subsequent novice performance. *Journal of Applied Psychology*, 86, 1232-1243.
- Hummel, H.G.K. Burgos, D., Tattersall, C., Brouns, F., Kurvers, H. & Koper, R. (2005). Encouraging contributions in learning networks using incentive mechanisms. *Journal of Computer Assisted Learning*, 21, 355-365.
- Kauffman, S. (1995). *At Home in the Universe: The Search for the Laws of Self-Organization and Complexity*. New York: Oxford University Press,
- Kester, L., Sloep, P., Brouns, F., Van Rosmalen, P., De Vries, F., De Croock, M., & Koper, R. (2006). Enhancing social interaction and spreading tutor responsibilities in bottom-up organized Learning Networks. *Proceedings of the IADIS International Conference on Web Based Communities 2006*, San Sebastian, Spain, 26-28 February 2006.
- King, A., Staffieri, A., & Adalgais, A. (1998). Mutual peer tutoring: Effects of structuring tutorial interaction to scaffold peer learning. *Journal of Educational Psychology*. 90, 134-152.
- Kollock, P. (1998). Design principles for online communities. *PC Update*. 15(5), 58-60.
- Kollock, P. & Smith, M. (1996). Managing the virtual commons: Cooperation and conflict in computer communities. In *Computer-mediated communication: Linguistic, social, and cross-cultural perspectives* (Ed. Susan Herings), 109-128. Amsterdam: John Benjamins.
- Koper, R., Pannekeet, K., Hendriks, M., & Hummel, H. (2004). Building communities for the exchange of learning objects: theoretical foundations and requirements. *ALT-J, Research in Learning Technology*, 12, 21-35.
- Kreijns, K. (2004). *Sociable CSCL environments. Social affordances, sociability, and social presence*. Unpublished master's thesis, Open University of the Netherlands, Heerlen, The Netherlands.

- Lewin, K., Dembo, T., Festinger, L., & Sears, P. (1944). Level of Aspiration. (pp. 333-378). In J. M. Hunt (Ed.). *Personality and the behavior disorders* Oxford: Ronald Press.
- Maturana, H. R. & Varela, F. J. (1992). *The tree of knowledge*. Shambala, Boston and London.
- McClure, B. A. (1998). *Putting a New Spin on Groups: The Science of Chaos*. Mahway, NJ: Lawrence Erlbaum Associates.
- Nichani, M. (2001). Communities of practice at the core. *Elearningpost* [On-line]. Available: <http://www.elearningpost.com/features/archives/002090.asp>
- Preece, J., Nonneke, B., & Andrews, D. (2004). The top five reasons for lurking: Improving community experiences for everyone. *Computers in Human Behavior*, 20, 201-223.
- Rovai, A. P. (2002). Building a Sense of Community at a Distance. *International Review of Research in Open and Distance Learning* [On-Line], 3(1). Available: <http://www.irrodl.org/content/v3.1/rovai.html>
- Slavin, R. E. (1995). When does cooperative learning increase student achievement? *Psychological Bulletin*, 94, 429-445.
- Sloep, P.B., Kester, L., Van Rosmalen, P., Brouns, F., Koné, M., & Koper, R. (submitted). Facilitating community building in learning networks through peer-tutoring in ad hoc transient communities.
- Strijbos, J. W. (2004). *The effect of roles on computer-supported collaborative learning*. Unpublished master's thesis, Open University of the Netherlands, Heerlen, the Netherlands.
- Thibaut, J. W. and Kelly, H. H. (1959). *The Social Psychology of Groups*. New York: John Wiley and Sons.
- Varela, F. J., Thompson, E., & Rosch, E. (1991). *The embodied mind: Cognitive science and human experience*. MIT Press, Cambridge, MA, USA.
- Wagner, D. & Berger, J. (1993). Status Characteristics Theory: The Growth of a Program (pp. 23-63). In J. Berger and M. Zelditch (Eds.), *Theoretical Research Programs: Studies in The Growth of Theory*. Stanford, CA: Stanford University Press.
- Weber, S. (2004). *The Success of Open Source*. Cambridge Mass.: Harvard University Press.
- Waldrop, M. (1992). *Complexity: The emerging science at the edge of order and chaos*. New York: Simon & Schuster.

Appendix 1. Activity diagram

The activity diagram describes the system of ad hoc transient communities into some detail. It provides a quite detailed flow through the process and suggests several algorithms for various steps.

For clarity sake, the handling of the request for support (tutee asking content question) and tutor selection are presented in two separate diagrams, but both processes will take place in parallel. Therefore some actions seem to be duplicated, but are not.

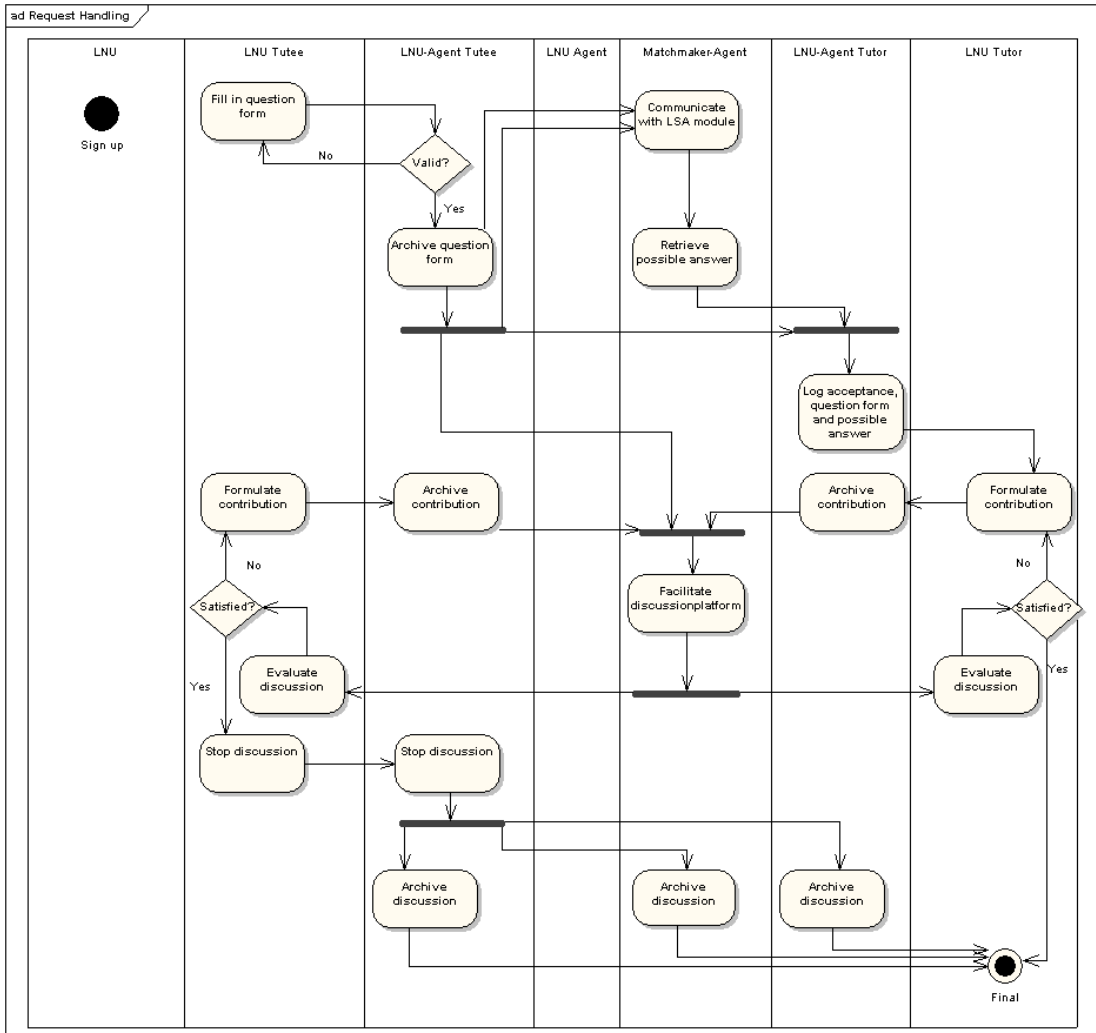


Figure 2: Request handling

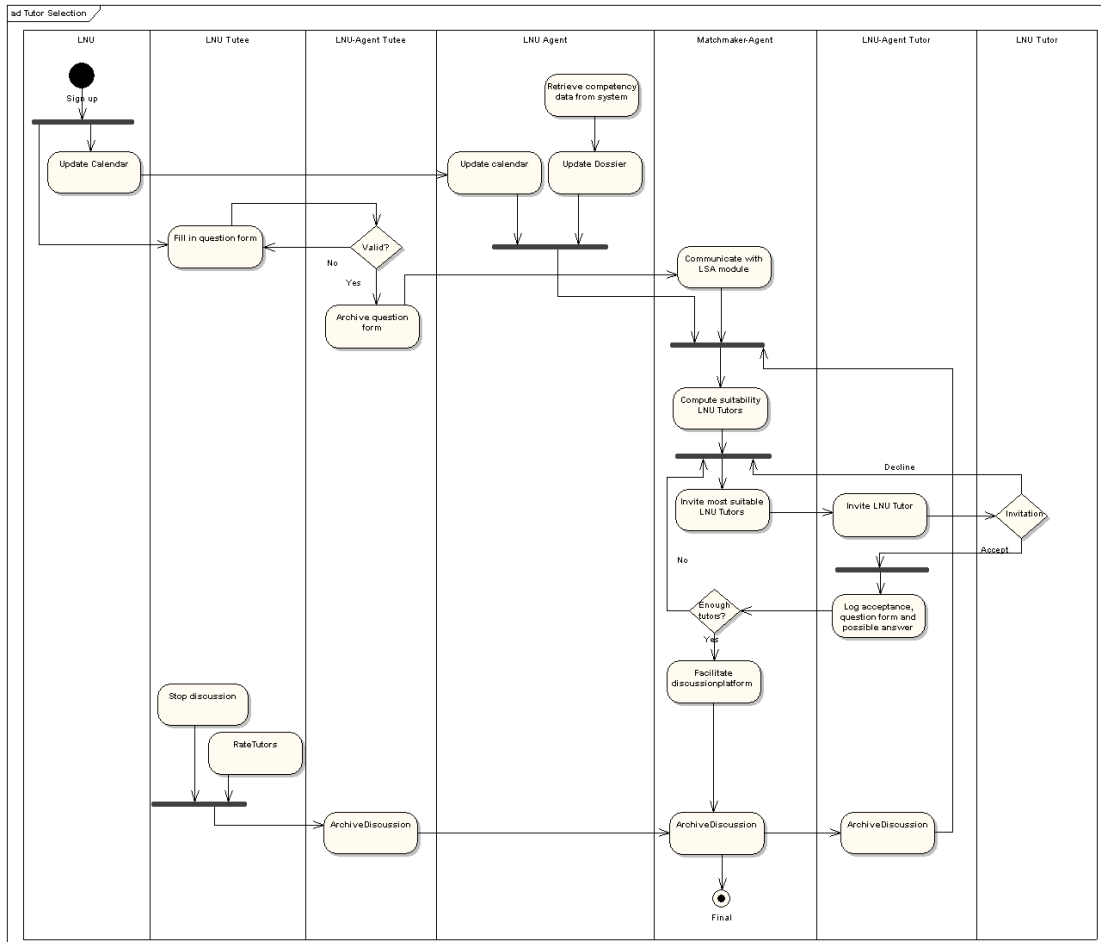


Figure 3: Tutor selection

The starting point of both diagrams is a user who has been enrolled in a learning network and has signed on.

1 Fill in question form

Some Learner asks a Content related Question (CQ).

This is where it all starts. Only questions that pertain to the content are permissible. Questions about procedures or administrative issues are not allowed.

- Procedural questions could of course be treated in a similar way. This is out of scope, though.
- A question asked will usually pertain to some Activity Node (AN) the Learner happens to study. This need not be the case, though. Learners may ask questions about several of the ANs they have studied thus far, even about ANs that are part of the LN, but weren't studied by the learners as the positioning system indicated they were part of his or her prior experience.

2 Valid?

Check whether the CQ is formally correct.

If not correct, it gets rejected immediately. Formally correct means whether the form used to submit the question has been filled in fully.

- This step may be skipped. It is mentioned here, because it might be useful.

- Since the CQ is being shown to some of the LNUs and to the LSA engine (see steps 4 and 6), it should at least be sufficiently detailed for both the tutor and LSA engine to work on. Perhaps a minimal number of meaningful words (i.e. words after stopping and stemming) could be used as a criterion. In this step, it could be checked whether the question is non-procedural. Perhaps a LSA algorithm could be used to establish the likelihood of it being procedural. Rather than reject the question outright, it could ask the learner whether the question really is content related or, after all, procedural. A standard collection of procedural documents could be used as a benchmark.
- In this step, the system could also check whether the question contains an indication of the AN the CQ belongs to and perhaps also the stage to which the learner has progressed within the AN (see 10).

3 Archive question form

Log CQ and AN

Log merely means adding the question to the requester's personal database for reuse by him or her later on. Having an overview of questions asked, linked to the forum thread in which they have been answered is a useful service to the user.

- Whether the log is going to be used in other ways depends on what additional functionality (in the form of use cases) the system should address. Only something as simple as the frequency distribution of questions asked (number of users asking 1, 2, 3, ... questions) already is an interesting community statistic; another one would be the number of questions per AN (or document within an AN). Obviously, these statistics would be computed by a system-wide agent, like the match maker.
- The registration could include the AN that the learner is currently studying: <CQ, AN>, perhaps even the stage he or she is in.

4 Communicate with LSA module

Carry out LSA on CQ

The purpose of this step is to find out the relevance of each AN for answering the CQ. Clearly, some ANs are more relevant than others. When trying to find a tutor to answer the CQ, the relevance of the ANs for the CQ is used to weigh each tutor's content and tutoring expertise (see steps 5 and 6): the more relevant an AN, the more a tutor's expertises should count. With the help of the LSA module, all documents in the entire Learning Network, arranged by AN, are compared with the question and the correlation coefficients are computed. All documents belonging to some AN are pooled. Then a listing is produced of <correlation, activity node ID> doublets. It is assumed that a high correlation points to a high relevance and vice versa. Relevance may be represented by the relevance vector (a column vector) $Ra_1 = [r_1 \dots r_j \dots r_a]$, where a is the number of ANs and r_j the relevance for the j -th AN

- The documents are pooled rather than, say, averaging the correlation coefficients of the various documents per AN. This is done to avoid effects like the following: an AN that contains one document with $r=0.9$ and 9 with $r=0.2$ scores lower (0.027) than an AN that contains one document only with $r=0.3$. For heterogeneous ANs, this is a serious problem as they get 'averaged out'. Alternatively, we could have chosen not to average the coefficients, but to use the maximum coefficient. It should be possible to figure out what operation on the document correlation coefficients yields the same result as the pooling of the documents. Not pooling them may be preferable in view of step 7, in which suggestions for answers to the CQ are extracted. This requires a resolution even at the level of paragraphs rather than the AN as a whole. But perhaps, this needs to be seen as two distinct steps.

5 Compute suitability peer tutors

Compute tutor suitability

All other LNUs, or rather their personal assistants, are now asked to evaluate their suitability to offer support to the learner, that is, to provide an answer to the CQ, however preliminary. This activity consists of 4 sub-activities that can be carried out in parallel. They are described below.

5a From an LNU's dossier, determine his or her content competency with respect to CQ.

Someone's competency level describes this person mastery of the learning objectives of some AN, hence the term 'content competency'. Competency levels vary between 0 and 1; 0 means not completed, 1 means completed. They can be represented by the vector $C1a = [c1 \dots ci \dots ca]$ where a is the number of ANs and ci the content competency on the i -th AN. These content competencies need to be adjusted so as to reflect their relevance to CQ. Content competency with respect to an AN irrelevant to CQ should be ignored (receive weight 0), content competency with respect to an AN highly relevant to CQ should be taken into account (say, receive weight 0.9). Therefore, each ci needs to be weighted by the relevant of AN_j for CQ. This is done by multiplying the competency (row) vector $C1a$ with the (column) vector $Ra1$. This gives:

$$C1a * Ra1 = (c1 * r1 + c2 * r2 + \dots + ca * ra)$$

$$(C1a * Ra1) / a = C$$

Division by a is done to make sure that $0 \leq C \leq 1$. C is called the consolidated content competency.

- Competencies are hard to measure, they can be measured through assessments, but that is a route we don't want to go for obvious reasons. Failing this option, our best measure of someone's competency on some AN_j is i) whether he or she has completed AN_j successfully, perhaps added with information on ii) how long ago that was (assuming that mastery fades with time).
- Someone who has just completed some AN_j is more competent than someone who has done this sometime ago; unless, of course, the latter person has pursued further studies that build on AN_j . That too should then be reflected in his or her dossier and could be taken into account.
- To add some more sophistication, those who are still studying a particular AN (and hence also are learners), may receive a content competency between 0 and 1.
- Content competency should take into account someone's previous experience, as reflected in his her portfolio (positioning!).
- To compute content competency, the following rules could be taken into account:
 - IF the tutor is not actively working on one or more of the relevant activity nodes THEN the tutor receives content competency 0.

5b From an LNU's dossier, determine his or her tutor competency with respect to CQ

Someone's content competency is related to but different than someone's tutoring competency. Someone who has good mastery of some subject doesn't necessarily make for a good tutor. And, vice versa, someone with average content competency may make an excellent tutor. Indeed, someone who has no content competency cannot be a tutor at all. So the system has to distinguish between content competency and tutor competency. In this step, the tutoring competency of some LNU is computed from his or her portfolio. As with content competencies, tutor competencies will vary over ANs. Via a procedure similar to the one followed in step 5a, we arrive at the consolidated tutoring competency vector T .

- For the time being, we assume that all LNUs have tutor competency 1, but this need be changed soon as tutoring competency is a crucial factor in arriving at satisfactory answers. A little more sophistication could be added by giving expert LNUs a value of 1 and giving peers a smaller value (0.5).
- A tutor's past performance should preferably be taken into account. It could be based on some sort of rating (kudos) given by learners whose content questions were answered in the past and who were asked to rate the answer.
- It may be assumed that more 'difficult' questions (according to some measure) require tutors with a higher level of tutor competency. This could be achieved in step 7, by giving tutoring competency and greater load than content competency. Lacking a 'difficulty measure', we will not take this into account and assume that all questions are of equal difficulty.

5c Restrict a tutor's eligibility

Some LNUs will have more expertise than others, either from previous experience (see positioning) or because of their history in the LN. Tutoring is a matter of making your expertise available to the community to answer questions of fellow LNUs with less expertise. The tutoring load may increase rapidly with increasing expertise. After all, experts by definition are able to answer many different questions, beginners only few. This is unequal spread of the tutoring load is undesirable. The effort an LNU is willing to spend on tutoring is limited and largely independent on his or her expertise. Asking too many questions would thus lead to the quick exclusion of the expert tutors from the community structure. The question therefore is how to spread the tutoring load evenly. There is an additional, pedagogical twist to this argument. If tutoring is an educationally valuable experience per se - and not just a matter of community service - then LNUs should act as tutors for CQs that relate to ANs they have mastered themselves just yet. For those ANs, the educational value is likely to be maximal. An LNU's eligibility is the degree to which a particular LNU is preferred over others because of experience and workload considerations. [An appropriate way to measure eligibility still needs to be developed]

- In relatively small communities, a random drawing could be used since the number of LNUs would be too small to make some ineligible. This could be done as a first approximation.

5d Determine tutor availability

Tutor availability should at least take into account a past tutor load. Someone who has answered many questions over the last few weeks should be exempted; someone who has answered few questions only so far should be preferred also so as to ensure that all LNUs get equal opportunity to perform a tutor role.

If we plot a frequency distribution of the number of ANs a tutor is competent for (both content and tutoring competency) against the number of tutors, in all likelihood few LNUs will be suited to answer almost all questions (as they have high competencies on all ANs) and many LNUs will be suited to answer a few only (as they are competent with respect to only one or two ANs). Even if we were to randomly distribute CQ requests over all competent tutors, the load would be unevenly spread over the tutors:

Those who are competent for more ANs will proportionally receive more requests. Since we want to spread the tutoring load evenly, this effect has to be compensated for, by making the tutor availability depend on the past tutor load.

Availability = A_i

- Preferably, a measure more sophisticated than the total workload over some period of time should be developed. It could for instance take into account both the number of questions and the time lapsed since they were first asked.

- If suitability computing is done centrally, in the interest of maximising the chance of at all obtaining a reaction to a request to participate, the tutor's online status (online versus offline), perhaps even the expected online status could be taken into account.
- Another rule that may be considered is:
 - IF the time a tutor has available for performing a tutoring role until the due date of the question is less than the time it takes to answer the question THEN the tutor is NOT available. (This presupposes a due date; due dates may be provided by the learner or by the system itself, which sets a time horizon before which the question has to be answered.)

The results of steps 5a through 5d - that is an LNU's consolidated content competency C_i , his or her consolidated tutoring competency T_i , eligibility E_i , and availability A_i - are now used to determine the tutor's overall suitability S_i .

$$S_i = k_1 * C_i * k_2 * T_i * k_3 * E_i * k_4 * A_i$$

The K s are introduced to weigh the relative effects of the various factors. For example, the effect of a tutor's content competency could be half the effect of his or her tutoring competency, etc. Furthermore, $0 \leq k \leq 1$.

- Here, consolidated content competency, consolidated tutor competency and eligibility and availability have all been expressed as numerical values between 0 and 1. This allowed us to compute the overall suitability S_i . But an alternative approach could have been followed. Suppose that an initial ad hoc community is formed after having computed the overall competency from the content competency and tutor competency. Eligibility or availability (both or either one) could now have been used to trim the size of this initial ad hoc community. Applying eligibility criteria already results to the elimination of some tutors, availability criteria could remove more, if needed. What approach is chosen depends on the ease with which the eligibility and availability vectors may be computed and the extent to which eligibility and availability could perhaps be captured more truthfully in a set of logical, if-then rules. This also results in different diagram. In the present case, a tutor's suitability is computed by his or her Personal Agent. In the proposed case, eligibility and availability considerations need to be pulled in by the Matchmaker Agent. This means more central processing (assuming that the PAs are client side agents. If they aren't, the argument changes.).

6 Invite most suitable LNU peer tutors

Invite tutors by suitability ranking

The individual tutor suitability S_i is now collected for every tutor. Their rank order from the largest to the smallest represents the order in which LNUs should be invited to participate in the ad hoc community.

- If suitabilities are computed on the client side, the chances that the LNUS are actually available immediately for participation in the ad hoc community, is maximised.

7 Retrieve possible answer

Carry out detailed LSA

A new LSA may well be needed to seed the ad hoc community with proto-answers. In step 4, all documents in any one AN were pooled. This is adequate if the objective is to create a relevance ranking of content and tutor competencies. If, however, the objective is to produce proto-answers even whole documents lack sufficient resolution.

- Perhaps documents need even to be split up down to their paragraph level, at which stage an LSA is performed. If thus sufficiently high and distinct correlations can be produced, these are bound to be very helpful. The literature and perhaps some experiments with actual materials should shed light in this matter.

8 Invite LNU peer tutor

Invite i-th tutor

The tutor gets an invitation (by e-mail, by notification, by sms, by instant message?). The message contains the CQ and the list documents within ANs that the LSA has shown to be relevant to the drafting of an answer.

- An expiration moment should be set on the invitation to join to avoid a stall of the community formation. This moment should be in the order of hours at most, its duration also depending on the foreseen community size (the larger, the shorter).

9 Invitation

Join the ad hoc community?

The tutor may either refuse or agree to join (by clicking a pertinent URL?). If a tutor were to either accept or reject after expiration of the invitation, he or she should receive a message to the effect that the invitation has expired.

- Should there be a mechanism available still to join the community? It seems wise to allow people in who are motivated to do so.

10 Log acceptance, question form and possible answer

Log tutor participation

Include the thread's ID (in the form of a permanent hyperlink) so that the tutor is always able to go back to this thread.

11 Enough tutors?

Enough tutors?

An optimal size of the community should be decided upon. It should not be too small, otherwise the learner has to wait too long for a response to appear and loses faith in the system; not too many, otherwise LNUs might have to spend too much time answering questions by fellow LNUs. A size of 5 tutors seems reasonable.

12 Facilitate discussion platform

Create ad hoc community as a forum thread

The tutors will be asked to draft an answer to the CQ. In the interest of community building, all tutors should be approached at the same time to form an ad hoc community centred on the CQ. We propose that some open source forum software (Colloquia?) be used so that, in case of need, the software can easily be adapted to the ASA system.

- Community members could be allowed to bring in others whom they know to have pertinent expertise. Clearly, the invitation should be accompanied with the LSA results and, if they accept the invitation, their participation should be properly logged.
- Organised by CQ, threads may be made available to the LN as a kind of FAQs. They should not become a new AN, as they do not qualify for a full-fledged AN. This gives the LN a history and thus adds to its identity. This, in turn, should enhance community formation.
- Threads may be made available for LSA analysis at a later stage. This will prevent that lessons learned in a thread are forgotten; it will also speed up and ease the drafting of answers by tutors.

13 Formulate contribution

Both tutors and tutees formulate contributions to the discussion in the community.

14 Archive contribution

The contributions of the tutees and tutors are archived by the LNU agent.

15 Evaluate discussion

After each contribution the tutees and tutors evaluate this contribution.

16 Satisfied?

When a tutee or tutor is not satisfied with the contribution he or she can add a new one to the community.

17 Stop discussion

When the tutee is satisfied with the answers given in the discussion in the community the discussion is stopped. (Tutors can never stop a discussion).

18 Archive discussion

When a discussion is stopped by the tutee, the discussion is archived.

19 Rate tutor

The tutee can rate the tutor about the manner in which support was provided. This rating can then be taken into account in step 5.

Update calendar

The LNU or his personal agent needs to provide availability data, for example via an electronic calendar in the system. This data is needed to be able to calculate tutor availability and react on due times (steps 5, 8, 10 and 11).

Update dossier

Performance data, such as completion of activity nodes has to be logged in the LNU's dossier. This data is required for step 4 and 5.

Retrieve competence data from system

For an effective system, parameters for tutor suitability, in particular content and tutor competency should be logged. This is required for step 5 and is related to step 19.

Name	Sign up to LN
Summary	The LNU identifies him/herself in the LN.
Actors	LNU
Assumptions	The LN is accessible.
Description	The LNU fills in his/her name to login to the system. The system checks the LNU ID and if the ID is known by the system the LNUs personal setting and data are loaded.
Exceptions	The login fails when the LNU ID is not known by the system.
Results	The LNU is logged in the LN.

Name	Update electronic calendar
Summary	The LNU updates his/her electronic calendar.
Actors	LNU
Assumptions	The LNU is logged in the LN and his/her electronic calendar is available.
Description	The LNU accesses his/her electronic calendar, checks it and updates it.
Exceptions	NA
Results	The electronic calendar of the LNU is updated.

Name	Ask CQ
Summary	The tutee fills in a question form and asks a question.
Actors	Tutee
Assumptions	The tutee is logged in and has access to a question form.
Description	The tutee fills in an electronic question form which is taken up by the system. The system checks the type of the question and if it is a content question than the question is accepted by the system for further processing.
Exceptions	Other questions than content questions are not accepted by the system.
Results	The tutee asks a content question and the content question is stored.

Name	Send question form to LSA module
Summary	The reformed question form is send to the LSA module.
Actors	Agent Matchmaker
Assumptions	A content question is asked.
Description	The reformed question form is send to the LSA module for further processing.
Exceptions	NA
Results	The LSA module is provided with input.

Name	Send question form to Agent Matchmaker
Summary	The question form is send to the Agent Matchmaker.
Actors	Agent Tutee
Assumptions	The question form is available and correct.
Description	The question form that is filled in by the Tutee is send to the Agent Matchmaker by the Agent Tutee.
Exceptions	NA
Results	The Agent Matchmaker is provided with the question form.

Name	Process CQ
Summary	The content question is processed by the Agent Matchmaker and the Agent Tutee.
Actors	Agent Matchmaker, Agent Tutee
Assumptions	The question form is send by the Agent Tutee to the Agent Matchmaker.
Description	The Agent Matchmaker receives a question form and reforms it into input for the LSA Module. The Agent Tutee provides the Tutee with a question form.
Exceptions	NA
Results	The content question is saved in a question form and reformed to LSA module input.

Name	Process output for community
Summary	The Agent Matchmaker processes the output of the LSA module.
Actors	Agent Matchmaker
Assumptions	The LSA generated output.
Description	The LSA Module provides the Agent Matchmaker with information on the ANs a content question belongs to. The LSA Module provides the Agent Matchmaker with shreds of documents in which the answer to the content question could be found.
Exceptions	The output of the LSA module is empty and as a result a new content question has to be formulated by the Tutee.
Results	Identification of the origin of the content question and input for the community in the form of shreds of relevant documents.

Name	Form community with tutee(s) and tutor(s)
Summary	Based on the LSA output, the electronic calendar and the electronic dossier suitable tutors are identified that form a community.
Actors	Agent Matchmaker
Assumptions	The LSA engine generated output. The electronic calendar and the electronic dossier are updated.
Description	The Agent Matchmaker matches the ANs to which a content question belongs to the ANs in a LNU portfolio to determine the LNUs content competence, based on the LSA output. The Agent Matchmaker uses the electronic calendar to determine the availability of the LNU and the electronic dossier is used to determine the LNUs tutor competence. The Agent Matchmaker calculates the most eligible LNUs that could act as a tutor.
Exceptions	No suitable tutor can be found.
Results	Suitable tutors are identified among the LNUs.

Name	Invite community members
Summary	Suitable tutor(s) are invited to join the community.
Actors	Agent Matchmaker
Assumptions	Sufficient users are registered in the learning network and at least one suitable tutor is identified.
Description	The Agent Matchmaker sends invitations to the tutee as well as the suitable tutor(s) to invite them to join the community.
Exceptions	NA
Results	The tutee and the suitable tutor(s) receive an invitation to join the community.

Name	Handle invitation
Summary	The invitation to join the community is accepted or declined.
Actors	Agent LNU, LNU
Assumptions	The LNU has logged in the LN.
Description	The invitation to join the community is accepted or declined by the LNU through the LNU agent.
Exceptions	NA
Results	The LNU does or does not join the community which is communicated to the LN by the Agent LNU. When the LNU declines the invitation, new tutors are invited (refer use case invite community members and form community).

Name	Send input discussion
Summary	The LSA module output is send to the community.
Actors	Agent Matchmaker
Assumptions	The LSA module generated output.
Description	The shreds of documents that contain possible answers to the content question are sent to the community by the Agent Matchmaker.
Exceptions	NA
Results	Shreds of documents that contain possible answers are available in the community.

Name	Process input discussion
Summary	The possible answers are made available to the Tutor through the Agent Tutor.
Actors	Agent Tutor
Assumptions	The LSA module generated possible answers.
Description	The possible answers are made available to the Tutor through the Agent Tutor.
Exceptions	NA
Results	The possible answers can be accessed by the Tutor.

Name	Ask for clarification CQ
Summary	The tutee asks for clarification of an answer provided by a tutor.
Actors	Tutee
Assumptions	The tutee asked a content question and received at least one answer to this content question.
Description	The system sends the answers from the tutors to the tutee. The tutee reads these answers and when they are not entirely clear to him/her he /she asks for clarification of these answers through interaction with the community.
Exceptions	NA
Results	A request for clarification is send to the community.

Name	Draft answers to CQ
Summary	The tutor generates or edits answers to the content question.
Actors	Tutor
Assumptions	A content question is put forward and a community is formed.
Description	The tutor uses the community to provide the tutee with a possible answer to his/her content question.
Exceptions	
Results	An answer to the content question is put forward in the community.

Name	Archive rounded up discussion
Summary	The rounded up discussion is archived by the agent tutee.
Actors	Agent Tutee, Agent Matchmaker, Agent Tutor
Assumptions	The tutee is satisfied with the answers discussed in the community.
Description	The rounded up discussion is archived by the Agent Tutee, the Agent Matchmaker, and the Agent Tutor.
Exceptions	NA
Results	The rounded up discussion is added to the LN and available through the Agent Tutee, the Agent Matchmaker, and the Agent Tutor.

Name	Rate tutor
Summary	The tutee indicates his impression of the tutor's suitability by providing a rating.
Actors	Tutee, Agent Tutee, Agent Tutor
Assumptions	The community is created and populated with tutee and tutors and discussion has taken place.
Description	The tutee indicates his impression of the tutor's suitability by providing a rating. This is stored both in tutee's and tutor's dossier by the tutee and tutor agents.
Exceptions	No rating is provided.
Results	The rating can be accessed by tutor and tutee. The rating can be taken into account when computing tutor competence and eligibility.

Name	Round up discussion
Summary	The tutee rounds up the discussion when his/her content question is satisfactorily answered.
Actors	Tutee
Assumptions	The tutee asked a content question and received at least one answer to this content question.
Description	In the community the content question and its answers are discussed by the tutee(s) and the tutor(s). When the tutee decides that the question is satisfactorily answered he closes the discussion. The system acts upon this decision signalling it to the archiving process.
Exceptions	NA
Results	A rounded up discussion is signalled by the system.

Name	Update dossier
Summary	The LNU updates his/her electronic portfolio and the system updates the dossier/electronic portfolio.
Actors	LNU, LNU Agent
Assumptions	The LNU is logged in the LN and his/her electronic portfolio/dossier is available.
Description	The LNU accesses his/her electronic portfolio, checks it and updates it. The system automatically updates the dossier/portfolio with regard to study progress (= completed ANs) and data related to tutor competence and eligibility (e.g. rating scores).
Exceptions	NA
Results	The electronic portfolio/dossier of the LNU is updated.

Appendix 3. Class diagram

The entities in the diagram are a mix of dynamic actors and more static data objects. Association classes indicate associations that are performed by the agents or LSA Engine. The association classes indicate several roles the agents take on or the LSA Engine has to perform. In fact, the specialisations are methods of the super class. Directions of associations and multiplicity are not yet indicated in the diagram. Processes like rating of contribution and tutor are not depicted.

- b. LearningNetwork > LNU
The LNUs (learning network users) form part of the LearningNetwork.
- c. LearningNetwork > ActivityNode
The ActivityNodes form part of the LearningNetwork.

ActivityNode

The ActivityNode contains the documents that comprise the unit of learning. The set of ActivityNodes form part of the LearningNetwork.

- a. LearningNetwork > ActivityNode
The ActivityNodes form part of the LearningNetwork.
- b. ActivityNode > Document
The ActivityNode contains documents detailing the unit of learning (activity description and resources).

Document

The Documents form part of the ActivityNode. Request, portfolio (Dossier) and PossibleAnswer also are Documents. Documents form the input and query DocumentSpace that is required for the LSAEngine, as well as the output from the LSAEngine.

- a. ActivityNode > Document
The ActivityNode contains documents detailing the unit of learning (activity description and resources).
- b. Document > DocumentSpace
The Documents from the ActivityNodes and Dossier form the input DocumentSpace for the LSAEngine.
The Request is a Document that is queried on the DocumentSpace.
The PossibleAnswer is a Document that is retrieved from the AN input DocumentSpace.
- c. Document > PossibleAnswer
The LSAEngine retrieves Documents as PossibleAnswers from the DocumentSpace.

DocumentSpace

The Documents in the LearningNetwork (from ANs, portfolio and request) form the input and query DocumentSpace for the LSAEngine as well as the output from the LSAEngine in the form of PossibleAnswers.

- a. ActivityNode > Document
The ActivityNode contains documents detailing the unit of learning (activity description and resources).
- b. Document > DocumentSpace
The Documents from the ActivityNodes form the DocumentSpace for the LSAEngine.

LNU

The population of a Learning Network consists of Learning Network Users (LNUs). Every LNU has a personal LNUAgent to assist the user with various actions and functions in the LearningNetwork. Personal data and progress information of a LNU are stored in a Dossier.

- a. LNU > Request
The LNU puts forward a request, i.e. a question that needs answering.
- b. LNU > LNUAgent

Every LNU has a personal LNUAgent that assists and represents the user. Several of the LNU actions are conducted via the LNUAgent, although the LNU might not be aware of this.

c. LNU > Dossier

Personal data and progress information are stored in a Dossier. The LNUAgent is responsible for keeping the Dossier up to date, although both the LNU (personal data and portfolio) and the system (progress information, completion of ActivityNodes, etc) can add to the Dossier.

LNUAgents

Each LNU has an LNUAgent that represents and assists the user in the LearningNetwork.

The LNU Agent has at least 4 functions.

1. NegotiateValidity
 - a. LNU > Request
The NegotiateValidity negotiates the formal validity of the request with the LNU and subsequently sends it to the LSA Engine.
 - b. NegotiateValidity > Matchmaker Agent.
NegotiateValidity sends the request to the Matchmaker Agent for placement in the community.
2. UpdateDossier
 - a. LNU > Dossier
UpdateDossier keeps track of the LNU's dossier information
3. SendTutorDossier
 - a. SendTutorDossier > Matchmaker Agent
FindTutor provides the Matchmaker Agent with specific dossier details of the LNU.
4. RecordParticipation
 - a. LNU > Community
RecordParticipation records behaviour and participation of the LNU in the Community. This information can be used to determine tutor competency.

Dossier

Personal data and progress information are stored in a Dossier. Portfolio data can be present in the dossier.

- a. LNU > Dossier
UpdateDossier keeps track of the LNU's dossier information.
- b. Dossier > MatchmakerAgent
The MatchmakerAgent asks the LNUAgent to provide the tutor competence data from the LNU's Dossier.

Request

A request refers to a request for support (e.g. a content related question) of the LNU to the peer LNUs.

- a. LNU > Request
A LNU formulates a Request for support.

LSAEngine

The LSAEngine is responsible for mapping a Request onto the Documentspace of the LearningNetwork to find relevant source ANs and to query the Documentspace for PossibleAnswers.

1. MapRequestOnAN
 - a. NegotiateValidity > MapRequestOnAN

- NegotiateValidity forwards formally valid requests to MapRequestOnAN.
- b. MapRequestOnAN > Request
MapRequestOnAN receives a request and determines the most relevant AN to which this request belongs in the Learning Network.
- c. MapRequestOnAN > Matchmaker Agent
MapRequestOnAN sends the most relevant ANs to the Matchmaker Agent.
- 2. Query
 - a. Query > Possible Answer
Query performs a query with the request on the document space to find possible answers to this request.
 - b. Query > Matchmaker Agent
Query sends the possible answers to the Agent Matchmaker for placement in the community.

Matchmaker Agent

The Matchmaker Agent is responsible for filling the community with content and for populating the community with actors.

- 1. Matchmaker Agent
 - a. Matchmaker Agent > Community (Request)
The Matchmaker Agent places the formally valid request in the community.
 - b. Matchmaker Agent > Community (Possible answer)
The Matchmaker Agent places the Possible Answer(s) in the community.
- 2. PopulateCommunity
 - a. PopulateCommunity > LNU Agent
PopulateCommunity maps the specific LNU dossier (i.e. ANs) with the request ANs and asks the LNUAgent to invite the LNU to the community as a peer tutor.
 - b. PopulateCommunity > LNU Agent
PopulateCommunity grants access, through the LNUAgent, to the LNU that poses the request to the community as tutee.

PossibleAnswer

The PossibleAnswer is the result or output of the LSAEngine when querying the Request onto the DocumentSpace and consists of (shreds of) relevant documents.

- a. PossibleAnswer > Query
The PossibleAnswer is the output of the LSAEngine when querying the Request onto the DocumentSpace to find documents from ActivityNodes that might represent possible answers.
- b. PossibleAnswer > Matchmaker agent
Query sends PossibleAnswers to the MatchmakerAgent.
- c. PossibleAnswer > Community
The PossibleAnswer forms input for the Community.

Community

The Community is comprised of LNUs both in a tutor role and in a tutee role and contains PossibleAnswers. The Community has as function to compile an Answer to the Request on the basis of the PossibleAnswer.

- a. Community > LNUAgent
The Community is populated with LNUs via PopulateCommunity and LNUAgent with the tutee and a set of, at least one, tutors.
- b. Community > MatchmakerAgent {only valid request}
The MatchmakerAgent places a valid Request into the Community.



- c. Community > MatchmakerAgent {send possible answers}
The MatchmakerAgents places PossibleAnswers into the Community to be discussed by the LNUs.
- d. Community > FAQ Items
The Community generates the FAQItems.

FAQ Item

A FAQItem is formed when a request is successfully answered by storing request and answer.

- a. FAQItem > Community
The FAQ Item is the output of the Community.