

A Low-cost Field-Programmable Pin-Constrained Digital Microfluidic Biochip

Daniel T. Grissom, Jeffrey McDaniel and Philip Brisk, *Member, IEEE*

Abstract—This paper introduces a field-programmable pin-constrained digital microfluidic biochip (FPPC-DMFB), which offers general-purpose assay execution at a lower cost than general-purpose direct addressing DMFBs and highly optimized application-specific pin-constrained DMFBs. The key cost driver for is the number of printed circuit board (PCB) layers, onto which the device is mounted; we demonstrate a scalable single-layer PCB wiring scheme for several FPPC-DMFB variations, for PCB technology with orthogonal routing capacity of at least three; for PCB technology with orthogonal capacity of two, more PCB layers are required, but the FPPC-DMFB retains its cost advantage. These results offer new insights on the relationship between PCB layer count, pin count, and cost. Additionally, to reduce the execution time of assays on the FPPC-DMFB, we present efficient algorithms for droplet routing, with and without contamination removal via wash droplets.

Index Terms—Digital Microfluidic Biochip (DMFB), Pin-constrained DMFB, PCB escape routing

I. INTRODUCTION

A field-programmable, pin-constrained digital microfluidic biochip (FPPC-DMFB), can execute any *assay* (biochemical protocol) after the device has been manufactured [12]. Prior general-purpose DMFBs were based on direct addressing, which provides independent control over each electrode in the device, at the cost of a high number of control pins, which requires expensive multi-layer printed circuit boards (PCBs). Prior pin-constrained DMFBs, allow each control pin to drive multiple electrodes [36], reducing the number of control pins and PCB layers; however, they have all been application-specific [9][18][19][20][21][23][24][25][28][36][38][40][41][42], which limits their usability. In contrast, the FPPC-DMFB is general purpose and can be implemented in one PCB layer. We demonstrate that the FPPC-DMFB is cheaper than direct addressing and prior application-specific pin-constrained DMFBs. Our results and analysis provide new understanding into the relationship between pin count, layer count, and the actual cost of the chip. We also demonstrate efficient droplet routing algorithms for the FPPC-DMFB, and establish the overhead of contamination removal via wash droplet routing.

A preliminary version of this article was presented at the IEEE/ACM 2013 Design Automation Conference (DAC), Austin, Texas, June 2013.

Daniel Grissom is with the Department of Computer Science, Azusa Pacific University, Azusa, CA 91702, USA. Jeffrey McDaniel and Philip Brisk are with the Department of Computer Science and Engineering, University of California Riverside, Riverside, CA 92521, USA. (email: dgrissom@apu.edu, jmcd001@ucr.edu, philip@cs.ucr.edu)




Fig. 1. (a) A DMFB is a planar array of electrodes; (b) cross-sectional view; (c) a droplet is transported from CE2 to CE3 by activating CE3, and then deactivating CE2 (white: activated; black: deactivated).




Fig. 2. Fundamental microfluidic operations form instruction set of a DMFB; these operations can be combined to form larger assays.

A. DMFB Technology Overview

1) Background: Physical Droplet Manipulation

DMFBs execute assays by manipulating nanoliter-sized droplets of fluid. DMFBs are based on a phenomenon known as electrowetting [31]. A DMFB, shown in **Fig. 1**, consists of top and bottom plates coated with a hydrophobic layer. The bottom plate contains an array of droplet-sized control electrodes, while the top plate has one conducting electrode that spans the entire array (**Fig. 1(a)**). Each droplet is sandwiched between the two plates and remains in place when its underlying electrode is activated. If a droplet is not centered on an activated electrode, it will unpredictably *drift* across the DMFB in; thus, an electrode underneath a droplet must be activated to store it in place. In **Fig. 1(b)**, a droplet centered on electrode CE2 overlaps neighboring electrodes CE1 and CE3. In **Fig. 1(c)**, activating CE3 pulls the droplet to the right, and deactivating CE2 centers the droplet over CE3.

Fig. 2 depicts the instruction set of a DMFB: droplet transport, splitting, merging, mixing, and storage. Sensor-based detection operations execute by moving a droplet to a detector (placed above an electrode) and storing the droplet in-place. Dispense and output operations are performed by I/O reservoirs on the perimeter of the DMFB.

[Type here]




Fig. 3. A typical microfluidic synthesis flow dictates that a microfluidic assay is represented in the form of a DAG; in Stage 1, its operations are *scheduled* and *placed* onto the DMFB array and droplets are *routed* between operation locations. In Stage 2, *pin-mapping* and *wire routing* are performed to eliminate unused electrodes and connect the electrodes to an external edge of the device to be driven by a microcontroller.




Fig. 4. The interference region (IR) for a droplet at (a) the beginning and (b) end of a droplet-actuation cycle.

2) Background: High-level Assay Synthesis

Synthesis is the process that maps a biochemical reaction (an *assay*), onto a DMFB, as shown in **Fig. 3**. An assay is specified as a directed acyclic graph (DAG), where vertices represent fluidic operations (e.g. mix, split, etc.) and edges represent precedence and droplet transfer between operations. The DAG in **Fig. 3** depicts a simple assay that inputs two droplets, mixes them, and outputs the resulting droplet.

The first step of synthesis is *scheduling*, which assigns start and stop time to each operation (e.g., the mix operation, *M1*, executes from time steps 1 to 4 in **Fig. 3**). Next, the *placement* step decides which DMFB locations perform each operation. For dispense (input) operations, a reservoir containing the appropriate fluid is chosen. “Reconfigurable” operations such as mixing and splitting can be performed anywhere on the chip; for example, in **Fig. 3**, the mix operation *M1* is placed in the 2x2 array of cells in the top-right corner of the DMFB; however, *M1* could be placed in any unoccupied 2x2 array of cells on the DMFB. The array of cells that denotes the location of an operation is called a *module*. Lastly, the *routing* step computes paths for droplet transport between operations that have been scheduled and placed. The router produces a list of electrodes to activate during each *droplet actuation cycle*, i.e., the time it takes to move a droplet from one electrode to the next. While computing droplet routes, the router must ensure that droplets do not interfere with one another while traversing the DMFB array [13][14][35]; to prevent droplets from accidental collisions, the router creates an interference region around a stationary droplet at the beginning of a droplet actuation cycle, as seen in **Fig. 4(a)**. As the droplet moves to an adjacent electrode, the interference region stretches to include all the electrodes surrounding the droplet’s initial and final electrode (**Fig. 4(b)**). Collisions are prevented by ensuring that no droplet enters the interference region of any other droplet (unless they are about to merge).




Fig. 5. Activating a pin on a (a) direct-addressing DMFB activates (white) exactly 1 electrode per pin; (b) pin-constrained DMFB activates 1+ electrodes per pin, depending on the pin layout.

3) Background: Pin Mapping and PCB Wire Routing

Synthesis typically targets an existing DMFB that has been designed and fabricated; alternatively, it can be coupled with pin mapping and wire routing phases to produce an optimized application-specific chip. The default pin mapper instantiates a *direct addressing DMFB*, in which each electrode is wired to an external electrical control pin, as shown in **Fig. 5(a)**. Direct addressable DMFBs offer the most flexibility in terms of droplet coordination; however, the large number of control pins can increase the two-dimensional area, and thus cost, of the PCB on which the DMFB is mounted. (**Fig. 6(a)**).

Pin mapping converts a direct-addressing DMFB into a *pin-constrained* DMFB by tethering multiple electrodes together so that a single control pin can activate them (**Fig. 5(b)**). In **Fig. 3**, the pin-mapper removes the 10 non-used electrodes and connects the remaining ones in to reduce the number of control pins from 15 to 7. Reducing the number of control pins reduces the overall cost of the device [24][36][38][40][42].

PCB wire routing establishes physical connections within the PCB from each external control input to the electrode(s) that it drives, as determined by the pin mapper. In **Fig. 1(b)**, electrodes reside on the lower substrate of the DMFB, while wire routing is performed within the PCB (shown in green in **Fig. 3** and **Fig. 6(b)**). PCB wire routing for pin-constrained DMFBs is a multi-terminal variant of the PCB escape routing problem [7][26][37]. Effective escape routers can reduce the number of PCB layers, and thus the overall cost of the device.

Past work on this topic has focused on the number of control pins and/or the number of PCB layers as a proxy for PCB cost. In this work, we estimate the actual cost of the PCB layout in dollars and discuss the various tradeoffs involved.

[Type here]




Fig. 6. A DMFB has PCBs (green layers) underneath the substrate containing the control electrodes that serve as the medium for wire-routing. A microcontroller sends signals to and interfaces with the DMFB via one or more integrated circuit (IC) clips. (a) A direct addressing DMFB is thought to require multiple layers of PCB, while (b) a pin-constrained DMFB performs is thought to require fewer PCB layers.

B. Contribution

The FPPC-DMFB is the first pin-constrained DMFB that is also general-purpose [12]. This paper introduces the FPPC-DMFB pin assignment scheme, describes a synthesis flow that targets it, and presents a PCB wire routing solution that minimizes its overall cost. Our results establish the feasibility of concurrent assay operations and droplet transport on the FPPC-DMFB, quantify the performance overhead of cross-contamination elimination via wash droplet routing, and show that a PCB designed for an optimized FPPC-DMFB can be significantly cheaper than a PCB designed for both direct-addressing and application-specific pin-constrained DMFBs.

II. RELATED WORK

Early pin mappers tried to minimize the pin count, but did not consider the impact of pin count reduction on the number of PCB layers. Under *array partitioning* [36], different groups of control pins are assigned to each partition, reducing or eliminating interferences among droplets that are transported concurrently. *Broadcast electrode addressing* [42] examines the electrode activation sequence produced by a synthesis tool and identifies electrodes that can share a control input. Luo and Chakrabarty [25] introduced a pin assignment scheme that facilitates interference-free and deadlock-free concurrent transport of up to two droplets. Several other papers optimize pin assignment in conjunction with other synthesis tasks, especially droplet routing [9][18][23][24][28][41].

Escape routing for PCBs routes known pins in a large array to the array perimeter [26][37]. For pin-constrained DMFBs, the escape routing problem must accommodate multi-terminal nets for control inputs that drive multiple electrodes. One paper has been published that focuses explicitly on escape routing for DMFBs [7], while another optimizes the PCB layout for multiple DMFBs that execute the same protocol concurrently in a lock-step [32]. Several papers have also been published that optimize pin assignment in conjunction with escape routing [19][20][21][38]; they optimize application-specific, not than general-purpose pin-constrained DMFBs.

The FPPC-DMFB is a pin-constrained virtual topology [11][13][14], which segregates the DMFB surface area into modules that perform assay operations (mixing, splitting, storage, detection, etc.) and a network of streets that transport droplets between modules and I/O reservoirs.

In a direct-addressing chip, virtual topologies limit the flexibility and reconfigurability of the device in order to facilitate fast online synthesis algorithms, which can respond to sensory feedback provided by the device in real-time [13][14]. Pin-constrained DMFBs exhibit limited flexibility and reconfigurability; imposing a virtual topology to achieve general-purpose operation is a favorable innovation [12].

Chang et al. [8] introduced a pin-constrained DMFB that shares some similarities with the FPPC-DMFB proposed here. Their device does not account for some of the finer details of module/device synchronization and I/O addressed in this paper (e.g., the ability to independently load/unload droplets into modules). It is also unclear if the layout and wiring solution is scalable to larger devices. In contrast, this paper presents a design variation of the FPPC-DMFB which can be routed in one PCB layer, and can scale to arbitrary numbers of operational and storage modules.

III. PIN ASSIGNMENT

The FPPC-DMFB employs a pin assignment scheme that enables all of the basic assay operations (Fig. 2) to execute in a conflict-free manner. Fig. 7 shows two similar, but different, 10×16 FPPC-DMFB layouts. For simplicity we first focus on the pin-optimized version in Fig. 7(a) to show the general characteristics of our FPPC-DMFB design.

Similar to virtual topologies [11][13][14], the FPPC-DMFB reserves specific regions for assay operations and others for routing. The FPPC-DMFB contains a vertical column of mixing modules on the left (blue/orange electrodes, Pins 10-20) and a vertical column of modules on the right (orange electrodes, Pins 31-36) that perform splitting, storage, and detection (which requires an external detector affixed above the module); we call these modules *SSD modules*.

White electrodes define droplet routing regions, which ensure full connectivity between all modules. I/O reservoirs can be placed anywhere along the top or bottom of the chip, as seen in Fig. 7. The green electrodes, Pins 21-30, allow droplets to enter/exit each module. An *interference region* (gray) surrounds each module to isolate droplets within it from droplets in the routing region or adjacent modules; these regions are not functional and do not contain electrodes.

[Type here]




Fig. 7. Pin assignment scheme for a 10×16 FPPC-DMFB which can accommodate 4 mixing modules and 6 split/store/detect (SSD) modules. Mixing pins are shared; the interference region is empty space and does not contain any electrodes. Holding and I/O electrodes are independently wired to single control pins for flexibility and programmability. Fluidic I/O reservoirs are connected to the top/bottom horizontal buses; (a) a pin-optimized FPPC-DMFB version with shared routing pins; (b) a route-optimized FPPC-DMFB version with independent routing pins.





Fig. 8. At least 3 repeatable pins are needed to move a droplet along a straight path without causing the droplet to split. Electrodes with bold borders indicate electrodes being activated next cycle.

The layout is designed for operation concurrency; mix and SSD modules can execute different operations that may start and stop at any time-step (i.e., droplets may enter/exit modules while other modules continue operating). The architecture is scalable and can be vertically lengthened or shortened to produce a DMFB with any desired number of modules.

A. DMFB Operations and Synchronization

1) Droplet Transport

The FPPC-DMFB facilitates droplet transfer between horizontal and vertical transport buses. Routable paths exist between all modules and I/O reservoirs on the chip's perimeter. Chips of arbitrary height can be instantiated without altering the wire-routing pattern (see Section IV.B). The mix and SSD module-hold electrodes (Fig. 7(a), Pins 17-20 and Pins 31-36, respectively) remain active during routing to ensure that droplets within the modules do not drift.



One droplet cannot enter a module while another is being routed or a split will occur
Fig. 9. Multiple droplets moving through the vertical bus will result in an unintentional split when one tries to enter a module.

Fig. 7 presents pin-mappings for pin-optimized and route-optimized designs of the FPPC-DMFB, which share the same topology (white electrodes). The route-optimized design has a higher cost (more control pins), but offers more opportunities to transport multiple droplets concurrently (Section VII).

Pin-optimized FPPC-DMFB: Fig. 8 shows that at least 3 pins are required to successfully transport a droplet along a straight path; this is called a 3-phase transport bus [33]. In Fig. 7(a), Pins 1-3 and Pins 7-9 control two horizontal buses; Pins 4-6 drive a vertical transport bus at the center of the array.

3-phase transport buses cannot hold droplets in the routing area while other droplets enter/exit a module. For example, Fig. 9 shows two droplets in the vertical routing bus: for the lower droplet to enter the lower mixing module, the DMFB must activate Pin 20, while simultaneously deactivating Pin 4, as the electrode underneath a droplet must be activated to hold it in place. Activating Pins 4, 5, or 6 to transport or hold the upper droplet, will inadvertently split the lower droplet. The supplemental section of ref. [12] elaborates on the futility of concurrent droplet routing with the pin-optimized design.

Route-optimized FPPC-DMFB: Replacing 3-phase transport buses with direct addressing buses rectifies the situation in Fig. 9. The route-optimized design increases the control pin count by 25 to facilitate concurrent droplet routing; the pin assignment for mixing and SSD modules is unchanged.

2) Droplet Dispensing and Output

I/O reservoirs are attached to the top and bottom horizontal transport buses, and have individually addressable electrodes, (red in Fig. 7) to allow droplets to enter/exit the chip.

3) Merging/Mixing

Fig. 10(a) shows a droplet (D_2) entering and exiting a mixing module (M_2) without conflicting with droplets in other modules (D_1 , D_3). On top, D_2 reaches the electrode adjacent to the mixing module (M_2); D_1 is in mixing module M_1 and D_3 is stored in SSD module $SSD1$. All SSD module electrodes are activated (Pins 24-26) to hold the stored droplets in place during mixing module I/O. Activating Pin 20 (M_2 's I/O cell) moves D_2 adjacent to M_2 . Activating Pin 16 draws D_2 into M_2 , while transporting D_1 to an adjacent cell within M_1 . Next, all mixer hold cells (Pins 17 and 18) move D_1 and D_2 to identical positions within M_1 and M_2 . The electrode sequence is reversed when a droplet exits a mixing module.

[Type here]




Fig. 10. Pin-activation sequence showing how a single droplet (D_2) can enter/exit (a) mix modules and (b) split/store/detect modules. Sequences are designed to allow a droplet to enter/exit any module without adversely affecting droplets (D_1 , D_3) in other modules.




Fig. 11. The electrode/pin activation sequence (from Cycle 1 to 4) that merges D_4 with D_2 (in M_2) to become D_5 (twice the volume) and re-sync with any other droplets in mix modules (i.e., D_1 in M_1).

Two droplets must merge before mixing, as shown in **Fig. 11**: droplet D_4 merges with droplet D_2 (in M_2) to become a new droplet, D_5 (with twice the volume). Once merged, D_5 is synced with D_1 back to the mixers' hold locations (Cycle 4, **Fig. 11**). Mixing can then begin, presuming that D_1 is merged.




Fig. 12. Pin-activation sequence for droplet splitting and storage using SSD modules: operations sequences are allow droplet D_2 to split and store without adversely affecting droplets D_1 , D_3 in other modules.

M_1 and M_2 perform concurrent synchronized mixing by activating *Pins 10-16*, in sequence, starting with *Pin 15* and continuing counterclockwise (i.e., *Pin 15*, *14*, *10*, *11*, *12*, *13*, *16*), followed by *Pin 17 and 18* together. Mixing can pause if one droplet needs to enter or exit any mixing module.

4) Storage, Detection, and Splitting

SSD modules perform storage and detection (if equipped with an external detector). Both operations require a droplet to enter an SSD module and remain in place. **Fig. 10(b)** illustrates a droplet entering/exiting an SSD module (SSD_3) without affecting droplets in other modules. All SSD hold electrodes are activated, except for SSD_3 's, which allows D_2 to enter. SSD_3 's I/O electrode is then activated, followed by its hold electrode, to complete the entrance. This sequence is reversed to let a droplet exit an SSD module.

Fig. 12(a)-(c) illustrates splitting. Droplet D_2 , which will be split, starts on a vertical transport bus by an SSD module's I/O cell; this cell remains activated throughout the split. Next, the I/O cell is activated, stretching D_2 . Then the SSD module's hold cell is activated and the I/O cell is deactivated, splitting D_2 between the hold cell (D_2), and the bus (D_4). D_4 can then be routed to an available SSD module for storage (**Fig. 12(d)**).

IV. HARDWARE LAYOUT OPTIMIZATION

A. Problem Formulation

This section introduces a deterministic and scalable co-optimized pin assignment and single-layer PCB wire routing solution for the FPPC-DMFB. The input is an architecture description, which includes the XY-dimensions of the chip, locations of I/O ports on the periphery of the horizontal buses, PCB feature sizes and a flag that indicates whether the user wants a pin-optimized or route-optimized FPPC-DMFB. The output is a programmable pin-mapping solution, which fits the user-provided dimensions, and a valid single-layer PCB wire routing solution, which adheres to the pin mapping solution.




Fig. 13. (a) The pin-mapping for a pin-constrained DMFB for a PCR assay [42] detailing a wire-routing solution for *Pin 1*; (b) A complete 4-layer, wire-routing solution (each layer is represented by a different color). NOTE: Gray cells do not contain electrodes.




Fig. 14. The original FPPC-DMFB [12] detailing the (a) pin-mapping and (b) 4-layer wire-routing solution (each color represents a separate layer); (c) layer 2 from (b) illustrates that *Pins 2 and 3* from the horizontal buses and *Pin 4* from the vertical buses prevent other pins from escaping; (d) layer 3 from (b) shows that *Pin 1* from the horizontal bus and *Pin 5* from the vertical bus prevent other pins from escaping.

B. Co-optimizing Pin Assignment and Wire Routing

As a motivating example, consider a pin assignment for a pin-constrained 15x15 assay-specific DMFB designed for the PCR assay [42]. **Fig. 13(a)** shows a 14-pin layout and highlights the wire routing solution for *Pin 1*; *Pin 1* drives 9 electrodes, many of which are on the perimeter of the chip. The wire routing solution for this one pin effectively blocks the ability to route additional wires into the chip on the same PCB layer. **Fig. 13(b)** shows a complete wire routing solution for all 14 pins; a total of four PCB layers are required.




Fig. 15. The single-layer wire-routing solutions for the enhanced FPPC-DMFB designs in this paper for the (a) pin-optimized version (**Fig. 7(a)**) and (b) route-optimized pin-mapping (**Fig. 7(b)**).

Fig. 14 presents two pin mapping and wire routing solutions for two FPPC-DMFB variants. **Fig. 14(a)** presents the original pin mapping architecture [12], and **Fig. 14(b)** shows the wire-routing solution obtained by the negotiated-congestion escape router [26]. This particular variant has three vertical buses (as opposed to the one central vertical bus shown in **Fig. 7**). Four PCB layers are required for routing, as shown in **Fig. 14(b)**.

Fig. 14(c-d) depicts two of these four wire-routing layers. Wires that connect to electrodes on the 3-phase buses must span the entire array, essentially blocking the ability of other wires to escape from the perimeter on the same PCB layer. To eliminate this problem for the pin-optimized version, we removed the two side buses and use separate three-phase buses (*Pins 1-3, 4-6, 7-9*) to control the three remaining buses, as shown in **Fig. 7(a)**. This yielded a single-PCB layer wire routing solution, shown in **Fig. 15(a)**. The route-optimized version also removes the periphery vertical buses, but uses individually addressable electrodes, as seen in **Fig. 7(b)**. **Fig. 15(b)** reveals the single-layer wire-routing solution, showing that the individually addressable bus electrodes are easily escapable and do not create large obstacles for other pins. We have implemented an algorithm to generate these escape routes for FPPC-DMFBs of varying vertical length.

Removing the left and right vertical buses may reduce the number of potential I/O locations; however, if extra I/O is required, the horizontal buses at the top and/or bottom can be extended; alternatively, mixing or SSD modules in the center of the chip could be replaced with an I/O reservoir attached to the central vertical bus. Another subtle detail is that an extra horizontal row is added between the top vertical bus and the topmost mixing and SSD modules; this extra space is needed to provide access for control wires that drive electrodes in the center of the chip to escape, as shown in **Fig. 15**.

The original design assumed that *Pins 7-13* (see **Fig. 14(a)**) could be shared by an arbitrary number of mixing modules, regardless of the height of the chip; however, because of the independently controlled module hold and I/O pins (*Pins 14-21* in **Fig. 14(a)**), there is not enough room to extend the shared pins indefinitely without introducing additional PCB layers to facilitate wire routing to these shared electrodes.




Fig. 16. The wire-routing model for the FPPC-DMFB (pin-optimized version shown) generalizes to an unlimited number of modules; each group of up to four mixing modules shares seven common pins as seen in FPPC-DMFBs with (a) eight mixing modules and (b) 5 mixing modules.

The solution is to limit the number of shared electrodes to groups of four continuous mixing modules (Fig. 15). For chips with more than four mixing modules (e.g., Fig. 16) the same layout and wiring pattern as in Fig. 15 is repeated. Fig. 16(a) shows two groups of four mixing modules, while Fig. 16(b) generalizes the scheme to an arbitrary number. This approach generalizes to the route-optimized design as well.

C. Escape Routing Details

The *orthogonal capacity* of a PCB is the number of wires that can route between two orthogonally adjacent electrodes. We have assumed an orthogonal capacity of 3 throughout this section [21][38]; this allows for a diagonal capacity of 6 (i.e., at most 6 wires can be routed between diagonally adjacent electrodes); ref. [37] provides more details on modeling these capacities. All escape routing results for architectures other than the FPPC-DMFB presented in this paper were obtained using an internally implemented multi-terminal variant of an escape routing algorithm based on negotiated congestion [26].

V. HIGH-LEVEL SYNTHESIS

A. Problem Formulation

The input to the high-level synthesis stage is an architecture description (array dimensions, I/O location, pin assignment), and a DAG representing an assay. After scheduling, placing, and routing the DAG, the output is a valid electrode activation sequence that executes all steps of the assay on the device.




Fig. 17. Split operations are converted to a split and two stores for synthesis.




Fig. 18. Cyclic routing dependencies can be broken by routing a droplet in the cycle to an SSD module, dedicated as a buffer. Arrows indicate that the droplet at the tail end is about to travel to the module at the head end.

B. Scheduling

Schedulers targeting direct addressing DMFBs treat the device as being reconfigurable, where any operation other than I/O or detection can be performed anywhere. When targeting the FPPC-DMFB, the number of mixing and SSD modules impose a resource limit. We modified list scheduling [15][34] and path scheduling [16] to target the FPPC-DMFB.

Modules in direct-addressing DMFBs can perform mixing and storage, and may store multiple droplets. Schedulers targeting direct-addressing DMFBs may route stored droplets from one module to another in order to free up modules to perform other operations [15][16][30]. Since SSD modules store at most one droplet, a scheduler targeting the FPPC-DMFB can eliminate these unnecessary routing transfers.

The scheduler reserves one SSD module to address routing deadlocks, as explained later in Section D. Thus, in Fig. 7, only 5 of the 6 SSD modules are available for general use.

C. Placement/Binding

Similar to placement algorithms targeting direct-addressing virtual topologies [13][15], we reduce placement to a binding problem, which is solved using the left-edge algorithm [22]. Synthesis software targeting the FPPC-DMFB does not bind a split operation to a module, as the split yields two immediate storage operations (Fig. 17). Instead, the software binds the children to the SSD modules directly.

D. Droplet Routing

A routing *sub-problem* refers to the set of droplets that must be routed just before each time-step begins [35][39]. We refer to the routes that transport droplets between operations in the original assay specification as *functional routes*. Routes are computed one-at-a-time, ignoring, for the moment, other droplets that also need to be routed during the same time-step: (1) to route a droplet from an input reservoir to module, the router computes a deterministic path over the horizontal and vertical buses, and applies the appropriate module input sequence when the droplet arrives (Section III.A); (2) a similar approach is taken to route droplets from modules to output reservoirs; (3) module-to-module routing uses the vertical column in the center of the chip, applying appropriate input/output sequences at the start/end of the route.

[Type here]




Fig. 19. A partial DAG in which four droplets must be routed (R1-R4) to begin time-step 10 (dotted arrows convey routes in previous time-steps).

Dependencies and Deadlock: Routing deadlock occurs when one or more droplets wait for resources to become available that will never become free; for example, Fig. 18(a) shows a cyclic dependency involving two droplets. To break the cycle, one droplet ($D3$) is routed to an empty SSD module ($SSD2$), as shown in Fig. 18(b). In Fig. 18(c), the dependency is broken, but droplet $D3$ must wait for $D1$ to complete its route. The scheduler reserves one SSD module as a *routing buffer* to break any cyclic dependencies that may result from binding.

Wash Droplet Routing: A droplet traveling on a DMFB may leave residue behind that contaminates other droplets. Wash droplets may be dispatched to clean electrodes that have been contaminated by a previous droplet; a future droplet can then travel over the cleaned electrode without being contaminated.

In Fig. 19, operations $Mix2$ and $Mix3$ each demand two droplets to be routed to their assigned modules at time-step 10; Fig. 19 depicts these routes as $R1$ - $R4$. Since $R1$ and $R2$ travel to the same destination to be merged, they cannot contaminate one another and no wash droplet is needed; the same is true for $R3$ and $R4$. Assume that $R1$ and $R2$ are routed first; if either $R3$ or $R4$ crosses over any of the same electrodes as $R1$ or $R2$, contamination could occur and must be prevented.

First, we designate dedicated wash fluid input and output ports on the top and bottom horizontal buses respectively. We then identify all operations starting at the given time-step (e.g., for time-step 10 in Fig. 19, $Mix2$ and $Mix3$). The first operation is examined and routes are generated to its assigned module (or output port) from all of its parents; without loss of generality, assume $Mix2$ is being examined, and thus, $R1$ and $R2$ are generated to transport droplets from the input reservoirs to $Mix2$'s mix module. Next, a wash droplet is generated and a route is computed to cover all of the electrodes covered by $R1$ and $R2$ on any of the transportation buses. Since the wash I/O ports are on opposite sides of the array, the entire vertical bus is cleaned each during each washing step. When considering $Mix3$, where routes are being cleaned that begin in a module ($R3$ begins in a SSD module, $R4$ begins in a mix module), the wash droplet enters and cleans each source module on its way down the vertical bus, cleaning all contaminated electrodes in both transportation buses and modules.

Concurrent Routing: As discussed in Section III the pin-optimized FPPC-DMFB must employ sequential routing to prevent unintended droplet splitting and drifting. Thus, wash droplet routes are inserted and executed after each operation executes its parent routes. In contrast, the route-optimized FPPC-DMFB employs individually addressable buses, which facilitates concurrent functional and wash droplet routing.




Fig. 20. A small portion of the pin-optimized design showing that shared module pins can cause (a) a droplet ($D2$) entering a module to be unintentionally split (Cycle 4) if a wash droplet (w) is cleaning another module; (b) droplets exiting ($D1$) and entering ($D2$) separate modules result in an unintended split (NOTE: In the route-optimized design, $D1$ will split while $D2$ will remain intact)

Similar to prior work by Grissom and Brisk [14], the concurrent router *compacts* the sequential routing solution (in time) so that multiple droplets can route in parallel without inadvertently colliding. Their algorithm compacts droplet routes one-at-a-time, delaying the starting time of each route routes until a legal solution is obtained. In the worse case, the compacted solution ends up being sequential.

To incorporate wash droplets and the FPPC-DMFB's virtual topology, several new compaction rules are introduced.

- 1.) No droplet may cross a contaminated electrode.
- 2.) Two mixing modules cannot be cleaned at the same time; a droplet may not enter or exit any mixing module while another is being cleaned.
- 3.) Two (or more) droplets may not simultaneously enter and exit two distinct mixing modules at the same time

If any of these rules is broken during compaction, stalls are inserted at the beginning of the route until legality is obtained.

To address *Rule 1*, a wash droplet must be routed before the current droplet can be compacted. *Rules 2* and *3* are specific to the FPPC-DMFB's shared mixing module pins (see Fig. 7). To illustrate *Rule 2*, Fig. 20(a) shows how a droplet entering a module during mixing may unintentionally split because of the simultaneous pin activations due to wash droplet traversal. Fig. 20(b) illustrates *Rule 3*: two droplets simultaneously entering and exiting mixing modules may inadvertently split.

Once all routes are compacted, we can reduce the length of the routing sub-problem. Typically, operations such as mixing are paused during the routing portion of an assay to allow all droplets to reach their destination and sync [13][14]; however, if all droplets still en route to their destinations are traveling to output reservoirs and are not on a mix electrode (shared, hold or I/O), then the next time-step may commence, overlapping a portion of the routing stage with an operational stage. The mix modules can start to activate their shared pins in a lockstep rotational pattern; these activations could interfere with any droplets concurrently exiting a mixing module, but not with droplets undergoing transportation on one of the buses.

VI. EXPERIMENTAL METHODOLOGY

A. Wire Routing Cost Analysis

Our experiments estimate the cost (in US dollars) of PCBs for different direct-addressing and pin-constrained DMFB architectures, along with several variants of the FPPC-DMFB. Prior work has reported pin-count and the number of PCB layers as a proxy for cost, but has not reported the actual costs. This makes it difficult to determine when it is profitable to increase the pin count if doing so reduces the number of PCB layers. Our experimental approach presents a more accurate picture of the tradeoffs involved here.

We use Advanced Circuits' online instant quote feature to estimate the cost of each PCB [2]. We provide the PCB length and width and the number of wire-routing layers. Vias connect multiple layers, and thus it is necessary to specify the via size, along with the wire trace spacing and size, which dictates the thickness of wire traces and the minimum spacing between tracing. All other metrics are left at their default values. Wire-length does not directly affect the cost, as long as the PCB is routable without increasing its area by adding extra space.

We assume that all DMFBs are driven by an Atmega 1284 microcontroller with 32 general purpose I/Os (GPIOs) to address the DMFB array [3]. If a DMFB has 32 or fewer control pins, no additional circuitry is needed; shift registers are required to drive addition pins in excess of 32. Shift registers can be daisy chained to feed an arbitrary number of additional inputs using 4 microcontroller signals to control the shift register chain: the serial data input (SER), shift register clock input (SCK), storage register clock input (RCK) and the reset input (SCLR). We assume the Fairchild 74VHC595MTC 8-bit shift register [4], which can be purchased from Mouser for \$0.14 per unit in quantities of 2,500 [6]. **Equation (1)** reports the number of shift registers required for a DMFB.

$$numShiftRegs = \begin{cases} \left\lceil \frac{numPins-28}{8} \right\rceil, & numPins > 32 \\ 0, & otherwise \end{cases} \quad (1)$$

The Atmega 1284 operates at 20MHz [3]; the *droplet actuation frequency* (i.e., the time it required to transport a droplet between two adjacent electrodes) of a typical DMFB is 100Hz [39]. With an estimate of 5 cycles per shift operation, the Atmega 1284 can load 400 pin values into the shift registers in 1% of the droplet actuation cycle, maintaining sufficient signal integrity for proper droplet transportation.

Equation (2) models the PCB wire routing cost as the sum of the PCB price estimate from Advanced Circuits and the shift registers to connect all pins. We do not include the cost of the microcontroller, which is constant, irrespective of the choice of DMFB. We do not consider the cost of circuitry to amplify the voltage produced by the microcontroller to levels appropriate to drive the DMFB. Typical actuation voltages are in the 50-70V range [29][31]; low voltage devices that operate at ~15V have also been reported [10][27].

$$Cost_{WR} = Cost_{PCB} + Cost_{SR} \quad (2)$$

The wire-routing cost is a function of the number of PCB layers, the PCB area, and wire trace width, as shown in **Equation (3)**. Assuming an orthogonal capacity of 3 or more, the FPPC-DMFB requires one PCB layer, as shown in **Fig. 14(f)** and **Fig. 16**. In all other cases, we use the negotiated congestion escape router [26] to determine the number of PCB layers that achieve a legal route. Using larger feature sizes tends to reduce the PCB cost estimate [2].

$$Cost_{PCBf}(numLayers, width_{PCB}, height_{PCB}, width_T) \quad (3)$$

Equation (4) computes the orthogonal capacity according to the metrics in **Table I** and diagram in **Fig. 21**. With an electrode pitch of 2 mm, wire trace width and spacing of 0.007 in, via width of 0.014 in and via contact width of 0.024 in, we estimate the orthogonal capacity to be $o_{cap} = \lfloor 3.12 \rfloor = 3$.

TABLE I
PCB FABRICATION PARAMETERS

Feature	Symbol
Electrode Pitch	$width_{ELEC}$
Via (Hole) Width	$width_V$
Via (Hole) Contact Width	$width_{VC}$
Wire Trace Width	$width_T$
Min. Space Between Wire Trace	$width_{TS}$
Shift Register Width [4]	$width_{SR}$
Shift Register Height [4]	$height_{SR}$
Spacing Between Shift Registers	$width_{SRS}$




Fig. 21. A top-down and cross-sectional view of a PCB showing dimensions for the electrode pitch (ELEC), via hole (V), via contact (VC), wire trace (T) and minimum wire trace spacing (VS) in a DMFB.




Fig. 22. The component layout for PCB size estimation. The electrode array is surrounded by a 0.5 inch perimeter of empty space. The PCB width is extended to add as many shift registers as necessary.

[Type here]

TABLE II
DMFB BENCHMARK DESCRIPTION AND NUMBER OF LAYERS PER
ORTHOGONAL CAPACITY

DMFB Characteristics						Number of Layers Per Orthogonal Capacity									
Name	WR Alg.	Array Dims.		# Elecs.	# Pins										
		X	Y			2	3	4	5	6	7	8	9	10	
ZHAO_PCR	[26]	15	15	62	14	4	4	4	4	4	5	4	4	4	4
ZHAO_INVITRO	[26]	15	15	59	25	4	4	3	4	4	4	4	4	4	4
ZHAO_PROTEIN	[26]	15	15	54	27	4	3	4	4	4	4	4	4	4	4
ZHAO_MULTI	[26]	15	15	81	32	5	5	6	6	6	5	4	5	5	5
LUO_PCR	[26]	15	15	62	22	5	5	5	5	5	5	5	5	5	5
LUO_INVITRO	[26]	15	15	59	21	5	5	5	5	5	5	5	5	5	5
LUO_PROTEIN	[26]	15	15	54	21	4	4	4	4	4	5	4	4	4	4
LUO_MULTI	[26]	15	15	81	27	6	6	6	6	6	6	6	6	6	7
PCR_DA	[26]	15	15	62	62	1	1	1	1	1	1	1	1	1	1
INVITRO_DA	[26]	15	15	59	59	1	1	1	1	1	1	1	1	1	1
PROTEIN_DA	[26]	15	15	54	54	1	1	1	1	1	1	1	1	1	1
MULTI_DA	[26]	15	15	81	81	1	1	1	1	1	1	1	1	1	1
FPPC_4_MODULE	FPPC	10	16	82	36	2	1	1	1	1	1	1	1	1	1
FPPC_8_MODULE	FPPC	10	30	146	65	2	1	1	1	1	1	1	1	1	1
FPPC_4_DA_BUS	FPPC	10	16	82	61	2	1	1	1	1	1	1	1	1	1
FPPC_8_DA_BUS	FPPC	10	30	146	104	2	1	1	1	1	1	1	1	1	1
FPPC_4_DA	[26]	10	16	82	82	1	1	1	1	1	1	1	1	1	1
FPPC_8_DA	[26]	10	30	146	146	1	1	1	1	1	1	1	1	1	1
15x15_DA	[26]	15	15	225	225	3	2	2	2	1	1	1	1	1	1
10x16_DA	[26]	10	16	160	160	2	2	1	1	1	1	1	1	1	1
10x30_DA	[26]	10	30	300	300	2	2	1	1	1	1	1	1	1	1

The left side gives a description of the 21 DMFB wire-routing benchmarks showing the wire-routing algorithm used (WR Alg.), array dimensions (Array Dims.), number of electrodes (# Elecs.) and number of control pins (# Pins). The right side shows the number of PCB layers yielded for each DMFB as a function of orthogonal capacity, which varies from 2 to 10; dark squares highlight the lowest orthogonal capacity which achieves the minimum number of PCB layers for each chip.

$$O_{cap} = \left\lfloor \frac{width_{ELEC} - width_{VC} - width_{TS}}{width_T + width_{TS}} \right\rfloor \quad (4)$$

Fig. 22 shows the layout for estimating the PCB dimensions ($width_{PCB} \times height_{PCB}$). In **Equation (5)**, the PCB height is the array’s height plus one inch; **Equation (6)** shows that the PCB’s width may need to account for space allocated for shift registers. The amount of extra space added to the PCB width depends on the number of shift registers, as computed by **Equation (7)**. Shift registers are stacked vertically until there is no more room, at which point additional columns are added.

$$height_{PCB} = height_A + 1 \quad (5)$$

$$width_{PCB} = width_A + width_{PCB_SR} + 1 \quad (6)$$

$$width_{PCB_SR} = \left\lceil \frac{numShiftRegs}{\lceil height_{PCB} / (height_{SR} + width_{SR}) \rceil} \right\rceil \times (width_{SR} + width_{SR}) \quad (7)$$

VII. EXPERIMENTAL RESULTS

We implemented the FPPC-DMFB and associated synthesis algorithms in a publicly available open-source software framework for DMFB compilation [17]. All experiments were performed using a 2.8GHz Intel Core i7 CPU and 4GB RAM running a 64-bit version of Windows 7.

A. Benchmarks

We extracted eight pin-constrained DMFB layouts from refs. [25] and [42]; these benchmarks are labeled *ZHAO_XXX* [42] and *LUO_XXX* [25], where “XXX” is one of the three

assays used in their experiments (*PCR*, *INVITRO*, *PROTEIN*) or a multi-functional chip that is co-designed to perform all three of those assays (*MULTI*). The electrode layouts are identical for both of these works; only the pin-assignment is different. We also created directly addressable versions of the corresponding electrode layouts, entitled *XXX_DA*.

We denote the FPPC as *FPPC_4_MODULE* for the 4-mixer version shown in **Fig. 14(e)** and *FPPC_8_MODULE* for the 8-mixer version shown in **Fig. 16(a)**. The routing-optimized FPPC-DMFB (**Fig. 7(b)**,) with 4- and 8-mixers are named *FPPC_4_DA_BUS* and *FPPC_8_DA_BUS* respectively.

We also consider two direct addressing DMFBs that use the 4- and 8-mixer FPPC-DMFB electrode layouts in **Fig. 14(e)** and **Fig. 16(a)**; these chips are named *FPPC_4_DA* and *FPPC_8_DA*, respectively. Lastly, we include results for three direct-addressing DMFBs having dimensions: of 15x15 (*15x15_DA*), 10x10 (*10x16_DA*), and 10x30 (*10x30_DA*).

PCB escape routing solutions for all FPPC-DMFBs, except for the directly addressable ones, were computed as described in **Section IV.B** (*FPPC*); all other DMFB wire routes were computed using a multi-terminal implementation of an escape router based on negotiated congestion [26]. If a pin’s wire net cannot be routed on a top-level layer, it is routed on a lower layer in its entirety using vias to connect the wire net to its corresponding electrodes and to the external driving pin.

B. PCB Layers & Orthogonal Capacity

We examine the relationship between orthogonal capacity, the number of control pins and the number of PCB layers. The right side of **Table II** shows the number of PCB layers required to route each DMFB as the orthogonal capacity varies from 2 to 10. The dark squares highlight the lowest orthogonal capacity that achieves the smallest number of layers for each benchmark: $16/21 \approx 76\%$ and $19/21 \approx 90\%$ of the benchmarks minimize the number of layers at orthogonal capacities of 3 or 4. For *ZHAO_XXX* and *LUO_XXX*, the minimum number of PCB layers ranges from 3 to 6.

As an example, consider *ZHAO_PCR*, which is shown in **Fig. 13**. This particular pin mapping requires long wires along the perimeter of the chip, which prevent other wires from escaping on the same PCB layer; thus, additional layers are necessary. Reducing the wire size to increase orthogonal capacity does not reduce the number of PCB layers; however, converting this chip to a direct-addressing design (*XXX_DA*) yields single-layer escape routing solutions for any orthogonal capacity, but increases the control pin count by 2x to 4.4x.

The FPPC-DMFBs shown in **Table II** can be routed in 1 or 2 PCB layers, even at the smallest orthogonal capacities. *FPPC_4_MODULE*’s pin count is comparable to *ZHAO_XXX* and *LUO_XXX*, indicating that the FPPC-DMFB offers the advantages of being both general-purpose and low cost.

C. Wire Routing Cost Analysis

We vary the electrode pitch (1, 2, 2.54 mm), and determine the parameters for each that minimize cost. We translate these results into PCB cost estimates for the DMFBs in **Table II**.

[Type here]

TABLE III
PRICE ESTIMATIONS FOR VARYING NUMBERS OF LAYERS AND PARAMETERS OF A 2"×2" PCB

Electrode Pitch		Advanced Circuit Metrics			oCap	2" × 2" Price (@ 2,500 QTY) with Varying Number of Layers				
mm	in	Trace Size/Space	Via Size	Via Contact Size		1	2	3	4	5
1	0.0394	0.005	0.010	0.013	2	\$1.20	\$1.20	\$1.73	\$1.81	\$2.09
		0.004	0.008	0.011	3	\$1.88	\$1.88	\$2.43	\$2.52	\$2.81
		0.003	0.009	0.012	4	N/A	N/A	N/A	N/A	N/A
		0.002	0.008	0.011	5	N/A	N/A	N/A	N/A	N/A
		0.008	0.028	0.038	2	\$0.99	\$0.99	\$1.41	\$1.46	\$1.79
2	0.0787	0.006	0.012	0.024	4	\$0.99	\$0.99	\$1.46	\$1.53	\$1.80
		0.005	0.013	0.023	5	\$1.10	\$1.10	\$1.63	\$1.71	\$1.99
		0.004	0.008	0.026	6	\$1.88	\$1.88	\$2.43	\$2.52	\$2.81
		0.012	0.030	0.040	2	\$0.99	\$0.99	\$1.41	\$1.41	\$1.79
		0.010	0.020	0.030	3	\$0.99	\$0.99	\$1.41	\$1.41	\$1.79
2.54	0.1000	0.008	0.018	0.028	4	\$0.99	\$0.99	\$1.41	\$1.46	\$1.79
		0.006	0.024	0.034	5	\$0.99	\$0.99	\$1.46	\$1.53	\$1.80
		0.006	0.012	0.022	6	\$0.99	\$0.99	\$1.46	\$1.53	\$1.80
		0.005	0.010	0.020	7	\$1.20	\$1.20	\$1.73	\$1.81	\$2.09
		0.004	0.022	0.032	8	\$1.38	\$1.38	\$1.93	\$2.02	\$2.31

The left side shows various metrics used for the Advanced Circuit PCB cost estimator [2] and resultant orthogonal capacity (oCap). The right side provides corresponding price estimates from the Advanced Circuit PCB cost estimator for a 2"×2" PCB with varying numbers of layers; dark rows represent the selected metrics for each electrode pitch.

TABLE IV
COST ESTIMATES FOR 21 DMFBs WITH 1 MM ELECTRODE PITCH AND ORTHOGONAL CAPACITY OF 2

DMFB Details					Cost (\$)			
DMFB Name	# Pins	# SR	Adjusted PCB Dim		# Layers	Board	SR	Total
			X (in)	Y (in)				
FPPC_4_MODULE	36	1	1.7638	1.6299	2	\$1.01	\$0.14	\$1.15
ZHAO_PCR	14	0	1.5906	1.5906	4	\$1.33	\$0.00	\$1.33
ZHAO_INVITRO	25	0	1.5906	1.5906	4	\$1.33	\$0.00	\$1.33
ZHAO_PROTEIN	27	0	1.5906	1.5906	4	\$1.33	\$0.00	\$1.33
LUO_PROTEIN	21	0	1.5906	1.5906	4	\$1.33	\$0.00	\$1.33
ZHAO_MULTI	32	0	1.5906	1.5906	5	\$1.55	\$0.00	\$1.55
LUO_PCR	22	0	1.5906	1.5906	5	\$1.55	\$0.00	\$1.55
LUO_INVITRO	21	0	1.5906	1.5906	5	\$1.55	\$0.00	\$1.55
LUO_MULTI	27	0	1.5906	1.5906	6	\$1.58	\$0.00	\$1.58
PROTEIN_DA	54	4	1.9606	1.5906	1	\$1.08	\$0.56	\$1.64
INVITRO_DA	59	4	1.9606	1.5906	1	\$1.08	\$0.56	\$1.64
FPPC_4_DA_BUS	61	5	1.7638	1.6299	2	\$1.01	\$0.70	\$1.71
PCR_DA	62	5	1.9606	1.5906	1	\$1.08	\$0.70	\$1.78
FPPC_8_MODULE	65	5	1.7638	2.1811	2	\$1.16	\$0.70	\$1.86
MULTI_DA	81	7	2.3307	1.5906	1	\$1.12	\$0.97	\$2.09
FPPC_4_DA	82	7	2.1339	1.6299	1	\$1.19	\$0.97	\$2.16
FPPC_8_DA_BUS	104	10	2.1339	2.1811	2	\$1.37	\$1.39	\$2.76
FPPC_8_DA	146	15	2.5039	2.1811	1	\$1.57	\$2.09	\$3.66
10x16_DA	160	17	2.8740	1.6299	2	\$1.37	\$2.36	\$3.73
15x15_DA	225	25	3.4409	1.5906	3	\$2.52	\$3.48	\$6.00
10x30_DA	300	34	3.6142	2.1811	2	\$2.94	\$4.73	\$7.67

Wiring costs for the 21 benchmarks, sorted in increasing order, showing the adjusted PCB dimensions (after adding room for shift registers (SR)), number of layers and resultant breakdown of costs (PCB and shift registers).

Metric Selection: The predominant factors that impact PCB costs are the wire-trace width/spacing ($width_T/width_{TS}$), smallest via size ($width_V$), PCB dimensions ($width_{PCB}$, $height_{PCB}$) and number of layers [2]. **Table II** shows that the number of layers depends on the orthogonal capacity, which depends on the trace and via metrics, as per **Equation (4)**.

Table III presents parameter combinations that yield a range of orthogonal capacities for 1 mm, 2 mm and 2.54 mm electrodes. Advanced Circuits' PCB cost estimator provides




Fig. 23. Two-layer wire-routing solution for *FPPC_4_MODULE* with an orthogonal capacity of 2; (a) Layer 1; (b) Layer 2.

the following wire trace sizes (all in inches): 0.0025, 0.003, 0.004, 0.005, 0.006, 0.007, 0.008, 0.010 and 0.012. To achieve each orthogonal capacity, we chose the lowest trace size and highest minimum via size and via contact size such that $width_V \geq 2 \times width_T$ and $width_{VC} \geq width_V + 0.01$; these requirements were relaxed for the smaller 1 mm electrodes (such that $width_{VC} \geq width_V + 0.003$) because there was not enough space to utilize such conservative size estimates.

For each electrode pitch, we select one set of parameters and report the orthogonal capacity. For 1 mm electrodes, estimations were not available for orthogonal capacities of 4 and 5, and we could not find parameters that could achieve an orthogonal capacity of 6; we conservatively opted for an orthogonal capacity of 2: even as the number of PCB layers increases to 4, the cost is less than a single-layer PCB with orthogonal capacity of 3. **Fig. 23** shows a two-layer solution for *FPPC_4_MODULE* with orthogonal capacity of 2.

For 2 mm and 2.54 mm electrode pitches, we select the metric set corresponding to orthogonal capacities of 4 and 6, respectively. These are the highest capacities before the price increases significantly for each electrode size and reductions in layer count do not typically occur at higher orthogonal capacities, as shown in **Table II**.

Table III reveals that, in general, as the feature sizes decrease, particularly the wire trace size, the fabrication costs increase, and that 1- and 2-layer solutions are identical for all cases; this is because PCBs can be printed on two sides. **Table III** reflects this observation, as prices tend to jump more significantly as each odd numbered layer is added, meaning that a new, physical dual-sided PCB layer must be added.

D. PCB Cost Results

Table IV presents cost results for each DMFB architecture in **Table II** assuming 1 mm electrode sizes and an orthogonal capacity of 2, as discussed in the preceding subsection. **Table IV** reports the number of pins and the subsequent number of required shift registers (SR), along with the PCB dimensions, which include extra space allocated for shift registers. The PCB dimensions are fed directly into the online cost estimator. The reported cost of the PCB is estimated under the

[Type here]




Fig. 24. Total wire-routing fabrication costs per DMFB board, including PCB manufacturing and shift register costs, for 21 DMFB designs utilizing 1 mm, 2 mm and 2.54 mm electrodes. The benchmarks are sorted, from left to right, in order of increasing average cost.

assumption of a 4-week delivery time (slowest) at a quantity of 2500 [2]; the shift registers cost \$0.139 apiece at a quantity of 2500 [6]. As described in **Equation (2)**, the total cost is that of the PCB plus shift registers.

In **Table IV** the DMFB chips are sorted in increasing order of cost, and the results show that *FPPC_4_MODULE* is the cheapest design by at least \$0.18 per board, compared to assay-specific chips *ZHAO_XXX* and *LUO_XXX*. The *FPPC_4_DA_BUS* design, which can transport multiple droplets concurrently because the 3-phase buses are replaced by direct-addressing buses, is \$0.56 more expensive than its pin-optimized counterpart. Given that droplet transport times are orders of magnitudes faster than operation times [34][39] and that the FPPC-DMFB with serial routing has proven competitive to pin-constrained assay-specific chips in terms of performance [12], we believe that *FPPC_4_MODULE* is the best overall solution when considering price and performance.

Table IV also shows that the PCB cost for many directly addressable DMFBs are less expensive than pin-constrained counterparts. For example, even though *PCR_DA* is 0.4 in longer than *LUO_PCR*, the PCB is \$0.47 cheaper because it requires fewer layers; however, the directly addressable version requires 5 shift registers, increasing the cost of the design to be \$0.23 more expensive than *LUO_PCR*.

This overall trend of shift-register costs is shown in the bottom half of **Table IV**; starting with *PROTEIN_DA* and looking downward, the number of pins, shift registers and total cost are non-decreasing. **Thus, it is important to minimize the pin-count; however, it must be balanced with a reduction in the number of PCB layers, as seen with our FPPC-DMFB, to achieve a comparably low-cost wire-routing solution.**

Fig. 24 presents the final cost estimates for each of the 21 DMFBs with 1 mm, 2 mm and 2.54 mm electrodes, using the parameters reported in **Table III**, and the shift register costs described previously. All estimates assumed for a 4-week delivery time and shipment quantity of 1000 because price estimates of 2500 could not be obtained online for all benchmarks. The benchmarks are sorted, from left to right, in order of increasing average cost. Although there is some variation when comparing the cost between different electrode sizes (e.g., from *MULTI_DA* to *LUO_INVITRO*, the 1 mm and 2 mm costs decrease, while the 2.54 mm cost increases), the

overall trend of increasing price follows the average for each electrode pitch. This indicates that electrode pitch minimally impacts PCB cost since the manufacturing parameters can be adjusted to compensate the orthogonal capacity to reduce the number of layers.

For several of the benchmarks, the 1 mm boards are more expensive than their 2 mm counterparts. Recall that shift registers are stacked vertically and that additional width is added in the X-dimension to accommodate new columns, as necessary. Since the 2 mm instances provide a greater PCB height ($height_{PCB}$) than the 1 mm instances, the 2 mm boards sometimes require less additional space ($width_{PCB,SR}$) to accommodate the extra shift registers, which indirectly offsets the initial area cost of using larger electrodes. This occurs, for example, for *FPPC_4_DA* and *FPPC_8_DA*. The extra cost for *15x15_DA* is due to an additional layer (3 vs. 2 layers) for the smaller 1 mm case, due to its lower orthogonal capacity.

E. Routing Performance

The top portion of **Table V** reports the time spent performing operations and routing on 9 common benchmark assays [12][14][16][34] for sequential (pin-optimized) and parallel (route-optimized) FPPC-DMFBs, with and without wash droplets; identical schedules were used for both chips. The width of the chip is increased to 11 electrodes, to make room for the two wash I/O reservoirs. All but one benchmark was performed with the standard 16-electrode height, as shown in **Fig. 7**, allocating 4 mix modules, 5 SSD modules and 1 routing buffer; the *ProteinSplit3* benchmark requires a larger 11x18 array with an extra SSD module. The results show that the parallel routing version (without wash droplets) yields a modest 1-10% savings (up to 4.29s on *ProteinSplit3*); however, when wash droplets are introduced, parallel routing on the route-optimized FPPC-DMFB reduces the overall completion time by 6-28% (up to nearly 18s on *ProteinSplit3*).

Using the preferred metrics for 2mm electrodes in **Table III**, the price for an 11x16 pin-optimized (sequential) and route-optimized (parallel) FPPC-DMFB are \$1.44 and \$2.00, respectively. Thus, as expected, the modest gains in performance for the route-optimized FPPC-DMFB come at the expense of cost. Ultimately, the end user should decide which tradeoff to take when considering their desired application and

TABLE V
ASSAY SCHEDULING/ROUTING RESULTS (SECONDS) FOR DIRECT ADDRESSING DMFB AND FPPC SEQUENTIAL/PARALLEL ROUTING (FPPC)

DMFB	Metric	Assay Benchmarks									
		PCR	InVitro1	InVitro2	InVitro3	InVitro4	InVitro5	ProteinSplit1	ProteinSplit2	ProteinSplit3	
FPPC	DMFB Size		11×16	11×16	11×16	11×16	11×16	11×16	11×16	11×16	11×18
	Schedule Length (s)		11	14	18	19	23	29	52	68	123
	No Wash	Sequential Routing (s)	1.52	2.32	3.31	4.89	6.44	8.66	2.17	5.05	10.82
		Parallel Routing (s)	0.89	1.53	2.12	2.81	3.73	5.04	1.48	3.24	6.53
		Overall Assay Savings	5%	6%	9%	9%	10%	1%	2%	3%	5%
	Wash	Sequential Routing (s)	3.88	5.88	8.53	12.66	16.68	22.42	5.99	13.53	29.74
		Parallel Routing (s)	2.17	2.07	3.08	4.85	5.88	8.07	2.28	5.63	11.83
Overall Assay Savings		19%	21%	25%	27%	28%	6%	10%	12%	11%	
Direct Addressing [14]	DMFB Size		15×11	15×11	15×11	15×11	15×11	15×11	15×11	15×11	15×11
	Schedule Length (s)		11	14	20	24	30	44	54	102	164
	No Wash	Parallel Routing (s)	0.44	0.64	1.12	1.4	1.95	2.43	1.22	2.73	5.28
		FPPC Route Savings	-102%	-139%	-89%	-101%	-91%	-107%	-21%	-19%	-24%
		FPPC Overall Savings	-4%	-6%	5%	14%	16%	27%	3%	32%	23%

Performance results showing schedule/route length in seconds. The top half shows parallel/sequential routing times with/without wash droplets and the savings gained (overall assay completion time) by using parallel routing in the FPPC-DMFB. The bottom gives results for a state-of-the-art direct addressing synthesis flow (no wash droplets) and details the savings obtained (overall assay completion time) when using the parallel FPPC-DMFB (no wash droplets).

whether cost or performance is more important.

The bottom portion of **Table V** gives results for a state-of-the-art direct addressing synthesis flow [14]. We use a similarly sized array (11×16 vs. 15×11), although the direct addressing device is wider since it is a better fit for its virtual topology. The virtual topology configuration of the direct addressing model dictates that a 15×11 DMFB can allocate four general-purpose modules, which can process mixes, splits or two simultaneous storage operations [14]. The results show that the direct addressing design is able to obtain better routing results because of its flexibility and extra electrodes; however, although the route saving percentages are high, the actual route savings do not grow larger than 2.5s. In contrast, as seen from the schedule lengths, the FPPC-DMFB is able to produce significantly better schedules (savings almost 40s with *ProteinSplit3*). Thus, the overall savings show that the FPPC-DMFB produces superior performance results as the assay size increases. Finally, using the preferred metrics for 2mm electrodes in **Table III**, the price for a 15×11 direct addressing DMFB is \$4.22, which is \$2.22 to \$2.78 more expensive than the 11×16 FPPC-DMFBs that perform the same assays.

VIII. CONCLUSION

This paper has extended the initial development of FPPC-DMFBs [12] with an enhanced design to facilitate more efficient wire routing, and has presented the first cost estimates, in terms of US dollars, for PCB fabrication for DMFBs. A complete synthesis flow, which addresses architectural issues that are specific to the FPPC-DMFB, has been presented, along with a detailed description of its general-purpose, as opposed to assay-specific, capabilities.

Our prior work has shown that the general-purpose FPPC-DMFB is competitive with direct-addressing and prior pin-constrained assay-specific DMFBs in terms of performance [12]; this work goes one step further by showing that the FPPC-DMFB is less expensive than previous pin-constrained designs, when optimized for cost and PCB design. Thus, the flexibility provided by the FPPC-DMFB is unmatched by prior pin-constrained DMFBs, which were optimized for

specific assays, and offers a significant advancement in terms of programmability at a lower overall per-unit cost.

IX. ACKNOWLEDGMENT

This work was supported in part by NSF Grant CNS-1035603. Daniel Grissom was supported by an NSF Graduate Research Fellowship. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the NSF.

REFERENCES

- [1] HEAD
- [2] www.4pcb.com
- [3] www.atmel.com/devices/atmega1284.aspx
- [4] www.fairchildsemi.com
- [5] <http://microfluidics.cs.ucr.edu>
- [6] www.mouser.com
- [7] J-W. Chang, S-H. Yeh, T-W. Huang and T-Y. Ho. An ILP-based routing algorithm for pin-constrained EWOD chips with obstacle avoidance. *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst*, 32(11), Nov. 2013.
- [8] J-W. Chang, S-H. Yeh, T-W. Huang and T-Y. Ho. Integrated fluidic-chip co-design methodology for digital microfluidic biochips. *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst*, 32(2), Feb. 2013.
- [9] S. Chatterjee, H. Rahaman, and T. Samanta. Multi-objective algorithm for efficient pin-constrained droplet routing technique in digital microfluidic biochip. In *Proc. of ISQED*, pages 252-256, Santa Clara, CA, USA, March 4-6, 2013.
- [10] M. Dhindsa, S. Kuiper, and J. Heikenfeld. Reliable and low-voltage electrowetting on thin parylene films. *Thin Solid Films*, 519, 2011.
- [11] E. Griffith and S. Akella. Performance characterization of a reconfigurable planar-array digital microfluidic system. *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst*, 25(2), Feb. 2006.
- [12] D. Grissom and P. Brisk. A field-programmable pin-constrained digital microfluidic biochip. In *Proc. of DAC*, article 46, Austin, TX, USA, 2013.
- [13] D. Grissom and P. Brisk. A high-performance online assay interpreter for digital microfluidic biochips. In *Proc. of GLSVLSI*, pages 103-106, Salt Lake City, UT, USA, May 3-4, 2012.
- [14] D. Grissom and P. Brisk. Fast online synthesis of digital microfluidic biochips. *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst*, 33(3), Mar. 2014.
- [15] D. Grissom and P. Brisk. Fast online synthesis of generally programmable digital microfluidic biochips. In *Proc. of CODES+ISSS*, pages 413-422, Tampere, Finland, Oct. 7-12, 2012.

[Type here]

- [16] D. Grissom and P. Brisk. Path scheduling on digital microfluidic biochips. In Proc. of DAC, pages 26-35, San Francisco, CA, USA, Jun. 3-7, 2012.
- [17] D. Grissom, et al. A digital microfluidic biochip synthesis framework. In Proc. of VLSI-SoC, pages 177-182, Santa Cruz, CA, Oct. 7-10, 2012.
- [18] T-W. Huang and T-Y. Ho. A two-stage integer linear programming-based droplet routing algorithm for pin-constrained digital microfluidic biochips. IEEE Trans. Comput.-Aided Design Integr. Circuits Syst., 30(2), Feb. 2011.
- [19] T-W. Huang, T-Y. Ho, and K. Chakrabarty. Reliability-oriented broadcast electrode-addressing for pin-constrained digital microfluidic biochips. In Proc. ICCAD, pages 448-455, San Jose, CA, USA, Nov. 7-11, 2012.
- [20] T-W. Huang, H-Y. Su, and T-Y. Ho. Progressive network-flow based power-aware broadcast addressing for pin-constrained digital microfluidic biochips. In Proc. of DAC, pages 741-746, San Diego, CA, USA, Jun. 5-10, 2011.
- [21] T-W. Huang, S-Y. Yeh and T-Y. Ho. A network-flow based pin-count aware routing algorithm for broadcast-addressing EWOD chips. IEEE Trans. Comput.-Aided Design Integr. Circuits Syst., 30(12), Dec. 2011.
- [22] F. J. Kurdahi and A. C. Parker, "REAL: a program for REGISTER ALlocation. In Proc. of DAC, pages 210-215, Miami, FL, USA, 1987.
- [23] C. C-Y. Lin and Y-W. Chang. Cross-contamination aware design methodology for pin-constrained digital microfluidic biochips. IEEE Trans. Comput.-Aided Design Integr. Circuits Syst., 30(6), Jun. 2006.
- [24] C. C-Y. Lin and Y-W. Chang. ILP-based pin-count aware design methodology for microfluidic biochips. IEEE Trans. Comp.-Aided Design Integr. Circuits Syst., 29(9), Sep. 2010.
- [25] Y. Luo and K. Chakrabarty. Design of pin-constrained general-purpose digital microfluidic biochips. IEEE Trans. Comp.-Aided Design Integr. Circuits Syst., 32(9), Sep. 2013.
- [26] Q. Ma, T. Yan and M.D.F. Wong. A negotiated congestion based router for simultaneous escape routing. In Proc. of ISQED, San Jose, CA, USA, Mar. 22-24, 2010.
- [27] H. Moon, S. K. Cho, R. L. Garrell, and C-J. Kim. Low voltage electrowetting on dielectric. Journal of Applied Physics 92(7), Sep. 2002.
- [28] R. Mukherjee, et al.. A heuristic method for co-optimization of pin assignment and droplet routing in digital microfluidic biochip. In Proc. of VLSI Design, pages 227-232, Hyderabad, India, Jan. 7-11, 2012.
- [29] J. H. Noh, J. Noh, E. Kreit, J. Heikenfeld, and P. D. Rack. Toward active-matrix lab-on-a-chip: programmable electrofluidic control enabled by arrayed oxide thin film transistors. Lab-on-a-Chip 12(2), Dec. 2011.
- [30] K. O'Neal, D. Grissom, and P. Brisk. Force-directed list scheduling for digital microfluidic biochips. In Proc. of VLSI-SoC, pages 7-11, Santa Cruz, CA, USA, Oct. 7-10, 2012.
- [31] M. G. Pollack, A.D. Shenderov, and R. B. Fair. Electrowetting-based actuation of droplets for integrated microfluidics. Lab on a Chip, 2:96-101, 2002.
- [32] S. Roy, D. Mitra, B. B. Bhattacharya, K. Chakrabarty. Congestion-aware layout design for high-throughput digital microfluidic biochips. ACM Journal on Emerging Technologies in Computing Systems, 8(3): article #17, Aug. 2012.
- [33] V. Srinivasan, V. Pamula, and R. Fair. An integrated digital microfluidic lab-on-a-chip for clinical diagnostics on human physiological fluids. Lab on a Chip, 4:310-315, 2004.
- [34] F. Su and K. Chakrabarty. High-level synthesis of digital microfluidic biochips. ACM Journal on Emerging Technologies in Computing Systems, 3(4): article #16, Jan. 2008.
- [35] F. Su, W. Hwang, and K. Chakrabarty. Droplet routing in the synthesis of digital microfluidic biochips. In Proc. of DATE, pages 323-328, Munich, Germany, Mar. 6-10, 2006.
- [36] T. Xu, W. L. Hwang, F. Su, and K. Chakrabarty. Automated design of pin-constrained digital microfluidic biochips under droplet-interference constraints. ACM Journal on Emerging Technologies in Computing Systems, 3(3): article #14, Nov., 2007.
- [37] T. Yan and M. D. F. Wong. Correctly modeling the diagonal capacity in escape routing. IEEE Trans. Comput.-Aided Design Integr. Circuits Syst., 31(2), Feb. 2012.
- [38] S-H. Yeh, J-W. Chang, T-W. Huang and T-Y. Ho. Voltage-aware chip-level design for reliability-driven pin-constrained EWOD chips. In Proc. of ICCAD, pages 353-360, San Jose, CA, USA, Nov. 5-8, 2012.
- [39] P-H. Yuh, C-L. Yang, and Y-W. Chang. BioRoute: a network-flow-based routing algorithm for the synthesis of digital microfluidic

biochips. IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst., 27(11):1928-1941, Nov. 2008.

- [40] Y. Zhao and K. Chakrabarty. Simultaneous optimization of droplet routing and control-pin mapping to electrodes in digital microfluidic biochips. IEEE Trans. Comput.-Aided Design Integr. Circuits Syst., 31(2), Feb. 2012.
- [41] Y. Zhao, K. Chakrabarty, R. Sturmer, and V. K. Pamula. Optimization techniques for the synchronization of concurrent fluidic operations in pin-constrained digital microfluidic biochips. IEEE Trans. Very Large Scale Integration Systems 20(6), Jun. 2012.
- [42] Y. Zhao, T. Xu, and K. Chakrabarty. Broadcast electrode-addressing and scheduling methods for pin-constrained digital microfluidic biochips. IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst., 30(7):986-999, Jul. 2011.



Daniel Grissom received the B.S. degree in computer engineering from the University of Cincinnati (UC), Ohio, in 2008 and the M.S. and Ph.D. degrees in computer science from the University of California, Riverside (UCR), in 2011 and 2014, respectively. He is now an assistant professor at Azusa Pacific University.

To date, he has published seven papers, three journals and filed one provisional patent. His research interests include languages, synthesis and hardware interfacing for digital microfluidic biochips.

Mr. Grissom was the recipient of the NSF graduate research fellowship (GRFP) award in 2010.



Jeffrey McDaniel received the B.S. degree in computer science and pure mathematics and M.S. degree in computer science from the University of California, Riverside (UCR), in 2011 and 2014 respectively. He is currently pursuing the Ph.D. degree in computer science at UCR.

To date, he has published four conference papers and one journal. His research interests include languages, synthesis and hardware interfacing for continuous-flow microfluidic biochips.



Philip Brisk received the B.S., M.S., and Ph.D. degrees, all in computer science, from UCLA in 2002, 2003, and 2006, respectively. From 2006-2009 he was a postdoctoral scholar in the Processor Architecture laboratory in the School of Computer and Communication Sciences at the École Polytechnique Fédérale de Lausanne, (EPFL), in Lausanne, Switzerland. He is now an assistant professor in the Department of Computer Science and Engineering at the University of California, Riverside (UCR).

His research interests include FPGAs, compilers, and design automation and architecture for application-specific processors. He was a recipient of the Best Paper Award at CASES, 2007 and FPL 2009. Dr. Brisk is or has been a member of the program committees of several international conferences and workshops, including DAC, ASPDAC, DATE, VLSI-SoC, FPL, FPT, and others. He has been the general (co-)chair of IEEE SIES 2009, IEEE SASP 2010, and IWLS 2011, and participated in the organizing committee of many other conferences and symposia as well.