An architecture and protocol for the management of **concurrency on DSoCs** Overview - Summer 2011



## Agenda

apple

SVP & Apple-CORE
Realizations & evolution
Key concepts (selected)
Current "hot" topics
Research directions



# SVP & Apple-CORE

# Origins



Igether = "Microgrid"
Igether = "Microgrid"

NWO Microgrids (2005-present): <u>Architecture concepts</u>, µTC and simulations

EU ÆTHER (2007-2009): Reconfigurability and resource abstractions



# The original SVP in a nutshell

- Onit of concurrency: <u>families</u> of threads
- The <u>cluster is the processor</u>: entire families created at once over multiple cores
- Both internal and external parallelism available in hw (ILP, SMT and multi-core parallelism)
- The <u>"create" primitive</u> can be used in programs to define families, "<u>sync</u>" to bulk synchronize
- Split-phase asynchrony assumed to be the norm for all operations

apple Ecore

Use requires <u>new programming model</u> (new language, new compilation methods)

#### SVP functional

concurrency



#### SVP replicated concurrency

Concurrency



#### SVP Pipelined concurrency

#### Concurrency & dependencies



## Parent-child relationship in heterogeneous tasks



## Moving up



Š

وَ core

EU Apple-CORE (2008-2011):

- Can the Microgrid be programmed efficiently?"; sets μTC as an "intermediate language"
- Goal: automatically parallelize existing code (C, SAC) while translating to  $\mu$ TC
- Goal: Show applicability to real-world applications

Compute kernels (performance languages)



Compute kernels (performance languages)



Compute kernels (performance languages)

> Compilers, OS and runtime support (for retargeting and resource mapping)

> > Many cores

apple

sore 💐

Compute kernels (performance languages)

> Compilers, OS and runtime support (for retargeting and resource mapping)

> > Many cores

apple

Score

Compute kernels (performance languages)

Compute kernels (retargetable concurrent SVP IR)

> Compilers, OS and runtime support (for retargeting and resource mapping)

> > Many cores

apple

Ň

a core

Compute kernels (performance languages)

Compute kernels (retargetable concurrent SVP IR)

> Compilers, OS and runtime support (for retargeting and resource mapping)

> > Many cores

apple

Ň

s core

Compute kernels (performance languages)

Compute kernels (retargetable concurrent SVP IR)

> Compilers, OS and runtime support (for retargeting and resource mapping)

> > Many cores

apple

Ň

s core

Compute kernels

(performance languages)

parallelizing compilers

Compute kernels

(retargetable concurrent SVP IR)

Compilers, OS and runtime support (for retargeting and resource mapping)

Many cores

apple

Ň

core

Compute kernels

(performance languages)

parallelizing compilers

Compute kernels

(retargetable concurrent SVP IR)

Compilers, OS and runtime support (for retargeting and resource mapping)

Many cores

apple

Ň

core

Ň





















# Duality in SVP (original view)

SVP is dual, a combination of a programming model (provides concepts and contracts) and an <u>architecture model</u> (provides abstract mechanisms)"

apple

The Microgrid and µTC-ptl (higher-level simulation) are implementations of SVP

# Moving up (cont.)

Cleaning up (2008-2009): Programming and architecture models stabilized, described and specified separately

Beefing up (2009-2010):

Compiler toolchains actually useable

 OS+compiler support for resource deadlocks, mutual exclusion, arbitrary inter-thread communication

# Apple-CORE foreground



#### Apple-CORE results

- Concurrency within cores provide good latency tolerance assuming enough threads
- Distribution of work across cores is technically easy and can be achieved at extremely <u>low</u> <u>latencies</u> (few cycles)
  - the difficult problem is placement decision
- We have designed a full <u>chip architecture</u>
   = custom in-order microthreaded RISC cores on a custom NoC with custom memory network

apple

ŝ

کے core



#### JUST BECAUSE YOU ARE UNIQUE DOES NOT MEAN YOU ARE USEFUL

#### Pause and reflection



- There is more to a usable system than a processor (or network thereof) and code generator
- Embedding dataflow + microthreading into a core is a disruptive architectural change

to easy to under-estimate compilation issues
 (µtc was a management failure, had to re-design mid-project)

The interface between hardware and software is not only a matter of language, OS syscalls and standard user library

too easy to under-estimate integration issues
 (an SVP-only system prevented reuse of benchmark code)

#### "Today"

Multiple Coarse time sharing Single sequential software scheduler processor applications Shift Extend responsibility languages Hardware Multi-Many granularity concurrency & cores concurrent code scheduling

Future architectures

apple **s** core

Ň×

"Tomorrow"

#### Today:

Multi-tasked applications Time and space sharing scheduler<u>s</u>

Multiple cores

Language interoperability Standard system services

Memory hierarchies

Portability, compatibility

Device drivers

I/O devices

Expectations

Standard practices

apple apple



# Evidence from the A-C experience

Component	Objects	Functions	Code words
Libc-dtoa-strtod	44	74	14528
Libc-math	165	238	13424
Libc-stdout	13	13	5040
Libc-dlmalloc	2	19	3600
Libc-string	26	26	1968
Libc-misc	9	28	1728
mginit	2	2	1120
Libmg-benchmarks	1	7	1072
Libmg-SEP	1	4	848
Libc-malloc-mg	2	9	832
Libsl-misc	3	12	736



Linguae francae

C/POSIX (not your textbook C: inline assembly, compiler extensions)
Java/JVM
CIL/CLR (.NET)
Everything else either <u>uses</u> or <u>implements</u> these

Pick your lingua franca = pick your <u>community</u>

# How to pitch novel architectures?

(Assuming the hardware is ready for use)

- Implement a compiler for a restricted dialect of C with minimal library support
  - tractable, but expect: "meh... what can you do with it?"
  - the bane of architecture research: underestimate sw integration
- Implement a complete GCC target then port an existing operating system with bootloader, hw drivers, file system, shell and utilities
  - <u>untractable</u>, unless new hw  $\approx$  old hw (to maximize software reuse)

apple 📲 s core

What else?

# Heterogeneous systems



- Run the linguae
   francae on an
   existing design
- Embed the existing design as component of a D-SoC
- Define standard
   protocols for control and communication
   between cores

ŝ

€ core

#### Evolution / transition

 Shift from core architecture design to system design and integration

- Innovate with protocols and language integration instead of machine primitives and programming models
   (hint: tcp/ip too heavy for low-latency on-chip interactions)
- New pitch:

"SVP is a protocol and architecture for the management of execution units on a chip"

ŝ

apple for score

#### Proposed base semantics

- Dataflow synchronization cells to support split-phase asynchrony between threads (int. concurrency) and components (ext. concurrency)
- NoC <u>messages</u> between and within cores with fixed semantics:
  - allocate/release to manage execution resources
  - create to trigger execution of a family
  - sync to bulk-synchronize a family
  - o rput/rget to access synchronizers remotely
- Multiple ISA integration strategies: <u>adaptation</u>, <u>extension</u>, <u>embedding</u>
- Responsibility for consistent use pushed up to the software stack, not constrained by model any more

apple

Ŵ

a core



# Key concepts (selected)

#### Granularity

- Code portability = distributing different computation grain sizes
- execution efficiency proportional to grain size, inversely to cost of managing concurrency
- Also have to consider the communication cost of distribution



#### Granularity

- Code portability = distributing different computation grain sizes
- execution efficiency proportional to grain size, inversely to cost of managing concurrency
- Also have to consider the communication cost of distribution



#### Granularity

- Code portability = distributing different computation grain sizes
- execution efficiency proportional to grain size, inversely to cost of managing concurrency
- Also have to consider the communication cost of distribution

This is easier than targeting a given granularity



#### Mutual exclusion

- Mutual exclusion is required to support controlled non-determinism
  - DSoC memory is asynchronous, only weakly consistent might not be usable to implement <u>atomics</u>
  - the SVP way is to define a "mutex place" that serialises all requests to create tasks at that place (Dijkstra's secretaries) – same protocol
- These can be used for multiple synchronization goals,
   e.g. I/O, semaphores, resource allocation, etc.

ŝ

apple apple

As such they are in the domain of systems programming i.e. the <u>concurrency engineer</u>



# Current "hot" topics

# Inter-thread communication



Communication channels beween threads:

OK for register-sized channels: use istructure hardware registers

Not OK for larger items: they must fit through memory ... which is asynchronous

Which memory model? How to use it in programs? Non-shared address spaces?

apple

a core

# Inter-thread communication



Š

کے core

Some directions:

Programming model: make all communication visible to the compiler

 Implementation using shared memory: synchronize memory updates via registers

Ø Distributed systems: explicit messaging

# Resource deadlocks



ŝ

کے core

Shortage of memory: not our problem (yet)

- Shortage of synchronizers: spread threads on multiple cores OK
- Shortage of thread entries: create less threads per core OK
- Shortage of family entries: <u>delegate</u> ... where to? or <u>sequentialize</u> ... how?

# Resource deadlocks



ŝ

core

Some directions:

- Rebalance concurrency trees early (by rewriting code for family creates)
- Delegation (to other places on the SoC)
- Automatic sequentialization, requires both:
   <u>feedback from hardware</u> about resource availability
  - <u>sequentializability</u> of concurrency pattterns in languages



#### Research directions

# Follow-ups to Apple-CORE

How to integrate the protocol in existing and future core architectures / ISAs

- in a way that minimize disruption in software stacks?

Is there a case for multiple ISAs in the same chip

- and how to manage multiple codes in the same program?

How to automate placement over multiple cores

- with a placement latency comparable to (or lower than) the operation latency?

apple Ecore

# Software engineering ADVANCE

What information should be available in programs, and how to represent resources, to allow efficient mapping and scheduling to computing resources?"

![](_page_56_Picture_2.jpeg)

# Software engineering ADVANCE

Ø Program:

Capture extra-functional requirements in programs

Virtualize hardware and modelize resources in <u>heterogeneous</u> environments
 Design <u>mapping</u> and <u>feedback</u> mechanisms

apple Ecore

![](_page_57_Picture_4.jpeg)

# ADVANCE at a glance

![](_page_58_Figure_1.jpeg)

Score 💐

![](_page_58_Picture_2.jpeg)

#### ADVANCE interactions

![](_page_59_Figure_1.jpeg)

# Software engineering issues (ADVANCE)

- Partner technology: S-NET (coordination via SISO streams and functional boxes)
- The position of SVP (WIP):
  - Simplify the analysis of <u>requirements</u> of a S-NET network implemented using our protocol
  - Simplify <u>mapping</u> by imposing <u>constraints on</u> <u>scheduling and communication patterns</u>
  - Establish <u>T/L/J models for both internal and</u> <u>external concurrency</u> on heterogeneous cores

#### Now what? Further discussion occurs now!

![](_page_61_Picture_1.jpeg)