

Subsuming Methods: Finding New Optimisation Opportunities in OO Software

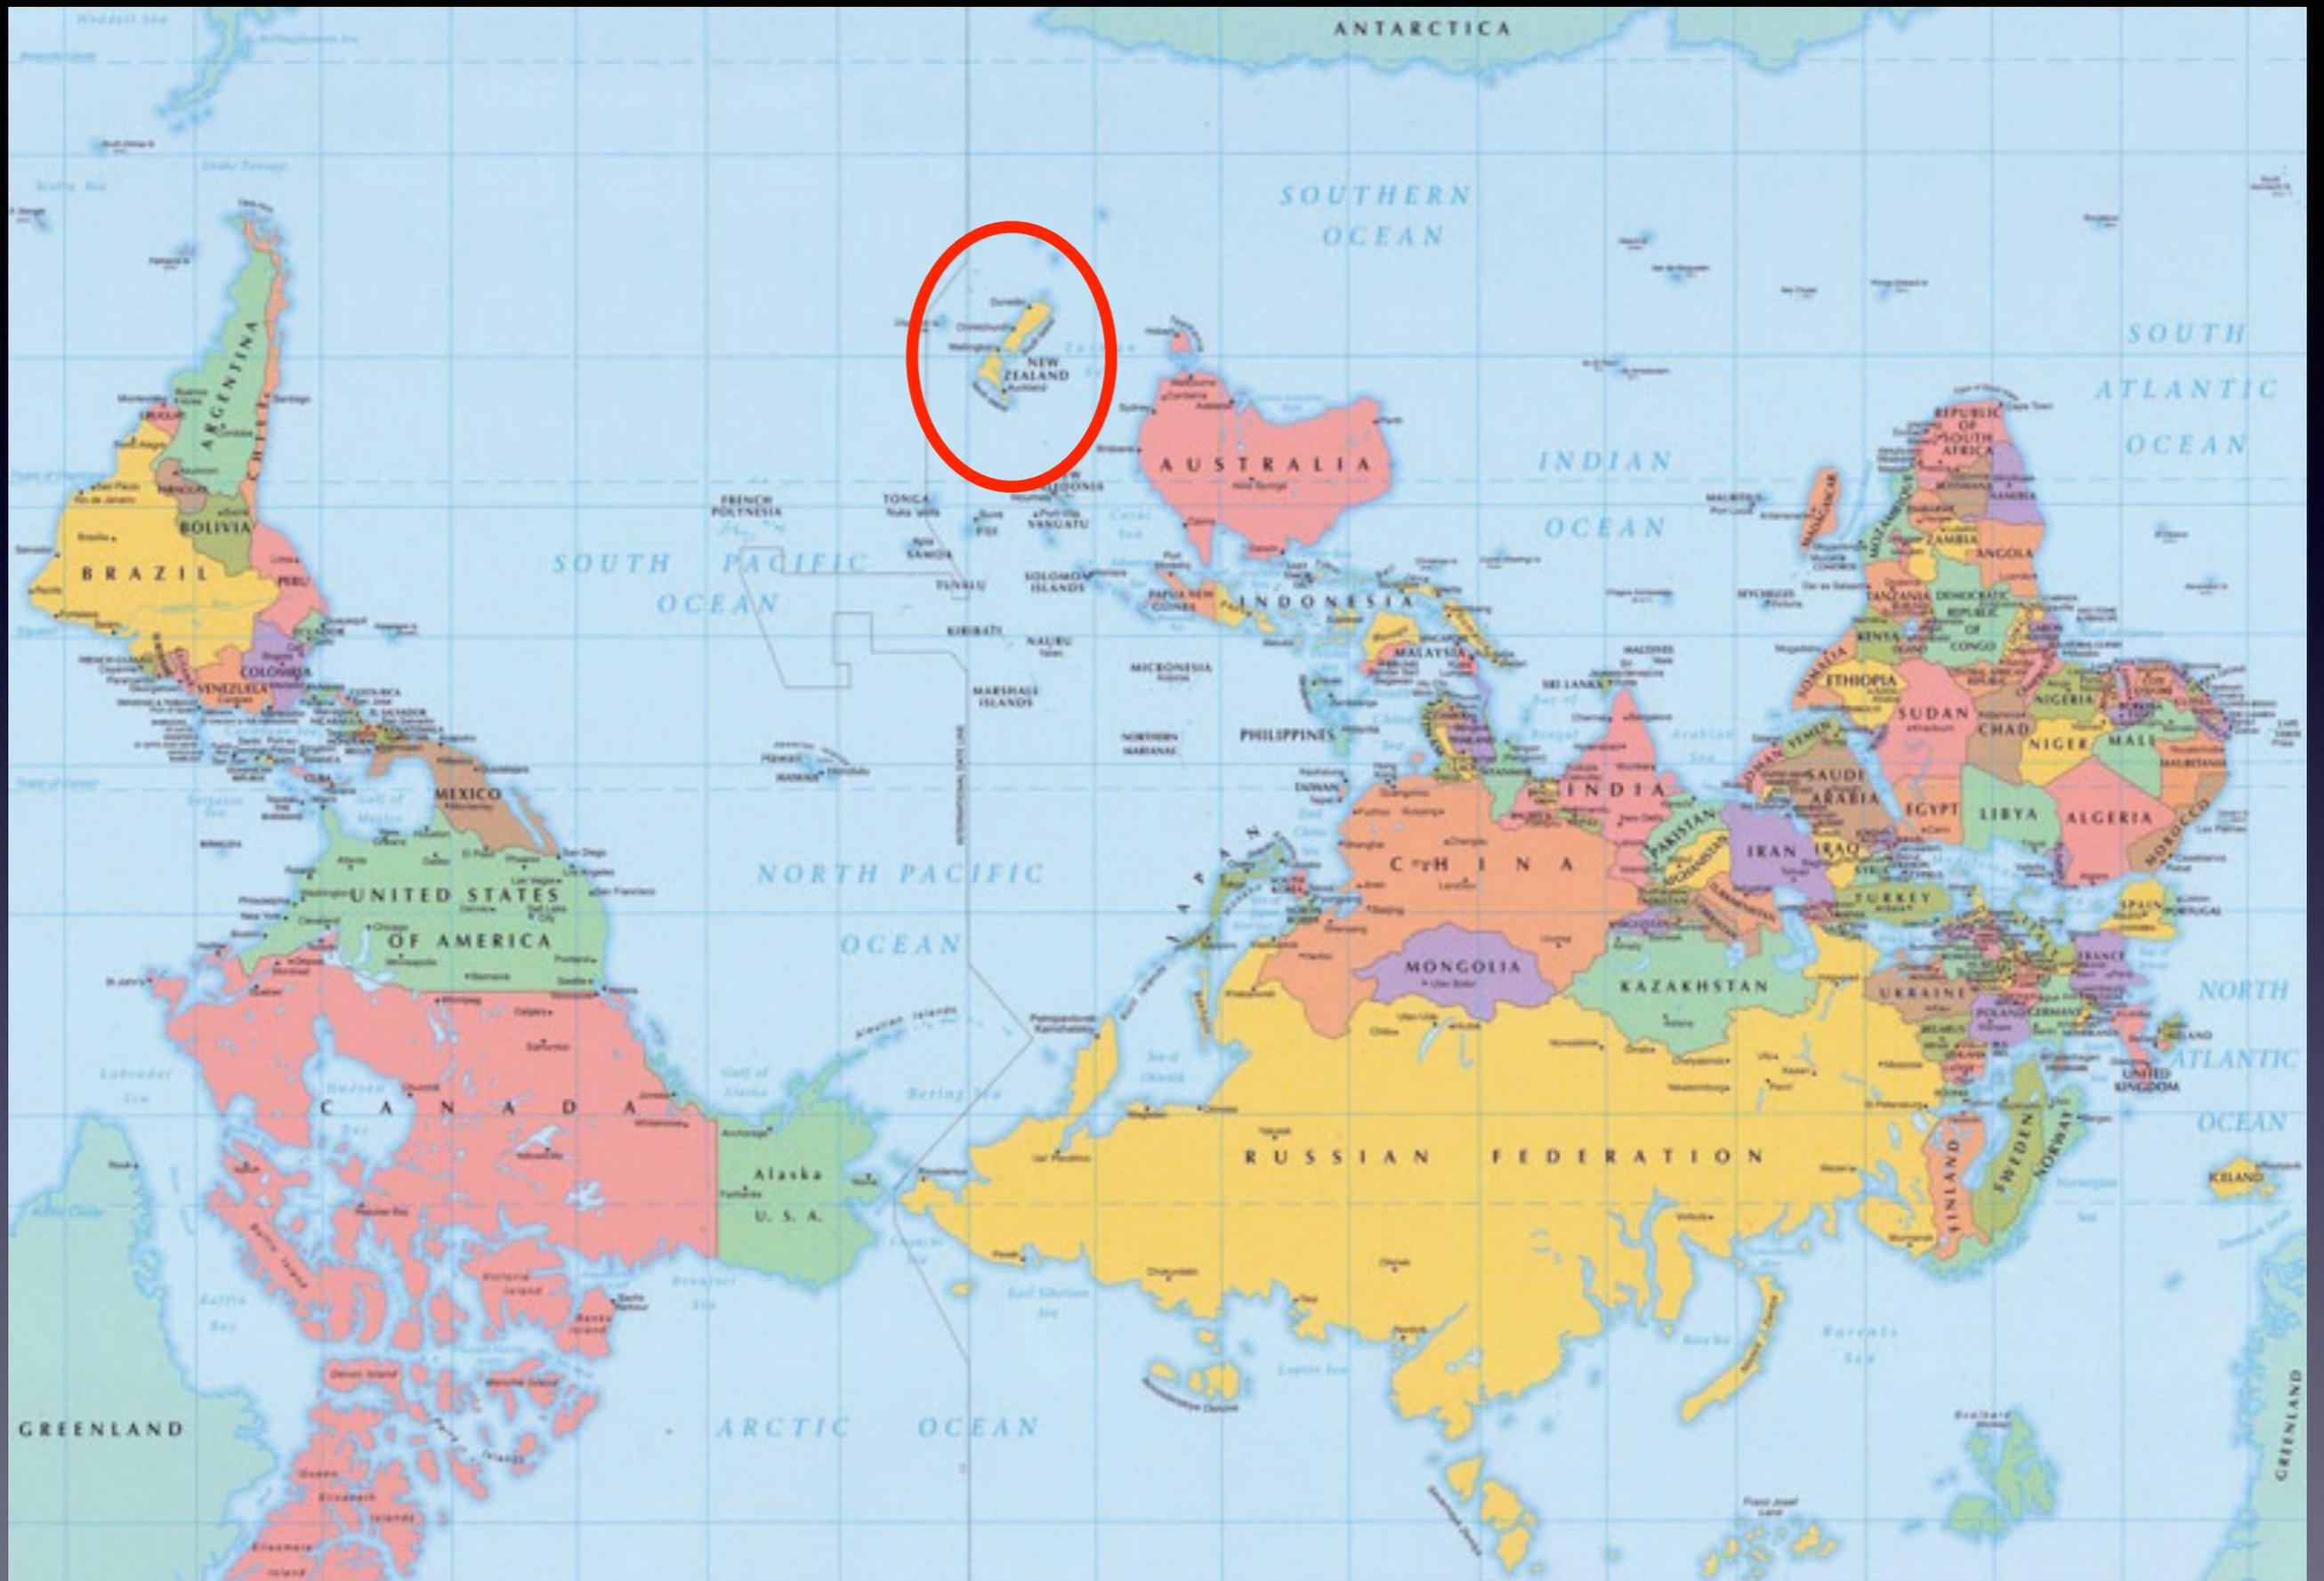
David Maplesden
John Hosking

Ewan Tempero
John Grundy

The University of Auckland
ICPE 2015

Performance is Important

- Cloud computing costs
- Resource constrained environments
 - Mobile applications
- Lost business - Amazon 100ms delay = 1% sales
- Failed projects



Focus

- Large-scale object-oriented software
 - Ubiquitous in industry
- Late-cycle empirical performance analysis
 - a.k.a. profiling and tuning
 - Complementary to model-based predictive methods
- *Analysis* - the neglected backward path

A 'modern' profiler

tradesoap-instrumented.jps - JProfiler 7.2.2

Thread selection: All thread groups Thread status: Runnable

Aggregation level: Methods View mode: Tree

- 43.8% - 1,838 s - inh. 79,590 μ s - 31,291 inv. org.apache.geronimo.jetty6.handler.ThreadClassLoaderHandler.handle
- 43.8% - 1,837 s - inh. 122 ms - 31,291 inv. org.apache.geronimo.jetty6.handler.TwistyWebAppContext\$TwistyHandler.handle
- 43.8% - 1,837 s - inh. 135 ms - 31,291 inv. org.apache.geronimo.jetty6.handler.TwistyWebAppContext.access\$101
- 43.8% - 1,837 s - inh. 142 ms - 31,291 inv. org.mortbay.jetty.webapp.WebAppContext.handle
- 43.8% - 1,837 s - inh. 91,867 μ s - 31,291 inv. org.mortbay.jetty.handler.ContextHandler.handle
- 43.7% - 1,833 s - inh. 114 ms - 31,290 inv. org.mortbay.jetty.servlet.SessionHandler.handle
- 43.6% - 1,829 s - inh. 100 ms - 31,290 inv. org.mortbay.jetty.security.SecurityHandler.handle
- 43.6% - 1,828 s - inh. 67,528 μ s - 31,290 inv. org.mortbay.jetty.servlet.ServletHandler.handle
- 43.5% - 1,822 s - inh. 92,449 μ s - 31,290 inv. org.apache.geronimo.jetty6.InternalJettyServletHolder.handle
- 43.5% - 1,822 s - inh. 145 ms - 31,290 inv. org.mortbay.jetty.servlet.ServletHolder.handle
- 43.5% - 1,821 s - inh. 140 ms - 31,290 inv. org.apache.geronimo.webservices.POJOWebserviceServlet.service
- 43.5% - 1,821 s - inh. 151 ms - 31,290 inv. org.apache.geronimo.webservices.WebServiceContainerInvoker.service
- 43.3% - 1,814 s - inh. 82,009 μ s - 31,290 inv. org.apache.geronimo.axis.server.AxisWebServiceContainer.invoke
- 43.2% - 1,812 s - inh. 278 ms - 31,290 inv. org.apache.geronimo.axis.server.AxisWebServiceContainer.doService
- 37.6% - 1,576 s - inh. 116 ms - 31,290 inv. org.apache.axis.handlers.soap.SOAPService.invoke
- 25.5% - 1,067 s - inh. 91,847 μ s - 31,290 inv. org.apache.axis.SimpleChain.invoke
- 25.4% - 1,066 s - inh. 73,157 μ s - 31,290 inv. org.apache.axis.SimpleChain.doVisiting
- 25.4% - 1,065 s - inh. 336 ms - 93,870 inv. org.apache.axis.strategies.InvocationStrategy.visit
- 21.1% - 882 s - inh. 114 ms - 31,290 inv. org.apache.axis.providers.java.JavaProvider.invoke
- 20.5% - 861 s - inh. 97,357 μ s - 31,290 inv. org.apache.axis.providers.java.RPCProvider.processMessage
- 19.5% - 817 s - inh. 47,026 μ s - 31,290 inv. org.apache.geronimo.axis.server.POJOProvider.invokeMethod
- 19.5% - 816 s - inh. 181 ms - 31,290 inv. java.lang.reflect.Method.invoke
- 9.8% - 409 s - inh. 113 μ s - 32 inv. org.apache.geronimo.samples.daytrader.soap.TradeWSAction.resetDaCapo
- 3.4% - 144 s - inh. 120 ms - 23,417 inv. org.apache.geronimo.samples.daytrader.soap.TradeWSAction.getQuote
- 2.6% - 109 s - inh. 16,821 μ s - 3,248 inv. org.apache.geronimo.samples.daytrader.soap.TradeWSAction.getHoldings
- 1.6% - 66,880 ms - inh. 4,040 μ s - 768 inv. org.apache.geronimo.samples.daytrader.soap.TradeWSAction.sell
- 1.5% - 62,256 ms - inh. 4,264 μ s - 808 inv. org.apache.geronimo.samples.daytrader.soap.TradeWSAction.buy
- 0.3% - 11,446 ms - inh. 9,399 μ s - 1,788 inv. org.apache.geronimo.samples.daytrader.soap.TradeWSAction.getAcco
- 0.2% - 6,998 ms - inh. 2,690 μ s - 512 inv. org.apache.geronimo.samples.daytrader.soap.TradeWSAction.login
- 0.1% - 3,386 ms - inh. 2,796 μ s - 564 inv. org.apache.geronimo.samples.daytrader.soap.TradeWSAction.logout
- 0.0% - 1,081 ms - inh. 548 μ s - 100 inv. org.apache.geronimo.samples.daytrader.soap.TradeWSAction.updateAccou
- 0.0% - 519 ms - inh. 277 μ s - 52 inv. org.apache.geronimo.samples.daytrader.soap.TradeWSAction.register
- 0.0% - 27,892 μ s - inh. 5 μ s - 1 inv. org.apache.geronimo.samples.daytrader.soap.TradeWSAction.initializeDaCapo
- 0.0% - 358 μ s - inh. 60 μ s - 5 inv. org.apache.geronimo.kernel.config.MultiParentClassLoader.loadClass
- 0.0% - 530 ms - inh. 530 ms - 31,290 inv. java.lang.Class.getMethod
- 0.0% - 164 ms - inh. 164 ms - 31,290 inv. java.lang.reflect.Method.getName
- 0.0% - 141 ms - inh. 141 ms - 31,290 inv. java.lang.reflect.Method.getParameterTypes

View Filters: Reset View Filters

Call Tree Hot Spots Call Graph Method Statistics Call Tracer

Aug 7, 2013 11:00:16 AM 45:21 Snapshot

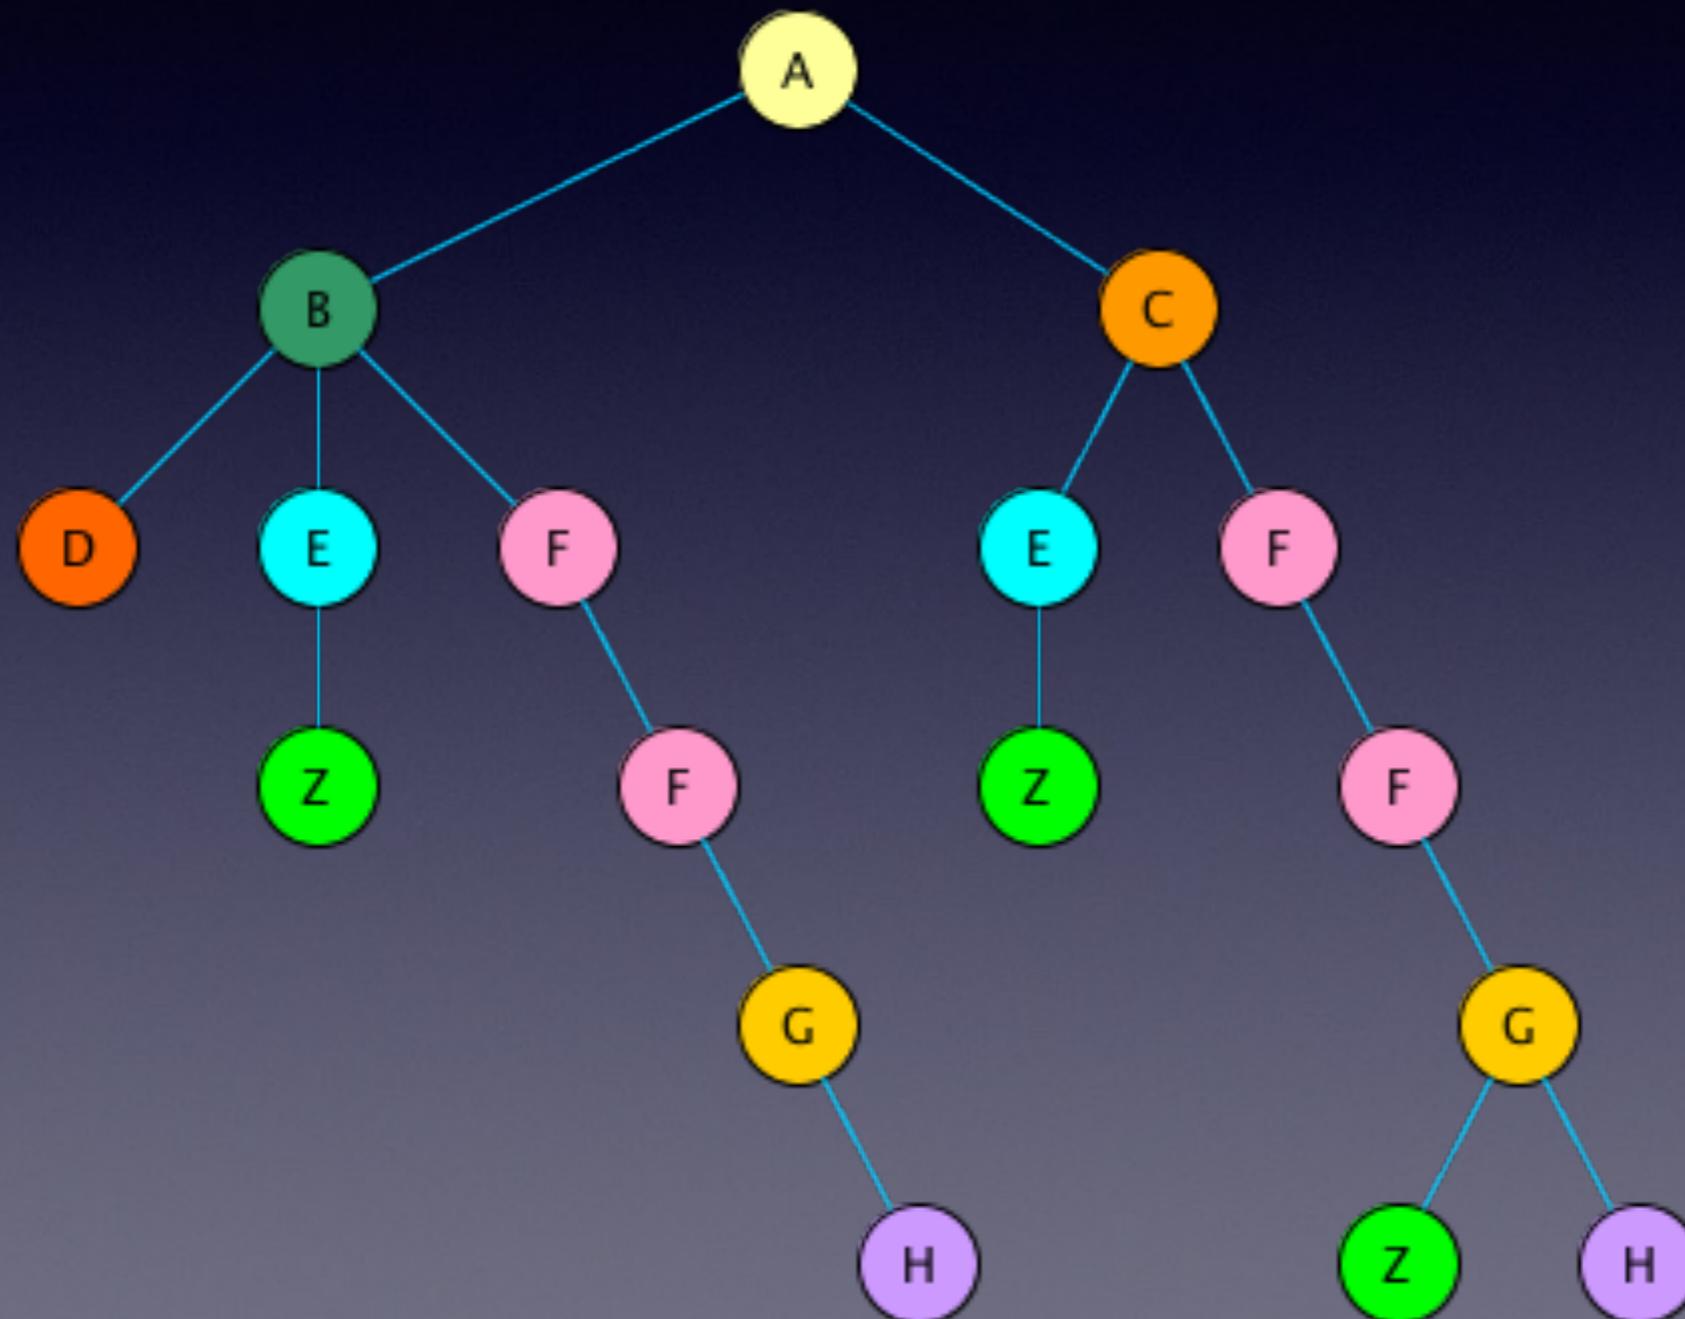
Challenges of OO software

- Numerous small methods
- Heavily layered architecture
 - Engineered for maintainability and reuse
 - Reusable frameworks, more abstractions
 - *Runtime bloat*
- Complex, thinly distributed, runtime behaviour
- Challenging to identify optimisation opportunities

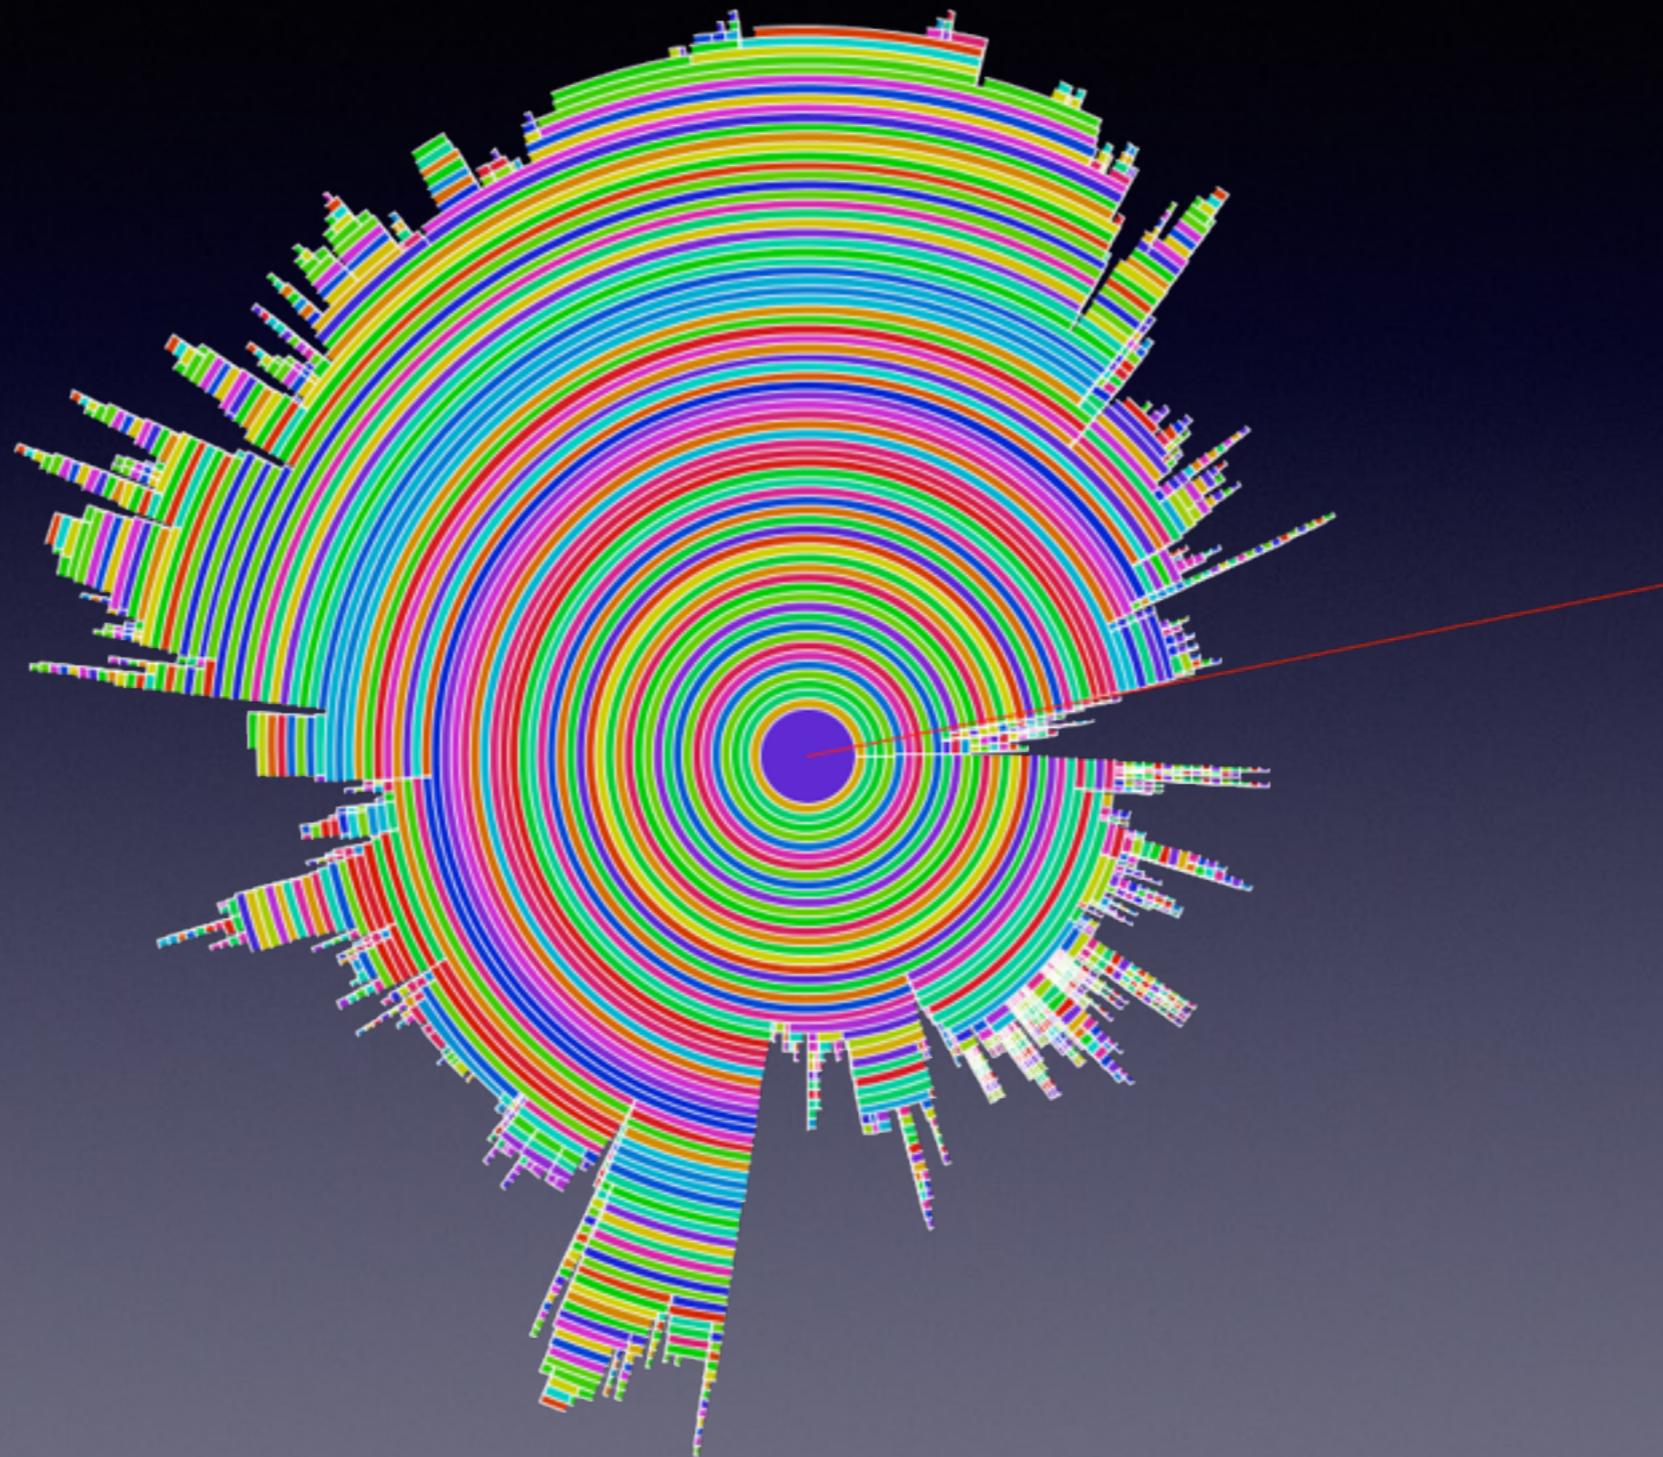
Hot methods - DaCapo fop benchmark

Method	Occurrences in CCT	% Exclusive Time
sun.misc.FloatingDecimal.dtoa	348	6.876
java.text.DigitList.set	374	5.245
java.text.DecimalFormat.subformat	374	3.110
org.apache.fop.fo.properties.PropertyMaker.findProperty	1501	2.461
java.lang.String.equals	4666	1.853
sun.nio.cs.US_ASCII\$Encoder.encode	568	1.788
sun.misc.FloatingDecimal.countBits	348	1.556
java.util.HashMap.hash	10663	1.506
java.util.HashMap.getEntry	6081	1.342
java.lang.String.indexOf	3343	1.295

Background - Calling Context Trees



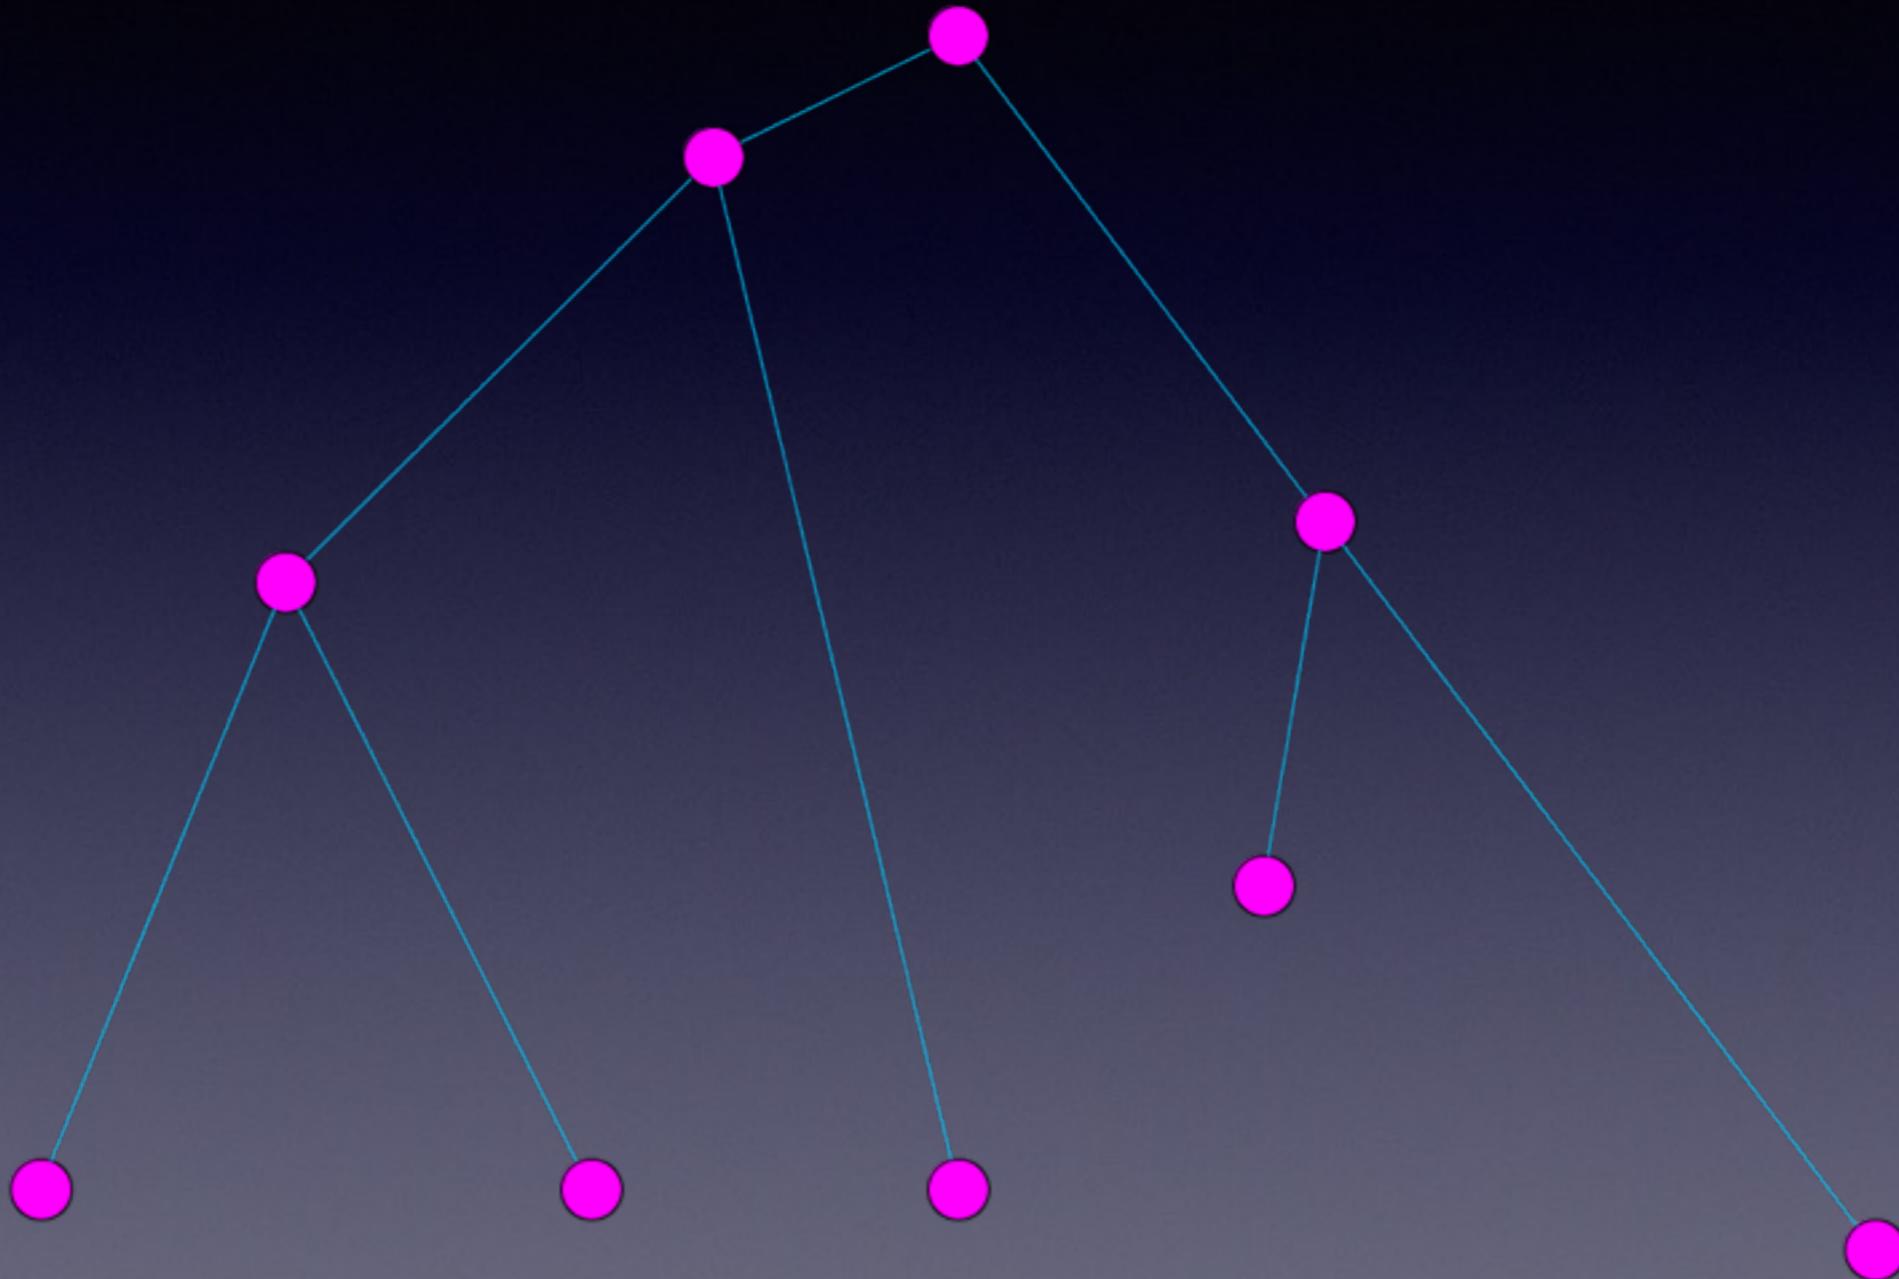
Calling Context Ring Chart



Key Idea

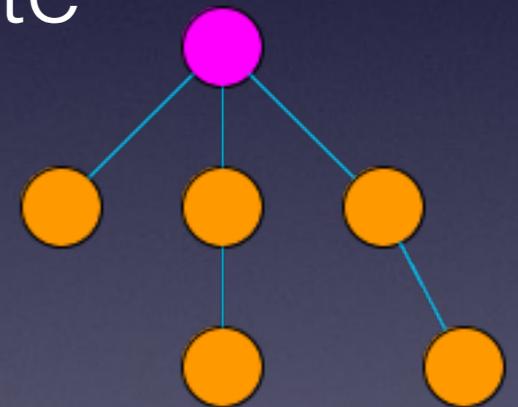
- CCTs aren't just random data
- There are patterns within the calling context tree
 - induced by the design of the software
 - compact
 - repeated in multiple locations
 - expensive when aggregated

Consolidated Tree



Subsuming Methods

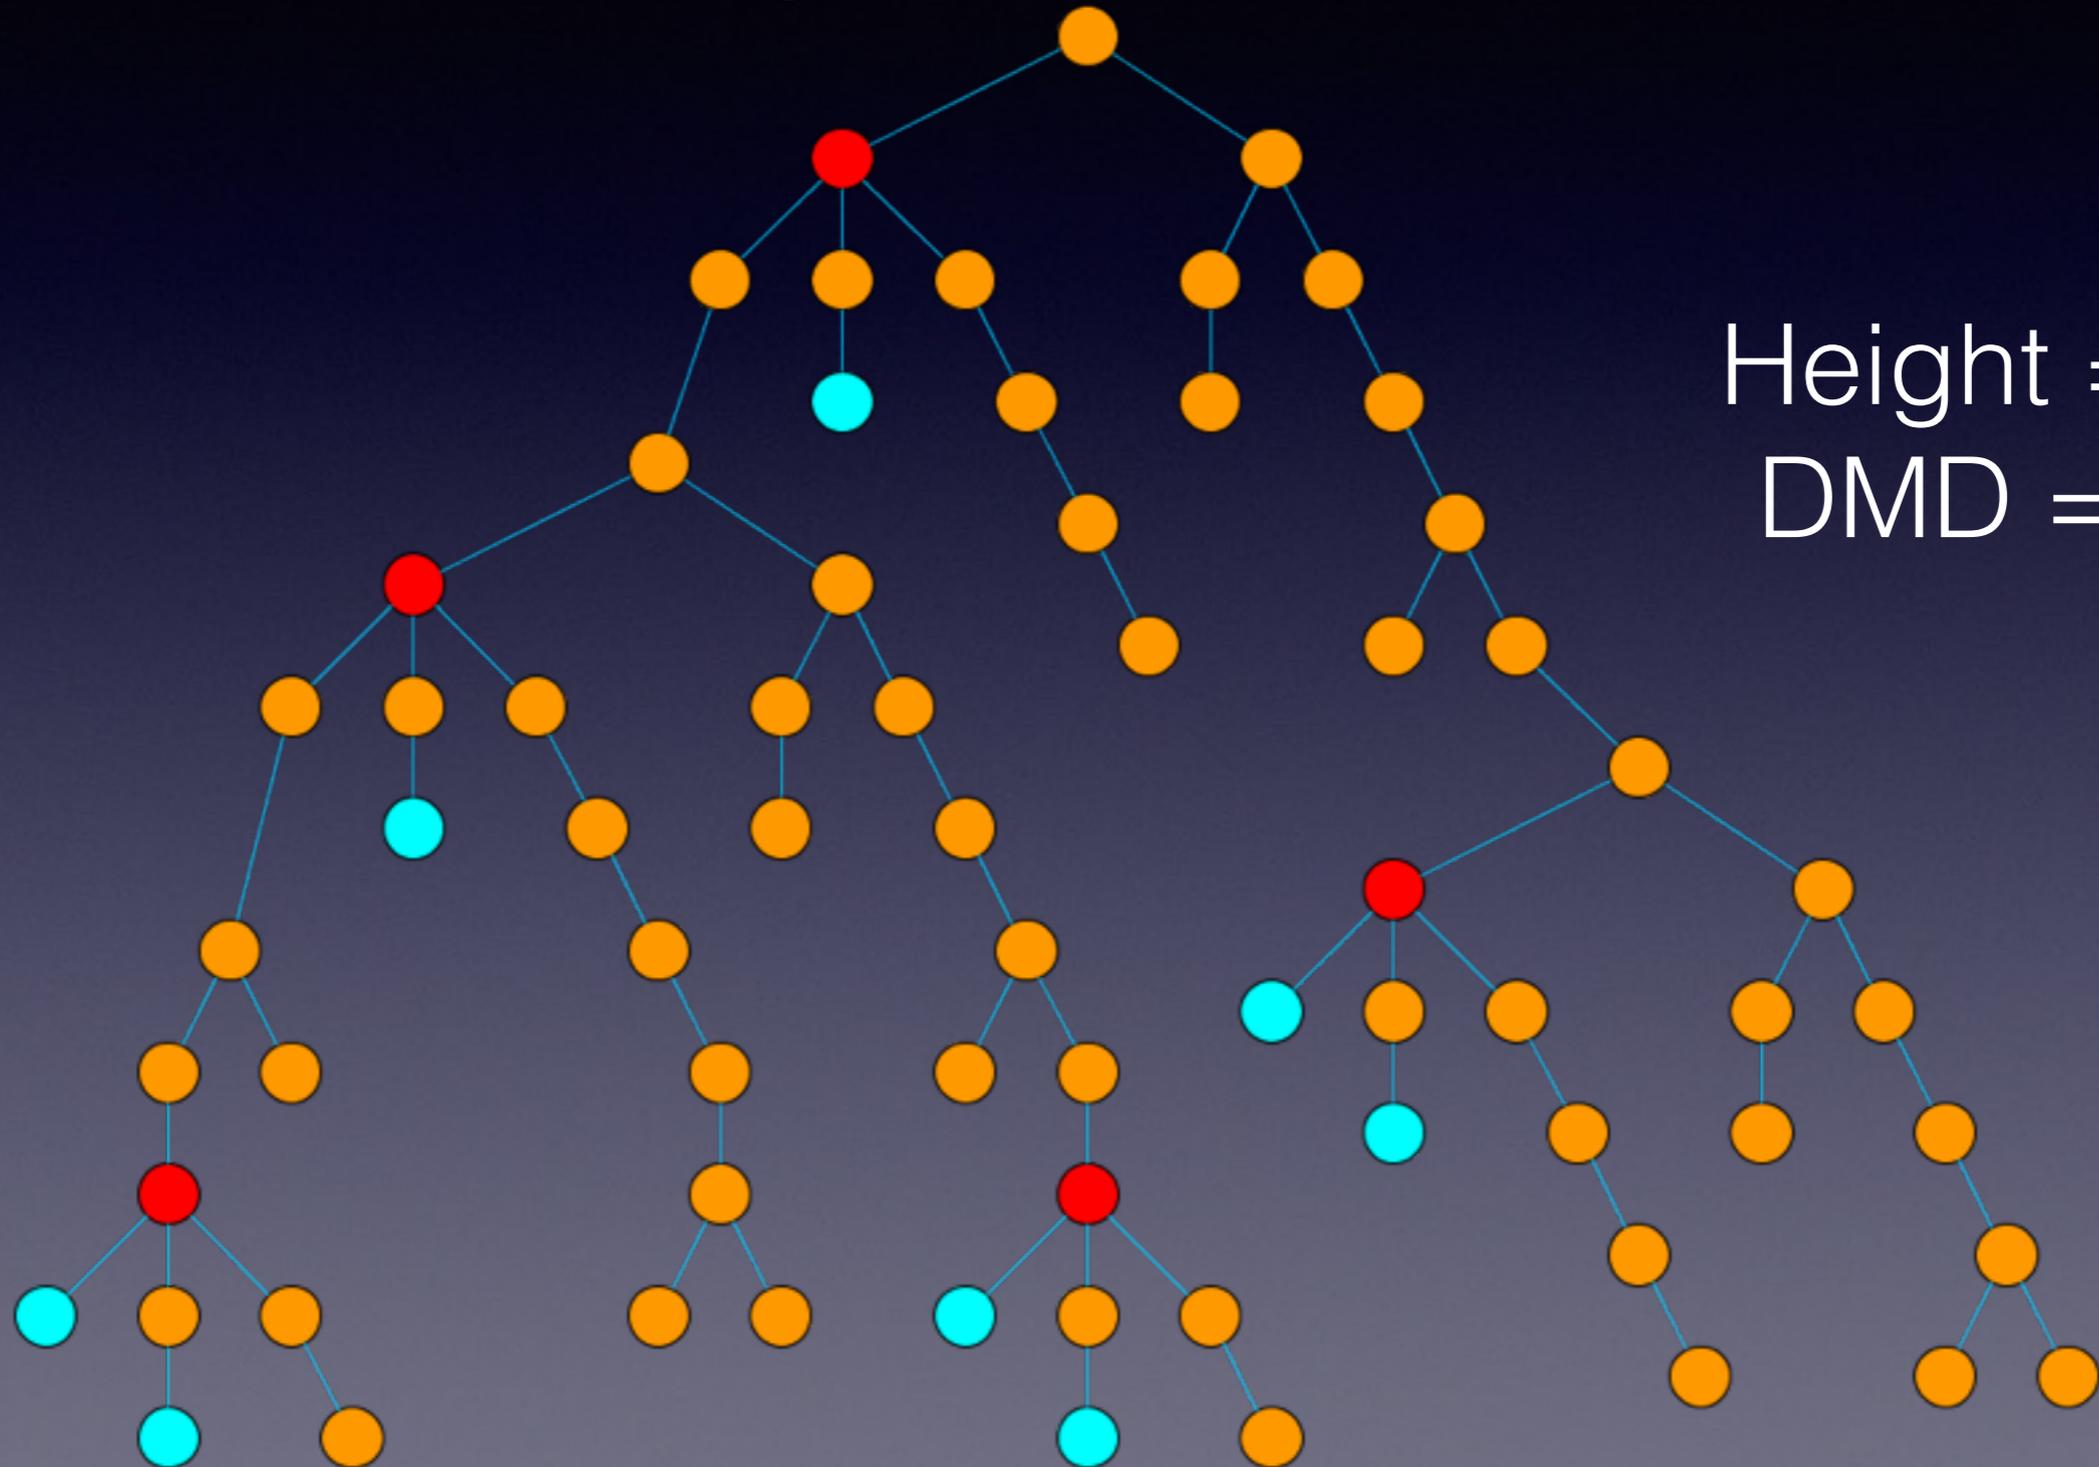
- Partition the CCT into areas of related functionality
- Induced time is very efficient to calculate
- Each subsuming method represents a repeated pattern
- How do we choose our subsuming methods?



Subsuming Attributes

- ‘Elementary’ methods - induce a limited range of behaviour
 - Approximated using height of method in CCT
 - Trivial case (height = 0) - makes no method calls - a leaf method
 - getters, setters, hashCode(), equals()
 - ~30% of all methods are leafs ~70% have height ≤ 4
- ‘Subordinate’ methods - called in a predictable manner
 - Every call to the method can be attributed to a (nearby) dominating method which is responsible for the invocation
 - Measured using novel metric - dominating method distance
 - Trivial case is when a method is only ever called from a single call site
 - ~70% of all methods have a single parent ~77% have DMD ≤ 4

Height and Dominating Method Distance



Height = 0
DMD = 2

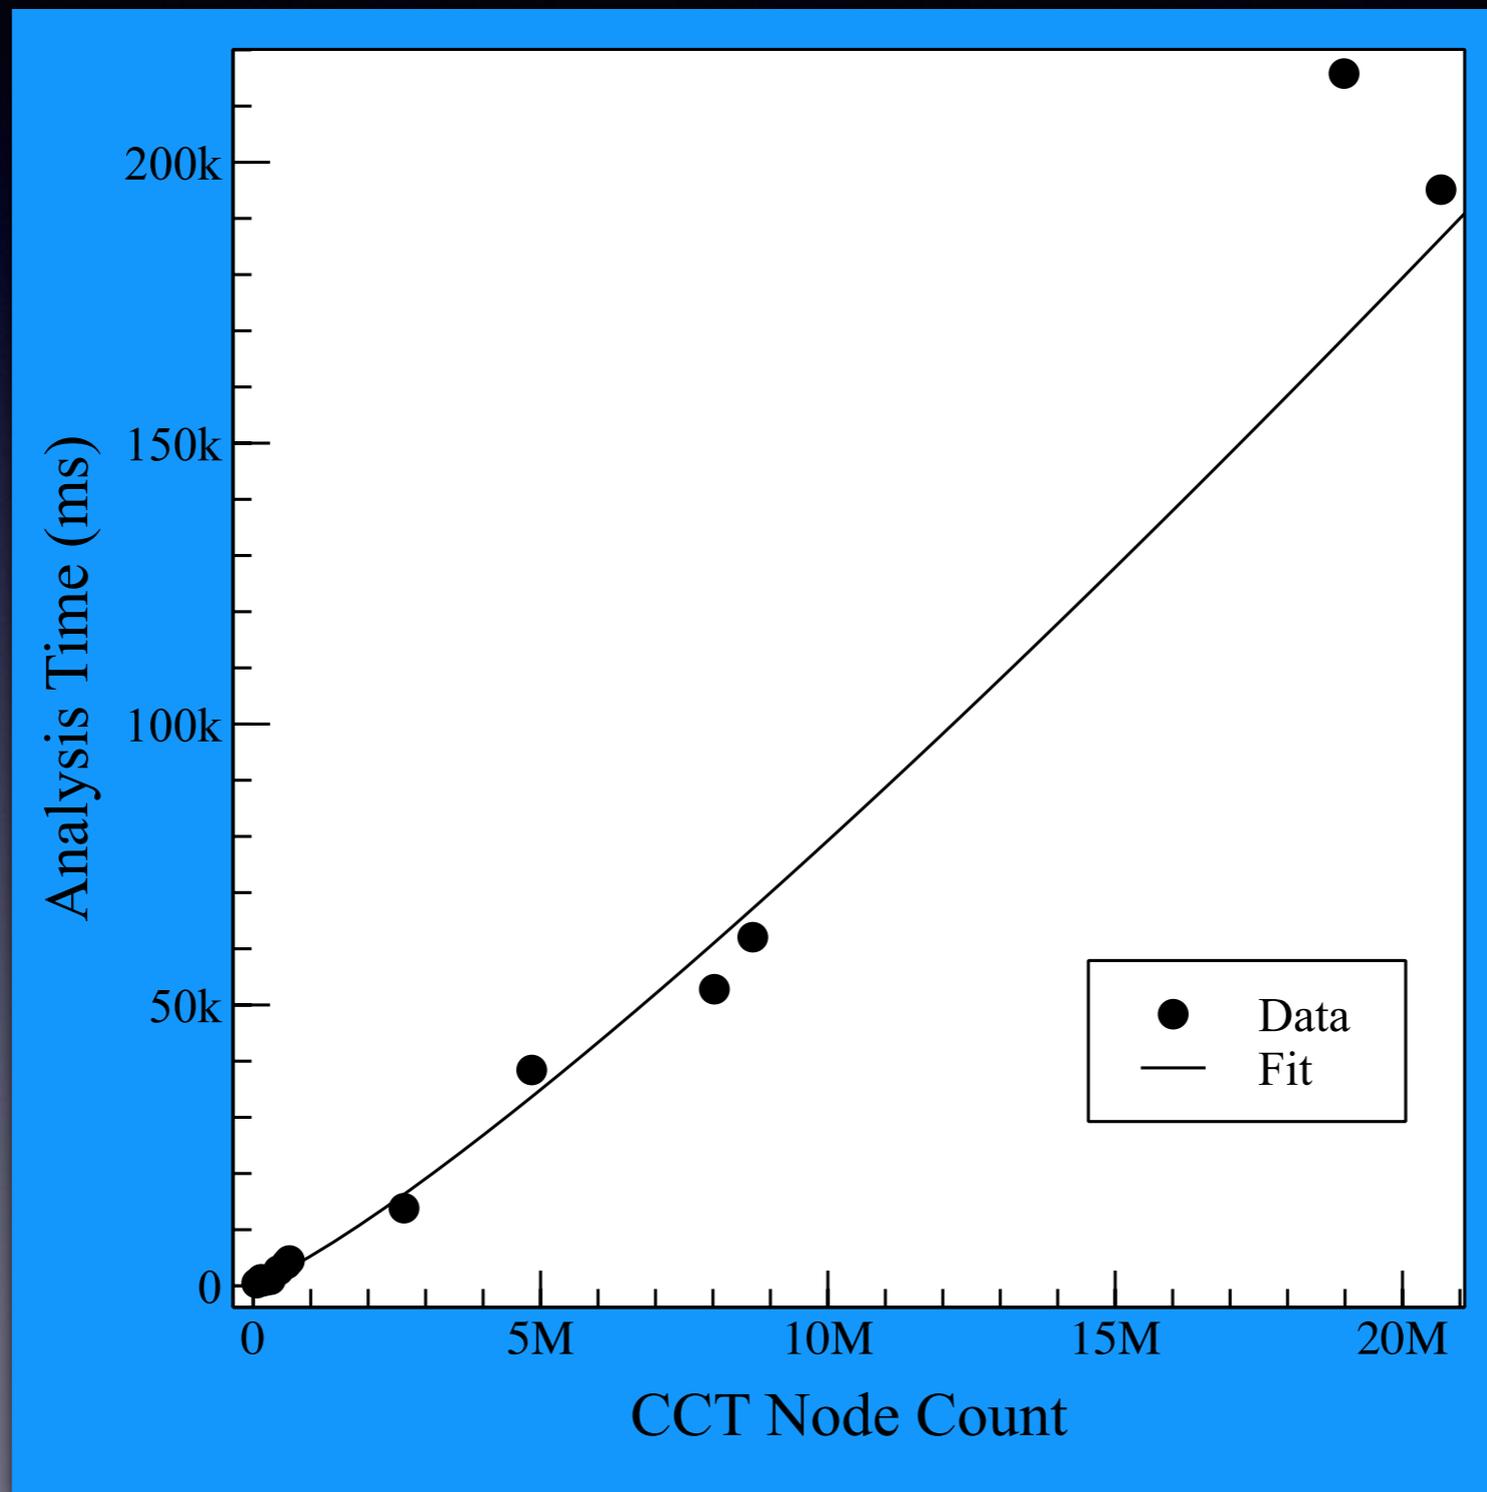
Experimental Evaluation

- DaCapo-9.12-bach (2009) benchmark suite
 - 14 different Java benchmark applications
- JP2 profiler to capture CCT profiles
 - Very consistent, reproducible results
- Apply our subsuming methods analysis
- 5 runs for each benchmark
- Constraints: height > 4 and DMD > 4

Results Summary

- Across the 14 benchmarks:
 - 6.12% of all methods were subsuming
 - 11.82% of nodes in the CCT were subsuming
 - => subsuming methods aggregation greatly simplifies profile information
 - 15 of the top 20 subsuming methods were not in the top 20 inclusive or exclusive cost methods
 - => new optimisation opportunities are identified
- <https://www.cs.auckland.ac.nz/~dmap001/subsuming/>

Results - Analysis Efficiency



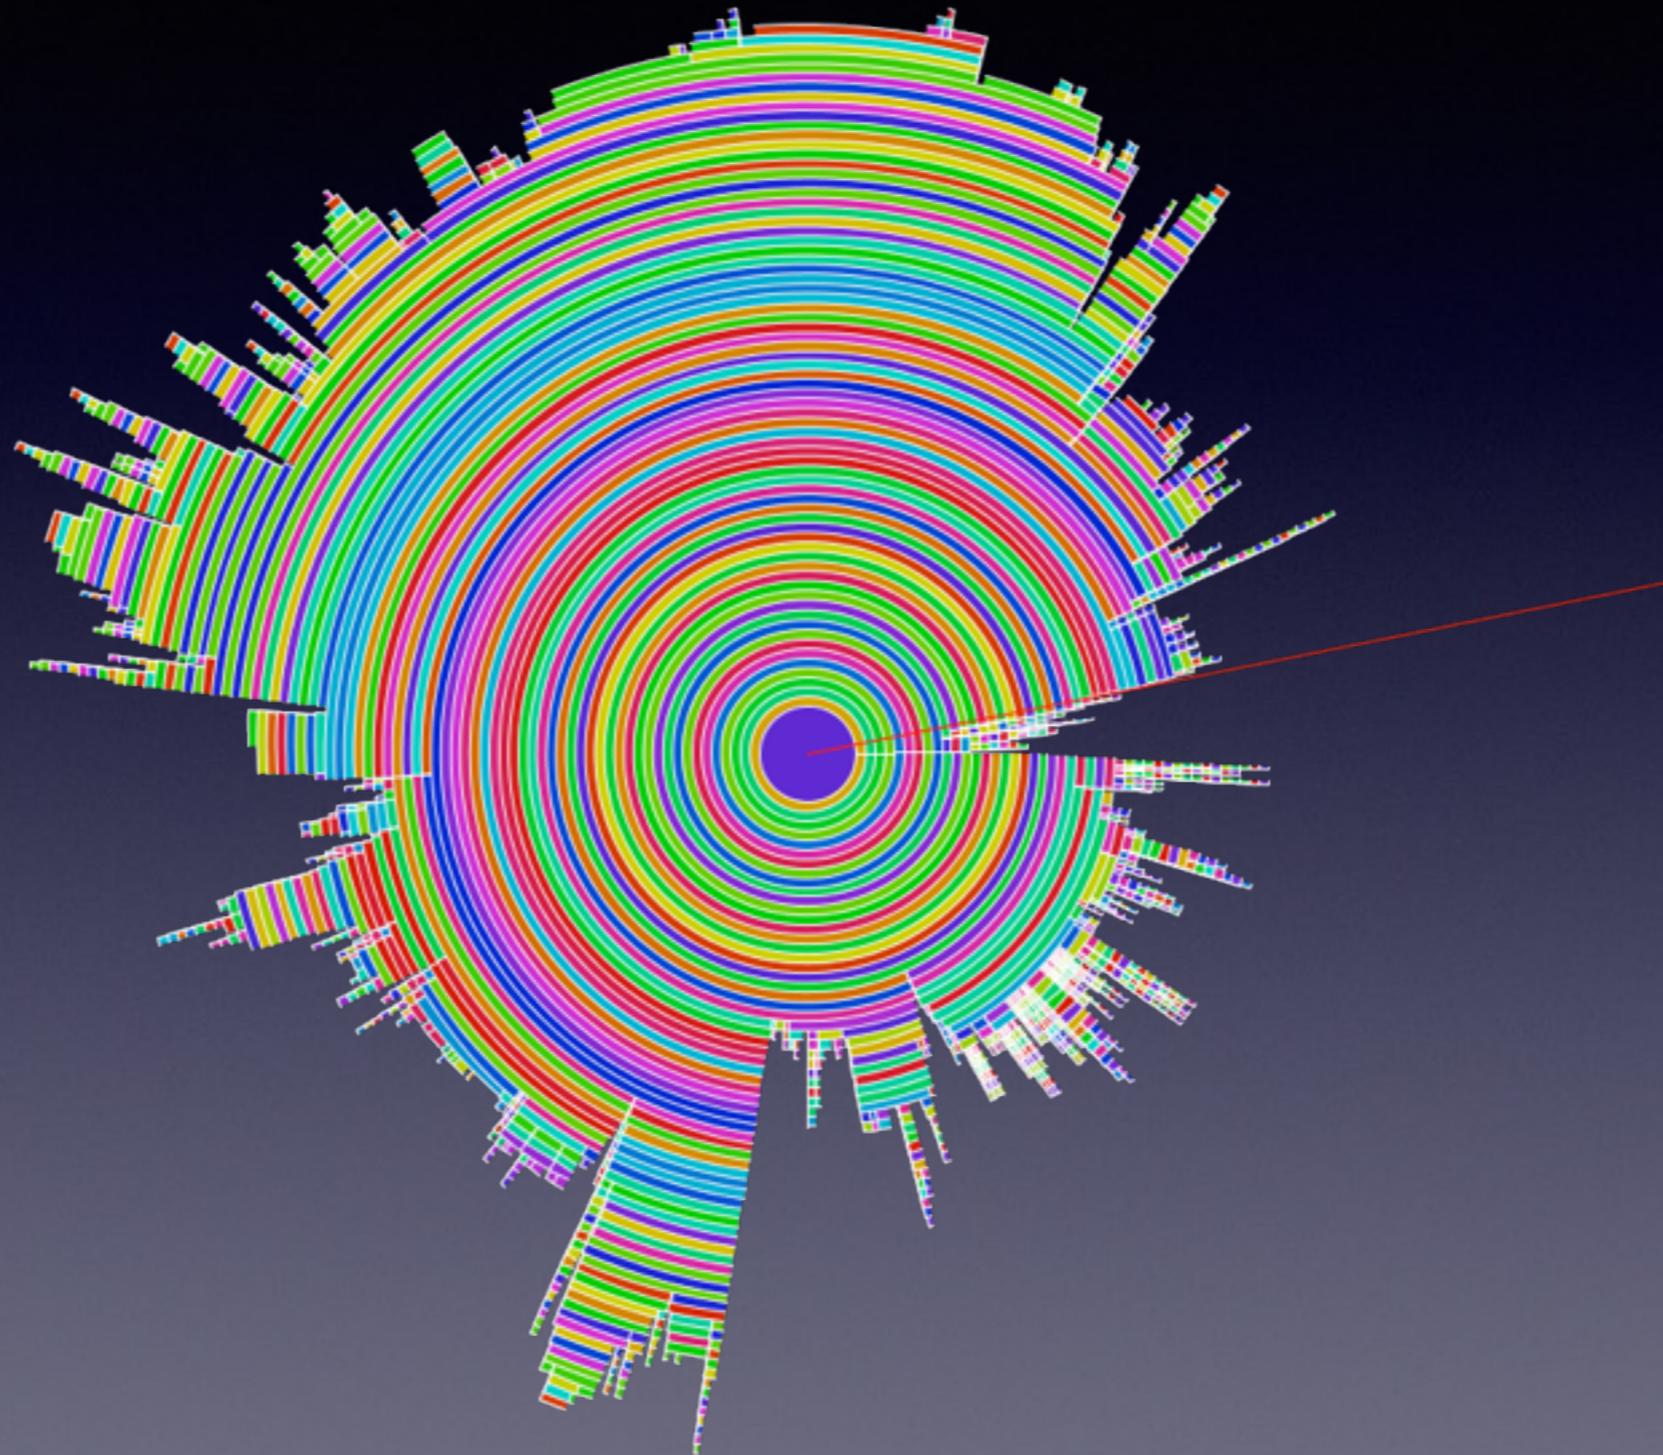
Results - DaCapo fop benchmark

	Full CCT	Subsuming CCT	Ratio
Nodes	628751	71430	11.36%
Height	111	25	22.52%
Unique Methods	6709	345	5.14%

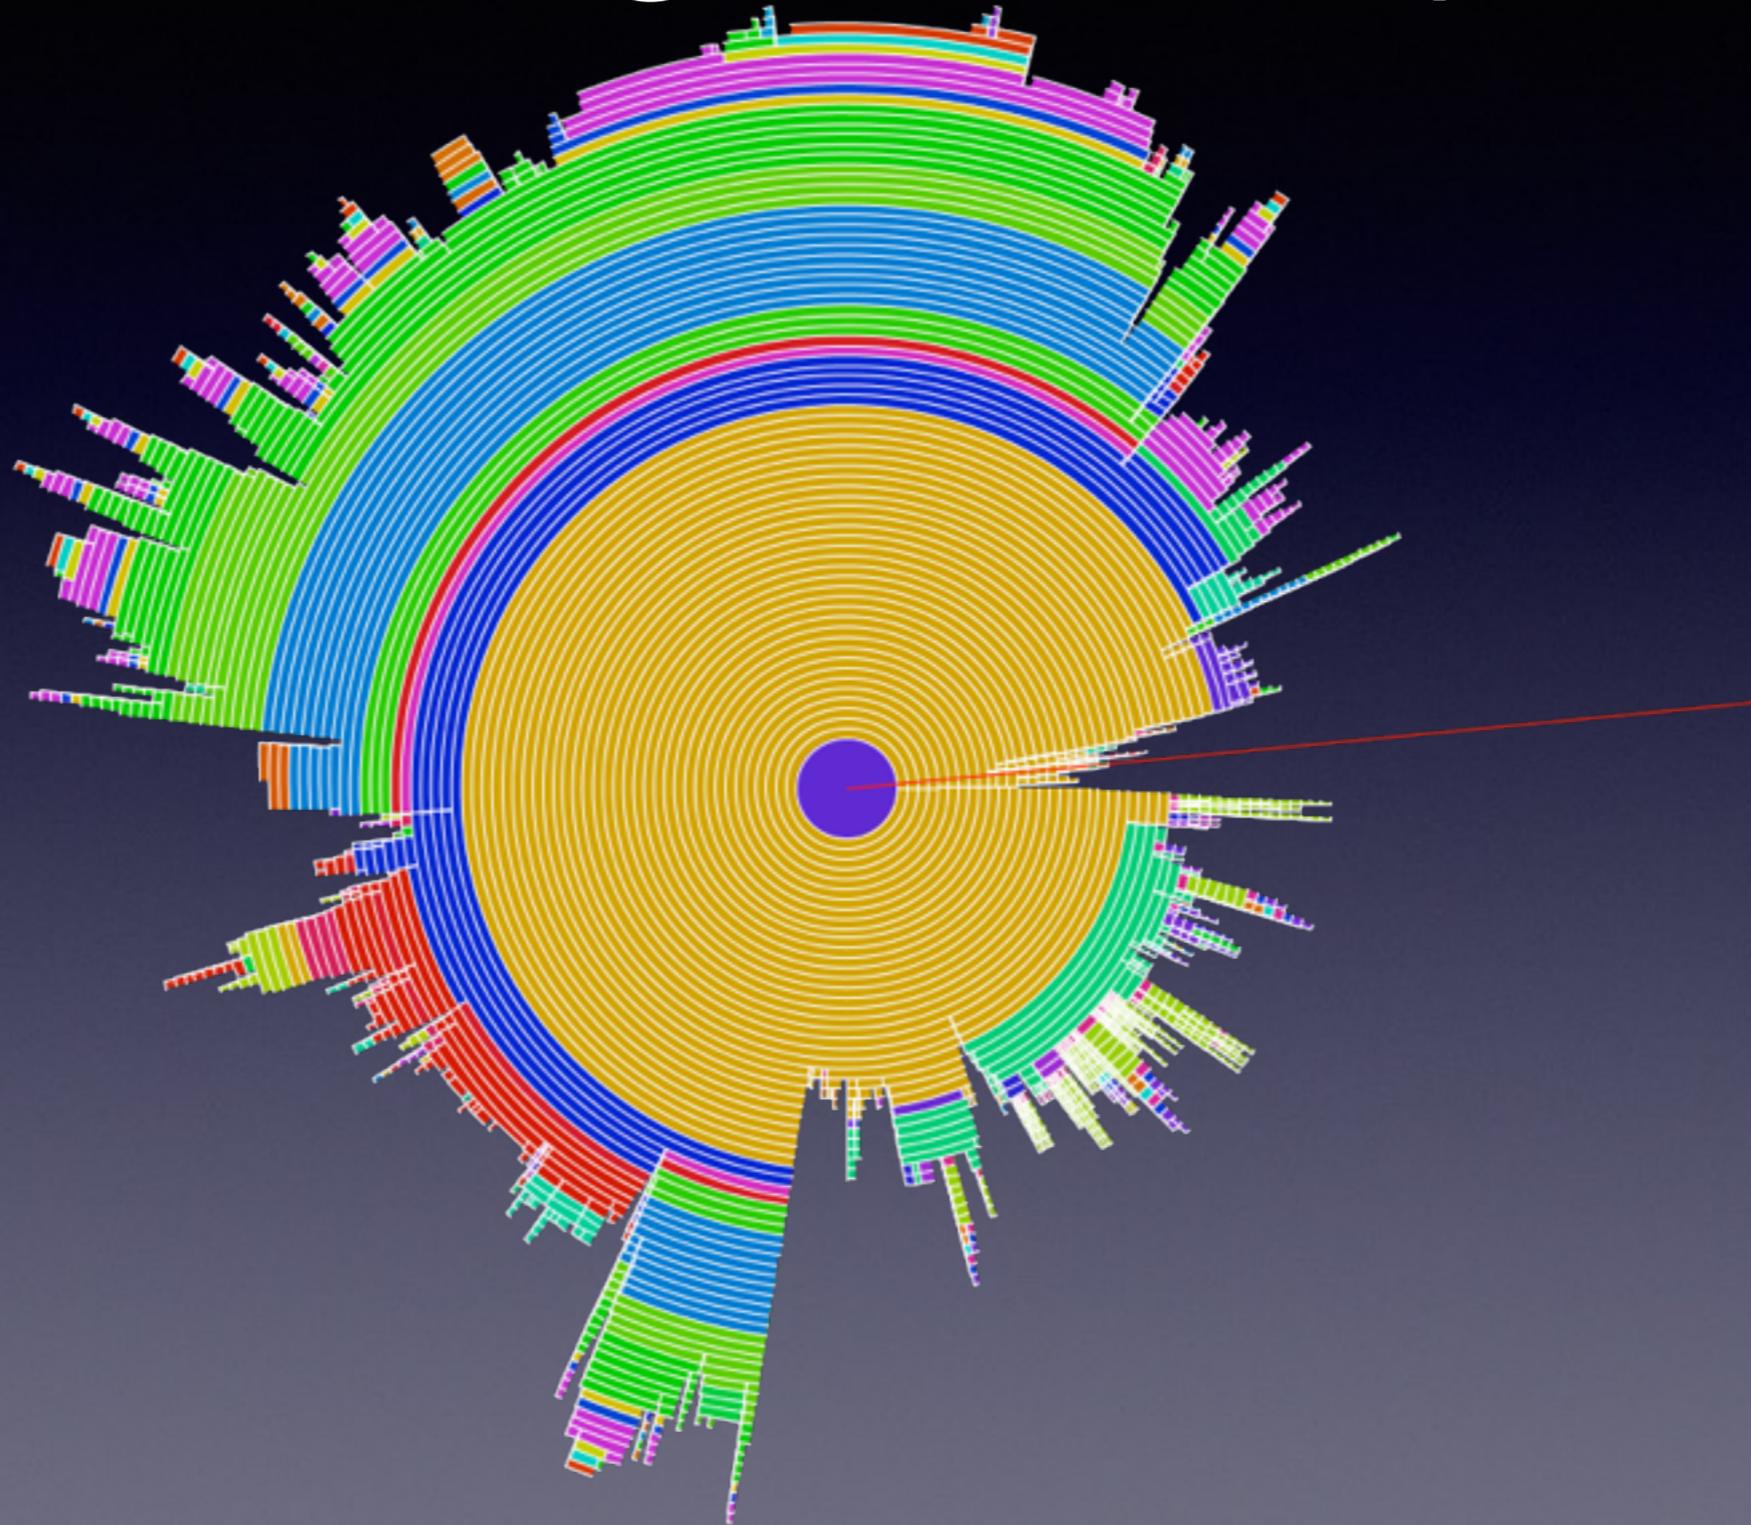
Top Subsuming Methods - DaCapo fop benchmark

Method	Occurrences in CCT	% Exclusive Time	% Induced Time
java.text.DecimalFormat.format	374	0.169	13.644
org.apache.fop.fo.StaticPropertyList.get	1691	1.228	8.884
sun.misc.FloatingDecimal.dtoa	348	6.876	8.871
org.apache.fop.layoutmgr.BlockStackingLayoutManager.getNextKnuthElements	12	0.068	6.381
org.apache.fop.render.AbstractRenderer.renderInlineArea	42	0.041	4.449
org.apache.fop.layoutmgr.BreakingAlgorithm.findBreakingPoints	16	0.002	4.340

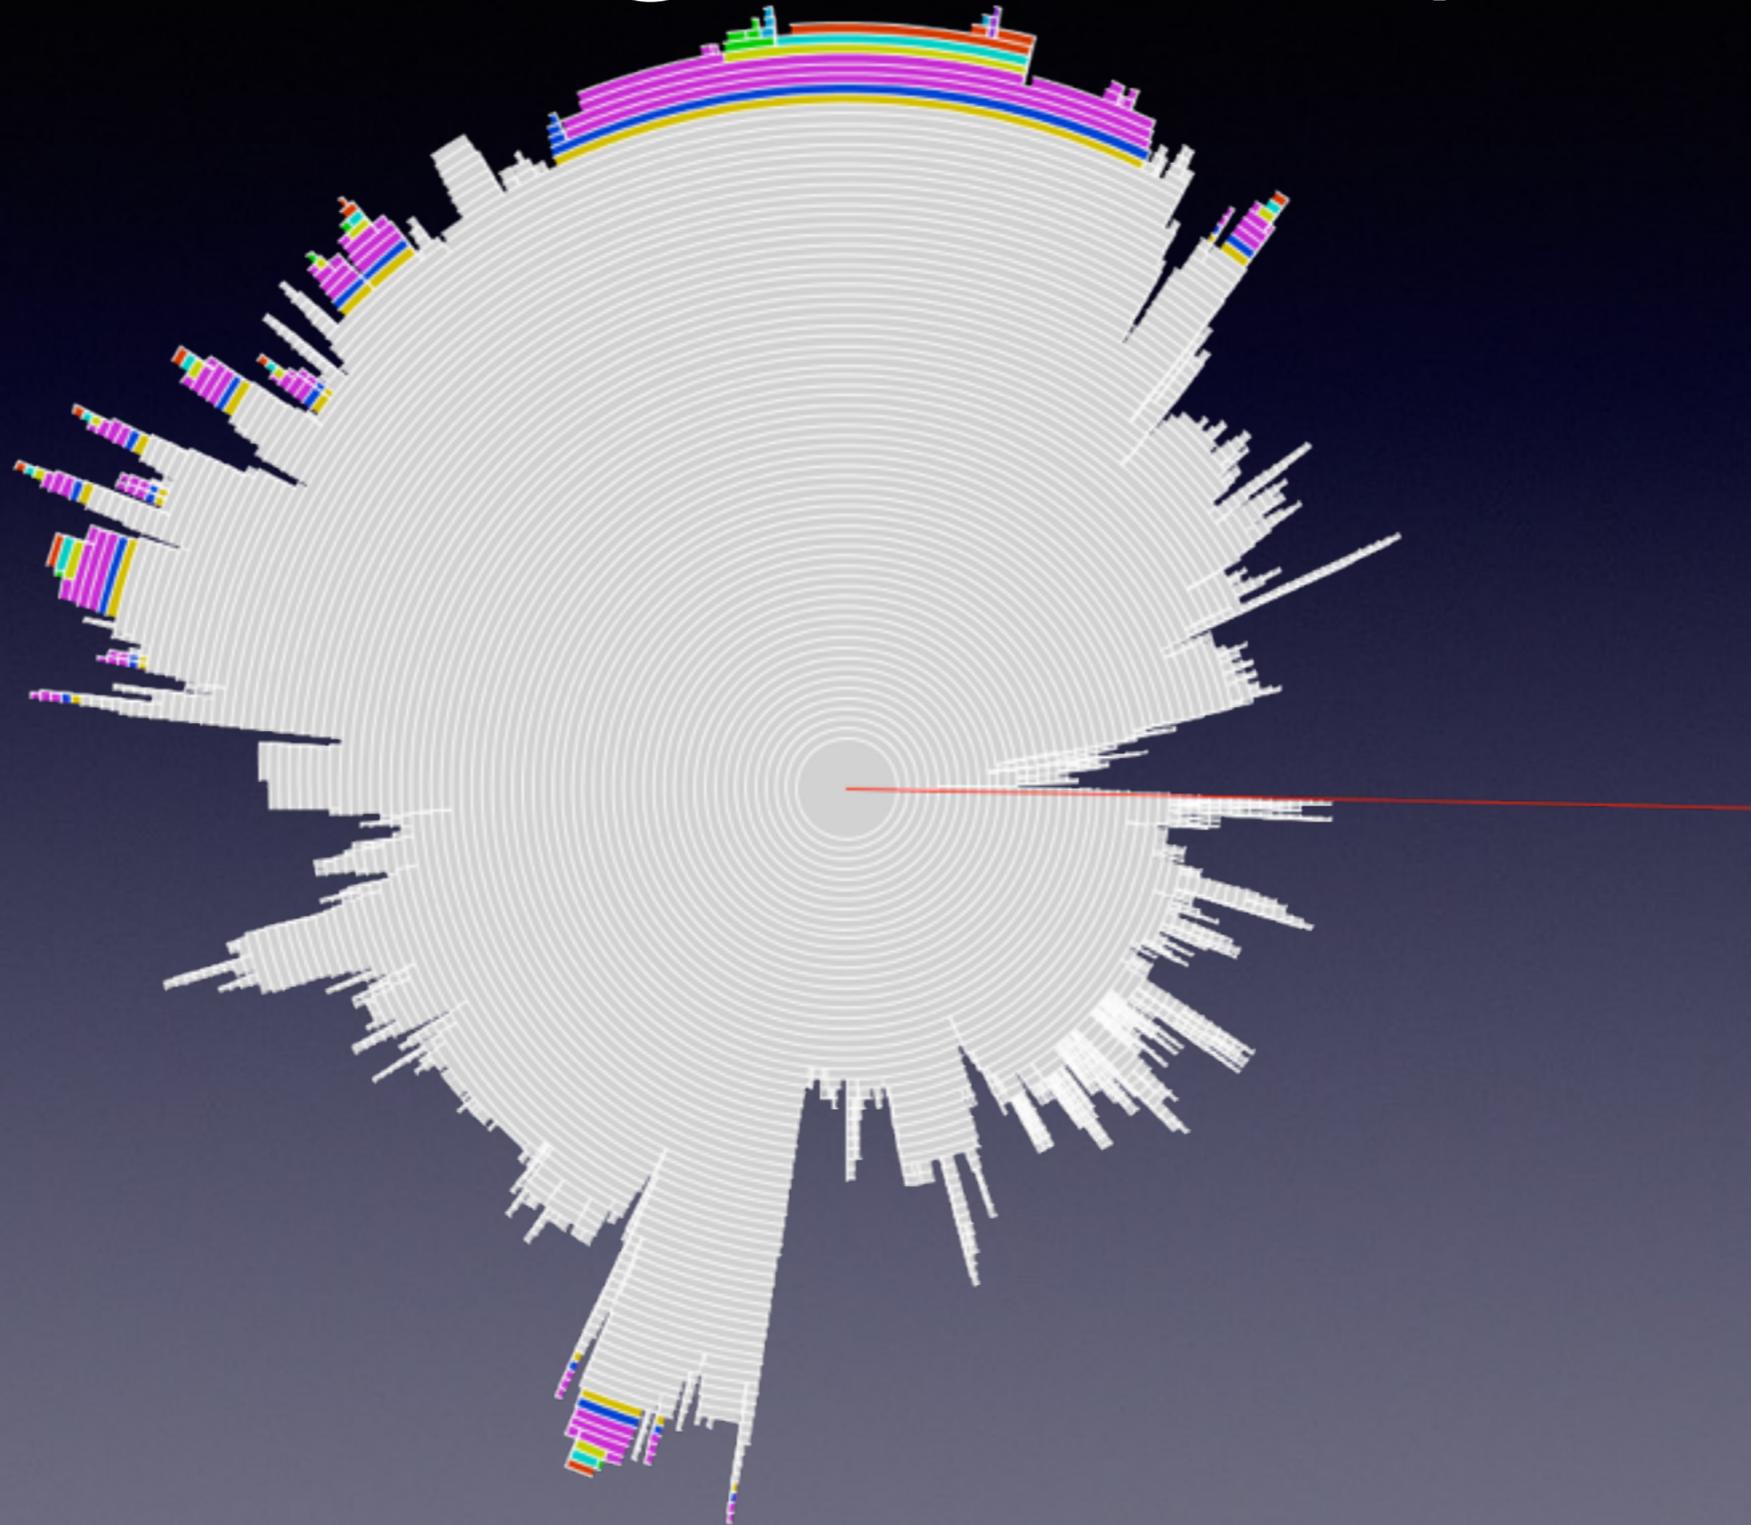
Calling Context Ring Chart



Subsuming - DaCapo fop



Subsuming - DaCapo fop



DaCapo fop - Improvements

- Top subsuming method - `java.text.DecimalFormat.format`
- Highly complex general purpose number formatter
- Called 98% of the time from one location to produce a very specific (and simple) 2 decimal place format
 - `org.apache.xmlgraphics.ps.PSGenerator.formatDouble`
- Accounts for 26% of the total benchmark cost
- Replace with a more specialised implementation
 - 22% reduction in benchmark cost

Summary

- Subsuming Methods
 - Empirical performance analysis approach
 - Helps identify repeated patterns within a CCT profile
 - Efficient offline analysis
 - Complementary to existing approaches
 - Applicable to a wide range of data
- Preliminary evaluation with DaCapo benchmark

Industry Case Study

- letterboxd.com
 - 125,700 registered members
 - 3.6 million requests per day
 - 54.8% reduction in CPU load
 - 49.6% reduction in response time
- Paper accepted at ICSE 2015 - SEIP track
 - “Performance Analysis using Subsuming Methods: An Industrial Case Study” - Maplesden et al

User Study

- Test the effectiveness of subsuming methods analysis in aiding software engineers
- Implemented as an on-line test and questionnaire
- Mid 2015
- If interested please volunteer!
 - Contact: david@maplesden.co.nz

Thank you!

Questions?

Related Work

- Very broad domain (100 venues in our SLR)
- Relevant work from HPC, Compiler, Programming Language domains
- Majority of approaches provide simple metrics
 - Lack of actionable feedback
- Very few approaches leverage static analysis
- Runtime bloat research focussed on data-flow

Systematic Mapping

- “Performance Analysis for Object-Oriented Software: A Systematic Mapping”
- Empirical methods focus
- Accepted for publication in TSE
- <http://dx.doi.org/10.1109/TSE.2015.2396514>

Runtime Bloat Research

- Tackle problem of excessive activity to achieve seeming simple functionality
- Data-flow centric approaches
 - Efficiency of data structures
 - Object pooling opportunities
 - Copy profiling
 - Reference propagation profiling

Existing Approaches

- Lin et al (2010). Towards Anomaly Comprehension: Using Structural Compression to Navigate Profiling Call-Trees.
 - Aggregation by package and class name
- Srinivas & Srinivasan (2005). Summarizing application performance from a components perspective.
 - Thresholding and filtering