

# The Good, the Bad, and the Ugly?

## Structure and Trends of Open Unix Kernels

Dr. Wolfgang Mauerer, Siemens AG, CT SE 2  
Corporate Competence Centre Embedded Linux  
wolfgang.mauerer@siemens.com

# Overview

- 1 Introduction
- 2 Evaluated Systems
- 3 Feature Comparison
- 4 Quantitative Measurements
  - Code Shape
  - Complexity
  - Size Matters
  - Dynamics
- 5 Summary

# Outline

- 1 Introduction
- 2 Evaluated Systems
- 3 Feature Comparison
- 4 Quantitative Measurements
  - Code Shape
  - Complexity
  - Size Matters
  - Dynamics
- 5 Summary

# Introduction

## Observations

- Users won't really notice different kernel underneath their systems (spark plugs!)
- Developers will mostly neither.

## Questions

Why so many different kernels?

How do they differ?

*Quantitative* assessment possible?

Which one's the best kernel?

# Introduction

## Observations

- Users won't really notice different kernel underneath their systems (spark plugs!)
- Developers (*except system and spark plug developers*) will mostly neither.

## Questions

- Why so many different kernels?
- How do they differ?
- *Quantitative* assessment possible?
- Which one's the best kernel?

## Introduction

### Observations

- Users won't really notice different kernel underneath their systems (spark plugs!)
- Developers (*except system and spark plug developers*) will mostly neither.

### Questions

- Why so many different kernels?
- How do they differ?
- *Quantitative* assessment possible?
- Which one's the best kernel?



## Introduction

### Observations

- Users won't really notice different kernel underneath their systems (spark plugs!)
- Developers (*except system and spark plug developers*) will mostly neither.

### Questions


- Why so many different kernels?
- How do they differ?
- *Quantitative* assessment possible?
- ~~Which one's the best kernel?~~



# Outline

- 1 Introduction
- 2 Evaluated Systems**
- 3 Feature Comparison
- 4 Quantitative Measurements
  - Code Shape
  - Complexity
  - Size Matters
  - Dynamics
- 5 Summary

## Evaluated Systems: Linux and OpenSolaris



- Started: 1991
- License: GPL
- Version Control: GIT



- Started: 2008 (1991)
- License: CDDL
- Version Control: Mercurial

## The BSD “family”



- Started: 1993
- License: BSD
- Version Control: Subversion



- Started: 1991
- License: BSD
- Version Control: CVS



- Started: 1996 (1991)
- License: BSD
- Version Control: CVS

# Outline

- 1 Introduction
- 2 Evaluated Systems
- 3 Feature Comparison**
- 4 Quantitative Measurements
  - Code Shape
  - Complexity
  - Size Matters
  - Dynamics
- 5 Summary

# Unbiased?

No flame wars intended. . .

## Feature Comparison

### Commonplaces

- Yes, OpenSolaris focusses on tracing, zones, and ZFS
- Yes, Linux still focusses on . . . everything
- Yes, NetBSD is really about multi-arch support
- Yes, OpenBSD is about security and paranoia
- Yes, FreeBSD would like Linux's device drivers

### Alternatives

Performance testing  
Tick list on Wikipedia

## Feature Comparison

### Commonplaces

- Yes, OpenSolaris focusses on tracing, zones, and ZFS
- Yes, Linux still focusses on . . . everything
- Yes, NetBSD is really about multi-arch support
- Yes, OpenBSD is about security and paranoia
- Yes, FreeBSD would like Linux's device drivers

### Alternatives

Performance testing  
Tick list on Wikipedia

## Feature Comparison

### Commonplaces

- Yes, OpenSolaris focusses on tracing, zones, and ZFS
- Yes, Linux still focusses on . . . everything
- Yes, NetBSD is really about multi-arch support
- Yes, OpenBSD is about security and paranoia
- Yes, FreeBSD would like Linux's device drivers

### Alternatives

Performance testing  
Tick list on Wikipedia

## Feature Comparison

### Commonplaces

- Yes, OpenSolaris focusses on tracing, zones, and ZFS
- Yes, Linux still focusses on . . . everything
- Yes, NetBSD is really about multi-arch support
- Yes, OpenBSD is about security and paranoia
- Yes, FreeBSD would like Linux's device drivers

### Alternatives

Performance testing

Tick list on Wikipedia

## Feature Comparison

### Commonplaces

- Yes, OpenSolaris focusses on tracing, zones, and ZFS
- Yes, Linux still focusses on . . . everything
- Yes, NetBSD is really about multi-arch support
- Yes, OpenBSD is about security and paranoia
- Yes, FreeBSD would like Linux's device drivers

### Alternatives

Performance testing

Tick list on Wikipedia

## Feature Comparison

### Commonplaces

- Yes, OpenSolaris focusses on tracing, zones, and ZFS
- Yes, Linux still focusses on . . . everything
- Yes, NetBSD is really about multi-arch support
- Yes, OpenBSD is about security and paranoia
- Yes, FreeBSD would like Linux's device drivers

### Alternatives

- Performance testing
- Tick list on Wikipedia

# Outline

- 1 Introduction
- 2 Evaluated Systems
- 3 Feature Comparison
- 4 Quantitative Measurements**
  - Code Shape
  - Complexity
  - Size Matters
  - Dynamics
- 5 Summary

# Quantitative measurements

## Quantitative != Useful

- Formal complexity vs. real complexity
- Line counting == bean counting?
- Reliability of direct comparisons?

## Question

Can measurements provide meaningful insights?

### A Tale of Four Kernels

Dionidis Spinellis  
Department of Management Science and Technology  
Athens University of Economics and Business  
Patission 76, GR-104 24 Athens, Greece  
dds@aubg.gr

#### ABSTRACT

The FreeBSD, GNU/Linux, Solaris, and Windows operating systems have kernels that provide comparable facilities. Interestingly, their code bases share almost no common parts, while their development processes vary dramatically. We analyze the source code of the four systems by collecting metrics in the areas of the organization, code structure, code style, the use of the C preprocessor, and data organization. The aggregate results indicate that across various areas and many different metrics, four systems developed using widely different processes score comparably. This allows us to posit that the structure and internal quality attributes of a working, non-trivial software artifact will represent first and foremost the engineering requirements of its construction, with the influence of process being marginal, if any.

#### Categories and Subject Descriptors

D.2.9 [Software Engineering]: Management—Software process models; D.2.9 [Software Engineering]: Metrics—Product metrics

#### General Terms

Measurement

#### 1. INTRODUCTION

Arguments regarding the efficacy of open source development processes often employ quantitative evidence [2], anecdotal evidence [3], or writing [1]. Although considerable research has focused on open source artifacts as a whole [9, 4], [3, 12], the direct comparison of artifacts with corresponding proprietary artifacts is an elusive goal. The recent open source code and the distribution of open source code to research institutions has opened a window of opportunity to compare the code of open

source code quality in large industrial-scale open source projects. OpenSolaris, and the Windows operating system are no significant exception between four large open source operating systems. An additional open source code system, written in C, is also included to provide a baseline for development.

2.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted, without fee, provided that copies are made or distributed for profit or commercial advantage and that the fee code for copying and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.  
© 2008 ACM 978-1-60558-079-1/08/0010...\$5.00  
Copyright 2008 ACM 978-1-60558-079-1/08/0010...\$5.00.

# Quantitative measurements

## Quantitative != Useful

- Formal complexity vs. real complexity
- Line counting == bean counting?
- Reliability of direct comparisons?

## Question

Can measurements provide *meaningful* insights?

### A Tale of Four Kernels

Dionidis Spinellis  
Department of Management Science and Technology  
Athens University of Economics and Business  
Patission 76, GR-104 34 Athens, Greece  
dds@aubg.gr

#### ABSTRACT

The FreeBSD, GNU/Linux, Solaris, and Windows operating systems have kernels that provide comparable facilities. Interestingly, their code bases share almost no common parts, while their development processes vary dramatically. We analyze the source code of the four systems by collecting metrics in the areas of the organization, code structure, code style, the use of the C preprocessor, and data organization. The aggregate results indicate that across various areas and many different metrics, four systems developed using widely different processes score comparably. This allows us to posit that the structure and internal quality attributes of a working, non-trivial software artifact will represent first and foremost the engineering requirements of its construction, with the influence of process being marginal, if any.

#### Categories and Subject Descriptors

D.2.9 [Software Engineering]: Management—Software process models; D.2.9 [Software Engineering]: Metrics—Product metrics

#### General Terms

Measurement

#### 1. INTRODUCTION

Arguments regarding the efficacy of open source development processes often employ quantitative evidence [2], anecdotal evidence [37], or writing [2]. Although considerable research has focused on open source artifacts as a whole [9, 41, 3, 12], the direct comparison of artifacts with corresponding proprietary counterparts is an elusive goal. The recent open source movement and the distribution of free and open source code to research institutions has provided a window of opportunity to compare the quality between the code of open

source and proprietary code. Here I report on code quality metrics for large industrial-scale open source systems: OpenSolaris, and the Windows operating system. I analyze the contribution of the source code to the overall size of the system between four large open source systems. An additional goal is to compare the code quality of open source code systems with that of code written in C and compiled with the same compiler, but developed using different


2

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted, without fee, provided that copies are made or distributed for profit or commercial advantage and that the fee code for copying and distribution appears on the first page. This request is granted on the condition that the fee code for copying and distribution appears on the first page. This request is granted on the condition that the fee code for copying and distribution appears on the first page. This request is granted on the condition that the fee code for copying and distribution appears on the first page.


Copyright 2006 ACM 978-1-60558-079-1/06/05...\$5.00.

# Code Shape

# Code Shape: Composition



# Code Shape: Composition





## Code shape: Measurements



### Measurements

- Line counts for
  - Input
  - Blanks
  - Comments
  - Preprocessor Statements
  - Declarative/executable instructions
- Measured per
  - File
  - Function
  - Structure/Union
- Measured over  $\approx$  5 years of history

# Code Shape: Linux



## Code Shape: FreeBSD



## Code Shape: Results

### Pattern: Documentation

Memory Management  $\approx$  Filesystems  $<$  Core Kernel  $\ll$  Device Drivers

### Pattern: Preprocessor Use

Comparable between subsystems of kernels

OpenBSD, NetBSD: More tendency towards preprocessor

### Exception: OpenSolaris

Very little preprocessor use

Documentation: core  $<$  fs  $\approx$  drivers  $<$  mm

## Code Shape: Results

### Pattern: Documentation

Memory Management  $\approx$  Filesystems  $<$  Core Kernel  $\ll$  Device Drivers

### Pattern: Preprocessor Use

- Comparable between subsystems of kernels
- OpenBSD, NetBSD: More tendency towards preprocessor

### Exception: OpenSolaris

Very little preprocessor use

Documentation: core  $<$  fs  $\approx$  drivers  $<$  mm

## Code Shape: Results

### Pattern: Documentation

Memory Management  $\approx$  Filesystems  $<$  Core Kernel  $\ll$  Device Drivers

### Pattern: Preprocessor Use

- Comparable between subsystems of kernels
- OpenBSD, NetBSD: More tendency towards preprocessor

### Exception: OpenSolaris

- Very little preprocessor use
- Documentation: core  $<$  fs  $\approx$  drivers  $<$  mm

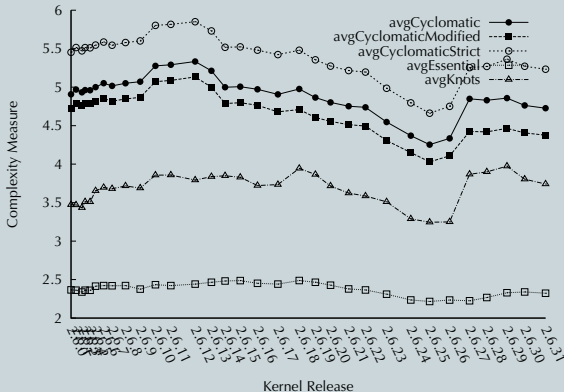
# Complexity

# Complexity Measures

## Complexities



- Cyclomatic complexity
- Essential complexity
- Maximal nesting
- Maximal essential knots
- Henry-Kafura Information Flow

# Complexity: Comparison of measures





- Similar picture for other systems
- *Long-term behaviour, not absolute numbers important*

# Complexity: NetBSD



## Discontinuities and Flips!

# Complexity: OpenBSD



## Discontinuities and Flips!

## Complexity: Interpretation

### Interpretation?

- Hard to attribute discontinuous jumps to specific causes
- Similar relations between subsystems
- Further averaging required

### Complexity ordering

Let  $A, B$  denote subsystems. Define ordering (with  $c \in ]0.5, 1]$ ):

$$A \sqsubseteq B \Leftrightarrow \frac{1}{\#\mathcal{M} \cdot \#\mathcal{R}} \sum_{r \in \mathcal{R}} \sum_{m \in \mathcal{M}} \Theta(m(B_r) - m(A_r)) \geq c$$

## Complexity: Interpretation

### Interpretation?

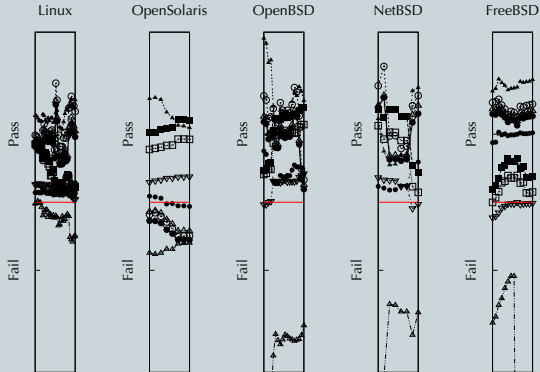
- Hard to attribute discontinuous jumps to specific causes
- Similar relations between subsystems
- Further averaging required

### Complexity ordering

Let  $A, B$  denote subsystems. Define ordering (with  $c \in ]0.5, 1]$ ):

$$A \sqsubseteq B \Leftrightarrow \frac{1}{\#\mathcal{M} \cdot \#\mathcal{R}} \sum_{r \in \mathcal{R}} \sum_{m \in \mathcal{M}} \Theta(m(B_r) - m(A_r)) \geq c$$

# Complexity Ordering: Results





Hypothesis

$$\text{Core} \sqsubseteq \text{MM} \wedge \text{MM} \sqsubseteq \text{FS}.$$

# Size Matters

# Size Matters: OpenBSD





## Results

Discontinuities attributable to specific causes

Per-subsystem consideration required

Linear vs. superlinear

# Size Matters: Linux





## Results

Discontinuities attributable to specific causes

Per-subsystem consideration required

Linear vs. superlinear

# Size Matters



## Results


- Discontinuities attributable to specific causes
- Per-subsystem consideration required
- Linear vs. superlinear

# Dynamics

## Dynamics of development


- Global analysis: Traditional method
- VCSs allow decomposition to patch level
- Good time resolution essential

# Dynamics: Linux




(Well known from lwn.net and the Linux foundation)

# Dynamics: Linux





(Well known from lwn.net and the Linux foundation)



# Dynamics: FreeBSD





# Dynamics: NetBSD




# Dynamics: OpenBSD




# Dynamics: OpenSolaris



# Dynamics: Comparison



## Dynamics: Comparison



## Conclusions


- VCS: Key influence on dynamics
- Linear vs. superlinear: Relevant?
- Maintainability criteria don't match reality




# Outline

- 1 Introduction
- 2 Evaluated Systems
- 3 Feature Comparison
- 4 Quantitative Measurements
  - Code Shape
  - Complexity
  - Size Matters
  - Dynamics
- 5 Summary**


## Summary


Single static snapshots 

Dynamics 

Metrics  

Version Control 

Maintainability Indices 

Quantitative similarities 

- Documentation, Preprocessor patterns
- Complexity relationships

# Thank you for listening!