

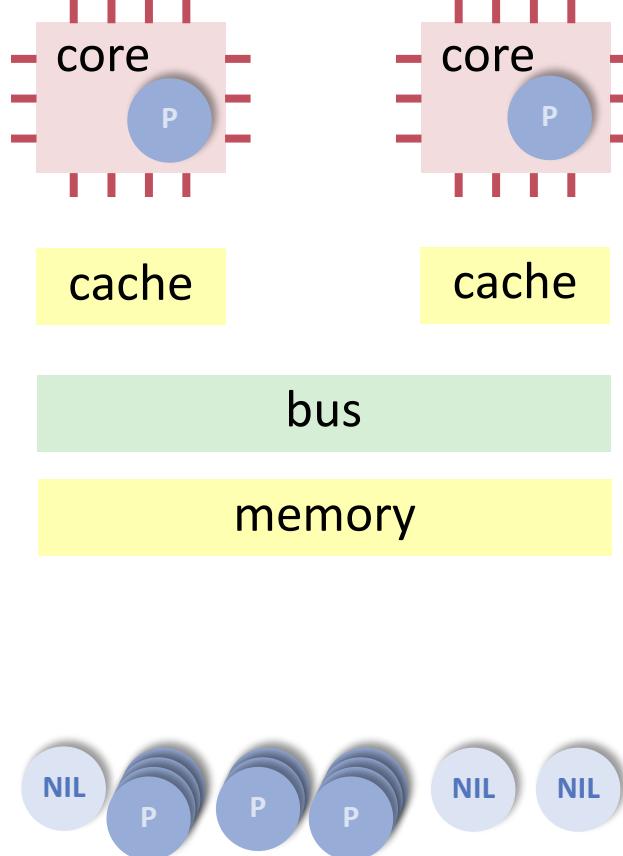
Active Cells:

A Programming Model for Configurable Multicore Systems

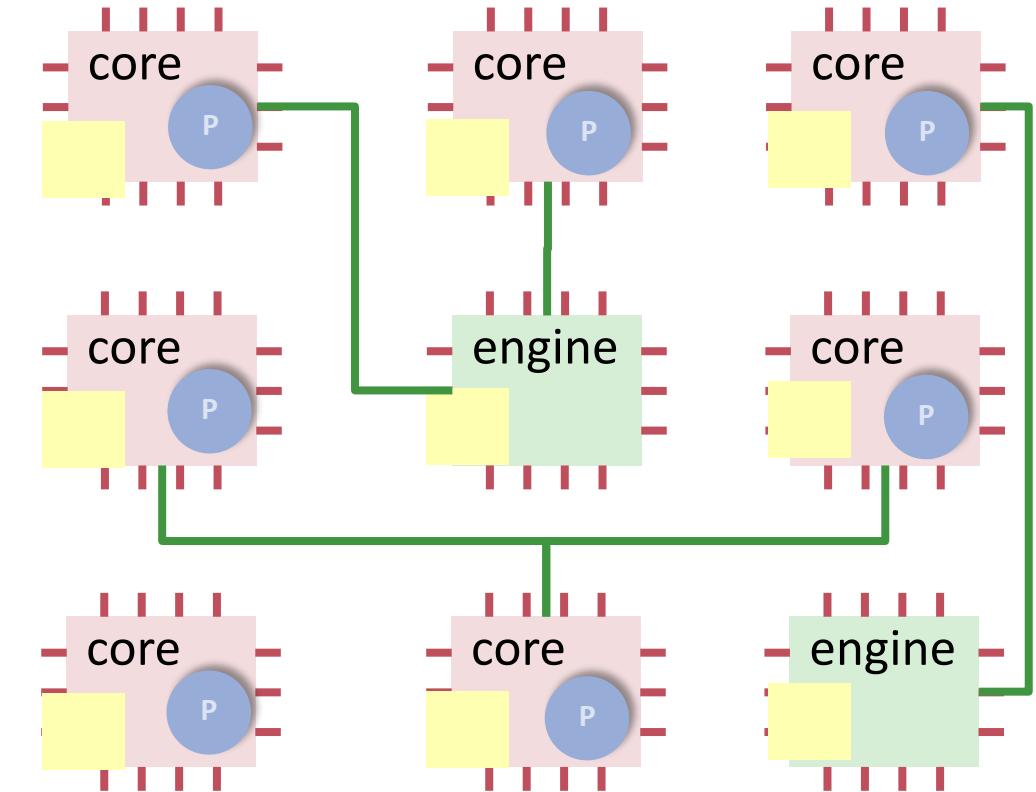
CASE STUDY 4: CUSTOM DESIGNED MULTI-PROCESSOR SYSTEM

Vision

General Purpose Shared Memory Computer



Application Specific Multicore Network On Chip



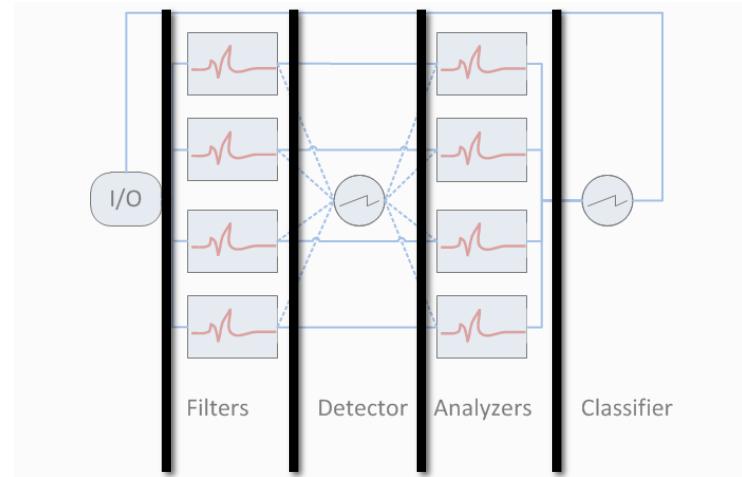
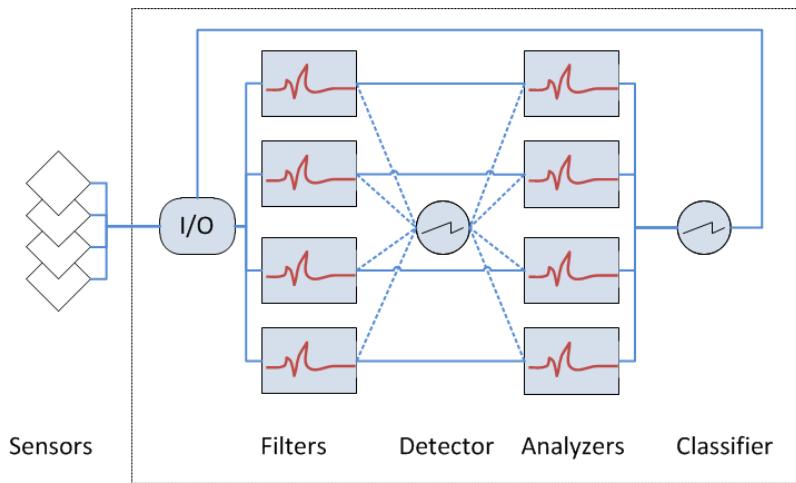
Motivation: Multicore Systems Challenges

- Cache Coherence
 - Shared Memory Communication Bottleneck
 - Thread Synchronization Overhead
-
- ⇒ Hard to predict performance of a program
- ⇒ Difficult to scale the design to massive multi-core architecture

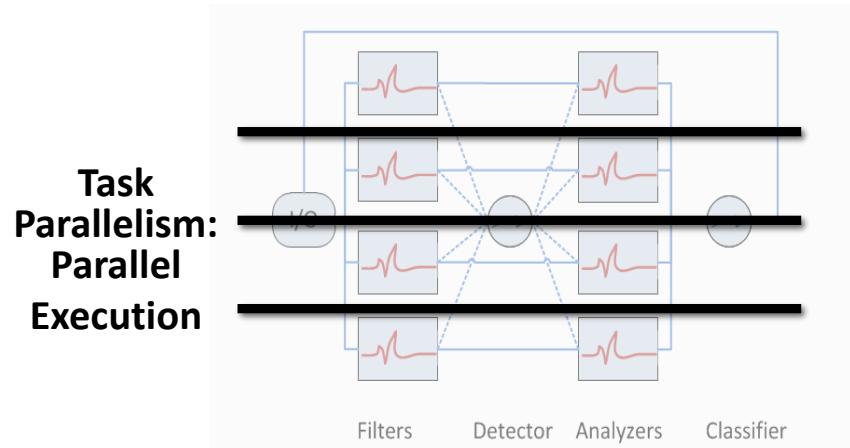
Operating System Challenges

- Processor Time Sharing
 - Interrupts
 - Context Switches
 - Thread Synchronisation
- Memory Sharing
 - Inter-process: Paging
 - Intra-process, Inter-Thread: Monitors

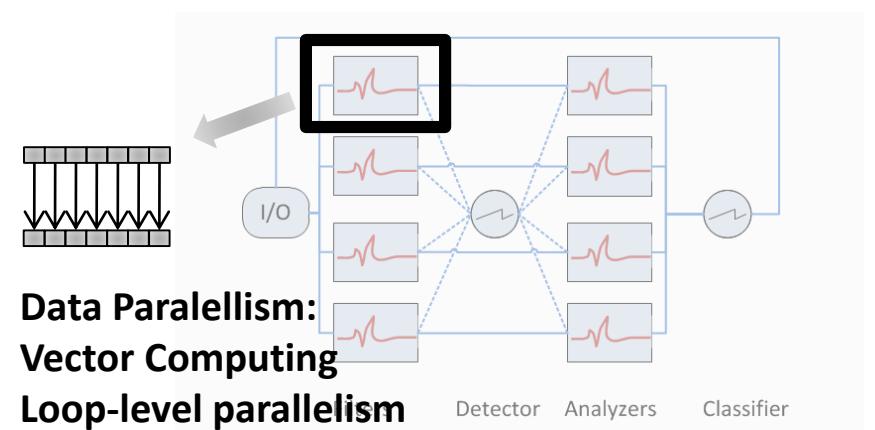
Focus: Streaming Applications



Stream-Parallelism: Pipelining



**Task
Parallelism:
Parallel
Execution**



**Data Parallelism:
Vector Computing
Loop-level parallelism**

4.1. HARDWARE BUILDING BLOCKS TRM AND INTERCONNECTS

TRM: Tiny Register Machine*

- Extremely simple processor on FPGA with Harvard architecture.
- Two-stage pipelined
- Each TRM contains
 - Arithmetic-logic unit (ALU) and a shifter.
 - 32-bit operands and results stored in a bank of 2×8 registers.
 - local data memory: $d \times 512 / 256$ (***) words of 32 bits.
 - local program memory: $i \times 1024 / 256$ instructions with 18/16 bits.
 - 7 general purpose registers
 - Register H for storing the high 32 bits of a product, and 4 conditional registers C, N, V, Z.
- No caches

* Invented and implemented by [Dr. Ling Liu](#) and Prof. Niklaus Wirth

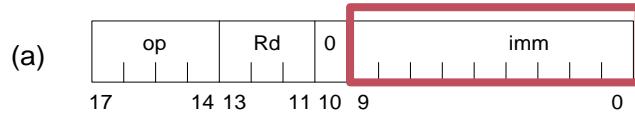
** Xilinx / Lattice Boards

TRM Machine Language

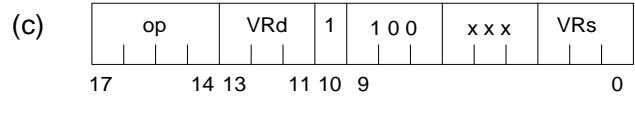
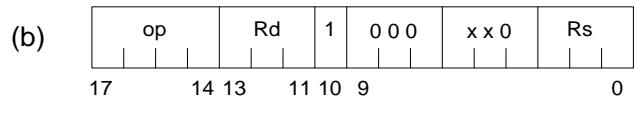
- Machine language: binary representation of instructions
- 18/16-bit instructions
- Three instruction types:
 - Type a: arithmetical and logical operations
 - Type b: load and store instructions
 - Type c: branch instructions (for jumping)

Encoding Overview

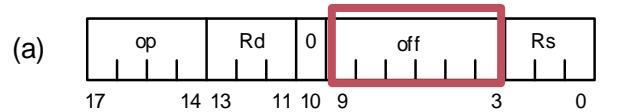
■ Register Operations



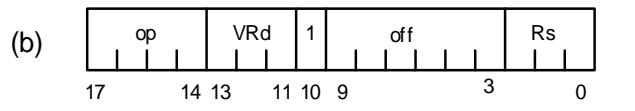
imm is zero extended to 32 bits



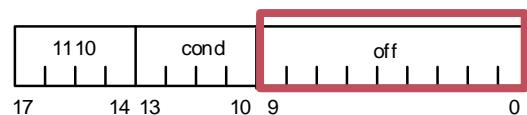
■ Load and Store



off is zero extended

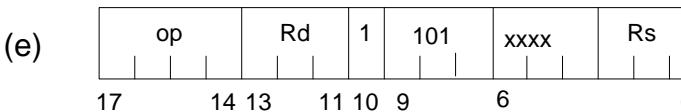
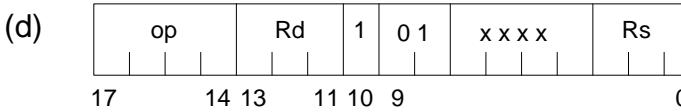
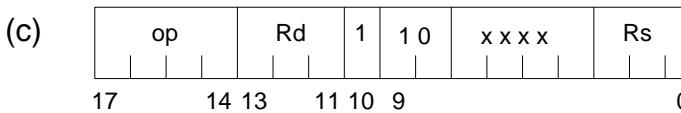
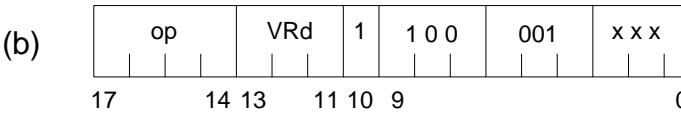
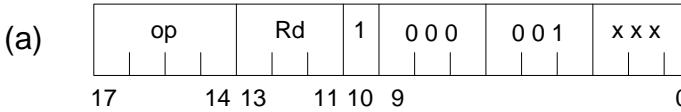


■ Conditional Branches

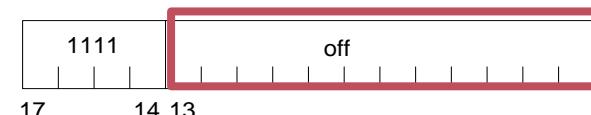


off is sign extended

■ Special Instructions



■ Branch and Link



off is 14-bit offset

TRM architecture

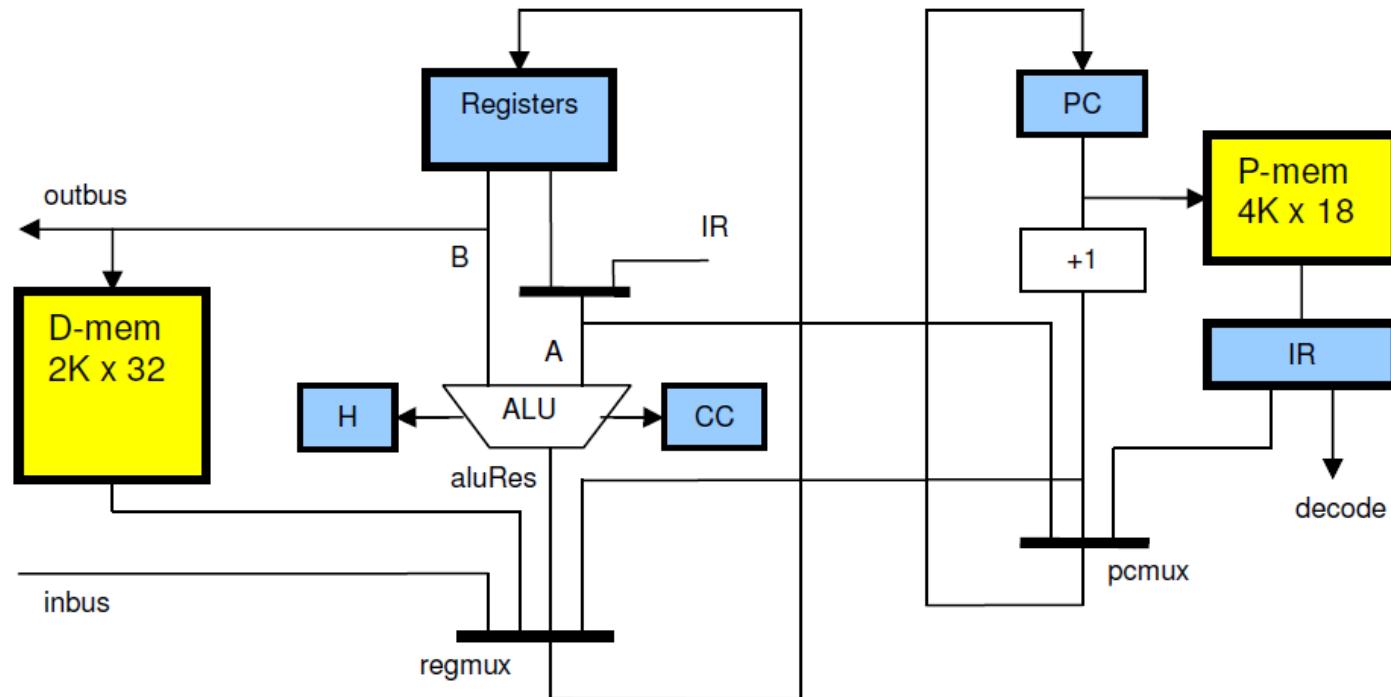


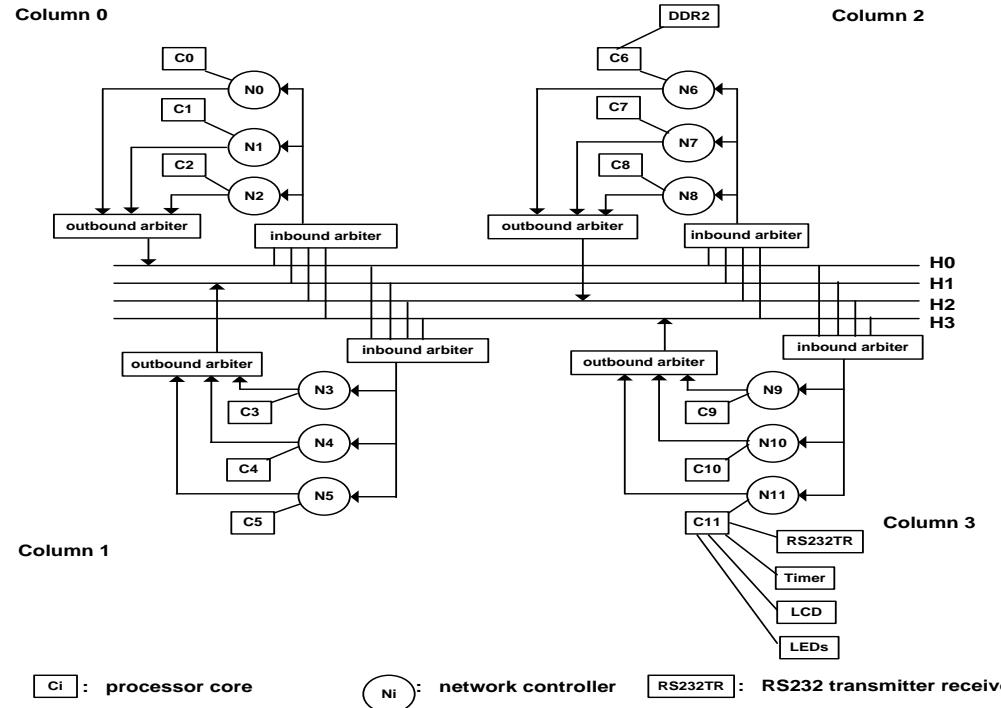
Figure from: Niklaus Wirth, *Experiments in Computer System Design*, Technical Report, August 2010
<http://www.inf.ethz.ch/personal/wirth/Articles/FPGA-relatedWork/ComputerSystemDesign.pdf>

Variants of TRM

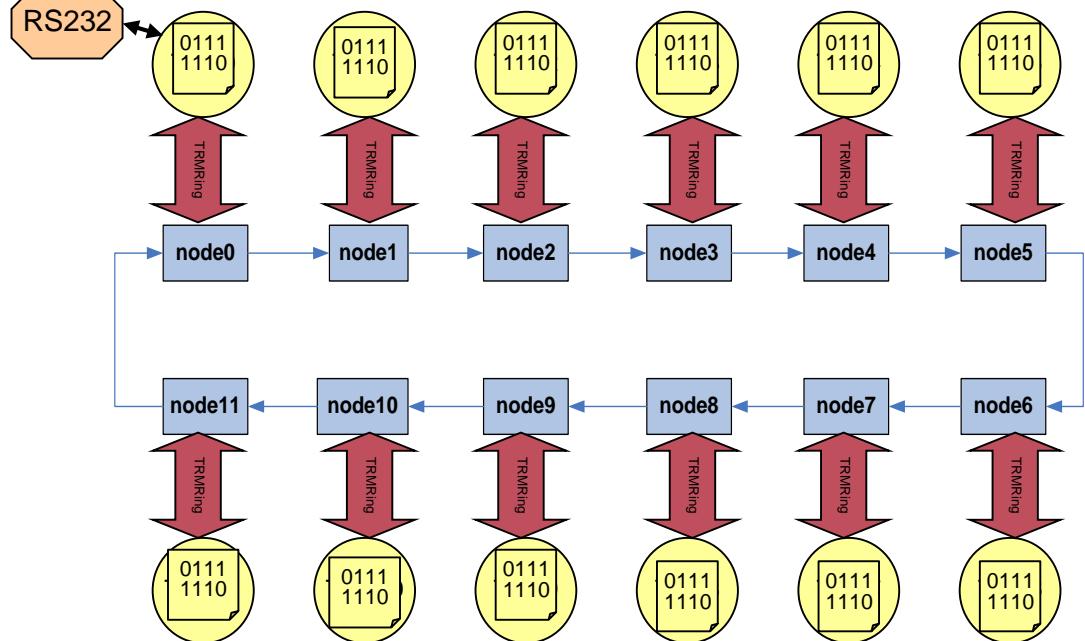
- FTRM
 - includes floating point unit
- VTRM (Master Thesis Dan Tecu)
 - includes a vector processing unit
 - supports 8 x 8-word registers
 - available with / without FP unit
- TRM with software-configurable instruction width
(Master Thesis Stefan Koster, 2015)

Initial Multicore Experiments

TRM12 Bus

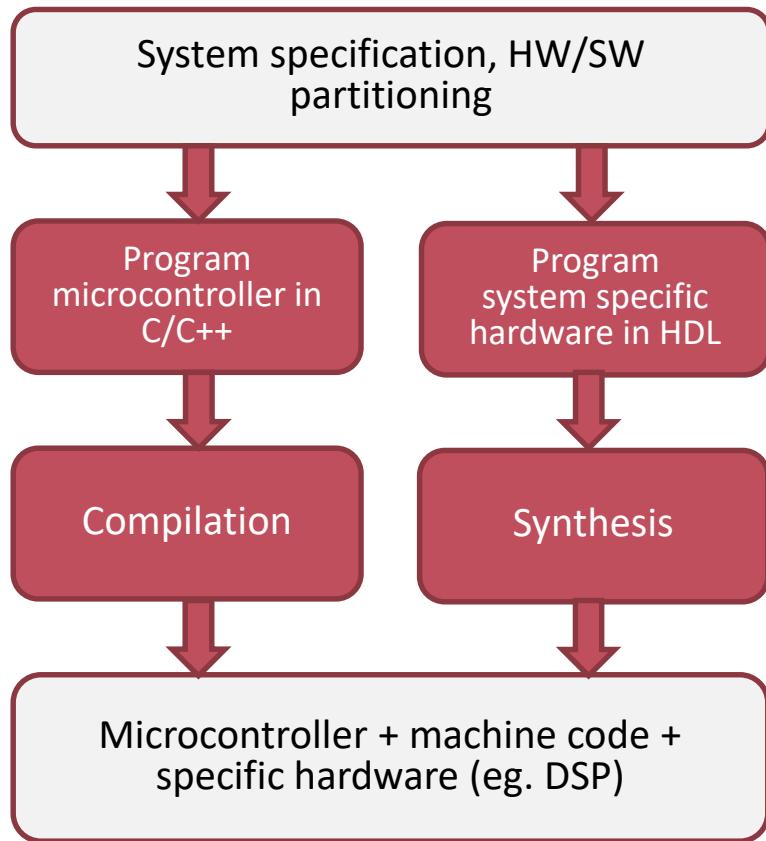


TRM12 Ring

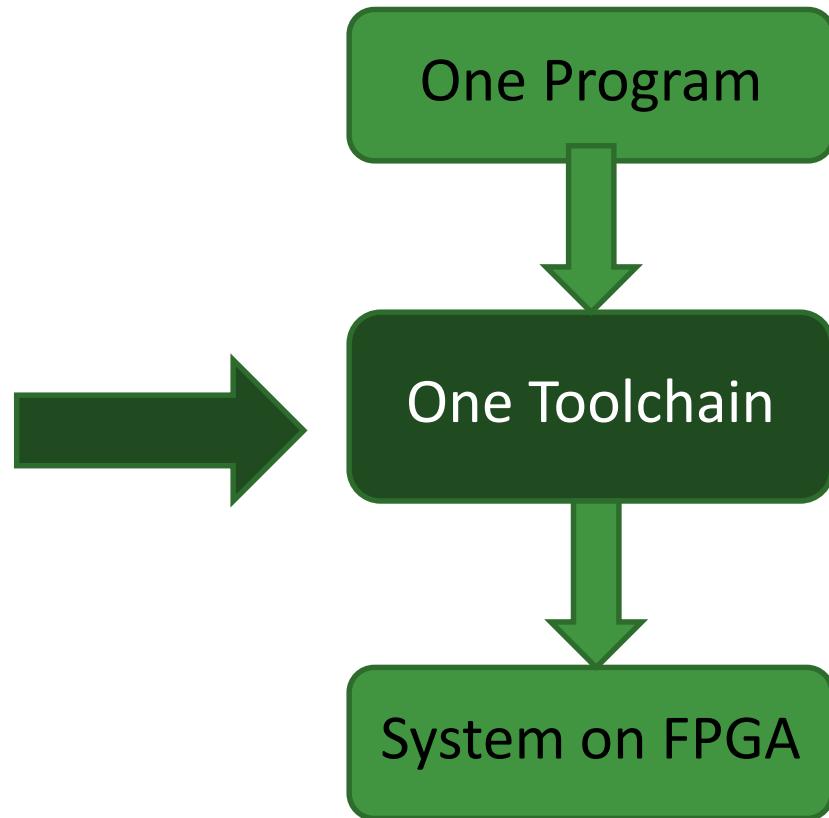


4.2. HARDWARE SOFTWARE CODESIGN

HW/SW co-design



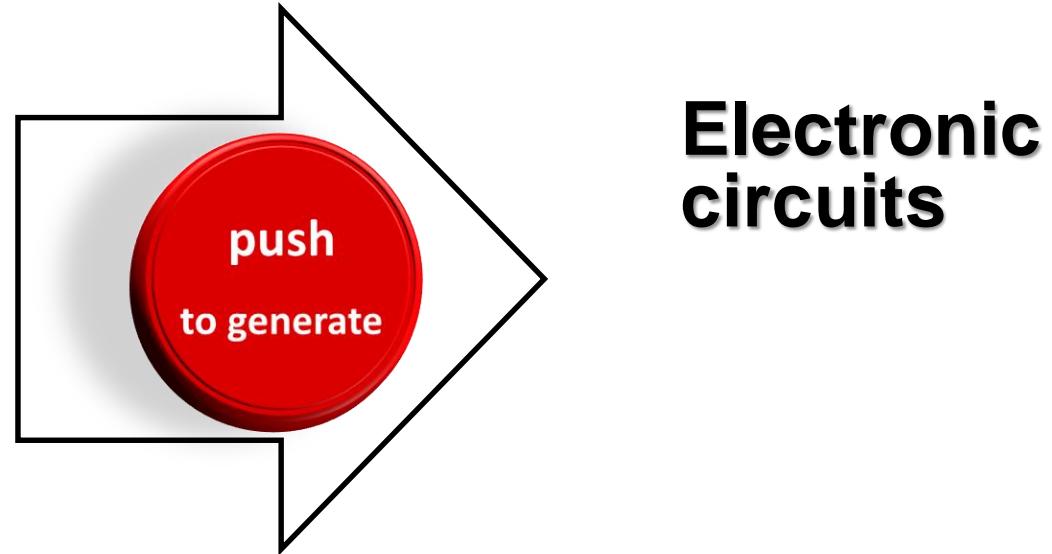
Traditional HW/SW co-design for embedded systems



Active Cells approach for embedded systems development

Vision: Custom System on Button Push

**System
design as
high-level
program
code**



Computing model
Programming
Language

Compiler,
Synthesizer,
Hardware Library,
Simulator

Programmable
Hardware
(FPGA)

Active Cells Computing Model

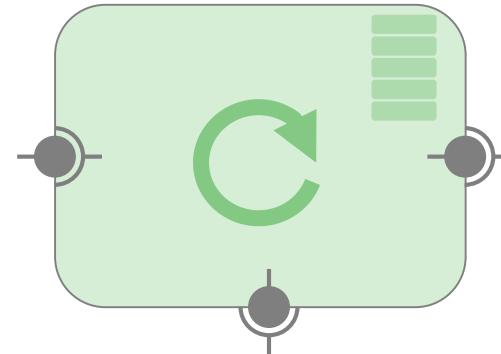
Inspired by

- Kahn Process Networks
- Dataflow Programming
- CSP (e.g. Google's Go)
- Actor Model (e.g. Erlang)

On-chip distributed system

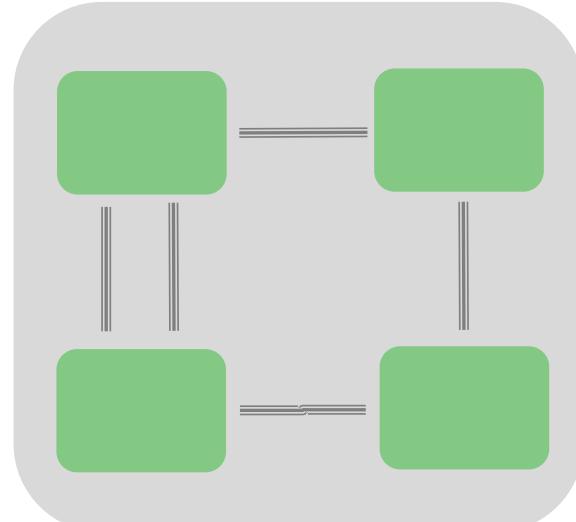
Cell

- Scope and environment for a running **isolated process**.
- Integrated **control thread(s)**
- Provides **communication ports**

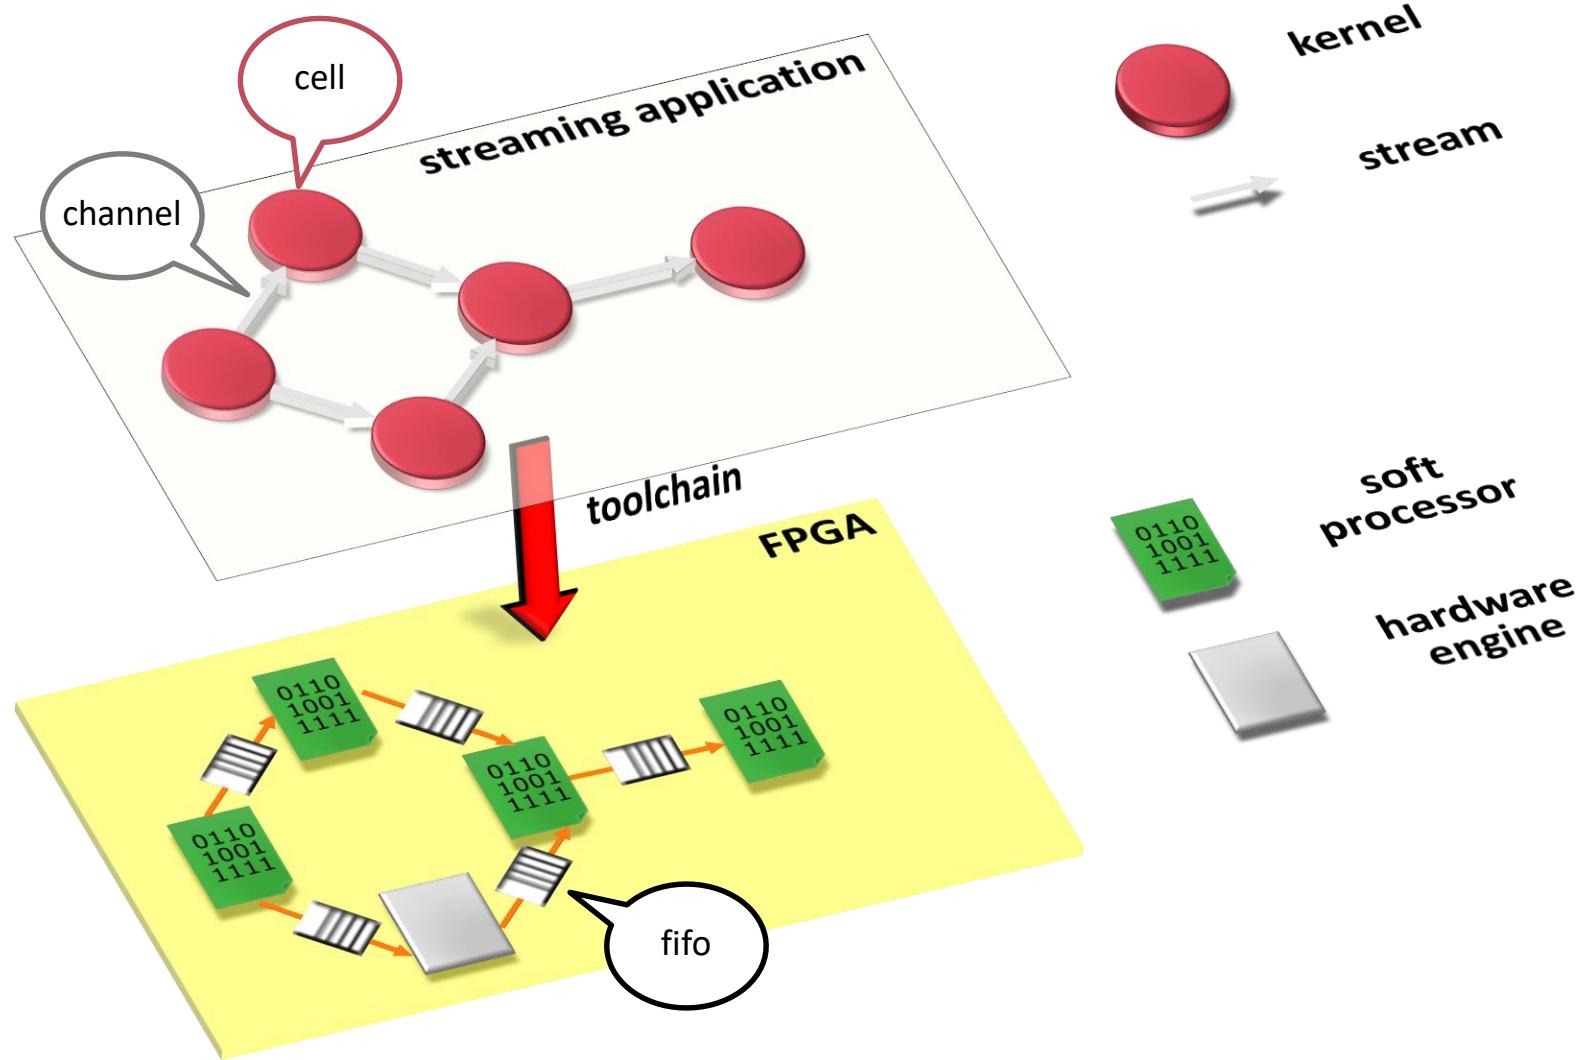


Net

- **Network** of communication cells
- Cells connected via **channels** (FIFOs)



Software → Hardware Map



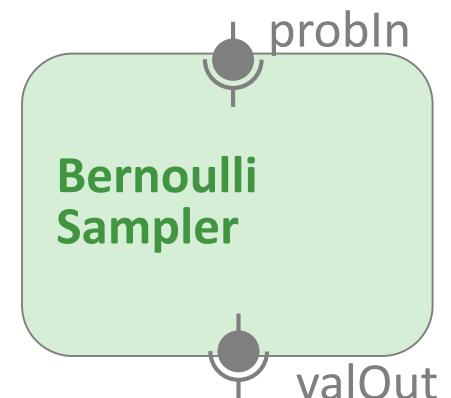
Consequences of the approach

- No global memory
- No processor sharing
- No peculiarities of specific processor
- No predefined topology (NoC)
- No interrupts

→ **No operating system**

Cell

```
non-typed communication ports  
type  
BernoulliSampler* = cell (probIn: port in; valOut: port out);  
var  
    r: Random.Generator;  
    p: real;  
begin  
    new(r);           blocking receive  
    loop  
        p << probIn;  
        valOut << r.Bernoulli(p);  
    end  
end BernoulliSampler;           asynchronous send
```



Cell Activity

Properties

Properties can influence both, the hardware synthesis and the generation of software code.

type

```
Controller = cell {Processor=TRM, FPU, DataMemory=2048, BitWidth=18}  
    (in: port in (64); result: port out);
```

...

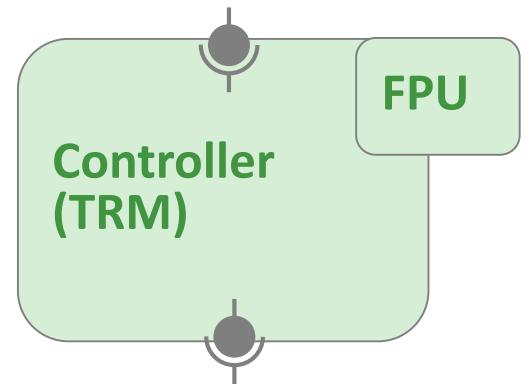
begin

```
(* ... controller action ... *)
```

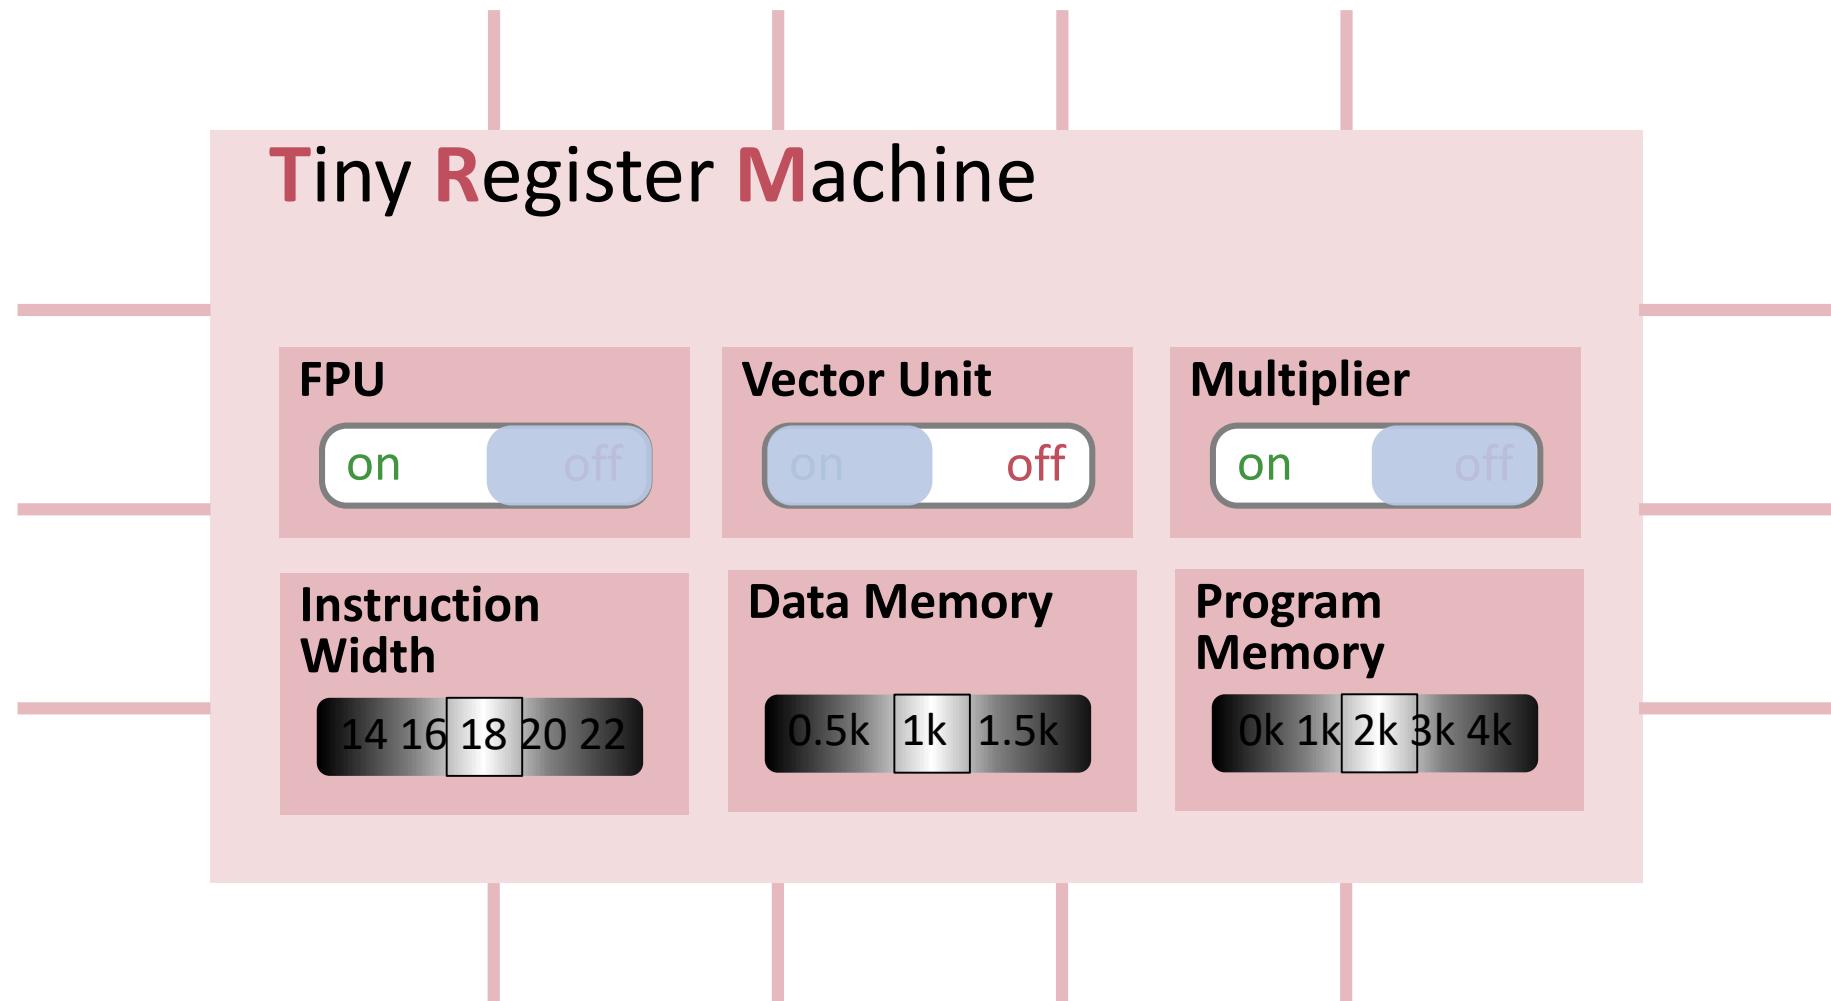
end Controller;

....

Port Width



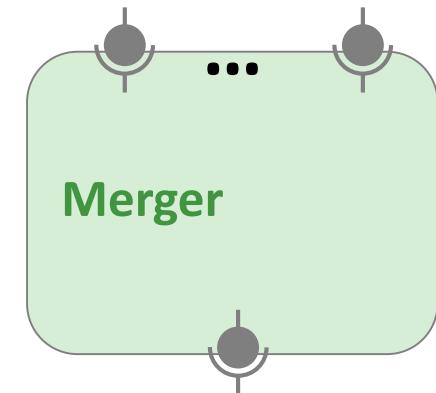
Configurable Processor on PL



Engine

Engines are prebuilt components instantiated as electronic circuits on a target hardware

```
type
Merger = cell {Engine, inputs=1}
    (ind: array inputs of port in; outd: port out);
var data: longint;
begin
    loop
        for i := 0 to len(in)-1 do
            data << ind[i]
            outd << data;
        end
    end
end Merger;
```



Unit of Deployment: (Terminal) Cellnet

Network* = CELLNET

VAR

```
    detector: PitchDetector;  
    tx : Engines.UartTx;  
    rx: Engines.UartRx;  
    sampler: I2SSampler;  
    rgb: RGB;
```

BEGIN

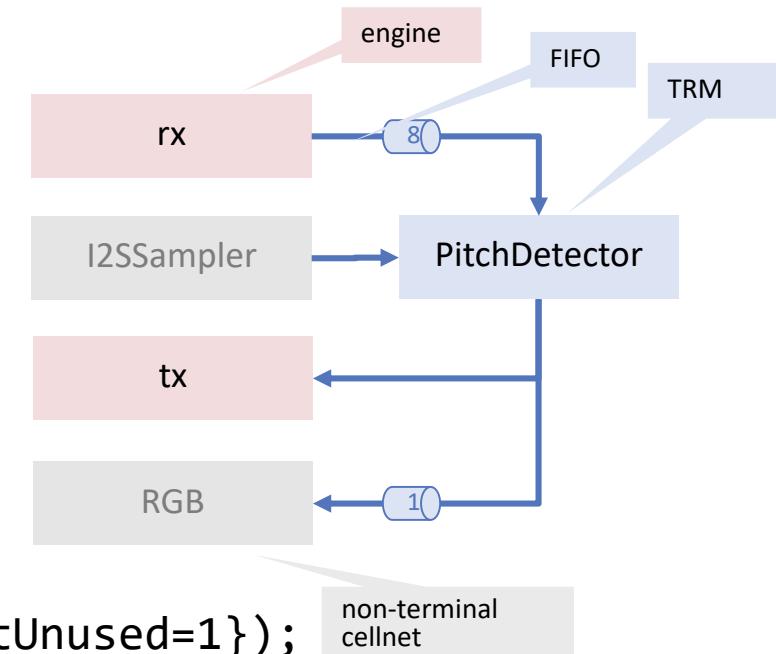
```
    NEW( detector );  
    NEW( tx {ClkDivisorWidth=16, InitClkDivisor=434, CtsPortUnused=1});  
    NEW( rx {ClkDivisorWidth=16, InitClkDivisor=434, RtsPortUnused=1});  
    NEW( rgb );  
    NEW( sampler );  
CONNECT(detector uartOut, tx.input);  
CONNECT(rx.output, detector uartIn, 8);  
CONNECT(sampler.output, detector.i2s);  
CONNECT(detector.colorOut, rgb.input, 1);
```

END Network;

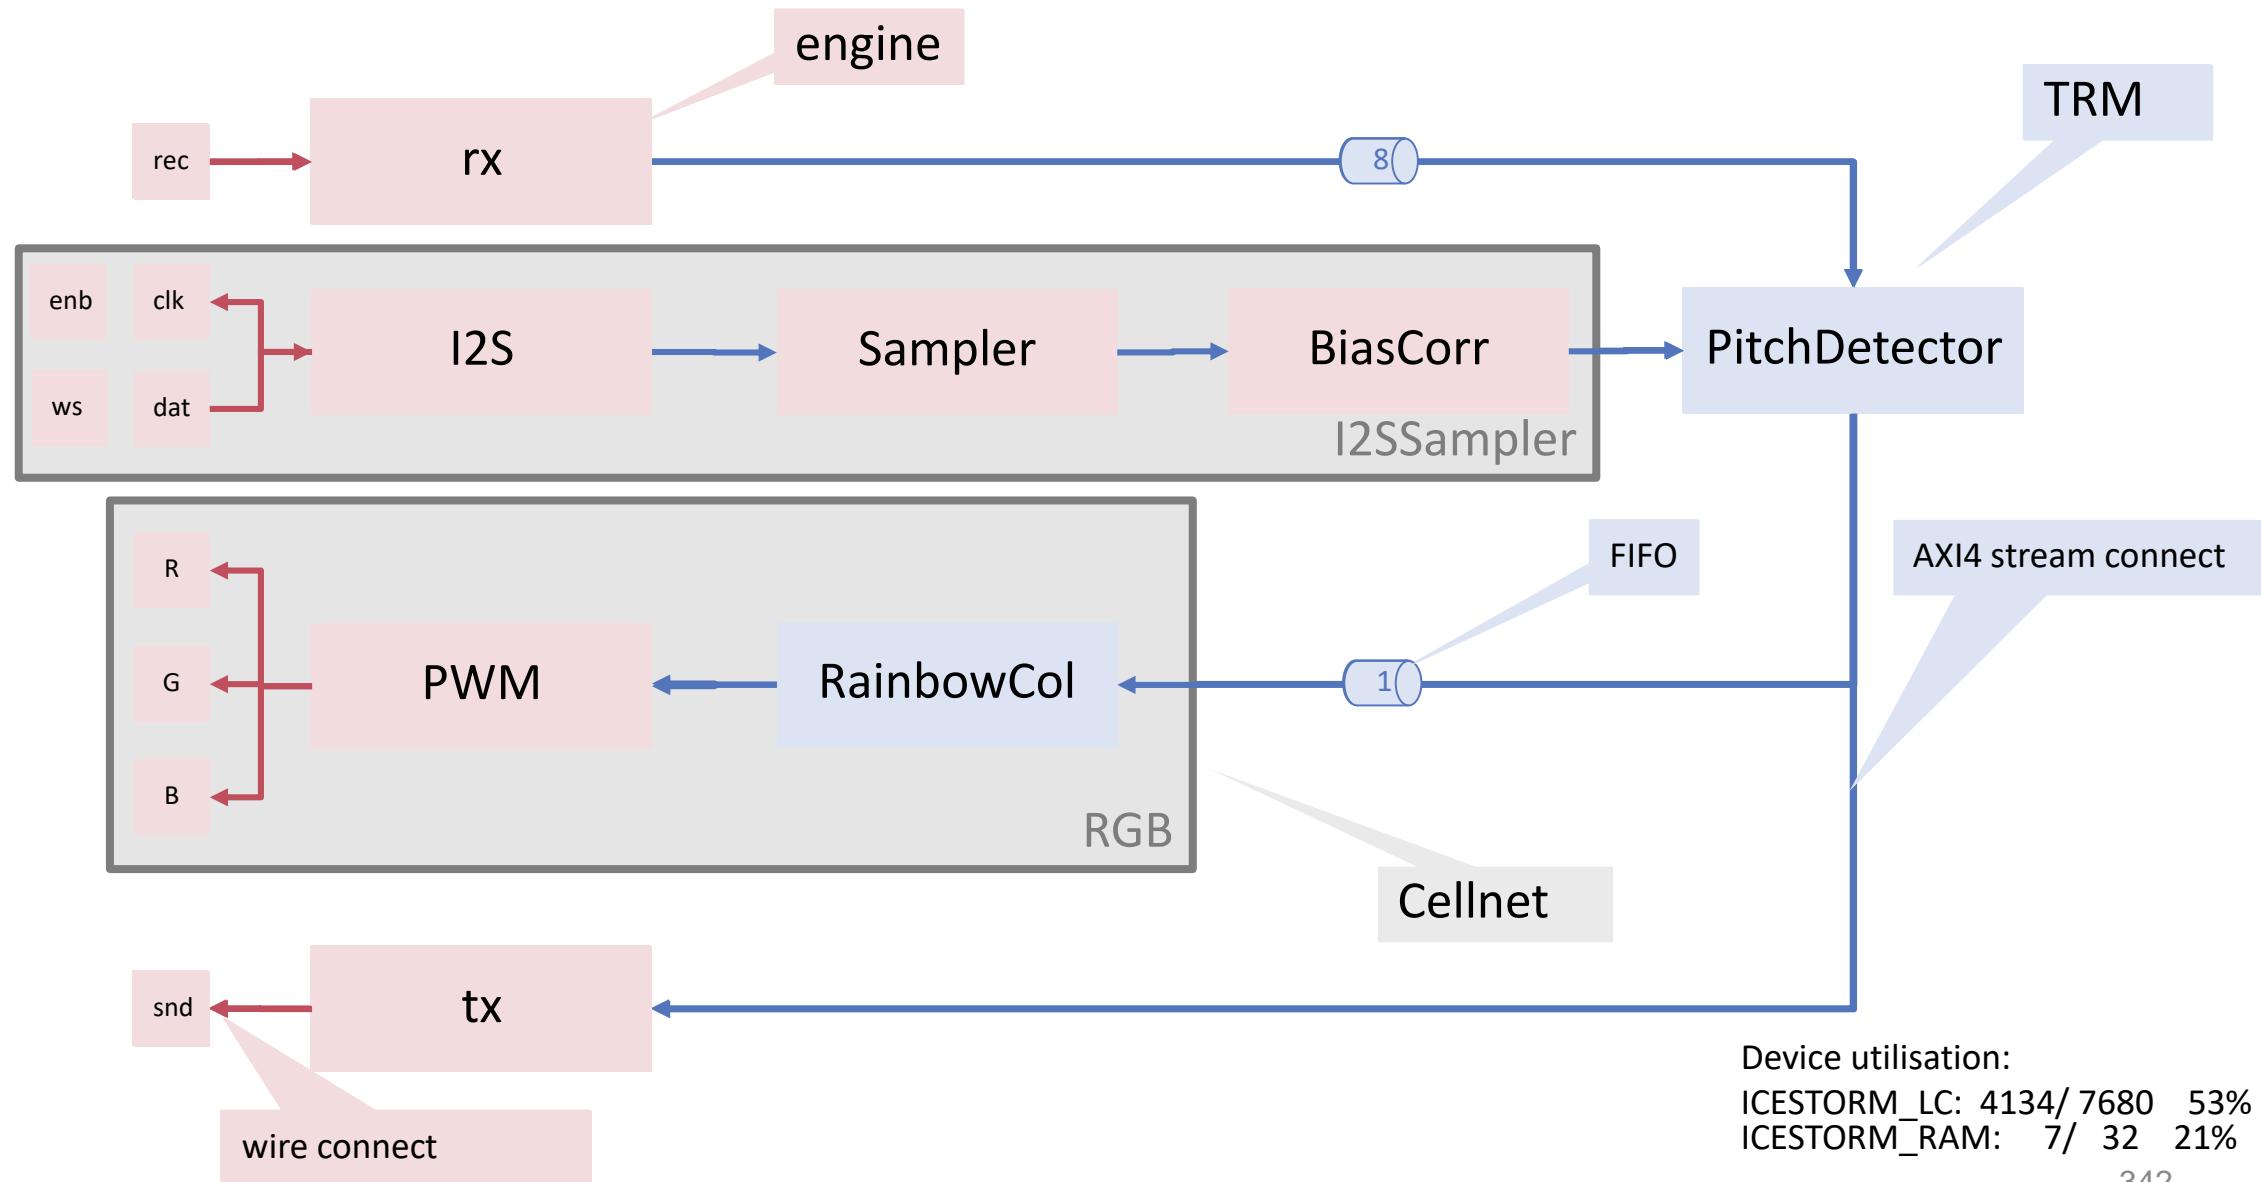
connection

properties

fifo depth (if any)



Hierarchical Composition: non-terminal Cellnets



Hierarchical Composition: non-terminal Cellnet

```
I2SSampler = CELLNET (output: PORT OUT);
```

VAR

```
    i2s: Engines.I2S;  
    sampler: Engines.Sampler;  
    biasCorrector: Engines.BiasCorrector;
```

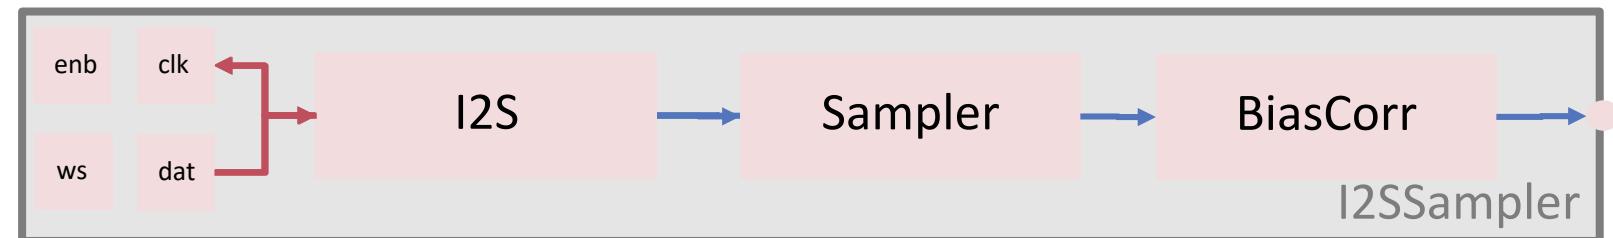
BEGIN

```
    NEW( i2s );  
    NEW( sampler );  
    NEW( biasCorrector );  
    CONNECT(i2s.output, sampler.input);  
    CONNECT(sampler.output, biasCorrector.input);  
    DELEGATE(output, biasCorrector.output);
```

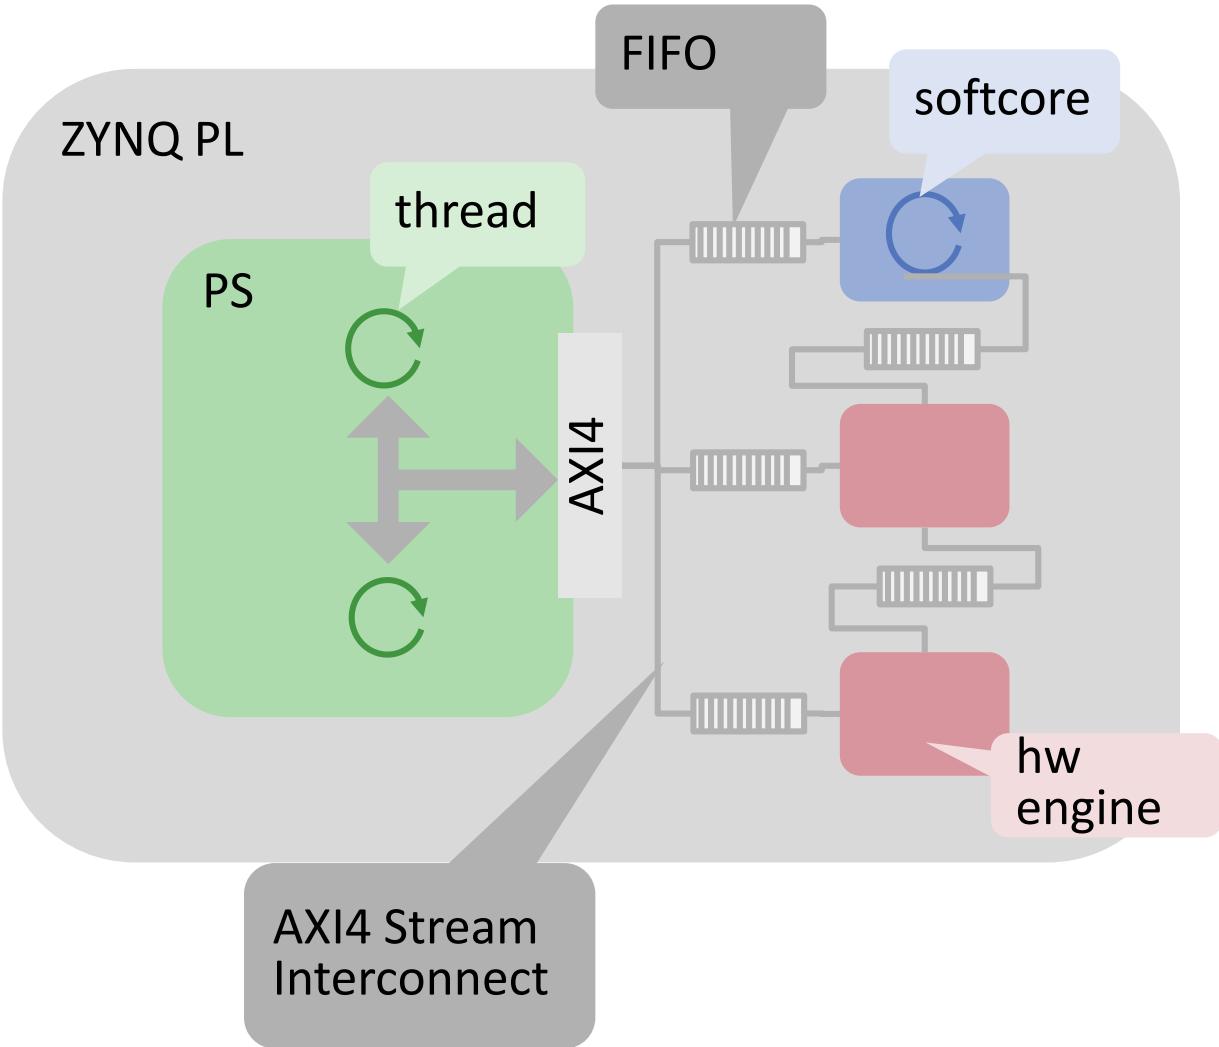
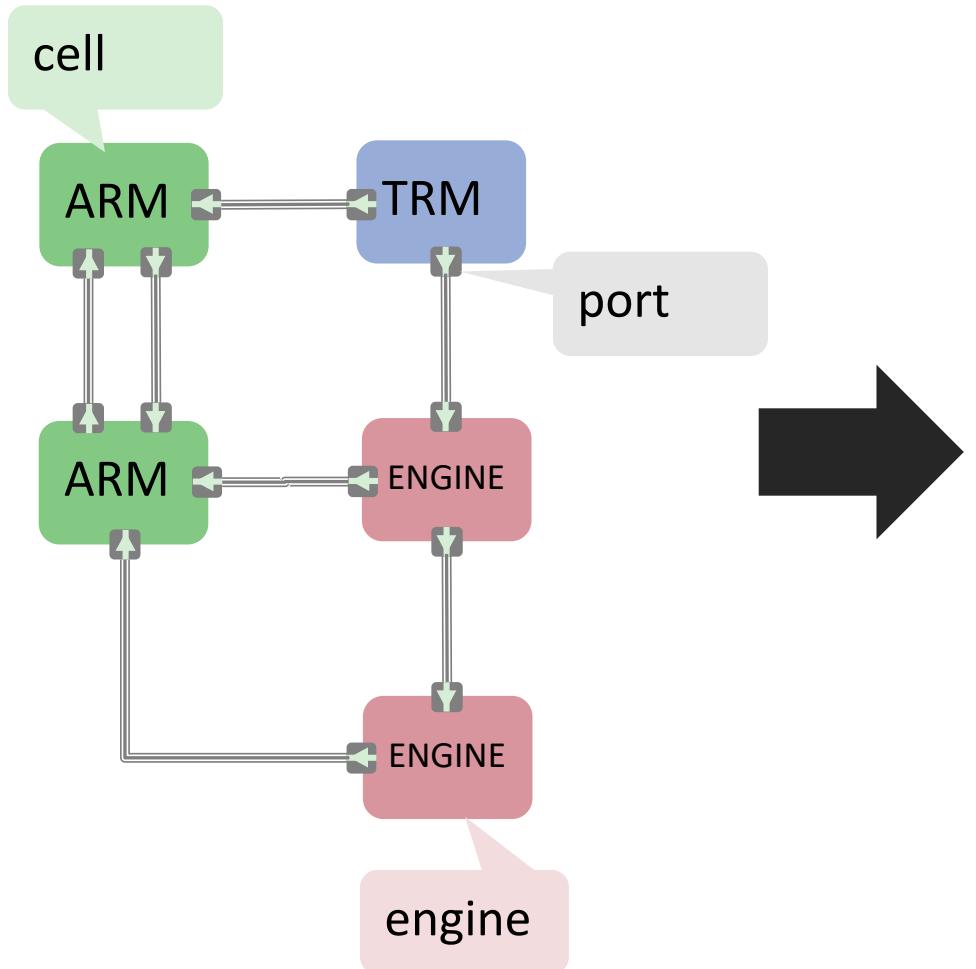
END I2SSampler;

ports

port delegation

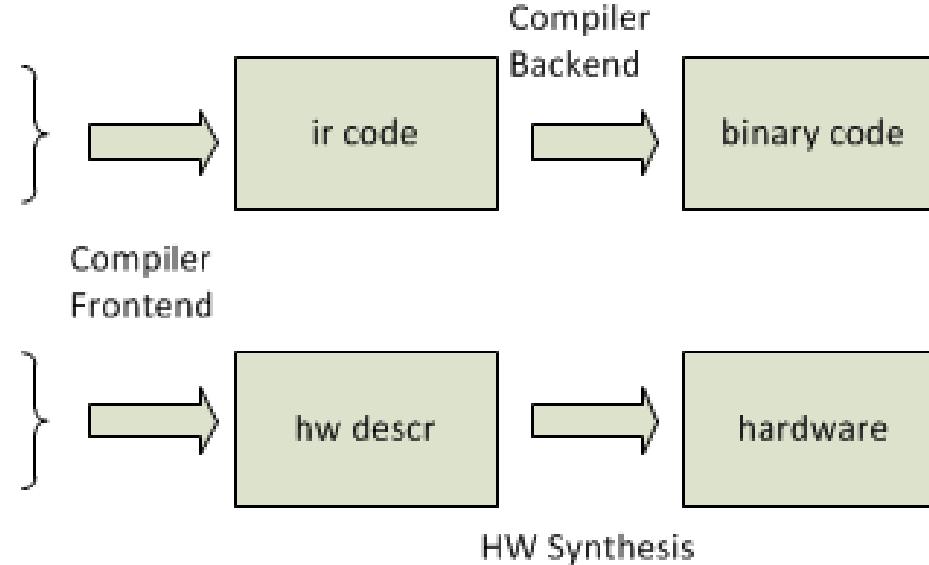


Software → Hardware Map



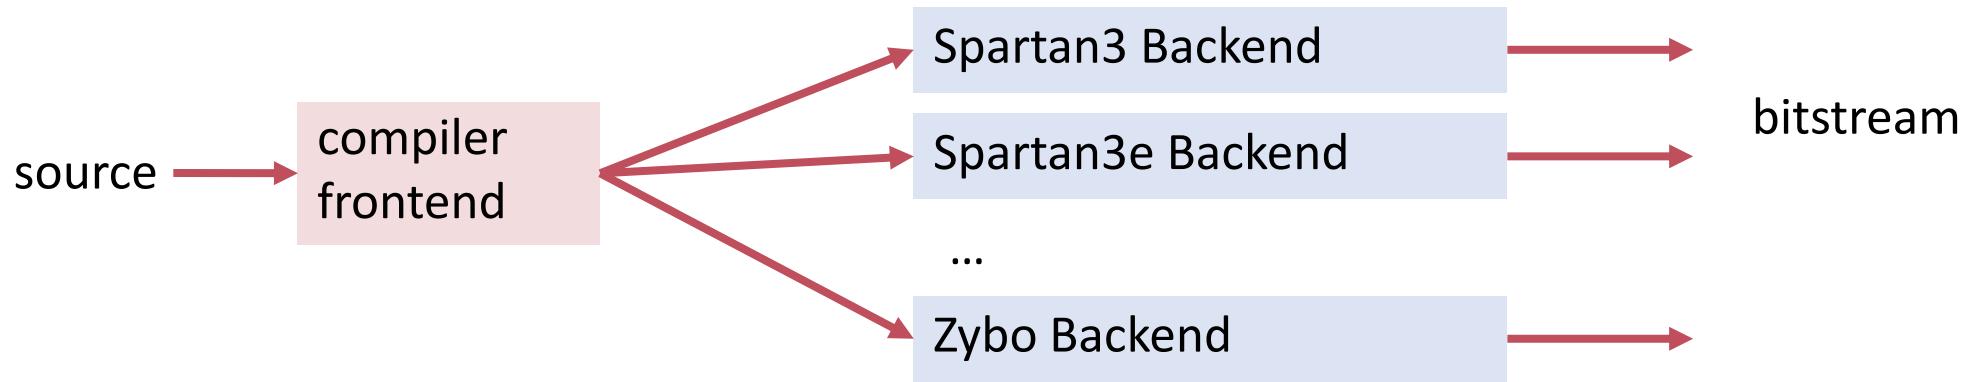
Hybrid Compilation

```
cellnet N;  
  
type A=cell(pi: port in; po: port out);  
var x: integer;  
begin  
... x << pi; ... po << x; ...  
end A;  
  
var a,b: A;  
begin  
... connect(a.po, b.pi)  
end N.
```



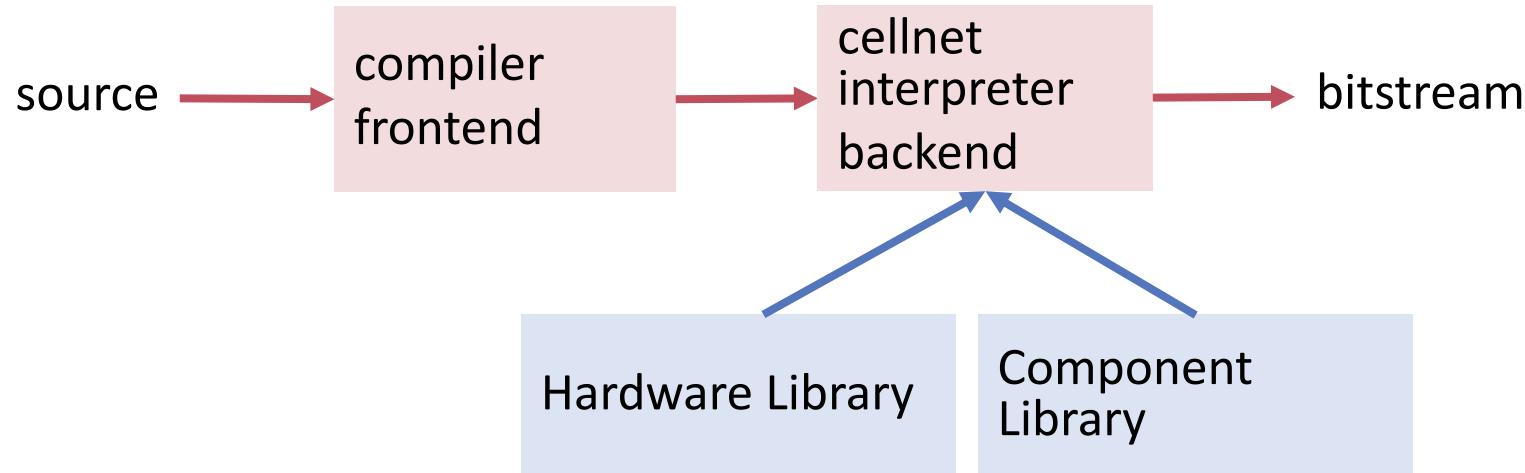
Code body	Role	Compilation method
Cell (Softcore)	Program logic	Software Compilation
Cell (Engine)	Computation unit	Hardware Generation
Cell Net	Architecture	Hardware Compilation

Implementation Alternatives (1)



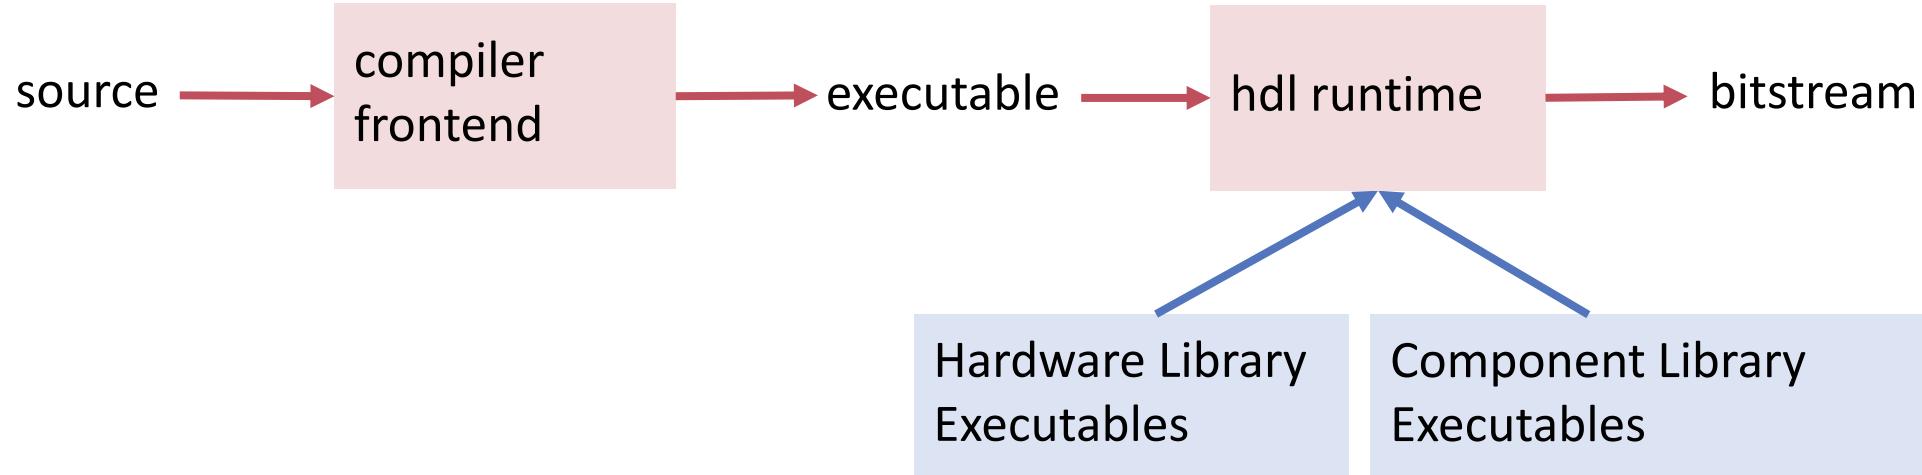
- + simple
- redundant
- not flexible
- hard to extend

Implementation Alternatives (2)



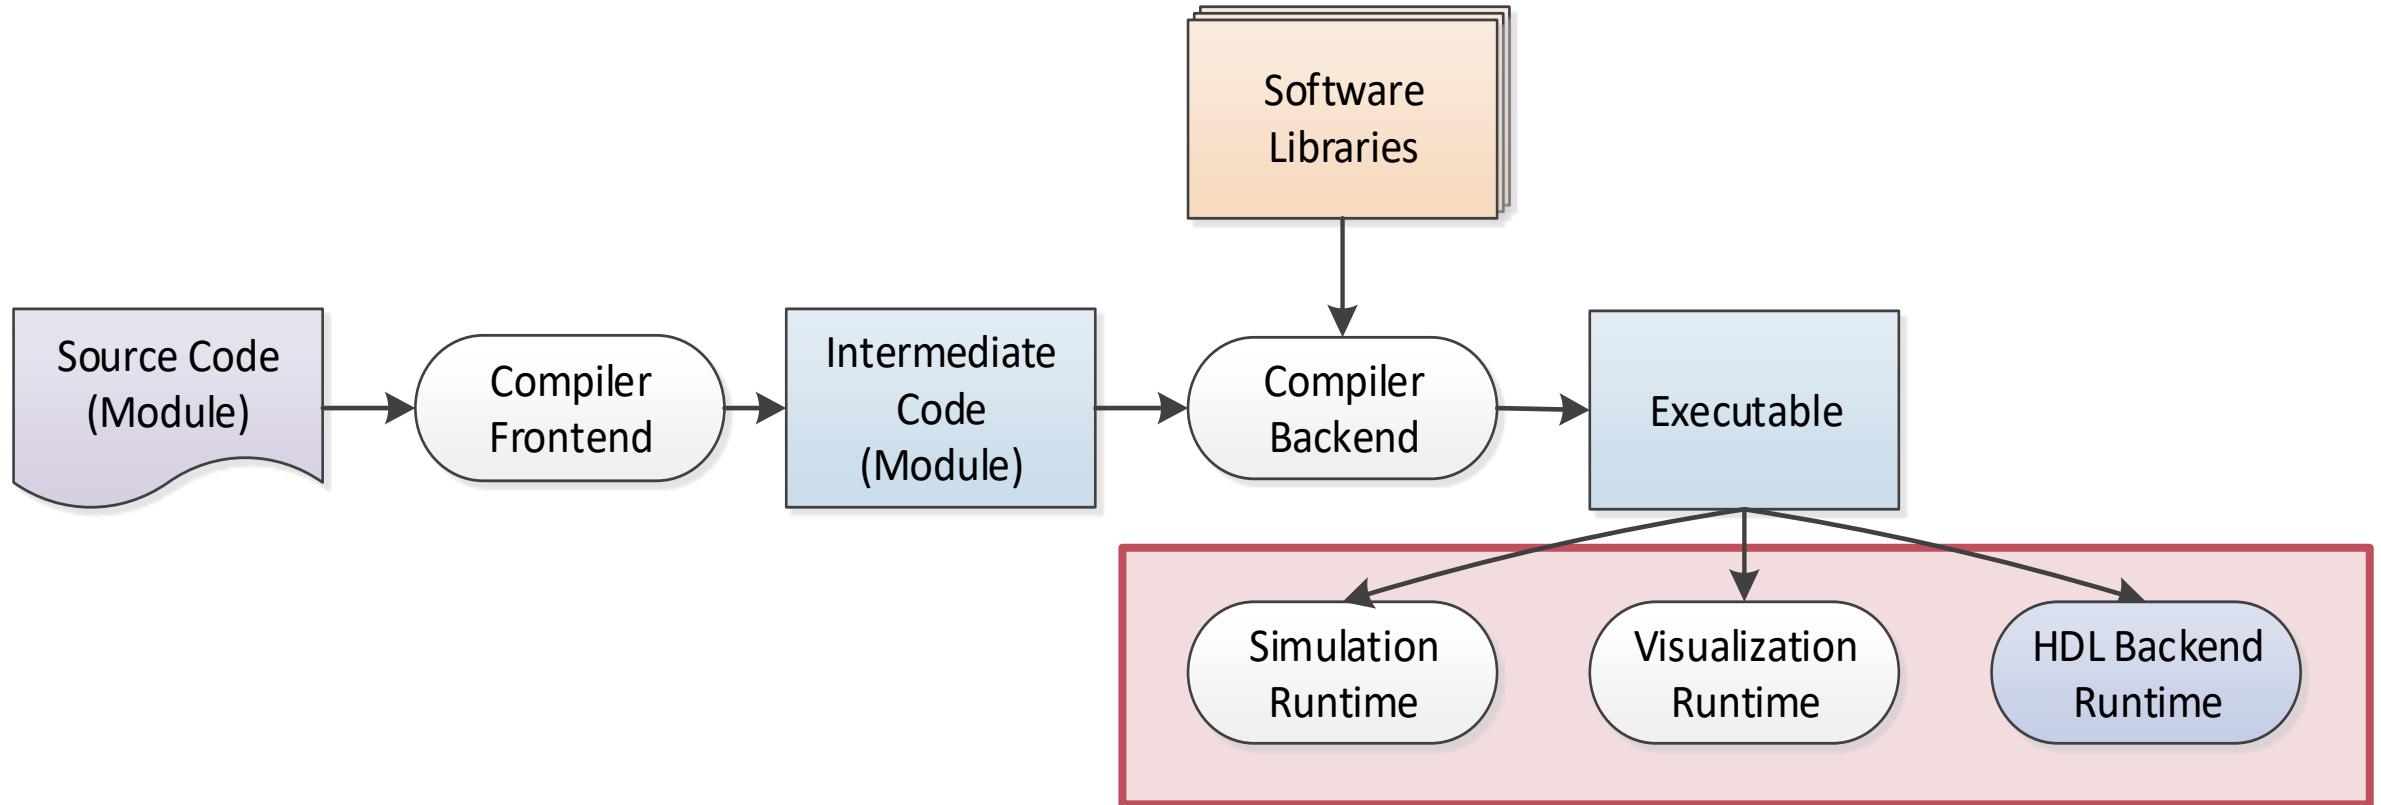
- + extensible
- + not redundant
- not simple
- static configuration limits flexibility

Implementation Alternatives (3)

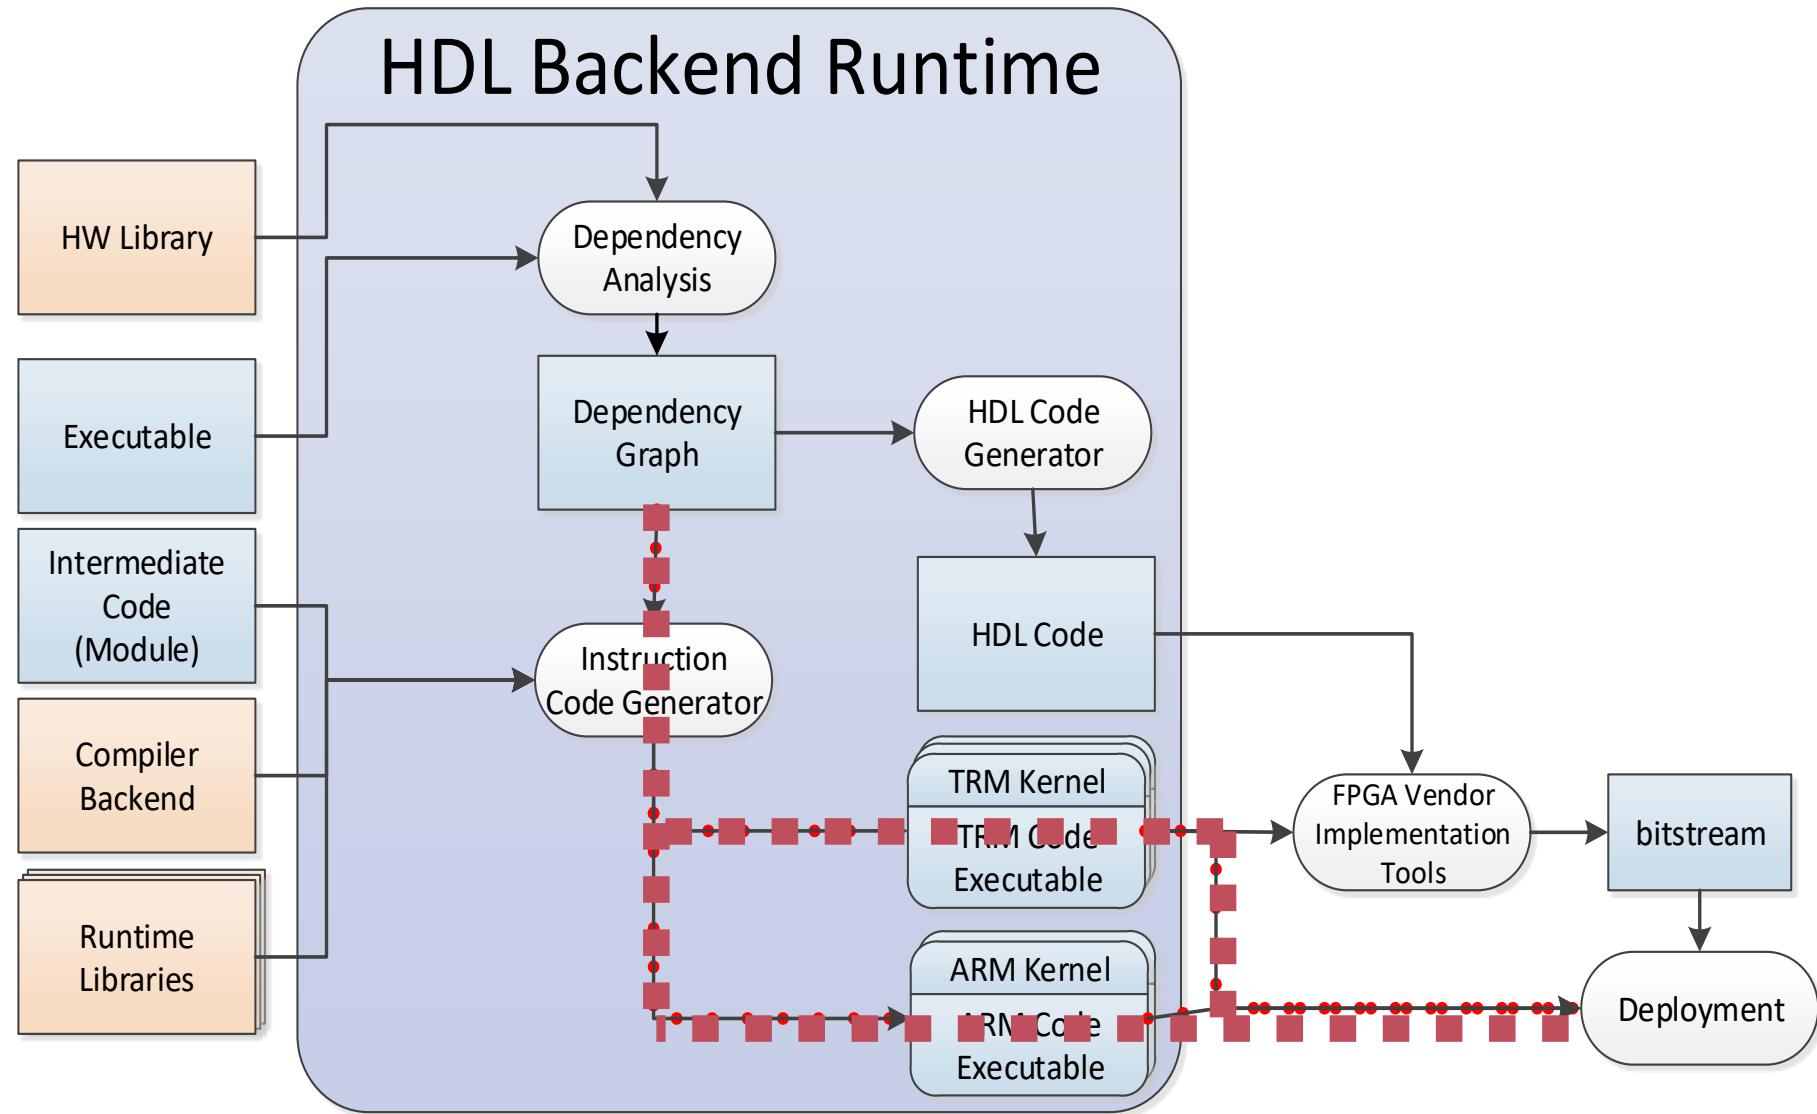


- + extensible
- + not redundant
- + simpler
- + simulation becomes execution

Frontend



Backend



Extensibility: Defining components and platforms

Software HLL Code

```
type Gpo =  
cell{Engine,DataWidth=32,InitState="0H"}  
(input: port in);
```

Build Command

```
AcHdlBackend.Build  
--target="ZyboBoard"  
-p="Vivado"  
CRBM.TestCellnet ~
```

Component Specification

Platform Specification

Hardware HDL Code

```
module Gpo  
#( parameter integer DW = 8 ... )  
(  
input aclk, input aresetn , input [DW-1:0] ...  
);
```

Hardware Platform and Tools

Hardware Types
Platform Instances
Vendor specific tools

Component Specification

module Gpo; import Hdl := AcHdlBackend; var c: Hdl.Engine; begin	Identification and description
new(c,"Gpo","Gpo"); c.SetDescription("General Purpose Output ... ");	Supported devices
c.SupportedOn("*"); (* portable *)	Parameters
c.NewProperty("DataWidth","DW",Hdl.NewInteger(32),Hdl.IntegerPropertyRangeCheck(1,Hdl.MaxInteger)); c.NewProperty("InitState","InitState",Hdl.NewBinaryValue("0H"),nil);	Ports
c.SetMainClockInput("aclk"); (* main component's clock *) c.SetMainResetInput("aresetn",Hdl.ActiveLow); (* active-low reset *) c.NewAxisPort("input","inp",Hdl.In,8); c.NewExternalHdlPort("gpo","gpo",Hdl.Out,8);	Dependencies
c.NewDependency("gpo","inp",true,false);	Configuration Actions
c.AddPostParamSetter("gpo",Hdl.HdlProperty.SetterWidthFromProperty("inp","DW")); c.AddPostParamSetter("gpo",Hdl.HdlProperty.SetterWidthFromProperty("gpo","DW")); Hdl.hwLibrary.AddComponent(c);	c.NewExternalHdlPort("gpo","gpo",Hdl.Out,8);
end Gpo.	

Target Platform Specification

```
module Basys2Board;
import Hdl := AcHdlBackend, AcXilinx;
var t: Hdl.TargetDevice;
    pldPart: AcXilinx.PldPart;
    ioSetup: Hdl.IoSetup;
    pin: Hdl.IoPin;
begin
    new(pldPart,"XC3S100E-4CP132");
    pldPart.SetJtagChainIndex(0);
    new(t,"Basys2Board",pldPart);

    new(pin,Hdl.In,"B8","LVCMOS33");
    t.NewExternalClock(pin,50000000,50,0); (* ExternalClock0 *)
    t.SetSystemClock(t.clocks.GetClockByName("ExternalClock0"),1,1);
    new(pin,Hdl.In,"G12","LVCMOS33");
    t.SetSystemReset(pin,true);

    new(ioSetup,"Gpo_0");
    ioSetup.NewIoPort("gpo",Hdl.Out,"U16,E19,U19,V19","LVCMOS33");
    t.AddIoSetup(i...`..);

    Hdl.hwLibrary.Add(t);
end Basys2Board.
```

FPGA Part

System Signals

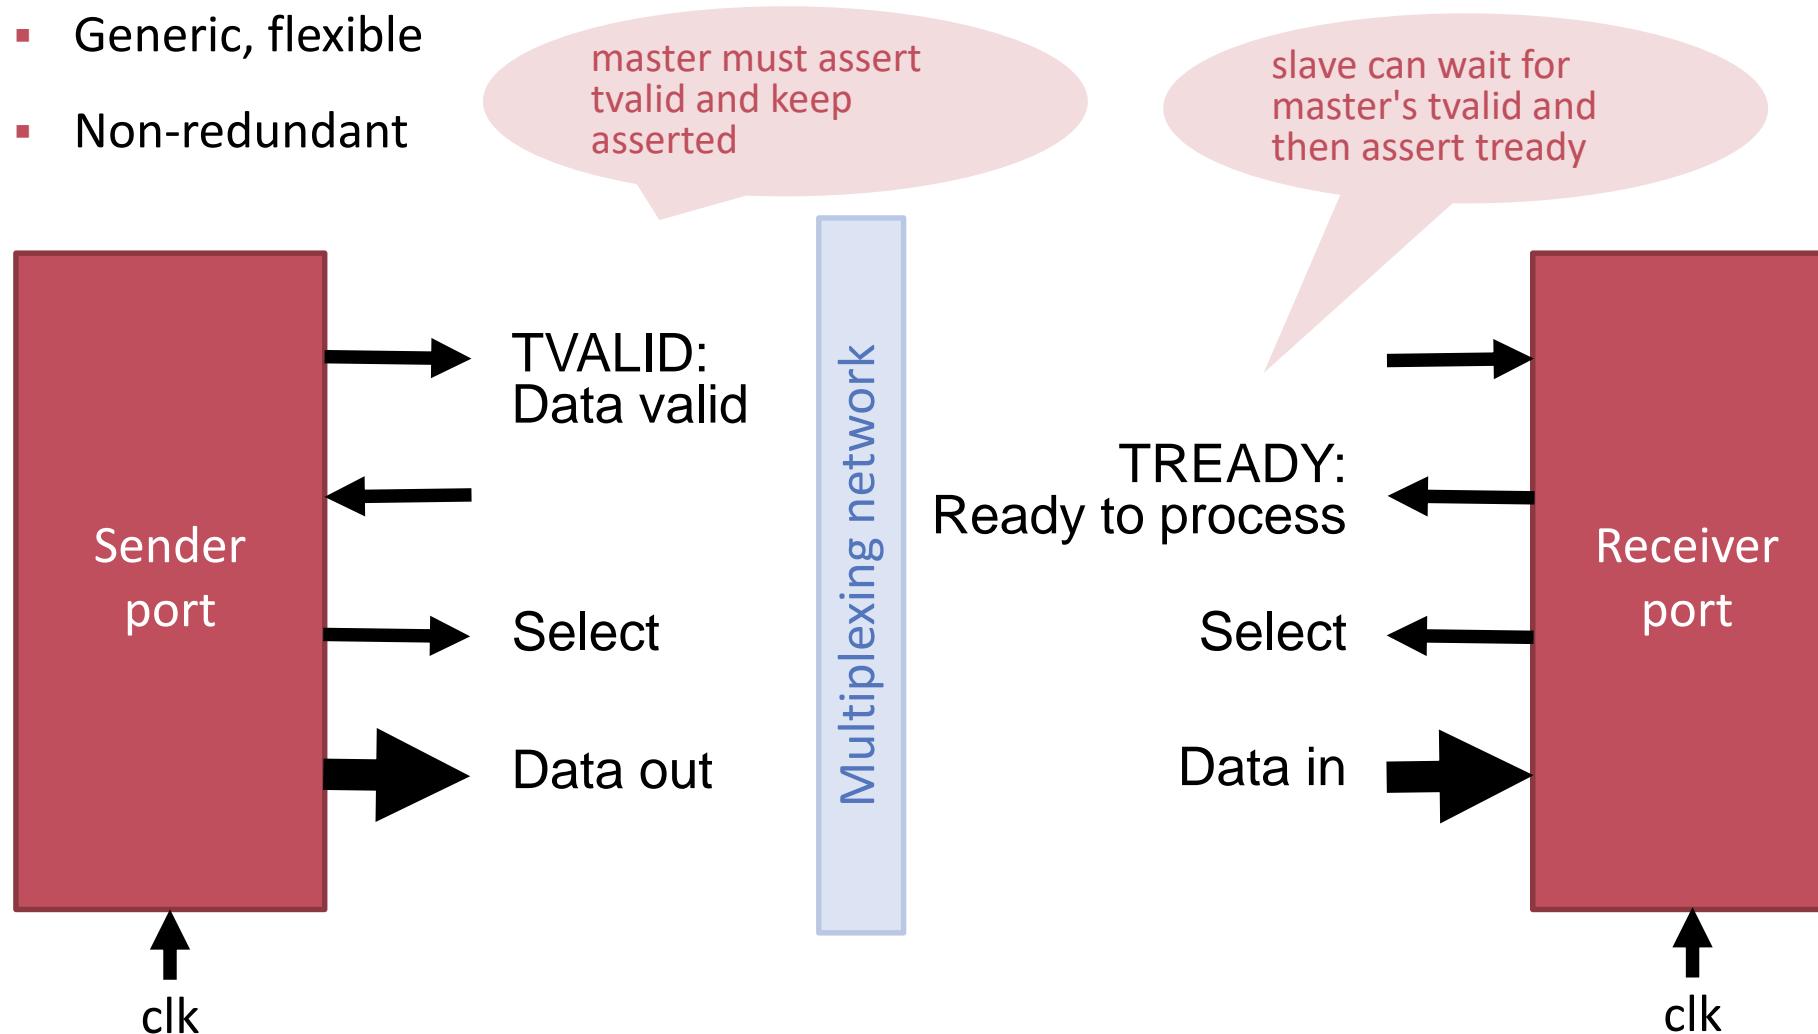
Mapping of external Ports

ioSetup.NewIoPort("gpo",Hdl.Out,"U16,E19,U19,V19","LVCMOS33");

Generic Peer-to-Peer Communication Interface

- Use of **AXI4 Stream** interconnect standard from ARM

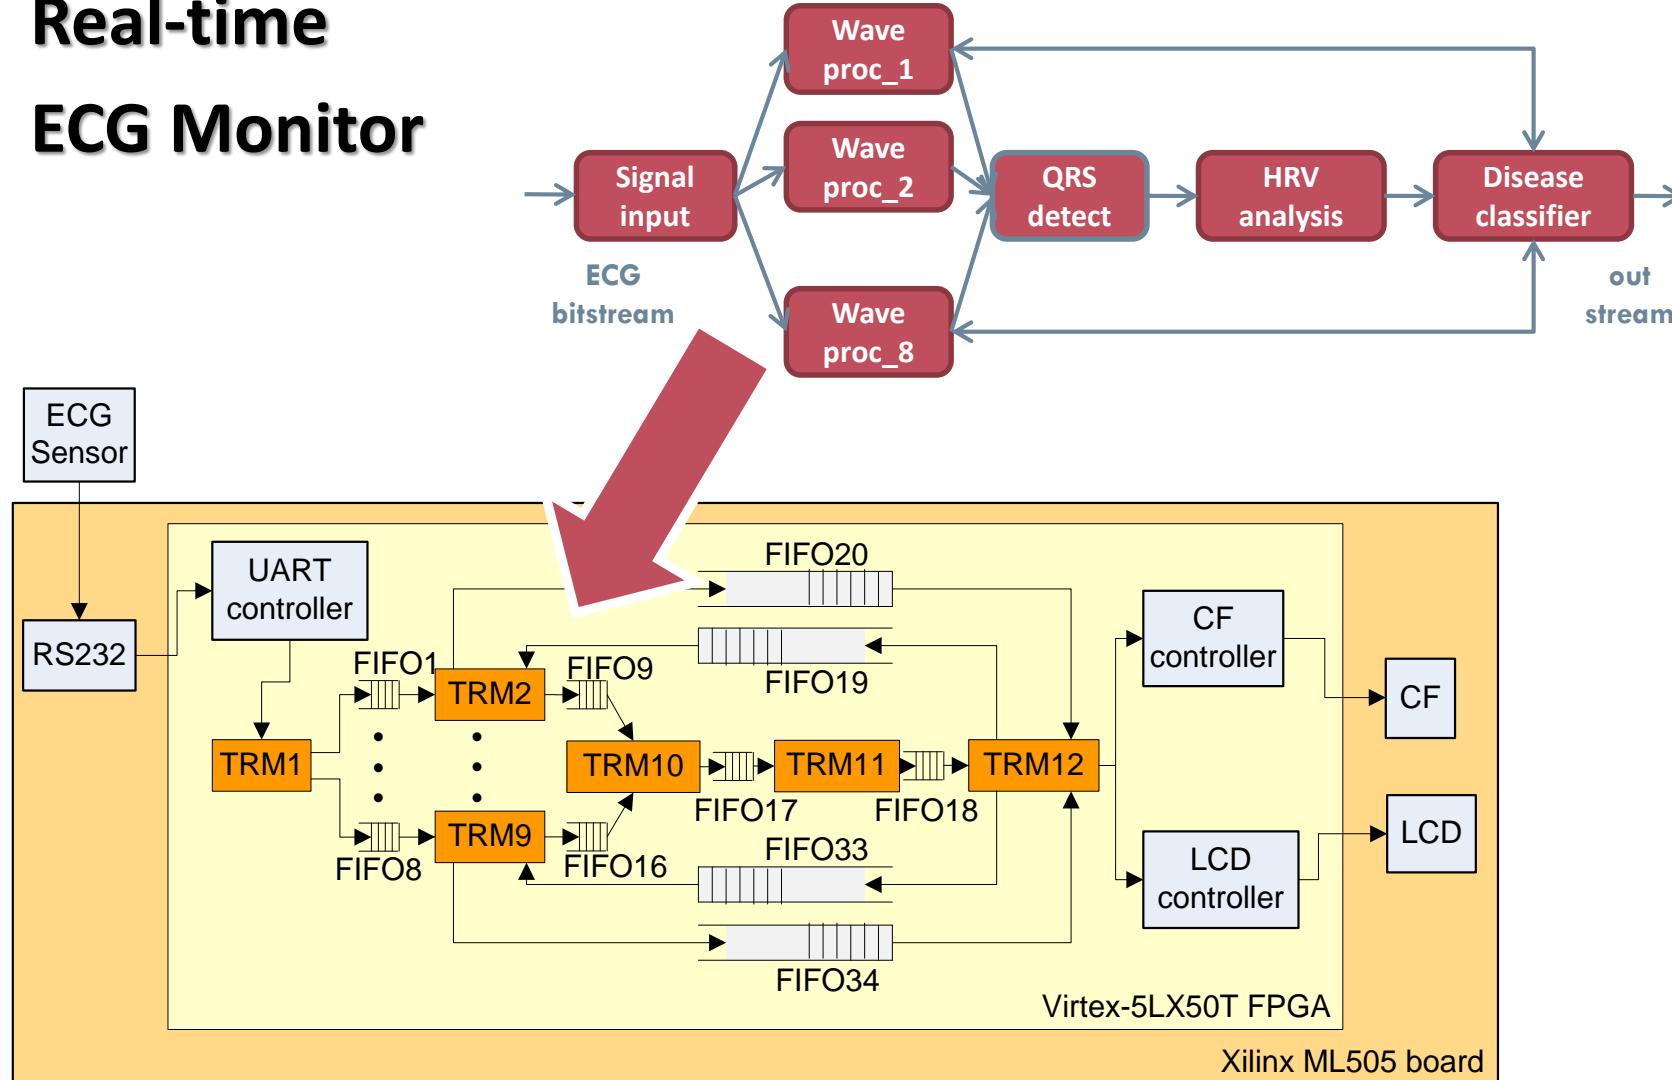
- Generic, flexible
- Non-redundant



Case Study 1: ECG

Focus: Resources and Power

Real-time ECG Monitor



Resources

- ECG Monitor*

#TRMs	#LUTs	#BRAMs	#DSPs	TRM load
12	13859 (48%)	52 (86%)	12 (25%)	<5% @116 MHz

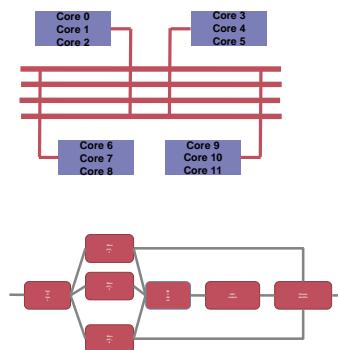
- Maximum number of TRMs in communication chain

FPGA	#TRMs	#LUTs	#BRAMs	#DSPs
Virtex-5	30	27692 (96%)	60 (100%)	30 (62%)
Virtex 6	500			

*8 physical channels @ 500 Hz sampling frequency
implemented on Virtex 5

Comparative Power Usage

- Preconfigured FPGA (#TRMs, IM/DM, I/O, Interconnect fixed) versus fully configurable FPGA (Active Cells)



System	Static Power (W)	Dynamic Power (W)
Preconfigured ("TRM12")	3.44	0.59
Dynamically configured	0.5	0.58



86% saving!

Case Study: Non-Invasive Continuous Blood Pressure Monitor

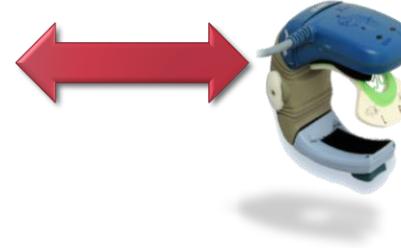
Focus: Development Cycle Time



A2 Host OS with GUI on ARM

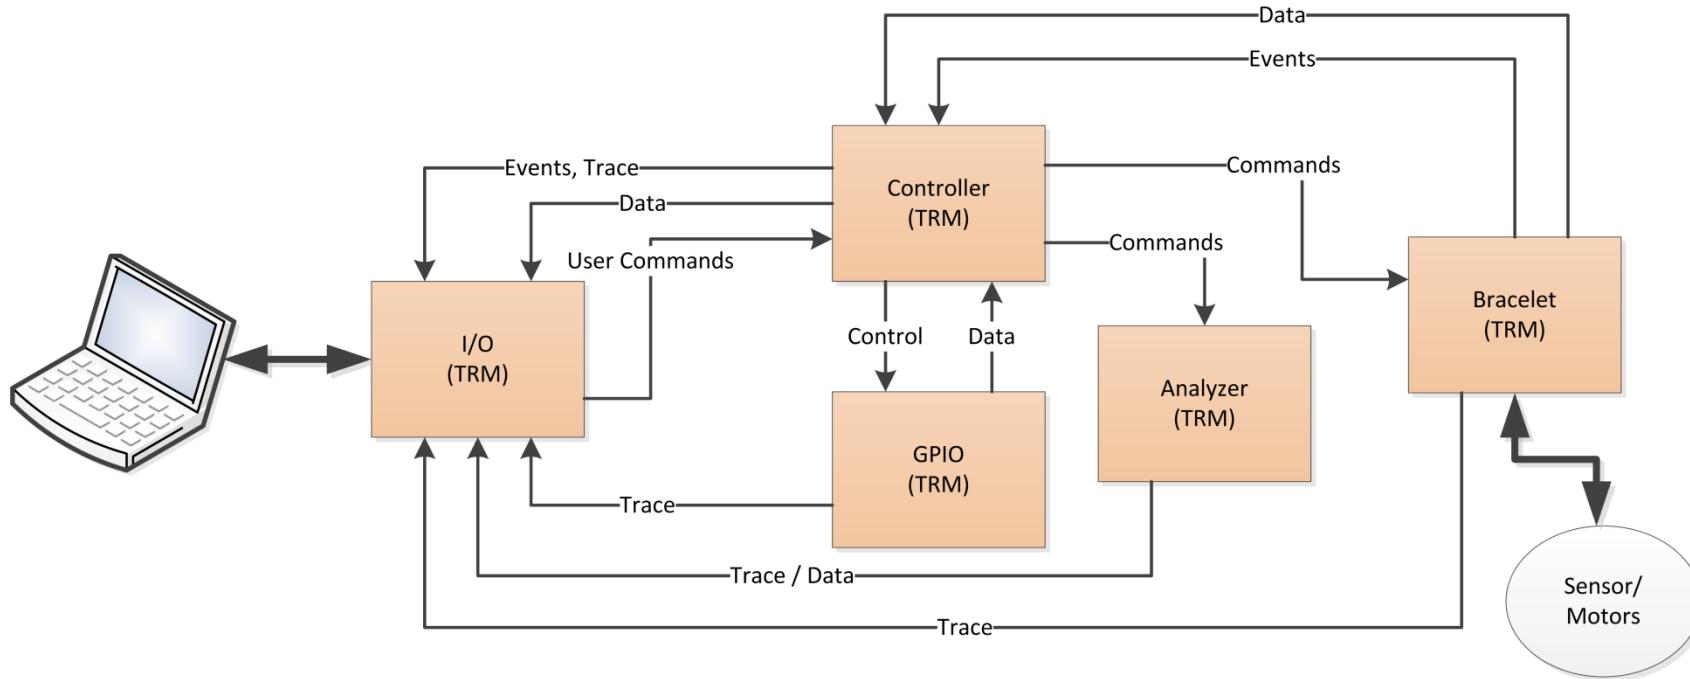


Sensor control and
medical algorithms on
Zynq PL



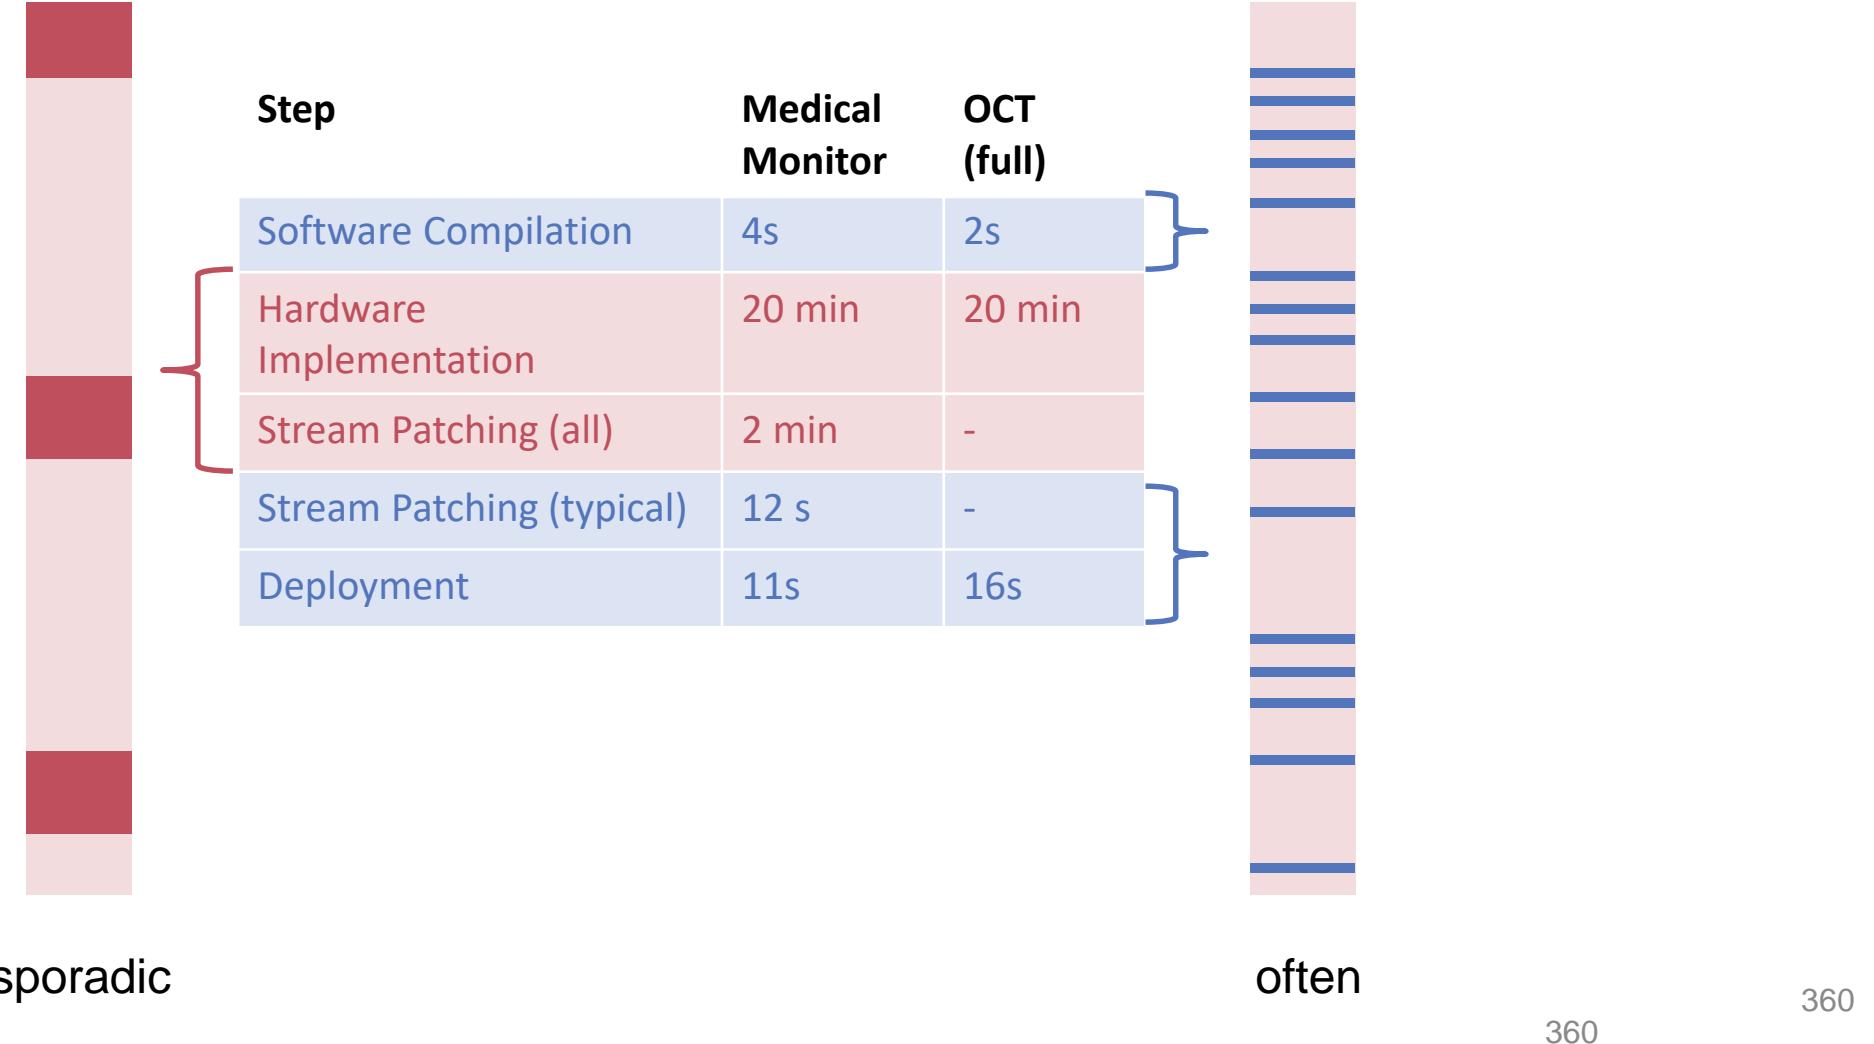
Sensors and Motors
on Bracelet

Medical Monitor Network On Chip



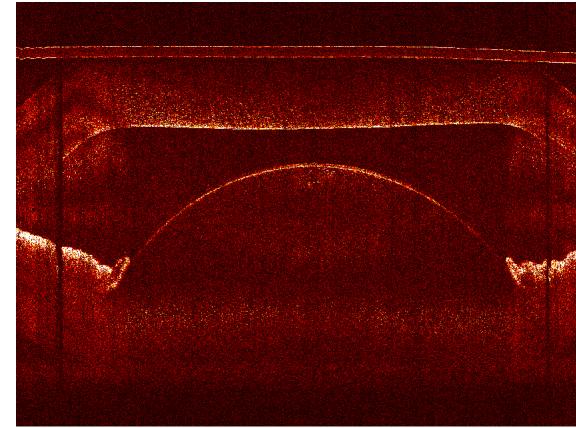
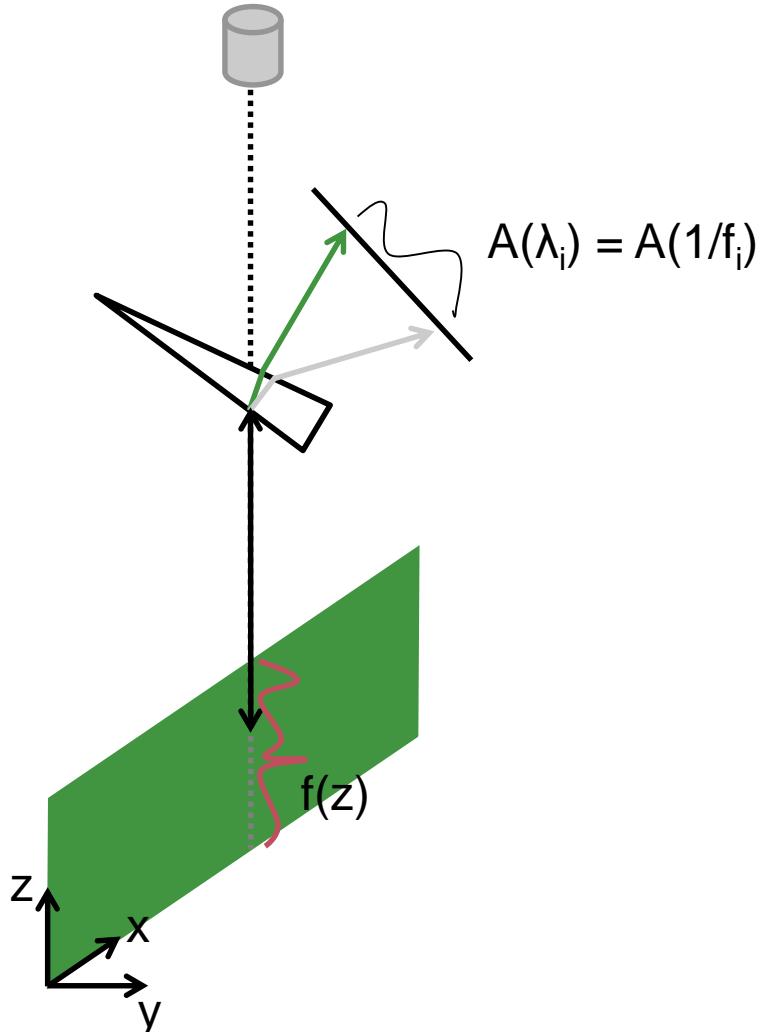
Dominated by TRM processors. Feedback driven. Not performance critical.

Development Cycle Times



Case Study 3: Optical Coherence Tomography

Focus: Performance



z-Axis Processing

1. Non uniform sampling

$$A(\lambda_i) \rightarrow \tilde{A}(f_i)$$

2. Dispersion compensation

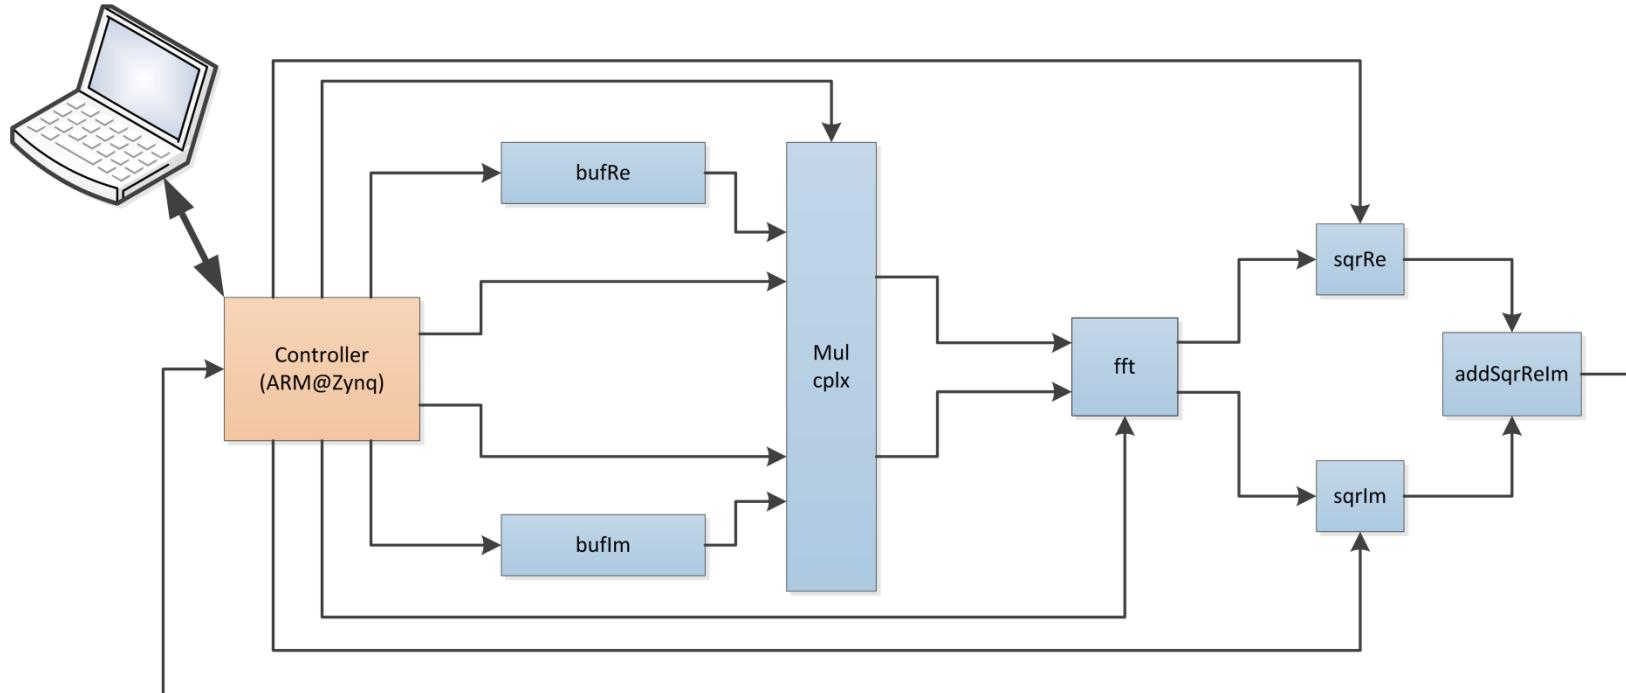
3. (Inverse) FFT

... for many lines x in a row (2d)

... and many rows y in a column (3d)

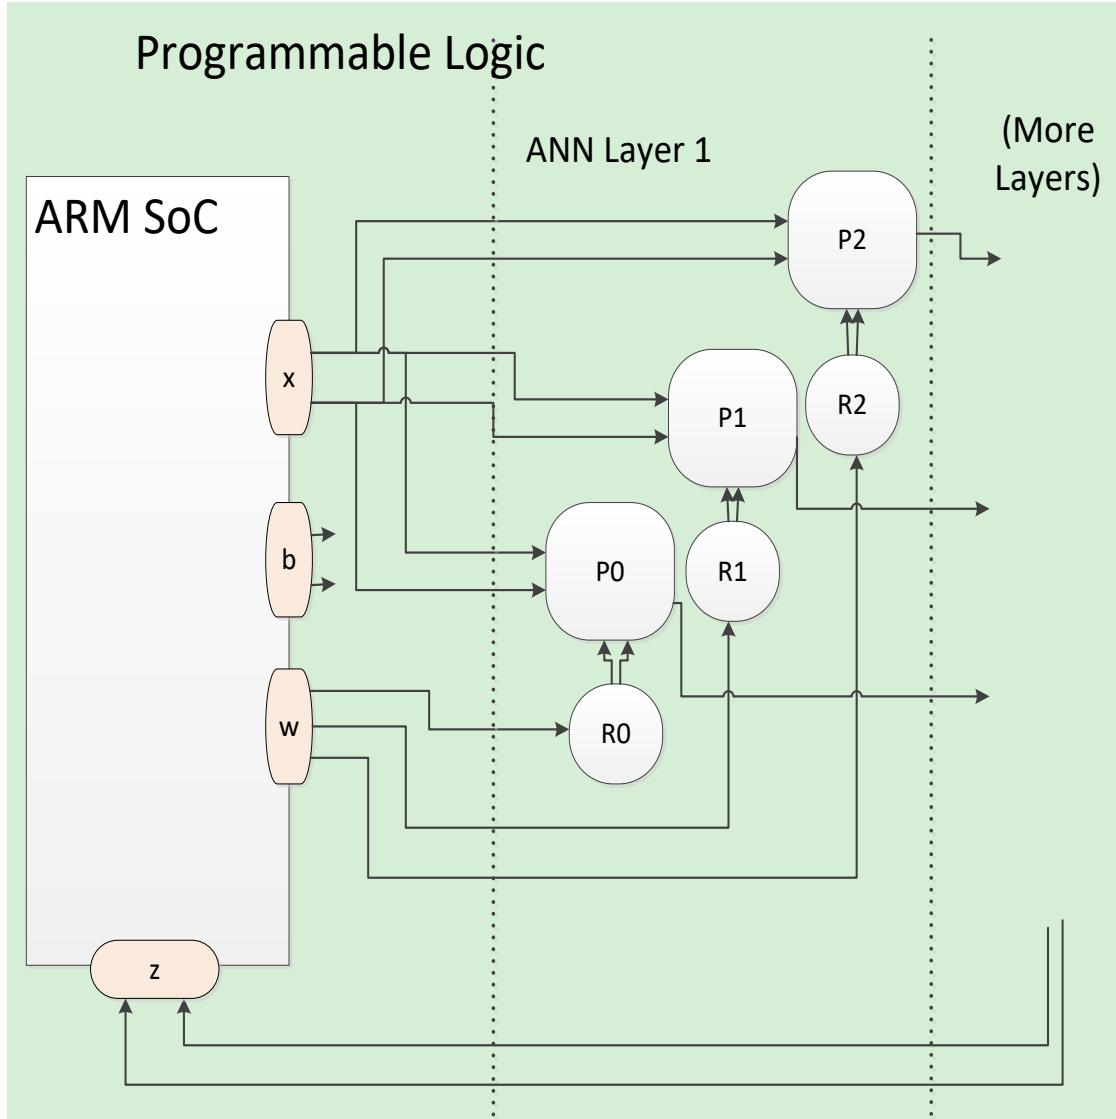
A component of OCT image processing

Dispersion Compensation

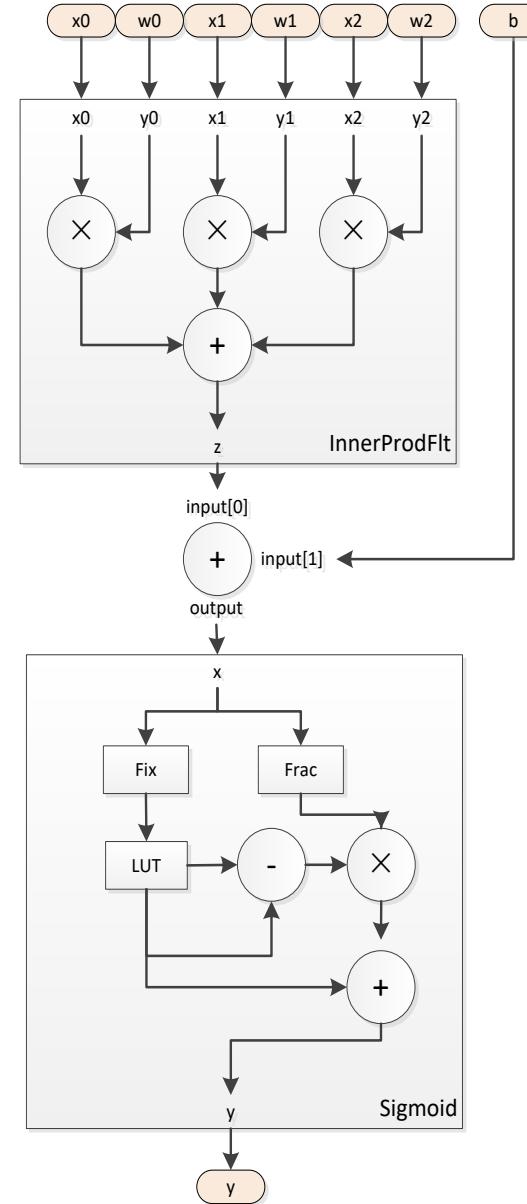


Dominated by Engines. Dataflow driven.

Case Study: ANN



Perceptron



Performance and Resource Usage

	Medical Monitor	Simple OCT	OCT	Perceptron
Architecture	Spartan 6 XC6SLX75	Zynq 7000 XC7Z020	Zynq 7000 XC7Z020	Zynq 7000 XC7Z010
Resources	28% Slice LUTs, 4% Slice Registers 80% BRAMs 24% DSPs	11% Slice LUTs, 6% Slice Registers 7% BRAMs 15% DSPs 1 ARM Cortex A9	17% Slice LUTS 8% Slice Registers 22% BRAMs 31% DSPs 1 ARM Cortex A9	83% Slice LUTS, 67% Slice Registers 13.33% BRAMs, 21% DSP 1 ARM Cortex A9
Clock Rate	58 MHz	118 MHz	50 MHz	147 MHz
Data Bandwidth	1.25 Mbit /s (in) 23 kB/s (out)	236 MWords/s (in) 118 MWords/s (out)	50 MWords/s (in) 50 MWords/s (out)	9.6 GBits/s in 9.6 GBits/s out
Performance	--	8.3 GFPOps* up to 32 GFPOps**	4.3 GFPOps*	4.9 GFlops
Power	~2W	~5W	~5W	2.1 W

** Fixed point operations, 32bit

* when instantiated 4 times

Conclusion

ActiveCells: Computing model and tool-chain for configurable computing

- Configurable interconnect → Simple Computing, Power Saving
- Embedding of task engines → High Performance
- Hybrid compilation → Quick Development
- Backend execution → Eased Flexibility and Extensibility

Mapping to Hardware – Simple Example

```
SimpleNetwork* = CELLNET
```

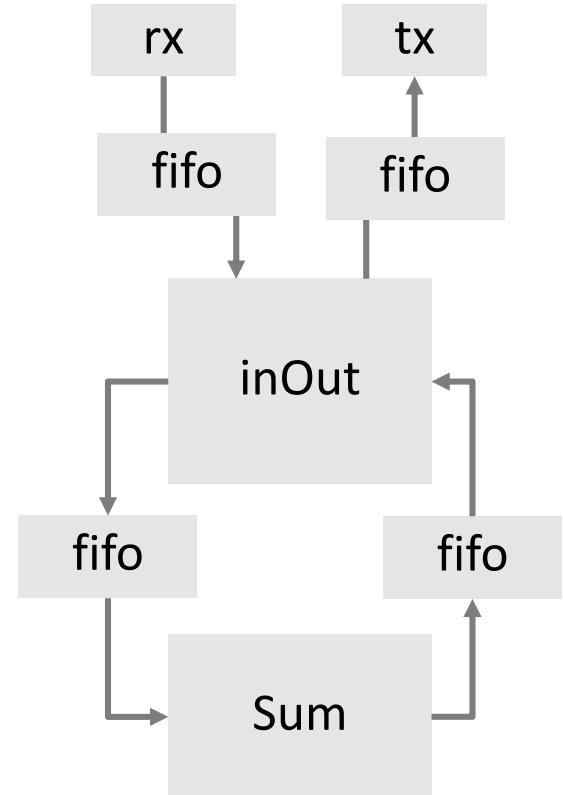
```
VAR
```

```
    sum: Sum; inout: InOut;  
    tx : Engines.UartTx;  
    rx: Engines.UartRx;
```

```
BEGIN
```

```
    NEW( tx {ClkDivisorWidth=16, InitClkDivisor=434, CtsPortUnused=1});  
    NEW( rx {ClkDivisorWidth=16, InitClkDivisor=434, RtsPortUnused=1});  
    NEW( sum );  
    NEW(inout);  
    CONNECT(rx.output, inout.uartIn,8);  
    CONNECT(inout.uartOut, tx.input,8);  
    CONNECT(sum.out, inout.in,8);  
    CONNECT(inout.out, sum.in,8);
```

```
END SimpleNetwork;
```



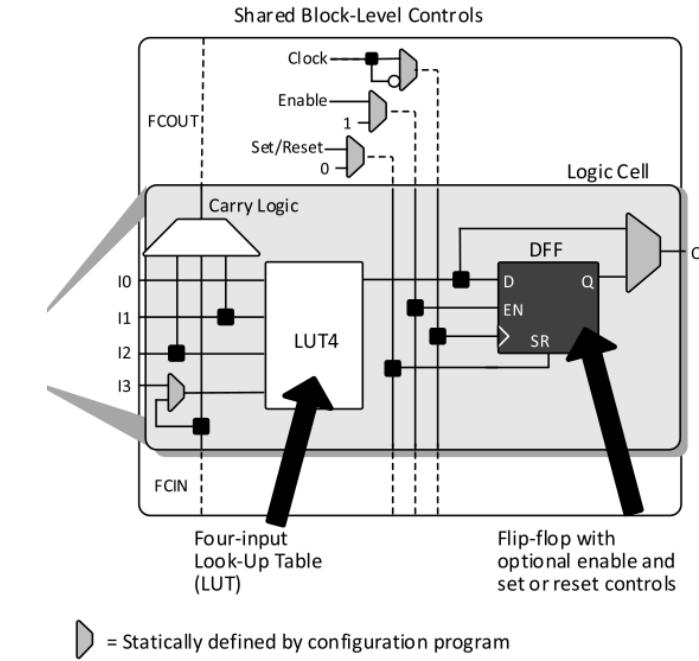
→ Blackboard / Code
inspection

Lattice iCE40 HX4K-tq144 Chip



Resources

LCells (4-input LUTs) +	7680
RAMs	32 x 4 KBit (32 x 256x16 bit)
DSPs	-

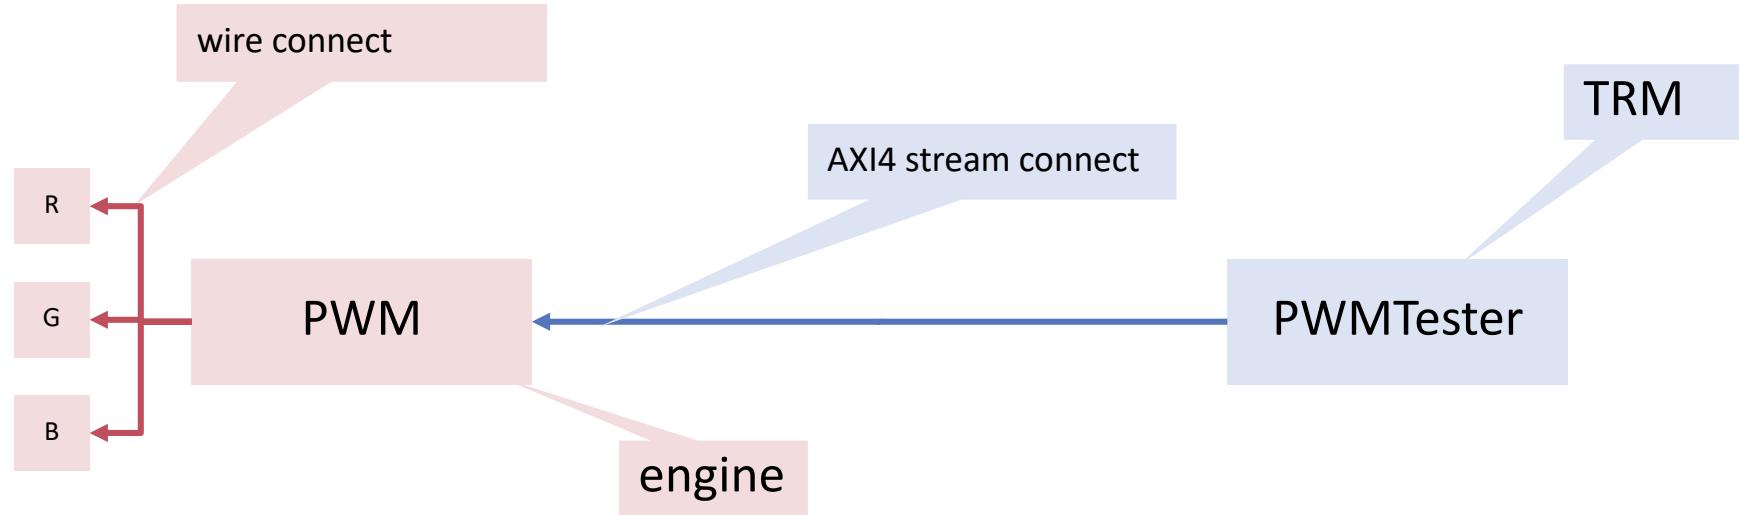


Resource Usage Scenarios

	TRM -> LED	TRM -fifo-> TRM -> LED	TRM –fifo-> TRM –fifo-> TRM -> LED	TRM –fifo-> TRM –fifo-> TRM –fifo-> TRM -> LED	TRM -> 8 fifo (8) -> TRM -> LED	TRM -> 8 fifo (1) -> TRM -> LED
LBs	1696 / 7680	3371/ 7680	5116/ 7680	6816/ 7680	3808 / 7680	3947 / 7680
RAMS	3 / 32	8 / 32	13 / 32	18 / 32	22 / 32	6 / 32
					(pnr failed for 50 MHz)	(pnr failed for 50 MHz)

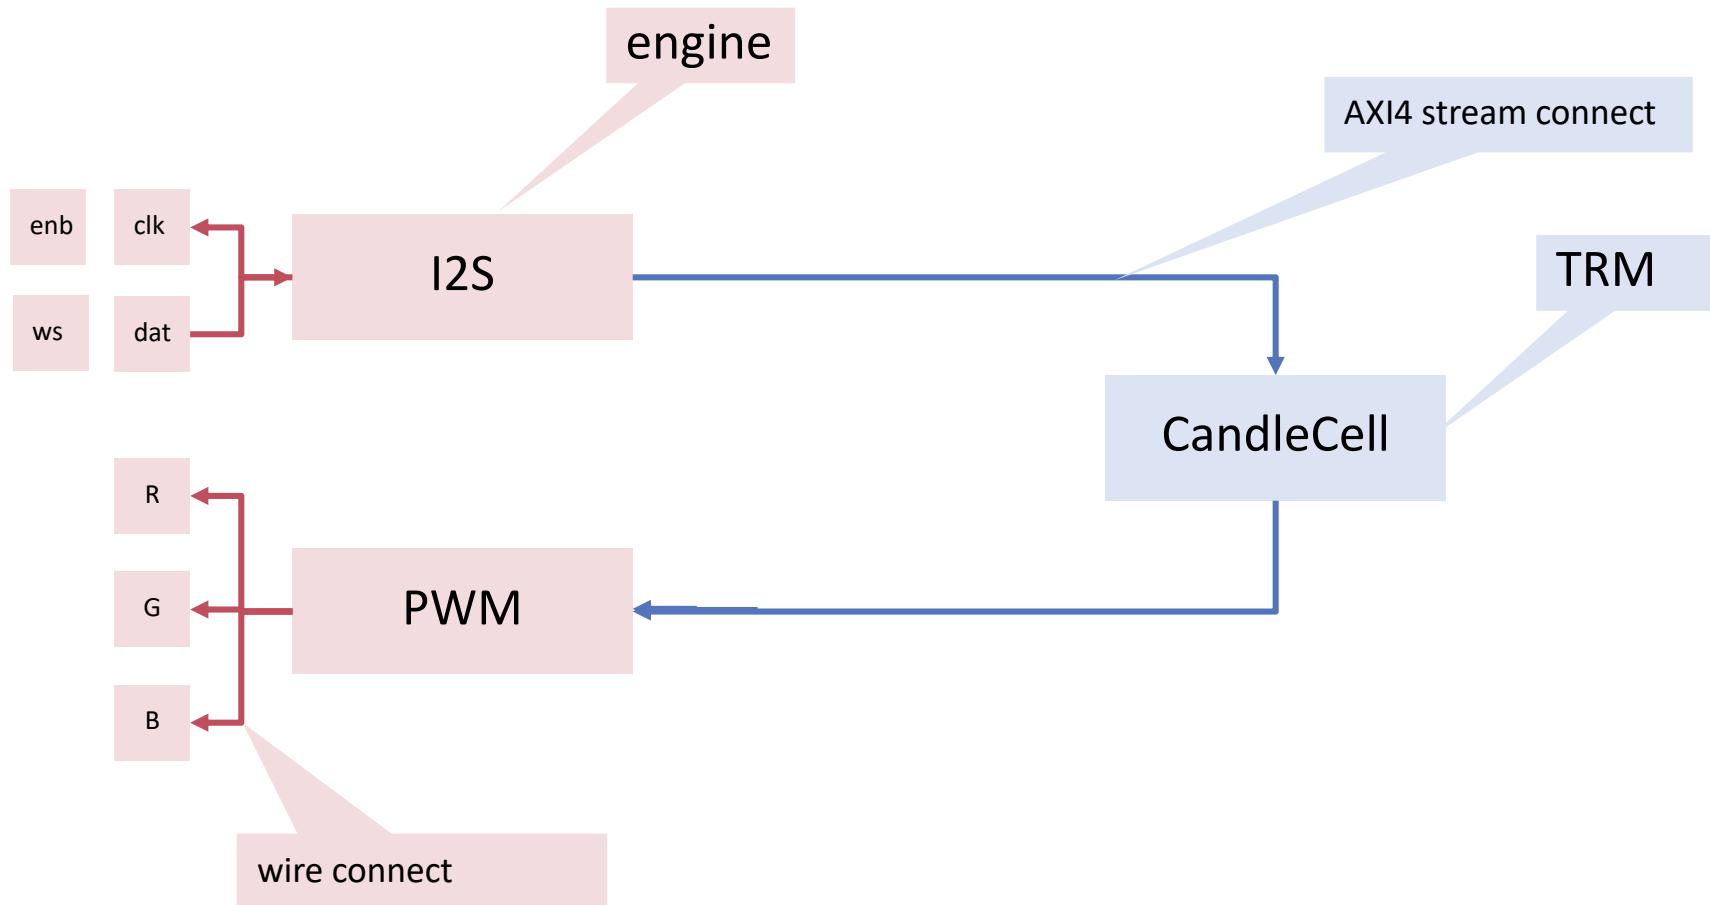
→ TRM ≈ 22% (LBs), 10% (BRAMs); Fifo ≈1% (LBs) + 7% (BRAMs)

Lab Scenario 0: PWMTester



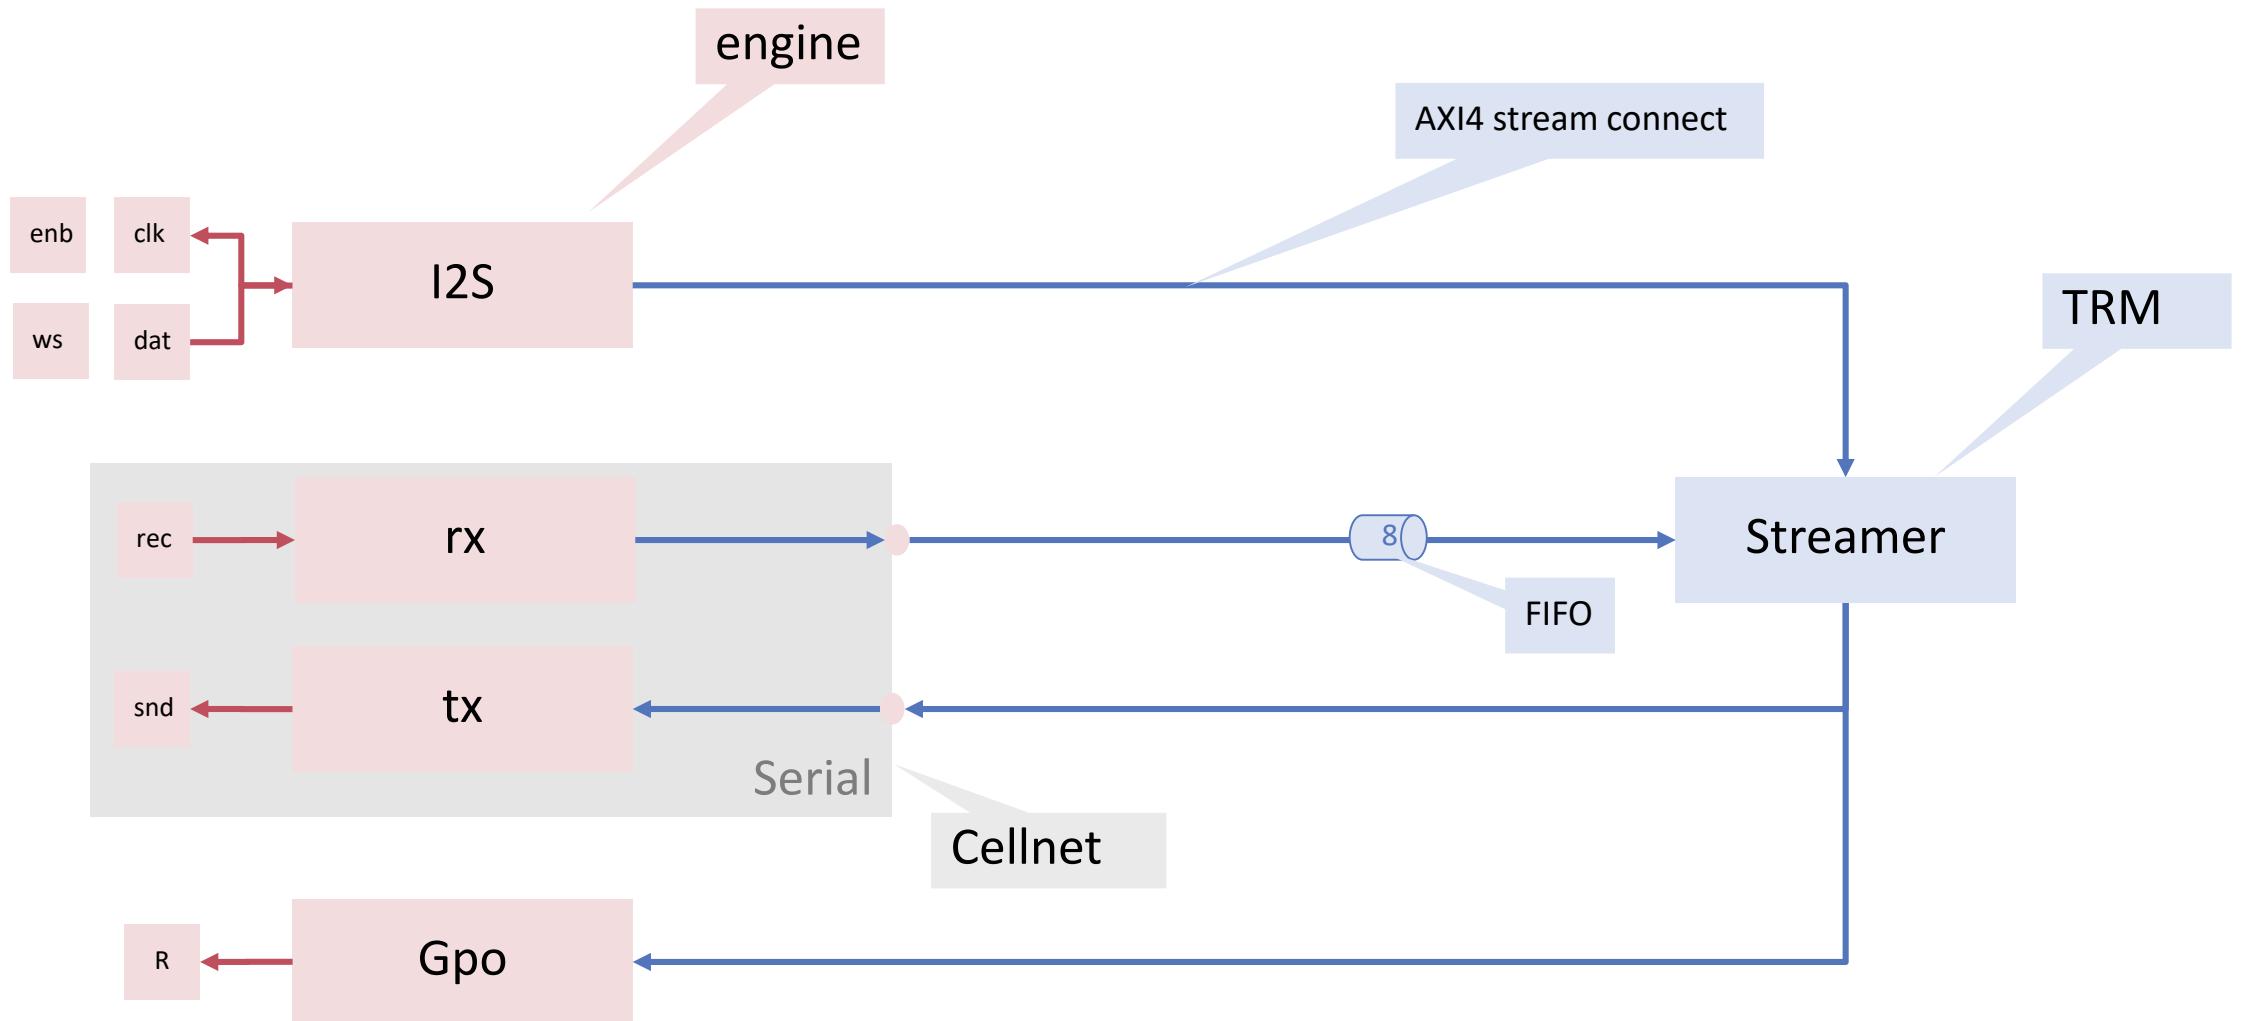
Device utilisation:
ICESTORM_LC: 1840 / 7680 23%
ICESTORM_RAM: 3 / 32 9%
369

Lab Scenario 1: Candle



Device utilisation:
ICESTORM_LC: 1949 / 7680 25%
ICESTORM_RAM: 3 / 32 9%
370

Lab Scenario 2: Sampler



Device utilisation:
ICESTORM_LC2099/ 7680 27%
ICESTORM_RAM: 4/ 32 12%
371

Lab Scenario 3:Tuner

