### Verilog to Routing (VTR): A Flexible Open-Source CAD Flow to Explore and Target Diverse FPGA Architectures

Vaughn Betz

With thanks to all VTR contributors, and particularly Kevin Murray for assistance in creating this presentation



### Agenda

- FPGAs: Basics & Utility
- Flexible CAD: Motivation and Challenges
- Verilog to Routing: Recent Progress
  - Architecture Generality
  - AIR: Faster, More Adaptive Routing
  - VTR 8.0: QoR Improvements
  - Reinforcement Learning for Adaptive Placement
- Summary



### **FPGA Basics & Utility**



### **FPGA Basics**

#### Programmable Hardware



Logic Element

#### Can implement arbitrary circuits

- Look-Up Table (LUT) implements logic
- Flip-Flop (FF) stores sequential state



### **FPGA Basics**

Logic Block (LB)

• Group of Logic Elements (LEs)





### **FPGA Basics**

Device





Mapping Applications to an FPGA

- Modern FPGAs are large (>10M LEs)
- Use Computer Aided Design (CAD) flow!
  - Designer: provides <u>behavioural</u> specification
  - Algorithms & tools: figure out the <u>details</u>





### Why Use an FPGA?

- 1. Flexible I/O
  - No off-the-shelf chip with right I/O
- 2. Power Efficiency
  - CPU/DSP/GPU performance per watt too low

45 nm CPU energy breakdown [1]



- 3. Programmability
  - Dedicated ASIC not practical (e.g. changing workload)



#### Widely Used Across Many Domains













#### CAD flow key to success

## Flexible, Open Source CAD: Motivation and Challenges

### I: FPGA Architecture Research

- Challenges:
- Need representative benchmarks
- Describe FPGA
  Architecture
- Need high quality CAD flow





### Why New FPGA Architectures?

### 1. Process Technology Changes

• Wire resistance, power limits, interposers, ...





### 2. Application Changes

- New standards, faster processing:  $4G \rightarrow 5G$ , ...
- Deep Learning  $\rightarrow$  new compute demands
- Embedded FPGAs
- 3. Innovative Architecture Ideas!







 VTR enhancements help Symbiflow (QuickLogic EOS S3, Xilinx Artix7), start-ups, architecture research, …

### **III: Program Novel Spatial Compute Architectures**

- Have an architecture? Build it!
- FPGA/FPGA+???/???
- Need a CAD flow!



Relmagine: FPGA + Imager (MIT Lincoln Lab) [1]



OpenFPGA (U of Utah) [2]

PRGA (Princeton) [3]

[1] https://www.darpa.mil/attachments/Final Compiled\_RelmagineProposersDay.pdf [2] OpenFPGA: a Complete Open Source Frameworkfor FPGA Prototyping, Open Source Design Automation, 2019 (https://sites.google.com/site/pegaillardon/research/openfpga) [3] PRGA: An Open-source Framework for Buildingand Using Custom FPGAs, Open 14 Source Design Automation, 2019 (https://prga.readthedocs.io)



## **VTR: Verilog to Routing**

### Verilog to Routing (VTR) Project



Open Source FPGA CAD Flow

Approach:

- Architecture Agnostic
- Data-driven FPGA Architecture description

16



### **VTR Challenges**





### Outline





2 AIR: Adaptive Incremental Router

<sup>3</sup>VTR 8 QoR, Run-time & CAD Enhancements



18

### Outline



### VTR 8 Capabilities & New Features

#### AIR: Adaptive Incremental Router

VTR 8 QoR, Run-time & CAD Enhancements

### Reinforcement Learning Enhanced Placement

K. E. Murray, et al., "VTR 8: High Performance CAD and Customizable FPGA Architecture Modelling", *ACM TRETS, 2020* 





Logic Block Size





Inter-Element Connectivity

- Supports both soft & hard blocks
- Arbitrary:
  - Netlist Primitives
  - Hierarchy
  - Connectivity





### **General Device Grids**







### **Arbitrary Grid Example**



22

### **Detailed Routing Architecture**





### **Customizable Routing Architectures**



### Switch-block Locations





#### Non-Configurable Edges



### **Fully Custom Routing Architectures**





- Load RR graph from file
  - VPR or non-VPR generated
- Full control of routing network
- Allows freezing RR graph (e.g. taped-out)

25

### **Analysis & Modeling**

- Run analysis independent of optimization
- Improved visualizations & reporting
- Improved area models
- Full featured timing analysis
  - Min & max timing models
  - Multi-clock primitives



Multi-clock Primitives



**Routing Utilization** 

### New & Enhanced!

Startpoint: FFC.Q[0] Endpoint : out:out1.outpad[0] Path Type : setup		
Point	Incr	Path
clock clk (rise edge)	0.000	0.000
clock source latency	0.000	0.000
clk.inpad[0] (.input)	0.000	0.000
FFC.clk[0] (.latch)	0.042	0.042
FFC.Q[0] (.latch) [clock-to-output]	0.124	0.166
out:out1.outpad[0] (.output)	0.550	0.717
data arrival time		0.717
clock virtual_io_clock (rise edge)	0.000	0.000
clock source latency	0.000	0.000
clock uncertainty	0.000	0.000
output external delay	0.000	0.000
data required time		0.000
data required time		0.000
data arrival time		-0.717
slack (VIOLATED)		-0.717

#### **Timing Reports**



Critical Path

### Outline





### 2 AIR: Adaptive Incremental Router

## VTR 8 QoR, Run-time & CAD Enhancements

#### Reinforcement Learning Enhanced Placement

K. E. Murray, et al., "AIR: A Fast but Lazy Timing-Driven FPGA Router", *ASP-DAC*, 2020, 1-7.

27

### **FPGA Routing Problem**

- Model interconnect network as Routing Resource (RR) graph:
  - Nodes: Conductors (wires/pins)
  - Edges: Configurable switches
- FPGA Routing Problem:
  - Find embedding of netlist in RR graph
  - Requires finding many non-overlapping trees in RR graph
  - Minimize timing and wiring (power)
  - Reasonable run-time
- Usually solved as single combined global/detailed routing stage



Routing Resource (RR) Graph

### **Negotiated Congestion Routing**



- Allow multiple nets to use same resources (*congestion*)
  - Nets negotiate for resources
  - Nets do not block each other



### **AIR: Adaptive Incremental Router**

Negotiated congestion router



### **Connection Router: Architecture Adaptation**

TotalCost(s, n, t) = PrevCost(s, n) + NodeCost(n) + ExpectedCost(n, t)



### **Adaptive Lookahead**

• Observation: Lookahead should adapt to routing architecture



### **Net Router: Lazy Routing**



AIR Routes Nets Lazily:

- Efficiently handle High-Fanout nets
- Incrementally re-route congested nets

### High-Fanout (HF) Routing

- Observation: Most routing of HF net irrelevant for a particular sink
  - But still consider as potential branch points (expensive)
- Optimization:
  - Only consider branch points spatially *nearby* target
  - Don't even look at others
  - Limit search to region near target



### **HF Router Expansion**



Traditional

Spatial Lookup

### **Incremental Routing**

- Observation: Most net connections routed legally
- Optimization:
  - Rip-up illegal sub-trees
  - Re-route *only* those sub-tree connections



- Challenge: May degrade critical path
  - Fix: Also rip-up delay sub-optimal connections

### **Evaluation Titan Flow & Benchmarks**



Titan Flow:

- Synth. & Tech-Map with
  Quartus
- Pack/Place/Route with Quartus or VPR

#### Titan Benchmarks:

- Target Stratix IV
- 23 designs, 90K-1.9M primitives
- Academic Routers:
  - VPR 7
  - CRoute
  - AIR
- Commercial Routers:
  - Intel Quartus 18.0

37

### **AIR Outperforms Academic Routers**



### **AIR is Faster than Quartus**





### Outline





### 2 AIR: Adaptive Incremental Router

## <sup>3</sup>VTR 8 QoR, Run-time & CAD Enhancements

#### Reinforcement Learning Enhanced Placement

K. E. Murray, et al., "VTR 8: High Performance CAD and Customizable FPGA Architecture Modelling", *In submission to ACM TRETS, 2020* 

### **Titan Benchmark Completion**

Tool	Routed Benchmarks	Unroutable Benchmarks	Device Too Small
VTR 7	14	9	-
VTR 8	23	0	-
Quartus 18.0	20	1	2



VTR 8 Titan Benchmark Implementations



### Packing & Placement Enhancements

Avoid highly interconnected (unnatural) clusters

- Avoid packing unrelated logic
- Include more indirect (e.g. transitive) connectivity measures
- De-prioritize high-fanout nets







Clustered with Unrelated Logic



Naturally Clustered

**Improved Placement Exploration** 

- Better macro/carry-chain swapping
- Improved swapping of sparse blocks





### VTR 8 Run-time & Memory Footprint



### VTR 8 QoR

#### Quartus 18.0





### Outline



![](_page_45_Picture_2.jpeg)

### 2 AIR: Adaptive Incremental Router

<sup>3</sup> VTR 8 QoR, Run-time & CAD Enhancements

### Reinforcement Learning Enhanced Placement

K. E. Murray and V. Betz, "Adaptive FPGA Placement Optimization via Reinforcement Learning", *MLCAD*, 2019

46

![](_page_45_Picture_8.jpeg)

## **Reinforcement Learning & CAD**

![](_page_46_Picture_1.jpeg)

### **CAD Tool Development: Human-in-the-loop**

Large solution space  $\rightarrow$  Use heuristics

![](_page_47_Figure_2.jpeg)

With human-in-the-loop:

•Slow

- •Simple heuristics, limited tool parameters (to keep tractable)
- •Tune for average case (can't investigate every benchmark design)

### **CAD Tool Development: Reinforcement Learning (RL)**

![](_page_48_Figure_1.jpeg)

- •Human out of the loop!
- •Learn better heuristics: exploit more information, more parameters
- •Online adaptation  $\rightarrow$  better than average case

![](_page_48_Picture_5.jpeg)

### **FPGA** Placement

![](_page_49_Picture_1.jpeg)

### **Simulated Annealing (SA) Placement**

- Modify placement by making 'moves'
- Accept/Reject move based on:
- Cost change
- Temperature (hill climbing)

![](_page_50_Figure_5.jpeg)

![](_page_50_Figure_6.jpeg)

![](_page_50_Picture_7.jpeg)

### **Move Generation**

Many possible types of moves!

![](_page_51_Picture_2.jpeg)

### **RL Move Generator**

- Actions: moved different block types Reward:
  - Accepted: -Δcost
  - Rejected: 0
- Agent:
  - Estimates value of actions
  - Selects action to take

![](_page_52_Figure_7.jpeg)

![](_page_52_Picture_8.jpeg)

Logic

### **Estimating Action Values**

Values of action are not stationary!

![](_page_53_Figure_2.jpeg)

### **Action Selection: Exploration vs Exploitation**

- ε-greedy: Mostly greedy (exploit), occasionally random (explore)
- ε: fraction of exploratory moves

![](_page_54_Figure_3.jpeg)

• VTR Benchmarks (10K-165K primitives), 3 seeds

![](_page_55_Figure_2.jpeg)

VTR Benchmarks (10K-165K primitives), 3 seeds

![](_page_56_Figure_2.jpeg)

VTR Benchmarks (10K-165K primitives), 3 seeds

![](_page_57_Figure_2.jpeg)

VTR Benchmarks (10K-165K primitives), 3 seeds

![](_page_58_Figure_2.jpeg)

### Conclusion

- RL-enhanced Simulated Annealing based FPGA Placer
  - RL agent controlled move generator
  - Learns on-line what types of moves are productive
  - Improves run-time/quality trade-offs
    - Particularly at low run-times
  - Much more to explore
    - More diverse placement perturbations / moves
    - More powerful agents

![](_page_59_Picture_9.jpeg)

![](_page_60_Picture_0.jpeg)

### Summary

- Open-source, flexible CAD crucial for
  - Evolution/innovation in spatial hardware
  - Enabling CAD research by the community
- VTR
  - Large, very active open-source effort
    - >12,000 commits, ~100 committers
  - Linked to other open-source efforts like Symbiflow, OpenFPGA
    - Mix and match your flow
  - Flexible, with good (but can be better!) result quality
  - Contributions welcome!

![](_page_61_Picture_11.jpeg)

### **Open source enables innovation across a wider community**

# Thanks!Questions?<br/>Email: vaughn@eecg.utoronto.ca

### **VTR:** verilogtorouting.org

![](_page_62_Picture_2.jpeg)

![](_page_62_Picture_3.jpeg)

![](_page_62_Picture_4.jpeg)

![](_page_62_Picture_5.jpeg)