

10 Sonnets Concerning FPT

Philip Leong

phwl@cse.cuhk.edu.hk

The Chinese University of Hong Kong






1. Introduction

Reconfigurable computing is known,
Otherwise by the two letters “RC”.
Enormous potential it will be shown,
Acceleration to a large degree.

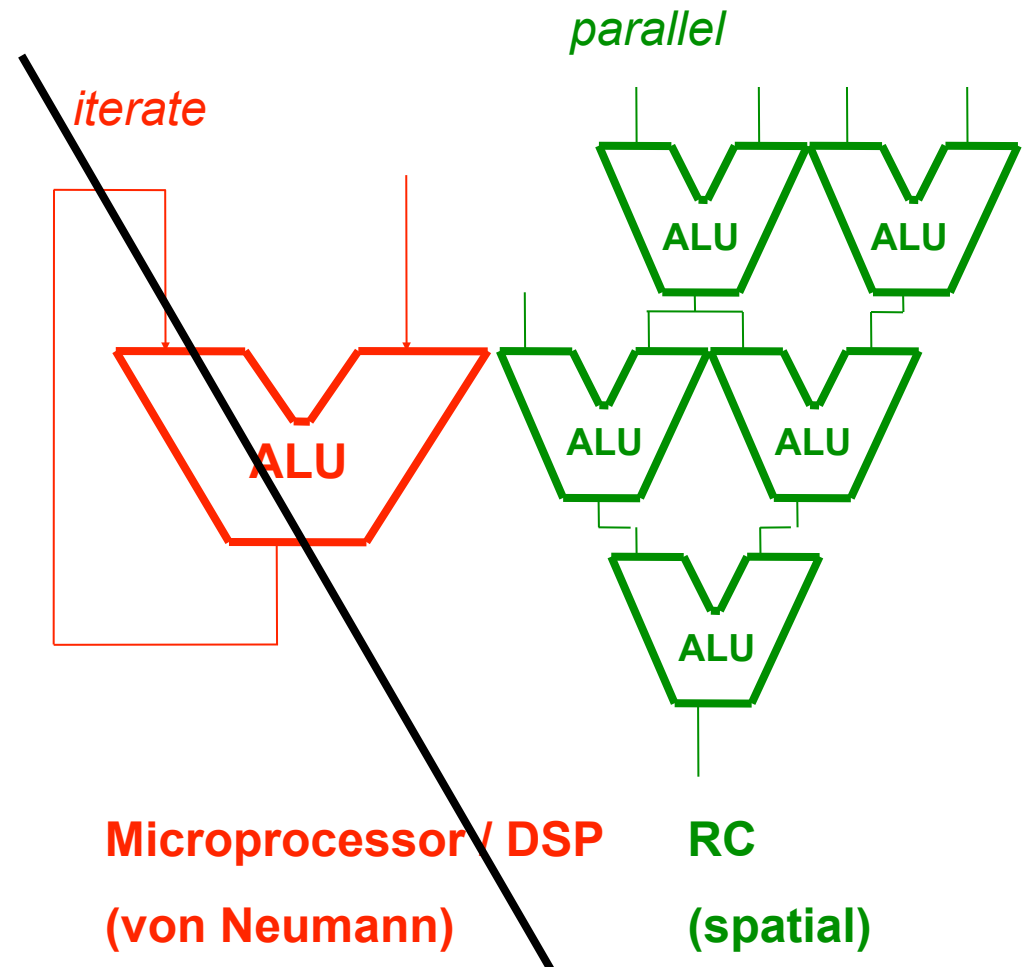
Parallelism is the basic trick
We apply to solve almost everything,
In spatial form we arrange our logic
Indeed to von Neumann we must not cling.

This talk, as presented in sonnet form
Will attempt to highlight developments;
We first describe what has become the norm
Then we'll present some new embodiments.

- 
- Introduction
 - Floating point FPGAs
 - Latency
 - Memory
 - Virtual logic
 - Applications
 - Conclusion

Introduction

- Spatial computing
- FP applications can greatly broaden scope
 - Architecture?
 - Other developments needed?
 - Potential areas for research





2. Fixed vs Floating

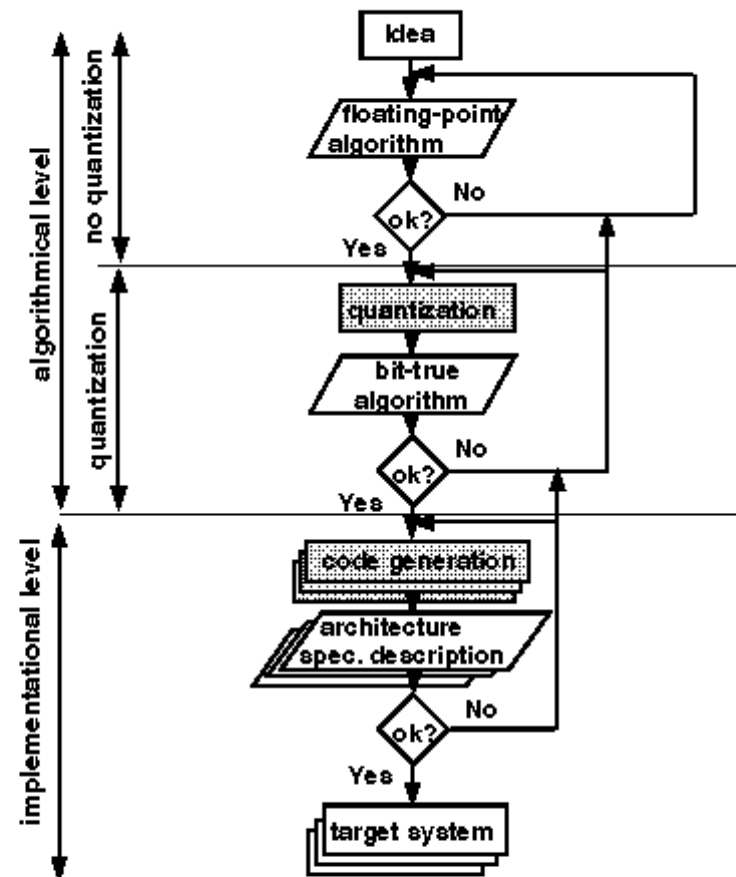
RC inevitably in the past,
Was done using fixed point arithmetic;
Optimised to be both small and quite fast,
To designers we are sympathetic.

For fixed point design, a difficult art,
Requires much patience, effort and time;
Quantization errors right from the start,
Overflow, underflow, we must decline.

IEEE 754 indeed,
Will address this most troublesome issue,
But embedded FPU cores will need,
Much extra memory bandwidth to accrue.

- Complex
- Difficult to verify
- Time consuming
- FP doesn't have these issues

Design Flow





Commercial FPGA
Weaknesses (for FP)

Large area, low clock
frequency

Run out of resources

Floating Point FPGA

Coarse-grain,
hardwired FPUs

Dynamic
reconfiguration




3. Coarse vs Fine Grain

Fine grain, coarse grain, using both is the trick,
Control and data are not quite the same;
Hard FPUs, density like ASIC,
Employing this scheme, performance we gain.

We found that coarse-grained blocks greatly reduce,
FP logic, compare and multiplex;
Using this scheme we found we could produce,
Twenty-fold smaller area at best.

Floating-point units are better in speed,
Which is another important factor;
We see that they are effective indeed,
Multiplier, adder and subtractor.



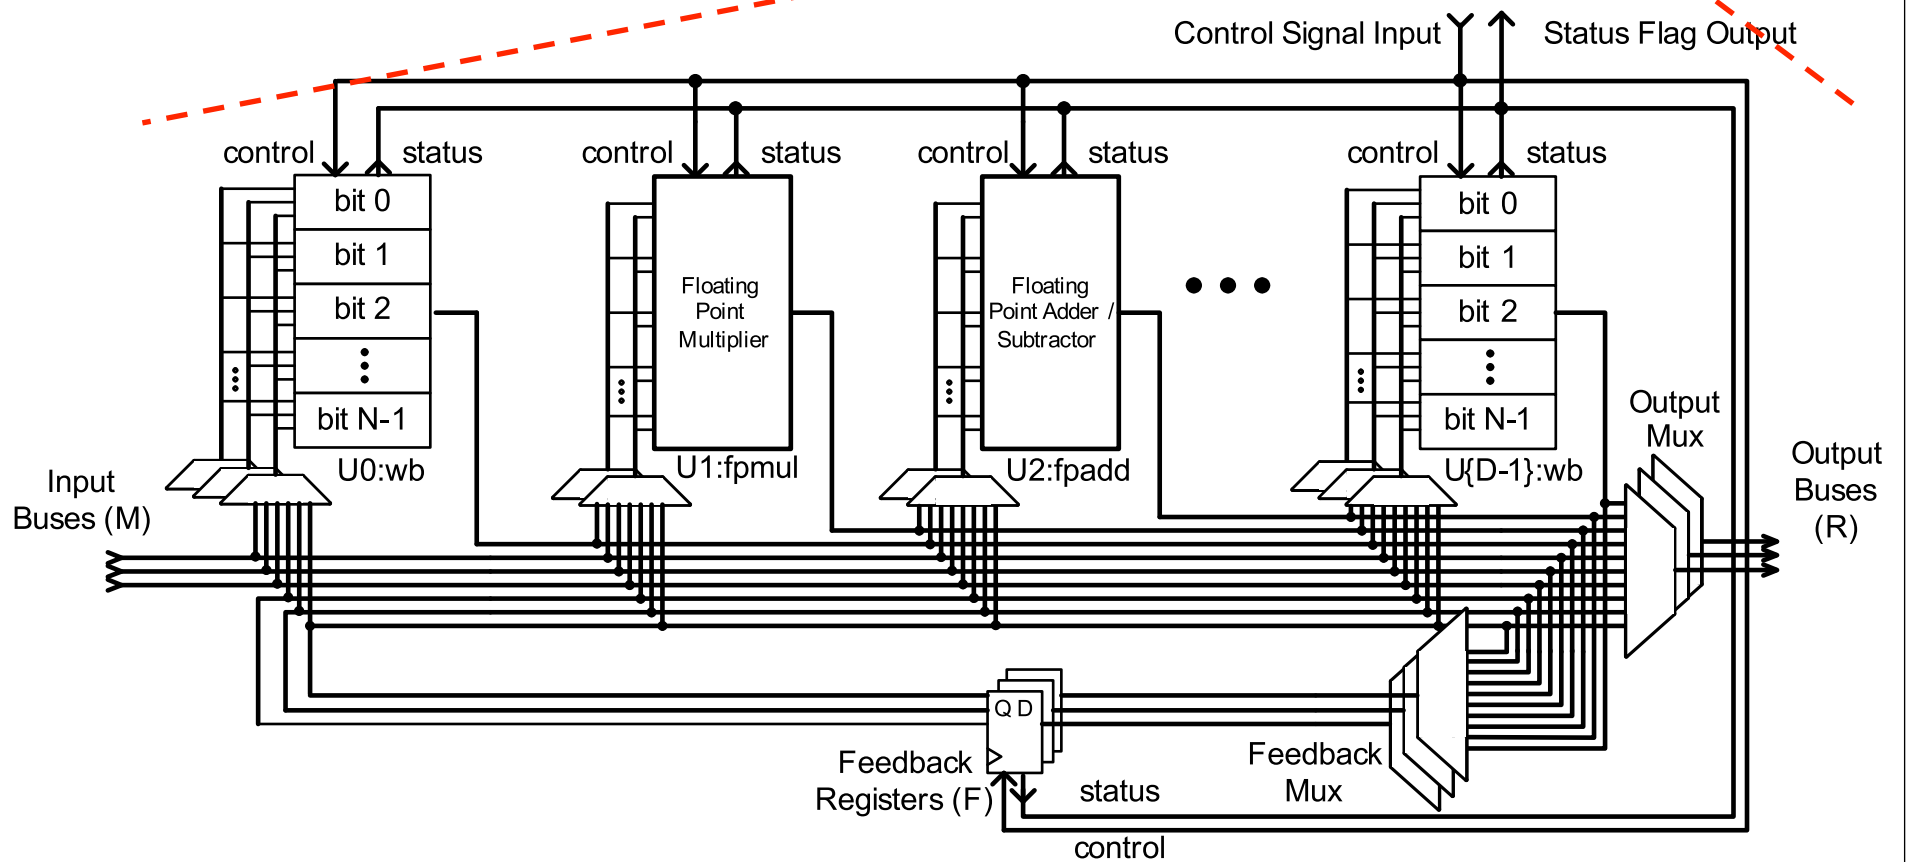



4. A Coarse-grained FP Fabric

The coarse-grained block is directional see,
Data in from the left, out on the right;
Transformed by blocks of which we have three,
We are able to customise each site.

The first is a 4-LUT for general use,
It is good for multiplex and compare,
A register too so we can produce,
Memory and latched values of output there.

In floating point blocks there are mult and add,
Single or double precision in size;
For speed and size the word block is not bad,
At compile time we parameterise.



D=9, M=4, R=3, F=3, 2 add, 2 mul: best density over benchmarks

	Floating Point hybrid FPGA		XC2V3000-6			
	Area (slices)	Delay (ns)	Area (slices)	Delay (ns)	Area (times)	Delay (times)
bfly	565	9.02	13733	24.57	24.3	2.72
dscg	661	10.11	9614	22.78	14.5	2.25
fir4	371	9.06	11290	23.68	30.4	2.61
mm3	642	8.90	8889	23.4	13.8	2.63
ode	545	9.74	8238	21.93	15.1	2.25
bgm	1810	10.00	30207	24.34	16.7	2.43

Coarse-grained vs XC2V3000
(floating point benchmarks)

Geometric
Mean

18.3 **2.48**





5. Latency

I/O is an unfortunate evil,
Without it our speedups would be higher;
We must have high bandwidth for retrieval,
Very low latency we aspire.

We hence desire a low latency bus,
Tightly coupled to the host CPU;
This is not a desire but a must,
So the data can come fast and be true.

Pilchard was an answer which used a DIMM,
In the memory bus speedups we did see;
Today some other connections are in,
Such as FPGA to FSB.

- Memory bus has high bandwidth and low latency
- Pilchard circa 2000





Acceleration technologies for high-performance computing



Front side bus RC systems from DRC, Xtremedata and Nallatech



6. Geneseo

A major issue with hardware we hate,
Is portability with host, we know;
It's great Intel has stepped up to the plate
With new technology "Geneseo."

This will be a standard which helps enhance,
RC interfaces to host machine,
Supporting coprocessors will advance,
PCI express to achieve our dream.

This will combine good portability,
With low latency, high bandwidth and more,
Without support for full coherency,
Interoperability furthermore.



7. Memory

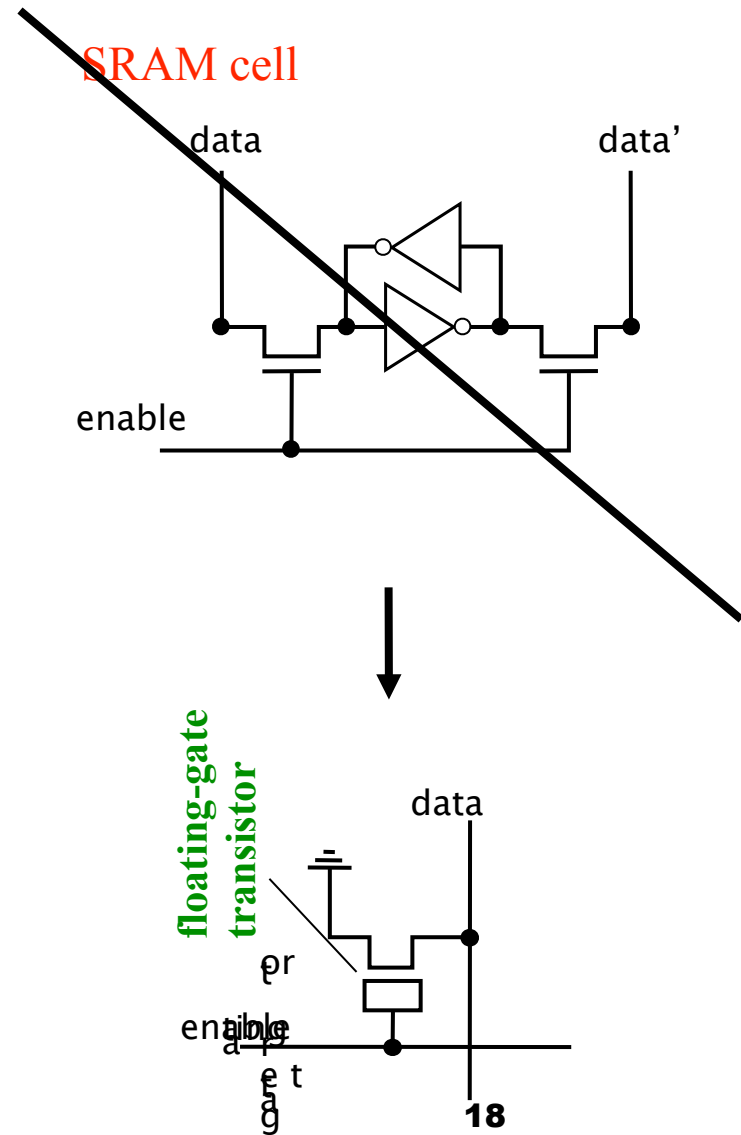
Memory is so expensive you know,
Five or six transistors to make a RAM;
Flash is a single transistor and so
A reduced area design we plan.

Unfortunately though flash very small,
Program and erase overhead is high
Currently large blocks make sense furthermore
No solution on which we can rely.

The great advantage is high density,
And has the static power problem beat,
FPGAs of such propensity;
To make most current designs obsolete.

- Use flash for FPGA storage

- Improved density
- Improved power
- Non-volatile






8. Virtual Logic

Virtual memory is used by everyone,
So why not the same scheme for RC;
To reuse logic would be lots of fun,
For coarse-grained blocks it makes more sense to me.

Hence designs would never run out of space,
With CPUs we could better compete;
Multiple contexts could do it with grace,
Reduced design time would make this complete.

Large applications we could demand page,
To designers complexity reduced;
I'm sure that this scheme would be all the rage,
Resulting in simpler designs produced.

- Obstacle to dynamic reconfiguration is in commercial architectures, we should figure out whether it is a good idea or not
- Coarse-grain has less state to reconfigure






9. Applications

We know that this technology is great,
But also need the right problems to solve,
To triumph over the micros we hate,
So what applications do we involve?

Well science, finance and oil are my bet,
All are computational in nature,
With parallel datapaths we will get,
Low power and performance much greater.

Unfortunately, we have other foes,
E.g. GPUs, multicores and Cell,
Easy to program, low power as well,
They all do run like a bat out of hell.

- 
- Oil and gas exploration
 - Financial engineering
 - Bioinformatics
 - Scientific computing
 - VLSI simulation and verification

10. Conclusion

Smaller, faster and less latency we,
All know what FPT really does need;
So include flash, coarse-grain and FSB,
To these research challenges we should heed.

For flash is small with low dissipation,
Coarse-grain gives us less routing and switches,
Tools for latency elimination,
RC applications bring the riches.

Logic devices were there from the start,
Along with some programmable routing,
Some future chips will be FP in heart,
For DSP and other computing.

