Hacker News vs. Slashdot
Reputation Systems in Crowdsourced Technology News

Robert Lehmann, Christoph Matthies
{robert.lehmann,christoph.matthies}@student.hpi.uni-potsdam.de

Abstract. In this paper we describe and contrast the reputation systems of the two technology discussion forums Slashdot and Hacker News, the former being very complex and well-studied, the latter smaller and simpler. We analyse these systems in regard to previously established design goals as well as our own criteria and political analogies. Furthermore, we explore how cheating is handled. In particular we focus on possible issues of these communities and propose solutions on how the reputation system could be improved.

Keywords: reputation systems, technology news, crowdsourcing

1 Introduction

In contrast to magazine-style technology-related news sites (e.g. news.cnet.com, heise.de) where editors write and publish stories, crowdsourced technology news sites rely on users to submit relevant items. These can be voted on by users or might be evaluated by editors of the site to decide which stories make it to the front page. Kuro5hin[1], the technology subreddit[2], as well as Slashdot[3] and Hacker News[4] are prominent examples of this system. The essence of these sites is discussion in the form of comments on stories. Users comment on the merit of the story, its implications and their thoughts. These comments are then evaluated by the community.

We chose Hacker News and Slashdot for a detailed study of the reputation systems as they share similar importance[5] and the source code to both is open source[6] and thus freely available for study.

Both sites share the concept of karma, an aggregate value representing the value of a user's submissions (stories as well as comments.) However, the reputation systems of these sites are very different in their complexity, making them an interesting case for comparison and analysing advantages and drawbacks.

[1] http://www.kuro5hin.org
[2] http://www.reddit.com/r/technology/
[3] http://slashdot.org/
[4] http://news.ycombinator.com/
[5] slashdot.org has an Alexa traffic rank of 2,024, compared to ycombinator.com which sits at a traffic rank of 2,506. (Jan 9, 2013)
[6] http://slashcode.com/, http://arclanguage.org, GPL-compatible licenses

2 Hacker News

Hacker News was started in 2007 by Paul Graham, of Y Combinator[7] fame, to counter the phenomenon of discussion forums suffering from Eternal September[8] sooner or later. It has more than 120,000 visitors per day[9] and claims to feature any content which "good hackers would find interesting."[10]

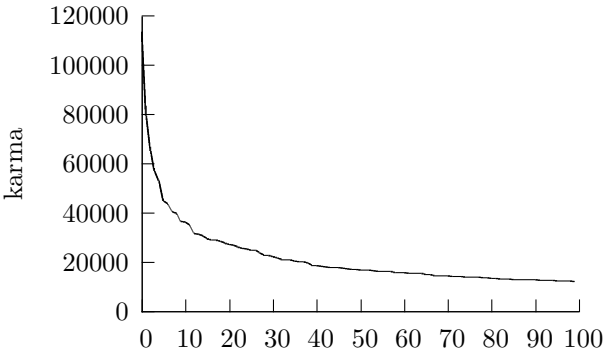

Fig. 1. Karma distribution among the top 100 posters as of Nov 13, 2012.

For every comment and story, each user can cast exactly one *upvote*. While this usually expresses endorsement of the posted content, the private *saved* site lists all comments a user has upvoted, adding a bookmarking quality to the feature. Other than increasing the overall quality of the site, there is no incentive to voting (but see below for an incentive to garner upvotes.)

Every user starts out with a karma of 1 and can not drop below zero; comments can only ever drop to a minimum score of -4.[11] For details on which restrictions apply to those calculations, cf. section 4.4.

Among the highest ranked members of the community, karma is distributed in a Zipfian fashion [fig. 1] and is only ever increasing, if at all, not decreasing. [2] They make up about 2.2 million karma (AUC.) A log-log fitting indicates that there is a total of 4.4 million karma on the site (interpolated AUC), held mainly by the top 300 members.

Users are rewarded for reaching certain karma levels by unlocking new features on the site such as downvoting. These limits are increased "over time to account for the karma inflation caused by an increased userbase." [6]

3 Slashdot

Slashdot was started in 1997 by Rob Malda as a "News for Nerds" website offering "Stuff that Matters." Slashdot is rather popular, averaging 5,300 comments daily.[12] It is also the origin of the "Slashdot effect," where linking to a smaller

[7] Y Combinator is a popular technology incubator which has also invested in Reddit and Disqus. `http://ycombinator.com`

[8] Eternal September is a reference to the perceived downfall of the Usenet when a lot of new, seemingly underqualified members joined it back in September 1993.

[9] `http://ycombinator.com/newsnews.html#5oct11` (Oct 5, 2011)

[10] `http://ycombinator.com/newsguidelines.html`

[11] See `deftem profile`, `def user-fields` and `lowest-score` in the source, respectively.

[12] `http://www.diceholdingsinc.com/phoenix.zhtml?c=211152&p=irol-newsArticle\&ID=1735911`

site from Slashdot results in an overload of the smaller site by the resulting traffic spike.

Starting from a very basic system, involving no moderation at all [3] Slashdot now uses a complex reputation system featuring two levels of moderation: the M1 level, which moderates comments and the M2 level which moderates M1 moderators.

Stories are submitted by the community and are selected by a group of paid editors to appear on the front page. Slashdot's democratic moderation system grants moderation points (which expire within 3 days) to users (M1) which they can use to moderate comments using either positive or negative adjectives (+1 or -1 to the comment score, respectively), influencing that poster's karma. The karma of a user influences the score his comments start with. A filter is employed to allow users to hide inadequate comments.

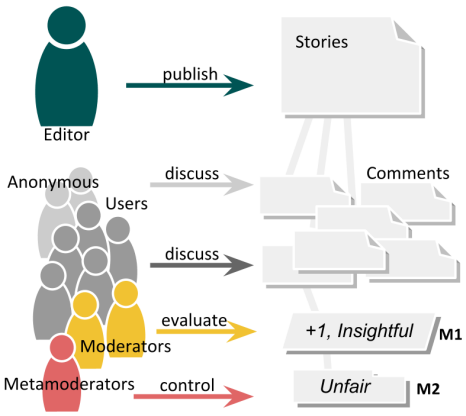

Fig. 2. Structure of Slashdot.

4 Comparison

Trust towards new users Being less complex, Hacker News allows all registered users to upvote, trusting them with a positive reputation statement only requiring a higher karma for downvoting or flagging. Because new users have a low karma to begin with, their comments will be sorted to the bottom (before they have received upvotes.)

Slashdot is more extreme in their approach: Users do not get to moderate under certain age and karma thresholds. Comments from unregistered users start at a score of 0, new users at +1 (compared to +2 for users with *Excellent* karma.)

Display of user karma Slashdot only displays karma levels[13] to the user himself, hiding it from all other users. For Hacker News, karma is publicly visible as a sum and average on profile pages and used for the top 100 leaderboard.

Display of comment scores Whereas Slashdot makes a comment's score explicitly visible by showing its score as well as the designated adjective[14] in the header of the comment, Hacker News only sorts comment threads by their scores.

[13] Per default, less than -10 karma is *Terrible*, -10 to -1 is *Bad*, 0 is *Neutral*, 1 to 12 is *Positive*, 12 to 25 is *Good* and more than 25 is *Excellent*.

[14] Either positive (*insightful, interesting, informative, funny, underrated*) or negative (*offtopic, amebait, troll, redundant, overrated.*)

Dealing with inadequate comments Slashdot allows each user to determine for themselves what they deem unqualified by making the comment filter adjustable; comments below the chosen threshold are folded or hidden. On Hacker News, comments with a negative score are greyed out.

4.1 Community Nature

Slashdot's community has always been described as cordial [1] and Hacker News has roughly the same direction: Every community member wants to promote valuable stories or comments, but that does not necessarily involve any strong feelings towards all other content posted on the site.

Audience demographics[15] indicate that both communities gravitate strongly towards single men but there are still observable differences. Slashdot's visitors are, relative the to the general internet population, pretty average. The older population (65+) seems slightly overrepresented; that belief is reinforced by the missing traffic from school locations. Hacker News attracts a well-educated, younger (18-34) crowd browsing from school or work.

These facts could be both, ex post facto or ipso facto, and do not necessarily say anything about the communities. Especially the age distribution could be skewed because Slashdot has been founded in a time where Internet was not yet a mass medium.

4.2 Design Goals for Distributed Moderation Systems

Lampe and Resnick [4, p. 7] have articulated the following four orthogonal design goals which apply to all distributed moderation systems, limit the solution space and must be traded off against each other.

Timeliness Both systems do not take measures to improve timeliness of moderation. This is well reflected in the huge and skewed delay on Slashdot. [4, p. 3]

Minimal Effort User reputation lends itself well to reduce the effort require to rate each and every comment: When in doubt (such as initially, at the time of submission,) a comment is ranked according to the poster's previous performance. Slashdot makes heavy use of this with its starting scores described in 4; Hacker News only factors in average comment scores — which are not entirely proportional with karma — when trying to evaluate a stalemate between two comments.

On Slashdot, moderation is the exception rather than the norm and is thus handled specially. On Hacker News, everybody is a moderator all the time. There is very little cognitive overhead in casting an upvote (and, once a user has surpassed the downvote threshold, for that as well.)

[15] Alexa Internet, Inc. offers such a service on `http://alexa.com/`.

Limited Influence Every moderator should have only limited influence within the reputation system; hijacking it should be hard.

For Hacker News this holds true if many votes were cast (which is easy because there is basically unlimited moderation points.) Downvoting a post with a high score is not going to change much; though evidence suggests that the global average is not too high.[16].

On Slashdot this is more of a problem due to the fact that comment scores are restricted to 7 levels (-1 to +5). It is therefore possible to demote a top comment (+5) to the lowest score (-1) with a cabal of only 6 moderators. To make matters worse, a comment demoted in such a way will be hidden by the default filter. Concerted action is unlikely due to the random election process though.

Accuracy Hacker News relies solely on the control and intelligence of the community as a whole to ascertain the value of a comment. If a user finds the ranking of a comment inappropriate he gets to respond by casting a vote in the opposite direction. This is only limited by the fact that not every user can downvote by default (see section 2.)

Slashdot introduced another whole layer of indirection (M2, see section 3.) Metamoderators rate M1 moderations, influencing their karma which in turn affects their chances of receiving moderation points in the future. Evidence further suggests [fig 3]), that only a third of all moderations deemed unfair by metamoderators get reversed by further moderations.

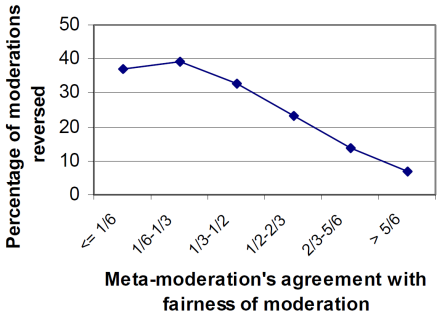

Fig. 3. Moderation reversals. [4, p. 4]

4.3 Political Discussion

Online communities tend to be their own little worlds featuring trust, friends, and foes. It therefore seems fitting to compare their nature to existing political systems. Both systems try to be meritocratic (the rule of the able), granting those who have proven their worth to the community more power and/or exposure.

Slashdot's moderation system seems very close as only regular, long-term users with high karma are eligible to moderate. Due to the involved complexity there are similarities to more obscure political systems as well:

[16] Among the top 100 users, the average comment score is only 4.95, which, funny enough, is even below Slashdot's +5 rating. (Jan 14, 2013)

- **Cronyism** (partiality to long-standing friends) is in effect as the user ID is shown on every comment and could thus bias ratings.[17]
- **Gerontocracy** (rule of the oldest) applies since only the oldest members become metamoderators with elevated rights.
- **Demarchy** (rule of the randomly selected) manifests itself through the random selection process.

Hacker News compares to a meritocracy in the respect that users with a high karma are granted bonuses (such as the ability to downvote.) Both reputation systems, however, lack the possibility of a citizen-driven impeachment process — there are no ways in which users can get rid of moderators in a democratic way (as is the case with Wikipedia's noticeboards.)

4.4 Cheating and Dealing with it

Protection mechanisms Hacker News silently blocks votes if it detects fraudulent behaviour. Its *karmabombing* prevention will only accept a limited number of downvotes between two users, thereby avoiding personal vendettas. The downvote ratio rejects excess downvotes; sockpuppet detection tries to filter dummy accounts.

Slashdot focusses on filtering posts automatically, such as rejecting posts from users who have recently been downvoted too much (their "trollishness" filter.)

Manual intervention Both systems leave leeway for the site administrators to punish malevolent individuals.

There are two scripts in the Slashdot code base called *bitchslap* and *modslap*, significantly reducing karma and revoking comment scores or moderation points respectively.

Hacker News offers its administrators to *blast* submissions (removing them and banning the poster) or *nuke* stories (blasting them and banning all future submissions from that host.) Users are never banned from visiting the site but are instead *hellbanned*, where they can participate normally but no other user will ever see their contributions.

5 Improvements

A problem for online discussion forums is that of *buried treasures*, i.e. content that should have a high score but does not. [4]

[17] User IDs are assigned sequentially, thus a low ID is highly valued and signals a veteran user (in fact, Malda is #1.) For Slashdot's 10 year anniversary a 3-digit Slashdot user ID was auctioned off in favor of the EFF. `http://w2.eff.org/effector/20/43.php`.

5.1 Attention Redirection

Lampe and Resnick proposed to adjust Slashdot's comment filter to redirect the attention of moderators to new as well as recently edited comments. Inaction of moderators infers that the comment is correct at its current score.

The same approach can be applied to Hacker News: Recent comments could be bumped to the top of the page, making them more visible and inviting moderation. The comments would need to appear without context as otherwise a whole, possibly long and, in terms of moderation, unimportant thread is unduly promoted. This proposal would increase the accuracy of moderations at the cost of more effort on the users side.

5.2 Power to Metamoderators

Currently on Slashdot, when a metamoderator (oldest 7,5% of users) deems a M1 moderation *Unfair* (as opposed to *Fair* or *Neutral*) the moderator loses karma but the moderation is not automatically reversed. By doing so, the accuracy (biased towards metamoderators) would improve, without causing any more effort for moderators.

Another issue with metamoderation is that there is no learning effect for moderators; making a public example of the worst (and best) moderations could help teaching community standards.

5.3 Karma Decay

Several attacks on Slashdot use a spray&pray approach to game the system: by posting low effort comments from different accounts, chances are — with negative ratings weighing much softer than positive ones — that one of them will eventually go over a certain amount of karma. Once an account has reached the *Excellent* threshold it is much more visible and has elevated rights. Even legitimate posters often relax and deprave after that hurdle.

Karma seems to be a largely negative factor and removing it altogether would shift the comment ratings from the user reputation to its individual merits

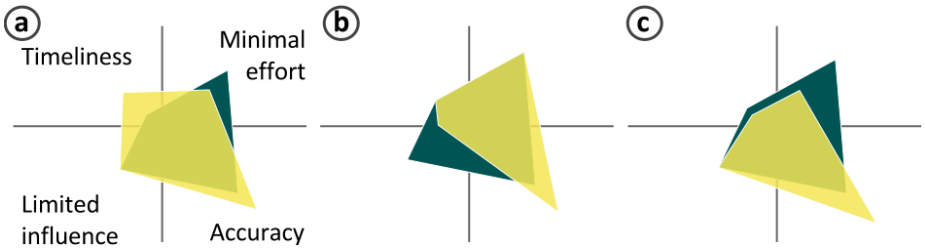

Fig. 4. Tradeoffs concerning Lampe's design goals (see section 4.2) when implementing (a) attention redirection, (b) power to the metamoderators and (c) karma decay.

(increasing effort but improving accuracy.) In order to encourage continuous positive contributions and still reward it, decaying karma over time is a good compromise.

6 Conclusion

The systems we have analyzed try to solve the problems of online communities in very different ways.

Slashdot tries to counter misbehaviour with different mechanisms, leading to an overall increase in complexity. It can be described as an experiment to find the best way to (more or less) democratically deal with a large community while trying to preserve quality posts. Evidence suggests [5] that while content published on the frontpage is good, this might not be caused by the reputation systems but rather the editors. Discussion is rated as only slightly above average, which further fuels the hypothesis that the system's complexity might overshoot the goal.

Hacker News, in contrast, is a seemingly simple system. While there is no conclusive evidence that this is a good thing per se, it seems to work quite well. Changes — e.g. raising the karma thresholds for rewards such as downvoting or flagging — are committed by the operators at their sole discretion and magically improve the system's performance.

References

1. Randy Farmer and Bryce Glass. *Building Web Reputation Systems.* Yahoo! Press, USA, 1st edition, 2010.
2. John Graham-Cumming. Forty two days of hacker news karma, September 2007.
3. Audun Josang, Roslan Ismail, and Colin A. Boyd. A survey of trust and reputation systems for online service provision. *Decision Support Systems 43(2)*, pages 618–644, 2007.
4. Cliff Lampe and Paul Resnick. Slash(dot) and burn: distributed moderation in a large online conversation space. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '04, pages 543–550, New York, NY, USA, 2004. ACM.
5. Clifford A. Lampe. *Ratings Use in an Online Discussion System: The Slashdot Case.* dissertation, The University of Michigan, 2006.
6. Jacques Mattheij. The unofficial hacker news FAQ, September 2011.