

University of East London Institutional Repository: <http://roar.uel.ac.uk>

This paper is made available online in accordance with publisher policies. Please scroll down to view the document itself. Please refer to the repository record for this item and our policy information available from the repository home page for further information.

To see the final version of this paper please visit the publisher's website. Access to the published version may require a subscription.

Author(s): Capiluppi, Andrea., Boldyreff, Cornelia., Beecher, Karl., Adams, Paul.

Article Title: Quality factors and coding standards - a comparison between open source forges

Year of publication: 2008

Citation: Capiluppi, A., Boldyreff, C., Beecher, K., Adams, P. (2008) 'Quality factors and coding standards - a comparison between open source forges.' Electronic Notes in Theoretical Computer Science 233 (27) 89-103 Proceedings of the International Workshop on Software Quality and Maintainability (SQM 2008).

Link to published version:

<http://dx.doi.org/10.1016/j.entcs.2009.02.063>

DOI: <http://dx.doi.org/10.1016/j.entcs.2009.02.063>

Publisher statement:

<http://www.elsevier.com/wps/find/authorsview.authors/authorsrights>

Replace this file with `prentcsmacro.sty` for your meeting,
or with `entcsmacro.sty` for your meeting. Both can be
found at the [ENTCS Macro Home Page](#).

Quality Factors and Coding Standards – a Comparison Between Open Source Forges

Andrea Capiluppi Cornelia Boldyreff Karl Beecher

*Centre of Research on Open Source Software – CROSS
Department of Computer Science
University of Lincoln, UK*

`{acapiluppi,cboldyreff,kbeecher}@lincoln.ac.uk`

Paul J. Adams

*Research and Development,
Sirius Corporation Ltd.,
Weybridge, UK*

`paul.adams@siriusit.co.uk`

Abstract

Enforcing adherence to standards in software development in order to produce high quality software artefacts has long been recognised as best practice in traditional software engineering. In a distributed heterogeneous development environment such those found within the Open Source paradigm, coding standards are informally shared and adhered to by communities of loosely coupled developers. Following these standards could potentially lead to higher quality software.

This paper reports on the empirical analysis of two major forges where OSS projects are hosted. The first one, the KDE forge, provides a set of guidelines and coding standards in the form of a coding style that developers may conform to when producing the code source artefacts. The second studied forge, SourceForge, imposes no formal coding standards on developers. A sample of projects from these two forges has been analysed to detect whether the SourceForge sample, where no coding standards are reinforced, has a lower quality than the sample from KDE.

Results from this analysis form a complex picture; visually, all the selected metrics show a clear divide between the two forges, but from the statistical standpoint, clear distinctions cannot be drawn amongst these quality related measures in the two forge samples.

Keywords: Coding Standards, Open Source Software, Complexity.

1 Introduction

Programming styles and coding standards form a set of formal rules which are internally shared among programmers, and enforced by software managers [?]. These rules reflect different concerns and affect areas of source code writing, with the aim of improving both the readability of source code and the maintainability of the underlying software system. They range from language-independent *typographic styles*, as the rules affecting how both the source code and comments are visually

structured and displayed [?,?], to general *programming practices* relative to specific programming languages (as C++ [?]), paradigms (such as the object-oriented paradigm [?]) or development approaches (such as the Agile methodologies [?]).

Research evidence has confirmed the presence of coding standards and programming styles in a wide spectrum of approaches: in commercial, proprietary, software [?], in Agile-driven systems [?] as well as in Open Source software [?]. Past empirical research also suggested that programming styles, applied in both creating new code or maintaining existing programs, have an impact on the quality of the resulting software [?].

In this paper it is argued that, in the presence of coding standards and programming guidelines, the resulting software will display a higher quality than software produced without such a coding framework. Specifically, a study of whether software produced within the Open Source paradigm has higher quality as long as they implement and enforce coding practices is provided. In order to detect the quality of the resulting software, three metrics were collected, both at the granular level of functions (or OO methods), the fan-in, the fan-out (both termed as “coupling”) and the relative cyclomatic complexity. A null hypothesis was formulated: software projects from the sample implementing coding styles and standards will display lower values of complexity and coupling than counterparts from a sample not enforcing the same standards. The first sample (“with treatment”) was randomly extracted from the KDE forge, already known for enforcing these standards and guidelines [?]. The other sample (“without treatment”) was extracted from the SourceForge repository¹. A statistical test [?] was used to evaluate the significance of the differences in the two samples.

This paper is structured as follows: **Section 2** provides a description of past and current research works which can be related to the present paper, and describes how it completes or enhances previous approaches. **Section 3** presents the definitions and the attributes used in this work. It also introduces and instantiates the GQM approach [?], tailoring a research question for the current study. **Section 4** produces an empirical hypothesis based on the presented research question, illustrating the *null* and the *alternative* hypotheses, as well as the complexity and coupling metrics used to evaluate them. In the same section, an overview of the samples from the two forges is also given. **Section 5** presents the results of the statistical tests linked to the hypothesis, and evaluates whether the null hypothesis should be rejected or not. **Section 6** illustrates both internal and external threats to validity to this empirical study, while **Sections 7 and 8** conclude this paper.

2 Related work and Background

The research areas related to this work can be divided into three main sections: metrics for software quality, literature on coding standards and styles and related practice including tool support. Regarding the first research area, there is an extensive software literature dealing with the measurement of software quality and associated characteristics such as complexity, comprehensibility, reusability and maintainabil-

¹ Even if not clearly stated, it may very well be the case that particular projects within SourceForge adhere to coding standards anyway. In the following, it is assumed that this effect is negligible.

ity. Measurements have been traditionally divided into development and design quality metrics (early in the product lifecycle), e.g. [?], and post-release metrics which can be applied to a finished software product. One of the most famous effort in the later area presented the results of the impact of coupling, cohesion and complexity on the development cost of object-oriented systems [?].

Quality measurement has also matured to the point at which a standard has been defined for this activity. The ISO standard for software quality measurement [?] defines the characteristics of software quality as: functionality, reliability, usability, efficiency, maintainability and portability. In the past, there have been other attempts to determine quality numbers based on source code metrics: the most well-known is the Maintainability Index [?], which is a composite metric to assess, at the system level, the relative maintainability. The Halstead Effort [?], the McCabe cyclomatic complexity, the lines of code and comments are averaged into a single number to represent a global index of maintainability and quality. As reported in other studies [?], the McCabe values typically follow a Power-law distribution for methods or procedures in a system: using the average for such distributions will hinder some of the characteristics of the distribution, such as its skewness. The Halstead Effort metric has been evaluated as highly correlated with the McCabe index [?], but also heavily criticized as an unreliable metric [?]. In this work, the McCabe indexes are evaluated both as distribution in each project, and considering the fraction of highly complex functions (i.e., whose index is larger than 10 [?]). The Halstead Effort as a metric is not considered in this work.

Works focusing on open source quality or success often measure endogenous metrics (often “people” or “process” metrics) such as the amount of developer activity, the number of developers, forum activity, version control etc. [?,?,?,?]. A related work examining product metrics by Stamelos et al [?] undertook an analysis of 100 open source programs written for Linux and took a variety of basic static measurements (such as program length, unconditional jumps etc.). These results were compared to “industrial standards” suggested by the measurement tool.

The use of coding standards is also well covered in the literature and can be partially enforced through automated checkers, such as the UNIX-based *lint*, developed for the C programming language in the late 1970s [?]. Similar tools were later made commercially available for other programming languages, like the C++ [?]. Within the FLOSS community, there are lint-like tools available and the tool CQual++ is available for both C and C++. The use of CQual++ on Debian packages to improve their code quality has been recently reported [?] although it should be noted that this application has taken place *after the fact*, rather than as a regular practice by the project’s developer community. Also, this tool is mainly focused on code improvement through the elimination of format string vulnerabilities. Within the KDE community, there is an explicit KDE Quality Team² and on its webpages, developers are guided towards information regarding Trolltech’s Qt application framework which is widely the core widget set within development of KDE applications. Coding standards can be automatically enforced by using Parasoft’s Insure++ in conjunction with Qt³.

² <http://quality.kde.org/>

³ <http://www.parasoft.com/jsp/products/home.jsp?product=Insure>

Each KDE project has the freedom to develop its own specific Coding Style; projects are recommended to follow the kdelibs coding style⁴. There is a coding style for kdelibs that is supported by a vim script in `kdesdk/scripts/kde-devel-vim.vim` that helps developers to keep the coding style correct. In addition to defaulting to the kdelibs coding style it will automatically use the correct style for Solid and kdepim code. Developers can add rules for other projects via the `SetCodingStyle` function.

In addition to these quality-related policies in KDE, there is also an automated quality assurance tool called the English Breakfast Network (EBN)⁵. The EBN is a tool for detecting and measuring aspects of quality within KDE as a whole. For example, the EBN measures code defects and errors in source documentation, such as spelling errors.

3 Definitions and GQM approach

The following concepts and attributes have been used to extract data from the two OSS forges, and to define the complexity and quality attributes (*i.e.* coupling) used to compare the samples. These metrics have been evaluated at the finest granularity level (*i.e.*, source functions or methods); in order to compare them on the system level, they should be elevated to system level [?]. Since the McCabe complexity [?] was evaluated, and the aggregation of this metric at the system level is complicated by its own definition (*i.e.*, the number of independent paths of an executable function, plus one), the rest of this study will analyze the software quality at the function level.

- (i) **Data points:** the selected KDE and SourceForge samples contain 50 projects each. The code of each project was downloaded from their respective Configuration Management System (either CVS or SVN) on the date of their latest available commit.
- (ii) **Source Function:** at the smallest level of granularity, both the functions of the procedural languages, and the methods of the Object-Oriented paradigm are considered as source functions, and treated as the basic unit measure for this study.
- (iii) **Complexity:** within the software engineering literature, this is a very broad term. To help in scoping it down, this paper consider complexity only from the point of view of the source functions. The attribute that will be used to empirically detect the complexity of source functions is the McCabe cyclomatic complexity. This is a measure of structural complexity, and it can be calculated from a graph representation of a function, with each executable statement being a node on the graph, and arrows between the nodes showing the execution pathways and decision or branching nodes. Cyclomatic complexity is calculated as the number of decision (or branching) nodes plus one. In this research McCabe complexity (Mc) was evaluated for all the functions of each software system studied.

⁴ <http://techbase.kde.org/index.php?title=Policies>

⁵ English Breakfast Network: <http://www.englishbreakfastnetwork.org>

- (iv) **Coupling**: it measures the degree to which each source function relies on other elements, that is, how interconnected is the code. Since this study is conducted at the function level, the union of all the function calls (and method invocations) form the network of couplings in a system. Each coupling can be uniquely categorized as inbound (*c_{in}*) or outbound (*c_{out}*) (or fan-in and fan-out), depending on the direction of the relative call. As an example, function ‘sign off’ (Figure 1) has two inbound (fan-in) and three outbound (fan-out) couplings. In this study we separately measured the number of inbound and outbound couplings of each function.




Fig. 1. example of ‘inbound’ and ‘outbound’ couplings for the function “sign off”

- (v) **Instability**: from past literature, instability is the ratio of fan-out coupling (*c_o*) to total coupling (*c_o* + *c_i*) such that $I = \frac{c_o}{(c_o + c_i)}$. This metric is an indicator of the resilience to change of software components (as in source functions) [?]. The range for this metric is 0 to 1, with $I = 0$ indicating the lowest instability for an element and $I = 1$ indicating a completely unstable element. In this paper it is argued that software systems should minimize the instability of its components in order to increase the quality of the underlying application.

3.1 KDE and SourceForge samples

The KDE forge ⁶ hosts a large number of OSS projects under a common name which together form both a desktop environment and associated application software, primarily for Unix-like operating systems. The KDE repository from which our samples are taken has somewhere in the order of 300 projects.

The SourceForge site ⁷ hosts more than 150,000 projects. However, it has been argued that a considerable proportion of SourceForge projects are to be considered as “tragedies” [?], owing to their failure to initiate a steady series of releases. No evidence of such a phenomenon in KDE is present, so to draw an accurate comparison, the sample from SourceForge has been extracted only from the pool of the “stable” projects (numbering approximately 20,000), *i.e.* those projects whose core developers labelled the status of the project with the tag “Production/Stable”.

It was shown in a previous work [?] that certain exogenous characteristics of KDE differ significantly from SourceForge; specifically, KDE projects are, on average, subjected to a greater amount of activity from a greater number of developers and have been developed over a greater period of time. In addition, KDE projects were also observed to be significantly smaller than their counterparts within SourceForge.

⁶ <http://www.kde.org/>

⁷ <http://sourceforge.net/index.php/>

3.2 GQM

The Goal-Question-Metric (GQM) method evaluates whether a goal has been reached, by associating that goal with questions that explain it from an operational point of view, and providing the basis for applying metrics to answer these questions [?]. The aim of the method is to determine the information and metrics needed to be able to draw conclusions on the achievement of the goal.

In the following, we applied the GQM method to first identify the overall goal of this research; we then formulate a number of questions related to the two OSS forges and their relative complexity (and conversely, their quality); and finally we collected adequate product and process metrics to determine whether the goal has been achieved.

- (a) Goal: The long-term objective of this research is to assess the presence of coding standards within distributed environments, and to evaluate its effectiveness towards the creation of software artefacts. In particular, the goal is to investigate whether OSS forges which reinforce coding standards achieve higher quality software.
- (b) Question: The purpose of this study is to establish differences between samples from KDE and SourceForge. Their complexity will be evaluated and a comparative research question will be evaluated via a direct comparison between the projects composing the two samples. Based on the complexity and coupling attributes defined above, the research question asks: in the presence of coding standards, will projects be less complex, and of higher quality, than of the counterpart forges not reinforcing any standards.
- (c) Metrics: Every project from both samples will have their functions evaluated in terms of McCabe cyclomatic complexity (to assess structural complexity) and the fan-in and fan-out metrics (to assess the coupling). As a compounded metrics, the value of *instability* will also be evaluated based on the definition given above. Once completed, these measurements will be summarised per project for both metrics, specifically into a median value, a maximum value and a variance.

4 Research Hypotheses

Two related hypotheses have been formulated based on the research question composing the GQM approach. The objective is to show that the samples from the two forges achieve different levels of complexity and coupling, as measured by the three attributes introduced above (mc , c_i , I). The directional tests will be also considered, in order to establish whether one forge achieves statistically higher quality results than the other: this means that a test will be run to check whether the KDE sample achieved a smaller level of complexity, or whether the KDE level of fan-in is larger than SourceForge. Based on common design principles, in fact, software designers and programmers should aim for low fan-out and high fan-in [?].

The statistical test have been used in this evaluation is the Wilcoxon test [?]: this is a non-parametric test, and the assumption of normality is not needed to run this test, as per the parametric family of tests. Nonparametric tests are more powerful in detecting population differences: since the normality assumptions are

not clearly satisfied when considering the distribution of complexity and coupling attributes, the Wilcoxon tests were applied.

In this paper, three hypotheses were considered, as also summarized in Table 1:

- (i) **Complexity:** the first hypothesis relates to the complexity, measured through the cyclomatic index. The null hypothesis assumes that the KDE and the SourceForge samples display different quality, i.e. different distributions of cyclomatic complexity. In addition, the directional test will check whether the KDE sample achieves a statistically smaller complexity than the SourceForge counterpart. The tests used will be both a simple and a directional Wilcoxon test: they will be applied on the median, mean, maximum and variance of the distribution of each observed project. These values are reported in the rightmost parts of Table 3 for the KDE sample, and Table 4 for the SourceForge sample.

	<i>Null (H_0)</i>	<i>Alternative (H_1)</i>	<i>Test</i>
complexity	KDE and SourceForge display similar levels of complexity	KDE and SourceForge display different levels	$mc_{kde} = mc_{sf}$
	KDE displays smaller complexity than SourceForge	KDE displays larger complexity than SourceForge	$mc_{kde} < mc_{sf}$
fan-in	KDE and SourceForge display similar levels of fan-in	KDE and SourceForge display different levels	$c.i_{kde} = c.i_{sf}$
	KDE displays larger fan-in than SourceForge	KDE displays smaller fan-in than SourceForge	$c.i_{kde} > c.i_{sf}$
instability	KDE and SourceForge display similar levels of instability	KDE and SourceForge display different levels	$I_{kde} = I_{sf}$
	KDE displays smaller instability than SourceForge	KDE displays larger instability than SourceForge	$I_{kde} < I_{sf}$

Table 1
Summary of the hypotheses to be tested

- (ii) **Fan-in:** the second set of hypotheses relates to the inbound coupling (or fan-in), measured through the inbound connections of each source function. The null hypothesis assumes that the KDE and the SourceForge show different distributions of the fan-in attribute. In addition, the directional test will check whether the KDE sample achieves a statistically higher fan-in than the SourceForge counterpart. The tests will be applied on the median, maximum and variance of the distribution of each project from KDE and SourceForge.
- (iii) **Instability:** the last set of hypotheses considers the compounded metrics presented above, which takes into account both the fan-in and the fan-out of each function. As above, the null hypothesis assumes that the KDE and the SourceForge show different distributions of the instability attribute. The added directional test instead compares the KDE and the SourceForge samples to detect whether the KDE one achieved a statistically lower instability. The tests will be applied on the median and variance of the distribution of each project from KDE and SourceForge. These values are reported in the central parts of

Table 3 for the KDE sample, and Table 4 for the SourceForge sample.




Fig. 2. Boxplots of measured attributes in the KDE and SourceForge samples

5 Results

This section summarizes the findings that were collected evaluating the research hypotheses. The attributes presented above were evaluated at the latest available change recorded in the CVS or SVN repositories: the couplings (fan-in and fan-out), instability and complexity of all the functions of the projects from KDE and SourceForge were calculated. For every project in the samples, the median and variance values of all the functions were used to compare the two forges. The maximum value was also used where appropriate.

As discussed above, a statistical test [?] allowed the null hypothesis to be either rejected or confirmed. The R programming language has been used to carry out these tests based on evaluation of the data extracted earlier from the respective forges [?]. A summary of the tests and their results is provided in Table 2 below: as visible, the majority of the paired tests have non-significant p-values, and therefore it is not possible to reject the null hypothesis.

5.1 Complexity

In this work the complexity was measured, at the functions level, via the McCabe cyclomatic index [?]. As results, this paper reports on both the actual distribution of these indexes, and the conditions on the fringes of this metric. In particular, the percentage of the complexity indexes larger than a threshold (i.e., 10, from [?]) was evaluated for each project, in order to detect whether one forge achieves lower quality in terms of highly complex elements. Visually, the boxplots in figure 2 show a difference of distribution: it can be seen how the KDE forge has a much compact distribution of these highly complex functions, with only one outlier around 25%. On the other side, the SourceForge projects have a wider boxplot distribution, with a larger median, and two outliers in its distribution.

Statistically, and as visible from Table 2, only one difference was found to be significant: the median of the distributions of the complexity is statistically different

Attribute	Metrics	W	p-value	
			1-tail	2-tail <i>a</i> : KDE > SF <i>b</i> : KDE < SF
Instability	median	693	<i>0.162</i>	<i>0.081^b</i>
	var	856	0.595	0.280 ^b
Fan-in	median	965	0.032	0.016^a
	max	590.5	0.044	0.978^a
	var	716.5	<i>0.058</i>	0.9715^a
McCabe	mean	767	0.73	0.36
	max	681.5	0.256	<i>0.128</i>
	var	872	0.49	0.247
	% ≥ 10	899	0.34	<i>0.17</i>

Table 2

Results of the Wilcoxon tests (p-values) when comparing the two forges – bold stands for high significance (95%), italics for weak significance

in the two sample, and also the directional test is significant: the median of the complexities in the KDE sample is lower than the SourceForge counterpart. In all the other tests, this significance was not achieved, hence making it impossible to reject all the null hypotheses: from the data collected in the two sample, it is not possible to conclude that the KDE sample achieves a higher quality than SourceForge, from the point of view of the complexity achieved.

5.2 Coupling

In this work the coupling of a function is measured primarily by its instability. Whilst functions within KDE projects were observed to be more stable on average, they did not achieve a statistically significant difference from functions of SourceForge projects.

As visible from Table 2, the median value of fan-in for KDE projects proved significantly higher than that of SourceForge projects, with H_0 of hypothesis 1.2 rejected for the test $c_{i_{kde}} < c_{i_{sf}}$ at $p = 0.016$. However, the SourceForge projects (which are typically much larger than KDE projects) possessed significantly higher values of both variance and maximum coupling. While the projects in SourceForge may contain some functions that are used by a larger number of clients, any randomly chosen function from a KDE project is more likely to have a greater fan-in than that of a SourceForge project. However, examining in more detail the median fan-in results (summarized in figure 2) reveals the marginality of difference: in both forges, most projects have a median fan-in value of 1 (35 KDE projects and 29 SourceForge projects).

6 Threats to Validity

Two main threats to validity have been identified; they are as follows:

- Size of populations – The SourceForge sample was taken from a considerably larger population than the KDE forge: approximately 20,000 “production/stable” SourceForge projects, compared around 300 within KDE. Consequently, the static sample size of 50 is not the same proportion of the population in both cases.
- Automated coupling analysis – Because of the large amount of analysis necessary (100 real-world non-trivial software projects) the evaluation of coupling was automated by analysis software, which is presently at a level of sophistication that has the following consequences:
 - Test suites, when included within the software package, are included in the analysis and so contribute to the perceived level of fan-in. It is arguable whether or not test suites should be considered “part of the software”;
 - The level of coupling, being measured by fan-in, is limited to being derived from a static view of the software; hence, dynamic coupling is not detected.

It is intended to refine the evaluation procedures for future analyzes to address these limitations in the analysis; and in future studies, the methods for selection of the sample populations will be given further consideration.

7 Conclusions

Within this paper the presence of coding standards and shared programming practices, and their effectiveness on the quality of the resulting software artefacts, were analyzed. In particular this work has argued that, within the Open Source paradigm, open collections of software projects could potentially benefit from higher quality software by openly sharing documents and guidelines on programming styles and coding standards. In order to assess this, two OSS forges (collections of software projects) were selected: the first (KDE) is a large container of applications which share the same graphical interface objectives, and guidelines are shared among developers to comply with existing coding standards. The second (SourceForge) is a wider-spectrum collection of projects, and no explicit effort is attempted towards a common framework of coding rules.

Two similar sets of 50 projects were randomly selected from two similar pools of projects (*i.e.* the “stable” set of projects) of the selected forges. Two attributes were measured, as negative proxies for software quality, both at the level of granularity of functions (or methods): cyclomatic complexity and inbound coupling (fan-in). A third characteristic was also measured, based on the definition of instability, composing the inbound and outbound couplings. A Goal-Question-Metrics approach was applied, and a research hypothesis derived based on the measured characteristics: a random sample of sample projects from a forge (KDE) where active actions are taken towards quality will display more quality than projects where such effort is absent (SourceForge).

For the coupling characteristics of the software’s functions (as measured by instability and fan-in) results showed only certain marginal improvements of KDE

over SourceForge. Instability was, on average, lower in KDE projects than those of SourceForge, but not significantly so. Functions within SourceForge projects had a greater variance of fan-in (with upper values exceeding those of KDE), However median fan-in values of KDE projects appeared significantly larger than SourceForge projects. Only a more detailed inspection revealed that the improvement could likely be interpreted as marginal only.

For the complexity characteristics of the software's functions, results showed that graphical differences could be observed in the distributions of the highly complex elements, resulting in SourceForge being more complex than KDE. Statistically, it was not possible to conclude that a significant difference could be observed in the complexity characteristics of the two forges. A general result was extracted from the metrics collected: whereas some results proved satisfactory (e.g. the statistical significance of the fan-in tests), other results contradict (complexity) these findings, creating a multi-faceted vision of software quality.

8 Further Work

The most important aim for future work on this topic is to refine the metrics and include new ones where it would increase understanding of the property being measured. In particular it is intended to expand the scope of metrics examining modular structure, by quantifying a project's individual structurally complex parts (modules with fan-in and fan-out values over a certain threshold) and by investigating the dynamic coupling.

It is also intended to carry out new work in addition to that outlined here. Firstly, we will investigate the evolution of the software from the perspective of these metrics by measuring their change over time and comparing the results from each forge. Secondly, having already identified a particular subset of the KDE forge that lays claim to high quality standards (the KDE PIM), it will be established if these claims are evidenced by any resultant effects on our measurements.

PLEASE LIST YOUR LASTNAME HERE

	functions	Instability (I)		McCabe (mc)			
		median	var	mean	max	var	$McCabe \geq 10$
ark	341	0.5	0.16	3.73	62	43.79	8.5%
dolphin	725	0.5	0.2	2.08	37	6.11	1.7%
kaddressbook	1379	0.44	0.21	2.29	56	12.75	2.5%
kamera	87	0.5	0.21	3.24	24	19.98	5.7%
kate	2544	0.5	0.18	2.94	160	33.61	5.0%
kbattleship	479	0.5	0.19	1.69	17	2.06	0.4%
kdebugdialog	22	0	0.27	2.41	7	3.11	0.0%
kfind	131	0	0.21	2.53	55	27.1	1.5%
kgamma	49	0.42	0.2	3.55	25	19.59	8.2%
khangman	91	0	0.22	2.08	22	8.27	3.3%
khtml	10253	0.5	0.19	3.06	782	115.29	4.7%
kioclient	11	0	0.22	3.55	18	23.67	9.1%
kjs	2053	0.67	0.19	2.65	309	88.22	3.3%
kjsembed	1231	1	0.2	2.1	102	12.43	1.0%
klinkstatus	941	0.37	0.22	1.61	17	2.63	0.9%
kmag	107	0	0.24	2.53	32	13.36	2.8%
kmailvt	116	0.89	0.2	3.74	17	15.38	10.3%
kmoon	37	0.25	0.18	3.11	43	51.38	5.4%
kmouth	340	0.33	0.18	2.53	17	6.55	2.9%
knetwalk	56	0	0.24	2.8	14	7.22	5.4%
knetworkconf	239	0	0.17	2.16	34	11.47	3.3%
knewsticker	67	0	0.26	1.78	16	4.15	1.5%
kpat	596	0.67	0.2	3.33	43	20.19	6.4%
kppp	716	0.5	0.16	2.49	84	18.94	3.6%
krfb	186	0	0.21	1.92	15	3.9	1.6%
krosspython	816	0.67	0.18	1.36	29	3.2	0.9%
ksim	1043	0.5	0.2	1.89	16	3.33	1.1%
ksquares	78	0.5	0.18	2.4	14	5.65	1.3%
kteatime	52	0.67	0.25	2.6	15	7.66	5.8%
ktnef	118	0	0.23	2.07	17	4.82	1.7%
kuiserver	102	0.5	0.22	1.89	21	9.03	3.9%
kxmlrpcclient	43	0.5	0.05	2	14	7.95	4.7%
lskat	242	0.5	0.18	2.82	25	11.62	5.4%
marble	1135	0.45	0.19	2.1	47	11.63	3.1%
nntp	20	1	0.2	7.3	25	34.96	25.0%
qtruby	1669	0	0.16	2.56	126	36.04	4.1%
shell	343	0.5	0.22	1.87	20	3.44	1.2%
solid	698	0.91	0.2	2.11	52	13.98	2.1%
sonnet	43	0.5	0.25	1.37	6	0.81	0.0%
umbrello	4683	0.63	0.17	3.21	94	30.83	6.6%

Table 3
Summary of instability and complexity attributes in sample KDE projects

PLEASE LIST YOUR LASTNAME HERE

	functions	Instability (I)		McCabe (mc)			
		median	var	mean	max	var	$McCabe \geq 10$
Aquila	48	0.75	0.25	3.81	38	36.92	8.3%
audiobookcutter	220	0.25	0.23	1.88	13	3.8	1.4%
Beobachter	223	0.5	0.22	1.53	9	1.09	0.0%
cotvnc	633	0.33	0.16	4.94	67	44.16	13.0%
cpia	370	0.5	0.12	8.36	96	203.26	24.1%
critical_care	2077	0.5	0.19	3.27	76	31.01	7.0%
csUnit	33	1	0.23	1.76	13	4.88	3.0%
expreval	327	0.42	0.22	1.72	51	15.23	1.2%
fitnesse	6944	1	0.18	1.31	40	1.07	0.1%
fn-javabot	492	0.25	0.23	2.48	43	14.54	4.1%
formproc	384	0	0.18	2.04	11	2.89	0.3%
fouever	1349	1	0.22	2.25	42	9.57	4.1%
freemind	2769	0.5	0.2	2.07	79	8.5	1.9%
galeon	4077	0.5	0.14	2.42	58	9.21	2.3%
hge	996	0.5	0.18	5.51	274	167.78	13.4%
icsDrone	33	0.5	0.11	7.48	108	346.82	18.2%
intermezzo	913	0.5	0.15	5.83	109	63.45	16.2%
j-trac	1154	0	0.22	1.55	19	2.32	0.9%
juel	824	1	0.19	1.93	69	9.98	1.7%
kpictorial	57	0	0.23	4.14	37	48.98	8.8%
mod-aspdotnet	93	0.5	0.21	2.8	18	9.97	6.5%
moses	10281	0.5	0.2	2.6	294	65.03	2.7%
nbcheckstyle	48	0.5	0.29	2.54	19	9.06	4.2%
neocrypt	68	0	0.26	2.96	16	12.64	8.8%
netstrain	26	0	0.13	5.04	23	30.52	23.1%
ozone	8288	0.89	0.2	1.96	65	5.64	1.6%
perpojo	201	0	0.21	1.58	9	1.82	0.0%
pf	58	0.75	0.17	2	66	127	24.1%
QPolymer	7039	1	0.19	1.2	65	2.54	0.6%
seagull	3256	0.5	0.2	5.86	344	183.12	14.4%
simplexml	70	0.33	0.13	5.53	26	31.38	14.3%
source	801	1	0.23	2.19	19	6.46	2.7%
swtjasperviewer	319	0.8	0.21	2.47	19	8.55	4.7%
txt2xml	165	0.63	0.2	1.65	10	2.03	1.2%
uniportio	42	0	0.17	3.5	23	18.79	4.8%
ustl	1801	0.6	0.16	1.43	16	1.24	0.3%
whiteboard	382	1	0.15	3.11	89	29	4.5%
wxactivex	259	0.5	0.19	2.53	36	18.93	5.0%
xmlnuke	127	1	0.2	2.47	20	8.7	3.9%
xqilla	4174	1	0.21	2.77	760	250.02	3.9%

Table 4
Summary of instability and complexity attributes in sample SourceForge projects