

Abstraction Layers for Scalable Microfluidic Biocomputers

William Thies¹, John Paul Urbanski²,
Todd Thorsen², and Saman Amarasinghe¹

¹ Computer Science and Artificial Intelligence Laboratory

² Hatsopoulos Microfluids Laboratory

Massachusetts Institute of Technology

{thies, urbanski, thorsen, saman}@mit.edu

Abstract. Microfluidic devices are emerging as an attractive technology for automatically orchestrating the reactions needed in a biological computer. Thousands of microfluidic primitives have already been integrated on a single chip, and recent trends indicate that the hardware complexity is increasing at rates comparable to Moore’s Law. As in the case of silicon, it will be critical to develop abstraction layers—such as programming languages and Instruction Set Architectures (ISAs)—that decouple software development from changes in the underlying device technology.

Towards this end, this paper presents BioStream, a portable language for describing biology protocols, and the Fluidic ISA, a stable interface for microfluidic chip designers. A novel algorithm translates microfluidic mixing operations from the BioStream layer to the Fluidic ISA. To demonstrate the benefits of these abstraction layers, we build two microfluidic chips that can both execute BioStream code despite significant differences at the device level. We consider this to be an important step towards building scalable biological computers.

1 Introduction

One of the challenges in biological computing is that the laboratory protocols needed to carry out a computation can be very time consuming. For example, a 20-variable 3-SAT problem required 96 hours to complete [1], not counting the considerable time needed for setup and evaluation. To automate and optimize this process, researchers have turned to microfluidic devices [2–10]. Microfluidics offers the promise of a “lab on a chip” system that can individually control picoliter-scale quantities of fluids, with integrated support for operations such as mixing, storage, PCR, heating/cooling, cell lysis, electrophoresis, and others [11–13]. Apart from being amenable to computer control, microfluidics drastically reduces the volumes of samples, thereby reducing costs and improving capture kinetics. Using microfluidics, DNA hybridization times can be reduced from 24 hours to 4 minutes [10] and the number of bases needed to encode information can be decreased from 15 bases per bit to 1 base per bit [1, 8].

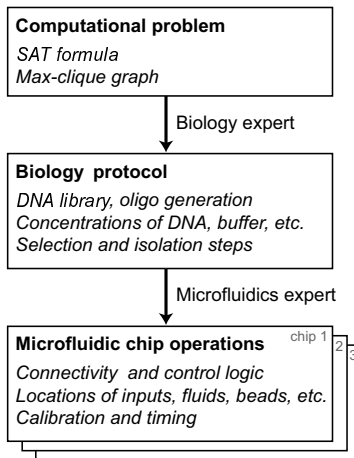


Fig. 1. Abstraction layers for DNA computing.

Thus has emerged a vision for creating a hybrid DNA computer: one that uses microfluidics for the plumbing (the control paths) and biological primitives for the computations (the ALUs). On the hardware side, this vision is becoming scalable: microfluidic chips have integrated up to 3,574 valves with 1,000 individually-addressable storage chambers [14]. Moreover, recent trends indicate that microfluidics is following a path similar to Moore’s law, with the number of soft-lithography valves per unit area doubling every 4.5 months [15].

On the software side, however, the microfluidic realm is lagging far behind its silicon counterpart. For silicon computers, the complexity and scale of the underlying hardware is masked by a set of well-defined abstraction layers. For example, transistors are organized into gates, which combine to form functional units, which together can implement an Instruction Set Architecture (ISA). The user operates at an even higher level of abstraction (e.g., C++), which is automatically translated into the ISA. These abstraction layers have proven critical for managing complexity. Without them, the computing field would have stagnated as every researcher tried to gain a transistor-level understanding of his machine.

Unfortunately, the current practice in experimental microfluidics is to expose all of the hardware resources directly to the experimentalist. Using a graphical system such as Labview, the user orchestrates the individual behavior of each valve in the microfluidic device. While this practice is merely tedious for today’s devices, it will soon become completely intractable—akin to programming a modern microprocessor by directly toggling each of a million gates.

In this paper, we present a system and methodology that uses new abstraction layers for scalable biological computing. As illustrated in Figure 1, our system consists of three layers. At the highest level, the programmer indicates the abstract computation to be performed—for example, in the form of a SAT formula. With some expertise in DNA computing and experimental biology, the computation can be transformed to the next layer: a portable biological proto-

col for performing the computation. The protocol is portable in that it does not depend on the physical implementation of the protocol; for example, it specifies fluid concentrations but not fluid volumes. Finally, the bottom layer specifies the operations needed to execute the protocol on a specific microfluidic chip. Each microfluidic chip designer provides a library that translates an abstract protocol into the specific sequence of valve actuations needed to execute that protocol on a specific chip.

These abstraction layers provide many benefits. Primarily, by using an architecture-independent description of the biological protocol (the middle layer), the application development can be decoupled from advances in the underlying device technology. Thus, as microfluidic devices come to support additional inputs, mixers, storage cells, etc., the existing suite of protocols can run without modification (much as C programs run without modification on successive generations of microprocessors). In addition, the protocol layer serves as a division of labor. Rather than requiring a heroic and brittle translation from a SAT formula directly to a microfluidic chip, a biologist provides a mapping to the abstract protocol while a microfluidics expert maps the protocol to the underlying device. The abstract protocol is also perfectly suited to simulation, thereby allowing the logical operations to be verified without relying on any physical implementation. Further, a portable protocol description could serve the role of pseudocode in technical publications, providing a precise account of the experimental methods used. Third-party protocols could be downloaded and executed (or called as sub-routines) on one's own microfluidic device.

In the long term, the protocol description language will support all of the operations needed for biological computing. However, as there does not yet exist a single microfluidic device that can encompass all the functionality (preparation of DNA libraries, selection, readout, etc.), this paper focuses on three fundamental primitives: fluid mixing, fluid transport, and fluid storage. We describe a programming system called BioStream (Section 2) that provides an architecture-independent interface for these operations. To show that BioStream is portable, we execute BioStream code on two fundamentally different microfluidic architectures (Section 3). We also present a novel algorithm for mixing fluids to a given concentration using the minimal number of simple on-chip mixing steps (Section 4). Our system represents a fully-functional, end-to-end demonstration of portable software on microfluidic hardware.

2 BioStream Protocol Language

We have developed a software system called BioStream for portable microfluidics protocols. BioStream is a Java library that virtualizes many aspects of the underlying hardware resources. While BioStream can be targeted by a compiler (for example, a DNA computing compiler that converts a mathematical problem into a biological protocol), it is also suitable for direct programming and experimentation by biologists. As such, the language provides several high-level abstractions to improve readability and programmer productivity.

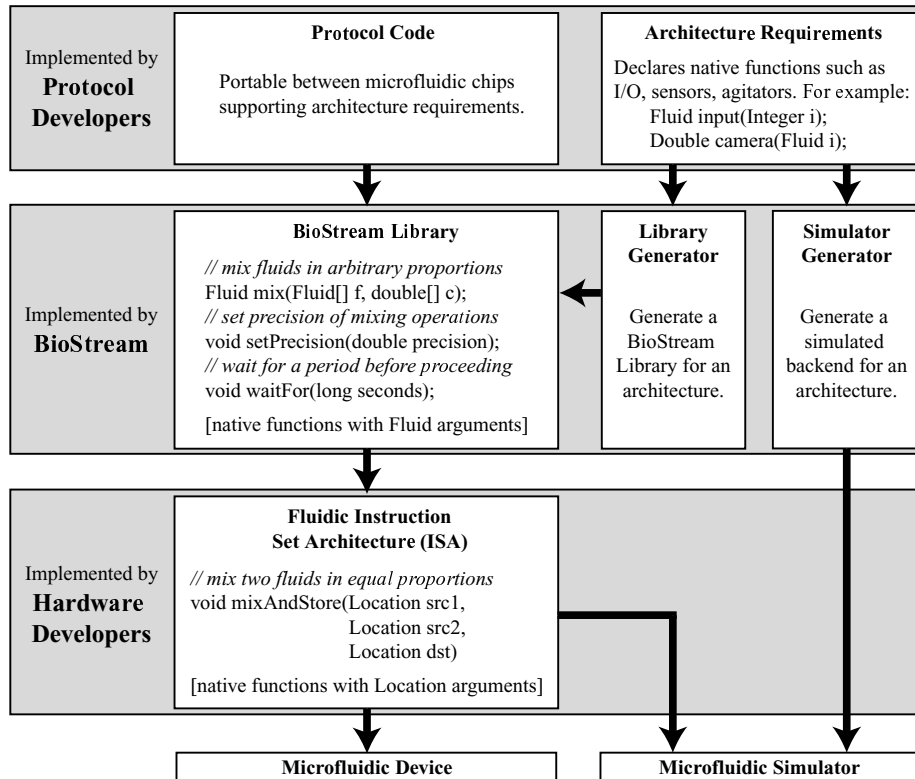


Fig. 2. Abstraction layers in the BioStream system.

2.1 Providing Portability

As shown in Figure 2, BioStream offers two levels of abstraction underneath the protocol developer. The first abstraction layer is the BioStream library, which provides first-class Fluid objects to represent the physical fluids on the chip. The programmer deals only with Fluid variables, while the runtime system automatically assigns and tracks the location of the corresponding fluids on the device. The library also supports a general mix operation for combining Fluids in arbitrary proportions and with adjustable precision.

The second abstraction layer, the Fluidic Instruction Set Architecture (ISA), interfaces with the underlying hardware. The fundamental operation is `mixAndStore`, which mixes two fluids in equal proportions and stores the result in a destination cell. (We describe how to translate the flexible mix operations in BioStream to a series of equal-proportion mixes in Section 4.) As all storage cells on the chip have unit volume, only one unit of mixture is stored in the destination; any leftover mixture may be discarded. As detailed in Section 3, this allows for a flexible implementation of `mixAndStore` on diverse architectures.

In addition to the abstractions for mixing, there are some architecture-specific features that need to be made available to the programmer. These “native func-

tions” include I/O devices, sensors, and agitators that might not be supported by every chip, but are needed to execute the program; for example, special input lines, cameras, or heaters. As shown in Figure 2, BioStream supports this functionality by having the programmer declare a set of architecture requirements. BioStream uses the requirements to generate a library which contains the same functionality; it also checks that the architecture target supports all of the required functions. Finally, BioStream includes a generic simulator that inputs a set of architecture requirements and outputs a virtual machine that emulates the architecture. This allows full protocol development and validation even without hardware resources.

The BioStream system is fully implemented. The reflection capabilities of Java are utilized to automatically generate the library and the simulator from the architecture requirements. As described in Section 3, we also execute the Fluidic ISA on two real microfluidic chips.

2.2 Example Protocol

An example of a BioStream protocol appears in Figure 3. This is a general program that seeks to find the ratio of two reagents that leads to the highest activity in the presence of a given indicator. Experiments of this sort are common in biology. For example, the program could be applied to investigate the roles of cytochrome-c and caspase 8 in activating apoptosis (cell death); cell lysate would serve as the indicator in this experiment [16]. The protocol uses feedback from a luminescence detector to guide the search for the highest activity. After sampling some concentrations in the given range, it descends recursively and narrows the range for the next round of sampling. Using self-directed mixing, a high precision can be obtained after only a few rounds.

The recursive descent program declares a `SimpleLibrary` interface (see bottom of Figure 3) describing the functionality required on the target architecture. In this case, a camera is needed to detect luminescence. While we have not mounted a camera on our current device, it would be straightforward to do so.

2.3 Automatic Fluid Regeneration

A distinguishing feature of BioStream code is the use of Fluid variables to represent samples on the device. The challenge in implementing this functionality is that physical fluids can be used only once, as they are consumed in mixtures and reactions. However, the programmer might reference a Fluid variable multiple times (e.g., variables A and B in the recursive descent example). BioStream supports this behavior by keeping track of how each Fluid was generated and automatically regenerating Fluids that are reused. This process assumes that the original steps employed to generate a Fluid (input, mixing, agitation, etc.) will produce an equivalent Fluid if repeated. While this assumption is a natural fit for protocols depending only on the concentrations of reagents, there are also non-deterministic systems (such as directed evolution of cells) to which it does not apply. We leave full consideration of such systems for future work.

```

import biostream.library.*;

public class RecursiveDescent {

    public static void main(String[] args) {
        String backend = args[0];

        SimpleLibrary lib =
            (SimpleLibrary)LibraryFactory.
                buildLibrary("SimpleLibrary", args[0]);
        run(lib);
    }

    private static void run(SimpleLibrary lib) {
        int ROUNDS = 10; int SAMPLES = 5;

        Fluid A = lib.input(new Integer(0));
        Fluid B = lib.input(new Integer(1));
        Fluid indicator = lib.input(new Integer(2));

        double center = 0.5, radius = 0.5;

        for (int i=0; i<ROUNDS; i++) {
            lib.setPrecision(0.1*(2*radius)/ SAMPLES);

            double bestActivity = -1; int bestJ = -1;
            for (int j=1; j<SAMPLES; j++) {

                double target = center+radius*
                    (1-2*(double)j/SAMPLES);
                Fluid sample = lib.mix(A, target, B, 1-target);

                Fluid test = lib.mix(indicator, 0.9, sample, 0.1);
                lib.wait(30);
                double act = lib.luminescence(test).doubleValue();

                if (act > bestActivity)
                    bestActivity = act; bestJ = j;
            }

            center = center+radius*(1-2*(double)bestJ/SAMPLES); // Zoom in by factor of 2 around best activity.
            radius = radius / 2;

            if (center < radius) center = radius;
            if (center > 1-radius) center = 1-radius;
        }
        System.out.println("Best activity: " + center);
    }

    interface SimpleLibrary extends FluidLibrary {
        Fluid input(Integer i);
        Double luminescence(Fluid f);
    }
}

```

*// The Recursive Descent protocol recursively zooms in on the ratio of fluids A and B that has the highest activity. It requires the following setup in the laboratory:
// - input(0) -- fluid A
// - input(1) -- fluid B
// - input(2) -- luminescent activity indicator*

*// Initialize the backend to use (for example, an actual chip or a microfluidic simulator) based on command-line input.
// Create an interface to the backend using the native functions declared in SimpleLibrary.*

*// Perform the protocol:
// Set number of rounds and samples per round.
// Assign names to the input fluids.*

*// Initialize center, radius of concentration range.
// Repeat for a number of rounds:
// Set absolute mixing precision to 10X more than the granularity of sampling.
// Repeat across concentrations in range:
// Obtain sample of the target concentration.
// Mix sample with indicator, wait, and measure activity.
// Remember highest activity.*

// Print concentration yielding highest activity.

*// Declare devices needed by RecursiveDescent:
// Require array of fluid inputs.
// Require luminescence camera.*

Fig. 3. Recursive descent search in BioStream.

	Driving fluid	Wash fluid	Mixing	Sample size	Inputs	Storage cells	Valves	Control lines
Chip 1	oil	N/A	rotary mixer	half of mixer	2	8	46	26
Chip 2	air	water	during transport	full mixer	4	32	140	21

Table 1. Key properties of the microfluidic chips developed. Chip 1 provides better isolation and retention of samples, while Chip 2 offers faster and simpler operation.

The regeneration mechanism works by associating each Fluid object with the name and arguments of the function that created it. The creating function must be a mix operation or a native function, both of which are visible to BioStream (the Fluid constructor is not exposed). BioStream maintains a valid bit for each Fluid, which indicates whether or not the Fluid is stored in a storage chamber on the chip. By default, the bit is true when the Fluid is first created, and it is invalidated when the Fluid is used as an argument to a BioStream function. If a BioStream function is called with an invalid Fluid, that Fluid is regenerated using its history. Note that this regeneration mechanism is fully dynamic (no analysis of the source code is needed) and is accurate even in the presence of pointers and aliasing.

The computation history created for Fluids can be viewed as a dependence tree with several interesting applications. For example, the library can execute a program in a demand-driven fashion by initializing each Fluid to an invalid state and only generating it when it is used by a native function. Dynamic optimizations such as these are especially promising for microfluidics, as the silicon-based control processors operate much faster than their microfluidic counterparts.

3 Microfluidic Implementation

To demonstrate an end-to-end system, we have designed and fabricated two microfluidic chips using a standard multi-layer soft lithography process [13]. While there are fundamental differences between the chips (see Table 1), both provide support for programmable mixing, storage, and transport of fluid samples. More specifically, both chips implement the `mixAndStore` operation in the Fluidic ISA: they can load two samples from storage, mix them together, and store the result. Thus, despite their differences, code written in BioStream will be portable between the chips.

The first chip (see Figure 4) isolates fluid samples by suspending them in oil [17]. To implement `mixAndStore`, each input sample is transported from a storage bin to one side of the mixer. The mixer uses rotary flow, driven by peristaltic pumps, to mix the samples to uniformity [18]. Following mixing, one half of the mixer is drained and stored in the target location. While the second half could also be stored, it is currently discarded, as the basic `mixAndStore` abstraction produces only one unit of output.

The second chip (see Figure 5) isolates fluid samples using air instead of oil. Because fluid transport is very rapid in the absence of oil, a dedicated mixing element is not needed. Instead, the input samples are loaded from storage and

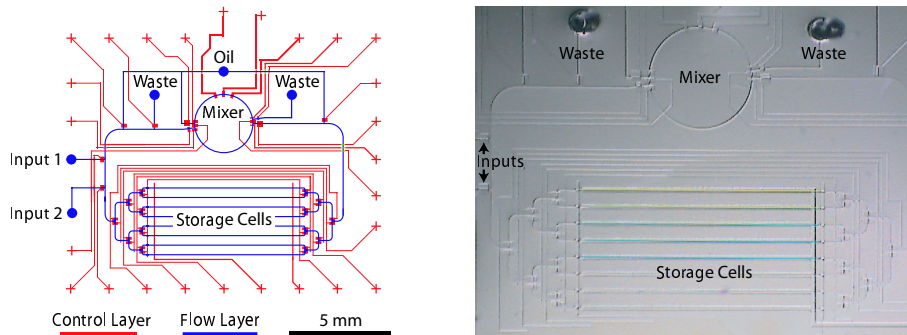


Fig. 4. Layout and photo of Chip 1 (driven by oil).

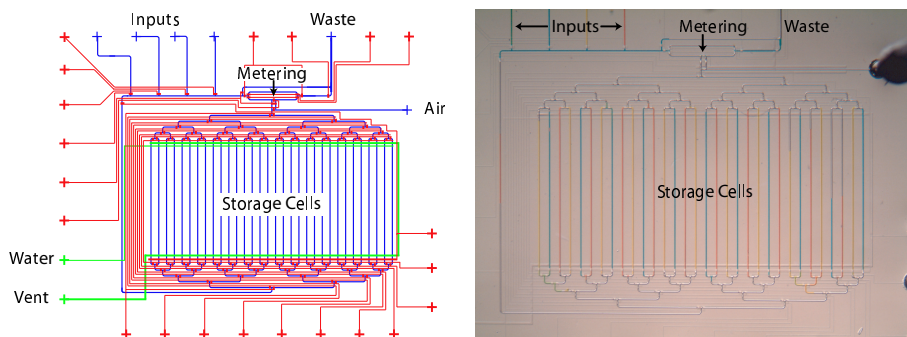


Fig. 5. Layout and photo of Chip 2 (driven by air).

aligned in a metering element; when the element is drained, the samples are mixed during transport to storage. Because the samples are in direct contact with the walls of the flow channels, a small fraction of the sample is lost during transport. This introduces the need for a wash phase, to clean the channel walls between operations. Also, to maintain sample volumes, the entire result of mixing is stored. Any excess volume is discarded in future mixing operations, as the metering element has fixed capacity.

To demonstrate BioStream’s portability between these two chips, consider the following code, which generates a gradient of concentrations:

```

Fluid blue = input(1);
Fluid yellow = input(2);
Fluid[] gradient = new Fluid[5];
for (int i=0; i<=4; i++) {
    gradient[i] = mix(blue, yellow, i/4.0, 1-i/4.0);
}

```

This code was used to generate the gradient pictured in Figure 4 and produces an identical result on both microfluidic devices. (The gradient shown in Figure 5 is different and was generated by a different program.)

4 Mixing Algorithms

The mixing and dilution of fluids plays a fundamental role in almost all bio-analytical procedures. Mixing is used to prepare input samples for analysis, to dilute concentrated substances, and to control reagent volumes. In DNA computing, mixing is needed for reagent preparation (e.g., DNA libraries, PCR buffers, detection assays) and, in some techniques, for restriction digests [19, 20] or fine-grained concentration control [21]. It is critical to provide integrated support for mixing on microfluidic devices, as otherwise the samples would have to leave the system every time a mixture is needed.

As described in the previous sections, our microfluidic chips support the `mixAndStore` instruction from the Fluidic ISA. This operation simply mixes two fluids in equal proportions. However, the `mix` command in BioStream allows the programmer to specify complex mixtures involving multiple fluids in various concentrations. To bridge the gap between these abstractions, this section describes how to obtain a complex mixture using a series of simple steps. We describe an abstract model for mixing, an algorithm for minimizing the number of steps required, and how to deal with error tolerances.

4.1 A Model of Mixing

The following definition gives our notation for mixtures.

Definition 1. A mixture \mathcal{M} is a set of substances S_i at given concentrations c_i :

$$\mathcal{M} = \{\langle S_1, c_1 \rangle \dots \langle S_k, c_k \rangle\} \\ \sum_{i=1}^k c_i = 1$$

For example, a mixture of 3/4 buffer and 1/4 reagent is denoted as $\{\langle \text{buffer}, 3/4 \rangle, \langle \text{reagent}, 1/4 \rangle\}$. We further define a *sample* to be a mixture with only one substance ($|\mathcal{M}| = 1$). For example, a sample of buffer is denoted $\{\langle \text{buffer}, 1 \rangle\}$, or just $\langle \text{buffer} \rangle$.

To obtain a given mixture on a microfluidic chip, one performs a series of mixes using an on-chip mixing primitive. While the capabilities of this mixer might vary from one chip to another, a simple *1-to-1 mixing model* can be implemented on both continuous flow and droplet-based architectures [18, 22]. In this model, all fluids are stored in uniform chambers of unit volume. The mix operation combines two fluids in equal proportions, producing two units of the mixture. However, since there may be some amount of fluid loss with every operation, the result of the mixture might not be able to completely fill the contents of two storage cells. Thus, the result is stored in only one storage cell, and the extra mixture is discarded.

The 1-to-1 mixing process can be visualized using a “mixing tree”. As depicted in Figure 6, each leaf node of a mixing tree represents a sample, while each internal node represents the mixture resulting from the combination of its




Fig. 6. Mixing tree yielding 3/4 buffer and 1/4 reagent.

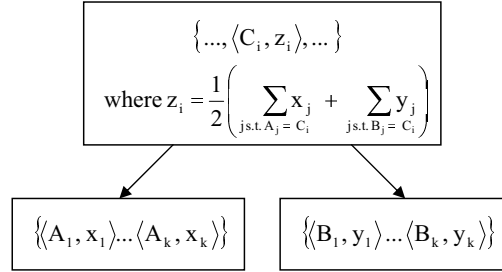


Fig. 7. Calculation of a parent mixture from child mixtures using a 1-to-1 mixer. For each substance, the resulting concentration is the average of the concentrations in the children.

children. Figure 7 illustrates that the mixture at an internal node can be calculated as the arithmetic mean of the components in child mixtures. In the 1-to-1 model, mixing trees are binary trees because each mix operation has two inputs. Evaluation of the tree proceeds from the leaf nodes upwards; the mixture for a given node can be produced once the child mixtures are available. The overall result of the operation is the mixture specified at the root node.

The following theorem is useful for reasoning about mixing trees. It describes the concentration of a substance in the overall mixture based on the depths of leaf nodes containing samples of the substance. The *depth* of a node n in a binary tree is the length of the path from the root node to n .

Theorem 1. Consider a mixing tree and a substance S . Let m_d denote the number of leaf nodes with sample $\langle S \rangle$ appearing at depth d of the tree. Then the concentration of S contained in the root mixture is given by $\sum_d m_d * 2^{-d}$.

Proof. A sample at depth d is diluted d times in the mixing process, each time by a factor of two. Thus it contributes 2^{-d} to the root mixture. Since each mix operation sums the concentrations from child nodes, the overall contribution is the sum across the leaf nodes at all depths: $\sum_d m_d * 2^{-d}$. \square

The following theorem describes the set of mixtures that can be obtained using a 1-to-1 mixer. Informally, it states that a mixture is reachable if and only if the concentration of each substance can be written as an integral fraction $k/2^d$.

Theorem 2. (1-to-1 Mixing Reachability) Consider a finite set of substances $\{S_1 \dots S_k\}$ with an unlimited supply of samples $\langle S_i \rangle$. Let \mathcal{R} denote the set of mixtures that can be obtained via any sequence of 1-to-1 mixes. Then:

$$\mathcal{R} = \left\{ \begin{array}{l} \{ \langle S_1, c_1 \rangle \dots \langle S_k, c_k \rangle \} \text{ s.t. } \exists p_i, q_i, d \in \mathcal{Z} : \\ LCM(q_1 \dots q_k) = 2^d \wedge \forall i \in [1, k] : c_i = \frac{p_i}{q_i} \end{array} \right\}$$

Proof. Available in an extended version of this paper [23].

It is natural to suggest a number of optimization problems for mixing. Of particular interest are the number of mixes and the number of samples consumed, as these directly impact the running time and resource requirements of a laboratory experiment. The following theorem shows that (under the 1-to-1 model) these two optimization problems are equivalent.

Theorem 3. *In any 1-to-1 mixing sequence, the number of samples consumed is exactly one greater than the number of mixes.*

Proof. By induction on the number of nodes, there is always exactly one more leaf node than internal node in a binary tree. The mixing tree is a binary tree in which each internal node represents a mix and each leaf node represents a sample. Thus there is always exactly one more sample consumed than there are mixes. \square

Note that this theorem only holds under the 1-to-1 mixing model, in which two units of volume are mixed but only one unit of the mixture is retained. For microfluidic chips that attempt to retain both units of mixture (such as droplet-based architectures or our oil-driven chip), it might be possible to decrease the number of samples consumed by increasing the number of mix operations.

4.2 Algorithm for Optimal Mixing

In this section, we give an efficient algorithm for finding a mixing tree that requires the minimal number of mixes to obtain a given concentration. For clarity, we frame the problem as follows:

Problem 1. (Minimal Mixing) Consider a finite set of substances $\{S_1 \dots S_k\}$ with an unlimited supply of samples $\langle S_i \rangle$. Given a reachable mixture $\{\langle S_1, p_1/n \rangle \dots \langle S_k, p_k/n \rangle\}$, what is the mixing tree with the minimal number of leaves?

Our algorithm runs in $O(k \lg n)$ time³ and produces an optimal mixing tree (with respect to this metric). The tree produced has no more than $k \lg n$ internal nodes.

The idea behind the algorithm, which we refer to as MIN-MIX, is to place a leaf node with sample $\langle S \rangle$ at depth d in the mixing tree if and only if the target concentration for S has a 1 in bit $\lg n - d$ of its binary representation. Theorem 1 then ensures that all substances have the desired concentrations, while fewer than $\lg n$ samples are used for each one.

Pseudocode for MIN-MIX appears in Figure 8. We illustrate its operation for the example mixture of $\{\langle A, 5/16 \rangle, \langle B, 4/16 \rangle, \langle C, 7/16 \rangle\}$. As shown in Figure 9, the algorithm begins with a pre-processing stage that allocates substances to bins according to the binary representation of the target concentrations. It then builds the mixing tree via calls to MIN-MIX-HELPER, which descends through the bins. When a bin is empty, an internal node is created in the graph and the procedure recurses into the next bin. When a bin has a substance identifier in it, the substance is removed from the bin and a corresponding sample is added

³ $\lg n$ denotes $\log_2 n$.

```

node Min-Mix
  (mixture { ⟨S1, p1/n⟩, ..., ⟨Sk, pk/n⟩ } ) {
  depth = lg(n)
  bins = new stack[depth+1]
  // pre-processing: build a stack of the
  // bitwise components of each concentration
  for i = 1 to k
    mask = 1
    for j = 0 to depth-1
      if (mask & pi ≠ 0) then
        bins[j].push(Si)
      endif
      mask = mask << 1
    endfor
  endfor
  return buildMixingHelper(bins, depth)
}


node Min-Mix-Helper(stack[] bins, int pow) {
  if bins[pow].empty() then
    node child1 = Min-Mix-Helper(bins,
                                  pow-1)
    node child2 = Min-Mix-Helper(bins,
                                  pow-1)
    return ⟨child1, child2⟩ as internal node
  else
    return bins[pow].pop() as leaf node
  endif
}

```

Fig. 8. MIN-MIX algorithm.

bin	2^{bin}	5A	4B	7C
4	16			
3	8			
2	4	A	B	C
1	2			C
0	1	A		C

(a)



(b)

(c)

Fig. 9. Example operation of MIN-MIX for the mixture $\{\langle A, 5/16 \rangle, \langle B, 4/16 \rangle, \langle C, 7/16 \rangle\}$. Part (a) illustrates the algorithm's allocation of substances to bins. The bin layout directly translates to a valid mixing tree, which appears in (b) with numbers indicating the order in which nodes are added to the tree. The mixing tree is redrawn in (c) for clarity.

as a leaf node to the graph. Figure 9 labels the order in which the nodes in the final mixing tree are created by the algorithm.

The following lemma is key to proving the correctness of MIN-MIX. We denote the n th least significant bit of x by $\text{LSB}(x, n)$. That is, $\text{LSB}(x, n) \equiv (x \ll n) \& 1$.

Lemma 1. *Consider the mixing tree t produced by $\text{MIN-MIX}(\{\langle S_1, p_1/n \rangle \dots \langle S_k, p_k/n \rangle\})$. A substance S_i appears at a depth d in t if and only if $\text{LSB}(p_i, \lg n - d) = 1$.*

Proof. If: It suffices to show that there is a substance added to the mixing tree for each LSB of 1 drawn from the p_i (that the substance appears at depth d is given by the only if direction.) Further, since $\text{bins}[j]$ is constructed to contain all substances i for which $\text{LSB}(p_i, j) = 1$, it suffices to show that a) all bins are empty at the end of the procedure, and b) the procedure does not try to pop from an empty bin. To show (a), use the invariant that each call to MIN-MIX-HELPER adds a total of 2^{-d} to the mixing tree, where d is the current depth; either a leaf node is added (which contributes 2^{-d} by Theorem 1) or two child

nodes are added, contributing $2 * 2^{-(d+1)} = 2^{-d}$. But since the initial depth is 0, the external call results in $2^0 = 1$ unit of mixture being generated. Since the bins represent exactly one unit of mixture (i.e., $\sum_j bins[j] * 2^{-j} = 1$), all bins will be used. To show (b), observe that MIN-MIX references the bins in order, testing if each is empty before proceeding. Thus no empty bin will ever be dereferenced.

Only if: When a substance is added to the tree from bins[j], it appears at depth $\lg n - j$ in the tree. This is evident from the recursive call in MIN-MIX-HELPER: it initially draws from bins[lg n] and then works down when the upper bins are empty. By construction, bins[j] contains only substances S_i with $LSB(p_i, j) = 1$. Thus, if S_i appears at depth d in the mixing tree, it was added from bins[lg n - d] which has $LSB(p_i, \lg n - d) = 1$. \square

The following theorem asserts the correctness of MIN-MIX.

Theorem 4. *The mixing tree given by MIN-MIX gives the correct concentration for each substance in the target mixture.*

Proof. Consider a component $\langle S, p/n \rangle$ of the mixture passed to MIN-MIX. Let m_d denote the number of leaf nodes with sample S at depth d of the resulting mixing tree. By Lemma 1, $m_d = LSB(p, \lg(n) - d)$. Using Theorem 1, this implies that the concentration for S in the root mixture is given by:

$$\begin{aligned} c &= \sum_d LSB(p, \lg(n) - d) * 2^{-d} \\ &= \sum_x LSB(p, x) * 2^{-(\lg(n)-x)} \\ &= \sum_x LSB(p, x) * 2^x / n \\ &= p/n \end{aligned}$$

Thus the concentration in the root node of the mixing tree is the same as that passed to MIN-MIX. \square

The MIN-MIX algorithm requires $O(k \lg n)$ time to find a mixing tree for mixture $(\{\langle S_1, p_1/n \rangle \dots \langle S_k, p_k/n \rangle\})$. The resulting mixing tree is optimal in that there does not exist a mixing tree that yields the same concentration using fewer mixes. Proofs of these properties are available in an extended version of this paper [23].

4.3 Supporting Error Tolerances

Thus far the presentation has been in terms of mixtures that can be obtained exactly with a 1-to-1 mixer, i.e., those with target concentrations in the form of $k/2^d$. However, the programmer should not be concerned with the reachability of a given mixture.

In the BioStream system, the programmer specifies a concentration range $[c_{min}, c_{max}]$ and the system ensures that the mixture produced will fall within the given range⁴. Such error tolerances are already a natural aspect of scientific

⁴ Alternately, BioStream supports a global error tolerance ϵ that applies to all concentrations.

experiments, as all measuring equipment has a finite precision that is carefully noted as part of the procedure. Given a concentration range, the system increases the internal precision d until some concentration $k/2^d$ (which can be obtained exactly) falls within the range.

5 Related Work

Several researchers have pursued the goal of automating the control systems for microfluidic chips. Gascoyne et al. describe a graphical user interface for controlling chips that manipulate droplets over a two-dimensional grid [24]. By varying parameters in the interface, the software can target grids with varying dimensions, speeds, etc. However, portability is limited to grid-based droplet processors. While the BioStream protocol language could target their chips, their software is not suitable for targeting ours.

Su et al. represent protocols as acyclic sequence graphs and map them to droplet-based processors using automatic scheduling [25] and module placement [26]. While the sequence graph is portable, it lacks the expressiveness of a programming language and cannot represent feedback loops (as in our recursive descent example). King et al. demonstrate a “robot scientist” that directs laboratory experiments using a high-level programming language [27], but lacks the abstraction layers needed to target other devices. Gu et al. have controlled microfluidic chips using programmable Braille displays [28], but protocols are mapped to the chip by hand.

Johnson demonstrates a special-purpose robotic system (controlled by Labview) that automatically solves 3-SAT problems using DNA computing [29]. Miniaturizing his benchtop devices could result in a fully-automatic microfluidic biocomputer. Livstone et al. compile an abstract SAT problem into a sequence of DNA-computing steps [5]. The output of their system would be a good match for BioStream and the abstraction layers proposed in this paper.

There are other microfluidic chips that support flexible gradient generation[30–32] and programmable mixing on a droplet array[33]. To the best of our knowledge, our chips are the only ones that provide arbitrary mixing of discrete samples in a soft lithography medium. A more detailed comparison of the devices is published elsewhere [17].

Fair et al. also suggest a mixing algorithm for diluting a single reagent by a given factor [34]. It seems that their algorithm performs a binary search for the target concentration, progressively approximating the target by a factor of two. However, since intermediate reagents must be regenerated in the search, this algorithm requires $O(n)$ mixes to obtain a concentration k/n . In contrast, our algorithm needs $O(\lg n)$ to mix two fluids.

6 Conclusions

Microfluidic devices are an exciting substrate for biological computing because they allow precise and automatic control of the underlying biological proto-

cols. However, as the complexity of microfluidic hardware comes to rival that of silicon-based computers, it will be critical to develop effective abstraction layers that decouple application development from low-level hardware details.

This paper presents two new abstraction layers for microfluidic biocomputers: the BioStream protocol language and the Fluidic ISA. Protocols expressed in BioStream are portable across all devices implementing a given Fluidic ISA. We demonstrate this portability by building two fundamentally different microfluidic devices that support execution of the same BioStream code. We also present a new and optimal algorithm for obtaining a given concentration of fluids using a simple on-chip mixing device. This algorithm is essential for efficiently supporting the mix abstraction in the BioStream language.

It remains an interesting area of future work to leverage DNA computing technology to target the BioStream language from a high-level description of the computation. This will create an end-to-end platform for biological computing that is seamlessly portable across future generations of microfluidic chips.

7 Acknowledgements

We are grateful to David Wentzlaff and Mats Cooper for early contributions to this research. We also thank John Albeck for helpful discussions about experimental protocols. This work was supported by National Science Foundation grant #CCF-0541319. J.P.U. was funded in part by the National Science and Engineering Research Council of Canada (PGSM Scholarship).

References

1. Braich, R.S., Chelyapov, N., Johnson, C., Rothmund, P.W.K., Adleman, L.: Solution of a 20-variable 3-SAT problem on a DNA computer. *Science* **296** (2002)
2. Farfel, J., Stefanovic, D.: Towards practical biomolecular computers using microfluidic deoxyribozyme logic gate networks. In: *DNA 11*. (2005)
3. Gehani, A., Reif, J.: Micro flow bio-molecular computation. *Biosystems* **52** (1999)
4. Grover, W.H., Mathies, R.A.: An integrated microfluidic processor for single nucleotide polymorphism-based DNA computing. *Lab on a Chip* **5** (2005)
5. Livstone, M.S., Weiss, R., Landweber, L.F.: Automated design and programming of a microfluidic DNA computer. *Natural Computing* (2006)
6. McCaskill, J.S.: Optically programming DNA computing in microflow reactors. *BioSystems* **59** (2001)
7. Somei, K., Kaneda, S., Fujii, T., Murata, S.: A microfluidic device for DNA tile self-assembly. In: *DNA 11*. (2005)
8. van Noort, D.: A programmable molecular computer in microreactors. In: *DNA 11*. (2005)
9. van Noort, D., Gast, F.U., McCaskill, J.S.: DNA computing in microreactors. In: *DNA 8*. (2002)
10. van Noort, D., Zhang, B.T.: PDMS valves in DNA computers. In: *SPIE International Symposium on Smart Materials, Nano-, and Micro-Smart Systems*. (2004)
11. Breslauer, D.N., Lee, P.J., Lee, L.P.: Microfluidics-based systems biology. *Molecular BioSystems* **2** (2006)

12. Erickson, D., Li, D.: Integrated microfluidic devices. *Anal. Chim. Acta* **507** (2004)
13. Sia, S.K., Whitesides, G.M.: Microfluidic devices fabricated in poly(dimethylsiloxane) for biological studies. *Electrophoresis* **24** (2003)
14. Thorsen, T., Maerkl, S., Quake, S.: Microfluidic large scale integration. *Science* **298** (2002)
15. Hong, J.W., Quake, S.R.: Integrated nanoliter systems. *Nature BioTechnology* **21**(10) (2003)
16. Allan, L., Morrice, N., Brady, S., Magee, G., Pathak, S., Clarke, P.: Inhibition of caspase-9 through phosphorylation at Thr 125 by ERK MAPK. *Nature Cell Biology* **5** (2003)
17. Urbanski, J.P., Thies, W., Rhodes, C., Amarasinghe, S., Thorsen, T.: Digital microfluidics using soft lithography. *Lab on a Chip* **6** (2006)
18. Chou, H., Unger, M., Quake, S.: A microfabricated rotary pump. *Biomedical Microdevices* **3** (2001)
19. Faulhammer, D., Cukras, A.R., Lipton, R.J., Landweber, L.F.: Molecular computation: RNA solutions to chess problems. *PNAS* **97**(4) (2000)
20. Ouyang, Q., Kaplan, P.D., Liu, S., Libchaber, A.: DNA solution of the maximal clique problem. *Science* **278** (1997)
21. Yamamoto, M., Matsuura, N., Shiba, T., Kawazoe, Y., Ohuchi, A.: Solutions of shortest path problems by concentration control. In: *DNA 7*. (2002)
22. Paik, P., Pamula, V., Fair, R.: Rapid droplet mixers for digital microfluidic systems. *Lab on a Chip* **3** (2003)
23. Thies, W., Urbanski, J.P., Thorsen, T., Amarasinghe, S.: Abstraction layers for scalable microfluidic biocomputers (Extended version). Technical Report MIT-CSAIL-TR-2006-034, MIT (2006) <http://hdl.handle.net/1721.1/32543>.
24. Gascoyne, P.R.C., Vykoukal, J.V., Schwartz, J.A., Anderson, T.J., Vykoukal, D.M., Current, K.W., McConaghy, C., Becker, F.F., Andrews, C.: Dielectrophoresis-based programmable fluidic processors. *Lab on a Chip* **4** (2004)
25. Su, F., Chakrabarty, K.: Architectural-level synthesis of digital microfluidics-based biochips. In: *ICCAD*. (2004)
26. Su, F., Chakrabarty, K.: Unified high-level synthesis and module placement for defect-tolerant microfluidic biochips. In: *DAC*. (2005)
27. King, R.D., Whelan, K.E., Jones, F.M., Reiser, P.G.K., Bryant, C.H., Muggleton, S.H., Kell, D.B., Oliver, S.G.: Functional genomic hypothesis generation and experimentation by a robot scientist. *Nature* **427** (2004)
28. Gu, W., Zhu, X., Futai, N., Cho, B.S., Takayama, S.: Computerized microfluidic cell culture using elastomeric channels and Braille displays. *PNAS* **101**(45) (2004)
29. Johnson, C.: Automating the DNA Computer to Solve n-Variable 3-SAT Problems. In: *DNA 12*. (2006)
30. Dertinger, S.K.W., Chiu, D.T., Jeon, N.L., Whitesides, G.M.: Generation of gradients having complex shapes using microfluidic networks. *Anal. Chem.* **73** (2001)
31. Neils, C., Tyree, Z., Finlayson, B., Folch, A.: Combinatorial mixing of microfluidic streams. *Lab on a Chip* **4** (2004)
32. Lin, F., Saadi, W., Rhee, S.W., Wang, S.J., Mittal, S., Jeon, N.L.: Generation of dynamic temporal and spatial concentration gradients using microfluidic devices. *Lab on a Chip* **4** (2004)
33. Pollack, M., Fair, R., Shenderov, A.: Electrowetting-based actuation of liquid droplets for microfluidic applications. *Applied Physics Letters* **77**(11) (2000)
34. Ren, H., Srinivasan, V., Fair, R.: Design and testing of an interpolating mixing architecture for electrowetting-based droplet-on-chip chemical dilution. *Transducers* (2003)