The case for the Three R's of Systems Research:

Repeatability Reproducibility & Rigor

Jan Vitek

Kalibera, Vitek. Repeatability, Reproducibility, and Rigor in Systems Research. EMSOFT 1

Science Done Bad

In 2006, Potti&Nevins claim they can predict lung cancer

In 2010, papers retracted, bancruptcy, resignations & investigation

Bad science ranging from fraud, unsound methods, to off-by-one errors in Excel

Uncovered by a repetition study conducted by Baggerly&Coombes with access to raw data and 2,000 hours of effort









Out of I22 papers in ASPLOS, ISMM, PLDI, TACO, TOPLAS

90 evaluated execution time based on experiments

71 of these 90 papers ignored uncertainty



Mytkowicz, Diwan, Hauswirth, Sweeney. Producing Wrong Data Without Doing Anything Obviously Wrong! ASPLOS'09



Out of I 22 papers in ASPLOS, ISMM, PLDI, TACO, TOPLAS

90 evaluated execution time based on experiments 71 of these 90 papers ignored uncertainty This lack of rigor undermines the results Yet, no equivalent to the Duke Scandal. Are we better? Is our research not worth reproducing? Is our research too hard to reproduce?

Reproduction

... independent researcher implements/realizes the published solution from scratch, under new conditions

Repetition

... re-doing the same experiments on the same system and using the same evaluation method

Is our research hard to repeat ? Is our research hard to reproduce ?

Break new ground in

hard real-time concurrent garbage collection

GC in 3 minutes

- Mutation
- Stop-the-world
- Root scanning
- Marking
- Sweeping
- Compaction

Phases

Mutation

- Stop-the-world
- Root scanning
- Marking
- Sweeping
- Compaction

thread#1

heap

thread#2

- Mutation
- Stop-the-world
- Root scanning
- Marking
- Sweeping
- Compaction

- Mutation
- Stop-the-world
- Root scanning
- Marking
- Sweeping
- Compaction

- Mutation
- Stop-the-world
- Root scanning
- Marking
- Sweeping
- Compaction

thread#1

heap

- Sweeping

thread#1

heap

- Sweeping

Phases

- Mutation
- Stop-the-world
- Root scanning
- Marking
- Sweeping

Compaction

Incrementalizing marking

Collector marks object

Application updates reference field

Compiler inserted write barrier marks object

Incrementalizing compaction

- Forwarding pointers refer to the current version of objects
- Every access must start with a dereference

original

copy

Obstacles

- No real-time benchmarks for GCed languages
- No clear competition, two GC algorithms claim to be best
- No accepted measurement methodology
- No open source experimental platform for comparison

Develop an open source experimental platform Picked the Real-time Specification for Java

First generation system, about 15 man/years Flew on a Boeing ScanEagle

Second generation system, about 6 man/years Fiji Systems Inc.

A Real-time Java Virtual Machine for Avionics. TECS, 2006

Observations

Results on noncompetitive systems not relevant

Much of work went into a credible research platform

HotSpot I.6 Server IBM J9 Sun Java RTS 2.1 IBM Metronome SRT Fiji VM CMR Fiji VM Schism/cmr/c Fiji VM Schism/cmr/a Fiji VM Schism/cmr/cw

SpecJVM98

Collision Detector Benchmark In Java, Real-time Java, and C (Linux/RTEMS)

Measure response time, release time jitter

Simulate air traffic control Hard RT collision detector thread Scalable stress on garbage collector

About I.5 man/years

A family of Real-time Java benchmarks. CC:PE 2011

Observation

Understanding what you measure is critical

Running on a real embedded platform and real-time OS, difference between Java & C small...

Good news?

No. The LEON3 lacks a FP unit, & the benchmark is FP intensive...

Step 3

Gain experience with the state of the art

Experiment with different GC techniques

GC in uncooperative environment **Brooks forwarding Object** replication **Object** handles

About 2 man/years

Accurate Garbage Collection in Uncooperative Environments. CC:P&E, 2009 Hierarchical Real-time Garbage Collection. LCTES, 2007 Replicating Real-time Garbage Collector. CC:P&E, 2011 Handles Revisited: Optimising Performance and Memory... ISMM, 2011

Observation

Trust but verify, twice.

- From workshop to journal, speed 30% better
- Good news?
- Later realized switching to GCC 4.4 slowed baseline (GCC didn't inline a critical function)
- Once accounted for this our speed up was 4%...
- A correction was issued...

Reproduce state of the art algorithms from IBM and Oracle

Metronome, Sun Java RTS

Choose measurement methodology

Existing metric (MMU) inadequate

About 3 man/years

Scheduling Real-Time Garbage Collection on Uniprocessors. TOCS 2011 Scheduling Hard Real-time Garbage Collection. RTSS 2009

Observation

Reproduction was difficult because of closed-source implementations & partial description of algorithms

Repetition was impossible because no common platform

Step 5

Develop a novel algorithm

Fragmentation tolerant Constant-time heap access

About 0.6 man/years

Schism: Fragmentation-Tolerant Real-Time Garbage Collection. PLDI 2011

Schism: objects

Avoid external fragmentation by splitting objects in 32-byte chunks

normal object

split object

Schism: arrays

For faster array access, array = variable sized spine
 + 32-byte chunk payload

normal array

spine

payload

In summary, 28 m/y reproduction .6 m/y novel work

Experimental platform	21	man/years
Benchmark	2	man/years
Implementing basic techniques	2	man/years
Reproduction of state-of-the art +measurement methodology	3	man/years

Implementing novel algorithm

0.6 man/years

$$\begin{aligned} & = (t_{7,1_{2})} \mathcal{E} \left(c_{0v} \left(F_{j_{n_{u}}} \right)^{2} \right) \\ & = (t_{2}) \mathcal{E} \left(c_{0v} \left(F_{j_{n_{u}}} \right)^{2} \right) \\ & = (t_{2}) \mathcal{E} \left(\mathcal{E} \left(F_{u} \right)^{2} \right) \\ & = (t_{2}) \mathcal{E} \left(\mathcal{E} \left(F_{u} \right)^{2} \right)^{2} \right) \\ & = E \left(\frac{1}{\prod_{k=3}^{n+1} n_{k}} \frac{1}{n_{2} - 1} \sum_{j_{n+1} = 1}^{n_{n+1}} \cdots \sum_{j_{2} = 1}^{n_{2}} \left(\frac{\overline{Y}_{j_{n+1} \dots j_{2} \bullet} - \overline{Y}_{j_{n+1} \dots j_{2} \bullet} \right)^{2}}{n_{1}} \right) \\ & = E \left(\frac{1}{\prod_{k=3}^{n+1} n_{k}} \frac{1}{n_{2} - 1} \sum_{k=1}^{n_{2}} (2z_{k} - \overline{2z})^{2} \right) = \underbrace{\sigma_{1}^{2}}_{n_{1}} + \sum_{k=2}^{n+1} \sigma_{k}^{2} - \sum_{k=3}^{n+1} \sigma_{k}^{2} - \underbrace{\sigma_{1}^{2}}_{n_{1} + 1 \dots j_{2}} \left((u_{1})_{j_{n+1} \dots j_{2}} \right) \right) \\ & = E \left(\frac{1}{n_{2} - 1} \sum_{k=1}^{n_{2}} (2z_{k} - \overline{2z})^{2} \right) = \underbrace{\sigma_{1}^{2}}_{n_{1}} + \sum_{k=2}^{n+1} \sigma_{k}^{2} - \sum_{k=3}^{n+1} (u_{1})_{j_{n+1} \dots j_{2}} \right) \end{aligned}$$

Cater for random effects, non-determinism

Repeat experiment runs, summarize results Threat to validity detectable by failure to repeat

Guard against bias

Use multiple configurations, hardware platforms Threat to validity detectable by failure to reproduce

Jain: The Art of Computer Systems Performance Analysis Lilja: Measuring Computer Performance, A Practitioner's Guide Evaluate Collaboratory, <u>http://evaluate.inf.usi.ch/</u>

Repeatability

Enable repetition studies Archival *Automate and archive* Disclosure

Share experimental details

Reproducibility

Community support for focused reproductions Open benchmarks and platforms

Reward system for reproductions Publish reproduction studies Regard them as 1 st class publications

The Correspondence Principle for Idempotent Calculus and some Computer Applications

Grigori L. Litvinov and Victor P. Maslov

1 Introduction

This paper is devoted to heuristic aspects of the so-called idempotent calculus. There is a correspondence between important, useful and interesting constructions and results over the field of real (or complex) numbers and atinial constructions and results over thempotent semiring, in the spirit or source the semicondense of the semicondense of the semisingle for several basic ideas, contractions and remults in Euroical Analysis and Mathematical Physics are discussed from this point of view. Thus the correspondence principle is a powerful heuristic to lot apply unexpected analogies and ideas borrowed from different areas of Mathematics and Theoretical Physics.

It is very important that some problems nonlinear in the traditional ensures (for example, the Bellman equation and its generalizations and the Hamilton-Jacobi equation) turn out to be linear over a suitable semiring; this linearity considerably simplifies the explicit construction of solutions. In this case we have a natural analog of the so-called superposition principle in Quantum Mechanics (see [1-3]).

The theory is well advanced and includes, in particular, new integration theory, new linear agebra, spectral theory and functional analysis. Applications include various optimization problems such as multicriteria decision making, optimization on graphs, discrete optimization with a large parameter media, optimial organization of parallel data processing, dynamic programming, discrete event systems, computer selects, discrete mathematics, mathematical logic and so on. See, for example, [4]–[4]. Let us indicate some application of these ideas in mathematical physica and hoppyires [63,77].

appreciations or these neess in manematical projects and nonposes [cov]=(10)= In this paper the correspondence principle is used to develop an approach to object-oriented software and hardware design for algorithms of idempotent calculus and scientific calculations. In particular, there is a regular method for constructing back-end processors and itechnical devices intended for an implementation of basic algorithms of basic algorithms.

Cambridge Books Online © Cambridge University Press

Paper

Program Committee

Key ideas

Artifact Evaluation Committee

Key ideas

Senior co-chairs

Authoritative site: http://www.artifact-eval.org/

Consistent with the Paper

Artifact

Complete

Easy to Reuse

Well Documented

Statistics from OOPSLA⁹13

2 AEC co-chairs
24 AEC members
3 reviews per AEC member
3 reviews per artifact

18 accepted

21 artifacts submitted

50 papers accepted

Artifact publication

Software, Artifact data, etc. key info

TreatJS: Higher-Order Contracts for JavaScript	— Title
Matthias Keil and Peter Thiemann	A - 1
Institute for Computer Science University of Freiburg Freiburg, Germany {keilr,thiemann}@informatik.uni-freiburg.de	Authors
- Abstract	
Treat system for JavaScript which enforces con- tract system for JavaScript which enforces con- tracts by run-time monitoring. Beyond providing the standard abstractions for building higher-order contracts (base, function, and object contracts). TreatJS is novel contributions are its guarantee of non-interfering contract execution, its systematic approach to building blogher onter tracts in the style of union and intersection types, and its notion of a parameterized contract sope which is the building block for composable run- time generated contracts.	- Abstract
1998 ACM Subject Classification D.2.4 Software/Program Verification Keywords and phrases Higher-Order Contracts, JavaScript, Proxies Divised Object Identifier, 10.4230/DAPTS 1.1.1	Matadata
Related Article Matthias Keil and Peter Thiemann, " <i>TreatJS</i> : Higher-Order Contracts for JavaScript", in Proceedings of the 29th European Conference on Object-Oriented Programming (ECOOP 2015),	
LtPics, Vol. 37, pp. 28–51, 2015. http://dx.doi.org/10.4230/LtPics.ECOOP.2015.28 Related Conference 29th European Conference on Object-Oriented Programming (ECOOP 2015), July 5–10, 2015, Prague, Czech Republic	(DOI, etc.)
1 Scope	
The artifact is designed to support repeatability of all the experiments of the companion paper, allowing users to test the contract system on a variety of benchmarks. In particular, it allows to include <i>TreatJS</i> in existing JavaScript code, to specify contracts by plain JavaScript functions, to construct contracts by an unrestricted combination of other contracts, and to enforce contracts in	Scope,
2 Content	content,
The artifact package includes: = the main source of <i>TreatJS</i> ;	lieenee
 a set of test cases to examine the feature of the contract system; modified version of the Google Octaine 2.0 benchmark suite; detailed instructions for using the artifact, provided as an index.html file. 	license,
© O Matthias Kell and Peter Thiemann; Bienened under Creative Commons Attribution 3.0 Germany (CC BY 3.0 DE) Dagstuhl Artifacts Series, Vol. 1, Isue 1, Artifact No. 1, pp. 01:1-01:2 Dagstuhl Artifacts Series Dagstuhl Artifacts Series Mattriacti Stenies Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany	etc.

TreatJS: Higher-Order Contracts for JavaScript

Matthias Keil and Peter Thiemann

Institute for Computer Science University of Freiburg Freiburg, Germany (keiır, thiemann)@informatik.uni-freiburg.de

---- Abstract

TreatJS is a language embedded, higher-order contract system for JavaScript which enforces contracts by run-time monitoring. Beyond providing the standard abstractions for building higher-order contracts (base, function, and object contracts), TreatJS's novel contributions are its guarantee of non-interfering contract execution, its systematic approach to blame assignment, its support for contracts in the style of union and intersection types, and its notion of a parameterized contract scope, which is the building block for composable run-time generated contracts that generalize dependent function contracts.

Treat.S is implemented as a library so that all aspects of a contract can be specified using the full JavaScript language. The library relies on JavaScript proxies to guarantee full interposition for contracts. It further exploits JavaScript's reflective features to run contracts in a sandbox environment, which guarantees that the execution of contract code does not modify the application state. No source code transformation or change in the JavaScript run-time system required. The impact of contracts on execution speed is evaluated using the Google Octane benchmark.

1998 ACM Subject Classification D.2.4 Software/Program Verification

Keywords and phrases Higher-Order Contracts, JavaScript, Proxies

Digital Object Identifier 10.4230/LIPIcs.ECOOP.2015.28

Supplementary Material ECOOP Artifact Evaluation approved artifact available at http://dx.doi.org/10.4230/DARTS.1.1.1

1 Introduction

A contract specifies the interface of a software component by stating obligations and benefits for the component's users. Customarily contracts comprise invariants for objects and components as well as pre- and postconditions for individual methods. Prima facie such contracts may be specified using straightforward assertions. But further contract constructions are needed for contemporary languages with first-class functions and other advanced abstractions. These facilities require higher-order contracts as well as ways to dynamically construct contracts that depend on run-time values.

Software contracts were introduced with Meyer's Design by Contract[™] methodology [39] that stipulates the specification of contracts for all components of a program and the monitoring of these contracts while the program is running. Since then, the contract idea has taken off and systems for contract monitoring are available for many languages [33, 1, 37, 32, 12, 22, 11, 10] and with a wealth of features [35, 31, 7, 20, 46, 16, 2]. Contracts are particularly important for dynamically typed languages as these languages only provide memory safety and dynamic type safety. Hence, it does not come as a surprise that the safety of the need to higher-order contract systems were devised for Scheme and Racket [24], out of the need to

 Austhalas Kell and Peter Thiemann: Leneaed under Contive Commons License CC-BY 20th European Conference on Object-Oriented Programming (ECOOP'15). Editor: John rang Boyland; pp. 28-51 Lenbinz International Proceedings in Informatics EIPICS Schisse Dagstuhl - Leibniz-Zentrum für Informatik, Dagstuhl Publishing, G

DOI cross-ref

> AEC badge

Artifacts as first-class citizens

[+] Jonathan Aldrich 🔺 🕹 😪 🖛

> Home > Persons

[-] Person information

affiliation: Carnegie Mellon University, Pittsburgh, USA

[-] 2010 – today 🔮

2015

[c84] E 🖄 🗟 🗟 Seph Lee, Jonathan Aldrich, Troy Shaw, Alex Potanin:
 A Theory of Tagged Objects. ECOOP 2015: 174-197

Conclusions

Develop open source benchmarks

- Codify documentation, methodologies & reporting standards
- Require executable artifacts
- Publish reproduction studies