

Fixing the 'Out of Sight Out of Mind' Problem

One Year of Mood-Based Microblogging in a Distributed Software Team

Kevin Dullemond, Ben van Gameren
Delft University of Technology

IHomer
The Netherlands
{k.dullemond, b.j.a.vangameren}@tudelft.nl

Margaret-Anne Storey
University of Victoria

Canada
mstorey@uvic.ca

Arie van Deursen
Delft University of Technology

The Netherlands
Arie.vanDeursen@tudelft.nl

Abstract—Distributed teams face the challenge of staying connected. How do team members stay connected when they no longer see each other on a daily basis? What should be done when there is no coffee corner to share your latest exploits? In this paper we evaluate a microblogging system which makes this possible in a distributed setting. The system, WeHomer, enables the sharing of information and corresponding emotions in a fully distributed organization. We analyzed the content of over a year of usage data by 19 team members in a structured fashion, performed 5 semi-structured interviews and report our findings in this paper. We draw conclusions about the topics shared, the impact on software teams and the impact of distribution and team composition. Main findings include an increase in team-connectedness and easier access to information that is traditionally harder to consistently acquire.

I. INTRODUCTION

The year is 2011 and in a young and small software company called IHomer a significant problem is being discussed. Because the default work location in the company is home the people work distributed from each other most of the time, and have always physically met up once a week on Tuesdays in order stay connected. However, one employee has been away on contract for a prolonged period now, preventing him from attending the weekly meet-ups and he is starting to feel more and more disconnected from the rest of the team. On the one hand he is starting to lose touch with what occupies his colleagues and on the other hand his colleagues weren't even aware he was feeling unhappy with the overall situation. Fast forward almost two years to the present day and the issue has been dealt with by the introduction of WeHomer, a microblogging environment in which people at IHomer share their activities and moods with each other, to stay current on each other's feelings, experiences and latest exploits and as a result stay connected as a team.

The story above is not unique as Software Engineering is becoming more and more distributed. This is caused by the increasing globalization of business [1], [2], [3] and the rising popularity of working from home [4]. At the same time, significant challenges are faced when collaborating in a distributed setting as reported in the well-known work of Olson and Olson [5]. More recently Nguyen et al. [6] have investigated whether the effects of distance on distributed communication delay and task completion reported in the

literature still exist. They report that advances have been made but also that more work is needed and there are still many open research questions.

We believe both microblogging and mood sharing are essential to alleviate challenges arising from this distributed nature. On the one hand microblogging is essential, since being able to exchange small elements of content makes people feel more connected with others, especially when people work distributed from their colleagues [7], [8], [9]. On the other hand mood sharing is essential, since being aware of the emotional state of your colleagues makes it possible to act accordingly and achieve better results in joint work [10]. Therefore, in this research, we aim to determine whether an environment in which one can both express himself and get a sense of how his team members feel is valuable to distributed software engineers. In this paper, we use a microblogging solution extended with mood indicators (MBMI) called WeHomer to learn from the use of such an environment. The main goal of the paper is:

"To understand how microblogging with mood-indicators helps distributed organizations in knowledge sharing"

Furthermore we have identified four research questions:

- What sort of topics are discussed in MBMI?
- What is the impact of the introduction of a MBMI on a software team?
- What is the impact of distribution on the use of a MBMI?
- How does team composition impact collaboration with MBMIs?

We structure this paper as follows: In the next section we discuss related work to this research. In section III we discuss the research site and methods of data collection and analysis. Subsequently we show descriptive statistics in section IV and present the most important findings in section V. Finally we discuss threats to the validity of our study in section VI and conclude upon our work and discuss future work in section VII.

II. RELATED WORK

Software Engineering is by nature a highly collaborative activity, and having access to knowledge about the context in

which you are working is essential to properly collaborate with others [11], [12]. In literature this kind of knowledge is commonly referred to as 'awareness' [13], [11]. When working co-located this information is exchanged relatively passively and unobtrusively [11], [14], so people are continuously aware of information related to their current context [15]. However, when people no longer share a physical work environment exchanging awareness information without technological support becomes unfeasible [15]. Therefore, the (Global) Software Engineering community has developed and studied a wide variety of tools, for example: Instant Messaging, email, issue management systems and configuration management systems.

According to Bly et al. [16] it is particularly important to recognize the need for informal interactions, spontaneous conversations, and general awareness of people and events when teams are geographically dispersed. Storey et al. advocate further research on understanding how social media plays a role in (Global) Software Engineering [17]. One of the potential implications they identified in their research, concerns the challenges which arise when teams are distributed across time zones and geographical locations, and lack informal mechanisms for communication. They emphasize social media is regarded as a mechanism that can support informal and serendipitous interactions across the team, and as such can alleviate these challenges. Finally, they present ten research questions, of which the following three are applicable to this study:

- 1) Can social media play an effective role in supporting coordination and task articulation?
- 2) What kinds of social media would increase informal communication, the flow of knowledge and awareness across team and project members?
- 3) What are the drawbacks from increased transparency in team projects? Does this lead to privacy concerns?

There are several social media tools available to software engineers which facilitate coordination, communication with and learning from other users, being informed about new developments and creation of informal documentation [17]. These tools can be characterized by an underlying 'architecture of participation'; systems that are designed for user contribution [18]. Such a design supports the creation of collective value, often as an automatic byproduct of an individual activity. Wikis, blogs and microblogs are some well-known examples of such social media solutions.

In our research we aim to determine whether an environment in which one can both express himself and get a sense of how his team members feel is valuable to distributed Software Engineering teams. The specific environment we used, called WeHomer, allows users to exchange small elements of content such as sentences, images and hyperlinks, and their corresponding emotion. In fact, we use a microblogging solution extended with mood sharing functionality. Several research projects have been conducted in the field of Software Engineering to increase the understanding of how and why people use microblogging solutions. We will consider three

of these user studies, namely studies on Twitter [7], Yammer [8] and BlueTwit [9], to identify similarities and differences between these microblogging solutions and WeHomer.

Firstly, Twitter is a publicly available microblogging service with which users can publicly share messages, limited to 140 characters. They can also indicate whether their messages are public or private. When messages are indicated to be public they are accessible to all users of Twitter. However, when messages are indicated to be private they are only accessible to those users who have subscribed and are explicitly authorized to the user his feed. Zhao et al. [7] identified several characteristics in the use of Twitter. Important examples are: (i) frequent small updates of personal life events enable users to stay aware of people they do not encounter on a daily basis and (ii) subscribing to people you personally know and selected enables users to get trustworthy and useful information.

Secondly, Yammer¹ is very similar to Twitter, the main difference being that Twitter is publicly available while Yammer is enclosed by organizational boundaries. Other differences are the absence of a character limit, the possibility to create private and public groups and the opportunity to add attachments to messages. Zhang et al. [8] also identified several characteristics of the use of Yammer. They give an indication that users use Yammer more for publishing news about their groups or business units than for news about themselves. Subsequently, they indicated that Yammer was used to have long conversations and discussions. Finally, they found that Yammer enables users to stay aware about what others are working on and to make new connections.

BlueTwit is also very similar to Twitter, however it differs on two points: (i) BlueTwit is only available within organizational boundaries and (ii) BlueTwit has a character limit of 250 instead of the 140 character limit of Twitter [9]. Ehrlich et al. [9] discussed characteristics of the use of BlueTwit, It enables: (i) having internal conversations about confidential information, (ii) staying aware of what others are working on and (iii) enhancing your reputation.

All of the above user studies on the usage of different microblogging solutions mentioned an important side effect of microblogging in general: people feel more connected with each other. This is especially the case when people work distributed from their colleagues since microblogging kept them connected to other colleagues and the company, and alleviated the feeling of isolation. WeHomer differs from BlueTwit in the sense that it drops the character limit, enforces subscription to all users automatically (manageable because of the small community), and adds mood sharing functionality.

Garcia et al. [10] introduced 'Emotional Awareness' and argue that it enables users to become aware of the emotional state of their collaborators and act accordingly to achieve better results in their joint work. Other research mainly focuses on electronic meeting systems in which each participant explicitly specifies his mood and changes in the average mood are visualized to all participants [19], [20]. WeHomer integrates

¹<http://yammer.com>

both microblogging functionality and the opportunity to express your current mood into a single environment. This is the main differentiator of WeHomer in comparison with other microblogging solutions.

III. RESEARCH SITE AND METHOD

A. Research Site

This research is carried out at IHomer, a Dutch Software Engineering company founded in August of 2008 in which it is common practice to work from home. IHomer currently employs 20 people, working on a variety of products, projects and contracts. The largest team consists of 7 people working on related projects, but the overall group is very close with personnel moving between teams and teams exchanging projects as needed. Even though it is common practice in the company to work from home, the employees try to get together once a week on Tuesdays to meet face-to-face at an office to stay connected. Sometimes this can be difficult however, for example when someone is away on a contract and has other obligations on Tuesdays. The company has grown over the past years and initially on Tuesdays everyone discussed what they were doing. This worked well until the company size reached 16, and then sub teams were formed to keep this face-to-face communication more tractable. Teams cluster according to various factors: projects and related technologies being two of them.

People at IHomer aim to work together closely and stay a very connected community. One of the core strategies to stay connected is the weekly face-to-face meeting. As mentioned above this is not always feasible which can become a practical problem if people are unable to attend the majority of the Tuesday meet-ups for a prolonged period. In order to cope with these issues the WeHomer system was developed by an employee of the company (not an author of this paper) and deployed in October 2011. It is a platform on which IHomers can share information about their day with their colleagues in order to stay connected and increase awareness. Users can share information about a new topic, called an entry or respond to an existing topic, called a comment (commenting was not supported in the first three months of our data analysis period). Comments are shown in chronological order grouped together under the entry to which they correspond. We use the term post to refer to something that is either an entry or a comment. Posts cover such items as what you are doing right now, what you have done, what you are going to do, how you feel about something and random thoughts.

Associated with each post is a happiness score ranging from 0 (totally unhappy) through 100 (utter bliss) depicting how the user feels about this post. In the user interface of WeHomer (see figure 1) the happiness index can be selected by use of a slider bar which shows one of 5 discrete emoticons corresponding to the level that is selected by the user². The exact integer value of the happiness index is not derivable by end users.

²Happiness index ranges: ':-(' = [0,20), ':-(' = [20, 40), ':-]' = [40, 60], ':-)' = (60, 80], ':-D' = (80, 100]

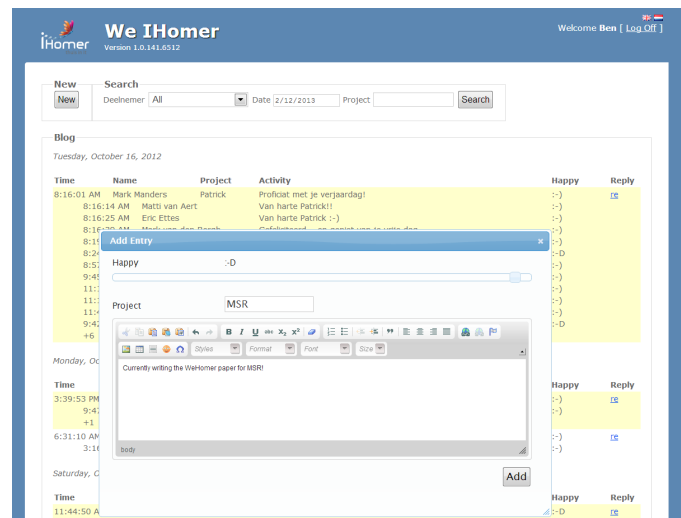


Fig. 1. WeHomer user interface

B. Method

The primary method of data collection we use is mining the WeHomer data between October 2011 and November 2012 by analyzing and subsequently coding the content. During this period there were a total of 1312 entries and 1189 comments. Because it is feasible to hand code each of these entries and comments sampling was unnecessary and we analyzed the content of all entries and comments. The coding was done by the first two authors and the coding set was arrived at in an iterative fashion. Firstly, a random (but consecutive) sample of 50 entries (including the corresponding comments) was selected and coded independently by both coders. Following this the two coders compared their codes and discussed their reasoning behind those codes. Based on this discussion the coders agreed on a joint coding set with which they independently coded another random (consecutive) sample of 25 entries (again including the corresponding comments) and discussed discrepancies in how they coded the sample. Based on this discussion they refined the coding set and did another iteration. After a total of three such iterations they decided the coding set was consistent between them and they could go ahead with the actual coding. To do this they divided the total data set in six ranges of approximately 200 entries and each coded three non-consecutive ranges.

Subsequently, based on what we found in the content analysis we conducted semi-structured interviews³ with five of the nineteen users in which we asked questions about what was unclear to us in the analysis and follow-up questions we had based on the analysis. To select which five people to interview we used purposive sampling in order to get an as complete view as possible. In the selection process we explicitly excluded authors of this paper. We did select the

³Interview structure can be found at: <http://aspic.nl/msr2013/docs/InterviewStructure.pdf>

original developer of the WeHomer system, the person with the highest number of posts, the one with the lowest number of posts and the two people with respectively the highest and lowest number of entries to number of posts ratios.

IV. DESCRIPTIVE STATISTICS

In this section we present information derived from the mined data to present the reader with an image of how the MBMI environment is used.

In figure 2 the weekly average number of entries, comments and posts is shown. In this figure it can be seen that there is significant and consistent use of the WeHomer system for the entire year we are investigating. Additionally we can see that at the start of the period there were no comments and considerably more entries than the remainder of the period. This is because at the start posting comments was not possible and people used entries to comment on another entry by referring to the entry they wished to comment on.

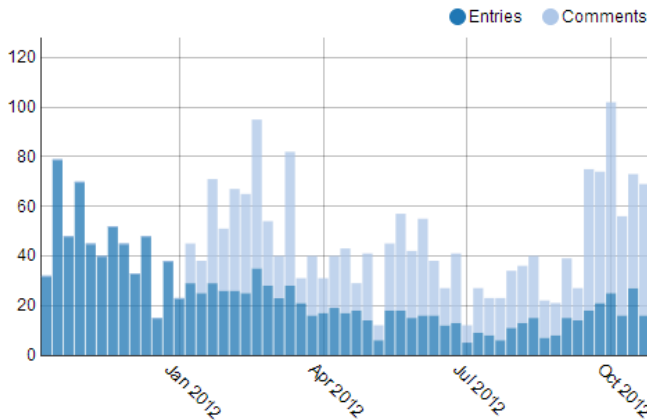


Fig. 2. The number of entries, comments and total posts per week

Subsequently, in table I we present the median length and the interquartile range of entries, comments and for entries and comments combined to give an indication of their respective lengths. So, for instance, we see that 75% of the entries are shorter than 143 characters.

TABLE I
AVERAGE LENGTH IN CHARACTERS

	Lower Quartile	Median	Upper Quartile
Entries	62	97	143
Comments	22	49	86
Posts	39	75	122

In figure 3 for the entire year the average happiness per week as well as the maximum and minimum happiness score are shown. In this figure it can be seen that the happiness fluctuates significantly and that in general the highest and lowest score for a week lie relatively far apart.




Fig. 3. The average, lowest and highest happiness score per week

Further, in WeHomer a default happiness score is automatically selected for each post a user makes, namely 70 on the range between 0 and 100. If a user doesn't manually select another happiness score this default score is used. Therefore it is interesting to investigate how often the users deviated from this default value. In table II we show separately for entries and comments how often this occurred. In this table we can see that a significant portion of the happiness scores were set at the default score. Therefore, we asked the five people we interviewed whether these values were chosen consciously. All the interviewees told us that even though they on occasion forgot to change the happiness score, in general they spent a minute to think which score to select, even if this is the default happiness score.

TABLE II
PERCENTAGE OF POSTS WITH THE DEFAULT HAPPINESS INDEX

	Default		Non-Default		Total	
Entries	791	60.3%	521	39.7%	1312	52.5%
Comments	931	78.3%	258	21.7%	1189	47.5%
Total	1722	68.9%	779	31.1%	2501	100%

Finally, in IHomer there exist 4 teams of people working on related projects. We compare the amount of directed communication between members of the same team and members of different teams to give an indication of how much collaboration was being done overarching the different teams. To do this we do the following: Firstly, we define commenting on an entry posted by a specific user as the utilization of a directed communication line between those two users. We do this even though all other users can see this communication because the communication is at least directed at that specific user. Subsequently we sum all of these utilizations of communication lines. An interactive visualization of the utilization of

communication lines between both team members and non-team members can be seen at <http://aspic.nl/msr2013/vis/sna>.

Following this we needed to compare the amount of communication within teams with the amount of communication outside of teams. We did this by doing the following for each team:

- 1) Count the total number of comments of the people in the current team on other people within the current team
- 2) Count the total number of comments of the people in the current team with the people outside the current team
- 3) Make the number found in step 1 relative to the team size (n) by dividing it by n*(n-1) (n*(n-1) to represent the total number of communication lines)
- 4) Make the number found in step 2 relative to the number of people outside of the team

A summary of these results is shown in table III:

TABLE III
UTILIZATION OF COMMUNICATION LINES WITHIN AND BETWEEN TEAMS

Team	Step 1	Step 2	Step 3	Step 4
Team 1 (5 people)	49	162	2.45	0.89
Team 2 (7 people)	386	194	9.19	1.47
Team 3 (5 people)	47	161	2.35	0.88
Team 4 (2 people)	24	165	12.00	0.61

In this table it can be seen that the utilization of communication lines within teams is considerably higher than communication lines crossing team boundaries.

V. FINDINGS

In this section we answer the research questions specified in section I. We structure the section based on the research questions, answering each research question in a separate subsection.

A. Topics

Research question 1 is: "What sort of topics are discussed in MBMI?" We answer this research question by first discussing the set of codings we used to code the data to give insight in the variety of topics discussed in the MBMI system. Subsequently we show the occurrence of each of the codings in both the entries and comments to give insight in the frequency each topic occurred. Finally we analyze these occurrences and make generalizations.

As described in section III we used an iterative bootstrapping process to construct the coding set. This coding set is structured in four major categories:

- 1) Nature
- 2) Form
- 3) Intention
- 4) Content

Each of these coding categories is further divided into sub-categories and actual codes. In appendix A we show this subdivision for each of the four major coding categories depicted as trees. In these trees leaf nodes depict actual codes

while non-leaf nodes depict sub-categories. In the appendix we also explain each of the codes and give examples.

Subsequently we applied all four coding categories⁴ to the entries and the first two coding categories to the comments. We only apply codes of the first two categories to the comments because comments are made in the context of an entry which makes it difficult to differentiate in how far the intention and content of a comment are dictated by the corresponding entry. We present the occurrence of each of the codes in the set of entries and the set of comments in figure 4 and figure 5 respectively.




Fig. 4. The frequencies of the codings for the entries




Fig. 5. The frequencies of the codings for the comments

⁴Note that the codes in the four categories are not mutually exclusive in their application to posts (e.g. a post can contain a positive and negative part)

In figure 4 it can be seen the most frequently occurring codes are 'Statement' (85.5%), 'Coordination' (73.9%), 'Positive' (58.5%), 'Work Planning' (54.9%) and 'Personal Information' (49.8%). In addition to this, four of the five people we interviewed indicated they have the tendency to post more positive things and to post more when they are in a positive mood. Further the consensus in the interviews about what they post is that it is everything they consider useful or interesting to their team members. One would suspect this to lead to a diverse list of topics to be shared on the medium and both the diversity of codes and the relatively distributed occurrence of these codes support this expectation. Finally, when asked to specify what they shared most, popular subjects mentioned are: personal information, project information with the intention to coordinate, technical information and prospects. This also corresponds with the actual data presented in figure 4.

Subsequently we discuss a deviation from the expected. An evident application for a MBMI system is asking questions. However, we found a relatively low amount of these. The total number of entries that are questions is 9.1% and the total number of comments that are questions is 7.5%. We compared these numbers to the results found in a study by Ehrlich et al. [9] on the investigation of the usage of Twitter and BlueTwit (an internal proprietary version of Twitter) in an organization for people that use both tools. To compare our number of questions to theirs we summarized the percentages for "Ask Question" and "Directed with Question" in their result set to yield a total number of questions of 6% for Twitter and 13% for BlueTwit (versus our 7.5% and 9.1% respectively). So, in the environment in our study relatively more questions were asked than in the Ehrlich setting for Twitter use, but less than with BlueTwit.

When asked about the amount of questions in the five interviews the respondents indicated that they asked a relatively low amount of questions on the MBMI system for two reasons. Firstly, the medium is asynchronous which makes it unpredictable when a question will be answered. Secondly, they indicated they usually knew who to contact (or at least knew who knew who to contact) and preferred contacting someone who would know the answer directly over asking it to the entire group. It was also discussed in the interviews the low amount of questions is likely to be specific to companies with a relatively small size, close personal connections and transparency between the team members because in such companies people will more easily know who knows what.

B. Impact on a Software Team

Research question 2 is: "What is the impact of the introduction of a MBMI on a software team?" The first main impact we found is that members of software teams feel more connected to each other when they are able to share activities and moods. We base this primarily on information from the interviews. Firstly, three out of the five people we interviewed explicitly reported feeling more connected to their colleagues since they were able to share their moods and activities within the team using the MBMI system. Additionally these interviewees also

reported being better able to understand their colleagues since using the environment and two of them reported they felt their colleagues understood them better as well. This finding corresponds to the results of the studies of Zhao et al. [7] and Ehrlich et al. [9] on microblogging in the work place. Zhao et al. [7] conclude: "Our results suggest that microblogging may help colleagues to know each other better as persons, that is in addition to professional relationships; this benefit is achieved by staying aware of small details about others' personal lives, interests, and current moods, which in turn creates more opportunities for exchanging acknowledgments and social support, generating new common ground, and creating and sustaining a feeling of connectedness."

Lesson Learned 1

Distributed software engineers feel more connected to each other when they are able to share activities and mood
--

As the second main finding, we found a MBMI makes information that is traditionally harder to consistently acquire more approachable and less volatile. This is based on both the content analysis and the interviews. In the content analysis we found that in particular the coding-categories entrepreneurial tasks (14.5% of all entries) and customer relations (9.3% of all entries) represent a considerable portion of the data. One of the interviewees indicated information about these types of activities is traditionally difficult to consistently gather. For instance information about "how to build a business" is often shared in face-to-face communication which makes it difficult to acquire at a later time (you will have to ask or try to recall) and the information is likely to be different from the original.

In the interviews people comment they consider it a strength of the system to be able to share non-time critical information: Information they would like to know about "eventually, but not necessarily within the next five minutes". Before WeHomer communicating this type of information was often postponed until a weekly face-to-face meeting or discarded altogether. One of the interviewees even said he found the system to offer benefits over meeting face-to-face on Tuesdays: "With WeHomer it is easier to stay current than by meeting people face-to-face on Tuesdays because then you don't get to talk to everyone".

Lesson Learned 2

A mood-activity environment makes information that is traditionally harder to consistently acquire more approachable and less volatile
--

Finally, we also found a MBMI system facilitates an unobtrusive way to express your personal feelings or thoughts to your colleagues. All of the interviewees mentioned they considered the low threshold the environment offered for sharing information with their colleagues an important strength. One of the interviewees said he felt he could "share knowledge and emotions like you are co-located". We can also see the environment offers a light-weight method to share personal information since over half of the entries (52%) contain personal information.

C. Impact on a Distributed Software Team

Research question 3 is: "What is the impact of distribution on the use of a MBMI?" Firstly, we found that people who work co-located with the majority of their team, share less activities and moods with those team members that are non-co-located. In our setting this behaviour presented itself as follows: while for the rest of the week the default work location is home, on Tuesdays people at IHomer try to work co-located at a central office as much as possible. However, at times this is unfeasible for specific team members, for instance due to being contracted at a customer location. In practice at least half of the team is present on Tuesdays the vast majority of the time, but it is rare for the entire team to be present. Therefore it is striking to see that on Tuesdays the number of entries and comments is significantly lower than on other days of the week (see figure 6). Finally, also in the interviews it was recognized that "on Tuesdays WeHomer is used very little". This is similar to what was reported in [21] on the deployment of a conversation overhearing tool in an industrial setting with a non-homogeneous geographical distribution.




Fig. 6. The number of entries and comments for each day of the week

Lesson Learned 3

In distributed Software Engineering teams, people who work co-located with the majority of their team, share less activities and moods with those team members that are non-co-located.

Paradoxically to what is discussed above, the interviewees indicated they do find it particularly valuable to share moods and activities with their distributed colleagues. One of the interviewees stated: "WeHomer is used very little when people work co-located on Tuesdays which makes things less transparent for people that cannot be there". The interviewees indicated they recognize the value in making sure the entire team stays connected, even when part of the team works co-located and part of the team works distributed. They recognize the value because they know from experience how difficult being the dislocated colleague can be. It is striking to see that even though they know it is important to help their distributed team members, they still struggle to do so.

Lesson Learned 4

In distributed Software Engineering teams, people who work co-located with the majority of their team, find it is particularly useful to share activities and moods with those team members that are non-co-located.

A useful insight on this is also shared by Fullerton of Stack Exchange in his blog post [22] about the lessons learned from three years of working in a distributed team. He states: "There's no halvesies in a distributed team. If even one person on the team is remote, every single person has to start communicating online. The locus of control and decision making must be outside of the office: no more dropping in to someone's office to chat, no more rounding people up to make a decision. All of that has to be done online even if the remote person isn't around. Otherwise you'll slowly choke off the remote person from any real input on decisions."

D. Impact of Team Composition on MBMI

Research question 4 is: "How does team composition impact collaboration with MBMIs?" Our main finding with respect to this research question concerns the regularity in which the members of a team use the MBMI: do all the members of the team use the MBMI system an equal amount of the time and for the same sort of topics? If this usage differs significantly between different team members the distribution of relevant information in the team will be unbalanced as well. This was the main challenge in the use of the MBMI system that came forward in the interviews. Examples of things interviewees said are: "The success of WeHomer is dependent on participation", "There is only a challenge for those not using it" and "If you don't participate you miss things". The challenge is threefold. Firstly, team members using the environment less than their colleagues run the risk of missing things. Secondly, a team member sharing information cannot be certain his team members know about this. Finally, since the MBMI system isn't used for everything, you cannot infer something did not happen because it is not available there. Therefore, to have a complete view of all valuable information about the mood and activities of team members software engineers need to consult other sources as well.

Lesson Learned 5

If the regularity with which team members utilize the mood-activity environment differs, the distribution of relevant information in the team will be unbalanced

Finally, in the interviews we talked about the type of teams for which they thought a MBMI system would be beneficial. Firstly, to use the system in the same (personal) fashion as it is being used at IHomer the teams need to be sufficiently involved. One of the interviewees said: "I need to know the people" while another said that "teams need to be homogeneous (shared interests, shared work)". They also mentioned that as a direct result of this, team size can become an issue but only if it leads to the teams members being less involved with each other.

On the other side of things the interviewees did consider a MBMI system beneficial to all companies. One interviewee said: "every organization needs a WeHomer because even a closed door is a barrier". They do believe however that the type of information that is being shared will be connected to the type of organization. For instance one of the interviewees explained that he considered the high amount of personal information and the diversity of the messages on WeHomer to be tied to the open character of IHomer and that he would expect a more traditional organization to share a larger portion of technical information instead.

VI. THREATS TO VALIDITY

Threats to external validity can exist at each of the levels of generalization in a study. In our study, a threat to external validity exists in the generalization of the single Software Engineering team to the population of all distributed Software Engineering teams. To be able to better generalize beyond the setting we performed the study in, the study should be repeated in other teams as well. With respect to the generalization of the sampled data to the population of IHomer our work is much less threatened. For the interviews we sampled 5 out of the 20 people in the team (25%) and for the content analysis we even coded 100% of the posts in WeHomer for the year of data we investigated.

Furthermore, there exist threats to construct validity in our study. Firstly, we attempted to mitigate threats to reliability by elaborately describing our research site and methods and making our coding set and interview design available. Next to this we also make all of our data available in anonymized form and make the tool available upon request. We do this to make both our data gathering methods and the analysis of our data, repeatable. Subsequently, a threat to construct validity is mono-operation bias. Because we only researched the application of MBMI environments with one specific tool, one could argue the results only apply to the use of that tool. The only mitigation we need to offer for this is the general nature of the tool itself. Basically any tool with which it is possible to share activities and moods in distributed teams will suffice and the WeHomer tool clearly fulfills these requirements. A final threat to construct validity in this study is that both the creation of the coding set and the coding of the posts was done by the first two authors who are also employees in the company at which the study was being performed. The advantage of this is that the researchers possess insight knowledge and can leverage this to code the data more accurately. A disadvantage is that the researchers might not be completely impartial due to their involvement in the setting. Overall, it is our opinion the advantages outweigh the disadvantages.

VII. CONCLUSIONS AND FUTURE WORK

The main contributions of this paper are the answers to the research questions. First, we answered what sort of topics are discussed in microblogging systems with mood indicators (MBMI) by presenting the nature, form, intention and content

of the posts in over one year of usage data and presenting the frequency at which these occurred. Based on this data we found that distributed software engineers primarily share positive posts with the intention to either coordinate or provide personal information. Furthermore, when compared to other corporate microblogging solutions we found a relatively low amount of questions.

Subsequently, we have shown there are two major impacts of the introduction of MBMI on a software team. Firstly, team members become more connected. The loss of teamness is a major and unresolved issue in the field of GSE and therefore advances in this area are significant. Secondly, the MBMI system made information that is traditionally harder to consistently access more approachable and less volatile

Further, on the impact of distribution of the software team on the use of MBMI, we found that the way people act when working co-located with the majority of their team is paradoxical to how they think they should act. On the one hand, people share less activities and moods with their distributed colleagues while on the other hand they do recognize the value in staying connected with the rest of the team. It is striking to see that even though they know it is important to help their distributed team members, they still struggle to do so.

With respect to the fourth research question, on the impact of team composition on collaboration with MBMI, we found that the distribution of relevant information in the team will be unbalanced if team members use the environment unequally.

Concerning future work we are particularly interested in researching the actual value of incorporating mood in microblogging systems. A way to accomplish this is to perform two user studies in similar software teams in which one of the teams receives access to a regular microblogging solution and the other team receives access to a system that is similar in every way except the addition of the possibility to share mood with team members.

REFERENCES

- [1] E. Carmel, *Global software teams: collaborating across borders and time zones*. Upper Saddle River: Prentice Hall PTR, 1999.
- [2] J. Herbsleb and D. Moitra, "Guest Editors' Introduction: Global Software Development," *IEEE Software*, vol. 18, no. 2, pp. 16–20, 2001.
- [3] J. Herbsleb, "Global Software Engineering: The Future of Socio-technical Coordination," in *Proceedings of the IEEE 2007 Workshop on the Future of Software Engineering*. IEEE Computer Society Press, 2007, pp. 188–198.
- [4] The Dieringer Research Group Inc., "Telework Trendlines 2009: A Survey Brief by WorldatWork," 2009.
- [5] G. Olson and J. Olson, "Distance matters," *Human-computer interaction*, vol. 15, no. 2, p. 139178, 2000.
- [6] T. Nguyen, T. Wolf, and D. Damian, "Global software development and delay: Does distance still matter?" in *Proceedings of the IEEE 2008 conference on Global Software Engineering*. IEEE, 2008, pp. 45–54.
- [7] D. Zhao and M. B. Rosson, "How and why people twitter: the role that micro-blogging plays in informal communication at work," in *Proceedings of the 2010 International Conference on Weblogs and Social Media*. ACM Press, 2009.
- [8] J. Zhang, Y. Qu, J. Cody, and Y. Wu, "A case study of micro-blogging in the enterprise: use, value, and related issues," in *Proceedings of the 28th international conference on Human factors in computing systems*. ACM, 2010, pp. 123–132.

[9] K. Ehrlich and N. S. Shami, "Microblogging inside and outside the workplace," in *Proceedings of the ACM 2009 international conference on Supporting group work*. The AAAI Press, 2010.

[10] O. García, J. Favela, and R. Machorro, "Emotional awareness in collaborative systems," in *String Processing and Information Retrieval Symposium, 1999 and International Workshop on Groupware*. IEEE, 1999, pp. 296–303.

[11] K. Schmidt, "The Problem with 'Awareness': Introductory Remarks on 'Awareness in CSCW,'" *Computer Supported Cooperative Work*, vol. 11, no. 3-4, pp. 285 – 298, 2002.

[12] A. Syri, "Tailoring cooperation support through mediators," in *Proceedings of the 1997 European Conference on Computer Supported Cooperative Work*. Kluwer Academic Publishers, 1997, pp. 157–172.

[13] P. Dourish and V. Bellotti, "Awareness and Coordination in Shared Workspaces," in *Proceedings of the ACM 1992 Conference on Computer Supported Cooperative Work*. ACM Press, 1992, pp. 107–114.

[14] J. Fogarty, S. Hudson, C. Atkeson, D. Avrahami, J. Forlizzi, S. Kiesler, J. Lee, and J. Yang, "Predicting human interruptibility with sensors," *ACM Transactions on Computer-Human Interaction*, vol. 12, no. 1, pp. 119–146, 2005.

[15] K. Dullemond, B. van Gameraen, and R. van Solingen, "Virtual open conversation spaces: Towards improved awareness in a GSE setting," in *Proceedings of the 2010 International Conference on Global Software Engineering*. IEEE Computer Society Press, 2010, pp. 247–256.

[16] S. A. Bly, S. R. Harrison, and S. Irwin, "Media spaces: bringing people together in a video, audio, and computing environment," *Commun. ACM*, vol. 36, no. 1, pp. 28–46, Jan. 1993. [Online]. Available: <http://doi.acm.org/10.1145/151233.151235>

[17] M.-A. Storey, C. Treude, A. van Deursen, and L.-T. Cheng, "The impact of social media on software engineering practices and tools," in *Proceedings of the FSE/SDP workshop on Future of software engineering research*. ACM, 2010, pp. 359–364.

[18] T. O'Reilly, "The architecture of participation," http://www.oreillynet.com/pub/a/oreilly/tim/articles/architecture_of_participation.html, 2004.

[19] S. Mora, V. Rivera-Pelayo, and L. Müller, "Supporting mood awareness in collaborative settings," in *Proceedings of the 2011 International Conference on Collaborative Computing: Networking, Applications and Worksharing*. IEEE, 2011, pp. 268–277.

[20] A. Fessel, V. Rivera-Pelayo, V. Pammer, and S. Braun, "Mood tracking in virtual meetings," in *Proceedings of the 2012 European Conference of Technology Enhanced Learning*. Springer, 2012, pp. 377–382.

[21] K. Dullemond and B. van Gameraen, "An industrial evaluation of technological support for overhearing conversations in global software engineering (in press)," in *Proceedings of the 2012 International Conference on Global Software Engineering*. IEEE Computer Society Press, 2012.

[22] D. Fullerton, "Why we (still) believe in working remotely," <http://blog.stackoverflow.com/2013/02/why-we-still-believe-in-working-remotely/>, 2013.

APPENDIX A: CODING SET

The used coding set is the following. Firstly we defined four major types of codes.

- 1) **Nature**
- 2) **Form**
- 3) **Intention**
- 4) **Content**

Each of these coding categories is further divided into sub-categories and actual codes. Below we will show this subdivision for each of the four major coding categories. In these trees leaf nodes depict actual codes while non-leaf nodes depict sub-categories. Below each figure we will explain the codes and give examples. Firstly for all entries and comments we coded the nature of the message (see figure 7). This depicts whether the content of the post can be considered **positive, negative or neutral**. So, for example a post stating a new

assignment has landed is positive while a post about a failed build or a sick family member is negative.




Fig. 7. The nature of an entry or comment

Secondly we also coded the **form** of the entries and comments (see figure 8). With the codings we mean the following:

- **Statement** - An assertion
 - **Answer** - Attempt to answer a question asked in a previous post
 - **Joke** - Something said or done to provoke laughter or cause amusement
 - **Compliment** - Expression of praise, commendation, or admiration
 - **Best Wishes** - Wishing something nice to someone (such as: "good luck" or "happy birthday")
 - **Standard** - A statement, but not an answer, joke, compliment or best wishes
- **Question** - Attempt to illicit an answer to some question




Fig. 8. The form of an entry or comment

We choose to only apply the final two coding categories (intention and content) to the entries and not the comments. We elect to do so because comments exist only in the context of the entry to which they belong. Because of this, comments are often quite brief and leave out much of this contextual information making it is infeasible for coders to consistently decide what intention and content is specifically applicable to that comment. The types of intention we distinguish are depicted in figure 9. With these codings we mean the following:

- **Sharing personal information** - Intention to share information that is about the poster and his/her personal life
- **Sharing work related information** - Intention to share information that is about the work of the poster

- **Coordination information** - When the intention is to coordinate with colleagues
- **Knowledge** - when the intention is to share factual knowledge
- **Social interaction** - Intention to follow social protocol for making and maintaining relationships with others




Fig. 9. The intention of an entry

Finally the content coding is shown in figure 10. With these codings we mean the following:

- **Information about a person**
 - **Health** - The poster's health
 - **Sentiment** - How the poster feels
 - **Personal Experience** - Information about an experience which is not primarily work-related (Travel, Family, Scenario)
- **Information about technology**
 - **Technical Knowledge** - Specialist information
- **Information about task articulation work** - Information about work which is done to support the core activities of IHomer (including infrastructure and planning)
 - **Work Planning** - When work will be done or is done
 - **Work Assignment** - Who will do certain work (Assignment, Expertise Finding)
 - **Supplies** - Information about supplies (Including for instance food)
 - **Non-Technical Infrastructure** - For instance the office or office equipment
 - **Technical Infrastructure Intern** - For instance IHomer's timesheet application
 - **Technical Infrastructure Extern** - For instance DNS, Skype, IDE and phonenumber
- **Information about customer relations** - Information about relations with customers and the process around this
 - **Relation** - Directly relating to the making and maintaining of the professional relationships with business relations
 - **Project Commissioning** - About transferring finished (or partially finished work) to the customer (including things such as training)
- **Information about entrepreneurial tasks** - Tasks directly related to the organization, operation and management of risk with respect to a business venture
 - **Prospects** - Opportunities for new work or projects
 - **Company Meeting**
 - **Applicants** - Hiring new people to work for IHomer
 - **Invoicing** - About actions needed to get paid by the customers (such as sending the actual bill)

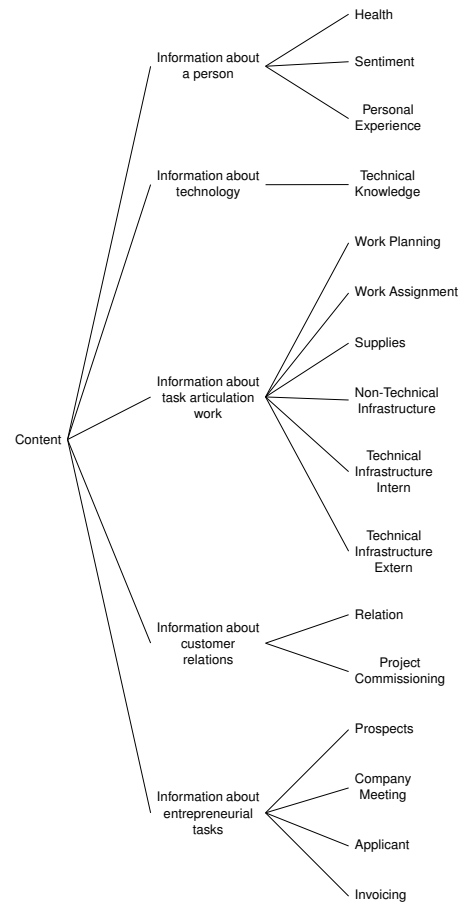


Fig. 10. The content of an entry