Statement on Reinforcement Learning for Chip Design

Anna Goldie & Azalia Mirhoseini - Mar 24, 2023

We appreciate the interest in our work on deep reinforcement learning (RL) for chip design, but we believe that a recent invited paper[1] at ISPD mischaracterizes it. Below, we summarize our work and offer a brief technical response.

First, some important context:

- **Our RL method has been used in production for multiple generations of Google's flagship AI accelerator (TPU)**, including the latest, and chips with layouts generated by our method have been manufactured and are currently running in Google datacenters.
    - Being physically manufactured at sub-10 nm is a level of validation that goes far beyond what is present in almost any paper.
    - ML-generated placements have to be significantly better than what Google engineers generate (i.e., superhuman), or they aren't worth the risk.
    - Our goal has always been to help chip designers do their jobs better and faster, and that is reflected in all of our design decisions, including evaluation metrics, benchmarks, and algorithms.
- We went through seven months of peer review with Nature (transparent peer review)
    - This involved extensive back and forth with three referees (2 experts in physical design, 1 expert in reinforcement learning).
- The results in our Nature paper were independently replicated by the TF-Agents team (announcement, reproduction led by Dr. Sergio Guadarrama)
- We open-sourced our code on 1/18/2022 (announcement).
    - 100+ forks and 500+ stars as of 3/18/2023 (Circuit Training Open-Source Repo)
    - Developing and open-sourcing this highly-optimized, distributed RL framework required an enormous amount of non-trivial engineering effort, and has applications beyond chip floorplanning or even EDA.

[1] Note that invited papers do not go through peer review.

    - Note that there was no requirement to open-source our code, and it is in fact unusual to do so in the field of commercial EDA.
        - AlphaGo, arguably the most famous RL paper in Nature, did not open-source its code.
        - Neither did last year's Best Paper at DAC, the top chip design conference.
        - Commercial EDA companies like Synopsys and Cadence do not even publish their methods (e.g., CMP), let alone open-source them.

    - We are very grateful to Google leadership for allowing us to share this code with the larger community, so that others can more easily build on our work.

- Following publication of our method, many RL for chip placement papers building on our work have been published at ML and EDA conferences, and there have been press releases and news articles from companies like NVIDIA, Synopsys, Cadence, and Samsung announcing their use of RL for chip design (listed below).

Our high-level response to technical aspects of the ISPD paper:

- The ISPD paper performed no pre-training for our method, Circuit Training (CT), meaning that the RL agent was reset each time it saw a new chip.
    - A learning-based method will of course take longer to learn and perform worse if it has never seen a chip before! See Fig. 4 of the Nature paper.
    - We pre-train on 20 blocks before evaluating on the held-out test cases in Table 1.

- Far fewer compute resources were applied to training CT than described in our Nature paper (half the number of GPUs and an order of magnitude fewer RL environments).

- Plots linked in the ISPD paper's accompanying document suggest that CT was not trained properly and was cut off while the RL agent was still learning.

    - Interdisciplinary research can be very impactful, but it is also challenging, because it requires the researcher to ramp up in an entirely new field. It took us years to ramp up on physical design, which was only possible due to the generosity of the TPU physical design team and experts like Prof. Andrew Kahng. We would have been happy to provide feedback and assistance on the RL training if the paper had been shared with us in advance of publication.

- RePlAce was state-of-the-art when we published the Nature paper (and arXiv preprint). We outperformed it then and we outperform it now, as shown in the ISPD paper's Table 1, even ignoring all of the above issues. (Incidentally, RePlAce, as noted in a footnote of the ISPD paper, is not able to produce any result at all for 2 out of the 6 test cases.)

- Although the study is entitled “Assessment of Reinforcement Learning for Macro Placement”, it does not compare against or even acknowledge any of the recent RL methods building on our work (link).

    - Instead, the ISPD paper compares CT against AutoDMP (ISPD 2023) and (presumably) the latest version of CMP, a black-box, closed-source commercial autoplacer. Neither of these methods were available when we released our paper in 2020.

- The ISPD paper focuses on the use of the initial placement from physical synthesis to cluster standard cells, but this is of no practical concern.

    - Physical synthesis must be performed before running any placement method. This is standard practice in chip design[2] and this is reflected in Figure 2 of the ISPD paper, which shows that physical synthesis is run before all placement methods presented.

    - As a preprocessing step, we reuse the output of physical synthesis (the previous step run for all methods) to perform clustering of standard cells. To be clear, our (macro placement) method does not place standard cells, which are already well-handled by prior methods (e.g., DREAMPlace).

    - In each RL episode, we provide the RL agent with a netlist of unplaced macros (memory components) and unplaced clusters of standard cells (logic gates), and the RL agent places these macros one-at-a-time onto a blank canvas.

    - We documented these details in our open-source repository nine months ago (link), and provide an API to perform this preprocessing step for any given netlist. The details of this preprocessing step have no bearing on our experimental results or conclusions[3] and we are frankly surprised that anyone would find them interesting.

We understand that a great deal of effort went into the ISPD paper and appreciate the community’s interest in our work. We also recognize that our method is not perfect, and that there almost certainly exist test cases for which the current iteration is less effective. Ultimately, we believe that this is just the beginning and that learning-based methods for chip design will have profound impact on both hardware and machine learning itself.

---

[2] https://www.synopsys.com/glossary/what-is-physical-synthesis.html

[3] In our Nature paper, we of course provide the same clustered netlist to each placement method, as shown in Extended Data Fig. 1.

Appendix

The list below is far from comprehensive (e.g., the Nature paper and arXiv preprint together have been cited ~400 times), but we think it's clear that progress is full-speed ahead:

- **On Joint Learning for Solving Placement and Routing in Chip Design**, Ruoyu Cheng & Junchi Yan. https://arxiv.org/abs/2111.00234. NeurIPS 2021.

- **Flexible Chip Placement via Reinforcement Learning**. Fu-Chieh Chang, Yu-Wei Tseng, Ya-Wen Yu, Ssu-Rui Lee , Alexandru Cioba, I-Lun Tseng, Da-shan Shiu, Jhih-Wei Hsu, Cheng-Yuan Wang, Chien-Yi Yang, Ren-Chu Wang, Yao-Wen Chang, Tai-Chen Chen and Tung-Chieh Chen. https://dl.acm.org/doi/pdf/10.1145/3489517.3530617. DAC 2022 Late Breaking Results.

- **The Policy-gradient Placement and Generative Routing Neural Networks for Chip Design**, Ruoyu Cheng, Xianglong Lyu, Yang Li, Junjie Ye, Jianye HAO, Junchi Yan https://openreview.net/pdf?id=uNYqDfPEDD8, NeurIPS 2022.

- **MaskPlace: Fast Chip Placement via Reinforced Visual Representation Learning**, Yao Lai, Yao Mu, Ping Luo, https://arxiv.org/abs/2211.13382, NeurIPS 2022.

Following the publication of our method, there have also been many press releases and news articles about major chip design companies using RL:

- **Synopsys Article on Reinforcement Learning:** https://www.forbes.com/sites/moorinsights/2020/04/20/using-ai-to-build-better-chips/?sh=3b605ad4306c

- **Synopsys Press Release on Using RL to Achieve Chip Design:** https://www.prnewswire.com/news-releases/synopsys-ai-driven-design-system-enables-renesas-to-achieve-breakthrough-in-productivity-301269870.html

- **Samsung using Synopsys's RL to design chips:** https://www.wired.com/story/samsung-ai-designed-chip-soon-others-too/

- **Cadence using RL:**
- https://www.zdnet.com/article/ai-on-the-bench-cadence-offers-machine-learning-to-smooth-chip-design/

- **NVIDIA R&D Chief on How AI is Improving Chip Design** with specific mention of RL to place the transistors**:** https://arxiv.org/abs/2107.07044, https://www.hpcwire.com/2022/04/18/nvidia-rd-chief-on-how-ai-is-improving-chip-design/

- **WIRED article on NVIDIA, MIT, and Google using RL for chip design:** https://www.wired.com/story/fit-billions-transistors-chip-let-ai-do/amp
- https://www.forbes.com/sites/karlfreund/2021/08/09/using-ai-to-help-design-chips-has-become-a-thing/
- https://www.nytimes.com/2021/05/07/technology/semiconductor-chip-innovation-boom.html
- https://analyticsindiamag.com/how-reinforcement-learning-is-advancing-chip-designing/
- https://gradientflow.com/applications-of-reinforcement-learning-recent-examples-from-large-us-companies/
- https://www.eenewseurope.com/news/ai-boost-standard-cell-layout-3nm
- https://www.emarketer.com/content/samsung-s-use-of-ai-chip-design-signals-industry-shift-away-dominant-suppliers
- https://www.extremetech.com/computing/325878-ai-powered-electronic-design-automation-tools-could-redefine-chipmaking
- https://inside.com/campaigns/inside-ai-2021-08-16-29025
- https://venturebeat.com/2021/07/25/with-post-pandemic-ai-weve-now-stepped-into-the-age-of-acceleration/
- https://www.eetasia.com/cadence-unveils-cerebrus/
- https://www.allaboutcircuits.com/news/entering-the-era-of-ai-ic-design-where-to-ees-fit-in/