

Sly and the RoarVM

Exploring the Manycore Future of Programming

Renaissance Project


<http://soft.vub.ac.be/~smarr/renaissance/>



Stefan Marr

Software Languages Lab, Vrije Universiteit Brussel

Smalltalk Dev Room, FOSDEM, 2012-02-05





Agenda

- Parallel Smalltalk and the RoarVM
 - [self demo] fork
 - Project
- Sly and the Renaissance Project
 - Demo
 - Context
 - Language and Use-Cases





Disclaimer: It wasn't me!

IBM Research

- **David Ungar**
- Sam S. Adams
- Doug Kimelman

Portland State University

- Andrew P. Black
- Max OrHai

Vrije Universiteit Brussel: Software Languages Lab

- Pablo Inostroza Valdera
- Stefan Marr
- Theo D'Hondt



The screenshot displays the Pharo IDE interface. At the top left is the Pharo logo. A 'Demo Steps' window contains the following text:

```
"Parallel processing is for speed.
Let's do a simple benchmark."
Transcript cr; show: 'Started.'.
[ 39 benchFib ] timeToRun.

"How to run something in Parallel?"
```

A terminal window shows the execution of a git clone and a make command, followed by a list of compiled files:

```
8. smarr@minerva: ~/FOSDEM/RoarVM/roarvm-code/vm/build (bash)
smarr@minerva:~/FOSDEM/RoarVM$ git clone --depth 1 git://github.com/smarr/RoarVM.git roarvm-code
Cloning into roarvm-code...
remote: Counting objects: 2949, done.
remote: Compressing objects: 100% (2019/2019), done.
remote: Total 2949 (delta 1165), reused 2316 (delta 752)
Receiving objects: 100% (2949/2949), 16.16 MiB | 879 KiB/s, done.
Resolving deltas: 100% (1165/1165), done.
smarr@minerva:~/FOSDEM/RoarVM$ cd roarvm-code/vm/build/
smarr@minerva:~/FOSDEM/RoarVM/roarvm-code/vm/build$ ./configure
smarr@minerva:~/FOSDEM/RoarVM/roarvm-code/vm/build$ make -j2
Checking Compiler for know problems
Compiling aio.c
Compiling B2DPlugin.c
Compiling BitBlitPlugin.c
Compiling bytemap.cpp
Compiling core_tracer.cpp
Compiling abstract_tracer.cpp
Compiling oop_tracer.cpp
Compiling error_handling.cpp
Compiling externals.c
Compiling FilePlugin.c
Compiling FloatArrayPlugin.c
Compiling interpreter_bytecodes.cpp
Compiling interpreter_primitives.cpp
Compiling LargeIntegers.c
Compiling Matrix2x3Plugin.c
Compiling measurements.cpp
```

A 'Workspace' window shows a transcript of the execution:

```
[ Transcript show: ' Started.'.
Transcript show: [ 34 benchFib ]
timeToRun asString, ' '. ] fork.
```

A 'Transcript' window is also visible, showing a similar transcript. A 'Multiaspect Kiviat' window displays a performance graph with the following data:

```
min: 0
max: 56.2K
total: 42.3K
avg: 21.1K
```


At the bottom of the IDE, a 'Run' button is visible.

RoarVM and Parallel Smalltalk?

Get the RoarVM

Smalltalk, Concurrency, and Parallelism






RoarVM Smalltalk




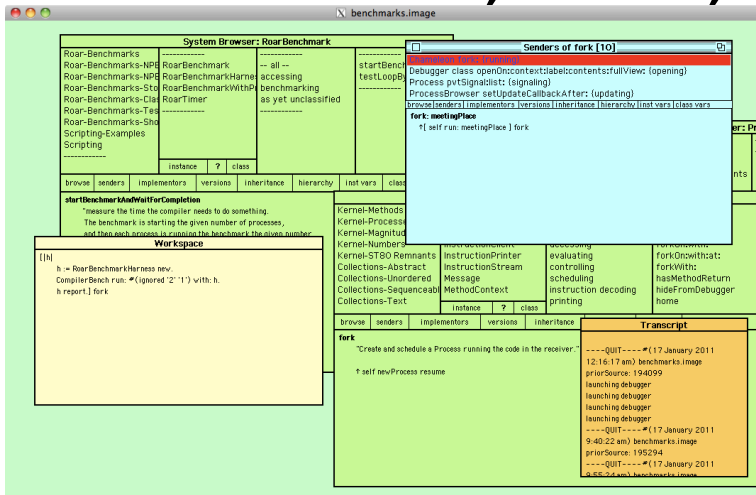
- Process class additions
 - yieldIfFewerCoresThan: anInt
 - coreMask, coreMask: anInt
 - onlyOnCore: anInt
 - hostCore
 - thisProcess instead of ~~activeProcess~~
- Uses standard Smalltalk facilities
 - [self do] fork
 - Semaphore class





RoarVM Status

- VM: rewritten Squeak, 20k SLOC C++
- MVC images run fine
- Squeak/Pharo: multicore instabilities
- Runs on: OSX, Linux, iOS, Tileria, up to 59 cores



Project Infrastructure

- Buildbot: Automate compilation, testing
 - Linux, OS X, Tiler
- Performance Tracking
 - Codespeed
 - ReBench
 - SMark
- <http://github.com/smarr/RoarVM>

25c2bfe49485... Stefan Marr


Added configure params to first commented line of generated makefile

- ManyCore-Tiler-release: tests 1 failure 1 skip - [stdio](#) [test.log](#) [problems](#)

64ba148f696d... Stefan Marr

Changed a loop that did a while to use while syntax for readability

- ManyCore-Tiler-release: tests 1 failure 1 skip - [stdio](#) [test.log](#) [problems](#)






Avoid the Concurrency Trap by Embracing Non-Determinism

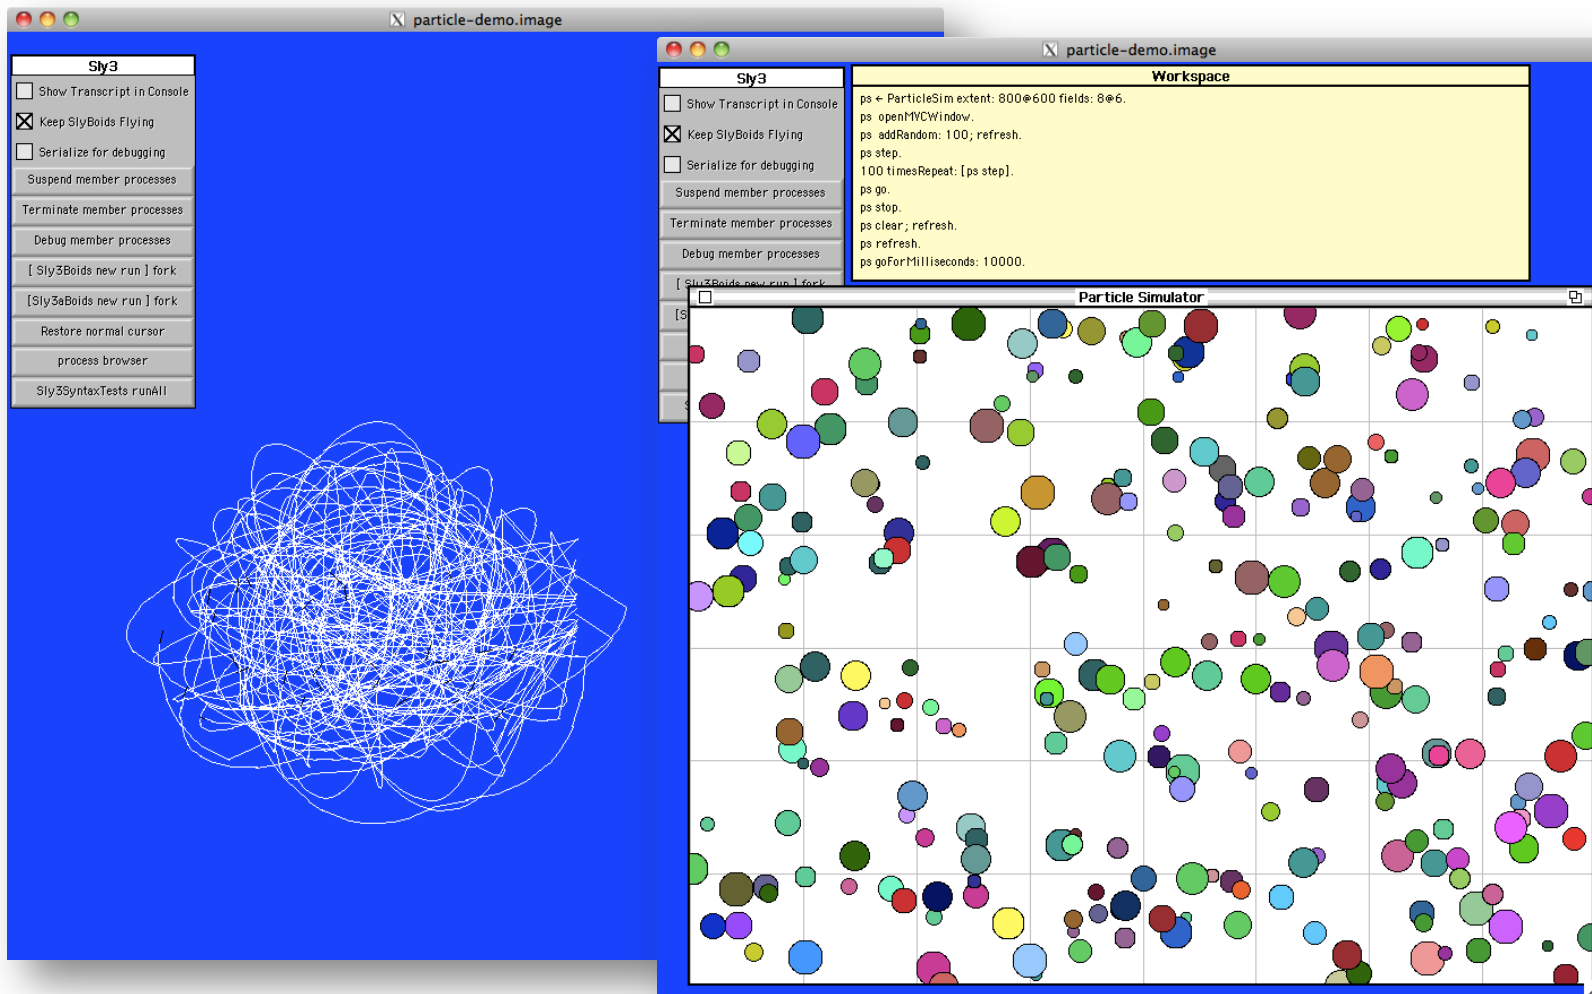
SLY AND THE RENAISSANCE PROJECT



Renaissance Project

- Manycore: > 100 cores
 - + non-uniform memory access
- Problem
 - Parallel systems are nondeterministic
 - Communication is expensive
 - Synchronization == bottleneck
- The vision
 - Embrace nondeterminism
 - Harness emergence
 - Confidence → Delay





Non-deterministic Boids and Particle Sorting

Flocking behavior of birds and particle bouncing as local-only synchronization strategies for scalability



Sly

- Ensembles: collections representing a whole, multi-part entities
 - e.g. a flock of birds
- Adverbs: modifiers to operands
 - individualLY, serialLY, randomLY: n
- Gerunds: reduction semantics
 - averagING, selectINGpredicate, ensemblING



Examples of Boids

- Every boid is its own thread, no synch.

computeCentroid

```
^ self flock boids averagINGserialLYposition
```

computeNeighbors

```
^ self flock boids selectINGisNear: self
```

matchVelocityWithNeighbors

```
^ self neighbors averagINGvelocity / 8.0
```

More details: <http://soft.vub.ac.be/~smarr/renaissance/sly3-overview/>



Possible Real Life Application






Image:
<http://blogs.technet.com/b/andrew/archive/2007/08/22/olap-cubes-and-multidimensional-analysis.aspx>



Possible Real Life Application

- Online analytical processing (OLAP)
 - Business intelligence, data mining
 - Query multidimensional data sets
 - Calculated data dimensions
- Gesture recognition
 - Many rules, low-precision input
- Race-and-repair: how to bound error?



CONCLUSION




Conclusion

- Sly: map/reduce-like programming model
 - Based on ensembles, adverbs and gerunds
- RoarVM
 - Research artifact, stable, used daily
 - No optimization of sequential performance
 - Tested up to 59cores
 - Shows weak scalability



Project, Sly Intro, Community



Multicore Programming The SLY3 Programming Language

Pablo Inostroza Valdera
<pinostro@vub.ac.be>

23 June 2011

1 Introduction


SLY3 is a Smalltalk dialect that features two concepts in order to deal with parallelism: *ensembles*, which are collections of objects that can receive parallel messages, and *modifiers*, which decorate messages in order to modify their parallel behavior (*adverbs*) or to specify how to handle reductions (*gerunds*). In this article we analyze SLY3's programming model.

SLY3 is a language being developed in the context of the Renaissance Project¹, at IBM. With SLY3, their authors David Ungar and Sam Adams explore abstractions for nondeterministic parallel programming. SLY3² is the experimental successor of LY, a javascript-like interpreted language that was implemented on top of Smalltalk [UA10]. The main idea behind the design of SLY3 is that multicore programming is commonly addressed by restricting the nondeterminism parts and trying to enforce strict determinism. As an example of this, Ungar and

[http://soft.vub.ac.be/
~smarr/renaissance/](http://soft.vub.ac.be/~smarr/renaissance/)

[http://soft.vub.ac.be/
~smarr/renaissance/
sly3-overview/](http://soft.vub.ac.be/~smarr/renaissance/sly3-overview/)

<http://github.com/smarr/RoarVM>
[@roarvm](https://twitter.com/roarvm)



Papers

- [1] J. Pallas and D. Ungar. Multiprocessor Smalltalk: A Case Study of a Multiprocessor-based Programming Environment. In PLDI '88, p. 268–277. ACM, 1988.
- [2] D. Ungar and S. Adams. Hosting an Object Heap on Manycore Hardware: An Exploration. In DLS'09, p. 99-110. ACM, 2009.
- [3] D. Ungar, D. Kimelman, and S. Adams. Inconsistency Robustness for Scalability in Interactive Concurrent-Update In-Memory MOLAP Cubes. In Inconsistency Robustness Symposium 2011.

