Guiding Global Placement with Reinforcement Learning

Robert Kirby
Nvidia
Santa Clara, CA, USA
rkirby@nvidia.com

Kolby Nottingham
Univeristy of California Irvine
Irvine, CA, USA
knotting@uci.edu

Rajarshi Roy
Nvidia
Santa Clara, CA, USA
rajarshir@nvidia.com

Saad Godil
Nvidia
Santa Clara, CA, USA
sgodil@nvidia.com

Bryan Catanzaro
Nvidia
Santa Clara, CA, USA
bcatanzaro@nvidia.com

Abstract—Recent advances in GPU accelerated global and detail placement have reduced the time to solution by an order of magnitude. This advancement allows us to leverage data driven optimization (such as Reinforcement Learning) in an effort to improve the final quality of placement results. In this work we augment state-of-the-art, force-based global placement solvers with a reinforcement learning agent trained to improve the final detail placed Half Perimeter Wire Length (HPWL).

We propose novel control schemes with either global or localized control of the placement process. We then train reinforcement learning agents to use these controls to guide placement to improved solutions. In both cases, the augmented optimizer finds improved placement solutions.

Our trained agents achieve an average 1% improvement in final detail place HPWL across a range of academic benchmarks and more than 1% in global place HPWL on real industry designs.

I. INTRODUCTION

In VLSI design high quality global placement results are highly correlated to the final quality (area, performance and power) of the physical design. As one of the first stages of the physical design flow, placement decisions affect the results of all downstream design stages. Recently DREAMPlace and ABCDPlace provide an extremely fast platform for running global and detail placement [1], [2]. This opens up the possibility of running large numbers of exploration placement runs to find the best possible placement solution. We leverage these advancements along with recent advances in reinforcement learning and introduce an augmented state of the art placement algorithm which learns new internal heuristics which provide higher quality final solutions.

State of the art force based academic placers use the Lagrangian relaxation technique to optimize the constrained objective function which takes into account both cell density and half-perimeter wire length (HPWL). Improvements to these algorithms (such as work in ePlace and RePlAce) frequently come as improvements to heuristic rules used during the course of the optimization [3], [4]. In this paper we investigate whether it is possible to use reinforcement learning to find better heuristics than the ones being used today. The main contributions of this paper are:

- We propose the use of large scale placement exploration accelerated by GPUs to fuel a data driven approach to placement.
- We propose two distinct methods of augmenting force-based global placement algorithms with reinforcement learning.
- We introduce a unique correlated sampling strategy for reinforcement learning algorithms acting on a two dimensional action space.
- We demonstrate that using these methods results in a 1% reduction in HPWL across a range of academic and industry benchmarks.

As far as we are aware this is the first attempt to use reinforcement learning to directly control the dynamics of a state of the art global placement algorithm.
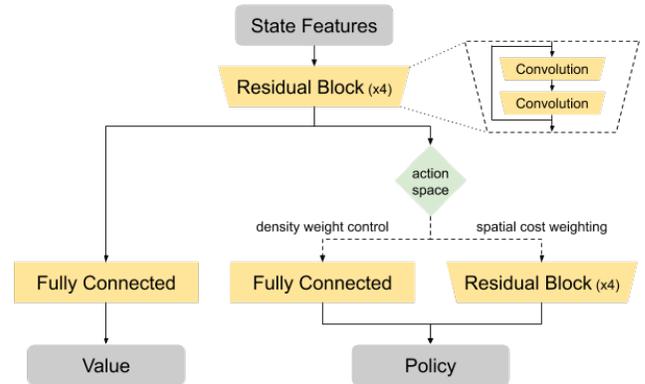

Fig. 1. Neural network architecture for trained policy and value functions.

The modern VLSI design flow is an iterative process. The placement stage of a given design is run many times often with incremental changes to the underlying design. A placement engine that can learn from previous iterations and apply learned strategies to provide better quality of results to later iterations is therefore desirable. To this end we also demonstrate that our learned policies retain on average 77% of their original benefit through hundreds of synthetic netlist edits. Given that the compute time for training our agents on a new partition is significant this ability to generalize to design changes is important.

II. BACKGROUND

A. Global Placement Optimization

The objective of the global placement problem is to find locations for design cells that minimize competing objectives. These objectives always include HPWL (Half Perimeter Wire Length) and cell density, and can optionally include other metrics such as routability, timing, etc. In this work we focus only on the HPWL and cell density objectives and leave extenstion to other objectives as possible future work.

Current state of the art academic placers model the netlist and cells as an electrostatic system [3], [4]. In these models cells are treated as point charges with the cell density cost calculated as potential energy of the system. This formulation allows for the use of the fast Fourier transform to efficiently and differentiably calculate the potential energy and therefore the density cost of a given placement. These methods then use Nesterov's method [5] to iteratively solve for a placement that minimizes both HPWL and density costs.

The optimization process of these approaches can be summarized as

$$\min_{\mathbf{x},\mathbf{y}} C = WL(\mathbf{x},\mathbf{y}) + \lambda D(\mathbf{x},\mathbf{y}) \qquad (1)$$

| Log of current HPWL |
| Log of ΔHPWL |
| Log of the current density weight (λ) |
| Current overflow as reported by DREAMPlace |
| Current cof (as defined in Algorithm 1) |
| Cell Density Map |
| Wire Density Map |
| Local HPWL (2x2, 4x4, 8x8, 16x16) |

where C is the overall cost, WL is the wire length cost function and D is the cell density cost function.

As part of the placement algorithm these solvers must choose a value (λ) to control the tradeoff between wire length and density costs (referred to as the density cost weight). Before RePlAce/ePlace, electrostatic-based solvers either chose a fixed or a gradually increasing value for the density cost weight. ePlace proposed a heuristic rule based on recent changes in wire length cost and RePlAce further suggested 'dynamic step size adaptation' which automatically adjusts the penalty based on the HPWL curve of a trial placement [3], [4]. Using this approach they were able to show an improvement in solution quality. Additionally, the authors were able to show that adding a local density cost function which adjusted the density weight based on local overflow statistics further improved the solution quality.

Many of the insights from the previous works highlight the importance of existing heuristics to global placement tools. As the authors of previous works were able to improve results by adding heuristic rules to control the dynamics of the placement solver, in this work we study instead training a reinforcement learning agent to either replace existing heuristics or leverage new controls in order to optimize for final solution quality.

B. Asynchronous Advantage Actor Critic

We frame global placement as a Markov Decision Process (MDP) with states $\mathcal{S}$, actions $\mathcal{A}$, transition function $\mathcal{T} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$, and reward function $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. We define $\mathcal{R}$, $\mathcal{S}$, and $\mathcal{A}$ for global placement in Sections IV, IV-A, and IV-B1 or IV-B2 respectively. $\mathcal{T}$ is defined by DREAMPlace as explained in Section IV.

In the reinforcement learning paradigm, we optimize a policy π : $\mathcal{S} \rightarrow \mathcal{A}$ to maximize expected discounted return $R = \mathbb{E}[\sum_{t=0}^{T} \gamma^t r_t]$ where T is the horizon length, γ is a discount factor between 0 and 1, and r_t is the reward at timestep t of a trajectory.

Policy gradient methods with continuous action spaces traditionally model π as a parameterized Gaussian and an underlying function provides the parameters of this distribution.

Policy gradient methods then optimize π directly by approximating its gradient using the objective function:

$$\nabla_\theta J(\pi_\theta) = \mathbb{E}_{\tau \sim \pi_\theta} \left[\sum_{t=0}^{T} \nabla_\theta \log \pi_\theta(a_t|s_t) A(\tau_{t:T}) \right] \quad (2)$$

where $s_t \in \mathcal{S}$ and $a_t \in \mathcal{A}$ are the state and action at timestep t of a sampled trajectory $\tau = [(s_0, a_0, r_0), ..., (s_T, a_T, r_T)]$. An agent following π_θ acts in an environment defined by $\mathcal{T}$ to collect τ. $A(\tau)$ is the advantage function and represents how much better or worse taking one action is compared to some baseline.

We choose the Asynchronous Advantage Actor Critic (A3C) algorithm to implement policy gradient reinforcement learning [6]. This method improves over vanilla policy gradient methods while

---

Algorithm 1 λ Scaling with Possible RL Control

1: **procedure** $\lambda\_Scaling$
2:      $\lambda \leftarrow (\sum grad\_wl)/(\sum grad\_density)$
3:      **for** $k = 0, last\_iteration$ **do**
4:          $p \leftarrow (HPWL_k - HPWL_{k-1})/\Delta HPWL_{ref}$
5:          **if** $RLEnabled$ **then**
6:              $cof \leftarrow RLPolicyOutput$
7:          **else**
8:              **if** $p < 0$ **then**
9:                  $cof \leftarrow cof_{max}$
10:             **else**
11:                $cof \leftarrow max(cof_{min}, pow(cof_{max}, 1 - p))$
12:             **end if**
13:          **end if**
14:          $\lambda \leftarrow \lambda \times cof$
15:      **end for**
16: **end procedure**

---

remaining straightforward to implement and modify. A3C defines $A(\tau)$ as:

$$A(\tau) = \sum_{i=0}^{n} \gamma^i r_i + \gamma^{n+1} V_\phi(s_{n+1}) - V_\phi(s_0) \quad (3)$$

where V_ϕ is a learned value function parameterized by ϕ and updated to approximate $\mathbb{E}_{\tau \sim \pi_\theta}[A(\tau)]$. This advantage function compares the return from a trajectory to the expected return from starting state s_0. The A3C advantage function makes use of an n-step Bellman target [7].

A3C also introduces an entropy term to the loss to aid in exploration and avoid early convergence. This is calculated with $\beta \nabla_\theta \mathbf{H}(\pi_\theta(\cdot|s_t))$ where $\mathbf{H}$ is the entropy function and β weights the entropy term's contribution to the total loss.

III. RELATED WORK

Some recent work has used reinforcement learning to try to solve related placement problems. [8] was able to learn a policy to explicitly place a smaller number of large macro cells before using a force based method to place the remaining cells. Other work such as [9] has studied using reinforcement learning for the assignment of logic elements to FPGA logic blocks. However, these differ from our work significantly as we investigate ways to directly improve the force-based method used to place smaller standard cells.

Many other works attempt to use machine learning to predict downstream problems during the placement stage to quickly identify potential problematic placement solutions [10] [11] [12]. These approaches are applications of supervised learning which provides potentially actionable information to other portions of the design flow. Our approach instead attempts to leverage reinforcement learning to learn a better placement algorithm.

The most relevant work is RePlAce/ePlace [3], [4] and DREAMPlace [1] which this work builds heavily upon. RePlAce and ePlace were able to push the state of the art in global placement optimization and DREAMPlace accelerated this placement using parallel processing on GPUs.

IV. METHOD

To create our reinforcement learning environment we modified DREAMPlace code to run placement in steps, yielding state information every 10 iterations and allowing the agent to observe this state and modify placement control parameters before continuing.
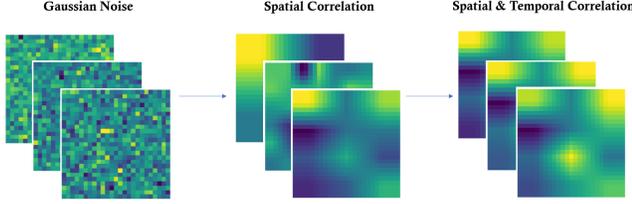
Fig. 2. Visualization of spatially and temporally correlated 2D exploration noise used during training

The reward, which is provided to the agent when the density target is reached, is the percent decrease in detail place HPWL when compared to the DREAMPlace algorithm run without agent interference (the baseline HPWL). In the event the placement process diverges the environment provides a fixed reward of -10.

A. DREAMPlace State

The state is presented to the policy network as a 3-dimensional tensor. The first two dimensions are spatial and the final channel dimension separates each individual feature. The features are listed in Table I. All scalar features are repeated across the first two dimensions. All features are clipped within their 10th and 90th percentile values and then normalized to zero-mean and unit-variance using the statistics of that feature during the baseline run.

B. DREAMPlace Actions

1) Density Weight Control: The first of two action spaces we define is referred to as 'density weight control'. It allows the agent to set the density weight coefficient (cof) used to adjust the value of the density cost weight (λ) during the placement algorithm. Algorithm 1 shows the original ePlace update rule for λ used in DREAMPlace with an additional condition added for our RL control. When enabled, the heuristic calculation of cof is instead replaced with the current output of the RL agent. This action space is therefore a single continuous value.

2) Spatial Cost Weighting: The second action space we define is referred to as 'spatial cost weighting'. This action space allows the agent to provide a 2D field which is used to scale the gradient of the placement objective function (C from Equation 1) for each cell i based on the cell's location, x_i, y_i, performing the below operation immediately after the gradient of the objective is calculated.

$$\frac{\delta C}{\delta x_i}_{new} = \frac{\delta C}{\delta x_i} \cdot a_t \left[\left\lfloor \frac{x_i}{d} \right\rfloor, \left\lfloor \frac{y_i}{d} \right\rfloor \right] \quad (4)$$

$$\frac{\delta C}{\delta y_i}_{new} = \frac{\delta C}{\delta y_i} \cdot a_t \left[\left\lfloor \frac{x_i}{d} \right\rfloor, \left\lfloor \frac{y_i}{d} \right\rfloor \right] \quad (5)$$

where a_t is the action at the current time step and d is the partition size divided by the dimension of the action space. This approach provides the agent the ability to resist or amplify cell movement in specific areas during the placement process.

C. Choice of Reinforcement Learning Algorithm

The choice of which RL algorithm to use for a given problem is an important one. Most RL algorithms fit into one of two categories, model-free or model-based, based on whether they attempt to directly model the environment dynamics. We focus on the 'model-free' category in this discussion (although we mention the possible application of 'model-based' approaches in section VI). Within 'model-free' RL we must decide between two styles, 'policy optimization' and 'Q learning'. In 'Q learning' an agent is trained to identify the expected future reward given a state-action pair ($Q_\theta(S, A)$). Alternatively 'policy optimization' methods directly train a policy function ($\pi_\theta(S)$) to maximize the expected future reward.

There are a couple of reasons we chose to explore the Actor-Critic ('policy optimization') framework for this specific problem. The first is that the complete state of the placement solver is difficult to represent and therefore a significant portion of the state of the placer remains hidden from the model. This makes approximating the true Q function difficult. The second is that the placement environment provides only a single reward at the terminal state. Because 'Q learning' methods train the Q function to match the sum of the immediate reward and $max_a Q(S, a)$, it can take many samples to incorporate sparse reward information into the model. Due to the 'on policy' nature of policy optimization methods such as Actor Critic, policy model updates can instead use the final observed reward to adjust the likelihood of all actions taken during a single sampled episode.

D. Actor-Critic Implementation

Our experiments use the A3C algorithm described above. We extend the asynchronous framework torchbeast with support for our multidimensional and continuous action spaces [13].

In our experiments we train both π_θ and V_ϕ simultaneously. We clip the gradient norm at 8 to increase training stability. To prevent the policy gradient loss from trending towards negative infinity during training we enforce a minimum value on the output of the policy standard deviation by adding a small fixed value to the network output.

Because our agent receives a sparse reward limited to terminal states, setting $\gamma = 1$ assures that early actions receive equal credit and that our agent has no incentive to reach terminal states too quickly.

Additional hyper-parameters can be found in Table II.

E. Neural Network Model Architectures

In deep reinforcement learning both π_θ and V_ϕ are approximated by deep neural networks. Figure 1 illustrates our network architecture. In our experiments both networks share a subset of their parameters in a trunk network (a 2D convolutional neural network with residual connections). The output of the trunk network is used as input to separate value and policy branches. The value branch is composed of a single fully connected layer. In the case of our 'density weight control' action space the policy branch is also a single fully connected layer. For the 'spatial cost weighting' action space the policy branch is a fully convolutional neural network similar to the main trunk.

F. Correlated Noise

Traditionally when drawing samples from the policy π for a continuous action space a sample is drawn from the unit normal distribution and then scaled and shifted by the parameters provided by the policy neural network. The entropy in this parameterized distribution allows the agent to occasionally "explore" new actions even if they were not highly likely under the given policy. For our 2D action space

TABLE III
BENCHMARK DETAIL PLACE HPWL ($\times 10^6$) RESULTS

Design	DREAMPlace Baseline		Spatial Cost Action			Density Weight Action		
	HPWL	Steps	HPWL (best)	Steps (best)	% Improvement (median,best)	HPWL (best)	Steps (best)	% Improvement (median,best)
Adaptec1	72.78	63	**72.39**	64	(0.51,0.53)	72.40	117	(0.51,0.52)
Adaptec2	81.98	68	81.42	73	(0.62,0.68)	**80.85**	172	(1.22,1.38)
Adaptec3	192.75	70	**190.13**	71	(1.25,1.35)	191.34	69	(0.62,0.73)
Adaptec4	173.46	73	172.66	76	(0.42,0.46)	**172.58**	151	(0.45,0.51)
NewBlue2	182.12	72	**179.33**	72	(1.53,1.53)	180.10	56	(1.04,1.10)
NewBlue3	255.62	74	254.93	81	(0.31,1.01)	**254.85**	148	(0.43,0.78)
NewBlue5	383.10	81	378.97	86	(1.04,1.07)	**378.25**	89	(1.15,1.26)
NewBlue6	441.10	76	439.26	81	(0.40,0.42)	**437.75**	121	(0.68,0.76)
NewBlue7	933.02	82	**929.39**	85	(0.35,0.38)	930.47	100	(0.27,0.30)
SuperBlue2	567.43	75	563.53	79	(0.67,0.69)	**562.83**	183	(0.78,0.81)
SuperBlue3	289.14	71	285.63	83	(1.17,1.21)	**284.61**	204	(1.56,1.57)
SuperBlue6	304.30	74	302.62	78	(0.50,0.55)	**301.10**	150	(0.99,1.05)
SuperBlue7	364.26	76	361.87	83	(0.53,0.65)	**359.63**	194	(1.22,1.27)
SuperBlue9	209.03	76	206.86	78	(0.98,1.03)	**205.45**	286	(1.61,1.71)
SuperBlue11	321.96	70	318.24	90	(1.01,1.16)	**316.96**	152	(1.35,1.55)
SuperBlue12	224.10	103	221.53	109	(0.96,1.15)	**220.86**	168	(1.41,1.44)
SuperBlue14	210.60	68	207.84	80	(1.07,1.31)	**207.27**	180	(1.54,1.58)
SuperBlue16	241.93	72	238.92	81	(1.06,1.24)	**238.64**	183	(1.34,1.35)
SuperBlue19	136.57	68	**133.58**	73	(2.14,2.18)	134.50	162	(1.48,1.52)
Mean	-	-	-	-	(0.87,0.98)	-	-	(1.03,1.12)

TABLE IV
INDUSTRY GLOBAL PLACE HPWL ($\times 10^6$) RESULTS

Design	DREAMPlace Baseline		Spatial Cost Action			Density Weight Action		
	HPWL	Steps	HPWL (best)	Steps (best)	% Improvement (median,best)	HPWL (best)	Steps (best)	% Improvement (median,best)
Design A	183.42	103	182.37	101	(0.23,0.57)	**182.22**	130	(2.00,2.53)
Design B	275.61	90	274.02	92	(0.30,0.58)	**268.34**	136	(1.29,2.64)
Design C	339.27	88	335.92	93	(0.81,1.18)	**331.15**	129	(1.22,2.39)
Design D	259.27	129	242.83	135	(5.97,6.34)	**236.66**	196	(4.32,8.72)
Design E	452.75	92	**445.32**	99	(1.60,1.64)	448.36	107	(0.75,0.97)
Design F	397.11	103	393.29	105	(0.86,0.96)	**391.80**	122	(1.23,1.34)
Design G	523.37	101	**519.71**	106	(0.63,0.70)	521.95	137	(0.19, 0.27)
Design H	367.87	117	365.16	119	(0.48,0.74)	**361.19**	134	(1.50, 1.82)
Design I	530.14	96	**517.73**	105	(2.32, 2.34)	524.30	116	(0.93,1.10)
Mean	-	-	-	-	(1.47,1.67)	-		(1.49,2.42)

TABLE V
INDUSTRY DATASET

Design	# Standard Cells	# Nets
Design A	1075356	1004645
Design B	1306336	1355267
Design C	1344574	1388626
Design D	1462400	1378091
Design E	1525405	1528489
Design F	1533034	1222882
Design G	1616198	1043353
Design H	1864867	1565958
Design I	2264505	2276034

care must be taken to ensure 'meaningful' exploration decisions are taken. If each value in the 2D action grid is IID there is too much high frequency noise (both spatially and temporally) in the actions for the agent to chance upon exploring coherent strategies. [14], [15] suggests sampling an Ornstein-Uhlenbeck process [16] to enforce temporal consistency in the sampled noise. This process models the velocity of a Brownian particle which is both temporally correlated and mean-reverting. We extend this to also include spatial consistency across the action space. To do this for each training episode we pick a random resolution lower than the action resolution and sample each pixel from an independent Ornstein-Uhlenbeck process. We then use bilinear upsampling to interpolate this lower resolution back to the size of original action space. This forces varying amounts of spatial and temporal consistency in the exploration noise. The noise sampled from this process is visually depicted in Figure 2. Without this correlated sampling our 'spatial cost' agent was unable to learn a policy that improved over the baseline HPWL.

G. Training Setup

Training is performed on a single DGX machine with 8 Tesla V100 GPUs. 7 GPUs are filled with actor instances (2 or 3 per GPU depending on design size) running our DREAMPlace based

Algorithm 2 Netlist Edit Algorithm

1: **procedure** MODIFYNETLIST($Netlist, NumEdits$)
2:    $n \leftarrow 0$
3:    **while** $n < NumEdits$ **do**
4:       $EditType \leftarrow RandomEditType()$
5:       **if** $EditType = AddNode$ **then**
6:          $N \leftarrow Netlist.AddNode(RandomSize())$
7:          $P \leftarrow$ Random integer $[2, 5]$
8:          **for** $i \leftarrow 0, P$ **do**
9:             ADDNET($N, Netlist$)
10:          **end for**
11:       **else if** $EditType = AddNet$ **then**
12:          $N \leftarrow Netlist.RandomNode()$
13:          ADDNET($N, Netlist$)
14:       **else if** $EditType = RemoveNode$ **then**
15:          $N \leftarrow Netlist.RandomNode()$
16:          $Netlist.RemoveNodeAndPins(N)$
17:          $Netlist.RemoveSinglePinNets()$
18:       **end if**
19:       $n \leftarrow n + 1$
20:    **end while**
21:    **return** InputNetlist
22: **end procedure**
23: **procedure** ADDNET($N, Netlist$)
24:    $P \leftarrow$ Random integer $[1, 4]$
25:    $NN \leftarrow Netlist.3HopNeighborhood(N)$
26:    $Nodes \leftarrow [N, NN.RandomNodes(P)]$
27:    $Pins \leftarrow Netlist.AddPinsToNodes(N)$
28:    $Netlist.AddNet(Pins)$
29:    **return** $Netlist$
30: **end procedure**

RL environment. The final GPU is used for running the training algorithm for the policy and value networks. Training is performed for 1.5e6 training steps (between 10,000 and 25,000 placement episodes depending on placement episode length) and the best solution found during training is reported. This training process takes between 24 and 48 hours depending on the size of the partition.

H. Inference Time

The addition of our agent to the placement process does have a small effect on runtime because we run our feature collection and agent neural network once every 10 placement steps. The addition of our agent adds a roughly 10% overhead to the DREAMPlace global placement runtime.

V. EXPERIMENTS

A. Data Set

We train our reinforcement learning policy on an array of open source placement benchmarks that have been previously used for various placement competitions. These include the ISPD '05 [17] and '06 [18] and DAC '12 [19] competition. We report non-scaled detail place HPWL values from DREAMPlace and ABCDPlace to obtain the 'DREAMPlace HPWL' detail place values for all partitions. In addition we run the same experiment on 9 industry designs. Due to memory limits running the GPU accelerated detail place step, we report improvement in the global place HPWL result.
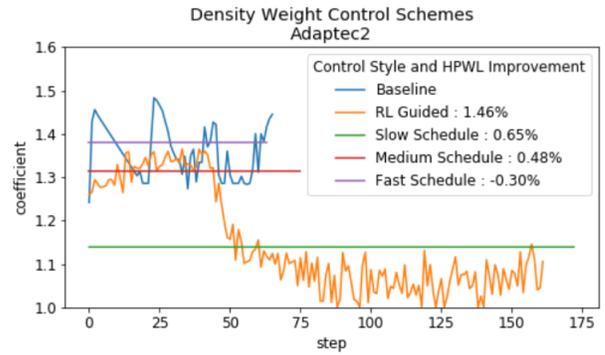

Fig. 3. Various Density Weight Control Schemes and resulting HPWL Improvement: RL Guided schedule outperforms static schedules with similar numbers of iterations.

B. Results

In Table III we present our results on open source benchmark designs. We report the max and median of the best placement solutions found across 3 independent training runs. The median 'density weight' action space solutions are on average 0.87% better than the DREAMPlace baseline solutions and the 'spatial cost' action space solutions are on average 1.03% improved over the baseline.

In Table IV we present results in a similar manner on nine designs taken from multiple industry workflows. On these designs the agent is able to improve by 1.47% and 1.49% on average using the 'spatial cost' and 'density weight' actions respectively. Notably the agent is able to improve the solution for Design D by more than 5%.

Our RL guided placement methods are able to find placement solutions for all designs which have a shorter HPWL than the DREAMPlace baseline. Interestingly, one action space does not outperform the other across all benchmarks. While the 'density weight' action does appear to be superior on most designs, it comes
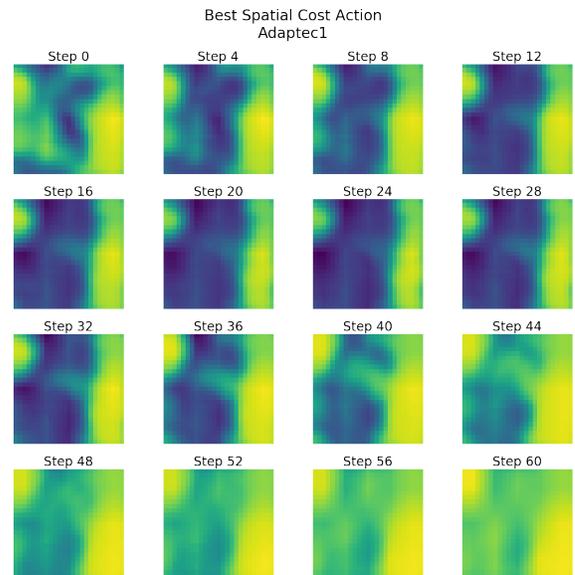

Fig. 4. Adaptec1 Best Policy (spatial cost): Action a from Equation 4 & 5 (low-high : blue-yellow). RL agent uses spatially and temporally varying control based on locations of macros (lower right) and logic clusters (upper right) to guide placement algorithm to more optimal solution.
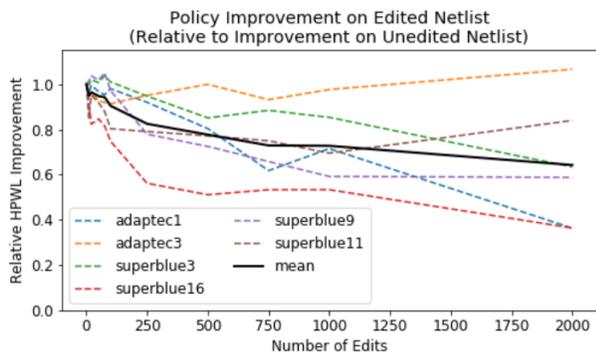
Fig. 5. Effect of Netlist Edits on Trained Spatial Cost Policies: With hundreds of netlist edits, trained policies still retain around 80% of their original improvement.

at the cost of an increase in the number of placement iterations to convergence.

C. Investigating Trained Agent Policies

1) Density Weight Control: As seen in Figure 3, the RePlAce rule for adjusting the density weight cost chooses a significantly different schedule than the baseline heuristic. Specifically, the agent tends to favor smaller increases in the density weight which leads to more placement iterations and longer time to converge to the target density. However, the choice of when to use small density weight updates and when to use larger ones seems to be both important and non-trivial because the final policies outperform both slow and fast static policies.

2) Spatial Cost Weighting: If we visualize the actions the trained 'spatial cost weighting' agent takes during placement (seen in Figure 4), we observe the network leverages both the current state of placement and placement structures such as the macro positions and high density logic clusters to control the evolution of the placement process. Interestingly, unlike the 'density weight' action space the number of placement iterations to convergence is not significantly larger than the baseline placement. The agent is able to improve results with only a small increase in the number of iterations.

D. Netlist Modification Test

In order to investigate how robust the learned policy is to changes in the netlist we performed the following test. We make synthetic edits to the benchmark netlists and rerun DREAMPlace with our learned policy to measure how the performance of our learned policy degrades as the netlist changes. Random edits which consist of adding and removing nodes as well as adding additional nets are performed according to Algorithm 2. As shown in Figure 5, we find that while the improvement of the learned policy degrades with these netlist edits we still see a significant improvement over the baseline placement scheme with hundreds of netlist edits. Specifically, on average we retain 78% of the trained benefit with 500 random edits and 73% of the benefit with 1000 random edits.

VI. Conclusion

In this work, we used GPU accelerated placement to train a reinforcement learning agent that augments a state of the art global placement algorithm. We have demonstrated this approach can learn heuristics that improve the final quality of results. Although we focus on improving half-perimeter wire length, this approach can

be extended to include many other downstream metrics such as congestion, timing or power.

Ideally, an agent trained with reinforcement learning would be able to perform placement on novel designs that were unseen during training. However, we observed that the training of a generalized agent is difficult. One major challenge to progress is a lack of open source datasets. Currently, open source datasets are too small to provide sufficient coverage of the distribution of all possible designs. We invite future work to tackle this challenge of generalization. As the placement environment is deterministic, exploration of model-based approaches to improve generalization is potentially promising. Meanwhile, since designs are often built iteratively, we believe our approach can learn strategies specific to each design, yielding benefits as the design is iteratively updated.

The result of our approach is more than 1% reduced HPWL on both widely used academic benchmarks and industry designs. While the training time is high this benefit is still significant to highly optimized designs and demonstrates the potential of data driven approaches to global placement optimization. To our knowledge this is the first attempt to integrate reinforcement learning directly into force based global placement algorithms.

References

[1] Y. Lin, S. Dhar, W. Li, H. Ren, B. Khailany, and D. Z. Pan, "Dreamplace: Deep learning toolkit-enabled gpu acceleration for modern vlsi placement," in *2019 56th ACM/IEEE Design Automation Conference (DAC)*, pp. 1–6, 2019.

[2] Y. Lin, W. Li, J. Gu, H. Ren, B. Khailany, and D. Z. Pan, "Abcd-place: Accelerated batch-based concurrent detailed placement on multi-threaded cpus and gpus," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, pp. 1–1, 2020.

[3] J. Lu, P. Chen, C. Chang, L. Sha, D. J. H. Huang, C. Teng, and C. Cheng, "eplace: Electrostatics based placement using nesterov's method," in *2014 51st ACM/EDAC/IEEE Design Automation Conference (DAC)*, pp. 1–6, 2014.

[4] C. Cheng, A. B. Kahng, I. Kang, and L. Wang, "Replace: Advancing solution quality and routability validation in global placement," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 38, no. 9, pp. 1717–1730, 2019.

[5] Y. E. Nesterov, "A method for solving the convex programming problem with convergence rate o$(1/k^2)$," *Dokl. Akad. Nauk SSSR*, vol. 269, pp. 543–547, 1983.

[6] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu, "Asynchronous methods for deep reinforcement learning," in *International conference on machine learning*, pp. 1928–1937, 2016.

[7] R. S. Sutton, "Learning to predict by the methods of temporal differences," *Machine learning*, vol. 3, no. 1, pp. 9–44, 1988.

[8] A. Mirhoseini, A. Goldie, M. Yazgan, J. Jiang, E. Songhori, S. Wang, Y.-J. Lee, E. Johnson, O. Pathak, S. Bae, A. Nazi, J. Pak, A. Tong, K. Srinivasa, W. Hang, E. Tuncer, A. Babu, Q. V. Le, J. Laudon, R. Ho, R. Carpenter, and J. Dean, "Chip placement with deep reinforcement learning," 2020.

[9] R. Manimegalai, E. Siva Soumya, V. Muralidharan, B. Ravindran, V. Kamakoti, and D. Bhatia, "Placement and routing for 3d-fpgas using reinforcement learning and support vector machines," in *18th International Conference on VLSI Design held jointly with 4th International Conference on Embedded Systems Design*, pp. 451–456, 2005.

[10] L. Chen, C. Huang, Y. Chang, and H. Chen, "A learning-based methodology for routability prediction in placement," in *2018 International Symposium on VLSI Design, Automation and Test (VLSI-DAT)*, pp. 1–4, 2018.

[11] A. F. Tabrizi, N. K. Darav, L. Rakai, A. Kennings, and L. Behjat, "Detailed routing violation prediction during placement using machine learning," in *2017 International Symposium on VLSI Design, Automation and Test (VLSI-DAT)*, pp. 1–4, 2017.

[12] A. F. Tabrizi, L. Rakai, N. K. Darav, I. Bustany, L. Behjat, S. Xu, and A. Kennings, "A machine learning framework to identify detailed routing short violations from a placed netlist," in *2018 55th ACM/ESDA/IEEE Design Automation Conference (DAC)*, pp. 1–6, 2018.

[13] H. Küttler, N. Nardelli, T. Lavril, M. Selvatici, V. Sivakumar, T. Rocktäschel, and E. Grefenstette, "Torchbeast: A pytorch platform for distributed rl," 2019.

[14] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," 2019.

[15] P. Wawrzynski, "Control policy with autocorrelated noise in reinforcement learning for robotics," *International Journal of Machine Learning and Computing*, vol. 5, no. 2, p. 91, 2015.

[16] G. E. Uhlenbeck and L. S. Ornstein, "On the theory of the brownian motion," vol. 36, p. 823, APS, 1930.

[17] G.-J. Nam, C. J. Alpert, P. Villarrubia, B. Winter, and M. Yildiz, "The ispd2005 placement contest and benchmarks suite," in *ISPD*, pp. 216–220, 2005.

[18] G.-J. Nam, "Ispd2006 placement contest: Benchmark suite and results," in *ISPD*, pp. 167–167, 2006.

[19] N. Viswanathan, C. Alpert, C. Sze, Z. Li, and Y. Wei, "The dac 2012 routability-driven placement contest and benchmark suite," in *DAC Design Automation Conference 2012*, pp. 774–782, 2012.