Device placement optimization using simulated reinforcement learning.

Alexander Frost for R244

Project Outline

- Device placement optimisation of computation graphs.
- Reinforcement learning approach based on the 2017 + 2018 papers by Mirhoseini et al., presented by Frank in last session.
- Simulated environment using Park.
- Extension: learning different representation of input using graph convolutional layers and principal neighbourhood aggregation.

What is Park?

- Park: An Open Platform for Learning-Augmented Computer Systems
- Aim: provide a standardised platform for RL in computer systems, similar to how OpenAI Gym is widely used for robotics and gaming
- Provides simulated environments for 12 systems tasks, including adaptive bitrate video streaming, circuit design, *and TensorFlow device placement*

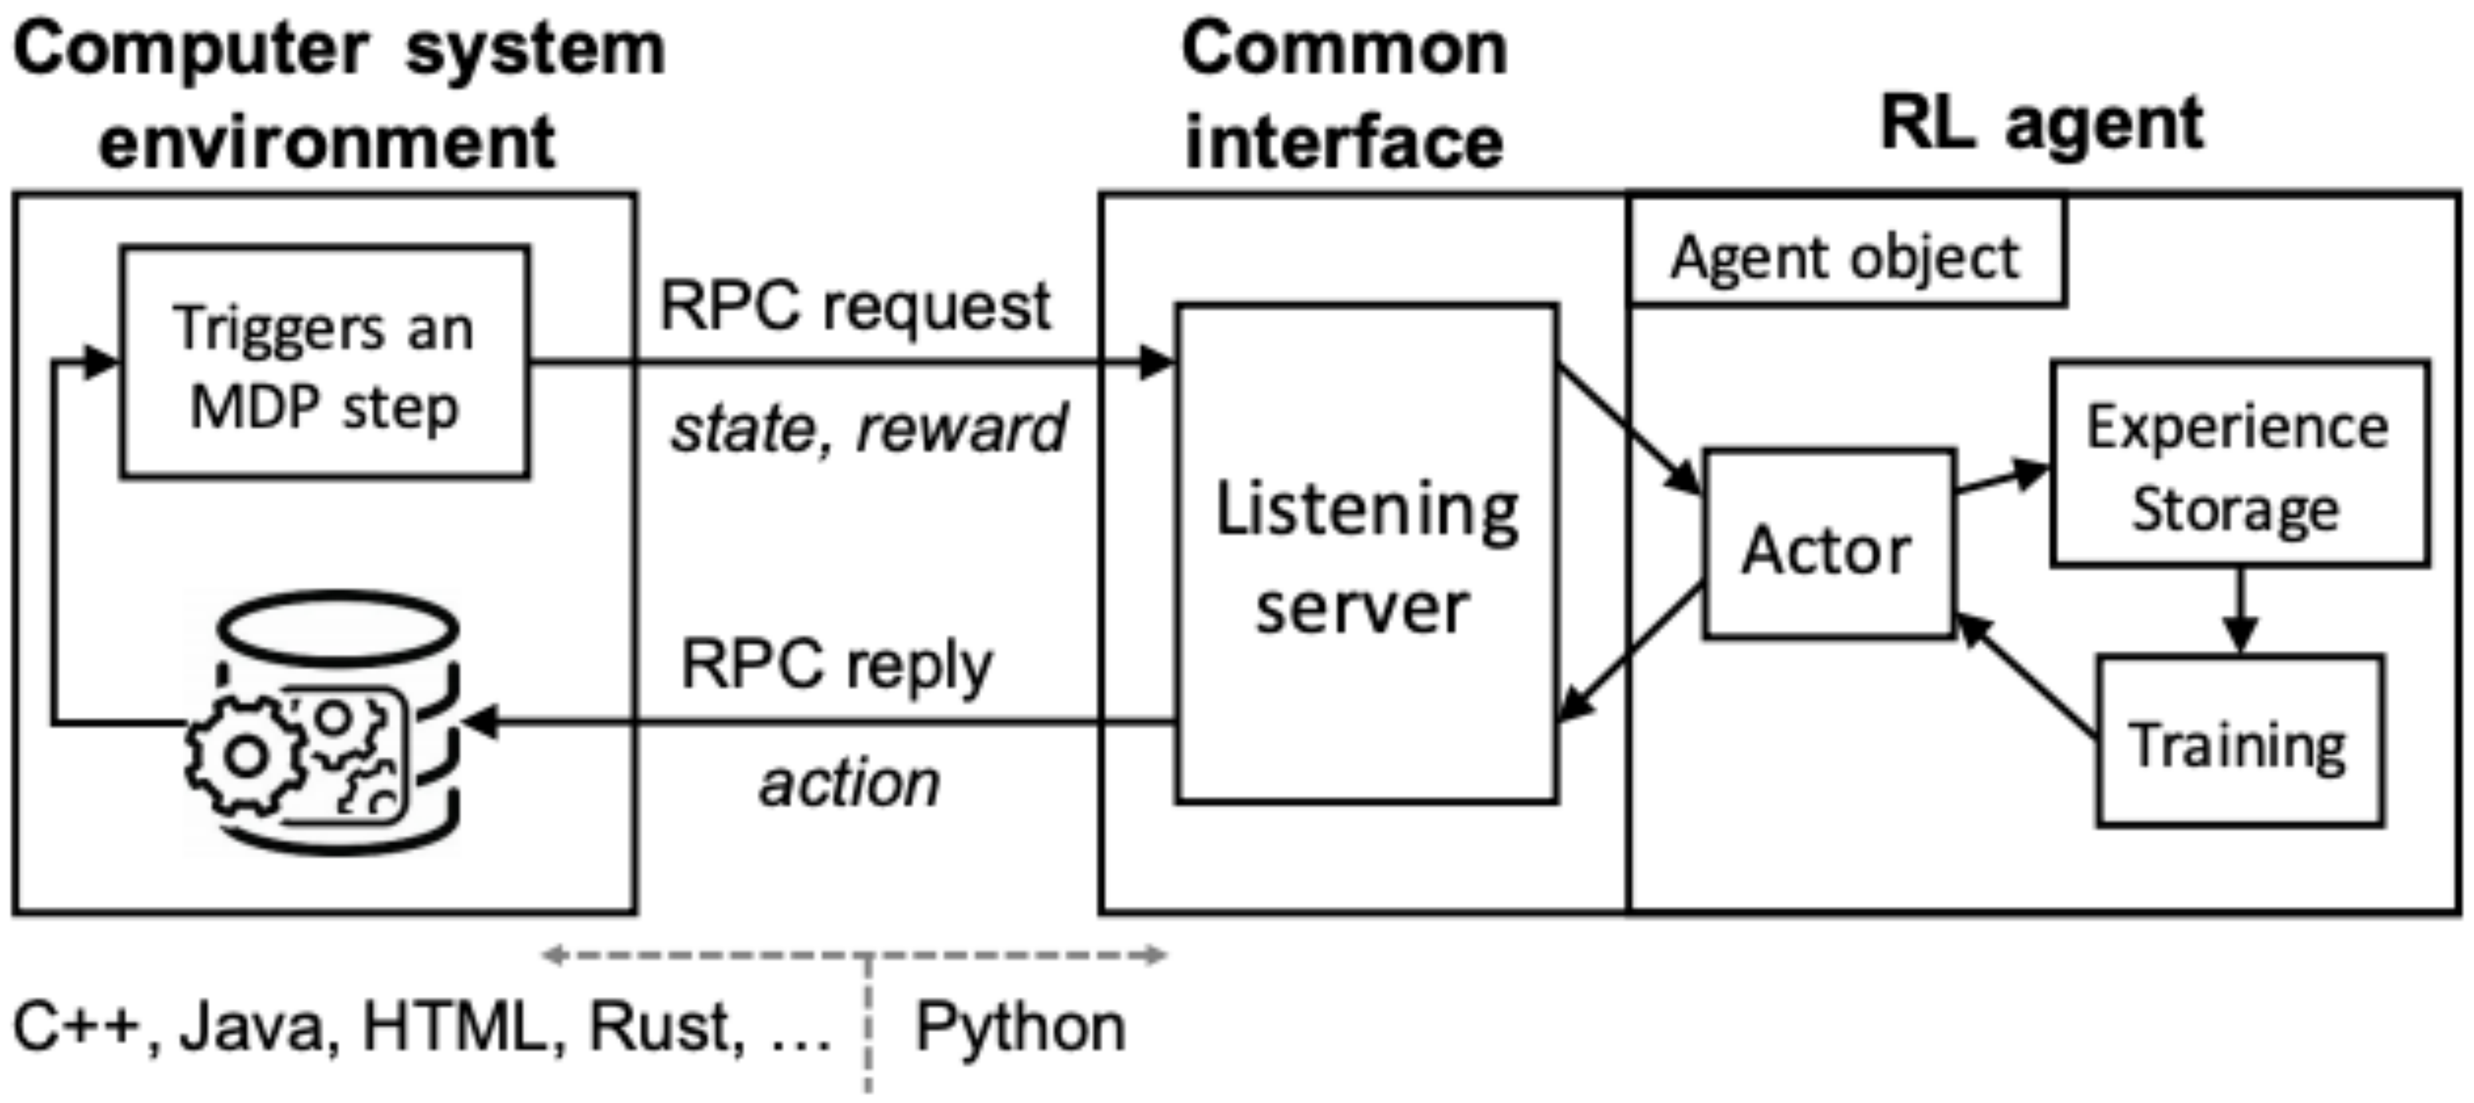

Device placement: original approach (2017)

1. Heuristically map operator nodes to colocation groups
2. Define ordering over nodes, embed into vector using encoder RNN
3. In same order, predict devices autoregressively using decoder RNN

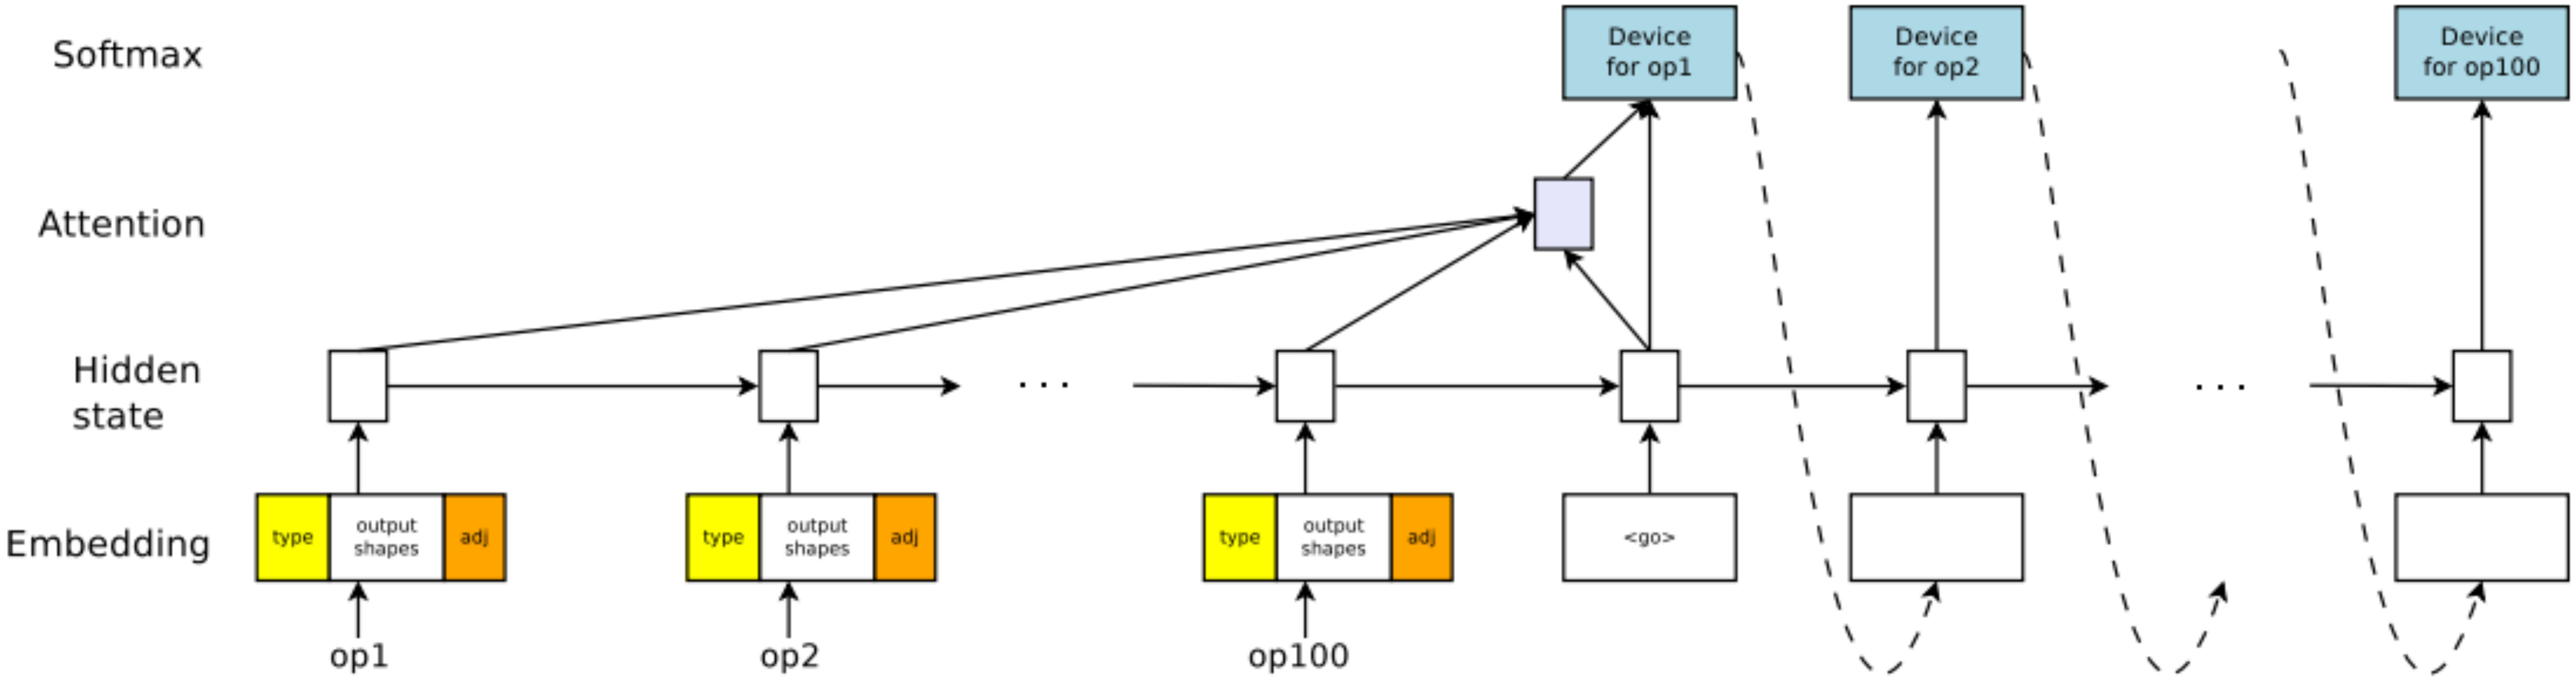

Revised approach (2018)

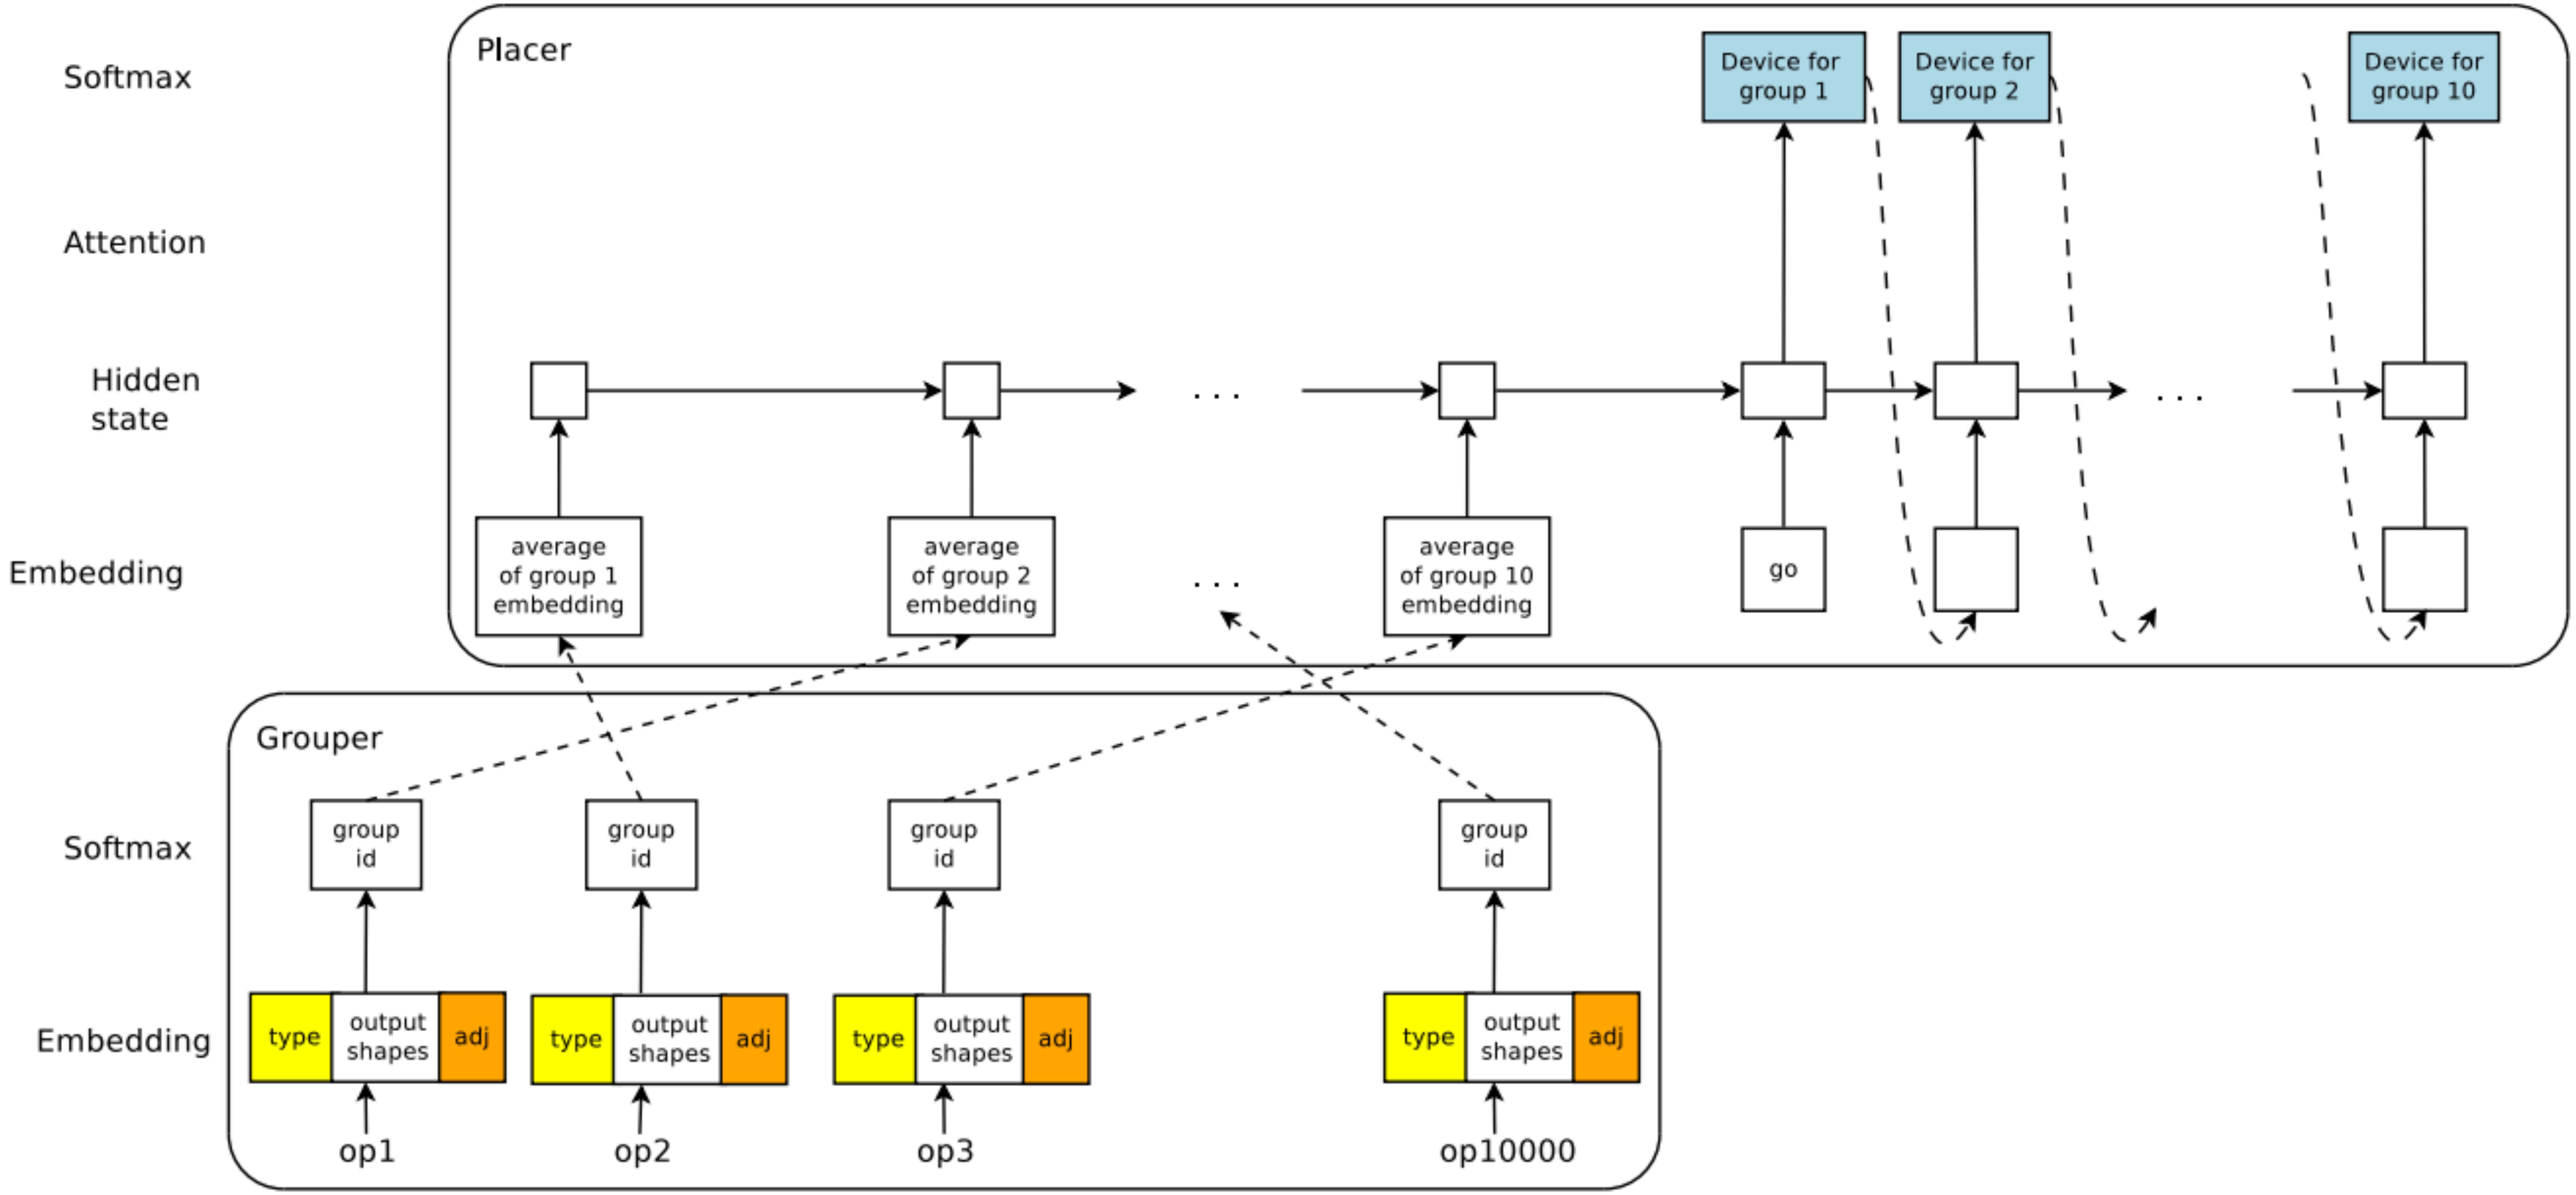

Revised approach (2018)

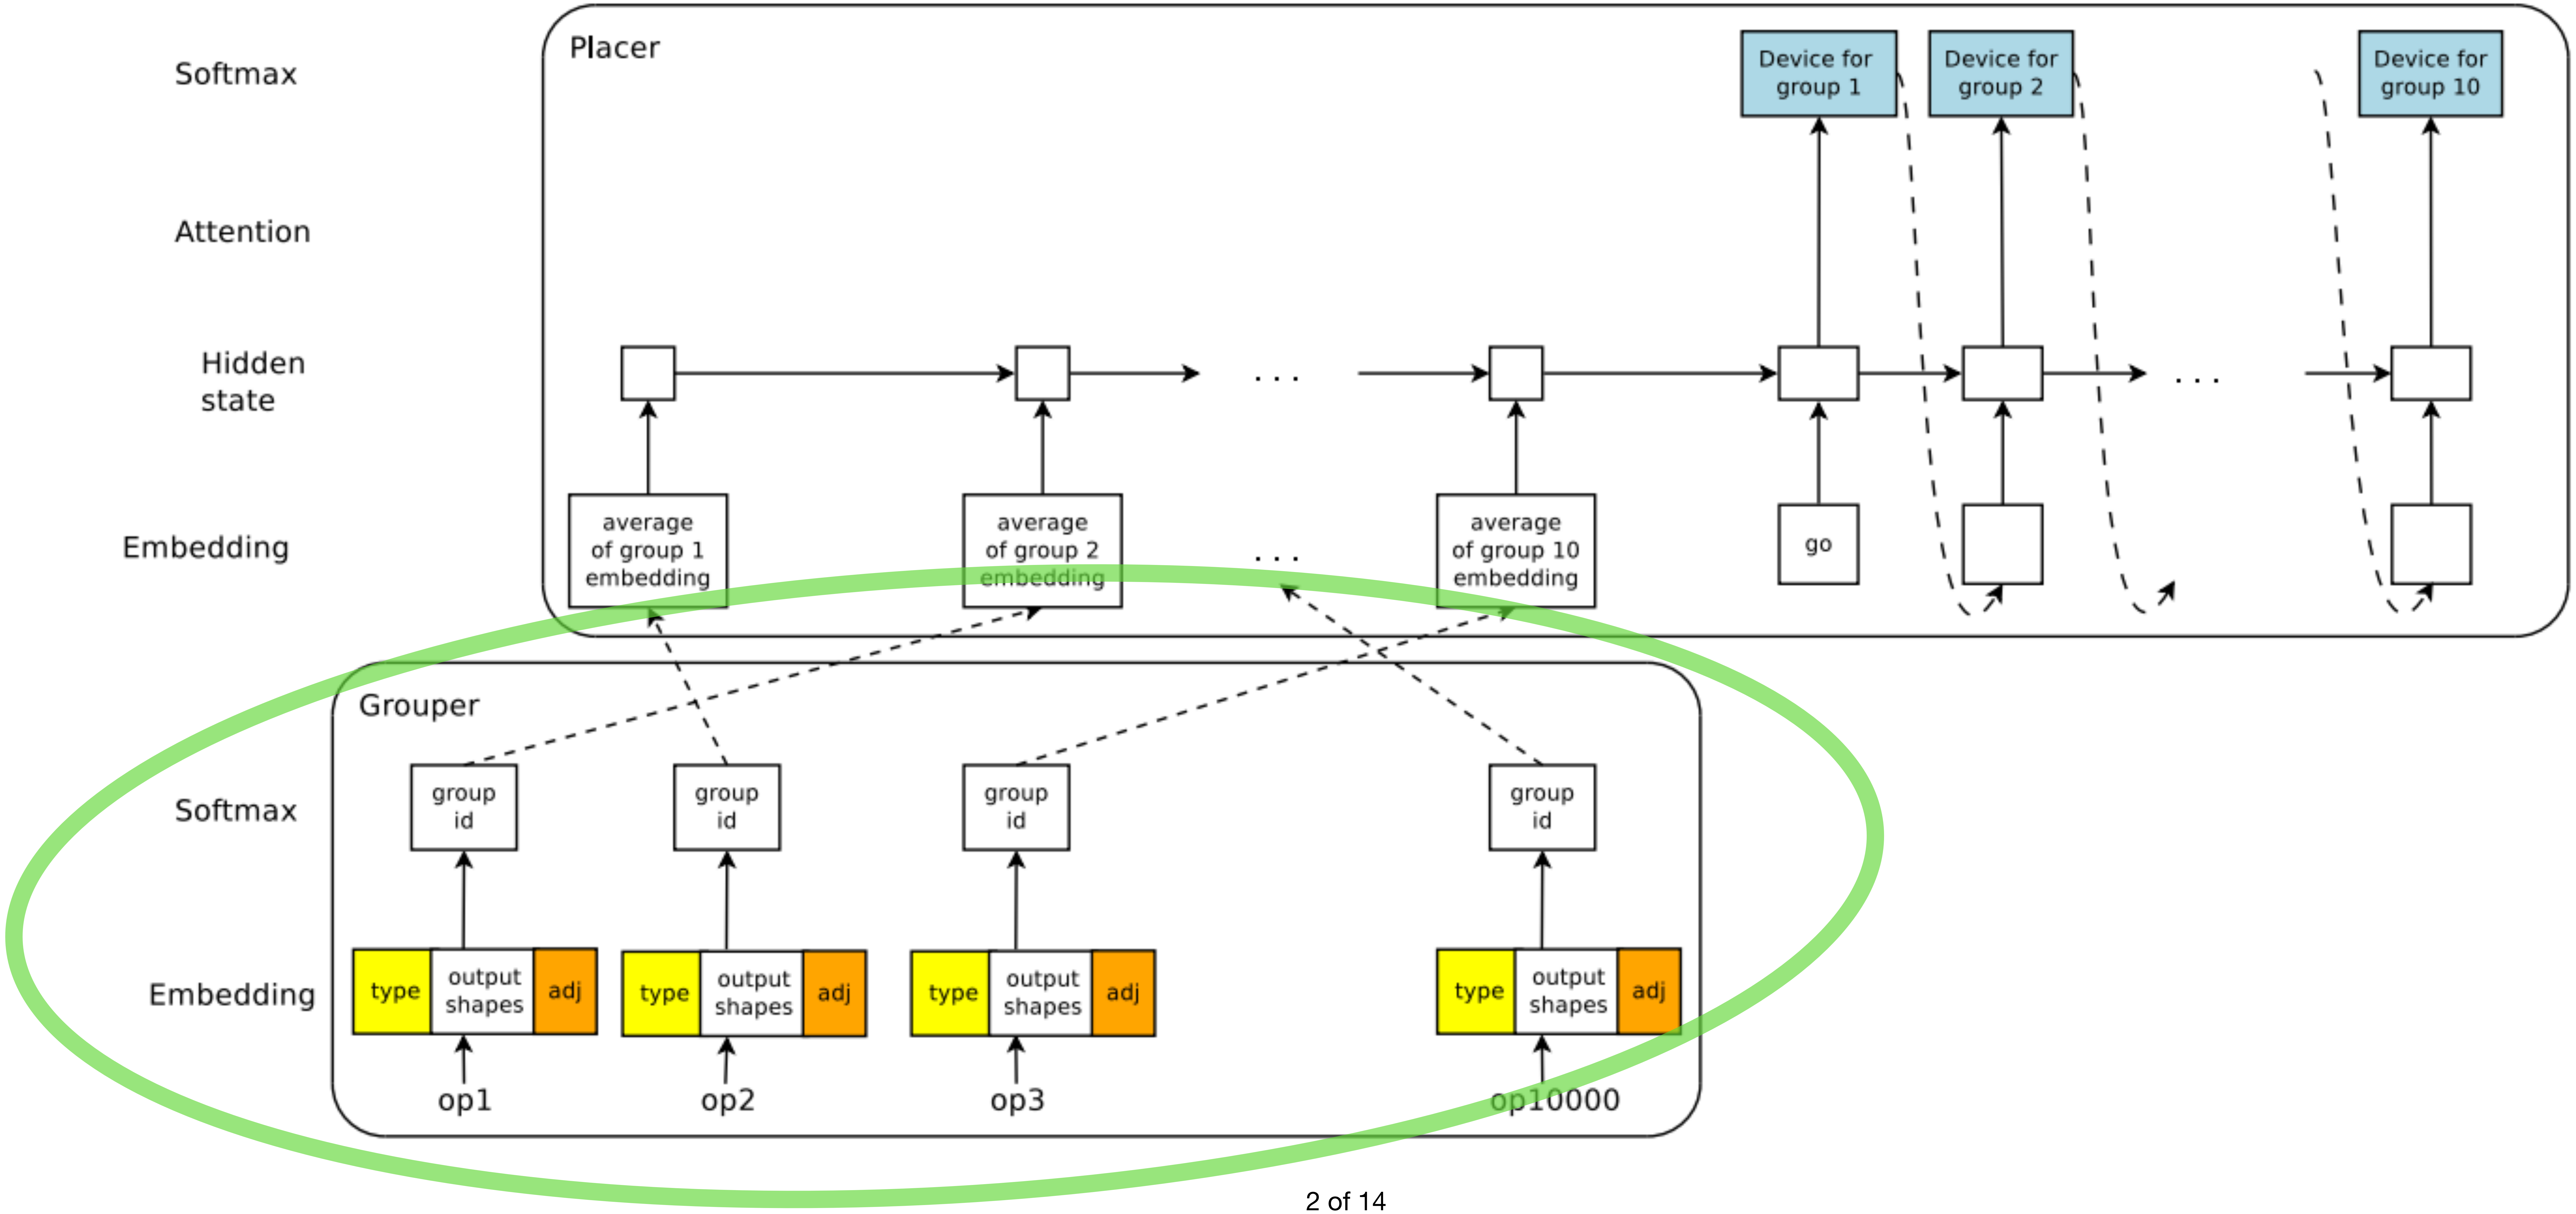

Principal Neighbourhood Aggregation

- Regular message passing aggregators (used in GNNs) cannot distinguish between some sets of neighbourhood messages in certain situations

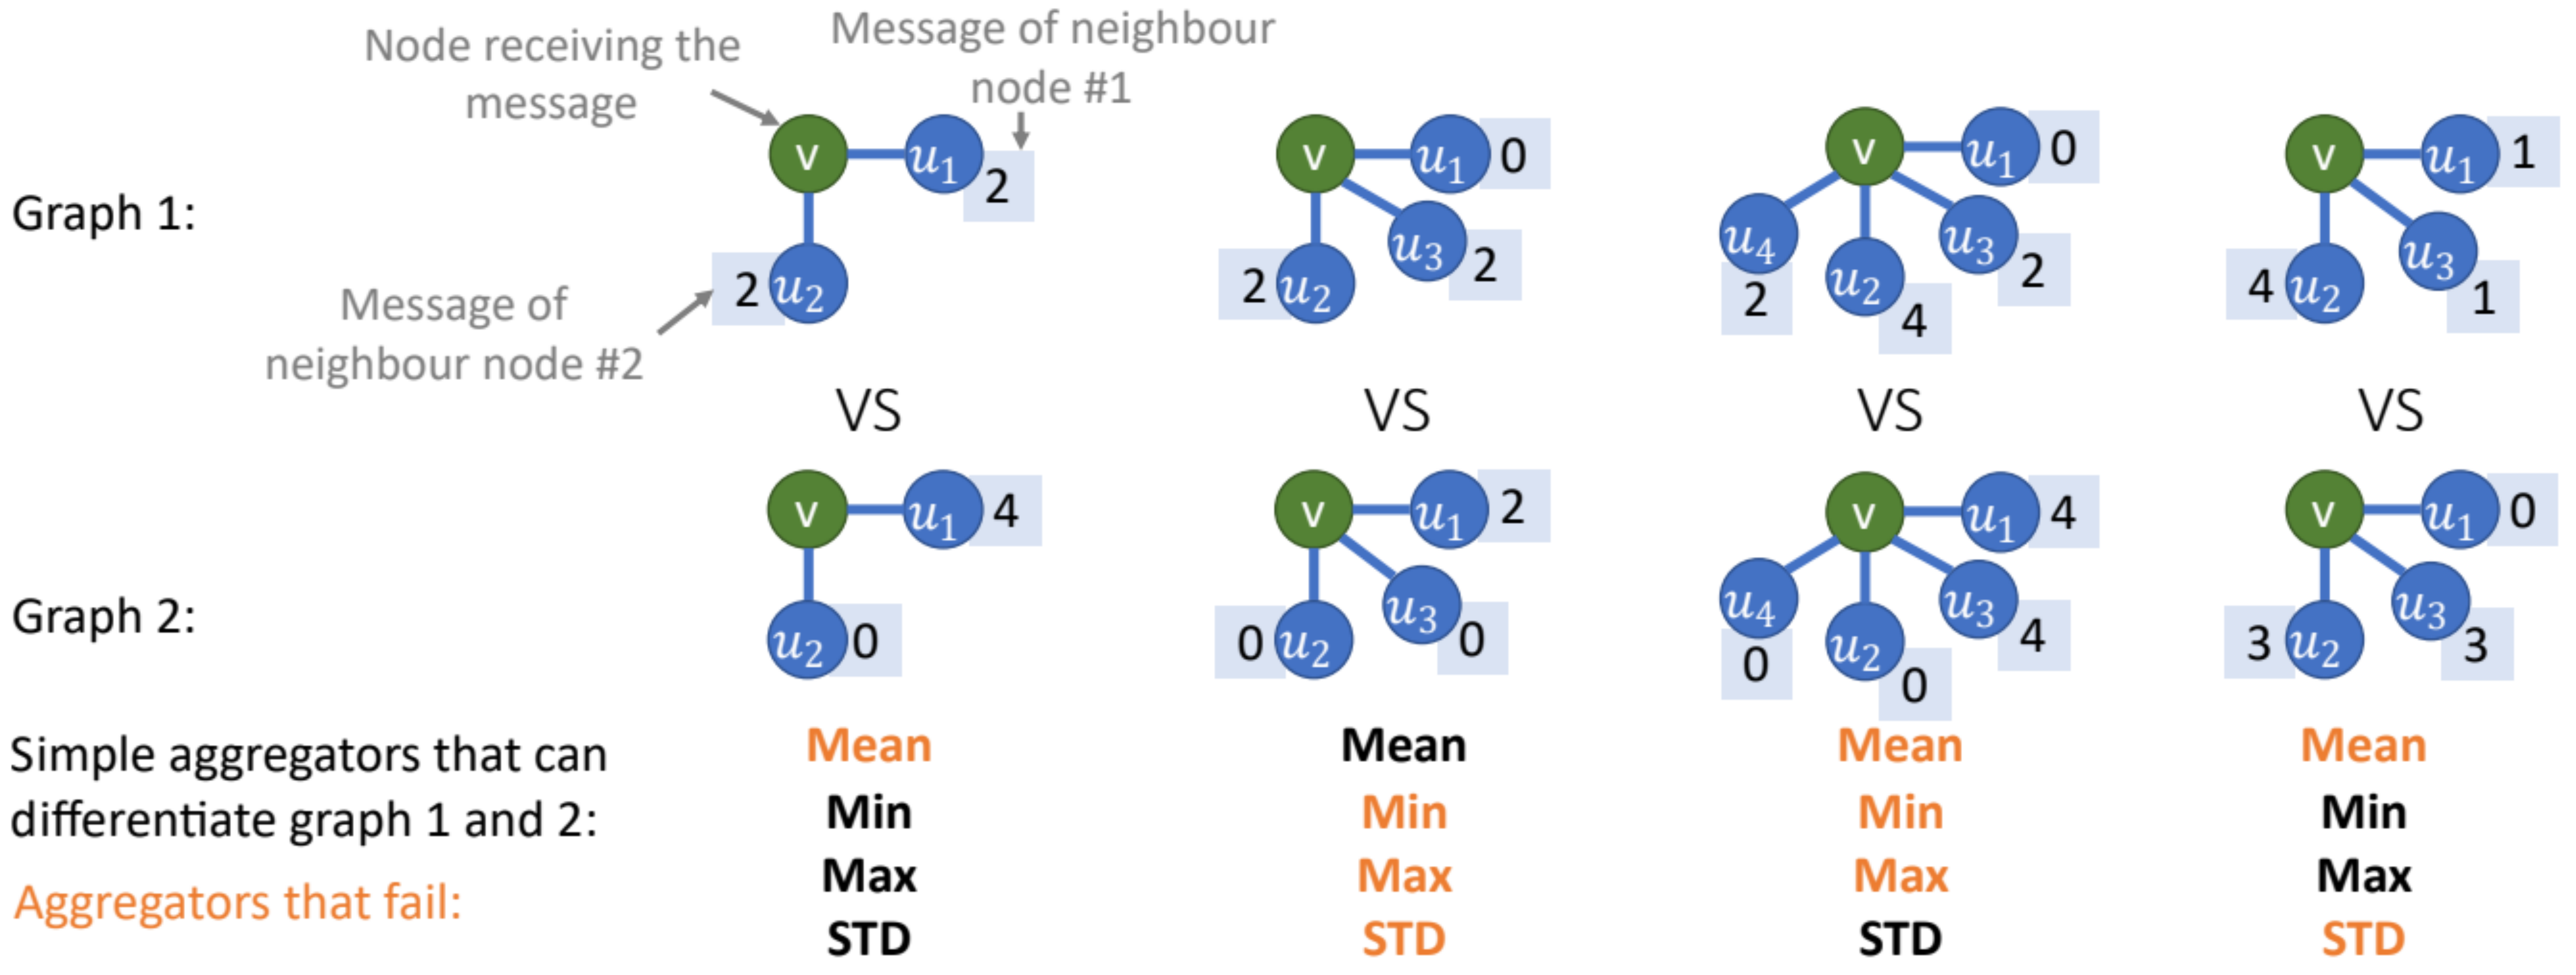

Principal Neighbourhood Aggregation

- PNA combines multiple aggregators (e.g. mean, std) with *degree scalers*
- Degree scalers amplify/attenuate signals from neighbouring neurons according to their degree (thereby incorporating their ‘influence’)
- Intuition: distinguishing two roughly equally sized input tensors from very differently sized ones should be informative for device placement tasks

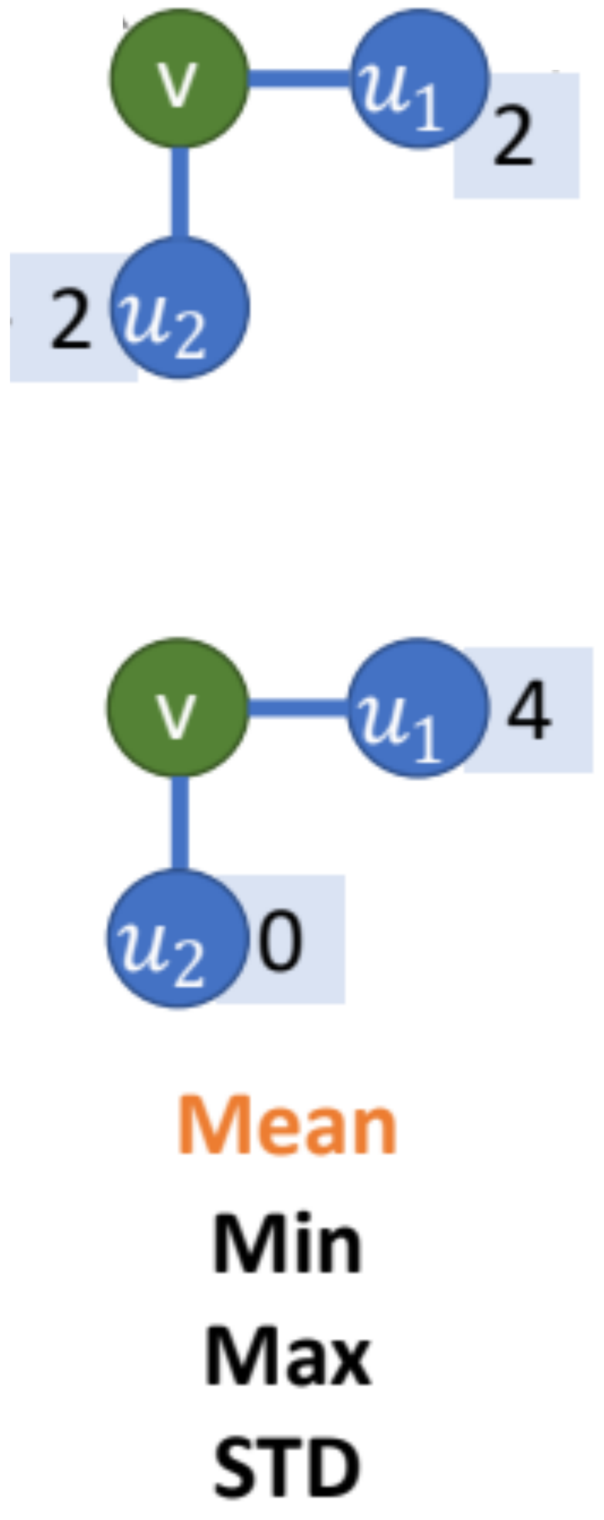

$$X_i^{(t+1)} = U\left(X_i^{(t)}, \bigoplus_{(j,i)\in E} M\left(X_i^{(t)}, E_{j\to i}, X_j^{(t)}\right)\right)$$

$$\bigoplus = \underbrace{\begin{bmatrix} I \\ S(D,\alpha=1) \\ S(D,\alpha=-1) \end{bmatrix}}_{\text{scalers}} \otimes \underbrace{\begin{bmatrix} \mu \\ \sigma \\ \max \\ \min \end{bmatrix}}_{\text{aggregators}}$$

References

- Mirhoseini, Azalia, Hieu Pham, Quoc V. Le, Benoit Steiner, Rasmus Larsen, Yuefeng Zhou, Naveen Kumar, Mohammad Norouzi, Samy Bengio, and Jeff Dean. "Device placement optimization with reinforcement learning." *arXiv preprint arXiv:1706.04972* (2017).

- Mirhoseini, Azalia, Anna Goldie, Hieu Pham, Benoit Steiner, Quoc V. Le, and Jeff Dean. "A hierarchical model for device placement." In *International Conference on Learning Representations*. 2018.

- Mao, Hongzi, Parimarjan Negi, Akshay Narayan, Hanrui Wang, Jiacheng Yang, Haonan Wang, Ryan Marcus et al. "Park: An open platform for learning-augmented computer systems." *Advances in Neural Information Processing Systems* 32 (2019): 2494-2506.

- Corso, Gabriele, Luca Cavalleri, Dominique Beaini, Pietro Liò, and Petar Veličković. "Principal neighbourhood aggregation for graph nets." *arXiv preprint arXiv:2004.05718* (2020).