

Beyond Wikipedia: Coordination and Conflict in Online Production Groups

Aniket Kittur Robert E. Kraut
Carnegie Mellon University
5000 Forbes Ave, Pittsburgh, PA 15213
{nkittur, robert.kraut}@cs.cmu.edu

ABSTRACT

Online production groups have the potential to transform the way that knowledge is produced and disseminated. One of the most widely used forms of online production is the wiki, which has been used in domains ranging from science to education to enterprise. We examined the development of and interactions between coordination and conflict in a sample of 6811 wiki production groups. We investigated the influence of four coordination mechanisms: intra-article communication, inter-user communication, concentration of workgroup structure, and policy and procedures. We also examined the growth of conflict, finding the density of users in an information space to be a significant predictor. Finally, we analyzed the effectiveness of the four coordination mechanisms on managing conflict, finding differences in how each scaled to large numbers of contributors. Our results suggest that coordination mechanisms effective for managing conflict are not always the same as those effective for managing task quality, and that designers must take into account the social benefits of coordination mechanisms in addition to their production benefits.

Author Keywords

Wiki, Wikipedia, coordination, conflict, social computing, collective intelligence, distributed cognition, collaboration, online production.

ACM Classification Keywords

H.5.3 [Information Interfaces]: Group and Organization Interfaces – Collaborative computing, Computer-supported cooperative work, Web-based interaction, K.4.3 [Computers and Society]: Organizational Impacts – Computer-supported collaborative work

General Terms

Measurement

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CSCW 2010, February 6–10, 2010, Savannah, Georgia, USA.
Copyright 2010 ACM 978-1-60558-795-0/10/02...\$10.00.

INTRODUCTION

Thousands of volunteers to online production groups have created artifacts that fundamentally affect our everyday lives, ranging from software that runs the internet to the largest encyclopedia of human knowledge ever written. One particular kind of production group, the wiki, has become an enormously popular way for people to generate, share, and make sense of information in a distributed manner. Wikipedia, an online encyclopedia driven by volunteer contributions, is a case in point: over two million articles have been created in the English language version alone. However, while Wikipedia is by far the most popular example of the wiki paradigm, it is by no means the only example. Wikis have been used in the enterprise (e.g., socialtext.com), in the classroom (e.g., [4][11]), in intelligence analysis (e.g., Intellipedia), in science (e.g., Scholarpedia, OpenWetware), and in many other contexts. More than 10,000 wikis use the popular MediaWiki software on which Wikipedia is built [16], some of them with very active user populations (e.g., the wiki for the computer game World of Warcraft has over one million users [18]).

This widespread use of wikis has led many researchers to study this method of online production. Most of these studies have focused on a single wiki, Wikipedia. The reasons for this are quite understandable: Wikipedia has achieved a tremendous level of popularity, consistently being ranked in the top 10 most trafficked sites on the internet [17], and provides a model for understanding successful distributed collaboration using the wiki medium. Other factors promoting its popularity in the research community include: that its entire editing history has been made publicly available; that it is available in many different languages; that it includes rich structures such as categories, links, quality assessments, and templates that can be profitably mined; and that its contributors can be relatively easily contacted for interviews and surveys.

However, a key unanswered question is how well findings from research on Wikipedia generalize to the use of wikis in other contexts. There are plausible reasons to believe that many findings could be highly context-specific. The Wikipedia community is eight years old and has evolved a complex set of explicit policies and informal norms to govern how content is produced and evaluated, how editors interact, how disputes are resolved, and myriad other details

that are critical to the functioning of the community. It also is comprised of a core group of highly dedicated editors, many of whom spend dozens of hours per week volunteering their efforts. They have developed a complex social network and informal reputation system. In contrast, most other wikis are relatively new and small, and thus may lack the complex social machinery that underlies the success of Wikipedia.

Just because two wikis are based on the same technical platform does not mean that they are similar in other ways. Factors such as the goals of the community, the norms governing behavior, the policies agreed upon, the nature of the work, and even the characteristics of individual contributors all may play a role in how production is organized. Because of these variations across wiki communities, the principles that have guided the success of Wikipedia may not generalize to these other wikis.

On the other hand, it is possible that some common processes are involved in the growth of all wiki communities. As communities grow, they encounter common coordination problems. As small groups grow, the informal norms once known by everyone need to be codified into explicit rules and enforced. Informal control structures are centralized into formal hierarchies. Subgroups are formed as the larger size undermines cohesion. Groups' needs for coordination and communication grow and are accompanied by process losses (i.e., the efforts of the group are lower than what the individuals could ideally accomplish [34]). While these group coordination processes are not specific to wikis, it is plausible that they underlie the evolution of wikis in the same way they do other kinds of online and offline groups.

In this paper we examine the generalizability to a large sample of wiki production groups previous findings from Wikipedia. We focus on comparisons to Wikipedia because it has been the subject of much research and thus rich quantitative and qualitative data for comparison. Within this context we limit our investigation to the development and effectiveness of coordination mechanisms, though we acknowledge that there are many other factors that could be of interest in future work. We also investigate the conditions under which conflict arises, and the effectiveness of coordination mechanisms in managing conflict at different scales.

Coordination in wiki work

Researchers have consistently found that Wikipedia's success depends upon more than a production process in which many independent contributors each do a little bit of work, causing the coverage and quality of the resulting articles to slowly and consistently grow [21]. To be successful, editors need to coordinate their work in both broad and detailed ways. By coordination, we mean "the act of managing interdependencies between activities" [28]. For example, editors must have a common view of what an encyclopedic article is. They must agree across the articles

in a genre (e.g., civil war battles) as to the needed elements. When writing a particular article, they must identify a structure that presents material with a logical flow, show the relationship between sections and paragraphs, produce grammatical sentences that are easy to read, include material that is at the appropriate level of detail, verify facts, show a neutral point of view and cite other sources appropriately. While coordination is important in all production groups, the coordination challenge is greater in many online production environments, including wiki communities. The fundamental problem of online production is managing interdependencies under adverse conditions imposed by the virtual environment. Virtual organizations may involve many more people than is common in a conventional organization. These contributors often have a tenuous and poorly defined relationship to the virtual organization, rarely undergo common training, often fail to share a common context or background and work without strong managerial direction.

According to Watson-Manheim and colleagues [40], collaborators in these virtual environments are seeking ways resolve their discontinuities by developing a common view of their task, their work processes and the work product. They have many coordination methods available to develop this common view, each with different strengths and weakness and suited for different production tasks. Here we examine the three major coordination mechanisms emphasized in classic research and theory on coordination in groups and organizations. First, direct, peer-to-peer communication is perhaps the most basic coordination mechanism. Collaborators develop a common vision by discussing the issues. Communication is an example of organic coordination or coordination by mutual adjustment, which organizational theorists argue is needed under conditions of uncertainty [28]. Second is group structure, consisting of role differentiation, division of labor and formal and informal management. In this more hierarchical view, the collaborators develop a common view of what is needed by relying on others, typically managers, to tell them. A third coordination method involves shared mental models. Shared mental models for task assignment are beliefs held in common among a set of editors about what should be done and who should be doing it. Shared mental models represent a blend of standardization and communication that allows editors to coordinate task assignment without explicit communication [31]. Although group members can develop shared mental models by spending enough time in a common environment (Goodman & Leyden, 1991), shared mental models can also be imposed on them through the use of standards, guidelines or policies. Below we discuss each of these coordination strategies in more detail as they relate to wiki work.

Communication

As discussed above, one important coordination mechanism found in virtually all organizations is interpersonal

communication [28]. In Wikipedia, direct communication is ubiquitous. Editors often discuss potential changes with each other on dedicated discussion (or “talk”) pages for articles before implementing them in the article itself [39]. This form of communication is especially prevalent while an article is in a formative stage: on average more than half of the edits to an article during its first week are to its talk page rather than to the content of the article, with this number quickly dropping as the article develops [23]. One possible reason for this is that the effectiveness of communication as a coordination mechanism (as measured by its effect on improving article quality) is strongest early in an article’s life, and drops quickly as the number of contributors increases. The effectiveness of communication as a coordination mechanism may also interact with factors such as who is communicating. Discussion among anonymous editors is associated with increased conflict, while discussion among registered editors is associated with reduced conflict [22].

Less studied, but also important, is communication between editors on users’ talk pages. Each editor can have their own user page, each of which has a corresponding talk page. Some of the edits to these pages are spill-over from article talk pages, with detailed discussion and negotiations about the article happening on the involved parties’ user talk pages. However, these pages additionally support different kinds of coordination that are not necessarily tied to the content of specific articles. For example, editors engage in activities such as requesting help from each other, notifying each other about shared interests (e.g., if an article that they both edited is being proposed for deletion), or discussing standards that should apply to multiple articles.

Group structure

Another mechanism for editors to coordinate is group structure. Workgroup structure is an important coordination method in offline organizations and can include mechanisms such as managerial control or division of labor [27][37]. In many peer production communities, a common workgroup structure is to have a small core of leaders set direction and do a disproportionate amount of the work. For example, the top four percent of the developers in the Apache server project contributed 88% of new lines of code [30]. Having the work concentrated in a small group of contributors can have a number of benefits. Communication needs are reduced, since many topics require discussion only by the core group. Similarly, consensus building can be easier since fewer stakeholders are involved. The core group is also more likely to develop a common view or shared mental models of the work that needs to be accomplished even without explicit communication. The core can also act as leaders by setting direction and creating a framework for the article, which allows more peripheral editors to contribute more effectively [23].

Policy and procedure

In addition to direct communication, coordination can occur through the development and use of policies and precedents. Wikipedia has developed an extensive policy system to govern areas ranging from the process of becoming an administrator to what constitutes an encyclopedic article to methods for managing conflict and consensus building to defining plagiarism [2][5][6][12].

CHARACTERIZING WIKI COMMUNITIES

In this paper we go beyond Wikipedia to examine the development of coordination and conflict in thousands of other online wiki communities. To do so we downloaded the full transaction history data from 6811 wikis publicly available from Wikia, an online wiki hosting platform. Although few Wikia wikis are explicitly encyclopedias, most aim for consolidating and sharing distributed information, such as Wookiepedia, a wiki with over 75,000 Star Wars related articles, or the Psychology Wiki, with more than 25,000 articles about psychology. The histories we downloaded included all revisions to the wikis up to November 2007, comprising 5,104,939 total pages and 23,411,652 total revisions made by 1,134,824 editors.

Using Wikia-hosted wikis provides a number of benefits for making comparisons to Wikipedia. All 6811 wikis are based on the same MediaWiki software platform that is used by Wikipedia, and are very similarly structured with separate namespaces for article pages and their associated discussion pages, user and user discussion pages, policy and project pages, and more. Furthermore, many of the norms, policies, and conventions used in many wikis are imported from Wikipedia itself, ranging from simple conventions such as marking reverts and vandalism in revision comments to high-level policies such as maintaining a neutral point of view.

Analyses

We provide a brief summary of the characteristics of wikis in order to provide an overview before discussing coordination and conflict. The distribution of users, pages, and revisions for all wikis is shown in Figure 1. The data are plotted on a log-log scale, and reflect a highly skewed distribution of contribution to wikis. For example, the top 10 wikis among the 6811 sampled account for over 35% of all revisions made to all wikis; the top 20, nearly 50%. This skewed contribution pattern is consistent with other online production groups, such as the skewed popularity of open-source software projects in SourceForge [8].

Figure 2 characterizes the lifespan of projects, showing the distribution of wikis ordered by the number of months from the first to the last edit. This distribution is also skewed, but less so than the previous figure (note that the y-axis is not log-scale, implying a distribution closer to exponential than power law). Still, we see a consistent pattern that very few communities survive to become highly successful with continued activity.




Figure 1. Distribution of pages, edits, and users for all wikis.

COORDINATION

We now turn to examining coordination in wiki communities, focusing on the types of coordination discussed earlier: (1) direct communication between participants in article talk and user talk pages; (2) group structure, and (3) policy and procedure. For each type of coordination we briefly review past findings in the context of Wikipedia and then present the results of our analyses on the larger set of wikis that we sampled in comparison.

Communication

Although communication has grown substantially in Wikipedia, the proportion of article talk edits has remained surprisingly stable. As described in [22] and shown in Figure 3, the proportion of article talk edits in Wikipedia has remained around 6-8% of all edits for most of the lifespan of Wikipedia¹. Meanwhile, user talk pages have been one of the fastest growing areas in Wikipedia, reflecting the growing importance of the types of coordination done on these pages. User talk grew by a factor of 78 from 2003 to 2005 [39], increasing to account for approximately 6-7% of all edits by 2008 [22]. We include the growth of user talk edits in Figure 3 to promote comparison between article-specific coordination and the coordination on user talk pages. Interestingly, the prevalence of article talk remains relatively constant over time, while the prevalence of user talk increases. An intriguing potential explanation for this pattern could lie in the types of coordination that article and user talk represent. If article-specific coordination scales with the number of articles, while user-specific coordination scales with the number of users, this could explain the difference in the growth curves of the respective kinds of coordination since the number of users has grown faster than the number of pages (see Figure 5). Another related explanation is that user coordination should scale super-linearly, since adding

¹ While there was a large jump in this proportion in 2001, it is difficult to make too much of this as the total number of edits during that year was quite small.

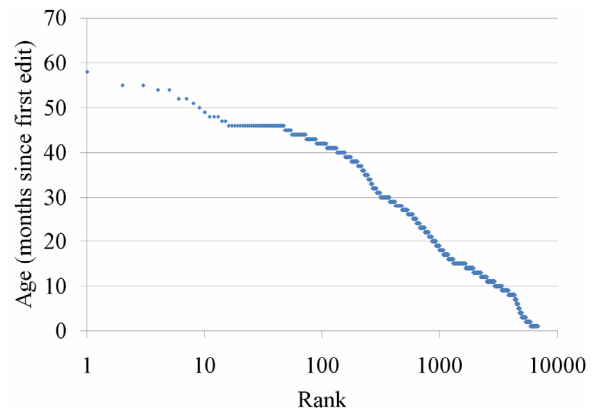


Figure 2. Distribution of ages (measured in months from first to last edit) for all wikis.

another contributor to a group increases the number of possible communication pairs exponentially, while article-specific communication is typically grounded in a specific topic which a limited number of contributors are involved in.

In Figure 4 we show the growth of communication in the diverse sample of wiki production systems we downloaded, broken out by whether the communication was on article talk or user talk pages. In our analyses we only include data from months in which at least one edit was made to an article in a wiki. There are a number of interesting results which we discuss below. First, the pattern of results is generally consistent with the growth of communication in Wikipedia. In both Wikipedia and the larger sample of wiki communities, the amount of communication ramps up considerably over time. The actual magnitudes are fairly comparable as well, with communication in later months accounting for about 10-15% of all edits in both Wikipedia and other wikis. Second, we see that after an initial spurt, the proportion of article talk remains relatively constant for the wiki sample, remaining largely between 5%-7% of all edits. This is quite similar to Wikipedia in both pattern and magnitude. In contrast, user talk has grown significantly over time to about 5% of all edits; in Wikipedia this number is slightly higher, between 6%-8%, but with a similar pattern. This is also consistent with our earlier hypothesis that user talk scales with the number of users while article talk remains relatively constant.

Taken together, these data suggest that the pattern of communication and coordination found in Wikipedia is not an anomaly; instead, coordination is a necessary task that a variety of wiki production communities invest an increasing amount of resources and effort into over time. Furthermore, it is remarkable that not only the general patterns appear similar in terms of growth and stabilization between Wikipedia and other wiki communities, but the *amount* of communication is also similar. This suggests the possibility that there may be common laws that govern coordination across a variety of community structures and domains which researchers may be able to discover through mining




Figure 3. Communication on article talk and user talk pages as a proportion of all edits in Wikipedia.

the rich archival history of thousands of online production groups.

Policy and procedure

Policies and procedures have not been quantified and studied to the same degree that communication has in Wikipedia. One way to quantify the importance of policies to examine Wikipedia namespace and its corresponding talk pages, where policies and procedures are edited and discussed. Butler et al. [6] identified 44 wiki pages in the “Wikipedia official policy” category and 248 in the “Wikipedia guidelines” as of 2007. The Wikipedia namespace also includes a variety of other kinds of coordination-oriented pages, such as requests for assistance, noticeboards for editors to bring issues to the attention of administrators, and discussions of whether to delete particular articles. The Wikipedia namespace is also home to Wikiprojects, which are subgroups in Wikipedia oriented around a certain topic or task (e.g., “Biology” or “Copyediting”) that organize and coordinate their members to improve topic-related articles or complete topic-related tasks [25]. Viegas et al. [39] provided an early snapshot of the growth of the Wikipedia namespace, in which many of the policies and procedures are edited and discussed, which grew by a factor of 68 between 2003 and 2005. A substantial portion of the work in the Wikipedia namespace is done to Wikiprojects, which increased to over 1% of all edits in 2007 [25].

We quantified policy and procedure work for each of the wikis in our sample by looking at the wiki namespace (the analog of the “Wikipedia” namespace) and its corresponding talk namespace. Although there is significant variability in the naming conventions and policy structures used across wikis, we believe looking at work within these namespaces is a reasonable simplifying assumption since both Wikipedia and the wikis examined here are based on the MediaWiki software namespace standards. The proportion of edits made to these namespaces for Wikipedia and all wikis are shown in Figure 6 and 7, respectively. In contrast to our findings on communication, as the figures

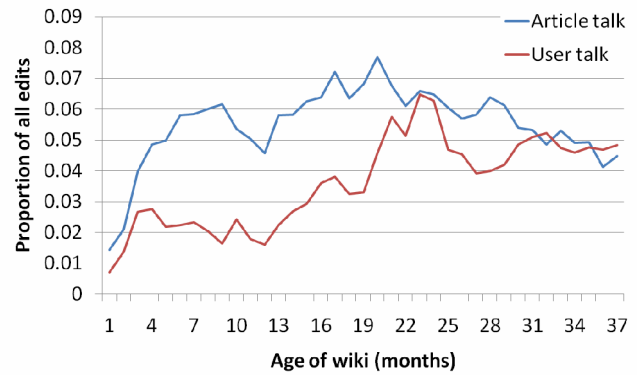


Figure 4. Communication on article talk and user talk pages as a proportion of all edits; aggregate data from all wikis.

indicate there is a striking difference between Wikipedia and other wikis. While policy and procedure edits grew up to 10% of all edits made in Wikipedia, their influence remains relatively small in other wikis, hovering around 3%. To determine whether we would see an increasing influence of policy and procedure edits just in wikis that had grown to a very large size, we redid our analysis using only wikis that eventually accumulated over 5000 edits, but the results were substantively identical. Another possible explanation for this difference could be the large number of articles that get deleted from Wikipedia every day; these “Articles for deletion” candidates get discussed within the Wikipedia namespace and could thus inflate its numbers. However, looking at edits to “Article for deletion” we find that substantial edits do not begin until late 2004 – when there was already a large proportion of policy and procedure edits – and only account for about 2% of all edits. Thus it appears that in the heavy use of policy and procedure, Wikipedia may indeed be an anomaly.

Group structure

Like many other online production groups, Wikipedia also has a small core of leaders who make most of the contributions, whether measured by edits or by words [21]. Although these core contributors are not managers, in that they have no formal authority to assign work, they set the direction of the work by their own efforts. Prior research shows that articles in which the workgroup structure




Figure 5. Average number of editors per page over time for all wikis.




Figure 6. Wikipedia and Wikipedia talk namespaces as a proportion of all edits in Wikipedia.

includes a small core of leaders are more likely to increase in quality than articles in which work is evenly distributed amongst contributors, and this effect becomes stronger as the number of involved contributors grows [23].

There are a number of possible ways to measure whether a small group of contributors is doing most of the work. One of the most commonly used is the Gini coefficient [1], typically used in economics or sociology to describe the inequality of a distribution (e.g., the degree to which wealth is unequally distributed within a country). The Gini coefficient is bounded between 0 and 1; when the work is equally distributed amongst contributors the coefficient will be closer to 0, while a coefficient near 1 reflects that most of the work is done by a small core of contributors. We computed the average Gini coefficient for each wiki for each month using the historical contributions of individual users; the aggregate data is shown in Figure 8². This figure shows the aggregate Gini for highly active wikis (as defined by those which eventually accumulate at least 5000 edits) versus those that did not reach that level of activity. There are two striking patterns that the figure makes evident. First, there is a large absolute difference between highly active and less active wikis. Highly active wikis have higher Gini coefficients (even at their inception), reflecting the presence of a core group of contributors who do a disproportionate share of the work. Second, the distribution of work becomes even more dominated by a core group

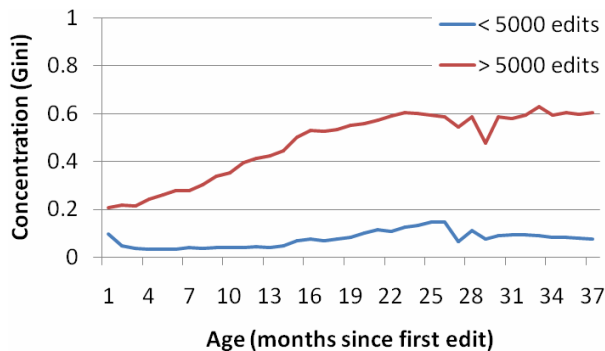


Figure 8. Average Gini inequality distribution for wikis which eventually reach 5000 edits (red) versus those that do not (blue) by age.




Figure 7. Wiki and wiki talk namespaces as a proportion of all edits; aggregate data from all wikis.

over time, as demonstrated by the increase in the Gini coefficient of highly active wikis but not less active wikis.

These results imply that having a core group of leaders is associated with greater editor activity, and that this effect becomes stronger with the age of the wiki. To test this explicitly we ran a multiple regression analysis in which we predicted the number of monthly edits from the degree of concentration (Gini coefficient) from the previous month, the age of the wiki, and their interaction, treating the wiki as a random effect. Because the number of edits had a highly skewed distribution, we used a log to the base 2 transformation to make the distribution more normal; thus, a unit increase or decrease in a variable coefficient should be interpreted as doubling or halving the predicted number of monthly edits, respectively. It is important to remember that our data are correlational and caution is warranted in making any claims about causality. However, to partially address this issue we employ a longitudinal analysis approach, in which we use metrics from one month to predict outcome measures in a subsequent month.

The results are shown in Table 1, and support the pattern shown in Figure 8. The degree of concentration had a large effect: going from a completely equal distribution of work to a completely concentrated distribution of work is associated with an approximately 2800% increase in the number of edits per month. The age of the wiki had no significant effect on the number of edits per month. However, we found a significant positive interaction between age and concentration, reflecting the growing importance of a core group of contributors over time.

Overall, these results support the hypothesis that much of the work, especially in active wikis, is done by a core group of contributors. This finding is consistent with studies in both Wikipedia and other groups such as open source software communities [8], and suggests that it is a core coordination characteristic of many successful wikis as well. Going beyond this, we further find that the importance and degree of concentration of edits in this core group appear to grow over time, and that this interaction is associated with greater editing activity. Together these findings suggest that workgroup structure, in the form of a

core group of contributors, has a remarkably strong degree of impact in online production groups.

CONFLICT

In systems where thousands of contributors must combine their own agendas and points of view, conflict is likely unavoidable. While in particular situations certain types of conflict can be productive (e.g., [19]), conflict generally has been found to lead to negative outcomes, both for the task and for the people involved [20]. Distributed groups are especially prone to conflict due factors including a lack of shared context, difficulties in sharing information, and reduced familiarity with other members [15].

How does conflict arise as an online production group grows? On the one hand adding new members, each with their own viewpoints and information, may cause conflict. On the other hand, these new members create new places for work to happen, such as new articles that did not previously exist. Having more places for work to happen could reduce conflict, as contributors would not have to interact with each other. To make an analogy with skiers on a mountain, conflicts (i.e., collisions) happen when there are many people going down the same slope. Two possible ways to reduce conflict are to reduce the number of people on the slope, or to create more (or bigger) slopes.

To test this hypothesis we analyzed the influence of the number of contributors (i.e., the skiers) and the size of the wiki (i.e., the slopes) on conflict. The measure of conflict is the number of reverts in the wiki per month. A revert involves returning an article to a previous state by throwing out subsequent work. At a minimum, reverts indicate task inefficiency. They implicitly indicate task conflict, in which different editors have different views of what should be written, with some of them overturning editing decisions made by others. Although reverts don't directly capture conflict relationship conflict (e.g., ill feelings between editors), prior literature suggests that task and relationship conflict are highly correlated (mean $r=.54$ [9]). Several studies in Wikipedia show that revert patterns can accurately reflect conflict [3][22][35].

Previous studies have also used different ways to measure revert activity, each with different strengths and weaknesses (for example, identifying reverts through identical text revisions misses partial reverts, in which editors revert only part of a revision), but on the whole most revert measures have been found to be fairly consistent with each other (e.g., [22][31][38]). Here we identify reverts as those in which editors explicitly mark the reverting revision in the

	Coef.	SE	P
Conc.	4.46	0.043	***
Age	0.0007	0.001	
Age X Conc.	0.039	0.002	***

Table 1. Regressions examining the influence of concentration (as measured by the Gini coefficient) on the number of monthly edits for all wikis. $p < .001 = *$**

comment field using the word “revert” or common variations (e.g., “rv”). This method is fairly conservative with low false positives and also captures partial reverts, though it will miss unmarked reverts.

We ran a multiple regression analysis predicting the number of reverts in a given month from the number of contributors and the total size of the wiki (in characters) from the previous month. We also controlled for the age of the wiki in months. The number of reverts, contributors, and wiki size were normalized by taking their logarithm base 2. The results of the analysis are shown in Table 2.

Consistent with our hypothesis, an increase in the number of contributors was associated with a significant increase in conflict, while an increase in the size of the wiki was associated with a significant decrease in conflict. This supports the idea that conflict is related not just to the number of contributors involved but also to the degree to which those editors are forced to interact; in other words, increasing the of contributor density of the wiki space is associated with an increase in conflict. However, by far the most important factor related to conflict in the wikis studied here was the number of contributors ($\beta \sim .7$), rather than the size or age of the wiki space ($\beta_s \sim .06$).

CONFLICT AND COORDINATION

In the previous section we showed that increasing numbers of contributors are associated with greater degrees of conflict. However, online production groups are highly dependent on recruiting and retaining additional contributors in order to grow and be effective. We turn now to examine the important question of how online production groups can manage conflict that occurs when the group grows in size in order to remain effective. Specifically, we examine the influence of the coordination techniques previously introduced as moderators of conflict as the number of contributors increases.

This analysis conceptually integrates our analyses on coordination and on conflict described earlier. Similar to the conflict analysis, we use reverts in a given month as our outcome measure; specifically, the proportion of edits in a month that are reverts. The coordination techniques we examined earlier – communication measure by article talk and user talk, policy and procedure, and concentration of workgroup structure – are used as predictor variables. We are interested both in how they direct influence conflict and how they moderate the increase in conflict associated with having a larger of number of editors involved, which we found in our previous analysis to be the most important predictor of conflict³.

We use a longitudinal analysis approach as in the previous analyses, in which we predict our outcome variable

³ While we could have also included the size of the wiki as a predictor, we did not in order to avoid multicollinearity issues as it is highly correlated with other predictors.

	Coef.	SE	P
Editors	0.3284	0.0017	***
Wiki size	-0.0131	0.0005	***
Age	0.0054	0.0002	***

Table 2. Regression predicting the number of reverts in a month from the number of editors and the size and age of a wiki from the previous month. $p < .001 = *$**

(proportion of reverts) in a given month from the values of our independent variables (e.g., number of editors and coordination methods) from the previous month. In addition, we also included the proportion of reverts from the previous month so that the analysis predicts *changes* in conflict during a time period. This controls for time-invariant unobserved qualities of a wiki, such as their topic or popularity, and allows us to make stronger causal claims than possible with a cross-sectional analysis while also avoiding problems with reverse causation.

The results are shown in Table 3. As expected our control variable -- the proportion of reverts from the previous month -- was correlated with the proportion of reverts in the current month, which also includes influence from time-invariant factors. Consistent with our previous conflict analysis, the number of editors in a month was significantly associated with the proportion of reverts in the following month, and the age of a wiki was also a significant though smaller factor.

Looking at the main effects of the coordination mechanisms, we found that the proportion of article talk was positively associated with conflict. It is difficult to distinguish whether article talk is a cause of conflict or a marker of conflict, since much conflict happens during the discussion of an article, but given that we are predicting conflict for a given month using the article talk from the previous month, this suggests that the more communication there is, the more likely it is for conflict to occur. This is consistent with [22], in which one of the strongest predictors of conflict in a machine learning model was the number of article talk edits in Wikipedia (although this work was purely correlational and hence subject to greater problems with reverse causation).

We also found that the main effect of concentration is associated with increased conflict. On the one hand, this is consistent with work in Wikipedia suggesting that core editors on a page are more likely to revert the contributions of others contributing to that page [13]. This may be a sign of territoriality [36], with committed editors “defending their turf”. Also, although not entirely comparable to concentration, machine learning models of conflict in Wikipedia suggest that having more editors involved in article talk is associated with less conflict [22], for example due to decreased territoriality or to intervention by more impartial mediators. On the other hand, having work highly concentrated in a core group of editors is also associated with improvement in article quality [23]. This suggests that

there may be a risk/reward situation with having a small group of editors doing most of the work: if they work well together, avoid territoriality, and help structure the work of less involved contributors they can be extraordinarily effective in improving quality; but if they do not, having a few highly committed editors who are highly committed only to their own personal viewpoints can be a source of high conflict. An intriguing possibility is that in some situations both may be true, if conflict between editors is constructive and helps to clarify arguments and improve the page.

The only coordination mechanism we studied that had a significant main effect of reducing conflict was policy and procedural work. On the one hand it appears that active policies and procedures can have a direct influence on managing conflict. However, this was also the only coordination mechanism that had a positive interaction with the number of editors, indicating that as wikis accumulate more editors, changes to their policies and procedures are associated with greater conflict in subsequent time periods.

The other coordination mechanisms -- article talk, user talk, and concentration of workgroup structure -- all had significant negative interactions with the number of editors, indicating that as wikis scale up, these coordination mechanisms increase in their effectiveness for managing conflict. This flies in the face of previous results which found that direct communication mechanisms did not scale up to many contributors when using increases in article quality as the outcome variable [23]. One interpretation of these findings is that apart from the benefit to the article itself, there may be a social benefit to communication between editors in reducing the likelihood of conflict between them. Communication may play an indirect role in improving quality by promoting shared mental models between contributors, which may enable them to work more effectively together on the present as well as future tasks

	Coef.	SE	P
Prop. reverts (previous)	0.0596	0.0040	***
Age	0.0003	0.0000	***
Editors	0.0079	0.0006	***
Prop. article talk	0.0241	0.0057	***
Prop. user talk	0.0054	0.0037	
Prop. policy + procedure	-0.0275	0.0031	***
Concentration	0.0107	0.0018	***
Editors X article talk	-0.0093	0.0027	**
Editors X user talk	-0.0074	0.0024	**
Editors X policy + proc.	0.0255	0.0024	***
Editors X concentration	-0.0072	0.0007	***

Table 3. Regression analysis predicting change in proportion of reverts in one month given coordination factors, number of editors, and control variables from the previous month.

[33]. This could also explain the beneficial effect of concentration of work in fewer contributors, which is another way of promoting shared mental models [33][37].

DISCUSSION

Wikis are a powerful and popular means for online production. Managing coordination and conflicts arising from breakdowns in coordination are important challenges for any online production group engaged in collaborative content co-creation. However, most of the research on coordination and conflict in wikis has focused on a single environment, Wikipedia. This paper analyzes the extent to which findings initially found in Wikipedia generalize to other wiki systems.

In the present research we found that many of the coordination methods and their patterns of growth in Wikipedia were also found in other wiki systems. Intra-article communication in both Wikipedia and other wiki systems grew and then held constant, while communication between editors in contexts not tied to specific articles started more slowly but continued to rise. These data suggest that article talk scales linearly as the size of the wiki grows, but user talk scales superlinearly. One possible explanation for this finding is that greater growth of users leads to superlinear increases in user talk as the number of possible communication pairs grows, but article talk scales linearly as the number of users involved in communicating about a specific topic is much smaller.

While direct communication was similar in Wikipedia and other wikis, they differed in their use of policies, procedures, and other mechanisms, such as projects, noticeboards and task lists, that Wikipedia uses heavily for imposing a common set of procedures on its members. The larger wiki sample did not use policies so heavily. It is not clear whether Wikipedia is unique in its heavy reliance on policies, whether other wikis will start to use policies more as they grow in size, or whether these difference reflect measurement problems (e.g., whether use of policies may be more difficult to detect in other wikis).

We also examined conflict, an issue that many online communities face. We found evidence that conflict depends not only the number of contributors involved, but also on the dependencies incurred by their contributions. When the size of the information space grows there is a decrease in conflict, as editors are not forced to interact as often. On the other hand, there is increased conflict when more editors join the community. This suggests that the *density* of the information space, as opposed to the absolute number of contributors, is a key determinant of conflict.

All of the coordination mechanisms we examined were effective for managing conflict in certain situations but not in others. Policies and procedures were associated with less conflict when smaller numbers of editors worked on content, but were associated with increased conflict when many editors were involved. Conversely, communication and

concentration were more associated with reduction of conflict as more editors were involved. In contrast, a previous study focusing on quality rather than conflict as an outcome measure found that communication was ineffective with many involved editors. This suggests that even if it is ineffective for the task-focused coordination of large numbers of editors, communication can play an important relationship-focused role in managing conflict between them. For example, sometimes it may be worth some extra communication overhead to avoid conflicts that lead to contributors getting angry, losing motivation, and quitting the community. Further research is needed to understand how task and relationship factors interact (as they are almost certainly not independent) and what designs and technologies may enable online production groups to have the best of both worlds.

The remarkable similarities between certain coordination patterns in Wikipedia and other wiki production groups suggest that there may be common laws that apply to a variety of communities working with different user populations in different domains. However, even with such a diverse sample caution is warranted for making generalizations. Many of the wikis studied here are influenced by the norms and conventions developed in Wikipedia. While these similarities increase the validity and tractability of direct comparisons with Wikipedia, they also indicate that these production communities are not examples of independent evolution, and thus may not be representative of other kinds of wiki systems.

One step towards a more general theory about wiki development may be to understand and model the structure of coordination at the process level. For example, if we can develop generative models of the network and communication structure of editors, we may be able to account for the increases in the amount and type of coordination that occurs. Given these models we may then predict what kinds of changes in that structure will lead to more optimal coordination patterns. Towards such a goal, we believe a potentially fruitful area of research may be combining large scale simulation models of network dynamics (e.g., [40]) with micro-level coordination dynamics (e.g., [14]). Research to test such models will likely need to go beyond observational and correlational studies to interventions and experiments, which could provide invaluable novel theoretical and practical results.

ACKNOWLEDGMENTS

This work was supported by NSF grants IIS0729286 and IIS0808711.

REFERENCES

1. Atkinson, A.B. On the measurement of inequality. *Journal of Economic Theory*, 2 (1970), 244-263.
2. Beschastnikh, I., Kriplean, T., & McDonald, D. W. *Wikipedian Self-governance in Action: Motivating the Policy Lens*. In *Proceedings of ICWSM* (2008).

3. Brandes, U. and Lerner, J. Visual analysis of controversy in user-generated encyclopedias. *Information Visualization 7* (2008), 34–48.
4. Bruns, A. & Humphreys, S. Wikis in teaching and assessment: The M/Cyclopedia project. In *Proceedings of WikiSym* (2005).
5. Burke, M. & Kraut, R.E. Mopping up: modeling wikipedia promotion decisions. In *Proceedings of CSCW* (2008).
6. Butler, B., Joyce, E., & Pike, J. Don't Look Now, But We've Created a Bureaucracy: the Nature and Roles of Policies and Rules in Wikipedia. In *Proceedings of CHI* (2008).
7. Collier, B., Burke, M., Kittur, A., & Kraut, R.E. Retrospective Versus Prospective Evidence For Promotion: The Case Of Wikipedia. In *Proceedings of the Academy of Management* (2008).
8. Conklin, M. (2004). Do the rich get richer? The impact of power laws on open source development projects. In *Proceedings of OSCON* (2004).
9. De Dreu, C.K.W. and L.R. Weingart, Task versus relationship conflict, team performance, and team member satisfaction: A meta-analysis. *Journal of Applied Psychology*, 88 (2003), 741-749.
10. Eisenhardt, K. M., Xin, K. R. Exploring the black box: An analysis of work group diversity conflict and performance. *Admin. Sci. Quart.* 44 (1999), 1–28.
11. Forte, A., & Bruckman, A. From Wikipedia to the classroom: exploring online publication and learning. *Proceedings of the 7th International Conference of the Learning Sciences* (2006).
12. Forte, A., & Bruckman, A. Scaling consensus: Increasing decentralization in Wikipedia governance. In *Proceedings of HICSS* (2008).
13. Halfaker, A., Kittur, A., Kraut, R. E., & Riedl, J. Quality, Experience and Ownership in WikiWork. In *Proceedings of WikiSym* (2009).
14. Herbsleb, J.D., & Mockus, A. Formulation and preliminary test of an empirical theory of coordination in software engineering. In *Proceedings of SIGSOFT* (2003).
15. Hinds, P. J., & Bailey, D. E. Out of sight, out of sync: Understanding conflict in distributed teams. *Organization Science* 14, (2003) 615-632.
16. http://meta.wikimedia.org/wiki/List_of_largest_wikis
17. <http://www.alexacom>
18. <http://www.wowwiki.com/Special:Listusers>
19. Jehn, K. A. A multimethod examination of the benefits and detriments of intragroup conflict. *Admin. Sci. Quart.* 40 (1995), 256–282.
20. Jehn, K. A., Chadwick, C., Thatcher, S. M. B. To agree or not to agree: The effects of value congruence, individual demographic dissimilarity, and conflict on workgroup outcomes. *Internat. J. Conflict Management* 8, (1997) 287–305.
21. Kittur, A., Chi, E., Pendleton, B. A., Suh, B., & Mytkowicz, T. Power of the few vs. wisdom of the crowd: Wikipedia and the rise of the bourgeoisie. *Alt.CHI* (2007).
22. Kittur, A., Suh, B., Chi, E., & Pendleton, B. A. He says, she says: Conflict and coordination in Wikipedia. In *Proceedings of CHI* (2007).
23. Kittur, A., & Kraut, R. E. Harnessing the wisdom of crowds in Wikipedia: Quality through coordination. In *Proceedings of CSCW* (2008).
24. Kittur, A., Lee, B., Kraut, R. E. Coordination in Collective Intelligence: The Role of Team Structure and Task Interdependence. In *Proceedings of CHI* (2009).
25. Kittur, A., Pendleton, B., & Kraut, R. E. Herding the Cats: The Influence of Groups in Coordinating Peer Production. In *Proceedings of WikiSym* (2009).
26. Kriplean, T., Beschastnikh, I., & McDonald, D. W. Articulations of WikiWork: Uncovering Valued Work in Wikipedia through Barnstars. In *Proceedings of CSCW* (2008).
27. Leavitt, H. Some effects of certain communication patterns on group performance. *Journal of Abnormal and Social Psychology*, 46 (1951), 38-50.
28. Malone, T. and K. Crowston, The interdisciplinary study of coordination. *ACM Computing Surveys*, 26, (1994), 87-119.
29. March, J. G., & Simon, H. A. *Organizations*. New York: Wiley (1958).
30. Mockus, A., Fielding, R. T., & Herbsleb, J. D. Two case studies of open source software development: Apache and mozilla. *ACM Transactions on Software Engineering and Methodology*, 11, 3 (2002), 309-346.
31. Mohammed, S. and B.C. Dumville, Team mental models in a team knowledge framework: Expanding theory and measurement across disciplinary boundaries. *Journal of Organizational Behavior Special Issue: Shared cognition*, 22 (2001), 89-106.
32. Priedhorsky, R., Chen, J., Lam, S., Panciera, K., Terveen, L., & Riedl, J. Creating, destroying, and restoring value in Wikipedia. In *Proceedings of GROUP* (2007).
33. Rouse, W.B., Cannon-Bowers, J. A., & Salas, E. The role of mental models in team performance in complex systems. *IEEE Transactions on Systems, Man, and Cybernetics*, 22 (1992) 1296-1308.
34. Steiner, I. D. *Group process and productivity*. New York: Academic Press (1972).
35. Suh, B., Chi, E., Pendleton, B. A., & Kittur, A. Us vs. Them: Understanding Social Dynamics in Wikipedia with Revert Graph Visualizations. In *Proceedings of VAST* (2007).
36. Thom-Santelli, J., Cosley, D., & Gay, G. What's Mine is Mine: Territoriality in Collaborative Authoring. In *Proceedings of CHI* (2009).
37. Tushman, M. Work characteristics and subunit communication structure: A contingency analysis. *Administrative Science Quarterly*, 24, (1979), 82-98.
38. Viégas, F. B., Wattenberg M., & Dave, K. Studying cooperation and conflict between authors with history flow visualizations. In *Proceedings of CHI* (2004).
39. Viégas, F. B., Wattenberg, M., Kriss, J., & van Ham, F. (2007). Talk before you type: Coordination in Wikipedia. In *Proceedings of HICSS* (2007).
40. Watson-Manheim, M., K. Chudoba, and K. Crowston, Discontinuities and continuities: A new way to understand virtual work. *Information Technology & People*, 15 (2002), 191-209.
41. Zhang, J., Ackerman, M.S., Adamic, L. CommunityNetSimulator: Using Simulations to Study Online Community Networks. In *Proceedings of Communities & Technologies* (2007).