Tool-Mediated Coordination of Virtual Teams in Complex Systems

Michael Gilbert, Mark Zachry
Human Centered Design & Engineering
University of Washington
Seattle, WA USA
{mdg, zachry} @uw.edu

ABSTRACT
Support for coordination in online spaces, specifically in peer production systems, has frequently been an after-thought. In the absence of such support, the users of such systems must work to find an emergent order that drives shared project goals and leads to equitable processes. In short, they must rely on the "wisdom of the crowds." As our study demonstrates, however, the reality is that often the system tools available for coordination, evaluation, and work articulation are not suitable to the task at hand. Our study, first, takes a theoretical approach to understanding how tool-mediated coordination functions within peer production systems. Secondly, we enumerate the methods available to identify automated and semi-automated tools that function within such systems by quantitatively and qualitatively analyzing trace interactions and their utility in Wikipedia over a year-long period. Finally, we identify potential vacuums where new design interventions have the greatest potential for enhancing peer-production systems.

Categories and Subject Descriptors
H.5.2 [**Information Interfaces and Presentation**]: User Interfaces – *Theory and methods*

General Terms
Measurement, Design, Theory

Keywords
Wikipedia, Coordination, Theory

1. INTRODUCTION
In the first week of July, 2014, the Wikipedia user AAlertBot made 237 individual edits to 51 WikiProject pages, notifying project members about articles that may need to be improved, updated, or deleted. In the same week, the user COIBot made 194 revisions notifying Wikipedia users of edits made with potential conflicts of interest. As well, 8,673 unique edits were made by 834 users, rolling back to previous revisions using the AutoWikiBrowser, a semi-automated Wikipedia editor. Each of these are examples of bot or tool-mediated edits, the proliferation of which has allowed Wikipedia to grow from a community of like-minded enthusiasts to one of the most visited and recognizable locations on the web today, facilitating a resilient socio-technical system in which individual efforts are amplified and coordination across multiple channels can be delegated to automated and semi-automated tools. In this study, we examine the extent to which such tools impact human activity, and explore how such tools might be designed in the future to better complement human goals.

Group coordination in online spaces has always posed unique challenges. Among them, group members most often are not collocated, frequently with individual team members not even being located in the same time zone. Communication among team members, consequently, can be difficult, sometimes only occurring in asynchronous fashion. Beyond this, knowledge of other group members, their talents, their desires, and their schedules and offline interests are frequently unknown. Policies and social norms are frequently only gleaned through immersion in the group—a process that can only be accomplished through continuous involvement over time, increasing the barrier to entry for new participants who may otherwise have been willing contributors. Rules and governance structures frequently exist, but they too are often impenetrable to new members, ensuring an entrenched minority controls the rules for the majority [17]. Given these myriad challenges and difficulties, it is indeed remarkable that such projects ever succeed. However some have succeeded; chief among them is Wikipedia.

Since launching in January of 2001 on the open-source MediaWiki platform, Wikipedia has engaged users in novel forms of collaboration and communication. Thousands of people have contributed to this new online encyclopedia, created with the goal of making freely available the sum of all human knowledge[1]. Over the last 14 years, Wikipedia has grown to be one of the most visited sites on the web, currently ranking number six in both global and US visits[2], with over 4.9 million articles comprised of over 480 million edits by 25 million registered users and over 136 million edits by anonymous users in the English Wikipedia alone since its inception[3].

Within the broad, diverse set of activities that make the encyclopedia possible, a key source of coordinated work is WikiProjects. Within Wikipedia, hundreds of groups of editors

OpenSym '15, August 19 - 21, 2015, San Francisco, CA, USA
 ISBN 978-1-4503-3666-6/15/08…$15.00
DOI: http://dx.doi.org/10.1145/2788993.2789843

[1] https://en.wikiquote.org/wiki/Jimmy_Wales

[2] From http://www.alexa.com, as of May, 2015

[3] Numbers collected from the Wikimedia Tool Labs servers for the English Wikipedia, as of May, 2015

have emerged to coordinate their shared efforts to develop content in areas of shared topical interest (e.g., cats, military history), or to coordinate work on meta-tasks, like cleaning the quality of mark-up text. WikiProjects are fundamentally then, "a group of contributors who want to work together as a team to improve Wikipedia."[4]

Recently, task-oriented projects, such as those dedicated to monitoring new articles or patrolling for spam have seen a fairly consistent increase in activity [16]. This increasing allocation of effort towards ensuring the validity and quality of the processes that create content for the encyclopedia is an indicator of the increasing importance of coordination in this ever-expanding encyclopedia, suggesting the need for additional tools to be designed to support that growth.

To this end, our study offers a quantitative and qualitative analysis of both automated and semi-automated coordination practices in WikiProjects. Our analysis identifies and classifies the myriad ways in which tool-mediated interactions impact coordination work within WikiProjects. In the analysis, tool-mediated interactions refer to the total set of revisions by automated and semi-automated tools, or bots, both within and external to the MediaWiki platform. We classify these tools through a directed content analysis [12] to determine the spectrum of work activities currently supported by existing tools and to define the types of structured coordination behavior that could be supported given additional, theory-driven insights. This analysis of existing tool-mediated practices was completed using Benkler's characteristics of commons-based peer production communities [1] and Malone and Crowston's coordination mechanisms for small group maintenance [4; 15].

2. BACKGROUND AND RELATED WORK

Our examination of the role of bots and semi-automated tools as coordination mechanisms in peer production systems begins at the intersection of three distinct, but related conceptual frameworks. First, we turn to the idea of coordination in commons-based peer production communities. We then turn to coordination theory, which has been instrumental in defining and modeling the productively synchronized actions of groups. Finally, we consider ways of theorizing tool use in peer production systems.

2.1 Structures and Processes within Commons-Based Peer Production Communities

A commons-based peer production community, according to Benkler [1], relies on the creativity, intelligence, and efforts of distributed individuals to work towards a common goal. In such communities, production is decentralized. Action within the community is dictated by the desires, talents, and attributes of community members rather than a central organizer, allowing for emergent goals and activities to drive production rather than a single coordinating vision. Beyond this, actors in such communities rely on social cues and motivations rather than explicit market exchanges to coordinate their activity.

These peer production communities typically exhibit three types of primary structural attributes [2]: first, the objects being acted upon should be modular, allowing for a clean division of potential effort; second, those objects should be relatively granular, with work units that are manageable for the intended population (individual, small group, or large collective) to operate on them effectively; and third, the mechanisms by which changes to the above objects are modified should be capable of being integrated back into the whole community in a simple, low-cost manner.

From the perspective of coordination activity, WikiProjects are modular work elements allowing editors to more effectively locate and define their efforts with regard to their stated membership in projects (Attribute 1). Beneath the project level are tasks, often organized within WikiProjects as the work of Task Forces [7], allowing groups to more effectively distribute and organize efforts in a granular fashion (Attribute 2)[5]. Finally, the efforts of these atomic work units can be integrated into the larger encyclopedia and the results of that effort can be recorded on project pages to aid coordination and task tracking for future efforts (Attribute 3), an integration and feedback mechanism that is currently not ideally implemented. The potential opportunity associated with this final attribute is expanded upon and addressed in later sections of this paper.

In coordination theory [4; 15], coordination refers to the means responsible for "managing dependencies between activities." Accordingly, if there are no dependencies in an activity, there is nothing to coordinate. Any activity that exists in isolation, with no portion of that activity depending on any other to successfully complete, does not require any coordinating power.

Given that coordination is "managing dependencies," it is possible to classify and describe such dependencies in any given system to delineate the types of coordination mechanisms that might be possible to ameliorate those dependencies. For example, if one activity is dependent on the successful completion of another activity, coordination mechanisms that facilitate sequential tracking, notifications, or peer review may be useful in driving an overall process to completion. A summary of Malone and Crowston's dependencies and alternative coordination mechanisms is presented in Table 1. Through the lens of coordination theory, we contend, it will be possible to more effectively define and explore dependencies in project activities and, through doing so, to both identify existing tool-mediated coordination activities within projects and to highlight potential interventions that could be made to improve those processes.

2.2 Mediating Tools in Peer Production Communities

In many leading-edge online environments a combination of bots, tools, and scripts substantially shape our interactions with the primary platform. Geiger has labeled this aggregated set of tools as "bespoke code" [10]. In a platform such as Wikipedia, this bespoke code comprises an estimated six million lines of code, an order of magnitude larger than the roughly 600,000 lines that make up the core MediaWiki platform. So beyond the code that drives the core of Wikipedia, beyond the configurations of namespaces, templates, and modules that shape that core, there exists a much greater ecosystem of code that creates and enacts Wikipedia as it is experienced. This includes the bots that patrol for vandalism, grammatical errors, and stylistic omissions, and the

[4] https://en.wikipedia.org/wiki/Wikipedia:WikiProject

[5] See also, http://en.wikipedia.org/wiki/Wikipedia:WikiProject_Council/Guide/Task_forces, for a more detailed description of WikiProject Task Forces

Table 1: Malone and Crowston's common dependencies and example alternative coordination mechanisms for dealing with them. Originally from (Malone & Crowston, 1994).

Dependency	***Alternative coordination mechanism***
Shared resources	"First come, first serve," priority order, budgets, managerial decision
Task assignments	(same as for "Shared resources")
Producer/Consumer relationships	
Prerequisite constraints	Notification, sequencing, tracking
Transfer	Inventory management (e.g. "just in time," "Economic order quantity")
Usability	Standardization, ask users, participatory design
Design for manufacturability	Concurrent engineering
Simultaneity constraints	Scheduling, synchronization
Task/Subtask	Goal selection, task decomposition

scripts that facilitate the ongoing maintenance of the site, the customizable interfaces, and even the bots that send automated greetings for newcomers to the community. Geiger argues [9; 10] that these bots and bespoke tools are increasingly ingrained in the systems that people interact within on a daily basis, and that their experiences in those systems are now increasingly shaped, expanded, or constrained by such tools. And, as these tools impact our perception of and means of interactions with a system, they require greater consideration.

Specifically, Geiger and Halfaker [9] show how pivotal these tools are to the healthy functioning of the community. They demonstrate that while robust and reliable quality control mechanisms exist within Wikipedia to combat the problems of spam and vandalism, the proper functioning of those tools are required to enable the com-munity to properly function at its current scale. Further, Geiger argues that these tools are no mere force multipliers [10], simply duplicating human efforts at greater speed and efficiency; instead, through their implementation and manifestation, they are integral to the core experience users have when they visit the site.

These bots, as tools that shape our perceptions of the online spaces we occupy, are worthy of greater attention and exploration. Despite this, it is evident that too little attention has been paid to the potential of tool-mediated systems to facilitate coordination in online spaces. Our study addresses this problem, charting the affordances and dependencies of an existing and active system from both an applied and theoretical perspective. The study explicates the operation and nature of coordination from a detailed analysis of the tools that enable such interactions. By extension, the study then also introduces the potential for new means of coordination within Wikipedia that can be measurably, empirically tested to further our knowledge of how online collaborative spaces could be more effectively designed.

3. DATA COLLECTION & ANALYSIS

This study makes three contributions to expanding our understanding of coordination in Wikipedia. First, we identify the scope of tool-mediated activity within WikiProjects using three distinct methods for data collection. Second, we classify the nature of those tools from the lens of commons-based peer production communities and coordination theory, expanding existing models to encompass the broader spectrum of functionality provided by these tools. Third, we identify potential design interventions suggested by the skew in existing versus potential coordination mechanisms fulfilled by existing tools to more effectively address community needs with respect to a modern, socio-technical view of coordination theory.

3.1 Identifying Automated and Semi-Automated WikiProject Edits

The landscape of Wikipedia contributions has become increasingly complex, with automated, semi-automated, and otherwise tool-mediated edits playing an increasingly important role in the maintenance and ongoing health of the encyclopedia. But, due to the nature of these myriad tools, the largely unstructured ecosystem of bots, user scripts, gadgets, and external applications available to contributors, it is often difficult to identify exactly which edits were tool-mediated[6]. As such, we take a three-pronged approach to identifying these tool-mediated edits with the goal of determining what coordination mechanisms are at work within the scope of WikiProjects.

The sections that follow introduce the method of identifying automated or semi-automated edits as well as present the findings of that method, each in turn. Data for each of the following sections were collected from the Wikimedia Toolserver[7] using custom Python scripts to aggregate and cache all WikiProject revisions in the English Wikipedia for a one year period, starting 05/01/2013. WikiProjects were identified using the category structure present on Wikipedia pages. Included in this data set are all edits to WikiProject pages, sub-pages, or any pages transcluded on either of the above, as well as the corresponding talk pages for each.

[6] Geiger, personal communication, https://lists.wikimedia.org/mailman/listinfo/wiki-research-l, "Kill the bots" thread from 5/18/2014

[7] Previously http://toolserver.org/, now available from the Wikimedia Tool Labs, http://tools.wmflabs.org/

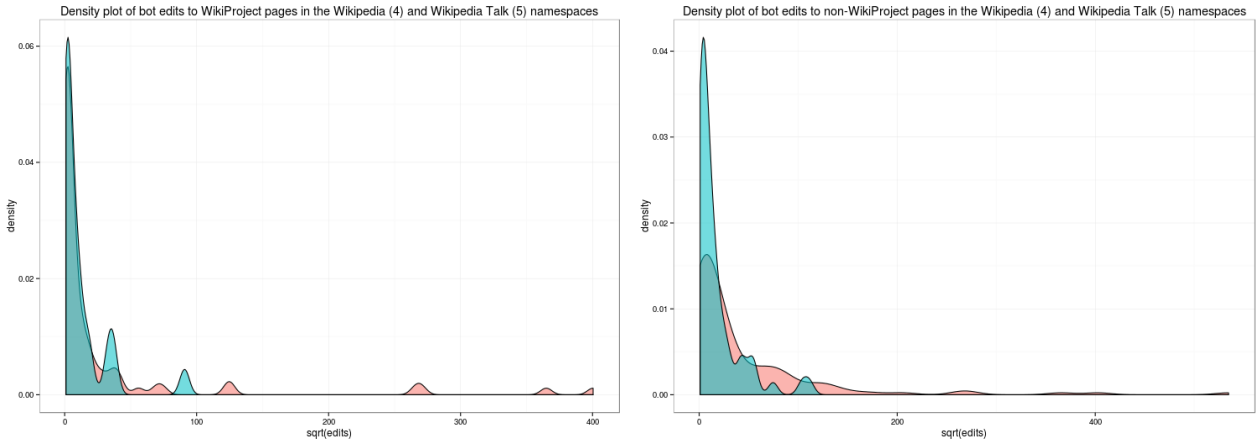

Figure 1. Density plot of bot edits to WikiProjects (left) and non-WikiProjects (right). Red represents the Wikipedia namespace (4) and blue represents the Wikipedia Talk namespace (5). X-axis is square root of edit counts.

3.1.1 Edits by users with the "bot" flag

Method: To be an officially recognized bot in Wikipedia, a unique bot account must be created and a request must be made with the Bot Approvals Group where the tasks that the bot aims to complete and the code that will enact those changes can be reviewed, vetted, and ultimately approved or rejected. Once a request is approved, the user account will be added to the "bot" user group, flagging that user as an official legitimate alternative account, capable of editing the encyclopedia in an automated or semi-automated fashion. Edits by these accounts are easily identifiable and provide a simple starting point to identifying the types of activities that automated and semi-automated users complete within the scope of WikiProjects.

Table 2. Non-bot edits to WikiProject pages identified by user and revision comment. Not represented are 19 other tools.

Method of identifying tool	**Tool**	**Edits**
rollback' WHERE rev_comment RLIKE "(Reverted ([0-9]+)?edits by \[\[Special:Contributions/[^|]+|[^]]+\]\] \(\[\[User talk:[^|]+|talk\]\]\) to last version by .+)";	awb	2121
undo' WHERE rev_comment LIKE "Undid revision%";	undo	1786
huggle' WHERE rev_comment LIKE '%WP:HG%' OR rev_comment LIKE '%WP:HUGGLE%';	twinkle	790
huggle' WHERE rev_comment RLIKE "(Message re\. \[\[[^]]+\]\])|(Level [0-9]+ warning re\. \[\[[^]]+\]\])";	huggle	176
twinkle' WHERE rev_comment LIKE '%WP:TW%';	rollback	1

Findings: In all, bot edits to WikiProject pages comprised 23.8% (507,754 edits) of the total edits to the Wikipedia and Wikipedia Talk namespaces, whereas the number of WikiProject pages, subpages, and corresponding Talk pages make up only 16.97% (218,611 WikiProject-related pages) of the total pages within those namespaces. This discrepancy, evident in the density plots in Figure 1, is largely the result of a select few bots' high level of activity across multiple projects. The top five bots by edit counts to WikiProject-related pages, for instance, make up 22.32% of the total bot edits, with all remaining edits by bots making up the remaining 1.48% of the total[8].

Initial analysis of the quantitative data clearly indicates two primary features. First, bot activity is higher among WikiProject-related pages than non-WikiProject pages, notably in the Wikipedia namespace (as opposed to the Wikipedia Talk namespace). Second, although there is a greater proportion of bot activity on WikiProject-related pages, the clear majority of that activity is coming from a very few, select bots. The nature of those edits, and the functions they aim to serve within the WikiProject space, will be explored further below.

3.1.2 Non-bot edits identified through standard strings in revision comments

Method: While it is official policy that bots go through the prior stated review and approval process before operation, there are still accounts that operate bots that have not been approved and therefore do not have the "bot" flag added to the user account, as well as semi-automated tools that facilitate edits on behalf of individual users. To identify these tools, and ultimately to identify the function that they serve within the project space, we additionally looked at all non-bot edits within the project scope, identifying strings that occurred frequently within the revision comments and tags that are added in known tool-mediated edits. For example, tags added in revision comments such as "WP:TW",

[8] Top five bots by total edit counts to WikiProject visible pages are AAlertBot (233,545 edits), WP 1.0 bot (134,001), COIBot (70,756), JL-Bot (18,723), and Cyberbot I (15,606).

"MWT|MWT", or "AWB|AWB" representing edits completed with Twinkle, Mike's Wiki Tool, and the Auto Wiki Browser, respectively.

Findings: While there are currently no absolute means of identifying tool-mediated edits within Wikipedia (see footnote 6 above), many of the more common tools used within Wikipedia leave identifiable traces through either user names (for instance, AntiVandalBot, OrphanBot, or PseudoBot) or in the revision comment left with an edit (for instance, for rollbacks, Huggle, Twinkle). The data set comprised of these automated and semi-automated edits was built using identifiers referenced in current Toolserver documentation[9], resulting in 346,008 total revisions to WikiProject-related pages during the investigation period. This analysis yielded only 4,874 total tool-mediated edits within the original set of total revisions. Foremost among the tool-mediated edits within the WikiProject space was AutoWikiBrowser, a semi-automated editor that allows users to make multiple edits more quickly and easily than if done by hand, simplifying the process of editing for many more tedious tasks. Second, among the top tool-mediated edits was "undo," a built-in functionality that enables users to easily, with one click, undo a prior revision on any Wikipedia page. This was followed by Twinkle, a tool that attempts to automate many frequent tasks such as listing a page for speedy deletion, tagging a page for protection, or jumping to the diff of previous revisions of the current page. Combined, these tools represent 96.3% of the tool-mediated revisions identified within the WikiProject space for the period under analysis; however, altogether these tool-mediated edits still only represented 1.41% of the total edits to pages within the WikiProject space. A breakdown of the tools identified is shown in Table 2.

3.1.3 Non-bot edits identified through successive revisions

Method: Finally, to identify tool-mediated edits that neither originate from a legitimate bot account nor contain any formulaic revision comment string, we analyzed all edits to pages within the scope of WikiProjects from individual users that occurred in rapid succession. As Geiger and Halfaker show [8], the majority of inter-edit time in edit sessions extends from seconds up to an hour, with the greatest frequency being around one minute. To increase the probability that we are retrieving only tool-mediated edits, therefore, we identified all accounts that made two or more successive edits to any page within the scope of a WikiProject in five seconds or less. Manual inspection of the results of this analysis removed any edits that ultimately did not appear to stem from an automated or semi-automated tool and were simply the result of a user rapidly committing multiple changes to the project.

Findings: Using the above dataset of all non-bot revisions to pages within the WikiProject space, including Talk pages and pages transcluded within WikiProjects, we reduced the initial set of 346,008 revisions to 25,851 that met the timing criteria. Through manual sorting and selection of frequently occurring strings we identified 26 strings within revision comments that accounted for 97% of the total revisions present in the reduced data set.

[9] List identifying frequent tool and bot edits through trace data at https://wiki.toolserver.org/view/MySQL_queries#Automated_tool_and_bot_edits

Table 3. Non-bot edits to WikiProject pages identified by concurrent edits which occurred <= 5 seconds apart from one another.

Revision Comment	**Edits**
[[WP:FWDS|FWDS]]	10873
archive OR archiving (case insensitive)	8261
moved page	1582
Updating user statistics	1475
new section"	805
AWB|AWB	644
"" (null comment)	447
WikiProject_Deletion_sorting	250
Other	1514

Among the top comments identified were "[[WP:FWDS|FWDS]]" (the trace of a deletion sorting tool), references to "archiving," "moved page," "Updating user statistics" (a trace left by an administrator completing batch updates), and "AWB|AWB" (the previously introduced AutoWikiBrowser). Table 3 shows a more detailed view of the top comment strings and the extent to which each was found within the data set, representing the top ~94% of strings by edit count. In all, these concurrent revisions on WikiProject-related pages made up 7.47% of the total number of edits, significantly more than the 1.41% identified through the analysis of non-bot tool-mediated edits shown above.

3.2 Classifying Automated and Semi-Automated Activity in the WikiProject Space

After determining if a bot or tool contributes to a project, we sought to classify how each of them contributes to coordination within the project space. We completed data analysis of automated and semi-automated edits to the project space in two primary ways. First, a mixed methods analysis of bots that have edited WikiProject pages, sub-pages, and pages transcluded within each of them was completed, which aimed to both describe the types of tool-mediated edits that are done to project pages and to show to what extent those edits may have an impact on the coordination of group members, discussed above. Second, we completed a content analysis of these bots, their functionality, and the boundaries of their utility to WikiProject members with particular focus paid to their potential, realized or not, to coordinate activity of those distributed project members.

Our analysis of WikiProject-contributing bots was completed using a directed content analysis approach [12] utilizing Malone and Crowston's [4] coordination mechanisms for small group maintenance as a starting framework. This directed content analysis allows for a more open inquiry into the specific mechanisms that these bots play in WikiProject coordination and, given their current usage within projects, facilitates the distinctions between coordination mechanisms that may be suggested by coordination theory [15] and those that are actually used in practice, highlighting the potential for a novel technological intervention into this space.

An inductive, iterative approach was taken towards identification of automated and semi-automated tool functionality. For each of

the 54 bots found to be active on project-related pages for the duration of the study period, we visited the bot User and User Talk pages, the bot's Request for Approval page[10], and the bot's contribution history to view both the expressly stated purpose for the bot as well as activity recorded through actual practice. For non-bot tool-mediated edits, we conducted a similar analysis of the documentation page, user manual, or revision history of the related tool to identify the types of coordination mechanism that may be fulfilled by any given tool. Additionally, as many of these automated and semi-automated tools provide multiple types of functionality, where it was identified that more than one coordination mechanism was fulfilled, a tool may be included in more than one category. Classification was completed through activity focused analysis [15], whereby the goal is to identify the mechanisms behind an interaction and then search for dependencies that mechanism resolves. For instance, as the primary mechanism behind AAlertBot is to deliver article alerts when pages move in or out of certain workflows (i.e., AfD, PROD, RfC, etc), the dependency that mechanism resolves is the Producer/Consumer relationship, prerequisite constraints, indicating notification, sequence, and tracking. The complete description of all dependencies used to classify tool-mediated activities was adopted from [15]. If a tool provided functionality not included in the coordination framework, we created a new top-level category to ensure that all tool-mediated interactions were accounted for in our updated classification framework, shown in Table 4.

4. DISCUSSION

In our analysis, we have confronted two primary research questions. First, how can coordination in online spaces be most effectively mapped and modeled? Second, how might the insight gained from such an analysis inform and direct the design of tools intended to facilitate continued coordination within such peer production systems? While myriad bots, tools, and scripts exist to ameliorate the act of contributing to these online spaces, and WikiProjects in particular, the tools currently in use within these spaces do not fulfill the spectrum of coordinative functions shown by prior research to ensure the optimal functioning of small groups. Further, we have supported this stance with an empirical analysis of the automated and semi-automated tools used within and beyond the scope of WikiProjects, highlighting the means in which the data was collected, the methods by which that data was interpreted. and the theoretically based justification for the suggestions that resulted—namely, that given the complex nature of the systems many of these tools are created to support, there exists the potential for more focused interventions to directly and positively impact the ability of these online groups to collaborate in a more successful and sustainable manner.

Benkler and Nissenbaum argue that the mechanisms by which individual efforts are integrated back into the whole should be efficient and low-cost. In WikiProjects, such a mechanism would provide a means by which individual editors would be able to maintain awareness of project activities, project needs, project members, and how they are able to act upon each of these needs. Currently there is no such available tool to facilitate these project-level interactions. The analysis completed in this work aims to fill this vacuum, to foster a peer production community that is resilient, adaptive, and transparent.

[10] https://en.wikipedia.org/wiki/Wikipedia:Bots/Requests_for_approval

Table 4: Prevalence of coordination mechanisms utilized in automated and semi-automated tools in the project space. Percentages given in top-level categories are global, sub-category percentages are relative to each category. Categories observed in practice but not included in the original coordination framework are included under the existing framework.

Shared Resources	**0**	**0.0%**
Task assignments	0	0.0%
Producer/Consumer relationships	**19**	**23.8%**
Prerequisite constraints	10	52.6%
Transfer	0	0.0%
Usability	9	47.4%
Design for manufacturability	0	0.0%
Simultaneity constraints	**1**	**1.3%**
Task/Subtask	**1**	**1.3%**
Activities not in the coordination framework		
Visibility/Awareness	**20**	**25.0%**
Force Multiplier	**39**	**48.8%**
Total mechanisms observed	**80**	**100%**

4.1 Discussion of Automated and Semi-Automated Tools

In [10], Geiger considers the coordinative power of these automated and semi-automated tools, noting that they are more than mere force multipliers. He argues that the functions these tools fulfill are not simply increasing the amount of work that individuals can complete in online spaces; rather, they fundamentally alter the ways that individuals can complete work in those spaces. We suggest that to serve a coordinative function within the space of an online team, a tool must meet two primary criteria.

First, the tool must in some way manage dependencies of those attempting to participate in group work. This definition extends from the taxonomy of dependencies as outlined by Malone and Crowston [15], including the producer/consumer relationship (notification of status, transfer of task along dependency chain), simultaneity constraints (effectively manage multiple group members working towards identical ends), and task/subtask dependencies (hierarchy of tasks required for successful completion). Notably, each of these also indicate that any tool-mediation that occurs will require trace evidence in the system in which they are deployed to effectively communicate the nature of the dependency, successfully coordinated or not, to the remaining team members. As we noted earlier, any activity completed in isolation would require no coordinating power or articulation.

Second, the tool must mediate action in a manner that extends beyond being a force multiplier of existing tasks. This requirement stems both from Geiger's [10] assessment, as well as the nature of the processes that these tools address. From the perspective of coordination theory, there is no transfer of information engendered by these tools that would indicate an increase in coordinative power; consequently, they do not fit

within the CT taxonomy. While these types of tools are plentiful, useful, and possibly one of the primary reasons that Wikipedia has continued to thrive, the ability to amplify existing effort by reducing the time to complete arduous tasks does not inherently impact the coordination of remaining group members.

4.2 Filling the Void: Implications for Theory-Based Design

While Wikipedia works as a peer production system, our analysis suggests that much more could be done to facilitate complex user interactions within the community of editors. Efforts to quantitatively measure the impact of automated and semi-automated tools have shown that, while these tools are present within WikiProjects, they still support only a small portion of the more arduous work of human edits. Further, our analysis of both the stated and given impact of those tools on project spaces in Wikipedia demonstrates that the primary functionality they provide falls outside what would be considered a coordination mechanism in prior framings of coordination theory.

Given this, we suggest two points for future research. First, while coordination theory may provide a valuable theoretical lens from which to analyze interactions in online spaces, it must be adjusted for explicating interactions in modern commons-based peer production communities. In particular, it needs to be more sensitive to the role of automated or semi-automated tools within such systems. And second, the future tools designed for use within these types of communities should be able to successfully mediate a full range of coordinative actions, thereby removing the greater human burden and facilitating future growth.

The primary evolutionary phases of commons-based peer production communities are disaggregation and aggregation [1]. While many of the automated and semi-automated tools in use within Wikipedia facilitate actions within either of these phases (the force multipliers above), few facilitate organization and interaction among them. For instance, consider an automated tool that provides task recommendations to Wikipedians based on their prior editing activity, or the editing activity of the Wikipedians that they have worked most closely with in the recent past. While prior framings of coordination theory accounts for dependencies such as *Task assignments*, in which work may be completed through managerial decree or organizational requirement, in modern commons-based peer production communities where tasks are self-selected based on personal interests and their completion is motivated by mechanisms outside organizational boundaries, a more salient dependency would require simple *awareness*. In this case, when then dependency is met, it would indicate that the contributor identified the list of potential tasks and was able to self-select which tasks to complete, or whether to attempt them at all. When such a dependency remains unmet, it is indicative that the contributor was unaware of the work that was being recommended.

This example, in which *awareness* of a tool-mediated interaction is the primary coordinative function that must be met to maintain an information flow that could lead to *Task selection* and potentially to completing that task (progressing from disaggregation to aggregation), is not accounted for in prior formulations of coordination theory. Accordingly, attempts at theory-driven design utilizing these prior formulations of the theory would be unable to identify the potential failure in the flow of information. We contend that greater attention to these theoretical foundations in the design of future tools will have a positive impact on these online communities as a whole, and the members who populate them.

5. CONCLUSION

The primary contributions of this work are two-fold. First, while significant research has been completed to understand the role of bots and other automated or semi-automated tools in Wikipedia (see, for instance, [8; 9]), there has been much less focused attention on the use of those bots and tools as a means of structured coordination within WikiProjects. Our study offers an examination of current practice with respect to the theoretical and practical underpinnings of what coordination is, how it is enacted, and what potential there is for future interventions into that coordinative practice. Second, given this more focused understanding of mediating tools for structured collaboration within WikiProjects, we have articulated the potential for design interventions to facilitate project goals and activity, providing a roadmap for developers and researchers to investigate small group functioning in online spaces and suggest more effective and directed interventions for those groups. Ultimately, our analysis shows that coordination theory has the capacity to identify interactions occurring within WikiProjects; however there are far more interactions being completed than the theory has the capacity to describe, suggesting that alternative formulations are required to adequately facilitate theory-driven design.

While the focus of this analysis of coordination in online teams has been on WikiProjects, we expect the knowledge gained from these investigations to be immediately relevant to other similarly distributed teams. What these investigations aim to accomplish is not only to improve coordination within WikiProjects, but to more clearly define what coordination is within online spaces where humans and automated or semi-automated tools interact, and to expand on our knowledge of its measurement, function, and evaluation.

6. ACKNOWLEDGMENTS

This research is supported by a National Science Foundation (NSF) grant, Enhancing Social Translucence in Systems to Support Virtual Teaming (#1162114).

7. REFERENCES

[1] Benkler, Y. (2002). Coase's Penguin, or, Linux and The Nature of the Firm. The Yale Law Journal, 112(3), 369–446. Retrieved from http://www.utexas.edu/law/journals/tlr/sources/Volume 92/Issue 2/Hemel/DSB/Hemel.fn211.01.SB.pdf

[2] Benkler, Y., & Nissenbaum, H. (2006).Commons-based Peer Production and Virtue*. Journal of Political Philosophy, 14(4), 394–419. Retrieved from http://onlinelibrary.wiley.com/doi/10.1111/j.1467-9760.2006.00235.x/full

[3] Cosley, D., Frankowski, D., & Kiesler, S. (2005). How oversight improves member-maintained communities. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '05 Human Factors in Computing Systems, 11. doi:10.1145/1054972.1054975

[4] Crowston, K. (2003). A taxonomy of organizational dependencies and coordination mechanisms. In T. Malone, K. Crowston, & G. Herman (Eds.), Organizing Business Knowledge: The MIT Process Handbook (pp. 85–108). Cambridge, MA: MIT Press. Retrieved from http://ccs.mit.edu/papers/CCSWP174.html

[5] Dourish, P., & Bellotti, V. (1992). Awareness and coordination in shared workspaces. Proceedings of the 1992 ACM Conference on Computer-supported cooperative work, (November), 107–114. Retrieved from http://dl.acm.org/citation.cfm?id=143468

[6] Forte, A., & Bruckman, A. (2008). Scaling consensus: Increasing decentralization in Wikipedia governance. Hawaii International Conference on System Sciences, Proceedings of the 41st Annual Conference, 15–17. Retrieved from ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4438860

[7] Forte, A., Larco, V., & Bruckman, A. (2009).Decentralization in Wikipedia governance. Journal of Management Information Systems, 26, 49–72. Retrieved from mesharpe.metapress.com/index/p44g2544m9471514.pdf

[8] Geiger, R., & Halfaker, A. (2013a). Using edit sessions to measure participation in Wikipedia. Proceedings of the 2013 Conference on Computer Supported Cooperative Work, 861–869. Retrieved from http://dl.acm.org/citation.cfm?id=2441873

[9] Geiger, R., & Halfaker, A. (2013b). When the levee breaks: without bots, what happens to Wikipedia's quality control processes? WikiSym '13 Proceedings of the 9th International Symposium on Wikis and Open Collaboration, 0–5. Retrieved from http://dl.acm.org/citation.cfm?id=2491061

[10] Geiger, R. S. (2014). Bots, bespoke, code and the materiality of software platforms. Information, Communication & Society, 1–15. doi:10.1080/1369118X.2013.873069

[11] Gilbert, M., Morgan, J., McDonald, D. W., & Zachry, M. (2013). Managing complexity: strategies for group awareness and coordinated action in Wikipedia. Proceedings of the 9th International Symposium on Open Collaboration, 1–8. Retrieved from http://dl.acm.org/citation.cfm?id=2491060

[12] Hsieh, Hsiu-Fang, & Shannon, Sarah E. (2005). Three approaches to qualitative content analysis. Qualitative Health Research, 15(9), 1277–88. doi:10.1177/1049732305276687

[13] Kittur, A., Chi, E., Pendleton, B. A., Suh, B., & Mytkowicz, T. (2007). Power of the few vs. wisdom of the crowd: Wikipedia and the rise of the bourgeoisie. World Wide Web, 1(2), 19. Retrieved from http://edouard-lopez.com/fac/ICPS - S7/Complexity/2008-Wikipedia-As-A-Complex-System/Power of the Few vs. Wisdom of the Crowd: Wikipedia and the Rise of the Bourgeoisie.pdf

[14] Krieger, M., Stark, E., & Klemmer, S. (2009). Coordinating tasks on the commons: designing for personal goals, expertise and serendipity. CHI '09 Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 1485–1494. Retrieved from http://dl.acm.org/citation.cfm?id=1518927

[15] Malone, T. W., & Crowston, K. (1994).The interdisciplinary study of coordination.ACM Computing Surveys, 26(1), 87–119. doi:10.1145/174666.174668

[16] Morgan, J., Gilbert, M., McDonald, D. W., & Zachry, M. (2014).Editing beyond articles: diversity & dynamics of team-work in open collaborations. Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing, 550–563. Retrieved from http://dl.acm.org/citation.cfm?id=2531654

[17] Shaw, A., & Hill, B. M. (2014).Laboratories of Oligarchy? How the Iron Law Extends to Peer Production. Journal of Communication, 1–24. doi:10.1111/jcom.12082

[18] Zhang, P, and Galletta, D. eds. (2006). Human-Computer Interaction and Management Information Systems: Foundations. Armonk, New York: M.E. Sharpe, Inc.