

EXHIBIT 1

1 ROBERT A. VAN NEST — #84065
rvannest@kvn.com
2 CHRISTA M. ANDERSON — #184325
canderson@kvn.com
3 KEKER & VAN NEST LLP
4 710 Sansome Street
San Francisco, CA 94111-1704
5 Telephone: (415) 391-5400
6 Facsimile: (415) 397-7188

7 DONALD F. ZIMMER, JR. (SBN 112279)
fzimmer@kslaw.com
8 CHERYL A. SABNIS (SBN 224323)
csabnis@kslaw.com
9 KING & SPALDING LLP
10 101 Second Street – Suite 2300
San Francisco, CA 94105
11 Telephone: (415) 318-1200
12 Facsimile: (415) 318-1300

13 Attorneys for Defendant
14 GOOGLE INC.

SCOTT T. WEINGAERTNER (*Pro Hac Vice*)
sweingaertner@kslaw.com
ROBERT F. PERRY
rperry@kslaw.com
BRUCE W. BABER (*Pro Hac Vice*)
bbaber@kslaw.com
KING & SPALDING LLP
1185 Avenue of the Americas
New York, NY 10036-4003
Telephone: (212) 556-2100
Facsimile: (212) 556-2222

IAN C. BALLON (SBN 141819)
ballon@gtlaw.com
VALERIE HO (SBN 200505)
hov@gtlaw.com
HEATHER MEEKER (SBN 172148)
meekerh@gtlaw.com
GREENBERG TRAUIG, LLP
1900 University Avenue
East Palo Alto, CA 94303
Telephone: (650) 328-8500
Facsimile: (650) 328-8508

15 **UNITED STATES DISTRICT COURT**
16 **NORTHERN DISTRICT OF CALIFORNIA**
17 **SAN FRANCISCO DIVISION**

19 ORACLE AMERICA, INC.

20 Plaintiff,

21 v.

22 GOOGLE INC.

23 Defendant.

Case No. 3:10-cv-03561-WHA

Honorable Judge William Alsup

**OPENING EXPERT REPORT OF DR.
OWEN ASTRACHAN**

24
25
26
27
28

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28

TABLE OF CONTENTS

- I. INTRODUCTION 4
- II. DOCUMENTS AND INFORMATION CONSIDERED 6
- III. BRIEF SUMMARY OF MY OPINIONS 6
- IV. BRIEF BACKGROUND ON JAVA..... 8
- V. DETAILED STATEMENT OF THE BASIS FOR MY
OPINIONS ON APIs 10
 - A. What is an API?..... 15
 - B. Useful Analogies 16
 - C. Purposes of an API..... 18
 - D. The Elements of an API 20
 - E. The Components of a Method Declaration 23
 - F. Organizing Related Methods into Packages 27
 - G. The Distinction Between An API and Its
Implementation..... 28
 - H. Sun and Oracle Also have Implemented and Distributed
APIs from Other Software..... 35
 - 1. Sun Implemented and Distributed APIs from
Previous Generations of Spreadsheets as Part of
StarOffice and OpenOffice.org 36
 - 2. Sun Implemented and Distributed APIs from
Linux as Part of the Solaris Operating System 39
 - 3. Oracle Implemented and Distributed APIs from
IBM as Part of the Oracle Database Server 44
 - I. Basic Example of Java Method Usage..... 46
 - J. The APIs at Issue are Methods of Operation 47
 - K. The Java API Package Names Are Dictated By Function 51

1 L. The Java API Class and Method Names Are Dictated by
 2 Function..... 55

3 M. The Java Names Are the Product of Mechanical Rules..... 64

4 N. The Organization of API Elements is Dictated By
 5 Function And Does Not Reflect Creativity 66

6 O. The Same Analysis Applies to the Native Files..... 67

7 P. Many API Elements are Drawn from the Public Domain
 8 and Are Not Original to Java..... 69

9 Q. The APIs at Issue Are Necessary For Basic
 10 Functionality and Interoperability 73

11 R. The APIs at Issue Are Demanded by the Industry..... 74

12 VI. THE ANDROID PLATFORM IS NOT VIRTUALLY
 13 IDENTICAL OR SUBSTANTIALLY SIMILAR TO THE
 14 JAVA PLATFORM..... 77

15 VII. ANDROID’S DOCUMENTATION OF THE APIs AT ISSUE
 16 IS NOT VIRTUALLY IDENTICAL OR SUBSTANTIALLY
 17 SIMILAR TO ORACLE’S DOCUMENTATION..... 80

18 VIII. THE TWELVE FILES OR PORTIONS OF FILES
 19 ALLEGED BY ORACLE TO HAVE BEEN COPIED ARE
 20 QUALITATIVELY AND QUANTITATIVELY
 21 INSIGNIFICANT AND THEY ADD NO OR VERY LITTLE
 22 VALUE TO ANDROID 82

23 A. TimSort Files 83

24 B. Security Test Files 85

25 C. Comments in CodeSourceTest.java and
 26 CollectionCertStoreParametersTest.java 88

27 EXHIBIT A: OWEN ASTRACHAN CV

28 EXHIBIT B: DOCUMENTS AND INFORMATION REVIEWED

1 EXHIBIT C: EXCEL AND STAROFFICE SPREADSHEET

2

3 EXHIBIT D: LX_BRAND SYSCALL TABLE

4

5 EXHIBIT E: SOURCE CODE FOR SLOCCOUNTER.PY AND

6

7 SLOCCOUNTERTOTAL.PY

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

1 **I. INTRODUCTION**

2
3 1. I am Professor of the Practice of Computer Science and Director of Undergraduate
4 Studies in the Computer Science Department at Duke University. I earned my AB degree
5 with distinction in Mathematics from Dartmouth College and MAT (Math), MS, and PhD
6 (Computer Science) from Duke University. I teach undergraduate computer science
7 courses using the Java, C++, and Python programming languages, and helped develop
8 broadly-used teaching materials, including a C++ textbook and Java language
9 programming exercises and documentation. I received a National Science Foundation
10 CAREER award in 1997 to incorporate design patterns in undergraduate computer
11 science curricula, an IBM Faculty Award in 2004 to support componentization in both
12 software and curricula, and was one of two inaugural NSF CISE Distinguished Education
13 Fellows in 2007 to revitalize computer science education using case- and problem-based
14 learning. My research interests have been built on understanding how best to teach and
15 learn about object-oriented programming, software design, and computer science in
16 general; and I am now working on developing a portfolio of substantial, interdisciplinary
17 problems that help explain how computer science is relevant to students in the social and
18 natural sciences. My qualifications and information regarding my prior testimony are
19 attached hereto as Exhibit A.

20
21
22 2. I am being compensated for my work in this litigation at the rate of \$300 an hour. My
23 compensation does not depend in any way on the outcome of this litigation.

24
25 3. I have been asked by Google to:
26
27
28

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28

- a. opine on whether the Java Application Programming Interface (“API”) specifications from which Oracle alleges certain parts of the Android platform are derived are methods of operation;
- b. opine on whether the allegedly infringing materials are driven by functional considerations, considerations of interoperability and efficiency, industry practice or demand, or drawn from public domain material;
- c. opine on whether the allegedly copyrighted Oracle works relating to the Java platform are virtually identical or substantially similar to the Android platform;
- d. opine on whether the allegedly copyrighted Oracle documentation relating to the Java APIs are virtually identical or substantially similar to Google’s documentation for the APIs in Android; and
- e. opine on whether the 12 files and/or portions of those files alleged by Oracle to include material literally copied by Google are qualitatively and/or quantitatively insignificant.

- 4. I understand that I may further be asked by Google to review submissions related to copyright issues from Oracle’s experts, and to provide my opinions on issues raised by any such submissions.
- 5. I understand that I may be called upon to testify in this case regarding my opinions and analyses set forth in this report. If called upon to testify, I may use various demonstratives, including tables or drawings, to assist in presenting my testimony.

1 **II. DOCUMENTS AND INFORMATION CONSIDERED**

2
3 6. My opinions are based on my relevant knowledge and experience, as well as review of
4 the documents and information identified in Exhibit B.

5 **III. BRIEF SUMMARY OF MY OPINIONS**

6
7 7. "Java" may refer to three very different things: the Java programming language, the Java
8 Application Programming Interfaces (APIs), or software source code that references and
9 implements the APIs. In this case, except for a very small number of files addressed in
10 Section VII below (12 files out of approximately 57,000), it is my understanding that
11 Oracle does not allege infringement of the software source code referencing the APIs.
12 Nor, I am informed, does it allege infringement because Android is written in the Java
13 language. I have been informed that Oracle's claim of infringement is based on Google's
14 creation of software source code written in the Java language that references and
15 implements Java APIs.

16
17 8. The Java language is a programming language. As with any language, it has a basic
18 syntax and grammar that must be followed for code written in the Java language to be
19 understood by a computer. Java's syntax includes such utilitarian features as spacing,
20 punctuation, and the meaning of a limited number of defined words and phrases. Starting
21 in the mid-1990s, Sun Microsystems widely promoted the use of the Java language by
22 developers, businesses and the general public, without restriction. It is my understanding
23 that neither Sun nor Oracle, which acquired Sun, claims that the use of the Java language
24 to write software is infringing.

25
26 9. As explained in more detail starting at paragraph 24, an API provides programmers with
27 a way to access the functionality of a software service. For example, most programming
28

1 languages provide a way to calculate the square root of a number. Java does this by way
2 of an API. In order, for example, to calculate the square root of 25.0, a Java programmer
3 includes the text `sqrt(25.0)` in the text of his program. The text “`sqrt(a)`” — where “a” is
4 replaced by the variable or number that the programmer for which the programmer wants
5 to calculate the square root — is the API for the square root “method.”
6

7 10. As explained in more detail starting at paragraph 52, APIs are implemented by software.
8 For example, the `sqrt(a)` API can be implemented in many different ways. First, there are
9 many different mathematical algorithms for calculating square roots (much as there are
10 many different ways to make an apple pie). Second, there are many different ways to
11 write the programming code that implements any given algorithm (just like two people
12 who write down the same recipe can describe the various steps using different words and
13 sentences). However, if one sets out to implement an API, the one part that cannot
14 change is the API itself. For example, if one wants to implement the `sqrt(a)` API, one
15 cannot change the method name from “`sqrt`” to, say, “`squareroot`.” As I will discuss
16 below, referencing an API requires the use of the API’s method declarations (names,
17 data, and data types). To implement an API, one must use all these exactly as the API
18 requires them to be used. If even a minor change is made, code that references the API
19 will fail to operate. That said, as noted above, there are potentially different ways to
20 write the underlying implementing code for a given API. It is my understanding that
21 except for the 12 files discussed below in section VII, Oracle has not identified any literal
22 code copying by Google in this case.
23
24

25 11. In this case, it is my understanding that Oracle contends that referencing the API, via
26 method names, data names, and data types, is infringing, but as I discuss below in Section
27 V.J, the APIs at issue in this case are purely functional. In addition, with respect to
28

1 Oracle's allegation of literal copying of 12 files out of 57,000 files in Android, as I
2 discuss below in Section VIII, it is my opinion that the allegedly copied material in these
3 files is qualitatively and quantitatively insignificant.

4
5 12. It is my opinion that the APIs at issue are methods of operation.

6 13. It is my further opinion that any similarity between the names of elements (such as
7 package, class and method names) in the implementations in these APIs in the Java and
8 Android platforms is driven by functional considerations. It is also my opinion that any
9 similarity between the organization of elements in the implementations in these APIs in
10 the Java and Android platforms is also driven by functional considerations.

11
12 14. It is my opinion that many of the names of elements of the Java API were drawn from
13 usage in other languages or platforms that pre-date Java.

14
15 15. It my opinion that Google's use of these APIs is necessary for interoperability and
16 efficiency, and/or driven by industry demand.

17 16. It is my opinion that the Android platform is not virtually identical or substantially
18 similar to the allegedly copyrighted works relating to the Java platform.

19
20 17. It is my opinion that any similarities between Google's documentation of the APIs at
21 issue and Oracle's documentation are driven by functional considerations and industry
22 practice regarding such documentation.

23
24 **IV. BRIEF BACKGROUND ON JAVA**

25 18. Oracle uses "Java platform" to mean a variety of interchangeable and overlapping
26 elements. These elements purportedly include the Java programming language itself, an
27 "object-oriented" programming language that uses syntax heavily based on prior
28

1 languages such as the “C” and “C++” programming languages. Also included in these
2 elements are a program known as a “compiler” that creates the “bytecode” in which Java
3 programs are executed; a virtual machine that executes the bytecode; and a set of core
4 libraries that facilitates the development of applications for the Java platform by
5 providing basic system or language functionalities. Java is a popular programming
6 language, and a variety of software, including internet services and mobile applications,
7 is written in the Java language.
8

9 19. Like any high-level programming language, the Java programming language contains
10 many rules of grammar and syntax that cannot generally be varied. For example, a
11 statement adding two numbers can only be written in certain ways, and the language
12 requires specific and precise key words to express such things as variable types (integers,
13 strings, or Booleans) and more complex object types such as dates or database queries.
14 In addition, the Java language, like many programming languages, employs key words
15 and operators (such as plus and minus symbols) that can only be used for specific
16 purposes and in specific ways; using them for other purposes will cause a program to fail
17 to function correctly. As a result, much of the structure and appearance of code written in
18 the Java programming language is dictated by these functional considerations.
19

20 20. Java’s specifications, including the specifications for the language API packages at issue
21 here, were published or made available in various forms, including in books and on the
22 Java website, starting with the release of version 1.0 in 1996. Several revisions have
23 been released since then, including version 1.5, which in my understanding is the most
24 recent version at issue in this case. Note that, at times, Java version 1.5 has also been
25 referred to as Java version 5.0. For consistency, I will refer to it as Java 1.5.
26
27
28

1 **V. DETAILED STATEMENT OF THE BASIS FOR MY OPINIONS ON APIS**

2
3 21. Based on my review of Oracle's responses to Google's interrogatories, I understand that
4 Oracle is claiming that Google's implementation of the Java API specifications for the
5 following packages infringe Oracle's copyrights. (It is my understanding that Oracle,
6 with the exception of the 12 files discussed in Section VIII, has not identified instances of
7 the copying of specific code.)

8 java.awt.font

9
10 java.beans

11 java.io

12
13 java.lang

14 java.lang.annotation

15
16 java.lang.ref

17 java.lang.reflect

18
19 java.math

20
21 java.net

22 java.nio

23
24 java.nio.channels

25 java.nio.channels.spi

26
27 java.nio.charset

28

1 java.nio.charset.spi
2
3 java.security
4
5 java.security.acl
6
7 java.security.cert
8
9 java.security.interfaces
10
11 java.security.spec
12
13 java.sql
14
15 java.text
16
17 java.util
18
19 java.util.jar
20
21 java.util.logging
22
23 java.util.prefs
24
25 java.util.regex
26
27 java.util.zip
28
29 javax.crypto
30
31 javax.crypto.interfaces
32
33 javax.crypto.spec
34
35 javax.net

1 javax.net.ssl
2
3 javax.security.auth
4
5 javax.security.auth.callback
6
7 javax.security.auth.login
8
9 javax.security.auth.x500
10
11 javax.security.cert
12
13 javax.sql
14
15 javax.xml
16
17 javax.xml.datatype
18
19 javax.xml.namespace
20
21 javax.xml.parsers
22
23 javax.xml.transform
24
25 javax.xml.transform.dom
26
27 javax.xml.transform.sax
28 javax.xml.transform.stream
29
30 javax.xml.validation
31
32 javax.xml.xpath

1 22. Based on my review of Oracle's responses to Google's interrogatories, I understand that
2 Oracle is also basing its infringement claim on the following native code implementations
3 of Java API classes:

4
5 java_lang_Class.c

6
7 java_lang_Object.c

8
9 java_lang_reflect_AccessibleObject.c

10
11 java_lang_reflect_Array.c

12
13 java_lang_reflect_Constructor.c

14
15 java_lang_reflect_Field.c

16
17 java_lang_reflect_Method.c

18
19 java_lang_reflect_Proxy.c

20
21 java_lang_Runtime.c

22
23 java_lang_String.c

24
25 java_lang_System.c

26
27 java_lang_Throwable.c

28
29 java_lang_VMClassLoader.c

30
31 java_lang_VMThread.c

32
33 java_security_AccessController.c

34
35 java_util_concurrent_atomic_AtomicLong.c

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28

sun_misc_Unsafe.c

23. Based on my review of Oracle's responses to Google's interrogatories, I understand that Oracle also bases its copyright claim on code and comments in the following files that allegedly have been copied from Oracle code or comments in Oracle's source code:

Allegedly copied test files:

AclEntryImpl.java

AclImpl.java

GroupImpl.java

OwnerImpl.java

PermissionImpl.java

PrincipalImpl.java

AclEnumerator.java

PolicyNodeImpl.java

Allegedly copied comments (but not source code):

CodeSourceTest.java

CollectionCertStoreParametersTest.java

Allegedly contain copied source code:

TimSort.java

ComparableTimSort.java

1 **A. WHAT IS AN API?**

2
3 24. An Application Programming Interface (API) is "a particular set of rules and
4 specifications that software programs can follow to communicate with each other. It
5 serves as an interface between different software programs and facilitates their
6 interaction, similar to the way the user interface facilitates interaction between humans
7 and computers." Wikipedia, *Application programming interface*,
8 [http://en.wikipedia.org/w/index.php?title=Application_programming_interface&oldid=43](http://en.wikipedia.org/w/index.php?title=Application_programming_interface&oldid=437864024)
9 7864024 (as of July 13, 2011, 00:30 GMT). An API provides a specified and
10 documented mechanism to invoke, operate, and interact with software services. The
11 interface itself is implemented by software, *i.e.*, by source code that is written to provide
12 the functionality of the interface.¹

13
14 25. APIs are used by software developers when writing software that will utilize the
15 functionalities operated through these communications. When used alone, the term API
16 can refer to either the set of rules and specifications, or to the software that implements
17 the rules and specifications and therefore is operated by the communication from another
18 program, as explained in more detail in paragraph 52. API can also refer to either a
19 specific "element," "component," or functionality within the API, or to a collection of
20

21
22 ¹ Newton's Telecom Dictionary, 25th Edition, defines API similarly as "Software that an application
23 program uses to request and carry out lower-level services performed by the computer's ... operating
24 system." Sun's Java glossary (*available at* <http://java.sun.com/docs/glossary.html>) also provides a
25 definition for API: "The specification of how a programmer writing an application accesses the behavior
26 and state of classes and objects." Although these definitions use different terminology, they are not
27 materially different from, and are in fact consistent with, the definition I have presented above.
28

1 them. This report will generally use API to mean such a collection, and “API elements”
2 or “API components” to refer to the individual mechanisms within the collection.

3
4 **B. USEFUL ANALOGIES**

- 5 26. APIs are similar to the interfaces that a computer user uses to operate software, like a
6 keyboard command or button. In each case, the person seeking to use the program does
7 something to inform the program being used that he wants a specific action to happen,
8 and then that action happens. No deep expertise or understanding of the inner workings
9 of the computer system is needed by the person seeking to use the program. For
10 example, a computer user might type the “Ctl+P” key combination or click an icon that
11 looks like a printer, and then, in the dialog box that appears, choose the file to be printed
12 and the number of copies that should be printed. Typing “Ctl+P” or clicking the icon
13 would invoke the underlying printing functionality, and (once the number of copies is
14 specified) cause the software to print the document that number of times. The user does
15 not need to have substantial understanding of the underlying printing mechanism, he just
16 needs to learn and remember the familiar “Ctl+P,” give the necessary information (*e.g.*,
17 number of copies he wants printed), and the computer takes care of the rest.
- 18
19
20 27. Similarly, when invoking or using an API in a software application, a programmer should
21 not need to review or understand the underlying implementation or source code for the
22 API, as that code has already been written. Like using “Ctl+P” to print, he only needs to
23 know the name and functionality of the API. In order to write software that prints, a
24 programmer would read and learn their chosen operating system’s API for printing, and
25 then invoke that API from their program, telling the API critical information like what
26 document to print and how many copies to print. Just like typing the “Ctl+P” command
27 or clicking the printer icon, using the name of the API element in the software invokes
28

1 the underlying functionality of the API and causes printing to happen, without the
2 programmer needing to have a deep knowledge of the particular mechanisms that allow
3 the API to function. Among other benefits, this means that a programmer can use an API
4 to create software that works on different printers (color, black and white, inkjet, etc.)
5 without knowing in detail how those different printers work, as long as the underlying
6 implementation supports them.
7

8 28. An API also can be analogized to the interface for driving a typical car. Every car has a
9 variety of elements that are part of the overall system of communication and operation
10 that a driver must understand and use in order to drive the car. These elements include
11 the gear lever, the turn signal stalk, the steering wheel, and the accelerator and brake
12 pedals. Some of these elements provide information from the driver to the car, while
13 others provide information from the car to the driver, and others do both things at the
14 same time. In combination, these elements form the interface to the “application” that is
15 the car, allowing a driver to “program” the car to do their bidding by using those
16 elements to operate the car.
17

18 29. Thus, for example, the driver of a car can make it accelerate by pressing the accelerator.
19 The further the driver presses the accelerator, the faster the car speeds up. The
20 accelerator can be thought of as an API for the car that makes the car go. Every car that
21 implements this API will have an accelerator, and each of them will share in common the
22 fact that the car speeds up faster the further the accelerator is pressed down.
23

24 30. So long as a driver understands this functionality — that the rate of speeding up a car
25 depends on how far the accelerator is pressed down — the driver need not know *how* this
26 happens. Similarly, if a driver understands how the other interfaces to the car’s
27
28

1 operation function (the turn signals, steering wheel, etc.), the driver need not know how
2 the engine, light bulbs, or transmission work.

3
4 **C. PURPOSES OF AN API**

5 31. The primary purpose of APIs is to allow one piece of software to speak to another piece
6 of software in a clearly defined, reusable, interoperable way. This simple goal has a
7 number of important ramifications and benefits.

8
9 32. Familiar interfaces make it simpler to use things and to use them more expertly. When
10 using a new car, most drivers do not think about how that particular steering wheel
11 works. For example, the wheels of the car might be turned by a rack and pinion system
12 or a recirculating ball system. But the steering wheel itself functions the same way
13 regardless of how the internal steering mechanism and system is designed, *i.e.*, when the
14 driver turns the wheel to the left, the car moves in the left direction. This interface — the
15 same, familiar steering wheel — facilitates using the car, regardless of which specific car
16 is being used. The same thing happens in software — using (or providing) a standard
17 API allows the users of that API (software programmers) to move between any software
18 platforms that provide the same API, because their familiarity and existing skills in using
19 that particular interface transition over.

20
21 33. A defined, fixed API allows different programs to substitute for each other, which gives
22 users the ability to move from one piece of software to another. In this way, APIs enable
23 user choice between competing software providers, and therefore help to promote
24 competition, innovation and choice in the software market. For example, if a software
25 platform provides a set of APIs, a subsequently created platform that implements those
26 same APIs (for example, by writing different source code to implement those APIs) can
27
28

1 improve competition and reliability because developers who use the platform would
2 already be familiar with the APIs and so would be more able to leverage their existing
3 knowledge and complementary software. In our printing analogy, another program that
4 provides the same commands (like “Ctl+P”) will be much easier for a user to switch to.
5 In fact, once a user is used to “Ctl+P,” he will often be confused if a program uses
6 something else to control printing.
7

8 34. APIs also help software programmers by insulating programmers from underlying
9 complexity. This is referred to by programmers as “encapsulation.” In the car example,
10 the internal steering system can be changed specifically because the familiar steering
11 wheel interface (the car’s “API”) has hidden these implementation details from the
12 driver. This shielding and simplification is an important part of what an interface
13 provides — most users do not need to know the details, which has in the past allowed car
14 manufacturers to switch from old technologies to new ones without introducing a new
15 learning curve for consumers. The concept that moving the wheel left turns the car left
16 remains the same, and that allows consumers to rely on this familiar concept, regardless
17 of which car they use.
18

19 35. APIs also help programmers and the industry by allowing software to be reused. This is
20 important; new code is difficult and expensive to write and test, and so individual
21 programmers and corporations like to reuse code as much as possible. Even if they
22 cannot reuse an entire program (say, because the two pieces of hardware are very
23 different, as they are between a desktop computer and a phone) they still prefer to reuse
24 as many parts of the software as possible. APIs help make this possible by allowing the
25 same basic functionalities to be provided and used in a replicable, but portable, way.
26
27
28

1 **D. THE ELEMENTS OF AN API**

- 2
- 3 36. An API typically consists of what programmers call *methods* or, equivalently, functions.
- 4 These two terms are synonyms, used somewhat interchangeably depending on the
- 5 programming language. Java programmers (and therefore this report) use the term
- 6 “methods.” Both represent the same thing — a piece of software that performs a specific
- 7 function and can be reused when needed. Methods are the primary mechanism by which
- 8 programmers invoke the functionalities provided by a software system.
- 9
- 10 37. In the user interface analogy, “Ctl+P” and the printer icon are methods. A user
- 11 interacting with a software program might use a menu or a menu shortcut to open a file,
- 12 or to save or print what has been opened, *e.g.*, in a word processing program. Often an
- 13 icon of a printer or a disk can be pressed to invoke the same functionality as choosing a
- 14 menu item or the menu shortcut. In all three cases — pressing the icon, choosing the
- 15 menu item, or typing a keyboard shortcut — the same underlying software is invoked and
- 16 causes a specific action — printing the file or saving it, for example. The functions or
- 17 methods in an API are directly analogous — just as the user might click on an icon or
- 18 press a sequence of keys to use a keyboard shortcut to invoke more complicated
- 19 operations such as printing or saving, the program calls the function or method to invoke
- 20 and control a more complicated service or feature provided by the underlying software.
- 21 For example, when a Java programmer wants to get the square root of 25, his program
- 22 will have to contain the following text:
- 23

24 `sqrt (25 . 0)`

25

26 This will cause the underlying system to do the math and tell the program that the answer

27 is “5.0”. Similarly, to get the absolute value of -25, the program must contain `abs (-`

28

1 25) This will cause the underlying system to do the math and tell the program that the
2 answer is “25.”

3
4 38. Related methods are often grouped together to make them easier to use, frequently into
5 groups called (depending on the programming language) libraries or packages. In Java,
6 these groupings are called packages. Because Java is what is known as an “object-
7 oriented” language, related methods are themselves encapsulated in a class, and then
8 related classes are encapsulated into a package or a sub-package. To put it a different
9 way, the API packages include subparts or files known as classes, and within these
10 classes are methods. As an analogy, one can think of menus (like “File,” “Edit,” etc.) as
11 “packages” of menu items, which organize the menu items so that they are grouped
12 together in reasonable groupings. For example, to print using the menu, a user needs to
13 know that Print is under File, rather than under Edit.

14
15 39. As an example from the Java API, the java.lang package (according to Oracle’s
16 documentation) “[p]rovides classes that are fundamental to the design of the Java
17 programming language.” One of these “fundamental” classes is the “Math” class, which
18 Oracle describes as containing “methods for performing basic numeric operations such as
19 the elementary exponential, logarithm, square root, and trigonometric functions.” The
20 actual methods contained with the class are listed in this chart:

Method Name	Functionality of the Method
abs	Returns the absolute value of the argument. (Four variants)
acos	Returns the arc cosine of the argument.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28

Method Name	Functionality of the Method
asin	Returns the arc sine of the argument.
atan	Returns the arc tangent of the argument.
atan2	Converts rectangular coordinates (two arguments) to polar coordinates.
ceil	Returns the smallest integer that is not less than the argument, <i>e.g.</i> , if the argument is 1.9, will return 2.
cos	Returns the cosine of the argument.
exp	Returns <i>e</i> raised to the power of the argument.
floor	Returns the largest integer that is not more than the argument, <i>e.g.</i> , if the argument is 1.9, will return 1.
IEEEremainder	Returns the remainder of two arguments as prescribed by the IEEE 754 standard.
log	Returns the natural logarithm of the argument.
max	Returns the greater of two arguments, <i>e.g.</i> , if the arguments are 3 and 4, will return 4. (four variants)
min	Returns the lesser of two arguments, <i>e.g.</i> , if

Method Name	Functionality of the Method
	the arguments are 2 and 3, will return 2. (four variants)
pow	Returns the value of the first argument raised to the power of the second argument.
random	Returns a random number between 0 and 1.
rint	Returns the closest integer to the argument.
round	Returns the closest number to the argument.
sin	Returns the sine of the argument.
sqrt	Returns the square root of the argument.
tan	Returns the tangent of the argument.
toDegrees	Returns the result of a conversion of the argument (an angle in radians) to degrees.
toRadians	Returns the result of a conversion of the argument (an angle in degrees) to radians.

E. THE COMPONENTS OF A METHOD DECLARATION

40. Every method has several important characteristics that collectively are referred to as the “method declaration.” The first is simply the method’s *name*. Method names describe the purpose of the method, so that a programmer can easily memorize and recognize the purpose from the method’s name, and vice-versa. A simple example of this is the method

1 in the table above named “abs,” so named because its function is to calculate the absolute
2 value of a number. To use a method, the programmer must know the method’s name. If
3 the programmer does not know the precise name, or knows only something similar, he
4 cannot use the method, because the software cannot guess at what the programmer meant.
5 For example, if a Java programmer writes “squareroot(25.0)” instead of “sqrt(25.0)”, this
6 will result in an error instead of calculating the square root of 25.0.
7

8 41. The second important characteristic of essentially every method is the set of *arguments*
9 that the method expects to receive when invoked. When the method is called, the
10 programmer typically provides information to the method that informs the software
11 exactly what the programmer wants to happen, just as a user must usually specify how
12 many copies he wants printed after he clicks the print button. The information provided
13 to the method is called an argument (or parameter), and a method is said to “accept” the
14 permitted arguments. The ability of a method to accept an argument is what allows a
15 general purpose method to act on specific data.
16

17 42. For example, think of the “plus” or “add” button on a calculator. This is a “general
18 purpose” button — it can add any numbers one can type in, not just one specific set of
19 numbers. If one thinks of the “plus” button on a calculator as a method, the numbers one
20 asks the calculator to add (say, 2 and 2) are the parameters to the “plus” button — those
21 parameters determine the specific outcome of the general purpose button. Similarly, the
22 number of copies one tells the print dialog (or print method) to print is also an argument
23 — they again tell the general purpose function (“print”) a specific behavior (“print two
24 copies.”) In the “abs” function mentioned previously, there is only one argument, and
25 that argument is simply a number, whose absolute value the program wishes to calculate.
26
27
28

- 1 43. These arguments or parameters must be defined when the API is first created, and are
2 typically limited. For example, it would not make sense to ask one's printer to print
3 "hippopotamus" number of copies of a document — that argument must be a number. In
4 fact, the definition of a method in many languages, including Java, will indicate what
5 "type" of argument a function will accept, such as an integer, a string, or another data
6 type. A steering wheel, similarly, can accept arguments of left, right, or any angle in-
7 between, but cannot accept "up" or "down." The functionality of each method constrains
8 what parameter(s) are acceptable, and if the proper parameters are not passed to the
9 method, any attempt to use the method will fail.
- 10
- 11 44. When a program uses a method and passes it the arguments, the method then typically
12 returns a *result* that the program can use for other purposes. This result is the final
13 important characteristic of the method, and is called the *return* (sometimes the *return*
14 *value*). In the calculator example, where plus is the method and the arguments are 2 and
15 2, the return value will be 4 — 2+2 returns 4. For the abs method, which computes the
16 absolute value of a number, when the argument is 2 or -2, the return value will be 2.
- 17
- 18 45. The purpose of the return value is to return information that can be used by the program
19 for other purposes. For example, after one asks one's calculator to add 2 and 2, the
20 calculator returns "4," which one can use as the first step in the next math problem one
21 intends to solve. Similarly, the return may be a message indicating the status of a method
22 — for example, a "print" method might return "OK" (telling the program that the printing
23 functionality has successfully completed) or "OUT OF PAPER" (telling the program that
24 the printing functionality has hit a snag). These status messages would in turn be handled
25 by other methods, possibly doing something like popping up an error message, or silently
26 concluding that all is well and allowing the user to continue with his work.
27
28

1 46. To summarize, each method declaration has three parts — the name, arguments, and
2 return. As a shorthand, programmers may refer only to the name of the method, but to
3 fully know what method they are describing, it is necessary to know all three parts of the
4 method declaration. They can be defined succinctly as:

5
6 **name:** the method name, which indicates its purpose and is used by a
7 programmer or program to call or invoke the method.

8 **argument:** the data on which the method acts. The data passed as an
9 argument to a method is often manipulated and referred to within
10 the method itself.

11
12 **return:** the result of calling the method with specific arguments. This is
13 “returned” to the programmer.

14
15 47. The documentation for a method will combine these pieces to form a reference for
16 programmers using the API. For example, the brief version of the documentation for the
17 abs method is:

18 `int abs(int a)` Returns the absolute value of an int value.

19
20 While this may not be easy for a non-programmer to understand, it is quite
21 straightforward to a programmer:

- 22
- 23 • The first part (“int”) shows that the return will be an “int” (short for an
24 “integer”; i.e., a number). This tells the programmer what type of result to
25 expect when using the function.
 - 26 • The second part (“abs”) is the name of the method.
- 27
28

- 1 • The part in the parentheses (“int a”) is the argument. Again, this uses “int”
2 to indicate that a single integer is expected; if the abs method is given
3 something other than an integer, such as “hippopotamus,” an error will
4 occur. (The letter “a” is a convenient name for the argument, and can be
5 changed without affecting compatibility.)
6
7 • The first, second and third parts discussed above comprise the “method
8 declaration.” The final part is a brief explanation of what the method
9 does.
10

11 In combination, this short statement will allow a programmer to know how to use “abs”
12 in their program to find the absolute value of a number. I will discuss documentation in
13 more detail in paragraph 145. This section is intended to explain how the various parts of
14 a method fit together.
15

16 **F. ORGANIZING RELATED METHODS INTO PACKAGES**

17 48. As noted in paragraph 38, most methods in an API are organized into packages of
18 functionalities that group related methods together. As with the method names
19 themselves, these packages are logically organized into functional groupings and named
20 so as to make it easy for programmers to remember and find the functionalities they need.
21 In Java, these groupings can be packages (the highest-level grouping, typically containing
22 many classes), sub-packages, or classes (the lowest level of grouping, typically
23 containing a handful of related methods). (*See, e.g.* “The Java Platform: A White Paper,”
24 Douglas Kramer, May 1996, *available at*
25 <http://java.sun.com/docs/white/platform/javaplatform.doc1.html>, and “Package
26 Members” in The Java Language Specification, Third Edition, *available at*
27
28

1 http://java.sun.com/docs/books/jls/third_edition/html/packages.html#7.1, for discussion
2 of packages, class, and methods.)

3
4 49. As an example of this grouping, take the “sqrt” method, which calculates the square root
5 of a number. This method is typically grouped together in a library or package with
6 other, related mathematical functions, such as “sin” and “cos” (short for the trigonometric
7 functions sine and cosine). Sqrt has been grouped with other math functions since at least
8 the Algol programming language in 1968, and is still grouped together with them in Java
9 and other modern languages, such as the Python and Ruby languages.

10
11 50. Methods that do not have related functionalities are not typically grouped together — for
12 example, a method that prints text to a screen would not typically be in the same class or
13 package as “sqrt.”

14
15 51. To use a particular method, the Java programmer has to know what class, and what
16 package, the method is in. A programmer calling the square root function or method in
17 Java, for example, needs to know that the method is in the Math class, and the name of
18 the method is sqrt. Frequently this is expressed in shorthand by combining the two
19 names, so that “sqrt” becomes “Math.sqrt.” The programmer also must know the other
20 key parts of the structure — specifically that the method takes one argument (a number)
21 and returns the square root of the argument. Once the programmer knows these things
22 the underlying functionality can then easily be invoked, allowing the programmer to
23 focus on the more complex task of writing their own software.

24 25 **G. THE DISTINCTION BETWEEN AN API AND ITS IMPLEMENTATION**

26
27 52. Every API, including the Java APIs at issue in this case, exists in two forms: the method
28 declaration of the API (comprising the elements mentioned above — name, arguments,

1 and return), and the implementation of the API (which includes those three elements as
2 well as the program logic that actually performs the steps necessary to accomplish the
3 purpose of the method). The method declaration can be combined with a brief, factual
4 explanation in the documentation so that developers can have a reference for the API,
5 much like a dictionary might help someone learn and reference a language. The method
6 declaration embodies the concept of the particular API.
7

8 53. Independent of, but related to, the API's method declaration is an implementation of the
9 API — the actual underlying source code that implements the API and allows the API to
10 function. An implementation will have some portions that are similar to the
11 documentation, because both the implementation and the documentation must also
12 include the exact method declaration, including all the elements of the declaration, such
13 as the arguments and return values. However, a given API may have more than one
14 implementation, *i.e.*, the underlying program logic for the implementation of the API will
15 differ from implementation to implementation. But each of the implementations must
16 necessarily share the elements of the method declaration — the package names and
17 related method elements — in order to interoperate with each other. If these elements are
18 not present and identical in different implementations of the same API, programmers will
19 not be able to use the same names and structures when using the API, since it is these
20 elements that allow each piece of the software to speak to each other. If these names or
21 structures are changed, software that references these names will fail to function, because
22 the software will not be able to find and access the functionality it needs. To see why
23 using the same names and structures is important, it may be useful to analogize this to
24 non-programming languages. It is only by having a common vocabulary of words like
25 “truck” that people can speak to each other. If the language is changed, even slightly —
26 as it is when a speaker of American English uses “truck” while a speaker of British
27
28

1 English uses “lorry” — then confusion arises. For programmers, similar confusion would
2 occur if two different implementations of the same API used different names, arguments,
3 and return values. For software, the result would be even worse than mere confusion,
4 since computers cannot guess at what the original software meant to say. Instead, faced
5 with a similar situation, software would fail to execute altogether. Different
6 implementations of the same API must use these same elements, in order to avoid
7 confusion, inefficiency, and incompatibility.
8

9 54. Other portions of the implementation not directly governed by the method declaration of
10 the API (*i.e.*, what I have referred to as the method’s program logic) will vary between
11 different implementations, *i.e.*, the source code comprising the different implementations
12 will be different. For example, different programming languages can be used to
13 implement a particular API. In the case of Android, both the Java programming language
14 and the C programming language were used to create code to implement the APIs at
15 issue. Any two given implementations, other than the parts required for compatibility
16 (*i.e.*, the elements of the method declaration), are not likely to be identical if they are
17 written by different programmers or companies. However, because they are constrained
18 by the API and practical considerations such as programming efficiency and the
19 underlying hardware, some portions may appear similar. For example, since cars are
20 constrained by the requirements to have a steering wheel, gas and brake pedals, and four
21 tires, they will often be similar “under the hood” to the untrained eye, featuring an
22 engine, drive train, and brakes. But an expert will be able to distinguish a V8 from a V6
23 or direct fuel injection from a carburetor. Similarly, the source code that is “under the
24 hood” of an API implementation may appear similar to another implementation of the
25 same API, in large part because of practical programming constraints. At the same time,
26 implementations can occasionally look quite different if there are specific reasons for
27
28

1 such differences — the digital equivalent of choosing between a fuel-efficient but slow
2 four-cylinder (or even a hybrid) versus a hungry but powerful V8.

3
4 55. Typically, the written form of the software is captured in a specification independent
5 from any specific implementation. The specification is a written document that describes
6 the API, including the method declaration (name, arguments, and return values) as well
7 as specifying any requirements that the code must meet. Despite capturing important
8 information about the API, it would be incorrect to say that the API is the specification,
9 or vice-versa. The analogies to written words may again be useful — just as one does not
10 say that the definition of a lion in a dictionary is, in fact, a lion, so the API's specification
11 is not the API, but rather a description of the API which may then take different forms.
12

13 56. This abstraction and conceptualization of the API is what makes it possible for new
14 implementations of APIs to be built. One of the key values of an API is that when
15 improvements are made “behind the scenes,” programmers who use the API do not need
16 to know that the change has occurred; they should only notice that the program is now
17 faster, more efficient, or more error-free. This can only happen because the programmers
18 (and the software they built) used the high-level abstraction represented by the API (*e.g.*,
19 the name) and did not work directly with the concrete, underlying implementation.
20

21 57. The full Android API documentation for the “abs” method (*available at*
22 <http://developer.android.com/reference/java/lang/Math.html#abs%28int%29>) can help
23 illustrate these issues:
24
25
26
27
28

1 public static int **abs**(int i)

2 Returns the absolute value of the argument.

3 If the argument is Integer.MIN_VALUE,
4 Integer.MIN_VALUE is returned.

5 Parameters

6 i the value whose absolute value has to be
7 computed.

8 Returns

9 the argument if it is positive, otherwise the
10 negation of the argument.

11 The documentation's first line, "public static int **abs** (int i)," is the *declaration* of the
12 method. The declaration is the formal statement of a method's structure, containing the
13 method's name, the list of arguments it accepts, and the type of result it returns (the
14 "return value"). I have already shown similar text in the short version of the
15 documentation shown in paragraph 47. Here, both the input and the result are numbers
16 ("int" is short for "integer."). There is effectively only one way to say this, and the only
17 thing the programmer chooses is the name of the method ("abs") and the shorthand for
18 the argument variable name ("i"). Presumably, Oracle named this method "abs" in part
19 to increase efficiency and ease of learning for programmers who were familiar with other
20 preexisting programming languages, since this name has been used for this function in
21 many older programming languages, such as C.

22 The Android source code that implements the "abs" method documented above is:²

25 ² Available at

26 [http://android.git.kernel.org/?p=platform/libcore.git;a=blob;f=luni/src/main/java/java/lang/Math.java;h=1da0b905c4](http://android.git.kernel.org/?p=platform/libcore.git;a=blob;f=luni/src/main/java/java/lang/Math.java;h=1da0b905c4af9aaf930adf5b8b80d92193b7c462;hb=HEAD#l100)
27 af9aaf930adf5b8b80d92193b7c462;hb=HEAD#l100.

```

1     package java.lang;
2     ...
3     /**
4     * Class Math provides basic math constants and operations such as
5     * trigonometric
6     * functions, hyperbolic functions, exponential, logarithms, etc.
7     */
8     public final class Math {
9         ...
10        /**
11        * Returns the absolute value of the argument.
12        * <p>
13        * If the argument is {@code Integer.MIN_VALUE}, {@code
14        Integer.MIN_VALUE}
15        * is returned.
16        *
17        * @param i
18        *         the value whose absolute value has to be
19        *         computed.
20        * @return the argument if it is positive, otherwise the
21        *         negation of the
22        *         argument.
23        */
24
25        public static int abs(int i) {
26            return i >= 0 ? i : -i;
27        }
28

```

The first line, “package java.lang;” is the name of the package of API elements in which the abs method resides, and indicates that this file contains a class which is part of that package. “public final class Math” is also part of this organization, reflecting the class which contains the abs method. Both of these lines (which appear above in black), in this exact form, must be present in order to accurately implement the API, so all implementations of the Math.abs function will contain these two lines. (The variable name, here “i”, is not part of the definition, and so can be different between different implementations without impacting compatibility.)

- 1 58. The lines of text that begin with asterisks (which appear above in blue) are programmer's
2 *comments*. Comments do not provide functionality to the software or affect the compiled
3 code that is distributed to users; instead, they document what the code does and explain it
4 to other programmers. In this case, they describe to a programmer the function of this
5 API, and may also contain information about how to use the API. These comments, in
6 turn, are used to automatically generate the documentation for the method.
7
- 8 59. Finally, the actual source code for the method is shown here in green and red. It starts by
9 repeating the declaration of the method — “public static int abs(int i)” (in green). It then
10 presents the program logic for the method — the single line “return i >= 0 ? i : -i;” (in
11 red). This red portion is what actually tells the computer how to perform the method's
12 functionality. In this case, the program logic could be stated in English as “if the number
13 we are given is greater than or equal to 0, return that number, and otherwise return that
14 number but with the opposite sign.” Because creating the absolute value is simple, this
15 program logic is brief, but for more complicated methods many more lines of program
16 logic may be needed.
17
- 18 60. Of this substantial amount of text that constitutes the implementation and documentation
19 of the abs method, other than the required organizational lines I discuss in paragraph 57,
20 only the single line “public static int abs(int i)” (the method's name and declaration,
21 underlined above) is identical between this implementation of abs (in Android) and
22 Oracle's implementation of abs (in the works at issue). This declaration identifies the
23 method, matching the declaration in the documentation and specification. Use of the
24 same declaration is necessary if the two implementations are to be compatible, and an
25 essentially identical declaration is in fact present in any implementation of
26 `java.lang.Math`.
27
28

1 61. Besides Oracle’s open source implementation of Java (typically referred to as OpenJDK),
 2 the non-profit GNU Project has written a Java implementation called GNU Classpath,
 3 and the non-profit Apache Foundation has written a Java implementation called Apache
 4 Harmony. *See, e.g.*, GNU Classpath documents at
 5 <http://www.gnu.org/software/classpath/docs/cp-hacking.html> and Apache Harmony
 6 documents at <http://harmony.apache.org/faq.html> and
 7 <http://harmony.apache.org/subcomponents/classlibrary/compat.html>. Compatibility
 8 between these implementations is desirable for a number of reasons (discussed in more
 9 detail in paragraph 33) — primarily the benefit to software developers and consumers that
 10 results from having choice and competition between API implementation providers. For
 11 example, for compatibility and standardization reasons, the “abs” function discussed
 12 earlier has the following identical method declaration not only in Java and Android, but
 13 also in Harmony and GNU Classpath:
 14

15		
16	Java:	<code>public static int abs(int a)</code>
17	Harmony:	<code>public static int abs(int i)</code>
18	GNU Classpath:	<code>public static int abs(int i)</code>
19	Android:	<code>public static int abs(int i)</code>

20 The similarities are not limited to the abs method. Each of these projects implements the
 21 API packages at issue in this case, using the same package, class, and method names.
 22

23 **H. SUN AND ORACLE ALSO HAVE IMPLEMENTED AND DISTRIBUTED** 24 **APIs FROM OTHER SOFTWARE**

25 62. One way of seeing the distinction between API and implementation is by noting that
 26 companies such as Sun and Oracle have, in the past, implemented pre-existing APIs.
 27
 28

1 **1. Sun Implemented and Distributed APIs from Previous Generations of Spreadsheets**
2 **as Part of StarOffice and OpenOffice.org**

3
4 63. Between 1999 and 2011, Oracle, and Sun Microsystems before it, developed and
5 distributed the StarOffice and OpenOffice.org “Calc” spreadsheet software, and this
6 software implemented and distributed APIs from previous generations of spreadsheets
7 created by other companies. As I will explain in this section, the Calc spreadsheet
8 software contains an API, and this API is in large part based on the APIs originally
9 developed for older spreadsheet software, including Visicorp’s Visicalc and Microsoft’s
10 Excel spreadsheet software. The implementation of the APIs in the Calc spreadsheet
11 program allows spreadsheet models developed in Excel, for example, to also be useful
12 and run in the StarOffice or OpenOffice programs.

13
14 64. Most spreadsheet programs provides “spreadsheet functions” that enable users to write
15 small programs — called “macros” — that manipulate data in a spreadsheet cell. For
16 example, the function “ABS” calculates the absolute value of the number in a given
17 spreadsheet cell, the function “AVERAGE” calculates the average value of the numbers
18 in multiple spreadsheet cells, and “NPV” returns the net present value of an investment.
19 These functions or macros are used by people using spreadsheets to create models
20 whether these people are professional software developers, engineers, lawyers,
21 investment bankers, scientists, or hobbyists. These functions and the macros that they are
22 used by constitute an API, because they are a mechanism that allows creation of written
23 programs that communicate with the spreadsheet software’s functionality.

24
25 65. Spreadsheets created by different software companies frequently use function names and
26 argument structures from older spreadsheet programs. For example, the first column of
27 the following table shows the names of all the spreadsheet functions that were supported
28

1 in VisiCalc in 1979 — VisiCalc’s API. The other columns of the table show the function
 2 names used to operate the same functionality in Lotus 1-2-3, Microsoft Excel, and
 3 OpenOffice Calc. (The additional API elements or functionalities added in the later
 4 programs are not shown in this table.) As the chart shows, the function names originally
 5 used by VisiCorp’s VisiCalc in 1979 were then used by Microsoft Excel 2003 and
 6 Oracle’s OpenOffice.org. This shows that the original VisiCalc API of 1979 is included
 7 to this day in the Microsoft and Oracle products, with only one exception (VisiCalc’s
 8 “ERROR” function, which has been replaced by #N/A or #VALUE in Excel and Calc).
 9

VisiCalc (1979)	Microsoft Excel (2003)	Oracle OpenOffice.org Calc (Today)
@ABS	ABS	ABS
@ACOS	ACOS	ACOS
@ASIN	ASIN	ASIN
@ATAN	ATAN	ATAN
@AVERAGE	AVERAGE	AVERAGE
@COS	COS	COS
@COUNT	COUNT	COUNT
@ERROR	#N/A or #VALUE!	#N/A or #VALUE!
@EXP	EXP	EXP
@INT	INT	INT
@LN	LN	LN
@LOG10	LOG10	LOG10
@LOOKUP	LOOKUP	LOOKUP
@MAX	MAX	MAX

VisiCalc (1979)	Microsoft Excel (2003)	Oracle OpenOffice.org Calc (Today)
@MIN	MIN	MIN
@NA	NA	NA
@NPV	NPV	NPV
@PI	PI	PI
@SIN	SIN	SIN
@SQRT	SQRT	SQRT
@SUM	SUM	SUM
@TAN	TAN	TAN

66. Attached as Exhibit C is a table showing the names of the spreadsheet functions in Microsoft Excel 2003 and Oracle's most recent version of OpenOffice.org Calc, which was prepared based on the publicly available documentation available for Microsoft Office Excel 2003 at <http://office.microsoft.com/en-us/excel-help/excel-functions-by-category-HP005204211.aspx> and for OpenOffice.org Calc at http://wiki.services.openoffice.org/wiki/Documentation/How_Tos/Calc:_Functions_listed_by_category. Five rows from Exhibit C are reproduced here for discussion purposes:

Microsoft Excel (2003)	Oracle OpenOffice.org Calc (Today)
AMORLINC	AMORLINC
AND	AND
	ARABIC
AREAS	AREAS
ASC	
ASIN	ASIN

1 67. Exhibit C shows that many of the functions that constitute the API of Microsoft Excel
2 2003 were also implemented in OpenOffice.org Calc. Functions on the left are
3 implemented in Excel, and functions on the right are implemented in Calc. In this sample
4 from Exhibit A, three of the functions (AMORLINC, AREAS, and ASIN) are
5 implemented in both spreadsheets, while ASC is implemented only in Excel and
6 ARABIC is implemented only in StarOffice.
7

8 68. As shown in Exhibit C, overall, of the 340 functions implemented in the Excel 2003
9 spreadsheet function API, 324 (95%) are also implemented in StarOffice.
10

11 **2. Sun Implemented and Distributed APIs from Linux as Part of the Solaris Operating** 12 **System**

13 69. As I will explain in this section, since 1999, the Solaris operating system, developed and
14 distributed by Oracle, and Sun Microsystems before it, has contained or been delivered
15 with APIs that are based on the APIs originally developed by the developers of the Linux
16 operating system. The implementations of these APIs in Solaris facilitates the use of
17 programs developed in Linux environments to run on Solaris machines.
18

19 70. The BrandZ project, also known as Solaris Containers, was a software system that Sun
20 implemented starting in 2004. BrandZ worked with other software, called a “brand,” to
21 translate a non-Solaris operating system’s functionality into the Solaris functionality, so
22 that software written for the other operating system would run on Solaris. Essentially,
23 each brand helped “translate” communications that used the other operating system’s
24 APIs into communication with similar Solaris APIs. In particular, Sun developed a brand
25 called the “LX Brand”. The purpose of the LX Brand software was to “enable[] Linux
26 binary applications to run unmodified on Solaris” (“BrandZ WebHome,” *available at*
27
28

1 hub.opensolaris.org/bin/view/Community+Group+brandz/WebHome). In order to
2 achieve this goal, several components of the Linux API are implemented by the LX
3 Brand software, including signals, system calls, and the “/proc” interface. (See “BrandZ
4 Overview” (available at
5 <http://hub.opensolaris.org/bin/download/Community+Group+brandz/WebHome/brandzo>
6 [vreview.pdf](http://hub.opensolaris.org/bin/download/Community+Group+brandz/WebHome/brandzo).)
7

8 71. For example, the “/proc” interface allows programs to interface with the Linux operating
9 system by reading and writing the contents of files in a special directory called “/proc.”
10 Reading and writing these files allows a program to discover the status of the operating
11 system and processes running on the operating system. A process can be a program or a
12 part of a program that the user is running and it can be part of the operating system, e.g.,
13 it might facilitate communication over the Internet, with a printer, or allow one program
14 to pass data to another program. In Linux environments and many Unix environments,
15 every process has a number that identifies it, the so-called Process Identifier or PID.
16 Processes also have names — for example the process that starts up the first process for
17 the operating system has PID one, but the name ‘init’ and a process designated for
18 cryptographic programs might have the name ‘crypto’ and would certainly have a
19 different PID than the ‘init’ process. One representative element of the /proc interface,
20 known as “/proc/[pid]/status,” allows a program to communicate with the operating
21 system about the status of a particular process. To initiate the communication, the
22 program asks for the contents of the /proc/[pid]/status file — that is a file whose name is
23 ‘status’ that is located in the directory/folder corresponding to the process identifier of a
24 process, e.g., /proc/523/status is the file that gives the status of process 523. The
25 operating system responds to a request for information about a particular process by
26
27
28

1 filling the file named 'status' with text that shows the status of the process. After such a
 2 request, the first eight lines of the /proc/[pid]/status file might, for example, look like this:

3
 4 Name: [the name of the process]
 5 State: [the status of the process]
 6 Tgid: [the "Thread Group ID" of the process]
 7 Pid: [the "Thread ID" of the process]
 8 PPid: [the "Thread ID of the process's parent"]
 9 TracerPid: [The "Thread ID" of the tracing process]
 10 Uid: [ID numbers of users involved in the process]
 11 Gid: [ID numbers of groups involved in the process]

12 72. The text on the lefthand side of the file (such as "Name: ") is part of Linux's API. These
 13 are always present in /proc/[pid]/status. The text on the right is the information about the
 14 specific process, and will be different each time /proc/[pid]/status is accessed. Changes
 15 to this layout (for example, changing "Name" to "ID" or "Reference") would break
 16 applications that use this API. As a result, if another operating system wanted to be
 17 compatible with this API, it would need to print "Name:," "State:," etc., in exactly the
 18 same manner as Linux prints it.
 19

20 73. The following chart shows the values — taken directly from the respective publicly
 21 available source code — which Linux and Sun's LX Brand use to create the text on the
 22 left hand side of the /proc/[pid]/status file. In each entry in the chart, "\t" means "tab",
 23 "\n" means end of line, and the "%s" is replaced by the relevant information for the
 24 particular process, so that "Name:\t%s\n" becomes
 25

26 Name: [the name of the process]

27 when the /proc/[pid]/process file is accessed.
 28

Linux	LX Brand
"Name:\t%s\n"	"Name:\t%s\n"
"State:\t%s\n"	"State:\t%s\n"
"Tgid:\t%d\n"	"Tgid:\t%d\n"
"Pid:\t%d\n"	"Pid:\t%d\n"
"PPid:\t%d\n"	"PPid:\t%d\n"
"TracerPid:\t%d\n"	"TracerPid:\t%d\n"
"Uid:\t%d\t%d\t%d\t%d\n"	"Uid:\t%u\t%u\t%u\t%u\n"
"Gid:\t%d\t%d\t%d\t%d\n"	"Gid:\t%u\t%u\t%u\t%u\n"
"FDSize:\t%d\n"	"FDSize:\t%d\n"
"Groups:\t"	"Groups:\t"

74. Each line of the chart is identical, and this demonstrates that the output of the /proc/[pid]/status API is the same between Linux and the LX Brand software, and therefore that (in this respect, at least) Linux and the LX Brand software are compatible. If these lines were different, then the resulting file would be different, and the LX Brand software would not be compatible with Linux. Other elements of the /proc interface are similarly implemented in the LX Brand software.

75. The LX Brand software also reimplements Linux kernel system calls. System calls are a part of an operating system's API; they allow programs written by users to access resources managed by the operating system, e.g., to read and write files to kill processes, or to allocate memory to use in a program. These resources are managed by the operating system, but programs written by users to run on the operating system need access to some of the resources to be able to run properly or at all. The LX Brand software provides an emulation function which translates the Linux system call to an

1 equivalent Solaris operating system call, for each of 317 Linux system calls. A relevant
2 fragment of the publicly available source code that performs this translation, listing each
3 Linux system call, and the LX Brand function that implements the system call, is attached
4 as Exhibit D. This source code file indicates that there were 178 Linux system calls that
5 were implemented as part of LX Brand (the other 139 system calls were either not
6 supported or were able to directly use the equivalent Solaris system calls without
7 translation). Each of the LX Brand implementations of the Linux system calls use the
8 same name as the relevant Linux system call, with “lx_” prepended to distinguish them.
9

10 76. For example, the Linux system call “futex” was introduced to Linux beginning in 2002,
11 and Solaris does not have a “futex” system call. In order to provide compatibility for
12 Linux software running on the LX Brand, the LX Brand software provides an
13 implementation of futex called lx_futex, which has essentially the same name as futex,
14 takes similar arguments, and behaves similarly. The actual program logic that
15 implements the LX Brand lx_futex function and the Linux system call are not similar,
16 suggesting that they were independently created.
17

18 77. The 177 other system calls implemented by the LX Brand follow the same pattern: the
19 Linux system call name, plus the lx_ prefix, is used to identify a function that takes
20 similar arguments and behaves similarly to the Linux system call for which the function
21 is named.
22

23 78. Sun’s “LX Brand” software implements only a subset of the Linux operating system API,
24 and so is not completely compatible with Linux. Sun’s overview presentation states that
25 it “support[s] a subset” of the /proc API and the “minimum needed” devices (“BrandZ
26 Overview” at 23 and 24) and the design document notes that the “CLONE_PARENT”
27 argument to the clone(2) system call is also only partially implemented (see “BrandZ
28

1 Design Doc”, section 3.5.1 (“Linux Threading”), available at
2 <http://hub.opensolaris.org/bin/view/Community+Group+brandz/design>). This partial
3 implementation still aids compatibility and programmer efficiency, because it is still
4 better for the programmer to use some of the APIs than to have to completely rewrite the
5 software to use new APIs.

6
7 79. Solaris has also incorporated specific APIs from the Linux C Library (“glibc”) into the
8 Solaris C Library. For example, the “uucopy()” system call, according to Sun’s BrandZ
9 Design Doc, “seems to be generically useful, so the uucopy() will be implemented in
10 [Solaris] libc” and, in fact, Solaris gained an implementation of the uucopy system call in
11 2006, shortly after BrandZ was introduced (*see* Solaris’s `common/syscall/uucopy.c`).

12
13 **3. ORACLE IMPLEMENTED AND DISTRIBUTED APIs FROM IBM AS PART OF**
14 **THE ORACLE DATABASE SERVER**

15
16 80. As I will explain in this section, the Oracle Database server distributed by Oracle since
17 1979 contains an implementation of the API originally developed by IBM for the
18 “System R” database.

19
20 81. The System R database’s SQL API was first described in an academic paper published by
21 IBM employees in 1974 (“SEQUEL: A Structured English Query Language,” DD
22 Chamberlin, et al.), and elaborated in a subsequent paper published in 1976.

23 82. The 1974 SEQUEL paper defined the following API elements or functionalities:

24
25 SELECT FROM

26 WHERE

27 GROUP BY

28

1 SUM
2 COUNT
3 AVG
4 MAX
5 MIN

6 83. IBM supplemented the functions in the 1974 paper in a subsequent paper published in
7 1976 (“System R: relational approach to database management,” M. M. Astrahan et al.),
8 adding several new elements or functionalities to the SQL API: HAVING, ORDER BY,
9 CURSOR, INSERT INTO, and DELETE.

10
11 84. Each of the API elements or functionalities referenced in paragraphs 82 and 83, and
12 defined in the 1974 and 1976 papers, were implemented by Oracle in 1979 and are still
13 present in current releases of the Oracle Database server. Because these API elements
14 are implemented in the Oracle Database server, a command using the API elements
15 “SELECT FROM ... WHERE ...” would also be able to operate, with minimal changes,
16 with current Oracle Database servers, as it did with the original IBM System R software
17 (*see* “Oracle SQL: The Essential Reference,” David C. Kreines (2000), Chapter 1,
18 “Elements of SQL,” *available at* <http://oreilly.com/catalog/orsqlter/chapter/ch01.html>).

19
20 85. For example, because the Oracle system implements the API elements or functionalities
21 defined in the 1974 paper, it will still execute commands written using the 1974
22 SEQUEL API. The 1974 paper gives this short command that uses elements defined in
23 the 1974 paper:

24
25 **SELECT** NAME
26 **FROM** EMP
27 **WHERE** SAL

28

```
1          SELECT    SAL
2          FROM      EMP
3          WHERE     NAME = B1.MGR ;
4
```

5 86. Because this command was written using API elements (in bold) originally defined by
6 IBM but later implemented in the Oracle Database server, this command should still
7 function in a modern Oracle Database server, and indeed, some sources report that this
8 exact command was used in early demonstrations of the Oracle database (*see* “Oracle
9 SQL: The Essential Reference,” David C. Kreines (2000), p. xiv and Chapter 1,
10 “Elements of SQL”).
11

12 **I. BASIC EXAMPLE OF JAVA METHOD USAGE**

13
14 87. When a programmer is writing an application, and wants to use a particular functionality,
15 he must invoke the functionality by using the appropriate method. If a programmer
16 writing in the Java programming language wants to use Java’s square root functionality
17 to find the square root of 25, he would do that by incorporating the following language in
18 his program:

```
19         double result = Math.sqrt(25.0);
20
```

21 88. The *argument* 25.0 is passed to the method Math.sqrt when the method is called, and
22 “5.0” is *returned* by the method. In this example, the return value is then stored in the
23 variable named “result” for use elsewhere in the program.
24

25 89. To write this example, a programmer who had never previously used Java would likely
26 have started by guessing that square roots were in the class “Math,” looking at that class’s
27 documentation, finding the familiar “sqrt” method, and then reading the documentation
28

1 for that method to understand what result is returned and what special cases need to be
2 considered in writing code. He would then write the fragment of code above, and in the
3 future, having learned to use this part of the API, he would not likely have to refer to the
4 documentation again.

5
6 90. Note that at no point does the programmer need to know how the program logic that is
7 invoked by the sqrt method actually calculates the square root — it could use a Newton-
8 Raphson method, logarithms, or another mathematical algorithm for calculating the
9 square root. As previously noted in paragraph 34, these details are “encapsulated” —
10 hidden behind the scenes. This focus on knowing and understanding the API name and
11 functionality, rather than understanding how the method’s underlying program logic
12 works behind the scenes, allows programmers to work more efficiently.

13
14 **J. THE APIS AT ISSUE ARE METHODS OF OPERATION**

15
16 91. I understand that section 102(b) of the Copyright Act states, “In no case does copyright
17 protection for an original work of authorship extend to any idea, procedure, process,
18 system, method of operation, concept, principle, or discovery, regardless of the form in
19 which it is described, explained, illustrated, or embodied in such work.” I also
20 understand that a method of operation has been described by the First Circuit as “the
21 means by which a person operates something, whether it be a car, a food processor, or a
22 computer.” I further understand that the Ninth Circuit, citing section 102(b) of the
23 Copyright Act, has stated that the functional requirements for compatibility are not
24 protectable. Under either of those definitions, as I will explain in more detail below, it is
25 my opinion that the Java API specifications at issue in this case are methods of operation.
26
27
28

1 92. As previously mentioned, in some computer languages, methods are referred to as
2 “functions.” Both the terms function and method suggest — correctly — that functions
3 and methods are literally a functional way to operate software. Once the method that is
4 part of the API is called and the right parameters are passed to it, the API invokes
5 functionality provided by the underlying software system. This “operates” the underlying
6 software system to create the return value, just as use of the car steering wheel makes the
7 car (through the steering system) take action to steer the car, the “plus” button makes a
8 calculator add two numbers, and the “print” command makes the operating system print a
9 paper copy of a document.
10

11 93. That an API is a functional method of operation is implicit in the definition of an API:
12 the entire purpose of an API is to allow one program to “interface” with another
13 “application.” This interfacing is not a social or creative chat, but a formal, functional
14 command from one program to another: “Do this thing for me, and report back when you
15 are done.” The program in command is using the API to operate the underlying program;
16 and the underlying program, likewise, is being operated by means of the API. In fact, it
17 is typically difficult, if not impossible, to operate the underlying system in any way
18 except through an API.
19

20
21 94. As demonstrated above in the example of Math.sqrt, using the name is necessary to
22 invoke the underlying functionality. It is also the *only* way to invoke the underlying
23 functionality — one cannot, for example, change “sqrt” to “square_root” in the example
24 above and still expect the code to work. Nor could one change the number or type of
25 arguments. Programs, unlike the human operators of our calculator and car analogies, are
26 not flexible — they must be fed precise information in order to operate.
27
28

- 1 95. Like “Ctl+P”, the print icon, or the steering wheel, APIs provide mechanisms that operate
2 underlying functionality, causing the libraries at issue to perform activities and return the
3 information requested by the programmer — that is to say, by the “operator” of the
4 software in question.
5
- 6 96. As an example of how APIs in Java are methods of operation, I will consider three
7 classes in the java.util package that allow developers, and the software they write, to
8 operate on dates and times: Date, Calendar and TimeZone. As the names convey, these
9 classes are used by programmers to create, manipulate, and use calendars and dates in
10 Java programs. As software is increasingly deployed throughout the world, it is
11 important that developers be able to simply create and manipulate dates and times in a
12 way that works across all cultures and time zones. These Java classes provide such an
13 API. These classes are used together, and the methods in these classes mirror the
14 functionality and operation non-programmers would expect if you needed to create and
15 use dates and times. For example, to use dates and times, first a programmer has to
16 record them. A Date allows that by representing a specific instant in time. While one
17 might expect this would be something like January 22, 2009, a Date actually represents a
18 specific millisecond on a specific day of a month of a year, and then provides methods
19 that translate that millisecond into a particular date, automatically translating (if
20 necessary) into other calendars (like the Chinese or Hebrew calendars). Having provided
21 a way to store the date, a programmer manipulating dates would likely want to be able to
22 perform a variety of actions on the date, and it is these actions that are most clearly
23 methods of operation. For example, a programmer might want to know whether one Date
24 comes “before” another, so that they could sort a list of files by time-of-modification, or
25 display a list of songs arranged by date of recording. Not surprisingly, the Date class
26 provides a method to test if one date comes before another called “before.” The code to
27
28

1 execute that test and determine if a Date A comes before a Date B is written as
2 “A.before(B).” This method executes the test and returns “true” if A is chronologically
3 before B, and returns “false” otherwise. The Date class also has methods to determine if
4 a date comes after another date and if two dates are equal. The names of these methods,
5 respectively, are after and equals, providing a clear example of how the form/name of the
6 methods follow their function. The Date class also provides a getTime() method that
7 returns the exact time in milliseconds. The Calendar class contains methods used to
8 create a calendar, e.g., for a specific time zone, year, and/or a specific international
9 location. The Calendar class has methods that allow the programmer to determine the
10 first day of the week, which is SUNDAY in the United States, but MONDAY in France,
11 for example. This method is Calendar.getFirstDayOfWeek() and it returns a value such
12 as Calendar.MONDAY or Calendar.SUNDAY — two values of the calendar class that
13 are clearly functional in representing days of the week. These classes are typically used
14 together with the TimeZone class, which provides convenient methods for creating and
15 using the timezones that occur in the world. For example, the method
16
17
18
19
20
21
22
23
24
25
26
27
28

97. Methods can also invoke more concrete functionality. For example, the class java.io.file gives programmers the ability to do operations on a file. This includes useful methods like “createNewFile” (which creates a file), “getName” (which gets the name of the file), and “delete” (which deletes the file). In each case, the programmer invokes the underlying functionality — such as creating or deleting a file — by using the name of the method. Once a file has been opened, the methods in the java.io class can be used to read and write the file; for example, use of the method name “readLine” will invoke software that reads a line from the file.

1 98. In each of these cases, whether operating on simple data like a date, or more complex
2 things like a file or web page, the method names are the way in which the underlying
3 functionality is invoked, providing information and taking actions as needed by the
4 program. The method name is, quite literally, the method of operation, just as the gas
5 pedal is the way that an engine is invoked.
6

7 **K. THE JAVA API PACKAGE NAMES ARE DICTATED BY FUNCTION**

8
9 99. It is my understanding that names are not entitled to copyright protection. However, even
10 if they were, it is my opinion that the Java API package names are short, fragmentary,
11 and functional. It is my understanding that short, fragmentary names that are dictated by
12 function are not protectable under copyright law.

13 100. It is my understanding that the following API packages (and, in some cases, certain
14 subpackages of these packages) are at issue in the case. In each case, the name of the
15 package and the basic organization of the classes and methods within each package are
16 merely descriptive of the functionalities in those packages.
17

18 **java.lang**

19
20 The java.lang package and its subpackages java.lang.ref, java.lang.reflect, and
21 java.lang.annotation are part of a group of classes that facilitate interacting with and
22 programming related to the Java language. The package name (“lang”) and contents of
23 the classes and methods in this package reflects this emphasis on the core Java language.
24

25 **java.math**

26 The java.math package provides the programmer with access to classes that facilitate
27 arbitrary precision arithmetic with integers, *e.g.*, integers with no upper or lower limit.
28

1 The package name (“math”) and contents of the classes in this package reflect this
2 underlying functionality.

3 **java.net**

4
5 The java.net package, and its extension javax.net and subpackage javax.net.ssl, provide
6 classes for the programmer to implement network connections at both a low- and high-
7 level. The package name (“net”) is short for “network” and therefore reflects this
8 underlying functionality.

9 **java.io and java.nio**

10
11 The java.io and java.nio packages group together classes for dealing with input and
12 output. Input and output are called “I/O” in long-standing programmer jargon, explaining
13 the name of the io package. (The nio package is so-named because it was an attempt to
14 present a “new” io (nio) package.) The java.nio hierarchy of classes also contains the
15 subpackages java.nio.channels, java.nio.channels.spi, java.nio.charset
16 java.nio.charset.spi. The nio package provides classes that facilitate more efficient
17 (faster) input and output. The nio package are designed to interact with each particular
18 operating system’s efficient I/O mechanisms, so that the Java programmer can use the nio
19 classes knowing that they will likely be faster than the java.io classes that were not
20 originally designed for efficiency.

21 **java.security and javax.security**

22
23 The java.security package, its subpackages java.security.acl, java.security.cert,
24 java.security.interfaces and java.security.spec, and its extensions, javax.security.auth,
25 javax.security.auth.callback, javax.security.auth.login, javax.security.auth.x500,
26 javax.security.cert, provide classes and functionality related to security, as the names
27
28

1 suggest.

2
3 **java.sql**

4 The java.sql package, and its extension, javax.sql, facilitates interacting with relational
5 databases or data in a format similar to that defined in such a database. These packages
6 are based on the “SQL” standard (Structured Query Language), a standard named and
7 defined in the late 1970s and early 1980s, and the name of the packages (“sql”) reflects
8 the name of this standard.
9

10 **java.text**

11 The java.text package facilitates writing software to handle text, dates, numbers, and
12 messages in a format that is independent of a particular natural language — allowing
13 programmers to cope more easily with languages other than English.
14

15 **java.util**

16 The java.util packages provides utilities and collection classes. Its subpackages
17 java.util.logging, java.util.jar, java.util.prefs, java.util.regex, and java.util.zip provide
18 utilities that are more specialized, *e.g.*, to deal with logging, archives (jar files), user
19 preferences, regular expressions, and zipped or compressed files, respectively. These
20 functionalities are diverse, but “utilities” are a traditional name for small, single-purpose
21 tools in the computing world, and so grouping these together under the name “util” is a
22 straightforward mapping of functionality to traditional naming.
23
24

25 **javax.crypto**

26 The javax.crypto package and its subpackages javax.crypto.interfaces and
27 javax.crypto.spec provide classes to write code that adheres to cryptographic protocols.
28

1 “Crypto” is common programmer shorthand for “cryptography” and so makes for a
2 natural mapping of name to functionality.

3
4 **javax.xml**

5 The package javax.xml and its subpackages javax.xml.datatype, javax.xml.namespace,
6 javax.xml.parsers, javax.xml.transform, javax.xml.transform.dom,
7 javax.xml.transform.sax, javax.xml.transform.stream, javax.xml.validation, and
8 javax.xml.xpath deal with XML—eXtensible Markup Language.
9

10 In some cases, a given package may require the functionality of another package in order to
11 function correctly, much like the upper floors of a building need the lower floors of the building
12 to remain standing. The final two packages listed below, while not themselves basic to the
13 functionality of modern operating systems, must be present in order for the previously listed
14 packages to operate correctly and provide their complete, intended functionality to users:
15

16 **java.awt.font**

17 The package java.awt.font allows programmers to interact with low-level font
18 information.
19

20 **java.beans**

21 The java.beans package facilitates software interaction with JavaBeans — traditionally
22 viewed as a reusable software component conforming to specific conventions so that the
23 component can be manipulated with visual and graphical tools.
24

25 101. Because these names all describe specific functionalities, and (as will be discussed in
26 paragraph 113) they are limited by design rules to short, fragmentary words and phrases,
27 there is no meaningful creativity in the package names.
28

1 **L. THE JAVA API CLASS AND METHOD NAMES ARE DICTATED BY**
2 **FUNCTION**

3
4 102. In the Android packages at issue, there are 472 public classes, 150 public abstract classes,
5 and 176 public interfaces. A public class is a class that is accessible to programmers who
6 are using the API; a private class is internal to the library and can only be used by other
7 parts of the library.

8
9 103. API element names, such as class names, must be factually descriptive of the underlying
10 functionality so that programmers can recognize, understand, and remember them when
11 reading and writing a program. While this is formally enshrined in the Java Language
12 Specification (see discussion in paragraph 113), the rule has more pragmatic roots that
13 date back to the earliest computer languages. The core reason that API component names
14 are short and reflect underlying functionality is that inventive and creative names only
15 loosely tied to the functionality would be difficult for programmers to remember.

16
17 104. For example, the method “sqrt” is short, simple, and memorable for programmers — and
18 possibly even for non-programmers; a reader of this report may not need to be reminded
19 more than a few times that “sqrt” means “square root.” It is technically possible to
20 instead call the square root method “Steve,” just as it would be possible to build a
21 calculator whose buttons use colors instead of numbers and mathematical symbols. But
22 such a calculator would be difficult to use; a user would have to memorize the colors and
23 their mapping to the underlying numbers and symbols. That would take some time and
24 effort — so much time and effort, in fact, that users are likely to use a traditional
25 calculator instead. Similarly, when the name of a method does not reflect the underlying
26 functionality (as in the case where a square root method is called “Steve”) the method
27

1 would become difficult to learn and remember. As a result, in practice all API element
2 names are simple and factually descriptive.

3
4 105. A method whose underlying functionality is to test to see if two file names are equivalent,
5 for example, could be called “equals” or “equivalent” but likely not much else. Even the
6 longest API element names, such as SQLNonTransientConnectionException, are still tied
7 to the underlying functionality. In that case, the name has three parts which demonstrate
8 the underlying functionality: the word “SQL” reflects that this relates to the SQL
9 database language (a language that pre-dates Java, and was not created by Sun or Oracle),
10 and the word “exception” reflects that this relates to an “exception” (similar to an error
11 message). Both of these terms have been used in software programming for over 30
12 years, predating Java by some time. A programmer would recognize that the third part of
13 the name (“Non-Transient Connection”) reflects underlying SQL functionality that is a
14 term of art from outside the Java language. Since the names of all of these underlying
15 concepts are fixed, or nearly so, the name of the method reflecting these concepts is also
16 necessarily inflexible. The fact that the name reflects the underlying functionality is not
17 merely convenient — it is practically required to allow the system to be comprehensible
18 to programmers.
19

20
21 106. Class names, like the method names discussed above, are highly functional, in many
22 cases showing only small variations directly related to the class functionality. For
23 example, consider the seven classes whose names end in Event as shown below. These
24 seven classes come from three different packages.
25
26
27
28

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28

Class Name	From Package	What is the Functionality?	Why is it in this Package?
HandshakeCompletedEvent	javax/net/ssl	A class describing an event that takes place once a Handshake is Completed.	“Handshakes” are part of the Secure Sockets Layer (SSL) networking protocol, and so this is grouped with other “net” and “ssl” features.
PreferenceChangeEvent	java/util/prefs	A class describing an event that takes place when a Preference Changes.	Utilities that user track Preferences must have a way to track what happens when the preferences change, and so this is grouped with “util” “prefs.”
NodeChangeEvent	java/util/prefs	A class describing an event that takes place when a Preference Node Changes.	“Nodes” are a common way to organize data. Since these nodes are used to store user preference data, information about the nodes (including

Class Name	From Package	What is the Functionality?	Why is it in this Package?
			changes to them) are grouped with other “prefs”-related functionality.
SSLSessionBindingEvent	javax/net/ssl	A class describing an event that takes place when an SSL Session Binds.	As part of the implementation of the SSL networking protocol, this is grouped with other “net” and “ssl” classes.
ConnectionEvent	javax/sql	A class describing an event that takes place when a Connection occurs.	Because this is an SQL Connection, it is grouped with other SQL methods.
RowSetEvent	javax/sql	A class describing an event that takes place when an SQL RowSet is changed.	Because this is an SQL Rowset, it is grouped with other SQL methods.
StatementEvent	javax/sql	A class describing an event that takes place when an SQL	Because this is an SQL Statement, it is grouped with other SQL

Class Name	From Package	What is the Functionality?	Why is it in this Package?
		Statement is changed.	methods.

These names are functional in specifying the purpose of the methods. The noun part of each name that precedes Event describes the event, but there is essentially no creativity in choosing the noun. For example, the HandShakeCompleteEvent describes an event that takes place after the hand-shaking protocol in making an SSL connection has been completed. Similarly, the RowSetEvent class describes an event that takes place when a “rowset” is changed in an SQL database. The names have been chosen not because of deep introspection or creativity on the part of the author, but by simply describing what functionality is contained in the class.

107. Another example of class names that conform to simple rules describing the underlying functionality are the 18 classes whose names end with InputStream and 15 that end with OutputStream. These classes are part of a variety of packages — some are grouped with other input and output functions in java.io and java.nio, but several are part of the packages java.util, java.util.zip, java.util.jar, java.security and javax.crypto. The table below shows thirteen of the InputStream classes below and their corresponding packages. Again the names for the classes are functional and limited by the responsibilities of each class: the LineNumberInputStream class reads data while keeping track of line numbers while the CipherInputStream uses a cryptographic cipher for reading data.

Class Name	From Package	What is the Functionality?	Why is it in this Package?
FileInputStream	java/io	A stream of inputs from a file.	This is part of the basic input/output (“io”) functionality.
PushbackInputStream	java/io	Adds the ability to push data back into an input stream.	This is part of the basic input/output (“io”) functionality.
ZipInputStream	java/util/zip	A stream of inputs from a zipped file.	Reading from a zipped file is part of the basic zip file (“zip”) functionality.
JarInputStream	java/util/jar	A stream of inputs from a “jar” file.	Reading from a jar file is part of the basic jar file (“jar”) functionality.
LineNumberInputStrea m	java/io	Adds the ability to count the line number to an input stream.	This is part of the basic input/output (“io”) functionality.
StringBufferInputStrea m	java/io	A stream of inputs from a string buffer.	This is part of the basic input/output

Class Name	From Package	What is the Functionality?	Why is it in this Package?
			("io") functionality.
CipherInputStream	javax/crypto	A stream of inputs that have been passed through a cipher.	Ciphers are part of cryptography, and so functionality to use ciphers is part of the cryptography ("crypto") package.
InflaterInputStream	java/util/zip	A stream of inputs from an "inflater" that "inflates" a zipped file.	Inflaters are part of "zipping" a file and so this is part of util/zip.
FilterInputStream	java/io	Adds the ability to filter to an input stream.	This is part of the basic input/output ("io") functionality.
ObjectInputStream	java/io	A stream of inputs from an object.	This is part of the basic input/output ("io") functionality.
DigestInputStream	java/security	Creates a "messages digest" of a stream of inputs.	Message digests are part of certain security routines, so

Class Name	From Package	What is the Functionality?	Why is it in this Package?
			this class is part of the java/security packages.
ByteArrayInputStream	java/io	A stream of inputs from a byte array.	This is part of the basic input/output (“io”) functionality.

108. Java class names are short, fragmentary words and phrases. This is true of other Java API elements, like method names, as well. Most Java API element names are short word phrases, typically of 1-3 words in length. (See paragraph 115 for more detailed statistics on method name length). One example is the class “PrintStream,” which (not surprisingly) adds printing functionality to output “streams.” Elements in the PrintStream class include the method named “append,” which appends the argument to the output stream, the method named “print,” which prints the stream, and the method named “close,” which closes the stream. In fact, every method in the PrintStream class, with only four exceptions, is one word. Two of the exceptions are two words — “setError” and “checkError,” which, as one would expect, set and check the error state of the output stream. The other exceptions are “printf” and “println” — abbreviations for “print formatted” and “print line.” Besides being brief and fragmentary, these abbreviations have been in use by programming languages since the 1970s, in Algol and C.

109. Other classes, such as the SecurityManager class, have slightly longer names. In this class, three-word phrases (such as “checkPackageDefinition”) are predominant and there

1 are some four-word phrases (such as “checkCreateClassLoader”). But even here, the
 2 naming follows a consistent pattern — 30 of the 40 methods are named
 3 check[SomeProperty], consistently describing their underlying functionality, which is to
 4 check the status of the property referred to by the method name. For example,
 5 “checkCreateClassLoader” checks to see if it is possible to create a new class loader.
 6

7 110. Because many classes need the same functionality, and the names of the methods in
 8 question are dictated by functionality or by rules (see next section), it is not surprising
 9 that many of the names are repeated. The most common names in Oracle’s
 10 implementation of Java 1.5 are:
 11

Method name	Number of Times Repeated	Functionality?
toString	194	Converts an object to a String.
equals	157	Tests to see if two objects are equal.
hashCode	147	Creates a “Hash Code” (a numeric representation) of a class.
run	139	Runs the code in the object.
read	96	Reads (typically to a stream of characters).
write	94	Writes (typically to a stream of characters).
remove	88	Removes something (exactly what is removed depends on the class).
get	74	Gets the value of an object.
close	72	Closes a stream.
size	68	Returns the number of items in a collection of items.
clear	61	Clears the content of the thing referenced.
clone	59	Clones the thing referenced.
TOTAL	1249	These 10 method names are used by roughly 1/6 of the methods in Oracle’s

		implementation of Java 1.5.
--	--	------------------------------------

111. The organization of the methods into classes, like the organization of classes into package, is driven by functionality and the requirement that programmers be able to efficiently find and use these methods. This is why, for example, all of the math functions listed in the table in paragraph 39 are in the same `java.lang.Math` class.
112. Because these names all describe specific functionalities, limited by design rules to short, fragmentary words and phrases, there is no meaningful creativity in the class or method names.

M. THE JAVA NAMES ARE THE PRODUCT OF MECHANICAL RULES

113. Java API element names frequently repeat certain key terms and patterns, following mechanical rules laid out in the Java Language Specification and elaborated over time by practice. The rules provide suggestions for structure and naming, stating, for example, that “[m]ethod names should be verbs or verb phrases, in mixed case, with the first letter lowercase and the first letter of any subsequent words capitalized.” Similarly, names of class types are to be “descriptive nouns or noun phrases.” Java Language Specification, First Edition, Section 6.8 “Naming Conventions,” *available at* java.sun.com/docs/books/jls/first_edition/html/6.doc.html#11186. Both of these rules are followed by all the examples shown in this report, except for those methods that are drawn from older programming languages (like “`sqrt`”).
114. Additional word patterns crop up repeatedly throughout the Java APIs. “`InputStream`” and “`ChangeEvent`,” cited above, are two examples affecting a few dozen names, but others go much further. For example, the Java Language Specification rules for method

1 names state that methods that return the value of a variable should start with “get,” and
2 method names that set the value of a variable should start with “set.” Other rules require
3 specific methods to be in many classes, such as “hashCode” and “toString.” In Oracle’s
4 implementation of Java 1.5, nearly one-third of the method names at issue (2,578 of the
5 7,796 methods) are determined by these rules, including roughly 2,000 that begin with
6 either “get” or “set,” and 164 named simply “equals.” Testing whether one thing is
7 equal to another thing is an extremely common operation for programmers, and so it
8 makes sense that many different methods for testing equality would exist. At the same
9 time, it makes sense to make sure that the operation has the same name everywhere — it
10 would unnecessarily complicate the learning process if it were “equals” in one place,
11 “sameAs” in another place, and so forth. These constraints yield the resulting name
12 (“equals”) — which is wholly functional, dictated by efficiency constraints and not
13 creativity.
14

15
16 115. An additional 2,347 method names were single words, like “run” or “add.” The
17 remaining 2,871 methods are not long or complicated — they are, on average, only 2.344
18 words “long” (counting a method name like locateURL as two words and findBestMatch
19 as three words). In Android, of the 9297 total methods, 3220 are unique methods, 2676
20 or 28.8% are one word names, 2909 are required names (like the “get” and “set”
21 examples above), leaving 3,712 other methods whose average word length is 2.41.
22

23 116. Following these mechanical rules and seeking to create consistency reduce the amount of
24 creativity and work necessary to write the API, and, more importantly, reduces the
25 amount of work necessary to learn and memorize the API. As a result, any good API
26 design will have naming rules like the Java API naming rules contained in the Java
27
28

1 Language Specification, which result in names that are functional and primarily dictated
2 by efficiency constraints.

3
4 117. It is also my opinion that the rules and naming conventions imposed by the Java
5 Language Specification further constrains any creativity associated with the API names.

6 **N. THE ORGANIZATION OF API ELEMENTS IS DICTATED BY**
7 **FUNCTION AND DOES NOT REFLECT CREATIVITY**
8

9 118. The same restrictions that apply to naming also apply to the organization of methods.
10 Just as the name must be tied to functionality so that the API is easy to find and
11 remember, the organization into related groupings must also reflect underlying
12 functionality so that programmers can discover and use the elements efficiently. For
13 example, methods related to security, such as `AccessController.checkPermission` and
14 `Signature.sign`, are most sensibly organized into packages primarily related to security —
15 `java.security` and `javax.security`.
16

17 119. The practical requirement that all API element names and package organizations be
18 related to the underlying functionality restricts the packages in which any class can be
19 placed, and restricts the classes in which any method can be placed. The fourth column
20 in the tables in paragraphs 106, and 107 show how and why various classes fit into their
21 respective packages, and various methods fit into their respective classes, which can help
22 demonstrate this. For example, `ZipInputStream` is a class which allows creation of an
23 “input stream” which can read from a compressed file (known as a .zip file). Because it
24 relates to .zip files, developers will look for it with other classes and methods related to
25 .zip files, and so it is grouped with them in an appropriately named subpackage called
26 `java.util.zip`.
27
28

1 **O. THE SAME ANALYSIS APPLIES TO THE NATIVE FILES**

2
3 120. It is my understanding that Oracle has alleged that a number of files written in the C
4 programming language are at issue. These files are part of the implementation of the
5 Java API. For example, the file name `java_lang_reflect_Array.c` (which I understand is
6 one of the files at issue) reflects the name of the `java.lang.reflect.Array` class. Similarly,
7 the C function names within the file `java_lang_reflect_Array.c` reflect method names
8 within the `java.lang.reflect.Array` class. As with the implementation files in the Java
9 language discussed above, these files written in the C language must use these names and
10 organizing principles in order to implement the API in a compatible and interoperable
11 manner. If Android could not use these names in these files, Android could not create a
12 compatible, efficient implementation of these APIs.

13
14 121. For example, the file `java_lang_Class.c` has these 29 C language functions:

15 `Dalvik_java_lang_Class_desiredAssertionStatus`
16 `Dalvik_java_lang_Class_classForName`
17 `Dalvik_java_lang_Class_getClassLoader`
18 `Dalvik_java_lang_Class_getComponentType`
19 `Dalvik_java_lang_Class_getDeclaredClasses`
20 `Dalvik_java_lang_Class_getDeclaredConstructors`
21 `Dalvik_java_lang_Class_getDeclaredFields`
22 `Dalvik_java_lang_Class_getDeclaredMethods`
23 `Dalvik_java_lang_Class_getInterfaces`
24 `Dalvik_java_lang_Class_getModifiers`
25 `Dalvik_java_lang_Class_getNameNative`
26 `Dalvik_java_lang_Class_getSuperclass`
27 `Dalvik_java_lang_Class_isAssignableFrom`
28 `Dalvik_java_lang_Class_isInstance`

1 Dalvik_java_lang_Class_isInterface
2 Dalvik_java_lang_Class_isPrimitive
3 Dalvik_java_lang_Class_newInstance
4 Dalvik_java_lang_Class_getSignatureAnnotation
5 Dalvik_java_lang_Class_getDeclaringClass
6 Dalvik_java_lang_Class_getEnclosingClass
7 Dalvik_java_lang_Class_getEnclosingConstructor
8 Dalvik_java_lang_Class_getEnclosingMethod
9 Dalvik_java_lang_Class_getGenericInterfaces
10 Dalvik_java_lang_Class_getGenericSuperclass
11 Dalvik_java_lang_Class_getTypeParameters
12 Dalvik_java_lang_Class_isAnonymousClass
13 Dalvik_java_lang_Class_getDeclaredAnnotations
14 Dalvik_java_lang_Class_getInnerClassName
15 Dalvik_java_lang_Class_setAccessibleNoCheck

16 Each of these functions corresponds exactly to a so-called native method in the class
17 java.lang.Class. For example the C language function

18 Dalvik_java_lang_Class_isAnonymousClass corresponds to the method
19 AnonymousClass in the class java.lang.Class; the C function

20 Dalvik_java_lang_Class_setAccessibleNoCheck corresponds to the method
21 setAccessibleNoCheck in the class java.lang.Class, and so on for each function in the file.

- 22 122. The other .c files share similar characteristics to their corresponding .java files. The
23 functions in the .c files typically correspond exactly to a corresponding public method in
24 the related .java file or to a private method used to implement the private method. For
25 example, in the file java_lang_Runtime.c ,the function
26 Dalvik_java_lang_Runtime_nativeLoad corresponds to the private method nativeLoad in
27 java.lang.Runtime which is used in the implementation of the public method
28

1 java.lang.Runtime.load. In my opinion, there is no expressiveness in the names used in
2 the C files, because they are directly derived from the functional names in the .java files
3 (as explained above), and are required for efficient implementation of those files.
4

5 **P. MANY API ELEMENTS ARE DRAWN FROM THE PUBLIC DOMAIN**
6 **AND ARE NOT ORIGINAL TO JAVA**

7 123. It is my understanding that names are not entitled to copyright protection. However, even
8 if they were, it is my opinion that many of the Java names are drawn from the public
9 domain.
10

11 124. Java, like many other programming languages, is based on features of previous well-
12 known languages, such as C and C++, including their grammar and syntax. *See, e.g.*,
13 http://java.sun.com/docs/books/jls/first_edition/html/1.doc.html (“the lexical structure of
14 Java . . . is based on C and C++”); *see also*
15 http://www.gotw.ca/publications/c_family_interview.htm (James Gosling, inventor of
16 Java, quoted as saying “You can go through everything in Java and say ‘this came from
17 there, and this came from there’”). Reuse of grammar and syntax from already-familiar
18 languages allowed developers to leverage their existing knowledge and more quickly
19 adopt Java. Similarly, authors of new programming languages often use old method and
20 class names when appropriate, in order to help developers reuse their skills and transition
21 to new languages, and to help make sure the ideas are time-tested. *See, e.g.*, James
22 Gosling’s “Feeling of Java” paper, *Computer*, Vol. 30, Issue 6, June 1997, where he
23 writes “Java feels very familiar to many different programmers because Sun had a very
24 strong tendency to prefer things that had been used a lot over things that just sounded like
25 a good idea.” As a result, many API element names in modern languages are drawn from
26 the public domain. For example, package names like java.io, java.util, and java.net
27
28

1 reflect industry shorthand for common functionality like input/output, utilities, and
 2 networking, respectively. These packages with common names then, in turn, contain
 3 methods whose naming reflects industry custom and representation of the underlying
 4 functionality of the method. Some examples of functions that are very similar in Java
 5 and the pre-existing C and C++ languages as a result of their functionality and industry
 6 custom are shown below:
 7

Name	Originated in?	Java equivalent	What is it?
char	At least C; <i>see</i> C Reference Manual, Dennis Ritchie, 1975 (<i>available at</i> http://www.cs.bell-labs.com/who/dmr/cman.pdf)	char	A data type holding a character; used repeatedly in method names, such as C++ “getchar” and Java “getChars.” Other data types, such as <i>int</i> and <i>double</i> , also date back to C and at least the 1970s.
int abs (int i)	At least C; <i>see</i> C Reference Manual.	public static int abs (int a)	A function, returning the absolute value of the argument.
printf()	At least C; <i>see</i> C Reference Manual.	printf() (part of the java.io package)	A function, printing a formatted string to the screen or other output device.

- 1 125. Indeed, many of the names and concepts in Java have been used by the industry for
2 decades. For example, the C Reference Manual references “int,” “double,” and “char,”
3 all used in Java. The “Bool” data type, which became “Boolean” in Java, dates back to at
4 least Algol in 1968, and is a direct reference to Boolean logic — invented in the 1800s.
5
- 6 126. Another example of this is the `java.util.regex` API package, which implements “regular
7 expressions” — a standardized way of testing if a given string of characters matches a
8 particular pattern. Regular expressions were first formalized in 1968 (“Programming
9 Techniques: Regular expression search algorithm,” Ken Thompson, *Communications of*
10 *the ACM*, Vol. 11, Issue 6, June 1968) and were known by the abbreviated name used by
11 Java (`regex`) at least as early as 1983 (*see, e.g.* [http://groups.google.com/group/net.lang.c/
12 browse_thread/thread/6409987225e13a31/50da7fdd143184bd?q=regex#50da7fdd143184
13 bd](http://groups.google.com/group/net.lang.c/browse_thread/thread/6409987225e13a31/50da7fdd143184bd?q=regex#50da7fdd143184bd)). `Java.util.regex` has two classes: `Pattern`, and `Matcher`. Method names in the `Pattern`
14 class are `compile`, `flags`, `matcher`, `matches`, `pattern`, and `split`, while method names in the
15 `Matcher` class include `matches`, `pattern`, `reset`, and `start`. Each of these names —
16 particularly the extremely common “pattern” and “matcher” — are used in publicly
17 available regular expressions software that predate `java.util.regex`, and all of them are
18 discussed in *Mastering Regular Expressions*, Jeffrey E. F. Friedl, O'Reilly and
19 Associates, 1997, which Oracle’s documentation for `java.util.regex` cites. Sun also used
20 third-party source code (Jakarta Regexp) to implement `java.util.regex`, and this code also
21 had references to many of these terms, including “compile,” “pattern,” and “match.”
22
- 23
- 24 127. Similarly, the `java.sql` and `javax.sql` packages are also based in part on pre-existing terms
25 widely used in the industry. The package names themselves are a reference to the SQL
26 standard, originally introduced in the academic literature as SEQUEL in 1974 (SEQUEL:
27 A structured English query language, Proc. ACM SIGFIDET Workshop, May 1974, pp.
28

1 249-264). The classes and methods frequently are named after SQL concepts, and in
2 particular (according to Sun's documentation at
3 http://jcp.org/aboutJava/communityprocess/first/jsr054/jdbc-3_0-pfd-spec.pdf) on the
4 X/Open SQL Call Level Interface (CLI), which dates to the first half of the 1990s
5 (*available at* <http://pubs.opengroup.org/onlinepubs/009654899/toc.pdf>). For example,
6 `java.sql` includes classes named "Array," "Blob," "Clob," and "Ref," which are the
7 names of data types from the SQL standard. Similarly, the SQL CLI standard defines a
8 method called "prepare" that operates on a `StatementText`. `Java.sql`'s `Connection` class
9 has a method called `prepareStatement` that has similar functionality. The SQL CLI
10 standard also uses "commit" and "rollback" to discuss specific actions that can be done to
11 a database, and `java.sql`'s `Connection` class has matching "commit" and "rollback"
12 methods that perform the actions discussed in the standard.
13

14
15 128. `Java.util.zip` is another example where the name of the package, and at least some API
16 element names within the package, are references to terms that substantially predate
17 Java's use of the terms. In this case, "zip" is a reference to the zip file format that has
18 been in use since before the creation of Java. Class names in this package include
19 "Adler32" (named after the Adler-32 algorithm invented by Mark Adler and licensed to
20 the public as part of the zlib library) and "CRC32" (named after the CRC-32 algorithm,
21 which dates back to the 1970s). Within the `java.util.zip` classes, method names include
22 "deflate," "inflate," and "setDictionary," which are very similar both to general industry
23 terms for these processes but also to the specific function names "deflate," "inflate," and
24 "deflatesetdictionary" that are in the publicly available open source library (zlib) that
25 predates, and is incorporated by, Java. (*See* <http://www.zlib.net/manual.html>)
26
27
28

1 programming language would not compile. Even more work on the part of the
2 programmer would be required if, for some reason, the organization, arguments, or return
3 values of the methods needed to be changed. It is not necessary to be a programmer to
4 understand how this could be jarring; changing the shortcuts for copy and paste from
5 Ctl+C and Ctl+V to something else would require every user of word processors to
6 change their behaviors, which is why those two keyboard commands have been used
7 unchanged across many programs since the first graphical user interfaces. As previously
8 mentioned in paragraph 35, this sort of compatibility and interoperability is important to
9 the industry, since (in the worst case) it allows programmers to reuse known, tested code
10 fragments — an important practice in the industry — and in the best case, where
11 compatibility is complete, it allows reuse of entire programs without modification.
12 Therefore, the APIs at issue were included in Android in order to allow Android to be
13 interoperable with existing code written in the Java programming language.
14

15
16 132. It is worth emphasizing that reuse of the API in this way does not mean that the
17 underlying program logic implementing the API was copied.

18
19 133. In my opinion, the Java APIs are necessary for functionality, interoperability and
20 programming efficiency.

21 **R. THE APIS AT ISSUE ARE DEMANDED BY THE INDUSTRY**

22
23 134. In my opinion, industry demand requires APIs that are compatible with Java, rather than
24 APIs that are similar to, but not compatible with, Java.

25
26 135. Industry and developer practice would tend to make it very difficult for Google to choose
27 a different API, or modify an existing Java API, when the Java language is used and
28 supported by Google. While curious developers do teach themselves new APIs, as a

1 general rule they prefer not to be forced to retrain on new APIs unless there is an
2 extremely good reason to do so. This is not simply a matter of losing the time spent
3 learning; learning new APIs also means buying new books, losing the ability to reuse
4 code fragments, and temporarily losing the fluency that comes with expertise in a
5 particular language's idioms and structures. Indeed, a significant goal of the discipline of
6 programming has always been to create reusable tools and build on what has been
7 developed before, and the value of sensible reuse of existing APIs has been understood to
8 be a significant part of this.
9

10 136. It is not just individual developers who strongly prefer existing languages. Companies
11 also prefer to write programs in existing languages. Doing so allows them to reuse
12 existing source code and tools; even where they cannot reuse entire programs, reuse of
13 fragments of code is very common. Such code reuse helps make software better by
14 allowing the reuse of tested, well-understood code, but it is only possible where platforms
15 allow the same APIs to be used. For example, existing code written in the Java language
16 which references the "abs" method would not run on a new platform unless the new
17 platform supported this method and the class and package in which this method resides
18 by implementing the API. As a result of these factors, the industry as a whole — both
19 programmers and the companies who employ them — strongly prefer to work with APIs
20 and API elements with which they are familiar.
21
22

23 137. It is not a coincidence that many software developers are familiar with Java APIs. Sun
24 went to great lengths to encourage developers to learn and use Java. This began when
25 Sun made the Java language, documentation, and API implementation available at no
26 charge in 1996, apparently with the intent to ensure that programmers throughout the
27 industry knew and had internalized the Java language, including these APIs. Sun also
28

1 worked extensively with educational institutions to help make the Java language a
2 common tool for introductory programming classes. For example, there was active
3 collaboration between Sun and the College Board in promoting Java as the language to be
4 used in the Advanced Placement Computer Science (AP CS) exams developed by the
5 College Board. These changes began just as the exam was switched from Pascal to C++
6 in 1999. Java was too large a language to be taught completely, so a subset of the classes
7 and methods were identified as being an important part of teaching computer science.
8 This led to the development of a group of classes that were part of the AP CS program
9 and that mirrored the Java API exactly, e.g., instead of using java.lang, the AP CS
10 program identified and used ap.java.lang with a corresponding documentation as part of
11 the subset. At a similar time, Sun collaborated with the BlueJ group headed by Michael
12 Kolling to develop an IDE (Integrated Development Environment) for novice
13 programmers that was simple, but that used best practices that were part of Sun's official
14 NetBeans IDE. The development of BlueJ led to a world wide adoption of BlueJ in many
15 colleges and high schools that continues today with the use of BlueJ and its derivative
16 GreenFoot which sees widespread adoption. In part as a result of that effort, I taught Java
17 to my students for many years, and continue to do so. This effort created a large base of
18 programmers who had learned the Java language and APIs, and would be reluctant to
19 retrain on other languages because of their investment in Java. It is my understanding
20 that Oracle, in its filings to the court, has claimed that the number of programmers who
21 have invested time and effort to learn Java is in the range of 6-7 million. The investment
22 of these programmers in learning the APIs is likely in the hundreds of millions of man
23 hours.
24
25
26

27 138. In part as a result of Sun's substantial efforts to encourage programmers to learn and use
28 the Java language, a large number of software applications have been written in the Java

1 language. The investment involved in creating these existing applications creates further
2 industry demand for compatibility that would allow use of this existing software or code
3 fragments on new platforms.

4
5 139. In conclusion, it is my opinion that the Java APIs have the following qualities: First, the
6 APIs represent concepts and methods of operation. Second, the API names, method
7 declarations, and organization are dictated by function or mechanical rules and do not
8 reflect creative expression. Indeed, as noted, many API element names existed in the
9 public domain and were not original to Java. Third, the APIs at issue are necessary for
10 interoperability and efficiency reasons, and their use is driven by industry demand.

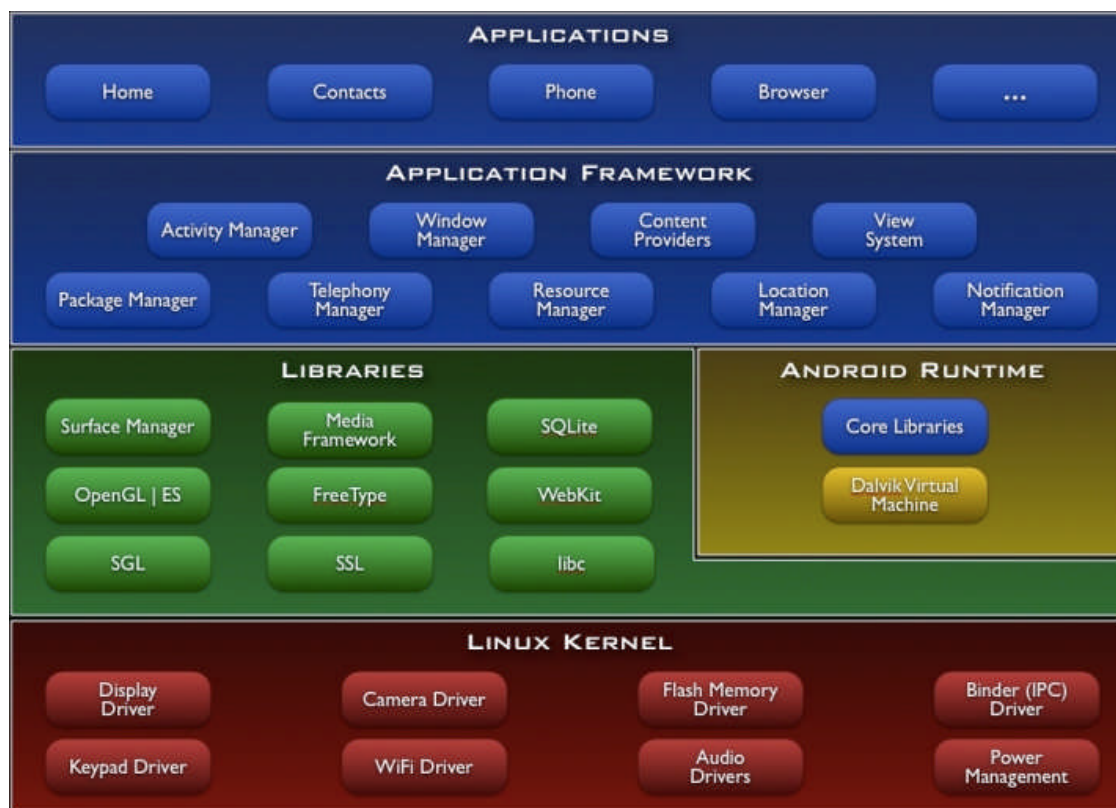
11
12 **VI. THE ANDROID PLATFORM IS NOT VIRTUALLY IDENTICAL OR**
13 **SUBSTANTIALLY SIMILAR TO THE JAVA PLATFORM**

14
15 140. The Android platform uses the unprotectable API elements at issue as part of a larger
16 overall framework that is substantially different from the Java platform, and more
17 appropriate for the mobile platform.

18
19 141. The Java API packages which have been implemented by Android are a small part of the
20 overall Android system, both in terms of the functionality they provide and the lines of
21 code involved. The 48 APIs at issue are roughly one-third of the Android Runtime Core
22 Libraries, which currently contains 168 API packages. The Core Libraries, in turn, are
23 themselves only a small part of the overall Android architecture — they are the small blue
24 box labeled “core libraries” in the yellow box at the right.³ Other elements include the

25
26
27 ³ Diagram taken from the Android website, *available at* <http://developer.android.com/images/system-architecture.jpg>.
28

1 Dalvik Virtual Machine, the Linux kernel, a web browser, and a variety of other libraries
 2 and systems that are not part of the Java platform. These other components, when
 3 combined with the Android Core Libraries, make for a complete mobile operating system
 4 — something substantially different in scope and ability than the Java platform.
 5



- 19
 20
 21
 22
 23
 24
 25
 26
 27
 28
142. While the diagram above is not “to scale” (boxes of the same size may represent software of different size and complexity), an analysis of the number of lines of source code in the API packages at issue, in the Android Runtime Core Libraries, and in the other components in the diagram suggests that, if anything, the diagram likely overrepresents the size of the APIs at issue.
143. Using a Python script SlocCounter.py (attached as Exhibit E) based on the “sloccount” tool, a commonly-used tool for measuring the size of the source code of large software projects, Android’s implementation of the APIs at issue in the “Gingerbread” release

1 constitutes 259,474 lines of code, in 1022 files. This is roughly 1.6% of the size of the
2 entire Android source code, which comprises 57,076 files and 15,347,169 lines of code,⁴
3 and roughly 15% of the 6,340 files and 1,713,087 lines of code⁵ in the overall Android
4 (Gingerbread) Runtime Core Libraries. Similarly, implementation of these APIs is a
5 small portion of Oracle's JDK 1.5 implementation of the entire Java API, constituting
6 315,570 lines of code out of 2,867,712 (11% of the total) and 1001 files out of 9521
7 (10.5%).
8

9 144. Because Android's implementation of the APIs at issue comprises only 1.6% of the entire
10 Android source code, Android as a whole is not virtually identical or substantially similar
11 to the Java platform. The Java APIs at issue are not only a small portion of the Android
12 platform, in my opinion the use of the Java APIs in the Android context is substantially
13 different from the use of those APIs in the Java platform, which creates a different
14 platform with different capabilities and functionalities targeting the emerging smartphone
15 market rather than the desktop.
16
17
18
19
20
21
22

23 ⁴ Numbers generated by running SlocCounterTotal.py (attached as Exhibit E) against a clean copy of
24 Android (obtained following the instructions available here: <http://source.android.com/source/downloading.html>)
25 and including only lines of code in .h, .c, .cpp, and .java files.

26 ⁵ Numbers generated by running SlocCounterTotal.py against the libcore/ and frameworks/base/core/
27 directories in a clean copy of Android, counting only lines of code in .h, .c, .cpp, and .java files.
28

1 **VII. ANDROID’S DOCUMENTATION OF THE APIS AT ISSUE IS NOT**
2 **VIRTUALLY IDENTICAL OR SUBSTANTIALLY SIMILAR TO ORACLE’S**
3 **DOCUMENTATION**
4

5 145. Every API has documentation. This documentation is a written description of the
6 functionalities provided by the API. The documentation is relied on by programmers
7 when they need more detail and context than what is available from the name and
8 organization of the method. Documentation also helps the programmer see what related
9 functionality is available in the same classes and packages. Like the naming and
10 organization of methods, good documentation of API components is extremely
11 constrained, because it must be factual and succinct. Just as a square root function should
12 not be named “Steve,” the documentation of the square root function, in order to be
13 maximally efficient for programmers who are referring to it, should be — and typically is
14 — factual and strictly descriptive of the underlying functionality.
15

16 146. Because of the practical requirement that documentation for an API strictly describe the
17 underlying functionality, and because programmers typically want to know the same
18 important pieces of information about a given API, there are not many ways to write
19 documentation for a specific API element. Two different authors documenting the same
20 API components would necessarily write very similar documentation because they are
21 describing the same functionality.
22

23 147. It may be useful to analogize writing documentation to writing a dictionary. If Merriam-
24 Webster defines a zebra as an “African mammal [] . . . related to the horse but
25 distinctively and conspicuously patterned in stripes of black or dark-brown” and
26 Dictionary.com defines a zebra as an “African mammal[] . . . each . . . having a
27 characteristic pattern of black or dark-brown stripes,” this does not mean that the second
28

1 dictionary copied from the first. The definitions are similar because they are describing
2 the same thing. It would be difficult to describe a zebra briefly without using words like
3 “Africa,” “stripes,” “black,” and “dark-brown.” In addition, even where there may be
4 multiple ways to describe things, the normal requirements that influence a dictionary
5 (*e.g.*, clarity and brevity) are present for both dictionaries. This results in dictionary
6 definitions that are similar even when written without copying.
7

8 148. Similarly, writing software documentation is significantly constrained, because the writer
9 must convey the same factual information briefly, accurately, and with clarity. As a
10 result of these constraints, skilled technical writers writing about the same API are likely
11 to come up with descriptions that appear very similar and contain very similar
12 descriptions of the critical features. The documentation for Java follows this pattern. For
13 example, the method in the class `java.io.PrintStream` whose prototype is “void write(int
14 OneByte)” takes one byte and writes (hence the name) that byte to the relevant “stream.”
15 Not surprisingly, Android’s brief documentation of this method states that this method:
16

17 Writes one byte to the target stream.

18
19 Oracle’s brief documentation for the same method states:

20 Write the specified byte to this stream.

21
22 The two different documentation writers have chosen different verb tenses for “write,”
23 and described the byte to be written slightly differently (“one byte” versus “the specified
24 byte”) but the descriptions are generally similar, not because they were copied from each
25 other but because they must accurately and briefly describe the underlying functionality.
26 Most documentation of the same method by two different authors will show the same
27 pattern of similarity, because the authors are aiming at efficient, factual statements.
28

1 149. In my opinion, the Android and Java documentations are not virtually identical, nor are
2 they substantially similar. Any similarity is dictated by the fact that they both document
3 the same functionalities, and does not imply or demonstrate that they both involve the
4 same creative expression. Indeed, the level of creativity reflected in the documentation is
5 very low because it is intended to be highly descriptive, and is generally circumscribed
6 by the nature of the functions described. It is my understanding that, in order to show
7 that the Android documentation for the APIs at issue infringes Oracle's copyrights
8 purportedly covering Oracle's documentation of the Java APIs, it must be shown that
9 original and creative elements of the Java documentation, if any, have been copied into
10 Android's documentation.
11

12 **VIII. THE TWELVE FILES OR PORTIONS OF FILES ALLEGED BY ORACLE TO**
13 **HAVE BEEN COPIED ARE QUALITATIVELY AND QUANTITATIVELY**
14 **INSIGNIFICANT AND THEY ADD NO OR VERY LITTLE VALUE TO**
15 **ANDROID**
16

17 150. In this section, I will analyze the 12 files that Oracle alleges were copied by Google.
18 These files represent, by number of files, 0.02% of Android and 0.13% of Oracle's
19 implementation of Java 1.5 (12 files out of 57,076 in Android and 9,521 in Oracle's
20 implementation of Java 1.5), and a similar percentage when measured by lines of code
21 (742 lines out of 15 million and 2.8 million lines of code for Android and Oracle Java
22 1.5, respectively.) It is my understanding that, aside from these 12 files, Oracle does not
23 allege that Google literally copied source code from Oracle.
24
25
26
27
28

1 **A. TIMSORT FILES**

- 2
- 3 151. It is my understanding that two Android files, TimSort.java and
- 4 ComparableTimSort.java, are at issue. These files are a tiny fraction of a percent of
- 5 Android and appear to have been donated by Google to Oracle, as explained below.
- 6 152. I have inspected three files: TimSort.java, ComparableTimSort.java, and Oracle's
- 7 implementation of the Array class contained in Arrays.java. Only a single method out of
- 8 the 11 methods in TimSort.java and 13 methods in ComparableTimSort.java is shared
- 9 between those two files and Arrays.java. This method is called rangeCheck, and I have
- 10 reproduced this method below:

Line #	Code
1	private static void rangeCheck(int arrayLen, int fromIndex, int toIndex) {
2	if (fromIndex > toIndex)
3	throw new IllegalArgumentException("fromIndex(" + fromIndex +
4	") > toIndex(" + toIndex+"");
5	if (fromIndex < 0)
6	throw new ArrayIndexOutOfBoundsException(fromIndex);
7	if (toIndex > arrayLen)
8	throw new ArrayIndexOutOfBoundsException(toIndex);
9	}

21 Quantitatively, this method is 9 lines of code out of 3,179 lines in the Oracle JDK 1.5

22 version of Arrays.java, 9 lines out of 924 lines in the latest Android version of

23 TimSort.java, and 9 lines out of the total 46,269 lines (0.5%) that compose Android's

24 implementation of the *single* java.util package at issue.

- 25
- 26 153. Qualitatively, the rangeCheck method is also trivial. It merely performs a simple,
- 27 utilitarian, and fairly mundane "sanity check," checking that certain arguments used

1 elsewhere are correct before they are used. The two arguments tested are the index to the
2 first element of interest in an array (“fromIndex”), and the index to the last element of
3 interest in the array (“toIndex”). They are compared against the number of elements in
4 the given array (“arrayLen”), against zero, and against each other to ensure that their
5 values are “in range” — that is to say that they are acceptable in the context of this code.
6

7 154. Because the first element should always come before the last element, if fromIndex is
8 larger than toIndex — line #2 in the code, “if (fromIndex > toIndex)” — that means that
9 the programmer made an error. If that happens, lines 3 and 4 of the method “throw an
10 exception” (in programming jargon, they indicate that an error has occurred).
11

12 155. Similarly, because an index should, by convention, never be less than zero, if fromIndex
13 is less than zero, this indicates another type of error. Line 5 tests for this (“if (fromIndex
14 < 0)”), and if the test fails, line 6 signals the error by throwing another exception.
15

16 156. Finally, the last element of interest (“toIndex”) must refer to an element that actually
17 exists, so if toIndex is greater than the total number of elements in the array (“arrayLen”)
18 (line 7), that also indicates an error, and the another exception is thrown by line 8.
19

20 157. This code was also necessary for API compatibility with other sort implementations in
21 Java. Because the “exceptions” thrown when an error occurs can be considered to be part
22 of the API of a method, using code extremely similar to this was necessary for TimSort to
23 be completely compatible with other sort implementations. If code extremely similar to
24 this, throwing the same exceptions under the same circumstances, had not been used, the
25 TimSort files could not have been accepted by Oracle into Java for use with Java 7.
26

27 158. It is my opinion that this single function is both qualitatively and quantitatively an
28 extremely small portion of the functionality of the TimSort.java and

1 ComparableTimSort.java files, which are in turn an extremely small portion of the overall
2 functionality of Java.

3
4 159. I have also reviewed the publicly available documents discussing the history of these two
5 files. It appears that they were written by Google employee Joshua Bloch. While the
6 files were first included in Android, it also appears that Google offered the files to Oracle
7 (*see*
8 [http://markmail.org/thread/xwyxemce75vvz33h#query:+page:1+mid:vnipd7bqzs5vxfjw+](http://markmail.org/thread/xwyxemce75vvz33h#query:+page:1+mid:vnipd7bqzs5vxfjw+state:results)
9 http://blogs.sun.com/mr/entry/jdk7_m5). As a result of this donation, these files are now
10 part of the most recent version of Oracle's implementation of Java.
11
12

13 **B. SECURITY TEST FILES**

14 160. It is my understanding that eight files in the directories
15 `"/support/src/test/java/org/apache/harmony/security/tests/support/acl/"` and
16 `"/support/src/test/java/org/apache/harmony/security/tests/support/cert/"` are at issue.
17 These files are all what are known as "test files," and as explained below they add
18 minimal if any material value to, and are a very small portion of, Android. Specifically,
19 the files at issue are:
20

21 AclEntryImpl.java

22 AclImpl.java

23 GroupImpl.java

24 OwnerImpl.java

25 PermissionImpl.java
26
27
28

1 PrincipalImpl.java

2 AclEnumerator.java

3
4 PolicyNodeImpl.java

5
6 161. It is good engineering practice to write companion software that tests the functionality of
7 the parts of the software that will be delivered to customers. This software (referred to as
8 “test software”) will do the software equivalent of factory testing, feeding the software
9 test inputs (such as “2 + 2”) and verifying that the correct output is returned (such as “4”).
10 Individual tests are known as “unit tests,” because they test a specific “unit” of the
11 software, and the files that are used to implement these unit tests are often referred to as
12 test files. Doing this testing in software, rather than manually, allows for the testing to be
13 done quickly and reliably. The testing can be done, for example, after every single
14 change to the software, rather than only every day or every week as might be required
15 with manual testing.
16

17 162. These particular files are in directories labeled “test,” and the structure of their source
18 code indicates that they are part of a test package, instead of a package which is part of
19 the publicly-available Android API. In addition, several of the files contain comments
20 indicating that they are for “verification” of a specific interface. Verification of an
21 interface is a common part of software testing. Finally, these files are referenced only in
22 trunk/dalvik/libcore/security/src/test/java/tests/security/acl/IAclTest.java and other files
23 in the same directory. This file, and others like it, appears to be the so-called “test
24 framework” — the software that runs the tests. This code is clearly structured to call
25 these files as tests, and not to use the files as part of the actual functionality of the
26 Android platform. In addition, it is worth noting that Android’s history shows that these
27
28

1 were removed from Android in January of 2011 and have not, as of this writing, been
2 replaced.⁶ This confirms that the files were immaterial.

3
4 163. These files are also in large part “dummy” files — instead of having complex logic, they
5 return certain, fixed values, which is a common practice in test files. For example, one
6 method in AclEntryImpl.java consists entirely of the following code:

```
7         public boolean isNegative() {  
8             return negative;  
9         }
```

10 164. In a real (not test) file, the “isNegative” method would do some complex logic to
11 understand whether the quality was negative. Here, because this is a “dummy” file used
12 for test purposes, no logic or work is done — instead, it simply immediately returns
13 “negative.” This “dummy” result would not do much good for real code, but in a test
14 environment, this predictability is useful — if a test shows that for some reason this
15 function is returning “positive,” then something is wrong and must be fixed. Dummy
16 files in general, and these files in particular, are not particularly creative — given the
17 functional constraint of the method names that they are testing, there is typically only one
18 efficient, reasonable way to write them.

19
20 165. In general, test files are written for internal use by developers prior to a product’s release.
21 They are typically not distributed as part of consumer products, for two reasons. First,

22
23 ⁶ See changes at

24 [http://android.git.kernel.org/?p=platform/libcore.git;a=commitdiff;h=6241c067e065065098eb50a7aef35a5](http://android.git.kernel.org/?p=platform/libcore.git;a=commitdiff;h=6241c067e065065098eb50a7aef35a58f78447a6)
25 [8f78447a6](http://android.git.kernel.org/?p=platform/libcore.git;a=commitdiff;h=6241c067e065065098eb50a7aef35a58f78447a6) and

26 [http://android.git.kernel.org/?p=platform/libcore.git;a=commitdiff;h=95d52b3b1446af2fed46f57efc1afb6c](http://android.git.kernel.org/?p=platform/libcore.git;a=commitdiff;h=95d52b3b1446af2fed46f57efc1afb6c679e8cc)
27 [679e8cc](http://android.git.kernel.org/?p=platform/libcore.git;a=commitdiff;h=95d52b3b1446af2fed46f57efc1afb6c679e8cc)
28

1 the testing is finished when the software is finalized since end users of consumer products
2 cannot fix any problems found by the testing. Test files are simply not material to the
3 customer experience. Second, distributing them takes up additional space and resources
4 that could be used for other purposes that actually provide direct functionality to
5 consumers, so test files can negatively impact the customer experience if distributed. As
6 a result of these factors, test files generally are not used by or distributed to consumers.
7 In my opinion, these test files are likely to be the same as others — not distributed to
8 consumers and not material to the consumer experience.
9

10 166. Quantitatively, these eight files represent an extremely small part of Android's and
11 Oracle's implementations of the Java APIs at issue. These codebases are roughly 1.7
12 million lines of code and 2.8 million lines of code, respectively, and so these eight files
13 represent less than 0.1% of the lines of code. They are not even a substantial portion of
14 the overall number of test files. My analysis of Android 2.3 suggests that there are at
15 least 142 test files comprising 48,376 lines of code. These eight accused test files are less
16 than 5% of this, measured in terms of lines of code.
17

18 167. Qualitatively, these eight files have no material effect on how Android implements the
19 Java APIs at issue, because they are not used by or distributed to consumers.
20

21 168. In my opinion, these test files are qualitatively and quantitatively insignificant to the
22 overall Android system.
23

24 **C. COMMENTS IN CODESOURCETEST.JAVA AND**
25 **COLLECTIONCERTSTOREPARAMETERSTEST.JAVA**

26 169. It is my understanding that the files `CodeSourceTest.java` and
27 `CollectionCertStoreParametersTest.java` are at issue. The comments in these files that are
28

1 at issue were not written by Google, add no value to Android, and are a very small
2 portion of Android.

3
4 170. I have inspected these four files (two each from Oracle's implementation and Android's
5 implementation), and the only things that appear to be the same between these classes are
6 certain comments. Of the 36 comments in Android's CodeSourceTest.java at the time of
7 the complaint, eight appear to be the same as comments in Oracle's implementation of
8 the CodeSource class. Of the 16 comments in Android's
9 CollectionCertStoreParametersTest.java, 12 appear to be the same as comments in
10 Oracle's implementation of the CollectionCertStore Parameters class. No source code
11 appears to have been copied, and the comments at issue appear to have been removed
12 from Android.⁷

13
14 171. Comments are used in software source code because they help programmers understand
15 the code that they are reading and potentially modifying. However, comments do not
16 become part of the final product that is shipped to the user of the software. If all
17 comments were removed, the functionality would be identical, and users would be
18 generally unaffected. These comments, like all comments generally, would not have
19 been distributed as part of any Android-based products; they would have only been
20 available to any programmers who downloaded the source code from the Android
21 website. It is my opinion, therefore, that these comments did not add material value to
22 the Android platform.
23

24
25
26
27 ⁷ The removal is documented here:

28 <http://android.git.kernel.org/?p=platform/libcore.git;a=commitdiff;h=a49d9caee4cd74c0d2cf83d79b8ecdc00453dff8>

1 172. In addition, these comments are largely very descriptive and functional. For example,
2 one of the copied comments appears to be from the following source code:

```
3     /**  
4     * Returns a formatted string describing the parameters.  
5     *  
6     * @return a formatted string describing the parameters  
7     */  
8     public String toString() {  
9         StringBuffer sb = new StringBuffer();  
10        sb.append("CollectionCertStoreParameters: [\n");  
11        sb.append(" collection: " + coll + "\n");  
12        sb.append("]");  
13        return sb.toString();  
14    }
```

15 173. The similar portion of the comment is “Returns a formatted string describing the
16 parameters.” This is a simple, declarative statement, which describes the source code
17 below it. Like other documentation discussed in paragraph 145, there are very few ways
18 to state this, because it is a simple, factual description of the operation of the public
19 method below it. The other comments at issue, with one exception, are very similar — a
20 single sentence factually describing the method in question. The one exception, slightly
21 longer and more detailed, is still only three sentences long.

22 174. I have investigated the history of these files, and have determined that these comments
23 were not created by Google. It appears that they were written by Intel employees who
24 donated the files containing the comments to an open source project called Apache
25 Harmony. My understanding is that Apache Harmony is an independent implementation
26 of the Java APIs, created by a non-profit foundation called the Apache Foundation. My
27 understanding is that Apache Harmony is licensed to the public under the terms of the
28 Apache License, which allow anyone to reuse the code with essentially no restrictions. It
is my understanding that, in compliance with the terms of the Apache License, Android
reused the Apache Harmony implementation of some parts of the Java APIs, and my
inspection of files in Android confirms this understanding. Android’s versions of these

1 two files are virtually identical to Apache Harmony's versions of the same files, strongly
2 suggesting that these two files (CollectionCertStoreParametersTest.java and
3 CodeSourceTest.java) were obtained by Google for Android under license from the
4 Apache Foundation.

5
6 175. The Apache Foundation makes its software records available to the public, documenting
7 the history of all files distributed by the Apache Foundation, including
8 CollectionCertStoreParametersTest.java and CodeSourceTest.java. By using these
9 records, it is possible to see when a particular file was created, when individual lines of
10 code were written or modified, and by whom. Using these records, I determined that
11 these comments have been present in these files since before the Android project. For
12 example, the Android file "CollectionCertStoreParametersTest.java" contained the
13 comment fragment "the default parameter values (an empty and immutable." This same
14 comment fragment is also present in the Harmony file
15 "CollectionCertStoreParametersTest.java." Finally, this same fragment is present in the
16 Oracle file "CollectionCertStoreParameter.java." By using the history features of the
17 Apache Foundation's source code storage tool, I verified that this same fragment and
18 matching surrounding text were present in that Harmony file when the file was initially
19 created by an Intel employee, Geir Magnusson, on Jan. 8, 2006.⁸ After the file was added
20 to Android, Google did not change the comment; the comment stayed the same from the
21
22
23
24

25
26 ⁸ See file history available at [http://svn.apache.org/viewvc/incubator/harmony/enhanced/classlib/trunk/java-](http://svn.apache.org/viewvc/incubator/harmony/enhanced/classlib/trunk/java-src/security2/test/common/unit/java/security/cert/CollectionCertStoreParametersTest.java?limit_changes=0&view=markup&pathrev=367016)
27 [src/security2/test/common/unit/java/security/cert/CollectionCertStoreParametersTest.java?limit_changes=0&view=](http://svn.apache.org/viewvc/incubator/harmony/enhanced/classlib/trunk/java-src/security2/test/common/unit/java/security/cert/CollectionCertStoreParametersTest.java?limit_changes=0&view=markup&pathrev=367016)
28 [markup&pathrev=367016](http://svn.apache.org/viewvc/incubator/harmony/enhanced/classlib/trunk/java-src/security2/test/common/unit/java/security/cert/CollectionCertStoreParametersTest.java?limit_changes=0&view=markup&pathrev=367016).

1 time that the file was licensed from Harmony until they were removed.⁹ The same is true
2 of the other comments in these two files, all of which appear to have originated with
3 Intel's contribution to Apache on Jan. 8, 2006. As a result, while I believe that these files
4 had some identical comments (before they were removed by Google), it appears that this
5 is related to choices made by Intel and not by Google.
6

7 176. It is also important to note that these two files represent an extremely small part of
8 Android's and Oracle's implementations of the Java APIs. These codebases are roughly
9 1.7 million lines of code and 2.8 million lines of code, respectively, and so these two files
10 represent less than 0.1% of lines of code.
11

12 177. In addition, these files are also test files, similar to the files discussed in the previous
13 section. In my opinion, these files are qualitatively and quantitatively insignificant to the
14 overall Android platform.
15

16
17 //

18
19 //

20
21 //

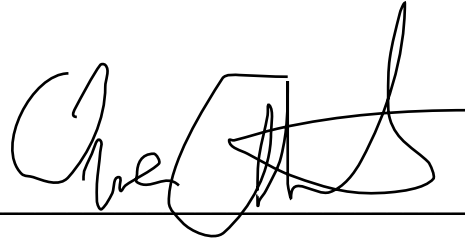
22
23
24
25 ⁹ See file history at

26 [http://android.git.kernel.org/?p=platform/libcore.git;a=history;f=luni/src/test/java/org/apache/harmony/security/tests/
27 java/security/CodeSourceTest.java;hb=a49d9caee4cd74c0d2cf83d79b8ecdc00453dff8](http://android.git.kernel.org/?p=platform/libcore.git;a=history;f=luni/src/test/java/org/apache/harmony/security/tests/java/security/CodeSourceTest.java;hb=a49d9caee4cd74c0d2cf83d79b8ecdc00453dff8)
28

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28

178. I reserve the right to update and refine my opinions and analyses in light of any additional materials or information that may come to my attention in the future, including additional contentions by Oracle as well as any rulings issued by the Court in this case. I also reserve the right to supplement my opinions and analyses as set forth in this report in light of any expert reports submitted by Oracle and in light of any deposition or trial testimony of Oracle's experts.

DATED: July 29, 2011



Owen Astrachan, Ph.D.

EXHIBIT A: OWEN ASTRACHAN CV

Owen L. Astrachan

Department of Computer Science
 Box 90129
 Duke University
 Durham, NC 27708-0129
 telephone: (919) 660-6522
 email: ola@cs.duke.edu

July 29, 2011

I. Education

Ph.D.	Computer Science	Duke University	1992
M.S.	Computer Science	Duke University	1989
M.A.T.	Mathematics	Duke University	1979
A.B.	Mathematics	Dartmouth College	1978

with distinction in Mathematics, Summa Cum Laude, Phi Beta Kappa

II. Professional Appointments

Duke University, Department of Computer Science

Professor of the Practice

July 2000 — .

Associate Professor of the Practice

July 1996 — July 2000.

Director of Undergraduate Studies

September 1993 — .

Assistant Professor of the Practice of Computer Science

1993 — 1996.

Lecturer in Computer Science

1991–1992. Developed and taught the introductory course for non-majors. Served on lab committee determining priorities for physical improvements.

Research Assistant

1991 (June–Aug). Research Assistant at SRI International, AI group, working for Mark Stickel on the design of intelligent and efficient automated reasoning systems.

Research Assistant

1988–1991. Research Assistant for Donald W. Loveland, investigating automated theorem proving. Designed and implemented an OR parallel theorem prover that runs on a BBN Butterfly GP-1000, TC2000 and on a network of Sun workstations.

Senior Graduate Instructor

1986–1988. Solely responsible for developing the curriculum and teaching the first course for majors in the Computer Science Department. Assisted with course in Operating Systems.

Teaching Assistant

1985–1986 Served as Teaching Assistant for the the first two courses for majors in Computer Science. Responsible for designing laboratory exercises and running recitation sections.

University of British Columbia, Computer Science Department

Visiting Scholar and Lecturer

Sept 1998 — June 1999 (on sabbatical from Duke)

Experience in Secondary Education

Math and Computer Science Teacher Durham Academy

1980–1985. Taught Advanced Placement Calculus, Advanced Placement Computer Science, Multivariable Calculus and Linear Algebra, Geometry, Introduction to Computer Programming, Finite Mathematics. Developed curriculum for Finite Math, AP Computer Science, and Multivariable Calculus.

Math Teacher Camp Lejeune HS, Camp Lejeune, NC

1978–1980. Taught honors Trigonometry, honors Geometry, Algebra I, Pre-Algebra.

III. Honors

2007, NSF, CISE Distinguished Education Fellow, *Interdisciplinary Problem- and Case-based Computer Science*, one of two inaugural CDEF awardees (see grants).

2004, IBM Faculty Award, *Issues in Large-scale Software Componentization*

2003, ACM International Collegiate Programming Contest (ICPC) Coaches Award.

2002, Richard K. Lublin Award for Distinguished and Motivating Teaching

2002, Nominated for Alumni Distinguished Teaching Award

2001, Nominated for Alumni Distinguished Teaching Award

1998, Outstanding instructor of Computer Science, University of British Columbia (teaching on sabbatical)

1997, NSF Career Award

1996, Nominated for Alumni Distinguished Teaching Award

1995, Duke University, Trinity College of Arts and Science: Robert B. Cox Distinguished Teaching in Science Award

1995, Sigma Xi

1994, Nominated for Alumni Distinguished Teaching Award

1978, A.B. degree awarded with distinction, summa cum laude, Phi Beta Kappa

IV. Publications**Journals:**

Owen Astrachan and Robert Dewar. CS Education in the U.S.: Heading in the Wrong Direction. *Communications of the ACM*. July 2009, v. 52, n. 7, pp. 41-45.

O.L. Astrachan and D.W. Loveland. The Use of Lemmas in the Model Elimination Procedure. *Journal of Automated Reasoning*, v. 19 n.1, August, 1997, pp. 117-141.

Owen Astrachan, Kim Bruce, Robert Cupper, Peter Denning, Scot Drysdale, Tom Horton, Charles Kelemen, Cathy McGeoch, Yale Patt, Viera Proulx, Roy Rada, Richard Rasala, Eric Roberts, Steven Rudich, Lynn Stein, Allen Tucker, Charles van Loan. Strategic Directions in Computer Science Education. *ACM Computing Surveys*. v 28, n 4, December 1996.

O.L. Astrachan. METEOR: Exploring Model Elimination Theorem Proving. *Journal of Automated Reasoning*. v.13 n.2, 1994, pp. 283-296.

Books:

Owen Astrachan. *A Computer Science Tapestry: Exploring Programming and Computer Science with C++*, Second edition. McGraw-Hill, 2000.

Owen Astrachan. *A Computer Science Tapestry: Exploring Programming and Computer Science with C++*. McGraw-Hill, 1997.

Owen Astrachan. *The Large Integer Case Study in C++*. The College Board, Advanced Placement Program, 1997.

Book Chapters:

- Owen Astrachan and Robert Duvall and Eugene Wallingford. Bringing Extreme Programming to the Classroom. in *Extreme Programming Perspectives*, Giancarlo Succi (ed.), Addison Wesley, 2002.
- O.L. Astrachan and D.W. Loveland. METEORs: High Performance Theorem Provers Using Model Elimination. in *Automated Reasoning: Essays in Honor of Woody Bledsoe* ed. R.S. Boyer, Kluwer Academic Press 1991.

Miscellaneous:

- O.L. Astrachan and Susan Horwitz and the Advanced Placement Computer Science Development Committee. *Communications of the ACM*. Technical Opinion: The First Course Conundrum. June, 1995. Pages 117-118.

Refereed Conferences:

- Owen Astrachan. Pander to Ponder, *SIGCSE Technical Symposium on Computer Science Education*, Chattanooga, TN, 2009
- Casey Alt, Owen Astrachan, Jeffrey Forbes, Richard Lucic, and Susan Rodger. Social Networks Generate Interest in Computer Science. *SIGCSE Technical Symposium on Computer Science Education*, Houston, TX, 2006.
- Owen Astrachan. Non-Competitive Programming Contest Problems as the Basis for Just-in-time Teaching. *Proceedings of Frontiers in Education*, October 2004.
- Owen Astrachan. Bubble Sort: An Archaeological Algorithmic Analysis. *SIGCSE Technical Symposium on Computer Science Education*, Reno, NV, 2003.
- Owen Astrachan and David Bernstein and Andrew English and Benjamin Koh. Development Issues for a Networked, Object Oriented Gaming Architecture (NOOGA) Teaching Tool. *Proceedings of Frontiers in Education*, November 2002.
- Owen Astrachan and Robert Duvall and Jeffrey Forbes and Susan Rodger. Active Learning in Small to Large Courses *Proceedings of Frontiers in Education*, November 2002.
- Owen Astrachan and Robert Duvall and Eugene Wallingford. Bringing Extreme Programming to the Classroom, *Proceedings of XPUniverse*, Raleigh, NC, July, 2001.
- Owen Astrachan. OO Overkill: When Simple is Better than Not, *SIGCSE Technical Symposium on Computer Science Education*. Charlotte, NC, February 2001.
- Charles Keleman, Allen Tucker, Peter Henderson, Kim Bruce, Owen Astrachan. Has Our Curriculum Become Math-Phobic?, *SIGCSE Conference on Integrating Technology into Computer Science Education (ITiCSE)*, June 2000.
- Owen Astrachan and Eugene Wallingford. Loop Patterns. *Pattern Languages of Programming (PLoP)*, Allerton Park, IL, August, 1998.
- Owen Astrachan. Hooks and Props as Instructional Technology. *SIGCSE Conference on Integrating Technology into Computer Science Education (ITiCSE)*, August 1998.
- Owen Astrachan, Geoffrey Berry, Landon Cox and Garrett Mitchener. Design Patterns: An Essential Component of CS Curricula. *SIGCSE Technical Symposium on Computer Science Education*. Atlanta, GA, February 1998.
- Owen Astrachan and Susan Rodger. Animation, Visualization, and Interaction in CS 1 Assignments, *SIGCSE Technical Symposium on Computer Science Education*. Atlanta, GA, February 1998.
- Owen Astrachan and Robert Smith and James Wilkes. Application-based Modules using Apprentice Learning for CS 2. *SIGCSE Technical Symposium on Computer Science Education*. San Jose, CA, February 1997, pp 233-237.

- Owen Astrachan and Trevor Selby and Joshua Unger. An Object-Oriented, Apprenticeship Approach to Data Structures using Simulation. *Frontiers in Education*, Salt Lake City, Utah, 1996, pp 130–134.
- Owen Astrachan and David Reed. AAA and CS1 : The Applied Apprenticeship Approach to CS 1. *SIGCSE Technical Symposium on Computer Science Education*. Nashville, TN, March 1995.
- Owen Astrachan and Claire Bono. Using simulation in an objects-early approach to CS1 and CS2. *OOPSLA Conference Proceedings, Educator's Forum*. Portland, Oregon, October 1994.
- O. L. Astrachan and D.W. Loveland. METEOR: Model Elimination Theorem Proving with Lemmas (system abstract). *Twelfth Conference on Automated Deduction (CADE-12)*. Nancy, France, 1994.
- Owen Astrachan. Self reference is a Thematic Essential. *SIGCSE Technical Symposium on Computer Science Education*. Phoenix, Arizona, March 1994.
- Owen Astrachan. METEOR: Exploring Model Elimination Theorem Proving. *Workshop on Theorem Proving with Analytic Tableaux and Related Methods*. Marseilles, France, April 1993.
- Owen L. Astrachan, Vivek Khera, and David Kotz. The Duke Internet Programming Contest: A Report and Philosophy. *SIGCSE Technical Symposium on Computer Science Education*. Indianapolis, IN, February 1993.
- Owen L. Astrachan and Mark E. Stickel. Caching and Lemmaizing in Model Elimination Theorem Provers. *Eleventh Conference on Automated Deduction (CADE-11)*. Saratoga Springs, NY, June 1992.
- Owen Astrachan. Finding a Stable roommate, job or spouse: a case study crossing the boundaries of Computer Science Courses. *SIGCSE Technical Symposium on Computer Science Education*. Kansas City, MO, March 1992.
- Owen Astrachan. Pictures as Invariants. *SIGCSE Technical Symposium on Computer Science Education*. San Antonio, TX, March 1991.
- Owen Astrachan. METEOR: Model Elimination Theorem Prover for Efficient OR-Parallelism. *AAAI Spring Symposium on Representation and Compilation in High Performance Theorem Proving: Titles and Abstracts*, ed. W.W. Bledsoe, M. Stickel, P. Lincoln, R. Overbeek, and D. Plaisted, Stanford, CA, March 1989.

Unrefereed Reports:

- Owen L. Astrachan and Donald W. Loveland *The Use of Lemmas in the Model Elimination Procedure*. Duke University Technical Report CS-1993-25.
- Owen L. Astrachan. *METEOR: Exploring Model Elimination Theorem Proving*. Duke University Technical Report CS-1992-22.
- Owen L. Astrachan. *Investigations in Model Elimination Based Theorem Proving*. Ph.D. Thesis. Also Duke University Technical Report CS-1992-21.
- Owen L. Astrachan and Mark E. Stickel. *Caching and Lemmaizing in Model Elimination Theorem Provers*. SRI International Technical Note 513, December 1991.

V. Service

Professional Service

- 2009 – *College Board HEAC: Higher Education Advisory Committee for Advanced Placement*. Provide oversight and advice regarding the program.
- 2009, July, NSF Review Panel: Broadening Participation in Computing (BPC)
- 2009 *NITRD: Networking and Information Technology Research and Development Program*, panelist at public forum for discussion of the 2009 Federal Strategic Plan
- 2008 – *NSF/College Board* joint group on the First Year in Computer Science (chair).
- 2008, May, NSF Review Panel *CPATH*
- 2007 – *College Board AP Computer Science Redesign Commission* Committee charged with examining and redefining the Advanced Placement Computer Science Program.
- 2006– ACM Ed-Council. One of 25 members providing leadership and governance to the ACM about educational activities and outreach.
- 2005, July, NSF Review Panel *Advanced Learning Technology Program Committee OOPSLA 2005, Educator’s Symposium*
- Program Committee OOPSLA 2004, Educator’s Symposium*
- Internet & Society Idea Exchange* Faculty Steering Committee for courses related to Internet and Society (oversight from Harvard and MIT, including 55 faculty from around the world).
- NSF Review Panel CRCD Program*
2004, reviewed proposals for the CRCD program in Computer Science at NSF.
- ACM/College Board Digital Library Project* 2004 – Advisory Committee to develop a project supporting Compute Science in high schools as part of the national NSF-sponsored digital library program.
- ACM/College Board JETT Steering Committee* 2002 – , Member of four-person steering committee providing oversight for joint ACM/College Board committee reviewing and approving national sites to host high school outreach programs for computer science.
- Illinois Math and Science Academy*
2002 – 2003, Member of three-person external review board for Mathematics/Computer Science at IMSA.
- College Board/High School Computer Science AP Computer Science*
2001 – 2002, Member of College Board *ad hoc* professional development committee to develop standards for training/educating high school teachers and workshop consultants in computer science.
- NSF Review Panel CISE Program*
2000, reviewed proposals for the CRCD program in Computer Science at NSF.
- Math and Computer Science Mathematical Association of America*
1999. Committee made recommendations to the MAA on the role of mathematics in computer science. Committee consisted of Alan Tucker, Charles Keleman, Dale Skrien, Charles van Loan, Peter Henderson, Kim Bruce, Owen Astrachan.
- Chair, Advisory Committee for AP Computer Science College Board*
1999–2000. Committee making recommendations on the use of new languages and curricula in Advanced Placement Computer Science. Committe consists of David Gries, Robert (Corky) Cartwright, Henry Walker, Ursula Wolz, Cay Horstmann, Fran Trees, Rich Kick.
- External Oversight Board North Carolina Central University*
1997 – 1998, oversee the growth and accreditation of the Computer Science Department.
- NSF Review Panel CISE Program*
1997, reviewed proposals to the Education Innovation program in the CISE directorate of NSF.

External Review Committee Oberlin College

1996, Member of external review committee to evaluate computer science department at Oberlin.

Program Committee CADE-13

1995/6, Member of program committee for CADE-13, Conference on Automated Deduction.

Chair, Advisory Committee for AP Computer Science College Board

1995–1996. Committee makes recommendations on the best use of C++ on the Advanced Placement Exam. The Committee is convened by the College Board, with representation from SIGCSE, the Computer Science Education special interest group of the ACM.

Judge, ACM International Programming Contest Association for Computing Machinery

1994–1997. One of six people responsible for developing problems and judging solutions for the ACM Programming Contest finals.

Chief Reader, Advanced Placement Computer Science Educational Testing Service

1989–1994. Responsible for developing grading standards and assigning scores for the AP exam in Computer Science. Assist with the development of the exam. Oversee, hire, and manage 70 University faculty consultants and High School educators in the grading of 10,000 AP exams taken by secondary students throughout the world.

Member, AP Computer Science Development Committee The College Board

1985–1989. Responsible for developing curriculum and devising tests for the AP exam in Computer Science.

Director, Duke Internet Programming Contest

1990 — 1994 . Co-founded a computer programming contest held in real-time over the Internet, involving 60 teams from 37 institutions in 5 countries (1990); 240 teams from 100 institutions in 9 countries (1991); 290 teams from 140 institutions in 14 countries (1992), 495 teams from 200 institutions in 20 countries (1993). Developed the problems used in the contest, designed solutions for the problems, and co-directed the administration of the contest.

IEEE Programming Contest

1994–1995. Coach of the Duke undergraduate IEEE Programming team. This competition is by invitation only to sixteen teams throughout the world. In 1994 Duke participated for the first time. In 1995 Duke won the contest.

ACM Programming Contest

1993–. Coach of Duke undergraduate ACM Programming Team. In 1994 the team won the mid-Atlantic regional contest and placed third in the world (first U.S. team) in the world finals. In 1995 the team won the mid-Atlantic regional contest and advanced to the world finals, finishing 22nd. In 1997 the team advanced to the world finals. In 1998 the team advanced to the world finals. (On sabbatical in 1999-2000.) In 2001, the team advanced to the world finals. In 2002 Duke had four teams in the top fifteen, and the top two in the Midatlantic regional contest; the top team advances to the world finals. In 2003 Duke advanced to the world finals. In 2004 Duke won the region (tied for first), advanced to the world finals and had three teams in the top 15 of the region; in the world finals Duke was one of four US teams that placed (above honorable mention) and was tied with Caltech and MIT for second among US teams. In 2005 Duke won the Region and advanced to the world finals receiving an honorable mention. In 2006 a Duke team advanced to the world finals, in 2007 Duke won the region and participated in the world finals, in 2008 Duke received an at-large bid to the world finals (to be held in 2009).

1989–1990. Member of Duke Programming team, 1989 and 1990 ACM International Programming Contests. Finished fourth in 1989 (Louisville, KY) and eighth in 1990 (Washington, DC) contest (world) finals.

Referee Activities

2010	SIGCSE	Technical Symposium on Computer Science Education
2009	SIGCSE	Technical Symposium on Computer Science Education
2008	SIGCSE	Technical Symposium on Computer Science Education
2007	SIGCSE	Technical Symposium on Computer Science Education
2006	SIGCSE	Technical Symposium on Computer Science Education
2005	ITiCSE	SIGCSE Conference on Integrating Technology into Computer Science Education (ITiCSE)
2005	SIGCSE	Technical Symposium on Computer Science Education
2004	ITiCSE	SIGCSE Conference on Integrating Technology into Computer Science Education (ITiCSE)
2004	SIGCSE	Technical Symposium on Computer Science Education
2003	SIGCSE	Technical Symposium on Computer Science Education
2001	SIGCSE	Technical Symposium on Computer Science Education
2000	SIGCSE	Technical Symposium on Computer Science Education
1996	SIGCSE	Technical Symposium on Computer Science Education
1995/6	CADE-13	Conference on Automated Deduction
1995	TABLEAUX '95	Workshop on Theorem Proving with Analytic Tableaux and Related Methods
	IJCAI	International Joint Symposium on Artificial Intelligence Automated Reasoning Track
	SIGCSE	Technical Symposium on Computer Science Education
1994	SIGCSE	Technical Symposium on Computer Science Education
1993	ILPS	International Logic Programming Symposium
	IJCAI	Automated Reasoning track
	ISMIS	International Symposium on Methodologies for Intelligent Systems
	SIGCSE	Technical Symposium on Computer Science Education
1992	CADE-11	Conference on Automated Deduction
	SIGCSE	Technical Symposium on Computer Science Education
1991	SIGCSE	Technical Symposium on Computer Science Education
1990	CADE-10	Tenth Conference on Automated Deduction

Duke Service

- 2009. Chair promotion committee for Jeffrey Forbes.
- 2008. Member of Office Education Committee (OEC) overseeing appointments of ROTC faculty.
- 2007-2009. Member of Academic Council.
- 2007-2010. University Committee on Admissions and Financial Aid (Academic Council).
- 2007–2009 Member of QEP (quality enhancement plan) University Committee to create a 75-page document outlining Duke's future as part of our 10-year SACS re-accreditation process.
- 2008 Search Committee, Dean to replace Robert Thompson, resulted in appointment of Lee Baker.
- 2007. Chair promotion committee for Susan Rodger.
- 2007-2009. University Commencement Committee.
- 2006-2007. Chair *ad hoc* Committee to Distinguish Trinity College Degrees. Will report on the status of the BA and BS degrees at Duke.
- 2006–. Member of Faculty Research Committee that decides on Trinity College competitive grants and awards to faculty for research and conference activity.
- 2005. Member of Committee on Departmental Support for Technology as part of the University committees on Strategic Directions.
- 2005. Member of Arts and Science search committee for the A&SIST Associate Dean.
- 2004-2005. Chair re-appointment committee for Prof. Richard Lucic
- 2004-. ISIS (Information Sciences and Information Studies) Faculty Steering Committee.
- 2004-. One of three faculty overseeing the development of teaching and learning, in conjunction with the CIT, as part of re-allocation of resources regarding the former center for teaching and learning.

2004-2007. Member Executive Committee of the Arts and Science Council (ECASC).
2004-2007. Member of ITAC, Committee on Information Technology at Duke.
Chair re-appointment committee for Prof. Richard Lucic 2002.
Member re-appointment committee for Prof. Jeffrey Forbes 2002.
Member Advisory Board for BlackBoard at Duke, 2002–
CITIE, IT skills committee, 2002.
Member Academic Integrity Council, 2001–
Member Executive Committee of the Arts and Science Council (ECASC) 2000–2003.
Interviewed candidates for A.B. Duke program, 2000, 2001
Member Board of Directors for Center for Instructional Technology (CIT), 1999–.
Arts and Science Council, 1999–2008.
Chair, Arts and Science Committee on Integrated Cluster Classrooms, 1999–2000.
Chair re-appointment/promotion committee for Prof. Robert Duvall, 2000.
Chair, University Planning Group on Instructional Technology, 1999–2000.
Search Committee, head of Career Development Center, 1998.
Member B.N. Duke Scholarship Committee, 1997
Member Core Team on evaluating use of Instructional and Information Technology, 1997.
ITAC Committee on Student Computing, 1996.
Chair search committee for Lecturer position in Computer Science, 1996.
Chair re-appointment/promotion committee for Prof. Susan Rodger, 1996.
Search Committee, Assistant Dean of Student Development, 1996.
Arts and Science Council, 1995–1996.
Member of Management Team (Center for Teaching and Learning) to develop an Exercise in Interactive Theatre for “Developing Teacher Knowledge”, 1996.
Director of Undergraduate Studies, 1993–.
Provost’s ad-hoc committee for Computer Technology and Education, 1995–1997.
Member of Steering Committee, Schulzbeger Interactive Learning Laboratory, Teaching and Learning Center, 1994–1996.
Faculty Advisor, DULUG: Duke University Linux User’s Group, 1995–.
Departmental major advisor 1992– (supervise 20 first majors, 19 second majors per yaer.)
Premajor advisor at Duke University, 1986– 1998(ten first year students each year, total of 20 per year)
Chaired search committee for Assistant Professor of-the-practice, 1994.
Committee on the role of teaching for graduate students 1990–1992

VI. Consultancies

Expert Witness, Alston and Bird

2010 — Retained by Nokia in “Certain Mobile Communications and Computer Devices and Components Thereof,” ITC Inv. No. 337-TA-704 (complaint filed Jan. 15, 2010), before the US International Trade Commission. Worked on Markman report, expert report, deposition and testimony.

Expert Witness, Alston and Bird

2008/2009 — Retained by Plaintiffs in Move, Inc., Nat’l Assoc. of Realtors, and Nat’l Assoc. of Home Builders vs. Real Estate Alliance Ltd et al., 2:07-CV-02185-GHK-AJW (filed Apr. 3, 2007) in the Central District of California. Worked on claim construction and expert report preparation.

Consultant, College Board

2008 – Oversee and help plan colloquium for college faculty (attended by 70 faculty) to understand current and future directions for AP Computer Science.

Google

2006 (six months) – Worked as an external contractor to help develop material used internally at Google for educating Google software engineers about Java and C++ programming.

Expert Witness, Womble, Carlyle, Sandridge and Rice

Software expert regarding work related to a contract dispute.

Consultant, AP Computer Science Educational Testing Service

1994 – Advise development committee on incorporating C++ and Java into the Advanced Placement Program. Wrote Pascal/C++ case studies for use on the exam. Provided workshops for high school and college consultants in making transition to object oriented programming. Critique free response questions that are part of the national exam (2002, 2004).

Case-Study author, AP Computer Science Educational Testing Service

1994 – Write the case study for use in the 1997-2000 AP exams. Write the code for the case study used in the 2001-2004 exams. A case is a “literate program”, a treatise on the design, development, and implementation of a programming solution to a problem.

VII. Panels/Conference Activities

CS Principles: Piloting a New Course at National Scale, SIGCSE, Dallas, 2011 (with Larry Snyder, Tiffany Barnes, Dan Garcia, Jody Paul, Beth Simon).

The CS10K Project: Mobilizing the Community to Transform High School Computing SIGCSE, Dallas, 2011 (with Jan Cuny, Chris Stephenson, and Cameron Wilson)

The CS/10K Project, CRA/Snowbird Conference, Snowbird UTAH, July, 2010.

Code as a Metaphor for Computational Thinking, CSTA/CSIT Symposium, Google, Mountain View, CA, July, 2010.

Re-imagining the First Year of Computer Science, SIGCSE, Milwaukee, WI, 2010 (with Lien Diaz, Chris Stephenson, Jan Cuny, Amy Briggs)

FOSS Workshop, Free and Open Source Software, SIGCSE, Chattanooga, TN, 2009, invited speaker.

Computational Thinking Panel, SIGCSE, Chattanooga, TN, 2009 (with Amber Settle, Susanne Hambrusch, and Joan Peckham).

Advanced Placement Computer Science: The Future of Tracking the First Year of Instruction, Special Session, SIGCSE, Chattanooga, TN, 2009 (with Henry Walker, Chris Stephenson, Lien Diaz, and Jan Cuny)

Nifty Assignments, Special Session, SIGCSE, Chattanooga, TN, 2009 (with Nick Parlante)

Innovating our Self Image Special Session, SIGCSE, Portland, OR, 2008 (with Peter Denning)

Teaching Tips We Wish They Told Us Before We Started Special Session, SIGCSE, Cincinnati, OH, 2007 (with Dan Garcia, Nick Parlante, Stuart Reges).

Resolved: Objects Early Has Failed SIGCSE, St. Louis, 2005, Special Session (with Stuart Reges, Kim Bruce, Michael Kölling, Elliot Koffman).

But it Looks Right: Bugs Students Don't See SIGCSE, Norfolk, 2004, Special Session (with David Ginat, Daniel Garcia, Mark Guzdial).

Colorful Illustrations of Algorithmic Design Techniques SIGCSE, Charlotte, 2001, Special Session (with David Ginat, Joseph Bergin, Dan Garcia).

Nifty Assignments in CS1 and CS2, Panelist, SIGCSE, Charlotte, 2001 (with Michael Clancy, Nick Parlante, Rich Pattis, Stuart Reges, Julie Zelenski).

FYI 2000: First Year Instruction, developed, organized, and chaired a workshop on first year instruction in computer science. The workshop had invited talks from both industry (Jon Bentley) and academia (Shriram Krishnamurthy, Richard Pattis) and was attended by more than 40 faculty from across the country. The workshop was sponsored by NSF, Microsoft, and the Duke Computer Science Department.

- Patterns in Computer Science*, (co-organizer with Eugene Wallingford), SIGCSE. Austin, TX. March 2000.
- The Future of Advanced Placement Computer Science (panel). *SIGCSE Technical Symposium on Computer Science Education*. Austin, TX, 2000. With Corky Cartwright, David Gries, Cay Horstmann, Richard Kick, Fran Trees, Henry Walker, Ursula Wolz.
- Nifty Assignments in CS1 and CS2*, Panelist, SIGCSE, New Orleans, 1999 (with Michael Clancy, Nick Parlante, Rich Pattis, Stuart Reges, Julie Zelenski).
- Incorporating Patterns into CS courses and Writing Patterns for CS Courses* (co-organizer with Eugene Wallingford). SIGCSE, New Orleans, March 1999.
- Future Directions in CS2 and Data Structures. Organized and ran the workshop that was held in conjunction with OOSPLA-98, Vancouver, CA, October 1998 (20 participants)
- Object Oriented Design. Invited Participant, OOPSLA-97, Atlanta, Georgia, 1997.
- Teaching Object-Oriented Programming: Practical Examples Using C++ and Java, Tutorial at PLDI 97, Las Vegas, Nevada, June 1997.
- Future Directions in Data Structures and CS2. Organized and ran two-day workshop held at Duke co-sponsored by NSF, 32 participants, March, 2000.
- Teaching Object-Oriented Design in the first year. Invited speaker and participant. OOPSLA-96, San Jose, CA, October, 1996.
- Strategic Directions in Computing Research (SDCR), working group in Computer Science Education. Sponsored by ACM, CRA, and NSF, Boston, MA, June, 1996.
- How to teach C++ in Introductory Courses, Tutorial part of PLDI, FCRC 1996, Philadelphia, PA, sponsored by SIGPLAN
- Formal Methods Considered [Help — Harm]ful: Engaging students in the first year. Exploring Formal Methods in the Early Computer Science Curriculum, Joint NSF/US Department of Education Workshop. September 16, 1995 (invited speaker).
- Developing an Object-Oriented Class Library. NSF sponsored workshop. Colgate University, June 1995. (invited participant)
- A Sorcerer's Apprentice Approach to using C++ in CS1. NECUSE (New England Consortium on Undergraduate Science Education) Workshop in Introductory Computer Science Curricula. January 1995.
- Measuring Performance of Automated Theorem Provers (with D. W. Loveland). *Twelfth Conference on Automated Deduction (CADE-12)*, Workshop on Evaluation of Automated Theorem Proving Systems. Nancy, France, 1994.
- Acquiring Object-Oriented Technology: A Bridge between Industry and Academia (invited participant). US West, Boulder Colorado, March 1994.
- OOP: An introduction for Secondary School Teachers. Workshop delivered to secondary school computer science teachers in Dallas, TX, August 1993.
- Simulation: A vehicle for exploring OOP. *Object-Oriented Curriculum Development Workshop*. NSF sponsored workshop. Colgate University, July 1993. (with Claire Bono)
- A Tapestry of Fundamental Ideas and Concepts in Computer Science: a Programming, Contextual View for Liberal Arts Students *L**3: Logic, Loops, and Literacy*. NSF sponsored workshop on computer science for non-majors, Brooklyn College, May 1993.
- Human Interaction with a High-Performance Theorem Prover. *International Joint Conference on Artificial Intelligence*. Workshop on Automated Theorem Proving. Chambéry, France, 1993. (with D.W. Loveland)
- Logic Programming Considered Harmful? *Joint Internatinal Conference on Logic Programming*. Program as a first language track, Washington DC, November 1993.
- Caching to reduce redundancy in Model Elimination Theorem Provers. *Joint Japanese-American workshop in Automated Theorem Proving*, Argonne National Labs, June 1991.

The METEOR implementations of the Model Elimination procedure. *Workshop on Proof Theory and Automated Theorem Proving*. Oberwolfach, Germany, April 1991. (with D.W. Loveland)

Online Exams in Advanced Placement Computer Science. *National Council of Teachers of Mathematics Conference*. San Francisco, April 1984.

VIII. Invited Lectures, Talks, and Workshops.

Code as a Metaphor for Computational Thinking, Harambeenet workshop, Durham, NC, July 2010.

CS Principles and the CS10K project, NSF, BPC Community Meeting, Los Angeles, February, 2010.

CCSC: Midwest, Keynote Speaker, "A New Way of Thinking about Computational Thining", St. Xavier University, Chicago, October 2009.

National Academies: Workshop on Computational Thinking February 2009.

Problem-Centric Learning, Sept 2008, Rochester Institute of Technology

What is Computer Science?, April 2008, NSF CPATH, Living in the Knowledge Society.

Problems in AP Computer Science, June 2008, Advanced Placement Computer Science Reading, Professional Development Night.

CPATH: Science and Computer Science Purde University, November 2007.

CPATH: Problem-based Learning Keynote, Workshop sponsored by Argonne Labs and Governors University, November 2007.

Microsoft Computational Thinking Summit Redmond, WA, September 2007.

Google Faculty Summit, Mountain View, CA, July 2007.

HarambeNet: Introducing Computer Science through Modeling and Analysis of Social Networks SIGCSE 2007 Workshop, with Jeffrey Forbes.

The Cruelty of Really Teaching Computer Science Redux, University of California, Riverside, January 2006.

The Cruelty of Really Teaching Computer Science Redux, University of British Columbia Computer Science Distinguished Lecturer Series, Fall 2005.

The Cruelty of Really Teaching Computer Science Redux, University of Washington Computer Science Distinguished Lecturer Series, Fall 2005.

The Cruelty of Really Teaching Computer Science Redux, Keynote talk at CCSC/SE, Consortium of Computing Sciences in Colleges, Southeast US, November 2005.

The Cruelty of Really Teaching Computer Science Redux, Keynote talk at CCSC/E, Consortium of Computing Sciences in Colleges, Eastern US, October 2005.

A Random Walk Through Computer Science, Invited/Keynote Talk at ACM/Student conference Reflections/Projections, University Illinois, Oct. 2004.

20 Years of Teaching Computer Science, invited talk at NSF Workshop for high school teachers, Stonehill College, October 2004.

Everything I Needed to Know about Programming and Computer Science I Learned from my Teachers, Keynote Talk, SIGCSE, Norfolk 2004.

Using Patterns in the First Year, invited tutorial and presentation as part of the 2000 Eastern Small College Computing Conference, University of Scranton, PA.

OO Design and Patterns, invited speaker at NSF-sponsored workshop for high school teachers at Stonehill College, June 2000.

Advanced OO Programming, invited speaker at NSF-sponsored workshop for high school teachers at Stonehill College, January 1999.

Object-Oriented Design and Programming. Three-day lecture/workshop co-taught with David Gries delivered to educators from business colleges in Denmark, October, 1998.

- Possible Futures for CS2 (panelist). *SIGCSE Technical Symposium on Computer Science Education*. Atlanta, GA, 1998.
- Teaching C++ in AP Courses: Four day workshop designed and delivered for the College Board, June and August, 1997.
- The First Computer Science Course and C++: Paradigm Lost or Regained*. DIMACS workshop on Discrete Mathematics, July, 1996.
- C++ in the Advanced Placement Program*. AP Computer Science Reading, Professional Night, Clemson, SC. June, 1996.
- Use of C++ for CS1 and CS2, Computing Science Conference, Philadelphia, PA, 1996.
- The First Year: Beyond Language Issues, (moderator and proposer), *SIGCSE Technical Symposium on Computer Science Education*. Philadelphia, PA, 1996.
- Advanced Placement and C++: Opening a Dialogue, (moderator and proposer), *SIGCSE Technical Symposium on Computer Science Education*. Philadelphia, PA, 1996.
- Object-Oriented Programming: How to “Scale Up” CS1. *SIGCSE Technical Symposium on Computer Science Education*. Phoenix, Arizona, 1994.
- Themes and Tapestries: A Diversity of Approaches to Computer Science for Liberal Arts Students. *SIGCSE Technical Symposium on Computer Science Education*. Phoenix, Arizona, 1994.
- Using Case Studies in Computer Science Courses. *SIGCSE Technical Symposium on Computer Science Education*. Phoenix, Arizona, 1994.
- On Computer Science and Teaching Computer Science with some perspectives from automated Reasoning. Bryn Mawr College, January 1993.
- Faster, Fairer and More Consistent Grading Techniques: Lessons From the Advanced Placement Reading *SIGCSE Technical Symposium on Computer Science Education*. Washington D.C., 1990.
- The Pleasures and Perils of Teaching Introductory Computer Science. *North Carolina Council of Teachers of Mathematics Conference*, November 1990.
- Teaching Recursion in Introductory Computer Science Courses. *North Carolina Council of Teachers of Mathematics Conference*. November 1986.

IX. Professional Affiliations

Member of ACM, IEEE Computer Society, SIGPLAN, SIGCSE, SIGACT, SIGCHI, SIGART, SIGSOFT, Sigma Xi.

X. Research Funding

- 2009, NSF IIS, Special Projects: *Using Computational Thinking to Model a New Course: Advanced Placement Computer Science: Principles*, \$2,093,450.00, funding to College Board, PI.
- 2008, NSF BPC, Computational Thinking and Fluency in the 21st Century. \$98,415. Submitted by College Board, PI.
- 2007, NSF CPATH, CISE Distinguished Education Fellow *Interdisciplinary Problem- and Case-based Computer Science*, \$250K over two years, one of two inaugural CDEF awardees in Computer Science.
- 2004 IBM Faculty Fellow, *Issues in Large-scale Software Componentization*, \$40K grant.
- 2002-2003 \$28K, IBM Echelon: Eclipse Help for Learning Online.
- 2002 \$400K. IBM SUR, Education/Research Grant for establishing COD (Cluster On Demand) (co-PI with Richard Lucic, Amin Vahdat, Jeff Chase).
- 2002 \$1,500, OOPSLA Educator’s Grant: Using Design Patterns (declined)
- 2001 \$1,500, OOPSLA Educator’s Grant: Using Design Patterns (declined)
- 2001 \$120,000 IBM Education/Research Grant for establishing a teaching/cluster classroom (co-PI with Susan Rodger, Richard Lucic, Kishor Trivedi, Amin Vahdat, Jeff Chase, the \$120,000 is just the education part of the grant, each of three groups was awarded a similar amount. This was part of a \$1.7 million SUR grant from IBM, some of which was software and related support.)

- 2000-2005, \$480,826, NSF, *Modules and Courses for Ubiquitous and Mobile Computing*, NSF CRCDCRCD 0088078, PI, co-PIs Prof. Carla Ellis, Prof. Amin Vahdat.
- 2000, \$1.19 million Microsoft *Interactive Research/Teaching Classroom*, with Jeffrey Vitter, Richard Lucic, Jeffrey Chase, Carla Ellis, Deitolf Ramm, Susan Rodger, Amin Vahdat. This includes \$750,000 in software support and the rest in equipment, construction, and staffing support.
- 1998 \$223,179 equipment *Establishing Interactive Collaborative Classrooms* Hewlett-Packard University Grants Program, co-PI with Susan Rodger
- 1998 \$1,500, OOPSLA Educator's Grant: Using Design Patterns
- 1998 \$50,000 Microsoft Educational Development Grant, co-PI with Susan Rodger and Jeffrey Vitter
- 1998–2001 \$150,306 U.S. Dept of Education GANN (co-PI with Jeff Chase, Carla Ellis, Alvin Lebeck, Jeffrey Vitter)
- 1997–1998 \$80,000 equipment (part of a \$1.6 million grant) from Intel supporting computer science education at Duke.
- 1997 \$1,200, OOPSLA Educator's Grant: Using Design Patterns
- 1997–2002, \$200,004, “Using and Developing Design Patterns”, National Science Foundation: CAREER Program, CAREER #9702550
- 1996, \$1,700, OOPSLA Educator's Grant: Using Design Patterns
- 1996–1997, \$119,382, “The Applied Apprenticeship Approach (AAA): An Object-Oriented/Object-Based Framework for CS2”, National Science Foundation Course and Curriculum Development, grant#DUE-9554910.
- 1996–2001, \$607,800, “CURIOUS: The Center for Undergraduate Education and Research: Integration Through Performance and Visualization”, NSF CISE Educational Infrastructure Program, grant #CISE-9634475 (co-PI with Prof. Susan Rodger)
- 1996–1998, \$13,080 “U.S.-Germany Cooperative Research to Enhance the Performance of the Model Elimination Proof Procedure” (co-PI with Prof. Donald Loveland), National Science Foundation INT-9514375
- 1995, \$1,000, OOPSLA Educator's Grant: Design Patterns in the Introductory CS Curriculum.

XI. Research Interests

Problem-centric Learning, Software Architecture, Object-oriented systems and languages, Computer Science Education, Networked and Distributed Computing, Automated Theorem Proving, Automated Reasoning, Parallel and Distributed Computing.

XII. Teaching

- 2002, Winner of the Richard K Lublin Award. Cited for “Ability to engender genuine intellectual excitement, ability to engender curiosity, knowledge of field and ability to communicate that knowledge, organizational skills, creative arrangement of course.”
- 1999, Winner of Outstanding Instructor of Computer Science Award while on sabbatical at University of British Columbia (teaching CPSC 252, a course on Data Structures for approximately 250 Engineering students).
- 1995, Trinity College, Robert B. Cox Distinguished Teaching in Science Award. Cited for “knowledge of field and ability to communicate it to students, openness to students, skill in organizing courses, commitment to teaching over time.”
- 1994, 1996, 2001, 2002 Nominated by undergraduates for outstanding teaching/faculty award (one of approximately 30 faculty nominated campus-wide each year).

Course and Curriculum Design/Implementation

CPS 53 – Program Analysis and Design I, Fall 1993

In 1993 I led the redesign of the core courses for majors (CPS 06, 08, and 100) to introduce object-oriented programming using C++ and an apprentice style of learning. This led to the development of new courses, described below. This redesign was a significant departure from the C-based courses that had been in place for three years requiring a large-scale change in philosophy as well as significant efforts in developing programming libraries to make the use of C++ feasible for students with no programming experience. This redesign has led to several publications and an NSF grant awarded in December of 1995.

CPS 100E - Program Analysis and Design II, Fall 1995 –

Created a new course for students with programming experience acquired elsewhere (not at Duke), replacing Computer Science 8 and accelerating students into the major. The course reviews material from the end of CPS 6 and then covers material from CPS 100. A formal laboratory component makes this possible. (with S. Rodger)

CPS 108 – Software Design and Implementation, Spring 1995 –

Designed and taught a new course required for all majors (in 1994 formalized the requirement). The course covers advanced object-oriented programming; introducing GUI programming using C++ and Java while emphasizing significant individual and team projects using object-oriented design, analysis, and programming.

In 1995 I established a new software design and engineering component of the curriculum via the course CPS 108. This curricular change led to an NSF CAREER grant for using and developing patterns in teaching software design and introductory programming.

CPS 149S - Problem Solving Seminar, Fall 1994

Created a new seminar course for problem solving, to prepare students for the ACM programming contest. Students worked previous contest problems once a week, and two mini-contests were held. Two teams participated in the regional contest with one team placing first.

CPS 182S – Technical and Social Analysis of Information and the Internet

Designed to meet the needs of Duke's Curriculum 2000. Satisfies, research and writing requirements and science technology and society requirement.

The development of technical and social standards governing the Internet and Information Technology in general. The role of software as it relates to law, patents, intellectual property, and IETF (Internet Engineering Task Force) standards. Written analysis of issues from a technical perspective with an emphasis on the role of software; but also on how standards relate to social and ethical issues.

In 2002 I designed and had approved CPS 182S, *Technical and Social Analysis of Information and the Internet* a course which R, W, and STS designations as part of Duke's Curriculum 2000 (research, writing, and science, technology and society, respectively). This course led to another, non-major's version of the course in 2008, to publications, and is part of the genesis for the new NSF CS Principles project.

CPS 004G – Programming for Bioinformatics

Designed as one course in a four-course, integrative and interdisciplinary program *The Genome Revolution: Society and Science* for first year students as part of Duke's FOCUS program. The course introduces programming in the context of solving problems from bioinformatics and computational biology.

In 2006 I used this course as a foundation, with work done by Alex Hartemink in our department on a more advanced course, to help spearhead and oversee the process leading to the approval of Duke's first interdepartmental and interdisciplinary minor: Computational Biology and Bioinformatics.

CompSci 82, Technical and Social Foundations of the Internet

In 2008 I used CompSci 182S (see above) as a model for developing CompSci 82, a course without the writing component, but with an Ethical Inquiry (EI) designation. This course was taught in the fall of 2008 to 239 students, the third largest enrollment for a one-section course at Duke.

In 2009 I taught this course to 345 students, the second largest course at Duke and the largest course that does not satisfy a major requirement. In 2010 I again taught the course to 345 students; the course was the largest single-section course taught at Duke.

Compsci 6, Introduction to Computer Science

In 2010 I oversaw the development of a new, introductory course in computer science: Compsci 6, the new description for the course follows.

Introduction practices and principles of computer science and programming and their impact on and potential to change the world. Algorithmic, problem-solving, and programming techniques in domains such as art, data visualization, mathematics, natural and social sciences. Programming using high-level languages and design techniques emphasizing abstraction, encapsulation, and problem decomposition. Design, implementation, testing, and analysis of algorithms and programs. No previous programming experience required.

This course is intended to appeal to a wider and more diverse audience than our previous version of the course. I developed the infrastructure for the course including developing materials used to teach Python, deciding on the libraries used, and developing the software infrastructure to support the use of Python. I worked with Robert Duvall to deliver the first version of this course in the fall of 2010.

Courses Taught

In the list below, CPS 08/53 corresponds to CS 1 and CPS 100/103 corresponds to CS 2 (courses were renumbered in 1994). CPS 206 is a graduate level programming-languages course. CPS 10 is a comprehensive/breadth first introduction to Computer Science for non-majors.

Date		Number	Title	Enrollment
2010	Spring	CompSci 149s	Problem Solving Seminar	12
		CompSci 100	Program Design and Analysis II	55
		CompSci 182s	Technical and Social Foundations of the Internet	18
	Fall	CompSci 149s	Problem Solving Seminar	8
		CompSci 006	Introduction to Computer Science	79
		CompSci 82	Technical and Social Foundations of the Internet	345
2009	Spring	CompSci 149s	Problem Solving Seminar	11
		CompSci 100	Program Design and Analysis II	29
		CompSci 182s	Technical and Social Foundations of the Internet	20
	Fall	CompSci 149s	Problem Solving Seminar	8
		CompSci 100	Program Design and Analysis II	41
		CompSci 82	Technical and Social Foundations of the Internet	345
2008	Spring	CompSci 149S	Problem Solving Seminar	7+
		CompSci 100	Program Design and Analysis II	33
		CompSci 82S	Technical and Social Foundations of the Internet	23
	Fall	CompSci 82	Technical and Social Foundations of the Internet 239	
		CompSci 100	Program Design and Analysis II	39
		CompSci 149s	Problem Solving Seminar	8+
2007	Spring	CompSci 149S	Problem Solving Seminar	8+
	Fall	CompSci 4G	Genomics Programming (FOCUS)	16
		CompSci 108	Software Design	44
		CompSci 149s	Problem Solving Seminar	5+
2006	Spring	CPS 182s	Technical and Social Analysis of the Internet	20
		CPS 100	Program Design and Analysis II	24
		CPS 149S	Problem Solving Seminar	4+
	Fall	CPS 004g	Introduction to Programming with Genomics (FOCUS)	11
		CPS 100	Program Design and Analysis II	36
		CPS 149S	Problem Solving Seminar	4+
2005	Fall	CPS 004G	Introduction to Programming with Genomics (FOCUS)	17
		CPS 108	Software Design and Implementation	32
		CPS 149S	Problem Solving Seminar	3
2005	Spring	CPS 182s	Technical and Social Analysis of the Internet	15
		CPS 149S	Problem Solving Seminar	4
2004	Fall	CPS 006G	Introduction to Programming with Genomics (FOCUS)	15
		CPS 100	Program Design and Analysis II	35
		CPS 149S	Problem Solving Seminar	7
2004	Spring	CPS 100	Program Design and Analysis II	35
		CPS 149S	Problem Solving Seminar	7
		CPS 108	Software Design and Implementation	45
2003	Fall	CPS 006X	Program Design and Analysis I (honors)	20
		CPS 182S	Technical and Social Analysis of the Internet	35
		CPS 149S	Problem Solving Seminar	7
2003	Spring	CPS 100	Program Design and Analysis II	68
		CPS 108	Software Design and Implementation	99

Date		Number	Title	Enrollment
2002	Fall	CPS 182s	Technical and Social Analysis of the Internet	43
		CPS 149S	Problem Solving Seminar	5
2002	Spring	CPS 100	Program Design and Analysis II	105
2001	Fall	CPS 108	Software Design and Implementation	64
		CPS 06	Program Design and Analysis I	105
		CPS 149S	Problem Solving Seminar	13
2001	Spring	CPS 108	Software Design and Implementation	124
		CPS 100	Program Design and Analysis II	114
		CPS 189S	CS Education Seminar	3
2000	Fall	CPS 100	Program Design and Analysis II	88
		CPS 149S	Problem Solving Seminar	14
		CPS 189S	CS Education Seminar	6
2000	Spring	CPS 100	Program Design and Analysis II	107
1999	Fall	CPS 108	Software Design and Implementation	115
		CPS 06	Program Design and Analysis I	173
1998	Fall	CS 252 (UBC)	Data Structures,CS2	163
1998	Spring	CPS 108	Software Design and Implementation	94
		CPS 196	Advanced Topics in OO Technology (seminar)	14
1997	Fall	CPS 100	Program Design and Analysis II	61
		CPS 108	Software Design and Implementation	63
		CPS 149S	Problem Solving Seminar	18
1997	Spring	CPS 100	Program Design and Analysis II	72
		CPS 108	Software Design and Implementation	65
1996	Fall	CPS 100E	Program Design and Analysis II	70
		CPS 108	Software Design and Implementation	58
		CPS 149S	Problem Solving Seminar	17
1996	Spring	CPS 100	Program Design and Analysis II	72
		CPS 108	Software Design and Implementation	40
		CPS 208	Software Design and Implementation	15

Date		Number	Title	Enrollment
1995	Fall	CPS 6	Intro. to Program Design/Analysis I (Team taught with S. Rodger)	121
		CPS 100	Program Design and Analysis II (Team taught with S. Rodger)	56
		CPS 100E	Program Design and Analysis II (Team taught with S. Rodger)	55
		CPS 149S	Problem Solving Seminar (Team taught with S. Rodger)	14
	Spring	CPS 108	Software Design and Implementation	38
		CPS 8	Intro. to Program/Design Analysis I	70
1994	Fall	CPS 8	Intro. to Program Design/Analysis I	63
		CPS 100	Program Design and Analysis II	45
	Spring	CPS 8	Intro. to Program Design/Analysis I	67
		CPS 100	Program Design and Analysis II	43
1993	Fall	CPS 8	Intro. to Program Design/Analysis I	39
		CPS 100	Program Design and Analysis II	29
	Spring	CPS 100	Program Design and Analysis II	48
		CPS 206	Programming Languages (graduate)	6
1992	Spring	CPS 10	Fundamentals of Computing	67
1991	Fall	CPS 10	Fundamentals of Computing	135
1987	Fall	CPS 51	Introduction to Programming	90
1987	Spring	CPS 51	Introduction to Programming	101
1986	Fall	CPS 51	Introduction to Programming	196

Thesis Advising

- 2004: Megan Murphy, *The Uses of Pair Programming in Introductory Computer Science Courses*, thesis for graduation with distinction.
- 2004: Megan Gessner, *Generation of Spanish Verb Conjugations*, thesis for graduation with distinction.
- 2002: Donald Onyango, *Comparison of Educational Tools*, Masters Thesis.
- 1998: Matthew Kotler, *An interactive CD-based Guide to Duke University*, undergraduate thesis for graduation with Highest Distinction.
- 1998: Eric Jewart, *Programming in CS 0*, undergraduate thesis for graduation with High Distinction.
- 1998: Tafawa Kesler, *GOODS: A Design and Class Building Tool*, undergraduate thesis for graduation with Distinction.
- 1997: Chih-ping Fu, *Towards a Java Bean Building and Using Environment*, Masters Thesis.

Independent Study

- Fall 2006, Spring 2007 *Computer and Mathematical Models of Insulin Pathways* Tiffany Chen, graduation with distinction, interdepartmental major, Computer Science and Biology (supervised from Biology by Fred Nijhout).
- Fall 2005, Industry/Academic software development with .NET technologies (with Glaxo-Smith-Kline)
- Spring 2005, Web-tools for Russian Vocabulary
- Spring 2004, Agile Methods and Programming for Spanish

Fall 2002, Patterns for Networked Games
Fall 2002, Online Grading
Fall 2001, Database-backed web sites: issues and solutions.
Fall 2000, A Framework for an online Calendar System supporting IETF standards
Spring 1998, Developing a CD-based guide to Duke, Advanced OO Design: A Class Browser
Fall 1997, Advanced Object Oriented Design, Interactive Web-based Journaling, Developing a CD-based guide to Duke
Spring 1997, Graphical Debugging
Spring 1997, Distance Learning using a Java Whiteboard
Fall 1996, On-line help by harvesting information with a GUI front end.
Fall 1996, Using C++ in High School Teaching.
Fall 1995, Graphics and Game programming for the Macintosh (6 students).
Summer 1995, Using C++ in High School Teaching
Spring 1995, A GUI/OO interface for air-quality modeling.
Spring 1995, Graphics Programming for the Macintosh.
Spring 1995, Object-Oriented Programming with Smalltalk.
Fall 1994, Implementing an Online Teacher Course Evaluation Book.
Spring 1994, An application-driven approach to foundations of computer science.
Summer 1994, The role of Computer Science for secondary school mathematics teachers.
Summer 1994, Computation Structures and Machine Organization.
Fall 1989, Graphical Display and Manipulation of Data Structures for the Macintosh.

EXHIBIT B: DOCUMENTS AND INFORMATION REVIEWED

- a. Oracle's First Amended Complaint
- b. Google's Answer to First Amended Complaint and Counterclaims
- c. Oracle's Supplemental Responses to Google's Interrogatories, Set No. 1
- d. The Android developer website at android.com
- e. The Oracle Java websites at java.sun.com and java.oracle.com, including
<http://java.sun.com/docs/white/platform/javaplatform.doc1.html>,
http://java.sun.com/docs/books/jls/first_edition/html/index.html,
<http://java.sun.com/docs/glossary.html>
- f. Source code, documentation, and discussion boards and blogs for Oracle's implementation of the APIs at issue, including
<http://download.oracle.com/javase/5/docs/index.html>,
<http://download.oracle.com/javase/1.4/docs/index.html>, and
<http://markmail.org/thread/xwyxemce75vvz33h#query:+page:1+mid:vnipd7bqzs5vxfjw+state:results>
- g. Source code and documentation for Android's implementation of the APIs at issue, including <http://developer.android.com/reference/packages.html> and
<http://android.git.kernel.org/?p=platform/libcore.git;a=history;f=luni/src/test/java/org/apache/harmony/security/tests/java/security/CodeSourceTest.java;hb=a49d9caee4cd74c0d2cf83d79b8ecdc00453dff8>
- h. Source code and documentation for Apache Harmony's implementation of the APIs at issue, including <http://svn.apache.org>, <http://harmony.apache.org/faq.html> and

<http://harmony.apache.org/subcomponents/classlibrary/compat.html>

- i. Source code and documentation for GNU Classpath's implementation of the APIs at issue, including <http://www.gnu.org/software/classpath/docs/>
- j. Wikipedia, *Application programming interface*, http://en.wikipedia.org/w/index.php?title=Application_programming_interface&oldid=437864024 (as of July 13, 2011, 00:30 GMT)
- k. Newton's Telecom Dictionary, 25th Edition
- l. Oracle SQL: The Essential Reference," David C. Kreines (2000)
- m. C Reference Manual, Dennis Ritchie, 1975, *available at* <http://www.cs.bell-labs.com/who/dmr/cman.pdf>
- n. ZLib Manual, *available at* <http://www.zlib.net/manual.html>
- o. Merriam-Webster Dictionary online
- p. Dictionary.com
- q. "Revised Report on the Algorithmic Language ALGOL 68", *available at* <http://www.fh-jena.de/~kleine/history/languages/Algol68-RevisedReport.pdf>
- r. "The C Family of Languages: Interview with Dennis Ritchie, Bjarne Stroustrup, and James Gosling," *Java Report*, 5(7), July 2000, *available at* http://www.gotw.ca/publications/c_family_interview.htm
- s. "The Feeling of Java," James Gosling, *Computer*, Vol. 30, Issue 6, June 1997
- t. *Mastering Regular Expressions*, Jeffrey E. F. Friedl, O'Reilly and Associates, 1997
- u. "Programming Techniques: Regular expression search algorithm," Ken Thompson,

Communications of the ACM, Vol. 11, Issue 6, June 1968

- v. "SEQUEL: A structured English query language," Proc. ACM SIGFIDET Workshop, May 1974, pp. 249-264
- w. JDBC 3.0 Specification, available at
http://jcp.org/aboutJava/communityprocess/first/jsr054/jdbc-3_0-pfd-spec.pdf
- x. X/Open: Data Management: SQL Call Level Interface (CLI), available at
<http://pubs.opengroup.org/onlinepubs/009654899/toc.pdf>
- y. Linux kernel 2.4 source code as available at
<http://git.kernel.org/?p=linux/kernel/git/stable/linux-2.4.37.y.git>, including
<http://git.kernel.org/?p=linux/kernel/git/stable/linux-2.4.37.y.git;a=blob;f=fs/proc/array.c;h=335226246dcafa18864e87c2f7be68f48a50b924;hb=HEAD>
- z. Solaris source code and revision history as available at
<http://cvs.opensolaris.org/source/xref/onnv>, including
<http://cvs.opensolaris.org/source/history/onnv/onnv-gate/usr/src/uts/common/syscall/uucopy.c> and
<http://cvs.opensolaris.org/source/history/onnv/onnv-gate/usr/src/uts/common/brand/>; and
also as archived at <http://hg.genunix.org/onnv-gate.hg/>; including
<http://hg.genunix.org/onnv-gate.hg/rev/4c5722bc28dc>
- aa. BrandZ documentation at
<http://hub.opensolaris.org/bin/view/Community+Group+brandz/>
WebHome and related web pages, including
<http://hub.opensolaris.org/bin/download/Community+Group+brandz/>

WebHome/brandzoverview.pdf and

<http://hub.opensolaris.org/bin/view/Community+Group+brandz/design>

- bb. “Fuss, Futexes and Furwocks: Fast Userlevel Locking in Linux,” 2002, *available at* <http://kernel.org/doc/ols/2002/ols2002-pages-479-495.pdf>
- cc. Futex(2) manual page, *available at* <http://www.kernel.org/doc/man-pages/online/pages/man2/futex.2.html>
- dd. “Excel functions (by category),” <http://office.microsoft.com/en-us/excel-help/excel-functions-by-category-HP005204211.aspx>
- ee. “Calc Functions listed by category,” http://wiki.services.openoffice.org/wiki/Documentation/How_Tos/Calc:_Functions_listed_by_category
- ff. Visicalc Reference Card, *available at* <http://www.bricklin.com/history/refcard1.htm>
- gg. “Oracle® Database Application Developer's Guide - Fundamentals,” *available at* http://download.oracle.com/docs/cd/B14117_01/appdev.101/b10795/toc.htm
- hh. “System R: relational approach to database management,” M. M. Astrahan et al; ACM Transactions on Database Systems; Vol. 1, Issue 2, June 1976.

EXHIBIT C: EXCEL AND STAROFFICE SPREADSHEET FUNCTION NAMES

Microsoft Excel 2003	Oracle Open Office Calc (Today)
ABS	ABS
ACCRI NT	ACCRI NT
ACCRI NTM	ACCRI NTM
ACOS	ACOS
ACOSH	ACOSH
	ACOT
	ACOTH
ADDRESS	ADDRESS
AMORDEGRC	AMORDEGRC
AMORLI NC	AMORLI NC
AND	AND
	ARABI C
AREAS	AREAS
ASC	
ASI N	ASI N
ASI NH	ASI NH
ATAN	ATAN
ATAN2	ATAN2
ATANH	ATANH
AVEDEV	AVEDEV
AVERAGE	AVERAGE
AVERAGEA	AVERAGEA
BAHTTEXT	BAHTTEXT
	BASE
BESSELI	BESSELI
BESSELJ	BESSELJ
BESSELK	BESSELK
BESSELY	BESSELY
BETADI ST	BETADI ST
BETAI NV	BETAI NV
BI N2DEC	BI N2DEC
BI N2HEX	BI N2HEX
BI N2OCT	BI N2OCT

BINOMDIST	BINOMDIST
CEILING	CEILING
CELL	CELL
CHAR	CHAR
CHIDIST	CHIDIST
CHIINV	CHIINV
	CHISQDIST
	CHISQINV
CHITEST	CHITEST
CHOOSE	CHOOSE
CLEAN	CLEAN
CODE	CODE
COLUMN	COLUMN
COLUMNS	COLUMNS
COM	
COMBIN	COMBIN
	COMBINA
COMPLEX	COMPLEX
CONCATENATE	CONCATENATE
CONFIDENCE	CONFIDENCE
CONVERT	CONVERT
	CONVERT_ADD
CORREL	CORREL
COS	COS
COSH	COSH
	COT
	COTH
COUNT	COUNT
COUNTA	COUNTA
COUNTBLANK	COUNTBLANK
COUNTIF	COUNTIF
COUPDAYBS	COUPDAYBS
COUPDAYS	COUPDAYS
COUPDAYSNC	COUPDAYSNC
COUPNCD	COUPNCD
COUPNUM	COUPNUM
COUPPCD	COUPPCD

COVAR	COVAR
CRI TBI NOM	CRI TBI NOM
CUMI PMT	CUMI PMT
	CUMI PMT_ADD
CUMPRI NC	CUMPRI NC
	CUMPRI NC_ADD
	CURRENT
DATE	DATE
DATEVALUE	DATEVALUE
DAVERAGE	DAVERAGE
DAY	DAY
	DAYS
DAYS360	DAYS360
	DAYSI NMONTH
	DAYSI NYEAR
DB	DB
DCOUNT	DCOUNT
DCOUNTA	DCOUNTA
DDB	DDB
	DDE
DEC2BI N	DEC2BI N
DEC2HEX	DEC2HEX
DEC2OCT	DEC2OCT
	DECI MAL
DEGREES	DEGREES
DELTA	DELTA
DEVSQ	DEVSQ
DGET	DGET
DI SC	DI SC
DMAX	DMAX
DMI N	DMI N
DOLLAR	DOLLAR
DOLLARDE	DOLLARDE
DOLLARFR	DOLLARFR
DPRODUCT	DPRODUCT
DSTDEV	DSTDEV
DSTDEVP	DSTDEVP

DSUM	DSUM
DURATION	DURATION
	DURATION_ADD
DVAR	DVAR
DVARP	DVARP
	EASTERSUNDAY
EDATE	EDATE
EFFECT	EFFECT_ADD
	EFFECTIVE
EOMONTH	EOMONTH
ERF	ERF
ERFC	ERFC
ERROR.TYPE	ERRORTYPE
EUROCONVERT	
EVEN	EVEN
EXACT	EXACT
EXP	EXP
EXPONDIST	EXPONDIST
FACT	FACT
FACTDOUBLE	FACTDOUBLE
FALSE	FALSE
FDIST	FDIST
FIND	FIND
FINDB	
FINV	FINV
FISHER	FISHER
FISHERINV	FISHERINV
FIXED	FIXED
FLOOR	FLOOR
FORECAST	FORECAST
	FORMULA
FREQUENCY	FREQUENCY
FTEST	FTEST
FV	FV
FVSCHEDULE	FVSCHEDULE
	GAMMA
GAMMADIST	GAMMADIST

GAMMAI NV	GAMMAI NV
GAMMALN	GAMMALN
	GAUSS
GCD	GCD
	GCD_ADD
GEOMEAN	GEOMEAN
GESTEP	GESTEP
GETPI VOTDATA	
GROWTH	GROWTH
HARMEAN	HARMEAN
HEX2BI N	HEX2BI N
HEX2DEC	HEX2DEC
HEX2OCT	HEX2OCT
HLOOKUP	HLOOKUP
HOUR	HOUR
HYPERLI NK	HYPERLI NK
HYPGEOMDI ST	HYPGEOMDI ST
I F	I F
I MABS	I MABS
I MAGI NARY	I MAGI NARY
I MARGUMENT	I MARGUMENT
I MCONJUGATE	I MCONJUGATE
I MCOS	I MCOS
I MDI V	I MDI V
I MEXP	I MEXP
I MLN	I MLN
I MLOG10	I MLOG10
I MLOG2	I MLOG2
I MPOWER	I MPOWER
I MPRODUCT	I MPRODUCT
I MREAL	I MREAL
I MSI N	I MSI N
I MSQRT	I MSQRT
I MSUB	I MSUB
I MSUM	I MSUM
I NDEX	I NDEX
I NDI RECT	I NDI RECT

INFO	INFO
INT	INT
INTERCEPT	INTERCEPT
INTRATE	INTRATE
IPMT	IPMT
IRR	IRR
ISBLANK	ISBLANK
ISERR	ISERR
ISERROR	ISERROR
ISEVEN	ISEVEN
	ISEVEN_ADD
	ISFORMULA
	ISLEAPYEAR
ISLOGICAL	ISLOGICAL
ISNA	ISNA
ISNONTEXT	ISNONTEXT
ISNUMBER	ISNUMBER
ISODD	ISODD
	ISODD_ADD
ISPMT	ISPMT
ISREF	ISREF
ISTEXT	ISTEXT
IT	
JIS	
KURT	KURT
LARGE	LARGE
LCM	LCM
	LCM_ADD
LEFT	LEFT
LEFTB	
LEN	LEN
LENB	
LINEST	LINEST
LN	LN
LOG	LOG
LOG10	LOG10
LOGEST	LOGEST

LOGI NV	LOGI NV
LOGNORMDI ST	LOGNORMDI ST
LOOKUP	LOOKUP
LOWER	LOWER
MATCH	MATCH
MAX	MAX
MAXA	MAXA
MDETERM	MDETERM
MDURATI ON	MDURATI ON
MEDI AN	MEDI AN
MI D	MI D
MI DB	
MI N	MI N
MI NA	MI NA
MI NUTE	MI NUTE
MI NVERSE	MI NVERSE
MI RR	MI RR
MMULT	MMULT
MOD	MOD
MODE	MODE
MONTH	MONTH
	MONTHS
MROUND	MROUND
MULTI NOMI AL	MULTI NOMI AL
	MUNI T
NA	NA
NEGBI NOMDI ST	NEGBI NOMDI ST
NETWORKDAYS	NETWORKDAYS
NOMI NAL	NOMI NAL
	NOMI NAL_ADD
NORMDI ST	NORMDI ST
NORMI NV	NORMI NV
NORMSDI ST	NORMSDI ST
NORMSI NV	NORMSI NV
NOT	NOT
NOW	NOW
NPER	NPER

NPV	NPV
OCT2BIN	OCT2BIN
OCT2DEC	OCT2DEC
OCT2HEX	OCT2HEX
ODD	ODD
ODDFPRICE	ODDFPRICE
ODDFYIELD	ODDFYIELD
ODDLPRICE	ODDLPRICE
ODDLYIELD	ODDLYIELD
OFFSET	OFFSET
OR	OR
PEARSON	PEARSON
PERCENTILE	PERCENTILE
PERCENTRANK	PERCENTRANK
PERMUT	PERMUT
PHONETIC	
	PERMUTATIONA
	PHI
PI	PI
PMT	PMT
POISSON	POISSON
POWER	POWER
PPMT	PPMT
PRICE	PRICE
PRICEDISC	PRICEDISC
PRICEMAT	PRICEMAT
PROB	PROB
PRODUCT	PRODUCT
PROPER	PROPER
PV	PV
QUARTILE	QUARTILE
QUOTIENT	QUOTIENT
RADIANS	RADIANS
RAND	RAND
RANDBETWEEN	RANDBETWEEN
RANK	RANK
RATE	RATE

RECEIVED	RECEIVED
REPLACE	REPLACE
REPLACEB	
REPT	REPT
RIGHT	RIGHT
RIGHTB	
ROMAN	ROMAN
ROUND	ROUND
ROUNDDOWN	ROUNDDOWN
ROUNDUP	ROUNDUP
ROW	ROW
ROWS	ROWS
	RRI
RSQ	RSQ
RTD	
SEARCH	SEARCH
SEARCHB	
SECOND	SECOND
SERIESSUM	SERIESSUM
	SHEET
	SHEETS
SIGN	SIGN
SIN	SIN
SINH	SINH
SKEW	SKEW
SLN	SLN
SLOPE	SLOPE
SMALL	SMALL
SQL.REQUEST	
SQRT	SQRT
SQRTPI	SQRTPI
STANDARDIZE	STANDARDIZE
STDEV	STDEV
STDEVA	STDEVA
STDEVP	STDEVP
STDEVPA	STDEVPA
STEYX	STEYX

	STYLE
SUBSTITUTE	SUBSTITUTE
SUBTOTAL	SUBTOTAL
SUM	SUM
SUMIF	SUMIF
SUMPRODUCT	SUMPRODUCT
SUMSQ	SUMSQ
SUMX2MY2	SUMX2MY2
SUMX2PY2	SUMX2PY2
SUMXY2	SUMXY2
SYD	SYD
TAN	TAN
TANH	TANH
TBILLREQ	TBILLREQ
TBILLPRICE	TBILLPRICE
TBILLYIELD	TBILLYIELD
TDIST	TDIST
TEXT	TEXT
TIME	TIME
TIMEVALUE	TIMEVALUE
TINV	TINV
TODAY	TODAY
TRANSPOSE	TRANSPOSE
TREND	TREND
TRIM	TRIM
TRIMMEAN	TRIMMEAN
TRUE	TRUE
TRUNC	TRUNC
TTEST	TTEST
TYPE	TYPE
UPPER	UPPER
VALUE	VALUE
VAR	VAR
VARA	VARA
VARP	VARP
VARPA	VARPA
VDB	VDB

VLOOKUP	VLOOKUP
WEEKDAY	WEEKDAY
WEEKNUM	WEEKNUM
	WEEKNUM_ADD
	WEEKS
	WEEKSINYEAR
WEIBULL	WEIBULL
WORKDAY	WORKDAY
XIRR	XIRR
XNPV	XNPV
YEAR	YEAR
YEARFRAC	YEARFRAC
	YEARS
YIELD	YIELD
YIELDDISC	YIELDDISC
YIELDMAT	YIELDMAT
ZTEST	ZTEST

EXHIBIT D: LX_BRAND SYSCALL TABLE

From lx_brand/common/lx_brand.c:

```
static struct lx_sysent sysents[] = {
    {"nosys",      NULL,          NOSYS_NULL,      0},      /* 0 */
    {"exit",      lx_exit,       0,              1},      /* 1 */
    {"fork",      lx_fork,       0,              0},      /* 2 */
    {"read",      lx_read,       0,              3},      /* 3 */
    {"write",     write,         SYS_PASSTHRU,   3},      /* 4 */
    {"open",      lx_open,       0,              3},      /* 5 */
    {"close",     close,        SYS_PASSTHRU,   1},      /* 6 */
    {"waitpid",   lx_waitpid,    0,              3},      /* 7 */
    {"creat",     creat,         SYS_PASSTHRU,   2},      /* 8 */
    {"link",      lx_link,       0,              2},      /* 9 */
    {"unlink",    lx_unlink,     0,              1},      /* 10 */
    {"execve",    lx_execve,     0,              3},      /* 11 */
    {"chdir",     chdir,         SYS_PASSTHRU,   1},      /* 12 */
    {"time",      lx_time,       0,              1},      /* 13 */
    {"mknod",     lx_mknod,      0,              3},      /* 14 */
    {"chmod",     lx_chmod,      0,              2},      /* 15 */
    {"lchown16",  lx_lchown16,    0,              3},      /* 16 */
    {"break",     NULL,          NOSYS_OBSOLETE, 0},      /* 17 */
    {"stat",      NULL,          NOSYS_OBSOLETE, 0},      /* 18 */
    {"lseek",     lx_lseek,      0,              3},      /* 19 */
    {"getpid",    lx_getpid,     0,              0},      /* 20 */
    {"mount",     lx_mount,      0,              5},      /* 21 */

```

```

{"umount",      lx_umount,      0,              1},      /* 22 */
{"setuid16",    lx_setuid16,    0,              1},      /* 23 */
{"getuid16",    lx_getuid16,    0,              0},      /* 24 */
{"stime",       stime,          SYS_PASSTHRU,   1},      /* 25 */
{"ptrace",      lx_ptrace,      0,              4},      /* 26 */
{"alarm",       (int (*)())alarm, SYS_PASSTHRU, 1},      /* 27 */
{"fstat",       NULL,           NOSYS_OBSOLETE, 0},      /* 28 */
{"pause",       pause,          SYS_PASSTHRU,   0},      /* 29 */
{"utime",       lx_utime,       0,              2},      /* 30 */
{"stty",        NULL,           NOSYS_OBSOLETE, 0},      /* 31 */
{"gtty",        NULL,           NOSYS_OBSOLETE, 0},      /* 32 */
{"access",      access,          SYS_PASSTHRU,   2},      /* 33 */
{"nice",        nice,           SYS_PASSTHRU,   1},      /* 34 */
{"ftime",       NULL,           NOSYS_OBSOLETE, 0},      /* 35 */
{"sync",        lx_sync,        0,              0},      /* 36 */
{"kill",        lx_kill,        0,              2},      /* 37 */
{"rename",      lx_rename,      0,              2},      /* 38 */
{"mkdir",       mkdir,          SYS_PASSTHRU,   2},      /* 39 */
{"rmdir",       lx_rmdir,       0,              1},      /* 40 */
{"dup",         dup,            SYS_PASSTHRU,   1},      /* 41 */
{"pipe",        lx_pipe,        0,              1},      /* 42 */
{"times",       lx_times,       0,              1},      /* 43 */
{"prof",        NULL,           NOSYS_OBSOLETE, 0},      /* 44 */
{"brk",         lx_brk,         0,              1},      /* 45 */
{"setgid16",    lx_setgid16,    0,              1},      /* 46 */
{"getgid16",    lx_getgid16,    0,              0},      /* 47 */

```

```

{"signal",      lx_signal,      0,          2},      /* 48 */
{"geteuid16",  lx_geteuid16,  0,          0},      /* 49 */
{"getegid16",  lx_getegid16,  0,          0},      /* 50 */
{"acct",       NULL,           NOSYS_NO_EQUIV, 0},      /* 51 */
{"umount2",    lx_umount2,    0,          2},      /* 52 */
{"lock",       NULL,           NOSYS_OBSOLETE, 0},      /* 53 */
{"ioctl",      lx_ioctl,      0,          3},      /* 54 */
{"fcntl",      lx_fcntl,      0,          3},      /* 55 */
{"mpx",        NULL,           NOSYS_OBSOLETE, 0},      /* 56 */
{"setpgid",    lx_setpgid,    0,          2},      /* 57 */
{"ulimit",     NULL,           NOSYS_OBSOLETE, 0},      /* 58 */
{"olduname",   NULL,           NOSYS_OBSOLETE, 0},      /* 59 */
{"umask",      (int (*)())umask, SYS_PASSTHRU, 1},      /* 60 */
{"chroot",     chroot,        SYS_PASSTHRU, 1},      /* 61 */
{"ustat",      lx_ustat,      0,          2},      /* 62 */
{"dup2",       lx_dup2,       0,          2},      /* 63 */
{"getppid",    lx_getppid,    0,          0},      /* 64 */
{"getpgrp",    lx_getpgrp,    0,          0},      /* 65 */
{"setsid",     lx_setsid,     0,          0},      /* 66 */
{"sigaction",  lx_sigaction,  0,          3},      /* 67 */
{"sgetmask",   NULL,           NOSYS_OBSOLETE, 0},      /* 68 */
{"ssetmask",   NULL,           NOSYS_OBSOLETE, 0},      /* 69 */
{"setreuid16", lx_setreuid16, 0,          2},      /* 70 */
{"setregid16", lx_setregid16, 0,          2},      /* 71 */
{"sigsuspend", lx_sigsuspend, 0,          1},      /* 72 */
{"sigpending", lx_sigpending, 0,          1},      /* 73 */

```

```

{"sethostname", lx_sethostname, 0, 2}, /* 74 */
{"setrlimit", lx_setrlimit, 0, 2}, /* 75 */
{"getrlimit", lx_oldgetrlimit, 0, 2}, /* 76 */
{"getrusage", lx_getrusage, 0, 2}, /* 77 */
{"gettimeofday", lx_gettimeofday, 0, 2}, /* 78 */
{"settimeofday", lx_settimeofday, 0, 2}, /* 79 */
{"getgroups16", lx_getgroups16, 0, 2}, /* 80 */
{"setgroups16", lx_setgroups16, 0, 2}, /* 81 */
{"select", NULL, NOSYS_OBSOLETE, 0}, /* 82 */
{"symlink", symlink, SYS_PASSTHRU, 2}, /* 83 */
{"oldlstat", NULL, NOSYS_OBSOLETE, 0}, /* 84 */
{"readlink", readlink, SYS_PASSTHRU, 3}, /* 85 */
{"uselib", NULL, NOSYS_KERNEL, 0}, /* 86 */
{"swapon", NULL, NOSYS_KERNEL, 0}, /* 87 */
{"reboot", lx_reboot, 0, 4}, /* 88 */
{"readdir", lx_readdir, 0, 3}, /* 89 */
{"mmap", lx_mmap, 0, 6}, /* 90 */
{"munmap", munmap, SYS_PASSTHRU, 2}, /* 91 */
{"truncate", lx_truncate, 0, 2}, /* 92 */
{"ftruncate", lx_ftruncate, 0, 2}, /* 93 */
{"fchmod", fchmod, SYS_PASSTHRU, 2}, /* 94 */
{"fchown16", lx_fchown16, 0, 3}, /* 95 */
{"getpriority", lx_getpriority, 0, 2}, /* 96 */
{"setpriority", lx_setpriority, 0, 3}, /* 97 */
{"profil", NULL, NOSYS_NO_EQUIV, 0}, /* 98 */
{"statfs", lx_statfs, 0, 2}, /* 99 */

```

```

{"fstatfs",      lx_fstatfs,      0,          2},      /* 100 */
{"ioperm",      NULL,            NOSYS_NO_EQUIV, 0},      /* 101 */
{"socketcall",  lx_socketcall,  0,          2},      /* 102 */
{"syslog",      NULL,            NOSYS_KERNEL,  0},      /* 103 */
{"setitimer",   lx_setitimer,   0,          3},      /* 104 */
{"getitimer",   getitimer,      SYS_PASSTHRU,  2},      /* 105 */
{"stat",        lx_stat,        0,          2},      /* 106 */
{"lstat",       lx_lstat,       0,          2},      /* 107 */
{"fstat",       lx_fstat,       0,          2},      /* 108 */
{"uname",       NULL,            NOSYS_OBSOLETE, 0},      /* 109 */
{"oldiopl",     NULL,            NOSYS_NO_EQUIV, 0},      /* 110 */
{"vhangup",    lx_vhangup,    0,          0},      /* 111 */
{"idle",        NULL,            NOSYS_NO_EQUIV, 0},      /* 112 */
{"vm86old",     NULL,            NOSYS_OBSOLETE, 0},      /* 113 */
{"wait4",       lx_wait4,       0,          4},      /* 114 */
{"swapoff",     NULL,            NOSYS_KERNEL,  0},      /* 115 */
{"sysinfo",     lx_sysinfo,     0,          1},      /* 116 */
{"ipc",         lx_ipc,         0,          5},      /* 117 */
{"fsync",       lx_fsync,       0,          1},      /* 118 */
{"sigreturn",   lx_sigreturn,   0,          1},      /* 119 */
{"clone",       lx_clone,       0,          5},      /* 120 */
{"setdomainname", lx_setdomainname, 0,          2},      /* 121 */
{"uname",       lx_uname,       0,          1},      /* 122 */
{"modify_ldt",  lx_modify_ldt,  0,          3},      /* 123 */
{"adjtimex",    lx_adjtimex,    0,          1},      /* 124 */
{"mprotect",    lx_mprotect,    0,          3},      /* 125 */

```

```

{"sigprocmask", lx_sigprocmask, 0,          3},      /* 126 */
{"create_module", NULL,          NOSYS_KERNEL, 0},      /* 127 */
{"init_module", NULL,          NOSYS_KERNEL, 0},      /* 128 */
{"delete_module", NULL,        NOSYS_KERNEL, 0},      /* 129 */
{"get_kernel_syms", NULL,      NOSYS_KERNEL, 0},      /* 130 */
{"quotactl",    NULL,          NOSYS_KERNEL, 0},      /* 131 */
{"getpgid",    lx_getpgid,     0,          1},      /* 132 */
{"fchdir",     fchdir,        SYS_PASSTHRU, 1},      /* 133 */
{"bdflush",   NULL,          NOSYS_KERNEL, 0},      /* 134 */
{"sysfs",     lx_sysfs,       0,          3},      /* 135 */
{"personality", lx_personality, 0,          1},      /* 136 */
{"afs_syscall", NULL,        NOSYS_KERNEL, 0},      /* 137 */
{"setfsuid16", lx_setfsuid16, 0,          1},      /* 138 */
{"setfsgid16", lx_setfsgid16, 0,          1},      /* 139 */
{"llseek",    lx_llseek,     0,          5},      /* 140 */
{"getdents",  getdents,      SYS_PASSTHRU, 3},      /* 141 */
{"select",    lx_select,     0,          5},      /* 142 */
{"flock",     lx_flock,      0,          2},      /* 143 */
{"msync",     lx_msync,      0,          3},      /* 144 */
{"readv",     lx_readv,      0,          3},      /* 145 */
{"writev",    lx_writev,     0,          3},      /* 146 */
{"getsid",    lx_getsid,     0,          1},      /* 147 */
{"fdatasync", lx_fdatasync,   0,          1},      /* 148 */
{"sysctl",    lx_sysctl,     0,          1},      /* 149 */
{"mlock",     lx_mlock,      0,          2},      /* 150 */
{"munlock",   lx_munlock,    0,          2},      /* 151 */

```



```

{"mlockall",    lx_mlockall,    0,    1},    /* 152 */
{"munlockall",  lx_munlockall,  0,    0},    /* 153 */
{"sched_setparam", lx_sched_setparam, 0,    2},    /* 154 */
{"sched_getparam", lx_sched_getparam, 0,    2},    /* 155 */
{"sched_setscheduler", lx_sched_setscheduler, 0, 3},    /* 156 */
{"sched_getscheduler", lx_sched_getscheduler, 0, 1},    /* 157 */
{"sched_yield", (int (*)( ))yield, SYS_PASSTHRU, 0},    /* 158 */
{"sched_get_priority_max", lx_sched_get_priority_max, 0, 1}, /* 159
*/

{"sched_get_priority_min", lx_sched_get_priority_min, 0, 1}, /* 160
*/

{"sched_rr_get_interval", lx_sched_rr_get_interval, 0, 2}, /* 161
*/

{"nanosleep",    nanosleep,    SYS_PASSTHRU,    2},    /* 162 */
{"mremap",       NULL,          NOSYS_NO_EQUIV,  0},    /* 163 */
{"setresuid16",  lx_setresuid16, 0,    3},    /* 164 */
{"getresuid16",  lx_getresuid16, 0,    3},    /* 165 */
{"vm86",        NULL,          NOSYS_NO_EQUIV,  0},    /* 166 */
{"query_module", lx_query_module, NOSYS_KERNEL, 5},    /* 167 */
{"poll",        lx_poll,        0,    3},    /* 168 */
{"nfsservctl",  NULL,          NOSYS_KERNEL,    0},    /* 169 */
{"setresgid16", lx_setresgid16, 0,    3},    /* 170 */
{"getresgid16", lx_getresgid16, 0,    3},    /* 171 */
{"prctl",       NULL,          NOSYS_UNDOC,     0},    /* 172 */
{"rt_sigreturn", lx_rt_sigreturn, 0,    0},    /* 173 */
{"rt_sigaction", lx_rt_sigaction, 0,    4},    /* 174 */

```

```

{"rt_sigprocmask", lx_rt_sigprocmask, 0,      4},      /* 175 */
{"rt_sigpending", lx_rt_sigpending, 0,      2},      /* 176 */
{"rt_sigtimedwait", lx_rt_sigtimedwait, 0,  4},      /* 177 */
{"sigqueueinfo", NULL,      NOSYS_UNDOC, 0},      /* 178 */
{"rt_sigsuspend", lx_rt_sigsuspend, 0,      2},      /* 179 */
{"pread64",      lx_pread64, 0,      5},      /* 180 */
{"pwrite64",     lx_pwrite64, 0,      5},      /* 181 */
{"chown16",      lx_chown16, 0,      3},      /* 182 */
{"getcwd",       lx_getcwd, 0,      2},      /* 183 */
{"capget",       NULL,      NOSYS_NO_EQUIV, 0}, /* 184 */
{"capset",       NULL,      NOSYS_NO_EQUIV, 0}, /* 185 */
{"sigaltstack", lx_sigaltstack, 0,      2},      /* 186 */
{"sendfile",     lx_sendfile, 0,      4},      /* 187 */
{"getpmsg",      NULL,      NOSYS_OBSOLETE, 0}, /* 188 */
{"putpmsg",      NULL,      NOSYS_OBSOLETE, 0}, /* 189 */
{"vfork",       lx_vfork, 0,      0},      /* 190 */
{"getrlimit",   lx_getrlimit, 0,      2},      /* 191 */
{"mmap2",       lx_mmap2, EBP_HAS_ARG6, 6},      /* 192 */
{"truncate64",  lx_truncate64, 0,      3},      /* 193 */
{"ftruncate64", lx_ftruncate64, 0,      3},      /* 194 */
{"stat64",      lx_stat64, 0,      2},      /* 195 */
{"lstat64",     lx_lstat64, 0,      2},      /* 196 */
{"fstat64",     lx_fstat64, 0,      2},      /* 197 */
{"lchown",      lchown,      SYS_PASSTHRU, 3},      /* 198 */
{"getuid",      (int (*)())getuid, SYS_PASSTHRU, 0}, /* 199 */
{"getgid",      (int (*)())getgid, SYS_PASSTHRU, 0}, /* 200 */

```

```

{"geteuid",      lx_geteuid,      0,      0},      /* 201 */
{"getegid",     lx_getegid,     0,      0},      /* 202 */
{"setreuid",    setreuid,      SYS_PASSTHRU, 0},      /* 203 */
{"setregid",    setregid,      SYS_PASSTHRU, 0},      /* 204 */
{"getgroups",   getgroups,     SYS_PASSTHRU, 2},      /* 205 */
{"setgroups",   lx_setgroups,  0,      2},      /* 206 */
{"fchown",     lx_fchown,     0,      3},      /* 207 */
{"setresuid",   lx_setresuid,  0,      3},      /* 208 */
{"getresuid",   lx_getresuid,  0,      3},      /* 209 */
{"setresgid",   lx_setresgid,  0,      3},      /* 210 */
{"getresgid",   lx_getresgid,  0,      3},      /* 211 */
{"chown",      lx_chown,      0,      3},      /* 212 */
{"setuid",     setuid,        SYS_PASSTHRU, 1},      /* 213 */
{"setgid",     setgid,        SYS_PASSTHRU, 1},      /* 214 */
{"setfsuid",   lx_setfsuid,   0,      1},      /* 215 */
{"setfsgid",   lx_setfsgid,   0,      1},      /* 216 */
{"pivot_root", NULL,          NOSYS_KERNEL, 0},      /* 217 */
{"mincore",    mincore,      SYS_PASSTHRU, 3},      /* 218 */
{"madvise",    lx_madvise,   0,      3},      /* 219 */
{"getdents64", lx_getdents64, 0,      3},      /* 220 */
{"fcntl64",    lx_fcntl64,   0,      3},      /* 221 */
{"tux",        NULL,          NOSYS_NO_EQUIV, 0}, /* 222 */
{"security",   NULL,          NOSYS_NO_EQUIV, 0}, /* 223 */
{"gettid",     lx_gettid,     0,      0},      /* 224 */
{"readahead",  NULL,          NOSYS_NO_EQUIV, 0}, /* 225 */
{"setxattr",   NULL,          NOSYS_NO_EQUIV, 0}, /* 226 */

```

```

{"lsetxattr",    NULL,                NOSYS_NO_EQUIV, 0},    /* 227 */
{"fsetxattr",    NULL,                NOSYS_NO_EQUIV, 0},    /* 228 */
{"getxattr",     NULL,                NOSYS_NO_EQUIV, 0},    /* 229 */
{"lgetxattr",    NULL,                NOSYS_NO_EQUIV, 0},    /* 230 */
{"fgetxattr",    NULL,                NOSYS_NO_EQUIV, 0},    /* 231 */
{"listxattr",    NULL,                NOSYS_NO_EQUIV, 0},    /* 232 */
{"llistxattr",   NULL,                NOSYS_NO_EQUIV, 0},    /* 233 */
{"flistxattr",   NULL,                NOSYS_NO_EQUIV, 0},    /* 234 */
{"removexattr",  NULL,                NOSYS_NO_EQUIV, 0},    /* 235 */
{"lremovexattr", NULL,                NOSYS_NO_EQUIV, 0},    /* 236 */
{"fremovexattr", NULL,                NOSYS_NO_EQUIV, 0},    /* 237 */
{"tkill",        lx_tkill,          0,                2},    /* 238 */
{"sendfile64",   lx_sendfile64,    0,                4},    /* 239 */
{"futex",        lx_futex,          EBP_HAS_ARG6,    6},    /* 240 */
{"sched_setaffinity", lx_sched_setaffinity, 0, 3},    /* 241 */
{"sched_getaffinity", lx_sched_getaffinity, 0, 3},    /* 242 */
{"set_thread_area", lx_set_thread_area, 0, 1},    /* 243 */
{"get_thread_area", lx_get_thread_area, 0, 1},    /* 244 */
{"io_setup",     NULL,                NOSYS_NO_EQUIV, 0},    /* 245 */
{"io_destroy",   NULL,                NOSYS_NO_EQUIV, 0},    /* 246 */
{"io_getevents", NULL,                NOSYS_NO_EQUIV, 0},    /* 247 */
{"io_submit",    NULL,                NOSYS_NO_EQUIV, 0},    /* 248 */
{"io_cancel",    NULL,                NOSYS_NO_EQUIV, 0},    /* 249 */
{"fadvise64",    NULL,                NOSYS_UNDOC,     0},    /* 250 */
{"nosys",        NULL,                0,                0},    /* 251 */
{"group_exit",   lx_group_exit,    0,                1},    /* 252 */

```

```

{"lookup_dcookie", NULL,          NOSYS_NO_EQUIV, 0},      /* 253 */
{"epoll_create", NULL,           NOSYS_NO_EQUIV, 0},      /* 254 */
{"epoll_ctl",    NULL,           NOSYS_NO_EQUIV, 0},      /* 255 */
{"epoll_wait",   NULL,           NOSYS_NO_EQUIV, 0},      /* 256 */
{"remap_file_pages", NULL,       NOSYS_NO_EQUIV, 0},      /* 257 */
{"set_tid_address", lx_set_tid_address, 0,    1},      /* 258 */
{"timer_create", NULL,           NOSYS_UNDOC,    0},      /* 259 */
{"timer_settime", NULL,           NOSYS_UNDOC,    0},      /* 260 */
{"timer_gettime", NULL,           NOSYS_UNDOC,    0},      /* 261 */
{"timer_getoverrun", NULL,        NOSYS_UNDOC,    0},      /* 262 */
{"timer_delete", NULL,           NOSYS_UNDOC,    0},      /* 263 */
{"clock_settime", lx_clock_settime, 0,    2},      /* 264 */
{"clock_gettime", lx_clock_gettime, 0,    2},      /* 265 */
{"clock_getres",  lx_clock_getres,  0,    2},      /* 266 */
{"clock_nanosleep", lx_clock_nanosleep, 0,    4},      /* 267 */
{"statfs64",      lx_statfs64,    0,    2},      /* 268 */
{"fstatfs64",    lx_fstatfs64,    0,    2},      /* 269 */
{"tgkill",       lx_tgkill,       0,    3},      /* 270 */

/* The following system calls only exist in kernel 2.6 and greater */
{"utimes",       utimes,          SYS_PASSTHRU,   2},      /* 271 */
{"fadvise64_64", NULL,           NOSYS_NULL,     0},      /* 272 */
{"vserver",      NULL,           NOSYS_NULL,     0},      /* 273 */
{"mbind",        NULL,           NOSYS_NULL,     0},      /* 274 */
{"get_mempolicy", NULL,          NOSYS_NULL,     0},      /* 275 */
{"set_mempolicy", NULL,          NOSYS_NULL,     0},      /* 276 */
{"mq_open",      NULL,           NOSYS_NULL,     0},      /* 277 */

```

```

{"mq_unlink",    NULL,          NOSYS_NULL,    0},    /* 278 */
{"mq_timedsend", NULL,          NOSYS_NULL,    0},    /* 279 */
{"mq_timedreceive", NULL,        NOSYS_NULL,    0},    /* 280 */
{"mq_notify",   NULL,          NOSYS_NULL,    0},    /* 281 */
{"mq_getsetattr", NULL,        NOSYS_NULL,    0},    /* 282 */
{"kexec_load",  NULL,          NOSYS_NULL,    0},    /* 283 */
{"waitid",      lx_waitid,    0,            4},    /* 284 */
{"sys_setatroot", NULL,        NOSYS_NULL,    0},    /* 285 */
{"add_key",     NULL,          NOSYS_NULL,    0},    /* 286 */
{"request_key", NULL,          NOSYS_NULL,    0},    /* 287 */
{"keyctl",      NULL,          NOSYS_NULL,    0},    /* 288 */
{"ioprio_set",  NULL,          NOSYS_NULL,    0},    /* 289 */
{"ioprio_get",  NULL,          NOSYS_NULL,    0},    /* 290 */
{"inotify_init", NULL,        NOSYS_NULL,    0},    /* 291 */
{"inotify_add_watch", NULL,      NOSYS_NULL,    0},    /* 292 */
{"inotify_rm_watch", NULL,      NOSYS_NULL,    0},    /* 293 */
{"migrate_pages", NULL,      NOSYS_NULL,    0},    /* 294 */
{"openat",      lx_openat,    0,            4},    /* 295 */
{"mkdirat",     lx_mkdirat,  0,            3},    /* 296 */
{"mknodat",     lx_mknodat,  0,            4},    /* 297 */
{"fchownat",    lx_fchownat, 0,            5},    /* 298 */
{"futimesat",   lx_futimesat, 0,            3},    /* 299 */
{"fstatat64",   lx_fstatat64, 0,            4},    /* 300 */
{"unlinkat",    lx_unlinkat, 0,            3},    /* 301 */
{"renameat",    lx_renameat, 0,            4},    /* 302 */
{"linkat",      lx_linkat,   0,            5},    /* 303 */

```

```
    {"symlinkat",    lx_symlinkat,    0,          3},    /* 304 */
    {"readlinkat",  lx_readlinkat,  0,          4},    /* 305 */
    {"fchmodat",    lx_fchmodat,    0,          4},    /* 306 */
    {"faccessat",   lx_faccessat,   0,          4},    /* 307 */
    {"pselect6",    NULL,           NOSYS_NULL,  0},    /* 308 */
    {"ppoll",       NULL,           NOSYS_NULL,  0},    /* 309 */
    {"unshare",     NULL,           NOSYS_NULL,  0},    /* 310 */
    {"set_robust_list", NULL,         NOSYS_NULL,  0},    /* 311 */
    {"get_robust_list", NULL,         NOSYS_NULL,  0},    /* 312 */
    {"splice",      NULL,           NOSYS_NULL,  0},    /* 313 */
    {"sync_file_range", NULL,         NOSYS_NULL,  0},    /* 314 */
    {"tee",         NULL,           NOSYS_NULL,  0},    /* 315 */
    {"vmsplice",    NULL,           NOSYS_NULL,  0},    /* 316 */
    {"move_pages",  NULL,           NOSYS_NULL,  0},    /* 317 */
};
```

**EXHIBIT E: SOURCE CODE FOR SLOCCOUNTER.PY AND
SLOCCOUNTERTOTAL.PY**


```

Jul 28, 11 15:57 Allclasses.py Page 1/2
'''
Created as part of work on expert report
for Google/Oracle for GreenbergTraurig

5 @author: ola
@copyright: owen astrachan, compsciconsulting
'''

import os,collections

10 acdict = collections.defaultdict(int)
aperclass = collections.defaultdict(int)
aset = set()
apack = {}

15

jcdict = collections.defaultdict(int)
jperclass = collections.defaultdict(int)
jset = set()
20 jpack = {}

public_ids = [ "public class",
               "public abstract class",
               "public interface",
25 "protected class",
               "protected",
               "public" ]

def do_one(onepath,cset,cpack):
30     if not onepath.endswith(".java"):
         return True
     if onepath.endswith("package-info.java"):
         return True
     f = open(onepath)

35     pcount = 0
     first = True
     public = False
     for line in f:

40         line = line.strip()

         if first and line.startswith("public class "):
             #print "class",onepath,line
             base = os.path.basename(onepath)
             cset.add(base)
             cpack[base] = onepath
             break

50     return public

def topcount(basepath,packname,cset,cpack):
    parts = packname.split(".")
    pathize = '/'.join(parts)
    packagepath = os.path.join(basepath,pathize)
    for top in os.listdir(packagepath):
        top_path = os.path.join(packagepath,top)
        if os.path.isdir(top_path):
            #print "*** %s is a directory in %s" % (top,packagepath)
            pass
        else:
            c = do_one(top_path,cset,cpack)
            if not c:
                #print "no public",top_path,top
                pass
            #print "%s has %d public" % (top_path,c)

65 def revs(s):
    return s[::-1]

70 def analyze():

    jpath = "/Users/ola/expert/google/ESOURCE/j2se/src/share/classes"
    apath = "/Users/ola/expert/google/SOURCE/libcore/luni/src/main/java"

```

```

Jul 28, 11 15:57 Allclasses.py Page 2/2

75     packages = [ "java.awt.font",
                  "java.beans",
                  "java.io",
                  "java.lang",
                  "java.lang.annotation",
                  "java.lang.ref",
80                  "java.lang.reflect",
                  "java.math",
                  "java.net",
                  "java.nio",
                  "java.nio.channels",
                  "java.nio.channels.spi",
85                  "java.nio.charset",
                  "java.nio.charset.spi",
                  "java.security",
                  "java.security.acl",
                  "java.security.cert",
90                  "java.security.interfaces",
                  "java.security.spec",
                  "java.sql",
                  "java.text",
                  "java.util",
95                  "# java.util.concurrent",
                  "# java.util.concurrent.atomic",
                  "# java.util.concurrent.locks",
                  "java.util.jar",
                  "java.util.logging",
100                  "java.util.prefs",
                  "java.util.regex",
                  "java.util.zip",
                  "javax.crypto",
105                  "javax.crypto.interfaces",
                  "javax.crypto.spec",
                  "javax.net",
                  "javax.net.ssl",
                  "javax.security.auth",
110                  "javax.security.auth.callback",
                  "javax.security.auth.login",
                  "javax.security.auth.x500",
                  "javax.security.cert",
                  "javax.sql",
115                  "javax.xml",
                  "javax.xml.datatype",
                  "javax.xml.namespace",
                  "javax.xml.parsers",
                  "javax.xml.transform",
120                  "javax.xml.transform.dom",
                  "javax.xml.transform.sax",
                  "javax.xml.transform.stream",
                  "javax.xml.validation",
                  "javax.xml.xpath"
125     ]

    for pack in packages:
        topcount(apath,pack,aset,apack)
        topcount(jpath,pack,jset,jpack)

130

    names = sorted(jset,key=revs)
    for i,name in enumerate(names):
        print "%d\t%s" % (i,name)

135

    for i,name in enumerate(names):
        print "%d %s" % (i,jpack[name])

140 if __name__ == "__main__":
    analyze()

```


Jul 28, 11 15:57	APIanalyzer.py	Page 3/3
150	"java.nio.channels.spi", "java.nio.charset", "java.nio.charset.spi", "java.security", "java.security.acl", "java.security.cert", "java.security.interfaces", "java.security.spec", "java.sql", 155 "java.text", "java.util", #" java.util.concurrent", #" java.util.concurrrent.atomic", #" java.util.concurrent.locks", 160 "java.util.jar", "java.util.logging", "java.util.prefs", "java.util.regex", "java.util.zip", 165 "javax.crypto", "javax.crypto.interfaces", "javax.crypto.spec", "javax.net", "javax.net.ssl", 170 "javax.security.auth", "javax.security.auth.callback", "javax.security.auth.login", "javax.security.auth.x500", "javax.security.cert", 175 "javax.sql", "javax.xml", "javax.xml.datatype", "javax.xml.namespace", "javax.xml.parsers", 180 "javax.xml.transform", "javax.xml.transform.dom", "javax.xml.transform.sax", "javax.xml.transform.stream", "javax.xml.validation", 185 "javax.xml.xpath"] c = 0 for pack in packages: 190 topcount(apath,pack,acdict,aperclass,aset) c += 1 topcount(javapath,pack,jcdict,jperclass,jset) #topcount(gnupath,pack,acdict,aperclass,aset) 195 print "%d packages analyzed" % (len(packages)) print "\nJava Analysis" report(jcdict,jperclass) print "\nAndroid Analysis" report(acdict,aperclass) 200 print "\n-----" print "common package/private" inter = jset&aset for name in inter: 205 print name print "\nAndroid\n-----" for name in aset: print name 210 print "\nJava\n-----" for name in jset: print name 215 if __name__ == "__main__": analyze()	

Jul 28, 11 15:57	ClassCounter.py	Page 1/2
	''' Created as part of work on expert report for Google/Oracle for GreenbergTraurig 5 @author: ola @copyright: owen astrachan, compsciconsulting ''' import os,collections 10 acdict = collections.defaultdict(int) aperclass = collections.defaultdict(int) aset = set() apack = {} 15 jcdict = collections.defaultdict(int) jperclass = collections.defaultdict(int) jset = set() 20 jpack = {} public_ids = ["public class", "public abstract class", "public interface"] 25 logger = open("classlog","w") def is_class(line): 30 for pub in public_ids: if line.startswith(pub): return True 35 pin = line.find("public") cin = line.find("class") if pin != -1 and cin != -1 and pin < cin: return True return False 40 def do_one(onepath,cset,cpack): if not onepath.endswith(".java"): return True if onepath.endswith("package-info.java"): 45 return True f = open(onepath) pcount = 0 first = True public = False 50 for line in f: line = line.strip() 55 if first and is_class(line): #print "class",onepath,line base = os.path.basename(onepath) cset.add(onepath) cpack[base] = onepath public = True 60 break return public 65 def topcount(basepath,cset,cpack): ''' 70 counting public classes and interfaces, use basepath as /java or /javax ''' for top in os.listdir(basepath): top_path = os.path.join(basepath,top)	

```

Jul 28, 11 15:57 ClassCounter.py Page 2/2
75     if os.path.isdir(top_path):
        #print "*** %s is a directory in %s" % (top,packagepath)

        print "recurse on ",top_path
        logger.write("recurse on "+top_path+"\n")
        topcount(top_path,cset,cpack)
80     else:
        c = do_one(top_path,cset,cpack)
        if not c:
            #print "no public",top_path,top
            pass
85     #print "%s has %d public" % (top_path,c)
def revs(s):
    return s[::-1]

def analyze():
90     jpath = "/Users/ola/expert/google/ESOURCE/j2se/src/share/classes"
    apath = "/Users/ola/expert/google/SOURCE/libcore/luni/src/main/java"

    for toplevel in [""]:
95         topcount(jpath+toplevel,jset,jpack)

    names = sorted(jset)
    for i,name in enumerate(names):
100        logger.write(str(i)+"\t"+name+"\n")
        print "%d\t%s" % (i,name)

    logger.close()

105 if __name__ == "__main__":
    analyze()

```

```

Jul 28, 11 15:57 PackageCounter.py Page 1/1
'''
Created as part of work on expert report
for Google/Oracle for GreenbergTraurig
5  @author: ola
  @copyright: owen astrachan, compsciconsulting
'''

import os,collections

10  acdict = collections.defaultdict(int)
    aperclass = collections.defaultdict(int)
    aset = set()
    apack = {}

15

    jcdict = collections.defaultdict(int)
    jperclass = collections.defaultdict(int)
    jset = set()
20  jpack = {}

    logger = open("packagelog","w")

25

def topcount(basepath,cset,cpack):
    '''
    looking for package names, find .java file, it's a package
30    '''

    jfound = False
    for top in os.listdir(basepath):
        top_path = os.path.join(basepath,top)
35        if os.path.isdir(top_path):
            #print "*** %s is a directory in %s" % (top,packagepath)

            #print "recurse on ",top_path
            #logger.write("recurse on "+top_path+"\n")
            topcount(top_path,cset,cpack)
40        else:
            if top_path.endswith(".java"):
                jfound = True

45        if jfound:
            #print ".java found in ",top_path
            cset.add(basepath)

50 def analyze():

    jpath = "/Users/ola/expert/google/ESOURCE/j2se/src/share/classes"
    apath = "/Users/ola/expert/google/SOURCE/" #libcore/luni/src/main/java"

55    packdirs = ["frameworks/base/core/java", "libcore/luni/src/main"]

    for dir in packdirs:
        #topcount(jpath,jset,jpack)
        topcount(apath+dir,aset,apack)
60        #topcount(apath,pack,aset)

    names = sorted(aset)
    for i,name in enumerate(names):
65        logger.write(str(i)+"\t"+name+"\n")
        print "%d\t%s" % (i,name)

    logger.close()

70 if __name__ == "__main__":
    analyze()

```

```

Jul 28, 11 15:57      PublicPrivateAnalyzer.py      Page 1/6
'''
Created as part of work on expert report
for Google/Oracle for GreenbergTraurig

5  @author: ola
   @copyright: owen astrachan, compsciconsulting
'''
import os,collections

10 acdict = collections.defaultdict(int)
   aperclass = collections.defaultdict(int)
   aprivdict = {}
   aset = set()
   amethnames = []

15 jcdict = collections.defaultdict(int)
   jperclass = collections.defaultdict(int)
   jprivdict = {}
   jset = set()
   jmethnames = []

   gcdict = collections.defaultdict(int)
   gperclass = collections.defaultdict(int)
   gprivdict = {}
25 gset = set()
   gmethnames = []

   afuncnlist = []
30 jfuncnlist = []
   gfuncnlist = []

   methnames = []

35 public_ids = ["public class",
               "public abstract class",
               "public interface",
               "protected class",
               "protected",
40               "public"]

def is_func(line):
   if "new" in line:
       return False
45   parts = line.split()
   if line.startswith("public") and line.find("(") >= 0 and line.find(")") >= 0:
       return True
   if line.startswith("private") and line.find("(") >= 0 and line.find(")") >= 0:
       return True
50   return False

def getClass(path):
'''
55   path ends with .java, return class name preceding .java including preceding .
   e.g., for java/lang/Arrays, return .Arrays
'''
   nm = path[:-5]
   index = nm.rfind("/")
60   return "."+nm[index+1:]

def do_one(packname, onepath, cdict, perclass, cset, funclist, privdict, methnames):
65   if not onepath.endswith(".java"):
       return True
   if onepath.endswith("package-info.java"):
       return True
   f = open(onepath)
70   class_name = getClass(onepath)

   pcount = 0

```

```

Jul 28, 11 15:57      PublicPrivateAnalyzer.py      Page 2/6
75   first = True
   public = False
   pubf = 0
   privf = 0
   for line in f:

80       line = line.strip()

       if is_func(line):
           methnames.append(line)
           if line.startswith("public"):
25         pubf += 1
           else:
               privf += 1
               nm = packname+class_name
               if not nm in privdict:
30                 privdict[nm] = []
               privdict[nm].append(line)

           if first and line.startswith("class "):
               #print "class", onepath, line
           base = os.path.basename(onepath)
           cset.add(base)

   pfound = False
   for pub in public_ids:
40         if line.startswith(pub):
             if first:
                 first = False
                 if line.find("public") >= 0 or line.find("protected") >= 0:
50                   public = True
                 else:
                     print "big problem", onepath, pub, line
                     if line.find("protected") < 0:
                         pcount += 1
                         cdict[pub] += 1
                         pfound = True
                         if line.find("class") >= 0 and line.find("extends") >= 0:
                             cdict["extends"] += 1
                             elif line.find("interface") >= 0 and line.find("extends") >= 0:
                             cdict["extends"] += 1
115                    break

   f.close()

120   perclass[pcount] += 1
   if pcount == 0:
       #print "%s = %d" % (onepath, pcount)
       pass

125   funclist.append((pubf, privf))
   return public

def topcount(basepath, packname, cdict, perclass, cset, funclist, privdict, methnames):
130   parts = packname.split(".")
   pathize = '/'.join(parts)
   packagepath = os.path.join(basepath, pathize)
   for top in os.listdir(packagepath):
       top_path = os.path.join(packagepath, top)
       if os.path.isdir(top_path):
135         #print "*** %s is a directory in %s" % (top, packagepath)
           pass
       else:
           c = do_one(packname, top_path, cdict, perclass, cset, funclist, privdict, m
ethnames)
           if not c:
               #print "no public", top_path, top
               pass
           #print "%s has %d public" % (top_path, c)

140   def func_stats(coll):
       low = 0
145

```

```

Jul 28, 11 15:57      PublicPrivateAnalyzer.py      Page 3/6

    word_total = 0
    wt_count = 0
    nonlow = 0
    getter = 0
150    setter = 0
    req = 0

    obj_names = ["toString", "hashCode", "notifyAll", "getClass"]

155    for nm in coll:
        if nm.islower():
            low += 1
            #print "\t lower",nm
        else:
160            wc = 0
            for i,ch in enumerate(nm):
                if ch.isupper() and i > 0 and nm[i-1].islower():
                    wc += 1

165            wc += 1
            #word_total += wc
            nonlow += 1

            if nm.startswith("get"):
170                getter += 1
            elif nm.startswith("set"):
                setter += 1
            elif nm in obj_names:
                req += 1
175            else:
                word_total += wc
                wt_count += 1

    print "total = %d, one = %d more = %d\n" % (nonlow+low, low, nonlow)
180    print "perc = %f avg = %f\n" % (1.0*low/(low+nonlow), 1.0*word_total/wt_count)
    print "non simple = %d\n" % (wt_count)

    print "getter = %d, setter = %d, req = %d, total = %d\n" % (getter, setter, req, req+getter+se
tter)

185 def funcalyze(methnames):
    all_names = set()
    names = []
    for meth in methnames:
        if meth.startswith("public"):
190            nameEnd = meth.find("(")
            if nameEnd == -1:
                print "error on ",meth
            else:
                name = meth[:nameEnd]
200            space = name.rfind(" ")
            mname = name[space+1:]
            all_names.add(mname)
            names.append(mname)

    print "total = %d, unique = %d\n" % (len(names), len(all_names))
    print "unique"
    func_stats(all_names)
    print "total"
    func_stats(names)

205    meth_counts = [(names.count(nm),nm) for nm in all_names]
    smc = sorted(meth_counts, reverse=True)
    print "top func occurrences"
    for pair in smc[:20]:
210        print pair

    return all_names

215

def report(cdect,perclass,funclist,privdict,methnames):

```

```

Jul 28, 11 15:57      PublicPrivateAnalyzer.py      Page 4/6

    uset = funcalyze(methnames)

220

    cttotal = 0
    for key in cdect:
        if key.find("public") < 0:
225            continue
        print "%s occurrences = %d" % (key,cdect[key])
        if key.find("class") >= 0 or key.find("interface") >= 0:
            cttotal += cdect[key]
230    print "public class/interface total = %d" % (cttotal)

    cttotal = 0
    for key in cdect:
        if key.find("protected") < 0:
235            continue
        print "%s occurrences = %d" % (key,cdect[key])
        if key.find("class") >= 0 or key.find("interface") >= 0:
            cttotal += cdect[key]
240    print "protected class/interface total = %d" % (cttotal)

    print "per class method counts"
    print "# methods\t#classes"
    total = 0
245    levels = collections.defaultdict(int)
    levlist = [0,1,6,11,16,21,51,101,100001]
    for method_count in sorted(perclass.keys()):
        print "%d\t%d" % (method_count,perclass[method_count])
        total += method_count*perclass[method_count]
250    for lev in xrange(1,len(levlist)):
        if levlist[lev-1] <= method_count < levlist[lev]:
            levels[lev] += perclass[method_count]

    print "-----"
    print "total methods = %d" % (total)
255    print "\n---summary---"
    total = 0
    for lev in xrange(1,len(levlist)):
        print "perclass from %d to %d = %d" % (levlist[lev-1],levlist[lev]-1,levels[le
v])
        total += levels[lev]
260    print "total = %d" % (total)

    print "size of funclist = %d" % (len(funclist))
    total = 0
    totalMeths = 0
265    totalPriv = 0
    for x in funclist:
        totalMeths += x[0] + x[1]
        totalPriv += x[1]
        if x[0] != 0 or x[1] != 0:
270            total += 100.0*x[0]/(x[1]+x[0])
    print "average = %f" % (total/len(funclist))
    print "total meths = %d" % (totalMeths)
    print "total private = %d" % (totalPriv)
    return uset

275 def analyze():

    apath = "/Users/ola/expert/google/SOURCE/libcore/luni/src/main/java"
    javapath = "/Users/ola/expert/google/ESOURCE/j2se/src/share/classes"
280    gnupath = "/Users/ola/expert/google/source-gnu/classpath-0.98"

    packages = ["java.awt.font",
                "java.beans",
                "java.io",
285                "java.lang",
                "java.lang.annotation",
                "java.lang.ref",
                "java.lang.reflect",
                "java.math",

```

```

Jul 28, 11 15:57      PublicPrivateAnalyzer.py      Page 5/6
290      "java.net",
      "java.nio",
      "java.nio.channels",
      "java.nio.channels.spi",
295      "java.nio.charset",
      "java.nio.charset.spi",
      "java.security",
      "java.security.acl",
      "java.security.cert",
      "java.security.interfaces",
300      "java.security.spec",
      "java.sql",
      "java.text",
      "java.util",
      #" java.util.concurrent",
305      #" java.util.concurrrent.atomic",
      #" java.util.concurrent.locks",
      "java.util.jar",
      "java.util.logging",
      "java.util.prefs",
310      "java.util.regex",
      "java.util.zip",
      "javax.crypto",
      "javax.crypto.interfaces",
      "javax.crypto.spec",
315      "javax.net",
      "javax.net.ssl",
      "javax.security.auth",
      "javax.security.auth.callback",
      "javax.security.auth.login",
320      "javax.security.auth.x500",
      "javax.security.cert",
      "javax.sql",
      "javax.xml",
      "javax.xml.datatype",
325      "javax.xml.namespace",
      "javax.xml.parsers",
      "javax.xml.transform",
      "javax.xml.transform.dom",
      "javax.xml.transform.sax",
330      "javax.xml.transform.stream",
      "javax.xml.validation",
      "javax.xml.xpath"
      ]

335      for pack in packages:
          topcount( javapath,pack, jcdict, jperclass, jset, jfunclist, jprivdict, jmethnames)
          topcount( apath,pack, acdict, aperclass, aset, afunclist, aprivdict, amethnames)
          topcount( gnupath,pack, gcdict, gperclass, gset, gfunclist, gprivdict, gmethnames)

340      print "%d packages analyzed" % (len(packages))
      print "\nJava Analysis"
      juset = report( jcdict, jperclass, jfunclist, jprivdict, jmethnames)
      print "\nAndroid Analysis"
      auset = report( acdict, aperclass, afunclist, aprivdict, amethnames)
345      print "\nGnuClasspath Analysis"
      report( gcdict, gperclass, gfunclist, gprivdict, gmethnames)
      print "\n-----"

      jmset = juset
350      amset = auset
      inter = jmset&amset
      aonly = amset-jmset
      jonly = jmset-amset
      print "android only count = ", len(aonly), len(amset)
355      print "java only count = ", len( jonly), len( jmset)
      print "android only"
      for i,n in enumerate(sorted(aonly)):
          print i,n
      print "java only"

```

```

Jul 28, 11 15:57      PublicPrivateAnalyzer.py      Page 6/6
360      for i,n in enumerate(sorted(jonly)):
          print i,n

      privlog = open("privatelog","w")
365      for pack in aprivdict:
          if pack in jprivdict:
              line = "package class private {0!s}\n".format(pack)
              print "package class private %s" % (pack)
              privlog.write(line)
370      for priv in aprivdict[pack]:
              line = "\tAndroid {0!s}\n".format(priv)
              privlog.write(line)
              #print "\tAndroid %s" % (priv)
              if priv in jprivdict[pack]:
375                  privlog.write("\t\talso in Java\n")
                  #print "\t\talso in Java"
              for priv in jprivdict[pack]:
                  if not priv in aprivdict[pack]:
380                      privlog.write("\tJava "+priv+"\n")
                      #print "\tJava %s" % (priv)
          privlog.close()

385

      #      print "common package/private"
      #      inter = jset&aset
      #      for name in inter:
390          #          print name
          #
          #      print "\nAndroid\n-----"
          #      for name in aset:
          #          print name
395          #      print "\nJava\n-----"
          #      for name in jset:
          #          print name

400

      if __name__ == "__main__":
          analyze()

```

```

Jul 28, 11 15:57      SlocCounter.py      Page 1/2
'''
Created as part of work on expert report
for Google/Oracle for GreenbergTraurig

5  @author: ola
   @copyright: owen astrachan, compsciconsulting
'''
import os,collections

10 def do_one(packname, onepath):
    '''
    return number of lines in onepath if a .java file
    '''

15     if not onepath.endswith(".java"):
         return 0
     if onepath.endswith("package-info.java"):
         return 0

20     f = open(onepath)

     lcount = 0
     for line in f:
25         lcount += 1

     f.close()
     return lcount

30 def topcount(basepath, packname):
    parts = packname.split(".")
    pathize = '/'.join(parts)
    packagepath = os.path.join(basepath, pathize)
    total = 0
    ftot = 0
35     for top in os.listdir(packagepath):
         top_path = os.path.join(packagepath, top)
         if os.path.isdir(top_path):
             #print "*** %s is a directory in %s" % (top, packagepath)
             pass
         else:
             c = do_one(packname, top_path)
             if c != 0:
                 ftot += 1
                 total += c
45     return (total, ftot)

def analyze48():

50     apath = "/Users/ola/expert/google/SOURCE/libcore/luni/src/main/java"
     javapath = "/Users/ola/expert/google/ESOURCE/j2se/src/share/classes"
     gnupath = "/Users/ola/expert/google/source-gnu/classpath-0.98"

55     packages = [ "java.awt.font",
                  "java.beans",
                  "java.io",
                  "java.lang",
                  "java.lang.annotation",
                  "java.lang.ref",
                  "java.lang.reflect",
60                  "java.math",
                  "java.net",
                  "java.nio",
                  "java.nio.channels",
                  "java.nio.channels.spi",
65                  "java.nio.charset",
                  "java.nio.charset.spi",
                  "java.security",
                  "java.security.acl",
70                  "java.security.cert",
                  "java.security.interfaces",
                  "java.security.spec",
                  "java.sql",

```

```

Jul 28, 11 15:57      SlocCounter.py      Page 2/2
"java.text",
75  "java.util",
     #" java.util.concurrent",
     #" java.util.concurrrent.atomic",
     #" java.util.concurrent.locks",
     "java.util.jar",
80  "java.util.logging",
     "java.util.prefs",
     "java.util.regex",
     "java.util.zip",
     "javax.crypto",
85  "javax.crypto.interfaces",
     "javax.crypto.spec",
     "javax.net",
     "javax.net.ssl",
     "javax.security.auth",
90  "javax.security.auth.callback",
     "javax.security.auth.login",
     "javax.security.auth.x500",
     "javax.security.cert",
     "javax.sql",
95  "javax.xml",
     "javax.xml.datatype",
     "javax.xml.namespace",
     "javax.xml.parsers",
     "javax.xml.transform",
100  "javax.xml.transform.dom",
     "javax.xml.transform.sax",
     "javax.xml.transform.stream",
     "javax.xml.validation",
     "javax.xml.xpath"
105 ]

     jpair = [0,0]
     apair = [0,0]
     gpair = [0,0]
110     for pack in packages:
         j = topcount(javapath, pack)
         a = topcount(apath, pack)
         g = topcount(gnupath, pack)
         jpair[0] += j[0]
115         jpair[1] += j[1]
         apair[0] += a[0]
         apair[1] += a[1]
         gpair[0] += g[0]
         gpair[1] += g[1]

120

     print "Java = %d %d\nAndroid = %d %d\nGnu = %d %d\n" % (jpair[0], jpair[1], apair[0], apair[1], gpair[0], gpair[1])

125

if __name__ == "__main__":
    analyze48()

```



```

Jul 28, 11 15:57      SlocCounterTotal.py      Page 1/2
'''
Created as part of work on expert report
for Google/Oracle for GreenbergTraurig

5 @author: ola
  @copyright: owen astrachan, compsciconsulting
'''
import os,collections

10 #global vars aren't used as global, but passed as parameters
   #in different calls

   #android vars

15 acdict = collections.defaultdict(int)
   aperclass = collections.defaultdict(int)
   aset = set()
   apack = {}

20 #java vars

   jcdict = collections.defaultdict(int)
   jperclass = collections.defaultdict(int)
25 jset = set()
   jpack = {}

   logger = open("sloclog","w")

30 idents = [ "AclEntryImpl.java",
             "AclImpl.java",
             "GroupImpl.java",
             "OwnerImpl.java",
             "PermissionImpl.java",
35 "PrincipleImpl.java",
             "AclEnumerator.java",
             "PolicyNodeImpl.java",
             "CodeSourceTest.java",
             "CollectionCertStoreParametersTest.java",
40 "TimSort.java",
             "ComparableTimSort.java" ]

def do_one(onepath,cset,cpack):
45     if onepath.endswith("package-info.java"):
         return 0

     endings = ["java", ".h", ".c", ".cpp"]
     ok = False
50     for e in endings:
         if onepath.endswith(e):
             ok = True
     if not ok:
55         return 0

     # for test files, uncomment below, find files with /test/ in path
     # dex = onepath.rfind("/")
     # if dex == -1:
60         print "trouble on ", onepath
     # base = onepath[:dex]
     # lst = base[base.rfind("/")+1:]
     # if lst != "test":
     #     return 0
65     # print onepath

     # for 12 files at issue in idents uncomment code
     # ok = False
70     # for ids in idents:
     #     if onepath.endswith(ids):
     #         ok = True
     # if not ok:

```

```

Jul 28, 11 15:57      SlocCounterTotal.py      Page 2/2
#         return 0
75 #
#         print onepath

         f = open(onepath)
         pcount = 0
80         for line in f:

             pcount += 1

85         #print pcount,onepath
         return pcount

90 def topcount(basepath,cset,cpack):
'''
   traverse directory structure looking for files and data via do_one
'''

95     fcount = 0
     lcount = 0

     for top in os.listdir(basepath):
100         top_path = os.path.join(basepath,top)
         if os.path.isdir(top_path):
             #print "**** %s is a directory in %s" % (top,packagepath)

             #print "recurse on ",top_path
105             #logger.write("recurse on "+top_path+"\n")
             res = topcount(top_path,cset,cpack)
             fcount += res[0]
             lcount += res[1]
         else:
110             c = do_one(top_path,cset,cpack)
             if c != 0:
                 lcount += c
                 fcount += 1

115         return (fcount,lcount)

def analyze():

120     jpath = "/Users/ola/expert/google/ESOURCE/j2se/src/share/classes"
     apath = "/Users/ola/expert/google/SOURCE/"

     #dirs = ["libcore", "frameworks/base/core"] #/luni/src/main/java"
     dirs = [""]
125     lcount = 0
     fcount = 0
     for dir in dirs:
         c = topcount(apath+dir,jset,jpack)
         fcount += c[0]
130         lcount += c[1]

     print "files=%d,lines=%d\n" % (fcount,lcount)

if __name__ == "__main__":
135     analyze()

```