“For the Record”

Andrew B. Kahng
CSE and ECE Departments, UC San Diego
abk@eng.ucsd.edu

To my colleagues in EDA and IC design:

With thanks to many who have engaged in discussions, written emails, asked probing questions, etc. – I would like to provide the following “for the record”.

- I served as peer reviewer (“Reviewer #3”)[1] for the Mirhoseini & Goldie et al. Nature paper, from November 2020 through March 2021. My reviews, and the authors’ rebuttals, are in the Peer Review File.

- After the Mirhoseini & Goldie et al. paper was accepted by Nature, I was requested[2] to provide a News and Views commentary. The commentary followed guidance[3] from Nature. It invited attention to, and replication of, a potentially groundbreaking work in the VLSI CAD field.

- I stand by what I wrote in both my reviews and my commentary. Your feedback, including criticism, is always welcome.

- The Nature paper challenges traditional wisdoms. Its generated macro placements were used in at least one taped-out product[4] at a sub-10nm technology node. The work clearly merits further evaluation and understanding of its intrinsic scientific and technical contributions. I hope that we all agree on these points.

- Further, I believe it is well-understood that the reported **methods** are not fully implementable based on what is provided in the Nature paper and the Circuit Training repository. While the RL method is open-source, key preprocessing steps and interfaces are not yet available. I have been informed that the Google team is actively working to address this. Remedying this gap is necessary to achieve scientific clarity and a foundation upon which the field can move forward.

- It has been a year since publication of the Nature paper and its statement of “Data and Code Availability”. I look forward to the research and technical community – including Google – coming together to accelerate availability of an open-source foundation.[5] The community’s participation and support in this quest – along with a calm, open technical dialogue – would be wonderful to see and is warmly invited.

Thank you and best wishes to all.

Notes.

[1] (1) With their second revision (3/26/21), the authors' rebuttal included, "Yes, your understanding of our method's role in the TPU-v5 tapeout is correct!".

(2) The second revision also added text that is seen in the published paper: "Data and Code Availability. The authors declare that the data supporting the findings of this study are available within the paper and its supplementary information files. The code used to generate this data will be made available by the corresponding authors upon reasonable request."

As a reviewer, I ascribed some (positive) weight to both (1) and (2). Indeed, the commitment made by authors in (2), in my opinion, goes beyond norms of documentation and reproducibility (data, code, scripts, etc.) seen in the VLSI CAD research literature.

All comments and criticisms are welcome.

[2] I agreed to write the commentary with the understanding that (A) the review-rebuttal thread would be accessible to readers, and (B) the authors had committed to making their code available. I was subsequently informed by the editor that "the authors of the paper opted out of transparent peer review". (I ascribed this to the presence of information designated confidential in the authors' rebuttal; see page 42 of the Peer Review File.) I am glad that as of 4/22/2022, both the Nature paper and the Peer Review File are publicly available and can speak for themselves.

[3] The editor's guidance: "The object would be to describe for a non-specialist audience what the authors have done; to comment on the paper's strengths, implications and shortcomings; and to sketch out future directions of research opened up by the findings. The deadline would be two to three weeks from today, and the word limit would be 800–900."

Again, all comments and criticisms are welcome.

[4] I.e., TPU-v5. Based on information shared by Google engineers in recent interactions, I believe this sentence to be unquestionably true.

[5] Efforts to this end (open testcases; source codes and scripts to fill in missing elements such as gridding or netlist clustering; format translators; reproduction of public Circuit Training results; ...) have begun, and may be tracked at https://github.com/TILOS-AI-Institute/MacroPlacement .

At the 59th Design Automation Conference in San Francisco, the program of the birds-of-a-feather meeting, "Open-Source EDA and Benchmarking Summit" (https://59dac.conference-program.com/session/?sess=sess294) on July 12, 2022, will include an update on these efforts.