

2003

Defining Open Source Software Project Success

Kevin Crowston

Syracuse University, School of Information Studies

Hala Annabi

Syracuse University, School of Information Studies

James Howison

Syracuse University, School of Information Studies

Follow this and additional works at: <http://surface.syr.edu/istpub>

 Part of the [Databases and Information Systems Commons](#)

Recommended Citation

Crowston, K., Annabi, H., & Howison, J. Defining Open Source Software project success. In *Proceedings of the International Conference on Information Systems (ICIS 2003)*, Seattle, WA, USA, December. doi: 10.1287/mnsc.1060.0550

This Article is brought to you for free and open access by the School of Information Studies (iSchool) at SURFACE. It has been accepted for inclusion in The School of Information Studies Faculty Scholarship by an authorized administrator of SURFACE. For more information, please contact surface@syr.edu.

DEFINING OPEN SOURCE SOFTWARE PROJECT SUCCESS

Kevin Crowston, Hala Annabi, and James Howison

School of Information Studies

Syracuse University

Syracuse, NY USA

crowston@syr.edu hpannabi@syr.edu

jhowison@syr.edu

Abstract

Information systems success is one of the most widely used dependent variables in information systems research. In this paper, we identify a range of measures that can be used to assess the success of open source software (OSS) projects. We identify measures based on a review of the literature, a consideration of the OSS development process and an analysis of the opinions of OSS developers. For each measure, we provide examples of how they might be used in a study of OSS development.

Keywords: Open source software, software development, system success

Introduction

Information systems (IS) success is one of the most widely used dependent variables in information systems research. Not surprisingly, much attention has been given to how best to measure it (e.g., DeLone and McLean 1992, 2002, 2003; Rai et al. 2002; Seddon 1997; Seddon et al. 1999). In this paper, we identify measures that can be applied to assess the success of open source software (OSS) projects based on a brief review of the literature, a consideration of the OSS development process and an analysis of the opinions of OSS developers. Since project success is likely to be a multidimensional construct, especially for OSS, our goal is to present a range of measures for future researchers to consider. As Seddon (1999) says, a “diversity of IS effectiveness measures is to be encouraged.”

The remainder of this paper is organized as follows. To ground our investigations, we first present the research context, namely OSS development. We then briefly review the literature on IS success to see what measures might be adopted and to identify problems in applying others. We then reconsider the underlying vision of system development to identify additional success measures that might be appropriate for OSS. For each measure, we provide examples of how they might be used in the context of a study of OSS. Finally, we compare our selection of measures to the opinions of OSS developers as expressed on SlashDot, a popular Web-based discussion board (<http://slashdot.org/>). The comparison suggests additional measures that might be incorporated to develop a fuller understanding of OSS project success. We conclude by making some suggestions for future research.

Research Context

OSS is a broad term used to embrace software that is developed and released under some sort of “open source” license. There are many licenses with a range of different features (Gacek et al., n.d.), but all allow inspection of the software’s source code. There are thousands of OSS projects, spanning a wide range of applications. Due to their size, success, and influence, the GNU/Linux operating system and the Apache Web Server are probably the most well known, but hundreds of other projects are

in widespread use, including projects on Internet infrastructure (e.g., sendmail, bind), user applications (e.g., the GIMP, OpenOffice), programming languages (e.g., Perl, Python, gcc), and games (e.g., Paradise). Much (though not all) open source software is also free software, in two senses: “free as in speech,” meaning that the code may be freely redistributed and reused in other OSS projects, and “free as in beer,” meaning that the software is available for download without charge. As well, many (though by no means all) OSS developers contribute to projects as volunteers without working for a common organization or being paid. As we will see, these two characteristics have implications for the applicability of certain measures of success.

It is important to develop measures of success for OSS projects for at least two reasons. First, having such measures should be useful for OSS project managers in assessing their projects. In some cases, OSS projects are sponsored by third parties, so measures are useful for sponsors to understand the return on their investment. Second, OSS is an increasingly visible and copied mode of systems development. Millions of users depend on OSS systems such as Linux (and the Internet, which is heavily dependent on OSS tools), but as Scacchi (2001) notes, “little is known about how people in these communities coordinate software development across different settings, or about what software processes, work practices, and organizational contexts are necessary to their success.” A recent EU/NSF workshop on priorities for OSS research identified the need both for learning “from open source modes of organization and production that could perhaps be applied to other areas” and for “a concerted effort on open source in itself, for itself” (Ghosh 2002). But to be able to learn from teams that are working well, we need to have a definition of “working well.”

Literature Review

OSS is a form of system development, so we begin our hunt for success measures in the Information Systems (IS) literature. Note, however, that we are not attempting an exhaustive review of this extensive literature, but rather are using the literature to identify success measures relevant to OSS. The most commonly cited model for IS success is the one developed by DeLone and McLean (1992, 2002, 2003), shown in Figure 1. This model suggests six interrelated measures of success: system quality, information quality, use, user satisfaction, individual impact, and organizational impact. Seddon (1997) proposed a related model that includes system quality, information quality, perceived usefulness, user satisfaction, and IS use. Taken together, these models suggest a number of possible measures that could be applied to OSS.




Figure 1. DeLone and McLean’s Model of IS Success

(Posted by permission, W. H. DeLone and E. R. McLean, “Information Systems Success: The Quest for the Dependent Variable,” *Information Systems Research* (3:1), 1992, p. 87. Copyright, The Institute of Management Sciences, now the Institute for Operations Research and the Management Sciences (INFORMS), 901 Elkridge Landing Road, Suite 400, Linthicum, MD 21090 USA.)

System and Information Quality

Code quality has been studied extensively in software engineering. This literature provides many possible measures of the quality of software including understandability, completeness, conciseness, portability, consistency, maintainability, testability, usability, reliability, structuredness, and efficiency (Boehm et al. 1976; Gorton and Liu 2002). To this list might be added the *quality of the system documentation*. Code quality measures would seem to be particularly applicable for studies of OSS, since the code is publicly available. Indeed, a few studies have already examined this dimension. For example, Stamelos et al. (2002) suggested

that OSS code is generally of good quality. Mishra et al. (2002) offer an analytic model that suggests factors contributing to OSS code quality, such as number of developers, mix of talent level, etc. The dimension of information quality seems to be less applicable since not many OSS systems include information (data) per se.

User Satisfaction

User satisfaction is an often-used measure of system success. For example, it is common to ask stakeholders if they felt a project was a success (e.g., Guinan et al. 1998). There is some data available regarding user satisfaction with OSS projects. For example, Freshmeat, a Web-based system that tracks releases of OSS (<http://freshmeat.net/>), collects *user ratings* of projects. Unfortunately, these ratings are based on a non-random sample (i.e., users who take the time to volunteer a rating), making their representativeness suspect. Furthermore, we have observed that the scores seem to have low variance: in a recent sample of 59 projects, we found that scores ranged only from 7.47 to 9.07. It seems likely that users who do not like a piece of software simply do not bother to enter ratings. There do not seem to be any easily obtainable data on the related measures of perceived ease of use and usefulness (Davis 1989). *Opinions expressed on project mailing lists* are a potential source of qualitative data on these facets, although again there would be questions about the representativeness of the data.

In principle, it should be possible to *survey users* to collect their satisfaction with or perceptions of the software. However, to do so properly poses a serious methodological problem. Because most OSS projects are freely distributed through multiple channels, the population of users is unknown, making it impossible to create a true random sample of users. In this respect, OSS differs greatly from information systems developed in an organizational setting that have a clearly defined user population. The situation is also different than for the Web, another non-traditional systems environment, because with a Web site, users are by definition the ones who visit the site, making the population effectively self-identifying. To achieve the same effect for OSS, although it may annoy some users, the best solution might be to build the survey into the software. For example, recent versions of the Mozilla Web browser include a program that offers to report crashes and collect other feedback.

Use

Although there is some debate about its appropriateness (DeLone and McLean 2003; Seddon 1997), many studies employ system use as an indication of information systems success. For software for which use is voluntary, as is the case for most OSS, use seems like a potentially relevant indicator of the project's success. Some interesting data are available. Avery Pennarun's *Debian Popularity Contest* (<http://people.debian.org/~apenwarr/popcon/>) collects statistics on the usage of software on Linux machines running the Debian distribution. Users install a program that collects and reports usage information daily and the resulting statistics show which packages have been installed and which of these have been recently used. Unfortunately, these data are collected from a nonrandom sample of machines, running a particular Linux distribution, so the results are likely not representative of use in the broader population.

Rather than measuring actual use, it may be sufficient to count the actual or potential *number of users* of the software, which we label *popularity* (Stewart and Ammeter 2002). For rare projects, these numbers can be directly measured. For example, Netcraft conducts a survey of Web server deployment (http://news.netcraft.com/archives/webserver_survey.html), which estimates the market share of the Apache Web server. Other projects that require some kind of network connection could potentially be measured in the same way (e.g., instant messaging or peer-to-peer file sharing clients), but this approach does not seem to be widely applicable.

A simple measure of popularity is the *number of downloads* made of a project. These numbers are readily available from various sites. Of course not all downloads result in use, so variance in the conversion ratio will make downloads an unreliable indicator of use. Furthermore, because OSS can be distributed through multiple outlets, on-line as well as offline (e.g., on CDs), the count from any single source is likely to be quite unreliable as a measure of total users. A particularly important channel is distributions such as RedHat, SuSE or Debian. Distributions provide purchasers with preselected bundles of software packaged for easy installation and are often sold on a CD-ROM to obviate the need to download everything. Indeed, the most popular software might be downloaded only rarely because it is already installed on most users' machines and stable enough to not require the download of regular updates. While the inclusion of software in a distribution might confound the usefulness of counts of downloads, the selection processes undertaken by distributions reflect expert judgment on the worth of the packages. Therefore *inclusion in distributions* is an independently important success measure.

Other sources of data reflecting on users are available. Freshmeat provides a *popularity* measure for packages it tracks, though a better name might be *interest*, as it is one step further removed from actual use. The measure is calculated as the geometric mean of subscriptions and two counts of page viewings of project information (for the precise details, see <http://freshmeat.net/faq/view/30/>). Similarly, SourceForge provides information on the number of *page views of the information pages* for projects it supports.

Finally, it may be informative to measure use from perspectives other than that of an end user. In particular, the openness of OSS means that other projects can build on top of it. Therefore, one measure of a project's success may be that many other projects use it. *Package dependency* information between projects can be obtained from the package descriptions available through the various distributions' package management systems. Analysis of source code could reveal the *reuse of code* from project to project (although identifying the true origin could be difficult, as SCO's claim of ownership of parts of the Linux kernel demonstrates).

Individual or Organizational Impacts

The final measures in DeLone and McLean's (1992) model are individual and organizational impacts for the users. Although there is considerable interest in the economic implications of OSS, these measures are hard to define for regular I/S projects and doubly hard for OSS projects, because of the problems defining the intended user base and expected outcomes. Therefore, these measures are likely to be unusable for most studies of individual OSS projects.

Summary

To summarize, existing models of information systems success suggest a range of potential success measures for OSS projects as shown in Table 1. However, a number of the measures are inapplicable, while others are difficult to apply in the OSS environment. We note that many of these measures are based on a vision of system development in an organization and do not take into account the unique characteristics of the OSS development environment. Therefore, in the next section, we examine the process of OSS development in more detail to identify additional potential measures of success.

Table 1. Success Measures Suggested by the Literature Review

Measure of Success	Indicators	Audience
<i>System and information quality</i>	<ul style="list-style-type: none"> • Code quality (e.g., understandability, completeness, conciseness, portability, consistency, maintainability, testability, usability, reliability, structuredness, efficiency) • Documentation quality 	Users, developers
<i>User satisfaction</i>	<ul style="list-style-type: none"> • User ratings • Opinions on mailing lists • User surveys 	Users, developers
<i>Use</i>	<ul style="list-style-type: none"> • Use (e.g., Debian Popularity Contest) • Number of users • Downloads • Inclusion in distributions • Popularity or views of information page • Package dependencies • Reuse of code 	Developers
<i>Individual and organizational impacts</i>	<ul style="list-style-type: none"> • Economic and other implications 	Users, developers

The Process of OSS Development

In this section, we reexamine the vision of systems development underlying DeLone and McLean's success model to identify additional measures that might be used for OSS project success. DeLone and McLean state that their model was built by considering "a process model [that] has just three components: the creation of a system, the use of the system, and the consequences of this system use" (2002), which we have shown graphically in Figure 2. We note that the measures included in the model focus on the use and consequences of the system (the right side of the figure), and do not open up either box in the process. While this focus may be appropriate given the traditional concern of information systems research with the organizational implication of IS, it seems to unduly restrict the range of measures considered.




Figure 2. Process Model Underlying the DeLone and McLean (1992) Model of Success

The choice of measures also seems to be influenced by the relative ease of access to the use environment compared to the development environment (especially true for packaged or commercial software). In the context of OSS, however, researchers are frequently faced with the opposite situation, in that the development process is publicly visible and the use environment is difficult to study or even identify. For both reasons, we believe that it will be useful to complement existing success measures with ones that take advantage of the availability of data on the development process. The following discussion examines such measures of success, some of which have been previously discussed in the information systems literature.

Measures of the Output of Systems Development

Two of the measures in the DeLone and McLean model concern the product of the systems development process, namely systems quality and information quality. We first consider possible additional measures of this process step.

First, given the large number of abandoned projects (Ewusi-Mensah 1997), simply *completing a project* may be a sign of success. However, many OSS projects are continually in development, making it difficult to say when they are completed. Faced with this problem, Crowston and Scozzi (2002) instead measured success as the *progress of a project* from alpha to beta to stable status, as self-reported on SourceForge.

Second, another commonly used measure of success is whether the *project achieved its goals*. This assessment is typically made by a comparison of the project outcomes with the formal requirements specifications. However, OSS projects often do not have such specifications. Scacchi (2002) examined the process of requirements engineering in open source projects and provided a comparison with traditional processes (e.g., Davis 1990; Jackson 1995). He argues that rather than a formal process, OSS requirements are developed through what he terms *software informalisms*, which do not result in agreed requirements documentation that could later be analyzed to see whether the project has met its goals. Scacchi's ethnography suggests that for OSS, goals will likely come from within through a discursive process centered on the developers. Therefore, a key measure for OSS may be simply *developer satisfaction* with the project, which could be measured by surveying developers. The developer community is much more clearly delineated than users, making such a survey feasible. Indeed, there have already been several OSS developer surveys (e.g., Ghosh 2002; Hertel et al. n.d.), although not on this topic specifically. Since in many projects there is a great disparity in the contribution of developers—a few developers contribute the bulk of the code (Mockus et al. 2000)—it may be desirable to weight developers' opinions in forming an overall assessment of a project.

Measures of the Process of Systems Development

In DeLone and McLean's (1992) process model, systems development is implicitly treated as a one-off event. However, for OSS projects (and indeed many other types of projects) development is instead an ongoing activity, as the project continues to release

“often and early” (Raymond 1998). In other words, an OSS project is characterized by a continuing process of developers fixing bugs, adding features, and releasing software. This characteristic of the OSS development process suggests a number of possible indicators of success.

Number of Developers

First, since many OSS projects are dependent on volunteer developers, being able to attract developers to a project on an on-going basis is important for their success. Thus the *number of developers* involved in a project could be an indicator of success. The number of developers can be measured in at least two ways. OSS development systems such as SourceForge list developers who are formally associated with each project. Examination of the mailing lists and other fora associated with projects can reveal the number of individuals who actively participate.

Level of Activity

More important than the sheer number of developers is their contribution to a project. Thus the *level of activity* of developers in submitting code and bug reports may be useful as an indicator of project success. For example, SourceForge computes and reports a measure of project activity based on the activities of developers. Researchers could also examine development logs for evidence of software being written and released.

Cycle Time

Another measure related to the group activity is *time between releases*. In OSS development, there is a strong community norm to “release early and release often,” which implies that an active release cycle is a sign of a healthy development process and project. For example, FreshMeat provides a “vitality score” (Stewart and Ammeter 2002) that assesses how recently a project has made an announcement of progress on the FreshMeat site (<http://freshmeat.net/faq/view/27/>).

In addition, detailed examination of bug-fixing and feature-request fulfillment activities might yield useful process data indicative of the project’s status. These processes involve interaction with the user community and can involve applying patches of contributed code supplied by non-core developers. Bug reports and feature requests are typically managed through a task-management system that records the developer and community discussion, permits labeling of priority items, and sometimes includes informal voting mechanisms to allow the community to express its level of interest in a bug or new feature.

We are currently exploring an analysis of the *time to close bugs* (or implement requested features) as a measure of project success. To collect data on bugs, we spidered the bug report data from SourceForge and extracted the length of time taken to fix bugs as well as other characteristics of each bug (priority, issue area, and explicit developer assignment). Our preliminary analysis of the bug and project level data suggests that this measure shows interesting variance between projects and so may be a useful measure of this dimension of success.

Project Effects on Projects

Finally, because the projects are on-going, it seems important to consider the impact of a project on the abilities of the project team itself and its ability to continue or improve the development process. As Shenhar et al. put it, “how does the current project help prepare the organization for future challenges?” (2001).

Employment Opportunities

Some literature on the motivation of OSS developers suggests that developers participate to improve their employment opportunities (e.g., Lerner and Tirole 2000). Thus, one can consider *salary* (Hann et al. 2002) or *jobs acquired* through the involvement in a particular project as possible measures of success. For example, Hann et al. (2002) found that higher status within the Apache Project was associated with significantly higher wages. Again, one might measure these indicators by surveying developers. While for a single developer, these measure are confounded with innate talent, training, luck, etc., aggregating across many developers and across time may provide a useful project level measure of success.

Individual Reputation

Similarly, literature also suggests that developers participating in OSS projects are rewarded with *reputation in the community*, and that this reputation is a sufficient reward for interaction. Kelty (2001) suggests that reputation might be measured through an analysis of credits located in source code (which he terms “greputation”). Alternative measures of OSS reputation might include the OSS communities’ implementation of a “Web of Trust” at the community site Advogato (<http://www.advogato.org/trust-metric.html>) where developer status is conferred through peer-review. Analyses of this kind of measure face the difficulty of tying the earning of reputation to the success of a particular project.

Knowledge Creation

Projects can also lead to *creation of new knowledge* for individuals as well as on the group level (Arent and Nørbjerg 2000). Through their participation in a project, individual developers may acquire new procedural and programming skills that would benefit them on future projects. This effect could be measured by surveying the developers for their perceived learning.

In addition, following Grant’s (1996) knowledge-based view of the firm, we view a firm (or in this case, a project) as a structure to integrate members’ knowledge into products. In this view, the project’s rules, procedures, norms and existing products are a reflection of knowledge being created by the project activities. This knowledge creation can be measured by observing and qualitatively analyzing changes in the written rules and procedures over time and may be reflected and transferred through the development of systems for OSS project support, such as SourceForge and Savannah. Analysis of the development of interactions and support systems closely linked to a project might give some insight into this aspect of project success.

Summary

In summary, consideration of the process of developing OSS suggests a number of additional measures indicative of success for these projects. These measures are summarized in Table 2. We note that as the measures move further back in the process model, they become increasingly removed from the user. As such, there may be a concern about their validity as measures of success: Is it a success if a project attracts developers but not users? Or if it develops high quality processes but not high quality code? We have two replies to this concern. First, the apparent disconnect may be an accurate representation of the reality of OSS projects, in which the developers frequently are the users. Second, the measures developed in this section should be viewed as complements to rather than replacements for the more conventional measures of success. Using a variety of measures will provide a richer picture of the status of a project. As well, because many of the measures seem likely to have measurement problems, adopting a portfolio of measures seems prudent.

Table 2. Measures Suggested by a Reexamination of the OSS Process

Measure of Success	Indicators	Audience
<i>Project output</i>	<ul style="list-style-type: none"> • Movement from alpha to beta to stable • Achieved identified goals • Developer satisfaction 	Developers
<i>Process</i>	<ul style="list-style-type: none"> • Number of developers • Level of activity (developer and user contributions, number of releases) • Time between releases • Time to close bugs or implement features 	Developers, users
<i>Outcomes for project members</i>	<ul style="list-style-type: none"> • Individual job opportunities and salary • Individual reputation • Knowledge creation 	Developers

Open Source Developer Opinions

In the previous two sections, we developed a list of possible success measures for OSS projects based on a review of the literature and consideration of a simple model of OSS development. To determine whether these measures had content validity as indicators of OSS project success and to identify additional possible measures, we sought input from OSS developers. The goal of the study was to assess the completeness and validity of our list of factors, rather than the relative importance of each one or the relations among them. Therefore, we chose a method to generate a range of ideas rather than one to support statistical inference. In this section, we discuss our data elicitation and analysis techniques and results from the analysis.

Methods and Data

We submitted a question to the moderators of SlashDot (<http://slashdot.org/>), a Web-based discussion forum. SlashDot attracts interest and participation from OSS developers and users. The question was judged to be of interest to the audience and posted on the front page of the site. This data elicitation technique was more like an on-line focus group (or perhaps the initial stage of a Delphi study) than a survey, as respondents were a nonrandom sample and could see and respond to earlier postings. This approach was chosen to match our goal of generating ideas about success measures, rather than testing a theory or making inferences from generalizable data. To elicit comments, the following question was posted on SlashDot on April 22, 2003 (<http://slashdot.org/article.pl?sid=03/04/21/239212>):

There have been a number of discussions on Slashdot and elsewhere about how good projects work (e.g., Talk to a Successful Free Software Project Leader), but less about how to tell if things are going well in the first place. While this may seem obvious, most traditional definitions of software project success seem inapplicable (e.g., profit) or nearly impossible to measure for most projects (e.g., market share, user satisfaction, organizational impact). In an organizational setting, developers can get feedback from their customers, the marketplace, managers, etc.; if you're Apache, you can look at Netcraft's survey of server usage; but what can the rest do? Is it enough that you're happy with the code? I suspect that the release-early-and-often philosophy plays an important role here. I'm asking not to pick winners and losers (i.e., NOT a ranking of projects), but to understand what developers look at to know when things are going well and when they're not.

The question received 201 responses within a few days.

Participants

Many of the individuals posting answers to our question identified themselves as developers or contributors to OSS projects. As a check on their qualifications, we searched SourceForge (a popular OSS development site) for information about the posters. Although SourceForge and SlashDot are separate sites, many developers have strong attachments to their user IDs and use the same one whenever possible, providing a possible link between the two systems. For example, it seems reasonable to expect that the user ID *Abcd1234* identifies the same individual on both systems. We identified the SlashDot IDs of 72 posters who provided useful responses (some responses were anonymous). Of these 72, 34 IDs matched a SourceForge ID exactly, and 6 could be matched with a bit of research (e.g., by matching the real name of the individual; real names are available for a few SlashDot posters and many SourceForge developers). Of the matched IDs, 16 and 3 respectively were listed as members of SourceForge projects (i.e., about half). A few other posters had pointers to non-SourceForge OSS projects on their SlashDot information page or information about their employment, generally as software developers. These data are not conclusive, but do suggest that a number of the contributors to the study had sufficient background as OSS developers to be able to comment knowledgeably.

Analysis

A transcript of responses was downloaded on April 26, 2003, and content analyzed by two coders. The content analysis process was carried out using Atlas.ti, a qualitative data analysis software package. Messages were coded using the thematic unit as the unit of analysis. Once a measure was identified within a message, the coder selected the text containing the measure and coded that text using categories from our coding scheme. A total of 170 thematic units were identified and coded in 91 responses (i.e., some postings contained multiple units; the remaining responses did not contain text addressing the question, e.g., a posting containing an advertisement). The content analysis process employed a mixture of deductive and inductive procedures. The initial content analytic scheme was based on the literature review described above. During the process of content analysis, additional

themes emerged from the data. Saturation was reached through the content scheme presented in Appendix A. The two raters agreed on the codes for 78 percent of the units. We felt that this level of agreement was sufficient for the purposes of the analysis (identification of measures to compare to the literature review), so we did not go on to refine the definitions of codes or retrain the coders to increase agreement.

Results

The outcome of the content analysis is summarized in Table 3. The codes were organized into a two-level hierarchy for presentation, with detailed codes (level 2 in the table) clustered into meta-categories (level 1). In all, 32 percent of the units included elements from the developers meta-category, indicating that the respondents felt that a project is successful if the developers are involved, satisfied, enjoyed the process, and that there is a variety of them. The users meta-category also had a large number of responses: 23 percent of units indicated that the poster felt a project was successful if it satisfies users (other than developers) and that the users are involved in discussions and bug reports. Involvement of both users and developers was frequently mentioned, accounting for 31 percent of the units. Project recognition codes were found in 11 percent of the units, exceeding the number of responses indicating use as a measure of success, which accounted for 5 percent of instances. Finally, the product's quality (13 percent) and process (13 percent) were suggested to be measures of success by respondents as well. Note that percentages are reported only for completeness. Given the nonrandom sample of contributors and the open data elicitation technique, the frequency of a response should not be interpreted as importance. The relative importance of these factors remains to be tested in future research.

Table 3. Results of the Content Analysis of SlashDot Responses

Level 1	Level 2	Frequency	Percentage
User	Satisfaction	14	8%
	Involvement	25	15%
Product	Meets requirements	9	5%
	Code quality	11	6%
	Portability	1	1%
	Availability	2	1%
Process	Activity	5	3%
	Adherence to process	10	6%
	Bug Fixing	4	2%
	Time	2	1%
	Age	1	1%
Developers	Involvement	16	9%
	Varied developers	2	1%
	Satisfaction	29	17%
	Enjoyment	8	5%
Use	Competition	4	2%
	Number of users	2	1%
	Downloads	3	2%
Recognition	Referral	3	2%
	Attention and recognition	9	5%
	Spin offs	6	4%
Influence		4	2%
	Total	170	

Discussion

Overall, the responses of the developers posting on SlashDot were in general agreement with the list of success measures we developed from the literature and our reexamination of the process. The analysis indicates that developers found their personal involvement, satisfaction, and enjoyment to be measures the success of a project, consistent with the view of OSS as “software that scratches an itch.” More interestingly, some new themes did emerge from the coding.

- First, a number of respondents suggested *recognition* (e.g., mention on other sites) as a measure of project success. Similarly, another suggested measure was the *influence* of the product or project’s process on other OSS groups and other commercial settings. These responses are consistent with the literature on OSS developers’ motivations that suggest recognition as a primary motivation for involvement.
- A second category that emerged was the level of *involvement of the users* as indicated by involvement of the users in submitting bug reports and participating in the project mailing lists. We had considered contributions from developers, but these responses reflect the fact that OSS projects are also dependent on help from users to identify problems and post suggestions.
- A final category that emerged from the data was the issue of *porting*. Developers consider porting of a product to different systems (especially to Windows) and requests for such ports as a measure of the success of the product. This theme might be considered a special case of popularity.

What was also surprising was what respondents did not say, in that respondents did not mention a few of the measures of success we had identified. For example, although several authors have suggested that developers are motivated by the chance to learn and perhaps get a better job, none of the respondents mentioned these factors. A possible explanation is the strong community norm that endorses altruism over expressions of self-interest, which may have restricted discussion in the non-anonymous and community-moderated SlashDot forum.

Conclusion

This paper makes a contribution to the developing body of empirical research on OSS by identifying a collection of success measures that might be applied to OSS. We have identified a range of possible measures by applying a popular model of IS success and by more detailed consideration of the vision of software development underlying that model. Furthermore, we have identified where data is available for particular measures. Finally, we checked our list of factors by comparing it to opinions of community members, which suggested additional measures beyond those we identified and raised questions about others.

We emphasize again that we do not view any single measure as the final word on success. As the measures focus on different aspects of the process, we expect that they will offer different perspectives on the process. Therefore, we suggest using a portfolio of measures, or perhaps developing synthetic measures that draw on different perspectives. Using multiple measures might be particularly interesting for examining how projects change their emphasis from one measure to another at different points in their evolution (Heo and Han 2003).

Acknowledgements

The authors gratefully acknowledge Chengetai Masango for his contribution to the content analysis process.

References

- Arent, J., and Nørberg, J. “Software Process Improvement as Organizational Knowledge Creation: A Multiple Case Analysis,” in *Proceedings of the 33rd Hawaii International Conference on System Sciences*, IEEE Computer Society Press, Los Alamitos, CA, 2000.

- Boehm, B. W., Brown, J. R., and Lipow, M. "Quantitative Evaluation of Software Quality," in *Proceedings of the 2nd International Conference on Software Engineering*, IEEE Computer Society, October 13-15, 1976, pp. 592-605.
- Crowston, K., and Scozzi, B. "Open Source Software Projects as Virtual Organizations: Competency Rallying for Software Development." *IEE Proceedings Software* (149:1), 2002, pp. 3-17.
- Davis, A. M. *Software Requirements Analysis and Specification*, PrenticeHall, Englewood Cliffs, NJ, 1990.
- Davis, F. D. "Perceived Usefulness, Perceived Ease of Use and User Acceptance of Information Technology," *MIS Quarterly* (13:3), 1989, pp. 319-340.
- DeLone, W. H., and McLean, E. R. "The DeLone and McLean Model of Information Systems Success: A Ten-Year Update," *Journal of Management Information Systems* (19:4), 2003, pp. 9-30.
- DeLone, W. H., and McLean, E. R. "Information Systems Success: The Quest for the Dependent Variable," *Information Systems Research* (3:1), 1992, pp. 60-95.
- DeLone, W. H., and McLean, E. R. "Information Systems Success Revisited," in *Proceedings of the 35th Hawaii International Conference on System Sciences*, IEEE Computer Society Press, Los Alamitos, CA, 2002.
- Ewusi-Mensah, K. "Critical Issues in Abandoned Information Systems Development Projects," *Communication of the ACM* (40:9), 1997, pp. 74-80.
- Gacek, C., Lawrie, T., and Arief, B. "The Many Leanings of Open Source," unpublished manuscript, Centre for Software Reliability, Department of Computing Science, University of Newcastle, Newcastle Upon Tyne, United Kingdom, no date.
- Ghosh, R. A. "Free/Libre and Open Source Software: Survey and Study," Report of the FLOSS Workshop on Advancing the Research Agenda on Free/Open Source Software, European Commission, Brussels, October 12, 2002 (available online at <http://www.infonomics.nl/FLOSS/report/workshopreport.htm>).
- Gorton, I., and Liu, A. "Software Component Quality Assessment in Practice: Successes and Practical Impediments," in *Proceedings of the 24th International Conference on Software Engineering*, IEEE Computer Society, 2002, pp. 555-558.
- Grant, R. M. "Toward a Knowledge-Based Theory of the Firm," *Strategic Management Journal* (17:Winter), 1996, pp. 109-122.
- Guinan, P. J., Cooprier, J. G., and Faraj, S. "Enabling Software Development Team Performance During Requirements Definition: A Behavioral Versus Technical Approach," *Information Systems Research* (9:2), 1998, pp. 101-125.
- Hann, I-H., Roberts, J., Slaughter, S., and Fielding, R. "Economic incentives for participating in Open Source Software Projects," in *Proceedings of the Twenty-Third International Conference on Information Systems*, L. Applegate, R. Galliers, and J. I. DeGross (eds.), Barcelona, December 2002, pp. 365-372.
- Heo, J., and Han, I. G. "Performance Measure of Information Systems (IS) in Evolving Computing Environments: An Empirical Investigation," *Information Management* (40:4), 2003, pp. 243-256.
- Hertel, G., Niedner, S., and Herrmann, S. "Motivation of Software Developers in Open Source Projects: An Internet-Based Survey of Contributors to the Linux Kernel," unpublished manuscript, University of Kiel, Kiel, Germany, no date.
- Jackson, M. *Software Requirements and Specifications: Practice, Principles, and Prejudices*, Addison-Wesley, Boston, 1995.
- Kelty, C. "Free Software/Free Science," *First Monday* (6:12), 2001.
- Lerner, J., and Tirole, J. *The Simple Economics of Open Source*, NBER Working Paper 7600, The National Bureau of Economic Research, Inc., Cambridge, MA, 2000.
- Mishra, B., Prasad, A., and Raghunathan, S. "Quality and Profits Under Open Source Versus Closed Source," in *Proceedings of the Twenty-Third International Conference on Information Systems*, L. Applegate, R. Galliers, and J. I. DeGross (eds.), Barcelona, December 2002, pp. 349-363.
- Mockus, A., Fielding, R. T., and Herbsleb, J. D. "A Case Study of Open Source Software Development: The Apache Server," in *Proceedings of International Conference on Software Engineering*, IEEE Computer Society, 2000.
- Rai, A., Lang, S. S., and Welker, R. B. "Assessing the Validity of IS Success Models: An Empirical Test and Theoretical Analysis," *Information Systems Research* (13:1), 2002, pp. 50-69.
- Raymond, E. S. "The Cathedral and the Bazaar," *First Monday* (3:3), 1998.
- Scacchi, W. "Software Development Practices in Open Software Development Communities: A Comparative Case Study (Position Paper)," paper presented at Making Sense of the Bazaar: First Workshop on Open Source Software Engineering, The 23rd International Conference on Software Engineering, Toronto, Ontario, 2001 (available online at <http://opensource.ucc.ie/icse2001/scacchi.pdf>, accessed September 2, 2003).
- Scacchi, W. "Understanding the Requirements for Developing Open Source Software Systems," *IEE Proceedings Software* (149:4), 2002, pp. 24-39.
- Seddon, P. B. "A Respecification and Extension of the DeLone and McLean Model of IS Success," *Information Systems Research* (8:3), 1997, pp. 240-253.
- Seddon, P. B., Staples, S., Patnayakuni, R., and Bowtell, M. "Dimensions of Information Systems Success," *Communications of the Association for Information Systems* (2:20), 1999.
- Shenhar, A. J., Dvir, D., Levy, O., and Maltz, A. C. "Project Success: A Multidimensional Strategic Concept," *Long Range Planning* (34), 2001, pp. 699-725.

Stamelos, I., Angelis, L., Oikonomou, A., and Bleris, G. L. “Code Quality Analysis in Open Source Software Development,” *Information Systems Journal* (12:1), 2002, pp. 43-60.

Stewart, K. J., and Ammeter, T. “An Exploratory Study of Factors Influencing the Level of Vitality and Popularity of Open Source Projects,” in *Proceedings of the Twenty-Third International Conference on Information Systems*, L. Applegate, R. Galliers, and J. I. DeGross (eds.), Barcelona, 2002, pp. 853e57.

Appendix A

Content Analytic Scheme

Level 1	Level 2	Description	Examples Responses
User	Satisfaction	Users of the product satisfied (code serves their need)	Every piece of software has an intended client, user or audience. Are the users happy, overall? Useful...to someone.
	Involvement	Users of the product are involved and interested, by submitting to mailing lists or bug reports or other forms of contribution	Second: I’ve had dozens of email’s asking for support as well as asking how to contribute. Traffic: both developer and user. Is there a relatively continuous level of input/interest in the project? If developers don’t want to develop, and users don’t want to use, it’s probably going nowhere, even if it’s the best thing since the BeOS. More activity in a mailing list usually indicates the size/success of a project. Small error in the makefile which causes something liek [sic] 50% of people [to] come back for help on compiling it. This gives me [a] pretty good estimate of how many people are actually using the package.
Product	Meets requirements	The product meets the requirement of design	How well does it fit the need for which it was designed?
	Code quality	Code structure and documentation is organized, clear, maintainable	Does what it is supposed to do, cleanly and efficiently, then by definition it is successful. It is well documented and maintained. Is the code maintainable? Take a look at the end product. Does it do what it’s supposed to without too many bugs? Stable in relation to the time invested.
	Portability	Software portable to and compatible with other systems and programs	Successful Open Source Software tends to have a greater scope of use than it’s [sic] original conception. The programs I find myself using are programs that can interact with each other in a modular fashion; whether that be throught [sic] a piped command, or simply support for “generic” file formats (such as XML, CSV, etc. etc). Win32 port: Win32 port of a project.
	Availability	The product is available through a number of avenues	Availability around the Internet (if your program is on say on most of the distributions).

Level 1	Level 2	Description	Examples Responses
Process	Activity	The project is active; fixing bugs, writing updates, documentation and releases	If there hasn't been an update to it in like 2 years, then chances are, unless it was perfect the first time around, it will fail.
	Adherence to process	The project has goals and objectives and have an established process that members adhere to	How is the process? Are there goals and are they being met? How is testing coverage and how often is testing being done? Milestones: establish concrete goals when you start the project, along with a timeline. Many goals, and projects evolve
	Bug fixing	Bug reports are attended to and fixed in reasonable time	Are issues being addressed in a timely manner?
	Time	How established is the software and how often do they release new features	Time is where you can measure your progress. This is where you can do things like determine milestones, develop feature lists and so on, then during the project you have a standard to compare yourself to.
	Age	How long has the group been active	
Developers	Contribution	There are a number of developers contributing to the project	Second: I've had dozens of email's asking for support as well as asking how to contribute . Traffic: both developer and user. Is there a relatively continuous level of input/interest in the project? If developers don't want to develop, and users don't want to use, it's probably going nowhere, even if it's the best thing since the BeOS.
	Varied developers	Developers from different projects, having different expertise contribute	Software developers
	Satisfaction	Developers satisfy their need to innovate and develop code	Open source is scratching an itch, right? Is the itch scratched? If yes, then it's [sic] a success.
	Enjoyment	Developers enjoy working on the code and with the group	Do you enjoy working on it? Then it's successful.
Use	In relation to competition	Replaced competitive products	So clearly and overwhelmingly superior to its predecessors that it supplanted them.
	Number of users	How many users are using the product in addition to the developers	Are there more people using the project than developers? If so, it's successful.. If a project has no user base, then it is doomed to fail. How is a project going to succeed without a user base.
	Downloads	How many downloads of the product	Can usually judge the success of your project by counting downloads from your site. First: I've had hundreds of downloads, and since I run this project on a Cable Modem connection, my ISP hasn't become unhappy :)

Level 1	Level 2	Description	Examples Responses
Recognition	Referral	Other sites, projects, organizations recognize and refer to the project	Links to site: and Third and finally (I think this one is a very good indicator): There are other websites out there that link to my site. Oh, and there's a fourth optional measure of success...more for bragging rights...my site is THE FIRST result when querying google with "Java X10." Use Google to see how often the name of the project comes up. Discussion in Google groups is also a good sign.
	Attention	The project attracted negative or positive attention from other institutions	You have been sued by a huge mega corp with a team of lawyers over patent infringement and the EFF comes to your rescue. Stallman demands that people call it GNU
	Spin offs	New projects or spins off original project	
Influence		Other projects adopt code or process from the project	Also adoption by other developers into the development group shows others are interested, so you must be doing something right.