

2023 IEEE 15th International Conference on ASIC

OpenPARF: An Open-Source Placement and Routing Framework for Large-Scale Heterogeneous FPGAs with Deep Learning Toolkit

Jing Mai¹, Jiarui Wang¹, Zhixiong Di³, Guojie Luo¹, Yun Liang¹, Yibo Lin¹

¹Peking University

²Southwest Jiaotong University

jingmai@pku.edu.cn



北京大学
PEKING UNIVERSITY



西南交通大学
Southwest Jiaotong University


October 26, 2023

- ① Introduction
- ② The OpenPARF Framework
- ③ Experimental Results
- ④ Conclusion & Future Work

Introduction

HDL

- ▶ Hardware design is modeled in a *Hardware Description Language* (HDL)

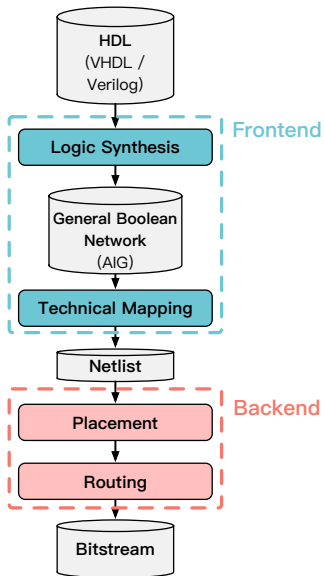


HDL

- ▶ Hardware design is modeled in a *Hardware Description Language* (HDL)

Frontend

- ▶ A FPGA "compiler" (synthesis tool) translates the HDL into a general Boolean network, e.g. *And-Inverter Graph* (AIG)




HDL

- ▶ Hardware design is modeled in a *Hardware Description Language* (HDL)

Frontend

- ▶ A FPGA "compiler" (synthesis tool) translates the HDL into a general Boolean network, e.g, *And-Inverter Graph* (AIG)
- ▶ which is then **mapped** to a FPGA technology tailored netlist



HDL


- ▶ Hardware design is modeled in a *Hardware Description Language* (HDL)

Frontend

- ▶ A FPGA "compiler" (synthesis tool) translates the HDL into a general Boolean network, e.g, *And-Inverter Graph* (AIG)
- ▶ which is then **mapped** to a FPGA technology tailored netlist

Backend

- ▶ The netlist components are **placed** on the FPGA layout



HDL

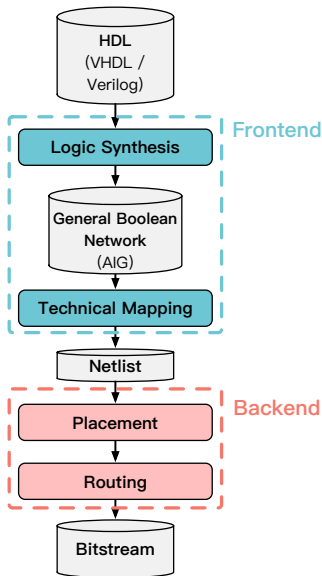
- ▶ Hardware design is modeled in a *Hardware Description Language* (HDL)

Frontend

- ▶ A FPGA "compiler" (synthesis tool) translates the HDL into a general Boolean network, e.g, *And-Inverter Graph* (AIG)
- ▶ which is then **mapped** to a FPGA technology tailored netlist

Backend

- ▶ The netlist components are **placed** on the FPGA layout
- ▶ and the connecting signals are **routed** through the interconnection network



HDL


- ▶ Hardware design is modeled in a *Hardware Description Language* (HDL)

Frontend

- ▶ A FPGA "compiler" (synthesis tool) translates the HDL into a general Boolean network, e.g, *And-Inverter Graph* (AIG)
- ▶ which is then **mapped** to a FPGA technology tailored netlist

Backend

- ▶ The netlist components are **placed** on the FPGA layout
- ▶ and the connecting signals are **routed** through the interconnection network
- ▶ A bitstream is finally generated for the FPGA configuration



FPGA P&R: the CORE of the FPGA backend CAD flow

- ▶ **Placement** significantly **determines** the final routability and timing performance¹

¹Shih-Chun Chen and Yao-Wen Chang (2017). “FPGA placement and routing”. In: *Proc. ICCAD*, pp. 914–921.

²Kevin E Murray et al. (2015). “Timing-driven titan: Enabling large benchmarks and exploring the gap between academic and commercial CAD”. In: *ACM TRETS 8.2*, pp. 1–18.

FPGA P&R: the CORE of the FPGA backend CAD flow


- ▶ **Placement** significantly **determines** the final routability and timing performance¹
- ▶ **Routing** is generally **the most time-consuming** step, accounting for 41-86% runtime²

¹Shih-Chun Chen and Yao-Wen Chang (2017). “FPGA placement and routing”. In: *Proc. ICCAD*, pp. 914–921.

²Kevin E Murray et al. (2015). “Timing-driven titan: Enabling large benchmarks and exploring the gap between academic and commercial CAD”. In: *ACM TRETS* 8.2, pp. 1–18.

FPGA P&R: the CORE of the FPGA backend CAD flow

- ▶ **Placement** significantly **determines** the final routability and timing performance¹
- ▶ **Routing** is generally **the most time-consuming** step, accounting for 41-86% runtime²
- ▶ FPGA P&R is deeply tied to the hardware architecture,
- ▶ and every FPGA manufacturer needs **tailored** P&R software → **Source of advantage!**



¹Shih-Chun Chen and Yao-Wen Chang (2017). “FPGA placement and routing”. In: *Proc. ICCAD*, pp. 914–921.

²Kevin E Murray et al. (2015). “Timing-driven titan: Enabling large benchmarks and exploring the gap between academic and commercial CAD”. In: *ACM TRETS 8.2*, pp. 1–18.



- ▶ We propose **OpenPARF**, an open-source academic FPGA P&R engine that supports complex industrial FPGA architectures with *state-of-the-art* (SOTA) placement and routing algorithms.



- ▶ We propose **OpenPARF**, an open-source academic FPGA P&R engine that supports complex industrial FPGA architectures with *state-of-the-art* (SOTA) placement and routing algorithms.
- ▶ We implement it with the deep learning toolkit `PyTorch`, running on both CPU and GPU platforms with highly flexibility and efficiency.

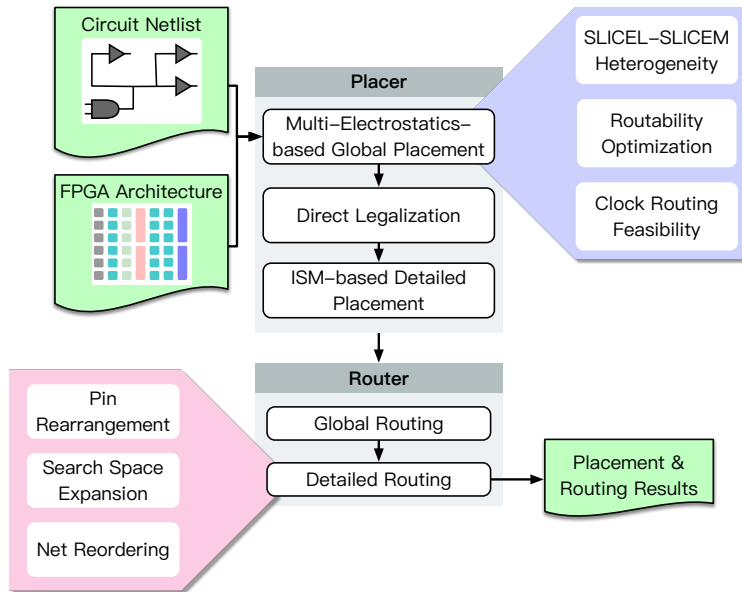


- ▶ We propose **OpenPARF**, an open-source academic FPGA P&R engine that supports complex industrial FPGA architectures with *state-of-the-art* (SOTA) placement and routing algorithms.
- ▶ We implement it with the deep learning toolkit `PyTorch`, running on both CPU and GPU platforms with highly flexibility and efficiency.
- ▶ We are capable of achieving superior placement results under various constraints such as **routability**, **clock feasibility**, and **SLICEL-SLICEM heterogeneity**



- ▶ We propose **OpenPARF**, an open-source academic FPGA P&R engine that supports complex industrial FPGA architectures with *state-of-the-art* (SOTA) placement and routing algorithms.
- ▶ We implement it with the deep learning toolkit `PyTorch`, running on both CPU and GPU platforms with highly flexibility and efficiency.
- ▶ We are capable of achieving superior placement results under various constraints such as **routability**, **clock feasibility**, and **SLICEL-SLICEM heterogeneity**
- ▶ We can reduce **0.4-12.7%** routed wirelength as well as more than **2×** speedup in placement efficiency compared with other SOTA academic P&R engines.

The OpenPARE Framework

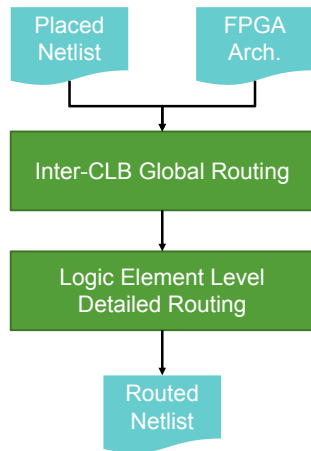


State-of-the-art P&R Algorithms

- ▶ **SOTA** multi-electrostatics-based global placement

State-of-the-art P&R Algorithms

- ▶ **SOTA** multi-electrostatics-based global placement
- ▶ **SOTA** two-stage CLB-level FPGA routing




State-of-the-art P&R Algorithms

- ▶ **SOTA** multi-electrostatics-based global placement
- ▶ **SOTA** two-stage CLB-level FPGA routing

More P&R Constraints

- ▶ SLICEL-SLICEM heterogeneity




State-of-the-art P&R Algorithms

- ▶ **SOTA** multi-electrostatics-based global placement
- ▶ **SOTA** two-stage CLB-level FPGA routing

More P&R Constraints

- ▶ SLICEL-SLICEM heterogeneity
- ▶ Routability optimization




State-of-the-art P&R Algorithms

- ▶ **SOTA** multi-electrostatics-based global placement
- ▶ **SOTA** two-stage CLB-level FPGA routing

More P&R Constraints

- ▶ SLICEL-SLICEM heterogeneity
- ▶ Routability optimization
- ▶ Clock routing feasibility
- ▶ ...



Constrained Optimization Formulation

- ▶ Nonlinear placers minimize wirelength

$$\min_{\mathbf{x}, \mathbf{y}} \tilde{W}(\mathbf{x}, \mathbf{y}) = \sum_{e \in E} \text{WL}_e(\mathbf{x}, \mathbf{y}), \quad (1)$$

- ▶ ... while subject to ePlace-series [Lu+, TCAD'15] density constraints for each object type $s \in S = \{\text{LUT}, \text{FF}, \text{DSP}, \text{BRAM}, \text{IO}\}$

$$\text{s.t. } \Phi_s(\mathbf{x}, \mathbf{y}) = 0, \quad \forall s \in S, \quad (2)$$

Augmented Lagrangian Method

- ▶ Constrained \rightarrow unconstrained³

$$\min_{\mathbf{x}, \mathbf{y}} \mathcal{L}(\mathbf{x}, \mathbf{y}; \boldsymbol{\lambda}) = \tilde{W}(\mathbf{x}, \mathbf{y}) + \sum_{s \in S} \lambda_s (\Phi_s + \frac{1}{2} \mathcal{C}_s \Phi_s^2) \quad (3)$$

- ▶ Update \mathbf{x} and \mathbf{y} by nonlinear optimization method (e.g., *Nestrov* method)
- ▶ ... and gradually increase $\boldsymbol{\lambda}$ to resolve the constraints

³ \mathcal{C}_s : penalty coefficient, $\boldsymbol{\lambda}$: Lagrangian multipliers.

⁴Jing Mai et al. (2022). "Multi-electrostatic FPGA placement considering SLICEL-SLICEM heterogeneity and clock feasibility". In: *Proc. DAC*, pp. 649–654.

Extended Electrostatic Fields

- ▶ Recall the multi-electrostatic-based placement flow ...


$$\min_{x,y} \tilde{W}(x,y) \quad \text{s.t.} \quad \Phi_s(x,y) = 0, \quad \forall s \in S \quad (4)$$

- ▶ Extend the electrostatic fields as

$$S = \{\text{LUTL}, \text{LUTM-AL}, \text{DSP}, \text{BRAM}, \text{IO}\} \quad (5)$$

Asymmetrical Demand and Supply Attributes

- ▶ Demand (cell type)
 - ▶ LUTL field: LUT, Distributed RAM, and SHIFT
 - ▶ LUTM-AL field: Distributed RAM and SHIFT
- ▶ Supply (site type)
 - ▶ LUTL field: SLICEL and SLICEM
 - ▶ LUTM-AL field: SLICEM



SLICEL-SLICEM Heterogeneity (II)



Asymmetrical coeffect of LUTL and LUTM-AL




- ▶ LUT can be placed in SLICEL or SLICEM
- ▶ Distributed RAM and SHIFT can only be placed in SLICEM

LUT
 SHIFT
 SLICEL Column
 SLICEM Column

		LUTL Field		LUTM-AL Field		Two Fields
		ρ_{LUT}	Φ_{LUT}	$\rho_{LUTM-AL}$	$\Phi_{LUTM-AL}$	$\Phi_{LUT} + \Phi_{LUTM-AL}$
Initial			-		-	-
Solution I ✓			Low		Low	Low
Solution II ✓			Low		Low	Low
Solution III ✗			Low		High	High
Solution IV ✗			Low		High	High


Inter-CLB level global routing

- ▶ Coarse-grained routing graph
- ▶ Provide inter-CLB routing topology



Logic element level detailed routing





- ▶ Fine-grained routing graph
- ▶ Generate final routing results
- ▶ **ILP-based tile assignment** is proposed to remedy congestion



Problem Formulation

- ▶ Route multiple nets inside a tile and its neighbor tile concurrently
 - ▶ No overflow vertices
 - ▶ Paths must be connected
 - ▶ No loop in the paths


● Net Source Vertex
 ● RRG Vertex
 ● Net Sink Vertex
 → Used Edges
 → Unused Edges



Integer Linear Programming (ILP) Modeling

1. No overflow vertex

$$\sum_{e,j} R_{e,j} \leq \text{cap}(v), e \in \text{FI}(v) \quad (6)$$




Integer Linear Programming (ILP) Modeling

1. No overflow vertex

$$\sum_{e,j} R_{e,j} \leq \text{cap}(v), e \in \text{FI}(v) \quad (6)$$

2. Each sink of each net is routed

$$S_{e,j,k} \leq R_{e,j}, k \in \text{SINK}(j) \quad (7)$$



Integer Linear Programming (ILP) Modeling

1. No overflow vertex


$$\sum_{e,j} R_{e,j} \leq \text{cap}(v), e \in \text{FI}(v) \quad (6)$$

2. Each sink of each net is routed

$$S_{e,j,k} \leq R_{e,j}, k \in \text{SINK}(j) \quad (7)$$

3. The signal is sent from source pin of each net

$$\sum_{e,j,k} S_{e,j,k} = 1, e \in \text{FO}(v), v = \text{SOURCE}(j), \forall k \in \text{SINK}(j) \quad (8)$$



Integer Linear Programming (ILP) Modeling

1. No overflow vertex

$$\sum_{e,j} R_{e,j} \leq \text{cap}(v), e \in \text{FI}(v) \quad (6)$$

2. Each sink of each net is routed


$$S_{e,j,k} \leq R_{e,j}, k \in \text{SINK}(j) \quad (7)$$

3. The signal is sent from source pin of each net

$$\sum_{e,j,k} S_{e,j,k} = 1, e \in \text{FO}(v), v = \text{SOURCE}(j), \forall k \in \text{SINK}(j) \quad (8)$$

4. The signal is received at each sink pin of each net

$$\sum_{e,j,k} S_{e,j,k} = 1, e \in \text{FI}(v), v = \text{SINK}(j, k) \quad (9)$$



ILP-based Tile Assignment (II)

Integer Linear Programming (ILP) Modeling

1. No overflow vertex

$$\sum_{e,j} R_{e,j} \leq \text{cap}(v), e \in \text{FI}(v) \quad (6)$$

2. Each sink of each net is routed

$$S_{e,j,k} \leq R_{e,j}, k \in \text{SINK}(j) \quad (7)$$

3. The signal is sent from source pin of each net

$$\sum_{e,j,k} S_{e,j,k} = 1, e \in \text{FO}(v), v = \text{SOURCE}(j), \forall k \in \text{SINK}(j) \quad (8)$$


4. The signal is received at each sink pin of each net

$$\sum_{e,j,k} S_{e,j,k} = 1, e \in \text{FI}(v), v = \text{SINK}(j,k) \quad (9)$$

5. There is a path from source pin to each sink pin and no loop

$$\sum_{e_{in}} S_{e_{in},j,k} = \sum_{e_{out}} S_{e_{out},j,k}, \quad (10)$$

$$e_{in} \in \text{FI}(v), e_{out} \in \text{textFO}(v), v \neq \text{SOURCE}(j), v \notin \text{SINK}(j)$$




Code Components

- 1. openparf**
The core placement and routing tool
- 2. openparf.ops**
A collection of operators that allow the implementation of various PR algorithms
- 3. openparf.placement**
A set of APIs for performing placement tasks
- 4. openparf.routing**
A set of APIs for performing routing tasks
- 5. openparf.py_utils**
Provides other utility functions for Python convenience

☰ README.md

OpenPARF

 OpenPARF is an open-source FPGA placement and routing framework build upon the deep learning toolkit [PyTorch](#). It is designed to be flexible, efficient, and extensible.

- [OpenPARF](#)
 - [More About OpenPARF](#)
 - [A Multi-Electrostatic-based FPGA P&R Framework](#)
 - [Reference Flow](#)
 - [Demo](#)
 - [Prerequisites](#)
 - [Build from Source](#)
 - [Install Dependencies](#)
 - [Install Gurobi \(Optional\)](#)
 - [Build with Docker](#)
 - [Docker Image](#)
 - [Using pre-built images](#)
 - [Building the image yourself](#)
 - [Running the Docker Image](#)
 - [Entering the Docker Container](#)
 - [Build and install OpenPARF](#)
 - [Get the OpenPARF Source](#)
 - [Install OpenPARF](#)
 - [Adjust Build Options \(Optional\)](#)
 - [Getting Started](#)
 - [ISPD 2016/2017 Benchmarks](#)
 - [Obtaining Benchmarks](#)
 - [Linking Benchmarks](#)
 - [Running the Benchmarks](#)
 - [Adjust Benchmark Options \(Optional\)](#)
 - [More Advanced Usages](#)
 - [Running Benchmarks in Batches](#)
 - [Vivado Flow for Placement Evaluation](#)
 - [Resources](#)
 - [Releases and Contributing](#)
 - [The Team](#)
 - [Publications](#)
 - [License](#)

Experimental Results

Implementation

- ▶ C++ & Python
- ▶ Build upon `PyTorch` for agile gradient computation

Machine

- ▶ Intel(R) Xeon(R) Gold 6230 CPUs (2.10 GHz, 40 cores)
- ▶ 512GB RAM
- ▶ One NVIDIA RTX 2080Ti GPU

Experiments Setup (II)

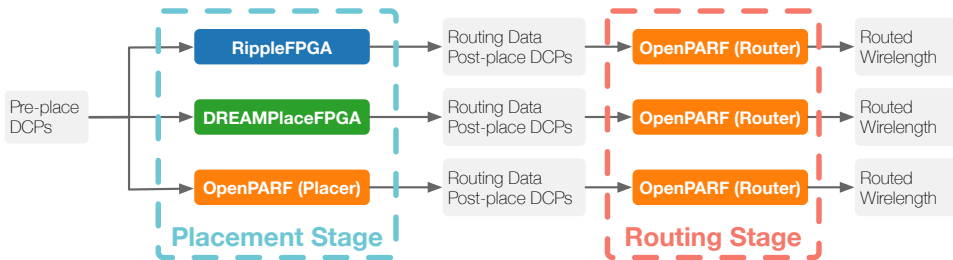
Benchmark Suite

- ▶ ISPD 2016 Routability-Driven FPGA Placement Contest [Yang+, ISPD'16]
- ▶ ISPD 2017 Clock-Aware FPGA Placement Contest [Yang+, ISPD'17]
- ▶ SLICEL-SLICEM Structure-Aware Industrial Benchmarks

Placers for Comparison

- ▶ RippleFPGA [Chen+, TCAD'18]
- ▶ DREAMPlaceFPGA [Rajarithnam+, ISPD'23]

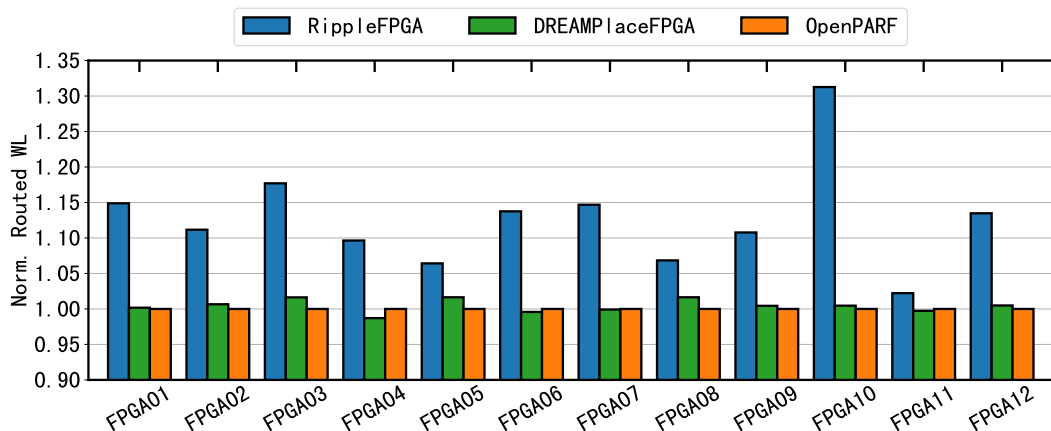
Evaluation Flow



Routed Wirelength Comparison on ISPD 2016

▶ 12.7% better than RippleFPGA

▶ 0.4% better than DREAMPlaceFPGA




OpenPARF significantly outperforms other placers on routed wirelength.

Placement Runtime Comparison on ISPD 2016

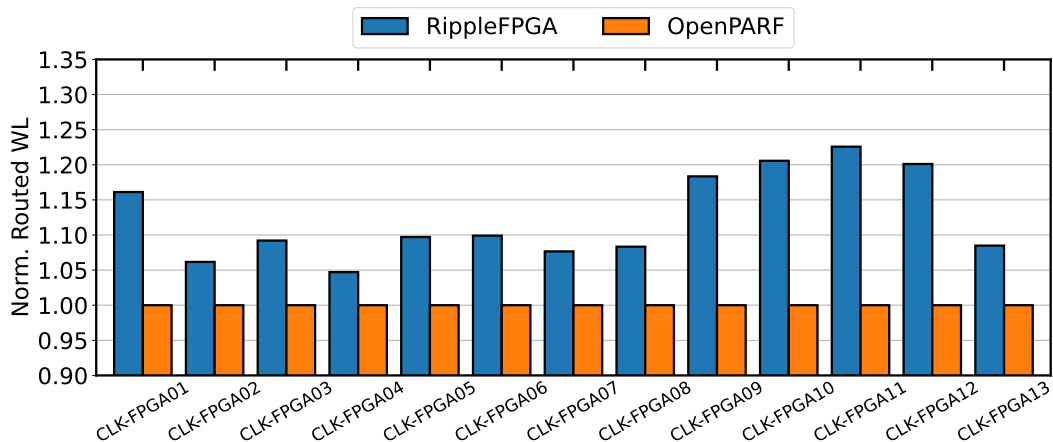
▶ 2.771× faster RippleFPGA

▶ 1.272× slower than DREAMPlaceFPGA



Routed Wirelength Comparison on ISPD 2017⁵


- ▶ 12.8% better than RippleFPGA



⁵DREAMPlaceFPGA is not applicable to this benchmark suite.

Placement Runtime Comparison on ISPD 2017

- ▶ 2.251× faster than RippleFPGA



OpenPARF show notable performance and efficiency on industrial benchmarks.

- ▶ 21K - 284K cells
- ▶ Distributed RAMs and SHIFTs (SLICEM tailored cell types)
- ▶ Cascaded DSPs, BRAMs and CARRYs

Design	#LUT/#FF/ #BRAM/#DSP	#Distributed RAM + #SHIFT	#Net	OpenPARF		
				PRT ⁶	RRT ⁷	RWL ⁸
IND01	17K/11K/0/13	9	52492	72.36	10	90
IND02	11K/10K/0/24	6	26678	77.82	15	100
IND03	109K/12K/0/0	0	121554	109.54	108	1021
IND04	29K/17K/0/16	218	60968	69.39	19	283
IND05	64K/191K/64/928	29K	371808	126.38	109	2360
IND06	112K/65K/21/0	0	221182	88.28	176	1593
IND07	40K/156K/89/768	26K	294075	140.33	68	1450

⁶Placement Runtime (Seconds).

⁷Routing Runtime (Minutes).

⁸Routed Wirelength.

Conclusion & Future Work



Conclusion

- ▶ **OpenPARF**: an open-source placement and routing framework for large-scale FPGAs



Conclusion

- ▶ **OpenPARF**: an open-source placement and routing framework for large-scale FPGAs
- ▶ We build `OpenPARF` upon the deep learning toolkit **PyTorch** for agile gradient computation and flexible programming interfaces



Conclusion

- ▶ **OpenPARF**: an open-source placement and routing framework for large-scale FPGAs
- ▶ We build `OpenPARF` upon the deep learning toolkit **PyTorch** for agile gradient computation and flexible programming interfaces
- ▶ We resolve the **SLICEL-SLICEM heterogeneity** by the SOTA asymmetrical multi-electrostatic FPGA placement algorithms [Mai+, DAC'22]



Conclusion

- ▶ **OpenPARF**: an open-source placement and routing framework for large-scale FPGAs
- ▶ We build `OpenPARF` upon the deep learning toolkit **PyTorch** for agile gradient computation and flexible programming interfaces
- ▶ We resolve the **SLICEL-SLICEM heterogeneity** by the SOTA asymmetrical multi-electrostatic FPGA placement algorithms [Mai+, DAC'22]
- ▶ We harness the **nested** Lagrangian relaxation methodology to resolve multiple placement objectives



Conclusion

- ▶ **OpenPARF**: an open-source placement and routing framework for large-scale FPGAs
- ▶ We build `OpenPARF` upon the deep learning toolkit **PyTorch** for agile gradient computation and flexible programming interfaces
- ▶ We resolve the **SLICEL-SLICEM heterogeneity** by the SOTA asymmetrical multi-electrostatic FPGA placement algorithms [Mai+, DAC'22]
- ▶ We harness the **nested** Lagrangian relaxation methodology to resolve multiple placement objectives
- ▶ We settle the large-scale irregular **CLB-level** FPGA routing problem by the SOTA two-stage negotiation-based routing algorithms [Wang+, ASP-DAC'23]



Conclusion

- ▶ **OpenPARF**: an open-source placement and routing framework for large-scale FPGAs
- ▶ We build `OpenPARF` upon the deep learning toolkit **PyTorch** for agile gradient computation and flexible programming interfaces
- ▶ We resolve the **SLICEL-SLICEM heterogeneity** by the SOTA asymmetrical multi-electrostatic FPGA placement algorithms [Mai+, DAC'22]
- ▶ We harness the **nested** Lagrangian relaxation methodology to resolve multiple placement objectives
- ▶ We settle the large-scale irregular **CLB-level** FPGA routing problem by the SOTA two-stage negotiation-based routing algorithms [Wang+, ASP-DAC'23]

Future Work

- ▶ GPU-accelerated legalization / routing



Conclusion

- ▶ **OpenPARF**: an open-source placement and routing framework for large-scale FPGAs
- ▶ We build `OpenPARF` upon the deep learning toolkit **PyTorch** for agile gradient computation and flexible programming interfaces
- ▶ We resolve the **SLICEL-SLICEM heterogeneity** by the SOTA asymmetrical multi-electrostatic FPGA placement algorithms [Mai+, DAC'22]
- ▶ We harness the **nested** Lagrangian relaxation methodology to resolve multiple placement objectives
- ▶ We settle the large-scale irregular **CLB-level** FPGA routing problem by the SOTA two-stage negotiation-based routing algorithms [Wang+, ASP-DAC'23]

Future Work

- ▶ GPU-accelerated legalization / routing
- ▶ Look-ahead Timing Prediction (LATP) for timing-driven placement



Conclusion

- ▶ **OpenPARF**: an open-source placement and routing framework for large-scale FPGAs
- ▶ We build `OpenPARF` upon the deep learning toolkit **PyTorch** for agile gradient computation and flexible programming interfaces
- ▶ We resolve the **SLICEL-SLICEM heterogeneity** by the SOTA asymmetrical multi-electrostatic FPGA placement algorithms [Mai+, DAC'22]
- ▶ We harness the **nested** Lagrangian relaxation methodology to resolve multiple placement objectives
- ▶ We settle the large-scale irregular **CLB-level** FPGA routing problem by the SOTA two-stage negotiation-based routing algorithms [Wang+, ASP-DAC'23]

Future Work

- ▶ GPU-accelerated legalization / routing
- ▶ Look-ahead Timing Prediction (LATP) for timing-driven placement
- ▶ Fence region-aware placement

THANK YOU!