What's *Your* Idea? A Case Study of a Grassroots Innovation Pipeline within a Large Software Company

Brian P. Bailey[1,2]
Department of Computer Science[1]
University of Illinois
Urbana, IL 61801
bpbailey@illinois.edu

Eric Horvitz[2]
Microsoft Research[2]
One Microsoft Way
Redmond, WA 98052
horvitz@microsoft.com

ABSTRACT
Establishing a grassroots innovation pipeline has come to the fore as strategy for nurturing innovation within large organizations. A key element of such pipelines is the use of an *idea management system* that enables and encourages community ideation on defined business problems. The value of these systems can be highly sensitive to design choices, as different designs may influence participation. We report the results of a case study examining the use of one particular idea management system and pipeline. We analyzed the content, interaction, and participation from three creativity challenges organized via the pipeline and conducted interviews with users to uncover motivations for participating and perceptions of the outcomes. Additional interviews were conducted with senior managers to learn about the objectives, successes, and unique nature of the pipeline. From the results, we formulate recommendations for improving the design of idea management systems and execution of the pipelines within organizations.

Author Keywords
Creativity, Idea Management, Innovation, Organizations.

ACM Classification Keywords
H.5.3 [Information Interface and Presentation]: Group and Organization Interfaces -- Evaluation/methodology.

General Terms
Design, Human Factors.

INTRODUCTION
While innovation may arise from unplanned contributions, most innovation results from the use of *intentional* methods within organizations [1]. Borrowing a response from a participant in our study, by *innovation*, we mean "the conversion of a good idea into competitive differentiation." Existing methods for promoting innovation within large organizations include establishing R&D labs, forming internal incubation teams, and allowing employees to allocate part of their work time to side projects. Despite the relative merits of each method, competition is driving large organizations to experiment with new strategies [1].

Figure 1: Screen shot of the idea management system studied. Authors and content of the ideas are intentionally blurred.

One emerging strategy is the establishment of a *grassroots innovation pipeline*. Grassroots innovation is a particular type of innovation where the ideas flow in a bottom-up manner, i.e., the ideas are generated by those who are least likely to have access to the resources to make them happen. This type of pipeline is generally executed in four phases; (1) pose a challenging business problem to the corporate community, (2) foster community ideation, (3) filter and refine the best ideas, and (4) launch or integrate the ideas into a product or product-centric pathway. The pipeline is motivated by the growing recognition that the collective wisdom of the corporate community is the best resource for innovation [20]. This type of pipeline has been executed in Microsoft, IBM, Dell, Whirlpool, and UBS, among others.

The pipeline is interesting for HCI because community ideation may be organized, supported, and tracked via an *idea management system*. This type of system supports submitting, discussing, scoring, and disseminating ideas, among other functions. Design choices may therefore affect ideation outcomes and perceptions and adoption within the organization. Design choices may be informed by research on brainstorming [6, 11, 19] and social software [3, 14, 17], but the scale, openness, and context of this class of system make it unique. A recent study documented the process and outcomes of IBM's innovation pipeline [2], but did not examine any aspect of its idea management system.

CHI 2010, April 10–15, 2010, Atlanta, Georgia, USA.

In this paper, we report the results of a study examining the use of one specific idea management system (see Figure 1) and pipeline executed within Microsoft Corporation, a large software company. The pipeline is still maturing and represents one of several explorations within Microsoft of systems, programs, and organizational structures aimed at enhancing idea capture and innovation. In our study, we analyzed the content, interaction data, and user participation from three creativity challenges organized via the pipeline.

Our data analysis was complemented by interviews (N=24) with recent and top contributors and challenge winners to uncover motivations for participating and perceptions of the outcomes. Additional interviews (N=8) were conducted with senior managers, business unit leaders who sponsored the challenges, and system designers to learn about the objectives, successes, and measures of the pipeline.

From all of the results, we formulate recommendations for improving the design of idea management systems and execution of the pipelines. We believe others can leverage these recommendations, along with the data and insights reported in this paper, to better anticipate and plan for their own implementation of a grassroots innovation pipeline.

Related Work

We discuss idea management systems and how our case study extends prior work on this topic. We describe how a grassroots innovation pipeline differs from other models of innovation within organizations. We also review how work on brainstorming may apply to idea management systems.

Idea Management Systems

An idea management system is a central element of a grassroots innovation pipeline and arguably represents a new class of collaborative system. Examples of this class of system include Dell's Idea Storm (http://ideastorm.com) and My Ideas at Starbucks (http://mystarbucksideas.com). While implementations may differ, these systems generally support submitting, commenting on, and scoring ideas as well as browsing, searching, and associating ideas. Ideas are typically displayed in list form, but researchers are exploring visual metaphors that scale better for larger numbers of ideas [15]. Such systems are often integrated within organizational pipelines for processing the ideas.

The design of an idea management system is complex and the choices made may affect the quantity and quality of ideas, scoring of ideas, and even who participates. This, in turn, may affect the outcomes of the pipeline. By studying the content, interaction, and participation of one particular idea management system, our goal is to produce lessons for improving the design of this broader class of system.

Studies of Innovation Pipelines

Many organizations, including Microsoft, IBM, Whirlpool, Starbucks and Dell have deployed variations of a grassroots innovation pipeline (and idea management systems) as a way to tap into the collective wisdom of their employees and customers. For example, IBM organizes Innovation Jams where employees and customers engage in online conversations centered on strategic business problems [2] and mines the conversations for strategic directions. As more organizations will likely implement similar efforts, there is growing need for lessons to guide their design.

To offer initial guidance, Bjelland and Wood conducted a study of IBM's Innovation Jam [2]. The study focused on documenting the process and outcomes and offered novel insights from a management perspective. But this study did not examine any aspect of the idea management system.

Our work significantly extends this direction by examining the user and organizational experience of one specific idea management system and formulating recommendations for improving such systems and execution of the pipelines.

Organizational Models for Innovation

Innovation is at the heart of successful competition in fast-paced markets typified by shifts in consumer preferences and expectations, political climates, and such rising themes as environmental concerns [7, 23]. One approach to gaining access to creative ideas and streams of innovation is maintaining an R&D lab for exploring forward-looking concepts. Such labs can provide great value to companies and to the scientific community but typically require a large sustained investment. Also, creating business propositions for forward looking concepts can be difficult, i.e., business innovation is not the same as *invention* [8]. A second method is forming one or more 'creative' groups charged with incubating new ideas. The ideas explored have nearer-term focus, but the task of innovation is delegated to a select few. A third method allows employees to allocate part of their work time to side projects and showcase them in company-sponsored venues. Everyone can participate but employees might pursue only those ideas achievable with limited time and resources, possibly inhibiting the big ideas.

Each method has its merits and limitations and companies choose the most appropriate methods, often as part of a larger portfolio, based on perceived effectiveness, available resources, and risk tolerance. But as the need for innovation persists, organizations must continue to explore new strategies [23]. This paper examines one such strategy: a grassroots innovation pipeline. Relative to the others, the pipeline is unique in its grassroots nature (ideas flow bottom up), scale and openness (anyone and everyone can contribute), and commitment to fund the best ideas.

Brainstorming and Supporting Systems

Brainstorming is a process for generating ideas for solving difficult problems [18]. Productivity is commonly measured by the quantity, diversity, and quality of ideas [6] while field work shows that selectivity, knowing which ideas to pursue, is also key [22]. Research has produced many lessons for building effective brainstorming systems. For example, enabling simultaneous input and anonymization of ideas mitigate the inhibitors of production blocking and evaluation apprehension [11, 24]. Interaction for structuring the idea space [19] and helping participants build on each other's ideas can also enhance creative outcomes [21].

An idea management system is similar to a brainstorming system and should therefore draw from this literature where appropriate. However, an idea management system differs in terms of scale (open to the entire company), the funding model (the stakes are real), and scope (supports discussing and scoring ideas). Therefore, it is less obvious how brainstorming principles apply for an idea management system in context of an innovation pipeline. For example, building on each other's ideas can be useful, but may be at odds with a strong sense of ownership over a carefully crafted idea and a pipeline that rewards only the author.

Our work seeks to understand how to improve the design choices for idea management systems by evaluating the content, interaction, and participation of one system. We leverage knowledge of brainstorming to help interpret and translate our results into actionable recommendations.

THE GRASSROOTS INNOVATION PIPELINE AND IDEA MANAGEMENT SYSTEM

Within the organization we studied, a business unit (which we shall refer to as the GI unit) was formed in 2007 with the charge of establishing a corporate grassroots innovation pipeline. The central motivations for creating this unit was growing appreciation that all employees in the company have good ideas that often extend beyond their specific job focus, there should a means for surfacing, sharing, and nurturing those ideas, and the best ideas should have the opportunity to affect products regardless of organizational boundaries. It was understood that this unit would likely produce more failures than successes, but those successes, when they occurred, would make the effort worthwhile.

The GI team prototyped an innovation pipeline (IP) and idea management system (IMS). Note that both IP and IMS are generic acronyms. The prototype drew upon lessons learned from several prior attempts at structuring grassroots innovation within the organization. Key lessons included creating a model for formally evaluating submitted ideas, directing employee creativity to specific business problems, and devising methods for facilitating absorption of ideas.

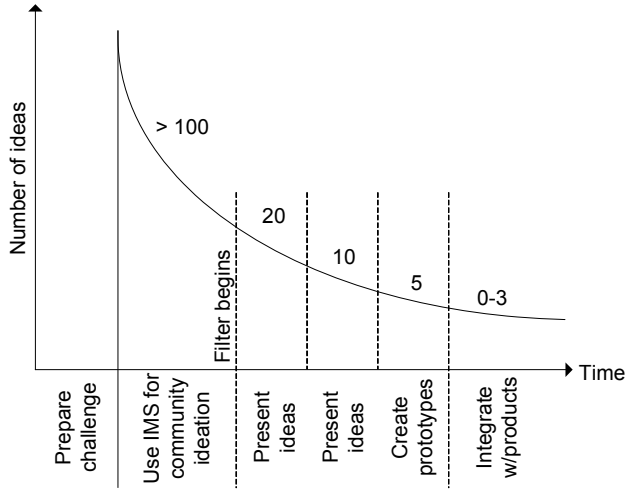

Figure 2: The grassroots innovation pipeline.

Shown in Figure 2, the IP consists of 4 phases; challenge preparation, community ideation, filtering and refining the best ideas, and integrating those ideas into products. To prepare a challenge, the GI team solicits proposals from internal business units grappling with complex problems. Units propose challenges through an informal process and the team evaluates them based on their scope, how forward thinking the solutions must be, and how receptive the unit appears toward absorbing new ideas. The selected challenge is then widely advertised via corporate e-mail, workshops, and posters and remains open for several months.

The IMS is a Web-based interactive forum that enables community ideation around challenges. Users can submit, comment on, and vote for ideas, as well as browse, search, and associate ideas. An idea is entered in narrative form, but additional media (e.g. videos, presentations, etc.) can be attached to it. When submitted, the author of an idea can place it within any open challenge. Comments can be attached to ideas and are shown in a typical discussion thread format. Similar to many online communities, a user can vote ideas up, but never down. Users can comment on and vote for as many ideas as desired. Once a challenge opens, community ideation occurs for about 2-3 months.

A challenge typically generates about one hundred ideas, which enter a filtering process (Figure 2). A unique aspect of the filter is that funding is pre-allocated for evolving selected ideas. The GI team filters the initial pool to about twenty and stakeholders from the business unit help winnow this set to about ten. Software developers on the GI staff create prototypes of these ideas and about five are chosen to move forward. From these, about three or fewer ideas are selected as the challenge winners. Each step in the filter requires a few weeks, though prototyping may require more time. At each step, the authors of ideas interact with the GI / stakeholder team via presentations and demos, and continue to refine their ideas. More or fewer ideas may pass through the filter based on the quality of the ideas, needs of the business unit, and resources available. The intention is to have the final prototypes either inform business unit leaders or form the basis of a new product or extension. The pipeline is repeated as often as resources and timelines allow and multiple challenges may be open at any time.

CASE STUDY: PURPOSE AND METHODOLOGY

To the best of our knowledge, there has not been prior research examining an idea management system in the context of an innovation pipeline. We therefore designed our study to answer several high-level questions:

- How much content is submitted (ideas, comments, and votes) and what is the quality of the content?
- What is the nature and degree of user participation? For example, where do users reside within the organization and what are the barriers to broader participation, if any?
- What is the user experience? For example, why do users contribute, what are the expectations, how are comments leveraged, what are the perceptions of voting, etc?

Can you briefly describe a recent idea that you contributed? Where did the idea come from?
What is your motivation for contributing ideas? What are the benefits in your opinion?
What did you expect would happen with your idea?
For what reasons do you review existing ideas before developing and posting your own idea?
How much influence do comments have on your ideas? Would you be willing to allow others to edit your idea?
What has been your experience with voting?
What are the strengths and limitations of the idea management system and overall innovation effort from your perspective?

Table 1: Sample of questions asked during user interviews.

- What have been the successes of the innovation pipeline, how is it measured, and what should the measures be?
- How can the core elements of the innovation effort, i.e., the idea management system and pipeline, be improved?

Our study examined these questions in the context of three creativity challenges organized by the GI team. The challenges addressed business problems in peer-to-peer advertising (P2P), identity-based system services (LIVE), and social computing (Social). The challenges were selected because they were recently conducted and open to the entire organization. For the P2P challenge, for example, ideas submitted included integrating advertising services within personal shopping lists, different means of transportation, and electronic communication systems.

Our methodology consisted of two parts. First, we analyzed the user interaction data logged from each challenge. This included analyzing the number and word length of ideas and comments, and the distribution of comments and votes across ideas. There was a total of 1491 users, 2211 votes, 488 comments, and 315 ideas. We also linked each user's alias to the corporate directory to extract where the person lies within the management chain (e.g. were they a senior or lower-level employee) and which business unit they belong (e.g. development, sales & marketing, research, legal, etc).

Second, we conducted semi-structured interviews (N=24) with recent and top contributors of ideas, comments, and votes. The sample included authors of ideas selected to move forward in the filter and the challenge winners as well as authors of ideas not selected. Table 1 shows a sample of questions asked. Questions were derived in part from early conversations with members of the GI team to understand key design problems and experiences with the IMS and IP.

We also conducted interviews (N=8) with senior managers who formed the GI unit, business leaders who sponsored the challenges, and GI team members. These interviews focused on learning the objectives, successes, and measures of the pipeline. All interviews were conducted in the user's workspace, lasted about one hour, and were audio recorded. Recordings were transcribed and coded to identify common themes. Users were compensated with a lunch coupon.

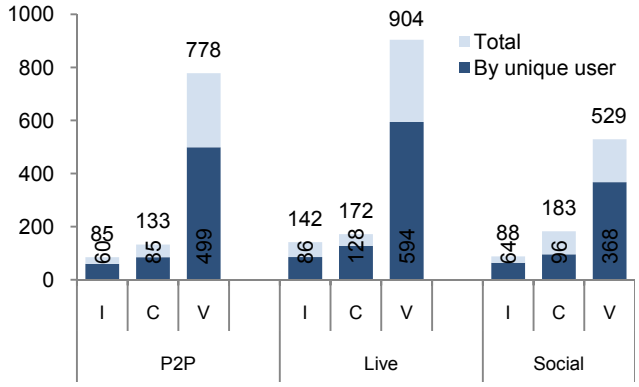

Figure 3: Amount of each contribution type (I=idea, C=comment, V=vote) per challenge. Top bars show the totals, while bottom bars show how many came from unique users.

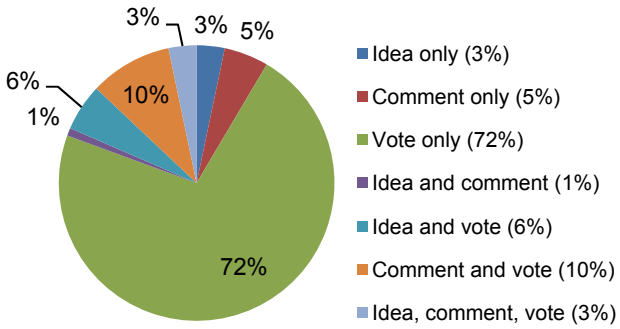

Figure 4: Distribution of users over each combination of contribution type, aggregated across challenges (N=1613).

RESULTS

In each subsection, we first describe the quantitative data and then draw from the interviews to help explain the data and add further insights. We begin with user participation.

User Participation

Figure 3 summarizes the total number of ideas, comments, and votes from each challenge. The top of each stacked bar shows the total while the bottom shows how many came from unique users. For example, for Social, there were 88 ideas, 183 comments and 529 votes; contributed by 428 users. Users contributing to the challenges had minimal overlap – 86% of users contributed to only one challenge, meaning each challenge tapped a different subset of the organization. Figure 4 shows a distribution of users over each combination of possible contribution types aggregated across challenges. It shows that most users (72%) only vote while the others (28%) offer at least one comment or idea.

An interesting question is whether parts of the organization, such as employees at the company's research division, contribute more than others. To answer this question, we indexed each user alias in the corporate directory and walked the user's management chain. We identified the user's highest-level manager under the CEO and the business area that executive oversees. Similar areas were merged to create a reduced set, e.g., Development was created by merging the Office and Windows development areas while Sales was created by merging Global Sales with

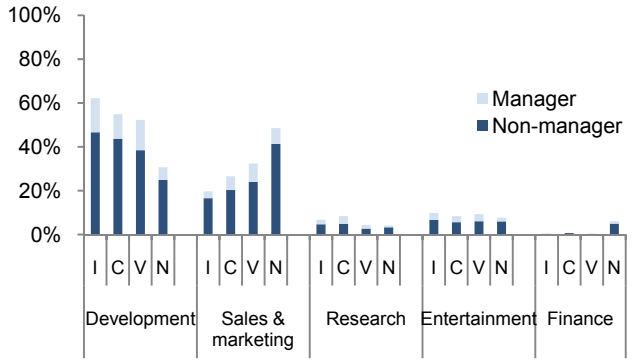

Figure 5: Participation by business area across the challenges. Normalization bars (N) show the percent of all employees in that area. Human Resources (I, C, V < 1%; N=1.7%) and Legal (I, C, V < 1%; N=1.3%) are not shown for brevity.

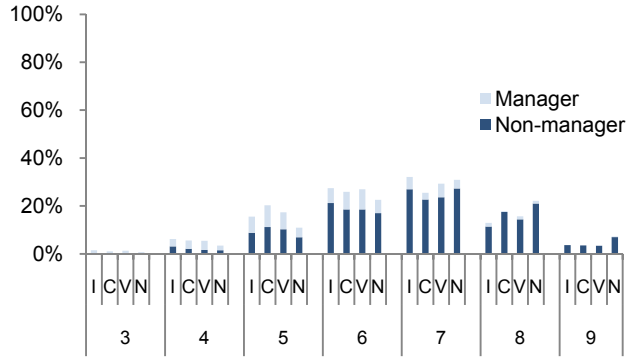

Figure 6: Participation by depth in the management chain across challenges. Normalization bars (N) show the percent of all employees at that depth. Depths 1 (I, C, V < 1%; N=.01%) and 2 (I,C,V < 1%; N=.12%) are not shown for brevity.

Windows Sales). We also extracted whether the user was a manager by checking if they had people reporting to them.

Results are summarized as a stacked bar chart in Figure 5. Each bar shows the percent contribution from users in that business area. The top of each bar shows the percent of that area that are managers and the bottom shows non managers. To the right is a normalization bar showing the total number of managers and non-managers in the area, and represents the percent of the total contributions that might be expected based on area size. For example as Sales accounts for 45% of the company, we might expect a priori that Sales would account for 45% of the ideas, comments, and votes.

Several interesting results are captured in the graph. First, Development contributed more ideas, comments, and votes than expected by area size (i.e. the bars are all higher for Development than its normalization, chi-square tests show $p<.05$ in each case). Second, though lower than its normalization bar, Sales contributed a surprising amount. Our interviews revealed that users from Sales felt their interactions with customers gave them unique insights and, since they have fewer technical skills, the pipeline was an ideal platform for surfacing and realizing their ideas (e.g. prototypes would be created by the GI staff) and gave them a basis for commenting and voting on ideas. Participation of Research was consistent with or slightly above expectation across challenges, but did not dominate the process. Finally, managers contributed consistent with expectation indicating that even with ostensibly tight schedules they were still willing to engage with ideas outside their own teams.

Another question was whether participation was coming from the deeper (grassroots) levels in the company. We extracted the depth of each user in the management chain (e.g. the second author was at depth four at the time, meaning he was four management levels below the CEO) and whether they were a manager or not. Results are summarized in Figure 6 and are interpreted analogous to Figure 5. If participation was not grassroots, for example, one would expect the graph to show over-representation at the smaller depths (left side) and under-representation at higher depths (right side). However, the graph shows the opposite, indicating the pipeline was receiving the largest participation from those employees furthest from decision makers about the allocation of resources for new ideas.

Returning to Figure 4, we believe that the numbers shown are reasonable given the pipeline was still maturing. But participation does seem low when one considers that the organization studied had about 95,000 employees at the time of this work. Our interviews revealed that one barrier to participation was the lack of clear incentive. For example even if their idea won the challenge, users were unsure if anything beyond personal satisfaction would come of it.

Ideas

Ideas spanned the range of the quality spectrum according to business unit leaders who evaluated them. One leader felt the value of the ideas came from their diversity; "There were definitely some that we were like, oh wow, that is a totally different way of looking at something ... even if we were skeptical, it was good to see the different perspectives." Another unit leader said "There were a few good ones, but a lot of them were not." From his perspective, the problem with many of the ideas was that they were too company-centric, ideas that would help only

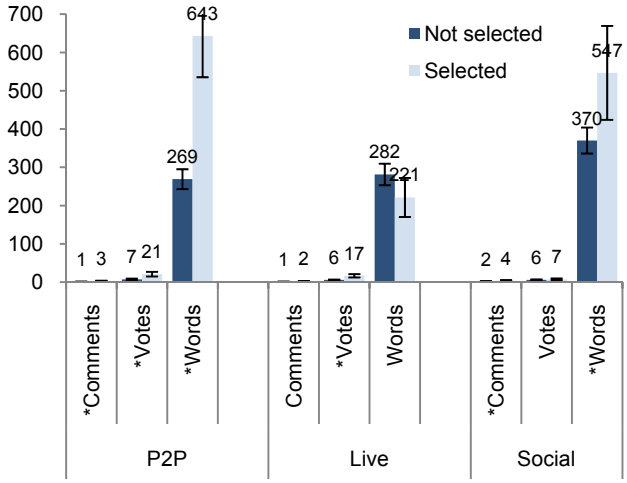

Figure 7: Comparison of comments, votes, and word length for not selected vs. selected ideas. * = significant at p<.05.

the company itself, as opposed to market-centric, showing how the idea makes sense in the market, why it is competitive, and how it would benefit consumers. This feedback reinforces the need for teaching *innovation* skills in the workplace [5]. Business leaders stated that the ideas selected as the winners were those already on the product roadmap, but still provided value as they had conveyed some of the scenarios better or offered a new perspective.

Figure 7 summarizes the average number of comments, votes, and words per idea for each challenge. It also compares ideas that were and were not selected to move forward (first step) in the filter. Ideas tended to be about a ½ page in length (about 500 words), and ideas with more words, comments, or votes were more likely to make it past the first filtering step, possibly because these ideas had more detail or appeared to have more 'energy' around them.

Users reported several motivations for contributing ideas. The most commonly cited motivation was the opportunity to see their idea *happen*. As one user said, "The choices that I had for all these ideas that come in my head was either to write them down, which I did for years, build a little prototype, but it doesn't go anywhere or, in this situation, it was an opportunity to take an idea … from my brain into potentially a product that millions of people are using." We note that despite the organization being a large software company, many employees do not work directly on software products, and even for those who do it can be very difficult to move one's own new idea through to a product.

A second motivation stemmed from users' feelings of being 'siloed' and seeing this as an opportunity to have their idea *heard*; "My motivation was to share the idea across the board so other people could see it and put their comments on it … to have a conversation around the idea rather than just have it in your head." A third motivation was cited as the desire to exercise one's own creativity; "One of the things I love to do is come up with new ideas … because you start thinking on a different perspective rather than what you are focusing on at work." Others reported submitting ideas because they wanted to learn new skills while others believed strongly in the innovation initiative and wanted to see it be successful.

We asked users about where their ideas came from. Users reported that ideas were almost never based on their current job focus. They already had an outlet for those ideas. Rather, their ideas germinated from their own personal experiences, desires, and frustrations with technology. One user stated directly "My ideas come from things that are bugging me." Users reported spending anywhere from a few hours to several weeks incubating ideas. Most would typically communicate their idea to trusted colleagues first, refine it, and then submit it. They did not want to put the idea into the system too soon for fear that others would begin commenting and voting on it before it was ready.

A very strong sense of ownership was felt over ideas, with several people describing an idea as their "child." Even if there were other similar ideas in the system, most users reported they would still enter their idea to show they had been thinking in this direction as well. Users were therefore apprehensive in their responses to our questions probing if they would be willing to allow community members to edit their idea (as in a Wiki). However, users were more open to this model if changes and rationale could be discussed first and they maintained ultimate control over any revisions.

Once an idea was submitted, few if any users expected their idea would win the challenge. Several users equated this to "winning the lottery." Rather the expectation was that a person on a team most related to the idea, whoever that person might be, would read and acknowledge the idea. Unfortunately, this expectation is typically not met because ideas in the system are not pushed to users in any way.

Comments

A goal of commenting within the IMS is helping authors understand how to refine and improve their ideas. Figure 8 shows the frequency of ideas receiving different numbers of comments, aggregated across challenges. The distribution is skewed; a few ideas receive many comments while most ideas receive only a few or none. On one hand, this result could mean the community is being discerning about which ideas are worth commenting on. This is partially supported in Figure 7 which shows that for two challenges, the ideas with more comments were more likely to be selected.

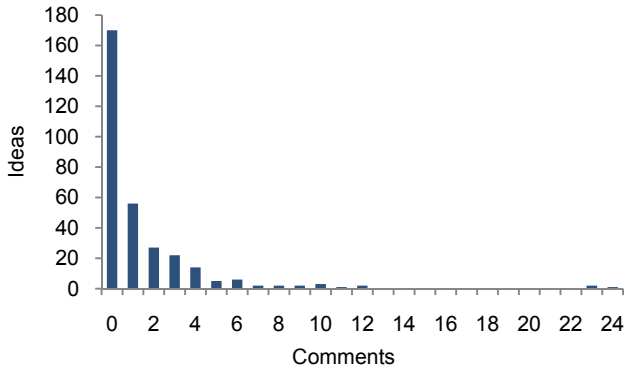

Figure 8: Frequency of ideas that received different numbers of comments, aggregated across challenges.

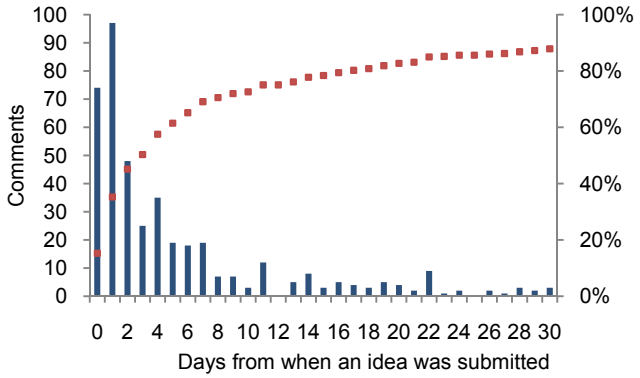

Figure 9: Distribution of comments over days from when an idea was submitted. Dotted line represents cumulative percent.

On the other hand, this result is less desirable because it does not meet the objective of helping authors understand how to improve their ideas. For example, the average length of a comment was 59 words (s.d.=68), which is about four sentences, and a large part of many comments was stating 'this is a good idea,' which does not help the idea to mature.

Though some ideas do prompt meaningful discussion, too many ideas receive too few comments or comments that are only cursory in nature. Our interviews revealed that users did not see the incentive for commenting. For example, one user described how he chose *not* to comment for fear of furthering someone else's idea at the expense of his own. Others chose not to make the effort to comment since only the author of the idea would benefit if it won the challenge.

Many comments tried to foster social connections between the author and others in the company working on similar topics. Links to ongoing product and research efforts related to the idea were also common. It could be useful for idea management systems to extract and attach a summary of this type of information to an idea.

The distribution of when comments are generated for ideas was analyzed and the results are summarized in Figure 9. Results show that most comments for an idea occur within a few days of it being submitted. This could be a function of the algorithm used for ordering the default view of ideas in the system - newest first. When submitted, an idea receives quick attention from the community, if at all, before being pushed out of view and out of focus by more recent ideas. A similar result was reported in [17] which found that most of the commentary in Slashdot occurs in the first few hours after a story is posted. The algorithm for ordering the default view of ideas may therefore be more effective if it weighs whether commenting for an idea has stalled or the idea has received few comments overall [17]. This may help direct community attention where it is most needed.

Votes

Figure 10 shows the distribution over votes for ideas. As with commenting, most ideas receive only a few votes and a few ideas receive many votes. One interpretation is the community is weighing the alternatives and voting only for those ideas judged to be of high quality. Some users shared this interpretation in the interviews. For example, one user felt the number of votes creates "the perception of a quality bar" while another felt if his idea did not receive many votes then "obviously people care for something else." However our interviews uncovered behavior that cast doubt on this interpretation and identified additional issues.

Most users reported that after submitting an idea, they would "market" the idea by mailing friends and colleagues within the company and asking them to vote for it. Because users can vote as often as desired, there is no cost to voting for an idea in response to such requests. As recognized by one user, the number of votes reflects "how much you are willing to market your idea" rather than its intrinsic quality.

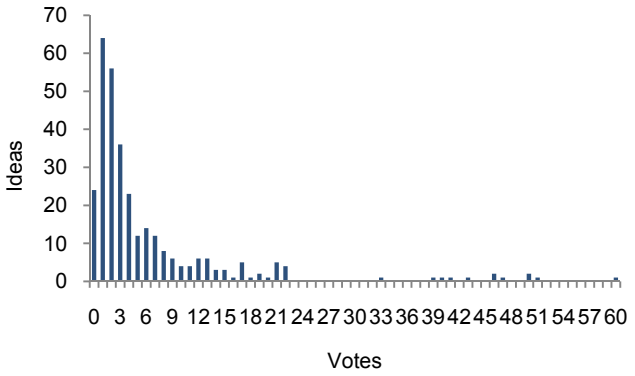

Figure 10: Frequency of ideas that received different numbers of votes, aggregated across challenges.

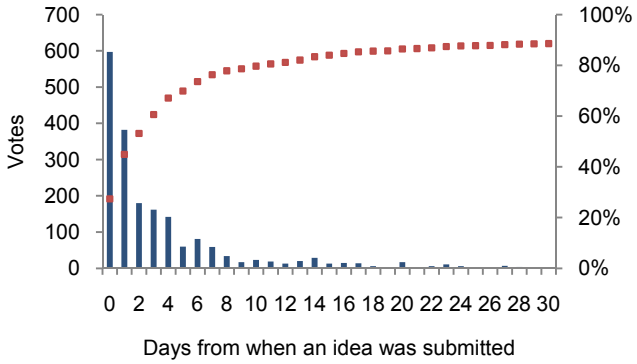

Figure 11: Distribution of votes over days from when an idea was submitted. Dotted line represents cumulative percent.

A second issue was concern about the meaning of a vote, which, in its current form, means the voter *likes* an idea. Business unit leaders argued this was inherently flawed because 'liking' an idea is different from saying it has business value. For example, as one leader bluntly stated "the problem is that there is no notion of 'relative importance.' Just because the employee thinks it's a great idea does not mean that the idea has business value." A better approach might be to rate ideas on relevant business dimensions (e.g. whether the idea opens a new market, how well it fits with current offerings, technical feasibility, etc.).

Third, many users realized that votes were not normalized against the number of views and therefore reported using the vote/view ratio as a better indicator of what people thought about the idea. For example, receiving 5 votes from 10 views may be better than receiving 20 votes from 100 views. Fourth, several users strongly disliked the concept of voting altogether. The reason is that "it [not receiving many votes] causes people to become disenfranchised with the site and following through with their ideas. The quick thumbs up/down is too immediate." This is particularly problematic for ideas that are actually good, but their value has not been adequately communicated or understood.

Finally, recall that ideas selected to move forward typically had more votes (Figure 7). This may be because seeing more energy for an idea in terms of more comments and/or votes influences decision makers during the filtering phase.

We also analyzed when votes are cast for ideas. Similar to comments, most votes for an idea, if any, occur within the first few days of idea submission (see Figure 11). This provides further evidence that the default ordering of ideas influences which ideas the community interacts with.

Length of Engagement

A goal of the pipeline is to engage an increasing number of employees and maintain that engagement over time. To test the level of engagement, we extracted the number of actions, including any idea, comment, or vote, made by each user. Results show that about 95% of users perform at most five interactions and this was true for each challenge. The pipeline still functions because there are new users participating each time, but this is probably not sustainable long-term. Interviews revealed that the lack of continued engagement was based in the absence of a compelling incentive and concern over taking time away from their main job focus. The latter indicates the potential value of a corporate policy describing when it is appropriate to pursue a grassroots idea and how much work time can be allocated.

Successes and Measures

One measure of success for this type of innovation pipeline, as one business leader mentioned, is "Counting how often something gets out there that wouldn't have happened without the process and impacts a huge number of people." On this measure to date, the pipeline has produced about one hundred prototypes and about six of those have been absorbed into product teams (it is not known if these have made it into a product yet). For example, one of the adopted prototypes was for a new mobile advertising platform. However, the fact that there *were* adoptions demonstrates grassroots innovation can be structured within organizations and yield new ideas of interest to product teams. As this pipeline was untested, we consider this to be a positive sign.

Our interviews revealed that additional measures of success may be warranted. For example, many users felt that the ideation skills gained from moving through the pipeline was one of the most valued outcomes. One user said "When they gave me all those resources, it was like basically a simulation of being a GPM or something, in a very small scale, I had one week and I had my graphic design, I had my architects, I had my people, and it was like 'let's execute.' It gave me management experience that definitely gave me a lot of insight. I think that's the best takeaway." Another user agreed, "What I am getting out of it … is learning skills that I don't really learn on the job. I am learning presentation skills, crafting messages skills, etc."

Senior managers recognized the value too, "We want to help people get smarter about their ideas … if we achieved nothing but a better educated, more thoughtful workforce that would make it worthwhile." Also, senior managers felt that, even without the challenges and funding model, it was important to have "a corporate forum for talking about ideas" and that "better measures may include understanding the people connections that get formed, how ideas influence other ideas, and how ideas get passed around the company."

RECOMMENDATIONS AND DISCUSSION

From our case study, we derived several recommendations for improving the design of idea management systems and execution of the pipelines in which they are embedded:

Foster meaningful participation. Despite the raw numbers, only a fraction of the corporate community participated in the pipeline. One reason is that users did not see a clear and compelling incentive that would justify diverting attention from their job focus. Users also perceived little incentive in commenting on or voting for ideas as only the author would benefit. As one user aptly stated "If you want more people to be involved, you need incentives." Though many types of incentive models exist, such as reputation systems [12], point-based rewards [10], and social rewards [16]), there was agreement the appropriate incentive here would be financial. For example, one model would be to offer a financial award for winning ideas distributed proportionally to those who contributed to them (author > commenter > voter). The number of comments and votes per user would need to be capped to force choices between ideas. With this model, for example, a user would not only have to choose which idea(s) to comment on, but it would be in her best interest to offer a substantive comment that improves the idea because if the ideas wins, she would benefit too. Tuning this type of model would require further research.

A related issue is that the pipeline currently executes in parallel with employees' main job focus and, if one wants to participate, must determine how to effectively balance their time between the two. Following the approach in [2], an alternative would be to define specific time period(s) in which employees could participate *as* their main job focus. This, combined with an effective incentive model, could spur participation that is both copious and meaningful.

Use business relevant criteria in the voting process and/or change the method of voting. Voting is intended to help filter the best ideas and scale the process [2]. But the model of voting within the system studied, which reflects models used in popular online communities (e.g. Digg.com), had several problems in this context. For example, votes do not take business value into account, votes may not be forcing value-centric choices between ideas (i.e. the number of votes per user is unlimited), and the author's passion for an idea may be diminished if too few votes are received.

There are several ways to address these problems. For example, to prompt choices, the number of votes per user could be capped in conjunction with using incentives (see prior recommendation) or the model of voting could be changed to a rank system [4]. Second, users could be asked to rate or rank ideas on dimensions relevant to the business (e.g., the idea opens a new market or enhances current offerings). Third, a system could make vote tallies available only for the highest rated ideas or only for the evaluation team. This may allow the top ideas to be identified without visibly reinforcing the lower scores of other ideas. Finally, voting could be framed as an 'expression of interest' for an

idea and the system could notify subscribers of updates, allowing subscribers to follow the idea at it matures.

Measure and appreciate outcomes beyond revenue. The success of the innovation pipeline is often perceived as a function of the products or revenue it generates [2]. But it was clear from our study that at least three additional outcomes should be appreciated and measured: workforce ideation skills, cross-pollination of ideas, and a feeling of contribution or being less "siloed." Acquisition of ideation skills or feelings of contribution could be measured, for example, by asking employees to rate their learning of or attitude toward these and similar elements on annual self-report surveys, common in many organizations. To measure cross-pollination, in addition to using a survey, one approach would be to measure the number of URLs included in internal documents and communications (email, IM, discussion posts, etc) that reference ideas in the system.

Structure the ideation phase of the pipeline. The pipeline studied overlaps entering, commenting, and voting for ideas in time. This creates undesired situations, e.g., ideas entered late in the ideation phase will not receive much attention from the community while ideas entered early may not be reconsidered if refined. As suggested in the brainstorming literature [13, 18], a more effective approach is to structure the ideation phase, creating defined periods for submitting, commenting, refining, and voting for ideas. This would help ensure each idea receives equal attention and has equal opportunity to be refined prior to being evaluated.

Continue with focused challenge problems. It is possible to execute the pipeline without challenge problems, e.g., users could submit ideas on any topic of their choice. This would require fewer resources, but it would likely attract fewer good ideas. Many users, particularly those whose ideas moved forward in the pipeline, reported they had been thinking about their idea for some time. But it was knowing there was a deadline and a team of people with relevant knowledge who would be evaluating the ideas that became the catalyst for the users to pursue and submit their ideas.

Raise the bar (slightly) for submitting ideas. Our analysis revealed that ideas span the quality spectrum, from whims to well-researched proposals. A consequence of having a large number of lower quality ideas is that it fosters the perception of "there is nothing good in there", it diminishes the credibility of the pipeline within the organization, and it absorbs resources from the community on the review of poor ideas. One solution is to place newly submitted ideas into a separate, non-public space and only allow those ideas that meet a certain threshold to pass into the public area. The threshold could be low such as only ensuring the idea addresses specific criteria (problem description, audience, etc.) or could be slightly higher such as requiring minimum discussion around the idea. This effort could be assigned to those who had ideas recently accepted, borrowing methods from [3]. Placing a low barrier to entry may help filter the whims without inhibiting ideas with more consideration.

Increase the lifetime of ideas. Ideas not selected to move forward in the pipeline remain accessible in the system but rarely receive further attention. As this represents the large majority of ideas, it creates the perception that the system is "a graveyard for ideas," as noted by one user. However, since there are many reasons why good ideas may not be selected, e.g., their value is not yet well understood or they are less relevant for the particular challenge, it would be useful to explore ways for increasing the lifetime of ideas. One way, as many users argued, is for those responsible for gathering product requirements to be asked to review relevant ideas in the system as part of their normal workflow. This would help expose them (and the product) to diverse perspectives and give authors the satisfaction of knowing their ideas were heard. A second way is for new challenges to be seeded with relevant ideas from prior challenges. A third way is to create tools for visualizing the idea space (e.g. to gauge what users across the company are thinking about) as this may identify strategic trends. Rather than abandon ideas once a challenge is complete, these enhancements may encourage authors to view the refinement of ideas as a long-term, valued effort.

Support the process of innovation within the system. As shown in Figure 2, the IMS is used for only part of the pipeline. A negative consequence is that users become unaware of the status of ideas as they move forward outside the system. This was a common source of frustration. Another possible consequence is creating the misperception that innovation is rapid and that it requires little effort, whereas it is typically a long and arduous process. One solution is to embed a macro-level process of innovation in the system. For example, borrowing the process from [8], the system could provide distinct spaces for 'opportunity', 'analysis', 'listening', 'focus', and 'leadership'. This would allow the pipeline to be captured as it unfolds, which helps to inform users as to what the innovation effort entails, and provides example materials that others can build upon.

Support users who want to advance their own ideas. Many users are passionate about their idea and even if the idea is not selected in the pipeline, they are willing to advance it on their own. However, because the IMS does not allow the capture of the progression of an idea beyond its initial form, such efforts unfold outside the system, thereby missing the benefits of community support [14]. To help users who wish to advance their own ideas, a starting point would be to implement an innovation process in the system (see prior lesson). This would create a template outlining the steps necessary for moving the idea forward and would make materials from other projects accessible as examples. The system could also help users recruit others with needed expertise (e.g. graphic design), seek community assistance at choice points (e.g. which toolkit or platform is best for prototyping the idea), and identify appropriate outlets for the idea (e.g. internal workshops, innovation fairs or lunch presentations). It could also serve as a resource for viewing what grassroots projects are being pursued within the

organization and their current status. However, facilitating grassroots projects must be done with caution, and there needs to be an explicit management buy in or more general corporate policy describing when it is appropriate to work on a project and how much work time can be allocated.

While these recommendations were derived from analyzing one particular idea management system and pipeline, they address issues of a more general nature, e.g., incentives, voting models, and idea lifetimes. Thus, we believe that organizations pursuing similar efforts should find value in the recommendations. But, realizing these and other design requirements in a system will require satisfying the goals of both lower-level employees and senior managers, which may not always align. System designers may therefore find participatory design [9] or other techniques that engage stakeholders in the design process particularly beneficial.

CONCLUSION AND FUTURE WORK

The use of an idea management system and pipeline offers a new approach for structuring and nurturing innovation within organizations. But as this approach is just emerging, there are few lessons to guide the design of these elements and little understanding of how they unfold in practice. From a study examining these elements within one organization, our research has made two contributions. First, from analysis of the content, interaction, and participation data and interviews with end users, we have formulated recommendations for improving the design of idea management systems and execution of the pipelines. Second, by sharing the perspectives and insights gained from senior managers, business unit leaders, and end users, others can better anticipate and plan for their own implementation of a grassroots innovation pipeline.

We see at least three promising directions for future work. First, we believe a longitudinal study would be valuable for understanding how the use of an innovation pipeline affects the idea culture of an organization (e.g., does it raise the exchange of ideas across unit boundaries?). Second, we wish to test the impact of different incentive models on the quantity and quality of participation. Finally, it would be interesting to study other types of grassroots innovation efforts (e.g., applying some percent of work time to side projects) and compare the findings with those reported here.

REFERENCES

1. Amar, A.D. and J.A. Juneja. A Descriptive Model of Innovation and Creativity in Organizations: A Synthesis of Research and Practice. *Knowledge Management Research & Practice*, *6* (4), 2008, 298-311.
2. Bjelland, O.M. and R.C. Wood. An Inside View of IBM's 'Innovation Jam'. *MIT Sloan Mgmt Review 50* (1), 2008, 31-41.
3. Cosley, D., D. Frankowski, L. Terveen and J. Riedl. Using Intelligent Task Routing and Contribution Review to Help Communities Build Artifacts of Lasting Value. *CHI*, 2006, 1037-1046.
4. Dasgupta, P. and E. Maskin. The Fairest Vote of All. *Scientific American*, *290* (3), 2004, 92-97.
5. Denning, P.J. The Social Life of Innovation. *Communications of the ACM 47* (4), 2004, 15-19.
6. Diehl, M. and W. Stroebe. Productivity Loss in Brainstorming Groups: Towards the Solution of a Riddle. *Journal of Personality and Social Psychology*, *53* (3), 1987, 497-509.
7. Dougherty, D. and C. Hardy. Sustained Product Innovation in Large, Mature Organizations: Overcoming Innovation-to-Organization Problems. *Academy of Management Journal*, *39* (5), 1996, 1120–1153.
8. Drucker, P. *Innovation and Entrepreneurship*. Harper Business, 1993.
9. Ehn, P. Scandinavian Design: On Participation and Skill. Schuler, D. and Namioka, A. eds. *Participatory Design. Principles and Practices*, Lawrence Erlbaum, 1993, 41-77.
10. Farzan, R., J.M. DiMicco, D.R. Millen, C. Dugan, W. Geyer and E.A. Brownholtz. Results from Deploying a Participation Incentive Mechanism within the Enterprise. *CHI*, 2008, 563-572.
11. Hymes, C. and G. Olson. Unblocking Brainstorming through the Use of a Simple Group Editor. *CSCW*, 1992, 99-106.
12. Josang, A., R. Ismail and C. Boyd. A Survey of Trust and Reputation Systems for Online Service Provision. *Decision Support Systems*, *43* (2), 2007, 618-644.
13. Kelley, T. *The Art of Innovation*. Doubleday, New York, 2001.
14. Kittur, A. and R. Kraut. Harnessing the Wisdom of Crowds in Wikipedia: Quality through Coordination. *CSCW* 2008, 37-46.
15. Krieger, M. and Y.Y. Wang. Ideas2ideas: Encouraging Constructive Ideation in an Online, Mass-Participation Brainstorming System. *UIST, Poster session*, 2008.
16. Kriplean, T., I. Beschastnikh and D.W. McDonald. Articulations of Wikiwork: Uncovering Valued Work in Wikipedia through Barnstars. *CSCW*, 2008, 47-56.
17. Lampe, C. and P. Resnick. Slash(Dot) and Burn: Distributed Moderation in a Large Online Conversation Space. *CHI*, 2004, 543-550.
18. Osborn, A.F. *Applied Imagination: Principles and Procedures of Creative Problem Solving*. Charles Scribner's Sons, New York, NY, 1963.
19. Prante, T., C. Magerkurth and N. Streitz. Developing CSCW Tools for Idea Finding - Empirical Results and Implications for Design. *CSCW*, 2002, 106-115.
20. Sandberg, J. Understanding Human Competence at Work: An Interpretive Approach. *Academy of Management Journal*, *43* (1), 2000, 9-25.
21. Shah, J., N. Vargas-Hernandez, J. Summers and S. Kulkarni. Collaborative Sketching (C-Sketch): An Idea Generation Technique for Engineering Design. *Journal of Creative Behavior*, *35* (3), 2001, 168-198.
22. Sharmin, M., C. Coats, B.P. Bailey and K. Hamilton. Understanding Knowledge Management Practices for Early Design Activity and Its Implications for Reuse. *CHI*, 2009, 2367-2376.
23. Sorensen, J.B. and T.E. Stuart. Aging, Obsolescence, and Organizational Innovation. *Administrative Science Quarterly*, *45* (1), 2000, 81-112.
24. Stefik, M., et al. Beyond the Chalkboard: Computer Support for Collaboration and Problem Solving in Meetings. *CACM*, *30* (1), 1987, 32-47.