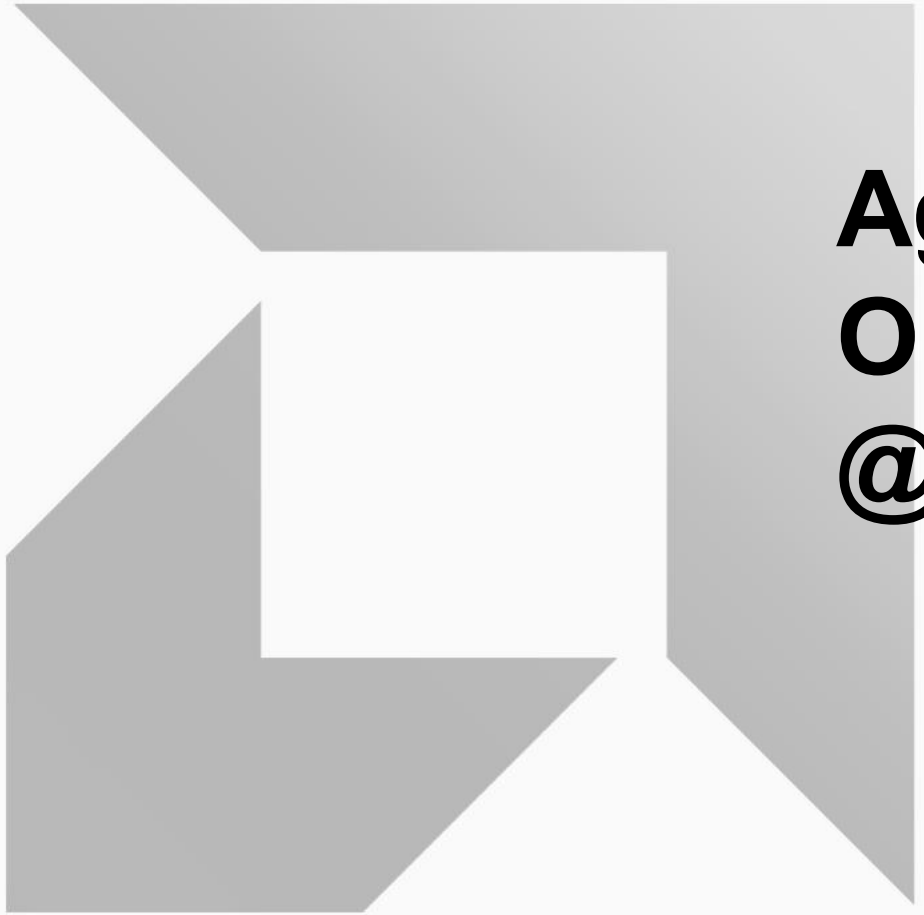

Agentic FPGA Backend Optimization Competition @ FPL'26

Chris Lavin, Eddie Hung, Preston Walker, and Alireza Kaviani

September 9, 2026

AMD
together we advance_

Ever wish you could magically make your FPGA design run faster?

Challenge

Fmax = 395MHz

Given a fully placed and routed design checkpoint…

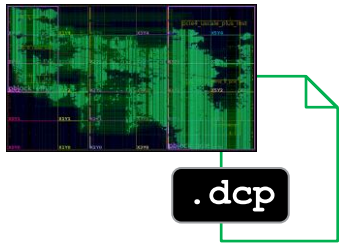

Fmax = 450MHz

…generate a new checkpoint that has improved Fmax as much as possible!

Constraints

- ☑ Fully Placed & Routed
- ☑ No Hold or Pulse Width Violations
- ☑ Functionally Equivalent

Budgets

1 Hour Wall-clock
runtime per benchmark

$1.00 USD LLM Credits
OpenRouter, per benchmark

8 vCPUs (4th Gen Epyc) / 32 GB RAM
AWS EC2 m7a.2xlarge

2025.1 Developer AMI
Ubuntu 22.02

How Much Intelligence for $1.00 USD?

RANK	MODEL	ARTIFICIAL ANALYSIS INTELLIGENCE INDEX	BALANCED TOKENS / $1	INPUT ONLY / $1	OUTPUT ONLY / $1
1	**Ling-3.0-flash** by inclusionai · $0.021 / $0.063 per 1M	37.8	**23.81M**	47.62M	15.87M
2	**DeepSeek V4 Flash 0731** by DeepSeek · $0.0798 / $0.1596 per 1M	51.8	**8.35M**	12.53M	6.27M
3	**Hy3 preview** by tencent · $0.063 / $0.21 per 1M	42.2	**7.33M**	15.87M	4.76M
4	**GPT-5.6 Luna** by OpenAI · $0.10 / $0.60 per 1M	52.3	**2.86M**	10.00M	1.67M
5	**MiMo-V2.5** by Xiaomi · $0.14 / $0.28 per 1M	38.0	**4.76M**	7.14M	3.57M

Source: OpenRouter Value Leaders, https://openrouter.ai/discover?lane=value, August 10th, 2026

Competition Registration Results

91
REGISTERED TEAMS

329
TOTAL PARTICIPANTS
269 members + 60 advisors

77
ACADEMIC INSTITUTIONS

26
COUNTRIES & REGIONS

Participant Roles

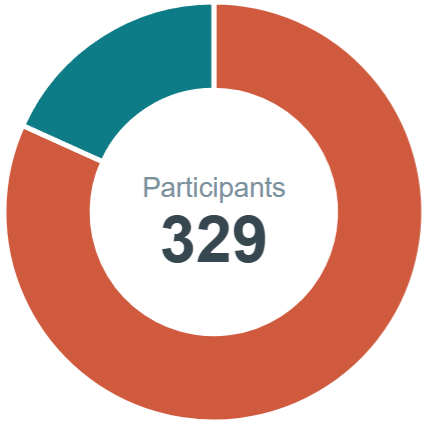

- Members: 269 (81.8%)
- Advisors: 60 (18.2%)

Affiliation Types

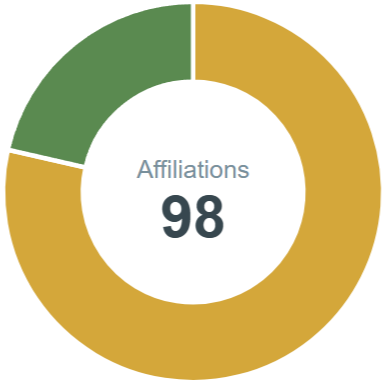

- Academic: 77 (78.6%)
- Industry / Lab / Independent: 21 (21.4%)

Worldwide Contest Participation: 91 Teams

Resources Provided to Contest Teams

91 Teams

Alpha
April
3rd

48 Teams

Beta
July
13th

25 Teams

Final
Aug
10th

5 Finalists

- **$80.00 USD OpenRouter API Key**
- **30 Hours AWS Validation Instance**
- **Vivado Developer Cloud Access**
- **Vivado Knowledgebase MCP Server**
- **Vivado Chat**

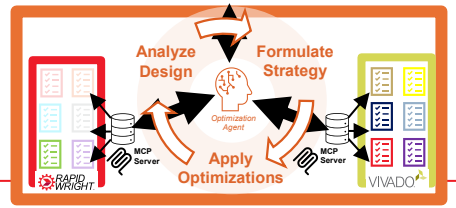

- **Open Source Starter Agent**
  - Vivado & RapidWright MCP Servers
  - 11k lines of code
- **30 Hours AWS Validation Instance**

- **$120.00 USD OpenRouter API Key**
- **30 Hours AWS Validation Instance**

Contest Cloud Infrastructure

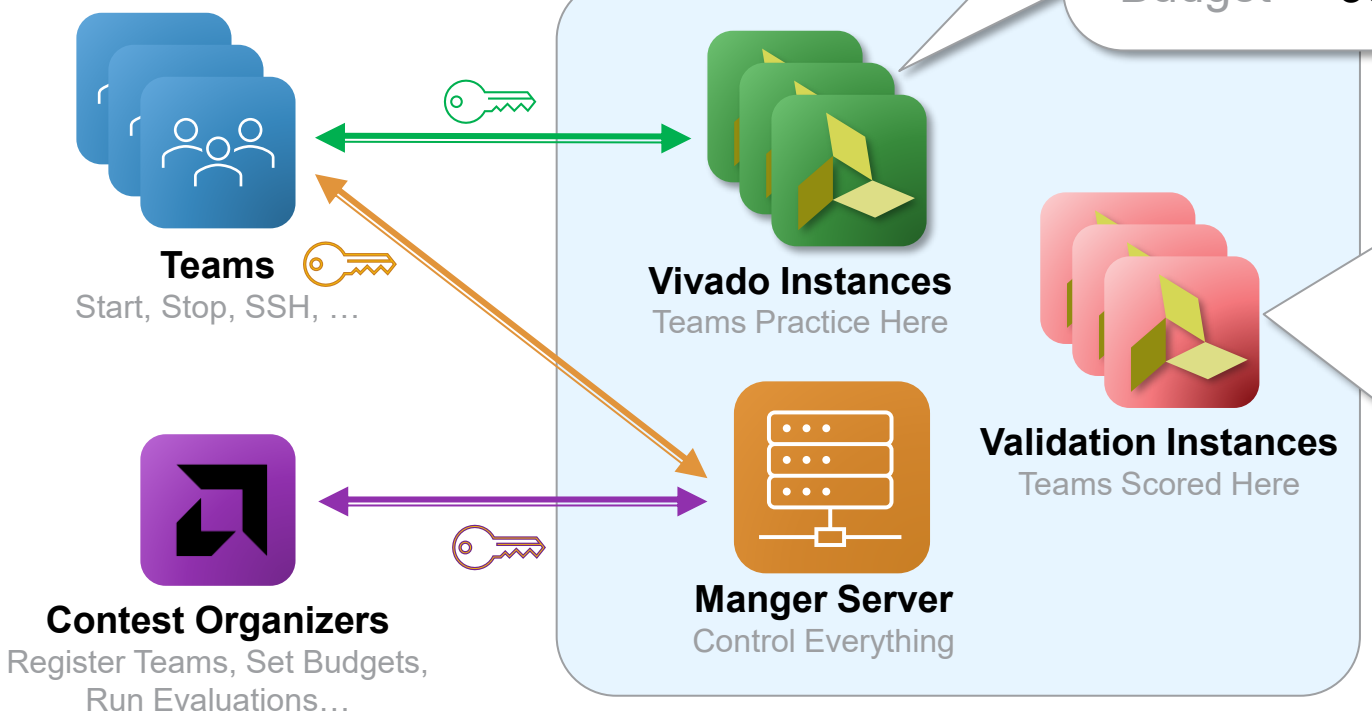

Vivado Instance Specs

Instance	M7a.2xlarge (8 CPU, 32 GB)
AMI	AMD Vivado ML 2025.1
Access	SSH (key-based)
Budget	30 hours per team, per phase

Validation Instance Procedure

1. Start instance, start memory watchdog
2. Run warmup Vivado script
3. Download, unpack team submission
4. Run team's 'make setup; 90 min cap
5. Setup trusted validator
6. Download trusted benchmark
7. Mint $1.00 USD OpenRouter key
8. Run 'make optimizer' kill after 1 hour
9. Kill any zombie processes
10. Use Vivado to validate timing, routed, DRCs of output DCP
11. Run validator flow
12. Publish scorecards, logs, etc.

Contest Benchmark Scoring

$$Benchmark\ Score = \alpha - 0.1\alpha(\beta + \gamma)$$

$$\alpha = \Delta Fmax\ Improvement\ (in\ MHz)$$

$$\beta = OpenRouter\ Cost\ (in\ \$USD)$$

$$\gamma = \frac{Wall\ Clock\ Runtime\ (in\ secs)}{3600\ secs}$$

TERMS

Team Ranking

1. For each benchmark, rank all teams based on their score
2. For each team, compute arithmetic mean of all benchmark rankings (lower is better)
3. Sort teams by mean ranking; assign prizes in ascending order

RapidWright Explores faster; Vivado Provides Trusted Fmax

VIVADO | REALIZE + VERIFY

Authoritative WNS + Fmax
→ Rank candidates on real timing

Complete physical optimization
placement · routing · timing · device rules
→ Legalize changes; catch DRC + hold failures

Global directives + phys-opt
→ Leverage existing optimizations

RAPIDWRIGHT | DIAGNOSE + EDIT

7.5× (18.3 → 2.4 s) faster DCP reads
→ Reload + inspect more candidates

20 ms path-spread analysis
→ Classify placement vs. routing problems immediately

0.17–2.5 s targeted actions
fanout edit · cell ECO · pblock analysis
→ Screen structural ideas before full implementation

Medians from the FPL'26 top-10 tool report; only DCP read is shown as a like-for-like speedup.

Best of 198 Vivado Runs* vs. Best of Contestants' Results

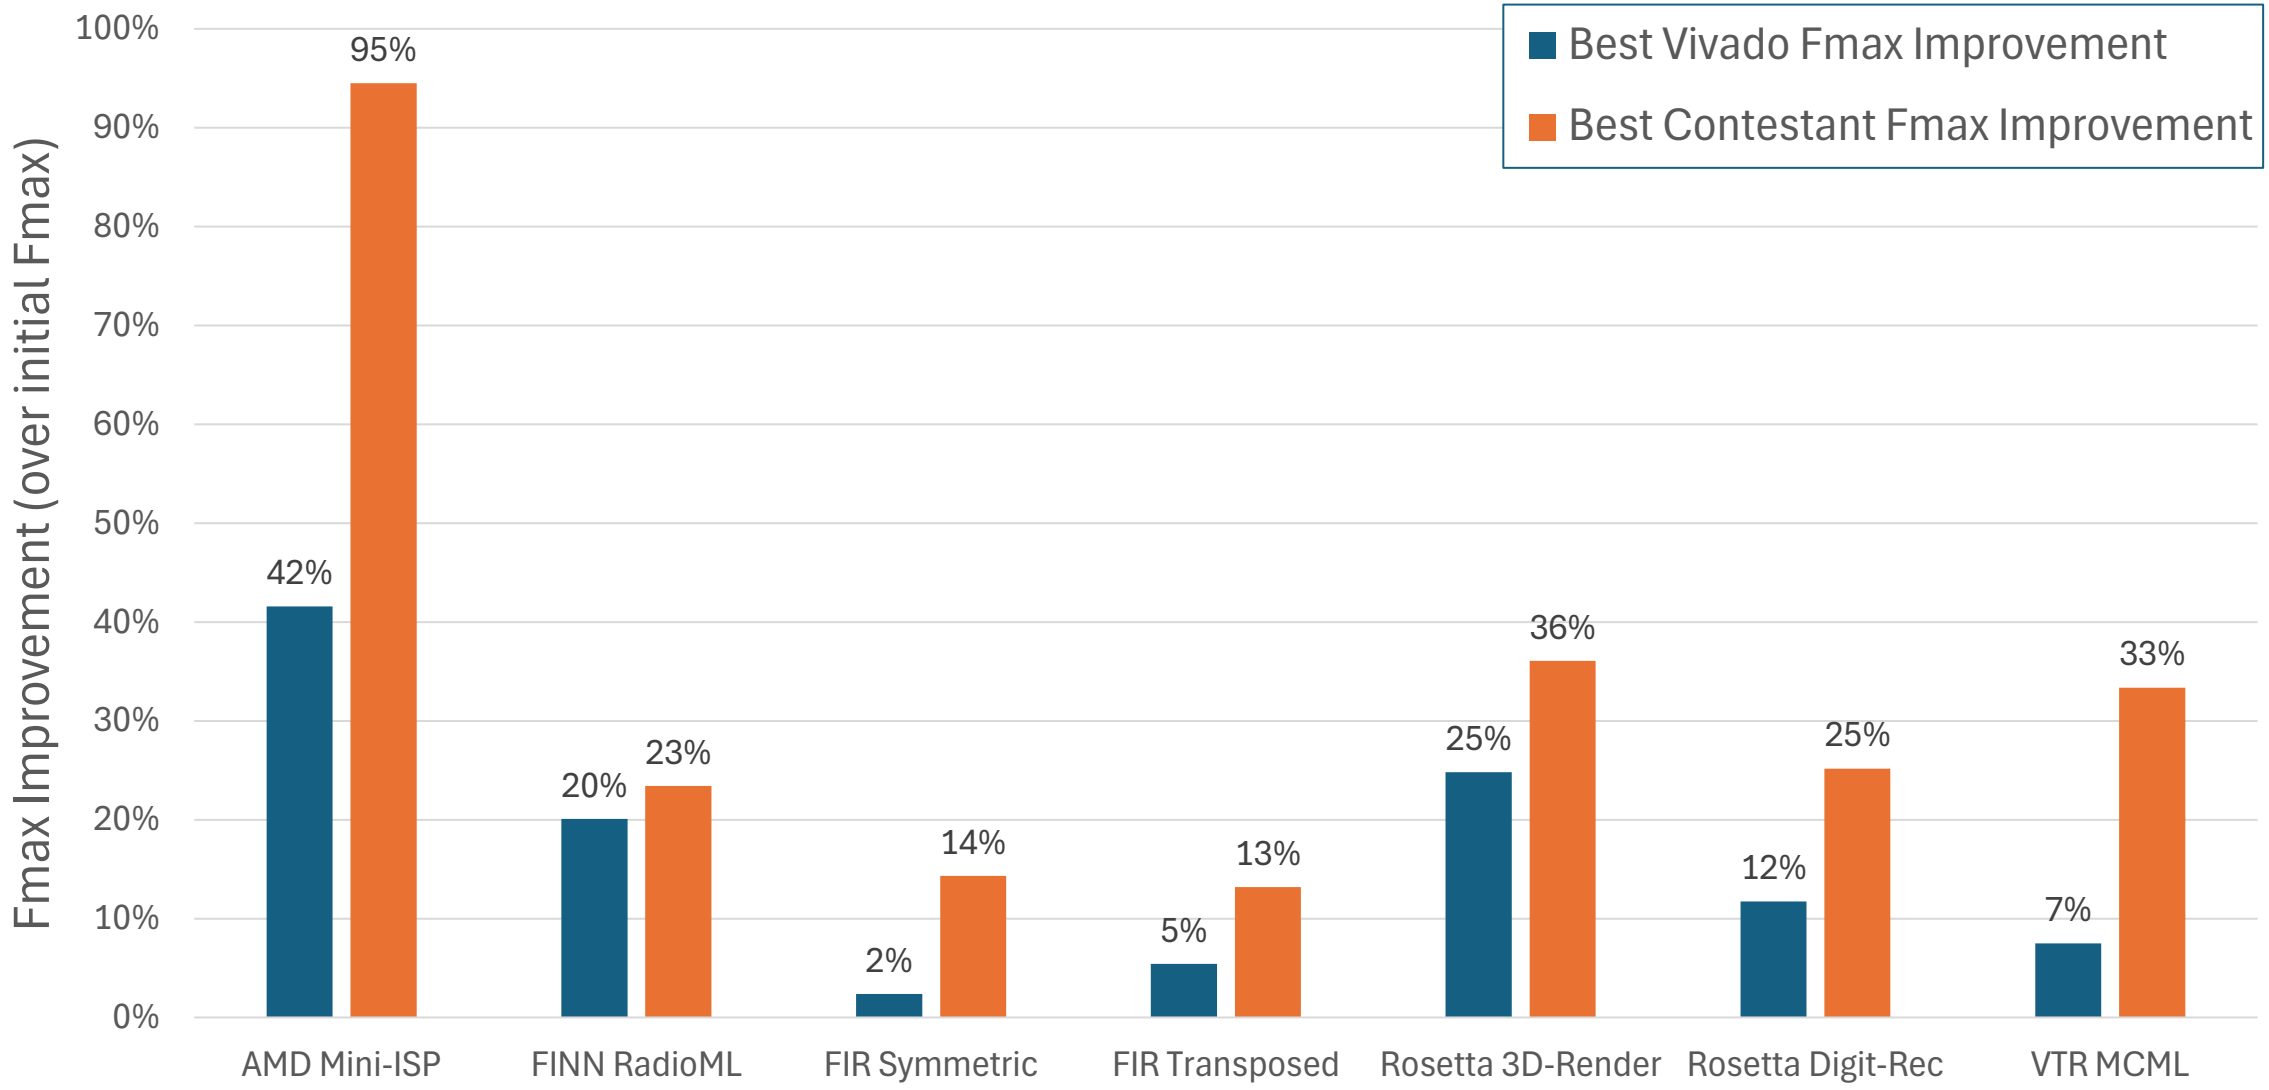

*198 independent Vivado runs sweeping: clock uncertainty {-2%, 0, 2%} x clock constraint {-10% to 10% in 2% increments} x placement directive {Default, Explore} x route directive {Default, Explore, AggressiveExplore}

Best of 198 Vivado Runs* vs. Best of Contestants' Results

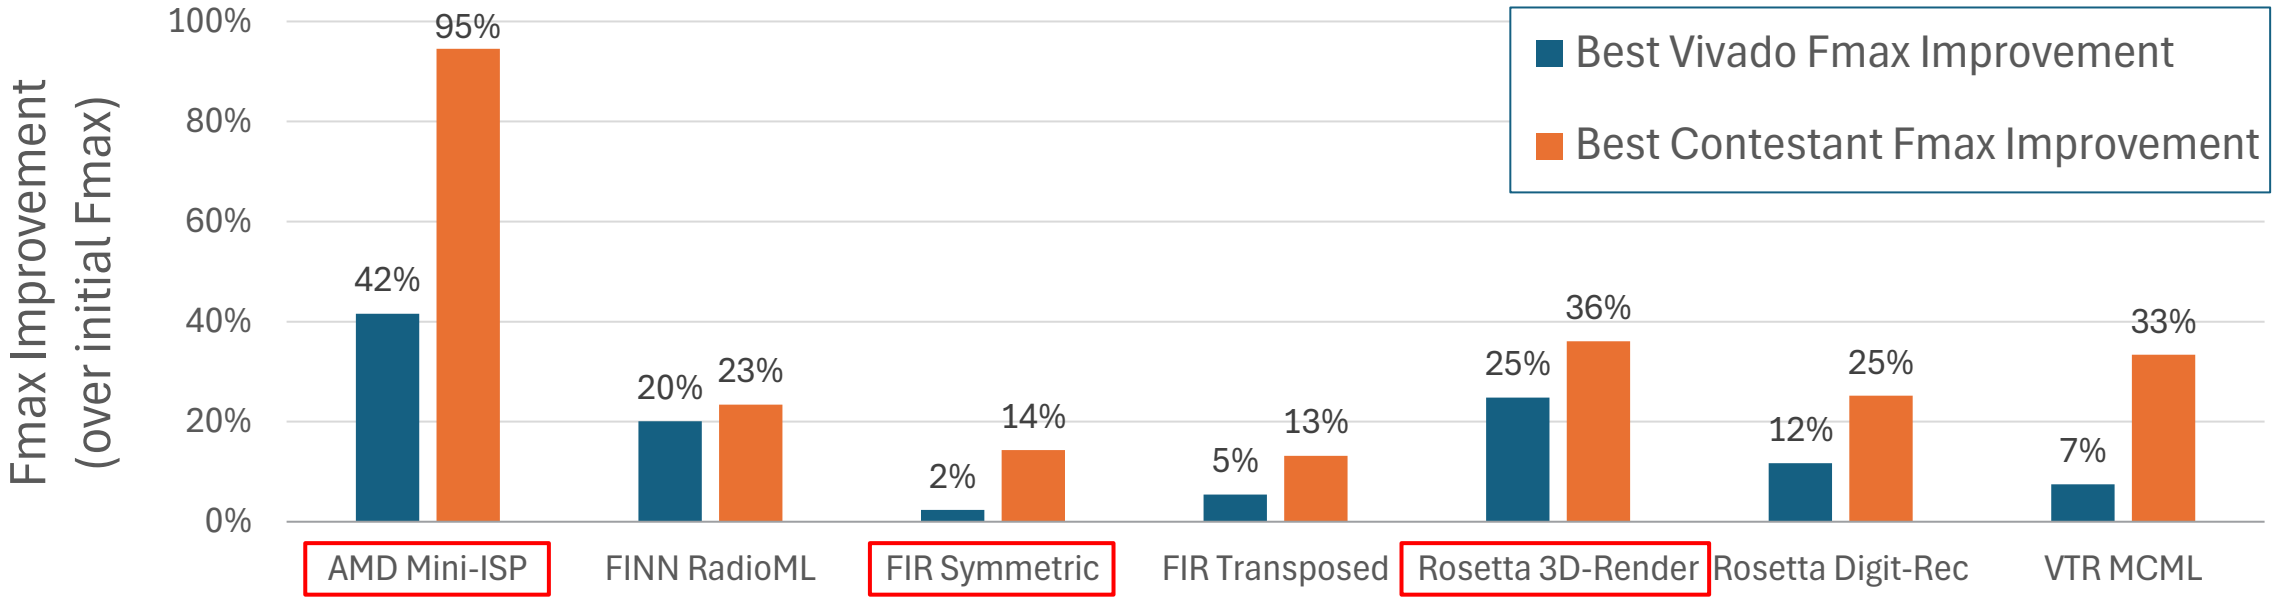

	AMD Mini-ISP	FINN RadioML	FIR Symmetric	FIR Transposed	Rosetta 3D-Render	Rosetta Digit-Rec	VTR MCML
Initial Fmax (MHz)	288.0	284.9	405.7	355.5	236.7	367.0	69.3
Best Vivado Fmax (MHz)	407.8	342.1	415.3	374.8	295.4	410.0	74.5
Best Contestant Fmax (MHz)	560.2	351.6	463.8	402.4	322.1	459.3	92.5
LUTs / FFs	3k / 4k	74k / 46k	0 / 0	0 / 0	14k / 5k	23k / 23k	43k / 15k
DSPs / BRAMs	40 / 12	0 / 25	456 / 0	680 / 0	3 / 0	0 / 161	105 / 142

= Previously unreleased

*198 independent Vivado runs sweeping: clock uncertainty {-2%, 0, 2%} x clock constraint {-10% to 10% in 2% increments} x placement directive {Default, Explore} x route directive {Default, Explore, AggressiveExplore}

What if Frontier Models Entered as Contestants?

Each frontier model given:

- Contest repo & rules
- 1 prompt & looped for ~80 hours

Prompt

"You are a meta-optimization agent…"
"…improve the contest agent to maximize contest score…"
"…follow all contest rules…"

Coding Harness	Frontier Model Tested	Approx. Cost
Claude Code	Opus 4.8 (high effort)	$2,200 USD
Codex	GPT 5.6 Sol (high effort)	$750 USD
Claude Code	Opus 5 (high effort)	$620 USD

Comparison of Frontier AI Built Team Submissions

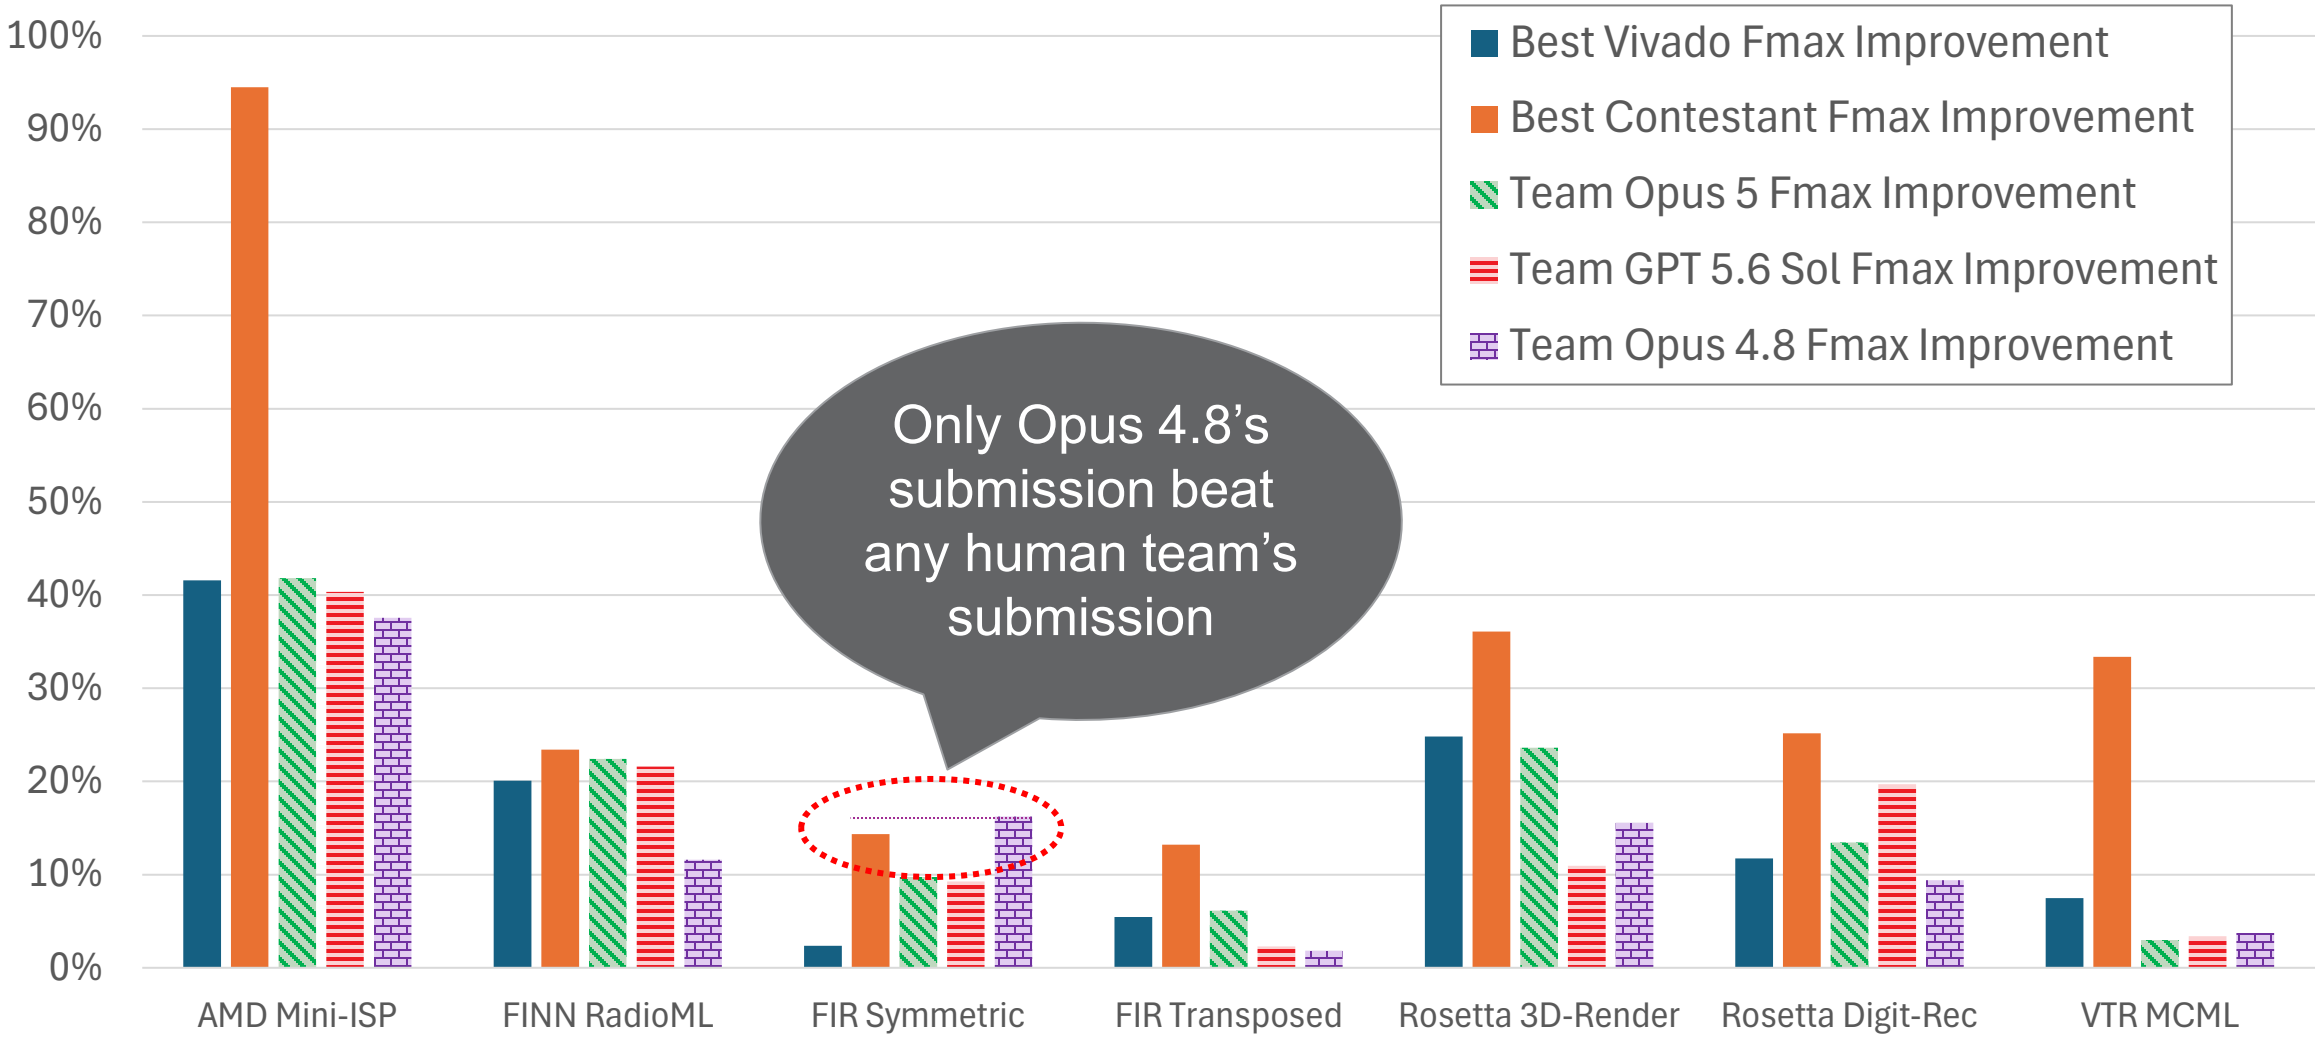

Top 5 Teams (Alphabetical Order) & Honorable Mentions

- **AER** | *Seoul National University[1], Kyunghee University[2]*
  - Sihun Lim[1], Sookwan Han[1], Sung Jun[1], and Haeun Ji[2] | **Advisor:** Jinho Lee[1]
- **ESCA** | *Korea University*
  - Jangho Lim, Junseok Kang, and Jinhyeok Choi | **Advisor:** Taeweon Suh
- **Fu_Do_Fdo** | *Fudan University*
  - KunLong Li and Huan Lin | **Advisor:** Zhiang Wang
- **MEEMAR** | *Ankara Yıldırım Beyazıt University*
  - Kamoliddin Fatkhiddinov | **Advisor:** Enver Çavuş
- **RouteAgents** | *University of Thessaly*
  - Georgios Chatzitsompanis | **Advisor:** Georgios Karakonstantis

Honorable Mentions

- **Fmaximizers** | *Delft University of Technology*
- **Novi_CAD** | *University of Novi Sad*

Agentic FPGA Backend Optimization Competition

Team AER | One-Slide Summary

Sookwan Han, Haeun JI, Jun Sung, Sihun Lim
Advising Professor: Jinho Lee

Phase 1

Autoresearch Meta-Harness with Sparse Human-Feedback

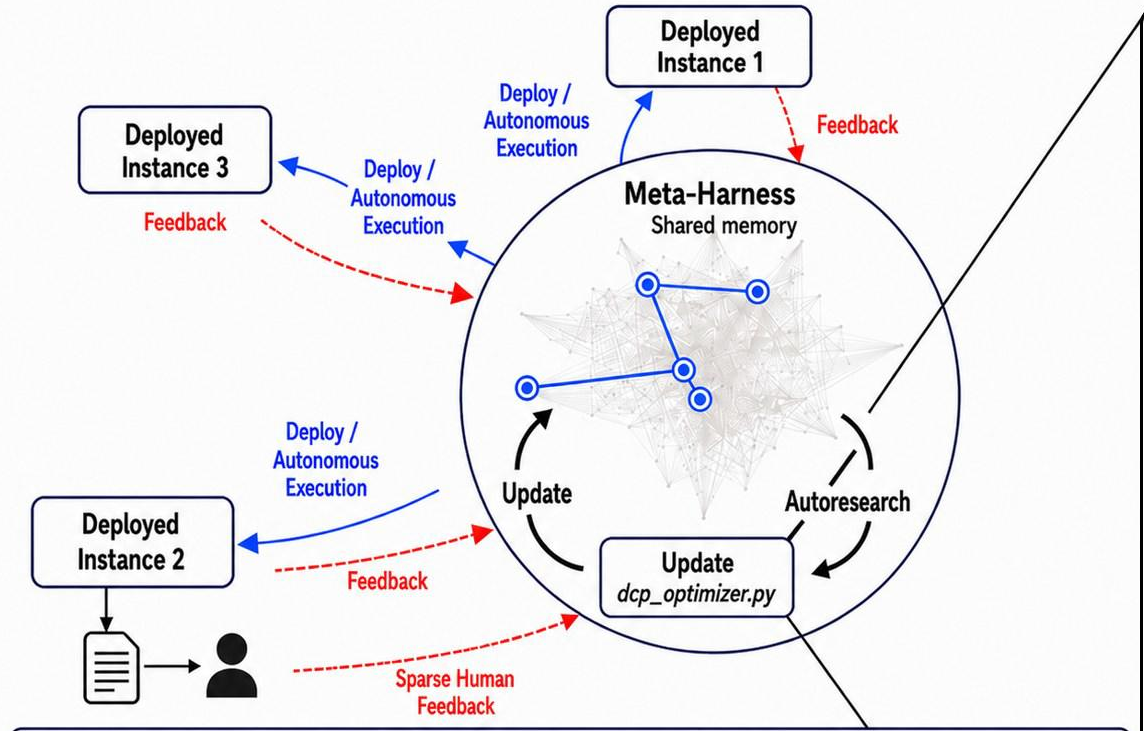

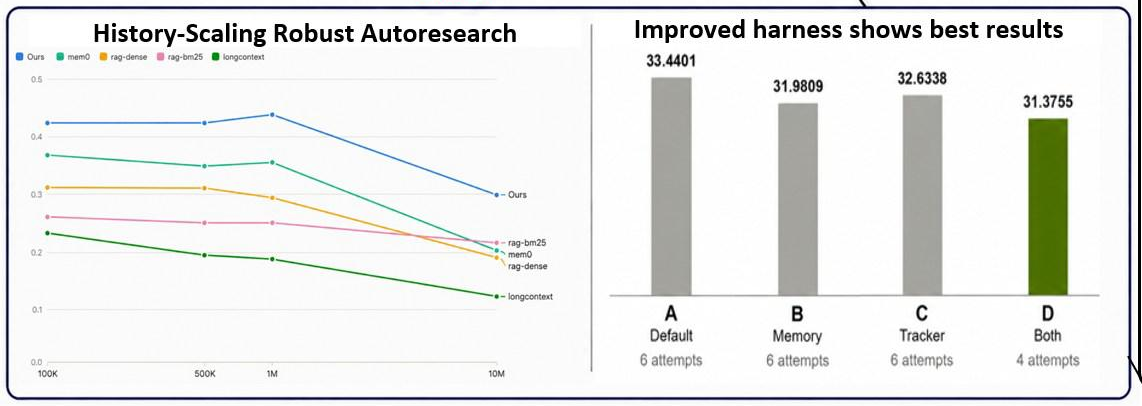

Phase 2

P1 Research-Informed Reconfiguration & Optimization

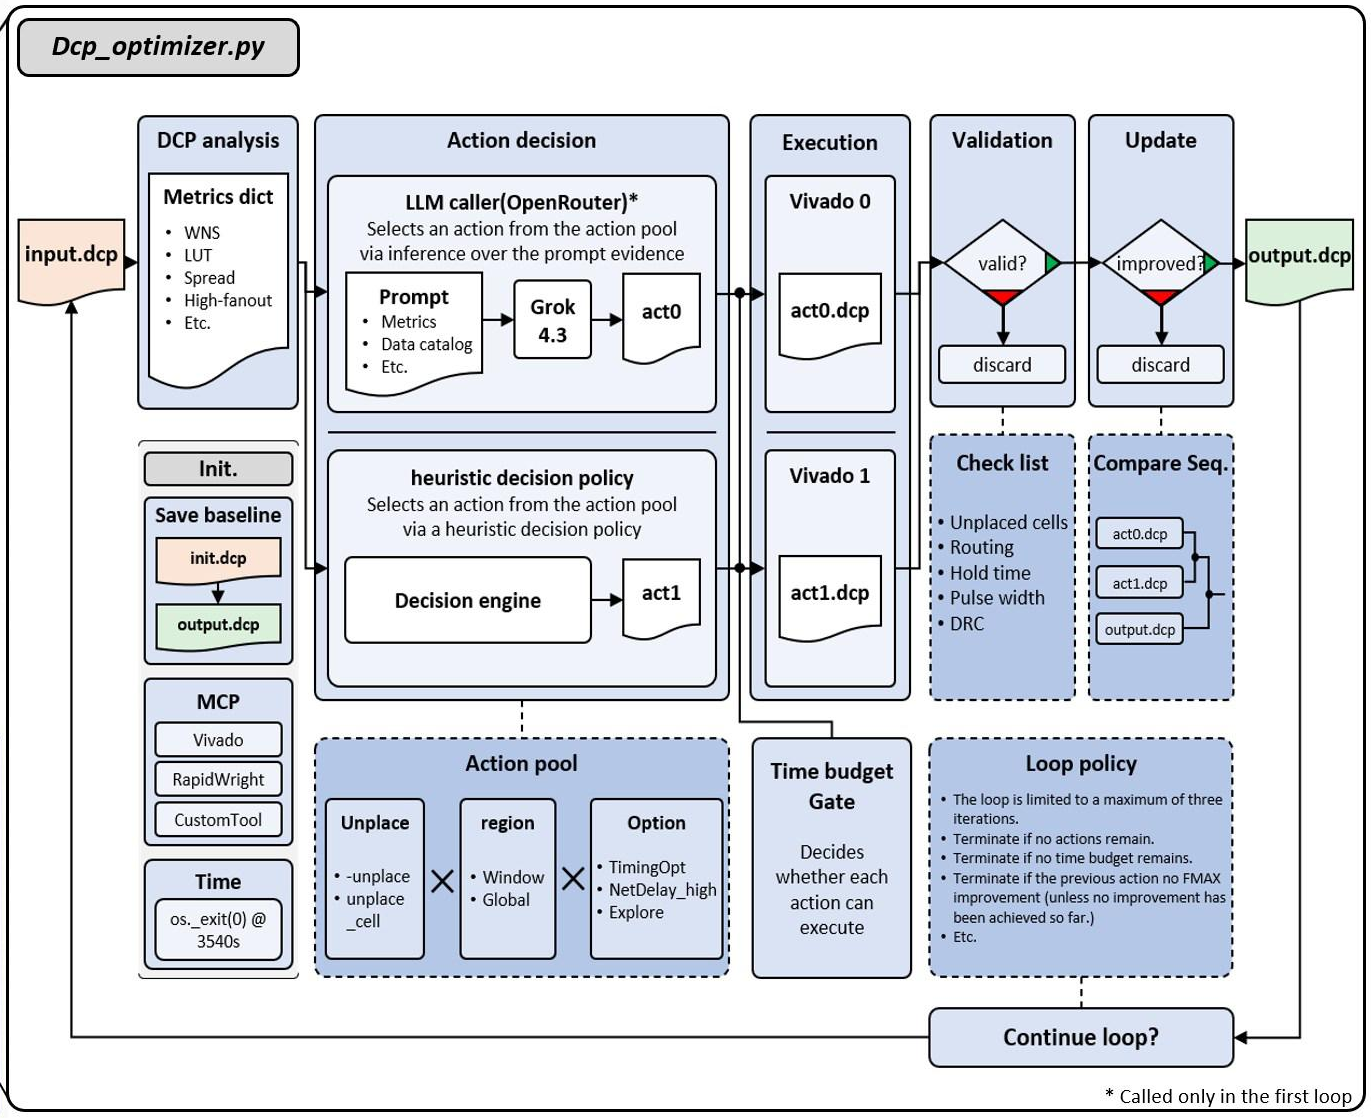

* Called only in the first loop

1. Optimization strategy · Parallel slots

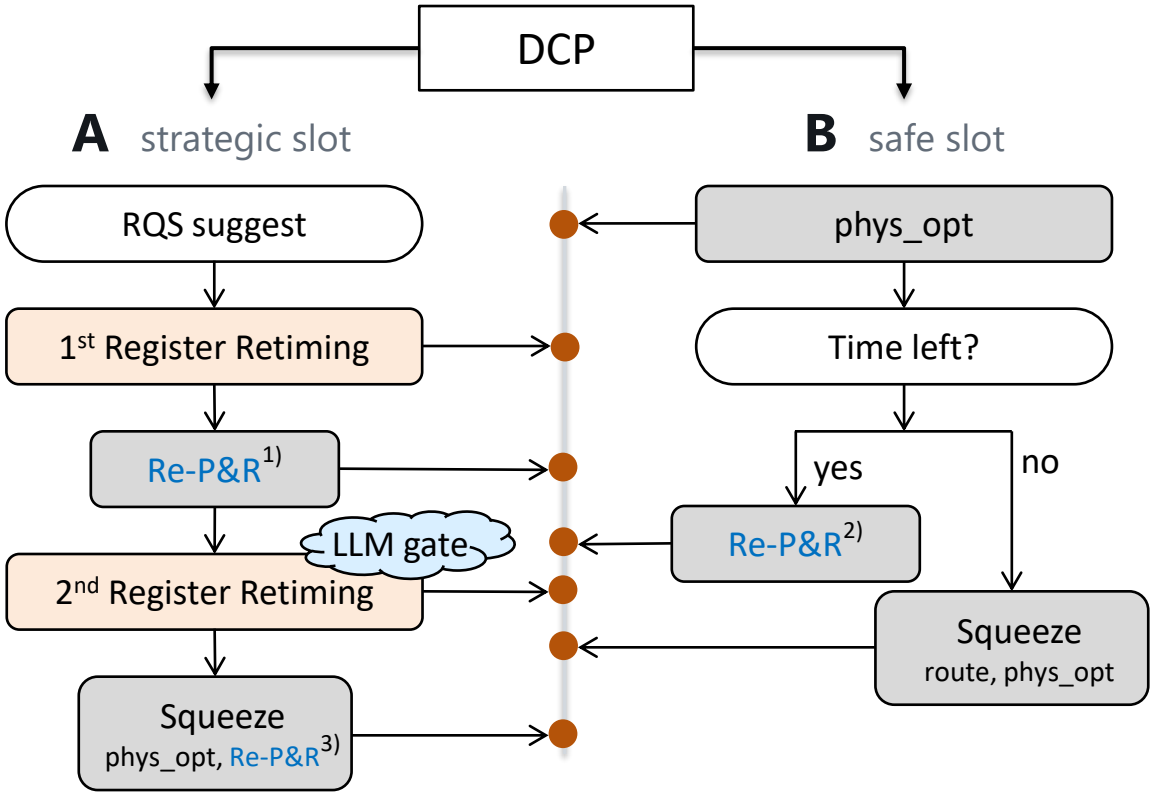

● publish if better than the best DCP (**keep-best**)

LLM gate prior benchmark, measured 1st retiming cost→go/no-go

Re-P&R Un-place, Un-route, and P&R

1) RQS-suggested directives
2) Explore directives – counter-bet to A
3) unused directives

2. Register Retiming

How it gets faster? the critical path borrows its neighbor's slack

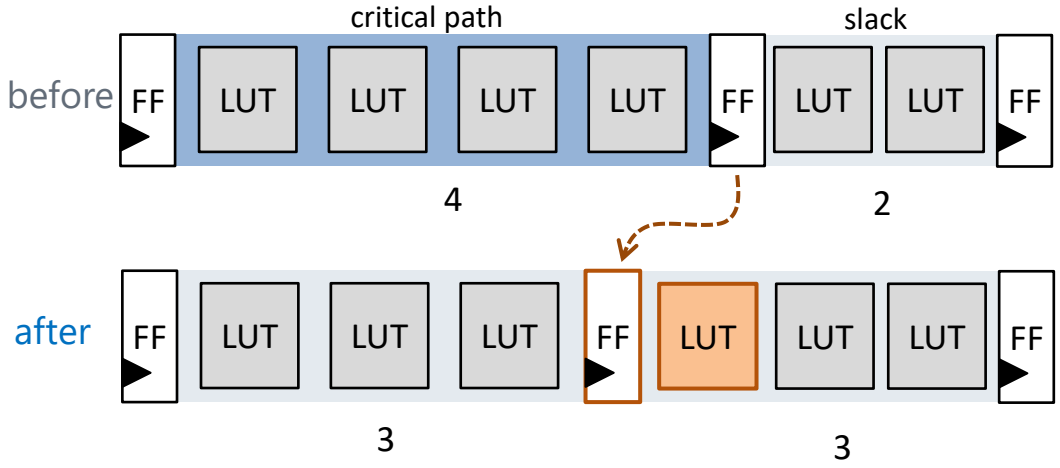

3. Official final scorecard · 7 hidden benchmarks

Benchmark	Strategy	Fmax in→out	Score
amd_mini-isp	A · 2nd Retiming	288.0 → 512.0	217.97
finn_radioml	B · Re-P&R² → phys_opt	284.9 → 345.9	56.69
rosetta_3d	A · Re-P&R¹ → phys_opt	236.7 → 290.4	51.28
rosetta_digit	A · Re-P&R¹ → phys_opt	367.0 → 419.1	47.04
fir_transposed	A · 1st Retiming → phys_opt	355.5 → 399.5	42.09
vtr_mcml	A · 1st Retiming → phys_opt	69.3 → 92.5	21.44
fir_symmetric	B · phys_opt only	405.7 → 422.8	16.32
Total			**452.83**

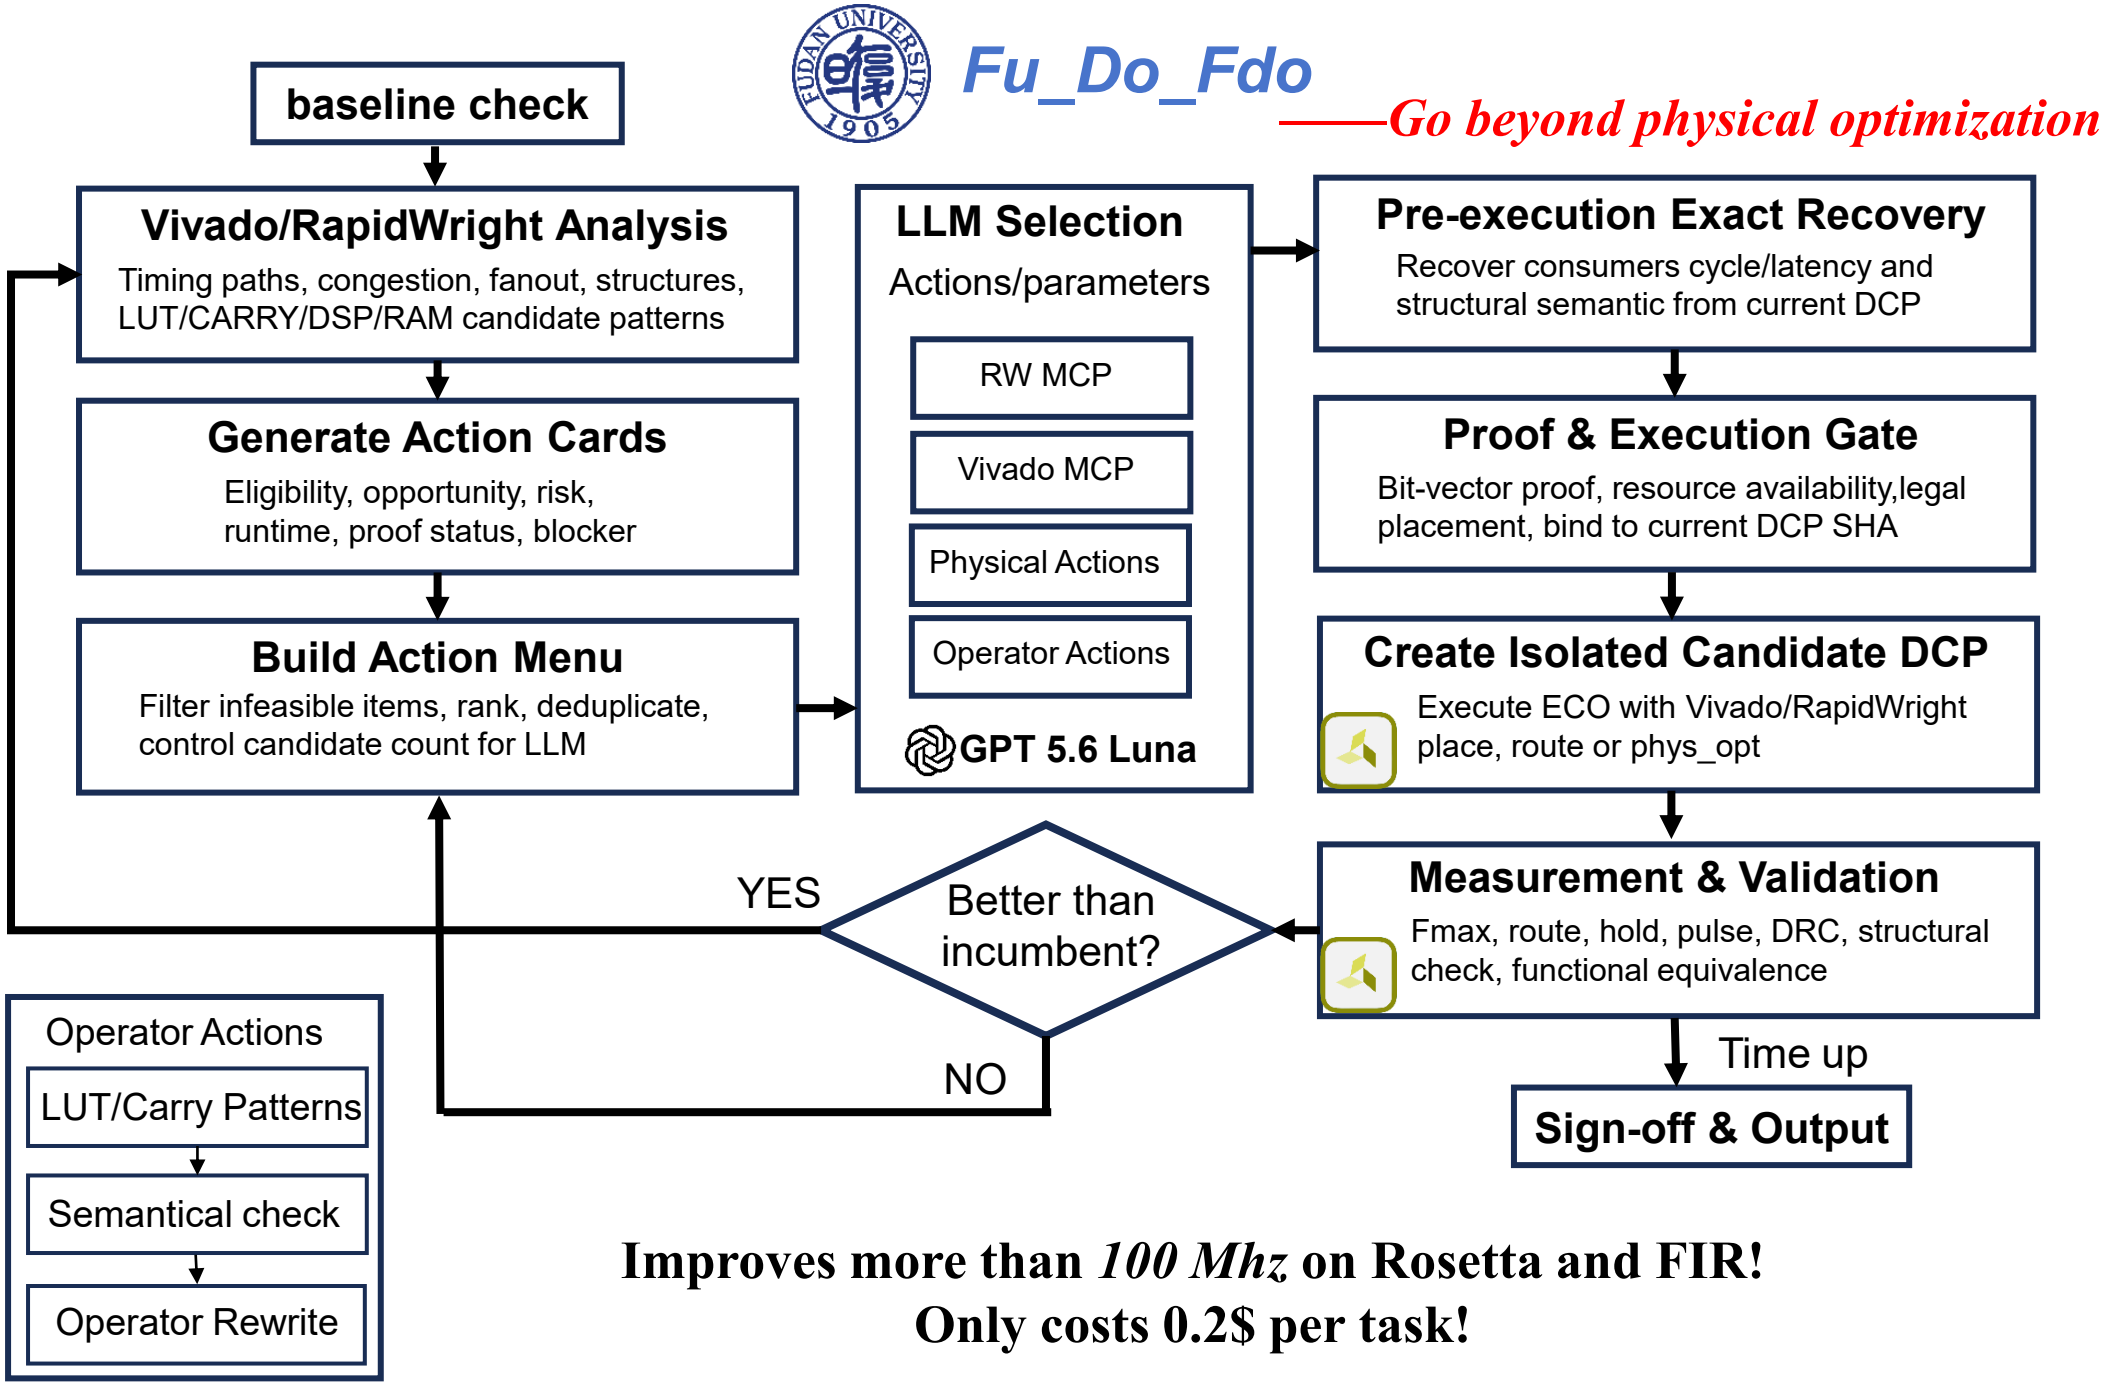
Fu_Do_Fdo
——Go beyond physical optimization
baseline check
Vivado/RapidWright Analysis
Timing paths, congestion, fanout, structures, LUT/CARRY/DSP/RAM candidate patterns
Generate Action Cards
Eligibility, opportunity, risk, runtime, proof status, blocker
Build Action Menu
Filter infeasible items, rank, deduplicate, control candidate count for LLM
LLM Selection
Actions/parameters
RW MCP
Vivado MCP
Physical Actions
Operator Actions
GPT 5.6 Luna
Pre-execution Exact Recovery
Recover consumers cycle/latency and structural semantic from current DCP
Proof & Execution Gate
Bit-vector proof, resource availability,legal placement, bind to current DCP SHA
Create Isolated Candidate DCP
Execute ECO with Vivado/RapidWright place, route or phys_opt
Measurement & Validation
Fmax, route, hold, pulse, DRC, structural check, functional equivalence
Better than incumbent?
YES
NO
Time up
Sign-off & Output
Operator Actions
LUT/Carry Patterns
Semantical check
Operator Rewrite
Improves more than 100 Mhz on Rosetta and FIR!
Only costs 0.2$ per task!

How a stochastic agent was made *unable to lose ground.*

Team MEEMAR · finalist
FPL'26 Agentic FPGA Backend Optimization Contest
Kamoliddin Fatkhiddinov, a one-person entry · advisor Prof. Enver Çavuş, AYBU
`xcvu3p-ffvc1517-2-e · Vivado 2025.1 + RapidWright · 8 threads`

Take a placed-and-routed checkpoint, return a faster one. Three deterministic stages do most of the work; the model is called last, and only when it can still help. **Nothing is kept unless it is legal and better than what came before.**

01

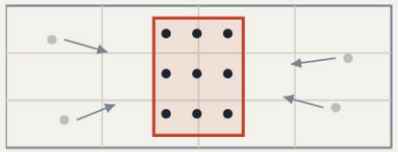

squeezed to ~50 % density, centred on the device

Floorplan prepass

device-centred, clock-region aligned, deterministic

+129.6 / +103.4 / +82.6 MHz

02

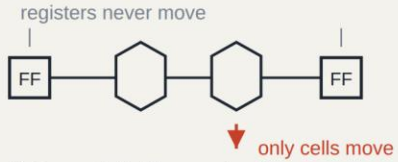

retiming excluded by construction

Physical optimisation

retiming-free by construction — it cannot fail validation

+6.7 / +2.7 MHz

03

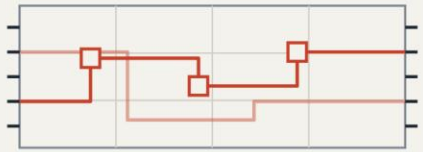

pin LOCs held · worst ~200 paths released

Surgical re-place, on demand

a rescue: it fires only on a weak floor of a large design

+57.7 / +52.3 MHz

04

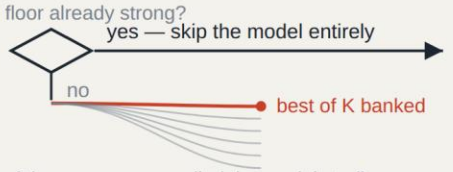

three of the seven never called the model at all

LLM phase, over MCP

its best behaviour is knowing when not to run

0 calls on 3 of 7 · 265 on the hardest

Every stage banks through one gate: legal *and* better, or not at all — single writer, atomic, never downward.

The output checkpoint is written before any tool server starts, so a crash costs the upside, never the benchmark.

amd_mini-isp	288 → 411	+42.5%
★ finn_radioml	285 → 350	+22.8%
★ rosetta 3d-render	237 → 290	+22.5%
★ rosetta digit-rec	367 → 429	+16.9%
fir_systolic_transp	355 → 374	+5.3%
vtr_mcml	69 → 72	+4.4%
★ fir_symmetric	406 → 414	+2.1%

a the model was skipped entirely — amd finished in 318 s of its 3600

b last gain banked at 3240 s, so the cap cost runtime, not megahertz

318.254

total score across seven benchmarks — zero disqualifications

4 / 7

designs we had never run — all clean, all positive

$0.35

model spend, on `gemini-3.1-flash-lite`

`score = α − 0.1·α·(β + γ)` — megahertz gained, charged alike for dollars and hours.

Every run routed, DRC-clean and simulation-equivalent. Most of Σα 332.8 MHz is deterministic: two repeated benchmarks reproduced our earlier measurement to three decimals.

Open source under Apache-2.0 — **github.com/randomuzbek/meemar-fpl26-optimizer**

alpha → beta → final · four and a half months, worked evenings

ROUTEAGENTS

LLM-GUIDED FREQUENCY OPTIMIZATION FOR POST-ROUTED DCPS

INPUT: A ROUTED .DCP TO BE MADE FASTER · ONE HOUR · ONE DOLLAR

FPL 26

Department of Electrical and Computer Engineering
University of Thessaly
Georgios Chatzitsompanis · Georgios Karakonstantis

THE PLAYBOOK

42 PLAYS, MINED OFFLINE FROM ATTRIBUTED RUNS ON THE PUBLIC DEVELOPMENT DESIGNS

EVERY PLAY IS STORED AS

MEASURED TRIGGER → THE PLAY → MEASURED EXCLUSION

the band it pays in **and the band it does not**

FROZEN AT RUN TIME
no writes, no retrieval: mined before the run and never edited while the clock is running

.dcp in
ROUTED
placed and routed checkpoint

DETERMINISTIC

1 SENSE
reports, path anatomy, utilization

2 DECIDE
the measurements pick the plays, and their order

3 RECIPE PASS
deterministic; can finish the run here

LLM-DRIVEN

4 LLM LOOP
READS logs, reports, the brief
CHOOSES the play **+ its operands**
HANDS BACK candidate **+ measurement**
34 tools + 6 macros · 1 h · under $1

DETERMINISTIC

5 POLISH LADDER
each step banks before the next risks

6 FINALIZE
a measured play-off picks the winner

.dcp out
VALID
routed · DRC setup · hold

THE GOVERNORS

EVERY STAGE ANSWERS HERE. NOTHING IS BANKED UNMEASURED

THE BANK
a worse attempt never replaces the best · every risk forks from an untouched copy

THE PRICED CLOCK
score = alpha * (1 - 0.1 * (beta + gamma))
the next second must out-earn its own tax

VALID-STATE ADMISSION
a result is banked only if it was measured on a design that is still fully routed

WHAT A PLAY LOOKS LIKE

1 OF 42 · THE TRIGGER AND THE EXCLUSION ARE BOTH COUNTED

ILS PARTIAL RUIN

UNPLACE ONLY THE WORST PATHS' CELLS, THEN REBUILD

PAYS WHEN
the critical path is spread across the die
improved 181 of 188 independent runs

DOES NOT PAY
when that same path sits co-located in one region
0 of 26, and those losses average -1.09 ns

SO THE PLAY IS FENCED AT 30 TILES OF PATH SPREAD
counted on runs the miner was never shown:
a boundary you cannot guess, only measure

INSIDE THE LLM LOOP

grok-4.3 VIA OPENROUTER · SWAP IN ANOTHER MODEL WITH ONE FLAG

READS
the brief, and everything it can measure
the router hands it an ordered, costed plan with the inapplicable plays blocked · reports, JSON, route status, and prior campaigns by fingerprint

CHOOSES
the play **and its operands**
which endpoints · *which* nets · *which* cells, named from a report it just produced: operands a directive-level interface never asks for.

HANDS BACK
a candidate **and its own measurement**
analyze, transform, route, measure, inside one call · the callable restores the bank itself when the number comes back worse

OFFICIAL RESULTS

OFFICIAL SCORECARD · SEVEN BENCHMARK DESIGNS · SCORED END TO END

369.9	7 / 7	0	$0.18
OFFICIAL SCORE	IMPROVED	INVALID	MOST SPENT ON ONE DESIGN

DELTA FMAX PER DESIGN (MHZ)

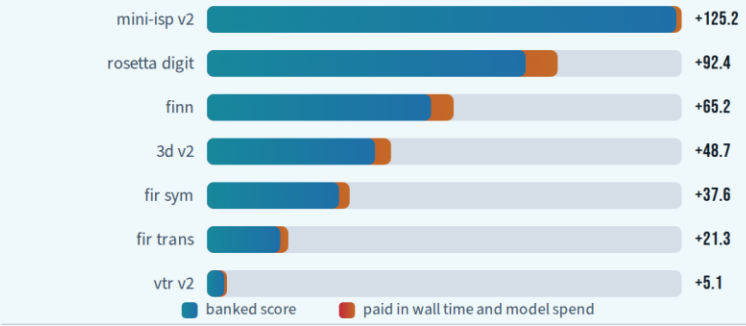

every design improved, and every output re-validated

Available: github.com/Geochatz3/routeagents-fpl26

This work is partially supported by the H.F.R.I IoTHealth Project (no. 25267) in the framework of the 3rd Call for H.F.R.I.'s Research Projects

Finalists & Honorable Mentions

Rank	Prize	Team	Mean Rank	AMD Mini-ISP	FINN RadioML	FIR Symmetric	FIR Transposed	Rosetta 3D-Render	Rosetta Digit-Rec	VTR MCML
1st	€ 3,000									
2nd	€ 2,000									
3rd	€ 1,000									
4th	€ 500									
5th	€ 500									
6th	Kria KD240									
7th	Kria KD240									

Finalists & Honorable Mentions

Rank	Prize	Team	Mean Rank	AMD Mini-ISP	FINN RadioML	FIR Symmetric	FIR Transposed	Rosetta 3D-Render	Rosetta Digit-Rec	VTR MCML
1st	€ 3,000									
2nd	€ 2,000									
3rd	€ 1,000									
4th	€ 500									
5th	€ 500									
6th	Kria KD240	**Fmaximizers** *Delft University of Technology*	8.86	1 (246.6)	19 (46.1)	4 (45.0)	7 (18.4)	19 (33.9)	7 (51.9)	5 (3.6)
7th	Kria KD240	**Novi_CAD** *University of Novi Sad*	9.57	10 (115.6)	7 (58.0)	3 (47.6)	3 (22.4)	3 (54.6)	21 (0.0)	20 (1.7)

Finalists & Honorable Mentions

Rank	Prize	Team	Mean Rank	AMD Mini-ISP	FINN RadioML	FIR Symmetric	FIR Transposed	Rosetta 3D-Render	Rosetta Digit-Rec	VTR MCML
1st	€ 3,000									
2nd	€ 2,000									
3rd	€ 1,000									
4th	€ 500	**Fu_Do_Fdo** *Fudan University*	7.43	22 (78.8)	9 (56.9)	1 (53.1)	1 (42.8)	1 (77.6)	11 (46.0)	7 (3.1)
5th	€ 500	**MEEMAR** *Ankara Yıldırım Beyazıt University*	7.71	4 (121.4)	2 (61.4)	18 (7.7)	11 (16.8)	7 (50.7)	3 (57.4)	9 (2.8)
6th	Kria KD240	**Fmaximizers** *Delft University of Technology*	8.86	1 (246.6)	19 (46.1)	4 (45.0)	7 (18.4)	19 (33.9)	7 (51.9)	5 (3.6)
7th	Kria KD240	**Novi_CAD** *University of Novi Sad*	9.57	10 (115.6)	7 (58.0)	3 (47.6)	3 (22.4)	3 (54.6)	21 (0.0)	20 (1.7)

Finalists & Honorable Mentions

Rank	Prize	Team	Mean Rank	AMD Mini-ISP	FINN RadioML	FIR Symmetric	FIR Transposed	Rosetta 3D-Render	Rosetta Digit-Rec	VTR MCML
1st	€ 3,000									
2nd	€ 2,000									
3rd	€ 1,000	**ESCA** *Korea University*	6.29	2 (218.0)	10 (56.7)	13 (16.3)	2 (42.1)	6 (51.3)	10 (47.0)	1 (21.4)
4th	€ 500	**Fu_Do_Fdo** *Fudan University*	7.43	22 (78.8)	9 (56.9)	1 (53.1)	1 (42.8)	1 (77.6)	11 (46.0)	7 (3.1)
5th	€ 500	**MEEMAR** *Ankara Yıldırım Beyazıt University*	7.71	4 (121.4)	2 (61.4)	18 (7.7)	11 (16.8)	7 (50.7)	3 (57.4)	9 (2.8)
6th	Kria KD240	**Fmaximizers** *Delft University of Technology*	8.86	1 (246.6)	19 (46.1)	4 (45.0)	7 (18.4)	19 (33.9)	7 (51.9)	5 (3.6)
7th	Kria KD240	**Novi_CAD** *University of Novi Sad*	9.57	10 (115.6)	7 (58.0)	3 (47.6)	3 (22.4)	3 (54.6)	21 (0.0)	20 (1.7)

Finalists & Honorable Mentions

Rank	Prize	Team	Mean Rank	AMD Mini-ISP	FINN RadioML	FIR Symmetric	FIR Transposed	Rosetta 3D-Render	Rosetta Digit-Rec	VTR MCML
1st	€ 3,000									
2nd	€ 2,000	**RouteAgents** *University of Thessaly*	5	3 (123.8)	5 (59.0)	6 (34.7)	6 (19.5)	12 (44.2)	1 (84.0)	2 (4.7)
3rd	€ 1,000	**ESCA** *Korea University*	6.29	2 (218.0)	10 (56.7)	13 (16.3)	2 (42.1)	6 (51.3)	10 (47.0)	1 (21.4)
4th	€ 500	**Fu_Do_Fdo** *Fudan University*	7.43	22 (78.8)	9 (56.9)	1 (53.1)	1 (42.8)	1 (77.6)	11 (46.0)	7 (3.1)
5th	€ 500	**MEEMAR** *Ankara Yıldırım Beyazıt University*	7.71	4 (121.4)	2 (61.4)	18 (7.7)	11 (16.8)	7 (50.7)	3 (57.4)	9 (2.8)
6th	Kria KD240	**Fmaximizers** *Delft University of Technology*	8.86	1 (246.6)	19 (46.1)	4 (45.0)	7 (18.4)	19 (33.9)	7 (51.9)	5 (3.6)
7th	Kria KD240	**Novi_CAD** *University of Novi Sad*	9.57	10 (115.6)	7 (58.0)	3 (47.6)	3 (22.4)	3 (54.6)	21 (0.0)	20 (1.7)

Finalists & Honorable Mentions

Rank	Prize	Team	Mean Rank	AMD Mini-ISP	FINN RadioML	FIR Symmetric	FIR Transposed	Rosetta 3D-Render	Rosetta Digit-Rec	VTR MCML
1st	€ 3,000	**AER** *Seoul National & Kyunghee University*	4.29	7 (116.8)	1 (62.5)	2 (53.0)	4 (21.5)	4 (53.4)	9 (50.7)	3 (4.1)
2nd	€ 2,000	**RouteAgents** *University of Thessaly*	5	3 (123.8)	5 (59.0)	6 (34.7)	6 (19.5)	12 (44.2)	1 (84.0)	2 (4.7)
3rd	€ 1,000	**ESCA** *Korea University*	6.29	2 (218.0)	10 (56.7)	13 (16.3)	2 (42.1)	6 (51.3)	10 (47.0)	1 (21.4)
4th	€ 500	**Fu_Do_Fdo** *Fudan University*	7.43	22 (78.8)	9 (56.9)	1 (53.1)	1 (42.8)	1 (77.6)	11 (46.0)	7 (3.1)
5th	€ 500	**MEEMAR** *Ankara Yıldırım Beyazıt University*	7.71	4 (121.4)	2 (61.4)	18 (7.7)	11 (16.8)	7 (50.7)	3 (57.4)	9 (2.8)
6th	Kria KD240	**Fmaximizers** *Delft University of Technology*	8.86	1 (246.6)	19 (46.1)	4 (45.0)	7 (18.4)	19 (33.9)	7 (51.9)	5 (3.6)
7th	Kria KD240	**Novi_CAD** *University of Novi Sad*	9.57	10 (115.6)	7 (58.0)	3 (47.6)	3 (22.4)	3 (54.6)	21 (0.0)	20 (1.7)

Ranking of AI Teams Against Human Teams

Rank	Prize	Team	Mean Rank	AMD Mini-ISP	FINN RadioML	FIR Symmetric	FIR Transposed	Rosetta 3D-Render	Rosetta Digit-Rec	VTR MCML
1st	€ 3,000	AER	4.714	8 (116.8)	1 (62.5)	3 (53.0)	4 (21.5)	4 (53.4)	10 (50.7)	3 (4.1)
2nd	€ 2,000	RouteAgents	6	3 (123.8)	7 (59.0)	9 (34.7)	7 (19.5)	13 (44.2)	1 (84.0)	2 (4.7)
3rd	€ 1,000	ESCA	7.286	2 (218.0)	12 (56.7)	16 (16.3)	2 (42.1)	7 (51.3)	11 (47.0)	1 (21.4)
4th	€ 500	Fu_Do_Fdo	8.429	25 (78.8)	11 (56.9)	2 (53.1)	1 (42.8)	1 (77.6)	12 (46.0)	7 (3.1)
5th	€ 500	MEEMAR	8.571	4 (121.4)	2 (61.4)	21 (7.7)	12 (16.8)	8 (50.7)	4 (57.4)	9 (2.8)
6		Team Opus 5.0	9	7 (116.8)	6 (59.3)	8 (36.1)	5 (20.5)	6 (52.3)	13 (45.1)	18 (1.9)
7		Fmaximizers	9.857	1 (246.6)	21 (46.1)	5 (45.0)	8 (18.4)	21 (33.9)	8 (51.9)	5 (3.6)
8		Novi_CAD	11.143	12 (115.6)	9 (58.0)	4 (47.6)	3 (22.4)	3 (54.6)	24 (0.0)	23 (1.7)
9		CUFO	11.571	19 (110.9)	8 (58.3)	14 (23.4)	13 (16.6)	12 (44.4)	5 (56.1)	10 (2.6)
9		Team GPT 5.6 Sol	11.571	11 (115.7)	5 (59.7)	7 (36.8)	20 (8.1)	23 (25.5)	2 (70.4)	13 (2.2)
19		Team Opus 4.8	16.429	21 (106.5)	23 (31.4)	1 (58.7)	22 (6.2)	20 (34.4)	17 (32.0)	11 (2.5)

Rank on Output Fmax Alone (Including Best Vivado & AI Teams)

Fmax Rank	Score Rank	Team	Mean Rank	amd_mini_isp	finn_radioml	fir_symmetric	fir_transposed	rosetta_3d	rosetta_digit	vtr_mcml
1	1	AER	5.285	8 (409.5)	1 (351.6)	3 (463.4)	4 (377.5)	6 (292.9)	10 (421.4)	5 (73.7)
1	2	RouteAgents	5.285	3 (413.2)	3 (350.1)	8 (443.3)	6 (376.8)	13 (285.4)	1 (459.3)	3 (74.5)
3	3	ESCA	7.714	2 (512.0)	14 (345.9)	16 (422.8)	2 (399.5)	8 (290.4)	11 (419.1)	1 (92.5)
4	4	Fu_Do_Fdo	8.571	26 (372.9)	10 (348.1)	2 (463.8)	1 (402.4)	1 (322.1)	12 (417.7)	8 (72.8)
5	5	MEEMAR	9.857	7 (410.5)	4 (349.9)	22 (414.2)	11 (374.4)	10 (289.6)	5 (428.8)	10 (72.4)
5	6	Team Opus 5	9.857	10 (408.5)	7 (348.8)	7 (445.2)	5 (377.4)	7 (292.6)	13 (416.3)	20 (71.4)
7	7	Fmaximizers	10.142	1 (560.2)	22 (335.8)	5 (455.4)	8 (375.8)	21 (274.1)	8 (424.3)	6 (73.4)
8	-	Best Vivado	11.428	12 (407.8)	17 (342.1)	21 (415.3)	9 (374.8)	4 (295.4)	15 (410.0)	2 (74.5)
9	8	Novi_CAD	12.142	15 (406.0)	11 (346.6)	4 (458.3)	3 (379.1)	3 (297.2)	25 (367.0)	24 (71.1)
16	9	Team GPT 5.6 Sol	14.714	18 (404.2)	12 (346.4)	8 (443.3)	21 (363.6)	25 (262.6)	3 (439.0)	16 (71.6)
22	19	Team Opus 4.8	17.857	24 (396.2)	24 (318.0)	1 (471.5)	23 (361.9)	22 (273.4)	19 (401.4)	12 (71.9)

Thank You to All Contest Participants!

- Acknowledgements
  - FPL organizers (Dirk Stroobandt & Evangeline Young)
  - AMD University Program (Mario Ruiz Noguera)
  - AMD RAD (Kevin Cheng)

- This is only the beginning of agentic FPGA optimization!
- Look for publications from contest teams in upcoming conferences
- All 5 finalists have (or will) open source their submissions

- Summary slides, repos, and 3-minute videos will be on contest website

https://xilinx.github.io/fpl26_optimization_contest

COPYRIGHT AND DISCLAIMER

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change and may be rendered inaccurate releases, for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

THIS INFORMATION IS PROVIDED 'AS IS." AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.


AMD

Contest Timeline 2026

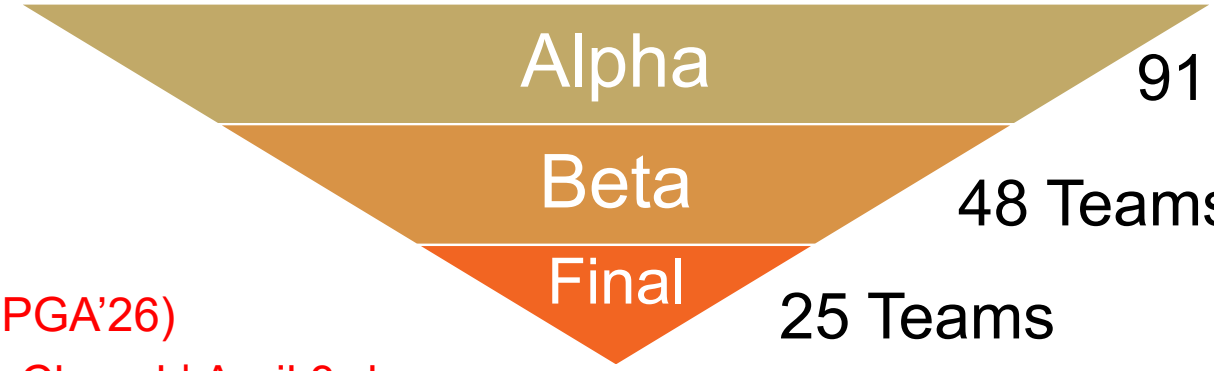

91 Teams

48 Teams

25 Teams

Announce | Mid February (@ FPGA'26)

Registration Closed | April 3rd

Alpha submission | May 5th

Beta Submission | July 13th

Final Submission | August 10th

Winners Announced @ FPL | September 9th

2026

Feb Mar Apr May Jun Jul Aug Sep Oct

Papers Accepted Notification – End May

FPL Papers Due – End March