



| | |
|--|---|
| <p>Seminar 10281 11. 7. - 16. 7. 2010 Dynamically Reconfigurable Architectures</p> | <p>Schloss Dagstuhl <i>Where Computer Scientists Meet</i></p> |
| <p>Reiner Hartenstein TU Kaiserslautern</p> | <p>Why we need to Reinvent Programmer Education</p> |

| | |
|--|---|
|  <p>http://hartenstein.de © 2009, reiner@hartenstein.de TECHNISCHE UNIVERSITÄT KAISERSLAUTERN</p> | <h1>Amid the Clamor</h1> |
| <p>Michael Wrinn, (keynote at SIGCSE2010): Suddenly, All Computing Is Parallel: Seizing Opportunity Amid the Clamor http://www.sigcse.org/sigcse2010/attendees/keynotes.php</p> |  <p>the proud era of von Neumann architecture passes into history.</p> <p>Foundational change of this magnitude will disrupt traditional habits throughout the discipline ...</p> <p>... especially how students are to be introduced ...</p> <p>© 2009, reiner@hartenstein.de 2 http://hartenstein.de</p> |



RC is the silver bullet

„Multicore computers shift the burden of software performance from chip designers to programmers.“



[J. Larus: Spending Moore's Dividend; C_ACM, May 2009]

To avoid future unaffordability of our cyber infrastructure we need a massive software to configware migration

The migration of the huge supply of legacy software saves much more energy than most proposals from the climate protection scene

... impossible without reinventing programmer education (embedded system scene has not sufficient manpower)

© 2009, reiner@hartenstein.de

3

http://hartenstein.de



New Programmer Education

What existing environments are a good starting point, or, could be included?

UML ? Access to existing libraries ?

New mix of skills needed, currently not available

Skills to detect overhead and bottlenecks

Fully understanding the architecture, its nodes, and its communication and memory architecture

Essential: intuitive awareness of locality

© 2009, reiner@hartenstein.de

4

http://hartenstein.de



Memory Mapping Issues

Must focus primarily on memory mapping issues
 μP and co-processors share memory on the node,
and interconnect many nodes with
NoC, GigE, Infiniband, or a custom interconnection
configuring the nodes in a distributed memory layout



Which parallelism model

The domain scientist, or RC-aware domain scientist, should be equipped with a sufficiently transparent parallelism model

practicing time to space mapping:

- to learn more locality awareness,
- efficient solutions from the beginning, avoiding a later re-design..



Parallelism Levels

Instruction-sequential to concurrent-only: **not** the silver bullet
ILP instruction level parallelism
I-PLP instruction process level parallelism

Twin-paradigm sequential to concurrent: **the way to go!**

DLP data object level parallelism
DSLIP data stream level parallelism
ILP instruction level parallelism
I-PLP instruction process level parallelism



Visible architecture

The programming model:
hardware view presented to the programmer:
Which hardware architecture parts are visible
and which are under the programmer's direct control.

RC programming model:
whether (and how) the programmer can control :
- data transfers between FPGA and onboard memory,
- between FPGA and μP memory, as well as
- between FPGA and μP itself.



Cray-XD1 Architecture

Cray-XD1 allows Opteron μP to access FPGA internal registers, internal and external memory.



The FPGA can access the μP memory.

However, HLL use disables some of these features.

several transfer modes between μP and the FPGA (depending on its initiator).

The μP can read from / write to the FPGA local memory space (i.e. internal registers, internal BRAMS, and external memory)

FPGA can read from / write to the μP local memory space.

© 2009, reiner@hartenstein.de

9

<http://hartenstein.de>



Cray-XD1 Architecture (2)

Cray-XD1 allows Opteron μP to access FPGA internal registers, internal and external memory.



most bandwidth-efficient transfer mode:
write-only mode (producer initiates the transfer):
burst (for large amount of data) or non-burst.

© 2009, reiner@hartenstein.de

10

<http://hartenstein.de>



How datastreams are processed

The graphic interface should transparently illustrate to its user, how data streams are running through the structures.

to the Domain scientist this is much more informative, than abstract **text-only notations like languages, which do not provide physical locality awareness**. The ease of understanding should be the main objective.

Here the user should not need to learn a textual language. Instead he should know the semantics of the types of available boxes. A simulator should enable "running" the structures sticked together.

© 2009, reiner@hartenstein.de

11

<http://hartenstein.de>



Graphic Representation

We need a modeling level less abstract than by textual languages

We need a graphic box interconnect scheme like from Mathworks.

Its supply of boxes should include (transport-triggered) data paths, memory blocks, asMs, queues, stacks, (de)multiplexers, interconnect resources, etc.

Users should be able to select and interconnect such "boxes" during a session at the graphical user interface.

The user determines the box dimension (path width, pipe length, etc.)

© 2009, reiner@hartenstein.de

12

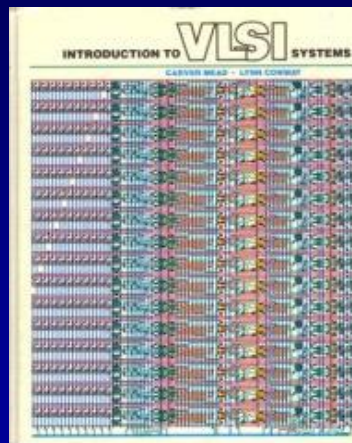
<http://hartenstein.de>



Qualified Programmer Population missing

Missing programmer population, tools and methodology:

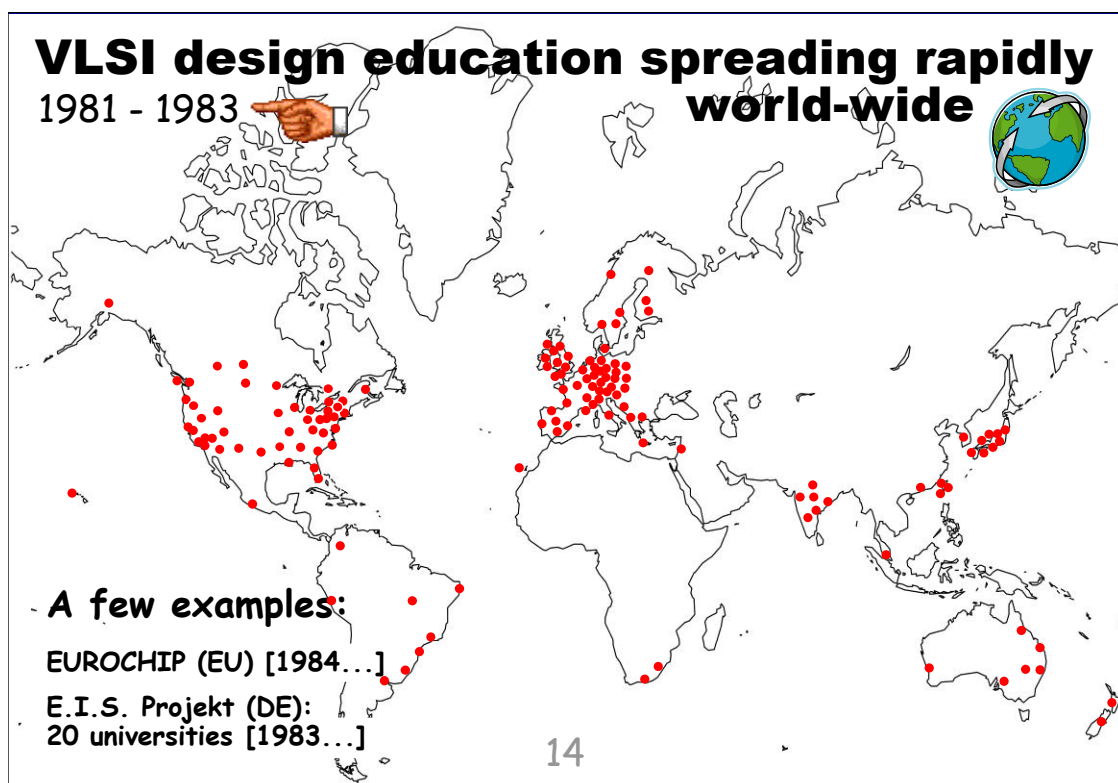
a scenario like before the Mead-&-Conway revolution



© 2009, reiner@hartenstein.de *) RC = Reconfigurable Computing

13

http://hartenstein.de





New Course Material

We need a new undergraduate text book

We need a graphic representation method and its semantics

We need a model-based undergraduate lab course running under this interface



New CS Curricula

CS curricula must represent today's computing reality by parallel, **reconfigurable**, and distributed computing.

We need to change Computer Architecture courses by introducing parallel, **reconfigurable**, and distributed computing in early undergraduate CS.

not yet here: complicated problems and in-depth analysis.

from the start, we expect students to consider the world of computing as filled with **hetero** systems.



thank you for your patience

© 2009, reiner@hartenstein.de

17

<http://hartenstein.de>



END

© 2009, reiner@hartenstein.de

18

<http://hartenstein.de>

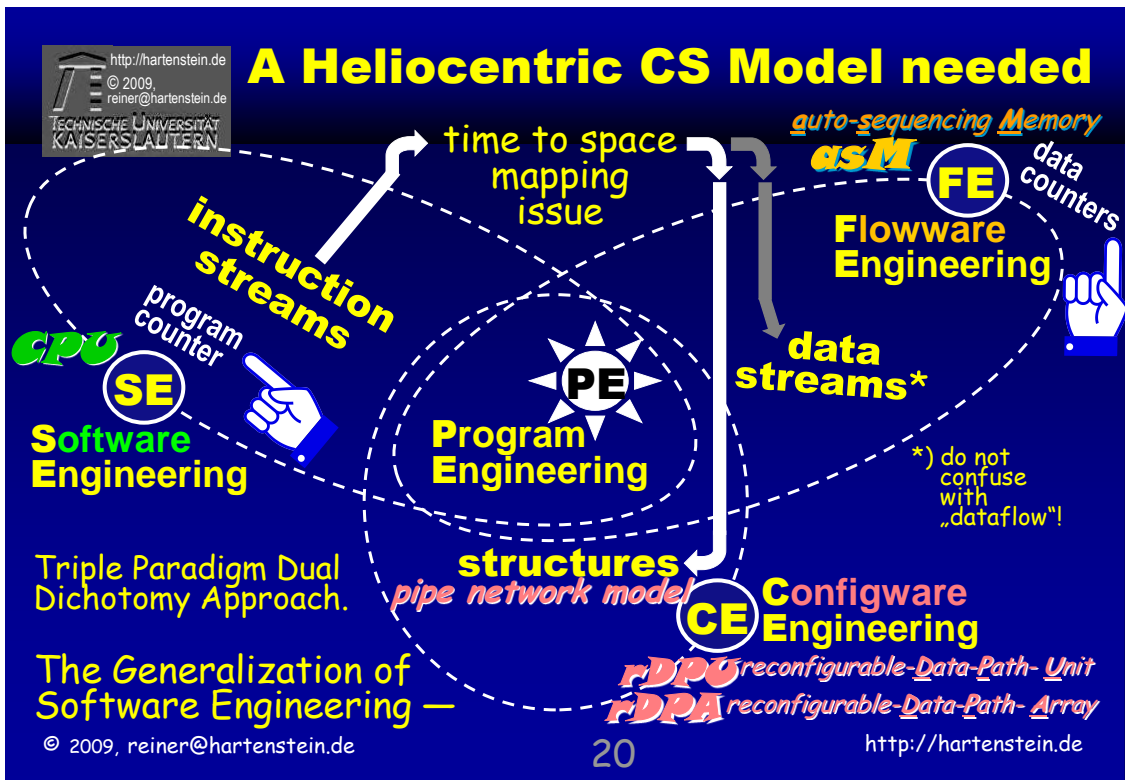


backup for discussion

© 2009, reiner@hartenstein.de

19

http://hartenstein.de



<http://hartenstein.de>
 © 2009, reiner@hartenstein.de
 TECHNISCHE UNIVERSITÄT
 KAISERSLAUTERN

POIIP: Loop turns into Pipeline

[1979]

© 2009, reiner@hartenstein.de 21 http://hartenstein.de

<http://hartenstein.de>
 © 2009, reiner@hartenstein.de
 TECHNISCHE UNIVERSITÄT
 KAISERSLAUTERN

Paradigm Dichotomy: an old hat

HDL scene ~1970: „decision box turns into demultiplexer“

time to space mapping

decision box:

demultiplexer:

“That's so simple! why did it take 30 years to find out?”

reductionists' tunnel view

PvOIIIP

© 2009, reiner@hartenstein.de 22 http://hartenstein.de



Double Dichotomy

1) Paradigm Dichotomy

von Neumann Machine
instruction stream
(Software-Domain)



Datastream Machine
data stream
(Flowware-Domain)

2) Relativity Dichotomy

time:
-Procedure
(Software-Domain)



space:
-Structure
(Configware-Domain)




What is a Dichotomy ?

Dichotomy = mutual allocation to two opposed domains such, that a third domain is excluded.

The dichotomy model as a didactic orientation guide to overcome the software/configware chasm

Dichotomy of paradigms (von Neumann /Anti machine): the „Twin Paradigm“



http://hartenstein.de
 © 2009, reiner@hartenstein.de
 TECHNISCHE UNIVERSITÄT
 KAISERSLAUTERN

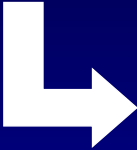
time-iterative to space-iterative

Often the space dimension is limited

n time steps, 1 CPU

a time to space mapping

1 time step, n DPUs



$n*k$ time steps, 1 CPU

a time to space/time mapping

n time steps, k DPUs

loop transformation methodology:
 70ies and later [survey: Diss. Karin Schmidt, 1994, Shaker Verlag]

e. g. example: bubble sort migration


Strip mining
 [D. Loveman, J-ACM, 1977]

← POTIP

© 2009, reiner@hartenstein.de

25

<http://hartenstein.de>



http://hartenstein.de
 © 2009, reiner@hartenstein.de
 TECHNISCHE UNIVERSITÄT
 KAISERSLAUTERN

Relativity Dichotomy

time domain:
 procedure domain
 (Machine Dichotomy)

2 phaseses:

- 1) programming instruction streams
- 2) run time

von Neumann Machine

space domain:
 structure domain

time ← *space*

3 phases:


- 1) reconfiguration of structures
- 2) programming data streams
- 3) run time

Anti Machine

© 2009, reiner@hartenstein.de

26

<http://hartenstein.de>



http://hartenstein.de
 © 2009, reiner@hartenstein.de
 TECHNISCHE UNIVERSITÄT
 KAISERSLAUTERN

time to space mapping


time domain:
procedure domain

time algorithm → space algorithm

program loop
n time steps, 1 CPU

complexity: $O(n^2)$

Bubble Sort
n × k time steps,
1 „conditional swap“ unit

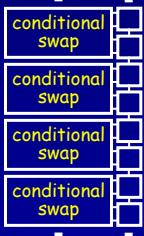


time algorithm → space/time algorithm s

space domain:
structure domain

complexity:
 $O(n)$ only

pipeline
1 time step, n DPUs




Shuffle Sort
k time steps,
n „conditional swap“ units

© 2009, reiner@hartenstein.de

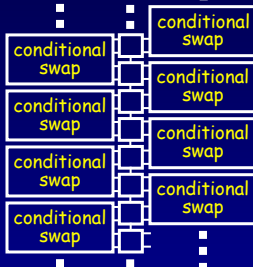
27

<http://hartenstein.de>



http://hartenstein.de
 © 2009, reiner@hartenstein.de
 TECHNISCHE UNIVERSITÄT
 KAISERSLAUTERN

Architecture instead of synchro



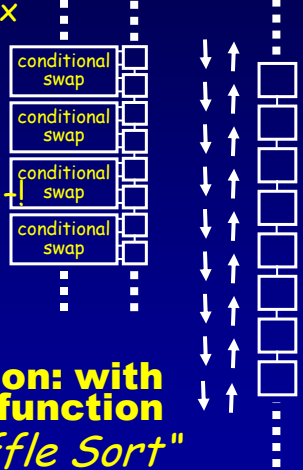
direct time to space mapping
accessing conflicts

Better Architecture instead of complex synchronisation:
 half the number of Blocks
 + up / down of data (shuffle function)

complexity: $O(n)$ only

modification: with shuffle-function
„Shuffle Sort“

Example



© 2009, reiner@hartenstein.de

28

<http://hartenstein.de>



CUDA: a programming language

Simple, elegant, code easy to maintain, mature compilers

The CUDA programming model forces the developer to identify the irreducible level of parallelism in his problem.

seems to be a better, more intuitive and extensible way to think about our problems.

the CUDA code was cleaner and more readable language for expressing parallelism

reducing the programming model from three levels to one level for a simpler more elegant solution



ISA Extensions

create instructions tailored to applications.

merge many simple operations into mega ops

automatic ISA extensions can be effective

manually created ISA extensions gives larger gains:
Tensilica reports speedups of 40x to 300x for FFT, AES and DES encryption

avoid store / communicate intermediate results



Too many HDLs

Table 1. Surveyed HLLs

| | | | | | |
|----|------------------|----|-------------------|----|-----------|
| 1 | Impulse-C | 14 | System-Studio | 27 | CoreFire |
| 2 | Handel-C | 15 | ConvergenSC | 28 | DSS |
| 3 | Mitriion-C | 16 | Transmogriifier-C | 29 | Forge |
| 4 | Dime-C | 17 | SPARK | 30 | CASH |
| 5 | System-C | 18 | Brass | 31 | C2Verilog |
| 6 | Catapult-C | 19 | DeFacto | 32 | Bach C |
| 7 | Carte-C | 20 | MATCH | 33 | SpecC |
| 8 | Streams-C | 21 | JHDL | 34 | Ocapi |
| 9 | AccelChip | 22 | Galadriel & Nanya | 35 | HardwareC |
| 10 | NAPA-C | 23 | Viva | 36 | Cones |
| 11 | SA-C | 24 | Ptolemy II | 37 | BDL |
| 12 | Trident Compiler | 25 | SysGen | 38 | Cocentric |
| 13 | CHIMPS | 26 | RC Toolbox | 39 | C2H |

© 2009, reiner@harte

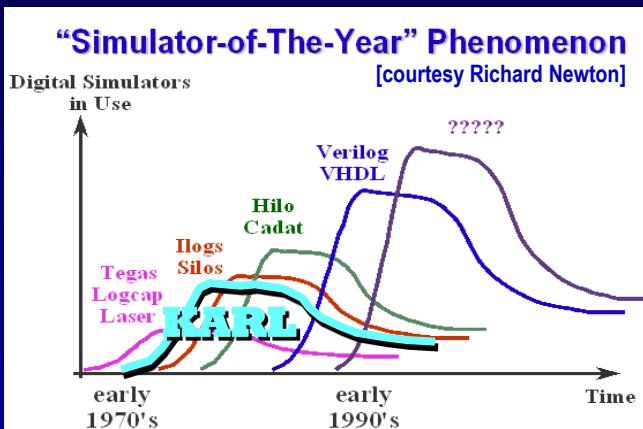
E. El-Araby et al.: Comparative Analysis of High Level Programming for Reconfigurable Comp Methodology And Empirical Study; Proc. SPL2007, Mar del Plata, Argentina, Febr. 2007



EDA the main bottleneck

The different HLL paradigms/approaches:

- imperative programming (Impulse-C)
- functional programming (Mitriion-C)
- schematic/graphical programming (DSPLogic)



© 2009, reiner@hartenstein.de

32

<http://hartenstein.de>