

## MUCH MODELS! MANY SIMPLICITY?

Raphael 'kena' Poss, University of Amsterdam

Putting Heterogeneous High-Performance Computing at the Fingertips of Domain Experts Shonan Village Center, Japan, November 17-20, 2015



Wisdom acquired:

C++ with templates = HPC + functional programming without the name Open standards essential for quality and against vendor lock-in





Programmers really do not understand caching



Assigned to work on the **unification of run-time back-ends** for a HPC **streaming dataflow coordination language** (Hi Clemens!)

Teaching: computer architecture

Wisdom acquired:

No coordination possible without operational semantics in PL ("pure FP for HPC is like trying to run with one leg")

**Programmers do care about operational semantics** But: only a little!

Programmers do not understand implicit caches and explicit memory











Diversity of ISA, system arch and micro-arch





How to explain complex parallel systems to 1st term undergrads so they can successfully write efficient & scalable code by their 2nd term?



- How to explain complex parallel systems to 1st term undergrads so they can successfully write efficient & scalable code by their 2nd term?
- What would Unix look like if designed for 21st century computers?



- How to explain complex parallel systems to 1st term undergrads so they can successfully write efficient & scalable code by their 2nd term?
- What would Unix look like if designed for 21st century computers?
- What should be the BASIC and C of our new computers?

#### WHAT'S AN AMM (ABSTRACT MACHINE MODEL)?

- An abstraction of a class of computer systems sufficient to reason about its behavior while executing a class of algorithms
- Example "pure AMMs" without a specific PL or code IR: Turing, Von Neumann, Harvard, P-RAM, Dataflow, BSP, Unix
- All modern PL AMMs are simulated upon a platform, using assumptions about the platform defined in the platform's AMM (usually VN/P-RAM)

#### **TERM RELATIONSHIPS**



"*Computing models*" = Union of all three

. . . . . . . . . . . . . . .

- Von Neumann / Harvard / P-RAM (threads, etc.)
  - ► Building blocks: N active sequential processors + 1 RAM
  - ► Perf. predictive power only when programmed via BSP
  - Performance scalability OR software composability, choose only 1!

- Von Neumann / Harvard / P-RAM (threads, etc.)
  - ► Building blocks: N active sequential processors + 1 RAM
  - ► Perf. predictive power only when programmed via BSP
  - Performance scalability OR software composability, choose only 1!
- Dataflow-like (incl. most RTS for lazy FP evaluation)
  - ► Building blocks: N on-demand, "passive" functional units + 1 tag soup
  - ► Perf. predictive power only for small tag stores and static graphs
  - Terrible at expressing priorities in mixes of data / control tasks
    Poor at controlling locality

- Von Neumann / Harvard / P-RAM (threads, etc.)
  - ► Building blocks: N active sequential processors + 1 RAM
  - ► Perf. predictive power only when programmed via BSP
  - Performance scalability OR software composability, choose only 1!
- Dataflow-like (incl. most RTS for lazy FP evaluation)
  - ► Building blocks: N on-demand, "passive" functional units + 1 tag soup
  - ► Perf. predictive power only for small tag stores and static graphs
  - Terrible at expressing priorities in mixes of data / control tasks
    Poor at controlling locality
- Especially missing: models for heterogeneous platforms, ability to predict/control extra-functional behavior other than time



(PLs with no AMM nor IR)



(PLs with no AMM nor IR)

















#### What we think we have:



#### What we actually have:



#### What we think we have:



#### What we actually have:



HLPL

e.g. SAC,

Haskell

#### What we think we have:



#### What we actually have:



#### What we think we have:

![](_page_31_Figure_2.jpeg)

#### What we actually have:

![](_page_31_Figure_4.jpeg)

#### WHERE WE COULD GO

![](_page_32_Figure_1.jpeg)

## **EVOLVING ABSTRACT MACHINE MODELS**

- ► Previously:
  - Co-evolution of AMM and architecture design independent from application code
  - ► OS/PL implementers "subject" to platform AMM
  - Platform AMM seeps through app. design and crystallizes there, makes translations of code to another AMM difficult
  - Back-compatibility with previous platform AMMs paramount

## **EVOLVING ABSTRACT MACHINE MODELS**

- ► Now:
  - Diversity of HLPLs each with own operational semantics and AMM
  - Platform AMM primary audience is OS/PL implementer, not app. dev.
  - app. dev only cares about platform AMM when overriding OS/PL intelligence
  - Hardware vendors listen to OS/PL demands about platform AMM in new architectures (cf. GPGPUs)
- Suggested methodology: Evolve arch. and AMM using HLPL design as guide!

#### **PROGRAMMING MODELS FOR 2010–2030**

. . . . . . . . . . . . . . . . . .

#### **PROGRAMMING MODELS FOR 2010-2030**

#### ► Surprise observation:

beginner programmers are comfortable coding with **process networks** (PNs) and conceptualizing / explaining data flows

- Suggests streaming calculus may be viable alternative to term rewriting / lambda calculus for base semantics
- As long as PL provides powerful PN assembly combinators: people hate wiring manually

#### **PROGRAMMING MODELS FOR 2010-2030**

#### ► Surprise observation:

beginner programmers are comfortable coding with **process networks** (PNs) and conceptualizing / explaining data flows

- Suggests streaming calculus may be viable alternative to term rewriting / lambda calculus for base semantics
- As long as PL provides powerful PN assembly combinators: people hate wiring manually
- Perhaps we could start with PN-based AMMs and build the rest upon this?

#### **ABSTRACT MACHINE MODELS FOR PNS**

- Example: what if we build OS/PLs upon PN semantics?
  - ► OK, then hardware must be good at supporting PNs
    - Surprise! already the case since ~1975 in-networks thanks to clusters and ~2005 on-chip thanks to NoCs
  - ► OK, then AMMs must be good at predicting behavior
    - ► Hmmm, VN/PRAM and DF just won't cut it. :(
    - We need new AMMs that model network topology, interconnect protocols and asymmetry in hw platforms (now working on this with AM<sup>3</sup>)
  - NB: Traditional distributed computing models insufficient: new need to separate cores from memory from accelerators

#### AM<sup>3</sup> – ABSTRACT MACHINE MODEL FOR MODERN (OR MANY-CORE) MACHINES

► Work in progress!

see DOI 10.1109/TPDS.2015.2492542

![](_page_39_Figure_3.jpeg)

- No preemption, no syscalls, no privilege separation in ISA, shared memory optional Network provides isolation / interfaces
- Uniform access interface to heterogeneous compute elements
  - Compute element = "virtual processor group", can be heterogeneous
  - ► Can be CUDA warp, functional unit, hardware thread group, iRAM
- Substitutes VN/P-RAM to implement OS or PL/RTS, sufficient to run Unix :)

![](_page_40_Figure_0.jpeg)

#### **OTHER AMM EVOLUTION DIRECTIONS**

- Surprise observation\*: often enough, programmers do know how to optimize better than the compiler/RTS
  - PL must provide syntax/semantics with opt-in overrides
  - HW + AMMs must provide abstractions for reasoning about (and enforcing) functional equivalence of operational mechanisms
  - (\* not only an observation, actually a theorem since 2001)
- Empirical result 2000-2015: people like run-time tuning!
  - Still incomplete support in HW + AMM: how to instrument an application systematically, cross-PL and cross-platform for comprehensive run-time tuning?

#### **PROPOSAL: THE "BASIC" LANGUAGE OF FUTURE COMPUTERS...**

- 1. has a <u>simple, easy to understand AMM</u> <u>that makes useful predictions</u> on heterogeneous platforms
- 2. provides main syntax to <u>declare intent</u> and annotations/typing to supply opt-in <u>elidable override mechanisms</u>
- has processes and dataflow links as main abstract building blocks expressed using combinators, not functions, data structures, processor logic or memory contents
- 4. provides run-time <u>tuning knobs</u> separate from alg. design to:
  - ► trade space vs. time costs
  - ► trade latency vs. throughput
  - trade performance vs. accuracy/reliability/determinism/accountability
  - trade data sharing/performance vs. isolation/privacy/segregation

# **THANK YOU!**

. . . . . . . .

Questions?