Supplementary Document: Parallel Droplet Control in MEDA Biochips using Multi-Agent Reinforcement Learning

Tung-Che Liang [1] **Jin Zhou** [1] **Yun-Sheng Chan** [2] **Tsung-Yi Ho** [3] **Krishnendu Chakrabarty** [1] **Chen-Yi Lee** [2]

Abstract

This supplementary document contains a detailed description of MEDA biochips, the CNN architecture for the MARL agents, the full set of MARL training processes, and the MARL evaluation results.

1. Microelectrode-Dot-Array Biochips

The MEDA biochip platform has been proposed in recent years to further advance microfluidics technology (Lai et al., 2015b; Ho et al., 2016). A MEDA biochip is composed of a two-dimensional microelectrode array that manipulates discrete fluid droplets. MEDA biochips manipulate nanoliter droplets using the principle of *electrowetting-on-dielectric* (EWOD) (Pollack et al., 2000). When driven by a sequence of control voltages, the microelectrode array can perform fluidic operations, such as dispensing, mixing, and splitting (Wang et al., 2011; Zhong et al., 2020b). Using MEDA biochips, bioassay protocols are scaled down to droplet size and executed through software-based control of nanoliter droplets. Figure 1(a) shows a fabricated MEDA biochip, where a droplet is present.

MEDA biochips manipulate nanoliter droplets using the principle of *electrowetting-on-dielectric* (EWOD) (Pollack et al., 2000). EWOD refers to the modulation of the interfacial tension between a conductive fluid and a solid electrode coated with a dielectric layer by applying an electric field between them. As shown in Figure 1(b), an imbalance of interfacial tension is created if an electric field is applied to only one side of the droplet; this interfacial tension gradient forces the droplet to move toward to the right.

Compared with traditional EWOD biochips, the most significant advantage of using MEDA biochips is that MEDA architecture supports real-time sensing of droplets (Wang et al., 2011; Ho et al., 2016). A MEDA biochip is composed of an array of identical microelectrode cells (MCs); see Figure 1(c). It has been fabricated at TSMC using a mainstream $0.35\,\mu m$ CMOS process. Each MC consists of a microelectrode, an electronic control circuit, and a sensing module that enables real-time sensing of droplets. An example of a sensing result (a two-dimensional Boolean array) obtained from a fabricated MEDA biochip is shown in Figure 1(d). The sensing results have been used for real-time point-of-care diagnostics such as glucose level test and follicle-stimulating hormone test (Lai et al., 2015a).

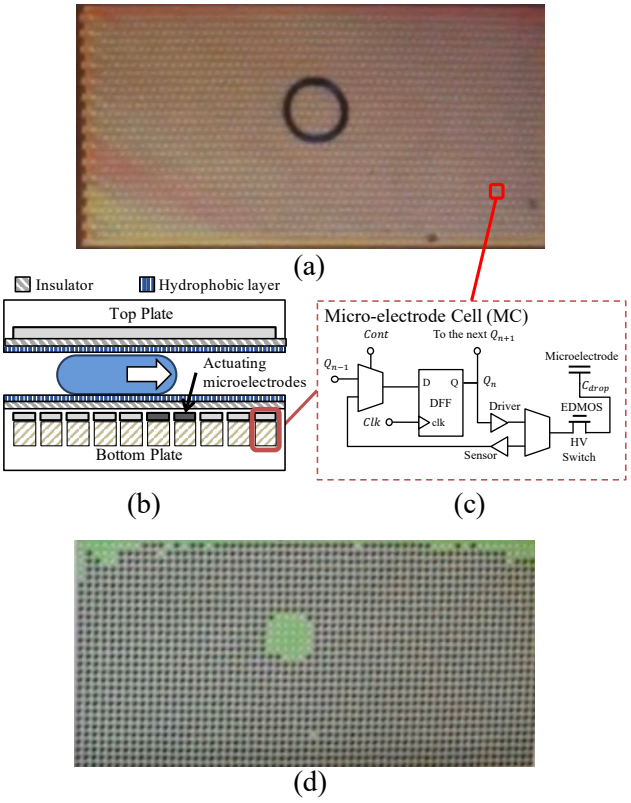

Figure 1. (a) Top-view of a fabricated MEDA biochip. A droplet is present on a 30×60 MEDA biochip. (b) Side-view of the MEDA biochip. (c) Schematic of the microelectrode cell. (d) Real-time sensing result indicates the droplet location.

[1]Department of Electrical and Computer Engineering, Duke University, Durham, NC, USA [2]Department of Electronics Engineering, National Yang Ming Chiao Tung University, Hsinchu, Taiwan [3]Department of Computer Science, National Tsing Hua University, Hsinchu, Taiwan. Correspondence to: T.-C. Liang <tl221@duke.edu>.

Proceedings of the 38^{th} International Conference on Machine Learning, PMLR 139, 2021.

As microfluidic biochips are being used for critical point-of-care diagnostics, reliability in these systems has become an important focus of research (Zhong et al., 2020a; Liang et al., 2020b). It has been reported that the unit cells (i.e., electrodes) of an EWOD biochip degrade over time (Verheijen & Prins, 1999; Chakrabarty, 2009; Drygiannakis et al., 2009). Electrode degradation results from charge trapping in the dielectric insulator (Dong et al., 2015); therefore, a degraded electrode cannot be observed using a CCD camera. Fluidic operations, such as droplet transportation, associated with the degraded microelectrode can fail, thereby compromising the integrity of the bioassay outcome. MEDA biochips, in particular, are more susceptible microelectrode degradation than other EWOD biochips because microelectrodes in MEDA biochips are charged during not only droplet actuation, but also during droplet sensing, i.e., a microelectrode in MEDA biochips is charged more frequently than in other EWOD biochips (Zhong et al., 2020a).

2. Neural Network for the MARL Agent

Various neural network architectures have been proposed over the past few years (Simonyan & Zisserman, 2014; Howard et al., 2017; Gatys et al., 2015). We consider fabricated MEDA biochips as test cases and evaluate the effectiveness of MARL-based adaptation using arrays of size 30×60 and 80×60 (Lai et al., 2015b). These array dimensions correspond to the sizes of fabricated MEDA biochips. Prior work has shown that while fully-connected neural networks are effective for the droplet-routing problem for small biochip instances (less than 100 electrodes), they do not converge for large biochips (Liang et al., 2020a). Therefore, we evaluated several convolutional neural networks (CNNs) (Lerer et al., 2020), and many of them are effective. However, because the network needs to be loaded on an affordable biochip platform, the computational resources on the associated controller may be limited compared to a server. For example, in (Willsey et al., 2019), the cyber-physical biochip system includes only a quad-core 1.2 GHz ARMv7 processor with 1 GB RAM, and it does not contain a GPU; therefore, large networks are not feasible in this application scenario. We tested several options for the number of hidden layers and number of neurons per layer. We found that a simple CNN, as described in Table 1, can solve the droplet-routing problem for large MEDA with more than $4,000$ electrodes.

3. Training Processes

We consider fabricated MEDA biochips as test cases and evaluate the effectiveness of RL-based adaptation using arrays of size 30×60 and 80×60 (Lai et al., 2015b). We evaluated three RL algorithms, i.e. double DQN, PPO, and ACER, using three training schemes, namely centralized, concurrent, and parameter sharing. We compare the MARL approaches with two baseline methods: 1) the single-agent RL framework in (Liang et al., 2020a) and 2) a static (offline) routing method in (Keszocze et al., 2017). We illustrate training processes for the MEDA-Env that contains 30×60 microelectrodes in Figure 2. For each training game of MEDA-Env, n_{rt} random routing tasks are generated, where $1 < n_{rt} \leq 4$. A training epoch contains $20,000$ timesteps. We observe that ACER is the most sample-efficient algorithm in decentralized learning. However, ACER encounters scalability issues in centralized learning, where reward becomes sparse in an exponentially growing action space.

Table 1. The convolutional neural network configuration used in this work.

Layer	Type	Depth	Activation	Stride	Padding
1	Convolution	32	ReLU	3	1
2	Convolution	32	ReLu	3	1
3	Max Pool	N/A	N/A	2	1
4	Convolution	64	ReLU	3	1
5	Convolution	64	ReLu	3	1
6	Max Pool	N/A	N/A	2	1
7	Convolution	128	ReLU	3	1
8	Convolution	128	ReLu	3	1
9	Max Pool	N/A	N/A	2	1
10	Fully-Connected	8	ReLU	N/A	N/A

Because double DQN and the single-agent model do not converge in all the 30×60 MEDA environments, we did not run these methods in 80×60 MEDA environments. Training processes of 80×60 MEDA biochips are illustrated in Figure 3. Unlike the results in 30×60 MEDA biochips, we observe that only ACER converges in decentralized learning. However, only PPO improves its performance in centralized learning. ACER is trapped in local minima in centralized learning.

4. MARL Evaluation

We evaluate the performance of the models in a realistic simulation setting, where microelectrodes degrade over time. The training processes in Section 3 show that models in concurrent schemes are more sample-efficient than the other schemes. Therefore, we used the PPO and ACER models that have been trained to achieve the same performance as that of the baseline (Keszocze et al., 2017). Figure 4 shows the simulation results when microelectrodes in MEDA-Env degrade over time, where $W = 30$ and $L = 60$. After the biochip is used for a while, we see that the performance of the baseline method degrades because the baseline method does not know which microelectrode is degraded and cannot dynamically change the routing paths. On the other hand, the MARL model "learns" the degradation process of the biochip and alter the routing paths accordingly. Therefore, the MARL model outperforms the baseline.

As shown in Section 3, ACER is the only algorithm that

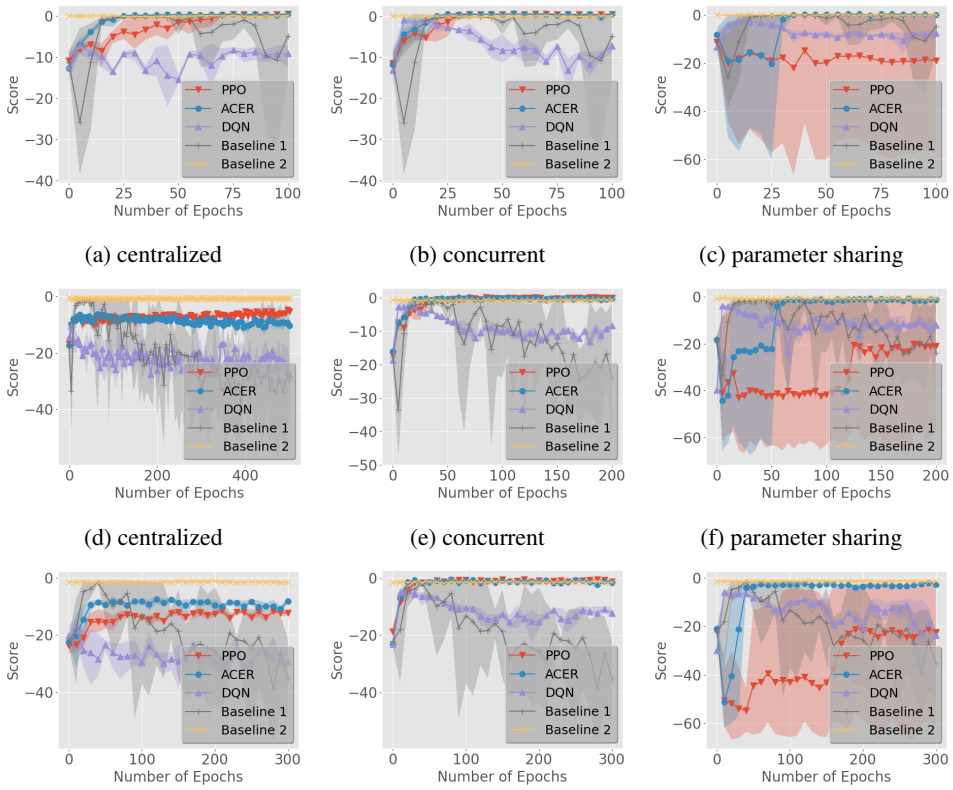

(a) centralized (b) concurrent (c) parameter sharing

(d) centralized (e) concurrent (f) parameter sharing

(g) centralized (h) concurrent (i) parameter sharing

Figure 2. Training process corresponding to different RL algorithms and training schemes in a 30 × 60 MEDA biochip. Score is the total reward that the MARL agents receive in a game. The performance is compared with two baseline methods: 1) a single-agent RL method (Liang et al., 2020a) and 2) a static (offline) routing method in (Keszocze et al., 2017). (a, b, c) Results with at most 2 concurrent routing tasks. (d, e, f) Results with at most 3 concurrent routing tasks. (g, h, i) Results with at most 4 concurrent routing tasks.

converges in concurrent schemes for a 80 × 60 MEDA biochip. Therefore, we used ACER models that have been trained to achieve the same performance as that of the baseline (Keszocze et al., 2017) in degrade mode. Figure 5 illustrates the simulation results when microelectrodes degrade over time in a 80 × 60 MEDA biochip. We also see that the MARL model outperforms the baseline, which is consistent with the performance of MARL models in a 30 × 60 MEDA biochip.

References

Chakrabarty, K. Design automation and test solutions for digital microfluidic biochips. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 57(1):4–17, 2009.

Dong, C., Chen, T., Gao, J., Jia, Y., Mak, P.-I., Vai, M.-I., and Martins, R. P. On the droplet velocity and electrode lifetime of digital microfluidics: Voltage actuation techniques and comparison. *Microfluidics and Nanofluidics*, 18(4):673–683, 2015.

Drygiannakis, A. I., Papathanasiou, A. G., and Boudouvis, A. G. On the connection between dielectric breakdown strength, trapping of charge, and contact angle saturation in electrowetting. *Langmuir*, 25(1):147–152, 2009.

Gatys, L., Ecker, A. S., and Bethge, M. Texture synthesis using convolutional neural networks. In *Advances in Neural Information Processing Systems 28*, pp. 262–270. Curran Associates, Inc., 2015.

Ho, Y., Wang, G., Lai, K. Y.-T., Lu, Y.-W., Liu, K.-M., Wang, Y.-M., and Lee, C.-Y. Design of a micro-electrode cell for programmable lab-on-cmos platform. In *Proceedings of the IEEE International Symposium on Circuits and Systems (ISCAS)*, pp. 2871–2874, 2016.

Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., and Adam, H. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *CoRR*, abs/1704.04861, 2017. URL http://arxiv.org/abs/1704.04861.

Keszocze, O., Li, Z., Grimmer, A., Wille, R., Chakrabarty, K., and Drechsler, R. Exact routing for micro-electrode-dot-array digital microfluidic biochips. In *Proceedings of*

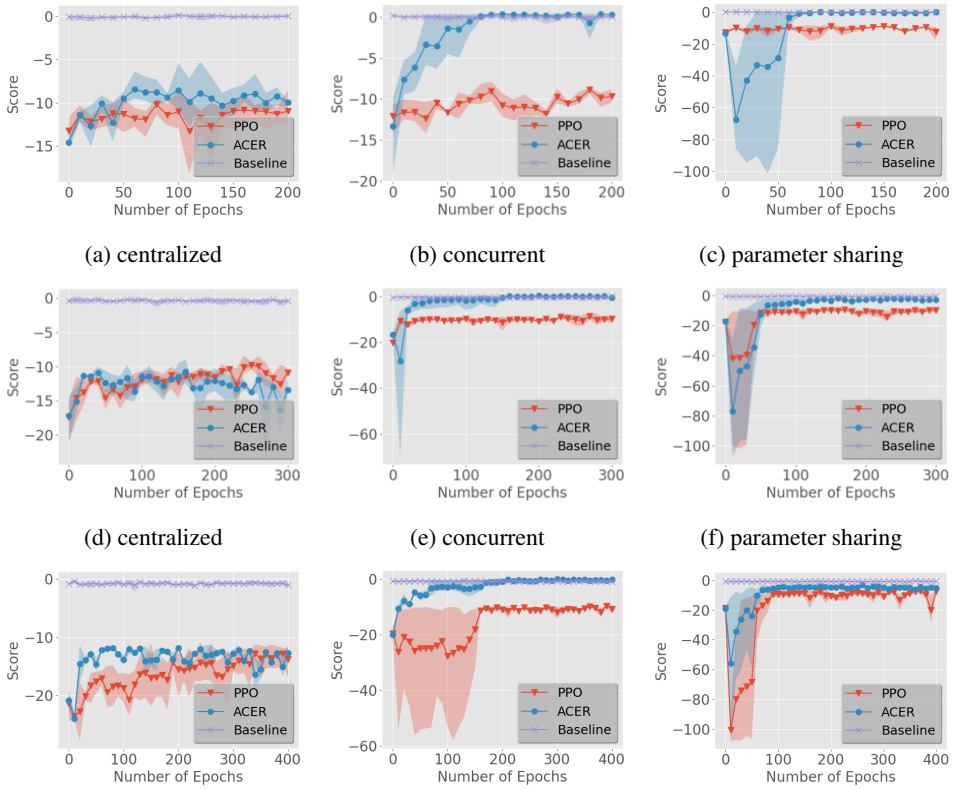

(a) centralized (b) concurrent (c) parameter sharing

(d) centralized (e) concurrent (f) parameter sharing

(g) centralized (h) concurrent (i) parameter sharing

Figure 3. Training process corresponding to different RL algorithms and training schemes in a 80×60 MEDA biochip. Score is the total reward that the MARL agents receive in a game. The performance is compared with a static (offline) routing method in (Keszocze et al., 2017). (a, b, c) Results with at most 2 concurrent routing tasks. (d, e, f) Results with at most 3 concurrent routing tasks. (g, h, i) Results with at most 4 concurrent routing tasks.

the Asia and South Pacific Design Automation Conference (ASP-DAC), pp. 708–713. IEEE, 2017.

Lai, K. Y.-T., Shiu, M.-F., Lu, Y.-W., Ho, Y., Kao, Y.-C., Yang, Y.-T., Wang, G., Liu, K.-M., Chang, H.-C., and Lee, C.-Y. A field-programmable lab-on-a-chip with built-in self-test circuit and low-power sensor-fusion solution in 0.35 μm standard cmos process. In *Proceedings of the IEEE Asian Solid-State Circuits Conference*, pp. 1–4. IEEE, 2015a.

Lai, K. Y.-T. et al. An intelligent digital microfluidic processor for biomedical detection. *Journal of Signal Processing Systems*, 78(1):85–93, 2015b.

Lerer, A., Hu, H., Foerster, J. N., and Brown, N. Improving policies via search in cooperative partially observable games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 7187–7194, 2020.

Liang, T.-C., Zhong, Z., Bigdeli, Y., Ho, T.-Y., Chakrabarty, K., and Fair, R. Adaptive droplet routing in digital microfluidic biochips using deep reinforcement learning. In *Proceedings of International Conference on Machine Learning*, 2020a.

Liang, T.-C., Zhong, Z., Pajic, M., and Chakrabarty, K. Extending the lifetime of meda biochips by selective sensing on microelectrodes. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 39(11): 3531–3543, 2020b.

Pollack, M. G., Fair, R. B., and Shenderov, A. D. Electrowetting-based actuation of liquid droplets for microfluidic applications. *Applied Physics Letters*, 77(11): 1725–1726, 2000.

Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

Verheijen, H. and Prins, M. Reversible electrowetting and trapping of charge: Model and experiments. *Langmuir*, 15(20):6616–6620, 1999.

Wang, G., Teng, D., and Fan, S.-K. Digital microfluidic

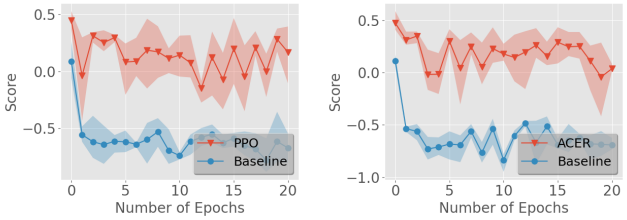

(a) PPO

(b) ACER

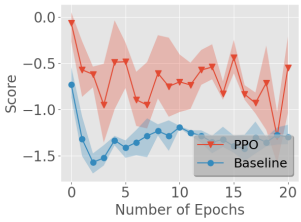

(c) PPO

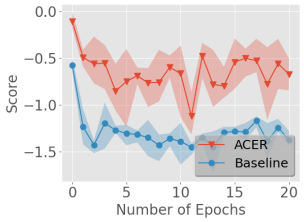

(d) ACER

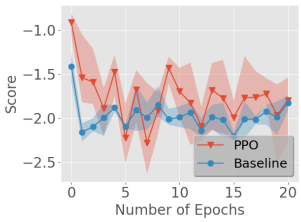

(e) PPO

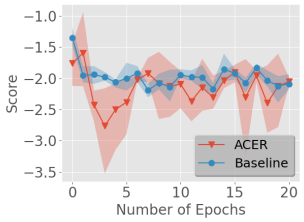

(f) ACER

Figure 4. Comparison between the MARL agents and the baseline method in degrade mode of a 30 × 60 MEDA. (a, b) Results with at most 2 routing tasks in a game. (c, d) Results with at most 3 routing tasks in a game. (e, f) Results with at most 4 routing tasks in a game.

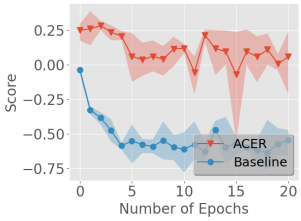

(a)

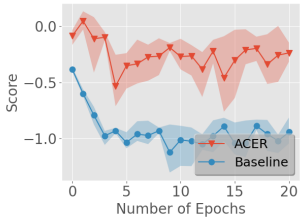

(b)

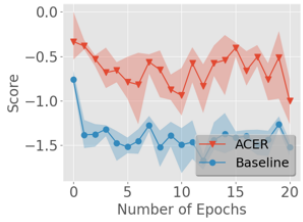

(c)

Figure 5. Comparison between the MARL agents and the baseline method in degrade mode of a 80 × 60 MEDA. (a) Results with at most 2 routing tasks in a game. (b) Results with at most 3 routing tasks in a game. (c) Results with at most 4 routing tasks in a game.

operations on micro-electrode dot array architecture. *IET Nanobiotechnology*, 5(4):152–160, 2011.

Willsey, M., Stephenson, A. P., Takahashi, C., Vaid, P., Nguyen, B. H., Piszczek, M., Betts, C., Newman, S., Joshi, S., Strauss, K., et al. Puddle: A dynamic, error-correcting, full-stack microfluidics platform. In *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pp. 183–197, 2019.

Zhong, Z., Liang, T.-C., and Chakrabarty, K. Enhancing the reliability of meda biochips using ijtag and wear leveling. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2020a.

Zhong, Z., Liang, T.-C., and Chakrabarty, K. Reliability-oriented ieee std. 1687 network design and block-aware high-level synthesis for meda biochips. In *Proceedings of the Asia and South Pacific Design Automation Conference (ASP-DAC)*, pp. 544–549. IEEE, 2020b.