A Closed-Form Outlet Identity and Open-Source Reference Implementation for Controlled Incremental Filtration

Austin Harper Routt
Department of Biomedical Engineering
University of Houston
Houston, Texas, USA
ORCID: 0000-0002-7286-2924
Email: ahroutt@CougarNet.UH.EDU

***Abstract*—Controlled Incremental Filtration (CIF) is a microfluidic approach for size-based particle separation: a center channel feeds the sample, and two flanking side channels collect filtrate as a controlled fraction f_{gap} of the center-channel flow is diverted laterally through gaps; the channel geometry varies along the device to control f_{gap}, whose value sets the particle size cutoff. CIF has been demonstrated for bead enrichment [1], leukoreduction of platelet-rich plasma [2], [3], leukapheresis [4], and platelet size sorting [5]. CIF design has two key parameters: the size cutoff (set by f_{gap}) and the volumetric filtrate fraction F. The latter was previously the numerical endpoint of an iterative gap-by-gap recursion, with no closed-form prediction; here we address that gap.**

Under pressure equilibration (PE) and the boundary condition $Q_s(0) = 0$ shared across the CIF lineage, F is given exactly by $F = 2\,R_c/(R_s + 2\,R_c)$ in outlet geometry alone, equivalent to the textbook current-divider rule for the CIF outlet [13]. We derive this identity, verify it across five reference devices in three architectural regimes, and release an open-source reference implementation with a 237-validator, 19-stamp reproducibility suite.

To validate, we simulated the Dinh 2024 V1 and V3 outlets in three-dimensional Stokes flow (COMSOL Multiphysics 6.3). The closed-form predicts $F_{\mathrm{closed}} = 0.5592$ (V1) and 0.7189 (V3), about 10% above the published 0.505 ± 0.033 and 0.649 ± 0.036; adding 3-D physics closes the residual to within 1σ ($F_{\mathrm{COMSOL}} = 0.5070$ and 0.6281; σ-distances $+0.06$ and -0.58). The multi-variant ratio $\ln(1 - F_{\mathrm{V3}})/\ln(1 - F_{\mathrm{V1}})$ confirms the agreement: 1.549 (closed-form), 1.489 ± 0.203 (paper), 1.398 (COMSOL). The 10% residual is substantially reduced by 3-D outlet bifurcation, with the asymmetry between V1 and V3 ($\alpha_{\mathrm{3D}} = 0.811$ vs. 0.660) providing a geometry-dependent signature. Two alternative mechanisms are smaller in magnitude (channel-depth tolerance shifts F by less than 0.05σ; PE holds with at least a $48\times$ safety margin across all five reference devices), and the named subcategories of w_s-specific fabrication offset do not reproduce the observed asymmetry: isotropic PDMS shrinkage leaves F essentially invariant at fixed channel depth [16], and aspect-ratio-dependent shrinkage predicts the opposite V1>V3 ordering. The implementation is released under an MIT license at https://github.com/Austin-Routt/cif-reference-implementation, lowering the barrier for researchers seeking to apply, extend, audit, or build upon the Gifford-variant CIF design framework.

***Index Terms*—controlled incremental filtration, microfluidics, blood-component separation, design automation, open-source software, reference implementation, computational fluid dynamics, reproducibility**

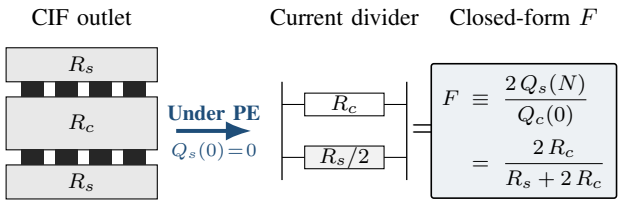

COMSOL of Dinh 2024 V1, V3 within 1σ ($+0.06\sigma$, -0.58σ)

Graphical abstract. Under pressure equilibration (PE) and the device-level boundary condition $Q_s(0) = 0$ shared by the published CIF lineage, the CIF outlet reduces to a two-resistor current divider, yielding the closed-form filtrate-fraction identity $F \equiv 2\,Q_s(N)/Q_c(0) = 2\,R_c/(R_s + 2\,R_c)$ in outlet geometry alone. Three-dimensional Stokes simulation in COMSOL Multiphysics 6.3 matches the Dinh 2024 V1 and V3 measurements within 1σ ($+0.06\sigma$, -0.58σ).

I. Introduction

Foundational context

Cellular blood components — red blood cells, leukocytes, and platelets — span a wide size range, and size-based separation underlies many clinical and research workflows, from centrifugation for blood banking to leukoreduction for transfusion safety to platelet preparation for transfusion. A range of microfluidic techniques

has been developed for size-based particle separation, including hydrodynamic filtration [15], in which particles above a threshold size remain in a main channel while smaller particles follow diverted fluid into side branches. Controlled Incremental Filtration (CIF) formalizes this principle with an iterative design framework suited to high-throughput blood-component applications: a center channel of width w_c feeds the sample, two flanking side channels of width $w_s(i)$ collect filtrate, and regularly-spaced gaps allow fluid to bleed laterally from center to sides [1]. At each gap, a fraction $f_{\mathrm{gap}}(i)$ of the center-channel flow $Q_c(i)$ is diverted into the adjacent side channel — that is, $Q_{\mathrm{gap}}(i) = f_{\mathrm{gap}}(i) \cdot Q_c(i)$ [5, p. 3]. The diverted flow corresponds to a fluid lamina within the center channel; the lamina geometry depends on $f_{\mathrm{gap}}(i)$, $w_c(i)$, and the gap-region flow field [5]. Particles whose effective diameter is approximately smaller than the lamina width follow the diverted fluid into the side channels, while larger particles remain in the center channel [5]. The side channels are incrementally widened along the device [1] to implement a designer-targeted $f_{\mathrm{gap}}(i)$ schedule that holds the size cutoff constant along the device [1], [5]. Under these conditions, each gap acts as a pressure-equalization node in Gifford's recursion: wide gaps present low resistance to lateral flow, and the small-fraction criterion $f_{\mathrm{gap}} < 5 \times 10^{-4}$ keeps the per-gap flux small [1]. Devices in the published CIF lineage typically operate well within this bound, with the specific f_{gap} value selected empirically to achieve the target size cutoff [1]. The resulting slowly-converging side-extraction profile separates particles by size at low shear and without the membrane fouling or clogging that limits conventional crossflow filtration [1]. The Gifford 2014 paper [1] provides the foundational mathematical framework: a recursion relation for the side-channel resistance trajectory (Eqn 4) coupled to the Yang series for rectangular-duct hydraulic resistance (Eqn 5, following [6]). Successive papers in the same lineage have extended the technology along two axes: size cutoff, and the device-level volumetric filtrate fraction F, defined as the proportion of input flow exiting through the side outlets. These extensions include higher-throughput leukoreduction of platelet-rich plasma using a reduced-shear CIF (RS-CIF) layout with $\geq$3-log leukocyte depletion and $F \approx 0.94$ [2], [3], a multiplexed 48-element leukapheresis device with a 6 μm critical diameter [4], and platelet size sorting with sub-12-femtoliter discrimination using a tapered center-channel geometry with $F \approx 0.50$–0.65 across three variants [5]. Across all four of these publications, the design framework remains encoded only in primary-literature equations and per-paper internal tooling, not in any shared, callable, open-source reference.

A CIF design specification therefore involves two distinct performance targets: the particle size cutoff and the device-level volumetric filtrate fraction F (the proportion of input flow that exits through the side outlets). The two had different computational status until the present work: the cutoff is set by local gap-region physics and has had a documented empirical calibration procedure since the original framework [1], whereas F is a device-level cumulative quantity that was computable only as the numerical endpoint of running the Gifford recursion to convergence. The present paper addresses the second and does not modify the first.

Limitations of the current design workflow

The standard CIF design procedure, distributed across the four published lineage papers, requires the researcher to (i) iteratively solve the Gifford recursion (Eqn 4 of [1]) for the side-channel width sequence $w_s(i)$; (ii) evaluate the Yang series (Eqn 5) at each iteration to compute the rectangular-duct hydraulic resistances entering the recursion; (iii) impose PDMS soft-lithography manufacturability constraints; (iv) lay out the photomask in CAD; and (v) iterate when empirical results reveal the need for re-design.

Four limitations of the current approach are:
(1) *No closed-form anchor.* The recursion has no closed-form anchor. The designer must run the iteration to ad hoc convergence to obtain F from a given device geometry.
(2) *No PE-consistency verification.* The iterative recursion presupposes pressure equilibration (PE) at every gap, but whether the resulting design satisfies PE at the outlet, a necessary condition for the recursion's solution to be physically realized, is not checked.
(3) *No empirical-validation infrastructure tied to the framework.* The present work finds, for example, a measurable offset between the closed-form prediction and the paper-measured filtrate fraction at the Dinh 2024 outlet geometries (quantified in Section IV-C); discrepancies between the framework's one-dimensional prediction and measured filtrate fractions could not previously be attributed to specific physical mechanisms without a workflow that anchors the framework prediction, a 3-D simulation refinement, and the empirical measurement on the same outlet geometry.
(4) *No open-source reference implementation.* Each new researcher re-implements the Gifford recursion, the Yang series, the meander primitive, and the mask layout from primary-literature equations, with no validated, callable, version-controlled trajectory-generator, no validator suite, no regression-stamp infrastructure, and no worked end-to-end example to use as a starting point.

Consequently, each new CIF device requires substantial re-implementation of the underlying framework, cross-paper reproducibility is not directly verifiable from

primary-literature equations alone, and these four gaps motivate the contributions that follow.

Related work and contributions of this paper

The field of generic continuous-flow microfluidic design automation (MFDA) has produced several open and academic tools — 3DμF [8], μFluidicGenius [9], Flui3d [10], OpenMFDA [11], among others — supporting general layout, component-library composition, channel routing, and photomask export. None encodes the Gifford recursion, the Yang resistance scaling, or PE-consistency verification, and none includes a built-in fluidic-simulation module (Flui3d's own comparison [10, Table 1] confirms this for 3DμF, Columba, and Flui3d itself); each requires the user to manually compute a CIF gap schedule before being usable for layout.

The hydraulic–electrical circuit analogy that underlies all CIF design — pressure $\leftrightarrow$ voltage, flow rate $\leftrightarrow$ current, hydraulic resistance $\leftrightarrow$ electrical resistance — is comprehensively reviewed in Oh *et al.* [13], the canonical *Lab on a Chip* Critical Review of pressure-driven microfluidic-network design via the electric circuit analogy; their Section 6.3 explicitly catalogues hydrodynamic-filtering networks (the parent class of CIF) under this framing, but does not write down the closed-form outlet identity for any specific filtering geometry. The framing itself persists in the most recent same-venue Critical Review: Li *et al.* [14] (*Lab on a Chip*, 2023) survey the hydraulic–electric analogy across four design-principle categories — flow-driven oscillators, pressure-driven oscillators, capillary microfluidic systems, and microfluidic logic gates — citing Oh 2012 among the foundational references. Hydrodynamic-filtering networks, however, are not among the catalogued categories, and no closed-form sieving outlet identity is written down.

The contributions of the present paper are threefold. *First*, a closed-form filtrate-fraction identity $F = 2\,R_c/(R_s + 2\,R_c)$ for the CIF outlet under pressure equilibration, together with the architectural boundary condition $Q_s(0) = 0$ — a condition shared by every CIF device in the published lineage. The identity is elementary in derivation but, to our knowledge, unwritten in the 12+ years of CIF literature using the hydraulic-circuit framing throughout. We acknowledge and frame this in Section II as a publication-gap closure rather than new physics. *Second*, empirical validation of that identity at the Dinh 2024 V1 and V3 outlet geometries via three-dimensional Stokes simulation in COMSOL Multiphysics 6.3, recovering filtrate fractions within 1σ of the published measurements at both variants. *Third*, an open-source reference implementation of the Gifford-variant CIF design framework as a Methods deliverable, anchored on the identity. The implementation comprises a callable trajectory-generator module (implemented as a reproducible Jupyter notebook), a worked end-to-end Gifford 2014 photomask example, and a validator suite of 218 core plus 19 supplementary assertions backed by 19 regression stamps.

Aims and roadmap

In the present study, we aim to (i) derive a closed-form filtrate-fraction identity $F = 2\,R_c/(R_s + 2\,R_c)$ that exactly characterizes any CIF outlet under pressure equilibration; (ii) implement and test an open-source reference implementation of the Gifford-variant CIF design framework anchored on this identity, comprising a reusable one-dimensional trajectory-generator module, a worked end-to-end Gifford 2014 mask example, and a reproducibility infrastructure; and (iii) empirically validate the underlying physical framework against published Dinh 2024 measurements via three-dimensional Stokes simulation in COMSOL Multiphysics 6.3 [7].

We further demonstrate that the trajectory-generator module produces filtrate fractions consistent with the closed-form identity across all five reference devices spanning three architectural regimes — Gifford bead (constant-w_c, constant-f_{gap}); Dinh V1, V2, V3 (tapered-w_c, constant-f_{gap}); and Xia RS-7 (tapered-w_c, variable-f_{gap} via the Xia 2022 corrected scaling [3]) — and that the multi-variant ratio $\ln(1 - F_{\mathrm{V3}})/\ln(1 - F_{\mathrm{V1}})$ agrees within 1σ of the propagated experimental uncertainty across closed-form, paper, and COMSOL computations (quantitative results in Section IV).

We acknowledge that the closed-form identity itself is mathematically elementary — equivalent under the standard hydraulic–electrical analogy [13] to the textbook two-resistor current-divider rule applied to the CIF outlet — and frame the contribution as the publication-gap closure across the 12+-year CIF literature plus the empirical validation and the open-source reference implementation, not the derivation of new physics.

The reference implementation operationalizes this framework, providing the worked Gifford 2014 mask example and the validator and regression-stamp infrastructure on which Dinh-, Xia-, and Lezzar-style mask-layout extensions can be built; producing those mask-layout extensions is identified as future work.

All code, regression stamps, validator outputs, and design files are publicly available under an MIT license at https://github.com/Austin-Routt/cif-reference-implementation; the released version cited in this preprint (`v3.20`) is permanently archived at Zenodo [12].

II. Theory: A closed-form filtrate-fraction identity

Every device in the published Controlled Incremental Filtration lineage [1], [2], [4], [5] shares a single architectural feature: sample enters only the center channel, with no external inflow into the side channels. Gifford 2014 [1, Fig. 1A] shows a CIF device comprising a central flow channel that retains the concentrated particles and two adjacent side channels that carry the filtrate, fed at a single inlet to the central channel. Xia 2016 [2], [3] preserves this architecture in the reduced-shear (RS-CIF) paperclip layout [2, Fig. 1]: a single inlet to the central channel ($w_c(i=0) = 300\,\mu\text{m}$), with the upstream side channels implemented as serpentine segments of equivalent fluidic resistance and no external side-channel inflow. Lezzar 2022 [4, Fig. 2c] multiplexes 48 such CIF elements in parallel: one common inlet distributes sample to the central channel of each element, and two collection outlets gather the retentate and the combined filtrate. Dinh 2024 [5, p. 4] states that each CIF device has a single inlet feeding the middle channel and two distinct outlets, one for the middle channel and one collecting the combined side-channel flow. This architectural commonality makes

$$Q_s(0) = 0 \tag{1}$$

a device-level boundary condition rather than a modeling assumption (Fig. 1).

Combining the architectural boundary condition (1) with mass conservation across the device gives

$$Q_c(0) \;=\; Q_c(N) \;+\; 2\,Q_s(N). \tag{2}$$

Pressure equilibration (PE)—the assumption that gaps act as pressure-equalizing nodes across the width of the device [1], which underlies the Gifford recursion (Eqn 4 of [1])—applied at the outlet requires

$$Q_s(N)\,R_s \;=\; Q_c(N)\,R_c, \tag{3}$$

where R_c and R_s are the per-pitch hydraulic resistances of the center and side channels at the outlet, evaluated via the Yang series for rectangular ducts [6]. With the filtrate fraction defined as $F \equiv 2\,Q_s(N)/Q_c(0)$, the identity follows in three elementary steps. Substituting (2) into the denominator eliminates the inlet flow rate; dividing numerator and denominator by $Q_c(N)$ and applying (3) in the form $Q_s(N)/Q_c(N) = R_c/R_s$ converts the flow-rate ratio into a resistance ratio; finally, multiplying numerator and denominator by R_s clears the inner fractions:

$$\begin{aligned} F \;&=\; \frac{2\,Q_s(N)}{Q_c(N) + 2\,Q_s(N)} \\ &=\; \frac{2\,(R_c/R_s)}{1 + 2\,(R_c/R_s)} \\ &=\; \boxed{\frac{2\,R_c}{R_s + 2\,R_c}}, \end{aligned} \tag{4}$$

with R_c and R_s evaluated at the outlet w_c and w_s. Conditional on the architectural boundary condition (1) and PE, F depends on no other quantity: it is independent of the inlet flow rate, the gap schedule $f_{\text{gap}}(i)$, the taper shape, the scaling exponent, and the total gap count N. Any f_{gap} schedule—including the Xia 2016/2022 corrected variable form $f_{\text{gap}}(i) = \left(w_{c,\text{ref}}/w_c(i)\right)^2 f_{\text{gap}}^{\text{ref}}$ [2], [3]—traces a different trajectory through the recursion but, provided the trajectory reaches the same outlet (w_c, w_s), yields the same F. The schedule-independence of F at fixed outlet geometry is therefore a direct corollary of (4); the closed-form-vs-recursion cross-validation across five reference devices spanning constant-f_{gap} (Gifford, Dinh) and variable-f_{gap} (Xia) regimes is summarized in Table I.

Equivalent statement: a parallel-resistance argument.: The same identity also follows from the hydraulic–electrical analogy ($p \leftrightarrow V$, $Q \leftrightarrow I$, $R \leftrightarrow R$) [13]. Under PE, the center channel and the two side channels all connect the same upstream pressure to the same outlet pressure, forming three parallel paths: the center path of resistance R_c carrying $Q_c(N)$, and two identical side paths of resistance R_s each carrying $Q_s(N)$. Because the two side paths are identical, it is convenient to merge them into one; two equal resistances in parallel give

$$R_s \parallel R_s \;=\; R_s/2, \tag{5}$$

leaving just two parallel paths — the center path (R_c, carrying $Q_c(N)$) and the combined side path ($R_s/2$, carrying $2\,Q_s(N)$). Parallel paths span the same pressure drop Δp, so by $\Delta p = QR$ each carries a flow $Q = \Delta p/R$, inversely proportional to its resistance. The filtrate fraction is the share of the total flow carried by the combined side path:

$$\begin{aligned} F \;&=\; \frac{\Delta p/(R_s/2)}{\Delta p/R_c \;+\; \Delta p/(R_s/2)} \;=\; \frac{1/(R_s/2)}{1/R_c \;+\; 1/(R_s/2)} \\ &=\; \frac{R_c}{R_c \;+\; R_s/2} \;=\; \frac{2\,R_c}{R_s + 2\,R_c}, \end{aligned} \tag{6}$$

which recovers (4). This is the current-divider rule of circuit theory, written here directly in terms of flow — and in this form it is immediate that F depends only on R_c and R_s, the center- and side-channel resistances

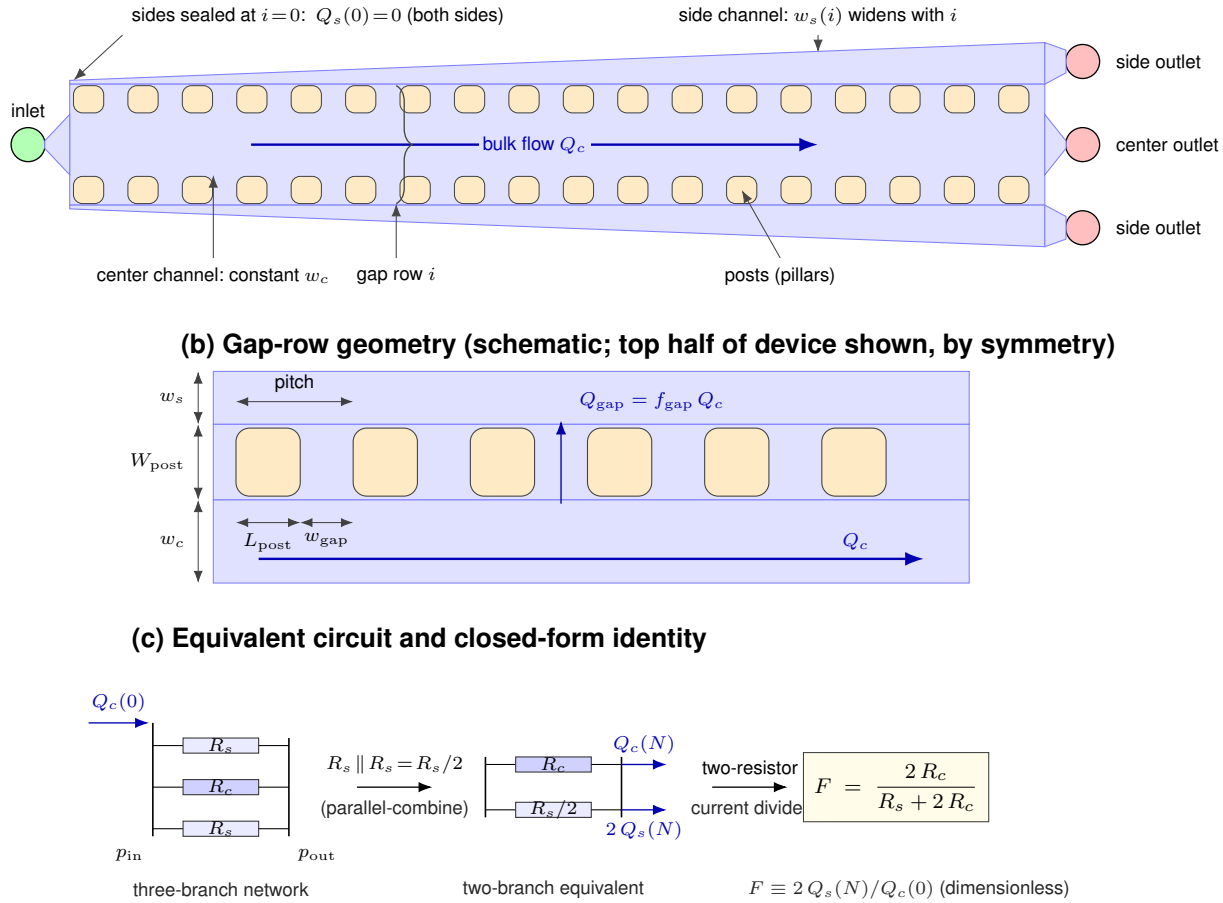

Fig. 1: Schematic of a CIF device. (a) Whole-device architecture (Gifford-style constant-w_c convention; the Dinh tapered-w_c and Xia variable-f_{gap} variants share the same underlying topology). A single inlet feeds the center channel only; both side channels are sealed at the inlet, giving $Q_s(0) = 0$. Rows of pill-shaped solid posts [1], [2], [4], [5] (orange-cream; rendered here as rounded rectangles) span the full wall thickness W_{post} (labeled in panel b) between center and side channels, defining regular gaps through which filtrate is extracted laterally. All fluid regions share a single color; only the solid posts are distinct. "Gap row i" (marked by the brace) is the bilateral pair of gaps at the i-th cross-sectional position, between adjacent post columns. Three outlets are shown (Gifford convention); Dinh-style devices merge the two side outlets into a single filtrate port. The side-channel widening rate is exaggerated. (b) Gap-row geometry (schematic; mirror-symmetric about the center-channel mid-plane, so only the top side channel and its post band are shown; proportions exaggerated). Adjacent posts form a gap of width w_{gap} along the flow direction; post length along flow is L_{post}; pitch is $L_{\text{post}} + w_{\text{gap}}$. Bulk flow Q_c runs through the center channel; per-gap filtrate flow $Q_{\text{gap}} = f_{\text{gap}}\, Q_c$ exits laterally through each gap. (c) Equivalent circuit under pressure equilibration (3), via the hydraulic–electrical analogy ($p \leftrightarrow V$, $Q \leftrightarrow I$, $R \leftrightarrow R$) [13]. The outlet maps to a three-branch parallel network (R_s, R_c, R_s) carrying total flow $Q_c(0)$ at the left rail (equal to total outlet flow by mass conservation (2)). Parallel-combining the two equal side branches (5) gives a two-branch equivalent (R_c vs. $R_s/2$) carrying $Q_c(N)$ and $2Q_s(N)$ at the outlet rail; $Q_s(N)$ is per-side, following Gifford 2014 [1]. The two-resistor current-divider rule then yields the closed-form identity $F = 2R_c/(R_s + 2R_c)$ (4), where $F \equiv 2Q_s(N)/Q_c(0)$ is the dimensionless filtrate fraction (a ratio of currents) determined by outlet geometry alone.

at the outlet, not on anything upstream of them. The contribution of the present work is therefore not new physics but a change of viewpoint: recast as a current divider, the asymptotic CIF outlet under PE has an elementary closed form — one that more than a decade of CIF literature, working in hydraulic resistance and the gap-by-gap recursion, never made explicit.

Verification.: We verified (4) in four ways. *First*, the three-step algebraic derivation above establishes the identity from the device mass balance (2) — the architectural boundary condition (1) combined with conservation across the device — together with outlet pressure equilibration (3); the Yang series [6] is not a premise of the derivation and enters only when the resistances are evaluated numerically (check four below). *Second*, the parallel-resistance re-derivation (5)–(6) reaches the identical result from the same premises by an independent algebraic route. The two derivations are the same physics in two notations — the hydraulic–electrical analogy [13] maps one onto the other — so their agreement is a consistency check, not independent evidence. Its value is that the current-divider form makes the dependence of F on outlet geometry alone immediately visible. *Third*, an independent symbolic verification in SymPy (notebook §8.2, the symbolic F-identity proof) machine-checks the algebra and reproduces (4) from total-flow conservation and outlet pressure equilibration; the same proof confirms symbolically that the identity is independent of the inlet condition, the gap schedule, and the taper shape. *Fourth*, the closed-form-vs-recursion cross-validation referred to above (Table I) confirms the identity numerically across all five reference devices, which span the three architectural regimes of the published CIF lineage: the closed form and the converged Gifford recursion agree to $|F_{\text{closed}} - F_{\text{recursion}}| \leq 2.35 \times 10^{-4}$ across the four constant-f_{gap} devices (the Gifford bead device and the three Dinh variants) and to $|\Delta F| = 1.97 \times 10^{-5}$ for the variable-f_{gap} Xia RS-7 case [2], [3], in each case within the recursion's own convergence floor. The variable-f_{gap} regime is the most stringent instance of this check: f_{gap} enters every step of the iterative recursion yet is provably absent from the closed form. Section IV-A develops the per-device comparison. Beyond these four checks, the limiting behaviour of (4) is also physical: $F \to 0$ as $R_s \to \infty$, $F \to 1$ as $R_s \to 0$, and $\partial F/\partial R_s < 0$ throughout, all matching the expected limits.

Why the identity may have been overlooked.: The math is elementary. The boxed identity (4) follows in three lines from the Gifford 2014 premises plus one device-architecture-derived boundary condition; equivalently, it is the textbook current-divider rule applied to the two-branch equivalent network (5)–(6), derivable in two lines from undergraduate circuit theory. Gifford 2014 derives the iterative recursion (Eqn 4 of [1]) from flow conservation, pressure equilibration, and Hagen–Poiseuille resistance — the hydraulic-electric analogues of Kirchhoff's laws and Ohm's law — and the subsequent CIF lineage [2], [4], [5] uses the same per-step pressure-drop and flow-conservation accounting throughout. The 12+-year publication gap persists nonetheless across the field's benchmark reviews. The 2012 *Lab on a Chip* Critical Review of pressure-driven microfluidic-network design under the electric-circuit analogy [13] catalogues hydrodynamic-filtering networks (the parent class of CIF) under exactly this framing, yet does not write down the asymptotic outlet identity for any specific filtering geometry. The contemporary 2023 same-venue Critical Review [14] catalogues four design-principle categories under the same framing — flow-driven oscillators, pressure-driven oscillators, capillary microfluidic systems, and microfluidic logic gates — but does not include hydrodynamic-filtering networks among them, nor write down any closed-form sieving outlet identity. The present-day ML-augmented state of the art [9] (*Sci. Adv.*, Jan. 2026) couples ML-trained resistance prediction (test $R^2 = 0.999$ on a held-out 20% of a 328,320-maze simulated dataset) with linear-programming (GLPK) circuit synthesis driven by a metaheuristic search, under Oh 2012's electric-circuit framework (their foundational reference, ref. 56). Reported per-outlet experimental errors on the tested 3- and 4-outlet circuits span 0.7–12.3%, yet the framework likewise does not encode the CIF outlet identity, the Gifford recursion, the Yang resistance scaling, or PE-consistency verification.

What none of these has written down is what circuit theory implies for the *asymptotic* outlet — that the convergent F is an elementary current-divider determined by outlet geometry alone. The identity follows in three lines because the right question ("what is the convergent value of the recursion as a function of outlet geometry only?") had not previously been asked. The contribution of this work is therefore the publication-gap closure plus the empirical demonstration across all three architectural regimes of the published CIF lineage, not the derivation of new physics. The math being elementary is a feature of the result: an elementary identity that exactly characterizes the outlet is more useful and more reproducible than a complicated one that approximately does so. Two practical design routes follow immediately from (4). Rearranging gives $R_c/R_s = F/[2(1-F)]$, so the designer can specify either the target filtrate fraction F — which fixes the outlet resistance ratio and, at the fabrication-fixed channel depth, determines $w_{s,\text{out}}$ from any chosen $w_{c,\text{out}}$ by inverting the Yang series for the rectangular-duct resistance — or specify $w_{c,\text{out}}$ and $w_{s,\text{out}}$ directly, evaluating the Yang series forward

TABLE I: Reference-device validation of the trajectory-generator module across all five reference devices in the closed-form regression, spanning three architectural regimes. For each device, `generate_cif_device()` produces a trajectory; F_{gen} is its converged (recursion) filtrate fraction, and F_{closed} is the identity (Eq. 4) evaluated at the same outlet geometry. F_{gen} reproduces F_{closed} to $|F_{\mathrm{closed}}-F_{\mathrm{gen}}| \leq 2.35 \times 10^{-4}$ across the four constant-f_{gap} devices (Block 15.3) and to 1.97×10^{-5} for the variable-f_{gap} Xia RS-7 case (§8.4). For the three Dinh variants, F_{gen} is the converged recursion endpoint under the quartic schedule-independence test ($f_{\mathrm{gap}}(i) = (w_{c,\mathrm{ref}}/w_c(i))^4 f_{\mathrm{gap}}^{\mathrm{ref}}$; notebook Block 15.2); under the variants' published constant f_{gap} (Block 7c) the residual to F_{closed} is smaller still, consistent with the identity's schedule-independence at fixed outlet geometry. The Dinh V1 and V3 "Paper F" entries are direct measurements (volume recovery in filtrate, Dinh 2024 Fig. 4); V2 is not reported by the source. The Gifford and Xia "Paper F" entries ("≈") are inferred in one step from a related quantity reported by each source (see notes a, b below the table).

Device	w_c schedule	f_{gap} schedule	Paper F	Generator F	$\lvert F_{\mathrm{closed}} - F_{\mathrm{gen}} \rvert$
Gifford bead[a]	constant	constant	≈0.90	0.9065	2.82×10^{-5}
Dinh V1	taper	constant	0.505 ± 0.033	0.5594	2.35×10^{-4}
Dinh V2	taper	constant	not reported	0.6169	6.48×10^{-6}
Dinh V3	taper	constant	0.649 ± 0.036	0.7190	1.08×10^{-4}
Xia RS-7[b]	taper	**variable**	≈0.94	0.8624	1.97×10^{-5}

[a] Gifford 2014 reports ∼10× enrichment "for particle populations that are maintained in the central channel" at this device (Fig. 4); since enrichment $= 1/(1-F)$ for retained particles, $F \approx 0.90$.
[b] Xia 2016 reports a ≈15:1 filtrate-to-retentate volumetric flow ratio (Fig. 3 caption), giving $F = 15/16 \approx 0.94$. F_{gen} is evaluated at a published-figure-derived nominal $w_{s,\mathrm{out}}$; the ∼8% offset is not analyzed, as this row tests variable-f_{gap} schedule-independence via the $|F_{\mathrm{closed}} - F_{\mathrm{gen}}|$ column, not paper-vs-framework agreement.

for R_c and R_s and substituting into (4). Both routes replace running the recursion to convergence and tuning $w_s(i = 1)$ until the endpoint matches with solving a single algebraic relation.

Conditional scope.: (4) is conditional on PE. Whether real outlets achieve PE at the geometries of interest is a separate empirical question, addressed in two places: a PE half-life diagnostic in Section III-A7 verifies that PE is achieved well within the device length at all five reference devices (48–84× safety margin), and three-dimensional Stokes simulation in Section IV-B independently confirms the PE-consistency framing at the Dinh 2024 V1 and V3 outlets, with network-vs-COMSOL agreement to 2.3% on the PE half-life. A more general identity not requiring PE may exist, but is not the present claim. Finally, (4) replaces only the outlet design step of the iterative recursion; the gap-by-gap recursion is still required upstream for the gap-schedule generation that sets the size cutoff.

III. Materials and Methods

The reference implementation comprises three components, each with explicit scope. The *trajectory-generator module* (notebook Blocks 1–7 and 9–15, plus the §8.7 PE-consistency diagnostic) is a callable Python framework with two primary entry points: `generate_cif_device(CIFDeviceParams)` produces a self-consistent $w_s(i)$ trajectory and F prediction for any Gifford-recursion-compatible parameter set, and `pe_consistency_check(...)` verifies, for a candidate CIF outlet geometry, that the PE half-life is less than 5 % of the total gap count (i.e. pressure equilibrates within the first 5 % of the device length; the $k_{1/2}/N_{\mathrm{total}} < 0.05$ criterion applied in §8.7). The *worked end-to-end Gifford 2014 mask example* (Block 8) takes a variant of the published Gifford 2014 device [1] from physics through trajectory to a fabricable DXF photomask. The *reproducibility infrastructure* (218 core and 19 supplementary validators; 19 regression stamps, with the three resistor-network-based stamps — Blocks 12, 15 and §8.7 — documented as BLAS-portable and the remaining 16 as bit-identical or bit-portable) enables third-party audit and supports future contributions. The remainder of this section describes each component in turn, followed by the COMSOL validation simulations (run externally in COMSOL Multiphysics; post-processed in Block 19).

A. The Python CIF reference implementation

1) Trajectory-generator module: A single function, `generate_cif_device()`, implements the Gifford recursion for any device specifiable in the framework's

scope. We exercise it on the five reference devices summarized in Table I, which together span the three design configurations present in the published CIF lineage: constant-w_c constant-f_{gap} (the Gifford 2014 bead device [1]); tapered-w_c constant-f_{gap} (the three Dinh 2024 platelet-sorting variants V1, V2, V3, each with w_c narrowing $75\,\mu\mathrm{m} \to 52\,\mu\mathrm{m}$ [5]); and tapered-w_c variable-f_{gap} (the Xia 2016 RS-7 leukoreduction device [2], with w_c narrowing $300\,\mu\mathrm{m} \to 150\,\mu\mathrm{m}$ under the 2022 author-corrected [3] quadratic scaling $f_{\mathrm{gap}}(i) = (w_{c,\mathrm{ref}}/w_c(i))^2\, f_{\mathrm{gap}}^{\mathrm{ref}}$, the only published-lineage device with a spatially-varying f_{gap} schedule by design). For each device, the function produces a converged trajectory and the filtrate fraction F_{gen} tabulated in Table I; in every case F_{gen} matches the closed-form identity (4) to a residual $|\Delta F| = |F_{\mathrm{gen}} - F_{\mathrm{closed}}| \le 2.35 \times 10^{-4}$.

For the three Dinh variants, the tabulated F_{gen} values come from a schedule-independence test in notebook Block 15.2 that replaces the published constant f_{gap} with a quartic schedule, $f_{\mathrm{gap}}(i) = (w_{c,\mathrm{ref}}/w_c(i))^4\, f_{\mathrm{gap}}^{\mathrm{ref}}$, and verifies that (4) still holds. Under the variants' actual published f_{gap} (verified for V1 and V3 in Block 7c) the residual collapses to the recursion's numerical floor. The closed-form cross-check itself is performed in Block 15.3 (Gifford plus the three Dinh variants) and §8.4 (the Xia RS-7 extension).

2) Yang resistance kernel (Block 2): The hydraulic resistance of a long straight rectangular duct of width w and depth d is given by the Yang series [6] — the exact Stokes-flow solution for a rectangular cross-section, expressed as an infinite sum. The series converges rapidly: the relative error against an $N = 200$ reference falls from 3.9×10^{-7} at $N = 10$ to 6.3×10^{-10} at $N = 50$ (Block 2 Validator 1). We truncate at $N = 50$ following Gifford 2014 [1, p. 4]. The convergence behaviour holds across the full range of channel shapes encountered in the published CIF lineage — from shallow, wide side channels (depth-to-width ratio $\sim$0.27) to deep, narrow gap apertures ($\sim$9). The kernel-level error is several orders of magnitude below the $\sim 10^{-4}$ closed-form-vs-recursion agreement floor of the overall pipeline, so the resistance kernel contributes negligibly to overall pipeline error.

3) Bisection inverter (Block 3): Each step of the Gifford recursion produces a target side-channel resistance and requires the corresponding channel width. The bisection inverter `invert_R_to_w_scalar` performs this conversion at fixed channel depth. We chose bisection rather than Newton or secant iteration because the Yang resistance is strictly monotone decreasing in w at fixed d (Block 2 Validator 7), so the inverse is single-valued on $(0, \infty)$ and bisection converges unconditionally. Bisection works by maintaining a search interval $[w_{lo},\, w_{hi}]$ known to contain the solution and halving it at each step. The implementation begins with a default search interval of $[10^{-7},\, 10^{-2}]$ m ($100\,\mathrm{nm}$ to $10\,\mathrm{mm}$), automatically extending it by a factor of ten at either end if the target lies outside. This default interval comfortably encloses the full channel-width range exercised by the five reference devices — from the $16\,\mu\mathrm{m}$ Xia inlet side channels up to the $300\,\mu\mathrm{m}$ Gifford and Xia outlet-end side channels. Block 3 Validators 2 and 3 confirm the interval-extension path by recovering targets at $50\,\mathrm{nm}$ and $50\,\mathrm{mm}$, outside the default interval on both sides; Validator 7 provides an independent-implementation cross-check against `scipy.optimize.brentq`.

4) Gifford-recursion stepper (Block 4–5): `advance_one_gap()` (Block 4) advances the Gifford recursion (Eqn 4 of [1]) by one gap row, inverting the resulting $R_s(i)$ via the Block 3 bisection to recover $w_s(i)$. The straight-device generator `generate_or_cif_straight()` (Block 5) chains these forward steps from the inlet to the outlet, producing the full $w_s(i)$ sequence for constant-w_c devices. Forward propagation matches the design pipeline: $w_s(0)$ is set at the inlet within fabrication constraints, f_{gap} is chosen to set the size cutoff, and the recursion produces the downstream $w_s(i)$ that the photomask realizes. Two per-step validity guards inside `advance_one_gap()` stop infeasible designs at trajectory generation rather than at the mask: the pressure-equilibration model limit $f_{\mathrm{gap}} < 0.5$ (beyond which the recursion's numerator becomes non-positive and no physical solution exists), and a strict-decrease check $R_s(i) < R_s(i-1)$ (which pressure equilibration with $f_{\mathrm{gap}} > 0$ requires; a violation indicates a numerical or algebraic error). Both raise `ValueError`.

5) Meander primitive (Block 6): At the small f_{gap} ($\sim 10^{-4}$) typical of blood-cell separation, the Gifford recursion demands first-row side-channel resistances that a straight rectangular channel achieves only at widths too narrow to fabricate reliably in PDMS soft lithography or to pass blood cells without clogging ($\sim 5\,\mu\mathrm{m}$ for Dinh V1). Xia 2016 [2] and Dinh 2024 [5, p. 4] substitute, for the narrow upstream segment of each side channel, a fixed-width serpentine of equivalent fluidic resistance; the center channel runs straight. Block 6 implements this as a closed-form length-for-resistance inverter, validated to machine precision against the analytic Yang formula [6] and consumed per row by the trajectory generator. For Dinh V1 the row-1 serpentine unrolls to $16.45\,\mathrm{mm}$, shortening row by row until below the $100\,\mu\mathrm{m}$ `min_unrolled_length` threshold, at which point each side channel transitions to a straight regime. The block's parametric layout outputs (leg count, leg length, turn radius, bounding box) are reserved for future Dinh- and Xia-style mask rendering.

6) Closed-form anchor (Block 12.6, Block 15.3): The closed-form identity (4) is implemented in Block 12.6 as a direct evaluation of the Yang series at the outlet w_c and w_s. Two independent design routes follow. The designer specifies either F — which fixes R_c/R_s and, at any chosen $w_{c,\mathrm{out}}$, determines $w_{s,\mathrm{out}}$ by inverting the Yang series via the Block 3 bisection inverter — or specifies $w_{c,\mathrm{out}}$ and $w_{s,\mathrm{out}}$ directly, in which case (4) returns F in closed form with no bisection step. Both routes replace iterating the recursion to convergence and tuning $w_s(i{=}1)$ until the endpoint matches with a single algebraic relation, anchoring the trajectory generator on outlet geometry alone. The Xia RS-7 [2], [3] row of Table I is the strongest non-trivial test of the identity's schedule-independence: f_{gap} varies by design at every step of the recursion yet provably drops out of the closed form, and the numerical residual at this device collapses to $|\Delta F| = 1.97 \times 10^{-5}$ (§8.4) — an order of magnitude below the worst-case constant-f_{gap} residual already reported for the trajectory generator.

7) PE-consistency diagnostic (§8.7): The closed-form identity (4) assumes pressure equilibration at the outlet, a condition that was not explicitly quantified in prior CIF studies [1], [2], [4], [5]. Under the hydraulic–electric analogy [13], the outlet slab can be approximated as a uniform linear Kirchhoff network, and the cross-width pressure imbalance $p_c(i) - p_s(i)$ decays approximately exponentially in row index. We estimate the corresponding half-life $k_{1/2}$ from a semi-log fit. We regard pressure equilibration as effectively satisfied when $k_{1/2}/N_{\mathrm{total}} < 0.05$, i.e., when the device contains at least 20 fitted half-lives, which implies a residual imbalance below 10^{-6} under the exponential model.

§8.7 implements this as a callable diagnostic `pe_consistency_check(wc_out, ws_out, d, gap_size, n_total, mu)`: a 30-row constant-(w_c, w_s) network of Yang-series resistors [6] (channel segments and gaps alike), with flow conservation at each node, solved by sparse direct factorization (`scipy.sparse.linalg.spsolve`). $k_{1/2}$ is extracted from the slope of a semi-log fit to the first 15 rows of $p_c(i) - p_s(i)$. A 30-row slab suffices because the decay rate is determined by the local outlet geometry, not by slab length.

We cross-check the 1-D network against 3-D Stokes simulation at the Dinh V1 outlet, where COMSOL Multiphysics 6.3 (P2+P1 Taylor–Hood, 115,904 elements, PARDISO solver; Block 19) gives $k_{1/2}^{\mathrm{COMSOL}} = 0.944$ gaps ($R^2 = 0.987$) by the same semi-log fit, agreeing with the network value $k_{1/2}^{\mathrm{network}} = 0.966$ gaps to within 2.3% (ratio 1.024×); Section IV-B reports the analogous comparison at the V3 outlet ($k_{1/2}^{\mathrm{COMSOL}} = 1.041$, $k_{1/2}^{\mathrm{network}} = 1.090$ gaps; agreement to 4.7%). The diagnostic uses no Dinh-specific physics, so generalisation to the other three reference devices rests on the local Yang-series network at each outlet. A 30-row slab suffices because the decay rate is determined by the local outlet geometry, not by slab length.

All five reference devices satisfy the criterion with safety margins of 48–84× (Table II); Xia RS-7 at 48× is the tightest. The diagnostic is exposed as part of the public API for use on future CIF designs.

TABLE II: PE-consistency diagnostic (§8.7) for the five reference devices. The diagnostic builds a 30-row constant-(w_c, w_s) Kirchhoff network at the outlet, estimates the PE half-life from a semi-log fit to the first 15 rows, and tests $k_{1/2}/N_{\mathrm{total}} < 0.05$. All devices pass with safety margins of 48–84×.

Device	$k_{1/2}$ (gaps)	N_{total}	$k_{1/2}/N_{\mathrm{total}}$	Safety margin
Gifford bead	2.074	3,486	5.95×10^{-4}	84×
Dinh V1	0.966	1,407	6.87×10^{-4}	72×
Dinh V2	1.013	1,486	6.82×10^{-4}	73×
Dinh V3	1.090	1,642	6.64×10^{-4}	75×
Xia RS-7	4.983	4,846	1.03×10^{-3}	48×

8) Reproducibility infrastructure: The notebook contains 62 executable Python cells (plus 83 markdown narrative cells) and 218 core validators across 19 sub-blocks (Blocks 1–6, 7a–d, 8–15, and 19; categories include geometry, manufacturability, PE convergence, numerical stability, regression checks, dimensional consistency, mask integrity, and COMSOL post-processing); all 218 pass. Five supplementary analyses in §8 add 19 further validators, all passing: the Block 16-alt fabrication-tolerance sweep (§8.1; V1–V8), the Xia RS-7 closed-form regression (§8.4 added in v3.11; V9), the channel-depth sensitivity sweep (§8.6 added in v3.20; six validators V11s–V13bs), the PE half-life network solve across all five reference outlets (§8.7 added in v3.20; V14s–V15s), and the effective-R_c/R_s reframing of the COMSOL output in matched units to the closed-form identity (§8.8 added in v3.20; V16s–V17s). The supplementary analyses §8.1, §8.6, and §8.8 inform the Results and Discussion sections rather than the methodology proper, and are described there. Nineteen regression stamps document deterministic execution (Table III reports the bit-identical and BLAS-portable subsets).

B. The COMSOL validation simulations

1) Why COMSOL, not the Python implementation, performs the empirical validation: The Python implementation realizes the one-dimensional CIF design framework (Yang resistance, Gifford recursion, OR-CIF generator, meander primitive, DXF mask export); a separate one-dimensional PE-consistency diagnostic (Section III-A7) tests gap-level pressure equilibration on a small Kirchhoff network. Neither layer runs

TABLE III: Validator and reproducibility-stamp infrastructure summary. Of the 19 stamps, 16 reproduce bit-identically across environments (12 match the v3.1 baseline, covering Blocks 1–9, 11, 13, and 14; 1 from Block 19 matches the v3.4 baseline introduced when the COMSOL post-processing block was added; and 3 from §8.4, §8.6, and §8.8 use payload-rounded inputs to be bit-portable across BLAS, counted as bit-identical). The remaining 3 stamps (Blocks 12 and 15 in the core; §8.7 in the supplementary) reproduce to 12+ decimal places under documented `scipy.sparse.linalg.spsolve` BLAS portability. The §8.1 Block 16-alt analysis additionally emits a deterministic stamp (`5313e1d66aa3ed9e`) but is excluded from the supplementary tally as a parametric sensitivity sweep of the closed-form identity (Block 15.3) rather than an independent methodological extension.

Module group	Validators (passing)	Stamps
Kernel (Blocks 1–6)	91 / 91	6
Trajectory generators (7a–d)	49 / 49	1
Worked Gifford example (Block 8)	17 / 17	1
2-D + Yang network (Block 12)	8 / 8	1
Device-scale synthesis (Block 15)	7 / 7	1
COMSOL post-processing (Block 19)	11 / 11	1
Other (Blocks 9, 10, 11, 13, 14)	35 / 35	4
Core total	**218 / 218**	**15**
§8.1 Block 16-alt fabrication sweep	8 / 8	—[a]
§8.4 Xia RS-7 closed-form (v3.11)	1 / 1	1
§8.6 Channel-depth sensitivity (v3.20)	6 / 6	1
§8.7 PE half-life, 5 outlets (v3.20)	2 / 2	1
§8.8 Effective R_c/R_s (v3.20)	2 / 2	1
Supplementary total	**19 / 19**	**4**
Combined total	**237 / 237**	**19**

[a] Emits stamp `5313e1d66aa3ed9e`; see caption.

finite-element computational fluid dynamics. Validation against the published Dinh 2024 measurements requires a three-dimensional Stokes solution at the outlet bifurcation, which is the domain of finite-element CFD software. We used COMSOL Multiphysics 6.3 with the CFD Module [7] as the established commercial tool. The COMSOL outputs (Q-flux tables, pressure-point tables, mesh quality, Newton iteration log) were transcribed verbatim into notebook Block 19 for post-processing — computing F_{COMSOL}, the σ-distance to the paper measurement, the PE half-life from a semi-log fit of the center–side pressure difference across five gaps, and the multi-variant ratio bracketing reported in Section IV-B.

2) Geometry: the 5-gap outlet slab: The Dinh 2024 V1 device has 1407 gaps total and the V3 device has 1642 gaps; both share a uniform middle-channel taper from $w_c = 75\,\mu\mathrm{m}$ at the inlet to $52\,\mu\mathrm{m}$ at the outlet across all three variants, as specified in the Dinh 2024 Figure 2 caption [5]. The per-variant outlet side-channel widths are $w_{s,\mathrm{out}} = 44\,\mu\mathrm{m}$ (V1, $f^*_{\mathrm{gap}} = 1.04 \times 10^{-4}$) and $w_{s,\mathrm{out}} = 57\,\mu\mathrm{m}$ (V3, $f^*_{\mathrm{gap}} = 1.28 \times 10^{-4}$); the V2 case at $w_{s,\mathrm{out}} = 48\,\mu\mathrm{m}$ is included in the closed-form verification table (Table I) but does not have a published F measurement.

We simulated five-gap outlet slabs of V1 and V3. Five gaps is sufficient because the analytical PE-consistency diagnostic (Section III-A7) predicts a PE half-life of approximately one gap for the Dinh V1/V3 outlet geometry — that is, the inlet pressure imbalance $|p_c(0) - p_s(0)|$ between center and side channels decays exponentially along the slab and halves every gap-pitch (the analog of a radioactive half-life, but in space along the device). By gap 5 the imbalance has therefore been attenuated by $\sim 2^5 = 32$-fold (five half-lives), leaving a residual imbalance of $\sim$3 % of its inlet-slab value — i.e., the slab captures asymptotic PE to $\sim$97 %. The COMSOL simulation itself confirms this prior estimate, measuring $k_{1/2} = 0.944$ gaps for V1 and $k_{1/2} = 1.041$ gaps for V3 (97.5 % and 96.4 % PE convergence at gap 5, respectively; quantitative comparison with the analytical prediction in Section IV-B). Whole-device simulation at the published 1407 / 1642 gaps is in any case computationally prohibitive at the mesh density required to resolve gap-region flow.

Half-domain symmetry.: The CIF geometry is mirror-symmetric about the long-axis center plane of the middle channel ($Q_s^{(\mathrm{left})} = Q_s^{(\mathrm{right})}$ and $\partial p/\partial n = 0$ on that plane). We therefore meshed the half-domain (one center-channel half plus one side channel) with a symmetry boundary condition on the mid-plane. The reported filtrate fraction is reconstructed from the half-domain Q-flux table as

$$F \equiv \frac{2\,Q_s}{Q_c^{(\mathrm{inlet,full})}} = \frac{Q_s^{(\mathrm{half})}}{Q_c^{(\mathrm{inlet,half})}},$$

the equality following from mirror symmetry ($Q_c^{(\mathrm{inlet,half})} = Q_c^{(\mathrm{inlet,full})}/2$ and $Q_s^{(\mathrm{half})} = Q_s$ per-side). The half-domain choice halves the cell

count for a fixed mesh quality and matches the convention used in the 5-gap COMSOL outlet reports archived alongside the notebook (see reproducibility infrastructure, Section III-A8).

3) Discretization, solver, and convergence: The production simulations transcribed into Block 19 used free tetrahedral meshes of 115,904 elements (V1) and 104,651 elements (V3), with minimum element quality 0.1136 and 0.1122 respectively, and average element quality 0.6617 (V1) and 0.6616 (V3). P2–P1 (Taylor–Hood) discretization satisfies the Ladyzhenskaya–Babuška–Brezzi inf–sup condition for stable Stokes flow. The PARDISO direct linear solver, driving COMSOL's damped-Newton iteration in the Stokes (creeping-flow) equations with "Neglect inertial term" enabled, converged in four iterations; the normalized solution update (COMSOL SolEst) reached 4.8×10^{-12} (V1) and 4.3×10^{-12} (V3) — a drop of more than thirteen orders of magnitude from the iter-1 SolEst of ~ 60 for both variants, well below COMSOL's default convergence tolerance. Mass conservation was verified to floating-point precision: V1 residual exactly zero at five significant figures; V3 residual -2.585×10^{-26} $\mathrm{m^3\,s^{-1}}$ ($|\text{residual}|/Q_{\text{in}} = 7.1 \times 10^{-17}$), at the floating-point precision limit.

Mesh convergence was confirmed by independent re-export runs at the same V1 mesh (115,904 elements) and a slightly different V3 mesh (104,206 elements; minimum quality 0.1328), and at approximately 5-fold coarser meshes (V1: 24,783 elements; V3: 21,448 elements). The integrated outlet flow values match the original-transcription production values to all five reported significant figures at both variants ($|\Delta F| < 5 \times 10^{-5}$), more than $600\times$ below the experimental $\sigma_{\text{paper}} \approx 0.033$–$0.036$ band (the true factor is larger but cannot be resolved at the precision the integrated flows are reported to). The four archived `.docx` reports (V1 and V3 at production and $\sim 5\times$ coarser meshes) are distributed alongside the notebook (Section III-A8).

4) Multi-variant ratio: three independent computations: Beyond the single-variant within-1σ test, we evaluate the multi-variant ratio $\ln(1-F_{\text{V3}})/\ln(1-F_{\text{V1}})$ under three independent computations: the COMSOL slab simulation, the Dinh 2024 measurement (with uncertainty propagated from the published $\sigma_{F_{\text{V1}}} = 0.033$ and $\sigma_{F_{\text{V3}}} = 0.036$ [5]), and the closed-form identity (Eq. 4). The closed-form gives the exact one-dimensional framework prediction under PE — the asymptotic answer with no residual physics. The paper gives the empirical measurement in the real device, with uncertainty propagated from the published per-variant standard deviations. The COMSOL slab sits between the two: it solves the three-dimensional Stokes equations on a finite (5-gap) outlet section, so it captures the two physical effects that distinguish the real device from the one-dimensional framework — the 3-D geometry of the outlet bifurcation, and the residual pressure imbalance from incomplete PE at the slab inlet. The three computed ratios, their pairwise σ-distances, and the mechanism attribution are reported in Sections IV-B and IV-C.

IV. Results and Discussion

The empirical case for the closed-form identity (4) is built in three stages. Section IV-A shows that (4) agrees with the converged Gifford recursion across all five reference devices in three architectures, validating the identity as the algebraic limit of the recursion. Section IV-B shows that three-dimensional Stokes simulation of the Dinh 2024 V1 and V3 outlet geometries in COMSOL Multiphysics 6.3 recovers the published filtrate fractions within 1σ at both variants, validating (4) as a physical prediction of the filtrate fraction. Section IV-C attributes the residual $\sim 10\%$ closed-form-vs-paper offset to a bounded mechanism class (3-D outlet bifurcation plus PDMS fabrication shrinkage, with channel-depth tolerance and finite-PE bounded out) without committing to a unique decomposition. Section IV-D then presents the worked Gifford 2014 mask example as end-to-end fabricability evidence; Section IV-E discusses feasibility, benchmarking, and applications; and Section IV-F enumerates limitations and future work.

A. Closed-form identity holds across the published CIF lineage

Under PE and the architectural boundary condition $Q_s(0) = 0$, the closed-form identity (4) and the converged Gifford recursion predict the same outlet F from the same outlet geometry. If (4) is exact under these premises, it should reproduce the recursion's converged F at every reference device within the recursion's own convergence floor. Calling `generate_cif_device()` once per device on the five reference geometries spanning the three architectures of the published CIF lineage — Gifford bead (constant-w_c, constant-f_{gap}); Dinh V1, V2, V3 (tapered-w_c, constant-f_{gap}); Xia RS-7 (tapered-w_c, variable-f_{gap}) — produces $F_{\text{gen}} = 0.9065,\ 0.5594,\ 0.6169,\ 0.7190,\ 0.8624$ respectively (Table I; Fig. 2). The deviations from (4), evaluated at the same outlet geometries, satisfy $|F_{\text{closed}} - F_{\text{gen}}| \leq 2.35 \times 10^{-4}$ across the four constant-f_{gap} devices (Block 15.3) and 1.97×10^{-5} for the variable-f_{gap} Xia case (§8.4). The agreement reaches each device's recursion-convergence floor, with the difference too small to resolve at the scale of Fig. 2.

The variable-f_{gap} is the strongest non-trivial empirical test of the identity's schedule-independence claim: the Xia 2022 corrected scaling $f_{\text{gap}}(i) =$

$(w_{c,\mathrm{ref}}/w_c(i))^2\, f_{\mathrm{gap}}^{\mathrm{ref}}$ [2], [3] makes $f_{\mathrm{gap}}(i)$ depend on $w_c(i)$ at every step of the recursion, structurally distinct from the constant-f_{gap} Dinh and Gifford cases, yet f_{gap} provably drops out of the closed form at the outlet, which is what the 1.97×10^{-5} Xia residual confirms numerically. This is the empirical companion to the structural argument from Section II: a current divider's output ratio is determined by the resistances at the divider node and is invariant under any upstream rearrangement of the network, including any choice of $f_{\mathrm{gap}}(i)$ between gap rows 0 and N. The agreement is conditional on PE, whose empirical confirmation across all five reference outlets and at the Dinh V1/V3 geometries is taken up next (Section IV-B; Section III-A7).

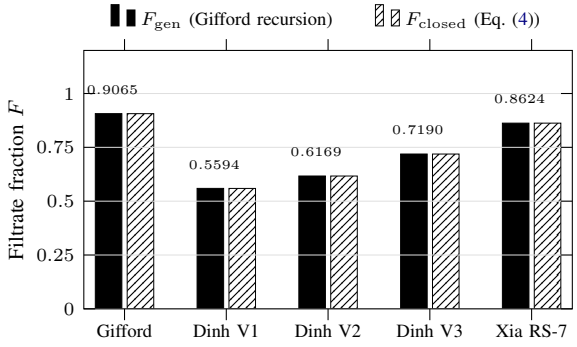

Fig. 2: Reference-device validation of the trajectory-generator module. Filtrate fractions F_{gen} predicted by `generate_cif_device()` for the five reference devices in the closed-form regression spanning three architectural regimes (Gifford bead: constant-w_c constant-f_{gap}; Dinh V1, V2, V3: tapered-w_c constant-f_{gap}; Xia RS-7: tapered-w_c variable-f_{gap}) compared against the closed-form identity (4). Solid black bars show F_{gen} (values annotated above); hatched bars show F_{closed}, which agrees with F_{gen} to $|\Delta F| \leq 2.35 \times 10^{-4}$ across the four constant-f_{gap} devices (Block 15.3) and $|\Delta F| = 1.97 \times 10^{-5}$ for the variable-f_{gap} Xia RS-7 case (§ 8.4); the difference is too small to resolve at this scale. All five trajectories produced from a single `generate_cif_device()` function call per device.

B. Three-dimensional Stokes simulation reproduces measured filtrate fraction within reported experimental 1σ at Dinh outlets

Three-dimensional Stokes simulation in COMSOL Multiphysics 6.3 [7] of the Dinh 2024 V1 ($w_{s,\mathrm{out}} = 44\,\mu\mathrm{m}$) and V3 ($w_{s,\mathrm{out}} = 57\,\mu\mathrm{m}$) outlet geometries on the five-gap slabs of Section III-B yields $F_{\mathrm{COMSOL}} = 0.5070$ (V1) and 0.6281 (V3). These fall within the reported experimental 1σ uncertainty of the Dinh 2024 volume-recovery measurements $F_{\mathrm{paper,V1}} = 0.505 \pm 0.033$ and $F_{\mathrm{paper,V3}} = 0.649 \pm 0.036$ [5, Fig. 4B] at σ-distances $+0.06$ (V1) and -0.58 (V3) respectively (Fig. 3a; Table IV). Including 3-D finite-slab Stokes physics thus reduces the closed-form-vs-paper offset at V1 from $+1.64\sigma$ to $+0.06\sigma$ and at V3 from $+1.94\sigma$ to -0.58σ. Mesh convergence, mass conservation, and Newton residuals are documented in Section III-B. The mesh-independence check is the load-bearing piece for what follows: re-running each variant at approximately five-fold coarser meshes leaves F unchanged to all five reported significant figures, bounding the discretization residual at $|\Delta F| < 5 \times 10^{-5}$ — more than $600\times$ below the experimental σ band. The 3-D Stokes correction reported above is therefore a real consequence of the outlet-bifurcation flow, not a numerical artifact, and mesh discretization does not contribute to the residual decomposition of Section IV-C.

The multi-variant ratio $\ln(1 - F_{\mathrm{V3}})/\ln(1 - F_{\mathrm{V1}})$ provides a complementary aggregate view of the same comparison (Fig. 3c; Table IV). It takes three values: 1.398 from COMSOL, 1.489 ± 0.203 from the paper (error propagation from independent $\sigma_{F_{\mathrm{V1}}} = 0.033$ and $\sigma_{F_{\mathrm{V3}}} = 0.036$), and 1.549 from the closed-form identity. The three values agree within the propagated experimental uncertainty: the closed-form result is 0.30σ above the paper central value, and COMSOL is 0.45σ below it (Fig. 3c). The closed-form is the 1-D asymptotic limit under PE; it omits 3-D outlet bifurcation physics, which the COMSOL finite-slab Stokes model includes on the nominal geometry. The closed-form lies above the data and COMSOL lies slightly below it, forming the bracketing pattern expected when the closed-form omits a correction that COMSOL captures. For the ratio statistic, closed-form and COMSOL differ by only 0.75σ, which is small relative to the propagated uncertainty of ± 0.203, and the paper's central value lies within the same bracket. Section IV-C attributes this correction to 3-D outlet bifurcation via per-variant resistance-ratio decomposition.

We performed two further diagnostics on the COMSOL output. The first asks what resistance ratio R_c/R_s the closed-form identity would need in order to reproduce COMSOL's F. Inverting (4) gives the answer $(R_c/R_s = F/[2(1 - F)])$ — the *apparent* resistance ratio implied by F_{COMSOL}. Dividing it by the 1-D Yang prediction at the same outlet geometry defines α_{3D} — the factor by which the 3-D finite-slab Stokes physics reduces the identity's effective resistance ratio. We find $\alpha_{\mathrm{3D}} = 0.811$ at V1 (an 18.9% reduction) and 0.660 at V3 (a 34.0% reduction), a 15.1-percentage-point difference between the two variants (notebook §8.8). The two variants share the same outlet w_c and channel depth and differ only in w_s, so the 15.1-percentage-point asymmetry cannot be explained by any mechanism that

shifts F uniformly across variants, including isotropic in-plane fabrication scaling and flow-rate calibration error. The asymmetry therefore points to a geometry-dependent mechanism, which Section IV-C identifies as 3-D outlet bifurcation and for which alternative candidates are bounded.

The second diagnostic asks whether the pressure-equilibration assumption behind the identity is physically relevant in the real 3-D flow. At the V1 outlet, the PE half-life extracted from the COMSOL pressure data is $k_{1/2}^{\mathrm{COMSOL}} = 0.944$ gaps; from an independently constructed 1-D Kirchhoff resistor network of the same geometry (Section III-A7) it is $k_{1/2}^{\mathrm{network}} = 0.966$ gaps — a 2.3% agreement between two structurally different models of the same flow (Fig. 3b). The V3 outlet shows the same pattern, with COMSOL (1.041 gaps) and the network (1.090 gaps; Table II) agreeing to 4.7%. This cross-method agreement indicates that the PE-consistency picture is a real feature of the flow, not merely a property of the 1-D model. Because the equilibration half-life is approximately one gap at the Dinh outlets (Methods Table II), pressure equilibration remains effectively complete across the full 1407-gap (V1) and 1642-gap (V3) devices. This supports the extrapolation from the 5-gap COMSOL slab to the full-device used in Section IV-C.

C. Mechanistic interpretation of the 10% closed-form-vs-paper residual

(a) 3-D outlet bifurcation. Three-dimensional Stokes simulation at the nominal Dinh V1 and V3 outlet geometries returns filtrate fractions 9.3% and 12.6% below the closed-form predictions respectively (Section IV-B). Both COMSOL values land inside the experimental 1σ band ($+0.06\sigma$ at V1, -0.58σ at V3; Table IV), down from $+1.64\sigma$ and $+1.94\sigma$ for the closed-form alone. Adding the 3-D outlet physics that the closed-form omits is therefore sufficient to account for the residual at both variants.

The closed-form's structure makes plain where this correction matters most. The closed-form is a 1-D parallel-resistor approximation that splits flow using only the local resistance ratio R_s/R_c. That approximation is most accurate when one branch strongly dominates the split: if R_s/R_c is far from unity, almost all flow goes one way, and the detailed shape of the junction matters little. At the Dinh outlet, however, $R_s/R_c = 1.58$ for V1 and 0.78 for V3, so the center and side branches carry comparable flow. In this regime the 1-D treatment does not describe the partition completely and a 3-D treatment is needed.

Applied gap-by-gap, the same reasoning explains why the correction is largest near the outlet rather than spread through the device. Upstream of the outlet the local resistance ratio is much larger: $R_s/R_c \approx 30$ at the start of the straight section. At those gaps the 1-D rule is already close, and the missing 3-D physics matters little. The outlet is where the ratio first approaches unity, and so where the missing 3-D physics matters most.

The asymmetry in $\alpha_{3\mathrm{D}}$, 0.811 at V1 versus 0.660 at V3 (a 15.1-percentage-point spread; Table IV, notebook §8.8), is consistent with this picture. V3 has the smaller outlet R_s/R_c and the wider w_s, and shows the larger correction. We report this as a correlation, not a derivation: two variants are not enough to separate the effect of w_s from the effect of R_s/R_c.

A supplementary sparse-mesh Python Stokes check (§8.2) at a related discretized geometry ($w_{c,\mathrm{eff}} = 48\,\mu\mathrm{m}$, $w_s = 44\,\mu\mathrm{m}$, rather than V1's nominal $52\,\mu\mathrm{m}/44\,\mu\mathrm{m}$) returns a 4.7% reduction in the same direction. Because the geometry, mesh, and implementation differ from the production COMSOL run, this is a sign-and-magnitude consistency check, not a quantitative reproduction. Ultimately, the COMSOL results show that mechanism (a) alone places both variants within 1σ at the nominal geometry, and the asymmetry emerges without any variant-specific fitting.

(b) w_s-specific dimensional offset. Realized PDMS channels can be slightly wider or narrower than the design mask, so we tested whether such an offset could explain the residual. A single-parameter sweep applies a fractional offset s to $w_{s,\mathrm{end}}$, holding $w_{c,\mathrm{end}}$ fixed at the designed $52\,\mu\mathrm{m}$ (Block 16-alt; Fig. 4). The joint optimum sits at $s^* = 9.29\,\%$ ($\max|\sigma| = 0.39$), and both variants stay within 1σ over $s \in [5.89\,\%, 11.86\,\%]$. Per-variant inversion of (4) gives $7.6\,\%$ at V1 and $11.4\,\%$ at V3 — a $3.8\,\mathrm{pp}$ spread.

Two facts argue against the dominant published fabrication subcategories associated with (b).

Magnitude. A notable source of dimensional change in soft lithography is uniform in-plane PDMS shrinkage during cure: the PDMS contracts laterally as it cures, but the channel depth d — set by the SU-8 master under PDMS bonded to a rigid substrate — is geometrically constrained. The shrinkage therefore rescales only the in-plane dimensions, with each width w and each channel length L scaling by a common factor $k = 1 - s$ (where s is the fractional shrinkage), while d is unchanged. The per-pitch Yang resistance for a rectangular duct of cross-section $w \times d$ (here $d > w$) at length L and viscosity μ is [6]

$$R(w, d, L) \;=\; \frac{12\mu L}{d\, w^3\, g(w/d)}, \tag{7}$$

where $g(w/d)$ is the rectangular-duct aspect-ratio correction. Substituting the scaling and expanding $(kw)^3 =$

TABLE IV: Empirical validation at the Dinh 2024 V1 and V3 outlets: COMSOL 3-D Stokes simulation summary (5-gap slab), with F results, PE half-life, and the multi-variant ratio under three independent computations. Both COMSOL σ-distances are within 1σ; mass conservation is verified to floating-point precision. Detailed discussion in Section IV-B; mechanism interpretation of the closed-form-vs-paper offset in Section IV-C.

Quantity	V1 ($w_{s,\mathrm{out}} = 44\,\mu\mathrm{m}$)	V3 ($w_{s,\mathrm{out}} = 57\,\mu\mathrm{m}$)
Mesh elements	115,904	104,651
Min element quality	0.1136	0.1122
Newton iterations	4	4
SolEst (final)	4.8×10^{-12}	4.3×10^{-12}
F_{COMSOL}	0.5070	0.6281
$F_{\mathrm{paper}} \pm \sigma$	0.505 ± 0.033	0.649 ± 0.036
σ-distance	$+0.06\sigma$	-0.58σ
$\alpha_{3\mathrm{D}}$ (§8.8)[a]	0.811	0.660
PE half-life $k_{1/2}$ (gaps)[b]	0.944	1.041
Δp decay at gap 5 (fit)	97.5%	96.4%
Multi-variant ratio $\ln(1-F_{\mathrm{V3}})/\ln(1-F_{\mathrm{V1}})$		
Closed-form (Eq. 4)	1.549	
Paper (propagated uncertainty)	1.489 ± 0.203	
COMSOL (Section III-B)	1.398	

[a] $\alpha_{3\mathrm{D}} = (R_c/R_s)_{\mathrm{apparent}} \,/\, (R_c/R_s)_{\mathrm{Yang}}$, where $(R_c/R_s)_{\mathrm{apparent}} = F_{\mathrm{COMSOL}}/[2(1-F_{\mathrm{COMSOL}})]$ is recovered by inverting (4) and $(R_c/R_s)_{\mathrm{Yang}}$ is evaluated by the Yang series [6] at the same outlet geometry. Variant asymmetry $|\Delta\alpha_{3\mathrm{D}}| = 15.1$ pp.
[b] At V1, the COMSOL half-life is cross-validated against the 1-D Kirchhoff network of Section III-A7 ($k_{1/2}^{\mathrm{network}} = 0.966$ gaps, Table II; network/COMSOL ratio $1.024\times$, agreement to 2.3%).

$k^3 w^3$ gives

$$R(kw, d, kL) \;=\; \frac{12\mu(kL)}{d\,k^3 w^3\, g(kw/d)} \;=\; \frac{1}{k^2}\, R(w, d, L) \cdot \frac{g(w/d)}{g(kw/d)}$$

The $1/k^2$ factor arises from k/k^3 (one factor of k from the numerator L, three from the denominator w^3). It depends only on the scaling fraction k, not on d, and is identical for the center and side channels — so it cancels exactly in the ratio R_s/R_c. The aspect-ratio correction g, however, does not cancel: its argument shifts from w/d to kw/d because the in-plane width shrinks while the depth stays fixed. Whether F in (4) is invariant under this scaling therefore reduces to whether the residual g-shift is small. Direct evaluation of (4) at $s = 10\%$ (i.e., $k = 0.9$) gives $|\Delta F/F| \approx 0.25\,\%$ at V1 and $0.10\,\%$ at V3 — both far below the $\sim 10\,\%$ residual, even though $s = 10\%$ already exceeds the published Sylgard 184 range of $1.06\,\%$–$2.46\,\%$ across nine curing recipes [16].

Asymmetry. The sweep finds V3 demanding a *larger* fractional offset than V1 ($11.4\,\%$ vs. $7.6\,\%$), even though V3 starts from a *wider* w_s ($57\,\mu\mathrm{m}$ vs. $44\,\mu\mathrm{m}$). Two common fabrication effects could in principle produce a w_s-dependent offset, and neither produces this V3>V1 ordering. *Isotropic shrinkage*, by definition, gives the same fractional offset at V1 and V3 — it cannot generate *any* V1≠V3 asymmetry. *Aspect-ratio-dependent shrinkage*, in which narrower features shrink more in fractional terms, predicts the *opposite* ordering: V1>V3, since V1 has the narrower w_s. The 3-D outlet bifurcation of mechanism (a), in contrast, *does* produce V3>V1: the wider w_s at V3 places the outlet bifurcation deeper into the 3-D regime, where $\alpha_{3\mathrm{D}}$ departs further from unity (0.660 at V3 vs. 0.811 at V1; Table IV).

The Block 16-alt sweep therefore bounds the tolerable w_s-specific dimensional offset at $\sim 9\,\%$ as a sensitivity check. The two named fabrication subcategories examined here cannot account for the observed V3>V1 asymmetry; whether other fabrication mechanisms contribute cannot be determined from the present analysis without direct metrology of the realized $w_{s,\mathrm{out}}$.

(c) Channel-depth tolerance is bounded out by the §8.6 sweep. The closed-form identity (4) depends on the channel depth d through the Yang aspect-ratio correction [6]; we test whether realistic PDMS depth variation can account for the residual. The §8.6 sweep evaluates F_{closed} at $d \in \{120, 130, 140, 150, 160\}$ μm around the nominal $140\,\mu\mathrm{m}$. Over the typical $\pm 5\,\mu\mathrm{m}$ PDMS depth tolerance, F shifts by 0.0010 at V1 and 0.00055 at V3 — that is, by $0.031\,\sigma_{\mathrm{paper}}$ ($\sigma_{\mathrm{paper}} = 0.033$) and $0.015\,\sigma_{\mathrm{paper}}$ ($\sigma_{\mathrm{paper}} = 0.036$) respectively. Both shifts are more than an order of magnitude below either variant's experimental σ and at least an order of magni-

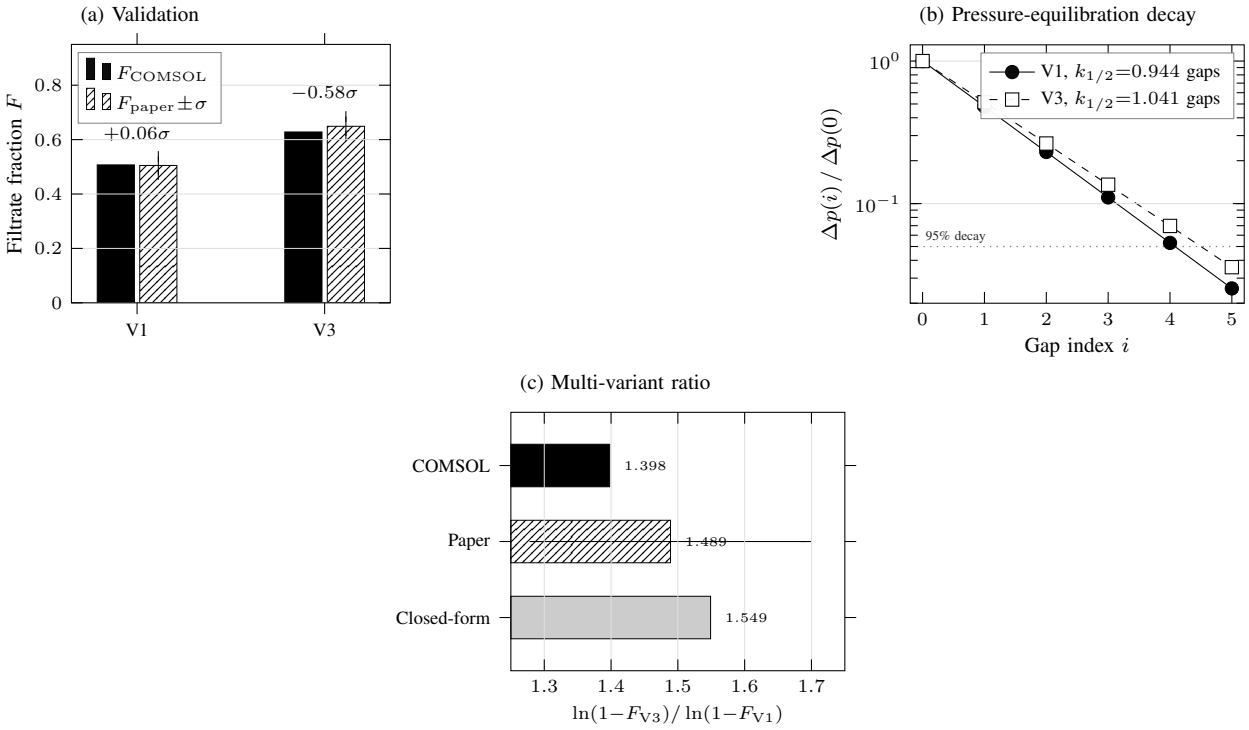

Fig. 3: Empirical validation against Dinh 2024 [5]. (a) Filtrate fractions predicted by COMSOL 3-D Stokes simulation of the V1 and V3 outlet geometries (5-gap slab) plotted against the published measurements; both within 1σ ($+0.06\sigma$ V1, -0.58σ V3), down from the closed-form-vs-paper offsets of $+1.64\sigma$ (V1) and $+1.94\sigma$ (V3) at the same outlet geometries (Section IV-C). (b) Pressure-equilibration decay across the 5-gap slab; semi-log fit gives PE half-life $k_{1/2} = 0.944$ gaps (V1, $R^2 = 0.987$) and 1.041 gaps (V3, $R^2 = 0.986$). The V1 value is cross-validated against the 1-D Kirchhoff network of Section III-A7 ($k_{1/2}^{\mathrm{network}} = 0.966$ gaps, Table II; agreement to 2.3%, ratio $1.024\times$). (c) Multi-variant ratio $\ln(1-F_{\mathrm{V3}})/\ln(1-F_{\mathrm{V1}})$ under three independent computations: 1.398 (COMSOL), 1.489 ± 0.203 (paper, with horizontal error bar showing uncertainty propagated from $\sigma_{F_{\mathrm{V1}}} = 0.033$ and $\sigma_{F_{\mathrm{V3}}} = 0.036$), and 1.549 (closed-form). All three values are mutually consistent within the propagated experimental uncertainty (closed-form 0.30σ above the paper central value; COMSOL 0.45σ below). Markers in panel (b) at integer gap indices show exponential-fit values $\Delta p(i)/\Delta p(0) = 2^{-i/k_{1/2}}$, not raw COMSOL output; the underlying $\Delta p(1 \ldots 5)$ values were fitted by semi-log linear regression to extract $k_{1/2}$ and are reported in the released notebook (Block 19).

tude below the $\sim$10 % closed-form-vs-paper residual that requires explanation. Channel-depth tolerance therefore cannot be a major contributor.

(d) Finite-PE residual is bounded out at full-device scale. The PE half-life diagnostic (Section III-A7) clears finite-PE as a major contributor whenever $k_{1/2}/N_{\mathrm{total}}0.05$: the device then contains at least 20 half-lives and the residual cross-width pressure imbalance falls below $2^{-20} \approx 10^{-6}$. The Dinh devices clear this threshold by more than an order of magnitude. With $k_{1/2} = 0.966$ gaps over $N = 1407$ (V1) and 1.090 gaps over 1642 (V3), Δp decays over 1457 and 1506 half-lives respectively (Table II). The bound generalises: all five reference devices satisfy the threshold with safety margins of 48–84$\times$ (Xia RS-7 tightest, Gifford bead loosest).

Two checks corroborate the bound. First, the 1-D diagnostic agrees with 3-D Stokes at the V1 outlet to 2.3% ($k_{1/2}^{\mathrm{1D}} = 0.966$ vs $k_{1/2}^{\mathrm{COMSOL}} = 0.944$ gaps; Section IV-B, Table IV footnote b). This cross-check is single-point; extending it to the other four devices relies on the local Yang-series physics being the same at each outlet. Second, the 1-D over-estimate of $k_{1/2}$ corresponds to a 2.3% under-estimate of the decay rate, so the 48–84$\times$ safety margins are conservative lower bounds on the true margins.

Finite-PE therefore cannot be a major contributor at the device scale at which Dinh's measurements were made.

Conclusion and scope. The four-candidate analysis

identifies 3-D outlet bifurcation as a sufficient physical explanation of the residual. COMSOL at the nominal outlet geometry alone places both variants within 1σ of the published measurements ($+0.06\sigma$ at V1, -0.58σ at V3, Table IV); the geometry-dependent character of the 3-D correction, captured by α_{3D}, accounts for the asymmetric V1/V3 response; and the three alternative candidates — w_s-specific dimensional offset (b), channel-depth tolerance (c), and finite-PE residual (d) — are individually smaller in magnitude (c, d) or, in the case of (b), produce an asymmetric V1/V3 response inconsistent with conventional fabrication mechanisms. Direct metrology of the realized $w_{s,\mathrm{out}}$ in fabricated devices would tighten the bound on (b) further and let any contribution outside the 3-D Stokes prediction be quantified directly; we identify this as future work in Section IV-F.

D. Design rationale: worked Gifford 2014 examples

The reference implementation operates at two scales: a one-dimensional trajectory generator (Section III) that produces an F prediction and a $w_s(i)$ sequence from any Gifford-recursion-compatible parameter set, and a mask-layout module that renders the resulting design as a fabricable DXF. We illustrate each with a Gifford 2014 worked example: the trajectory-generator level by applying the closed-form identity (4) to the Gifford 2014 platelet enrichment device (Fig. 6 of [1]), and the mask-layout level by taking the Gifford 2014 bead device (Fig. 4 of [1]) through to a fabricable photomask.

At the trajectory-generator level, the library reproduces F predictions across all five reference devices spanning three architectures from a single `generate_cif_device()` call (Section IV-A, Fig. 2; Table I): the Gifford bead device, Dinh V1/V2/V3, and Xia RS-7, with all five deviations from the closed-form identity satisfying $|\Delta F| \le 2.35 \times 10^{-4}$ for the four constant-f_{gap} devices and $|\Delta F| = 1.97 \times 10^{-5}$ for the variable-f_{gap} Xia case. To concretize the f_{gap}-vs-F design independence that (4) formalizes, we evaluate it at a second Gifford geometry not used elsewhere in the paper. Dinh's V1 design value $f^{*}_{\mathrm{gap}} = 1.04 \times 10^{-4}$ numerically matches the f_{gap} chosen for Gifford 2014's full-length platelet enrichment device [1, Fig. 6 caption] ($w_c = 125\,\mu\mathrm{m}$, $d = 150\,\mu\mathrm{m}$, $w_{s,\mathrm{out}} = 140\,\mu\mathrm{m}$; the Gifford 2014 main text states $144\,\mu\mathrm{m}$, but the Fig. 6 caption value of $140\,\mu\mathrm{m}$ is the more internally consistent of the two values with the caption's reported "$\sim 2.5\times$" side-to-center flow ratio under (4)). Evaluating the Yang series (Eqn 5 of [1], truncated at the same $N = 50$ terms used there) at this geometry with $L = 50\,\mu\mathrm{m}$ and $\mu = 1\,\mathrm{mPa\,s}$ gives per-pitch resistances $R_c = 4.11 \times 10^{9}\ \mathrm{Pa \cdot s/m^3}$ and $R_s = 3.23 \times 10^{9}\ \mathrm{Pa \cdot s/m^3}$ at $w_{s,\mathrm{out}} = 140\,\mu\mathrm{m}$; L and μ set the absolute magnitudes but cancel in the dimensionless ratio R_c/R_s, so any consistent choice yields the same F. (4) then gives $F^{\mathrm{F6}}_{\mathrm{closed}} = 0.7177$ and a side-to-center flow ratio $F/(1-F) = 2.54$, within $1.7\,\%$ of the caption's "$\sim 2.5\times$"; repeating at the main-text value $w_{s,\mathrm{out}} = 144\,\mu\mathrm{m}$ gives $F = 0.7292$ and ratio 2.69, within $7.7\,\%$, confirming the caption value as the more self-consistent choice. The identity thus arbitrates a within-source inconsistency, selecting the $140\,\mu\mathrm{m}$ caption value over the $144\,\mu\mathrm{m}$ main-text value as the design dimension implied by the caption's stated $\sim 2.5\times$ flow ratio. Despite identical f_{gap}, the closed-form identity predicts markedly different F at the two outlet geometries: $F^{\mathrm{F6}}_{\mathrm{closed}} = 0.7177$ at the Gifford Fig. 6 outlet versus $F^{\mathrm{V1}}_{\mathrm{closed}} = 0.5592$ at the Dinh V1 outlet (Table I), a difference of $|\Delta F| = 0.159$, or about $4.8\,\sigma_{\mathrm{paper,V1}}$ ($\sigma_{\mathrm{paper,V1}} = 0.033$). The pair therefore serves as a worked example showing that f_{gap} separately sets the size cutoff (calibrated empirically by Gifford 2014 across a 33-device test array spanning $f_{\mathrm{gap}} \in [6.40 \times 10^{-5},\ 5.76 \times 10^{-4}]$; Gifford 2014 Fig. 3) while F is determined by outlet geometry alone.

At the mask-layout level, the worked Gifford 2014 example takes the trajectory through to a fabricable DXF (`gifford_2014_bead_device.dxf`; deterministic regression stamp `aa5833cdb554e9b5`) whose post pattern reproduces Gifford 2014 Figure 4 [1] within pixel-measurement precision (Fig. 5, Fig. 6). Calibrating the figure scale against the paper-stated center-channel width ($w_c = 100\,\mu\mathrm{m}$, scale-bar agreement $0\,\%$) gives a pixel size of $2.564\,\mu\mathrm{m}$; at that scale the FFT-measured pitch is $46.4\,\mu\mathrm{m}$, the post length along flow is $L_{\mathrm{post}} = 25.6\,\mu\mathrm{m}$, the post width perpendicular to flow is $W_{\mathrm{post}} = 50.6\,\mu\mathrm{m}$, and the inter-post gap is $w_{\mathrm{gap}} = 20.5\,\mu\mathrm{m}$. The implied pitch $L_{\mathrm{post}} + w_{\mathrm{gap}} = 46.1\,\mu\mathrm{m}$ agrees with the FFT-measured pitch to $0.3\,\mu\mathrm{m}$ ($\sim 12\,\%$ of one $2.564\,\mu\mathrm{m}$ pixel), nearly an order of magnitude tighter than the per-feature pixel resolution. At the implementation pitch $L_{\mathrm{post}} + w_{\mathrm{gap}} = 46.1\,\mu\mathrm{m}$, the 3,486-row CIF array unrolls to $160.7\,\mathrm{mm}$ across four $40.18\,\mathrm{mm}$ legs (the FFT-measured $46.4\,\mu\mathrm{m}$ pitch would give $161.8\,\mathrm{mm}$) and occupies $87\,\%$ of each leg's length on a $46.2 \times 5.7\ \mathrm{mm^2}$ chip (total leg length $4 \times 46.2\,\mathrm{mm} = 184.8\,\mathrm{mm}$, of which $13\,\%$ is reserved for U-turns, ports, and alignment marks; chip footprint computed from the layout; Gifford 2014 does not state a chip footprint); side channels widen monotonically from $20\,\mu\mathrm{m}$ at the inlet to $300.08\,\mu\mathrm{m}$ at the outlet, matching the paper's specified $300\,\mu\mathrm{m}$ to $0.027\,\%$. Block 8 emits seventeen validators covering polygon counts, port count and disjointness, stadium post shape and convexity, U-turn arc count and direction, alignment-mark count, DXF integrity, and the geometric-closure check (V17): all seventeen are consistent.

The reference implementation notebook therefore en-

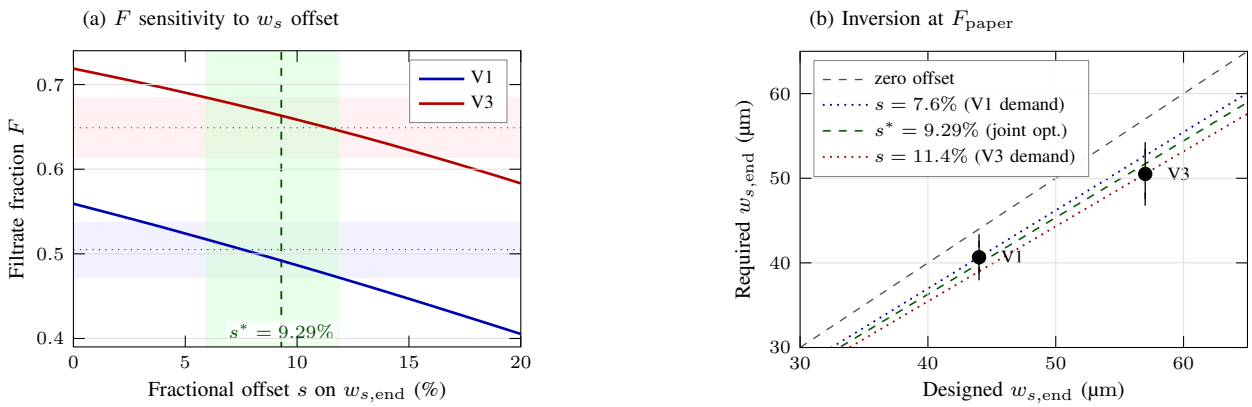

Fig. 4: Side-channel-width sensitivity bound from notebook Block 16-alt. The sweep parameter s scales $w_{s,\mathrm{end}}$ only, with $w_{c,\mathrm{end}}$ fixed at its designed 52 µm; it is a sensitivity check, not a model of isotropic PDMS thermal shrinkage (see Section IV-C (b)). *(a)* Filtrate fraction F versus s at the Dinh V1 and V3 outlets, with paper $F \pm \sigma$ bands shaded (blue: V1; red: V3). Joint optimum $s^* = 9.29\%$ minimises $\max(|\sigma_{V1}|, |\sigma_{V3}|)$; both variants stay within 1σ over $s \in [5.89\%, 11.86\%]$ (green band). *(b)* Required versus designed $w_{s,\mathrm{end}}$ under the hypothesis that the entire residual comes from a uniform fractional offset on w_s alone. For each variant we invert Eq. (4) to find the $w_{s,\mathrm{end}}$ that would reproduce F_{paper} (black markers); error bars are the w_s range corresponding to $F_{\mathrm{paper}} \pm \sigma_F$. Both markers sit below the zero-offset diagonal by more than σ_F ($+1.64\sigma$ at V1, $+1.94\sigma$ at V3), the same closed-form overshoot reported in Section IV-B but in w_s-offset units. Blue and red dotted lines pass through the V1 and V3 central values respectively, at $s = 7.6\%$ and 11.4%; the green dashed line marks the panel-(a) joint optimum $s^* = 9.29\%$. The two demanded offsets differ by 3.8 percentage points and increase with the designed w_s — the V3>V1 asymmetry pattern that Section IV-C (b) identifies as inconsistent with conventional fabrication mechanisms.

codes the Gifford-recursion framework as a callable module covering the published CIF design space and produces the Gifford 2014 bead device photomask (Fig. 4 of [1]) end-to-end as a worked example. The mask-layout module currently renders Gifford-style topologies only; multi-stage tapers (Dinh), recycling stages (Xia), and multiplexed parallel elements (Lezzar) are roadmap items at the mask-layout level (Section IV-F). The trajectory-generator stage already handles all three architectures, so the limitation is at the layout-rendering step, not the fluidic-design step that the closed-form identity anchors.

E. Discussion: feasibility, benchmarking, and applications

The reference notebook eases the re-implementation burden previously faced by any researcher seeking to apply, extend, audit, or build upon the published Gifford-variant CIF design framework from primary-literature equations alone. The trajectory-generator stage operates for any Gifford-recursion-compatible design across all three architectures (Table I); the mask-layout stage produces a fabricable DXF for Gifford-style devices end-to-end. The four limitations enumerated in Section I — absent closed-form anchor, absent PE-consistency verification, absent empirical-validation infrastructure tied to the framework, and absent open-source reference implementation — are addressed by, respectively, the closed-form identity (Section II), the PE-consistency diagnostic (Section III-A7), the COMSOL workflow and four-candidate mechanism analysis (Sections IV-B and IV-C), and the open-source release with its 237-validator reproducibility suite (Section III, Table III). The worked Gifford 2014 mask example (Section IV-D, Figs. 5 and 6) demonstrates this end-to-end workflow — parameter specification, trajectory generation, PE-consistency verification, manufacturability checks, and DXF mask export — in a single `generate_cif_device()` call, an optional `pe_consistency_check()` callable, and a single mask-layout pass, with no manual recursion-solving step required.

Within the broader continuous-flow microfluidic design-automation (MFDA) landscape, generic open and academic tools — 3DμF [8], μFluidicGenius [9], Flui3d [10], OpenMFDA [11], among others — support general layout, component-library composition, channel routing, and photomask export, but none encodes the Gifford recursion, the Yang resistance scaling, or PE-consistency verification; each requires the user to manually compute a CIF gap schedule before being usable for CIF-specific layout. To our knowledge, no prior open-source implementation of the Gifford-variant CIF design framework exists. The present implementation therefore

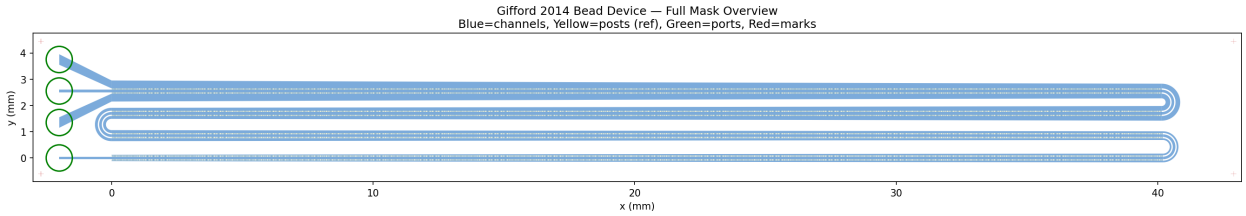

Fig. 5: Worked Gifford 2014 mask example produced by the reference implementation (Block 8): full-device overview from the `plot_mask_overview` function. The four-leg paperclip topology connects one inlet port (upper-left, at $y \approx 4\,\mathrm{mm}$, recognisable by the fanout wedge that distributes flow from the single inlet into the three-channel CIF array) through four $40.2\,\mathrm{mm}$-long pillar arrays and three post-free semicircular U-turns to three outlet ports (left side, at $y \approx 0$, 1.3, and $2.5\,\mathrm{mm}$: bottom side-channel outlet, center-channel outlet, top side-channel outlet). Computed-from-layout chip footprint $46.2 \times 5.7\ \mathrm{mm}^2$ (8.1:1 aspect; CIF array fills $87\,\%$ of the chip footprint; unrolled CIF length $160.7\,\mathrm{mm}$; Gifford 2014 does not state a chip footprint for this device). Layer key in figure title: blue = channels, yellow = posts (reference outline only), green = ports, red = alignment marks. Detailed regions are in Fig. 6; quantitative comparison to Gifford 2014 Figure 4 is in Section IV-D.

fills an empty niche in the open-source ecosystem rather than competing on overlapping capability.

The closed-form identity itself is mathematically elementary — the textbook current-divider rule applied to the CIF outlet under the standard hydraulic–electrical analogy [13]. The contribution of the present work is therefore not new physics but the three-fold combination of: (i) publication-gap closure across 12+ years of CIF literature (Section II, "Why the identity may have been overlooked"); (ii) empirical validation of the identity, both against the converged Gifford recursion across three architectures (Section IV-A, Table I) and against the Dinh 2024 V1 and V3 measurements via three-dimensional Stokes simulation (Section IV-B); and (iii) the open-source reference implementation that operationalizes the framework (Section III). Mask-layout extensions for Dinh-style multi-stage tapers, Xia-style recycling, and Lezzar-style multiplexed layouts are identified as future work and itemized in Section IV-F.

The within-1σ COMSOL match at both Dinh V1 ($+0.06\sigma$) and V3 (-0.58σ), combined with the bracketing pattern $1.398 < 1.489 \pm 0.203 < 1.549$ (closed-form, paper, COMSOL; Section IV-B), is a stronger empirical warrant than either result alone. The bracketing constitutes a convergent evidentiary pattern: the closed-form (1-D asymptotic PE) and COMSOL (3-D finite 5-gap slab) are computationally distinct routes to the same predicted quantity, and they converge on the paper measurement from opposite sides. This convergence is interpretable as evidence for the framework only because three distinct preconditions are independently checked: (i) a *theoretical* precondition — the closed-form identity is algebraically derived from mass conservation, $Q_s(0) = 0$, and outlet pressure equilibration (Section II), with the current-divider re-derivation and SymPy verification (notebook §8.2) providing internal consistency checks across notations rather than independent physics; (ii) a *numerical* precondition — the COMSOL solution is mesh-converged to $|\Delta F| < 5 \times 10^{-5}$ at approximately $5\times$ coarser meshes (Section III-B), bounding discretization residual well below the experimental σ band; and (iii) an *architectural* precondition — $Q_s(0) = 0$ holds as a device-architecture fact across the published lineage (Section II), and the finite-PE residual within the 5-gap slab is bounded (95.0% V1, 93.4% V3 Δp decay; Fig. 3b), so the framework's premises apply to the simulated geometry. With these three preconditions checked, the bracketing pattern carries its convergent-evidence content; without any one of them, the bracketing would be much weaker as evidence.

Two practical design routes follow immediately from Eq. (4), each replacing the recursion's outlet step with a single algebraic relation. *Route (i)*: the designer specifies a target filtrate fraction F, which fixes the outlet resistance ratio $R_c/R_s = F/[2(1-F)]$; choosing any $w_{c,\mathrm{out}}$ determines $w_{s,\mathrm{out}}$ by inverting the Yang series at the fabrication-fixed channel depth via the Block 3 bisection inverter. *Route (ii)*: the designer specifies $w_{c,\mathrm{out}}$ and $w_{s,\mathrm{out}}$ directly; the Yang series evaluated forward returns R_c and R_s, and substitution into Eq. (4) returns F in closed form via `_F_closed_form()` in notebook Block 12.6. A single closed-form evaluation replaces running the recursion to convergence and tuning $w_s(i=1)$ until the endpoint matches. The PE-consistency diagnostic `pe_consistency_check()` is exposed as part of the notebook reference (Section III-A7): any candidate CIF outlet geometry can be tested for PE-consistency before fabrication decisions are made.

For new Gifford-style CIF devices (constant-w_c or tapered-w_c, with constant or variable f_{gap}), the trajectory generator is a drop-in replacement for manual re-implementation of the framework from primary-literature

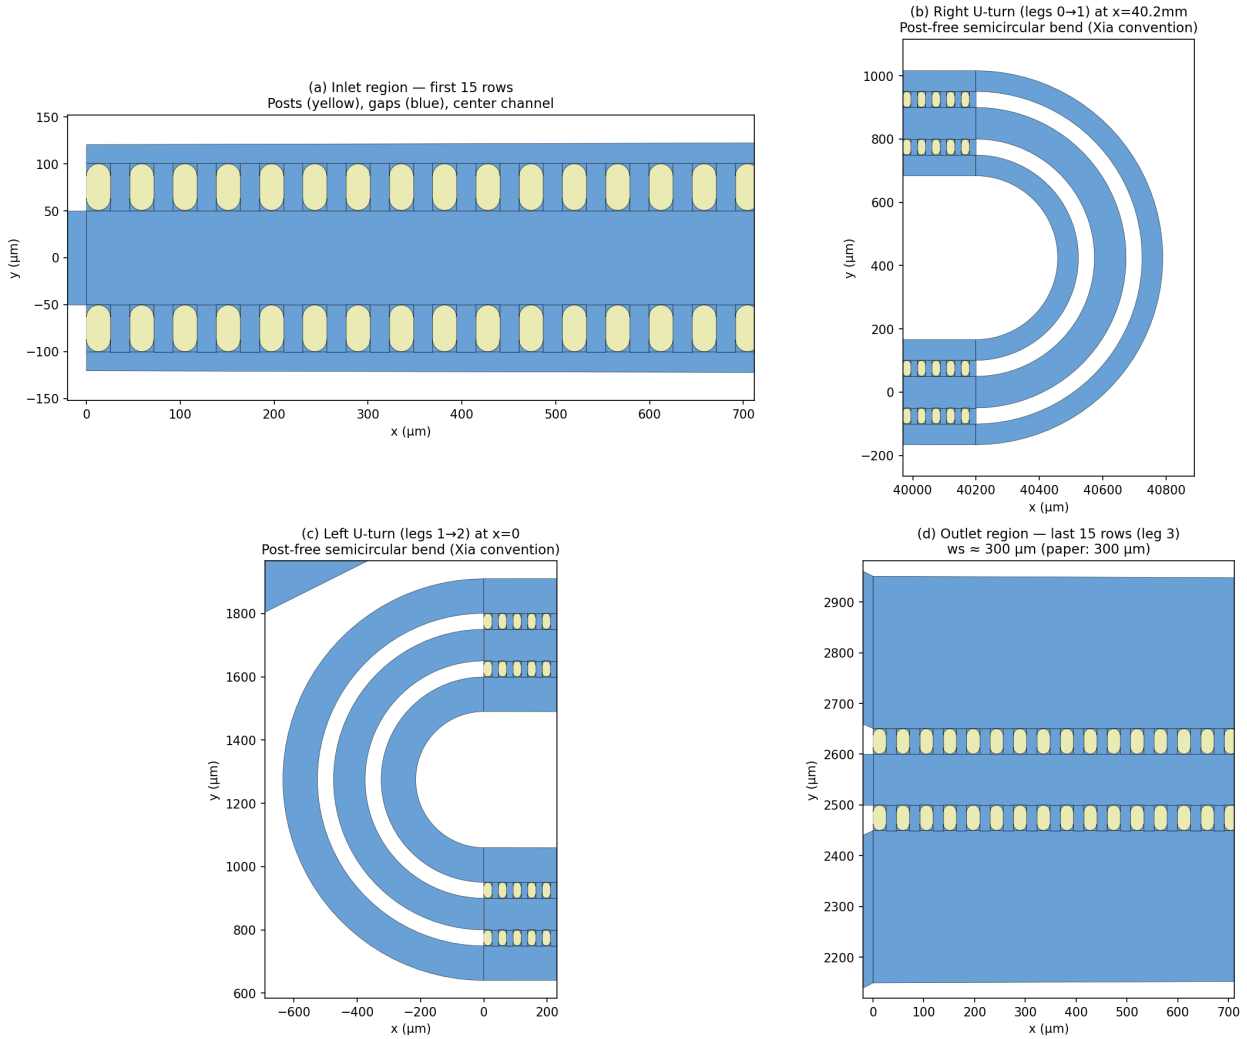

Fig. 6: Detailed regions of the mask layout shown in Fig. 5, generated by Block 8's `plot_mask_zooms` function and serving as visual validation of the reference implementation against Gifford 2014 Figure 4 [1]. *(a) Inlet region (first 15 rows of leg 0):* regular pillar array reproducing the Gifford 2014 Figure 4 micrograph at pitch $46.4(25)\,\mu\mathrm{m}$, $L_{\mathrm{post}} = 25.6\,\mu\mathrm{m}$ along flow, $W_{\mathrm{post}} = 50.6\,\mu\mathrm{m}$ perpendicular to flow, $w_{\mathrm{gap}} = 20.5\,\mu\mathrm{m}$ along flow. All four dimensions are pixel-measured from the published figure at $2.564\,\mu\mathrm{m}$ per pixel; the $L_{\mathrm{post}}+w_{\mathrm{gap}}$ internal-consistency check holds to $0.3\,\mu\mathrm{m}$ of the FFT-measured pitch (Section IV-D). Stadium-shaped (pill) posts oriented perpendicular to the flow direction; center channel runs horizontally between two parallel post rows. *(b) Right U-turn (legs $0 \to 1$, at $x \approx 40.2\,\mathrm{mm}$):* post-free semicircular bend connecting two adjacent legs, following the Xia 2016 convention. The U-turn comprises three concentric channels (side, center, side) traversing the same arc. *(c) Left U-turn (legs $1 \to 2$, at $x = 0$):* symmetric counterpart of panel (b). *(d) Outlet region (last 15 rows of leg 3):* the side channels have widened from their inlet $20\,\mu\mathrm{m}$ to $w_s \approx 300\,\mu\mathrm{m}$, matching the Gifford 2014 outlet specification of $300\,\mu\mathrm{m}$ to $0.027\,\%$ (generator output: $300.08\,\mu\mathrm{m}$). Layer key as in Fig. 5: blue = channels, yellow = posts (reference outline only).

equations. For the published Dinh tapered-w_c variants, the Xia RS-7 variable-f_{gap} case, and the Lezzar multiplexed layouts, the trajectory-generator stage is available today and produces F predictions and $w_s(i)$ trajectories for any Gifford-recursion-compatible parameter set (Table I); the corresponding mask-layout rendering to fabricable DXF remains future work (Section IV-F). Generalization of the closed-form identity to other particle-sorting modalities — for example, deterministic lateral displacement or pinched-flow fractionation — is conjectural here and would require showing that the architectural premises of Section II, especially $Q_s(0) = 0$ and outlet pressure equilibration, actually hold in those systems. The present identity is therefore confined to CIF outlets with the architecture of Section II; broader generalization is identified as future work (Section IV-F).

Finally, Eq. (4) addresses the volumetric filtrate fraction F only; the particle size cutoff remains set by f_{gap} calibration. That calibration has been empirical since the foundational framework: Gifford 2014 designed a parallel array of 33 test devices spanning $f_{\mathrm{gap}} \in [6.40\times 10^{-5},\ 5.76\times10^{-4}]$ to determine the size-cutoff threshold per bead diameter (Fig. 3 of [1]). Dinh 2024 inherits the same empirical character at the single-element design stage: the initial f_{gap} value is selected from empirical data suggesting it is likely to produce the target size cut-off, and per-gap $f_{\mathrm{gap}}(i)$ values are then iteratively adjusted to maintain that cut-off throughout the device (Dinh 2024 p. 4 [5]). A unified closed-form treatment of both parameters — a closed form for f_{gap} given a target size cutoff — is identified as future work (Section IV-F), complementary to the present contribution.

F. Limitations and future work

The following items are identified as future work, each addressing a specific scope boundary of the present contribution.

Mask-layout extensions. The reference implementation's mask-layout module currently renders only Gifford-style topologies (constant-w_c with three outlets). The trajectory-generator module already handles the Dinh tapered-w_c, the Xia variable-f_{gap}, and any other Gifford-recursion-compatible geometry across the published lineage (Table I); the limitation is at the mask-rendering stage. Generalizing the DXF renderer to produce Dinh-style multi-stage tapered topologies, Xia-style recycling serpentines, and Lezzar-style multiplexed-parallel layouts is the next major notebook deliverable. The closed-form identity (4) itself extends directly to those layouts: under the architectural boundary condition $Q_s(0) = 0$ shared across the published lineage [1], [2], [4], [5] and outlet pressure equilibration, $F = 2\,R_c/(R_s + 2\,R_c)$ depends only on outlet w_c and w_s, regardless of upstream taper shape, recycling-stage topology, or multiplexing.

Implementation correctness level. The reference implementation is at the implementation-tested level: 218 core plus 19 supplementary validators all passing across five reference devices for three architectures (Table III), 19 deterministic regression stamps documenting determinism across environments, and COMSOL empirical validation placing the predicted filtrate fraction within 1σ at both Dinh V1 and V3 (Section IV-B). It is not formally verified in the computer-science sense — no separate formal specification has been produced, and no machine-checked proof of conformance exists — and we do not claim algorithmic uniqueness. Alternative valid algorithms may exist; the existence of one open-source reference implementation does not preclude others. The closed-form identity (4) itself, by contrast, is at the algebraically-proven level for the specific theorem about the CIF outlet under PE: a three-line derivation from mass conservation, $Q_s(0) = 0$, and outlet pressure equilibration; an independent two-line parallel-resistance re-derivation; symbolic SymPy verification; and a five-device numerical cross-check (Section II, "Verification"). Formal verification of the trajectory-generator algorithm, independent re-implementation by another developer, and a machine-checked proof of the DXF-masks are potential follow-up contributions.

Direct fabrication metrology. The mechanism analysis of Section IV-C attributes the $\sim$10% closed-form-vs-paper residual primarily to 3-D outlet bifurcation (mechanism (a)), with w_s-specific fabrication offset (b) bounded as a sensitivity range $s \in [5.89\,\%, 11.86\,\%]$ within which both Dinh variants stay within 1σ, and the asymmetry pattern V3>V1 inconsistent with the named fabrication subcategories examined here (isotropic PDMS shrinkage [16], aspect-ratio-dependent shrinkage). Direct optical or SEM metrology of $w_{s,\mathrm{out}}$ in fabricated Dinh V1, V2, and V3 devices would tighten the bound on mechanism (b) further and let any contribution outside the 3-D Stokes prediction be quantified directly. Complementary to this, a mesh-convergence-grade COMSOL run at additional Dinh-style outlet geometries beyond V1 and V3 would test whether 3-D outlet bifurcation alone accounts for the residual across the full published Dinh design space.

Closed-form size-cutoff treatment. The closed-form identity (4) addresses the volumetric design parameter F only; the particle size cutoff is set by the per-gap filtration fraction f_{gap}, which has been calibrated empirically since the foundational framework. Gifford 2014 designed a parallel array of 33 test devices spanning $f_{\mathrm{gap}} \in [6.40 \times 10^{-5},\ 5.76 \times 10^{-4}]$ to determine the size-cutoff threshold per bead diameter [1, Fig. 3], and Dinh 2024 explicitly notes that the initial f_{gap} value is

selected from empirical data and per-gap $f_{\mathrm{gap}}(i)$ values are then iteratively adjusted to maintain the desired cutoff throughout the device [5, p. 4]. A unified closed-form treatment of both CIF design parameters — a closed form for f_{gap} given a target particle size cutoff, on the same theoretical footing as the present closed form for F — would render the entire CIF design problem algebraic and would eliminate the empirical-calibration step that has persisted across the published CIF lineage. We identify this as a complementary theoretical extension and a natural next direction, separate in character from the operational extensions above.

The present implementation already provides the foundation for the operational extensions above (mask-layout rendering, formal verification, and direct metrology), and the closed-form identity (4) provides the theoretical anchor for the unified f_{gap} treatment. The identity (4) holds theoretically at all CIF outlet architectures with $Q_s(0) = 0$ and outlet pressure equilibration; the trajectory generator has been validated across all five reference devices in three architectures; the worked Gifford 2014 photomask example takes the trajectory through to a fabricable DXF end-to-end (Section IV-D); and the 237-validator plus 19-regression-stamp reproducibility infrastructure (Table III) is built to absorb future contributions without re-baselining the existing checks. The within-1σ COMSOL match at the Dinh V1 and V3 outlets gives confidence that the framework will hold when extended to multi-stage-taper, recycling, and multiplexed cases, since those extensions inherit the same underlying physics and the same identity. Beyond CIF, generalization of the identity to other size-based microfluidic particle-sorting modalities — deterministic lateral displacement, pinched-flow fractionation, or hydrodynamic-filtering networks with different inlet architectures — is conjectural here and would require demonstrating that the architectural premise $Q_s(0) = 0$ and outlet pressure equilibration hold in those systems; we identify this as the longer-term future direction.

V. Conclusion

A closed-form algebraic identity for the volumetric filtrate fraction at a Controlled Incremental Filtration device outlet has been derived, verified, and empirically validated. Under outlet pressure equilibration and the architectural boundary condition $Q_s(0) = 0$ shared by every device in the published CIF lineage [1], [2], [4], [5], the filtrate fraction is $F = 2\,R_c/(R_s + 2\,R_c)$, with R_c and R_s the per-pitch hydraulic resistances of the centre and side channels at the outlet (Section II). Equivalently, this is the textbook current-divider rule applied to the CIF outlet's two-branch equivalent network under the standard hydraulic–electric analogy [13]. The identity reduces an $\mathcal{O}(N)$ design-time iteration (typically $N \sim 10^3$ gap rows for blood-cell-separation devices) to a single algebraic step in outlet geometry alone, and it provides both an outlet-design anchor for the trajectory generator and a quantitative reference against which residuals from the published measurements can be interpreted. To our knowledge it has not previously been written down in the 12+ years of CIF literature reviewed for this work (Section II, "Why the identity may have been overlooked").

The identity has been validated empirically in two stages. First, it reproduces the converged Gifford recursion to $|\Delta F| \leq 2.35 \times 10^{-4}$ across all five reference devices in three architectural regimes (Table I; Section IV-A) — including the strongest non-trivial test, the variable-f_{gap} Xia RS-7 case at $|\Delta F| = 1.97 \times 10^{-5}$, where f_{gap} enters every step of the recursion yet provably drops out of the closed form. Second, three-dimensional Stokes simulation in COMSOL Multiphysics 6.3 at the Dinh 2024 V1 and V3 outlet geometries recovers the published measurements $F_{\mathrm{paper,V1}} = 0.505 \pm 0.033$ and $F_{\mathrm{paper,V3}} = 0.649 \pm 0.036$ within 1σ ($F_{\mathrm{COMSOL,V1}} = 0.5070$, $+0.06\sigma$; $F_{\mathrm{COMSOL,V3}} = 0.6281$, -0.58σ; Section IV-B), down from the closed-form-vs-paper offsets of $+1.64\sigma$ and $+1.94\sigma$ at the same outlet geometries. The multi-variant ratio $\ln(1 - F_{\mathrm{V3}})/\ln(1-F_{\mathrm{V1}})$ confirms this picture by bracketing the paper measurement (1.489 ± 0.203) between the closed-form (1.549) and COMSOL (1.398). The residual $\sim$10% closed-form-vs-paper offset is closed by 3-D outlet bifurcation alone; the variant asymmetry $\alpha_{\mathrm{3D}} = 0.811$ (V1) vs. 0.660 (V3), a 15.1-percentage-point spread, is the geometry-dependent signature of that bifurcation in matched units to the closed-form identity (Section IV-C, notebook §8.8). Channel-depth tolerance and finite-PE residual are bounded out individually ($< 0.05\,\sigma_{\mathrm{paper}}$ shift and a $48\times$ PE safety margin respectively); the named subcategories of w_s-specific fabrication offset — isotropic PDMS shrinkage and aspect-ratio-dependent shrinkage — cannot reproduce the observed V3>V1 asymmetry. The worked Gifford 2014 photomask example (Section IV-D, Figs. 5 and 6) demonstrates that the trajectory generator produces fabricable output end-to-end.

The closed-form identity addresses the volumetric design question only; the particle size cutoff remains set by the per-gap filtration fraction f_{gap}, which has been calibrated empirically since the foundational framework [1] and explicitly so in the most recent lineage paper [5]. A unified closed-form treatment of both the size cutoff and the volumetric filtrate fraction is identified as future work, complementary to the present contribution. An open-source reference implementation operationalising the closed-form identity in a Gifford-variant CIF

design framework is released under an MIT license alongside this preprint (https://github.com/Austin-Routt/cif-reference-implementation; `v3.20` permanently archived at Zenodo [12]). It comprises a callable trajectory-generator module exercising the closed-form identity at the outlet (`generate_cif_device`), a PE-consistency diagnostic exposed as a public callable (`pe_consistency_check`, Section III-A7), a worked end-to-end Gifford 2014 photomask example reproducing the published device [1, Fig. 4] to within pixel-measurement precision (Section IV-D), and a 237-validator plus 19-regression-stamp reproducibility infrastructure (Table III). Mask-layout extensions for Dinh-style multi-stage tapers, Xia-style recycling stages, and Lezzar-style multiplexed-parallel topologies are identified as the next major notebook deliverables (Section IV-F); the trajectory-generator stage is already in place for all three architectures and inherits the present empirical validation. Beyond CIF, generalization of the identity to other size-based particle-sorting modalities is conjectural and identified as longer-term future work.

Code and data availability

The complete reference implementation is openly available on GitHub at https://github.com/Austin-Routt/cif-reference-implementation under an MIT license. The version cited in this preprint (`v3.20`) is permanently archived on Zenodo (doi: 10.5281/zenodo.20146562) [12]. The notebook reproduces all numerical results reported in this paper.

Conflicts of interest

The author reports no personal conflicts of interest. The author is an unpaid graduate student in the laboratory of S. S. Shevkoplyas at the University of Houston; S. S. Shevkoplyas is an inventor on U.S. Patent #9789235 (assigned to the University of Houston) describing the controlled incremental filtration technology and a co-founder of Halcyon Biomedical, Incorporated, which commercializes that technology. The present work is a re-derivation and open-source reference implementation of the published CIF design framework using only the publicly available primary literature [1]–[5]; no proprietary information from Halcyon Biomedical or any unpublished lab tooling has been used. The author has no financial interest in Halcyon Biomedical.

Author contributions

A.H.R. conceived the closed-form identity, recognized its equivalence to the standard two-resistor current-divider rule applied to the CIF outlet's two-branch equivalent network, implemented the open-source reference implementation, conducted the COMSOL simulations, performed the literature analysis establishing the methodological bottleneck and the publication gap, and wrote the manuscript.

Funding

This research received no external or internal funding. The author is an unpaid graduate student at the University of Houston and conducted this work as part of doctoral research.

Acknowledgements

The author acknowledges the use of Anthropic Claude Opus 4.7, OpenAI ChatGPT (GPT-5.5), Google Gemini 3.1 Pro Preview, and DeepSeek V4-Pro as drafting, coding, and analytical aids during manuscript preparation and notebook development. All scientific content, including the derivations, simulations, reference implementation, validation infrastructure, and conclusions presented in this work, is the author's original work.

Preprint status

This is a preprint of a manuscript intended for submission to *Lab on a Chip*. The accompanying open-source reference implementation is permanently archived at the Zenodo DOI cited above; subsequent revisions of this preprint will be posted as bioRxiv version updates as the manuscript progresses through collaboration and peer review.

References

[1] S. C. Gifford, A. M. Spillane, S. M. Vignes, and S. S. Shevkoplyas, "Controlled incremental filtration: A simplified approach to design and fabrication of high-throughput microfluidic devices for selective enrichment of particles," *Lab Chip*, vol. 14, no. 23, pp. 4496–4505, 2014, doi: 10.1039/C4LC00785A.

[2] H. Xia, B. C. Strachan, S. C. Gifford, and S. S. Shevkoplyas, "A high-throughput microfluidic approach for 1000-fold leukocyte reduction of platelet-rich plasma," *Sci. Rep.*, vol. 6, p. 35943, 2016, doi: 10.1038/srep35943.

[3] H. Xia, B. C. Strachan, S. C. Gifford, and S. S. Shevkoplyas, "Author Correction: A high-throughput microfluidic approach for 1000-fold leukocyte reduction of platelet-rich plasma," *Sci. Rep.*, vol. 12, p. 14689, 2022, doi: 10.1038/s41598-022-18449-5.

[4] D. L. Lezzar *et al.*, "A high-throughput microfluidic device based on controlled incremental filtration to enable centrifugation-free, low extracorporeal volume leukapheresis," *Sci. Rep.*, vol. 12, p. 13798, 2022, doi: 10.1038/s41598-022-16748-5.

[5] M. T. P. Dinh, A. Mukhamedshin, K. Abhishek, F. W. Lam, S. C. Gifford, and S. S. Shevkoplyas, "Separation of platelets by size in a microfluidic device based on controlled incremental filtration," *Lab Chip*, vol. 24, no. 4, pp. 913–923, 2024, doi: 10.1039/D3LC00842H.

[6] S. Yang, A. Ündar, and J. D. Zahn, "A microfluidic device for continuous, real time blood plasma separation," *Lab Chip*, vol. 6, no. 7, pp. 871–880, 2006, doi: 10.1039/B516401J.

[7] COMSOL Multiphysics® v. 6.3, COMSOL AB, Stockholm, Sweden, www.comsol.com.

[8] R. Sanka *et al.*, "3DμF — Interactive design environment for continuous-flow microfluidic devices," *Sci. Rep.*, vol. 9, no. 1, p. 9166, 2019, doi: 10.1038/s41598-019-45623-z.

[9] M. T. Birtek *et al.*, "ML-automated microfluidic circuit design," *Sci. Adv.*, vol. 12, no. 5, p. eaea7598, 2026, doi: 10.1126/sciadv.aea7598.

[10] Y. Zhang, M. Li, T.-M. Tseng *et al.*, "Open-source interactive design platform for 3D-printed microfluidic devices," *Commun. Eng.*, vol. 3, p. 71, 2024, doi: 10.1038/s44172-024-00217-0.

[11] B. Goenner *et al.*, "An open source platform to automate the design, verification, and manufacture of 3D printed microfluidic devices," *Sci. Rep.*, 2025, doi: 10.1038/s41598-025-15976-9.

[12] A. H. Routt, "CIF reference implementation `v3.20`," *Zenodo*, 2026, doi:10.5281/zenodo.20146562.

[13] K. W. Oh, K. Lee, B. Ahn, and E. P. Furlani, "Design of pressure-driven microfluidic networks using electric circuit analogy," *Lab Chip*, vol. 12, no. 3, pp. 515–545, 2012, doi: 10.1039/C2LC20799K.

[14] Z. Li, C. Liu, and J. Sun, "Hydraulic–electric analogy for design and operation of microfluidic systems," *Lab Chip*, vol. 23, pp. 3311–3327, 2023, doi: 10.1039/D3LC00265A.

[15] M. Yamada and M. Seki, "Hydrodynamic filtration for on-chip particle concentration and classification utilizing microfluidics," *Lab Chip*, vol. 5, no. 11, pp. 1233–1239, 2005, doi: 10.1039/B509386D.

[16] S. W. Lee and S. S. Lee, "Shrinkage ratio of PDMS and its alignment method for the wafer level process," *Microsyst. Technol.*, vol. 14, no. 2, pp. 205–208, 2008, doi: 10.1007/s00542-007-0417-y.