

Listening to big data

Or, philately will get you everywhere

Mike Godfrey

Software Architecture Group

University of Waterloo



"Big data"

- Three Vs
 - Volume, Velocity, Variety
- Why?
 - Enhanced decision making, insight discovery, and process optimization
- Common problems:
 - Capture, curation, storage, search, sharing, transfer, analysis, and visualization

Overview

- Is clone analysis / empirical SE a Big Data problem?
 - ... and should we care?
- Looking hard for the Big Picture
 - And why sometimes that can be a bad idea
- Let's go swimming with the data!
 - Some experiences and some advice

(More data + simple algorithms)
>> (complex algorithms)

- Fantastic talk by Peter Norvig of Google:
 - "The unreasonable effectiveness of data"
<http://www.youtube.com/watch?v=yvDCzhbjYWs>
- "Every time I fire a linguist, my scores get better."
 - [Fred Jelinek, paraphrased]
- But does that work for clone detection / ESE too?
 - Should we all use N-gram algorithms?

Data quality



(Big data + simple algorithms)?

- NLP, for example, analyzes unstructured prose
 - Much variation: intent, word ordering, relationships, ...
 - NLP often does some pre-processing e.g., stemming
- ESE examines development artifacts with lots of internal structure + external linkage, implicit and explicit
 - Source code text, including comments
 - Version control meta-data
 - Bug reports
 - ...
- When you have reliable structure, exploit it!
 - Yes?
 - So maybe big ESE data isn't really big data ...

Looking for the Big Picture

[Trials and Errors: Why Science is Failing Us](#)

Wired Magazine, December 2011

by [Jonah Lehrer](#)

Looking for the Big Picture

A selective attention test

"I used to think that the brain was the most wonderful organ in my body. Then I realized who was telling me this."

— Emo Philips

<http://www.youtube.com/watch?v=vJG698U2Mvo>

Tim Minchin



<http://www.upworthy.com/this-is-the-most-inspiring-yet-depressing-yet-hilarious-yet-horrifying-yet-heartwarming-grad-speech>

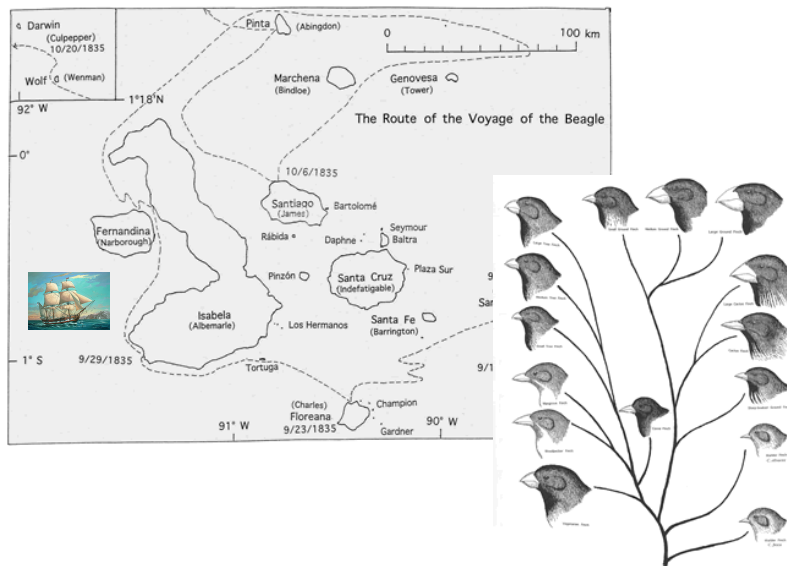
*"Physics is the only real science.
The rest are just stamp collecting."*



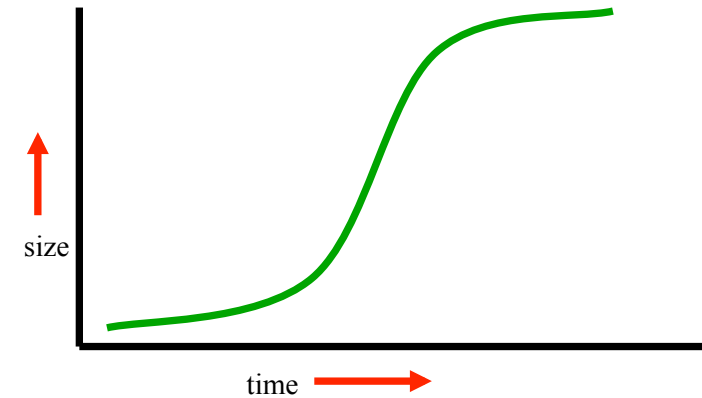
Ernest Rutherford (1871-1937)

*Father of atomic physics
Nobel prize for ... chemistry*

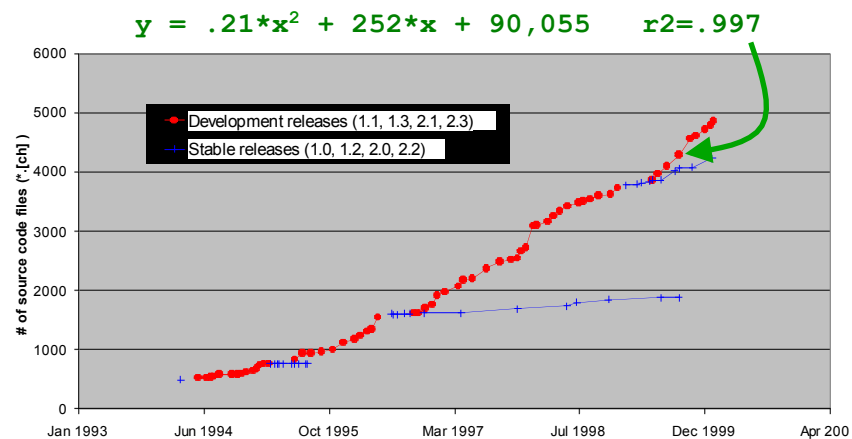




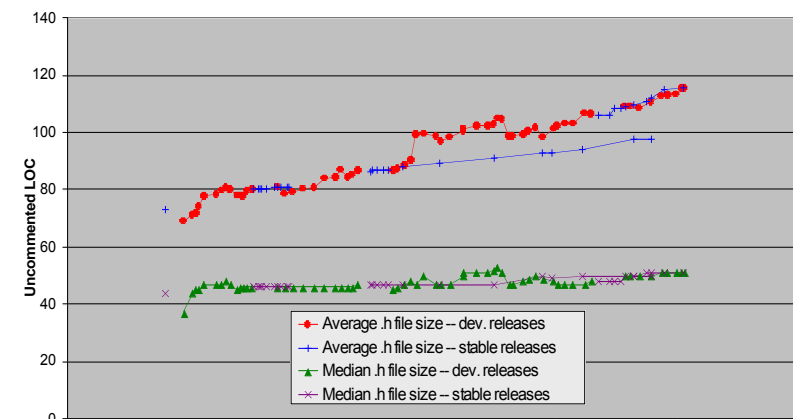
The "S" curve of successful growth



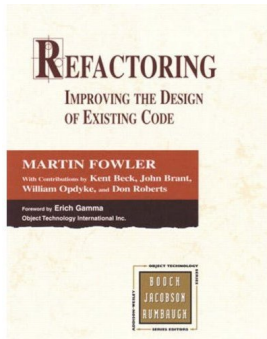
Linux kernel:
Growth of kernel src tree (# of files)



Linux kernel:
Average / median .h file size



Source code cloning



“Number one in the stink parade is duplicated code. If you see the same code structure in more than one place, you can be sure that your program will be better if you find a way to unify them.”

– “Bad Smells”

[Beck/Fowler in *Refactoring*]

'Cloning considered harmful' considered harmful

1. Forking
 - Hardware variation
e.g., Linux SCSI drivers
 - Platform variation
 - Experimental variation
2. Templating
 - Boilerplating
 - API / library protocols
 - Generalized programming idioms
 - Parameterized code
3. Post-hoc customizing
 - Bug workarounds
 - Replicate + specialize

Cloning harmfulness: Two open source case studies

Group	Pattern	Apache		Gnumeric	
		Good	Harmful	Good	Harmful
Forking	Hardware variation	0	0	0	0
Forking	Platform variation	10	0	0	0
Forking	Experimental variation	4	0	0	0
Templating	Boiler-plating	5	0	6	7
Templating	API	0	0	0	9
Templating	Idioms	0	12	1	1
Templating	Parameterized code	5	12	10	34
Customizing	Replicate + specialize	12	4	15	16
Customizing	Bug workarounds	0	0	0	0
Total		36	28	32	67

Apache httpd 2.2.4 - 60 Tokens

Gnumeric 1.6.3 - 60 Tokens

What to do?

- *Swim* with the data
- *Be* the gorilla in the mist
- *Look for lumps* under the carpet & ask "Why?"

Helper function #1

```
public int countZeros (int n) {
    if (n < 10) {
        return 0;
    } else if (n % 10 == 0) {
        return 1 + countZeros (n / 10);
    } else {
        return countZeros (n / 10);
    }
}
```

```
rascal> int i = fact(1000);
int: 4023872600770...
rascal> countZeros(i);
int: 472
```

Helper function #1, v1.1

```
public int countTrailingZeros (int n) {
    if (n < 10) {
        return 0;
    } else if (n % 10 == 0) {
        return 1 + countTrailingZeros (n / 10);
    } else {
        return 0 ;
    }
}
```

```
rascal> countTrailingZeros(i);
int: 249
```

Helper function #2 — Let's play around

```
public void printLastTwenty (int n){
    for(int i <- [n-19..n]) {
        println ("<i>! has <countTrailingZeros(fact(i))> "
            + "trailing zeros.");
    }
}
```

```
rascal>printLastTwenty(1000);
981! has 243 trailing zeros.
982! has 243 trailing zeros.
983! has 243 trailing zeros.
984! has 243 trailing zeros.
985! has 244 trailing zeros.
986! has 244 trailing zeros.
987! has 244 trailing zeros.
988! has 244 trailing zeros.
989! has 244 trailing zeros.
990! has 245 trailing zeros.
```

```
991! has 245 trailing zeros.
992! has 245 trailing zeros.
993! has 245 trailing zeros.
994! has 245 trailing zeros.
995! has 246 trailing zeros.
996! has 246 trailing zeros.
997! has 246 trailing zeros.
998! has 246 trailing zeros.
999! has 246 trailing zeros.
1000! has 249 trailing zeros.
ok
```

Looking for lumps

```
public void findLumps (int n) {
    int iMinusOneFactZeros = 0;
    for (int i <- [1..n]) {
        int iFactZeros = countTrailingZeros(fact(i));
        int diff = iFactZeros - iMinusOneFactZeros;
        if (diff >= 1) {
            println ("<diff> more zeros at <i>!");
        }
        iMinusOneFactZeros = iFactZeros;
    }
}
```

```
rascal>findLumps(1000);
1 more zeros at 5!
1 more zeros at 10!
1 more zeros at 15!
1 more zeros at 20!
2 more zeros at 25!
1 more zeros at 30!
1 more zeros at 35!
```

```
1 more zeros at 40!
1 more zeros at 45!
2 more zeros at 50!
1 more zeros at 55!
1 more zeros at 60!
1 more zeros at 65!
1 more zeros at 70!
2 more zeros at 75!
```

Looking for lumps

```
1 more zeros at 80!  
1 more zeros at 85!  
1 more zeros at 90!  
1 more zeros at 95!  
2 more zeros at 100!  
1 more zeros at 105!  
1 more zeros at 110!  
1 more zeros at 115!  
1 more zeros at 120!  
3 more zeros at 125!  
1 more zeros at 130!  
...  
1 more zeros at 245!  
3 more zeros at 250!  
1 more zeros at 255!
```

```
1 more zeros at 495!  
3 more zeros at 500!  
1 more zeros at 505!  
...  
1 more zeros at 620!  
4 more zeros at 625!  
1 more zeros at 630!  
...  
1 more zeros at 985!  
1 more zeros at 990!  
1 more zeros at 995!  
3 more zeros at 1000!  
ok
```

Helper function #3, v1.1

```
public void findLumps2 (int n, int tao) {  
    int iMinusOneFactZeros = 0;  
    for (int i <- [1..n]) {  
        int iFactZeros = countTrailingZeros(fact(i));  
        int diff = iFactZeros - iMinusOneFactZeros;  
        if (diff >= tao) {  
            println("<diff> more zeros at <i>!");  
        }  
        iMinusOneFactZeros = iFactZeros;  
    }  
}
```

- We can parameterize the threshold to look for jumps of 2, 3, or 4 zeros

Looking for lumps

```
rascal>findLumps2(1000,2);  
2 more zeros at 25!  
2 more zeros at 50!  
2 more zeros at 75!  
2 more zeros at 100!  
3 more zeros at 125!  
2 more zeros at 150!  
2 more zeros at 175!  
2 more zeros at 200!  
2 more zeros at 225!  
3 more zeros at 250!  
2 more zeros at 275!  
...  
2 more zeros at 950!  
2 more zeros at 975!  
3 more zeros at 1000!  
ok
```

```
rascal>findLumps2(1000,3);  
3 more zeros at 125!  
3 more zeros at 250!  
3 more zeros at 375!  
3 more zeros at 500!  
4 more zeros at 625!  
3 more zeros at 750!  
3 more zeros at 875!  
3 more zeros at 1000!  
ok
```

```
rascal>findLumps2(1000,4);  
4 more zeros at 625!  
ok
```

$$5^0 = 1$$

$$5^1 = 5$$

$$5^2 = 25$$

$$5^3 = 125$$

$$5^4 = 625$$

$$5^5 = 3125$$

An analytic solution

Let N be a positive integer.

Let $k = \text{floor}(\log_5 N)$

Start a counter at zero, call it nz

We want to examine $i \leftarrow [1..N]$

If i is not divisible by 5, ignore it

If i is divisible by 5, add 1 to nz

If i is also divisible by 25, add 1 more

...

If i is also divisible by 2^k , add 1 more

Final functions

```
public int predictZeros (int N) {
    int k = floorLogBase(N, 5);
    int nz = 0;
    for (int i <- [1..N] ){
        int p5 = 1;
        for (int j <- [1..k]) {
            p5 *= 5;
            if (i % p5 == 0) {
                nz += 1;
            } else {
                break;
            }
        }
    }
    return nz;
}
```

```
public void verifyTheory (int N) {
    int checkInterval = 100; // for printing
    bool failed = false;
    for (int i <- [1..N]) {
        ifact=fact(i);
        int p = predictZeros(i);
        int c = countTrailingZeros(ifact);
        if (p != c) {
            failed = true;
            println ("Found a counter example at i=<i>");
            break;
        } else {
            if (i % checkInterval == 0) {
                println ("<i>! has <p> trailing zeros");
            }
        }
    }
    if (!failed) {
        println ("The theory works for i: 1..<N>");
    }
}
```

Time to celebrate!

```
rascal>verifyTheory(10);
The theory works for i: 1..10
ok
rascal>verifyTheory(100);
100! has 24 trailing zeros
The theory works for i: 1..100
ok
```

```
rascal>verifyTheory(1000);
100! has 24 trailing zeros
200! has 49 trailing zeros
300! has 74 trailing zeros
400! has 99 trailing zeros
500! has 124 trailing zeros
600! has 148 trailing zeros
Found a counter example at i=625
predicted zeros = 155
observed zeros = 156
ok
```

Looking under the hood

```
// I wrote these little wrappers.  
  
// Log for an arbitrary base  
public real logB(real a, real base) {  
    return log(a) / log(base);  
}  
  
public real floor (real a) {  
    return toReal(round (a - 0.5));  
}  
  
public int floorLogBase (int a, int b) {  
    return toInt(floor(logB(toReal(a), toReal(b))));  
}
```

```
rascal>floorLogBase(625,5);  
int: 3  
rascal>logB(625.0,5.0);  
real: 3.9999999999999998757330130880776320985295476764801684.....
```

A bad fix (that kinda works)

```
// I wrote these little wrappers.  
  
// Log for an arbitrary base  
public real logB(real a, real base) {  
    return log(a) / log(base);  
}  
  
public real floor (real a) {  
    return toReal(round (a - 0.5 + 0.00001));  
}  
  
public int floorLogBase (int a, int b) {  
    return toInt(floor(logB(toReal(a), toReal(b))));  
}
```

A better, exact solution

```
// Also change predictZeros to call this version  
public int floorLogBase2 (int a, int b) {  
    int remaining = a;  
    int ans = 0;  
    while (remaining >= b) {  
        ans += 1;  
        remaining /= b;  
    }  
    return ans;  
}
```

```
rascal>verifyTheory(1000);  
100! has 24 trailing zeros  
200! has 49 trailing zeros  
300! has 74 trailing zeros  
400! has 99 trailing zeros  
500! has 124 trailing zeros  
600! has 148 trailing zeros
```

```
700! has 174 trailing zeros  
800! has 199 trailing zeros  
900! has 224 trailing zeros  
1000! has 249 trailing zeros  
The theory works for i: 1..1000  
ok
```



Lessons?

- Explore the terrain, take notes, build intuition, develop theories, test them
 - Refine, repeat
 - Double check
- Build infrastructure with natural “break points”
 - Understandable >> fast, esp. in the beginning
 - The correct way >> the easy way,
 - The correct way may be pretty easy too
- Document and later challenge your assumptions
 - Are you measuring what you think you are measuring?

What history taught me

- Study what you already have and understand
 - Often, your intuition is golden
 - Take it apart and see how it works (e.g., Linux study)
- Challenge pre-conceived notions
 - Create testable hypotheses + evaluate them (e.g., cloning)
- Software archives contain lots of rich data
 - But need to process, link, mine the artifacts
- Need to continually re-examine reasonableness of assumptions
 - Don't blindly trust the numbers; dig and validate!

Listening to big data

Or, philately will get you everywhere

Mike Godfrey

Software Architecture Group

University of Waterloo

