Designing to Maximize Community Inertia

Samuel Blake

Luke Parrott

ABSTRACT
As administrators try to engineer online communities and socially-oriented applications, it is frequently desired that the community of users maintain some relatively consistent standards of contribution, even as users join and leave the community. We propose a system that can help maintain this property of community inertia using a combination of user feedback and automatic comparisons of new content with older content that meets the desired community criteria. Our system classifies users in such a way that new users are not significantly barred from contribution, but they can be judged by existent users who meet community standards. Regular users can also be judged relative to the rest of the community, and can be encouraged to participate or discouraged depending on their behavior. Such a system, we assert, would allow community administrators to be able to more easily maintain desired behavior from the community at large as well as manipulate the community's structure.

1. INTRODUCTION
While the Internet is most famed for its ability to incubate and evolve new and original communities, many of these communities, once established, have a desire to restrict the community to certain classes of members. A knowledge-sharing site, such as a question-and-answer application or a discussion forum, may want to ensure that its active participants are knowledgeable and helpful in answering questions or debating issues. A social news site like Digg may want to encourage users who regularly provide interesting or hard-to-find content. As long as there is some standard of contribution that makes some users more valuable to the community than others, there is a clear benefit to be had in trying to maintain these standards. We define *community inertia* to be the maintenance of a set of community users who maintain the desired contribution standards.

In order to maintain community inertia, we need a system where new users are expected to behave similarly to existing users. In the style of operant conditioning, we reward users who behave as desired (by the standards that already exist in the community) and punish[1] those who do not. Unlike most systems that exist, where such rewards are generally meted out by hand (such as an editor position on Wikipedia or a ban from a discussion forum), our system strives to operate in a fully automatic fashion, making it fully scalable and presenting minimal burden to system administrators.

In our system, the desired behavior is measured in two ways: feedback from users that the system has designated as contributors of an acceptable level and automatic content rating based on comparisons with past content. The reward or punishment is meted out by classifying users based on these ratings of their content. Users of different classes have different abilities and influences within the community, which either provides them an incentive to maintain standards of contribution, to improve up to a certain standard, or to stop contributing to the community if their contributions do not conform to expectations. Regardless of how they are influenced, even if not at all, our system allows users to easily ignore those who are not meeting the standards of the community.

To restrict the problem to a feasible scope, we design our system around the concept of a question-and-answer site, but the basic functionality can be expanded to any community that has some way of measuring, either quantitatively or qualitatively, the value of a user's contributions.

In addition, for the purposes of this paper, we will frequently use the concept of "quality" as a substitute for "similarity to expected community norms." While not strictly correct, since quality is (by definition) qualitative and similarity is more quantitative, we assume that most community administrators would use a system such as the one we describe primarily as a quality assurance tool.

It is also worth noting that some aspects of the design of our system closely mirror the contribution viewing and rating mechanisms on Slashdot. This is actually something of a coincidence; neither author had used Slashdot before this system was designed, and the similarity was only noticed after the fact. We presume that the similarities are the result of similar initial beliefs being taken to their logical design conclusions. As a result, we find the similarities encouraging, as it suggests that the design is likely to be plausible in implementation. There are important distinctions, however, which are discussed in detail in a later section.

[1]"Punish" might in fact be too strong a word. In order to keep barriers to entry minimized for potentially valuable contributors, the only "punishment" one can actually receive in our system is not actually experienced by the user; he or she is aware of the "punishment," but it is only *other* users whose experience actually changes.

2. BACKGROUND AND JUSTIFICATION

Traditionally, question-and-answer applications have been created to facilitate the act of knowledge sharing amongst a community of Internet users. Though many may dispute the validity and quality of answers, participants of these systems believe this to be valuable knowledge, and continue to exchange it [1].

Past research has found that users in contribution-based systems naturally fall into different classes [6]. Generally, these class divisions are based on users' tendencies to either consume or produce content (or some combination of the two). We believe that, in the context of the system we are proposing, a preferred user is one who produces a reasonably large amount of high-quality content. Pure consumers of a question-and-answer system (those who only read or only ask questions) generally contribute nothing to the system's overall usefulness. As such, we feel that this justifies isolating users who meet acceptable levels of contribution and nurturing them.

Beyond these naturally-occurring classes of users, the accepted standards that we wish to maintain within a system also appear to evolve of their own accord. Research has shown that, as people use a discussion group, they can be assigned to distinct clusters based on their actions and viewpoints [5]. However, on the edges of these clusters exist users whose views are not seen as worthy of response from the core members of any group. This phenomenon is referred to as boundary maintenance, and the ignored user can be implicitly or explicitly expelled from the network [5]. While users can coexist even while having differing viewpoints, a congruent user community will tend to eject users who do not meet the standards of appropriate behavior [5]. This is significant to our system's design because these standards evolve naturally and are sufficiently polarizing that the overall quality of a community is often directly related to how well those standards are met.

The perceived value of an author is frequently the prime determinant of a piece of content's perceived value, often regardless of the actual value of the content [4]. As a result, it is generally important in content-based systems that information about the author of content is made readily available. However, many systems tend more to reveal information about the content itself rather than about the author that posted it. This forces users to focus their attention on the layout of the system rather than on the quality of the poster [4]. It has also been shown that knowledge of an author's reputation has a significant impact on how users will rate that author's content, which again suggests that this information is crucial [4]. This familiarity can extend either from explicit experience with an author or from an author's reputation. To this end, our system emphasizes the classification of *users as contributors of content*, not the measurement of the pieces of content themselves.

In the domain of product reviews, it has been observed that people rely heavily on the textual quality of reviews when deciding if a reviewer is trustworthy or not [3]. This quality, extended over time, results in an author receiving a reputation. Reputations help guide users toward the highest quality opinions, while also providing a standard to judge ratings when computing overall scores for rated items, leading to much more reliable scores. In a secure quality system, newcomers must achieve a certain level of participation and reputation in order to have a meaningful level of impact on the community, thus reducing the likelihood of cheating ratings or spamming. It is also important that maintaining a reputation should not be static, but instead should be dynamic and ongoing. If a user drops below the threshold that earned them their reputation, that reputation should be either removed or decreased, given that either the user's behaviors or the communities evolving standards [5] have changed. In our system, a user's reputation (as defined by his classification) allows other users to make certain assumptions about the quality of his contributions, which can make interactions more efficient.

Further research suggests that content quality, specifically in the domain of answers to questions, can be found by looking at both the content of the answer as well as any knowledge known about the contributor, such as their expertise [7]. This expertise can be automatically inferred by noting their experience with the system and how other users acknowledge them. Beyond this, the quality of the answers can further be determined by looking at the answer's length, the number of answers that this user has provided so far, and the number of times this answer has been rated highly by other users [7]. We expand this idea in our system by using information about a user's past quality (his classification) to determine the visibility of his content to others.

Past research has found that it is quite possible for accurate predictions about authors and answers to be made early. In one study, it was shown that users who will someday become highly regarded in a discussion group receive preferential treatment from the moment they first join and post within the group, before there is any chance to develop meaningful relationships with other members [2]. This information is useful in making predictions about users as they join a group, and especially useful for seeking out future quality posters. This could easily become an important tool for filtering the vast amount of information found in large groups, and provide means for moving the more engaged members of a group into smaller, more manageable sub-groups [2]. We can take this to suggest that using a ratings system focused on already-established users can be a very accurate predictor of which new users will become valuable contributors later. It is also important to expose the existing reputation or ranking system to new users, if only to point out the advantages of following the norms and conventions currently enforced [2].

We also know that automatically detectable patterns exist in this domain. For instance, categories of questions that favor factual answers tend to have shorter thread lengths than those requiring opinions or more thought-provoking answers [1]. Also, it has been observed that users who focus their efforts on certain categories produce better answers overall [1]. Looking at even the most basic metrics like reply length and number of competing answers can help greatly in predicting which answer will be considered best [1]. These findings suggest that automatic scoring of answers can potentially be useful.

This approach has been taken in other domains as well. It has been shown that surface features of a post are helpful for examining its quality by looking at its length and other structural features, even down to the level of punctuation and formatting [8]. Relevance and originality are also good metrics to use, as we want to avoid repeat posts while keeping them in line with the conversation at hand.

3. DESIGN PROPOSAL

3.1 Design Overview

Given that there is a benefit in filtering users based on their contributions [3], we have chosen to make this the basis of our design. In our system, users are classified based on the quality of their contributions as determined by feedback from other users and by similarity to past content. Once users are classified, it affects how they can be viewed; users considered to be valuable contributors are fully visible, whereas contributions from less valuable (either new or low-quality) users can only be viewed through a set of restrictions. We attempt to minimize barriers to entry so as to not discourage new users from becoming active members of the community; as a result, the restrictions placed by the system generally do not directly impact the user who earns them.

3.2 Detailed Design Specifications

The system we propose has four major components: (1) classification of users into categories, (2) dynamic user visibility, (3) a user feedback system (ratings), and (4) automatic rating calculation.

3.2.1 User Classes

The most significant aspect of our proposed system is that it automatically classifies users into three categories. We feel this is justified based on research in [6] and [5], which suggests that users fall into categories similar to those we define anyway. Specifically, [6] posits two naturally-occurring groups, *Content Producers* and *Producer & Consumers*, which we are attempting to isolate from the users of the system at large (albeit with the additional criterion of quality).

Classification of users is based on two metrics: the number of answers the user has submitted (denoted as a) and the percentage q of the user's answers that have a positive rating. (The rating system is discussed in more detail later.) Each metric has an associated threshold, the value of which is specific to the implementing community: a_{min}, or the number of answers that constitute enough to cross from a "new" user to an "experienced" user, and q_{min}, the minimum percentage of answers required to be of positive quality. The user classes are then defined thusly:

- New User ($a < a_{min}$): A New User is one who is new to the site and has not yet submitted enough content to be accurately measured. New Users have the full capabilities to ask and answer questions, but external visibility of their content is restricted (as discussed in detail in the next section). New Users are essentially considered as undergoing a trial phase, during which they prove their worth to the community.
- Restricted User ($a \geq a_{min}$, $q < q_{min}$): A Restricted User is a user who has submitted a reasonable amount of content but does not meet the minimum standards of quality. The term "restricted" is something of a misnomer, because, like New Users, Restricted Users, still possess the full ability to ask and answer questions. However, the external visibility of their content is again restricted. For all practical purposes, New and Restricted Users are treated identically; the only difference is in the amount of content they have submitted.
- Full User ($a \geq a_{min}$, $q \geq q_{min}$): A Full User is a user who the system identifies as a valuable contributor. Full Users have full external visibility of their content, and have the additional ability to provide feedback into the system.

3.2.2 User Visibility

Once a user has been classified, it impacts how visible his answers are to users of the site. The expectation is that, primarily, Full Users will read answers provided by other Full Users; as a result, we have opted to make Full Users' answers always visible, while answers by New and Restricted Users are collapsed by default, as in Figure 1. Only the beginning of the answer is visible, and users have the option to expand the answer to view it in its entirety. If multiple collapsed answers are adjacent, they are further collapsed into a single answer set, which can also be expanded if desired.

Note in Figure 2 that it is possible for users to permanently expand the New and Restricted Users' answers. Our hope is that more active users, wanting to take on the role of moderator rather than simply producing and consuming answers, will use this mechanism to make it easier to rate the answers of New and Restricted Users. Thus, we refer to this as the *moderator view.* To further facilitate moderator view, we have made the New and Restricted Users' answers a different color, in order to draw attention to them if they are expanded. Users who prefer to simply read from Full Users (the "good" answers) retain that option, and users who activate moderator view can deactivate it at any time.

While we created moderator view with the intent that it be used by Full Users, it is an option accessible for all users. In the worst case, this would allow all users to have moderator view active, in which case the site would be the same as a question-and-answer site without this system. However, we have made the default behavior that moderator view is deactivated, with the expectation that users will not activate it unless they have an express interest in moderating answers.

3.2.3 User Feedback

As mentioned in the previous section, users in the Full User class have access to the rating system. For any contributed answer, a Full User can give a positive, negative, or neutral rating. An answer is considered to be "good" if it has more positive ratings than neutral and negative ratings combined, or:

$$r_{positive} > r_{negative} + r_{neutral}$$

While a bit strict, this enforces a high standard of conformity. A user's quality value q is then the percentage of

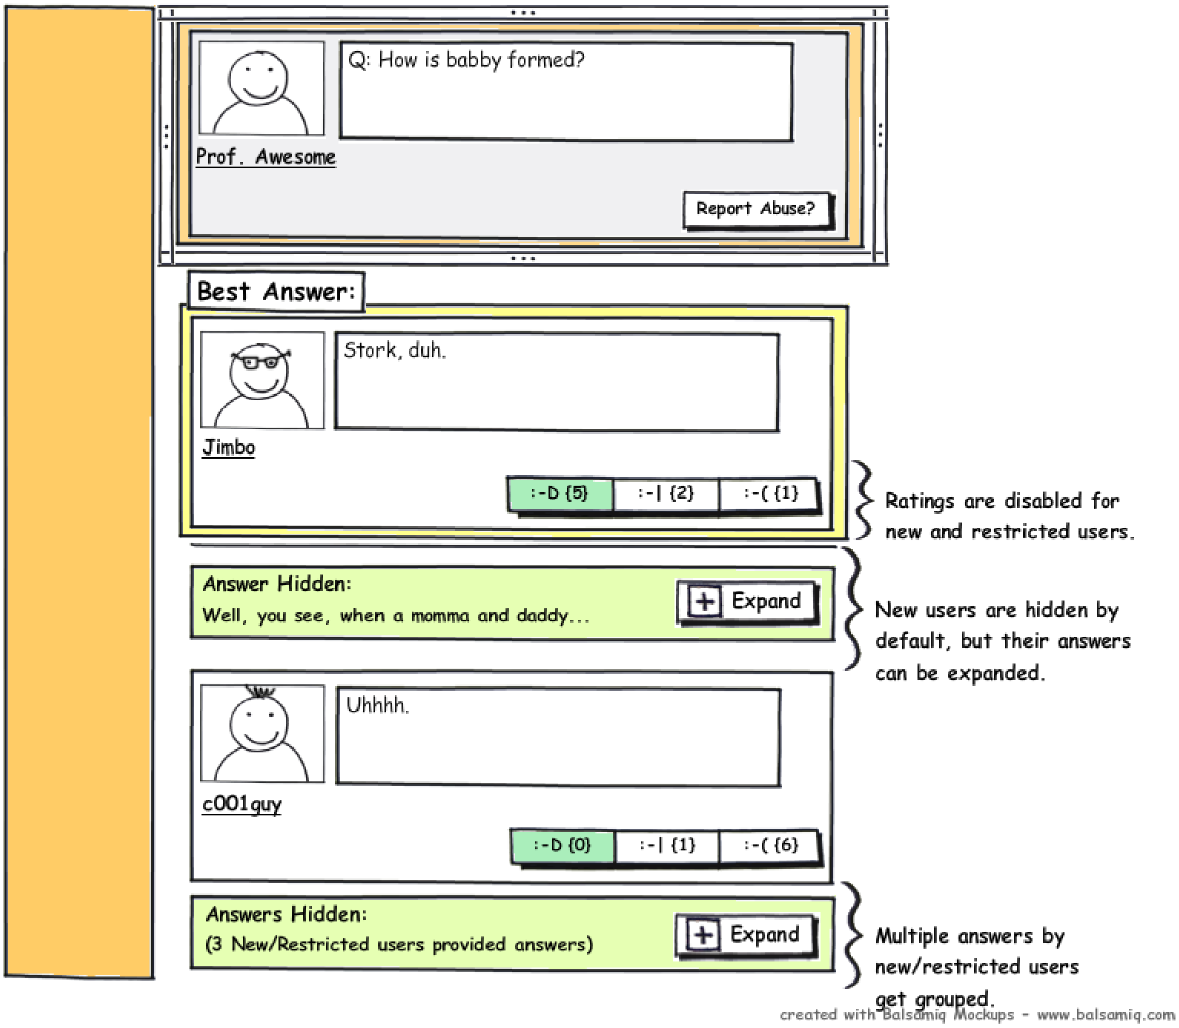

Figure 1: The standard view for all users, with Full Users' answers shown and New and Restricted Users' answers collapsed.

the user's answers for which the above inequality is true. If $q > q_{min}$ (the user is a "good" contributor), then the user attains Full User status. We feel that this is an accurate measure based on research in [2], which shows that long-time users of a community naturally tend to show preferential treatment to users who have a high probability of becoming long-time users themselves. We extrapolate this to suggest that Full Users can accurately predict other potential Full Users, and will rate them appropriately.

Users in the New or Restricted User classes do not have the ability to vote on answer quality. We justify this design choice based on the findings in [3], which suggest that users with strong reputations within a community provide more useful ratings. However, these users are still able to see the rating tools; they are simply disabled for these users. This is intended to act as a reminder that the user needs to contribute better content if he wants to receive full use of the application.

We have explicitly chosen not to provide a rating mechanism for questions. All users, regardless of classification, have equal ability to ask questions (within whatever boundaries the application sets), and asking questions has no impact on one's status. This is to try and avoid barriers to entry for new users, who are likely using the site primarily to ask questions.

3.2.4 *Automatic Scoring*

Because we assume that Full Users are representative of the community's desired users, the rating system should be very accurate, even given variance in opinion and taste. However, there is a more significant issue, which is data sparsity. Data sparsity would be a problem in any case, but restricting which users can rate content just makes the problem worse. In under-represented subsections of the site, there simply may not be enough ratings for content to meet the minimum rating threshold. Even if there are enough users, there is also the time delay to consider, specifically the time between the posting of content and the point at which it has enough

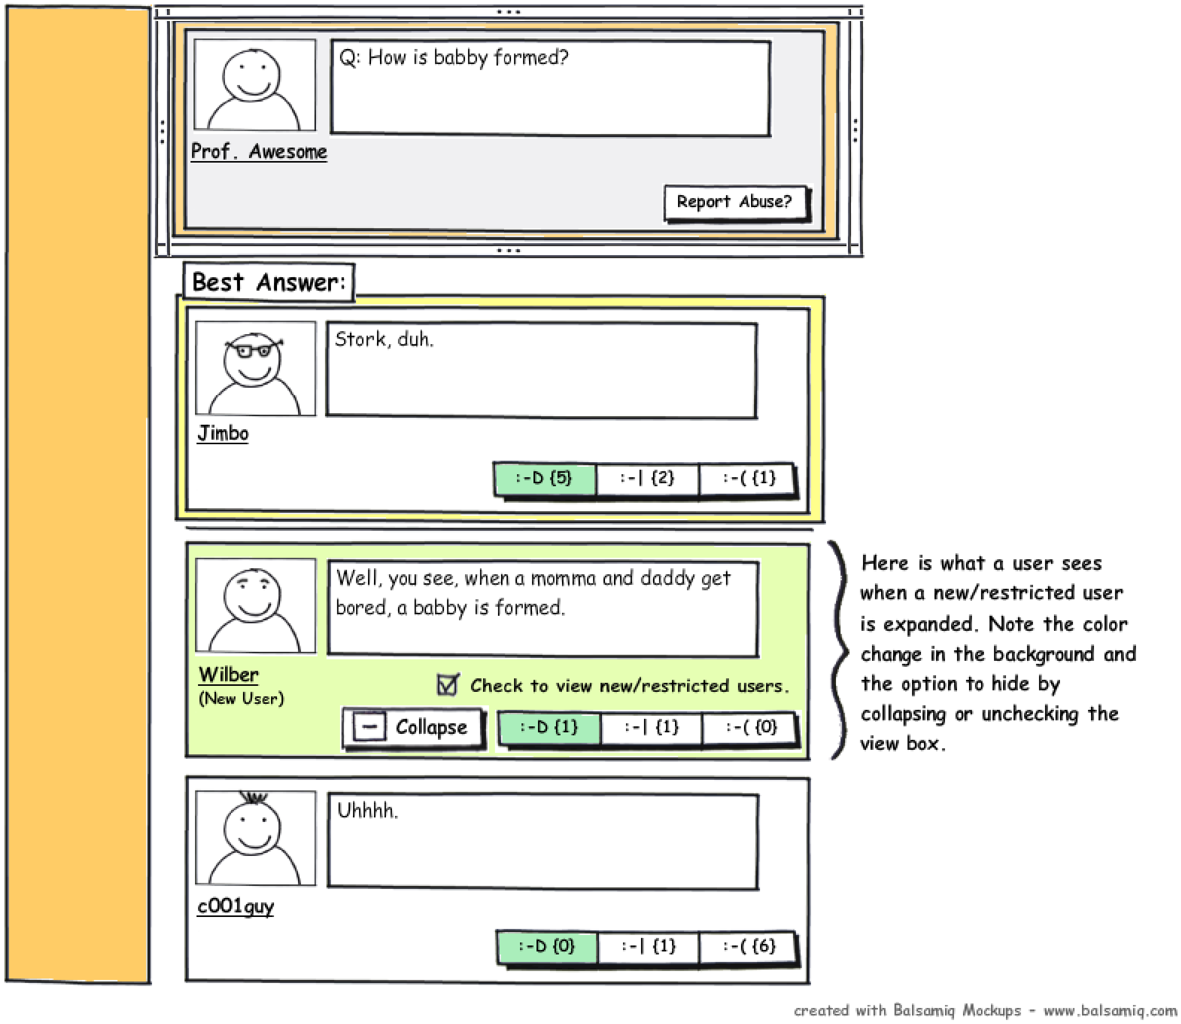

Figure 2: An expanded view (the "moderator view"), showing the answers provided by New and Restricted Users. Users have the option of checking the provided checkbox so that all answers are always expanded.

(r_{min}) ratings. Finally, it is especially hard for New Users to receive the ratings they need to be promoted to Full Users, since their answers are hidden by default.

As a fallback solution, we propose a mechanism that automatically compares new content to older content and attempts to make an intelligent guess at what ratings it is likely to receive. This type of automatic scoring has been shown to be effective on Slashdot [8] and on Yahoo! Answers [7], so we can reasonably infer that it will work well in our system as well.

We propose that, for this system, the automatic scoring model used in [8] is most appropriate for our purposes. This model is designed specifically to be language-independent and to take into account the linguistic standards of the community. Other than minor modifications (the model was designed for forums, so references to "forums" and "threads" would have to be replaced with "categories" and "questions"), we feel that the model as it is proposed in [8] performs the task that we expect of it.

Again, this solution is primarily intended to fill the gaps left by time-delay and data sparsity issues. Once a piece of content has received r_{min} ratings, this score can be ignored and the actual user rating can be applied.

4. ADVANTAGES

There are a number of advantages to using the system that we have proposed.

- As suggested earlier, this system helps create community inertia. Thus, if a community is at a state that its administrator feels is desirable, he can implement this system. Older members will quickly be divided into their classifications, and then new users can be filtered as they join.
- Negative behaviors are discouraged. If a user posts content that his fellows disapprove of, it will be downvoted and he will be classified as a Restricted User. While this does not eliminate his ability to use the

site as a question-asking tool, it will reduce the user's tendency to provide low-quality answers.

- Positive behaviors are encouraged. If a user posts content that other users feel is valuable, he will be given full powers to use the site and to be visible. By creating a distinction, Full Users are more likely to appreciate the power that they have that others do not, and they will be encouraged to maintain it.

- An inertial system as we have described can significantly reduce the number of users who contribute strictly based on the "gimmick" factor of the site. If a user has no intention to contribute more than the minimum number of answers required to become a Full User, he will probably not bother to submit answers at all. While this obviously reduces the number of potential answerers, it ensures that those who do answer questions are those with a genuine, long-term interest in doing so.

- It has been shown that when long-standing community users are given the ability to decide which new users will be promoted to "veteran" status, it enhances the focus of interaction into more manageable units [2].

5. WEAKNESSES

As with any system, the proposed solution is not perfect, and it is only proper that we explicate its weaknesses.

- While this system is easily added in an environment that is already stable, it is far more difficult if the environment is in an undesirable state or is brand new. The process of "seeding" the community and initial datasets such that the correct standards are created is very complex.

- Because this system is designed to discourage users who do not provide content that meets the standards of the community, there will by definition be a higher rejection rate than a comparable site without this system. However, we contend that, in many cases, a smaller number of quality contributors is preferable to a large number of average users, and this system is intended to be used when that preference outweighs the desire not to turn people away.

- By the same argument, this system also restricts growth. While we have striven to make the barriers for entry as minimal as possible, they still exist and as such growth rates will be lower than if the system was not used. Again, one can argue that this is preferable.

- This system supports community inertia, but of course it cannot guarantee it. The quality of the community can still degrade slowly, as community standards shift downward. This is arguably desirable behavior, as the standards match those perpetrated by the *community* rather than its administrators. If rigid, absolute standards are desired, the system would need to be modified to accommodate them.

- As discussed previously, this system is applicable only to communities where user-submitted content can be measured by some quality metric. In systems like MovieLens (content based primarily on opinion) and Facebook (content based on self-identification and social networking), no such metric exists, so this system would not be useful and would probably be detrimental.

- There is the obvious potential for users who fail to meet content standards to simply make new accounts that do not have negative ratings. This is a bit of a pain, as it might artificially inflate the number of accounts compared to the number of actual users. However, since New Users and Restricted Users have identical privileges, the only real motivation to create a new account is if you intend to try and reform and provide higher-quality content, in which case there is no reason why we would not want the user to have the ability to start fresh.

6. DIFFERENCES FROM EXISTING SYSTEMS

There are two existing systems which are most similar to the system we propose: Yahoo! Answers and Slashdot.

6.1 Yahoo! Answers

Yahoo! Answers (`http://answers.yahoo.com/`) is by far the most popularly-used question-and-answer community currently operating. It essentially uses an open-door approach to attract as many users as possible, which as a result has made it very large. Yahoo! Answers was one of our launching points in trying to determine how popular question-and-answer implementations could be improved by filtering users. While the core functionality is the same, the improvements made by our system illuminate the differences:

- Philosophically opposite of our system, Yahoo! Answers places almost no restrictions on its users. Yahoo! Answers's user base is extremely large, and that base would probably be severely reduced were restrictions placed on users who do not contribute at the desired level of quality. Our system proposes these restrictions early, so that the community norms evolve such that restriction of low-performing users is expected and desired.

- Yahoo! Answers uses a point system. Users receive points by answering questions and spend points to ask questions. Our system has no such mechanism. We feel that question-asking will be naturally restricted simply by virtue of the fact that the community is being restricted. If users are expected to provide high-quality answers to questions, we expect that users will also tend to provide higher-quality questions simply to match the overall tone of the community. We also do not wish to provide too much incentive for simply answering a question regardless of quality, as we expect that to simply inflate the number of low-quality answers (as can be observed on Yahoo! Answers).

- Yahoo! Answers has a simple thumbs-up/thumbs-down rating system for rating answers. All users have equal access to this rating system, unlike our system in which rating is restricted to Full Users so that the rating most accurately reflects community inertia. We also have included a neutral rating, which we feel allows for a more accurate rating distribution.

- In the Yahoo! Answers community, users gain levels along a scale, indicating their level of participation. However, this level of participation does not necessarily correlate with value to the community, as a user can gain levels by providing a large number of unhelpful answers. In our system, there is no level system; either a user is considered valuable or he is not. We feel this will both suppress lower-quality contributions as well as eliminate undue competition between Full Users.

To summarize, the main difference between Yahoo! Answers and our proposed system is that our system imposes restrictions on users to try to isolate those who meet community standards; Yahoo! Answers, on the other hand, has no such restrictions.

6.2 Slashdot

Slashdot (`http://slashdot.org/`) is a news site that focuses primarily on technology and open-source issues. Slashdot's community is the set of users who choose to comment and discuss the various news stories. This forum-style discussion is somewhat different in structure from question-and-answer, but there are many overlapping needs in terms of contribution quality. As a result, Slashdot has developed a very complex set of rules that determine post visibility and user reputation. While we did not develop our system with Slashdot in mind (as mentioned previously, neither of the authors were familiar with Slashdot's moderation system prior to designing the system), in trying to improve on classical question-and-answer models we have developed mechanisms that are quite similar to some of those used by Slashdot. This is likely the result of similar design principles; namely, trying to develop systems that filter content in such a way as to make the most useful content the most visible. However, while many of the underlying principles are the same, there are again some key differences:

- Slashdot filters content; content is visible based on whether it meets some standard based on their automatic scoring mechanism and user feedback. Our system filters users; users are classified based on the measurement of their content, and that content is then visible as a result of the user's classification.
- Slashdot, like our system, gives users the ability to permanently expand hidden posts so that they can be moderated. However, on Slashdot moderators are chosen randomly; in our system they are chosen by virtue of their classification. In addition, we have designed the moderator view in such a way that we expect those who take on the moderation task to be able to provide large amounts of feedback, rather than the limited amount of feedback allowed under the Slashdot system.
- Slashdot uses their automatic scoring mechanism to provide a "seed score" for posts; other users than modify this seed score up or down. In our system, the automatic score is hidden from users so as not to unduly bias ratings. The score is only used until the minimum number of votes has been reached, at which point it is discarded. Our belief is that, given the goal of community inertia, the perceived value of a post given its author is more important than the absolute value of a post.
- Slashdot uses a "karma" system, which essentially acts as a gradient measure of a user's value. We have chosen to have a strictly defined classification system with a very deliberate ceiling. Once a user is a Full User, there is no motivation to go any higher. We have made this distinction in order to reduce the competition inspired between users, which could easily lead to unnecessary content.
- Our system intentionally isolates and distinguishes New Users, in the hopes that Full Users will attempt to provide feedback and nurture them, particularly those acting as moderators.

The most significant difference is the distinction between *content-based* filtering and *user-based* filtering. Our system is designed to create a community of users who are relatively congruent. Slashdot, on the other hand, is designed to create a set of *content* that is relatively congruent. While similar mechanisms can be used to achieve both goals, the differences are crucial. (Which approach is superior is an open question.)

7. EXTENSIONS

As with any system, this system as we have described it could be extended in a number of ways.

- While we have designed the system around the concept of a question-and-answer application, it can be applied to many different environments. Discussion forums, for instance, would require only the removal of the concept of a question (which is not really relevant to user classification anyway) for the system to work. In fact, any application where the primary user contribution is some form of text post would require very little modification. The system could also be adapted for other domains, such as a knowledge repository like Wikipedia (grant certain privileges based on edit rates and quality) or a photo-sharing site like Flickr (make photos from Full Users more likely to be seen by search/browsing).
- The automatic scoring process is designed to make calculations based on comparisons to older content; however, it is worth noting that this concept is very similar to predictions in a recommender system. It might be possible to use recommender techniques to predict content scores rather than the methods we have proposed.
- Rather than using predominantly user feedback with automatic scoring as a fallback metric, it would be conceivable to modify the system so that it takes both into account all of the time.

- While not an extension *per se*, the thresholds used to classify users (like a_{min} and q_{min}) can be tweaked to match the needs of different communities, or even different subsets of the same community.

8. CONCLUSION

Our goal with this system is to recognize and exploit an inherent distinction in online communities. Communities that have quality-measurable content, like knowledge repositories or photo-sharing sites, are very different from quality-indifferent sites, like social networking or rating sites. Even more significant are communities where the average quality of content defines the quality of the community itself; this is most generally true in knowledge-sharing systems like wikis or question-and-answer sites. When there is a significant benefit to having a higher average quality, it is clear that a community designer or administrator should try to make the average quality as high as possible, at least given the standards of the specific case.

Our system encourages communities to maintain a particular level of "quality," or adherence to some arbitrary set of standards, by rewarding users who meet these standards and punishing those who do not. As a result, we believe that the community will adapt itself to the standards put in place, either by changing user behaviors or by filtering out users who do not meet the desired standards. "Good" users who meet the expectations of the community will be treated as full-ranking members, and will even receive some artificial elevation of their status (given the new "lower class"), which we expect to have a positive effect on their contribution rate. Users unable or unwilling to meet community expectations will still be able to use the community for its intended purpose, i.e. to ask questions and receive answers, but will have to improve their contribution quality if they want to become committed members.

In short, we feel that community inertia is an important goal for systems where contribution quality is definitive to the system's quality. By encouraging inertia, the system retains stability and its internal norms and expectation structures. The system we have designed, we believe, promotes this concept of community inertia and will have a positive effect on the overall quality of contributions and the strength of the community.

9. REFERENCES

[1] L. A. Adamic, J. Zhang, E. Bakshy, and M. S. Ackerman. Knowledge sharing and yahoo answers: everyone knows something. In *WWW '08: Proceeding of the 17th international conference on World Wide Web*, pages 665–674, New York, NY, USA, 2008. ACM.

[2] L. Backstrom, R. Kumar, C. Marlow, J. Novak, and A. Tomkins. Preferential behavior in online groups. In *WSDM '08: Proceedings of the international conference on Web search and web data mining*, pages 117–128, New York, NY, USA, 2008. ACM.

[3] M. Chen and J. P. Singh. Computing and using reputations for internet ratings. In *EC '01: Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 154–162, New York, NY, USA, 2001. ACM.

[4] A. T. Fiore, S. L. Tiernan, and M. A. Smith. Observed behavior and perceived value of authors in usenet newsgroups: bridging the gap. In *CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 323–330, New York, NY, USA, 2002. ACM.

[5] J. W. Kelly, D. Fisher, and M. Smith. Friends, foes, and fringe: norms and structure in political discussion networks. In *dg.o '06: Proceedings of the 2006 international conference on Digital government research*, pages 412–417, New York, NY, USA, 2006. ACM.

[6] M. Maia, J. Almeida, and V. Almeida. Identifying user behavior in online social networks. In *SocialNets '08: Proceedings of the 1st workshop on Social network systems*, pages 1–6, New York, NY, USA, 2008. ACM.

[7] M. A. Suryanto, E. P. Lim, A. Sun, and R. H. L. Chiang. Quality-aware collaborative question answering: methods and evaluation. In *WSDM '09: Proceedings of the Second ACM International Conference on Web Search and Data Mining*, pages 142–151, New York, NY, USA, 2009. ACM.

[8] N. Wanas, M. El-Saban, H. Ashour, and W. Ammar. Automatic scoring of online discussion posts. In *WICOW '08: Proceeding of the 2nd ACM workshop on Information credibility on the web*, pages 19–26, New York, NY, USA, 2008. ACM.