Thesis Summary

This document summarises the Master's thesis titled *PCB Component Placement using Reinforcement Learning (RL)*. First, the topic is introduced with a brief motivation and problem description. Next, the literature review outcomes are described before investigating state-of-the-art approaches. Lastly, our novel placement methodology and an accompanying solution are presented.

Printed Circuit Boards (PCBs) are at the centre of all electronic systems and provide the surface onto which electronic components are soldered and the routing infrastructure for wiring them together. They are three-dimensional structures with many internal copper layers along which wires are etched spanning both horizontally and vertically. The PCB design process involves translating the logical topology of a circuit (e.g. a schematic) into a manufacturable geometrical representation and is divided into component placement followed by routing. Automated placement techniques have had some success [1, 2], but they often lack the intuitive understanding of human engineers. Partly for this reason, such tools did not gain traction, and presently it is a predominantly manual process. The thesis investigates the current state-of-the-art [3] and proposes a novel end-to-end machine-learning approach to optimise the placement of components on a PCB. We aim to create an AI-assisted workflow that enhances productivity by allowing designers to focus on higher-level tasks, reducing design time while leveraging the differentiating benefits of customised solutions.

The placement task accepts a circuit netlist describing the logical representation of the circuit accompanied by geometrical information of the individual components (e.g. dimensions) [4]. The task is at least NP-Complete [5] and is concerned with identifying the best spatial location and orientation for all the components given solution constraints (e.g. no overlaps) and optimisation goals (e.g. minimised wirelength) [6]. Placement techniques can be classified as constructive and iterative. The former order the circuit netlist according to a criterion (e.g. component area) and sequentially place the constituents onto an empty layout region. The latter starts with a randomly initialised placement that is iteratively improved until a terminal condition is reached.

PCB and Integrated Circuit (IC) placement processes share many commonalities, but problem size is one aspect they greatly differ. As a result, research on the former is sparse [7, 8] and limited primarily to particular placement co-optimisation tasks [9-12]. By contrast, since IC (digital) design is only feasible with automated tools, it has a thriving research community with contributions from academia [14-19] and industry [20-23]. Since the 1960s, four placement categories have emerged, namely, partitioning-based methods inspired by graph theory [24, 25], black-box optimisation methods [2, 26-29], analytic placement [16, 17, 30-38] and presently, learning-based solutions [3, 39, 40]. Metaheuristics proved effective while offering flexibility in defining the objective function. However, driven by stochastic decision processes, they were not feasible on circuits exceeding 1e5 elements. Analytic placement appeared as a more scalable alternative and is presently considered state-of-the-art [14, 15, 22]. It requires a differentiable objective that is optimised using numerical techniques. Recently hybrid [39, 40] and end-to-end solutions [3] for floorplanning using RL have been proposed, albeit they are still in their infancy. RL offers attractive solutions to such problems, particularly for its ability to represent vast state spaces and generalise to unseen similar ones.

The thesis initially investigates the state-of-the-art constructive placement methodology by Mirhoseini et al. [3] and proposes a novel formulation for iterative placement. Mirhoseini et al. proposed using a neural network to represent the problem state in a compressed manner. They offer a novel edge-based graph neural network to automatically extract features from the circuit netlist and, together with task-relevant metadata, they predict a placement quality metric as a linear

combination of wirelength, congestion and density. After removing the final prediction layer, the authors encoded the problem state and used PPO [48] to train RL policies for predicting placement probabilities over a discretised layout region. Using this approach as a guideline, we trained a neural network to predict circuit wirelength and subsequently trained policies using TRPO [47] and PPO in a similar way. Using unseen circuits, we achieved an accuracy of 69.5% for the graph-level wirelength prediction task, albeit subsequent RL policies were significantly outperformed by 41% after establishing a baseline with Simulated Annealing (SA) [42, 43]. While the authors provided innovative ideas, we were concerned about important details as well as their evaluation procedures. Recent literature shared similar concerns [40] and even debunked some of the original claims [50].

Learning from the previous limitations, we formulated the iterative PCB placement task as a Markov Decision Process (MDP). First, we studied its mechanics in a constrained environment (single-component approach), then pooled our findings and adapted the setup to yield general solutions (multi-component approach). Concerning the former, the agent represents a single component. The goal, starting from an arbitrary location on the PCB, is to orient the component within a fixed optimised layout while minimising wirelength and avoiding overlap in the terminal state. The observation space mainly captures the perceived surroundings and direction information related to movement. We propose two fundamentally distinct reward signals to motivate the desired behaviour. First, we attempt to mimic the expert by using expert positioning as the goal. Secondly, we motivate self-improvement using problem-related performance metrics. We extensively study the problem by investigating a variety of RL algorithms [46-50], discrete and continuous action spaces, and environment features such as optimal episode length, step size and replay buffer size, especially in cases of adaptive reward signals. Our overarching aim was to learn fundamental placement techniques applicable to unseen circuits. While we learned a lot from these experiments, we concluded that expert-generated data introduce inconsistencies that prevent generalisation. These inconsistencies leaked into our problem setup through the reward signal when mimicking the expert and through the observation space arising from the fixed portion of the circuit.

To address data inconsistencies, we adapted the training process to suit a multi-component setup that simultaneously places all components except one. In other words, we lock a single component to serve as an anchor (typically the main IC) and for every step in the episode, we invoke the policy on all the moveable components (all except the anchor), each time sampling an updated environment state. This training process collects highly diverse data points because every component in the circuit contributes a different perspective into the problem. In pursuit of generalisation, we use a dataset of nine circuits sampled from real-world applications having up to 12 components. We train on six and evaluate on the remaining three. We establish a baseline using SA [42, 43] and compare placements in terms of post-routing wirelength [44] after 600 iterations. For each circuit, we average four evaluations starting from different conditions. Over the three unseen circuits, our best configuration outperforms SA by up to 21%. Additionally, we observed that optimised placements generated by our method can converge over an order of magnitude faster, suggesting that leveraging experience over stochastic decision-making is beneficial. Quantitatively the policies exhibit fundamental placement techniques (i.e. witnessed taking actions to minimise wirelength) and emergent collaborative or competitive features conditional on the reward scheme.

The thesis proposed a novel MDP formulation for iterative placement and delivered a general RL solution. Although further optimisations and additional key features are necessary, this work offers a promising direction towards automating PCB component placement. A demonstration is available on GitHub, while the complete thesis and associated publication are available on my website.

References

[1] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, ‘Optimization by Simulated Annealing’, *Science*, vol. 220, no. 4598, pp. 671–680, 1983. doi: 10.1126/science.220.4598.671.

[2] C. Sechen and A. Sangiovanni-Vincentelli, ‘The TimberWolf placement and routing package’, *IEEE J. Solid-State Circuits*, vol. 20, no. 2, pp. 510–522, Apr. 1985, doi: 10.1109/JSSC.1985.1052337.

[3] A. Mirhoseini *et al.*, ‘A graph placement methodology for fast chip design’, *Nature*, vol. 594, no. 7862, pp. 207–212, Jun. 2021, doi: 10.1038/s41586-021-03544-w.

[4] H. Kaeslin, *Digital Integrated Circuit Design From VLSI Architectures to CMOS Fabrication*, 1st ed. Cambridge University Press, 2008.

[5] M. R. Garey, D. S. Johnson, and L. Stockmeyer, ‘Some simplified NP-complete graph problems’, *Theoretical Computer Science*, vol. 1, no. 3, pp. 237–267, 1976, doi: 10.1016/0304-3975(76)90059-1.

[6] A. B. Kahng, J. Lienig, I. L. Markov, and J. Hu, *VLSI Physical Design: From Graph Partitioning to Timing Closure*. Cham: Springer International Publishing, 2022. doi: 10.1007/978-3-030-96415-3.

[7] T. Badriyah, F. Setyorini, and N. Yuliawan, ‘The implementation of Genetic Algorithm and Routing Lee for PCB design optimization’, in *2016 International Conference on Informatics and Computing (ICIC)*, Mataram, Indonesia: IEEE, 2016, pp. 148–153. doi: 10.1109/IAC.2016.7905706.

[8] C.-K. Cheng, C.-T. Ho, and C. Holtz, ‘Net Separation-Oriented Printed Circuit Board Placement via Margin Maximization’, in *2022 27th Asia and South Pacific Design Automation Conference (ASP-DAC)*, Taipei, Taiwan: IEEE, Jan. 2022, pp. 288–293. doi: 10.1109/ASP-DAC52403.2022.9712480.

[9] A. Alexandridis, E. Paizis, E. Chondrodima, and M. Stogiannos, ‘A particle swarm optimization approach in printed circuit board thermal design’, *ICA*, vol. 24, no. 2, pp. 143–155, Mar. 2017, doi: 10.3233/ICA-160536.

[10] F. S. Ismail, R. Yusof, and M. Khalid, ‘Optimization of electronics component placement design on PCB using self organizing genetic algorithm (SOGA)’, *J Intell Manuf*, vol. 23, no. 3, pp. 883–895, Jun. 2012, doi: 10.1007/s10845-010-0444-x.

[11] Puqi Ning, Fei Wang, and K. D. T. Ngo, ‘Automatic layout design for power module’, *IEEE Trans. Power Electron.*, vol. 28, no. 1, pp. 481–487, Jan. 2013, doi: 10.1109/TPEL.2011.2180739.

[12] P. Ning, H. Li, Y. Huang, and Y. Kang, ‘Review of power module automatic layout optimization methods in electric vehicle applications’, *Chin. J. Electr. Eng.*, vol. 6, no. 3, pp. 8–24, Sep. 2020, doi: 10.23919/CJEE.2020.000015.

[13] X. Gao *et al.*, ‘Interactive Analog Layout Editing With Instant Placement and Routing Legalization’, *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 42, no. 3, pp. 698–711, Mar. 2023, doi: 10.1109/TCAD.2022.3190234.

[14] 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 *Proceedings of the 56th Annual Design Automation Conference 2019*, Las Vegas NV USA: ACM, Jun. 2019, pp. 1–6. doi: 10.1145/3316781.3317803.

[15] C.-K. Cheng, A. B. Kahng, I. Kang, and L. Wang, ‘RePlAce: Advancing Solution Quality and Routability Validation in Global Placement’, *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 38, no. 9, pp. 1717–1730, Sep. 2019, doi: 10.1109/TCAD.2018.2859220.

[16] R. Hao, Y. Cai, and Q. Zhou, ‘Intelligent and kernelized placement: A survey’, *Integration*, vol. 86, pp. 44–50, Sep. 2022, doi: 10.1016/j.vlsi.2022.05.002.

[17] I. L. Markov, J. Hu, and M.-C. Kim, ‘Progress and Challenges in VLSI Placement Research’, *Proc. IEEE*, vol. 103, no. 11, pp. 1985–2003, Nov. 2015, doi: 10.1109/JPROC.2015.2478963.

[18] G. Huang *et al.*, ‘Machine Learning for Electronic Design Automation: A Survey’, *ACM Trans. Des. Autom. Electron. Syst.*, vol. 26, no. 5, pp. 1–46, Sep. 2021, doi: 10.1145/3451179.

[19] D. S. Lopera, L. Servadei, G. N. Kiprit, S. Hazra, R. Wille, and W. Ecker, ‘A Survey of Graph Neural Networks for Electronic Design Automation’, in *2021 ACM/IEEE 3rd Workshop on Machine Learning for CAD (MLCAD)*, Raleigh, NC, USA: IEEE, Aug. 2021, pp. 1–6. doi: 10.1109/MLCAD52597.2021.9531070.

[20] B. Khailany *et al.*, ‘Accelerating Chip Design With Machine Learning’, *IEEE Micro*, vol. 40, no. 6, pp. 23–32, Nov. 2020, doi: 10.1109/MM.2020.3026231.

[21] C.-T. Ho *et al.*, ‘NVCell 2: Routability-Driven Standard Cell Layout in Advanced Nodes with Lattice Graph Routability Model’, in *Proceedings of the 2023 International Symposium on Physical Design*, Virtual Event USA: ACM, Mar. 2023, pp. 44–52. doi: 10.1145/3569052.3578920.

[22] A. Agnesina *et al.*, ‘AutoDMP: Automated DREAMPlace-based Macro Placement’, in *Proceedings of the 2023 International Symposium on Physical Design*, Virtual Event USA: ACM, Mar. 2023, pp. 149–157. doi: 10.1145/3569052.3578923.

[23] Y. Zhang, H. Ren, and B. Khailany, ‘GRANNITE: Graph Neural Network Inference for Transferable Power Estimation’, in *2020 57th ACM/IEEE Design Automation Conference (DAC)*, San Francisco, CA, USA: IEEE, Jul. 2020, pp. 1–6. doi: 10.1109/DAC18072.2020.9218643.

[24] A. E. Dunlop and B. W. Kernighan, ‘A Procedure for Placement of Standard-Cell VLSI Circuits’, *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 4, no. 1, pp. 92–98, Jan. 1985, doi: 10.1109/TCAD.1985.1270101.

[25] M. A. Breuer, ‘A class of min-cut placement algorithms’, in *Papers on Twenty-five years of electronic design automation*, ACM, Jun. 1988, pp. 142–148. doi: 10.1145/62882.62896.

[26] J. P. Cohoon and W. D. Paris, ‘Genetic Placement’, *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 6, no. 6, pp. 956–964, Nov. 1987, doi: 10.1109/TCAD.1987.1270337.

[27] J. P. Cohoon, S. U. Hegde, W. N. Martin, and D. Richards, ‘Floorplan design using distributed genetic algorithms’, in *[1988] IEEE International Conference on Computer-Aided Design (ICCAD-89) Digest of Technical Papers*, Santa Clara, CA, USA: IEEE Comput. Soc. Press, 1988, pp. 452–455. doi: 10.1109/ICCAD.1988.122547.

[28] J. P. Cohoon, S. U. Hegde, W. N. Martin, and D. S. Richards, ‘Distributed genetic algorithms for the floorplan design problem’, *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 10, no. 4, pp. 483–492, Apr. 1991, doi: 10.1109/43.75631.

[29] Maogang Wang, Xiaojian Yang, and M. Sarrafzadeh, ‘Dragon2000: standard-cell placement tool for large industry circuits’, in *IEEE/ACM International Conference on Computer Aided Design. ICCAD - 2000. IEEE/ACM Digest of Technical Papers (Cat. No.00CH37140)*, San Jose, CA, USA: IEEE, 2000, pp. 260–263. doi: 10.1109/ICCAD.2000.896483.

[30] P. Spindler, U. Schlichtmann, and F. M. Johannes, ‘Kraftwerk2—A Fast Force-Directed Quadratic Placement Approach Using an Accurate Net Model’, *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 27, no. 8, pp. 1398–1411, Aug. 2008, doi: 10.1109/TCAD.2008.925783.

[31] M.-C. Kim, J. Hu, D.-J. Lee, and I. L. Markov, ‘A SimPLR method for routability-driven placement’, in *2011 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, San Jose, CA, USA: IEEE, Nov. 2011, pp. 67–73. doi: 10.1109/ICCAD.2011.6105307.

[32] M.-C. Kim, N. Viswanathan, C. J. Alpert, I. L. Markov, and S. Ramji, ‘MAPLE: multilevel adaptive placement for mixed-size designs’, in *Proceedings of the 2012 ACM international symposium on International Symposium on Physical Design*, Napa California USA: ACM, Mar. 2012, pp. 193–200. doi: 10.1145/2160916.2160958.

[33] M.-C. Kim and I. L. Markov, ‘ComPLx: A Competitive Primal-dual Lagrange Optimization for Global Placement’, in *Proceedings of the 49th Annual Design Automation Conference*, San Francisco California: ACM, Jun. 2012, pp. 747–752. doi: 10.1145/2228360.2228496.

[34] M.-K. Hsu *et al.*, ‘NTUplace4h: A Novel Routability-Driven Placement Algorithm for Hierarchical Mixed-Size Circuit Designs’, *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 33, no. 12, pp. 1914–1927, Dec. 2014, doi: 10.1109/TCAD.2014.2360453.

[35] J. Lu *et al.*, ‘ePlace: Electrostatics Based Placement Using Nesterov’s Method’, in *Proceedings of the 51st Annual Design Automation Conference*, San Francisco CA USA: ACM, Jun. 2014, pp. 1–6. doi: 10.1145/2593069.2593133.

[36] T. Lin, C. Chu, and G. Wu, ‘POLAR 3.0: An ultrafast global placement engine’, in *2015 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, Austin, TX, USA: IEEE, Nov. 2015, pp. 520–527. doi: 10.1109/ICCAD.2015.7372614.

[37] J. Gu, Z. Jiang, Y. Lin, and D. Z. Pan, ‘DREAMPlace 3.0: multi-electrostatics based robust VLSI placement with region constraints’, in *Proceedings of the 39th International Conference on Computer-Aided Design*, Virtual Event USA: ACM, Nov. 2020, pp. 1–9. doi: 10.1145/3400302.3415691.

[38] P. Liao, S. Liu, Z. Chen, W. Lv, Y. Lin, and B. Yu, ‘DREAMPlace 4.0: Timing-driven Global Placement with Momentum-based Net Weighting’, in *2022 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, Antwerp, Belgium: IEEE, Mar. 2022, pp. 939–944. doi: 10.23919/DATE54114.2022.9774725.

[39] Y.-H. Huang *et al.*, ‘Routability-Driven Macro Placement with Embedded CNN-Based Prediction Model’, in *2019 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, Florence, Italy: IEEE, Mar. 2019, pp. 180–185. doi: 10.23919/DATE.2019.8715126.

[40] Q. Xu *et al.*, ‘GoodFloorplan: Graph Convolutional Network and Reinforcement Learning-Based Floorplanning’, *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 41, no. 10, pp. 3492–3502, Oct. 2022, doi: 10.1109/TCAD.2021.3131550.

[41] P. Crocker, *Physically Constrained PCB Placement Using Deep Reinforcement Learning*. Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, 2021. Accessed: Jul. 20, 2023. [Online]. Available: https://dspace.mit.edu/handle/1721.1/139247

[42] C. Holtz, D. J. Merrill, and M. Woo, ‘SA-PCB: Simulated Annealing-based Placement For PCB Layout’, *GitHub repository*. GitHub, 2020.

[43] D. J. Merrill, ‘Hungry for Fully Automated Design of Embedded Systems?’, University of California, San Diego, 2021. Accessed: Jul. 20, 2023. [Online]. Available: https://escholarship.org/uc/item/55d5c7s8

[44] T.-C. Lin, C. Holtz, Yenyi, and D. J. Merrill, ‘Printed Circuit Board (PCB) router’, *GitHub repository*. GitHub, 2020.

[45] S. Fujimoto, H. van Hoof, and D. Meger, ‘Addressing Function Approximation Error in Actor-Critic Methods’. arXiv, Oct. 22, 2018. Accessed: Jun. 04, 2023. [Online]. Available: http://arxiv.org/abs/1802.09477

[46] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, ‘Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor’. arXiv, Aug. 08, 2018. Accessed: Jun. 04, 2023. [Online]. Available: http://arxiv.org/abs/1801.01290

[47] J. Schulman, S. Levine, P. Moritz, M. I. Jordan, and P. Abbeel, ‘Trust Region Policy Optimization’. arXiv, Apr. 20, 2017. Accessed: Jun. 04, 2023. [Online]. Available: http://arxiv.org/abs/1502.05477

[48] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, ‘Proximal Policy Optimization Algorithms’. arXiv, Aug. 28, 2017. Accessed: Jul. 19, 2023. [Online]. Available: http://arxiv.org/abs/1707.06347

[49] T. P. Lillicrap *et al.*, ‘Continuous control with deep reinforcement learning’. arXiv, Jul. 05, 2019. Accessed: Jun. 03, 2023. [Online]. Available: http://arxiv.org/abs/1509.02971

[50] C.-K. Cheng, A. B. Kahng, S. Kundu, Y. Wang, and Z. Wang, ‘Assessment of Reinforcement Learning for Macro Placement’. arXiv, Mar. 27, 2023. Accessed: Jun. 16, 2023. [Online]. Available: http://arxiv.org/abs/2302.11014