

Diversity among Enterprise Online Communities: Collaborating, Teaming, and Innovating through Social Media

Michael Muller*, Kate Ehrlich*, Tara Matthews†, Adam Perer*, Inbal Ronen‡, and Ido Guy‡

IBM Research

*One Rogers Street, Cambridge, MA, USA 02142 {michael_muller, katee, adam_perer}@us.ibm.com

†650 Harry Road, San Jose, CA, USA 95120, tlmatt@us.ibm.com

‡Haifa University, Mount Carmel, Haifa, Israel, {inbal, ido}@il.ibm.com

ABSTRACT

There is a growing body of research into the adoption and use of social software in enterprises. However, less is known about how groups, such as communities, use and appropriate these technologies, and the implications for community structures. In a study of 188 very active online enterprise communities, we found systematic differences in size, demographics and participation, aligned with differences in community types. Different types of communities differed in their appropriation of social software tools to create and use shared resources, and build relationships. We propose implications for design of community support features, services for potential community members, and organizations looking to derive value from online groups.

Author Keywords

Communities, Online communities, Teams, Virtual Teams, Idea Labs, Communities of practice

ACM Classification Keywords

H.5.3 Group and organizational interfaces/CSCW

General Terms

Design, Human factors

INTRODUCTION

There is a growing body of research into the adoption and use of social software in an enterprise including studies of blogs [10,15,16,46], microblogs [11,30,48], wikis [14,34], shared bookmarks [25], shared files [40,43,45] and social network sites [8,40], adding to earlier studies of email [32]. However, there is much less known about how these technologies are adopted and used by groups in general and by communities in particular [22,36,39].

About a decade ago, online communities engaged in knowledge- and information- sharing largely through discussions and conversations [44]. The advent of social soft-

ware augments traditional community technologies, adding blogs, wikis, and social file-sharing services that offer other forms of collaborative engagement and information sharing [10,11,14,25,34,37,41,43]. At the same time, social software tools have become increasingly meshed in a broader social media environment, so that a tool aimed to support enterprise communities may appear in the same software suite as a tool for virtual teams (e.g., [36,39]) and a tool for enterprise social networking (e.g., [8]). The new structures and use patterns add to the diversity that has already been studied among online communities (see Related Work, below). The new capabilities allow us to discover the types of collaborations that are currently taking place in tools aimed to support communities, and to explore new classifications that reflect these appropriations.

Although most studies of online communities have been conducted with public internet communities, studies of social software tools tell us that enterprises provide a different context for interaction [10,11,27,28,34,41]. With respect to communities, we suggest there are several critical differences. Firstly, an enterprise provides a shared context in addition to the context of the community, which can contribute to a level of trust and common ground [8,24]. Secondly, enterprise communities are likely to be business-focused, leading to different content and perhaps styles of discussion [24,36]. Thirdly, companies, which require authenticated access and use of real names, eliminate anonymity and provide greater transparency [8].

Our goal is to delineate different patterns of tool use as a way of understanding the different functions and practices of enterprise communities. In this paper, we report on five emergent structures and usage patterns when enterprise groups appropriate an online communities application. Specifically, the study described in this paper examined the patterns of use of social software in enterprise online communities across five different types of communities: Communities Of Practice (COPs), Teams, End-user technical support, Idea Labs, and Recreation. We will also consider the implications of this new diversity in structure and usage for the design of metrics and services.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies

bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI'12, May 5–10, 2012, Austin, Texas, USA.

RELATED RESEARCH

Online communities form in response to many human and organizational needs, including sharing information, pursuing passions, building skills, working toward shared goals, finding sociality, engaging in commerce, and curiosity [1,2,4,6,9,13,18,19,24,26,33,35,37,44]. Researchers from different disciplines have provided many ways of analyzing communities by attributes or by types.

Types of Communities

Lazar and Preece considered communities in terms of their social attributes, supporting technology, and relationship to physical communities [1]. Wenger considered more functional aspects, similar to the *attributes* of Lazar and Preece, such as communities of practice, formal work-groups, project teams, and informal networks [44]. Several researchers pursuing specific theoretical agendas have provided analyses that dealt with social network analysis variables [17], division of labor [5], and social qualities [18].

Several researchers described communities in terms of what members were seeking, such as knowledge platforms, networking with fellow enthusiasts, and finding solutions to problems [20]; or interest, relationship, fantasy, and transaction [13]; or educational opportunities, professional opportunities, or topical interests [6].

Business-oriented researchers focused on factors that shape the community, such as the source of initiative in business-oriented communities [35]; or types of sponsoring organization [7]. Marketing typologies have taken a more pluralistic position, organizing their analyses in terms of customer goals, business (provider) goals, governance models, and demographics [26,33]. Other business researchers have focused on technical aspects, such as discussion-based, task-based, virtual worlds, or “hybrid” [42].

The development of global virtual teams has also suggested new ways of organizing people and resources in enterprises [23,35], with unique aspects of coordination and time-frame that may affect the success of those teams [36,39].

Many of these reports relied on interviews, surveys, or inspections of internet community webpages. By contrast, our analysis is narrower but also deeper. We used data from a production “Communities” application in IBM, the internal deployment of IBM’s Connections product. We observed quantitatively how a workforce appropriated the Communities application in their daily work.

Social Software

Online communities have traditionally relied on communication and discussion technologies such as electronic bulletin boards, real-time chat, email message archives and perhaps members’ individual web pages. However, contemporary social software suites afford additional shared authoring via blogs and wikis and information sharing via shared bookmarks [25], shared files [31,43], and wikis [14].

Groups form unique combinations of these tools, depending on their goals, tasks, timeframes, and work dynamics [22].

“Landscape Analysis”: Community Usage and Patterns

In a preliminary, unpublished study, we used methods based on [3] to perform a “landscape analysis” of the use of the Communities application during April 2010. This analysis showed great diversity of community configurations, and many overlaps among the names and descriptions that organizers gave to their communities. One of many examples occurred around the concept of communities of practice: similar configurations carried names or labels such as “community of interest,” “center of excellence”, and “network”. We also became aware that work-group teams had begun to use the Communities application, sometimes in configurations of as many as 140 teams.

Employees and organizations were appropriating the Communities application in unanticipated ways. Similar results, across a more diverse set of enterprise social media, were provided in [23]. Our landscape analysis showed an enormous diversity in community types and characterizations.

Research Questions

The initial observations from our “landscape analysis” led us to believe that online enterprise communities required a different typology from the ones previously used. Thus, our first goal was to develop a typology. Secondly, the focus of social software on sharing information and developing relationships led us to ask the following research questions:

RQ1: Do different types of communities differ in their *Human Capital* – i.e., number of people, participation rate?

RQ2: Do different types of communities differ in their *Intellectual Capital* – i.e., what they share with each other, and which social software tools they use in that sharing?

RQ3: Do different types of communities differ in their *Relational Capital* – i.e., the relative contribution made by people, and the connections they make through the tools?

RESEARCH SETTING

The study took place in IBM, a global multinational company that had been using, and encouraging the use of a suite of social software tools called “IBM Connections Communities” that can include discussion forums, blogs, wikis, shared bookmarks, shared files, feeds and activities. Employees who participate in these Communities come from diverse corporate organizations, including Sales, Marketing, Consulting, Development (hardware and software), Support, and Research. Participation is in general voluntary, with two exceptions: (a) People in certain teams may have been “strongly encouraged” to participate in specific communities; (b) Some community *leaders* included their community work in their job responsibilities.

Community Roles

The Communities application differentiates participants in terms of “owners” and “members”. There is no upper or

lower bound on the number of owners or members in a community; in practice, some communities in the company-wide dataset have only one member (probably failed experiments), and a few communities have 20,000 or more members. Many communities have an associated “leadership team” of owners who may share responsibilities in initiating and managing the community.

Owners have additional access rights, including the ability to change the role of members and to add members or owners directly. They can also delete any entry, and customize the community's website.

Community Tools and Media

The default tools in a community are shared bookmarks, shared feeds, and shared discussion forums. Bookmarks of internet or intranet pages with descriptive tags can be shared with the rest of the community. Feeds are RSS feeds whose output is public. Discussion forums allow users to start a thread with a topic on which other users can respond. Blogs allow authorship of blog entries, on which other users can comment. Wikis allow collaborative editing of web pages. Files, such as documents and presentations, can be uploaded by a user, and shared with other members of the community. Activities [30] allow a more structured set of shared materials, including collaborative task management.

Initiating a Community

Any employee may create a community by clicking on a “Start a Community” button in the Communities application. The community creator specifies the name of the community and its level of privacy (public or private), and provides a description and optional tags. The Communities application provides no data field for a “community type”, although some owners may include “CoP”, “Network”, “Circle”, “Team”, or “Idea Lab” in the title and/or text description of their community or provide some indication of the function of the community in the description.

The content of all public communities, including the membership list, is visible to any employee, although only community members have permission to create content. The membership list includes members and owners under their real names, along with basic contact information and job function. Further information on any member is available from their general profile, which is public to all employees.

Study Timeframe

The Communities application has been in operation since April 2007. At the time of our study (July, 2011), there were approximately 25,000 communities with membership open to any employee, and just under 10,000 additional communities whose contents were visible to all employees, but which required the permission of an owner for membership (i.e., for privileges to create content in the community). There were also an estimated 25,000 fully private communities that were inaccessible to our analysis. Over 75% of employees had joined at least one community over the four years in which the service had operated.

DATA COLLECTION

Communities

To reduce the overwhelming diversity of community descriptions that we had encountered in the landscape analysis, we targeted the top 250 public communities as defined by the total number of contributions of any type during the 2 months prior to the data collection for our study. Since our focus was on the use and appropriation of technology by community members, we wanted communities that were mature enough to have reached a stable set of behaviors, and active enough to provide us with sufficient data for analysis and that we might ask people about their *recent* activities and experiences (see Self-Reported, below). We removed communities that could subsequently not be found (e.g., deletion or change-of-name) or which were duplicates ($n = 18$), or in which the content was generated by software rather than humans ($n = 16$). We also omitted communities in which there was no English content ($n = 17$), or for other technical reasons ($n = 11$), leaving us with a corpus of 188 communities.

In those 188 communities, 96,877 employees from 107 countries had joined at least one community; 31,159 people (32%) had joined more than one community in the dataset. Among the 96,877 members, 12,190 (13%) from 74 countries had made at least one contribution to at least one community; among the contributors, 2053 (17%) had made contributions to more than one community. The number of members per community ranged from 10 to 8750 people (mean = 828, median = 192). The duration of membership ranged from 1 to 1206 days (mean = 190, median = 88).

There were 1078 employees who served as owner in at least one community in our sample. The number of owners per community ranged from 1 to 49 (mean = 6.57). The number of communities per owner (in the role of owner) ranged from 1 to 10 (mean = 1.15).

Community Type

Consistent with our earlier landscape analysis, the 188 communities studied in this paper were created by diverse employees in many different organizations. Some appeared to be large communities of practice for employees with a common discipline, but who were embedded in teams whose majority of members had a different discipline (e.g., [24,27]). Other communities appeared to be large, multi-country teams with executive leadership and a specific mission, such as selling to financial customers. Other communities appeared to be focused on solving specific problems in a relatively short time-frame (e.g., [28]). Informal interviews had convinced us that different types of communities might have different needs. We therefore focused our investigation on a quantitative study of the diversity among communities, with a goal of understanding which *types* of communities were present. Future research will provide an analysis of success factors associated with types of communities.

Our earlier landscape analysis suggested several major categories of communities. We pursued those categories in our quantitative study. As part of a larger survey, we asked the community owners to choose from the following choices to describe their community: *Community of Interest*, *Community of Practice*, *Center of Excellence*, *Team*, *Idea Lab*, and *Recreation*. Surveys were sent to 1296 owners. We received 469 responses (36% response rate) representing 162 communities (86% of the communities).

Three researchers independently coded each community based on the owners' responses; we resolved the few disagreements through discussion. We created a combined category, COP (Community of Practice), because the owners seemed to have difficulty distinguishing between Community of Practice, Community of Interest and Center of Excellence. Additional analysis revealed that several communities were providing technical support (e.g., [47]). Although this community-type had not been offered as a choice in the survey, we added it to the coding protocol. In summary, we used the following community types:

- **COP.** A group of people with a common interest or practice who share information and/or network. Typical COPs in our study tended to focus on well-defined methods for notifying members of events and content. There was little evidence of “deliverables” (work products) or other team-like attributes.
- **Team.** Communities working on a shared goal for a particular client, project, or business function. Typical teams in our study tended to focus on particular clients, industries, or products, with well-defined “deliverables” and deadlines. (Teams have previously been thought of as an *alternative* structure to communities [36,39,42]. In the Communities application, teams adopted the technology and language of communities, so we include them as a *type of community* in our analysis.)
- **Technical Support.** Providing technical support for a particular technology. Typical Tech Support communities in our study tended to focus on the technology itself, rather than a product based on the technology, or a client who used the technology.
- **Idea Lab.** Communities in which members brainstorm around a set of questions or issues for a limited period of time, usually as part of a client engagement. Typical Idea Labs in our study tended to have a very tight focus in terms of topic (i.e., a specific client) and timeframe (24 to 72 hours in duration).
- **Recreation.** Communities devoted to recreational activities unrelated to work.

Three coders examined the owner reports and the community itself to assign a unique type to each community. Cronbach alpha measuring the level of agreement between

the 3 coders was 0.84 ($F = 11.53$, $p < 0.001$). Coders resolved the few cases of uncertainty by discussion.

Shared Resources

The data in each community were extracted by crawling the Communities database through a published Application Program Interface (API). This is a common method used for large-scale studies of social media usage within an organization [8,0,11,12,14,15,25,29,40]. For each community, we extracted

- a list of owner and member IDs
- a count of all bookmarks, feeds, and files
- a count of all forum entries (topics and responses)
- a count of each activity entry (summarized as a “root” object or as a response to that root)
- a count of each blog entry (posts vs. comments)
- a count of each wiki (original pages vs. revisions)

Each data item included a timestamp, and the ID of the person who created or modified the item. (From a privacy perspective, we note that the contents of all of these communities were visible to any employee, and we never associated any data result with personal information.)

Self-Reported Community Participation

Beyond the formal data (community type and contents), we were also interested in the experiences of the owners and members of the communities. As part of a larger survey, we asked a stratified sample of members to rate their Visit Frequency – i.e., how often did the respondent visit the community? (Scale of *0/never* to *5/daily*.)

1342 members completed the survey (19% response rate). However, 51 members did not provide the community name or the name could not be resolved, and therefore their responses could not be used. An additional 170 members reported never visiting the community they were responding to. We omitted these data, leaving a sample of 1121 responses for 156 communities (83% communities).

RESULTS

Our results center around the question of how different types of communities use technology and whether there are similar or different patterns of usage depending on the type of community (COP, Team, Technical support, Idea Lab, Recreation). For these business communities, we organized our analysis in terms of the business concepts of Human Capital, Intellectual Capital, and a newer concept, of Relational Capital.

Unless stated otherwise, all analyses were done using SPSS GLM, with LSD tests between levels of each effect ($p < .05$). Any results reported in this paper achieved at least that level of significance. To guard against non-normality, we confirmed our analyses with log-transformed data. The results were in substantial agreement.

Human Capital

Counting People

We use the term *Human Capital* to refer to the people in the communities. In this analytic category, we differentiate the roles and activities of members and owners within a community.

Table 1 shows the mean number of *members* and *owners* for each type of community. *Participation rate* represents the proportion of unique people, out of the total community membership (including both members and owners), who made at least one contribution to the community – i.e., it is a summary measure of the visible or “public” signs of participation. By contrast, *visit frequency* is people’s self-reports of how often they visited the community, including view participation – the invisible or “non-public” aspect of participation (“lurking”) [38]. We also include a calculation of the *age* of the community, which can be an important explanatory variable.

The communities differed with respect to the number of owners and members, with COPs and Tech Support communities the most populous; these results were significant ($p < .02$) in a simple ANOVA. However, the community differences for members became non-significant when we included age as a covariate. We considered age, because communities that are in existence for a longer period may have an opportunity to gain more members (“more days, more people”). By this age-corrected analysis, there are no differences in overall size of the communities. The number of owners, however, is generally not affected by age, because owners are usually assigned when the community is formed.

We considered three explanations for the large number of owners of COPs. First, this effect may reflect a greater effort required for running a COP. Alternatively, these findings may reflect the tendency for COPs to fall outside of the formal team structure of workplaces [24,44]. Leaders were, in effect, volunteering their time to keep the COPs vital, and may have deliberately shared the workload for the community in order to protect time and energy for their primary jobs. A third possibility is that COPs could have a more diverse membership, with a social need to include people with different backgrounds among the leadership. These hypotheses may become future research topics.

Visible and Invisible Participation

Communities also differed in terms of participation rate (percent of people making visible contributions) and in terms of Visit Frequency; these effects remained robust ($p < .01$), even with community age as a covariate. Idea Labs, which are primarily venues for quick brainstorming [28], have the highest participation (50%), but this is not significantly different from Recreation communities or Teams. Teams had significantly greater participation rates than COPs, but were not significantly different from Tech communities. Future research should explore the nature of

	COP	Team	Tech	IdeaLab	Rec
N.	91	73	16	5	3
Members ***	1154.82	416.47	1086.13	259.20	76.67
Owners ***	8.43	5.14	4.19	4.60	1.33
Part. Rate ***	21%	37%	31%	50%	44%
Visit Freq. ***	2.95	2.80	2.78	1.27	1.00
Age (days) ***	571.49	382.96	494.56	53.00	455.00

*** $p < .001$

Table 1. Human Capital results for each community type: Mean numbers of people and relevant ratios.

these Tech communities, which seem to have some attributes relating to both Teams and COPs.

Overall, these communities are large (average size = 828 people), well established (all except Idea Labs, have been in operation for over a year) and have a higher participation rate than the 10%, which has been reported for online communities (e.g., [21]). These data confirm that communities with a high level of activity are likely to spread the work of contributing over several people rather than concentrating in just a few people.

Now that we have a preliminary idea of how many people use each type of community, and with what degree of participation, it is time to look at the materials they share with one another in the communities.

Intellectual Capital

We use the term *Intellectual Capital* to refer to the structured and unstructured knowledge that is created in communities. Companies are often concerned to catalogue this knowledge for retrieval and re-use. We analyze the Intellectual Capital of a community as the mass of contributions that the owners and members share with one another. These contributions could take the form of bookmarks, feeds, forum discussions, blogs, files, activities, and wikis. We discovered that some wikis were automatically updated through RSS feeds, so we have excluded them from our analyses of *human* actions related to the shared contents of the communities.

Use of any of the types of shared objects in a community was optional. With all of these choices, which types of objects were used in each type of community?

Invoking the Tool (Resource) Types

Table 2 shows the percentage of communities in each type, which made use of each of the resource types. Differences within each row were tested via a 2x5 chi-squared test.

Some resources, like bookmarks, and forums, appear to have achieved a ceiling effect, in which nearly all of each

	COP	Team	Tech	IdeaLab	Rec.
Bookmarks	88%	89%	88%	80%	100%
Feeds*	67%	53%	50%	0%	33%
Forums	96%	92%	100%	100%	100%
Activities	28%	30%	25%	0%	33%
Blogs**	75%	58%	63%	0%	33%
Files***	93%	99%	63%	80%	100%
Wikis*	90%	88%	75%	100%	33%

* p<.05, ** p<.01, *** p<.001 – differences within each row

Table 2. Intellectual Capital results for each community type: Percentage of each community type that contained at least one of each type of shared resource.

community type used that resource. Forums were used by a larger percentage of authors than any other tool (Anova and least-significant difference, $p<.001$). Feeds were used in all community types except Idea Labs, which typically involve an active brainstorm of only 1-3 days. We assume that Idea Labs are not used long enough to warrant a feed, whose value tends to grow over time.

Blogs had the same pattern as feeds: zero use in Idea Labs, and perhaps for the same reason. The files feature was used in most of the communities, with the partial exception of the Tech communities. As will be seen, the Tech communities put most of their effort into the forums, so perhaps files were too cumbersome for their relatively light-contribution-cost style of collaboration. Finally, although wikis were used in most of the communities in all types except Recreation, many communities populated their wikis via an automated mechanism. We refrain from over-interpreting the wikis results until a future time when we understand communities' use of the automatic features better.

These results show that communities made use of the rich features of contemporary social media, going well beyond the traditional discussion forum patterns of past communities. The next questions are about the extent of usage of these resources within each type of community.

Using the Tools (Resources)

Table 3 shows the contributions per resource type, in each type of community. As a first analysis, we have counted user-created entries at the lowest level. Thus, we count each forum topic, and each forum response, and we report the total. Similarly, we count each blog post and each blog comment, and we report the total. We do the same for the tree-structured (thread-structured) contents of the activities. For wikis, we count each original page, and each revision, and report the total. Of course, different metrics schemes are possible, such that a discussion topic would be reported as a single data item, no matter how many responses accrued to it. We hope to explore and compare these different ways of counting social resources in a future paper.

Because these data are *usage-sensitive*, they provide a somewhat different picture from the preceding data, which were based on whether a resource type was used at all.

Unlike the Human Capital analyses, the analyses of the number of shared resources were complex to think about. From the perspective of an organization, one of the important aspects of a community of any type is its productivity – i.e., *how many of each type of resource does each type of community contain?* We provide a summary of the number of each type of resource, for each type of community, in Table 3A.

However, from the perspective of a researcher, a more important question is to compare the shared resources in proportion to the *opportunity for those resources to be produced* – i.e., are there differences in the creation of certain type of resources among the community types? We address this question via a summary of the number of each type of resource in each type of community, *divided by the number of people in the community* (i.e., where each person represents an “opportunity” to contribute a resource), in Table 3B. Because many people were non-contributors (“lurkers”), these results tended to be relatively small numbers.

For the organization-oriented analyses, we conducted a separate analysis of variance for the data in each row of Table 3A. We also conducted analyses of covariance, using the age of the community and the number of people in the community. Forums and files were the principal differentiators of community types. There was significantly more forum activity among Tech and Recreation communities than COPs, Teams, and Idea Labs. Among files, there were more items in COPs and Teams than among Idea Labs and Recreational communities.

These results make it clear that different types of communities steward different configurations of shared resources. An organization might want to consider what kinds of outcomes it hopes to “harvest” from its communities, and it might want to select the type(s) of communities to initiate based on those “harvesting” goals.

A researcher may be more interested in the number of resources weighted by the opportunity to create a resource. Each person in a community represents an “opportunity” to create a resource in that community. Therefore, the data in Table 3B are expressed as resources/persons.

The results are rather different from the organizational perspective of Table 3A. Forums are again a significant differentiator, but in the population-weighted analysis of Table 3B, the Recreation communities produce far more forum items *per person* than the other types of communities. Wikis were also a differentiator in this representation; however, as we noted above, some of the wikis were populated by an automated process, so we will not attempt to interpret these results until we have more fully investigated that process and its rationale.

A	COP	Team	Tech	IdeaLab	Rec.
Bookmarks	75.56	34.03	37.19	4.80	5.00
Feeds	4.21	2.12	1.94	0.00	0.67
Forums**	313.26	257.34	859.13	113.00	817.33
Activities	31.74	34.30	43.25	0.00	0.33
Blogs	76.82	48.12	52.81	0.00	1.67
Files***	39.80	39.71	21.00	0.80	5.67
Wikis	133.37	140.49	102.50	6.00	0.67
Total ^a	541.40	415.63	1015.31	118.60	830.67

B	COP	Team	Tech	IdeaLab	Rec.
Bookmarks	0.275	0.425	0.078	0.102	0.981
Feeds	0.017	0.023	0.007	0	0.011
Forums***	0.781	2.202	1.803	1.448	13.412
Activities	0.152	0.387	0.151	0	0.003
Blogs	0.309	0.612	0.114	0	0.028
Files	0.302	0.567	0.071	0.017	0.094
Wikis***	0.694	1.649	0.244	0.100	0.011
Total ^{a***}	1.837	4.034	2.225	1.566	13.630

* p<.05, ** p<.01, *** p<.001 – differences within each row
^a Excludes Wikis

Table 3. Intellectual Capital results for each community type. (A) From a researcher’s perspective: Mean contributions divided by the number of people in each community. (B) From the organization’s perspective: Mean contributions per tool per community type.

Again, we have found different patterns of production of shared resources in different types of communities. We also note that there was a nearly significant difference for files (weighted by community size), and we speculate that, in a larger sample, Teams might have had significantly higher scores in the per-person production of files.

In summary, different community types are associated with different patterns of resource type utilization (Table 2), different patterns of shared resource production (Table 3A), and different per-person patterns of shared resource production (Table 3B). The different community types appear to use different types of collective resources.

Relational Capital

We propose the term *Relational Capital* as a measure of the opportunity for one person to interact with another person through the resources that we described in the preceding subsection, potentially leading to stronger social ties. We are studying several different algorithms for measuring interaction opportunities in social media (e.g., [29]); Relational Capital is our candidate metric for communities. For this paper, we will use a method based on combinatorics within a discussion thread (or equivalent structure). Suppose UserA creates a new forum topic, and UserB writes a response, and UserC writes a response to UserB, and UserA also writes a response to UserB. We now have a thread containing items written by a total of

	COP	Team	Tech	IdeaLab	Rec.
Forums***	972.35	466.24	4027.44	561.20	1838.33
Activities	222.26	132.78	126.00		1.00
Blogs	27.20	36.91	15.25		0
Wikis	77.52	78.09	323.83	0	0
Total*	1299.33	714.20	4492.52	561.2	1839.33

* p<.05, ** p<.01, *** p<.001 – differences within each row

Table 4. Relational Capital for each community type, for the types of shared resources that support a thread-like structure of responses. Each entry is a count of the pairwise opportunities for one person to interact with a colleague.

three people. We assume that each of the three people has interacted (responded to) the others, and thus we can compute a combinatoric estimate of “interaction opportunities” for each unique pair among the k users among three people of $3 \times 2 / 2 = 3$ interaction opportunities (i.e., using the combinatoric of $k [k - 1] / 2$).

We analyzed Relational Capital in forums (above), blogs (through blog-post and blog-comment), activities (through their node-and-leaf structure [30]), and wikis (through their linear revision sequences). For each of these resources, we identified the root item for each thread; we found the associated items within the thread; we counted the number of unique authors of those items; and we calculated the per-thread relational capital based on the above formula.

Table 4 presents a summary of the calculated interaction opportunities. While each of the four media potentially supports the creation of relationships through post-and-response, these community types carry most of their sociality in the forums. The Tech communities show the greatest development of pairwise connections in the forums, significantly more than COPs or teams.

This result appears to be similar to the Intellectual Capital results. However, that result was based on the *number of items* of each type, whereas the Relational Capital results were based on the *relationships of the persons* who had created those resources. In principle, it would be possible for a single person to create a large number of resources within a single thread (high *Intellectual Capital*). However, if no one else created a resource in the same thread, then the Relational Capital would be very low.

DISCUSSION

Out of the 35,000 public and invitational communities in an enterprise Communities application, we selected 188 of the most active communities for detailed study. Our initial landscape analysis had shown us that the Communities application was being used in diverse ways. For a first in-depth analysis, we pre-selected four types of communities for comparison. In the course of coding the 188 communities, we discovered a fifth type of community, and added that to our classification, resulting in categories of

Communities of Practice, Idea Labs, Recreational Communities, Teams, and Technical Support communities.

We examined the five types of communities for differences, using three conceptual lenses: Human Capital (people and participation), Intellectual Capital (shared resources), and Relational Capital (interaction opportunity). We found highly significant differences through each of those conceptual lenses.

Each of the three lenses provides its own distinct emphasis. The view based on Human Capital showed significantly more owners in COPs than other communities, and a trend toward more members as well. Despite the smaller number of people in Idea Labs, their intense time-constrained brainstorming process produced the highest participation rates. And yet the COPs, Teams, and Tech communities had the highest overall visit frequency.

The view based on Intellectual Capital showed differential patterns of feature use, with a surprisingly high productivity of forum entries in the Tech communities. We also noted different patterns depending on whether we took an organizational view (emphasizing the resource available to the company) vs. a research view (emphasizing the contributions normalized for community size and age).

Finally, the view based on Relational Capital showed differential accumulation of social connections among employees in different types of communities.

These results provide new ways of thinking about typologies of online enterprise communities. Previous research categorized communities in terms of broad distinctions, such as networking vs. knowledge-sharing vs. solution-finding [20] or relationships vs. “fantasies” vs. transactions [13,35] or types of opportunities sought by members [6] or by community sponsors [7,26,33]. Our analyses show specific categories of enterprise activities (similar to [44]), and suggest new features to support those categories (see Implications for Design, below). Other typologies focused on the technological features in each category of community [42]. Our analyses show that, even with the same technologies available in all communities, owners and members will make differential use of those resources to achieve strikingly different organizational forms and outcomes.

Limitations

Our study involved the technology-informed workforce at IBM, which has been an early adopter of enterprise social software. Other companies, such as Microsoft, MITRE, and Yammer, are also pursuing enterprise social software, and the space of commercial products in this area is growing. We therefore think that our results will become increasingly useful to other companies, and the growing body of researchers who study social software in organizations.

Our selection strategy focused on the most active communities. This may have biased our sample toward

more successful communities. We plan to test our results on new samples based on the largest communities and the longest-active communities.

We have used the word “communities” because the users themselves chose that word – perhaps because it was suggested by the name of the application. We note that IBM Connections Communities provides a richer feature set than earlier discussion-only communities, and thus may eventually require a new analytic category.

Implications for Design

Community Metrics. We found that different types of communities comprise different numbers of people; seem to require different extents of effort from their leaders; make use of different resources; and support different configurations of relationships (see also [24]). These results suggest that future research into success metrics for communities should analyze each community type separately.

Community Design. It may be useful to develop templates for different types of communities, so that a new community can achieve a faster launch by assembling the right configuration of resources, with appropriate goals for that type of community. This kind of support could be important for community-founders who are not technically savvy, and/or who want to accomplish a particular work goal, rather than to explore and configure a technology. Our preliminary landscape analysis showed that there was much more diversity than the five categories that we studied in this paper. Therefore, the community-templating service should be flexible and customizable, to allow users to appropriate the Community technology to innovate new genres. The Idea Labs provide a strong example: These communities made novel use the discussion forums, transforming the Communities application from the intended use for long-term, informal exchanges in a COP, into a brief-duration, highly-focused problem-solving environment [28] – without changing a single line of source code.

Community Recommendation Service. Different types of communities offer different benefits [2,6, 13,20,24]. People looking for opportunities to contribute (or for resources to “consume”) may need advice about how to choose communities to join. The different qualities of the different community types could become features for a community-recommendation service similar to [12], which could be valuable for new employees or employees with new job responsibilities.

Organizational Design. Organizations may want to use the information in this paper to choose the community types that they need to achieve certain objectives (e.g., [7,35,44]). An organization that values file-based resources, for example, might foster Teams and COPs. By contrast, an organization that wanted to create discussions or greater connections among its workers, might focus its resources on creating Tech support and Recreation communities. Organ-

izations concerned about low-sized, under-performing communities might use our three families of metrics to find highly-similar communities that could be combined to reach a form of critical mass for community vitality.

Implications for Theory

As social applications become more important for getting work done, there is a need to understand the abstract qualities of different types of social applications, and the different configurations of those applications that can support business, governmental, non-profit, and civic organizations. Earlier work to distinguish types of communities used interview, survey, or inspection data. Different analyses have focused on business needs [1,7,26,35], human strivings [19,31,37,38,44], social motivations [2, 20], and technologies [41].

We add to that literature with a more quantitative investigation that begins to integrate several of the conceptual domains of the former analyses, combining human action (Human Capital), business needs as expressed through shared resources (Intellectual Capital), and social dynamics (Relational Capital). Our work can contribute precision regarding types of communities, which can be used in conjunction with broader models of Collaboration Personas [22] and other explorations of the organizational adoption of social media [10,11,14,16,34,38,43,45]. Stronger methods for classifying communities can clarify agendas for determining organizational benefits (e.g., [29]), and crucial attributes of successful communities.

Our work points toward new investigations into large-scale enterprise COPs and enterprise virtual teams [36,39]. Previous work suggested that COPs could become vibrant information exchanges among a large percentage of members of a “minority discipline,” using an online venue for sharing aspects of their practice [2,24,44]. Comparing COPs and Teams, our results suggest a less participative pattern in COPs, in which a small percentage of people make the COP contributions. We had previously thought of COPs as diffuse networks of co-contributing practitioners [24,27], but our results suggest a small active core and a larger group of non-public participants [38].

We were surprised that Teams made such strong use of the Communities application. The IBM Connections product contains an “Activities” application that was intended for collaborative task management [30], and that permits a more formal structuring of a group’s shared resources. We would have expected virtual teams to prefer the task-focused, knowledge-curatorial affordances of Activities. Yet there were no significant differences in the use of the Activities features within Communities. Furthermore, we found that virtual teams had a higher participation rate in the Communities application than COPs, for whom it had been designed. We do not yet understand why Teams chose the Communities application. Our findings may show that, in enterprises, the distinction of teams vs. communities [36,

39] may be blurring. Perhaps the less formal *social* affordances of Communities were more important than the more formal *structured* capabilities of Activities (see [39] for review of social needs in virtual teams). Future research should re-examine concepts of virtual teams and communities in the context of the new enterprise social media.

Summary and Contributions

We have shown three lenses through which to analyze the differences among social software entities in organizations, and we have applied those lenses to show systematic differences among active, productive groups in an enterprise online communities application. Our work may inform the design of online communities, the design of services for people seeking online communities, and the design of organizations. We hope that our work will contribute to on-going efforts to describe and model the organizational use of online communities, and of collections of social software for use by collaborative groups.

ACKNOWLEDGMENT

The work of some of the authors was partially supported by the EU-FP7 project ROBUST (grant no. 257859).

REFERENCES

1. Armstrong, A., & Hagel III, J., Real profits from virtual communities. *Mckinsey Quart*, 3, 126-141, 1995.
2. Bogdan, C., & Bowers, J., Tuning in: Challenging design for communities through a field study of radio amateurs, *Proc Communities & Technologies 2007*, Springer, 439-461.
3. Bos, N., Zimmerman, A., Olson, J., Yew, J., Yerkie, J., Dahl, E., & Olson, G., From shared databases to communities of practice: A taxonomy of collaboratories, *JCMC 12(2)*, 652-672, 2007.
4. Bressler, S. E., & Grantham, C. E., *Communities of commerce: Building Internet business communities to accelerate growth, minimize risk, and increase customer loyalty*. McGraw-Hill, 2000.
5. Butler, B., Sproull, L., Kiesler, S., & Kraut, R., Community effort in online groups: Who does the work and why? In Weisband, S., & Atwater, L., (eds.), *Leadership at a distance: Research in technologically supported work*, Earlbaum, 2008.
6. Carlén, U., Typology of online learning communities, *Proc. Int. Conf. Net Learning 2002*.
7. Denison, T., Hardy, G., Johanson, G, Stillman, L., & Schauder, D., Community networks: Identities, taxonomies and evaluations, *Proc Elec. Netwk. 2002*.
8. DiMicco, J., Millen, D.R., Geyer, W.G., Dugan, C., Brownholtz, B., & Muller, M. (2008), Motivations for social networking at work, *Proc. CSCW 2008*.
9. Dubé, L., Bourhis, A., & Jacob, R., Towards a typology of virtual communities of practice, *Int. J. Info. Knowl. Mgmt. 1*, 69-93, 2006.
10. Efimova, L., & Grudin, J., Crossing boundaries: A case study employee blogging, *Proc. HICSS 2007*.
11. Ehrlich, K. & Shami, N.S. 2010. Microblogging inside and outside the workplace. *Proc. ICWSM '10*.

12. Freyne, J., Jacovi, M., Guy, I., & Geyer, W., Increasing engagement through early recommender intervention, *Proc. RECSYS 2009*.
13. Hagel, J., & Armstrong, A., *Net Gain: Expanding markets through virtual communities*, Harvard Business School Press, 1997.
14. Holtzblatt, L.J., Damianos, L.E., & Weiss, D., Factors impeding wiki use in the enterprise: A case study, *Proc. CHI 2010*.
15. Huh, J., Jones, L., Erickson, T., Kellogg, W.A., Bellamy, R.K.E., & Thomas, J.C., BlogCentral: the role of internal blogs at work, *CHI 2007 Extended Abstracts*.
16. Jackson, A., Yates, J., and Orlikowski, W. (2006). Corporate Blogging: Building community through persistent digital talk. *Proceedings of HICSS*.
17. Jones, Q., & Rafaeli, S., Time to split, virtually: 'Discourse architecture' and 'community building' create vibrant virtual publics. *Elec. Mrkts*, 10(4), 214-223, 2000.
18. Komito, L., The net as a foraging society: Flexible communities. *Info. Soc.* 14, 97-106, 1998.
19. Lazar, J., & Preece, J., Classification schema for online communities, *Proc. AISAC 1998*.
20. Liu, C.-C., Identifying the value types of virtual communities based on the Q method, *Int. J. Web Based Commun.* 7(1), 2011.
21. Ludford, P.J., Cosley, D., Frankowski, D., & Terveen, L., Think different: Increasing online community participation using uniqueness and group dissimilarity, *Proc. CHI 2004*.
22. Matthews, T., Whittaker, S., Moran, T., Yuen, S. (2011). Collaboration personas: A new approach to designing workplace collaboration tools. *Proc. CHI 2011*.
23. Matthews, T., Whittaker, S., Moran, T., Yuen, S., Judge, T. (2011), Productive interrelationships between collaborative groups ease the challenges of dynamic and multi-teaming, *CSCW Journal Special Issue on Collective Intelligence in Organizations*. DOI: 10.1007/s10606-011-9154-y.
24. McDermott, R. (1999). Learning across teams: The role of communities of practice in team organizations. *Knowledge Management Review*, May/June 1999.
25. Millen, D.R., Feinberg, J., & Kerr, B., Social bookmarking in the enterprise, *Proc. CHI 2006*.
26. Mittilä, T. and Mäntymäki, M., Classification of Online Communities. *Proc. Frontiers of e-Business Research 2003*. <http://www.ebrc.info>.
27. Muller, M., & Carey, K., Design as a minority discipline in a software company: toward requirements for a community of practice, *Proc. CHI 2002*.
28. Muller, M., & Chua, S., Brainstorming for Japan: Rapid distributed global collaboration for disaster response, *Proc. CHI 2012*.
29. Muller, M.J., Freyne, J., Dugan, C., Millen, D.R., & Thom-Santelli, J., Return On Contribution (ROC): A metric for enterprise social software. *Proc. ECSCW 2009*.
30. Muller, M.J., Geyer, W., Brownholtz, B., Wilcox, E., & Millen, D.R., One-hundred days in an activity-centric collaboration environment based on shared objects, *Proc. CHI 2004*.
31. Naaman, M., Boase, J., & Lai, C.-H., Is it really about me? Message content in social awareness streams, *Proc. CSCW 2010*.
32. Orlikowski, W., Learning from Notes: Organizational issues in groupware implementation, *Info. Soc.* 9(3), 237-250, 1993.
33. Plant, R., Online communities, *Tech. Soc.* 26, 51-65, 2004.
34. Poole, E.S., Grudin, J., A taxonomy of wiki genres in enterprise settings, *Proc. WikiSym 2010*.
35. Porter, C.E., A typology of virtual communities: A multi-disciplinary foundation for future research, *JCMC* 10(1), 2004.
36. Powell, A., Piccoli, G., & Ives, B., Virtual teams: A review of current literature and directions for future research, *DATABASE Adv. Info. Sys.* 35(1), 6-36, 2004.
37. Preece, J., *Online communities: Designing usability, supporting sociability*, Wiley, 2000.
38. Preece, J. and Maloney-Krichmar, D., Online Communities. In J. Jacko and A. Sears, A. (Eds.) *Human-Computer Interaction Handbook*, Erlbaum, 596-620, 2003.
39. Sarker, S., & Sahey, S., Understanding virtual team development: An interpretive study, *JASIS* 3, 247-285, 2002.
40. Shami, N.S., Muller, M., & Millen, D.R., Browse and discover: Social file sharing in the enterprise, *Proc. CSCW 2011*.
41. Skeels, M.M., & Grudin, J., When social networks cross boundaries: A case study of workplace use of Facebook and linkedin, *Proc. GROUP 2009*.
42. Stanoevska-Slabeva, K., & Schmid, B.F., A typology of online communities and community supporting platforms, *Proc HICSS 2001*.
43. Vaida, S., Edwards, W.K., Newman, M.W., Grinter, R.E., & Ducheneaut, N., 'Share and share alike: Exploring the user interface affordances of file sharing, *Proc CHI 2006*.
44. Wenger, E., McDermott, R., & Snyder, W., *Cultivating communities of practice: a guide to managing knowledge*. Harvard Business School Press, 2002.
45. Whalen, T., Toms, E.G., & Blustein, J. (2008). Information displays for managing shared files, *Proc CHIMIT 2008*.
46. Yardi, S., Golder, S.A., Brzozowski, M., Blogging at work and the corporate attention economy, *Proc. CHI 2009*.
47. Yardi, S., & Poole, E.S., Please help!: Patterns of personalization in an online tech support board, *Proc. C&T 2009*.
48. Zhang, J., Qu, Y., Cody, J., & Wu, Y. 2010. A case study of micro-blogging in the enterprise: use, value, and related issues. *Proc. CHI 2010*.