

Timing-aware routing in the RapidWright framework

Leo Liu, Nachiket Kapre
leo.liu@uwaterloo.ca, nachiket@uwaterloo.ca



Background

RapidWright: Open source project for accessing low-level resources for Xilinx FPGAs

Advantage: design generation without FPGA CAD tools

Weakness: no timing knowledge of FPGA resources; hard to build timing-driven tools


Problem Statement

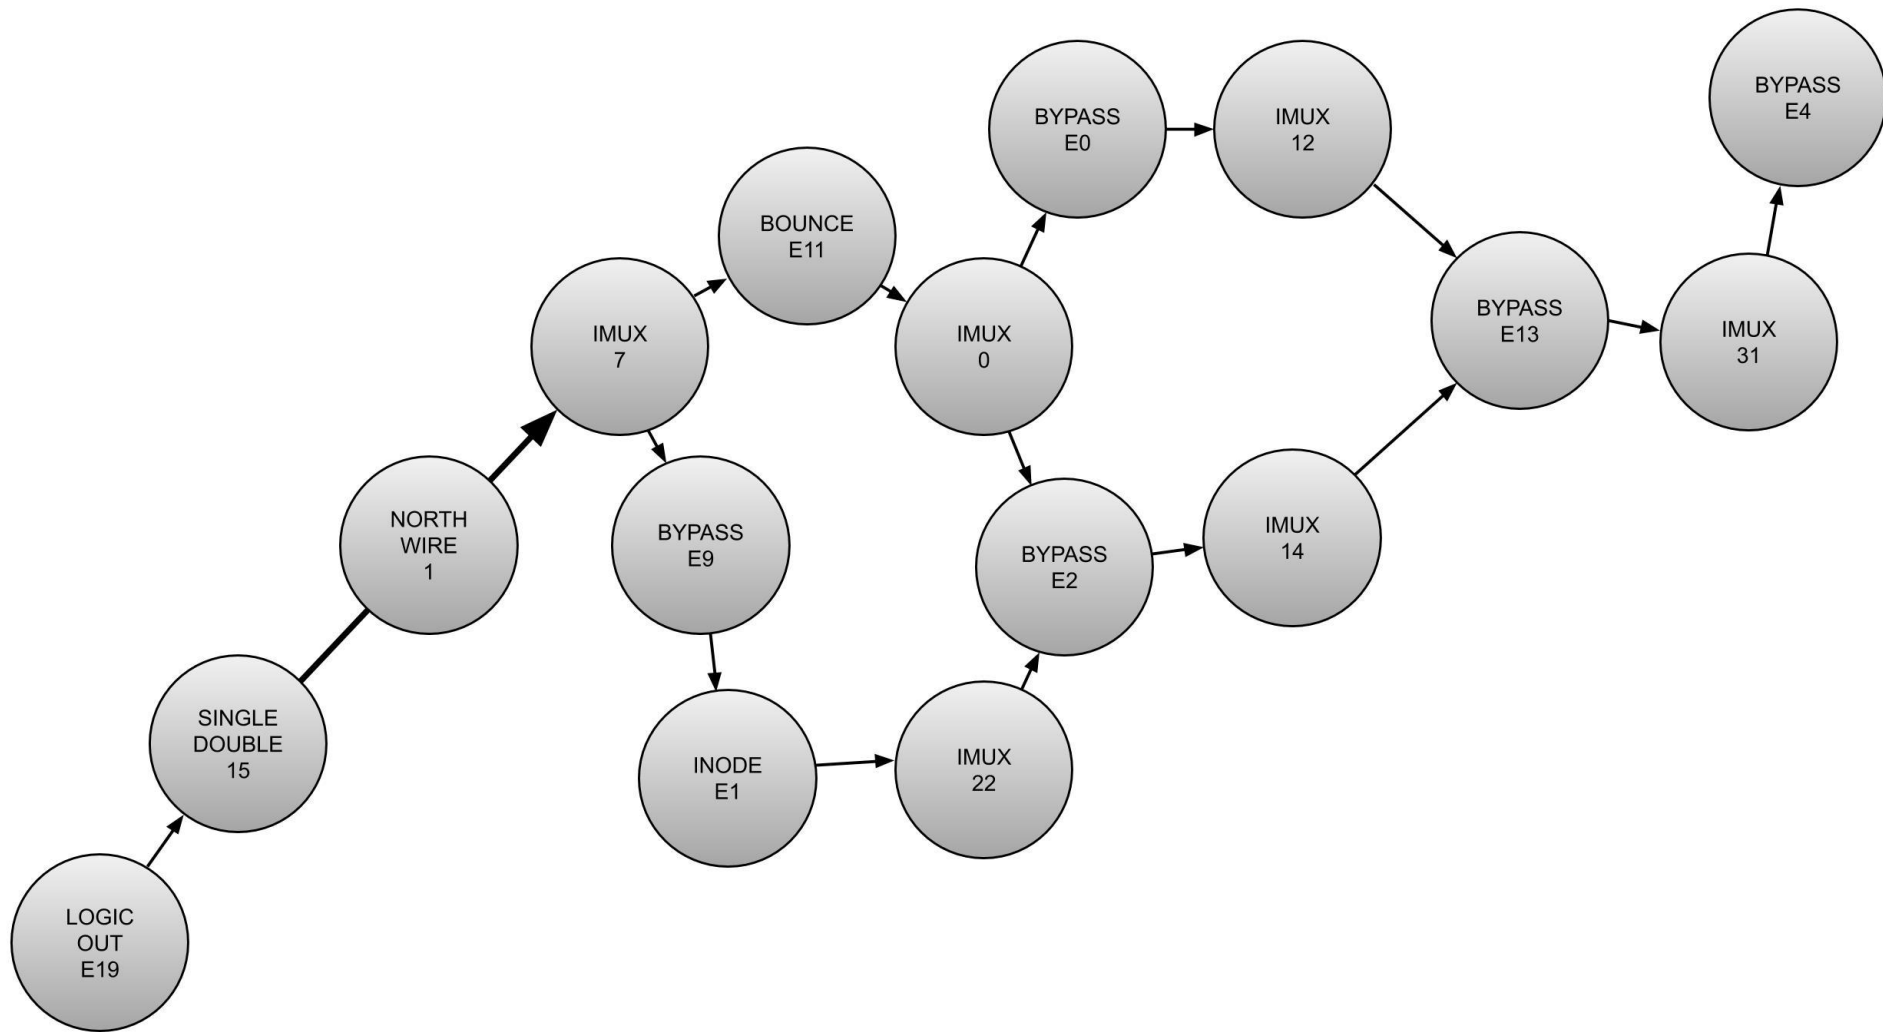
We used RapidWright to design **RapidRoute**, a fast router for building communication networks

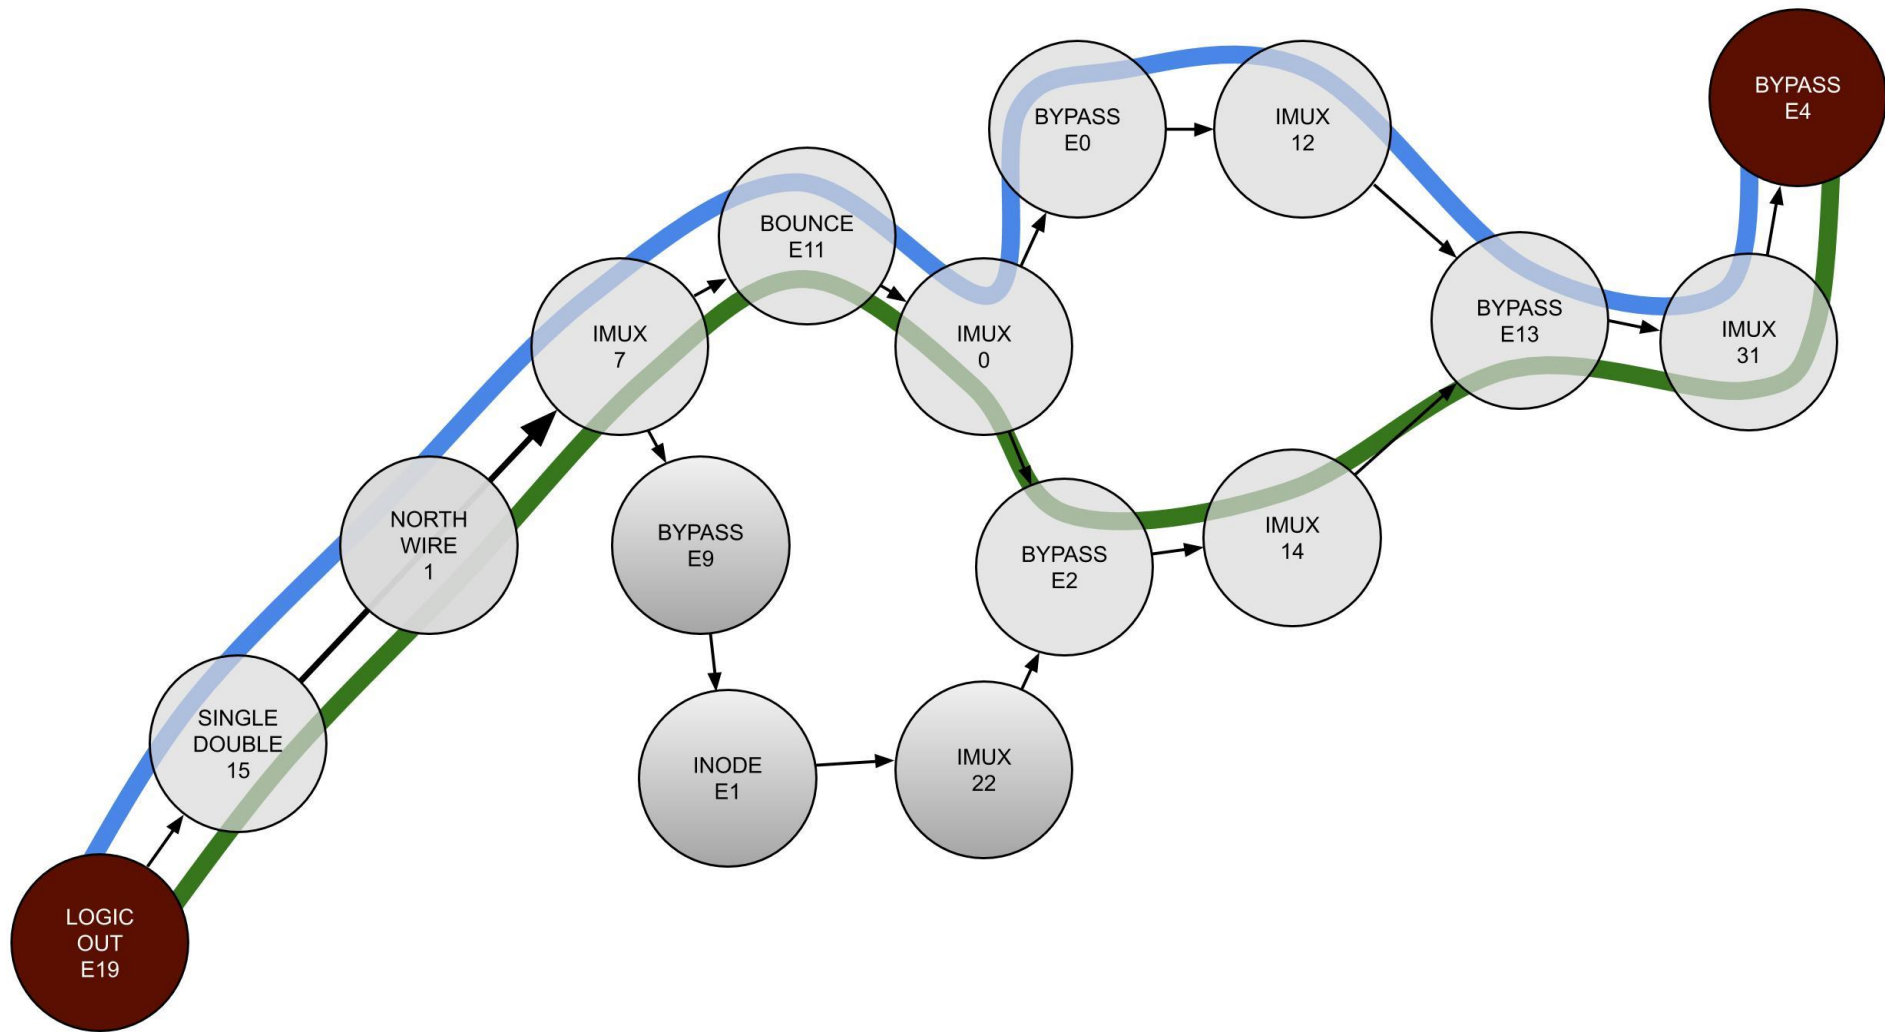
Problem: RapidWright does not allow RapidRoute to be timing-driven:

- Routing algorithms cannot optimize for shortest path
- Consistently loses to Vivado


Timing Slack







■ BLUE:
11 nodes
0.887ns



■ GREEN:
11 nodes
0.815ns

Solution

Build our own timing library

Main ideas:

- Extract relevant timing data using both RapidWright and Vivado
- Integrate extracted data into RapidRoute algorithm

Key Claim

We can extract fine-grained timing information
of Xilinx FPGA routing resources

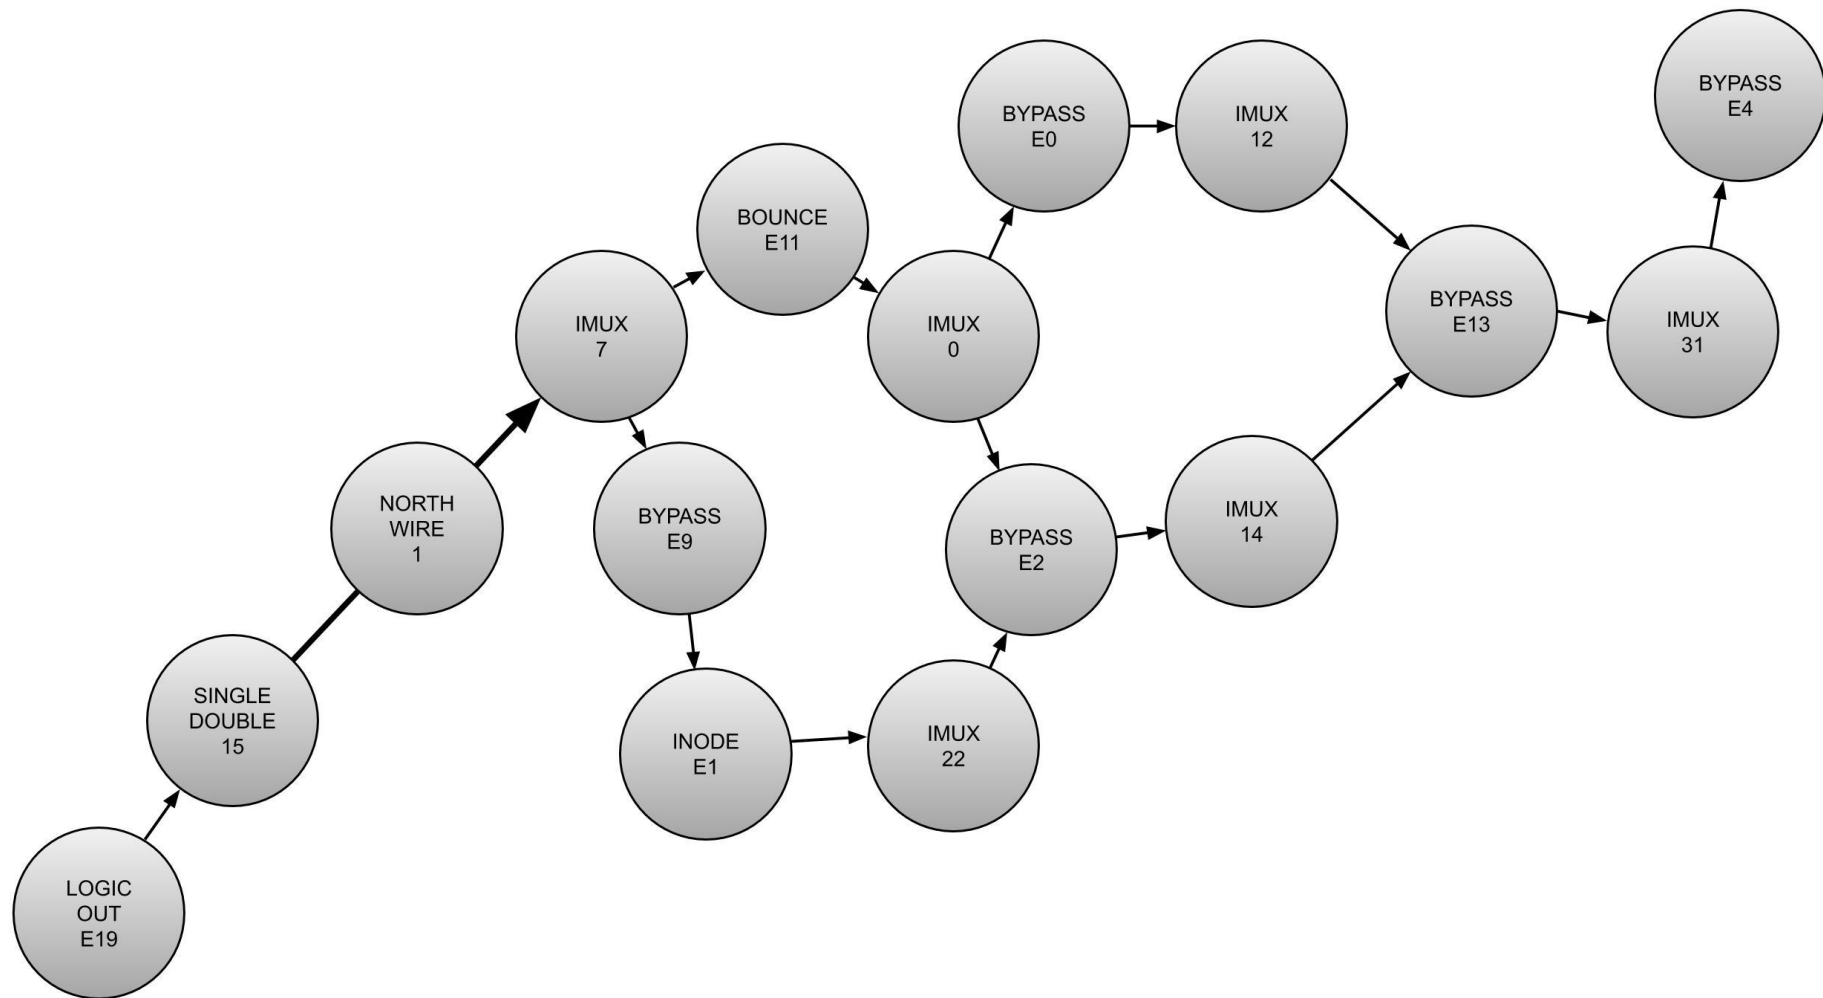
- Library allows RapidRoute to **match Vivado performance**
- **Less than 10 mins** of one-time analysis
- **Extremely lightweight**
 - RapidRoute retains its routing speed
 - Low memory overhead

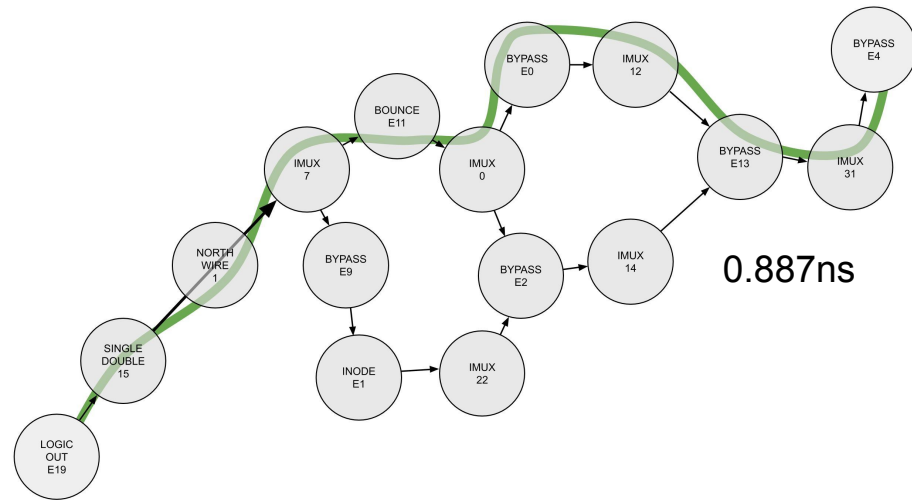
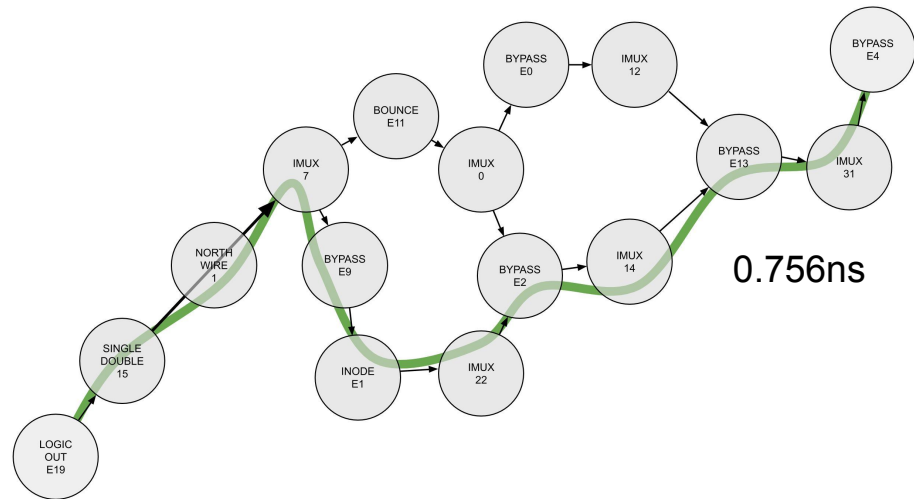
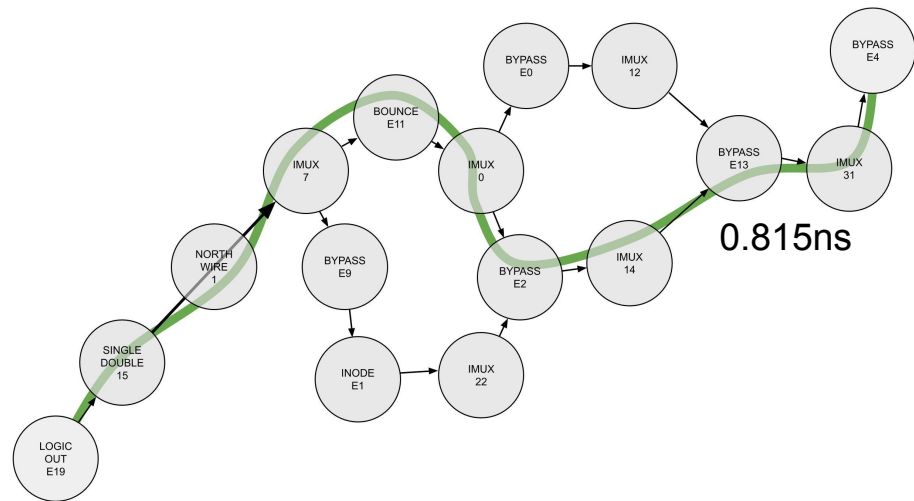
Main Approach

1. Build many calibration designs with RapidWright
2. Load designs in Vivado for timing feedback
3. Organize into a linear system

$$\begin{array}{c} \text{node usage for each design} \\ \left[\begin{array}{cccccc} 1 & 2 & 5 & 0 & 0 & \dots & 1 \\ 5 & 3 & 0 & 3 & 1 & \dots & 5 \\ \vdots & \vdots & \vdots & \vdots & \vdots & & \vdots \\ 0 & 4 & 1 & 1 & 0 & \dots & 0 \end{array} \right] \begin{bmatrix} t_1 \\ t_2 \\ \vdots \\ t_m \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_m \end{bmatrix} \end{array}$$

m designs Vivado path delays (ns)





$$\begin{aligned} &\text{LOGIC_OUT_E19} + \text{SNG_DBL_15} + \text{NN1} + \text{IMUX_7} \\ &\quad + \text{BNCE_E11} + \dots = \quad 0.815\text{ns} \end{aligned}$$


$$\begin{aligned} &\text{LOGIC_OUT_E19} + \text{SNG_DBL_15} + \text{NN1} + \text{IMUX_7} \\ &\quad + \text{BNCE_E11} + \dots = \quad 0.887\text{ns} \end{aligned}$$

$$\begin{aligned} &\text{LOGIC_OUT_E19} + \text{SNG_DBL_15} + \text{NN1} + \text{IMUX_7} \\ &\quad + \text{BYPASS_E9} + \dots = \quad 0.756\text{ns} \end{aligned}$$

Opportunities in Timing Extraction

Symmetry: symmetrical elements on FPGA have similar timing characteristics

Narrow usage: RapidRoute only targets communication overlays



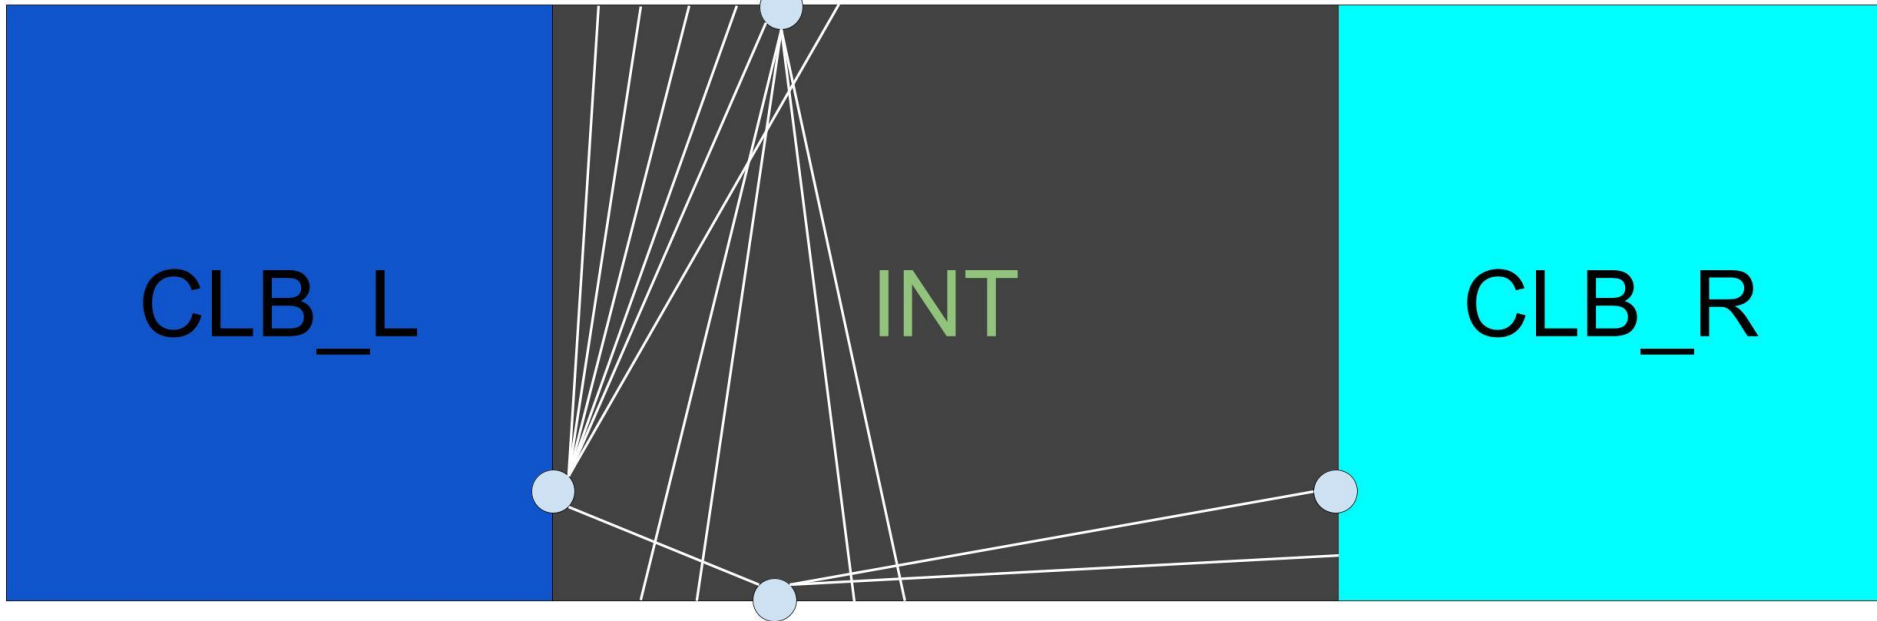
[illegible]

Node

CLB_L

INT

CLB_R



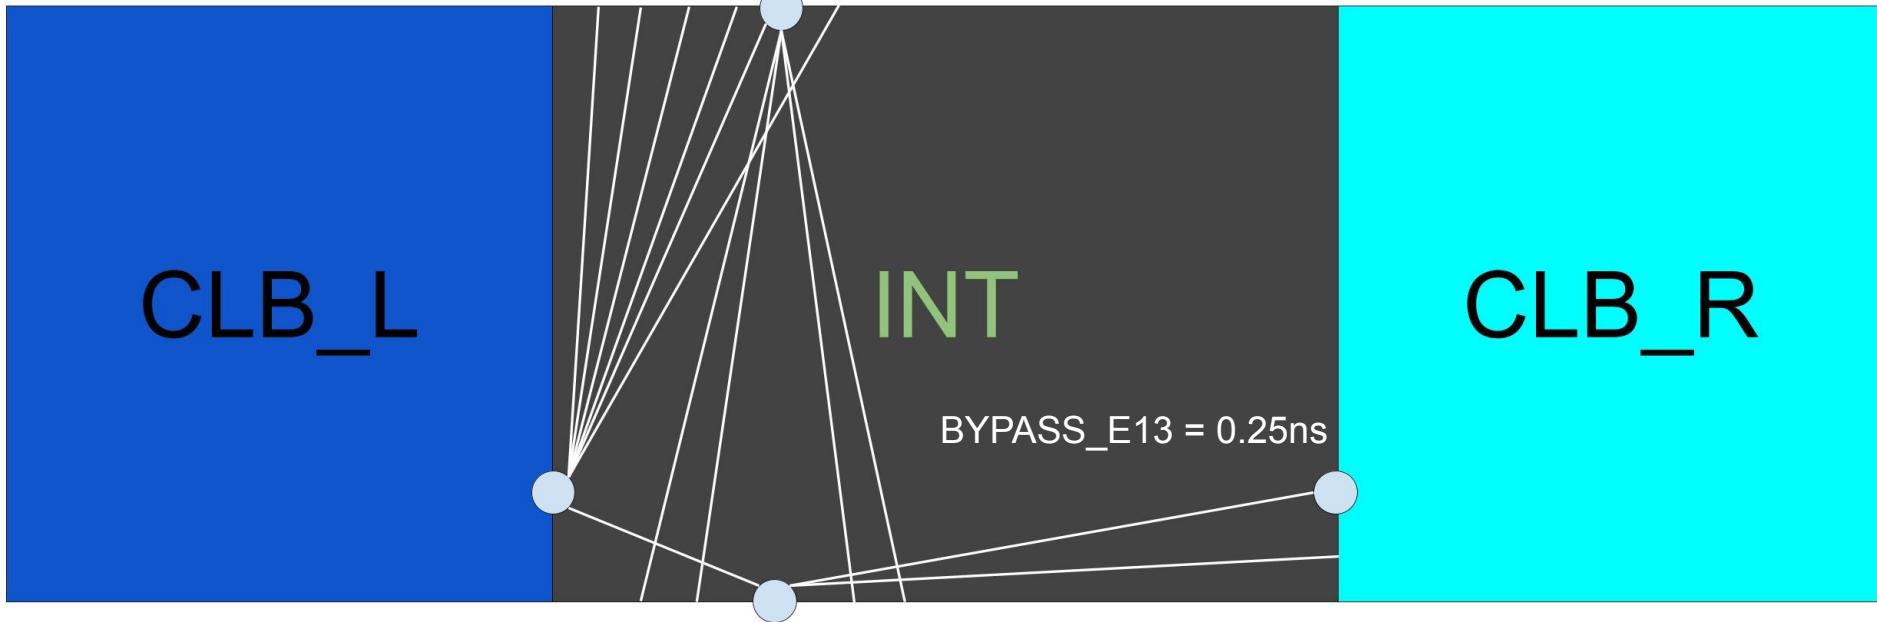
Node

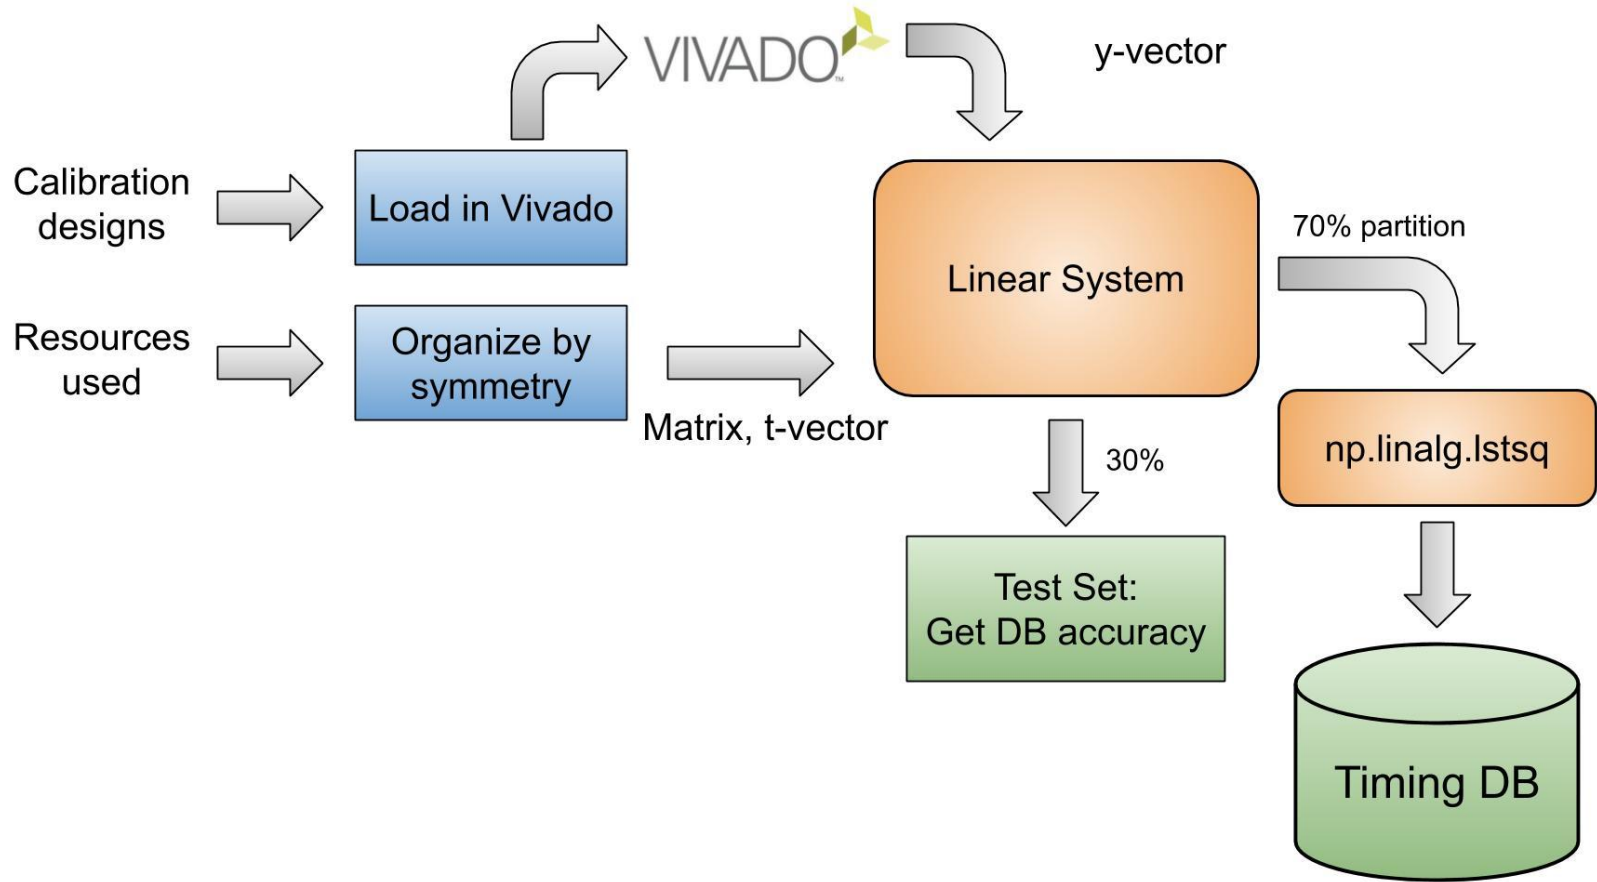
CLB_L

INT

CLB_R

BYPASS_E13 = 0.25ns

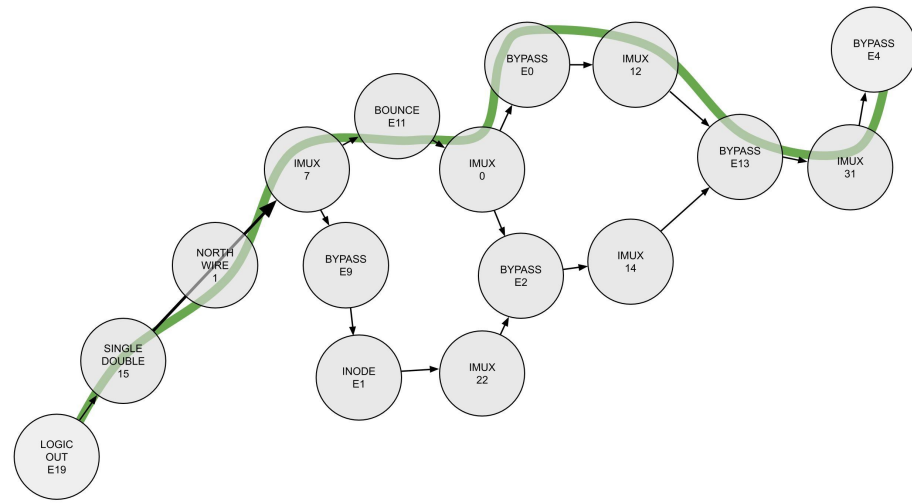
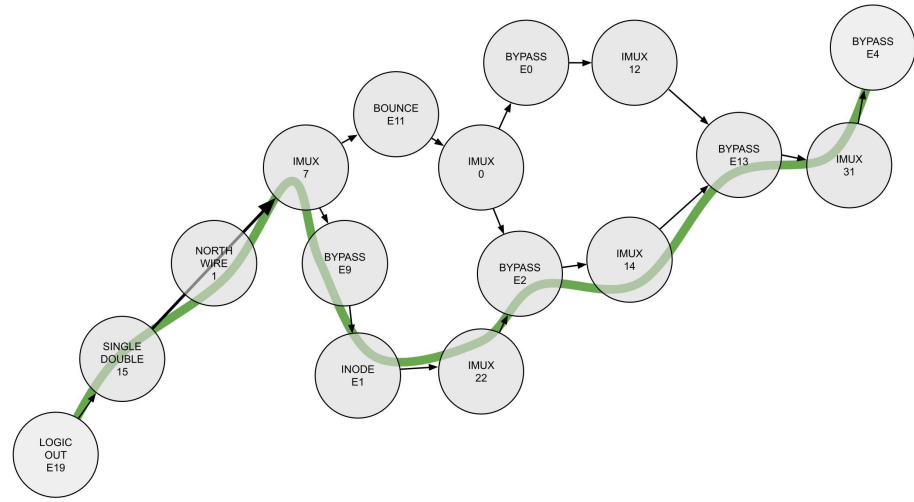
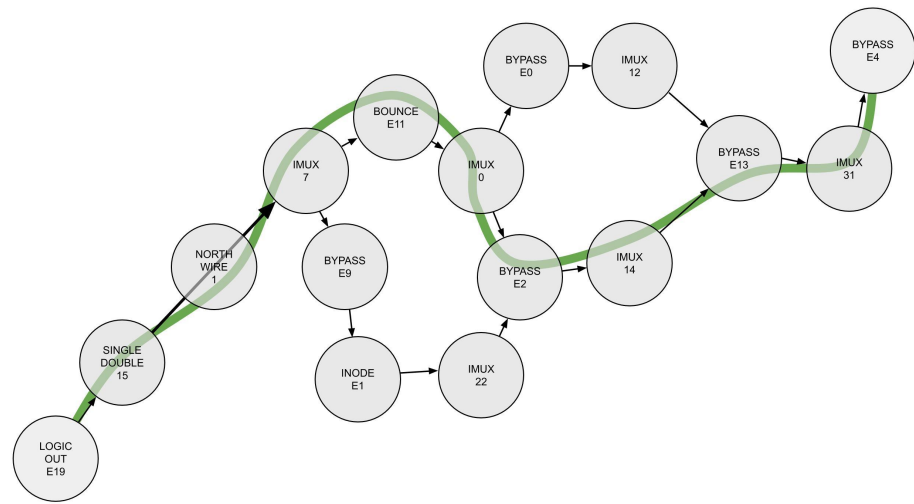


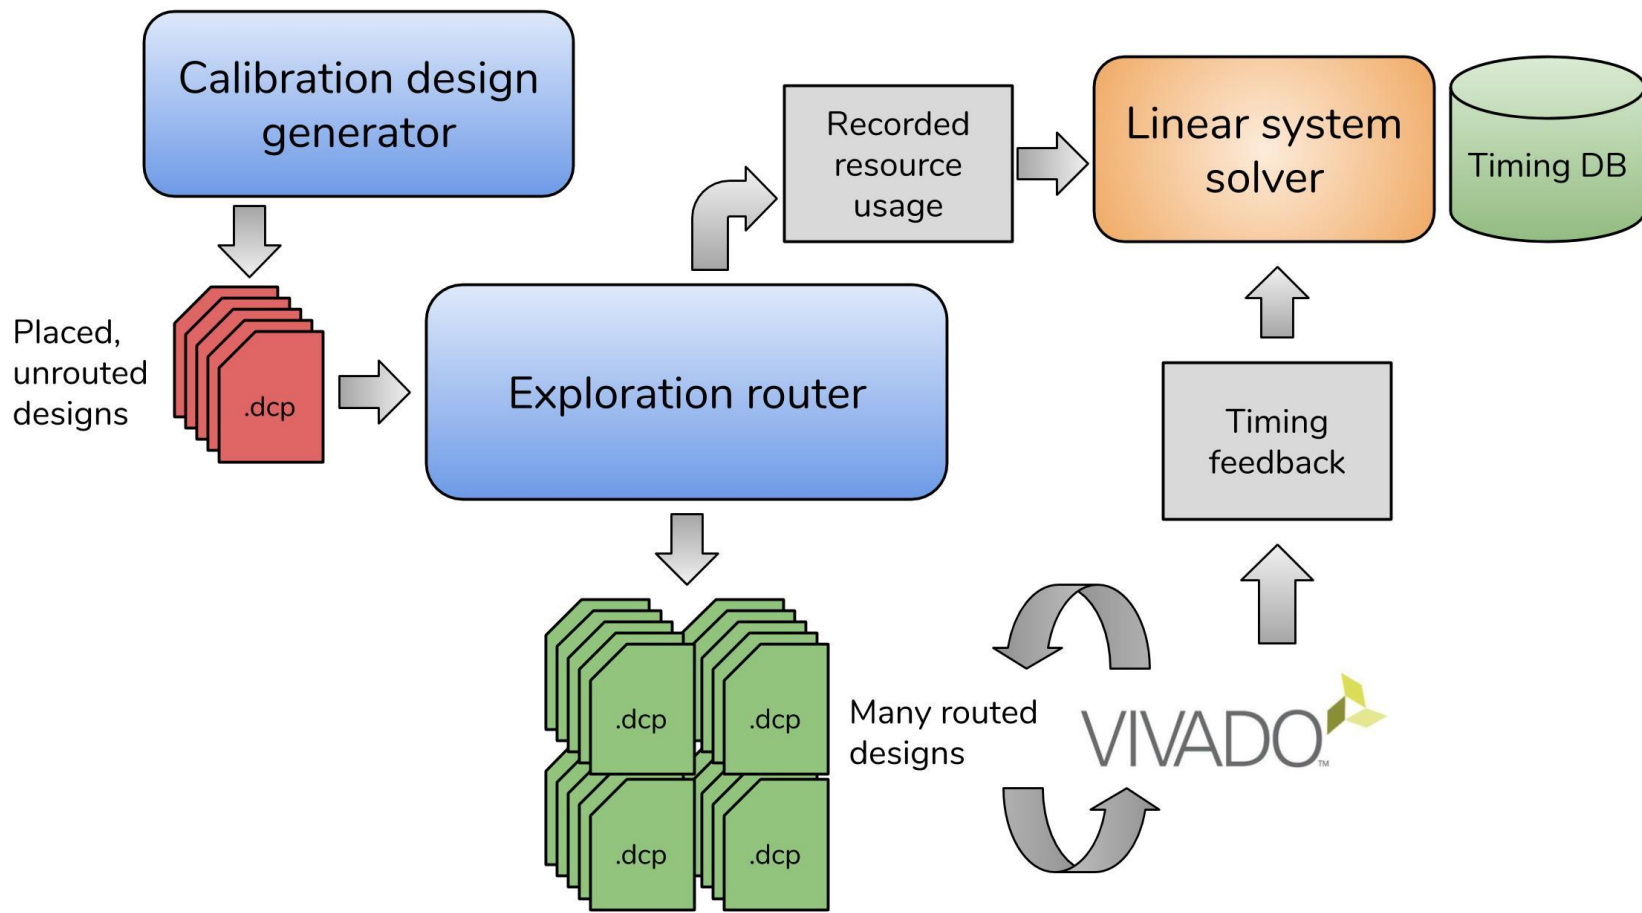


Calibration Designs

Designs: **1-bit signal**, with **changing start and end nodes**

1. For each design, **route in different ways**
2. Write out each routing result into DCP
3. Track which nodes are used for each route





Experimental Setup

- Metrics:
 - **Timing prediction accuracy**
- Calibration designs
 - **Single-bit routes of arbitrary displacement**
 - **Various devices and speed grades**
- Compare methods:
 - Partition **70% training, 30% testing** of all calibration runs

Experimental Setup

- Devices:
 - Ultrascale XCKU115 (-3, -2, -1 speed grades)
 - Ultrascale+ XCKU5P (-3, -2, -1 speed grades)
- Vivado: 2018.3
- RapidWright: 2018.3.3-beta
- Hardware: Intel Xeon E5-1630


Timing Accuracy

Measuring accuracy:

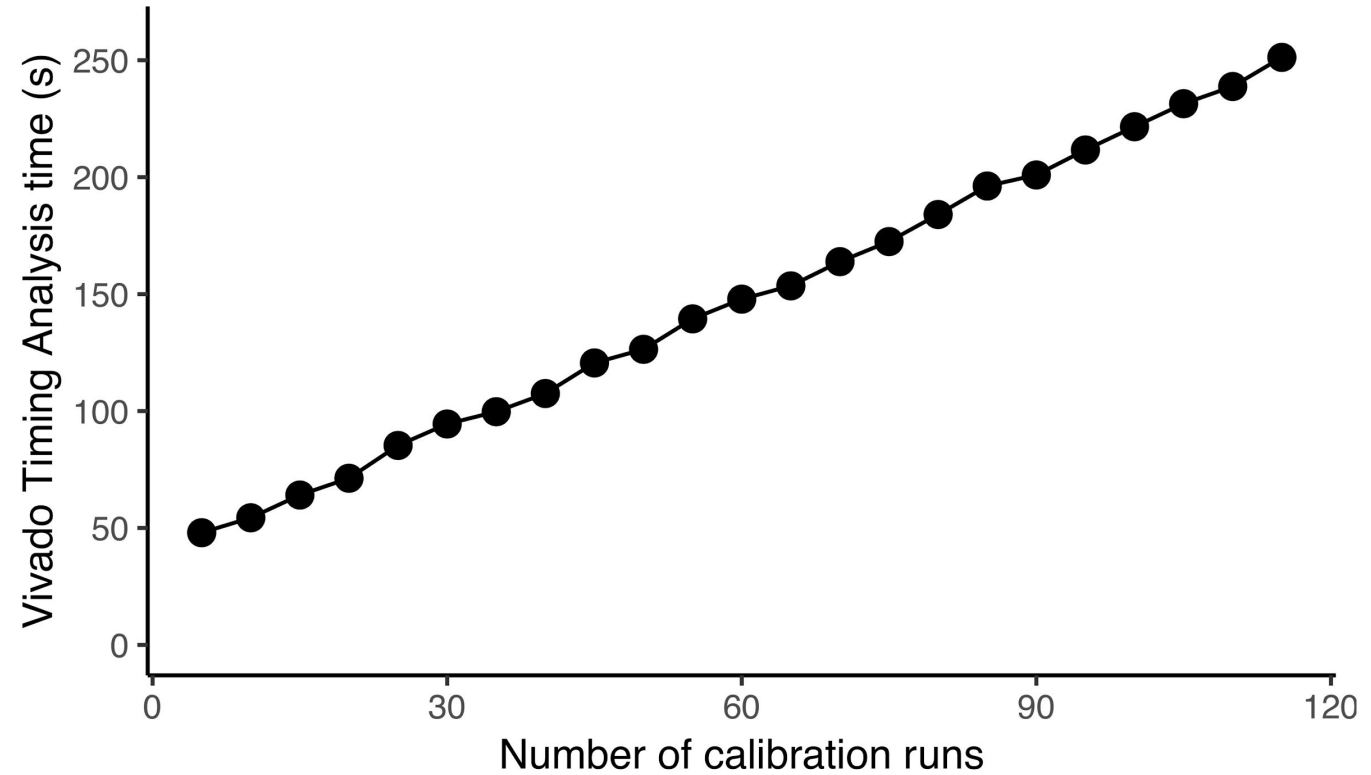
We check datapath
prediction errors of 30%
partition.


X-axis: size of 70% partition

Y-axis: average prediction
error



Vivado Runtime





Additional Notes

- Output timing database is extremely small (< 100KB)
- Majority of timing extraction solver runtime is due to Vivado query wait times

Integrating with RapidRoute

- RapidRoute accepts timing database file(s) as input, overwriting default heuristic
- Heuristic has nearly identical computing cost as default heuristic


Experimental Setup

- Metrics:
 - **Timing performance**
 - **Routing runtimes**
- Communication structures:
 - **1D rings, 2D torii, 2D meshes**
- Compare methods:
 - **RapidRoute default, RapidRoute+Timing, Vivado**

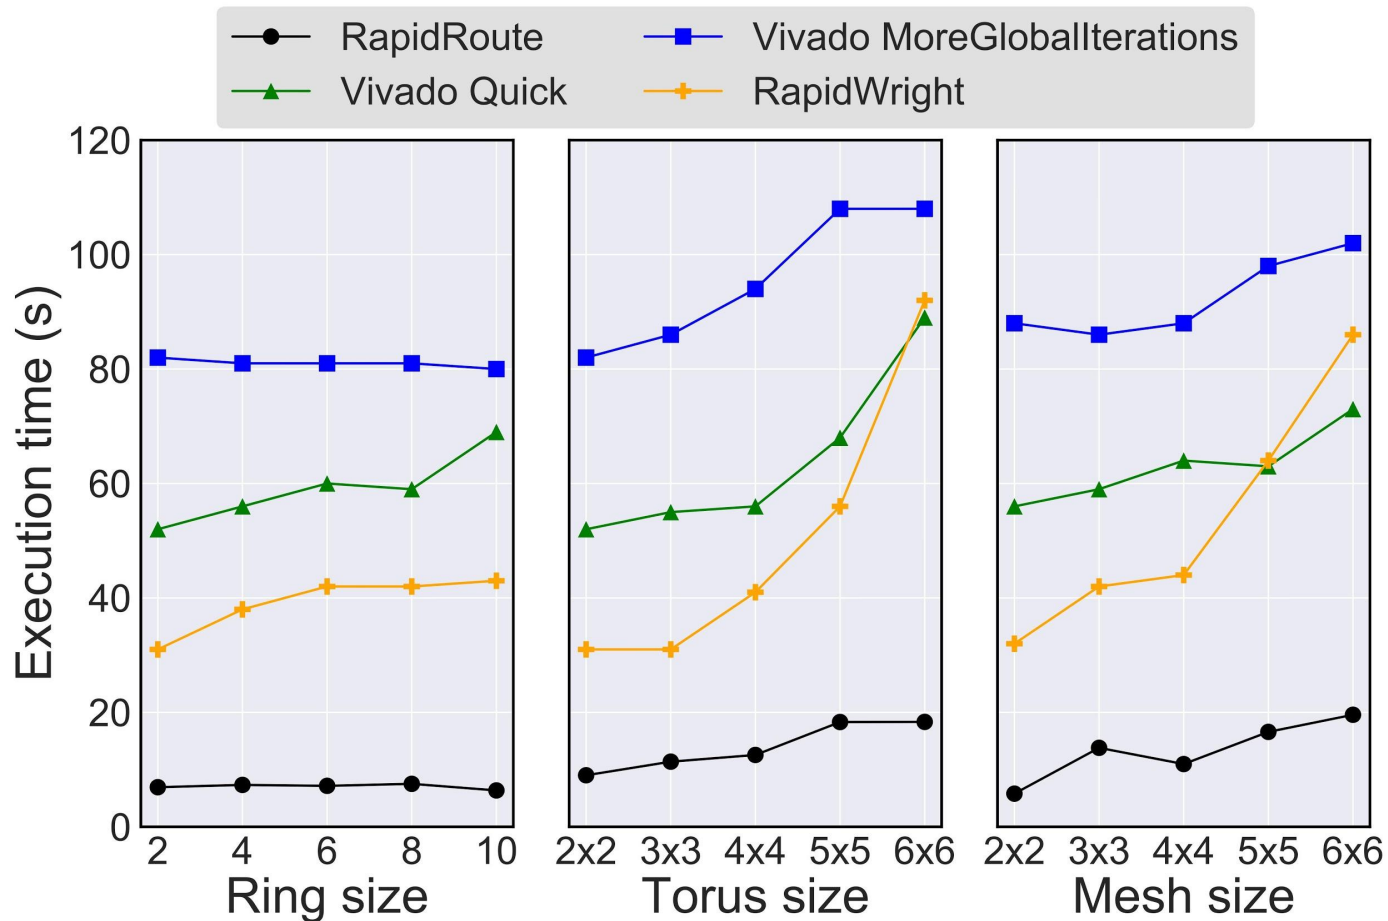
Experimental Setup

- Device: Ultrascale XCKU115 `xcku115-flva1517-3-e`
- Vivado: 2018.3
- RapidWright: 2018.3.3-beta
- Hardware: 32-core 2.6GHz Intel Xeon

Timing Results



Routing Runtime



Conclusion

- We developed a **timing extraction tool**, which is **light-weight** and **highly-accurate**
- Timing results expected to be within 1% error margin
- Total calibration phase takes minutes
- Extremely lightweight output and usage

Improved RapidRoute

- RapidRoute retains a **5-8x** routing speed advantage over Vivado
- RapidRoute now gains **competitive timing performance** on communication overlay designs