Learning from Open Source:
Applying Internet-Based Collaboration
to Virtual Teams

Mason Alexander Hale
August 5, 2004
http://flowdelic.org

Abstract

The nature of work is increasingly becoming geographically distributed. As economic barriers to working at a distance decline, more work is being done by individuals and teams who are remote to one another. As teams have become more and more dispersed, businesses and managers have struggled to make these remote working relationships productive. At the same time, the Internet has given rise to several communities and organizations which have been successfully, and productively, collaborating, despite being more geographically dispersed than their corporate counterparts.

In this paper, I study the workings of several successful Internet-based collaborative communities to identify what it is that enables them to succeed, even thrive, despite the highly-dispersed nature of their collaboration. This research reveals that while the practices and tools used by the referenced communities are important to their success, the most critical difference lies much deeper, in the economic-basis of their organizational structure. This economic mode of production, described by Yochai

Benkler as “commons-based peer-production”, is studied to answer two key questions: First, does this economic model, in itself, encourage more successful virtual teamwork? Second, is the peer-production model of collaboration fundamentally tied to the open source model or can it also be applied in a commercial context to create proprietary products?

The commons-based peer-production model is found to naturally promote a strong focus on quality. The factors that lead to this natural promotion are explored. The paper concludes with an analysis of the potential for adopting the peer-production model and/or practices and tools from open collaboration in a commercial context.

Table of Contents

Introduction 1

Effects of Distance on Productivity 2

Motivation 3

Approach 3

Part I: Open Collaboration on the Internet 5

Apache Software Foundation 5

Wikipedia 11

The {fray} 18

Comparison and Summary 23

Part II: The Nature of Internet-Based Collaboration 25

Measuring Success 25

Netscape to Mozilla: A Transition to Open Source 28

The Attractiveness of Quality 30

Compensation, Management and Motivation 33

Part III: Application in a Commercial Context 41

Application of Practices and Tools 41

Application of the Peer-Production Model 45

Summary 48

Conclusion 50

References 52

Vita 56

Introduction

Today, for many types of products, the distance between the place where the goods are produced and where they are sold or used, is—economically speaking—irrelevant. The rise of widely-available cheap and fast telecommunications coupled with the increasingly information-based nature of goods and services has made it much easier and more attractive for companies to locate the production of the goods and services in a different part of the world from where they are consumed.

The cost of long-distance telecommunications has steadily declined for at least the last decade.[1] At the same time and perhaps more significantly, the Internet has become a major carrier of global communications, via technologies like e-mail, instant messaging and the world-wide-web. The increasing use of the Internet further reduces the economic importance of proximity, because geographic distance generally has no impact on the cost of communications that travel over the Internet.[2] With the introduction of new technologies like Voice over Internet Protocol (VoIP), which bypasses the traditional telecommunications networks by routing telephone calls over the Internet and thereby avoiding distance-based charges, the impact distance has on the cost of communications is expected to become increasingly insignificant.

The nature of work is also changing, shifting from the manufacture of physical goods, to the exchange of information, ideas and culture. Information and culture in the form of books, newspapers, magazines, music and movies had, until recently, required being fixed to a physical medium before being sold. These products can now be produced, distributed and sold entirely in digital form.

The net effect is that information-based products, such as software, creative writing, research and expert opinion, can effectively be produced anywhere and then be inexpensively and nearly instantaneously transported anywhere else on the globe. This has resulted in more flexibility for some workers, where telecommuting has given them

[1] U.S. Department of Labor, Bureau of Labor Statistics; Consumer Price Index for long-distance telecommunications (1997-2004).

the freedom to work from home, even if they live hundreds or thousands of miles from their place of employment. It has also made it economically attractive for companies to move production of these goods and performance of these services to those locations where the work can be done at a lower cost. The types of work affected by this shift, a practice known as off-shoring, are predominantly in information technology and engineering fields, but also include a range of information-rich work, including tax preparation services and radiological evaluations (Maher 2004).

Arguing whether or not this phenomenon is economically beneficial or good public policy is not purpose of this paper. Whether for good or ill, the fundamental economics of information-based products and services has changed. Working at a distance now often costs less than working in close proximity. Communication costs no longer pose a barrier to distributing teams across the world, whether to bring together the best resources from different fields of expertise or place production where it can be done for the least cost. Given this reality, we should expect more work to be geographically distributed, and by extension we should expect more people to find themselves in distanced worked relationships with their peers and managers.

EFFECTS OF DISTANCE ON PRODUCTIVITY

In most cases, geographic separation makes establishing trust and interpersonal relationships more difficult. Creating and maintaining trust between leaders and team members is critical to the success of virtual teams. Research suggests that while the importance of trust may not be any greater in a virtual setting (Staples & Ratnasingham 1998), establishing and maintaining trust is more difficult across distance (Jarvenpaa & Leider 1998). Leaders of virtual teams face additional challenges in communicating effectively with distanced team members and combating feelings of isolation (Lipnack & Stamps 1997). In addition, increased distance has been found to decrease cooperation and persuasion, while promoting deceptive behavior (Bradner & Mark 2002). Without face-to-face interaction, interpersonal relationships are slower-forming and more fragile;

---

[2] For either the sender or receiver. The network provider(s) over whose networks the communications travel may, in fact, incur additional incremental costs for communications that travel over a greater distance.

hence virtuality may serve to erode relationships, thereby decreasing the social capital—the value of interpersonal relationship networks—in the organization (Prusak & Cohen 2001).

Despite the drawbacks associated with distanced work relationships, there exist numerous examples of productive collaborations among globally-dispersed virtual teams. Recent research suggests that virtual teams can, in fact, be *more* productive than co-located teams, when diversity is successfully exploited and communication technologies are effectively used (Majchrzak et al. 2004). The Internet provides numerous examples where wildly successful collaborations among hundreds, even thousands, of participants have been made. These include, perhaps most famously, the creation of the open source Linux operating system software. Perhaps less known, but no less successful, are projects such as Wikipedia, a free online encyclopedia with over 790,000 articles in more than a dozen languages written and edited by a virtual volunteer army of scholars, students and self-appointed experts; or {fray} a tightly-knit online community of writers, editors and artists that has succeeded in creating close relationships through storytelling, commentary and shared experiences.

MOTIVATION

The evidence that virtual teams can be just as successful as proximate teams, perhaps even more so, is encouraging, given the trend to more far-flung work relationships. Given the success demonstrated by these online collaborations, is there something to be learned from them that could help make other types of distanced work relationships more productive? Answering that question is the motivation for this paper.

APPROACH

This paper is organized into three parts. In the first part, I present descriptions of successful online collaborations as a series of case studies. For each example, I explain why the example was chosen and describe its basic purpose and accomplishments. For each group, I describe the organizational structure and governing principles. The processes, practices and tools are described, with particular attention paid to those that affect or influence communication or management.

In the second part, I explore the underlying nature of the type of collaboration illustrated by the earlier examples; to better define the nature of their success and understand what about these examples makes them successful. Here I ask the questions "how is success defined?" and "does this model of collaboration lead to more successful projects?" In this section, the economic basis of open source style collaboration is examined.

The third part examines the applicability of the practices and tools to virtual teams operating in a more traditional, proprietary commercial organization. Through this examination I identify dependencies that exist between certain practices and the fundamental economic model, and describe how these practices are naturally promoted by the underlying economics of the organizational form.

Part I: Open Collaboration on the Internet

In this section, I present three case studies of successful Internet-based collaborations. The projects described are all examples of community-driven collaboration conducted in open over the Internet. All the communities presented here invite their users to participate in achieving the community's goals. They are also all non-commercial entities. The work created by the community is made available free of charge, and (in most cases) free of copyright restrictions. Contributors to these projects are almost always volunteers. These projects are examples of a specific social and economic model of collaboration, which I will call "open network-based collaboration" or just "open collaboration" in this paper. The nature of this collaboration model is explored in detail in Part II.

For all their similarities, the examples presented here were chosen primarily for their differences. These examples cover a wide range of scales and purposes. The communities range in size from a (relatively) small and intimate circle of friends to a booming online virtual metropolis. The work created by these communities ranges from highly-technical to academic to expressive and highly-creative.

APACHE SOFTWARE FOUNDATION

The Apache Software Foundation (ASF) is a non-profit organization, formed to foster the development of open source software. The terms "open source" or "free software," describe software that is distributed in human-readable source form, thereby revealing to anyone with the interest and ability the internal workings and implementation details of the software code. The term "free" describes the unrestricted nature of license terms, not the price, although free or open software is most often distributed free of charge as well. The open source model of software development has been used to create all types of software. The count of active open source projects currently numbers in the tens of thousands (SourceForge 2004).

The ASF was born from the success of it's first and foremost accomplishment, the open source Apache web server. Apache is the world's most popular web server software, a title it has held since 1996. As of July 2004 it is used to power 67% of web sites on the Internet (Netcraft 2004). The Apache Software Foundation was incorporated in 1999 to support open, collaborative software development by providing hardware, business infrastructure and legal protection to the projects it sponsors (Apache Software Foundation 2004).

The Apache Software Foundation serves as a compelling study subject because (a) it has been very successful, (b) it has well documented, formalized organizational structures and processes, and (c) unlike some other notable open source software projects, it does not feature any highly visible individual leaders, hence its success is not attributed to the talents of any one person (O'Reilly 2004) (Apache Software Foundation 2004).

Organizational Structure and Governance

The Apache Software Foundation describes itself as a meritocracy. Governing authority is derived based upon the merit, or recognized value, of one's contribution to the foundation and its projects. This governing principle is explicitly described as the "chain of merit." Contributors achieve progressively higher levels of status and responsibility in the organization based on the amount and quality of contributions they make.

The base status level is "user", a label which explicitly acknowledges the active role users of the Apache software play in its development. Users who actively participate in discussions via e-mail, or otherwise contribute by reporting bugs, suggesting features and fixes or writing documentation, may be nominated and elected to the next status level, "committer". A committer is given the ability to directly make changes to the software code, voting rights in community-related issues, and the ability to nominate other active users for committership. A committer's status is indicated to the community in the form of an @apache.org e-mail address. An active committer can in turn be nominated for membership in the Apache Software Foundation itself. As a member, one

becomes a shareholder in the foundation and is given the right to nominate and elect board members and to nominate new projects for incubation.

Election to membership status in the Apache Software Foundation is a prestigious honor within the open source software community. ASF members take great pride in their membership; one member has described it as “one of the great honors of my life” (Leung 2004).

The ASF sponsors many projects, each of which is overseen by a Project Management Committee (PMC). The PMC is the interface between the project and the ASF. The PMC reports on project progress and requests allocation of ASF resources on behalf of the project. Projects themselves enjoy a great deal of latitude in their internal management, although they are expected to follow a core set of ASF principles.

The Board of Directors is elected by the ASF membership and is the decision making body for the Apache Software Foundation as a whole. The Board of Directors coordinates activities between internal projects and external organizations. Board approval is also required for a new project to be accepted into the ASF incubation program (described below).

Key Processes

Voting is used within ASF for a variety of purposes. Within a project, these include decisions related to the software itself, such as what bugs should be fixed before the next release and which features should be added to the software. The voting process is designed to create consensus. Rather than seeking a majority opinion, the voting method used is designed to identify objections to any proposal so that they can be addressed. Among the foundation membership, votes are taken to elect the board of directors and to decide important issues facing the foundation as a whole.

The ASF has a formal incubation process to help promising technologies develop into full-fledged projects. To enter the Apache Incubator program, a prospective project must start with a working code base, be sponsored by an ASF member or officer, and win approval of the Incubator PMC. Incubation project candidates are selected based not on technical merit but on the likelihood that they will attract a community of co-developers.

A project accepted into the ASF Incubator program is evaluated during an incubation period of pre-determined length. A project successfully exits the Incubation process if it attracts a diverse base of committers and demonstrates an ability to operate according to the principles of the ASF (Apache Software Foundation 2004).

Key Tools and Practices

ASF projects are typically comprised of a team of programmers, usually volunteers working in their spare time, none of whom are co-located with any other member of the team. Further, the geographic separation can be extreme with core team members often working on different continents and many time zones away from one another.

Mailing Lists

The primary mode of communication among ASF projects is via project specific electronic mailing lists. E-mail is favored because it is asynchronous (allowing greater participation across multiple time zones) and can be easily archived, providing a permanent written record that can be referenced by future volunteers (Apache Software Foundation 2004).

Web sites, Wikis and Blogs

Each project has a web site where project status updates are posted and the software and documentation are made available for download. This web site serves as the primary distribution point for the software created by the project.

A special type of web site, called a wiki, is used by many projects to collect best practices and develop documentation. A wiki is a read/write web site that allows anyone to add, update or delete any content on the site. The wide-open nature of the wiki enables the community keep it up-to-date with very little overhead. Users, even anonymous users, can correct errors when they see them, and even add content if they find it missing. Vandalism does occur on wikis, but is not as prevalent as one might think. The wiki software provides defensive measures to combat persistent vandals. Perhaps the strongest

deterrent is the fact that a web site designed to be freely changed offers no challenge to someone wanting to deface it.

Another important class of web site is the weblog, or blog. A blog generally takes the form of an online journal written from the perspective of a single person or small group of people. While the ASF does not provide blogs to its members, many members and contributors to Apache projects have them. This form of web site provides a distinctly personal flavor that enables people to learn more about each other without ever meeting in person. Blogs also provide an important "soapbox" function by enabling the blog author to publish opinions, comments and observations. Through cross-linking between blogs, discussion threads form connecting posts to their rebuttals.

Code Repositories and Bug Databases

The software code repository itself is an important communications medium, providing a single place to get the latest version of any file and to see what changed from version to version. Read access to the code repository is available to anyone, but only commiters are granted write access. Additionally, most projects make use of issue and defect-tracking software to track open issues and bugs. One of the most basic forms of communication between a project and its user community comes in the form of bug reports submitted by users.

Chat and Instant Messaging

The primary means of synchronous communication is chat or instant messaging. Internet relay chat, or IRC, creates a virtual channel or "room" in which multiple people can communicate with one another by sending text messages back and forth in real time. IRC channels are often used by users seeking a quick solution to a problem and to conduct real-time online meetings. Some groups have adopted formal protocols for conducting meetings via IRC that rival the parliamentary procedures of legislative bodies.[3] Instant messaging, like IRC, is text-based and real-time but is most often used as a point-to-point communication medium between two people. Text-based chat is

generally preferred to voice communications, because (a) it is less expensive, (b) it provides a written record of the communication and (c) it facilitates easier communication between non-native speakers.

Conferences

While it is not a technology, an important medium for communication between ASF members, users and contributors is the ApacheCon conference held by the Apache Software Foundation. This conference, and others, gives contributors the valuable opportunity to meet and interact in person.

Summary

The Apache Software Foundation is one of the premier open source software development organizations and serves as an exemplary model for open source in general. The ASF differs from some other major open source projects in that there is no single prominent leader at the helm. This feature indicates, at minimum, that the presence of a strong, visionary leader is not a prerequisite for success in the open source model. Also noteworthy is the high level of importance the ASF places on cultivating a large and diverse base of contributors. The ASF explicitly seeks diversity in order to insure the long term survivability of its projects and to benefit from a wider variety of perspectives and backgrounds, thereby improving the collaboration itself (Apache Software Foundation 2004). This focus on diversity agrees with the findings of a recent study of virtual teams in a corporate context, where diversity was identified as a key success factor (Majchrzak et al. 2004).

The ASF is a self-proclaimed meritocracy and it has developed formal processes to identify and recruit contributors from its user population based on the value, or "merit", of their contributions. Central to this process are multiple, well-defined tiers of participation, each with gradually more attendant power and responsibility. It is important to note that the ASF is very selective in this process, the nomination and election process serves as a filter, allowing only the most qualified and committed

[3] This observation is primarily based on conversations and correspondence with James Craig related to his experience participating in IRC-based meetings as part of the W3C Accessibility working group. A similar,

participants to reach the higher levels of its power structure. The selectivity of the process is reflected in the great honor its members ascribe to attaining higher levels of status.

While the ASF does feature some hierarchy, especially at higher levels of its power structure, the decision-making process is rarely driven by top-down directives. Instead, decision-making, especially at the project level, is mostly achieving through group consensus. Votes, when they are taken, are not decided by a simple majority. Any negative vote by any committer can essentially "veto" an issue. Hence the consensus-driven process is focused on addressing concerns and achieving resolution.

The communications technologies used in the ASF, and the way they are used is also revealing. E-mail discussion lists serve as the communications "backbone" for most ASF projects. This observation is at odds with one study that found e-mail to be a poor means of communication for virtual teams (Majchrzak et al. 2004). This discrepancy may be explained by observing that the ASF relies on e-mail primarily for group-wide communications, whereas the study cited does not indicate whether the negative conclusion it reached regarding the efficacy of e-mail was based on person-to-person or group-wide communication. I suggest the study may have focused on individual communication, and hence the conclusion reached may be more reflective of the efficacy of directed communications (relative to broadcast communications) than of e-mail generally as a mode of communication.

WIKIPEDIA

Wikipedia is a free online encyclopedia, the content of which is written and edited entirely by volunteers collaborating over the Internet. As its name suggests, Wikipedia is a "wiki", a special form of web site that makes no distinction between readers and authors, allowing anyone to add, change or remove content on the site. Originally invented by Ward Cunningham in 1995, the distinguishing feature of a wiki is the presence of an "edit this page" link or button that leads to a simple interface for changing the content of the page. Because any user can delete any content created by any other user, controversial or erroneous material is unlikely survive successive edits. Through

---

but less detailed account of an Apache Software Meeting is described in Leung (2004).

this give-and-take dynamic, the wiki format naturally forces consensus building among multiple authors (Viégas, Wattenberg & Dave 2004).

Wikipedia started as an experimental extension of Nupedia, a commercial online encyclopedia that employed paid writers and editors and used a traditional expert-review driven editorial process. After facing strong resistance from Nupedia's editorial staff, Wikipedia was launched publicly as a separate project on January 15, 2001. Wikipedia grew quickly, each month attracting new contributors and adding approximately 1,500 new articles per month (Wikipedia "History of Wikipedia"). As of July 2004, the community had grown to some 19,000 Wikipedians who had authored more than 790,000 articles in more than a dozen languages, making Wikipedia easily the largest public wiki in existence (Wikipedia "Statistics").

Compared to a print encyclopedia, the English language Wikipedia covers more territory with more than 290,000 articles, far more than the 85,000 articles found in 2002 edition of *Encyclopædia Britannica* (Wikipedia "Size Comparisons"). The English-language version is the largest Wikipedia, followed by German, French and Japanese language versions. The depth and accuracy of Wikipedia articles is inconsistent and much lower in quality, on average, compared to professionally-produced encyclopedias. However the system is self-correcting—given enough to contributions over enough time, all errors and omissions will eventually be repaired (Wikipedia "Replies to Common Objections").

Wikipedia was chosen as a study example primarily because of the extreme size and diversity of its community. Wikipedia, by its nature, attracts contributors from a much broader and heterogeneous cross-section of disciplines than would be found in a programmer-centric open source software project. The diversity of viewpoints combined with the free-for-all editing mode of the wiki, inevitably leads to conflict. Wikipedia has developed a sophisticated array of processes, tools and cultural standards aimed at diffusing controversy and promoting consensus. It is this consensus-building environment that most makes Wikipedia an attractive subject of study.

Organizational Structure and Governance

The central guiding principle of Wikipedia is the editorial policy known as the "Neutral Point of View" or NPOV. This principle seeks to eliminate personal bias from articles by describing all sides of an issue from a neutral perspective. The goal is explicitly not objectivity, but rather "presenting conflicting views without asserting them" (Wikipedia "Neutral Point of View"). While this policy itself has, at times, been the subject of debate among Wikipedians, having a clearly defined editorial policy provides the foundation that has enabled the community to reach consensus even when describing hotly-debated topics such as abortion, Scientology, God and same-sex marriage.

Wikipedia defines several different classes of user permissions that enable additional levels of capability on the site. The most basic permission level is "Anonymous", assigned to any visitor who has not created an account or logged in. Anonymous users can still edit any unprotected page, but anonymous edits attract closer scrutiny from other Wikipedians than do edits from registered users (Viégas, Wattenberg & David 2004). Any user can self-register to create an account, which enables additional features including the ability mark edits as "minor", create watchlists (explained later), upload images and rename pages. Edits made by registered users are automatically logged with the user's account name (Wikipedia "Why create an account?").

Any registered user can be nominated, or nominate themselves, to be promoted to Administrator status. Administrator status is awarded based on a consensus of the community, using a process described in a later section (see: "Request for Adminship"). Administrators are given the ability to delete and undelete pages, to protect pages from editing by any non-administrator user, and to block the editing privileges of another user. In practice, these administrative powers are used sparingly.

The infrastructure used to operate Wikipedia is provided by the non-profit Wikimedia Foundation, which is governed by a board of trustees. The board is comprised of five members including Jimbo Wales, Wikipedia's Founder, two trustees appointed by Wales, and two elected user representatives: one who represents the interests of

registered users, and another who represents the interests of all users (Wikimedia "Board of Trustees").

Wikipedia's power structure is a complex mixture of anarchy, despotism, democracy, meritocracy, plutocracy and technocracy (Wikimedia "Power structure"). It is anarchic in that any user, even malicious anonymous users, can edit any page. Also, because Wikipedia relies heavily on community consensus to make decisions, guidelines and policies are constantly evolving as the community grows and changes. Wikipedia's "despot" is its founder Jimbo Wales. Wales has paid for all of Wikipedia's operations and retains the right to veto all community decisions. Wikipedia is democratic in the sense that voting has been used to decide contentious issues. However, voting is more often used to measure consensus than to establish rule by majority opinion. Wikipedia is described as a meritocracy because quality is "the abiding goal" of Wikipedia, and so respect for the quality of a contributor's work increases her standing and influence in the community. A plutocracy is ruled by the wealthy. Wikipedia is described as a "weak plutocracy" because it implicitly gives more influence to those with the resources to commit significant time and energy to the creation and editing content. In addition, the Wikimedia Foundation has started to accept donations, which is expected to increase the influence of donors. Finally, Wikipedia is a technocracy because it relies heavily on the software and hardware used to run the Wikipedia web site. Therefore the developers who write the software and maintain the servers hold significant potential power to drastically alter the way Wikipedia works.

Key Processes

Request for Adminship

Registered users seeking administrative privileges, do so following a formal Request for Adminship process. The candidate's username is added to a specially designated page along with a written justification for the request. Comments from other members supporting and opposing the adminship request are collected on the same page. Members set their own standards for supporting adminship requests, but a common threshold is several hundred positive contributions over a 3 to 4 month period (Wikipedia

"Requests for Adminship"). If consensus supports the request after a week of consideration, the request is granted.

Vote for Deletion

Although any administrator has the ability to delete a page, except in cases of blatant vandalism or spamming, pages are rarely deleted without first gathering consensus for deletion through the Vote for Deletion process. Similar to the Request for Adminship process, a reference to the offending page is added to a specially designated page where comments for and against the suggested deletion are made (Wikipedia "Votes for deletion").

Neutrality and Accuracy Disputes

Disputes over the neutrality or accuracy of a page can be brought to the attention of the larger community via the NPOV Dispute and Accuracy Dispute pages (Wikipedia "NPOV Dispute") (Wikipedia "Accuracy Dispute"). These types of disputes generally arise when two or more authors are working on the same page and cannot come to agreement on the page content. These sorts of confrontations sometimes degrade into "edit wars" where each author alternatively adds or removes the content the other disagrees with. The availability of formal content dispute resolution processes deters this sort of counterproductive behavior by providing an alternative that is focused on consensus-building rather than controversy.

General Dispute Resolution

Beyond content-related disputes, a general dispute resolution process is defined to enable members to resolves all kinds of conflicts, including personal attacks between members and recurring inappropriate behavior on the part of any user. Rules of etiquette dictate that formal dispute resolution processes should only be invoked when more direct, informal attempts at resolution have failed. A user can request public comment on any issue via the Request for Comment page. Here the dispute is described and community members weigh in with their opinions. If a user fails to recognize or adhere to the consensus of the community, or if no clear consensus is created, the dispute can be

referred to mediation or arbitration processes. A mediator works with both parties to resolve the disagreement, but does not make binding decisions. Disputes that cannot otherwise be resolved are escalated to the Arbitration Committee which holds the power to make binding decisions affecting Wikipedia content and its users, including the ability to ban users from the site (Wikipedia "Dispute resolution").

Key Tools and Practices

It is unsurprising that the primary means of communication between Wikipedia contributors is through Wikipedia itself. Wikipedians make use of a range of communication options both internal and external to Wikipedia. Here, we look at the combination of tools used and how they are used to support the community objectives.

History, Recent Changes and Watchlists

Every page on Wikipedia preserves a history of the edits made to it. This version history allows any user to see what changes have been made, when they were made and by whom. This feature also enables users to revert a page to any previously saved state. The ability to revert a page is critical to effectively combat vandalism. A non-minor change to a page causes it to appear on a global "recently changed pages" list that is frequently monitored by Wikipedia contributors. In addition, a registered user can add any page to his watchlist, which sends the user an e-mail notification anytime a watched page changes. Through the combination of these mechanisms, Wikipedians keep close watch on content changes, and react quickly if a page they are monitoring is vandalized or drastically changed. As a result, Wikipedia is able to quickly react to and recover from vandalism. One study found that in the case of one frequent form of vandalism, the damage was repaired in most cases in less than 3 minutes (Viégas, Wattenberg and David 2004).

Meta Pages

In addition to the main article content on Wikipedia, there is, in addition, a large amount of content about the content and about Wikipedia itself. This meta-content includes the "talk" pages attached to each content page, where discussions about the

article content take place. The talk pages provide a social space where opinions can be expressed and disagreements over article content can be hashed out. Individual "user" pages, associated with each registered user account, can provide more detailed information about users, their backgrounds and motivations. User pages are editable by anyone, and so provide a space where users can pose questions, provide suggestions and give praise to their fellow contributors. Finally, there is a separate, parallel web site, Meta-Wiki that serves as a space for community building and education.

E-mail and Chat

Wikipedia hosts several public mailing lists for Wikipedia-related discussions and announcements. Registered users can also send e-mail messages to each other without exchanging e-mail addresses via an e-mail relay feature built into the Wikipedia software. A #wikipedia channel is available for users of Internet Relay Chat (IRC), who want to communicate in real-time (Wikipedia "Mailing lists").

Meet-ups

Wikipedians are encouraged to meet in person, and these face-to-face meetings are facilitated via a Meetup scheduling page where meeting places and times are proposed and set (Wikipedia "Meetup").

Summary

Wikipedia is an extreme example of large-scale collaboration. Its nineteen thousand plus contributors may very well be the single largest virtual team ever assembled. Given that the Wikipedia community exists, in essence, on a computer, it is interesting to note how infrequently technology is used to enforce community policies. The primary solutions to conflict resolution are process and culture-based, not technology-based. For example, rather than creating fine-grained security controls in the software to prevent unwanted deletions, Wikipedia lets social convention and consensus guide user behavior via the Request for Deletion process. Nothing in the software *requires* that this process be used, although adding such a feature to the software would likely be trivial. Why does Wikipedia choose imprecise consensus-driven processes over

more automated solutions? Surely an automated solution would be more consistent and expedient. Therefore, one must conclude the benefits of consensus-building somehow outweigh the negatives of slowness and inconsistency, otherwise why wouldn't an automated solution be used? I suggest that Wikipedia's consensus-driven processes provide three key benefits.

Establish Trust – first and foremost, relying on consensus to make decisions places a great deal of trust in the community. Wikipedia already places a great deal of trust in its users by allowing anyone to edit any page and consensus-driven decision making serves to deepen that sense of trust.

Increase Involvement – seeking the opinions of contributors encourages them to become more involved, by not only contributing content but by shaping the community itself.

Promote Commitment – as individual contributors see their opinions and actions shaping the community, it naturally strengthens their personal identification as a member of the community, and thereby increases their commitment to it.

I suggest that these three benefits together are critical to the creation of a sense of community, and by extension to Wikipedia's success. The importance of establishing trust among virtual teams has been well documented (Jarvenpaa & Leider 1998). I suggest that by explicitly trusting its contributors, Wikipedia more easily earns their trust in return, thereby overcoming skepticism and encouraging participation. Wikipedia encourages deeper participation by giving contributors the opportunity to increase their involvement by seeking administrator status and participating in the various dispute resolution processes. The end result, I suggest, is that Wikipedians feel personally invested in their community. They are able to look around and say, with pride, "I helped build this." It is that sense of commitment and pride that gives Wikipedia lasting power.

THE {FRAY}

The {fray} is an online community web site (www.fray.com) created by Derek Powazek in the fall of 1996. The central feature of the {fray} is the collection of personal

stories shared by its members with the rest of the community. The stories published on the {fray}, approximately one per month, are known for their honesty and intimacy. However, the fundamental purpose of the {fray} is not to collect stories, but rather to use the sharing of personal stories as a vehicle to foster interpersonal relationships and a sense of community. Whereas previous examples were of communities formed to achieve ulterior objectives, such as creating software or an online encyclopedia, the {fray} exists, explicitly, to create a community.

The {fray} presents an interesting study subject because it is so directly focused on community-building. The creator of the {fray}, Derek Powazek, has also authored a book, *Design for Community*, which provides an additional level of insight into the decisions that went into creating the {fray} and lessons learned from it and other community-focused sites. The {fray} also presents an interesting contrast to the other examples presented thus far because the community it serves is less technical, and its goals are more defined by the closeness of its community than by its size.

Organizational Structure and Governance

The {fray} is clearly the brainchild of its creator, Derek Powazek. Powazek personally designed the web site, and is a frequent contributor of stories from his own life experience. While the stories presented on the fray.com web site were contributed by some eighty-plus authors, the site itself remains under the control of its founder. A separate legal entity, the non-profit Fray Organization, has been formed to organize live, in-person events and to sell fray-related merchandise. The governing structure of the Fray Organization is not publicly disclosed.

Key Processes, Tools and Practices

Feature Story Contributions

Readers of the {fray} are encouraged to contribute their own stories to the site. Submitted features are expected to be true stories, describing personal experiences in 1,000 words or less. Anonymous contributions are not accepted, authors are required to submit stories under their real names. The criteria and process used to evaluate stories is

not described. Unlike some other prominent community-driven sites (e.g. www.kuro5hin.org), submitted story contributions are not made available to the community until after they have been reviewed and accepted. When a story is accepted, a complementary design is created to enhance its presentation on the web site. Design submissions are also accepted, but fray.com reserves the right to accept story submissions without accepting the accompanying design. While contributors are not explicitly required to relinquish their copyrights on submitted works, the site displays a notice reserving all copyrights by the {fray} on all contents (fray.com 2004).

Comments and Posted Stories

Every story on the {fray} ends with a question, to which readers are invited to respond by adding their own personal stories. For example, a story about dealing with the death of a close friend might end with the question "Have you ever lost a best friend?" Readers respond by clicking the hyperlinked word "you" which leads to the first page of comments posted by other readers (if any). Another link labeled "add your experience" is situated at the bottom of each comment page; clicking this link leads to a simple form allowing new comments to be posted to the site. Readers are not required to register in order to respond to a story, and respondents are not required to provide their real name or e-mail address when commenting, although {fray} readers often do. In his book, Powazek recommends allowing anonymous commentary to encourage participation, but he also points to a lack of anonymity as a sign of a successful community site (Powazek 2004, 144). The link to the comment form is intentionally "buried" beneath multiple links to filter out low quality comments. Powazek claims that the more clicks required to reach the posting form, the better the posts will be. The additional effort required to reach the posting form serves to screen out troublemakers who are not genuinely engaged in the story (Powazek 2004, 52).

The amount and quality of commentary associated with a feature story is seen as a key indicator of its quality and emotional evocativeness. Powazek describes the quality of the commentary on the {fray} as "consistently on topic and honest. The site is known for

the quality of its posts, and threads can go on for months, with hundreds of posts and no dip in quality" (Powazek 2004, 54).

Announcements Mailing List

The {fray} encourages readers to subscribe to an electronic mailing list to receive e-mail notifications when new stories are added as well as announcements about fray-related events. Unlike the mailing lists used by the Apache Software Foundation and Wikipedia, this is not a discussion mailing list; subscribers cannot also send messages to the list. The list is one-way. It is used to send announcements from fray.com to the {fray} community.

Fray Day and Fray Café Events

Perhaps the most impressive achievements of the {fray} and its community are the live storytelling events it sponsors. Fray Day is an annual celebration commemorating the birth of {fray} with parties held in multiple cities around the world. Fray Day 7 was celebrated in thirteen cities in October 2003, and attracted thousands of attendees. A similar event, Fray Café, was first held in Austin, Texas in 2001 in conjunction with the SXSW Music, Film and Interactive festival. Fray Café events are now held at various times throughout the year in different cities. Recordings of live {fray} events are made available for sale on the fray.com web site, with all proceeds going to the Fray Organization.

Summary

The {fray} is well respected for the closeness of its community, as demonstrated by the honesty and openness of the stories and commentary that form the basis of the site and the success of its well-attended live events. However, relative to other open collaboration examples presented thus far, the {fray} is decidedly less open. While the fray does welcome and encourage content contributions, it does not openly invite its users to help shape the community itself. Unlike our previous examples, the {fray} has no tiered membership structure whereby users can gradually take on greater levels of responsibility and receive higher levels of recognition and status in the community.

There is also very little content on {fray} about the {fray} itself. This is interesting because, apart from sharing stories, members of the community have little involvement in actual building of the community. Users can send comments about the site via an online comment form, but these comments are not shared with the rest of the community. Users mostly get to know each other through reading each other's stories. Direct contact can be made via e-mail, if a user provides an e-mail addresses, but otherwise the site does little to encourage or facilitate direct communication between community members. It is worth noting, however, that many {fray} participants also publish personal weblogs, thereby providing an additional avenue for community members to learn about and interact with one another.

The role intellectual property rights play is also interesting. A copyright notice is prominently displayed on the site reserving all rights on the contents of the site. One would assume this includes the stories themselves, which comprise nearly all the site content, yet the copyright notice is conspicuously absent from the feature story pages. The selective placement of a copyright notice on the site, at minimum, inserts some ambiguity into who owns the content. In addition, recordings of live performances from fray events are sold online, presumably with the permission of the performers. That the ownership of the creative works published and sold on the site is ambiguous is, in my opinion, not so much a sign of malfeasance as it is a symptom of the trust contributors place in the {fray} and its founder. If contributors were concerned that their work would be exploited, then one would expect intellectual property ownership issues to be more explicitly addressed to assuage those concerns.

The {fray} ultimately is a very simple web site. There are few technological bells and whistles. There is very little organizational structure or formalized process. Yet, despite its simplicity, it succeeds. By simply encouraging people to tell stories about themselves, the {fray} has successfully built a thriving virtual community. I suggest that through its minimalism, the {fray} demonstrates the potency and importance of personal communication in building trust and establishing a community. This conclusion supports earlier research which found that social, non-task oriented communications, promoted the formation of trust in virtual teams (Jarvenpaa & Leider 1998).

COMPARISON AND SUMMARY

These examples of Internet-based communities presented thus far offer a glimpse into a diverse cross-section of network-based open collaboration. The Apache Software Foundation is a leading open source software development organization, and was shown to use an array of communication technologies and formalized processes to nurture and develop promising software technologies. Wikipedia, the free online encyclopedia, is an example of a complex virtual society that is remarkable both in its sheer magnitude and extreme openness. In the case of Wikipedia, its preference for consensus-driven manual conflict resolution processes, over more expedient, but less personal, automated solutions was found to be a key contributor to its success. Finally the {fray}, the smallest and least open of the three, presents an example of the power of personal stories in building a close-knit community.

All three examples recruit contributors from their respective user populations. The ASF and Wikipedia each define a system of tiered status levels, where each level represents a higher level of responsibility and influence in the community. At {fray} a distinction exists between authors of featured stories and readers who respond by posting stories of their own, but this distinction is not formalized.

The decision-making processes of both the ASF and Wikipedia are largely consensus-driven. The ASF, through its board of directors and project management committees, utilizes more hierarchy than the other examples to manage its projects. Wikipedia also utilizes hierarchy through its board of trustees and arbitration committee, but the hierarchical management structure is mostly used as a "means of last resort," to resolve conflicts when all other consensus-based approaches have been exhausted. The decision-making processes of the {fray} are less transparent. Most decisions, including acceptance or rejection of story submissions, are apparently made personally by Derek Powazek, the site's creator.

Not surprisingly, all the examples make heavy use of Internet-based technologies to facilitate communication. The ASF and Wikipedia both make use of a wide-range of technologies, including e-mail, web sites, and chat, among others. Both exhibit a preference for asynchronous, written modes of communication. This preference is

explicitly explained as being more accommodating to non-native speakers and users situated many time zones away from one another (Apache Software Foundation 2004). The {fray} provides less communications infrastructure to its community, relying primarily on comments posted to the web site and a one-way electronic mailing list for communication.

It is worth noting that all three examples provide some means for community members to exchange personal information with one another. At the Apache Software Foundation, personal information is exchanged primarily through messages posted to project-specific discussion lists and to personal weblog web sites. Wikipedia provides a more explicit avenue for the sharing of personal information by providing each registered user with a "user page" that is intended both as a personal workspace and a place where users can describe themselves and their interests to the rest of the community. In this respect, {fray} provides the extreme example; it is focused primarily on creating strong interpersonal bonds by encouraging people to share intimate, deeply personal experiences with one another through storytelling.

Part II: The Nature of Internet-Based Collaboration

In the previous chapter, examples of successful open collaborative projects and their practices were described. In this chapter, I delve deeper into the nature of this type of collaboration in order to understand what makes the example projects successful, particularly with regard to effective collaboration among virtual teams.

MEASURING SUCCESS

A good place to start is to understand the nature of the success itself. In what ways were the examples presented successful?

For the Apache Software Foundation, success is measured by the level of involvement. Graduation from the incubation process occurs when a budding project has attracted a large and diverse group of committers. The success of the flagship Apache web server, for example, is measured based the number of servers running the software, a metric that equates to more users, which ultimately equates to more project involvement. Success is measured in terms of contributors and users.

Wikipedia closely tracks several metrics related to the number of contributors, activity and public interest. The total article count is closely monitored and milestones, such as reaching 200,000 articles, are noted and publicized. A more accurate measure of community activity, though, is number of edits, which reflects ongoing improvements to existing articles in addition to new articles. Articles with a high number of edits, if they did not serve as the battleground in an "edit war", are considered more likely to be thorough, accurate and well written. Controversy breeds higher edit rates, but even in these cases, the editing process is designed to work toward what community consensus defines as a "good" article. When the edit frequency on a controversial topic tapers off, it is indicative that consensus has been reached. The contributor population is closely tracked and segmented into new contributors, active contributors (10+ edits per month) and very active contributors (100+ edits per month). Public interest is gauged based on

number of press and Usenet mentions, search engine rankings, and web site traffic (Wikipedia "Statistics").

At {fray} success is measured in terms of participation. Great stories are identified by the number of comments posted in response to them. The success of in person events is measured in terms of audience size and reaction.

All of these measures of success are different, or at least incomplete, when compared to the traditional success measures of project management. A standard definition of a project is "a unique set of activities meant to produce a defined outcome within an established time frame using specific allocation of resources" (Harvard Business School 1996). The success of a traditional project is bounded by results, time and resources. Further, these dimensions are most often at odds with one another, requiring trade-off decisions between the three parameters. In fact, the discipline of project management itself is defined as maximizing the effectiveness of this trade-off decision (Harvard Business School 1996).

With open collaboration projects, this tension between results, time and resources is nearly non-existent. Whereas in traditional project management the focus is on minimizing the amount of effort expended to complete the project, an open source project has the exact opposite focus, seeking to maximize the number of person-hours that can be applied to the problem at hand (Raymond 2000). Rather than viewing human contributors as scarce and expensive resources to be allocated sparingly, the number of contributors is instead a measure of project success, the more the better. This constraint-free dynamic was explicitly cited by the Apache Software Foundation as one of the fundamental advantages of the open source model.

> Being no conservative resource at stake (money, energy, time), the group was happy to have new people coming in and help, they were only filtering the people that they believed committed enough for the task and matched the human attitudes required to work well with others, especially in disagreement.

The availability of a large, diverse pool on contributors provides a number of benefits to an open source project. Perhaps the primary benefit, dubbed Linus' Law, is that "given enough eyeballs, all bugs are shallow." (Raymond 2000) Simply stated, given

a larger and more diverse group of contributors, defects are more easily detected and fixed.

Open source projects are also noticeably free of the deadline pressures common to most other types of projects. In his influential essay, "The Cathedral and the Bazaar," Eric S. Raymond describes the fundamental differences between conventional, proprietary software development and open source software development. Raymond describes the anti-deadline policy of open source as the "wake me when it's done" scheduling methodology, one that eschews hard deadlines in favor of delivering software only when it is finished. Raymond also notes an alternate, less common, fixed-deadline delivery method, where releases are made at regular intervals, but no commitment is made to the feature set of each release. Either method avoids the all too common situation where the programmer is faced with both an immutable feature list and an unmovable deadline, and hence sacrifices quality in attempt to meet both objectives.

The lack of resource constraints and deadlines leaves the project results, the quality and completeness of the work, as the undisputed, overriding objective of the project. This naturally creates a results-oriented environment that is less conflicted and therefore easier to manage. If the essence of project management is effectively making trade-off decisions between results, time and resources, those decisions become much easier and hence less formal management is required.

Despite the focus on results, time and resources, most conventional software development projects are failures by this measure. The average software development project overshoots its original estimated schedule by 25 to 100 percent (McConnell 1996). If the same measure—the ability to meet original schedule estimates—were used to evaluate the success of open source software development projects, how would they fare? This question is difficult to answer, as open source projects commonly operate without setting fixed deadlines, hence a comparison of estimated schedules to actual schedules is impractical. Anecdotal evidence suggests adoption of an open source methodology mandates a quality-first focus in order to attract and retain project contributors. If quality is lacking, deadlines and resources will be sacrificed to correct the deficit, even in cases where significant deadline pressure is present.

NETSCAPE TO MOZILLA: A TRANSITION TO OPEN SOURCE

Through the latter half of the 1990's Netscape and Microsoft where locked in a heated battle for the dominant position in the web browser software market. This battle, referred to ominously as "the browser wars," was played out not only in the features of the software itself, but in licensing deals with computer hardware makers, and ultimately in the legal system. The U.S. Justice Department brought suit against Microsoft in 1999 for allegedly violating the terms of an anti-trust consent decree it signed in 1994, by using the market-dominance of its Windows operating system to coerce hardware manufacturers into making Microsoft's Internet Explorer software the default web browser on the computer systems they sold.

In January 1998, Netscape announced plans to make its Navigator browser software open source, a bold move intended to "harness the creative power of thousands of Internet developers" (Netscape 1998). At the time of this announcement, Netscape's browser software held the dominant, although tenuous, position with an estimated 60% share of the web browser market (CNN 1998). However, Microsoft was closing in fast and had Netscape on its heels. By giving away its Internet Explorer browser, Microsoft forced Netscape to follow suit, effectively cutting off its oxygen supply by drying up revenues from its flagship Navigator software product (Lea 1998). Netscape was suffering financially and had held its first round of layoffs in the weeks prior to its open source announcement, leading some to view the announcement as an act of desperation (Zawinski 1999).

Netscape formed the Mozilla project to take over development of the Netscape web browser software, adopting an open source development methodology. By most accounts, this transition did not go smoothly. In an open letter of resignation dated March 31, 1999, Jamie Zawinski, a Netscape employee and one of the project's lead developers, publicly aired his frustrations with the project.

> …And here we are, a year later. And we haven't even shipped a beta yet. In my humble but correct opinion, we should have shipped Netscape Navigator 5.0 no later than six months after the source code was released.

From Zawinski's letter we can observe that the measure of success he held himself and the project to, was shipping a new version of the software within a specific

schedule, six months. This expectation no doubt was formed based on Netscape's established, rapid pace of shipping a new version of its software approximately every six months (Iansiti & MacCormack 1996). Zawinski states that his goal was to form "a network-collaborative project in which Netscape was but one of many contributors." This also had not happened by the time of his resignation.

> For whatever reason, the project was not adopted by the outside. It remained a Netscape project. …The truth is that, by virtue of the fact that the contributors to the Mozilla project included about a hundred full-time Netscape developers, and about thirty part-time outsiders, the project still belonged wholly to Netscape – because only those who write the code truly control the project.

On June 5, 2002 the Mozilla Project shipped version 1.0 of its web browser software.[4] The release came four and one-half years after Netscape's initial announcement, and more than three years after Jamie Zawinski's resignation. By this late date, however, the browser wars were long over, with Microsoft the victor and its Internet Explorer browser accounting for more than 95% of all traffic on the web (McMillan 2004).

The Mozilla software that eventually shipped bore little resemblance to the original Netscape Navigator code that started the project. The Netscape code was ultimately scrapped and the browser was created anew, an almost total rewrite from the ground up. The decision to start from scratch was made because the existing code was "too complicated and crufty and hard to modify," a situation that created a barrier to attracting new contributors (Zawinski 1999). The Mozilla project's decision to abandon the Netscape code base drew sharp criticism from software industry veterans (Spolsky 2000), but it serves as a stark example of the uncompromising quality-driven focus naturally promoted by the open source methodology. Perhaps because the code is written out in the open for anyone to see, it becomes unpalatable to deliver code, or even to contribute to code, that is poorly architected and badly written. Whatever the rationale, the team's decision to start over, despite knowing full well the massive additional effort and schedule impact such a decision would incur is clear evidence that quality concerns overwhelmed considerations of resources and time.

Today, there are indications all the extra work is beginning to pay off. Recent measures of browser market share indicate the Mozilla browser is slowly starting to take market share away from Microsoft Internet Explorer (McMillan 2004). The Mozilla project, once derided as colossal failure, is now held up as an open source success story (Sims 2000). The project, now governed by the non-profit Mozilla Foundation, is supported by an active and diverse group of corporate and individual sponsors, and has attracted the participation of thousands of coders and testers. In addition, Mozilla has emerged as a development platform, spawning hundreds of extension projects building new functionality on top of the Mozilla software code, porting it to new hardware devices, and embedding it into new applications (Baker 2004).

While the Mozilla project is now enjoying some long awaited success, one can easily imagine that if Netscape management had known in advance how long and arduous the process would be to reach the first browser release, the project would never have gotten off the ground. The Mozilla project certainly didn't succeed in stopping Microsoft from taking control of the web browser market. Within a year of the formation of the Mozilla project, Netscape was acquired by AOL. Netscape does distribute a Netscape-branded browser based on the Mozilla software, but as a company and brand, Netscape now exists primarily as a portal web site in AOL's portfolio of web properties and serves as the brand for AOL's lower-cost Internet access services.

THE ATTRACTIVENESS OF QUALITY

What lessons can we extract from the Mozilla story? First and foremost, that it is in the fundamental nature of this model to put quality above all else. Project success, even survival, is wholly dependent on the ability of the community to attract and retain users and contributors. Without contributors, there is no one to do the work. Without users there is no reason to do the work in the first place. Users care very much about how "good" something is, and very little about how much time and effort went into making it good. In order to attract the users needed to sustain the project, the open collaboration model naturally pushes quality concerns to the forefront.

---

[4] http://www.mozilla.org/releases/mozilla1.0.html

Focusing on quality also serves the self-interests of individual contributors. Because users are primarily attracted by quality and project contributors are recruited from the user population, contributors are naturally inclined to focus on quality because it *serves their own interests as users*. Within the meritocracy-based communities that support open-collaboration projects, high-quality work proves the contributor's merit and thereby improves her standing in the community.

The importance of quality is also evident in the way success and failure is measured. Quality, by its very nature, is largely subjective and hard to quantify. Thus success metrics that count the number of users or contributors act as surrogates to direct measures of quality. Like counting the number of cars in a restaurant parking lot, counting the number of people who use a product or service and voluntarily participate in its creation, provides a good indicator of its quality. Similarly, when open collaboration projects flounder and fail, they do so not because they run out of money or time, but because the users go somewhere else. Raymond noted this dynamic at work in open source software projects.

> Anyway, in a world of cheap PCs and fast Internet links, we find pretty consistently that the only really limiting resource is skilled attention. Open-source projects, when they flounder, essentially never do so for want of machines or links or office space; they die only when the developers themselves lose interest.

The extended development timeframe of the Mozilla browser prompts the question: Is the open source model slower than conventional software methodology? Not necessarily. The GNU/Linux project, for example, has proven surprisingly quick to respond to security threats, and to adapt to new platforms such at the 64-bit processor architecture (Raymond 2000). Even critics of the Mozilla project do not blame the open source methodology for its failure to release new browser software within four years (Spolsky 2000) (Zawinski 1999). However, I argue that while open source is not inherently slower, the adoption of the open source methodology was the major cause of the delayed Mozilla release. I suggest the core problem, and the primary source of Jamie Zawinski's frustration, was that the quality-focused open source model was adopted by a team of schedule-focused Netscape engineers. That a team of one hundred full-time Netscape developers, who had previously cultivated a reputation for rapid software

releases, found themselves unable to ship software in a timely fashion suggests a fundamental change had occurred. I propose quality became a gating factor, an issue that had to be addressed before project could progress. Unfortunately, the Netscape code was of poor quality and ultimately had to be discarded before the project could attract more contributors. The lesson here is that participants who bring expectations of tight shipping schedules to an open source project may be sadly disappointed if the inherent, fundamental quality objective is not already met.

While I argue that quality acts as a powerful attractive force, drawing new users and collaborators to projects and thereby inducing projects to adopt a quality-driven focus, I am not suggesting that quality is the *only* thing that attracts people to products or projects. There are many other considerations, beyond quality, that factor into individual and group decision-making behavior. For example, the size of a community may, in itself, attract people who are less interested in the project objectives than in simply connecting with other people and belonging to a community.

Just as in the marketing of products, quality plays an important role, but does not necessarily insure success. There are many factors, intrinsic and extrinsic to the product, that may influence its acceptance by the marketplace. For example, in the 1980's the technically superior Betamax video format lost to the competing VHS format, because the industry wanted to support a single video format and lower-priced VHS video players had attracted a larger installed base of users.

I suggest dynamics similar to those seen in product marketing are also present in the formation and growth of open collaboration communities. Larger communities may attract more members simply by virtue of being large than due to the quality of the collaboration or the work it produces. The quality of a product should not, therefore, be measured by the size of the community that built it. In fact, I expect that as a community grows, its need to focus on quality in order to attract new collaborators will diminish as other attractive forces, such as size, begin to take effect. In addition, as the pressure to focus on quality is lessened, a large community attracts new participants who may be attracted more by the size of the community than by the opportunity to work on something of high quality. In particular, one would expect large communities to attract

higher numbers of miscreants such as vandals, trolls (users who purposefully post inflammatory comments to stir up trouble) and scam artists.

COMPENSATION, MANAGEMENT AND MOTIVATION

When people are paid for their contributions, the open collaboration environment that naturally promotes quality and participation is fundamentally disrupted. This disruption occurs at two levels: (1) at the managerial level, in the way resources are selected and managed, and (2) at the individual level in terms of human motivation.

Managing Abundance and Scarcity

When resources are monetarily free there is no upper limit on the number who can contribute. The level of participation is only limited by the ability of the project to attract and retain interested, motivated contributors, and its ability to efficiently aggregate the individual contributions. Project success is derived not from efficient use of scarce resources, but through the effective direction of abundant resources.

When people are paid for their work, the opposite occurs. Instead of being abundant and free, resources become scarce and costly. As a result, management focuses on carefully controlling the use of the resources to reduce costs. When considering trade-offs between time, resources and results, resource scarcity shifts more focus to the time and resources side of the equation. Quality is no longer *naturally promoted* in this scenario, and managers who wish to make quality a top priority must do so through conscious and concerted effort. This is not to say compensation and quality are incompatible—they are not. Rather, compensation naturally creates a need to manage for efficient resource utilization. The need for efficiency can sometimes come in conflict with the need for quality, resulting in the need to make trade-off decisions that would not have been required in situations where resources are free.

Compensation also serves to discourage open criticism. Fearful of retaliation, employees generally do not feel free to openly express dissenting opinion or to disagree with their managers or company policy. As a result, conflicts often linger below the surface, unspoken and unresolved. Worse, constructive criticism and debate that might lead to dramatic improvements never happen.

Human Motivation and Coase's Penguin

Intuitively, one would expect the addition of monetary compensation, all other things being equal, to increase motivation and thereby increase productivity. It follows then that volunteer-based projects would have difficulty attracting large numbers of motivated, qualified contributors. Empirically, this has clearly not been the case in the studied examples. We can conclude then that the *absence* of monetary compensation did not inhibit motivation or participation. The question then is does the presence of monetary compensation somehow discourage motivated participation and effective collaboration?

In his paper, "Coase's Penguin: or Linux and *The Nature of the Firm*", Yokai Benkler proposes an economically-grounded model to explain how open collaboration projects attract participants and direct their work. In *The Nature of Firm* (1937), economist Ronald Coase first described an economic model to explain the fundamental differences in the ways markets and firms direct behavior. Markets influence investor behavior through price signals to balance supply and demand. Firms direct employee behavior through hierarchical authority and contracts. Benkler extends Coase's model to include a third model, which he calls "commons-based peer production."

Benkler defines commons-based peer production to include open source software development, as exemplified by the Apache Software Foundation, and other forms of collaboration including Wikipedia, Amazon.com product reviews and NASA's "clickworkers" program where volunteers on the Internet were enlisted to identify craters from images of Mars' surface, among others.

The primary advantage of the peer-production model is that it is better able to identify and allocate human creativity to the pursuit of project objectives. This capability relies on "very large aggregations of individuals independently scouring their information environment in search of opportunities to be creative in small or large increments" (Benkler 2002, 7). Key factors affecting success in the peer-production model are:

Access to a very large pool of potential contributors

Granular division of work into units of various size

Effective communication of available work opportunities

Efficient integration of individual contributions

The central premise of the peer-production model is that individuals are best able to ascertain their own ability and suitability to perform a specific task. Therefore, by dividing work into different size units and advertising contribution opportunities to a large pool of potential contributors, the best available resources will be allocated to the tasks through self-selection. Peer-production works by assimilating many, many contributions, large and small, from a very large base of self-selected contributors. The viability of peer-production to create any given information-based product is limited by the capability to efficiently integrate work created independently by many contributors into the final product (Benkler 2002, 70). This integration capability must also be able to filter out unwanted contributions generated by incompetent or malicious participants.

It is important to note that Benkler does not describe peer-production as requiring voluntary contributions. However, there are cases where property rights and compensation can undermine the effectiveness of the model. Specifically, Benkers argues that assertion of intellectual property rights automatically limits the size of the contributor pool to those with access to the protected property.

Further, the willingness of contributors to voluntarily donate their time and effort to a project hinges on the output of that project being a "commons," that is a shared resource made freely and equally available to all. If, on the other hand, the project output is to become private property, contributors will naturally expect to be compensated for their effort. Contributors are assumed to be altruistic only when their work benefits the common good.

When contributors are monetarily compensated, it introduces transaction costs which, in turn, raise the minimum threshold for the size of contributions. Paying someone any amount of money incurs transaction costs associated with accounting for the payment and complying with applicable tax and labor regulations. If the value of very small contributions, such as writing a single encyclopedia entry, does not cover the transaction costs, then accepting very small contributions ceases to make economic sense. The net effect is compensation requires an increase in the minimum acceptable contribution size

to cover the transaction costs associated with paying for the work. This further shrinks the contributor pool by limiting it to those with the time to make larger contributions.

Stated simply, the peer-production model does not preclude paying contributors for their contributions, but monetary rewards do appear to limit the scale of problem to which the model can be applied. The most extreme examples of the model in action, such as Wikipedia, where countless edits from thousands of contributors are assimilated into a cohesive, encyclopedic work, can only occur in environments free of the transaction costs associated with compensation. Large communities of voluntary contributors, in turn, will only arise when the work output is a commons, free of property rights restrictions. Thus, the efficacy of the peer-production model, especially on large-scale problems, is inexorably linked to intellectual property rights, and strongly favors scenarios where property rights are relaxed or nonexistent.

To explain the factors that influence individual behavior in this mode of production, especially factors that motivate contributions to peer-production projects without monetary compensation, Benkler presents a mathematical model to explain the combinatorial effects of types of rewards on human motivation. He defines three types of rewards: (M) monetary rewards, (H) hedonic, or self-gratifying, rewards, and (SP) socio-psychological rewards, including social status. When faced with a choice between different courses of action, a rational decision-maker is assumed to decide based on the combined, perceived value (R) of all three reward types. Thus,

$$R = M + H + SP$$

Benkler shows that the presence of monetary rewards can negatively affect the perceived socio-psychological benefit of a given course of action, depending on the situation. For example, an act of love takes on negative SP value if it is made in exchange for money. In these cases where M is negatively correlated with SP, the addition of monetary rewards can, in fact, reduce the total perceived utility of a given course of action. This effect is captured in the variable *p* (for "professional" or "prostitute") which modifies the value of SP and can be either positive or negative. A positive value indicates that monetary rewards enhance the socio-psychological value of a course of action, as is the case among professional athletes or Hollywood actors. A negative value indicates a

negative correlation between money and the socio-psychological value of the option, as is the case with prostitution or bribery. Larger *p* values indicate a stronger interaction between monetary and socio-psychological rewards.

Monetary rewards may also influence the value of SP to the degree that some participants are paid while others are not, or where some are paid more than others. Benkler denotes this factor as *jalt* (jealosy/altruism). A positive value indicates a situation where altruism is culturally rewarded. Negative values capture situations where one party is jealous of the rewards of another.

Finally, the incremental motivational benefit of monetary rewards generally decreases as the monetary rewards increase and as material needs become satiated. This effect is denoted as the variable *s*, a modifier of M.

Taking into account the combinatorial effects between M and SP rewards, and the satiation effect of increasing amounts of M, Benkler's motivational model becomes:

$$R = M_s + H + SP_{p,jalt}$$

From this model we can observe that in cases where values of *p* or *jalt* are large (either positive or negative), the impact monetary rewards (M) have on the perceived socio-psychological rewards (SP) is more pronounced. We can also observe that when monetary rewards are *not* present and the value of *jalt* is neutral or positive, these interactions do not occur and the model can be simplified to:

$$R = H + SP$$

That the value of *jalt* is positive or neutral is critical, because a negative value would indicate a situation where one perceives potential for exploitation or inequity. This would be the case, for example, if one is not being compensated and perceives the benefits of one's contribution will be appropriated by another party for private financial gain. It is the influence of the *jalt* factor in the equation that captures the way property rights limit the upper scale of the peer-production model.

In relation to the examples presented in Part I, we can see that while all three examples solicit voluntary contributions, only the Apache Software Foundation and Wikipedia explicitly make the end work product available under terms free from property rights restrictions. In contrast, the {fray} explicitly, if ambiguously, claims copyright

ownership over the work produced. I suggest that in the case of {fray}, the assertion of property rights is a limiting factor on the level of contributions received, but that this effect is somewhat dissipated by the presence of a strong trust relationship that serves to offset fears of exploitation or inequity.

When the model is simplified to remove the presence of monetary rewards and their associated interactions with other motivating factors, one can clearly see that such projects are fully reliant on hedonic and socio-psychological rewards to motivate their contributors. Hedonic, or self-gratifying, rewards are internally and individually determined and are thus are difficult to externally influence or provide. Therefore, one would expect projects that provide greater socio-psychological rewards to attract more contributions.

Socio-psychological rewards are, in themselves, complex and difficult to unravel. Abraham Maslow described these as a hierarchy of needs, where basic needs for physiological well-being and safety must be met before higher order needs for love, esteem and self-actualization can be satisfied (Maslow 1970). Open collaboration projects that recruit volunteers from the Internet do so by appealing to individuals seeking to fulfill these higher order needs.

Participation in group collaboration, whether virtual or in-person, serves the human need for love and belongingness, to connect with and be accepted by other people. From the examples in Part I, we see this need being served in all the examples, perhaps most explicitly by the deep, emotional connections encouraged by the {fray}.

The esteem needs, including self-esteem, are met by providing opportunities for recognition and indicators of status. For example, the @apache.org e-mail address of a committer within the Apache Software Foundation provides a visible status indicator, and the great honor of being elected to membership in the Apache Software Foundation provides both public recognition and personal self-esteem benefits.

The need for self-actualization is “the desire to become more and more what one is, to become everything that one is capable of becoming.” This need is served in the example projects by merely providing opportunities to work on challenging problems. I

suggest the desire to fulfill the need for self-actualization is the primary and strongest attractor of participants to open collaboration projects.

In "The Cathedral and the Bazaar" Raymond concludes:

> I want to suggest what may be a wider lesson about software, (and probably about every kind of creative or professional work). Human beings generally take pleasure in a task when it falls in a sort of optimal challenge zone; not so easy as to be boring, not to hard to achieve. A happy programmer is one who is neither under-utilized nor weighed down with ill-formulated goals and stressful process friction. …It may well turn out that one of the most important effects of open source's success will be to teach us that play is the most economically efficient mode of creative work.

Describing open source software development as "play" is an important distinction that drives to the heart of the motivation problem. As Benkler has shown, multiple, inter-related factors affect motivation, but ultimately the people who voluntarily work on these types of projects do so because they *enjoy* the work and take pleasure in the creative act itself. Whether called "hedonic gain" or an "urge to create," the fact that loving what you do is a strong motivational force, should come as no surprise. Yet it is easily lost among the other, more easily controlled concerns. As Raymond's description of a happy programmer implies, enjoyment cannot be manufactured, but can be easily undermined by through boredom, unrealistic goals, undue stress and other factors. The best a manager can do to foster joy in the work environment is to remove those obstacles that undermine it.

The human need for self-actualization, to become better and to enjoy what you do, is always present—whether in a paid or voluntary position. While there is some evidence to suggest that monetary rewards, in themselves, actually undermine personal motivation by turning play into work (Kohn 1993), I contend that compensation undermines personal enjoyment by concealing its absence. Enjoyment is difficult to measure or observe; it is a distinctly personal and internalized experience. Therefore it is outwardly difficult to discern whether someone is doing a job because they enjoy doing it or because they are paid to do it.

I suggest it is this enjoyment in the work and growing sense of pride and accomplishment that keeps contributors involved and projects going in the absence of

monetary rewards. Among open collaboration projects, if the joy is lost and the work turns to drudgery, the collaborators will lose interest and leave, and the project will fail. This, I suggest, creates a self-correcting environment where dips in motivation are more readily observed and quickly remedied to insure project survival. As a result, in the same way the need to attract contributors naturally promotes a focus on quality; the need to retain contributors naturally puts significant focus on keeping them happy.

Part III: Application in a Commercial Context

In Part I, I described several examples of Internet-based open collaboration. Part II delved deeper the model itself and examined its economic and psychological underpinnings. Here, I look at the feasibility and approaches for applying positive attributes of open collaboration to virtual teams working in a commercial context. By "virtual teams working in a commercial context" I mean teams comprised of employees, contractors and/or vendors, who are geographically situated apart from each other, and who are working under the direction of a commercial enterprise. These are paid contributors working for a company that expects to own and exercise property rights over the end work product.

I will examine the application of open source principles and practices from two distinctly different directions. First, I will identify possibilities for adopting the practices and tools common to open collaborative environments, without also adopting the underlying economic model. In doing so, I will identify interrelations between certain practices and underlying principles embedded in the economic model. Secondly, I will explore potential ways that a commercial enterprise might incorporate the peer-production model itself into its operations, either directly or indirectly.

APPLICATION OF PRACTICES AND TOOLS

Communications

At the surface level of communication practices and tools, there is much that can be learned from open collaboration projects and applied to virtual teams. But first one must assess how closely a given virtual team resembles an Internet-based example of open collaboration. For example, is the team widely distributed across cultures and time zones? Is there a language barrier? Do different contributors have different levels of involvement, or is every team member exclusively dedicated to this project? The answers to these questions will help guide the selection of appropriate communications tools and practices.

In the examples studied, we found a range of technologies in use, including e-mail, chat and different types of web sites. In addition, there was a strong stated preference for using asynchronous, written communications such as e-mail. This preference contradicts the finding of a study of successful, commercial virtual teams, where e-mail was found to be a poor means of communication, in favor of conference calls and collaborative web sites (Majchrzak et al. 2004).

I previously suggested this discrepancy could be explained by distinguishing between direct person-to-person e-mail communications and group discussion-oriented mailing lists. This lack of agreement may be further explained by the size of the teams involved, where the greater the size of the team, number of locations, and geographic separation (especially across time zones) increases the logistical complexity of getting everyone on the phone at the same time. In the case of a large Internet-based collaboration where each contributor is in a separate location and the team may be spread across multiple continents, conducting a phone call in real-time may simply be impractical. However, another real-time medium, Internet Relay Chat is commonly utilized, indicating that the higher cost of international phone calls is the more likely deterrent. It will be interesting to see if open source groups begin to switch from text-based chat to voice communications with increased adoption of technologies that route voice communications over the Internet and avoid long-distance charges.

Still, the use of written forms of communication provides two benefits that should not be overlooked (1) it automatically provides a written record of the communication that can be easily archived and searched (2) it helps bridge language barriers by accommodating non-native speakers who are able to communicate more clearly in writing.

Another noteworthy feature of open collaboration projects is that the "hub" of communications is typically a web site, or some other virtual workspace, rather than a physical location such as a home office or headquarters. Locating the central repository of information in virtual space, where everyone has equal access to the information, and making sure all status information and other important communications are kept there helps to alleviate feelings of isolation and "being out of the loop" that can occur when a

large group of team members work in close physical proximity, while also working with a smaller remote team. The previously cited study corroborates this conclusion, finding that 83% of the teams it studied relied on virtual workspaces for communication, and tended to use teleconferences to resolve disagreements (Majchrzak et al. 2004).

The creation and utilization of a central communications hub may also serve to offset the additional communication overhead associated with larger teams. Where interpersonal communication is required to coordinate tasks, the additional effort associated with communication increases as n (n - 1) / 2, where n is the total number of team members (Brooks 1975). This formulation assumes that each team member needs to individually communicate with every other team member. However, if a centralized communication hub, such as a wiki or other virtual workspace, is used, each team member only needs to direct communications to the hub in order to effectively reach the rest of the team. Here it is important that the communications be organized to mirror the organization of the work product, so that communication about a specific portion of the work can be made in context to that portion. For example, the "talk" pages associated with each encyclopedia entry on Wikipedia provides a centralized and contextually-relevant workspace for communications about a given entry.

Finally, perhaps the most valuable communication related finding that can be applied to virtual teams is the importance of incorporating personal communications to build trust and foster cohesion among virtual team members. The extremely personal stories shared via {fray} are most likely too personal for a professional corporate environment, but they demonstrate the powerful role personal communications play in building a sense of community. Perhaps a better model to emulate would be the "user" pages of Wikipedia, by providing individual team members with a dedicated space for self-expression.

Decision-making

Valuable lessons can be extracted from the decision-making processes used by the examples in Part I, but adopting or emulating these processes may pose significant challenge in a hierarchically driven organization.

At was described at length in Part II, the free availability of abundant resources combined with a need to attract participants to voluntary contribute to a project creates an environment where quality concerns are naturally pushed to the forefront. Within a commercial context, we cannot expect resources to be free, and so trade-off decisions between results, resources and time will likely need to be made in a commercial context that would have been unnecessary in an open collaboration environment. Therefore a manager will need to exert conscious effort to maintain a quality focus. The literature on quality management is extensive, and will not be repeated here, but suffice to say that in many cases, a quality-focus is highly desirable.

Further, in addition to trade-off decisions between resources, results and time, a manager will likely need to deal with resource allocation, adding yet another layer of decision-making overhead. Recall that in a peer-production model, resource allocation decisions are made by self-selection rather than managerial direction. The key advantage of the peer-production model, as described by Benkler, is the ability for agents to self-select assignments to which they are particularly well suited. The salient distinction here is that self-selection process is inherently qualitative; it seeks the *best* resource from the available pool, not just a number of resources. In a corporate context, where resources are generally marshaled by managers, this introduces the potential for information loss, reflected in the gap between the manager's knowledge of the abilities and availability of all resources, and each resource's knowledge of his or her own abilities and availability. Short of adopting a model that allows resources to select their own assignments, this observation highlights the need for managers to offset this information loss and make better resource allocation decisions by being better informed of the capabilities of available resources.

In Part I, consensus-driven approaches were shown to be the dominant method used for decision-making and conflict resolution among the example projects. With Wikipedia in particular, consensus-driven methods were preferred even in cases where policies could more efficiently be enforced via automation. In Part II, I concluded that open collaboration environments naturally induce a focus on keeping contributors happy in order to retain their continued interest and involvement. I suggest that a preference for

consensus-driven approaches is a symptom of this natural need to promote contributor happiness. More specifically, I propose that consensus-building activities invite deeper involvement and commitment, by encouraging a greater sense of personal investment in a community through direct involvement in shaping it via participation in the decision making process. The nature of the relationship between consensus-driven decision-making and motivation is a subject worthy of further study.

For a manager wanting to adopt successful practices from open source methodology, I encourage taking a close look at consensus-driven approaches as a way to potentially increase motivation and commitment and to distribute the overhead of decision-making. However, caution is advised as consensus-building can require more time and effort, and can, in some cases, devolve into "design by committee" resulting in over-compromised, lower-quality decisions.

Motivation

In Part II, significant space was dedicated to the analysis of the complex social and psychological factors that influence motivation. The salient observations for a manager of virtual teams are: (1) recognition that monetary rewards are only part of the equation and (2) the presence of monetary rewards can mask the absence of other, more powerful, motivational influences, including personal enjoyment. The practice to be emulated then is not the removal of monetary rewards, but the addition of socio-psychological rewards such as public recognition. In addition, to maximize motivation and commitment, managers should act to eliminate factors like stress and unrealistic deadlines that serve to undermine personal enjoyment of the work itself.

APPLICATION OF THE PEER-PRODUCTION MODEL

There has been significant interest on the part of businesses to find a way to somehow harness or otherwise co-opt the success of the open source development model. There have been some notable successes here, such as RedHat, the leading distributor of the Linux open source operating system, and OSDN, the publisher of SlashDot and other sites that provide news to the open source community, which in turn are built using open source technologies. But there have been failures too, as evidenced by Netscape, whose

attempt to hold Microsoft at bay by creating the open source Mozilla project failed to come to fruition in time to save the company.

For most commercial software companies, the open source model of licensing and distribution is fundamentally at odds with existing business practices. The implementation details of the software are highly valued intellectual property – hence the source code is protected as a closely guarded trade secret. For the traditional software company, distribution of source code effectively eliminates the commercial value of the software, because it potentially enables anyone, even competitors, to modify, re-package and/or re-distribute the software at no cost.

The question then is: can a company adopt an open source style product development methodology without also giving up control of its intellectual property rights over the work produced by it? In answering this question, one must make a distinction between emulating the practices and tools used by open source communities, which I have described in the previous section, and adopting the underlying economic and motivational framework, which I equate with Benkler's peer-production model. In this section I will explore the potential for the latter option.

In his paper titled "The Open Source Paradigm Shift" Tim O'Reilly states "this cultural shift may have had its first flowering with open source software, but it is not intrinsically tied to the use of open source licenses and philosophies." At first glance, this statement appears to disagree with Benkler's assertion, described in detail in Part II of this paper, that a linkage exists between property rights and monetary incentives which ultimately serves to limit the number of potential, willing participants. Closer examination, however, reveals that O'Reilly is making the point that open source style development practices do not *require* the relaxed property rights restrictions of open source-style licensing agreements in order to function. On this point, Benkler does not disagree, but he goes further to show that the assertion of property rights, while not incompatible with the model, do significantly undermine its effectiveness.

Citing examples from Microsoft and HP, O'Reilly goes on to assert that "given enough developers and a network to connect them, open source-style development behavior emerges." His point here is that Microsoft and HP are very large companies

with large pools of resources that can be applied to a given problem. I would argue here that the definition O'Reilly uses for open source style development is somewhat more relaxed than Benkler's definition of commons-based peer-production. Whereas Benkler describes the self-selection resource allocation model as a fundamental component of peer-production, in his examples, O'Reilly does not describe scenarios where work is broken down into highly-granular units to be performed by a multitude of self-selected contributors. While the examples given by O'Reilly may be instances where practices and tools first popularized by open source development have been adopted by commercial companies, they do not meet Benkler's definition of commons-based peer-production.

Benkler draws a clear line between direct and indirect appropriation of benefits created in the peer-production model. Direct appropriation occurs when property rights are asserted over the end product of the collaborative effort. It is this case where efficiency of the peer-production model is generally reduced by the effects of property rights. However, there is a case where direct appropriation can occur successfully, and on a large scale; this is when the size of an individual contribution is small enough that monetary or socio-psychological incentives necessary to bring it about are trivial. This is the case, for example, with customer-written product reviews on Amazon.com. Collectively, these reviews significantly enhance the value of Amazon.com and give it an edge over its competitors. Although Amazon technically owns these product reviews—one would expect them to vigorously fight an attempt by a competitor to co-opt them—they are given freely to Amazon by its customers because writing a review requires little effort, and doing so delivers the socio-psychological benefits of helping one's fellow consumers and a modicum of individual recognition. In effect, these socio-psychological benefits serve as the necessary payment to attract the contributions.

When it comes to *indirect appropriation* of the benefits of peer-production, Benkler's and O'Reilly's descriptions and examples begin to converge. Indirect appropriation occurs when a peer-produced work product is incorporated into a commercially-delivered product or service, but no property rights are asserted over the portion that was created via peer-production. This would occur, for example, when a

software-based service, like the Google search engine, is created using open source software and tools, like the Linux operating system.

O'Reilly argues that the real potential for the commercial exploitation of open source methodology lies with these indirect forms of appropriation, although he does not make the same distinction between direct and indirect appropriation as Benkler.

Benkler describes the possible viability of a "cooperative monetary appropriation" model, where the community that produces the product retains ownership rights over it and sells it, distributing the proceeds to each contributor commensurate with his or her level of contribution. In addition, Benkler suggests the potential for hybrid-models where peer-production attributes are incorporated into a firm or market economic model. Although Benkler does not cite it as an example, Google Answers may be an example of such a hybrid model. Google Answers is a web site where users post questions they need answered, and attach a price, say $20, they are willing to pay for a well-researched answer. This creates a market in which researchers compete to provide acceptable answers and collect the monetary bounties. At the same time, it embodies a peer-production model where a self-selected group of contributors is gradually building content in small-sized increments.

SUMMARY

In summary, there are many lessons to be learned from the way open collaboration is conducted. These include methods for effectively communicating with large, geographically distributed populations of collaborators and enhancing productivity and commitment by stimulating social and psychological motivational drivers. Managers should be aware, however, that adopting these practices may not come as naturally outside the peer-production model. In particular, resource allocation and trade-off decisions that occur naturally within the model will instead require conscious management effort outside of it.

While it may be relatively easy and worthwhile to adopt practices and tools proven via open collaboration environments, the benefits of adopting the surface-level features pale in comparison to the potential benefits that can be derived by tapping in the

peer-production model itself. As demonstrated by companies like Amazon.com and Google, the benefits of peer-production can be harnessed for competitive and financial advantage. Indirect appropriation may be the most directly applicable means for a commercial company to utilize the peer-production model, but potential also exists for direct appropriation and hybridized models.

Conclusion

In this paper I present case studies of three successful, yet diverse, Internet-based open collaboration projects. For each example, the governing structure, processes, practices and tools are analyzed to better understand the individual factors that contribute to its success. Success factors identified include:

A tiered system of social status that encourages increasing levels of user participation and responsibility

Consensus-driven decision-making processes, which serve to establish trust and increase levels of commitment in the community

Sharing of personal information to promote interpersonal bonds and community cohesion

The social, psychological and economic underpinnings of Internet-based collaboration are examined in depth. I suggest that because open collaboration projects effectively compete for the attention of abundant and freely available human resources, and these resources are primarily attracted to projects by their quality, this results in a naturally-occurring, strong focus on quality, over competing concerns of resource utilization efficiency and time. This inherent focus on quality is shown to be a gating factor that must be met before other project concerns can be addressed. I speculate that this attractive power of quality gradually dissipates as a community grows in size.

An economic model devised by Yokai Benkler, commons-based peer-production, is described and is used to explain the complex role monetary rewards play in human motivation, specifically with regard to open collaboration projects. Non-monetary, social and psychological rewards are revealed to be central to attracting and retaining the interest and participation from contributors in open collaboration projects. I suggest this phenomenon results in the creation of a self-correcting work environment that puts significant focus on keeping contributors happy by supporting their needs for personal enjoyment and self-actualization.

I examine the potential for open collaboration to be applied in a commercial context. This examination considers separately the adoption of practices and tools

commonly used on open collaboration projects, and the adoption or appropriation of the more fundamental economic model. I find that potential exists for virtual teams to benefit by emulating some of the communication and motivational practices of open collaboration projects, although this benefit is likely to be decreased by additional management overhead required in a commercial context that is unnecessary in an open collaboration context. Finally, I suggest that the peer-production economic model can be harnessed in commercial contexts for competitive and financial gain via a variety of approaches, including direct, indirect and hybrid appropriation models.

References

Apache Software Foundation (2004) "How the ASF Works!" http://www.apache.org/foundation/how-it-works.html (accessed: May 3, 2004).

Baker, Mitchell (June 14, 2004) "Time flies when you're having fun." *Mitchell's Blog*. http://weblogs.mozillazine.org/mitchell/archives/2004/07/time_flies_when_1.html (accessed: July 18, 2004).

Benkler, Yochai (2002) "Coase's Penguin, or, Linux and *The Nature of the Firm*." *Yale Law Journal.* (112:369).
(available from: http://www.benkler.org/CoasesPenguin.html).

Bradner, Erin and Gloria Mark (2002) "Why Distance Matters: effects on cooperation, persuasion and deception." *Proceedings of the 2002 ACM conference on Computer supported cooperative work*. New Orleans.

Brooks, Frederick P., Jr. (1975) *The Mythical Man Month*. Reading, Massachusetts: Addison-Wesley Publishing Company, Inc.

CNN (October 9, 1998) "Behind the numbers: Browser market share." http://www.cnn.com/TECH/computing/9810/08/browser.idg/ (accessed: July 18, 2004).

Fray.com (2004) "Contribute to Fray." http://www.fray.com/is/you/
(accessed: July 22, 2004).

Harvard Business School (1996) "Project Management Manual." revised: August 24, 2001. Cambridge.

Iansita, Marco and Alan MacCormack (1996) "Living on Internet Time: Product Development at Netscape, Yahoo!, NetDynamics, and Microsoft." Harvard Business School. rev. June 30, 1999.

Jarvenpaa, S. and D. E. Leider (1998) "Communication and trust in global virtual teams." *Journal of Computer-mediated Communication* (3:4).

Kohn, Alfie (1993) "For Best Results, Forget the Bonus." *The New York Times*. October 17, 1993.

Lea, Graham (October 28, 1998) "USA versus Microsoft: the first week." *BBC News*. http://news.bbc.co.uk/1/hi/special_report/1998/04/98/microsoft/201714.stm
(accessed: July 12, 2004).

Leung, Ted (May 22, 2004) "Apache Software Foundation Members Meeting." *Ted Leung: On the air*. http://www.sauria.com/blog/2004/05/22#941 (accessed: July 18, 2004).

Lipnack, Jessica and Jeffrey Stamps (1997). *Virtual Teams: Reaching across space, time, and organizations with technology.* New York: John Wiley & Sons.

Majchrzak, Ann, Arvind Malhotra, Jeffrey Stamps and Jessica Lipnack (2004) "Can Absence Make a Team Grow Stronger?" *Harvard Business Review.* May 2004.

Maher, Kris (2004) "Now in Offshoring's Sights: High-Level Professionals." The *Wall Street Journal Online*. http://www.careerjournal.com/myc/survive/20040324-maher.html (accessed: June 14, 2004).

Maslow, Abraham (1970) *Motivation and Personality*. 2nd Ed. Harper & Row.

McConnell, Steve C. (1996) *Rapid Development.* Redmond, Washington: Microsoft Press.

McMillan, Robert (July 9, 2004) "Mozilla Gains on IE." *PC World*. http://www.pcworld.com/resource/printable/article/0,aid,116848,00.asp (accessed: July 14, 2004).

Netcraft (July 2004) "July 2004 Web Server Survey." http://news.netcraft.com/archives/2004/07/01/july_2004_web_server_survey.html (accessed: July 7, 2004).

Netscape Communications Corporation (January 22, 1998) "Netscape Announces Plans to Make Next-Generation Communicator Source Code Available Free on the Net." (available from: http://wp.netscape.com/newsref/pr/newsrelease558.html).

O'Reilly, Tim (May 2004) "The Open Source Paradigm Shift." *O'Reilly Network.* http://www.oreillynet.com/pub/a/oreilly/time/opensource/paradigmshift_0504.html (accessed: July 12, 2004).

Powazek, Derek M. (2002) *Design for Community: The art of connecting real people in virtual places.* Indianapolis: New Riders.

Prusak, Laurence and Don Cohen (June 2001) "How to invest in social capital." *Harvard Business Review.* June 2001.

Raymond, Eric S. (2000) "The Cathedral and the Bazaar." *Thyrus Enterprises*. version 3.0. (available from: http://www.tuxedo.org/~esr/).

Sims, David (March 14, 2000) "Mozilla Interview: Brendan Eich and Mitchell Baker." *O'Reilly Network*.

http://www.oreillynet.com/pub/a/network/2000/03/10/mozilla.html (accessed: July 18, 2004).

SourceForge (2004) "SourceForge.net Statistics." http://sourceforge.net/ (accessed: July 11, 2004).

Spolsky, Joel (April 6, 2000) "Things You Should Never Do, Part I." *Joel on Software*. http://www.joelonsoftware.com/articles/fog0000000069.html (accessed: July 2, 2004).

Staples, D. Sandy and Pauline Ratnasingham (1998) "Trust: The Panacea of Virtual Management?" *Proceedings of the international conference on Information systems*. Helsinki, Finland.

Viégas, Fernanda B., Martin Wattenberg, and Kushal Dave (2004) "Studying cooperation and conflict between authors with history flow visualizations." *CHI 2004.* Vienna, Austria.

Wikimedia (July 21 2004, 05:50 UTC) "Board of Trustees." from *Meta: a wiki about Wikimedia.*
(retrieved from: http://meta.wikimedia.org/wiki/Board_of_Trustees).

Wikimedia (June 5 2004, 14:14 UTC) "Power structure." from *Meta: a wiki about Wikimedia.*
(retrieved from: http://meta.wikimedia.org/wiki/Power_structure).

Wikipedia (July 9 2004, 01:07 UTC) "Accuracy Dispute." from *Wikipedia: The Free Encyclopedia.*
(retrieved from: http://en.wikipedia.org/wiki/Wikipedia:Accuracy_dispute).

Wikipedia (July 15 2004, 23:28 UTC) "Dispute resolution." from *Wikipedia: The Free Encyclopedia.*
(retrieved from: http://en.wikipedia.org/wiki/Wikipedia:Dispute_resolution).

Wikipedia (July 11 2004, 23:29 UTC) "History of Wikipedia." from *Wikipedia: The Free Encyclopedia.*
(retrieved from: http://en.wikipedia.org/wiki/History_of_Wikipedia).

Wikipedia (July 24 2004, 02:42 UTC) "Mailing lists." from *Wikipedia: The Free Encyclopedia.*
(retrieved from: http://en.wikipedia.org/wiki/Wikipedia:Mailing_lists).

Wikipedia (July 25 2004, 16:21 UTC) "Meetup." from *Wikipedia: The Free Encyclopedia.* (retrieved from: http://en.wikipedia.org/wiki/Wikipedia:Meetup).

Wikipedia (July 27 2004, 12:45 UTC) “Neutral point of view.” from *Wikipedia: The Free Encyclopedia.* (retrieved from: http://en.wikipedia.org/wiki/WP:NPOV).

Wikipedia (July 27 2004, 11:34 UTC) “NPOV Dispute.” from *Wikipedia: The Free Encyclopedia.*
(retrieved from: http://en.wikipedia.org/wiki/Wikipedia:NPOV_dispute).

Wikipedia (July 27 2004, 07:04 UTC) “Replies to common objections.” from *Wikipedia: The Free Encyclopedia.*
(retrieved from: http://en.wikipedia.org/wiki/Wikipedia:Replies).

Wikipedia (July 28 2004, 01:14 UTC) “Requests for Adminship.” from *Wikipedia: The Free Encyclopedia.* (retrieved from: http://en.wikipedia.org/wiki/WP:RFA).

Wikipedia (July 7 2004, 20:33 UTC) “Size Comparisons.” from *Wikipedia: The Free Encyclopedia.*
(retrieved from: http://en.wikipedia.org/wiki/Wikipedia:Size_Comparisons).

Wikipedia (July 27 2004, 23:12 UTC) “Votes for deletion.” from *Wikipedia: The Free Encyclopedia.*
(retrieved from: http://en.wikipedia.org/wiki/Wikipedia:Votes_for_deletion).

Wikipedia (July 7 2004, 12:01 UTC) “Why create an account?” from *Wikipedia: The Free Encyclopedia.*
(retrieved from: http://en.wikipedia.org/wiki/Wikipedia:Why_create_an_account%3F).

Zawinski, Jamie (March 31, 1999) “nomo zilla: resignation and post mortem.” http://www.jwz.org/gruntle/nomo.html (accessed: July 18, 2004).

Vita

Mason A. Hale is Chief Technologist at frog design, inc., an international design services firm, a position he has held since 1998. Prior to joining frog design, he founded and operated a technology consulting firm, which specialized in building electronic commerce-based applications. From 1994 to 1996, he held the position of Director of Internet Services at WebEdge, Inc., a technology consulting firm that produced a series of popular Internet-related conferences. In 1994, Hale served as a member of the City of Austin Internet Task Force, where he helped define policy for Internet usage among city employees and in the library system, and launched the Austin City Connection, the official web site of the City of Austin and the second official web site published by a municipality in the United States.

Mason Hale studied film and film production at the University of Texas at Austin where he was awarded a Bachelor of Science in Radio-Television-Film in 1992. He now lives in Austin, Texas with his wife, Colleen, and their three children.