

A Hybrid ECO Detailed Placement Flow for Improved Reduction of Dynamic IR Drop

Andrew B. Kahng

University of California, San Diego
La Jolla, California, USA
abk@ucsd.edu

Bodhisatta Pramanik

University of California, San Diego
La Jolla, California, USA
bopramanik@ucsd.edu

Mingyu Woo

University of California, San Diego
La Jolla, California, USA
mwoo@ucsd.edu

ABSTRACT

With advanced semiconductor technology progressing well into sub-7nm scale, voltage drop has become an increasingly challenging issue. As a result, there has been extensive research focused on predicting and mitigating dynamic IR drops, leading to the development of IR drop engineering change order (ECO) flows – often integrated with modern commercial EDA tools. However, these tools encounter QoR limitations while mitigating IR drop. To address this, we propose a hybrid ECO detailed placement approach that is integrated with existing commercial EDA flows, to mitigate excessive peak current demands within power and ground rails. Our proposed hybrid approach effectively optimizes peak current levels within a specified “clip” – complementing and enhancing commercial EDA dynamic IR-driven ECO detailed placements. In particular, we: (i) order instances in a netlist in decreasing order of worst voltage drop; (ii) extract a clip around each instance; and (iii) solve an integer linear programming (ILP) problem to optimize instance placements. Our approach optimizes dynamic voltage drops (DVD) across ten designs by up to 15.3% compared to original conventional flows, with similar timing quality and 55.1% less runtime.

ACM Reference Format:

Andrew B. Kahng, Bodhisatta Pramanik, and Mingyu Woo. 2024. A Hybrid ECO Detailed Placement Flow for Improved Reduction of Dynamic IR Drop. In *Great Lakes Symposium on VLSI 2024 (GLSVLSI '24)*, June 12–14, 2024, Clearwater, FL, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3649476.3658727>

1 INTRODUCTION

As technology nodes scale down, IR drop becomes a critical, blocking step in design signoff. Below 7nm, high-end silicon designs, particularly those with aggressive clock targets, tend to consume more power and occupy more area. For correct timing analysis and closure of these designs, dynamic IR drop simulations are performed and a *timing derate factor* is then applied to all instances affected by IR drops. Notably, as dynamic IR drop increases, the *timing derate factor* increases, potentially leading to slower timing (*fmax*) – or greater power and area (from upsizing) – in the design outcome [16]. Therefore, to optimize power, performance, area and

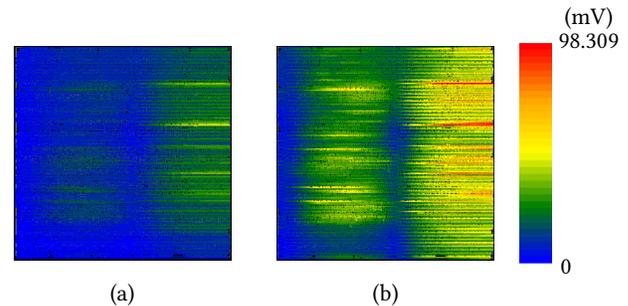


Figure 1: Discrepancy between (a) static IR drop and (b) dynamic IR drop heatmaps for *aes_cipher_top* on 7nm technology, generated with a commercial tool.

cost (PPAC) in these advanced designs, it is crucial to effectively mitigate dynamic IR drops.

To tackle the challenge of dynamic IR drop mitigation, designers often strengthen the power delivery network (PDN) to supply additional current with lower resistance to instances with the worst IR drops. However, this method is not always viable for high-end designs. The PDN consumes routing resources, and adding more PDN can fail when there is a shortage of these resources [16]. Another strategy is to consider static IR drop during various stages (power-driven placement, clock tree synthesis (CTS), routing) of the physical design flow. However, there can be significant discrepancies between static and dynamic IR drop, as illustrated in Figure 1. Static IR drop models cannot fully capture the complexities and transient characteristics of dynamic IR drops, risking both underestimations and overdesign before the layout is final.

An alternative approach to mitigate dynamic IR drop involves applying detailed placement using an ECO flow in a post-route opt (PRO) stage. [7] proposes an ECO detailed placement method to address dynamic IR drop at the PRO stage. Their approach employs two integer-linear programming (ILP)-based strategies: (i) adjusting the *y*-axis placement to minimize the peak current waveforms for power and ground rails, and (ii) minimizing *x*-displacement within each row. However, the approach can potentially lead to suboptimal solutions since the *x* and *y* coordinates are not optimized concurrently. Commercial tool flows (an example is [12]) offer dynamic IR drop-driven ECO detailed placement features. However, current tools only imperfectly mitigate dynamic IR drop, leaving quality of results (QoR) on the table. With this as motivation, we propose a *hybrid* ECO detailed placement approach that integrates an ILP-based detailed placer with a commercial EDA tool, and effectively mitigates peak current demands within power and ground rails. Our contributions are summarized as follows.



This work is licensed under a Creative Commons Attribution International 4.0 License.

GLSVLSI '24, June 12–14, 2024, Clearwater, FL, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0605-9/24/06
<https://doi.org/10.1145/3649476.3658727>

- We propose an ILP-based detailed placer that can be used at the post-route opt (PRO) stage. Our approach uses a hybrid ECO methodology that further optimizes IR drops beyond the capabilities of traditional ECO detailed placement flows.
- To achieve scalability, we extract “clips” and formulate an integer linear programming (ILP) problem to minimize cell displacement and maximum peak currents for each power and ground rail. In contrast to [7], our method concurrently optimizes both x and y coordinates within a clip.
- Our proposed approach is designed to complement and integrate seamlessly with existing ECO detailed placement flows of commercial EDA tools. This compatibility makes it particularly valuable in addressing real-world design challenges.
- Experimental results show that our hybrid approach achieves up to 15.3% reduction in DVD compared to conventional flows, with similar timing quality and 55.1% less runtime.¹

This paper is organized as follows. Section 2 reviews related works. Section 3 discusses a linearity assumption for the demand current waveform. Section 4 explains our hybrid ECO detailed placement approach, and Section 5 gives details of our ILP formulation. Section 6 presents experimental results and Section 7 concludes our paper.

2 RELATED WORKS

Placing instances to mitigate IR drop is known to be an NP-hard problem [3]. Hence, various heuristics have been proposed in the literature. We discuss two categories of previous works: IR drop prediction and IR drop optimization.

IR drop prediction. [6] proposes an XGBoost-based machine learning model to predict the static IR drop at each power node in a design. This avoids rerunning the IR drop tool for incremental changes in the placement or PDN. [4] proposes an XGBoost-based machine learning model to predict dynamic IR drop, thus speeding up dynamic IR drop analysis in each ECO iteration. [11] improves upon the work of [4] by proposing a CNN-based machine learning model to predict dynamic IR drop. Compared to previous approaches, the authors *preprocess* design dependent information before it is fed to the ML model. A recent work [7] improves upon [6], [4] and [11], and predicts both static and dynamic IR drop. The authors use feature engineering and introduce a Random Forest (RF)-based regression model for IR drop prediction.

IR drop optimization. [8] proposes an IR drop-aware global placer, which considers the worst voltage drop across all cells in the design. The authors approximate this voltage drop with a log-sum-exp function and integrate it with other placement objectives (wirelength, congestion and timing) in the APlace placer [9]. [3] improves upon [8] by introducing *power spreading forces* in the analytical placement framework of NTUPlace3 [2]. The power spreading forces push cells with voltage drop violations to new locations that improve the voltage drop. [1] proposes an IR drop mitigation approach that is applicable at the CTS stage. The authors decompose the peak current minimization problem into many smaller subproblems. Solving each subproblem reduces the local peak current of each via-stack in the on-chip PDN. By greedily scheduling useful skews from unused

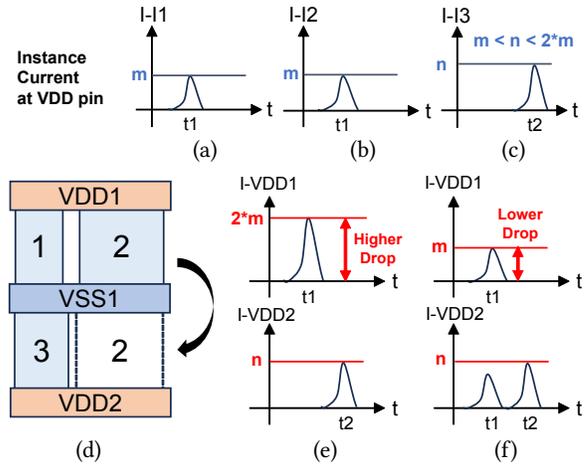


Figure 2: Illustration of IR drop aware detailed placement. Measured demand current waveform at the VDD pin of instances (a) 1, (b) 2, and (c) 3. (d) Example of cell movement: instance 2 moves from first row to second row. (e) Current waveforms at VDD1 and VDD2 before cell movement. (f) Current waveforms at VDD1 and VDD2 after cell movement. The cell movement reduces the IR drop.

timing slacks, the authors demonstrate that their approach can substantially reduce peak IR drop and peak current. [5] proposes a dynamic programming based detailed placer that maximizes power staple insertions to reduce IR drop. The authors show that their approach can achieve significant reductions in IR drop with similar WNS when compared to a standard flow. [7] proposes a machine learning (ML)-based IR drop-aware detailed placer that uses predicted IR drop hotspots to guide cell movements. The authors show that their approach can reduce IR drop without timing degradation.

3 LINEARITY OF THE DEMAND CURRENT WAVEFORM

[7] introduces the concept of minimizing the maximum peak current values using detailed placement. Following [7], in this work, we assume that after moving an instance by one row either upward or downward, the total demand current waveform [13] seen by power and ground rails remains unchanged. Figure 2 illustrates the overarching concept of *current linearity* in this context. The figure is divided into the following parts for clarity.

- Parts (a), (b) and (c) show examples of current waveforms measured at VDD pin of instances 1, 2 and 3, respectively. In these examples, instances 1 and 2 reach their peak currents, denoted as m , at timestamp t_1 . Instance 3 reaches its peak current, denoted as n , at a later timestamp t_2 . Here, we assume that $m < n < 2 * m$ and $t_1 < t_2$.
- Part (d) shows an example of placed instances. In this placement, instances 1 and 2 are connected to the VDD1 power rail, and instance 3 is connected to the VDD2 power rail.
- Part (e) shows a scenario prior to any cell movement. Here the VDD1 rail has a combined maximum demand peak current of $2 * m$ at timestamp t_1 . This is due to the accumulation of currents from instances 1 and 2. Meanwhile, the VDD2 rail

¹This work is not intended to be, and should not be considered as, benchmarking. We do not perform any benchmarking of commercial EDA tools.

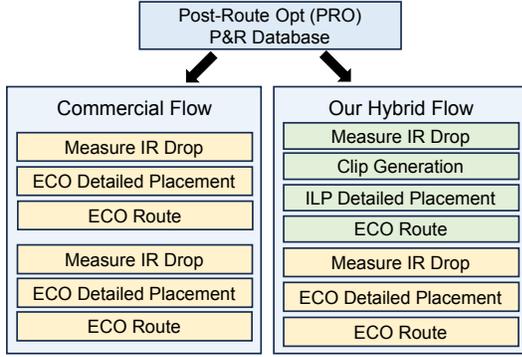


Figure 3: Conventional ECO DP flow and our hybrid DP flow.

shows a maximum demand peak current of n at timestamp $t2$, solely attributed to instance 3.

- Part (f) depicts the scenario after cell movement. The VDD1 rail now shows a maximum demand peak current of m at timestamp $t1$ because only instance 1 remains on this rail. On the other hand, the VDD2 rail still has a maximum demand peak current of n at timestamp $t2$. This is due to the accumulation of currents from instances 2 and 3. Notably, the maximum peak current on the VDD2 rail remains the same as before (n), since the peak currents of instances 2 and 3 occur at different timestamps ($t1$ and $t2$, respectively).

Our proposed detailed placement methodology is based on the above concepts, and focuses on minimizing the peak demand current over all power and ground rails. Cell moves during the detailed placement are limited to enable an assumption of current linearity.

4 PROPOSED METHODS

In this section, we introduce our hybrid ECO detailed placement flow (Figure 3). We first define three terminologies.

- **Clip:** A *clip* is a window of layout extracted from the original layout of a netlist.
- **PDN zone:** A *PDN zone* is a region that is formed by adjacent power and ground connections on the bottom metal layer.²
- **Clip overlap violations:** Clip overlap violations refer to geometric overlaps between multiple clips. More details are presented in a later subsection.

Overall flow of our hybrid ECO detailed placer. The details of our proposed hybrid detailed placement flow are given in Algorithm 1. Our flow comprises two steps: (i) ILP-based detailed placement optimization and (ii) commercial ECO detailed placement. Our flow starts by measuring the dynamic IR drop at every cell instance (Line 1). We next determine the filtered instance set FN according to the input DVD threshold Th , and effective DVD [13] of all instances (Line 2). We also define the clips set C and PDN zones set P to avoid clip overlap violations as shown in Figure 5 (Lines 3–4). If an instance i in the routed layout has a *DVD hotspot*, i.e., i has greater effective DVD than the input DVD threshold Th , we generate a clip c at the DVD hotspot, using i 's location and

²PDN supplies the power from the top metal to the bottom. As a result, the connections in the bottom metal layer induce IR drop independent PDN zones. More details are presented in the *Clip and PDN zones extraction* subsection.

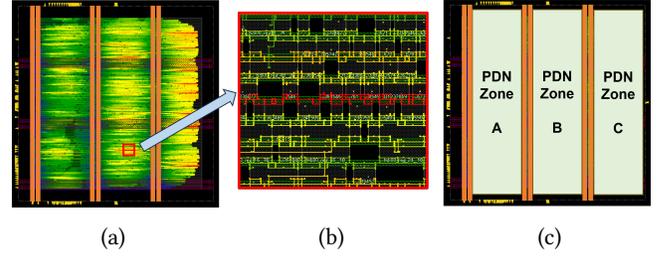


Figure 4: Illustration of *clip* and *PDN zones* from *mpeg2_top*. (a) Heatmap of dynamic IR drop with PDN; (b) an extracted clip from the red square in (a); and (c) *PDN zones* defined in (a).

c 's dimensions, cw and ch (Line 7). If c does not have clip overlap violations, we update the clip set C (Lines 8–9). After all clips are generated, we run the ILP-based ECO DP with inputs $maxDisp$ and RL (Lines 12–15); details are given in Section 5 below. We run ECO route to restore routing correctness following the relocation of instances by the ILP-based ECO DP (Line 16). To further improve the dynamic IR drop, we run the commercial ECO DP on the ILP-generated solution (Line 17).

Algorithm 1: Our Proposed Hybrid Detailed Placer

Input: Instance set N after RouteOpt stage of P&R,
DVD threshold Th , clip width cw , clip height ch ,
max displacement $maxDisp$, max row limit RL

Output: IR drop-mitigated placements

- 1: Run dynamic IR drop simulation with DVD threshold Th
 - 2: Define filtered instance set $FN \leftarrow i \in N$ if effective DVD of $i \geq Th$
 - 3: Define clips set $C \leftarrow \phi$
 - 4: Define PDN zones set P from layout
 - 5: // Create clip set C from FN
 - 6: **for** $i \in FN$ **do**
 - 7: Generate clip c , centered at i with dimensions (cw, ch)
 - 8: **if** c does not have clip overlap violations using C and P **then**
 - 9: Insert c to C
 - 10: **end if**
 - 11: **end for**
 - 12: **for** $c \in C$ **do**
 - 13: Create ILP formulation of c with $maxDisp$, RL and solve
 - 14: Update instance locations from c
 - 15: **end for**
 - 16: Run ECO Route
 - 17: Continue commercial IR drop-aware detailed placement
-

Clip and PDN zones extraction. [10] extracts clips for better scalability of the ILP solver. We adopt a similar idea. We extract a clip around any instance that violates the target Dynamic Voltage Drop (DVD) threshold. Figure 4 gives an illustration: (a) shows a heatmap from dynamic IR drop simulation of the *mpeg2_top* design; (b) shows an extracted clip from the red square in (a); and (c) shows how we define three PDN zones. In Figures 4(a) and (c), it is observed that the voltage drops across various *PDN zones* are independent of each other. This is due to the vertical PDN stripes (depicted in orange in Figure 4(a)), that are responsible for establishing the power and ground connections. Therefore, we only consider the y -axis overlap when the clips lie within the same *PDN zone*.

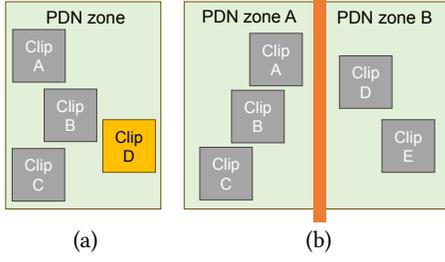


Figure 5: Illustration of clip overlap violations. (a) Overlap violations occur at clip D because of a y -axis overlap with clips B and C. (b) No overlap violations.

Table 1: ILP constants and variables.

Name	Range	Var/Const	Meaning
p_{irg}^k	{0,1}	Const	Instance i of k^{th} placement is located in (r, q) grid.
λ_i^k	{0,1}	Variable	Instance i of k^{th} placement is used.
d_i^k	Int	Const	Cell displacement for instance i of k^{th} placement.
r_i^k	Int	Const	Row displacement for instance i of k^{th} placement.
I^P	Float	Variable	Max peak demand currents of all power rails.
I^G	Float	Variable	Max peak demand currents of all ground rails.
B_{pt}^P	Float	Const	Noise of p^{th} power rail at timestamp t .
B_{gt}^G	Float	Const	Noise of g^{th} ground rail at timestamp t .
W_{it}^P	Float	Const	Instance i 's power pin's current at timestamp t .
W_{it}^G	Float	Const	Instance i 's ground pin's current at timestamp t .

Table 2: ILP sets definitions.

Name	Meaning
N	Set of instances
R	Set of rows
Q	Set of columns
K_i	Set of placements for an instance i
T	Set of timestamps from DVD simulation
P	Set of rows attached to power rails
G	Set of ground rails for VSS
V_{ip}^P	Set of placements for an instance i , attached to p^{th} power rail
V_{ig}^G	Set of placements for an instance i , attached to g^{th} ground rail

Clip overlap violations. During clip generation, it is crucial that there are no overlaps between the clips. Overlaps along the y -axis can disrupt the optimization process for peak currents during the detailed placement.³ Figure 5 presents an example. Figure 5(a) shows a clip overlap violation scenario where clip D overlaps with clips B and C along the y -axis. Figure 5(b) demonstrates a scenario without any clip overlap violations. Despite clip D overlapping along the y -axis with clips A and B, clips A and B are situated in different PDN zones. We see that taking PDN zones into account enables generation and processing of more clips, which can lead to more effective reduction of IR drop.

5 PROPOSED ILP FORMULATION

We now describe our Integer Linear Programming (ILP) formulation, adapted from [10]. Table 1 presents definitions of all terminologies, where *Var* denotes a variable and *Const* refers to a constant value. Table 2 defines the various sets used in the ILP formulation. Our ILP optimizes

³We apply our ILP to clips with no overlap violations and that lie in a given PDN zone.

$$\min OBJ = \alpha \cdot OBJ_{ird} + \beta \cdot OBJ_{disp} \quad (\alpha \gg \beta) \quad (1)$$

where OBJ_{ird} denotes the objective function related to IR drop mitigation, and OBJ_{disp} denotes the objective function related to the displacement of instances during ECO detailed placement. α and β denote the weights between IR drop and displacement, respectively, as shown in Equation 1. Our goal is to effectively mitigate the IR drop while simultaneously minimizing the perturbation caused during the detailed placement process. To accomplish this, we assign a higher value to α compared to β ($\alpha \gg \beta$).⁴ We now explain the two components of our ILP objective function in more details.

5.1 IR Drop Mitigation

The objective of mitigating the IR drop is to minimize the maximum peak current across every power rail (I^P) and ground rail (I^G) (Equation 2).

$$OBJ_{ird} = I^P + I^G \quad (2)$$

$$I_{pt}^P = B_{pt}^P + \sum_{i \in N} \sum_{k \in V_{ip}^P} \lambda_i^k \cdot W_{it}^P \leq I^P \quad \forall p \in P, \forall t \in T \quad (3)$$

$$I_{gt}^G = B_{gt}^G + \sum_{i \in N} \sum_{k \in V_{ig}^G} \lambda_i^k \cdot W_{it}^G \leq I^G \quad \forall g \in G, \forall t \in T \quad (4)$$

The expected current value within the p^{th} power rail at timestamp t (I_{pt}^P) is defined in Equation 3. B_{pt}^P represents the background noise within the p^{th} power rail at timestamp t . The term $\sum_i W_{it}^P$ indicates the sum of all current values contributed by the instances attached to the p^{th} power rail at the same timestamp. The expected current value for a ground rail at a given timestamp is defined similarly (Equation 4).

5.2 Cell Displacement Minimization

The objective of minimizing the total displacement during ECO detailed placement in the post-route opt (PRO) stage is captured in Equation 5.

$$OBJ_{disp} = \sum_{i \in N} \sum_{k \in K_i} d_i^k \cdot \lambda_i^k \quad (5)$$

$$\sum_{i \in N} \sum_{k \in K_i} r_i^k \cdot \lambda_i^k \leq RL \quad (6)$$

Minimizing total displacement during the PRO stage is beneficial for the following reasons:

- *Reduced ECO routing runtime:* Minimizing cell disturbance avoids excessive cell movements that can potentially lead to increased ECO routing runtime.
- *Stability in optimization results:* Excessive cell movement can also lead to divergent or less predictable results, after optimization by commercial tools. Keeping cell displacement to a minimum helps ensure more reliable optimization outcomes.

⁴In our experiments, we set α and β to 10^{12} and 1, respectively. This can be changed if different units are applied.

For each possible placement of instance i , represented by the variable λ_i^k , there is a unique cell and row displacement value denoted as d_i^k and r_i^k , respectively. These values quantify the extent of movement for each cell from its original position. To improve scalability, we use a threshold displacement value $maxDisp$, and generate the λ_i^k variables only when the displacement does not exceed the threshold $maxDisp$, i.e., $d_i^k \leq maxDisp$. This constraint reduces the ILP runtime by reducing the total number of variables and constraints in our ILP formulation.⁵ To further reduce the detailed placement perturbation and preserve the assumption of demand current linearity, we introduce a row limit, denoted as RL . This limit is set to control the total number of row swaps allowed for each clip (Equation 6). Consequently, the total number of instances that can be row-swapped is bounded by $RL \cdot \#Clips$.

5.3 Detailed Placement Baseline

For our ILP-based detailed placement, we adopt the single-cell-placement (SCP) model, which has demonstrated superior performance in comparison to other detailed placement models (S and RQ models) [10].⁶ The SCP model comprises multiple variables λ_i^k for each instance i , corresponding to the total number of possible placements ($|K_i|$). The overall complexity of the model scales as $O(|N||R||Q|)$.

$$\sum_{i \in N} \sum_{k \in K_i} p_{irq}^k \cdot \lambda_i^k \leq 1 \quad \forall r \in R, \forall q \in Q \quad (7)$$

$$\sum_{k \in K_i} \lambda_i^k = 1 \quad \forall i \in N \quad (8)$$

The detailed placement solution must have at most one site occupation for every cell and its associated placement (Equation 7). For each instance i , only a single placement must exist as enforced by Equation 8.

6 EXPERIMENTAL RESULTS

We use *IBM ILOG Cplex 12.10* as our ILP solver and run all experiments on a server with an Intel Xeon(R) Gold 6148, 2.40GHz (40 cores) CPU and 256GB memory. We use five testcases: *aes_cipher_top*, *des_perf*, *jpeg_encoder*, *mpeg2_top*, and *vga_enh_top* from *OpenCores* [15]. Table 4 shows a summary of the five designs used in our experiments.⁷ These are based on the 7nm FinFET technology node. In Table 4, *Init FP Util* and *Target CP* denote the initial floorplan utilization and target clock periods, respectively. In our experiments, we use: (i) *Synopsys Design Compiler R-2020.09* [17] to synthesize from RTL; (ii) *Cadence Innovus v21.11-s130_1* [19] to run the P&R flow (e.g., floorplan, placement, clock tree synthesis, and routing); and (iii) *Cadence Voltus v21.14-s111_1* [18] to measure the dynamic IR drop at the PRO stage. We use an unnamed commercial tool for the IR drop-aware ECO detailed placement. We build the power mesh on metal layers M7 to M9, characterized by a width of 1.0 μ m,

⁵In preliminary experiments, we vary the $maxDisp$ parameter between 0.15 μ m and 0.75 μ m and measure ILP runtime. These early studies indicate that setting $maxDisp \leq 0.25\mu$ m achieves reasonable runtime and QoR.

⁶The S model is known as the site occupation model, and the RQ model refers to the row and column occupation model. Refer to Section 3 of [10] for more details.

⁷We performed studies using a total of ten designs, but for conciseness report on just five here. We make available the remaining results along with runscripts to enable reproduction of our work, in our open-sourced repo [14].

a pitch of 15.5 μ m and an offset of 0.05 μ m. We set 7ns as the dynamic IR drop timestamp interval. The dimensions of the clip size are set to 1.4 μ m by 1.2 μ m, with a maximum displacement ($maxDisp$) of 0.25 μ m. The parameters α and β are set to 10^{12} and 1, respectively, and $Row Limit$ (RL) in Equation 6 is set to 5.⁸

6.1 Hybrid ECO DP Results

Table 3 shows a comparison of DVDs on our designs. We compare between: (i) conventional approaches, denoted as C (one commercial ECO DP iteration) and $C+C$ (two commercial ECO DP iterations), and (ii) our proposed approaches, denoted as O (one ILP ECO DP) and $O+C$ (a combination of our approach followed by the commercial method). Our study does not include comparisons with the previous state-of-the-art [7] due to unavailability of their closed-source implementation. [5] adds power staples for mitigating IR drop whereas we perform ECO DP/DR – so we also omit [5] from our comparisons. In most cases, the hybrid $O+C$ approach demonstrates superior performance across three metrics: worst DVD, average DVD of the top 5 worst cases, and average DVD of the top 10 worst cases. In these comparisons, the minimum counts for each category are highlighted in bold across the four ECO DP approaches we consider (C , $C+C$, O , $O+C$).

Our $O+C$ hybrid flow achieves the best results in 17 out of 30 instances⁹ for the worst DVD, in 16 instances for the average of the top 5 worst DVDs, and in 15 instances for the average of the top 10 worst DVDs. These results are visually represented in Figure 6, where half of the data points correspond to Table 3. A notable example is the *jpeg2* case, under a DVD threshold of 100 mV. Here, our proposed $O+C$ hybrid method achieves a 15.3%¹⁰ better improvement in DVD reduction compared to the conventional ECO DP flow ($C+C$).

6.2 ECO DP Characteristics

Table 5 presents overall characteristics for the ECO detailed placer and our ILP-based DP. We report the total number of moved instances ($MInsts$), total displacement ($Disp$), and runtime (RT). $\#Clips$ denotes the total number of clips generated during our detailed placement flow, $CInsts$ denotes the total instances considered during the ILP-based detailed placement, and $MRIInsts$ denotes the total row-swapped moved instances during the ILP-based detailed placement. Given that we set RL to 5 as per Equation 6 in Section 5.2, the count of $MRIInsts$ is strictly bounded by $5 * \#Clips$. Since our method avoids unnecessary cell movements, it results in lower values of $MInsts$ and $Disp$ compared to the commercial detailed placement. In terms of runtime, our proposed hybrid methods ($O+C$) achieve an average of 55.1% less runtime compared to two sequential commercial ECO detailed placements ($C+C$). Since our ILP detailed placement does not excessively perturb the placement compared to the conventional method, the hybrid detailed placement can generate better results with small runtime overheads compared

⁸In preliminary experiments, we vary the RL parameter between 1 and 15 and measure the correlation between the expected peak current before DP and the actual measured peak current after the *ecoRoute* and DP stages. These early studies indicate that the highest correlation is obtained when we set RL to 5.

⁹Tables that we present here show 15 instances = 5 testcases \times 3 thresholds T . Results for the other 15 instances are available in [14].

¹⁰The improvement is calculated as $(\frac{value-baseline}{baseline})$.

Table 3: DVD QoR comparison between the conventional commercial flow and our flow. Th = input DVD threshold. C = commercial dynamic IR-driven detailed placer. O = our ILP-based DP. $C + C$ = running commercial dynamic IR-driven detailed placer two times. $O + C$ = running our detailed placer followed by commercial dynamic IR-driven detailed placer. Nominal voltage = 650 (mV). To obtain a range of IR drop-violating instances, we sweep $Th = 100/80/65$ (mV). Best results are marked with bold font.

Design	Th (mV)	Worst DVD					Average of Top Worst 5 DVDs					Average of Top Worst 10 DVDs				
		Baseline	Conventional		Our Proposed		Baseline	Conventional		Our Proposed		Baseline	Conventional		Our Proposed	
		PRO	C	C+C	O	O+C	PRO	C	C+C	O	O+C	PRO	C	C+C	O	O+C
aes2	65		141.63	138.81	136.12	138.51		140.758	136.908	134.078	137.886		138.282	134.879	132.535	136.601
	80	134.23	138.41	121.05	136.08	134.32	133.072	136.98	120.694	134.024	132.576	131.855	135.189	120.144	132.462	131.643
	100		133.34	140.63	135.88	124.39		132.492	137.648	133.83	123.978		131.482	134.736	132.299	123.231
des2	65		118.25	118.16	119.85	118.4		109.632	108.63	117.322	109.758		105.04	102.732	115.694	105.882
	80	119.7	117.14	117.14	119.85	117.34	117.344	108.716	108.358	117.262	106.452	115.653	104.752	103.059	115.613	101.585
	100		108.67	108.67	119.85	108.51		103.986	103.986	117.26	105.502		101.345	101.345	115.617	103.15
jpeg2	65		136.93	140.54	153.68	153.49		135.054	137.712	152.67	152.324		133.355	136.429	150.138	147.428
	80	152.27	150.53	153.65	153.5	150.14	151.832	146.82	150.226	152.484	147.446	150.215	142.786	147.976	149.959	144.863
	100		159.04	159.77	153.38	136.39		155.414	155.52	152.312	135.148		150.705	151.884	149.738	132.537
mpeg2	65		149.63	153.65	150.43	139.75		147.79	151.75	148.96	138.236		144.161	149.555	146.399	136.464
	80	150.63	149.76	146.57	150.45	144.51	149.122	147.916	144.002	148.976	142.404	146.423	144.42	142.306	146.41	138.499
	100		141.76	141.57	150.49	139.84		139.52	139.65	149.02	136.368		137.367	136.523	146.474	132.517
vga2	65		102.86	103.04	110.81	110.81		101.618	101.754	109.76	109.76		98.007	99.36	107.472	107.472
	80	105.79	107.05	105.19	106.18	95.38	105.32	105.706	104.336	105.84	94.948	104.395	104.799	102.419	105.219	93.688
	100		105.62	105.62	106.23	100.77		104.074	104.074	105.872	99.514		101.837	101.837	105.247	99.02
Best Count		N/A	5	7	1	17	N/A	5	8	1	16	N/A	5	9	1	15
Average		145.437	139.03	138.54	145.83	136.59	143.73	136.19	135.33	144.06	133.95	141.26	133.10	131.98	141.27	131.16

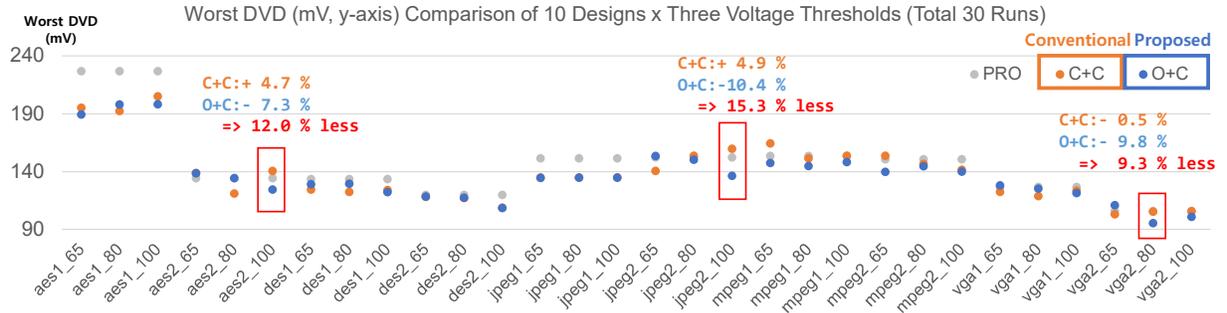


Figure 6: Visualization of Table 3 with the worst DVD (mV, y-axis) cases. $O+C$ shows up to 15.3% improvement compared to $C+C$.

Table 4: Designs [15] used in our experiments (7nm node).

Design	#Insts	#Nets	Init FP Util (%)	Target CP (ns)
aes2	11,657	11,920	80	0.8
des2	13,812	13,935	70	0.5
jpeg2	50,245	50,269	70	0.35
mpeg2	11,744	11,875	60	0.45
vga2	74,521	74,618	70	0.6

to C), and even achieves better DVD results compared to $(C+C)$. For a relatively larger design, $vga2$ with threshold = 100 (mV), our approach reduces the ECO DP runtime by 55.9% (1212s to 534s) and reduces the IR drop from 105.62 mV to 100.77 mV (see Table 3).

Table 6 shows a timing comparison of the four ECO DP approaches (C , $C+C$, O , $O+C$). In this table, W represents the worst negative slack (WNS), T denotes the total negative slack (TNS), and N indicates the number of failing endpoints (NFE). The bold font is used to signify the minimum NFE achieved across the C , $C+C$, O , and $O+C$ approaches. Our ILP-based detailed placement considers minimizing instance displacement in its objective. Consequently O tends to show less timing perturbation in most cases compared to

PRO. This is particularly evident when comparing the average values between $C+C$ and $O+C$, and suggests that greater cell movement often correlates with increased timing degradation.¹¹ Compared to C and $C+C$, our approach also reduces NFE count by up to 83%. Based on this timing comparison, we conclude that our proposed methods, O and $O+C$, perturb the timing less than C and $C+C$.

7 CONCLUSION

We have presented a novel hybrid ECO detailed placement method that combines commercial flows with an ILP-based detailed placer to reduce the maximum dynamic current demands within power and ground rails. Our hybrid method optimizes maximum current peaks and follows commercial EDA detailed placement optimizations, outperforming conventional approaches. Our approach uses a commercial tool to measure the voltage drop and extract current waveforms, then formulates an ILP optimization to refine cell placements. The hybrid method reduces maximum dynamic voltage drops by up to 15.3% with 55.1% less runtime compared to the

¹¹The ranking of the number of moved instances is as follows: $C+C > O+C > C > O$. This ordering aligns well with the observed changes in timing, specifically in terms of the average NFE and TNS as presented in Table 6.

Table 5: Overall characteristics of detailed placement. Th = input DVD threshold. $MInsts$ = #moved instances. $MRInsts$ = #row-changed instances. $Disp$ = total displacement. RT = runtime. $\#Clips$ = #clips generated during ILP DP. $CInsts$ = #instances across all clips.

Design	Th (mV)	Conventional						Our Proposed								
		C			C+C			O			O+C					
		MInsts	Disp (μ m)	RT (s)	MInsts	Disp (μ m)	RT (s)	#Clips	CInsts	MInsts	MRInsts	Disp (μ m)	RT (s)	MInsts	Disp (μ m)	RT (s)
aes2	65	3297	1935.996	109	3684	2842.872	219	14	1300	228	62	37.728	35.944	3575	2188.38	181
	80	2638	1081.716	104	2858	1548.198	217	11	1040	182	50	31.278	29.306	2603	972.708	179
	100	1126	263.436	95	1283	382.932	198	6	570	107	28	18.228	8.014	1132	288.384	159
des2	65	1723	536.754	152	1933	745.302	500	29	1390	244	93	36.81	18.771	1760	571.134	175
	80	487	110.94	139	543	130.998	287	13	697	124	46	17.592	5.52	520	110.832	167
	100	79	14.394	134	79	14.394	296	4	208	35	13	4.8	1.991	95	17.58	172
jpeg2	65	8384	3015.366	421	9522	4063.86	783	107	6360	567	284	96.678	25.103	8563	3046.194	375
	80	3575	1261.254	432	4173	1766.322	893	55	3075	312	162	54.588	24.677	3788	1322.106	384
	100	1245	378.084	424	1412	490.59	860	29	1439	181	99	32.328	19.299	1270	388.782	386
mpeg2	65	2395	856.392	157	2947	1255.536	324	34	1344	116	61	21.024	19.618	2392	850.53	180
	80	1286	445.008	151	1536	601.146	357	22	800	63	33	11.658	9.179	1324	457.092	163
	100	531	162.834	164	637	221.322	377	12	364	41	20	7.614	6.781	540	168.018	165
vga2	65	8759	2606.238	539	10308	3365.496	1089	131	6679	466	239	79.578	20.25	8700	2604.69	114
	80	1346	382.404	637	1502	437.568	1269	45	2348	198	98	33.222	19.8	1321	356.184	506
	100	32	6.942	603	32	6.942	1212	3	152	5	5	1.2	1.678	36	7.536	534
Average		2395.3	885.93	292.93	2764.4	1234.83	599.93	33.5	1726.43	175.47	78.97	29.01	15.52	2453.1	898.7	269.6

Table 6: Timing comparison between the conventional commercial flow and our proposed flow. Th = input DVD threshold. W = Worst negative slack (WNS). T = Total negative slack (TNS). N = #failing endpoints (NFE).

Design	Th (mV)	Baseline			Conventional						Our Proposed					
		PRO			C			C+C			O			O+C		
		W	T	N	W	T	N	W	T	N	W	T	N	W	T	N
aes2	65				-0.061	-0.467	29	-0.101	-1.131	58	0.000	0.000	0	-0.097	-0.438	28
	80	0.000	0.000	0	-0.053	-0.062	4	-0.091	-0.312	19	0.000	0.000	0	-0.012	-0.033	4
	100				0.000	0.000	0	0.000	0.000	0	0.000	0.000	0	0.000	0.000	0
des2	65				0.000	0.000	1	-0.001	-0.001	1	-0.013	-0.013	1	-0.014	-0.014	1
	80	0.002	0.000	0	0.002	0.000	0	0.001	0.000	0	0.002	0.000	0	0.001	0.000	0
	100				0.002	0.000	0	0.002	0.000	0	0.002	0.000	0	0.002	0.000	0
jpeg2	65				-0.034	-1.453	209	-0.032	-1.830	223	-0.012	-0.091	36	-0.035	-1.676	250
	80	0.000	0.000	0	-0.033	-0.384	62	-0.033	-0.754	89	-0.012	-0.061	22	-0.023	-0.478	83
	100				-0.007	-0.031	11	-0.008	-0.036	14	0.000	-0.001	4	-0.004	-0.022	17
mpeg2	65				0.000	0.000	0	0.000	0.000	0	0.000	0.000	0	0.000	0.000	0
	80	0.000	0.000	0	0.000	0.000	0	0.000	0.000	0	0.000	0.000	0	-0.001	-0.001	1
	100				0.000	0.000	0	0.000	0.000	0	0.000	0.000	0	0.000	0.000	0
vga2	65				-0.004	-0.007	4	-0.016	-0.077	9	-0.013	-0.042	11	-0.013	-0.059	12
	80	0.002	0.000	0	0.002	0.000	0	0.002	0.000	0	-0.013	-0.037	8	-0.012	-0.020	4
	100				0.002	0.000	0	0.002	0.000	0	0.002	0.000	0	0.002	0.000	0
Average		0.002	0.000	0.10	-0.018	-0.471	35.20	-0.028	-0.778	43.00	-0.001	-0.011	5.77	-0.021	-0.510	37.63

original conventional flow, across ten designs. Our future work aims to develop a new methodology that eliminates the need for hybrid approaches. Reducing the timing perturbation, although challenging, remains another potential direction for improvement. Gaining a stronger understanding of the commercial ECO detailed placement flow will also be beneficial.

REFERENCES

- [1] L. Bhamidipati, B. Gunna, H. Homayoun and A. Sasan, "A Power Delivery Network and Cell Placement Aware IR-drop Mitigation Technique: Harvesting Unused Timing Slacks to Schedule Useful Skews", *Proc. ISVLSI*, 2017, pp. 272–277.
- [2] T.-C. Chen, Z.-W. Jiang, T.-C. Hsu, H.-C. Chen and Y.-W. Chang, "NTUplace3: An Analytical Placer for Large-scale Mixed-size Designs with Preplaced Blocks and Density Constraints", *IEEE TCAD* 27(7) (2008), pp. 1228–1240.
- [3] Y. Chuang, P. Lee and Y. Chang, "Voltage-drop Aware Analytical Placement by Global Power Spreading for Mixed-size Circuit Designs", *IEEE TCAD* 30(11) (2011), pp. 1649–1662.
- [4] Y. Fang, H. Lin, M. Sui, C. Li and E. J. Fang, "Machine-learning-based Dynamic IR Drop Prediction for ECO", *Proc. ICCAD*, 2018, pp. 1–7.
- [5] S. I. Heo, A. B. Kahng, M. Kim, L. Wang and C. Yang, "Detailed Placement for IR Drop Mitigation by Power Staple Insertion in Sub-10nm VLSI", *Proc. DATE*, 2019, pp. 830–835.
- [6] C. T. Ho and A. B. Kahng, "IncPIRD: Fast Learning-based Prediction of Incremental IR Drop", *Proc. ICCAD*, 2019, pp. 1–8.
- [7] X.-X. Huang, H.-C. Chen, S.-W. Wang, I. H.-R. Jiang, Y.-C. Chou and C.-H. Tsai, "Dynamic IR-Drop ECO Optimization by Cell Movement with Current Waveform Staggering and Machine Learning Guidance", *Proc. ICCAD*, 2020, pp. 1–9.
- [8] A. B. Kahng, B. Liu and Q. Wang, "Supply Voltage Degradation Aware Analytical Placement", *Proc. ICCAD*, 2005, pp. 437–443.
- [9] A. B. Kahng, S. Reda and Q. Wang, "APlace: A General Analytic Placement Framework", *Proc. ISPD*, 2005, pp. 233–235.
- [10] S. Li and C.-K. Koh, "Mixed Integer Programming Models for Detailed Placement", *Proc. ISPD*, 2012, pp. 87–94.
- [11] Z. Xie, H. Ren, B. Khailany, Y. Sheng, S. Santosh, J. Hu and Y. Chen, "PowerNet: Transferable Dynamic IR Drop Estimation via Maximum Convolutional Neural Network", *Proc. ASP-DAC*, 2020, pp. 13–18.
- [12] Cadence Innovus Rapid Adoption Kit. <http://www.cadence.com>
- [13] Cadence Voltus IC Power Integrity Solution User Guide. <http://www.cadence.com>
- [14] ILP-ECO DP Sources and Full Tables with Ten Designs (aes1-2, des1-2, jpeg1-2, mpeg1-2, and vga1-2). <https://github.com/ABKGroup/Dynamic-IR-Drop-ECO-DP>
- [15] OpenCores: Open-Source IP Cores. <https://opencores.org>
- [16] Qualcomm, personal communication, September 2023.
- [17] Synopsys Design Compiler, version R-2020.09. <https://www.synopsys.com/>
- [18] Cadence Voltus IC Power Integrity Solution, version 21.14-s111_1. <https://www.cadence.com>
- [19] Cadence Innovus Implementation System, version 21.1. <https://www.cadence.com>