

How the
JVM spec
came to
be

*James
Gosling*



Page 1 of 33

or: Just how General
Purpose is it?

Page 2 of 33

The Prehistory of Java

- What shaped my thinking
- Something of a geek autobiography

Page 3 of 33

First Learned

- FOCAL
 - build stuff fast!
- PDP-8 assembler
 - build fast stuff!
for the day :-)

Page 4 of 33

Eventually

- Fortran, Cobol, Basic, CDC assembler, ...
 - jobs
 - Multics Pascal compiler
 - VAX Cobol (never shipped as product)
- Pascal, Simula, Lisp, Algol, IBM assembler, ...
 - when I finally went to school

Page 5 of 33

Designing my first scripting language: YORKMT

- Started life as a magtape copy program
- Feature creep added data correction, imaging, device handling, ...
- Too many changes!
- Experience with TECO suggested adding scripting. Success!
- Users did unexpected things: instrument calibration, diagnostics, ...

Page 6 of 33

Lessons along the way

- Threading is cool
 - Hydra, Cm*, SMP BSD, Hibbards Algol68
- ANDF is easy (!)
 - PERQ->VAX code migration (the PhD thesis I should have written)
- The power of trustworthy pointers
 - Multics, Hydra, Cm*, Pascal, Mesa
- Optimizers can do surprising things
 - BLISS, MUMBLE (the other thesis I should have written)
- Code-as-algebra is awesome!
 - PhD thesis, ACE

Page 7 of 33

Unix Emacs

- Reinforced YORKMT lessons
 - Power of scripting
 - People do the damndest things (everything gets stretched to its limits)
- Added
 - The power of the community

Page 8 of 33

NeWS

- Networked Extensible Window System (predated X11)
- A window system based on scripting.... PostScript (!!)
- More insanely unusual users
 - (never kid yourself into believing you know what will be done with what you've built)
- "Technical" success, but two nagging issues:
 - Scripting performance
 - (cope with user surprise scale)
 - Security

Page 9 of 33

GreenTalk

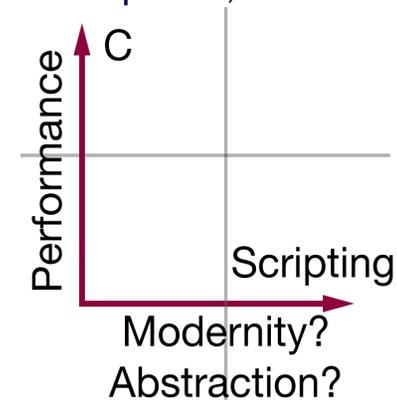
Page 11 of 33

Trivia Question:

What was Java called, before it was called "oak"?

Page 10 of 33

Perceptions, circa 1990

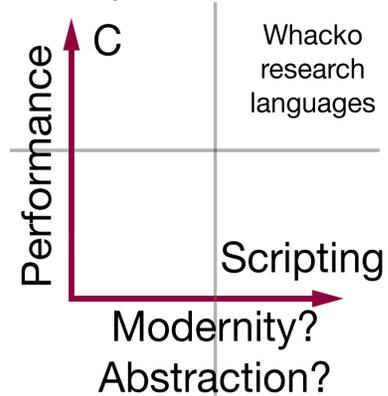


Page 12 of 33

This dichotomy always bugged me

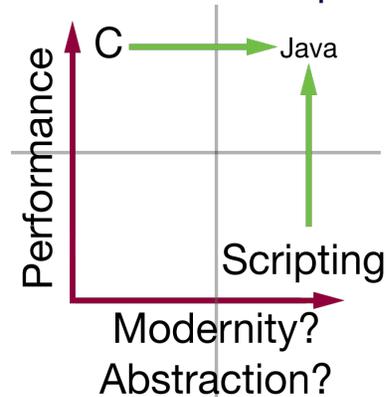
This perception always bugged me

Perception of Academia



A Big (but quiet) Goal:
How close could I get to a "scripting" feel, without giving up too much?

A subversive compromise



Page 17 of 33

The original concept

- Was all about building networks of things, orchestrated by a scripting language
- (Unix shells, AppleScript, ...) but I was real nervous

Page 19 of 33

A Wolf in Sheeps Clothing

- C syntax to make developers comfortable
- Lisp/Smalltalk/Mesa/Simula/ ... to get the job done

Page 18 of 33

Commercial software priority list, 1990

- Compatibility (V2V)
- Performance
- Portability
- Reliability
- Networking
- Multithreading
- Security

Page 20 of 33

Consumer Electronics priority list

- Security (safety)
- Networking (connectivity)
- Portability (cost: CPU du jour)
- Reliability
- Performance
- Multithreading
- Compatibility (rewrite for every product)

Page 21 of 33

Shake and Stir

- Security
- Reliability
- Networking
- Portability
- Compatibility
- Multithreading
- Performance

Page 22 of 33

BUT

- Performance that sucks
- Sucks

Page 23 of 33

Design space is gravitationally lumpy

"Lisp is a black hole. If you try to design something like Lisp, it gets sucked in and becomes Lisp"

Guy Steele, hallway conversation, 1980ish

But orbital dynamics are chaotic

- Byte coding is another black hole

Page 24 of 33

ANDF

- Architecturally Neutral Distribution Format
- One of the Holy Grails of the early 1990s
- Got bogged down by politics and complexity
- Tended to be persistent forms of ASTs
- Which said too much about how compilers were build
- Food Fight

Page 25 of 33

The Big Restriction:

- pointer/object integrity
 - ie. no fraud
- Some languages depend on fraud:
 - C/C++/Objective C

Page 27 of 33

Virtual Machines to the Rescue

- PERQ experience said performance could be good
- verifier invented to enforce security issues
 - imposes restrictions on control&data flow
 - that are also useful for compilers
- Says "just enough"
- Abstracts-away AST issues

Page 26 of 33

The Other Big Restriction

- Reactionary method calls
 - Simple simula-esque object model
 - The more complex object models didn't seem to be worth their weight
- A building block
- Tricks with interfaces

Page 28 of 33

Why a stack machine?

- (why not?)
- No preconceived register architecture
- Compact
- "free" live/dead analysis
- reverse polish encoding of AST

Page 29 of 33

Primitive Types

- They're all about performance
- Design Goal:
a=b+c
- in one instruction

Page 31 of 33

Control Flow

- Goto statement
 - Caused flow analysis problems
 - Poof! Dead!
- JSR/RET
 - oops!

Page 30 of 33

Java performance

- The myth:
 - it's interpreted, therefore slow
- The truth:
 - it's highly optimized
 - dynamic metrics
 - generally beats C/C++
 - 2% LINPACK, +4% SciMark
 - often beats Fortran (matrices are the issue)
 - GC is a lot faster than malloc/free
- Dynamic compilation beats static
- No need to be afraid of methods!
 - who needs opcodes?

Page 32 of 33

Thanks!

