

# Inverter Redistribution through Self-Dual and Self-Anti-Dual Function Transformation

Jingren Wang<sup>1</sup> Guangyu Hu<sup>2</sup> Shiju Lin<sup>1</sup> Hongce Zhang<sup>1</sup>

<sup>1</sup>MICS, HKUST(GZ), Guangzhou, China

<sup>2</sup>IIP-MICS, HKUST, Hong Kong, China

May 30, 2025



微电子学域  
**MICROELECTRONICS THRUST**  
功能枢纽  
FUNCTION HUB



- 1 Motivation
- 2 Background
- 3 Proposed Method
- 4 Experiments
- 5 Conclusion

# The Problem: Responsibility on workflow

Imaging a scenario that Alice work on pure logic synthesis and Bob works on mapping. One day Bob complained to Alice: “Why you have obviously redundant inverters and why you don’t have a distribution guideline for them?” Both of them blame each other for “Never asked.”

- This was originally a workflow guideline problem rather than an optimization technique.

## Key insight:

Give inverter distribution guideline *before* mapping can standardize the workflow and clear the responsibility.



EDA Engineering Life.

# The Problem: Lack of control due to separate heuristics

Actual drawback on inverter control due to different heuristic optimization methods.

- Inverter number might silently increase without monitoring.
- Optimizations are independent of each other. They make arbitrary different decisions on inverters in different scenarios.

## Key insight:

An “early stopping” is needed also for inverters in such scenario.

```
(base) → Downloads abc -c "source /Users/jingrenwang/Downloads/Github/abc/abc.rc; read_truth -xf ./2026_IWLS_Contest_Benchmarks_260121/ex270.truth;d;d; strash; &get; &st; &dc4; &put; share; strash; &get; &st; &dc4; &put; share; compress2rs; compress2rs; compress2rs; compress2rs; print_stats;"
ABC command line: "source /Users/jingrenwang/Downloads/Github/abc/abc.rc; read_truth -xf ./2026_IWLS_Contest_Benchmarks_260121/ex270.truth;d;d; strash; &get; &st; &dc4; &put; share; strash; &get; &st; &dc4; &put; share; compress2rs; compress2rs; compress2rs; compress2rs; print_stats;"
```

```
Loading resource file "/Users/jingrenwang/Downloads/Github/abc/abc.rc".
NegEdges = 428 lev = 10
(base) → Downloads abc -c "source /Users/jingrenwang/Downloads/Github/abc/abc.rc; read_truth -xf ./2026_IWLS_Contest_Benchmarks_260121/ex270.truth;d;d; strash; &get; &st; &dc4; &put; share; strash; &get; &st; &dc4; &put; share; compress2rs; compress2rs; compress2rs; compress2rs; compress2rs; print_stats"
ABC command line: "source /Users/jingrenwang/Downloads/Github/abc/abc.rc; read_truth -xf ./2026_IWLS_Contest_Benchmarks_260121/ex270.truth;d;d; strash; &get; &st; &dc4; &put; share; strash; &get; &st; &dc4; &put; share; compress2rs; compress2rs; compress2rs; compress2rs; compress2rs; print_stats;"
```

```
Loading resource file "/Users/jingrenwang/Downloads/Github/abc/abc.rc".
NegEdges = 430 lev = 10
(base) → Downloads █
```

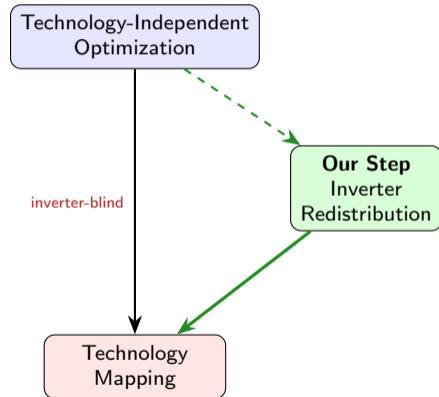
# The Problem: Inverter Blindness in AIG Optimization

Given facts:

- AIG-based synthesis treats inverters as **complemented edge attributes**
- Cost functions (level, node count) **ignore** inverters entirely
- Inverter distribution evolves **randomly** during optimization
- At technology mapping → inverters **suddenly matter** for delay & area



## Key insight:

Can we reshape inverter distribution *before* mapping to standardize the working flow, remove redundancies and improve results?





# Minimal Reproducible Example

Network structure visualized by ABC  
The network contains 6 logic nodes and 0 latches.



(a)

Network structure visualized by ABC  
The network contains 6 logic nodes and 0 latches.



(b)


- Both sides map the **same function** from AIG to standard cells; dashed edges are complemented
- Left: baseline; Right: our method identifies a **self-dual subfunction**, complements all cut inputs + root fanout
- Result: fewer complemented edges on the **critical path** → better delay

# Self-Dual and Self-Anti-Dual Functions

## Self-Dual

$$f(x_1, \dots, x_n) = \bar{f}(\bar{x}_1, \dots, \bar{x}_n)$$


Complementing *all* inputs and the output preserves the function.



## Self-Anti-Dual

$$f(x_1, \dots, x_n) = f(\bar{x}_1, \dots, \bar{x}_n)$$

Complementing all inputs alone preserves the function.



# Self-Dual and Self-Anti-Dual Functions

## Self-Dual Example: 3-input Majority

$$\text{maj}(a, b, c) = ab + bc + ac$$

- Complement all inputs:

$$\text{maj}(\bar{a}, \bar{b}, \bar{c}) = \bar{a}\bar{b} + \bar{b}\bar{c} + \bar{a}\bar{c}$$

- Negate output (De Morgan):

$$\overline{\bar{a}\bar{b} + \bar{b}\bar{c} + \bar{a}\bar{c}}$$

$$= (a + b)(b + c)(a + c) = ab + bc + ac \checkmark$$

## Self-Anti-Dual Example: 2-input XOR

$$a \oplus b = a\bar{b} + \bar{a}b$$

- Complement all inputs:

$$\bar{a} \oplus \bar{b} = \bar{a}b + a\bar{b}$$

- Equals original:

$$\bar{a} \oplus \bar{b} = a \oplus b \checkmark$$




## Three-step process per node:

- 1 **Detect** — simulate **truth tables** to identify self-dual / self-anti-dual subfunctions on tree cuts
- 2 **Evaluate** — apply policy to decide if redistribution is beneficial
- 3 **Transform** — complement cut leaves (and root fanout for self-dual)

**Key property:** The AND nodes in the original AIG are **never modified**.

# Detection Algorithm



- $f = \bar{f}^d$  (**truthEq**): phase differs  $\rightarrow$  Self-Dual; same phase  $\rightarrow$  Self-Anti-Dual
- $f = \neg\bar{f}^d$  (**isComp**): phase matches  $\rightarrow$  Self-Dual; phase differs  $\rightarrow$  Self-Anti-Dual
- Masking vector `uMask` handles cuts with fewer than 5 leaves


# AIG Transformation

## Self-Anti-Dual transformation:



- Complement edges at **cut leaves** (only those feeding **internal nodes**)

## Self-Dual transformation:

- Same as self-anti-dual, **plus** complement **fanout edges** of the root node



# Tree Cuts vs. DAG Cuts



DAG cut leads to equivalence failure

5 tree cuts (dashed) of node *po146*

Tree cuts are safe and independent. DAG cuts are **prohibited**.

# Node and Cut Selection

## Which nodes to process:

- 1 DAG nodes and tree nodes on the **primary graph** (no choices)
- 2 **Choice candidates** in the choices network
  - Choice representatives are **skipped** — the primary structure is already near-optimal

## Which cuts to use:

- Only **tree cuts** — guaranteed inside MFFC (Maximum Fanout Free Cone)
- DAG cuts are excluded due to equivalence **risk**(Opportunity?)

# Redistribution Policy

## Critical Path

- Redistribute only if  $\# \text{complemented edges} > \# \text{regular edges}$  on critical path
- $\Rightarrow$  **Net reduction** of inverters on critical paths
- Targets **delay improvement**

## Non-Critical Path

- Redistribute if inverter count  $\geq$  half of all edges
- Allows **zero-gain** redistribution
- Creates favorable conditions for mapping

## Zero-Gain Redistribution

- Tree-structured XORs (self-anti-dual) with inverter ratio = 0.5
- No local improvement, but **enables inverter absorption** during mapping, e.g., XNOR2 (44.4 ps) over XOR2 (53.5 ps)
- Cumulative delay improvement: **3–4%** on arithmetic circuits

# Experimental Setup

- **Platform:** Apple M4 Pro, 24 GB, macOS Tahoe 26.2
- **Framework:** ABC — entirely on AIG representations
- **Verification:** CEC (combinational equivalence checking) after redistribution
- **Benchmark:** EPFL combinational suite (random control + arithmetic)

## Synthesis flows compared:

**OptFlow** `st; &get; &if -g; &dch; &put`

**Baseline** `(r2rs)8 (OptFlow)8 map; topo; stime`

**Ours** `(r2rs)8 (OptFlow)8 rd_inv -s -c map; topo; stime`


**Key:** Only a *single* round of `rd_inv` needed — AND nodes unchanged, level slack preserved.

# Results: Arithmetic Logic

Benchmark	Delay Diff	% Improvement	Gate Diff	INV_X1 Diff
<i>sqrt</i>	-5289.82 ps	<b>3.86%</b>	+429	-429
<i>square</i>	-17.69 ps	<b>1.60%</b>	-29	-15
<i>hyp</i> <sup>16</sup>	-10437.94 ps	<b>1.46%</b>	-361	-1959
<i>voter</i>	-13.93 ps	<b>1.09%</b>	-22	-10
<i>multiplier</i>	-26.84 ps	<b>0.77%</b>	-175	-20
<i>log2</i>	-73.87 ps	<b>0.85%</b>	-101	+17
<i>div</i>	-61.47 ps	<b>0.10%</b>	-86	+8
<i>sin</i>	-1.55 ps	<b>0.04%</b>	-4	+34

- Random control logic: **no change** — inverter distribution unaffected
- **No degradation on any benchmark**
- Arithmetic circuits are **highly sensitive** — structured XOR chains and self-dual subfunctions
- *sqrt*: 3.86% delay reduction with only +29 area cost

# Analysis: Self-Dual vs. Self-Anti-Dual Distribution



SD/SAD function count



Critical vs. non-critical

- **Self-anti-dual** (tree XORs) **vastly outnumbers** self-dual — dominant contributor
- Self-dual: found exclusively on **critical paths** → direct delay reduction
- Self-anti-dual: found exclusively on **non-critical paths** → enables inverter absorption

# Conclusion

## Summary

- Proposed a **delay-oriented inverter redistribution** preprocessing step using self-dual and self-anti-dual function properties
- Operates on AIG tree cuts — **preserves original structures of AND nodes**
- Equivalence guaranteed by tree-cut isolation and compositional closure

## Key Results

- **No degradation** on any benchmark
- Arithmetic benchmarks: up to **3.86% delay reduction** (*sqrt*)
- Average delay reduction: **0.49%**
- Zero-gain redistribution on non-critical paths enables inverter absorption during mapping

# Thank You

Jingren Wang

`jwang929@connect.hkust-gz.edu.cn`