Towards Secure Checkpointing for Micro-Electrode-Dot-Array Biochips*

Mohammed Shayan, *Student Member, IEEE*, Tung-Che Liang, , *Student Member, IEEE*, Sukanta Bhattacharjee, *Member, IEEE*, Krishnendu Chakrabarty, *Fellow, IEEE*, and Ramesh Karri *Fellow, IEEE*

***Abstract*—Biochemical experiments such as diagnostics must be precise and trusted, and provide quick time to results. This has been enabled by automated digital microfluidics; however, it also exposes these experiments to security threats. Previous work has shown that the critical challenge in securing digital microfluidic devices is the lack of sensing resources. The micro-electrode-dot-array (MEDA) is a next-generation digital microfluidic biochip platform that supports fine-grained control and real-time sensing of droplet movements. These capabilities permit continuous monitoring and checkpoint-based validation of assay execution on MEDA. This paper presents a class of 'shadow attacks' that abuse the timing slack in the assay execution. State-of-the-art checkpoint-based validation techniques cannot expose the shadow operations. We overcome this limitation by introducing extra checkpoints in the assay execution at time instances when the assay is prone to shadow attacks. We achieve this by identifying the conditions that enable shadow attacks. We use these conditions to minimize the number of checkpoints required to guarantee the correctness of bioassay implementation. Our simulation results confirm the effectiveness and practicality of the defense.**

I. Introduction

A digital microfluidic biochip (DMFB) is a two-dimensional electrode array used to manipulate discrete fluid droplets. When driven by a sequence of control voltages, the electrode array can implement fluid operations such as dispensing, mixing, and splitting that can be used to build complex protocols such as immunoassays [2], [3], and cell-based assays [4], [5]. DMFBs are revolutionizing point-of-care diagnosis as evident by the commercialization of the first United States Food and Drug Administration approved Baebies SEEKER DMFB platform [6]. Baebies SEEKER provides a high throughput laboratory solution for screening diseases in a newborn child.

Security and trustworthiness of DMFBs are important as these are used in safety-critical applications like point-of-care diagnosis and environmental monitoring of chemical, biological and nuclear weapons [7], [8]. DMFBs are susceptible to attacks such as actuation tampering and mis-calibration leading to disastrous assay outcomes [9], [10], [11]. Cyber-physical integration of DMFBs enables online monitoring of assay execution [12], [13], [14]. *Checkpointing* is a technique to validate in real-time the location of droplets against a golden droplet map. Checkpointing requires CCD-based image capture and analysis and real-time optical detection of droplets. Checkpoint-based validation is limited by either the image processing capabilities or the number of optical detectors [12].

Recently the Micro-electrode-dot-array (MEDA) DMFB has been developed. MEDA has a "sea-of-electrodes" (micro-electrodes) that can be dynamically grouped to act as an actuator for droplet movement [15], [16]. Each micro-electrode is integrated with activation circuitry and sensing modules which allow fine-grained control and real-time sensing of a droplet [17]. The sensor data specifies the droplet location, size, and shape—the *droplet map*. The real-time droplet map can be compared with the golden droplet map for assay validation. From a security perspective, MEDA is promising as it overcomes the resource constraints of a traditional DMFB.

The fine-grained control and sensing in MEDA cuts both ways for MEDA security. It aids seamless monitoring of the entire biochip. And it also aids an attacker in launching stealthier attacks that were not feasible on a traditional DMFB [18]. The MEDA biochips support movement of droplets with multiple size and shape, which have different speeds. It also supports extraction of smaller droplet from larger droplet, i.e., aliquot operation. An attacker can exploit these properties to launch stealthy attacks. Previous work has shown how aliquot droplets can be used to fine-grained manipulation of glucose assay outcomes [18]. In this paper, we focus on how differential speeds in smaller (faster) and larger (slower) droplets can be used to launch stealthy manipulations called shadow attack. In the next subsection, we describe one such shadow attack on MEDA biochip.

A. Motivation

MEDA has enhanced droplet manipulation capabilities like the ability to operate on different droplet sizes and shapes and allows diagonal movement of droplets. Droplets are moved on a MEDA biochip by activating a set of micro-electrodes, and the speed of movement of a droplet depends on its size and shape [17]. The small droplets move multiple steps in the time it takes for the larger droplet to move a single step. We call this differential the "time slack" in an actuation cycle. The actuation cycle consists of scanning in the micro-electrode actuation pattern, actuation of the micro-electrodes,

*This research is supported in part by the Army Research Office under grant number W911NF-17-1-0320, NSF Award numbers CNS-1833622 and CNS-1833624, NYU Center for Cyber Security (CCS), and NYU Abu Dhabi Center for Cyber Security CCS-AD. A preliminary version of this paper has appeared in the proceedings of ICCAD 2018 [1].

M. Shayan, and R. Karri are with the Department of Electrical and Computer Engineering, New York University, Brooklyn, NY, 11201 USA e-mail: (mos283@nyu.edu,rkarri@nyu.edu).

S Bhattacharjee is with the Indian Institute of Technology Guwahati, India e-mail: (sukantab@iitg.ac.in)

K. Chakrabarty and T-C. Liang are with the Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708 USA (e-mail: krish,tung.che.liang@duke.edu).

Manuscript received August 25, 2019.

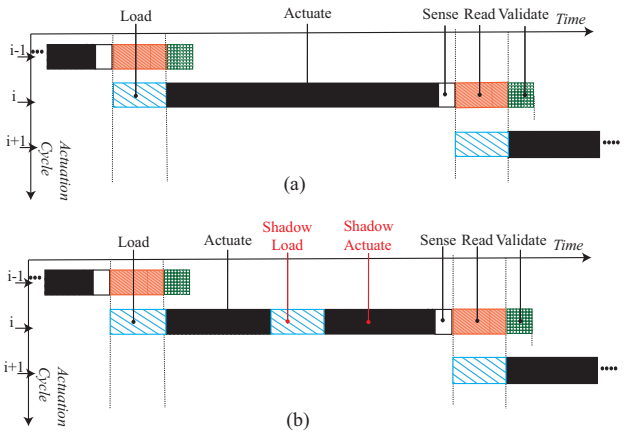

Fig. 1: (a) The baseline MEDA actuation cycle. (b) The modified MEDA actuation cycle with a "shadow operation" embedded in the time slack (in red).

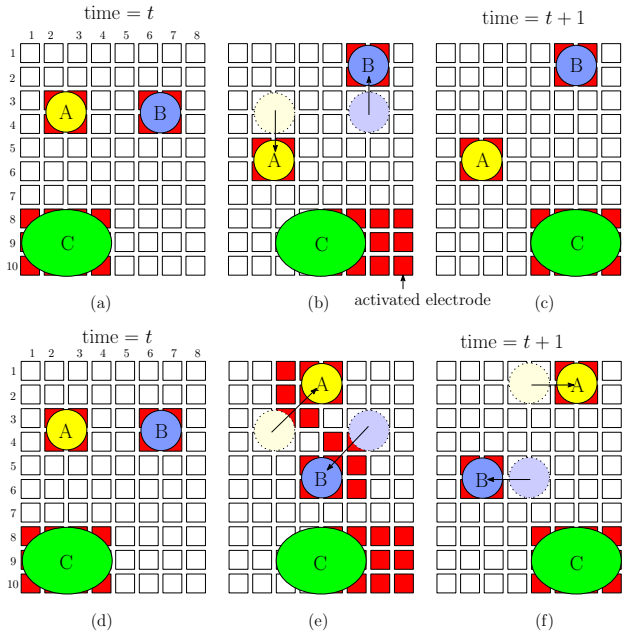

Fig. 2: Droplet transport on a MEDA biochip: (a) the initial state, (b) the intermediate state, and (c) the final state of an actuation cycle. A shadow operation during the droplet transport: (d) the initial state, (e) the intermediate state, and (f) the shadow state at the end of an actuation cycle.

and scanning out the sensed data as shown in Fig. 1(a). At the end of an actuation cycle, a droplet map can be created from the scanned out sensor data. This can be checked against the golden droplet map to detect malicious operations. An attacker can exploit the time slack in an actuation cycle to load multiple actuation sequences depending on the amount of slack (ref. Fig. 1(b)). Using the extra actuation sequences, an attacker can manipulate the droplets (e.g., interchange the location of two droplets) while preserving the golden droplet map at the end of the actuation cycle. We call these the "shadow" attacks.

Example 1. *Consider a* 10×8 *MEDA biochip in which three droplets (a large* 3×4 *size droplet and two smaller* 2×2 *size droplets) are moved in an actuation cycle. The MEDA biochip has an electrode pitch size of* $50\,\mu\text{m}$ *and the spacing between the plates is* $50\,\mu\text{m}$*. The three droplets can be the same reagent of varying concentrations in a sample preparation protocol, which have similar viscosity and interfacial tension. Fig. 2(a, c) show the initial and final states of the golden actuation cycle. The larger droplet moves slowly relative to the smaller ones because of the greater resistance and viscous drag that it runs into. Let the smaller droplet* (2×2) *move at an average speed of* $1.3\,\text{mm}/\,\text{sec}$ *and the larger droplet has an average speed of* $1\,\text{mm}/\,\text{sec}$*. Fig. 2(b) presents the transitional state in which the two smaller droplets reach their destinations. Without loss of generality, we estimate that the smaller droplets move two times faster than the larger droplet. Hence, the smaller droplets travel two steps in the time it takes for the larger one to finish one step. This sets up a timing slack for the smaller droplets in the actuation cycle.*

An attacker can manipulate the timing slack to carry out shadow transport operations on the smaller droplets. The attacker can uphold the golden droplet map at the conclusion of the actuation cycle. Figs. 2(e)-(f) show one possible malicious droplet movement that interchanges the destinations of the two smaller droplets while keeping the droplet-map at the end of the actuation cycle (Fig. 2(c, f) have identical maps).

A straightforward fix to the shadow attack shown in Example 1 is to cut down the time slack to zero. This entails tightening the actuation cycle time (t_{cycle}) to equal the time taken by the smaller droplet to advance one step (t^{min}_{cycle}). The larger droplet takes extra cycle(s) to reach its destination. At the end of each cycle, the droplet map is validated. Reducing the actuation cycle time builds up the number of cycles, and there is an analogous rise in the number of checkpoints and hence the space to store the golden droplet maps. The largest droplet that was manipulated on the MEDA is 30×30 [19]. If the actuation cycle time is t^{min}_{cycle}, this can lead to $30\times$ more actuation cycles and $30\times$ increase in the storage.

B. Contributions

We investigate new class of shadow attacks in the MEDA biochips. We develop conditions that may lead to the shadow attacks. We employ these conditions to determine the actuation cycles in which the shadow attacks cannot be introduced. We use these conditions to prune the number of checkpoints required to monitor the assay implementation. The key contributions are:

- We identify a new class of shadow attacks on MEDA biochips that exploit the timing slack in the droplet actuation cycles to launch the shadow operations.
- We retrospectively re-classify the attack space broadly into spatial and temporal attacks, and each class is divided into sub-classes of attacks.
- We derive the necessary conditions to launch different classes of shadow attacks.
- We present a defense that inserts extra-checkpoints in the time instances that are susceptible to shadow attacks.
- Further, we optimize the defense by pruning the checkpoints where attacks cannot be introduced.

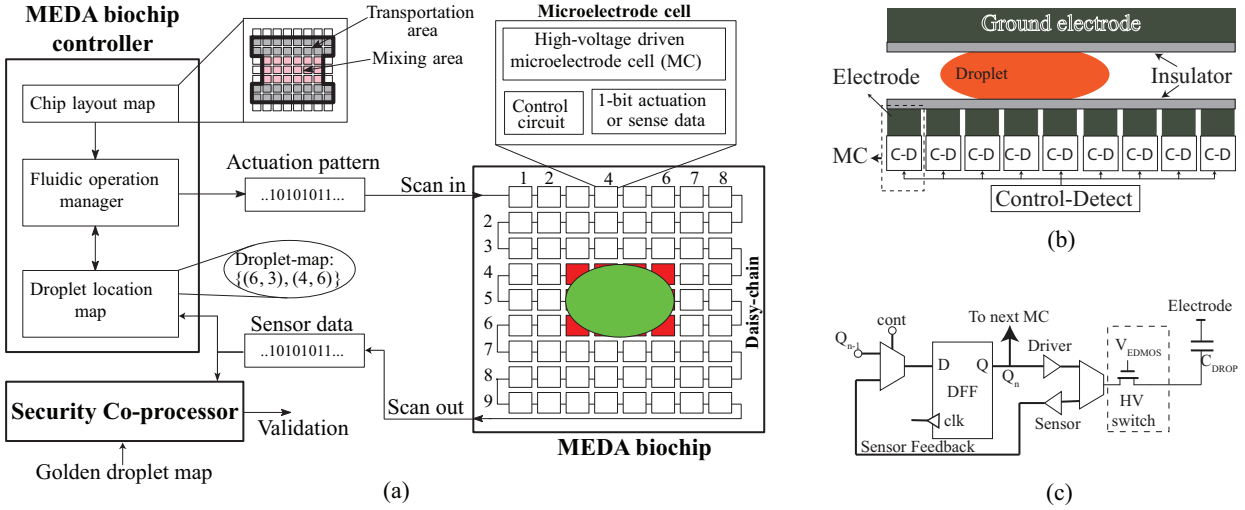

Fig. 3: (a) MEDA cyberphysical system, which includes the MEDA biochip, the controller, and the security co-processor. (b) Side view of the MEDA biochip. (c) A circuit schematic of the sensor and control module comprising a micro-electrode cell (MC). Q_n denotes the n^{th} cell in the scan-chain, receiving an input from the $(n-1)^{th}$ cell denoted by Q_{n-1}.

- We simulate the practicality of the attack analysis and defense on a case study of sample preparation assay and eight real-life benchmarks.

C. Roadmap of the Paper

In Section II, we give a sketch of the MEDA cyberphysical system. In Section III, we ouline the threat model and attack space for MEDA biochip. In Section IV, we define various shadow attacks and study their primary and side effects. In Section V, we describe our defense against shadow attacks and pruning of checkpoint list. In Section VI, we demonstrate the attacks and defense on a real-life bioassays and discuss trade-offs. Section VII concludes the paper.

II. Background

Traditional DMFBs suffer from several disadvantages that limit their scalability and reconfigurability: 1) droplet size is constrained by the electrode size, 2) droplet volume control is limited, and 3) sensors must be integrated post-fabrication [20]. In this section, we describe how MEDA overcomes these drawbacks. In other words, we describe the essential features of MEDA biochip. In the rest of the paper, we refer to "traditional DMFBs" as "DMFBs."

A. MEDA Biochip Architecture

A MEDA biochip platform consists of the following components: 1) a two-dimensional array of identical microelectrode cells (MCs) and 2) a biochip controller which consists of a chip layout map, droplet location map, and the fluidic operation manager, as shown in Fig. 3(a). Each MC includes a high voltage driven microelectrode, an actuation circuit, and a sensing circuit for real-time sensing of the droplet under the microelectrode, as shown in Fig. 3(c). Depending on the application, MCs are grouped to form a virtual chip layout that contains reservoirs, mixers, and fluidic paths. The chip layout map stores this configuration data. The droplet location map stores the real-time locations of the droplets provided by the sensors in the MCs. The fluidic operation manager converts a fluidic instruction to an MC actuation pattern which is then shifted into the MCs. The MCs loaded with logic '1' are actuated by connecting high-voltage. After the actuation, the sensor is enabled, and the data is shifted out, creating the droplet location map, as shown in Fig. 1(a). The load, execute, and sense steps form an actuation cycle.

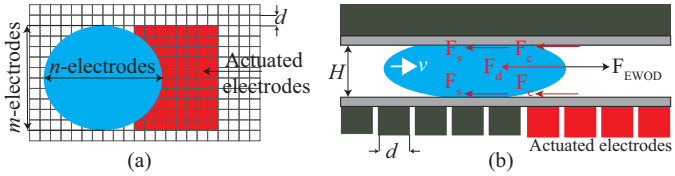

Fig. 4: Droplet transportation (a) top view and (b) side view.

B. Droplet Velocity on MEDA

The MEDA biochip manipulates fluids in discrete quantities based on the electrowetting-on-dielectric (EWOD) principle that the contact angle between a droplet and substrate can be controlled through the application of a suitable electric potential [21], [15], [16].

The motion of a $m \times n$ sized droplet on a MEDA biochip with electrode size d and plate spacing H is a result of the following forces:

1) When an electric potential is applied between a droplet and an electrode, the droplet is subjected to EWOD (F_{EWOD}) force due to the modification of interfacial tension. This force is proportional to the unit capacitance of the electrode (C_{unit}), the square of the applied voltage (V), and the effective droplet contact length (L_{eff}). Note that $F_{EWOD} = 0.5 C_{unit} V^2 L_{eff}$ [20].
2) In order to move, a droplet movement needs to overcome contact-line pinning forces that are caused by contact

TABLE I: MEDA biochip hardware parameters [17]

Parameter	Value
Unit capacitance (C_{unit})	35.6×10^{-6} F
Filler fluid density (ρ_f)	760 kg/m^3
Droplet viscosity (μ_d)	1.9 $\times 10^{-3}$ Pa.s
Drag coefficient C_D	30
Micro-electrode (cell) pitch, d	50 $\times 10^{-6}$ m
Plate spacing, H	50 $\times 10^{-6}$ m
Proportionality constant ζ	40 N.s/m^2

angle hysteresis. This can be modeled by a force (F_c) proportional to the droplet velocity (v), and its perimeter ($2(m+n)d$). It has been shown in the literature that $F_c = 2\zeta(m+n)dv$, where ζ is the proportionality constant [22]. Please note that the authors in [22] intuitively assumed that the contact-line pinning force varies linearly with velocity; subsequently, this was experimentally validated.

3) A droplet is subjected to a viscous drag force (F_s) between the droplet and the top and bottom plates. This force is directly proportional to the droplet viscosity (μ_d), its velocity (v), and its area (mnd^2); it is also inversely proportional to the gap (H) between the top and bottom plates. It has been shown that $F_s = 12\mu_d vmnd^2/H$ [23].
4) Further, the droplet movement is resisted by the drag from the filler medium. This force (F_d) is proportional to the square of droplet velocity (v), the projected droplet area in the direction of the velocity ($L_{eff} \cdot H$), and the filler fluid density (ρ_f). We know from the literature that $F_d = 0.5C_D\rho_f v^2 L_{eff}H$, where C_D is the drag coefficient [23].

Therefore, the net force (F_{net}) on the droplet is given by:

$$F_{net} = F_{EWOD} - F_c - F_d - F_s \quad (1)$$

The droplet is initially stationary. When the neighboring electrodes are actuated, the droplet is transported to the right, as shown in Fig 4. After the droplet is completely moved to the neighbouring electrodes, it is stationary again [23]. The random pinning forces causes contact angle hysteresis that resists a droplet movement. The actuated voltage needs to be more than a threshold to move a droplet [23]. Previous work on droplet dynamics assumes that droplet moves on a biochip with an average (fixed) velocity. Experimental data agrees with such an approximation and it helps to simplify the analysis [20]. The average velocity is derived by solving the equation for net force (Equation 1). Consider a droplet of size $m \times n$ moving in x-direction, as shown in Fig. 4, i.e., effective length $L_{eff} = md$. The average velocity of the droplet is given by the positive root of the following quadratic equation [20]:

$$\left(0.5C_D\rho_f mdH\right) \cdot v^2 + \left(12\mu_d mnd^2/H + \zeta 2(m+n)d\right) \cdot v - \left(0.5C_{unit}V^2 md\right) = 0 \quad (2)$$

Example 2. *Consider a MEDA biochip with the parameters shown in Table I, as reported in [20]. On this MEDA biochip, a droplet of size* 5×5 *has a velocity of* $1.1\,\mathrm{mm/s}$*. On the other hand, a larger droplet of size* 16×16 *size has an average velocity of* $0.6\,\mathrm{mm/s}$*. The analytical solutions for the droplet velocities have been experimentally verified [20].*

C. MEDA versus DMFB

MEDA biochips have a large number of microelectrodes which are 10 to 20 times smaller than DMFB electrodes [25]. MEDA biochip supports precise and flexible control of the droplet size and shape, whereas DMFB support only droplets of single size and shape. MEDA biochips also support diagonal movement apart from the vertical and horizontal movement, whereas DMFB only supports the latter [20]. MEDA also introduces specific fluidic operations such as *aliquot operation*. Finally, MEDA enables fine-grained sensing in contrast to absence of integrated sensing in DMFB [26].

D. Online Monitoring

The assay execution is validated using sensor data at the end of each actuation cycle. The real-time droplet location map is compared against the golden map to validate the assay execution. The time duration between successive checkpoints (CPs) is constrained by the actuation cycle time. To ensure that the biochip controller and the validation are not simultaneously compromised, the golden droplet map resides in a secure co-processor that is physically separated from the MEDA controller, as shown in Fig. 3(a) [12].

III. MEDA Security

In the previous section, we described the MEDA-specific features. In this section, we focus on the security implications of these features. We describe the threat model and attack space for the MEDA biochip.

A. Threat Model

The responses to the ensuing four questions clarify the threat model for biochip security:

1) Who presents a risk and why?: The attacker, – who is in a remote location or near the biochip – could be a competitor seeking to bring disrepute to the biochip designer [27]. The proximity attacker can be an insider seeking to harm the end-user by manipulating the biochip results [28] or by denying service (DoS attack).

2) What are the attacker capabilities?: The attacker is able to read, reverse engineer, and arbitrarily modify the actuation sequence on the MEDA controller. Moreover, the attacker is aware of the checkpoint based defense, i.e., knows that the golden droplet maps are used to validate the assay execution at the end of each actuation cycle.

3) How does one launch an attack?: The cyberphysical biochip system comprises controllers, software, network interface, etc. The designer can source them from third-party vendors and integrated by the biochip designer [10]. One can connect the biochip to a network for software updates and to process the results online. Informed by this biochip supply chain, the attacker can launch as attack as follows:

TABLE II: Digital microfluidic attack space.

Attack space								
Types	Macro [18]			Micro [18]	Shadow [1]			
Sub-types	Add	Modify	Skip	Aliquot	Swap	Merge	Aliquot	Split
Effects	High result deviations			Minute result deviation	Variable result deviation between CPs			
Corresponding defenses								
DMFB	Probable security by randomized CPs [12]			Not applicable	Not applicable			
MEDA	Provable security by graph reconstruction [24]			Variability-aware droplet-map comparison [18]	Shadow-attack-aware extra checkpoint insertion			

1) Exploit the in-built hardware or software Trojan to access the biochip controller [12].
2) Use malware to gain control of the network and to manipulate the control software or stored actuation sequence [29].
3) Compromise the biochip by inducing faults in the controller or actuators using electrical probes or lasers [30].

4) What are the constraints on an attacker?: The attacker manipulates the results of the biochip in a stealthy and untraceable way. To do this, the attacker has to evade detection by the sensors. The defender can monitor the biochip using CCD camera [31]. The operator is assumed trustworthy and that the MEDA platform is not tampered with. The objective of the attacker is to launch shadow operations within the available slack for a droplet within an actuation cycle and remain undetected by the security co-processor.

B. Attack Space

We define two new classes of actuation tampering attacks for MEDA biochip: granular and shadow attacks. The attack space and their corresponding state-of-the-art defenses is shown in Table II. These are briefly described as follows:

1) Granular attacks: MEDA biochips offer fine-grained droplet control. The attacker can leverage this to launch attacks of varying granularity.

Macro-droplet attacks are malicious modifications to the actuation sequence that operate on whole droplets. This can be done by performing one of the following actions: *add*, *modify*, *skip* operation(s) specified in the assay. This results in a major deviation in the droplet map and is similar to the DMFB threat models discussed in earlier work [9], [12], albeit for an arbitrarily sized droplet.

Micro-droplet attacks extract a small droplet from a larger droplet using an aliquot operation [26]. These attacks are difficult to detect since droplets naturally lose volume through cutting or evaporation during movement [15]. Furthermore, the movement of micro-droplets is fast compared to large droplets. Therefore, micro-droplet attacks are stealthy and easy to launch.

2) Shadow attacks: As the MEDA biochip supports droplets of varying sizes and velocities, attacks can exploit slack in faster droplets to execute extra operations. Let t_{cycle} be the actuation cycle period, and t_{load} be the time required to load the actuation sequence through the scan chain. A droplet of size $m \times n$, subjected to a maximum p transport operations along *x*-direction within the available slack is given as:

$$t_{cycle} - t_{op} = p \cdot t_{load} + p \cdot \frac{m \cdot d}{v} \tag{3}$$

where v is the average droplet velocity, d is the MEDA biochip pitch and t_{op} is the time taken to perform a droplet operation before its transportation. Attacker can use these p operations to insert malicious implementation.

C. Defenses

To secure the DMFBs, various checkpointing schemes that either check randomized cells [13], or check static cells in the droplet neighbourhood [12]. The performance of these schemes can be further improved by the use of pin-constrained architectures [32]. However, these schemes could only offer probabilistic guarantees owing the sensor resource constraints. MEDA biochips overcome this limitations through integrated fine-grained sensing. These sensor results (droplet-maps) can be used to reconstruct the sequencing graph and thereby verify the implementation [24]. However, the enhanced droplet manipulation abilities of MEDA over DMFB open doors for newer attacks such as micro-droplet and shadow attacks. A variability-aware droplet-map comparision scheme was proposed to detect MEDA-specific micro-droplet attacks [18]. To validate a MEDA operation, real-time sensor results are compared with the corresponding expected (golden) droplet map. The comparison between real-time and golden droplet map needs to account for the expected natural variations in droplet size and shape. This is achieved by dividing the droplet map into an three regions: ideal droplet occupancy region, guard-band region, and empty region. The occupancy region and empty region are monitored for security validation. A drawback of previous work is that they do not consider shadow attacks that insert malicious operations by taking advantage of the available slack between checkpoints.

IV. Shadow Attacks

The study of MEDA features and their security implications reveal that the fine-grained and flexible control enables MEDA-specific attacks: micro-droplet and shadow attacks. Micro-droplet has been demonstrated in the previous work [18]. In this section, we formalize the definition of shadow attack and describe its sub-classes.

A. Shadow Attack - Formal Definition

The state of the MEDA biochip at time t consists of the droplet map and a unique identifier associated with each droplet on the MEDA biochip at that time instant. The state at time t, $S^t = \langle L_S^t, I_S^t, \delta_S^t \rangle$, where L_S^t, I_S^t, and δ_S^t is the droplet map, and the set of all droplet identifiers at time t,

and a function that maps the droplet identifier to the location. At time t, the *shadow state* $S^t = \langle L_S^t, I_S^t, \delta_S^t \rangle$ of the golden state $G^t = \langle L_G^t, I_G^t, \delta_G^t \rangle$, if $L_S^t = L_G^t$, $|I_S^t| = |I_G^t|$, and $\delta_S^t \neq \delta_G^t$. Hence, the shadow state has a similar droplet map to the golden state. An assay is a sequence of state transitions wherein a state S^t transitions to state S^{t+1} on a fluidic operation F^t. i.e., $S^t \xrightarrow{F^t} S^{t+1}$. If the state S^t transitions to state $\tilde{S}^{t+1}$ with the same droplet map as S^{t+1} on a different set of fluidic operations $\tilde{F}^t$, then $\tilde{F}^t$ is a shadow operation. $S^t \xrightarrow{\tilde{F}^t} \tilde{S}^{t+1}$, where $L_S^{t+1} = L_{\tilde{S}}^{t+1}$.

Example 3. *Fig. 2(c) shows a MEDA biochip golden state at the end of an actuation cycle. The golden state at $t+1$ is given by $G^{t+1} = \langle L_G^{t+1}, I_G^{t+1}, \delta_G^{t+1} \rangle$, where droplet identifier set is $I_G^{t+1} = \{A, B, C\}$, the droplet map is $L_G^{t+1} = \{((10,5),(8,8)), ((6,2),(5,3)), ((2,6),(1,7))\}$, and the mapping function is $\delta_G^{t+1}(A) = ((6,2),(5,3))$, $\delta_G^{t+1}(B) = ((2,6),(1,7))$, $\delta_G^{t+1}(C) = ((10,5),(8,8))$. Each droplet location is denoted by the co-ordinates of the bottom left and the top right MCs. Fig. 2(f) shows a shadow state $S^{t+1} = \langle L_S^{t+1}, I_S^{t+1}, \delta_S^{t+1} \rangle$, where $L_S^{t+1} = L_G^{t+1}$, $I_S^{t+1} = I_G^{t+1}$, and $\delta_S^{t+1}(C) = \delta_G^{t+1}(C)$, $\delta_S^{t+1}(B) = \delta_G^{t+1}(A)$, $\delta_S^{t+1}(A) = \delta_G^{t+1}(B)$, after malicious droplet manipulation. The golden state G^{t+1} and the shadow state S^{t+1} have an identical number of droplets and droplet map, but droplet locations are different.*

B. Attack Classification

We classify the shadow attacks into following sub-classes:

1) *Swap:* where droplets interchange their respective locations as shown in Example 1 and Figs. 2(d)-(f).
2) *Split-merge:* where two droplets are split, and the resulting child droplets of one fluid are merged with the child droplets of the other fluid. This results in two droplets that are a mix of two droplets, as shown in Figs. 5(d)-(f).
3) *Aliquot-merge:* where an aliquot droplet is extracted from two droplets, and these aliquot droplets are mixed with the other droplet. This is shown in Figs. 5(g)-(i).
4) *Merge-split:* where two droplets are merged and then split, resulting in contaminated droplets, as shown in Figs. 5(j)-(l).
5) *I/O-swap:* where a droplet can be swapped with a different droplet from an inlet and the swapped droplet pushed out to an outlet, as shown in Figs. 6(d)-(f).

We describe these attacks in detail next.

1) Swap: A shadow transport operation swaps two (or more) droplets while maintaining the droplet map at the end of an actuation cycle. For this, droplet paths should be close and the timing slack should be considerable. Let L_A^t (L_B^t) and L_A^{t+1} (L_B^{t+1}) be the location of two equal sized droplets A (B) at cycle t and $t+1$, respectively. The maximum distance traveled by the droplet A (B) in an actuation cycle t_{cycle} is $d_{max}^A = v^A \cdot t_{cycle}$ ($d_{max}^B = v^B \cdot t_{cycle}$), where v^A (v^B) is the average speed of the droplet A (B). The shadow swap is viable if the locations L_A^t and L_B^t are such that the distance between L_A^t to L_B^{t+1} and L_B^t to L_A^{t+1} can be traversed in less than an actuation cycle, as shown in Figs. 2(d)-(f). Example 1 illustrated the shadow transport. The necessary condition for shadow swap is:

$$\left(|L_A^t - L_B^{t+1}| < v^A \cdot t_{cycle}\right) \quad \wedge$$
$$\left(|L_A^{t+1} - L_B^t| < v^B \cdot t_{cycle}\right) \tag{4}$$

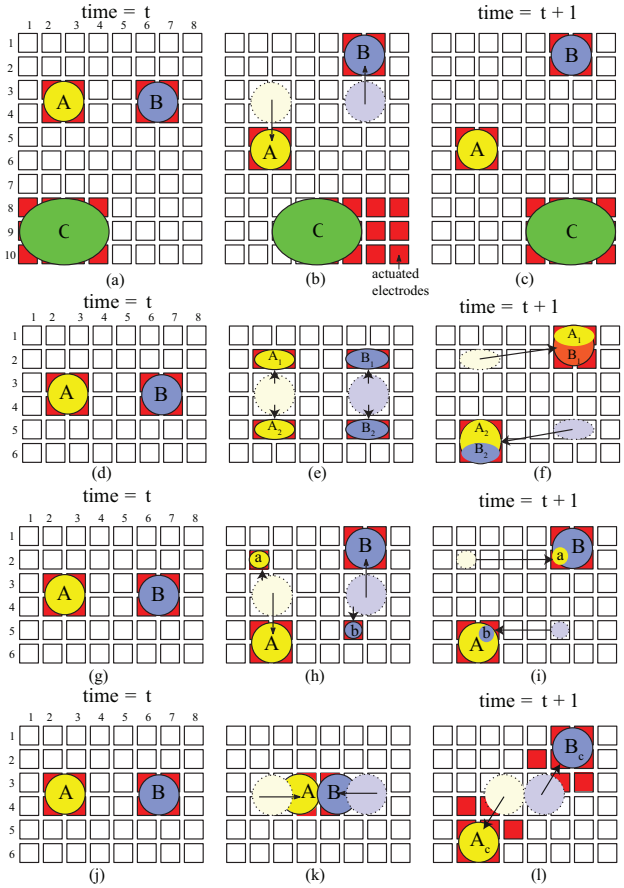

Fig. 5: Illustration of shadow operations. (a)-(c): Snapshots of the golden execution. (d)-(f): A split-merge splits droplets A and B and then merges child droplets. (g)-(i): An aliquot-merge extracts aliquot droplets from A and B and merges it with the other droplets. (j)-(l): A merge-split contaminates droplets A and B by merging them and splitting.

2) Split-merge: According to Eq. 1, the average velocity of a droplet depends on its size. A smaller droplet can travel further in a time duration. A split operation is performed by actuating electrodes at opposite ends of a droplet (refer Fig. 5(e)) to create two equal-sized child droplets. The child droplets coming from different parent droplets are mixed, contaminating both. The following example illustrates the split-merge attack.

Example 4. *Figs. 5(a)-(c) show the golden snapshot of the MEDA biochip in which droplets A and B move quickly to their destinations and wait for the larger droplet to reach its destination (see Example 1). Figs. 5(d)-(f) show the malicious split-merge that an attacker may perform. The two droplets are split into equal halves (A_1;A_2 and B_1;B_2). A_1 is transported*

and merged with B_1 and B_2 is transported and merged with A_2. These shadow operations retain the golden droplet map (Fig. 5(c)) at time $t+1$, i.e., the MEDA states in the Fig. 5(c) and Fig. 5(f) have the exact droplet maps.

Let $L^*_{A_1}$ ($L^*_{B_1}$) and $L^*_{A_2}$ ($L^*_{B_2}$) be the locations of the child droplets after splitting the droplet A (B). '*' denotes the creation time of the children droplets and it can happen any time between t and $t+1$. Let the time taken to split is t_{split}, and the velocity of the child droplets are v^{A_i} and v^{B_i}, for $i = 1, 2$. The maximum distance the child droplet A_i can be transported is given by:

$$d^{A_i}_{max} = v^{A_i} \cdot (t_{cycle} - t_{split}), \qquad \text{for } i = 1, 2 \tag{5}$$

Similar equations can be derived for the droplets B_i. In order to preserve the golden droplet map, droplets A_1 and B_2 should be transported and merged with droplets B_1 and A_2, respectively. All these operations must be completed within the actuation cycle, yielding the following necessary condition for the split-merge attack is given by:

$$\begin{aligned} &\Big(|L^*_{A_1} - L^{t+1}_B| < v^{A_1} \cdot (t_{cycle} - t_{split})\Big) \quad \wedge \\ &\Big(|L^*_{B_2} - L^{t+1}_A| < v^{B_2} \cdot (t_{cycle} - t_{split})\Big) \end{aligned} \tag{6}$$

Here, we assume that the actuation voltage V is greater than the threshold voltage required to split a droplet in t_{cycle}. The split attack is feasible when there is enough timing slack available to perform droplet split and droplet transportation as captured by the condition in (6).

3) Aliquot-merge: An aliquot operation extracts a small droplet from a large droplet that can move fast [33]. The aliquot droplets of different droplets can be swapped and merged, resulting in contaminated droplets, as shown in Fig. 5(g)-(i). This attack enables shadow operation involving droplets of various sizes.

Let L^*_a (L^*_b) be the location of an aliquot droplet a (b), created from the droplet A (B). '*' denotes the creation time of aliquot droplets, between t and $t+1$. An aliquot operation has four steps, i.e., four actuations [33]. This leaves $(t_{cycle} - t_{aliquot})$ time for swapping aliquot droplets, where $t_{aliquot}$ denotes the time to create the aliquot droplet. The maximum distance an aliquot droplet can travel in an actuation cycle is $d^a_{max} = v^a \cdot (t_{cycle} - t_{aliquot})$, where v^a is the velocity of the aliquot droplet a. Analogously, $d^b_{max} = v^b \cdot (t_{cycle} - t_{aliquot})$, where v^b is the velocity of the aliquot droplet b. The necessary condition for shadow aliquot operation is:

$$\begin{aligned} &\Big(|L^*_a - L^{t+1}_B| < v^a \cdot (t_{cycle} - t_{aliquot})\Big) \quad \wedge \\ &\Big(|L^*_b - L^{t+1}_A| < v^b \cdot (t_{cycle} - t_{aliquot})\Big) \end{aligned} \tag{7}$$

4) Merge-split: An attacker can manipulate the droplet route such that different droplets merge, followed by a split. This contaminates both droplets. Let L^t_A (L^t_B) and L^{t+1}_A (L^{t+1}_B) be the locations of droplet A (B) at cycle t and $t+1$, respectively. The maximum distance traveled by the droplets is $d^A_{max} = v^A \cdot (t_{cycle} - t_{split})$ and $d^B_{max} = v^B \cdot (t_{cycle} - t_{split})$, where v^A (v^B) is the average speed of droplet A (B). If the droplets merge and split at the location L^*_{mid}, then the necessary condition for this operation is:

$$\begin{aligned} &\Big(|L^t_A - L^*_{mid}| + |L^*_{mid} - L^{t+1}_A| < v^A \cdot (t_{cycle} - t_{split})\Big) \quad \wedge \\ &\Big(|L^t_B - L^*_{mid}| + |L^*_{mid} - L^{t+1}_B| < v^B \cdot (t_{cycle} - t_{split})\Big) \end{aligned} \tag{8}$$

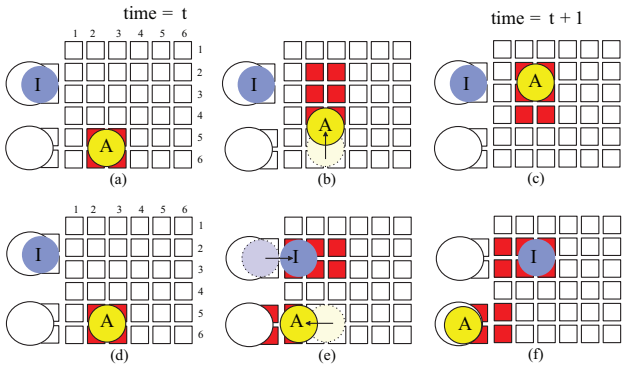

Fig. 6: Illustration of shadow I/O-swap attack. (a)-(c): Snapshots of the golden execution of droplet A from cycle t to $t+1$. (d)-(f): A swap of droplet A and droplet I, and subsequent output of droplet A through outlet O.

5) I/O-swap: An attacker can swap a droplet on the biochip with a different droplet from an inlet. The replaced droplet on the biochip can be pushed out through an outlet port, thus, maintaining the droplet map. This can be seen a special-case of the swap attack which considers droplet insertion. Let L^t_A and L^{t+1}_A be the location of droplet A at cycle t and $t+1$, respectively. Let L_I be the location of an inlet that produces equal sized droplet I and L_O be the location of an fluid outlet. The maximum distance traveled by the droplet A (I) in an actuation cycle t_{cycle} is $d^A_{max} = v^A \cdot t_{cycle}$ ($d^I_{max} = v^I \cdot t_{cycle}$), where v^A (v^I) is the average speed of the droplet A (I). The shadow I/O-swap is viable if the locations L^t_A, L_I, L_O are such that the distance between L_I to L^{t+1}_A and L^t_A to L_O can be traversed in less than an actuation cycle, as shown in Figs. 6(d)-(f). The necessary condition for shadow I/O-swap is:

$$\begin{aligned} &\Big(|L_I - L^{t+1}_A| < v^I \cdot t_{cycle}\Big) \quad \wedge \\ &\Big(|L^t_A - L_O| < v^A \cdot t_{cycle}\Big) \end{aligned} \tag{9}$$

C. Sufficient Condition for Shadow Attacks

The expressions derived in (4)-(9) capture the necessary conditions for the shadow attacks. These are not sufficient condition for a meaningful attacks. For example, if droplet A and B are same fluids, then the attacks do not lead to any meaningful impact to the assay. Similarly, if the droplet A and B is a waste droplet being discarded from the biochip, then the shadow attacks do not lead to any tampering of final results. In other words, the sufficiency of the attack expressions (4)-(9) varies with the bioassay and its synthesis. To keep our discussion generic, we limit ourselves to the necessary conditions. This also streamlines the defense mechanism.

V. Detecting Shadow Attacks

Having described the various types of shadow attacks on MEDA biochip, we now describe a defense mechanism that inserts extra checkpoints to mitigate these attacks. This leads to increase in the number of checkpoints. Next, we prune the number of checkpoints by revoking checkpoints that are not susceptible to any attacks from the baseline defense (checkpoint at each cycle).

A. Shadow Attack Aware Checkpoints

We recommend a defense to go through the shadow operation possibilities in each actuation cycle and shrink the duration of those actuation cycles susceptible to shadow attacks. We use Eq. (4)-(8) to measure the susceptibility of an actuation cycle. For each actuation cycle in the golden actuation sequence, we establish the number of droplets and their sizes, which determines their velocity. We use these parameters in Eq. (4)-(8) to check the condition of each shadow operation between all pairs of droplets appearing in the actuation cycle. Algorithm 1 outlines the defense.

A bioassay is executed on the MEDA platform by synthesizing the assay specification into a sequence of actuation cycles [17]. The actuation sequence has clearly delineated actuation cycle boundaries. A security co-processor checks the golden map against the droplet map at the end of each actuation cycle. Synthesis algorithms for the MEDA biochip are oblivious to shadow operations [34], [17], [35]. Hence, the shadow operations can subvert the bioassay. The defense increases the number of actuation cycles in the golden actuation sequence, which increases the checkpoints and the space needed to store the golden maps.

Algorithm 1: Defend Against Shadow Attacks With Extra Checkpoints.

Input: Actuation sequence
Output: Shadow attack resistant actuation sequence

```
for each actuation cycle in the actuation sequence do
    t_cycle = time of the current actuation cycle
    t = t_cycle
    for each droplet pairs A, B
        in the current actuation cycle do
        // Swap attack -- Eq. 4
        if Expression 4 is true for A and B then
            Calculate min cycle time to avoid swap.
            t = min{t, t_min}
        // Split-merge attack -- Eq. 6
        if Expression 6 is true for A and B after splitting
         then
            // Any one or both A and B can
               be split
            Calculate min cycle time to avoid split-merge.
            t = min{t, t_min}
        // Aliquot-merge attack -- Eq. 7
        if Expression 7 is true for A and B then
            Calculate min cycle time to avoid
             aliquot-merge.
            t = min{t, t_min}
        // Merge-split attack -- Eq. 8
        if Expression 8 is true for A and B then
            Calculate min cycle time to avoid merge-split.
            t = min{t, t_min}
    for each pair A, Inlet I, Outlet O
        in the current actuation cycle do
        // IO-swap attack -- Eq. 9
        if Expression 9 is true for A, I, O then
            Calculate min cycle time to avoid I/O-swap.
            t = min{t, t_min}
    if t < t_cycle then
        Split current actuation cycle into two cycles of
         lengths t and (t_cycle − t) respectively.
```

B. Pruning of Checkpoints

We next minimize the number of checkpoints by dropping those time-steps that are not susceptible to any attacks. Consider a case where droplet occupancy be spare in a given time-step t. Let the time interval between current time-step ($t_{current}$) and previous checkpoint time-step (t_{prev_check}) be $t_{slack} = t_{current} - t_{prev\_check}$. The attack condition is diagnosed by evaluating expressions (4), (6), (7), (8), and (9) by substituting $t_{cycle} = t_{slack}$. If in this time interval no attack is possible, i.e., the conditions 4-9 are not satisfied for $t_{cycle} = t_{current} - t_{prev\_check}$. Then, time-step $t_{current}$ can be safely dropped from the checkpoint list. We use this observation to minimize the number of checkpoints. We initialize the checkpoint list with all time-steps, considered as baseline. We then iterate over each assay time-step $t = 1, 2, \cdots, T$; Time-step $t_{current}$ is dropped from the list if it is safe; Else we update the latest checkpoint time-step as $t_{prev\_check} = t_{current}$ and move to the next time-step $t_{current} + 1$. The proposed checkpoint pruning methodology is shown in Algorithm 2.

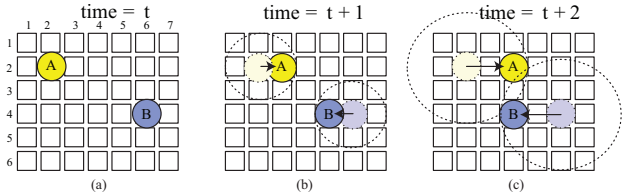

Fig. 7: Biochip snapshots with droplets A and B at three consecutive time-steps. The dotted circle represents the distance droplet can traverse from its latest checkpoint.

Example 5. *Consider biochip snapshots shown in Fig. 7. Let time-step t be in the checkpoint list, i.e., $t \in CP$. At time-step $t + 1$, the droplets A and B have a slack of one time-step from the previous checkpoint t. The area that these droplets can traverse in the given slack is showcased by a dotted circle which is centered at the previous checkpoint droplet location. As shown in Fig. 7 (b), it is not possible for droplets A and B to be swapped or contaminated in the given time slack. Hence, the time-step $t+1$ can be dropped from the checkpoint*

Algorithm 2: Pruning of Checkpoint List

Input: Actuation sequence
Output: List of time-steps for checkpoints.

```
// Initialize the checkpoint list
List CP = t ∈ [1, T];
// Initialize the checkpoint pointer
t_prev_check = 0;
for each actuation cycle in the actuation sequence do
    t_current = time of the current actuation cycle
    // Slack available for attacker
    t_slack = t_current − t_prev_check
    for each droplet and/or dispense port pairs A, B
        in the current actuation cycle do
        // Shadow attack check
        if ((4) or (6) or (7) or (8) or (9) is true with
          t_cycle = t_slack) then
            // Update checkpoint pointer
            t_prev_check = t_current;
        else
            // Prune checkpoint list
            Delete t_current from CP;
Return(CP)
```

list CP. Now, at time-step $t+2$, the available slack from previous checkpoint is two time-steps (previous checkpoint is at t). The area that the respective droplets can traverse is shown in Fig. 7 (c) by the dotted circle. It is now feasible for droplets A and B to be contaminated by merging in the given time slack. In order to avoid this, we retain the checkpoint at $t+2$, i.e., $t+2 \in CP$. This conservatively enforces the proper behavior of the droplets.

C. Guarantee of Detection

In this work, we use a conservative approach of guarding against the necessary conditions associated with shadow attacks. These necessary conditions may not lead to a meaningful attack, as explained in Section IV-C. However, our approach guarantees the neutralizing of the modeled threats. In other words, we choose to have false positives to eliminate false negatives. Therefore, there is no possibility of detection escape. We prove this claim below:

Theorem 1: The checkpoint list CP generated by Algorithm 2 guarantees defense against all the modeled threats in Section IV.

Proof: We prove the theorem by contradiction. Suppose droplets A and B are susceptible to shadow attack at $t_{current}$ i.e., $t_{current} \notin CP$.

Let t_{prev_check} be the previous checkpoint time-step, then the slack available to the attacker is $t_{slack} = t_{current} - t_{prev\_check}$. According to our assumption either (4) or (6) or (7) or (8) or (9) is true with $t_{cycle} = t_{slack}$). Therefore, Algorithm 2 adds $t_{current}$ to the checkpoint list CP i.e., $t_{current} \in CP$. This contradicts the assumption and the theorem follows. □

VI. Simulation Results

In this section, we analyze the shadow attack on the sample preparation assay case on MEDA and showcase the defense on this and other real-life bioassay benchmarks.

A. Case Study: Sample Preparation

The attacks described in Section IV can be best demonstrated on a bioassay that uses multiple-sized reagent droplets. For this purpose, we choose a sample preparation protocol that uses a mixing model that mixes various sized droplets. Sample preparation is the process of mixing two or more input reagents in a desired volumetric ratio. Sample preparation algorithms targeting various biochip platforms have been widely studied [36]. The traditional DMFBs allow for a (1:1) mixing model, whereas MEDA enables mixing of different size droplets — *multiple mixing model*. In other words, we choose a protocol (multiple mixing model) that is unique to MEDA biochip to demonstrate attacks (shadow attacks) that are unique to MEDA biochips. The case study described here is a generic example and its findings hold true for other protocols as well.

We consider a mixing graph (Fig. 8) to generate a sample with the mixing ratio of {sample:buffer = 125:131} [37]. We synthesized the mixing graph for a 40×40 MEDA biochip with a square microelectrode of $50\,\mu\mathrm{m}$ side. The other hardware parameters are shown in Table I. We assume each mix operation takes 100 actuation cycles, based on previous experimental demonstration [38]. It takes 438 actuation cycles to complete the sample preparation. Fig. 9 shows the snapshots of the MEDA biochip in different actuation cycles. The MEDA biochip manipulates droplets of dimensions 8×8 (largest), 4×8, and 4×4 (smallest) to implement the sample preparation and their speed can be calculated as $0.8\,\mathrm{mm/sec}$, $1\,\mathrm{mm/sec}$, and $1.3\,\mathrm{mm/sec}$, respectively [17].

The time taken by the largest (smallest) droplet to move one step, i.e., 8 (4) microelectrodes, at $0.8\,\mathrm{mm/sec}$ ($1.3\,\mathrm{mm/sec}$) is $0.5\,\mathrm{s}$ ($0.17\,\mathrm{s}$). If the actuation cycle time is set to t_{cycle}^{max}, which is the time taken by the largest droplet to move one step ($0.5\,\mathrm{s}$), the execution takes 438 cycles to finish. This implementation is prone to shadow attacks due to the time slack between different size droplets. A naïve defense can be to reduce the available slack to 0. This can be done by constraining the cycle time with that of the faster droplet (4×4 droplet takes $0.17\,\mathrm{s}$ to move one step) rather than the slower droplet (8×8 droplet takes $0.17\,\mathrm{s}$ to move one step). However, in a given bioassay, the total distance traveled by the droplets remains the same. Therefore, the number of shorter cycles is increased to: $(438 \times 0.5)/0.17 = 1702$.

The difference in droplet sizes introduces a time slack in the actuation cycles. An attacker can used this slack to launch shadow attacks. Fig. 10(a) shows a snapshot of the MEDA biochip (actuation cycle 115) with three droplets: sample S_2, buffer B_2, waste W_1 and intermediate R_1. We describe the possible attacks on it.

Swap: In the cycle 115 (Fig. 10(a)), droplets B_2 and R_1 are expected to be at $L_{B_2}^{115} = \{(8,34),(10,36)\}$ and $L_{R_1}^{115} = \{(14,30),(16,32)\}$, respectively. Droplets B_2 and R_1 are of

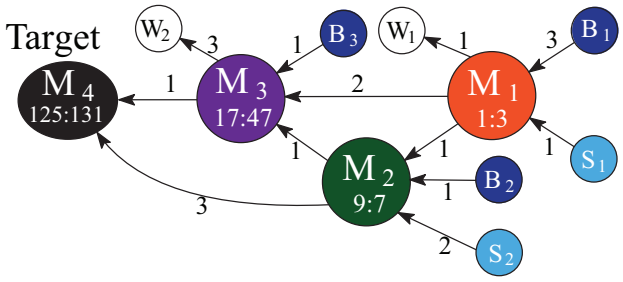

Fig. 8: A mixing graph that dilutes a sample (S) with buffer (B) in the ratio of {sample:buffer = 125:131} using a multiple mixing model. S_i is a sample droplet, B_i is a buffer droplet, and M_i is the mixing operation. The number on the edges denotes the relative size of the droplet.

the size 4×4 each and the maximum distance traveled by the droplets in $0.5\,\text{s}$ (t_{cycle}) is $d_{max} = 0.5 \cdot 1.3 \cdot 10^3 = 650\,\mu\text{m}$. Droplets of buffer B_2 and the intermediate droplet R_1 can be swapped as the radial distance between the pair of droplets is $500\,\mu\text{m}$, which is less than d_{max}, as shown in Fig. 10(b).

Merge-split: Droplet paths of B_2 and W_1 can be modified such that these droplets merge. This can be achieved by moving droplet B_2 diagonally from $L^{114}_{B_2} = \{(8,38),(10,40)\}$ to $\{(12,34),(14,36)\}$ and droplet W_1 from $L^{114}_{W_1} = \{(20,30),(22,32)\}$ to $\{(16,34),(18,36)\}$, as shown in Fig. 10(c). By this movement, the droplets come in contact and result in contamination. The droplets are then routed to $L^{115}_{B_1} = \{(8,34),(10,36)\}$ and $L^{115}_{W_1} = \{(24,30),(26,32)\}$. The total distance traveled by the droplets B_2 and W_1 are $484\,\mu\text{m}$, and $627\,\mu\text{m}$ respectively (which are less than the $d_{max} = 650\,\mu\text{m}$).

Split-merge: Droplet S_2 is split into two 4×4 droplets (S^1_2, S^2_2). The split operation of a 8×4 droplet takes $t_{split} = 0.09\,\text{s}$ [17], [23]. The resulting child droplets are displaced by one electrode ($100\,\mu\text{m}$). After the split operation, the maximum distance the child droplets can travel is $d_{max} = (0.5 - 0.09) \cdot 1.3 \cdot 10^3 = 533\,\mu\text{m}$. One child droplet S^2_2 is swapped with W_1 and W_1 is merged with S^1_2, i.e., S^2_2 is moved from $\{(30,38),(32,40)\}$ to $\{(24,30),(26,32)\}$, S^1_2 is moved from $\{(22,38),(24,40)\}$ to $\{(16,30),(18,32)\}$ and W_1 is moved to $\{(20,38),(22,40)\}$ resulting in a merge, as shown in Fig. 10(d). The distance traveled by the child droplets of S and W_1 are $500\,\mu\text{m}$ and $400\,\mu\text{m}$ respectively, which are less than $d_{max} = 533\,\mu\text{m}$.

Aliquot-merge: An aliquot operation requires four steps. The total time for the aliquot operation can be calculated as $t_{aliquot} = 0.36\,\text{s}$, wherein each step the aliquot droplet is moved across an electrode ($100\,\mu\text{m}$) [26]. This leaves $0.14\,\text{s}$ for transportation. In this remaining time slack, an aliquot droplet can traverse less than two electrodes. Therefore, an aliquot-merge attack is not possible.

I/O-swap: The inlet ports are located at $\{(40,6),(40,32)\}$, and outlet ports are located at $\{(2,30),(22,2)\}$. When each time-step is being monitored, then I/O-swap attack is possible with the given I/O port locations. However, if the checkpoints are pruned (as shown next), the available slack increases making such attack possible. For example, if the time-steps 328 to 332 are not being monitored, then the sub-droplet M_3 being routed from $\{(20,8)\}$ to $\{(30,16)\}$ is susceptible to I/O-swap attack. The sub-droplet can be replaced with a similar size droplet from inlet B located at $\{(40,6)\}$, and subsequently the sub-droplet can be routed to outlet at $\{(22,2)\}$.

Our analysis shows that in cycle 115 the assay implementation is prone to swap, split-merge, and merge-split attacks. Similarly, other cycles are analyzed, and a total of eight cycles (114-118 and 221-225) were found to be vulnerable to shadow attacks.

In order to avoid these attacks, we insert one extra checkpoint (CP) in each of the cycles. Table III shows the results of our analysis and compares the cost of our solution with the shadow oblivious execution (t^{max}_{cycle}) and the naïve defense (t^{min}_{cycle}). Our proposal secures the assay implementation with eight additional CPs. This incurs a memory overhead of 1.8% to store the droplet maps corresponding to the extra CPs. The naïve solution of changing the cycle time to t^{min}_{cycle} incurs a memory overhead of 100%. Our solution saves the information of 8 cycles that are prone to shadow attacks. In these cycles, the security co-processor performs the extra CP to validate the assay implementation. This validation incurs no time cost as the security co-processor validates in parallel with the assay execution.

TABLE III: Shadow attacks on the sample preparation.

	Shadow oblivious	Naïve defense	Proposed solution
No. of CPs	438	1702	446
Swap attack	Y	N	N
Split-merge attack	Y	N	N
Merge-split attack	Y	N	N
Aliquot-merge attack	N	N	N
I/O-swap attack	N	N	N

B. Experiments on Other Assays

We experiment to test the cost of our solution for five bioassays. We simulated the execution of the assay on a DMFB biochip and utilize these results to estimate the implementation cost for a MEDA biochip. We simulated the assay using the open-source UCR Static Simulator [39]. To emulate the high density of microelectrodes in a MEDA biochip, we broke up each DMFB electrode into 4×4 microelectrodes [34], [40]. Since these assays involve mixing of two unit droplets, the largest droplet is of 8×4 size. We consider a MEDA biochip with a square microelectrode of $100\,\mu\text{m}$ side. The average speeds of 8×4 and 4×4 sized droplets are $1\,\text{mm/sec}$, and $1.3\,\text{mm/sec}$ [17]. We use these parameters to analyze susceptibility of the bioassay executions and to estimate the cost of the defense. Next, we use the pruning mechanism of Algorithm 2 to reduce the total number of checkpoints required and report the benefits.

Table IV exhibits the results of analysis of the shadow attacks and the defense for five real-life bioassay implementations. Our evaluation reveals that all bioassays are vulnerable to shadow attacks. The recommended defense increases the

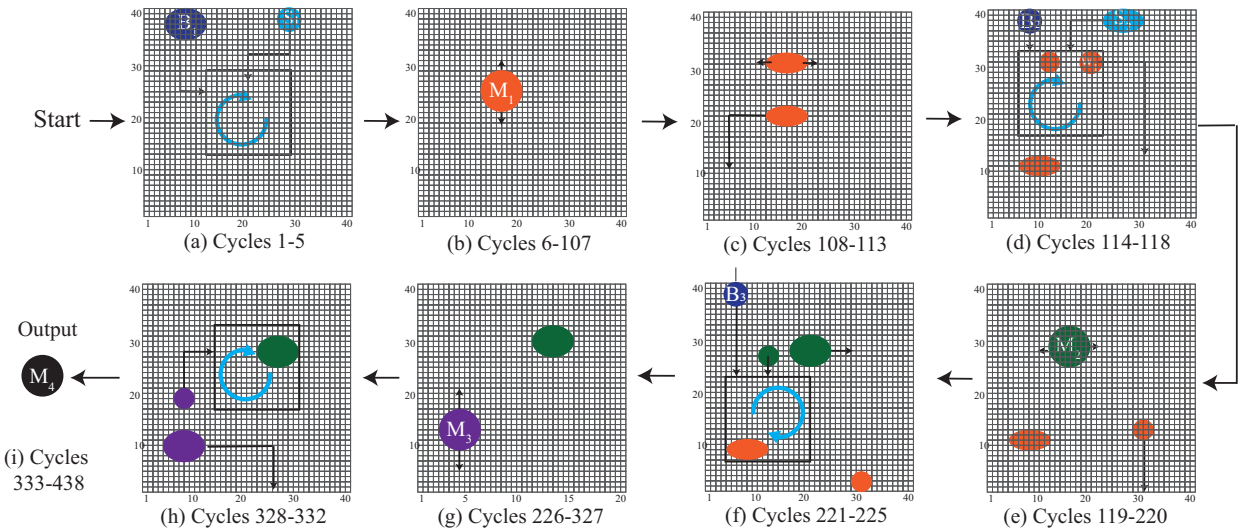

Fig. 9: Snapshots of the MEDA biochip implementing the mixing graph shown in Fig. 8. (a-b) The mix operation M_1, which takes 100 cycles (6-106). (c) The mixed droplet M_1 is split into three (droplets used in mix M_2, mix M_3, and a waste droplet W_1). (d-e) The mix operation M_2, which is performed in 119-220 cycles. (f) The resulting mixed droplet M_2 is split into two droplets that are used in M_3 and M_4 operations. (g-h) Mix operation M_3 is performed, and the resulting droplet is split into a droplet for M_4 operation besides yielding a waste droplet M_4. (i) The output M_4 mix operation is performed in cycles 333-438.

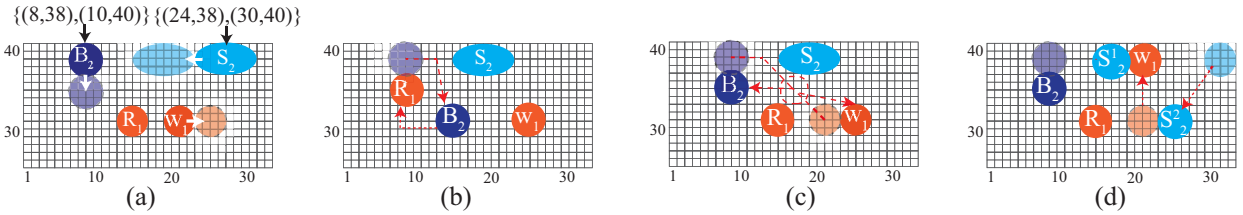

Fig. 10: (a) Snapshot of a portion of the MEDA biochip at the beginning of cycle 115. (b) Swapping of droplets B_2 and R_1. (c) Merging of droplets B_2 and W_1, followed by a split. (d) Splitting of S_2 and swapping of the child droplet S_2^2 with W_1.

TABLE IV: Shadow attacks and defense on real-life benchmarks.

Assay	Baseline	Adding extra CPs for shadow defense			CP list pruning of safe time-steps		
	#CPs Shadow oblivious	#Vulnerable cycles	#CPs Proposed	Storage overhead	#Safe cycles	#CPs Proposed	Storage overhead
PCR	969	29	998	3.0 %	659	339	-65%
InVitro 1	1317	33	1350	2.5 %	792	558	-57.6%
InVitro 2	3382	124	3506	3.7 %	1775	1792	-47%
Protein 1	5440	305	5745	5.6 %	3022	2723	-49.9%
Protein 2	24151	1648	25799	6.8 %	12719	13470	-44.2 %
Remia	34	25	59	73.5 %	7	52	52.9%
Dilution	36	30	66	83.3 %	5	61	69.4%
PCR Stream	70	70	140	100 %	0	140	100%

memory storage expense by less than 7% (to store the golden droplet maps). The overhead is shown to be insignificant with CP pruning. The bioassays have numerous safe cycles that are not susceptible to any of the attacks which can be dropped from the checkpoint list. This leads to reduction of more than 45% in the number of checkpoints. This shows that the proposed defense is very practical and extensible.

Further, we applied the shadow attack analysis to three smaller sample preparation bioassays. The analysis reveals that these smaller bioassays are more prone to shadow attacks. These bioassays are targeted for smaller biochips compared to earlier considered bioassays. For example, the Remia sample preparation bioassay is targeted for 5×5 biochip and takes only 40 cycles to complete [41]; whereas, PCR bioassay is targeted for 13×15 biochip and completes in 969 cycles. This means that for these smaller bioassays the droplets are in closer proximity, leading to more chances of shadow attacks. The defense increases the memory storage by more than 70%. The pruning mechanism only has a marginal effect on the overall number of checkpoints.

VII. Conclusion

We have demonstrated advance class of *shadow attacks* on MEDA biochips, which thwart a state-of-the-art validation technique. The attacker utilizes the time slack to introduce shadow operations. We derived the necessary conditions for such attacks and applied them to determine the unsafe periods. Additional checkpoints alleviate this risk in an extensible way. These additional checkpoints can be easily accommodated by pruning the safe cycle. The overall checkpoint overhead shows that our proposal is not only secure but also has lower overhead (by at least 44%) than the baseline defense.

Variations of these attacks are conceivable. Examples include multiple split operations followed by a switch of the ensuing smaller droplets. However, these are conditional on the availability of ample slack. The relative ratio of the largest and smallest sized droplets determines the slack. In practical bioassays this ratio is small (about 5) [2], [9], [42], [43]. This makes the multiple split-swap attacks impractical. The aliquot-merge analysis also supports this assertion that sophisticated attacks like splitting multiple times before swapping are not feasible in a realistic assay due to limited timing slack. Therefore, we studied five practical attacks to streamline the defense and its evaluation.

References

[1] M. Shayan, S. Bhattacharjee, T.-C. Liang, J. Tang, K. Chakrabarty, and R. Karri, "Shadow attacks on meda biochips," in *Proc. Intl. Conf. on Comput.-Aided Design*, 2018, pp. 73:1–73:8.

[2] N. Vergauwe, D. Witters, F. Ceyssens, S. Vermeir, B. Verbruggen, R. Puers, and J. Lammertyn, "A versatile electrowetting-based digital microfluidic platform for quantitative homogeneous and heterogeneous bio-assays," *J. Micromechanics Microengineering*, vol. 21, no. 5, p. 054026, 2011.

[3] A. H. C., K. Choi, R. P. Luoma, J. M. Robinson, and A. R. Wheeler, "Digital microfluidic magnetic separation for particle-based immunoassays," *Anal. Chem.*, vol. 84, no. 20, pp. 8805–8812, 2012.

[4] J. He, A. Chen, J. Lee, and S. Fan, "Digital microfluidics for manipulation and analysis of a single cell," *Intl. J. Molecular Sciences*, vol. 16, no. 9, pp. 22 319–22 332, 2015.

[5] M. Ibrahim, K. Chakrabarty, and K. Scott, "Synthesis of cyberphysical digital-microfluidic biochips for real-time quantitative analysis," *IEEE Trans. on CAD of Integr. Circuits and Syst.*, vol. 36, no. 5, pp. 733–746, 2017.

[6] (2016) FDA advisors back approval of Baebies' SEEKER analyzer for newborns. [Online]. Available: http://baebies.com/fda-advisors-back-approval-baebies-seeker-analyzer-newborns

[7] B. Bhattacharjee and H. Najjaran, "Size dependent droplet actuation in digital microfluidic systems," in *Proc. SPIE, Micro- and Nanotechnol. Sensors, Syst., Applications*, vol. 7318, 2009.

[8] R. B. Fair, A. Khlystov, T. Tailor, V. Ivanov, R. Evans, V. Srinivasan, V. Pamula, M. Pollack, P. Griffin, and J. Zhou, "Chemical and biological applications of digital-microfluidic devices," *IEEE Des. Test. Comput.*, vol. 24, no. 1, pp. 10 – 24, 2007.

[9] S. S. Ali, M. Ibrahim, O. Sinanoglu, K. Chakrabarty, and R. Karri, "Security implications of cyberphysical digital microfluidic biochips," in *Proc. Intl. Conf. Comput. Design*, 2015, pp. 483–486.

[10] S. S. Ali, M. Ibrahim, J. Rajendran, O. Sinanoglu, and K. Chakrabarty, "Supply-chain security of digital microfluidic biochips," *IEEE Computer*, vol. 49, no. 8, pp. 36–43, 2016.

[11] S. Bhattacharjee, J. Tang, M. Ibrahim, K. Chakrabarty, and R. Karri, "Locking of biochemical assays for digital microfluidic biochips," in *Proc. European Test Symp.*, 2018, pp. 1–6.

[12] J. Tang, M. Ibrahim, K. Chakrabarty, and R. Karri, "Secure randomized checkpointing for digital microfluidic biochips," *IEEE Trans. on CAD*, vol. 37, no. 6, pp. 1119 – 1132, 2018.

[13] J. Tang, R. Karri, M. Ibrahim, and K. Chakrabarty, "Securing digital microfluidic biochips by randomizing checkpoints," in *Proc. Intl. Test Conf.*, 2016, pp. 1–8.

[14] P. Roy and A. Banerjee, "A new approach for root-causing attacks on digital microfluidic devices," in *Proc. Asian Hardware-Oriented Sec. Trust Conf.*, 2016, pp. 1–6.

[15] G. Wang, D. Teng, and S. Fan, "Digital microfluidic operations on micro-electrode array architecture," in *Proc. Intl. Conf. on Nano/Micro Engineered Molecular Syst*, 2011, pp. 1180–1183.

[16] G. Wang, D. Teng, Y. Lai, Y. Lu, Y. Ho, and C. Lee, "Field-programmable lab-on-a-chip based on microelectrode dot array architecture," *IET Nanobiotechnology*, vol. 8, no. 3, pp. 163–171, 2014.

[17] Z. Li, T. Ho, K. Lai, K. Chakrabarty, P. Yu, and C. Lee, "High-level synthesis for micro-electrode-dot-array digital microfluidic biochips," in *Proc. Design Automation Conf.*, 2016, pp. 1–6.

[18] M. Shayan, J. Tang, K. Chakrabarty, and R. Karri, *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, 2018.

[19] Z. Li, K. Y. T. Lai, K. Chakrabarty, T. Y. Ho, and C. Y. Lee, "Droplet size-aware and error-correcting sample preparation using micro-electrode-dot-array digital microfluidic biochips," *IEEE Trans. Biomed. Circuits Syst.*, vol. 11, no. 6, pp. 1380–1391, 2017.

[20] Z. Li, K. Lai, P. Yu, K. Chakrabarty, T. Ho, and C. Lee, "Droplet size-aware high-level synthesis for micro-electrode-dot-array digital microfluidic biochips," *IEEE Trans. Biomed. Circuits Syst.*, vol. 11, no. 3, pp. 612–626, 2017.

[21] T. Ho, K. Chakrabarty, and P. Pop, "Digital microfluidic biochips: recent research and emerging challenges," in *Proc. Intl. Conf. Hardware/Software Codesign Syst Synthesis.*, 2011, pp. 335–344.

[22] H. Ren, R. Fair, M. Pollack, and E. Shaughnessy, "Dynamics of electrowetting droplet transport," *Sensors and Actuators B: Chemical*, vol. 87, no. 1, p. 201206, 2002.

[23] J. H. Song, R. Evans, Y.-Y. Lin, B.-N. Hsu, and R. B. Fair, "A scaling model for electrowetting-on-dielectric microfluidic actuators," *Microfluidics and Nanofluidics*, vol. 7, no. 1, pp. 75–89, Jul 2009.

[24] T.-C. Liang, M. Shayan, K. Chakrabarty, and R. Karri, "Execution of provably secure assays on meda biochips to thwart attacks," in *Proc. ASP Design Automation Conf.*, 2019, pp. 51–57.

[25] K. Yi-Tse Lai, Y.-T. Yang, and C.-Y. Lee, "An intelligent digital microfluidic processor for biomedical detection," *J. Signal Process. Syst.*, vol. 78, no. 1, pp. 85–93, 2015.

[26] Z. Zhong, Z. Li, and K. Chakrabarty, "Adaptive error recovery in meda biochips based on droplet-aliquot operations and predictive analysis," in *Proc. Intl. Conf. Comput. Aided Design*, 2017, pp. 615–622.

[27] U.S.Government., "Increase in insider threat cases highlight significant risks to business networks and proprietary information," 2014. [Online]. Available: https://www.ic3.gov/media/2014/140923.aspx

[28] U. S. Service, CERT, C. Magazine, and Deloitte, "2011 cybersecurity watch survey: How bad is the insider threat?" 2011.

[29] R. Langner, "Stuxnet: dissecting a cyberwarfare weapon," *IEEE Security Privacy*, vol. 9, no. 3, pp. 49–51, 2011.

[30] J. Tang, M. Ibrahim, K. Chakrabarty, and R. Karri, "Security implications of cyberphysical flow-based microfluidic biochips," in *IEEE Asian Test Symp.*, 2017, pp. 115–120.

[31] Y. Luo, K. Chakrabarty, and T.-Y. Ho, "Error recovery in cyberphysical digital microfluidic biochips," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 32, no. 1, pp. 59–72, 2013.

[32] J. Tang, M. Ibrahim, K. Chakrabarty, and R. Karri, "Tamper-resistant pin-constrained digital microfluidic biochips," in *Proc. Design Automation Conf.*, 2018, pp. 1–6.

[33] R. B. Fair, "Digital microfluidics: is a true lab-on-a-chip possible?" *Microfluid Nanofluid*, vol. 3, no. 3, pp. 245–281, 2007.

[34] O. Keszocze, Z. Li, A. Grimmer, R. Wille, K. Chakrabarty, and R. Drechsler, "Exact routing for micro-electrode-dot-array digital microfluidic biochips," in *Proc. ASP Design Automation Conf.*, 2017, pp. 708–713.

[35] S. Chakraborty, C. Das, and S. Chakraborty, "Securing module-less synthesis on cyberphysical digital microfluidic biochips from malicious intrusions," in *Proc. Intl Conf. VLSI Design*, Jan 2018, pp. 467–468.

[36] B. Bhattacharya, S. Roy, and S. Bhattacharjee, "Algorithmic challenges in digital microfluidic biochips: Protocols, design, and test," in *Proc. Intl. Conf. Applied Algorithms*, 2014, pp. 1–16.

[37] S. Bhattacharjee, S. Poddar, S. Roy, J. Huang, and B. B. Bhattacharya, "Dilution and mixing algorithms for flow-based microfluidic biochips," *IEEE Trans. on CAD of Integr Circuits Syst.*, vol. 36, no. 4, pp. 614–627, 2017.

[38] P. Paik, V. K. Pamula, and R. B. Fair, "Rapid droplet mixers for digital microfluidic systems," *Lab Chip*, vol. 3, pp. 253–259, 2003.

[39] D. Grissom, C. Curtis, S. Windh, S. Phung, N. Kumar, Z. Zimmerman, K. ONeal, J. McDaniel, N. Liao, and P. Brisk, "An open-source compiler

and pcb synthesis tool for digital microfluidic biochips," *Integration: The VLSI Journal*, vol. 51, pp. 169–193, 2015.

[40] Z. Li, K. Y. Lai, P.-H. Yu, K. Chakrabarty, M. Pajic, T. Ho, and C. Lee, "Error recovery in a micro-electrode-dot-array digital microfluidic biochip," in *Proc. Intl. Conf. Comput. Aided Design*, 2016, pp. 1–8.

[41] J.-D. Huang, C. Liu, and T. Chiang, "Reactant minimization during sample preparation on digital microfluidic biochips using skewed mixing trees," in *Proc. Intl. Conf. on Comput.-Aided Design*, 2012, pp. 377–383.

[42] K. Choi, A. Ng, R. Fobel, D. Chang-Yen, E. Yarnell, E. Pearson, M. Oleksak, A. Fischer, R. Luoma, J. Robinson, J. Audet, and A. Wheeler, "Automated digital microfluidic platform for magnetic-particle-based immunoassays with optimization by design of experiments," *Analytical Chemistry*, vol. 85, no. 20, pp. 9638–9646, 2013.

[43] Y. Zhao, T. Xu, and K. Chakrabarty, "Integrated control-path design and error recovery in the synthesis of digital microfluidic lab-on-chip," *ACM J. on Emerg. Technol. in Comput. Syst.*, vol. 6, no. 3, pp. 11:1–11:28, 2010.