

How to trigger emergence and self-organisation in Learning Networks

Citation for published version (APA):

Brouns, F., Fetter, S., & van Rosmalen, P. (2009). How to trigger emergence and self-organisation in Learning Networks. In R. Koper (Ed.), *Learning Network Services for Professional Development* (1 ed., pp. 57-72). Springer-Verlag Berlin Heidelberg. https://doi.org/10.1007/978-3-642-00978-5_5

DOI:

[10.1007/978-3-642-00978-5_5](https://doi.org/10.1007/978-3-642-00978-5_5)

Document status and date:

Published: 27/05/2009

Document Version:

Peer reviewed version

Document license:

CC BY-NC-SA

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: 07 Nov. 2024

Open Universiteit
www.ou.nl



Learners invariably have a lot of questions. These questions are of varying natures; some concern administrative or organisational matters, such as how do I enrol in this course, when is the exam; others are domain-related, such as what topic should I study next, or are content-related and arise when learners are studying a particular topic. In more traditional educational settings, there are provisions to cater for these questions. Administrative and organisational questions can be dealt with by the student administration services. Other types of questions are dealt with by the teachers, and tutors (De Vries et al. 2005). Turning to the teacher when a learner has a question, or asking a fellow student is more obvious in traditional settings, because in these circumstances learners usually know who partakes in the course. And there often are more opportunities, as there usually is a lot of face-to-face contact. In Learning Networks as we define them, it is more cumbersome for learners to obtain an answer to their question. The traditional support structures are lacking. In these circumstances, it is even highly likely that tutors are unavailable for several subject domains, e.g. because it is not economically viable, or because the resource provider did not provide for them. For certain questions people can be directed towards (online) resources, search engines or FAQs, but not all questions can be resolved in this manner. Some questions, such as content related questions, or requests for comments on a paper, or discussions on a certain topic can be dealt with only by contacting somebody else. Often people do not know whom to contact for a particular request.

This is where ad-hoc transient communities come in. In this circumstance an ad-hoc transient community will be set up to assist a learner who has a particular question. The participants of the ad-hoc transient community will consist of peers with the required expertise to support the learner in finding an answer to a specific question. The peers are selected by the peer support Learning Network Service, based on certain selection criteria that suit the goal and adhere to the guidelines for heterogeneity, roles and goal as described in Chapter 4.

Let's illustrate this with the case introduced in Chapter 2 Bas Timmer is one of the members of the vintage motorcycles Network. He currently works as a mechanic, but really would like to run his own garage. He already started with a business plan and already engaged in several lively discussions about the plan with his fellow Network members Jannie and Jessica. These discussions made him realise that technical skills alone are not sufficient to run a garage. He needs to learn how to manage a business and update his accounting and customer skills. Bas decides to join another Learning Network, on business administration, and starts with accounting. One of the first things to learn is how to make up a balance sheet. His maths is not very strong and he cannot figure out how to do the credits and debits properly. In particular, the assets accounts are confusing him. Just reading about it is not taking him any further and he decides to ask somebody for clarification. Because he effectively is on his own, he decides to call in the help of other people in the Learning Network. However, he does not know whom to contact and turns to the peer support service of the Learning Network. Bas accesses the support request form. The request form guides him in detailing his question.

The Learning Network Service creates an ad-hoc transient community, facilitated with relevant communication and collaboration tools, to which Bas and relevant peers with complementary expertise get access, with the goal to reach jointly at a satisfactory answer for Bas' problem.

5.2 Addressing the Solution

Our first attempt at implementing a peer support Learning Network Service focuses on content-related questions. The main assumption is that for the ad-hoc transient community to be effective we have to select suitable and eligible peers. The data on which to base this selection should be taken from a participant's profile. This requires that the relevant data is recorded into the profile, and that the Learning Network stores a record of every participant's progress through the Learning Network preferably into the participant's profile. Chapters 3 and 4 already indicate some of the data that should be stored. Below we explicate the data actually used in this implementation. In addition, the Service requires access to all documents that constitute the resources in the Learning Network.

5.2.1 Steps Involved

The model contains the following steps (Van Rosmalen 2008; Van Rosmalen et al. 2006; Van Rosmalen et al. 2008b) as illustrated in Fig. 5.1.

Formulating a support request A question form is available to the participant who wants to pose a question (tutee). This form contains guidelines on how to phrase the question in such manner that it provides maximum input to the peer tutors. Additional information related to the question, such as urgency, date by which an answer is expected or required, subject domain of the question (if known) or the learning activity that raised the question, should also be provided on the form. Information related to the identity of the tutee can be determined automatically when the participant has to authenticate for the Learning Network or the form, else the form should request this data as well.

Defining the context of the support request Usually somebody will have a question while studying a learning activity and the question will be related to that learning activity, but this does not have to be so. When somebody is studying more learning activities at the same time, in particular when these activities are closely related, it is not always possible to exactly pinpoint the activity that gave rise to the question. Sometimes a question relates to more than one learning activity. Or the question might be related somehow to the subject domain, but not di-

rectly to content in one of the learning activities. An indication of the context can be taken from the form. When a question form is available from within a certain learning activity, this can be automatically set on the form, but the tutee should be able to override this. When the form is also available outside the context of the learning activity, e.g. as a generic service at the Learning Network level, the form should contain a field in which the learner indicates the subject domain or learning activity to which the question pertains. Instead of forcing the learner to indicate always a learning activity, the option ‘Do not know’ can be used when the learner does not know exactly which learning activity the question relates to. The context of the question is rather important because it provides the basis on which suitable peers are selected. Therefore, we not only rely on the information provided by the learner but also have a mechanism in place to determine the learning activities that are relevant to the question. We use language technologies to assist in analysing the questions, which after all are formulated in natural language. Latent Semantic Analysis (LSA) is quite often used in an educational context (Van Bruggen et al. 2004; Landauer et al. 1998). LSA is used to find documents that are related because concepts are used instead of searching for the exact keyword. This allows people to use more of natural language in search queries instead of having to specify the correct keywords. We apply LSA to determine correlations between the text of the question and the documents in the Learning Network. Because every document belongs to a single learning activity, we can extrapolate to the learning activities the question is related to.

Fig. 5.1. The five steps involved in peer tutoring ad-hoc transient communities

Identify potential peer tutors This step involves finding and selecting suitable peer tutors and deciding upon the optimal number of peer tutors. The peer tutors are selected based on four criteria: content competence, tutor competence, availability and eligibility, as explained below. Most of the data required for the selection algorithms, such as activities taken, competences achieved, and learning goal, are already stored in Learning Networks and can be taken from a learner’s profile, as well as system’s logging data. Other data might require additional services that are not always present in a Learning Network; availability, for example, can be retrieved from a calendar, rating services need to be added to allow members to rate tutor’s ability and quality of contributions.

Support for creating the answer The previous step returns a ranked list of suitable peers. A number of these peers are then invited to assist in answering the question and are asked to join a collaborative workspace. The invitation contains the question and guidelines. The invitation or notification of the invitation preferably is sent out via e-mail, to ensure a quick response. The ad-hoc transient community is supplied with the necessary communication and collaboration facili-

ties (such as a wiki or discussion forum), that allow the peers to jointly discuss and phrase an answer to the question. The question is made available, together with guidelines on how to proceed. To lower the peers' required effort, the system retrieves relevant text fragments related to the question from the resources in the Learning Network and makes those available in the ad-hoc transient community. These documents should assist the peers in getting a quick overview of documents relevant to the questions, and could also contain part of the answer or be the basis for the formulation of the answer.

End of process After some time the peers should be able to formulate the answer to the question. The tutee decides when the process is ended; hopefully when the answer is satisfactorily. There may be several other reasons for a tutee's decision to end the process: because the peer tutors fail to provide an answer the tutee can agree with, because the process did not get started at all, or maybe just because it looks like no decent answer can be found in the near future, or they run out of time. Whatever the reason, it is up to the tutee to end the process. When the process is ended, the tutee should rate the peers and the product. The ad-hoc transient community then is dismantled and ceases to exist.

5.2.2 Peer Selection Criteria

We identified four criteria for the selection of peers who can act as potential tutors when it comes to answering content related questions. These criteria are designed to ensure quality and economy. Only suitable or competent peers should be selected, otherwise it will not be possible to get an acceptable answer to the request or query. In addition, peers should be able to meet the request and there should be a fair spread in workload. The weighted sum of these four criteria results in a rank-ordered list of suitable tutors. More details on the algorithms can be found in (Van Rosmalen et al. 2008b).

- *Content competence.* First, the peer tutor should be sufficiently competent and knowledgeable in the subject domain in order to answer the question.
- *Tutor competence.* The peer tutor should have sufficient tutoring capabilities to be able to support the tutee.
- *Tutor eligibility.* The peer tutor should be eligible to answer the question.
- *Availability.* And of course, the peer tutor should be available to respond to the request within a certain timeframe.

Content competence is a measure of how well a peer has mastered the content relevant to the question. This information can be obtained from the member's pro-

file in which evidence for competences attained is stored or in which the learning activities of the member together with completion status and competence level are recorded. However, competence assessment is not trivial, and even when the profile contains some information regarding competence, the interpretation of these data is not straightforward, let alone determining the competence proficiency level the learner obtained. An alternative would be to consider completion status of learning activities. Usually learning activities are designed for a certain learning objective or competence, and an assessment is used to ascertain whether the learner has met the objective. The result of the assessment can be taken as an indication of proficiency level. In our model, the content competence is expressed as a weighted sum of completion status of all learning activities that are relevant to the question. For a more detailed measure, time since completion can be taken into account.

Tutor competence indicates the ability of a peer to act as a tutor and is a measure of tutor capabilities. It refers to the capability to satisfactorily support the learner who has questions in such way that the tutee is satisfied with the process and the product. By nature, the staff tutor would be the most suitable tutor. Similarly, we can assume that learners who have progressed furthest in the Learning Network have obtained the most content competence and therefore are likely to be the most suitable peer to select. In addition, the peers that have been selected before or were more often selected, are likely to be better tutors because they gained experience in previous situations. Tutoring competence is difficult to measure because the evidence is hard to record. Another way to obtain a measure of tutoring competence is with a rating mechanism, in which the members rate both the tutoring process and the final answer. Or it could be derived from the frequency, size and quality of the contributions.

Tutor eligibility is a measure that is required to ensure a proper spread of workload amongst the members of the Learning Network. To avoid that the same people are selected time and again, the tutee should be compared to the potential peer tutors and only those peers that are nearest to the tutee should be selected. The nearest peers are those members that have a similar learning path, or at least have progressed as far. This can be measured by looking at the goal, i.e. the learning activities a member needs to complete to obtain the chosen competence and the completion status of the learning activities in the learner's profile. Peers who have completed the same learning activities (or get as close as possible to the tutee's profile) are most eligible.

Availability refers to the actual availability of the peers. Is the member present and active in the Learning Network, taking into account absence for holidays, days off, illness, etc., and workload in the sense of study-load, i.e. studying for exams or preparing an assignment, and previous participation (acting as tutor)? This measure is also time-dependent: more recent workload should have a higher weight or affect availability more than workload in the past.

For other types of support requests, these four criteria equally apply. The underlying algorithms might vary or be based on different parameters. In some cases the weight of the four criteria might need to be adjusted. Content competence is much less relevant when answering questions related to administrative matters. And when asked to answer a question about time and place of the assessment content and tutoring competence hardly are required. Nevertheless, the same principles apply and the selection criteria are still involved.

5.2.3 A Tutor Locator: Example Implementation of the Model

The model we described above is implemented in a first prototype, *ATL A Tutor Locator*. With it, we intended to test and validate our assumptions about ad-hoc transient communities and to validate the steps and selection criteria. We distinguish a request component, a population component, a community component, an indexing component, and a database to store relevant data (Kester et al. 2007) as depicted in Fig. 5.2. The request component covers the first two steps of the model (formulate and define context of support request) and provides the basis for the third step, identification of suitable peer tutors. The population component handles step 3 and the logistics required for step 4. The community component finalises the logistics and provides the facilities for the peers to enter the peer support process, up to the finalisation of the process (step 5). The indexing component precedes these steps and indexes all documents in the Learning Network. Finally the database stores personal data, data on learning activities and resources in the Learning Network, as well as results of the several components.

Fig. 5.2. The main components: Moodle constituting the Learning Network with learning activities, resources, question form and wiki; indexing component to create the corpus; the request component to define the context; the population component to select suitable peers and set up the wiki; and the community component for peers to discuss and formulate an answer. [Adapted from Van Rosmalen et al. 2008b].

The Learning Network was implemented in Moodle, an open source platform for online learning environments (<http://moodle.org>). The four components have been developed as add-ons to Moodle or as an application that could be called from Moodle. The sources are made available under the BSD licence and can be downloaded from Sourceforge (<http://sf.net/projects/asa-atl>) or from our repository (<http://hdl.handle.net/1820/960>). For certain parameters in each of the steps

we had to make decisions on how to implement and apply the model. These are described below.

Indexing Component

This component makes the preparations that are required for our model to operate on. First, it retrieves all resources that are used in the learning activities. These resources are split into better manageable text-fragments and saved as text documents. An index is created to record which document belongs to which learning activity; every document can occur in a single activity only. The documents constitute the text corpus that is required for LSA. The corpus is pre-processed by removing common words ('stopping'). LSA is then applied to the corpus to set up the term versus document matrix. This matrix is required to be able to match the question to the text corpus and find the documents to which the question correlates. We use the General Text Parser software (GTP) (Giles et al. 2003) to carry out LSA algorithms.

Request Component

This component provides the question form, assisting in formulation of the question, and determines the context of the question. The question form was developed as an add-on to Moodle. The members should be able to access the question form from anywhere in the Learning Network and from all learning activities in the Learning Network. The question form contains some instructions and guidelines for proper formulation of the question. There is some form of validation of the question, by checking whether the question is long enough. When the question is very short, a message appears, explaining that the question is not formulated well enough for the peer to understand what the question is about and the learner is prompted to reformulate the question. The learner also has to indicate from which learning activity the question originates (if known) and has to provide details such as within which time limit an answer is required. Of course the learner identity is also required. In our case, identity is known automatically because the Learning Network requires authentication.

When the form is submitted several algorithms are invoked. First we verify the context of the question by using LSA to match the question to the documents of the learning activities. The documents with the top three (this number is configurable) highest correlations to the question are retrieved. Via the index in the database (as prepared in the indexing component), we can determine to which learning activity these documents belong. We then calculate a weighted sum of the correlation for each of these learning activities.

Population Component

This component is developed to select the most suitable peers to populate the ad-hoc transient community. There are two main processes, the selection process and the invitation process.

Selection process In the selection process the content competence, tutor competence, eligibility and availability are calculated to return a rank-ordered list of suitable peers. The suitability is calculated as a weighted sum (value between 0 and 1) of the four criteria. The content competence is considered to be the most important and is included with a weight of 1; the eligibility and availability each with a weight of 0.5. In our situation, tutoring competence was not available, and thus received weight = 0. We calculate the other measures only for those learners with a positive value (> 0) for content competence. If the content competence algorithm fails, because there are no other learners with a positive competence for the relevant activities, we fall back to a ‘random’ selection in which we only take availability into account. The tutor suitability measure is a dynamic measure, and will be different every time the criteria are calculated. That is why the tutor suitability for every learner for every particular question is stored in the database.

Content competence Preferably, content competence is based on assessment of competence and proficiency level attained. In our situation we did not have access to learners’ profiles nor could we conduct some form of competence assessment, because no data on competences were available. Most online learning environments however, offer the possibility to log members’ actions, including whether the learner has accessed certain learning activities. Moodle does not provide a profile, but logs whether learners accesses a particular learning activity. Every learning activity in Moodle can conclude with an assessment. Individual assessment grades can be retrieved from the database. In the present implementation the content competence algorithm is based on the completion status of the learning activity. As explained before, via LSA we can determine which learning activities are relevant for the question (via the correlation between the question and documents from the learning activities). This is required to determine content competence. For every learner we calculate the content competence for each of these learning activities, as a weighted sum of the correlation and competence proficiency level for each of the learning activities. The correlation is provided by the LSA component. The content competence for each learning activity is determined by the score of an assessment, which is added to every learning activity. When the learner never accesses the learning activity, the content competence for that learning activity is set to 0; when the learner has started with the learning activity but not yet obtained a sufficient score for the assessment, content competence is set to 0.3; when the learner successfully passes the assessment, content competence is set to 1. A more accurate measure of content competence should take into account the score (or competence proficiency level), time expired since completion of the learning activity, and time required to study (time lapsed between start and successful completion).

Tutor competence Tutor competence is a weighted sum of previous contributions, derived from logging and ratings by other peers. This is based on how actively the member participated in previous questions and how answers to previous questions have been rated.

Eligibility Eligibility is a measure that indicates which peer is most near or similar to the learner asking the question. It is calculated over all learning activities in the Learning Network, but excluding those that correlate with the question. So the learning activity from which the question originates is omitted from the equation. For all other learning activities we need the content competence (0, 0.3 or 1) for every member. When the peer has the same content competence as the tutee, the value becomes 1, else 0. The total eligibility for the peer is the sum of all values divided by the number of learning activities involved. This is calculated for every peer.

Availability Availability depends on presence and workload (e.g. currently to busy studying, preparing for exam, work), but should also take past workload (being involved in answering questions) into account. Data on presence can be taken from a calendar, but requires a calendar that members maintain accurately. Lacking this, we calculate availability as a relative measure on the basis of the number of invitations sent for the learning network, the number of invitations accepted by the member, the average number of invitations per member over a certain time period. This is set against a threshold M . A member can be at the left or right side of the threshold. Depending on where he is in the range, availability varies from 0, 0.25, 0.5, and 0.75 to 1. For a member who has contributed more than average, availability is set to 0 when the member contributed more than M above the average or to 0.25 when contribution is above average, but not more than M above average, while for a member who has not been very active, availability is considered to be high (1). The value for availability is set to 0.5 when the member contributed on average. We had to build in a certain period of time (set number of hours or days) over which to calculate availability to allow for members who did not respond (in time) and omit those for that set period from the selection algorithm. There are several reasons why a member does not respond to an invitation: not present, not available, time pressure, taking exams, holidays, etc. When insufficient members are available the invitation process fails. To prevent that these unresponsive learners are included in new invitation cycles, we exclude those learners for a certain pre-defined period of time.

Invitation process When the list of suitable peers has been calculated, the invitation cycle can start. First it checks whether the suitable peers have any outstanding invitations. To prevent a lot of invitations not being accepted because they get stuck in this loop, peers with outstanding invitations are removed from the invitation list. Invitations are sent out by e-mail to the peers selected, and contain the question as formulated by the tutee, the invitation to assist in finding the answer and instructions about how to proceed and accept.

A new page was added to Moodle to allow the peers to either accept or decline the invitation. When the peers accepted they are asked to provide some information on their alleged competence level.

The invitation cycle is complete when the specified minimum number of peers has accepted. In this case, another e-mail is sent to the tutee and peers to inform

them that the invitation has been successful. The message contains the name of the peer who has accepted as well as instructions on where and how to access the ad-hoc community. It is possible that an insufficient number of peers accept, either because they decline or because they do not respond at all. After all, they might not be present when the invitation is sent. When none of the peers accepts, a new invitation cycle can be started, this time omitting the peers who declined from the selection list. Again, the list of suitable peers is checked against learners with outstanding invitations. When this cycle also fails, the tutee receives an e-mail message that unfortunately no peers can be found to assist with the question. The required group size (2), the minimum number (1) of peers that have to accept, the time period peers can take to respond (2 days) and how many invitation cycles (2) should be run, are configurable parameters.

Community Component

Finally, the ad-hoc transient community is set up. When the required number of people has accepted the invitation, a wiki is created. At that time, another e-mail message is sent to inform the tutee and peers that sufficient peers (in our case at least 1) have accepted the invitation and a hyperlink to the wiki is provided. (A forum might be used instead, but is more suitable for discussion than jointly writing an answer.) The wiki module of Moodle had to be modified to add instructions on how to use the wiki, guidelines on how to respond to the question, and to add the text of the question, and the three relevant documents. Access to the wiki is granted only to the tutee and selected peers.

A few more modifications to Moodle had to be made. In Moodle, wikis normally are added to a learning activity and are available to all learners. We had to provide an entrance to the wiki that is restricted to the members involved in a particular ad-hoc transient community. The link to the wiki is included in the invitation e-mail, but members need an alternative in case the e-mail gets lost. The tutee has to be able to end the process, and rate the contributions. In addition, all members require an overview of the status of the ad-hoc transient communities they are involved in. This was accomplished by adding a couple of additional pages to Moodle. The first page provides the member with an overview of the status of the ad-hoc transient communities they are involved in, either because they invoked them or because they accepted an invitation of somebody else. Via this page, the members can enter the wiki. The tutee also can change the status of the wiki to 'Completed'. When the tutee changes the status, he is asked to rate the contribution. After the status has changed to 'Completed' access to the wiki is closed, but the wiki and all contributions are stored.

Table 5.1. The peer selection algorithms [Adapted from Van Rosmalen et al. 2008b]

Explanation	Algorithm	Weights
<p>T_{s_L} Tutor suitability of learner L</p> <p>Value between 0 (not suitable at all) and 1 (very suitable).</p> <p>Weights to adjust relative importance of the four measures.</p> <p>Calculate only when content competence > 0, to assure a minimum knowledge level.</p> <p>To assure presence, learners with available time = 0 are removed from the list.</p>	$T_{s_L} = \left(\frac{(WT \times T_L) + (WE \times E_L) + (WA \times A_L) + (WC \times C_L)}{(WT + WE + WA + WC)} \right)$	<p>WT = 0</p> <p>WE = 0.5</p> <p>WA = 0.5</p> <p>WC = 1</p>
<p>T_L Tutor competence</p> <p>Value between 0 and 1.</p> <p>Tw_1 and Tw_2 to adjust the relative importance of T_e (on average how active the learner behaved in previous questions) and T_r (on average how previous answers were rated).</p>	$T_L = ((Tw_1 \times T_e) + (Tw_2 \times T_r)) / (Tw_1 + Tw_2)$	Not available, since WT=0

E_L Eligibility

Value between 0 and 1.

E_L is relative to L_q (tutee).

Calculated over all learning activities (AN) that do not relate to the question.

$$E_L = \left(\frac{\sum_{i=1, \dots, N \text{ \& all } i | \text{ ANi is not question related}} (\text{score}(\text{ANi}_L) = \text{score}(\text{ANi}_{Lq}))}{(N - \# \text{question-related ANi's})} \right) /$$

AN score can be:

0 – not started

0.3 – started

1.0 – assessment completed successfully

A_L Availability

Value between 0 and 1.

M (max_extra_workload) and T_p (time period over which workload is calculated).

Depends on relative past workload, comparing number of times a member is involved relative to other members in given time period.

A_L = one of [0, 0.25, 0.5, 0.75, 1]

0.5 if L contributed on average; 0.25 if L contributed above average but not more than M above average; 0 if L contributed more than M above average

$M=1$

C_L Content competence

Value between 0 and 1.

D : number of documents to calculate LSA correlations for.

D_i : number of documents provided.

W_{ANi} based on correlation between question and documents in ANi.

C_{ANi} : content competence for ANi, taking into account score, time expired since completion and study time of ANi.

$$C_L = \frac{(W_{AN1} \times C_{AN1}) + (W_{AN2} \times C_{AN2}) + \dots + (W_{ANn} \times C_{ANn})}{(W_{AN1} + W_{AN2} + \dots + W_{ANn})} /$$

$D=3$

$D_i=3$

C_{ANi} based on score of ANi

5.3 Conclusion

Applying the peer tutoring ad-hoc transient communities in an actual Learning Network provided promising results. The model was evaluated in an experimental setting (Van Rosmalen et al. 2008b), involving over 100 learners divided in two groups. In the experimental group, the following weights applied: 1.0 for content competence, 0.5 for availability, and 0.5 for eligibility. In a newly started Learning Network, learners have not yet developed any content competence. Because content competence is an imperative constituent of the model, this could result in ‘cold start’ problems: the selection does not return any suitable persons because none has the required content competence. To prevent that from happening, we resort to the same selection mechanism as used for the control group. In the control group, we only tried to distribute the questions evenly across all members by taking availability and past workload into account. Thus, the selection criteria were applied with weight 1.0 for availability and 0 for all other criteria.

We looked at the number of questions solved successfully. This was determined by the learner asking the question and by two expert tutors. For the system to be effective at least 50% of the questions had to be solved by the peer tutors. To determine whether our selection criteria made any difference, we compared the number of invitations accepted, the time needed to answer the question, the number of questions answered and the quality of the answers in the experimental group with those in the control group. We also looked at whether the peers used the three documents and whether they perceived them to be useful.

More questions were asked and above all a much larger percentage of the questions was answered in the experimental group (71%) than in the control group (45%). More importantly, a single invitation cycle was sufficient in around 80% of the questions in the experimental group, while in the control group at least 2 invitation cycles were required for at least half of the questions. In addition, in the experimental group questions were answered within a day and a half, while the peers in the control group took more than 2 days. In the control group, a larger proportion of the learners was involved, as side-effect of the model that tries to distribute workload evenly among the learners when no other measures are available.

Initially we allowed peers a period of six days to respond to an invitation. Due to the limited number of members in the Learning Network and the small number of invitations being sent out, this resulted in many outstanding invitations not being answered, while the same members were repeatedly invited for multiple questions. Therefore we had to decrease this response time to 2 days to prevent the selection from failing. We did not have proper availability data, other than past workload. Having members indicate the time and days when they are available, or determining when somebody is online can alleviate this problem.

The peers were asked to rate their own competence in answering the question when they accepted or declined an invitation. And the tutee rated the answer. In the experimental group, the peers were more outspoken in their ratings: more either absolutely accepted or declined, while in the control group the ratings were more evenly distributed. This indicates that the selection criteria returned the more suitable peers.

Although the supplied documents were used in a small proportion of the questions, members tended to rely more on course content, prior knowledge or non-specified resources; the learners in the control group used the documents more often. In particular, when asked to rate the usefulness of the documents they rated them at a higher level. Apparently the documents were more useful to the learners in the control group as the learners were randomly selected. The usefulness of the documents also depends on the type of question, the complexity of the domain and the previous knowledge and content competence of the peers. When the selection delivers a good match between question and suitable peers, the documents might not be needed.

There were some drawbacks to the implementation. The participants were not familiar with wikis and did not engage in a joint writing process. Instead they used the wiki more like a forum. Care should be taken to either use a facility participants are familiar with or to provide elaborate instructions on its purpose and proper way of use. The wiki did not notify peers when the wiki was changed; so some members complained that they wanted to elaborate on their question but no longer got a response from the invited peers. Nor were the people involved notified when the tutee closed the question or when it was successfully answered. A clear notification system seems imperative. People at least need to be able to see what is changed and preferably are notified about any changes. The peers in the ad-hoc transient community need additional communication facilities, for example by providing e-mail or Skype contact details. These kinds of data can be added to a personal profile the learners can share. In the personal profile the learners can indicate what information is public and what is private.

As discussed, the model was tested in an experimental settings and the learners were aware of this. This might have prompted a more extensive use of the system. The experiment ran only for a short period of time, while a Learning Network won't have a fixed duration and learners will not all start and end at the same time. Additionally, the Learning Network domain was rather small and of a relatively simple level. In more real-life Learning Networks incentives for participation should be made clear and policies such as described in Chapter 3 should be applied to ensure prolonged participation.

Interestingly and importantly, all members indicated that answering questions is a good investment of time. The mostly selected reasons given were: "I'm aware that other students also have questions" and 'It improved my knowledge and understanding'. In particular the learners in the experimental group gave more positive feedback and indicated that they would like to use ad-hoc transient communities in other Learning Networks (Van Rosmalen et al. 2008a). Furthermore the members remarked that they wanted facilities that allow them to contact the peers

outside the context of the peer tutoring ad-hoc transient community. This can be taken as a first step towards enhancing social interactions in the Learning Network as a whole.

We can conclude that ad-hoc transient communities are well-suited for answering content-related questions by involving peers. Moreover, the model implementation offers sufficient evidence that ad-hoc transient communities contribute to improving the community aspect of the Learning Network by enhancing the social interactions. The most important aspects to be covered by the service are:

- Base the peer selection criteria should at least be based on eligibility and past workload. Important for the eligibility measure is the content competence relative to the other members.
- As the eligibility measure depends on an accurate measure of content competence, record data for this in a learner's profile; logging at least measures of content competence, e.g. completion of activity, date activity started and completed.
- Ensure that the question can be related to the correct activity or competence by using appropriate technologies.
- Allow learners to decline only a set number of invitations.
- Obtain valid availability data, e.g. by using a calendar.
- Allow members to contact the peers outside the context of the ad-hoc transient community, e.g. by adding the peer to one's list of contacts, or friends, noting reasons for doing so.

References

- Van Bruggen, J. et al.: Latent semantic analysis as a tool for learner positioning in learning networks for lifelong learning. *Br. J. Educ. Tech.* **35**(6), 729-738 (2004)
- De Vries, F.J. et al.: Identification of critical time-consuming student support activities in e-learning. *ALT-J Res. Learn. Technol.* **13**(3), 219 - 229 (2005)
- Giles, J.T. et al.: GTP (general text parser) software for text mining. In: *Statistical data mining and knowledge discovery* (CRC Press, Boca Raton 2003) pp 455-471
- Kester, L. et al.: Matchmaking in learning networks: Bringing learners together for knowledge sharing. *Interact. Learn. Environ.* **15**(2), 117 - 126 (2007)
- Landauer, T.K. et al.: Introduction to latent semantic analysis. *Discourse Process.* **25**, 259-284 (1998)
- Van Rosmalen, P.: *Supporting the tutor in the design and support of adaptive e-learning* Unpublished PhD thesis (Open Universiteit Nederland, Heerlen 2008)
- Van Rosmalen, P. et al.: A model for online learner support based on selecting appropriate peer tutors. *J. Comput. Assist. Learn.* **24**(6), 483-493 (2008a)

60 Francis Brouns et al.

Van Rosmalen, P. et al.: Knowledge matchmaking in learning networks: Alleviating the tutor load by mutually connecting learning network users. *Br. J. Educ. Tech.* **37**(6), 881-895 (2006)

Van Rosmalen, P. et al.: A learner support model based on peer tutor selection. *J. Comput. Assist. Learn.* **24**(1), 74-86 (2008b)

