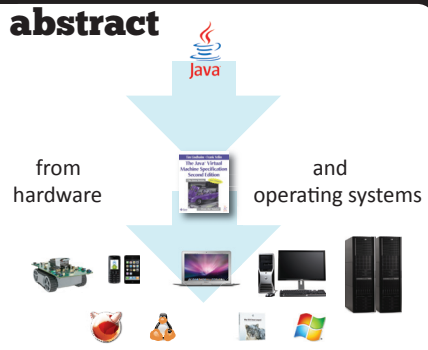


Many-Core Virtual Machines

Decoupling Abstract From Concrete Concurrency

Virtual Machines



for multiple languages



- powered by fast JIT compilers, and great GCs
- foundation for multi-language VMs
- allow to reuse existing infrastructure
- require huge investments
- reuse is economically necessary

Abstraction by ILs

- VM Intermediate Languages (ILs)
- often defined as bytecodes
- expressive abstraction for various target languages
- state of the art is very diverse

	Abstraction Model	#Register	Execution Mode	length in Bytes	#OpCodes	
CLI	Bytecode stack	0	variable	>= 1	217	
CPython	Bytecode stack	0	switch	variable 1 or 3	102	
Dalvik VM	Bytecode register	@threaded	variable	>= 2	218	
Dis VM	Bytecode memory-to-memory	0	variable	1 - 33	158	
Erlang	Bytecode register	1024	threaded, JIT	fixed	4	148
JVM	Bytecode stack	0	variable	>= 1	201	
Lua	Bytecode register	255	switch	fixed	4	38
Mozart	Bytecode register-memory	@threaded	variable	4 - 24	97	
Parrot	Bytecode register	@switch, threaded, JIT	variable	>= 4	>1200	
PHP	Bytecode register-memory	@threaded	fixed	76	136	
Rubinius	Bytecode stack	0	JIT	variable	4-16	89
Ruby 1.8	AST stack	0	switch	-	105	
Ruby 1.9	Bytecode stack	2	threaded	variable	>= 32	77
Self	Bytecode stack	0	JIT	fixed	1	17
Squeak	Bytecode stack	0	switch, threaded	variable 1 or 2	71	
TraceMonkey	Bytecode stack	1	threaded, JIT	variable	>= 1	234
VB	AST	-	JIT	-	38	

but...

concurrency support is limited

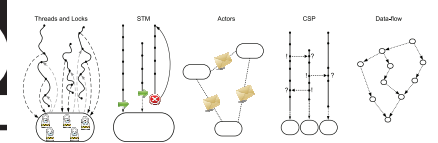
- VM support is minimal
- only one specific concurrency model is supported
- only few ILs provide notion of concurrency
- no comprehensive abstraction

Model	Threads/Locks	CSP	Actors	Data-flow
IL Support	Marginal	High-level	High-level	Marginal
Stdlib	Low/high-level	High-level	High-level	High-level

Stefan Marr, Michael Heule, and Theo D'Hondt
Intermediate Language Design of High-Level Language Virtual Machines:
Towards Comprehensive Concurrency Support
In Proc. of the 3rd Workshop on Virtual Machines and Intermediate Languages, ACM, October (2009)

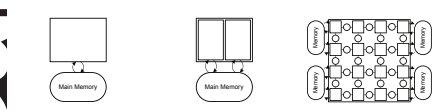
a VM has to: decouple abstract concurrency models

- abstract concurrency models are defined by languages or libraries
- used by application developers

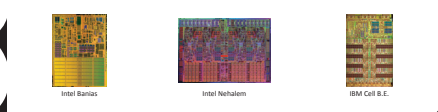


- wide range of models supported by VM is necessary
- implementing unsupported models on top is hard
 - restrictions hinder efficient implementation
- support at VM-level allows reuse and optimization

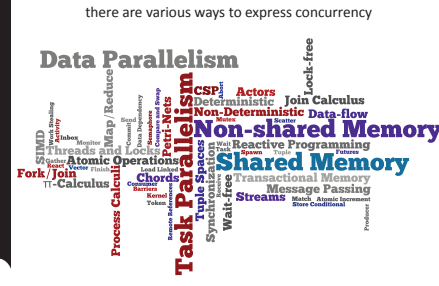
and concrete concurrency models



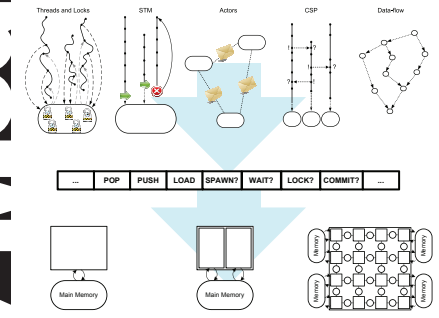
- Single-Core**
 - preemptive OS threads
 - instruction-level parallelism
 - VM challenges
 - deep cache hierarchies
 - cache-consciousness required
- Multi-Core**
 - uniform memory access
 - native support for thread-level parallelism
 - and cache coherency
 - locality and cache hierarchy must be considered
 - avoid cache thrashing
- Many-Core**
 - non-uniform memory access architectures
 - can have explicit core-to-core communication
 - very diverse designs
 - with/without cache coher.
 - explicit inter-core com.



what to include in ILs?



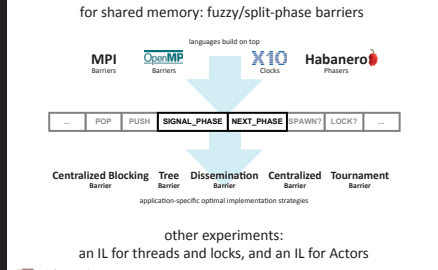
how to combine the various models?



will result in a methodology



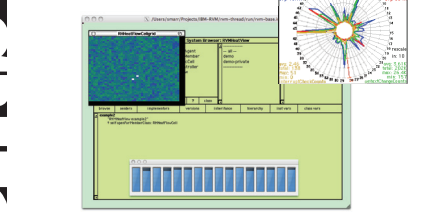
extending ILs



other experiments: an IL for threads and locks, and an IL for Actors

our VM

- A Smalltalk VM for many-core systems
- runs on the 64-core TILE64 chip
- runs on standard Intel systems
- supports Linux and OS X



In cooperation with David Ungard and Sam Adams from **IBM Research**

our hardware

for the validation of our research we can use a wide spectrum of state of the art technology

- Apple MacPro**
 - 2 CPUs
 - 4 cores per CPU
 - 2 threads per core
 - complex cache and memory hierarchy
- TILEPro64**
 - 64 cores
 - explicit core-to-core communication
 - small caches
 - shared coherent memory

