

# CBTune: Contextual Bandit Tuning for Logic Synthesis

Fangzhou Liu ♠ Zehua Pei ♠ Ziyang Yu ♠ Haisheng Zheng ♡  
Zhuolun He ♠♡ Tinghuan Chen ♣ Bei Yu ♠

♠ The Chinese University of Hong Kong

♡ Shanghai Artificial Intelligence Laboratory

♣ The Chinese University of Hong Kong, Shenzhen

Mar. 25, 2024



# Outline

1 Introduction

2 CBTune

3 Evaluation

# Introduction

**Synthesis Flow** is a set of synthesis transformations that apply iteratively to the AIG.

- Use typical ABC commands like *balance(b)*, *rewrite(rw)*, *refactor(rf)*, *resubstitute(rs)*, ...

**Synthesis Flow** is a set of synthesis transformations that apply iteratively to the AIG.

- Use typical ABC commands like *balance(b)*, *rewrite(rw)*, *refactor(rf)*, *resubstitute(rs)*, ...

## Problems

- Exponential solution space.
- Same synthesis flow works differently across designs.

# Introduction

**Synthesis Flow** is a set of synthesis transformations that apply iteratively to the AIG.

- Use typical ABC commands like *balance(b)*, *rewrite(rw)*, *refactor(rf)*, *resubstitute(rs)*, ...

## Problems

- Exponential solution space.
- Same synthesis flow works differently across designs.

## Limitations of Recent Works

- Timing-intensive and resource-demanding.
- Lack of transferability; Training is hard to converge.

# Bandit Model Definition

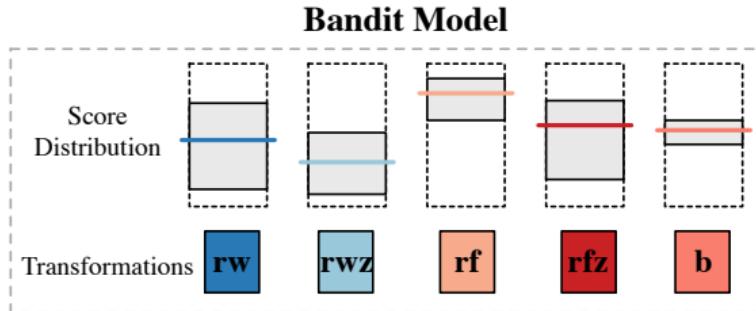


Figure: Bandit Model for Synthesis Flow Generation.

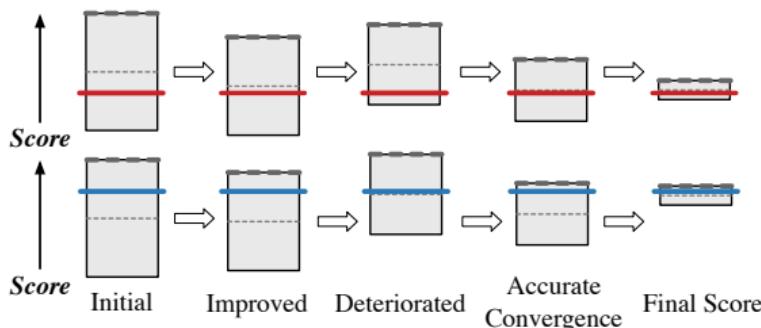


Figure: Score Iterations for Each Arm in Bandit Model.

- **Bandit Model**

A Decision-Making Framework focused on maximizing rewards by balancing exploration and exploitation.

- **Upper Confidence Bound**

$$UCB_i(t) = \hat{x}_i(t) + \sqrt{\frac{2 \ln(t)}{T_{i,t}}} \quad (1)$$

# CBTune

# CBTune Framework

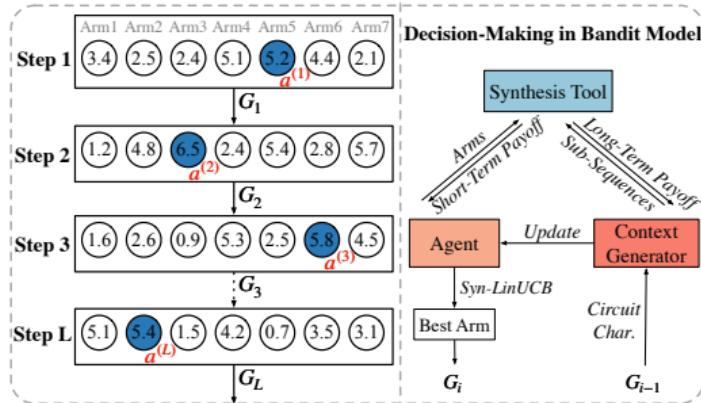


Figure: CBTune Framework Overview.

- **Action Space**

$\mathcal{A} = \{resub, resub-z, rewrite, rewrite-z, refactor, balance\}$ , denoted as  $\mathcal{A} = \{a_1, a_2, a_3, \dots, a_6\}$ .

- **Reward**

The short-term payoff (a single-step) of each arm: The number of AIG nodes or 6-LUTs.

- The highest-scoring one  $a^{(i)}$  is selected at each step, and the final synthesis flow has the ordered sequence of transformations  $a^{(1)}, a^{(2)}, \dots, a^{(L)}$ .

## Contextual Generator

- Observe features of  $a \in \mathcal{A}$ :  $\mathbf{x}_{t,a} = [\mathbf{x}_a^c, \mathbf{x}_{t,a}^l] \in \mathbb{R}^d$ ;

Feature	Description	Example
Circuit Characteristics ( $\mathbf{x}^c$ )	AIG information extracted by applying $a_j^{(i)}$ ( $j \leq n$ ) to $G_{i-1}$ using yosys and ccirc.	#Number of wires/cells/nots, #Maximum delay, #Number of combinational nodes, #Number of high degree comb, #Fanout distribution .....
Long-term Payoff of the Arm ( $\mathbf{x}^l$ )	Random DSE result: Employ "arm" $a^{(i)}$ as the first transformation to produce $m$ random subsequences of length $l$ , and then utilize these subsequences to obtain synthesis results.	Arm: <i>rewrite (rw)</i> ; $l = 5$ ; $m = 3$ ; $\{\text{rw}, \text{rw}, \text{rfz}, \text{rf}, \text{rs}\} \rightarrow \text{Nodes: 28010, Level: 66}$  Arm: <i>refactor (rf)</i> ; $l = 4$ ; $m = 2$ ; $\{\text{rf}, \text{b}, \text{rf}, \text{rw}\} \rightarrow \text{Nodes: 28324, Level: 67}$

## Agent (Syn-LinUCB) [1]

- Update hyperparameter  $\alpha$  by  $\alpha = 1.0 + \sqrt{\frac{\log(2.0/\delta)}{s_a}}$  (2)

- Update the decision parameter by  $\boldsymbol{\theta}_a = \mathbf{A}_a^{-1} \mathbf{b}_a$  (3)

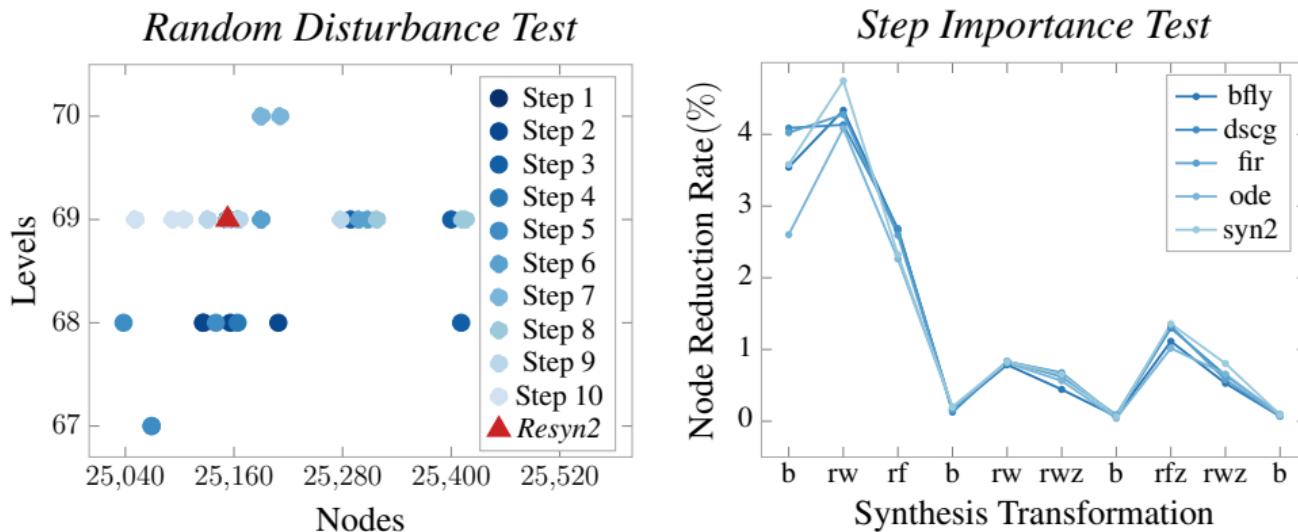
- Calculate the weighted context  $\mathbf{x}_{t,a}^w = \mathbf{x}_{t,a} \mathbf{w}$  (4)

- Update score by  $p_{t,a} = \boldsymbol{\theta}_a^\top \mathbf{x}_{t,a}^w + \alpha \sqrt{\mathbf{x}_{t,a}^{w\top} \mathbf{A}_a^{-1} \mathbf{x}_{t,a}^w}$  (5)

- Update the parameters  $\mathbf{A}_{a_t}$  and  $\mathbf{b}_{a_t}$  of the chosen arm  $a_t$  by  $\mathbf{A}_{a_t} = \mathbf{A}_{a_t} + \mathbf{x}_{t,a_t} \mathbf{x}_{t,a_t}^\top$ ,  $\mathbf{b}_{a_t} = \mathbf{b}_{a_t} + r_{t,a} \mathbf{x}_{t,a_t}$  (6)

[1] Lihong Li et al. (2010). "A Contextual-bandit Approach to Personalized News Article Recommendation". In: *The Web Conference*.

# Optimization Techniques



**Figure:** Analysis of Transformation Effectiveness in the Synthesis Flow.

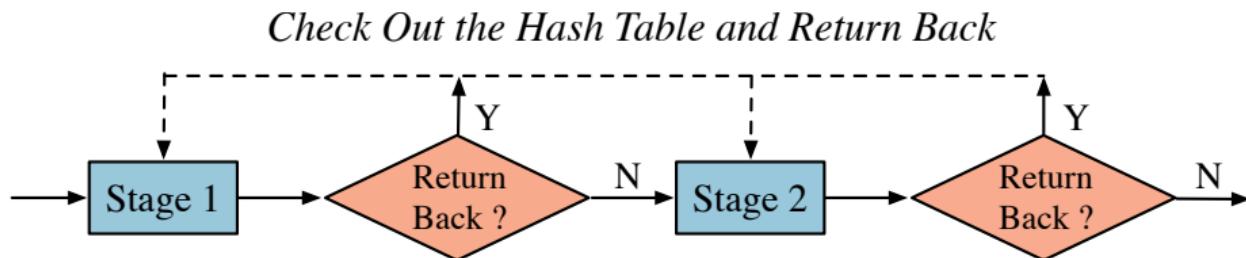


Figure: The Return-Back Mechanism in CBTune.

- ① Tuning Iteration Count; Early Stop
- ② Return-back Mechanism

# Evaluation

# AIG Node Optimization

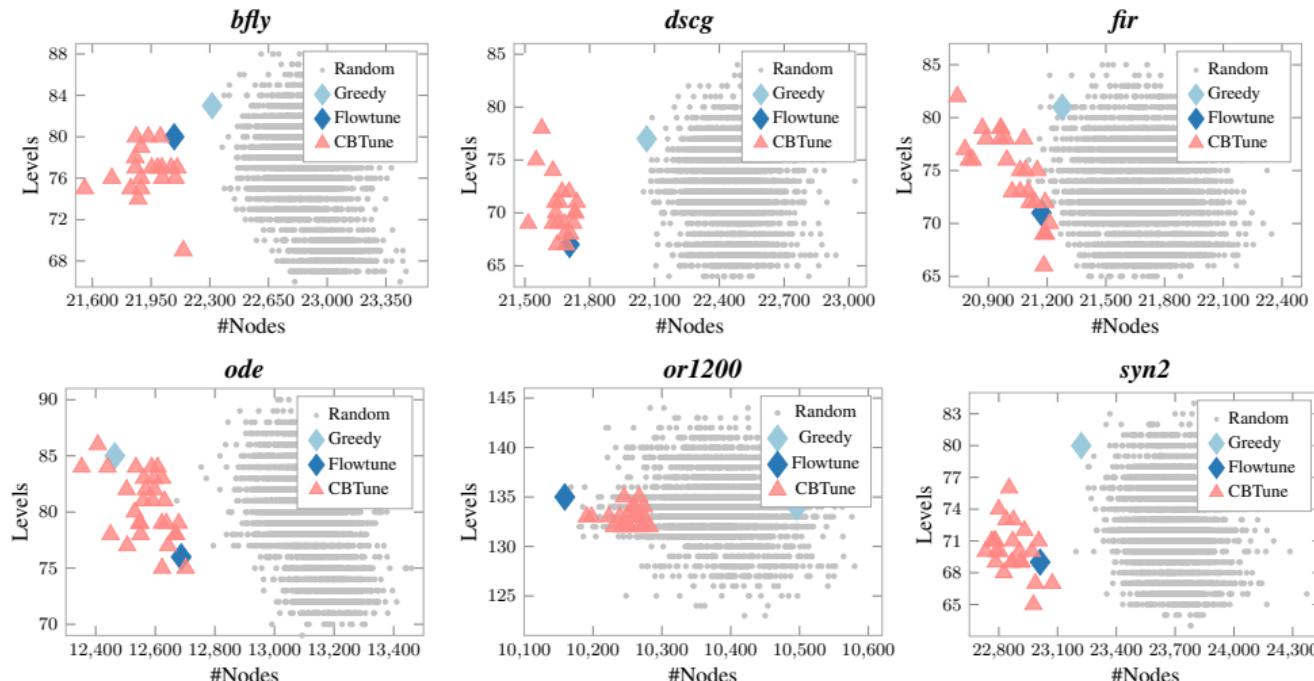


Figure: Performance Comparison between CBTune & FlowTune.

# 6-LUTs Optimization



Benchmark	Initial	Greedy	FlowTune <sup>[2]</sup>		CBTune		
	#LUTs	#LUTs	#LUTs	$\tau(m)$	#LUTs	#LUTs	$\tau(m)$
<i>bfly</i>	9019	8269	8216	76.47	<b>7962</b>	8086.03	<b>29.63</b>
<i>dscg</i>	8534	8313	8302	77.15	<b>7981</b>	8119.84	<b>30.44</b>
<i>fir</i>	8646	8385	8094	74.23	<b>7820</b>	7977.38	<b>27.6</b>
<i>ode</i>	5244	5316	5096	34.83	<b>4920</b>	5046.71	<b>17.32</b>
<i>or1200</i>	2776	2748	2747	20.08	<b>2731</b>	2754.07	<b>15.62</b>
<i>syn2</i>	8777	8669	8603	81.33	<b>8234</b>	8360.53	<b>31.67</b>
GEOMEAN	6631.20	6464.69	6364.89	54.04	<b>6166.39</b>	6271.82	<b>24.48</b>
Ratio Avg.	1.000	0.975	0.960	1.000	<b>0.930</b>	0.946	<b>0.453</b>

Table: Performance Comparison Between CBTune & FlowTune.

<sup>[2]</sup> Walter Lau Neto et al. (2022). “FlowTune: End-to-End Automatic Logic Optimization Exploration via Domain-specific Multi-armed Bandit”. In: IEEE TCAD.

# 6-LUTs Optimization

Benchmark	Initial	Greedy	DRiLLS <sup>[3]</sup>		RL4LS*		CBTune	
	#LUTs	#LUTs	#LUTs	$\tau(m)$	#LUTs	$\tau(m)$	#LUTs	$\tau(m)$
<i>max</i>	721	697	694	32.58	687.8	54.34	<b>684.25</b>	<b>6.01</b>
<i>adder</i>	249	<b>244</b>	<b>244</b>	24.05	<b>244</b>	10.05	<b>244</b>	<b>5.97</b>
<i>cavlc</i>	116	115	112.2	26.02	111.3	3.22	<b>111</b>	<b>2.37</b>
<i>ctrl</i>	29	<b>28</b>	<b>28</b>	24.25	<b>28</b>	2.85	<b>28</b>	<b>0.59</b>
<i>int2float</i>	47	46	42.6	21.7	42.3	2.81	<b>40</b>	<b>2.76</b>
<i>router</i>	73	67	70.1	22.01	69.5	3.07	<b>68.11</b>	<b>2.32</b>
<i>priority</i>	264	146	<b>133.4</b>	23.32	142.9	5.9	138.86	<b>3.41</b>
<i>i2c</i>	353	291	292.1	25.17	289.32	7.55	<b>283.11</b>	<b>3.61</b>
<i>sin</i>	1444	1451	1441.5	51.15	<b>1438</b>	20.1	1441.67	<b>9.71</b>
<i>square</i>	3994	3898	3889.4	130	3889	72.88	<b>3882.11</b>	<b>25.99</b>
<i>sqrt</i>	8084	4807	4708	147.64	4685.3	196.15	<b>4607</b>	<b>36.51</b>
<i>log2</i>	7584	7660	7583.6	198.6	7580.1	125.28	<b>7580</b>	<b>41.27</b>
<i>multiplier</i>	5678	5688	5678	180.84	<b>5672</b>	187.81	5679.75	<b>29.08</b>
<i>voter</i>	2744	1904	1834.7	84.43	<b>1678.1</b>	330.48	1682.25	<b>11.46</b>
<i>div</i>	23864	4205	7944.4	259.75	7807.1	482	<b>4180.91</b>	<b>25.58</b>
<i>mem_ctrl</i>	11631	<b>10144</b>	10527.6	229.33	10309.7	1985.84	10242.57	<b>45.81</b>
GEOMEAN	926.59	732.69	753.49	59.48	748.34	34.54	<b>712.83</b>	<b>8.37</b>
Ratio Avg.	1.000	0.791	0.813	1.000	0.808	0.581	<b>0.769</b>	<b>0.141</b>

**Table:** Performance Comparison Between CBTune & NN-Enhanced RL. \* Last10 in RL-PPO-Pruned<sup>[4]</sup>.

<sup>[3]</sup> Abdelrahman Hosny et al. (2020). “DRiLLS: Deep Reinforcement Learning for Logic Synthesis”. In: Proc. ASPDAC.

<sup>[4]</sup> Guanglei Zhou and Jason H Anderson (2023). “Area-Driven FPGA Logic Synthesis Using Reinforcement Learning”. In: Proc. ASPDAC.

**THANK YOU!**