

A theory of linear typings as flows on 3-valent graphs

Zeilberger, Noam

DOI:
[10.1145/3209108.3209121](https://doi.org/10.1145/3209108.3209121)

Document Version
Peer reviewed version

Citation for published version (Harvard):
Zeilberger, N 2018, A theory of linear typings as flows on 3-valent graphs. in *LICS '18 : Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science*. Symposium on Logic in Computer Science, Association for Computing Machinery (ACM), New York, NY, USA, pp. 919-928, Thirty-Third Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), Oxford, United Kingdom, 9/07/18.
<https://doi.org/10.1145/3209108.3209121>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:
Checked for eligibility: 17/05/2018

© {Owner/Author | ACM} 2018. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in LICS '18 Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, <http://dx.doi.org/10.1145/3209108.3209121>

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

A theory of linear typings as flows on 3-valent graphs

Noam Zeilberger

University of Birmingham

April 30, 2018


Abstract

Building on recently established enumerative connections between lambda calculus and the theory of embedded graphs (or “maps”), this paper develops an analogy between typing (of lambda terms) and coloring (of maps). Our starting point is the classical notion of an abelian group-valued “flow” on an abstract graph (Tutte, 1954). Typing a linear lambda term may be naturally seen as constructing a flow (on an embedded 3-valent graph with boundary) valued in a more general algebraic structure consisting of a preordered set equipped with an “implication” operation and unit satisfying composition, identity, and unit laws. Interesting questions and results from the theory of flows (such as the existence of nowhere-zero flows) may then be re-examined from the standpoint of lambda calculus and logic. For example, we give a characterization of when the local flow relations (across vertices) may be categorically lifted to a global flow relation (across the boundary), proving that this holds just in case the underlying map has the orientation of a lambda term. We also develop a basic theory of rewriting of flows that suggests topological meanings for classical completeness results in combinatory logic, and introduce a polarized notion of flow, which draws connections to the theory of proof-nets in linear logic and to bidirectional typing.

1 Introduction

The study of graphs embedded on surfaces, or *maps*, has a long history, much of it linked with the rich history of the Four Color Problem (now the Four Color Theorem, or 4CT) [41, 15]. Formally, 4CT is a statement about maps, namely that *every bridgeless planar map has a proper face-4-coloring*.

A mathematician who made great contributions to the study of maps and colorings was Bill Tutte, including his observation of a duality between *chromatic polynomials* and *flow polynomials* [44], which turns on the very natural notion of an *abelian group-valued flow* over a graph. By a well-known reduction going back to Tait [38], 4CT is equivalent to the statement that *every bridgeless planar 3-valent map has a proper edge-3-coloring*,



and within Tutte’s theory this may be reformulated as the statement that every such map has a *nowhere-zero \mathbb{V} -flow*, where $\mathbb{V} \cong \mathbb{Z}_2 \times \mathbb{Z}_2$ stands for the Klein Four Group [48].

A separate line of work that Tutte began in the 1960s (also originally motivated by 4CT, see [50, Ch. 10]) was the *enumerative* study of maps, establishing some remarkably simple formulas for the number of (rooted) planar maps of a given size satisfying varying constraints. Though Tutte’s approach to 4CT was ultimately side-stepped by the Appel-Haken proof [1], enumeration of maps remains a very active area of combinatorics, with links to wide-ranging domains such as algebraic geometry, knot theory, and mathematical physics.

This is the preliminary arXiv version of a paper to be presented at LICS 2018. The content is the same as the conference version, with the addition of two appendices containing proofs of results (Appendix A) and an extended example (Appendix B).

family of lambda terms	family of rooted maps	OEIS
linear	3-valent (of genus $g \geq 0$)	A062980
planar	planar 3-valent	A002005
unitless linear	bridgeless 3-valent ($g \geq 0$)	A267827
unitless planar	bridgeless planar 3-valent	A000309
β -normal linear/ \sim	(all maps of genus $g \geq 0$)	A000698
β -normal planar	planar	A000168
β -normal unitless linear/ \sim	bridgeless ($g \geq 0$)	A000699
β -normal unitless planar	bridgeless planar	A000260

Table 1: Known correspondences [3, 56, 53, 54, 55, 7] between families of lambda terms and rooted maps, as combinatorial classes. Here “ \sim ” stands for an equivalence relation defined in [53], and the rest of the terminology is explained in Sections 3 and 4.1. (Indices on the right refer to the Online Encyclopedia of Integer Sequences [34].)

The ideas developed in this paper sprang from the recent discovery of a host of surprising links between map enumeration and lambda calculus: see Table 1. Bridges between lambda calculus and graph theory trace back to the pioneering works of Statman [35] and Girard [14], but these combinatorial connections go even further in suggesting that very concrete objects of study are actually shared: for example, it turns out that many of the sequences first computed by Tutte also count natural families of lambda terms!

Although the correspondences listed in the lower half of Table 1 for the moment rest mainly at the level of enumeration, there is a simple bijection between linear lambda terms and rooted 3-valent maps (originally described in [3]) that may be restricted to account for the entire upper half of the table. This was used in [54] to state a lambda calculus reformulation of 4CT, essentially by turning the existence of a nowhere-zero \mathbb{W} -flow into a *typing* problem for a certain family of terms. In turn, that motivated the development of a more general theory of linear typings-as-flows on 3-valent maps, and this paper represents a preliminary sketch of such a theory.

For example, rather than limiting types to abelian groups, flows over linear lambda terms are naturally valued in a more general algebraic structure consisting of a preordered set equipped with an “implication” operation and unit element satisfying composition, identity, and unit laws: what we call an *imploid* (and is otherwise known as a “thin skew-closed category”). Considering imploid-valued flows over arbitrary well-oriented 3-valent maps with boundary, a natural question is when can the local flow relations (across vertices) be lifted to a global flow relation (across the boundary). We will see that this holds just in case the map is equipped with the canonical orientation of a lambda term, giving a new perspective on the interplay between linear typing and compositionality.

The rest of the paper is structured as follows. In Section 2 we establish some elementary properties of imploids, leading up through the construction of quotients. In Section 3, after a review of topics in graph theory, we introduce the basic definition of imploid-valued flows and nowhere-unit flows over well-oriented 3-valent maps. We recall the bijection between rooted 3-valent maps and linear lambda terms based on the *topological orientation* of a map with boundary, and use it to prove the above-mentioned characterization of the global flow condition. We also briefly explain how the definitions may be recast algebraically in terms of a certain quotient imploid we call the *fundamental imploid* of a well-oriented 3-valent map (analogous to the fundamental quandle of a knot). In Section 4, we discuss rewriting of flows across operations such as β -reduction and η -expansion, and prove a *topological completeness* theorem which is closely related to classical completeness results in combinatory logic. Finally, in Section 5 we briefly describe a polarized notion of flow, pointing out connections with linear logic proof-nets, as well as with bidirectional typing.

Acknowledgments. The ideas described in this paper have been in gestation for a while and have benefited from interaction with numerous individuals and seminar audiences, as well as from the friendly and supportive atmosphere within the Theory Group at Birmingham. I am especially grateful to Jason Reed for a long series of email exchanges on map coloring and lambda calculus that (in addition to being highly enjoyable) really brought these ideas into focus.

2 The elementary theory of imploids

2.1 Preliminary definitions and examples

We recall some standard terminology and conventions. A *preorder* on a set P is a binary relation which is reflexive and transitive, typically indicated \leq or \leq_P . We write \equiv for the induced equivalence relation $a \equiv b := (a \leq b) \wedge (b \leq a)$, and say that the preorder is a *partial order* when $a \equiv b$ implies $a = b$. A *preordered set* is a set equipped with a preorder (P, \leq) , and we may also sometimes write P to stand for the preordered set with \leq left implicit. To any preordered set P is associated an *opposite* P^{op} with the same underlying set of elements but the opposite preorder, $a \leq^{\text{op}} b$ iff $b \leq a$. To any pair of preordered sets P and Q is associated a *product* $P \times Q$ with underlying set of elements given by the cartesian product and with the componentwise ordering $(a_1, b_1) \leq_{P \times Q} (a_2, b_2)$ iff $a_1 \leq_P a_2$ and $b_1 \leq_Q b_2$. Finally, we write $f : P \rightarrow Q$ to indicate that f is an order-preserving (a.k.a. *monotonic*) function from P to Q (i.e., that $a \leq_P b$ implies $f(a) \leq_Q f(b)$ for all $a, b \in P$), and we say that f is *left adjoint* to $g : Q \rightarrow P$ (or g is *right adjoint* to f) just in case $f(a) \leq_Q b$ iff $a \leq_P g(b)$ for all $a \in P, b \in Q$.

Definition 2.1. An **imploid** is a preordered set P equipped with an operation \multimap (called *implication*) which is *contravariant* in its left argument and *covariant* in its right argument,

$$\frac{a_2 \leq a_1 \quad b_1 \leq b_2}{a_1 \multimap b_1 \leq a_2 \multimap b_2} \quad (\text{imp})$$

together with a distinguished element $I \in P$, satisfying *composition*, *identity*, and *unit laws*:

$$b \multimap c \leq (a \multimap b) \multimap (a \multimap c) \quad (\text{comp})$$

$$I \leq a \multimap a \quad (\text{id})$$

$$I \multimap a \leq a \quad (\text{unit})$$

A **non-unital imploid** is given by the same except that we do not require the element I and omit axioms (id) and (unit).

It is worth emphasizing that what we call an imploid is simply the preorder restriction of what Street has recently referred to as a *skew-closed category* [36], corresponding to a slight relaxation of Eilenberg and Kelly's original definition of a (non-monoidal) closed category [12]. Although the development we give here is limited to preordered sets (which include our motivating examples), it is likely that many of these constructions could be lifted with care to the more general context of skew-closed categories.

Definition 2.2 (cf. [36]). In any imploid P we have that $a \leq b$ entails $I \leq a \multimap b$, by (id) and (imp). We say that P is **left normal** if this entailment is reversible, that is, if $I \leq a \multimap b$ entails $a \leq b$.

Example 2.3. Any Heyting algebra defines a left normal imploid with $a \multimap b := a \supset b$ and $I := \top$. More generally, any (unital) quantale [52] gives an example of a (left normal, unital) imploid.

Example 2.4. Any group can be seen as a left normal imploid under the discrete order ($a \leq b$ iff $a = b$), where I is the unit element of the group and implication can be defined by either right division $a \multimap b := b \cdot a^{-1}$ or left division $a \multimap b := a^{-1} \cdot b$.

Definition 2.5 (cf. [12, 4]). We say that a (unital or non-unital) imploid is **symmetric** if it satisfies the law of exchange:

$$a \multimap (b \multimap c) \leq b \multimap (a \multimap c) \quad (\text{exch})$$

Proposition 2.6. The law of double-negation introduction

$$a \leq (a \multimap b) \multimap b \quad (\text{dni})$$

implies the law of exchange. Conversely, (exch) implies (dni) under assumption of left normality.

Example 2.7. In the above examples, the imploid associated to a group/quantale is symmetric whenever the underlying multiplication operation is commutative.

2.2 Skew monoids, upsets and downsets

One of the motivations for the recent study of skew-closed categories (the category-theoretic version of imploids) is their close connection to *skew-monoidal categories* [37]. We will refer to the order-theoretic versions of the latter as *skew monoids*.

Definition 2.8. A (left) **skew monoid** is a preordered set M equipped with an operation \bullet which is covariant in both arguments,

$$\frac{a_1 \leq a_2 \quad b_1 \leq b_2}{a_1 \bullet b_1 \leq a_2 \bullet b_2} \quad (\text{mul})$$

as well as a distinguished element $I \in M$, satisfying semi-associativity, left unit, and right unit laws:

$$\begin{aligned} (a \bullet b) \bullet c &\leq a \bullet (b \bullet c) && (\text{assocr}) \\ I \bullet a &\leq a && (\text{lunit}) \\ a &\leq a \bullet I && (\text{runit}) \end{aligned}$$

A **non-unital skew monoid** is given by the same except that we do not require I and omit (lunit) and (runit).

One simple relationship between imploids and skew monoids is via adjunction: any family of right adjoints to the partially instantiated multiplication operations $- \bullet b : M \rightarrow M$ of a skew monoid M can be extended to an operation $\multimap : M^{\text{op}} \times M \rightarrow M$ satisfying the imploid laws, and dually, any family of left adjoints to the partially instantiated implication operations $b \multimap - : P \rightarrow P$ of an imploid P can be extended to an operation $\bullet : P \times P \rightarrow P$ satisfying the skew monoid laws.¹ We will make use of a related duality between the *downsets* of a skew-monoid and the *upsets* of an imploid.

Definition 2.9. A subset R of a preordered set A is said to be **downwards closed** (or a **downset**), written $R \sqsubset^{\downarrow} A$, if $b \in R$ and $a \leq b$ implies $a \in R$. Dually, it is said to be **upwards closed** (or an **upset**), written $R \sqsubset^{\uparrow} A$, if $a \in R$ and $a \leq b$ implies $b \in R$. (We sometimes use mirror notation “ $R \ni a$ ” to denote the elementhood relation in an upset. Also, we sometimes write $R \sqsubset A$ to indicate that R is a subset but emphasize that it is not necessarily closed with respect to the order on A .)

Recall that every element $x \in A$ of a preordered set induces both a *principal downset* $x^{\downarrow} := \{a \mid a \leq x\}$ and a *principal upset* $x^{\uparrow} := \{a \mid x \leq a\}$, and that these define a pair of faithful embeddings $(-)^{\downarrow} : A \rightarrow \hat{A}$ and $(-)^{\uparrow} : A \rightarrow \check{A}$, where \hat{A} denotes the set of downsets of A partially ordered by inclusion, and \check{A} the set of upsets partially ordered by reverse inclusion (an order-preserving function $f : P \rightarrow Q$ is said to be *faithful* when $f(a) \leq_Q f(b)$ implies $a \leq_P b$). The following constructions amount to a “skew” variation of the well-known *Day construction* (cf. [8, 36]).

Proposition 2.10. If $P = (P, \leq, \multimap, I)$ is an imploid, then \check{P} can be given the structure of a skew monoid as follows (for all $R, S \sqsubset^{\uparrow} P$):

$$\begin{aligned} R \bullet S \ni p &\iff \exists q. R \ni q \multimap p \wedge S \ni q \\ I &:= I^{\uparrow} \ni p \iff I \leq p \end{aligned}$$

Observation 2.11 (cf. [4]). Let P be an imploid. Then

1. $a \leq (a \multimap b) \multimap b$ for all $a, b \in P$ iff $R \bullet S \supseteq S \bullet R$ for all $R, S \sqsubset^{\uparrow} P$; and
2. $a \multimap (b \multimap c) \leq b \multimap (a \multimap c)$ for all $a, b, c \in P$ iff $(R \bullet S) \bullet T \supseteq (R \bullet T) \bullet S$ for all $R, S, T \sqsubset^{\uparrow} P$.

Proposition 2.12. If $M = (M, \leq, \bullet, I)$ is a skew monoid, then \hat{M} can be given the structure of an imploid as follows (for all $K, L \sqsubset^{\downarrow} M$):

$$\begin{aligned} m \in K \multimap L &\iff \forall n. n \in K \Rightarrow m \bullet n \in L \\ m \in I &:= I^{\downarrow} \iff m \leq I \end{aligned}$$

¹A categorical version of this fact is proved by Street [36], who mentions that it resolves a nagging asymmetry in the traditional setting (cf. [8]), where the existence of left adjoints is not enough to ensure that an Eilenberg-Kelly closed category can be given the structure of an ordinary monoidal category.

Definition 2.13. An order-preserving function $f : P \rightarrow Q$ between two imploids is said to be a **(lax) homomorphism** if it weakly preserves the imploid structure in the sense that $I \leq_Q f(I)$ and $f(a \multimap b) \leq_Q f(a) \multimap f(b)$ for all $a, b \in P$. It is said to be **strong** if it preserves this structure up to equivalence, $I \equiv_Q f(I)$ and $f(a \multimap b) \equiv_Q f(a) \multimap f(b)$.

Proposition 2.14. The composite $(-)^{\uparrow\downarrow} : P \rightarrow \hat{P}$ is a strong homomorphism of imploids.

2.3 Deductive closure, dni and imploid quotients

Given that any group can be seen as an imploid (Example 2.4), it is natural to wonder what is the imploid analogue for *subgroups*. In fact, there are at least two different natural substructures of an imploid that could be considered as generalizations of the group-theoretic concept, one starting from the view of a subgroup as the image of an injective homomorphism, the other from the view of a (normal!) subgroup as the kernel of a surjective homomorphism.

Definition 2.15. A subset $R \sqsubset P$ of an imploid P is said to be a **subimploid** if 1) $I \in R$, and 2) $a \in R$ and $b \in R$ implies $a \multimap b \in R$.

Definition 2.16. An upset $R \sqsubset^\uparrow P$ of an imploid P is said to be **deductively closed** (or a **dedupset**) if 1) $I \in R$, and 2) $a \in R$ and $a \multimap b \in R$ implies $b \in R$.

It is easy to check that for any group G viewed as a discrete imploid, a subset $H \subseteq G$ is a subimploid iff it is a dedupset iff it is a subgroup. However, in general these two notions are quite different. Since imploid quotients will play an important role in this paper, we spend the rest of the section on elaborating the second definition, beginning with the following simple observation.

Observation 2.17. A dedupset of P is the same thing as a comonoid in \check{P} (relative to the skew monoid structure defined in Prop. 2.10), i.e., an upset $R \sqsubset^\uparrow P$ such that $R \supseteq I$ and $R \supseteq R \bullet R$.

Corollary 2.18. The **deductive closure** $!R$ of an upset R is given by the formula $!R := \bigwedge_{n \geq 0} R^{\bullet n}$, where \bigwedge denotes the meet in \check{P} (corresponding to union of subsets), and where $R^{\bullet n}$ denotes the left-associated product $R^{\bullet 0} = I$, $R^{\bullet n+1} = R^{\bullet n} \bullet R$.

Definition 2.19. The **induced relation** of an upset $R \sqsubset^\uparrow P$ is a binary relation $\#_R$ on the elements of P defined by $a \#_R b$ iff $a \multimap b \in R$.

Proposition 2.20. If R is a dedupset, $\#_R$ is a preorder extending \leq .

Given an imploid P and a dedupset $R \sqsubset^\uparrow P$, a natural candidate for the *quotient imploid* is given by $P/R = (P, \#_R, \multimap, I)$ (i.e., by the same underlying set and operations considered relative to a coarser order). . . but the problem is that this does not necessarily define an imploid. Although the three imploid axioms obviously remain valid (since $\#_R$ is an extension of \leq), and one can even verify that the implication $a \multimap b$ is monotone in b relative to $\#_R$, nothing guarantees that it is also antitone in a . To ensure that implication restricts to an operation of type $(P/R)^{\text{op}} \times P/R \rightarrow P/R$ we impose a further condition on dedupsets, which is the precise analogue of the restriction to *normal* subgroups in the construction of group-theoretic quotients.

Definition 2.21. We say that an upset $R \sqsubset^\uparrow P$ is **dni-closed** if $a \in R$ implies $(a \multimap b) \multimap b \in R$ for all $b \in P$.

Proposition 2.22. Any upset R has a **dni-closure** (i.e., a dni-closed upset containing R and maximal wrt. \supseteq).

Proposition 2.23. Let $R \sqsubset^\uparrow P$. The following are equivalent:

1. R is dni-closed.
2. $R \bullet S \supseteq S \bullet R$ for all $S \sqsubset^\uparrow P$.
3. R satisfies the following closure conditions:
 - (a) $a \in R \Rightarrow I \multimap a \in R$
 - (b) $a \multimap b \in R \Rightarrow \forall c \in P, (b \multimap c) \multimap (a \multimap c) \in R$

4. The induced relation $\#_R$ satisfies the following rules:

$$\frac{a \in R}{I \#_R a} \quad \frac{a_2 \#_R a_1 \quad b_1 \#_R b_2}{a_1 \multimap b_1 \#_R a_2 \multimap b_2}$$

Corollary 2.24. 1) If R and S are dni-closed then so is $R \bullet S$; 2) if R is dni-closed then so is $!R$.

Proposition 2.25. If $R \sqsubset^\uparrow P$ is deductively closed and dni-closed, then $P/R := (P, \#_R, \multimap, I)$ is a left normal imploid.

Proposition 2.26. Let $f : P \rightarrow Q$ be any order-preserving function, and $S \sqsubset^\uparrow Q$ any upset.

1. The inverse image of S along f is an upset $f^{-1}(S) \sqsubset^\uparrow P$.
2. If f is a homomorphism of imploids and S is deductively closed, then $f^{-1}(S)$ is deductively closed.
3. If f is a strong homomorphism and S is dni-closed, then $f^{-1}(S)$ is dni-closed.

Proposition 2.27. The upset I^\uparrow is deductively and dni-closed.

Definition 2.28. Let $f : P \rightarrow Q$ be a homomorphism of imploids. The **kernel** of f is the upset $\ker f \sqsubset^\uparrow P$ given by the inverse image of the unit, $\ker f := f^{-1}(I_Q)$.

Proposition 2.29. Let $f : P \rightarrow Q$ be a homomorphism. Then $\ker f \sqsubset^\uparrow P$ is deductively closed, and dni-closed if f is strong.

Proposition 2.30. Let $f : P \rightarrow Q$ be a strong homomorphism of left normal imploids. Then f is faithful iff $\ker f = I$.

Proposition 2.31 (Universal property of the quotient). For any imploid P and dni-closed dedupset $R \sqsubset^\uparrow P$, the function acting as the identity on elements defines a strong homomorphism of imploids $[-] : P \rightarrow P/R$ whose kernel is R . Moreover, for any other left normal imploid Q and lax (respectively, strong) homomorphism $f : P \rightarrow Q$ such that $R \subseteq \ker f$, there exists a unique lax (respectively, strong) imploid homomorphism $\bar{f} : P/R \rightarrow Q$ such that $f = \bar{f} \circ [-]$.

Corollary 2.32. For any collection C of ordered pairs of elements of an imploid P , we can define the **quotient of P modulo the relations** $[a \leq b]_{(a,b) \in C}$ as P/\tilde{R}_C , where \tilde{R}_C is the deductive closure of the dni-closure of the upwards closure of the set $\{a \multimap b \mid (a,b) \in C\}$.

Note that although our construction of the imploid quotient only defines a coarser preorder on the existing elements of P , we can always obtain a partially ordered set by considering the image of P/R under the $(-)^{\uparrow\downarrow}$ embedding (Prop. 2.14). In the case of a group G seen as a discrete imploid, quotienting by a dni-closed dedupset (= normal subgroup) $H \triangleleft G$ corresponds to introducing an equivalence relation on the elements of G , namely $a \equiv_H b$ iff $b \cdot a^{-1} \in H$. Applying $(-)^{\uparrow\downarrow}$ then corresponds to taking equivalence classes, and what results is just the usual construction of the group-theoretic quotient as the group of cosets of a normal subgroup.

3 Imploid-valued flows on 3-valent maps

3.1 Background: graphs, orientations, flows, maps

In this paper we take **graph** to mean finite, undirected graph with loops and/or parallel edges [33, 49]. Formally, such a graph can be considered as a diagram $e \curvearrowright A \begin{array}{c} \xrightarrow{s} \\ \xrightarrow{t} \end{array} V$ where e is a fixpoint-free involution such that $t = s \circ e$. Elements of the set V are called *vertices* and elements of A are called *arcs*, while the functions s and t return the *source* and *target* vertex of an arc. The involution e matches each arc x with an *opposite arc* $-x := e(x)$, and the unordered pair orbit(e, x) = $\{x, -x\}$ is called an *edge*. We write $E := \text{orbits}(e)$ for the set of edges. The *degree* of a vertex v is the cardinality of the set $\{x \mid s(x) = v\}$,


or equivalently of the set $\{x \mid t(x) = v\}$. A graph is said to be *trivalent* (or *3-valent* or *cubic*) if every vertex has degree three. A graph is *connected* if it is neither empty nor the sum of two smaller graphs. A *bridge* in a connected graph is an edge whose removal disconnects the graph. More generally, given a non-empty subset of vertices $V^+ \subseteq V$, the set of edges $C(V^+) = \{\text{orbit}(e, x) \mid s(x) \in V^+ \wedge t(x) \notin V^+\}$ with one end in V^+ and the other outside is called a *cut* (a bridge is a cut containing a single edge). A connected graph is *bridgeless* (or *2-edge-connected*) if it has no bridges; contrarily, it is a *tree* if every edge is a bridge.

An **orientation** of a graph corresponds to the selection of one arc from every edge, or in other words to the choice of a subset $A^+ \subseteq A$ of arcs such that $A = A^+ \uplus -A^+$. Given an orientation A^+ , we define the *inputs* of a vertex as the set $\text{in}(v) := t^{-1}(v) \cap A^+$ and the *outputs* as the set $\text{out}(v) := s^{-1}(v) \cap A^+$. We say that a trivalent graph is *well-oriented* (cf. [28]) if it is oriented so that every vertex either has two inputs and one output (we refer to this as a *negative* vertex), or one input and two outputs (we refer to this as a *positive* vertex). Note that a connected trivalent graph can always be well-oriented by considering any *spanning tree* (we describe this more systematically in Section 3.3).

Let $\Gamma = (V, A, s, t, e)$ be a connected graph equipped with an orientation $A^+ \subseteq A$, and let G be any abelian group. A **group-valued flow** (or G -flow) [44, 22] on Γ (relative to A^+) is a function $\phi : E \rightarrow G$ satisfying the equation

$$\sum_{x \in \text{in } v} \phi(x) = \sum_{x \in \text{out } v} \phi(x) \quad (\text{Kirchhoff's law})$$

at every vertex v . (Observe that the commutativity condition on G is necessary for the equation to be well-defined.) A *nowhere-zero flow* is a flow ϕ such that $\phi(x) \neq 0$ for all $x \in E$. For example, below on the left we show an orientation of the complete graph K_4 with a nowhere-zero \mathbb{Z} -flow, and on the right another graph with a \mathbb{Z}_3 -flow which is *not* nowhere-zero:



Although the notion of flow is defined relative to a given orientation, a (nowhere-zero) flow for one orientation can be transformed into a (nowhere-zero) flow for any other simply by negating the values assigned to some edges. Also, it is easy to prove that a graph cannot admit a nowhere-zero flow unless it is bridgeless, as a corollary of the more general fact that the net flow across any cut is always zero. Finally, it is worth mentioning that many questions about flows on general graphs can be reduced to questions about flows on trivalent graphs (cf. [22]).

In this paper we take **map** to mean cellular embedding of a connected graph Γ into a connected, compact oriented surface [27, 13]. Formally, such an embedding is determined up to orientation-preserving homeomorphism of the underlying surface by the purely combinatorial data of an additional permutation $v : A \rightarrow A$ on the arcs of Γ , assuming that $V \cong \text{orbits}(v)$, and that s factors via the function $x \mapsto \text{orbit}(v, x)$ sending an arc to its v -orbit. The *faces* of a map are then defined as the orbits of the permutation $f := (e \circ v)^{-1}$, and the *genus* g of the underlying surface can be determined from its *Euler characteristic* $\chi := |\text{orbits}(v)| - |\text{orbits}(e)| + |\text{orbits}(f)| = 2 - 2g$. (A *planar map* is a map of genus $g = 0$.) The triple of permutations (v, e, f) (or equivalently the pair (v, e)), which up to isomorphism determines the graph, the surface, and the embedding, is sometimes referred to as a “combinatorial” map. Every combinatorial map also has a *dual map* $(v, e, f)^* := (f^{-1}, e, v^{-1})$ in which the role of vertices and faces is reversed. For example, any *trivalent map* (i.e., a map whose underlying graph is trivalent) on a given surface induces a dual *triangulation* of the same surface, and vice versa. One of the reasons trivalent maps in particular arise as natural objects of study is that they have close connections to the *modular group* $\text{PSL}(2, \mathbb{Z}) \cong \langle v, e \mid v^3 = e^2 = I \rangle$ (cf. [33, 24, 51]).

A **rooted map** is a map equipped with a distinguished root arc $x_0 \in A$, considered up to root-preserving isomorphism. The study of rooted maps was initiated by Tutte in a series of papers on the combinatorics of planar maps [46, 45, 47], taking advantage of the fact that rooted maps have no symmetries and so are easier to count. A rooted map can also be seen as a *map with marked boundary*. While the classical theory of combinatorial maps [23, 49] is formulated in terms of surfaces without boundary (such as the sphere or torus), it is possible to consider boundaries as distinguished faces




Figure 1: Some small examples of rooted 3-valent maps.

representing “holes” in the surface [13]. After removing these faces what is left is an *open graph* in the sense that some edges have only one end attached to a vertex and the other attached to the boundary: see Figure 1 for such depictions of **rooted trivalent maps** as trivalent maps with a marked boundary. Observe that the first two diagrams in Figure 1 represent *two different embeddings* of the same underlying graph: the first into a surface of genus 0 (the open disc), the second into a surface of genus 1 (the open disc with a handle attached – the crossing in the diagram should be thought of as “virtual”, arising from the projection of this higher genus surface down to the page). In contrast, the second pair of diagrams represent *two different rootings* of the same underlying map: if we forget the marking of the root, then a 180° rotation of the disc witnesses an isomorphism between the two diagrams.

It’s not unreasonable to think of the boundary of a rooted trivalent map as a single “external vertex” of arbitrary positive degree², or at least as a *cut* across which values can flow between its interior and its exterior. In fact, the extra structure of the vertex permutation that comes with a combinatorial map naturally enables a more general notion of flow valued in arbitrary (not necessarily abelian) groups, but one important point of divergence with the theory of flows on abstract graphs (see [16] for a discussion) is that in the case of a non-planar map, the local condition on vertices does not automatically extend to arbitrary cuts. After formulating the appropriate definitions, our main results in this section characterize when an *imploid-valued flow* is guaranteed to have such a global extension property, relating the flow at each trivalent vertex to the flow across the boundary.

3.2 Imploid-valued flows

Notation. Suppose given a well-oriented 3-valent map T . We write $[x, y, z]^+ \in T$ to indicate that T contains a positive vertex with output x , input y , and output z as listed in counterclockwise order. Dually, we write $[x, y, z]^- \in T$ to indicate that T contains a negative vertex with input x , output y , and input z as listed in clockwise order.

Definition 3.1. An **imploid-valued flow** on a well-oriented 3-valent map T is a function $\phi : E \rightarrow P$ assigning each edge a value in some left normal imploid P , such that the relation

$$\phi(x) \multimap \phi(y) \leq \phi(z) \tag{3-flow⁺}$$

holds at every positive vertex $[x, y, z]^+ \in T$, and the relation

$$\phi(z) \leq \phi(x) \multimap \phi(y) \tag{3-flow⁻}$$

holds at every every negative vertex $[x, y, z]^- \in T$. (These relations are summarized visually in Figure 2.)

Notice that in Figure 2 we have used colors to help visually distinguish positive vertices (blue) from negative vertices (red), as we will continue doing throughout the paper.

Definition 3.2. A flow $\phi : E \rightarrow P$ is said to be **nowhere-unit** if $I \not\leq \phi(x)$ for all $x \in E$, where I is the unit of P .

Proposition 3.3. A (nowhere-unit) flow $\phi : E \rightarrow P$ may be pushed forward along any (faithful) strong homomorphism $f : P \rightarrow Q$ to obtain a (nowhere-unit) flow $f\phi : E \rightarrow Q$ defined by post-composition.

²This is dual to Tutte’s original treatment of rooted planar triangulations [46], which included an external face of unbounded degree.


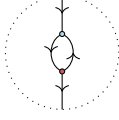



Figure 2: Defining relations for imploid-valued flows on well-oriented 3-valent maps.

Example 3.4. The oriented “bubble”  admits a nowhere-unit P -flow just in case P contains a pair of elements a and b such that $I \not\leq a$ and $I \not\leq b$ and $a \leq b$.

Definition 3.5. Let T be a 3-valent map with boundary ∂T . We say that an arc is an **input of T** (respectively, **output of T**) if its source (resp., target) lies in ∂T (otherwise, the arc is internal to T).

Definition 3.6. Let T be a well-oriented rooted 3-valent map. We say that T is **globally well-oriented** if its orientation contains exactly one output of T , whose target is the root.

For example, here are two different global well-orientations of the third rooted map in Figure 1:



Notation. We write $\partial T = [x_0; x_1, \dots, x_n]$ to indicate that x_0 is the unique output of a globally well-oriented map T , followed by inputs x_1, \dots, x_n in clockwise order around the boundary.

Notation. Let P be any imploid, $\vec{a} = a_1, \dots, a_n \in P$ a list of elements, and $b \in P$ a distinguished element. We write $\vec{a} \rightarrow b$ for the right-associated implication defined inductively by $\cdot \rightarrow b := b$, $(\vec{a}, a) \rightarrow b := \vec{a} \rightarrow (a \rightarrow b)$.

Definition 3.7. Let T be a globally well-oriented 3-valent map, with $\partial T = [x_0; x_1, \dots, x_n]$. We say that a flow ϕ on T satisfies the **global flow condition** if the following relation holds:

$$I \leq (\phi(x_1), \dots, \phi(x_n)) \rightarrow \phi(x_0) \quad (\text{global flow})$$


Before we move on to study the notion of imploid-valued flow, let’s record the following easy observation, which relates it to the classical notion.

Proposition 3.8. Let T be a globally well-oriented 3-valent map, and let G be an abelian group, seen as a discrete symmetric imploid. Then a (nowhere-unit) flow $\phi : E \rightarrow G$ on T is the same thing as a group-valued (nowhere-zero) flow on the underlying graph of T . As a consequence, any flow $\phi : E \rightarrow G$ necessarily satisfies the global flow condition.

Remark 3.9. Since deciding the existence of a proper edge-3-coloring for an abstract cubic graph is NP-complete [20], the problem of deciding for an arbitrary well-oriented 3-valent map T and finite imploid P whether T admits a nowhere-unit P -flow is likewise NP-complete, taking $P = \mathbb{V}$. (Of course the problem might be easier in the case of particular imploids or classes of maps. For example, every bridgeless planar 3-valent map admits a nowhere-zero \mathbb{V} -flow [1].)


3.3 Topological orientations as linear lambda terms

As mentioned in Section 3.1, every connected trivalent graph can be well-oriented. Indeed, given the extra data of an embedding and a rooting, there is a canonical way of reconstructing such an orientation, which we refer to as the *topological orientation* of a rooted 3-valent map. The definition rests on the fact that any rooted 3-valent map T must have one of the following three schematic forms:




In other words, if we remove the vertex incident to the root, either 1) the map becomes disconnected; or 2) it stays connected; or 3) there was no such vertex in the first place.

Definition-Proposition 3.10. *Every rooted 3-valent map may be globally well-oriented by its **topological orientation**, defined as follows by induction on the number of vertices:*



As some examples, here are the topological orientations of the four rooted 3-valent maps displayed in Figure 1:



To be a bit more formal, let $\Theta(n)$ denote the set of isomorphism classes of rooted 3-valent maps with n non-root arcs incident to the boundary. One way of constructing a rooted 3-valent map is to glue a pair of maps T_1 and T_2 by their roots onto a fresh trivalent vertex (case 1), corresponding to a natural family of operations $@ : \Theta(n_1) \times \Theta(n_2) \rightarrow \Theta(n_1 + n_2)$. Another way is to pick one of the non-root arcs on the boundary of T_1 and glue it together with T_1 's root onto a fresh trivalent vertex (case 2), corresponding to a natural family of operations $\lambda_i : \Theta(n + 1) \rightarrow \Theta(n)$ for every $1 \leq i \leq n + 1$. Together, these two operations generate *all* rooted 3-valent maps starting from the trivial rooted map (case 3), and since they naturally extend to operations on globally well-oriented maps (where $@$ introduces a negative vertex and λ_i a positive vertex) this explains the definition of the topological orientation.

Now, we can also observe that $\Theta(n)$ has the structure of a *symmetric operad* [30], meaning that there is a natural family of *composition* operations $\circ_i : \Theta(m + 1) \times \Theta(n) \rightarrow \Theta(m + n)$ ($1 \leq i \leq m + 1$) together with an action of the symmetric group S_n on $\Theta(n)$, satisfying appropriate axioms of associativity, unitality, and equivariance. Composition corresponds to grafting the root of one map onto a boundary arc of another, while the (free) action of the symmetric group corresponds to permuting the boundary arcs.

At this point, the reader with a background in lambda calculus may recognize that our description of $\Theta(n)$ as a symmetric operad exactly mirrors the syntactic structure of **linear lambda terms** with n free variables, reading $@$ as *application* $T_1(T_2)$ and λ_i as *abstraction in the i^{th} variable* $\lambda x_i.T_1$, and interpreting grafting by *substitution* and the symmetric action by *variable exchange*. Indeed, this operadic perspective (cf. [21]) is one way of understanding the one-to-one correspondence between (isomorphism classes of) rooted 3-valent maps and (α -equivalence classes of) linear lambda terms: the latter may be understood as *complete invariants* of rooted 3-valent maps, corresponding to their topological orientations [3, 54]. For instance, under this correspondence, the first two of the four examples above (the topological orientations of the rooted planar tetrahedron and rooted toric tetrahedron) correspond to the *B combinator* $\lambda x.\lambda y.\lambda z.x(yz)$ and *C combinator* $\lambda x.\lambda y.\lambda z.(xz)y$, respectively, in the sense of classical combinatory logic [19] (see [54, Example 1]).

Finally, let us draw attention to the fact that $\Theta(n)$ also has several significant *suboperads*, corresponding to natural subfamilies of maps and subsystems of lambda calculus. By restricting to maps constructed using $@$ and the operation λ_{n+1} we obtain the (non-symmetric) operad $\Theta_0(n)$ of **planar** (i.e., genus 0) rooted 3-valent maps. Note this corresponds to the restriction on linear lambda terms that variables are used in the order they are abstracted (i.e., the forbidding of exchange). By restricting the domain and codomain of the operations $@$ and λ_i to maps with at least one non-root boundary arc, we obtain the operad $\Theta^2(n)$ of **bridgeless** (i.e., 2-edge-connected) rooted 3-valent maps. This corresponds to the

restriction on linear lambda terms that they have no closed subterms, which we refer to as being **unitless**.³

3.4 Topological flows are global

The connection to lambda calculus suggests another way of understanding the global flow condition. In the case of a topological orientation, the problem of building a flow on a rooted 3-valent map may be recast as one of constructing a *linear typing derivation* for the corresponding lambda term. From this the global flow condition follows by an easy proof theory-style argument (we also give a more conceptual explanation in Section 4).

Proposition 3.11. *Let T be a topologically oriented rooted 3-valent map, and let $\partial T = [x_0; x_1, \dots, x_n]$. Then T has a flow ϕ such that $\phi(x_1) = a_1, \dots, \phi(x_n) = a_n$ and $\phi(x_0) = b$ iff the judgment $x_1 : a_1, \dots, x_n : a_n \vdash T : b$ is derivable in the following type system, where the boxed rule is only needed in the non-planar case:*

$$\frac{\Gamma \vdash T_1 : c \quad \Delta \vdash T_2 : a \quad (c \leq a \multimap b)}{\Gamma, \Delta \vdash T_1(T_2) : b} \quad \frac{\Gamma, x : a \vdash T_1 : b}{\Gamma \vdash \lambda x. T_1 : c} \quad (a \multimap b \leq c)$$

$$\frac{x : a \vdash x : a}{\boxed{\frac{\Gamma, y : b, x : a, \Delta \vdash T : c}{\Gamma, x : a, y : b, \Delta \vdash T : c}}}$$

Lemma 3.12. *Let T be a (non-planar) rooted 3-valent map equipped with its topological orientation. Then any flow ϕ on T valued in an arbitrary (symmetric) imploid P satisfies the global flow condition.*

Corollary 3.13. *Let T be a (non-planar) rooted 3-valent map equipped with its topological orientation. If T has a nowhere-unit flow ϕ valued in a (symmetric) imploid P , then T is bridgeless.*

3.5 Non-topological orientations can violate global flow

Conversely, there is no such guarantee for non-topological orientations of rooted 3-valent maps, and indeed, for any such orientation we can always exhibit an explicit *counterexample* in the form of an assignment $\phi : E \rightarrow P$ valued in a specific (symmetric left normal) imploid P , such that ϕ satisfies the local relations (3-flow⁺) and (3-flow⁻) but violates the global flow condition. For this purpose, consider the imploid $P = \hat{2}$ consisting of three linearly ordered elements $0 < 1 < 2$ with the implication $a \multimap b$ defined as follows:

$a \backslash b$	0	1	2
0	2	2	2
1	0	1	2
2	0	0	2

Observe that $\hat{2}$ is isomorphic to the imploid of downsets associated to the unique idempotent skew monoid with two elements $2 = (\{1, 2\}, \leq, \max, 1)$.

Lemma 3.14. *Let T be a rooted 3-valent map equipped with a well-orientation A^+ containing the root x_0 as an output, and the remaining boundary edges x_1, \dots, x_n as either inputs or outputs. If A^+ is non-topological then there is a $\hat{2}$ -flow ϕ such that $\phi(x_0) = 0$, and $\phi(x_i) = 1$ or 2 for all $1 \leq i \leq n$.*


As a corollary, we obtain the following characterization:

Theorem 3.15. *Let T be a globally well-oriented 3-valent map. The following are equivalent:*

1. T is topologically oriented, i.e., has the orientation of a linear lambda term.
2. Every $\hat{2}$ -flow on T satisfies the global flow condition.
3. Every P -flow on T satisfies the global flow condition, for any symmetric imploid P .

³ In other words, a unitless term is one that can be constructed in the absence of the empty (unit) context. This property (with a very minor technical variation) was called being “indecomposable” in [54].

Example 3.16. A pair of non-global $\hat{2}$ -flows on non-lambda terms:



Theorem 3.15 says in a sense that the global flow condition acts as a “correctness criterion” in the terminology of linear logic [14]. (Indeed it has similarities with de Groote’s algebraic criterion for intuitionistic proof-nets [9]. We elaborate on the relationship with proof-nets a bit more in Section 5.) It may be surprising that such a small imploid is powerful enough to distinguish topological orientations from non-topological ones, although this is consistent with the fact that the topological orientation of a rooted 3-valent map can be computed efficiently, in a single depth-first traversal [3].

3.6 Fundamental imploids and universal flows

Following a familiar pattern of abstract nonsense, it is possible to bundle the notion of an imploid-valued flow into that of the *fundamental imploid* of a well-oriented 3-valent map.

Definition 3.17. The **fundamental imploid** of a well-oriented 3-valent map T is the left normal imploid $P[T]$ freely generated from the edges of T modulo the relations in Figure 2.

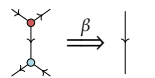
The function $[-] : E \rightarrow P[T]$ sending each edge to the corresponding generator of the fundamental imploid tautologically defines a flow, and by the universal property of the quotient, any other flow $\phi : E \rightarrow P$ uniquely extends to a strong homomorphism of imploids $\bar{\phi} : P[T] \rightarrow P$ such that $\phi = \bar{\phi}[-]$. Moreover, ϕ is nowhere-unit just in case $\ker \bar{\phi}$ does not contain a generator. The *fundamental symmetric imploid* $\tilde{P}[T]$ can be defined similarly, with analogous properties for flows valued in symmetric left normal imploids.

Although this abstract definition of the fundamental imploid (reminiscent of the *fundamental quandle* of a knot [25]) allows us to express flow concepts in a more uniform language, it doesn’t immediately provide us much help in understanding the space of possible flows over a given 3-valent map. To get a more concrete handle on this space, in the rest of the paper we develop a computational perspective on imploids and flows that is inspired by their connections to combinatory logic and type theory.

4 Rewriting and pullback of flows

4.1 Background: beta reduction and eta expansion

All of the computational power of the lambda calculus lies in the rule of β -reduction $(\lambda x.T_1)(T_2) \rightarrow T_1[T_2/x]$, and as Church originally showed, the problem of determining if a general lambda term has a β -normal form is undecidable [6]. On the other hand, if one imposes linearity the rule becomes much more tractable: every linear lambda term has a β -normal form, and the problem of computing it is complete for polynomial time [29]. Graphically, β -reduction corresponds to the operation of

“unzipping” a pair of trivalent vertices of opposite polarity: ; we refer to the matching pair of vertices as a β -redex. This rule can in principle be applied whenever such a configuration appears in a well-oriented 3-valent map, but the fact that it corresponds to β -reduction of lambda terms means that it *preserves topological orientation*. The graphical rule also manifestly preserves planarity, and β -reduction correspondingly restricts to an operation on planar terms.⁴ Finally, that β -reduction

⁴The precise computational complexity of β -normalization for planar terms is a natural question, which is open as far as I am aware. (Mairson’s proof of PTIME-hardness for linear lambda calculus [29] is based on an encoding of boolean circuits that uses non-planarity in an essential way.)




Figure 3: Flow relations for 1-valent and 2-valent vertices.


restricts to an operation on unitless terms implies it preserves 2-edge-connectedness when restricted to topological orientations (although it can lead to disconnected maps when applied to non-topological orientations). Dual to β -reduction is the less computationally interesting (but still logically important) rule of η -expansion $T \rightarrow \lambda x.T(x)$. Graphically, this rule corresponds to “bubbling” an oriented edge,




4.2 Pullback of flows

It is not hard to check that an imploid-valued flow can always be *pulled back* along a β -reduction or an η -expansion, in a suitable sense. To make this statement more precise, it is useful to first liberalize the notion of flow on a trivalent map to allow for arbitrary subdivision of edges by 2-valent vertices: we assume these to be well-oriented (one input, one output) and to satisfy the natural flow relation shown at the right in Figure 3. Edge subdivision provides us an additional degree of flexibility when relating one flow to another, but it is always possible to recover a flow on a strictly 3-valent map by choosing any of the component values along a subdivided edge.

Consider again the rule of β -reduction, now written in reverse:



Here we have annotated the rule as it acts in the backwards direction on flows, as a β -expansion taking a pair of subdivided edges and “re-zipping” them into a β -redex. That this is a well-defined operation on flows reduces to the *totality* of the implication $a \multimap b$ and its monotonicity properties (imp). Dually, pullback along η -expansion



may be justified by (imp) and *uniqueness* of implication.

Formally, these pullback operations on flows may be analyzed in terms of fundamental imploids as follows. The rules of β -reduction and η -expansion both lift to *strong homomorphisms* $P[T_L] \rightarrow P[T_R]$ from the fundamental imploid of the map on the left-hand side to that of the right-hand side. By the universal properties of $P[T_L]$ and $P[T_R]$, pulling a flow $\phi : E_R \rightarrow P$ back along these operations reduces to pre-composing $\bar{\phi} : P[T_R] \rightarrow P$ with the corresponding homomorphism $P[T_L] \rightarrow P[T_R]$. Moreover, these homomorphisms are *boundary-preserving* in the sense that they fix all of the generators in $\partial T_L = \partial T_R$, which implies that the corresponding transformations can be applied locally anywhere inside a larger


flow. On the other hand, observe that nothing guarantees we can push a flow $\phi : E_L \rightarrow P$ forward along $P[T_L] \rightarrow P[T_R]$, and it is easy to come up with counterexamples to such a principle for β -reduction (e.g., taking $P = \mathbb{Z}_2$, $a_1 = b_2 = 1$, $b_1 = a_2 = 0$ in the first diagram above).

4.3 Imploid moves and topological completeness


All this suggests a more powerful theory of rewriting for well-oriented maps, where each rule $T_L \Rightarrow T_R$ corresponds to a principle for pulling back flows on T_R to flows on T_L . To gain the full benefits of such a theory, it is natural to further generalize the notion of flow to allow for 1-valent vertices, with the relations shown in Figure 3; we refer to maps containing only vertices of positive degree ≤ 3 as *essentially trivalent maps*.

Definition 4.1. A transformation $T_L \Rightarrow T_R$ between a pair of well-oriented essentially 3-valent maps with the same boundary $\partial := \partial T_L = \partial T_R$ is called a **(symmetric) imploid move** if it is realizable by a strong homomorphism $P[T_L] \rightarrow P[T_R]$ (respectively, $\tilde{P}[T_L] \rightarrow \tilde{P}[T_R]$) fixing every element in ∂ .

Proposition 4.2. The following are imploid moves:



Proposition 4.3. The following are symmetric imploid moves:



The rules above (many interderivable) do not give a complete set of generators for imploid moves. However, they are *topologically complete* in the following sense.

Proposition 4.4. All of the moves listed in Propositions 4.2 and 4.3 preserve topological orientation.

Theorem 4.5. Let V_n be the n -**spine** defined inductively by $V_0 := \downarrow$, $V_{n+1} := \begin{matrix} V_n \\ \downarrow \\ \text{blue vertex} \\ \downarrow \\ V_n \end{matrix}$, and let $V'_n := \begin{matrix} V_n \\ \downarrow \\ \text{red vertex} \\ \downarrow \\ V_n \end{matrix}$. Starting from the V'_n , the imploid moves in (the first three rows of) Prop. 4.2 generate all rooted essentially 3-valent planar maps with their topological orientation. With the addition of (any of) the symmetric imploid moves in Prop. 4.3, they generate all rooted essentially 3-valent maps of arbitrary genus.




Figure 4: Defining relations for polarized flows.

As an immediate corollary of Theorem 4.5 we get another proof of Lemma 3.12: any flow on a topologically oriented map T (with $n + 1$ boundary arcs counting the root) can be pulled back to a flow on V'_n while preserving the boundary, and so the global flow condition on T may be read off directly from the local flow conditions on V'_n .

Although the link may seem astonishing at first, this topological completeness theorem is closely related to the classical result in combinatory logic that the combinators B , C , and I form a complete basis for linear lambda terms [19, §9F], as well as its planar restriction stating that B and I form a complete basis for planar lambda terms. It may be further refined by restricting to *non-unital imploid moves* (that is, moves not involving 1-valent vertices), which generate all bridgeless essentially 3-valent maps with their topological orientation (i.e., unitless linear lambda terms).

Tantalizingly, these completeness results also appear connected to a basic but motivating result in the theory of *knotted trivalent graphs* (KTGs), which in one formulation states that any KTG (and hence any knot) can be generated from the planar tetrahedron \triangleleft and crossed tetrahedron \triangleleft using unzip, bubbling, connect sum, and the unknot. (See [42, Theorem 1] and [2, Appendix]. Indeed, Thurston’s article inspires our terminology of “unzipping” and “bubbling” for β and η , backing an analogy made by Buliga [5].) This strong formal similarity suggests it could be worthwhile to develop a more refined treatment of the exchange law as a *braiding* on linear lambda terms (cf. [31]), moving up a dimension from 3-valent maps to KTGs. For example, it may be interesting to relate imploid flows to *qualgebra colorings* of KTGs [28].

Late in the development of the theory of imploid-valued flows described here (and motivated by the parallel connections discussed in [55]), I discovered with excitement that it has much in common with the “graphic theory of associativity” proposed by Tamari in a relatively obscure conference publication [40], which built on seeds planted thirty years earlier in his thesis [39]. Tamari’s approach can be seen as slanted towards monoids rather than imploids, but is in a sense more foundational, beginning with the minimalistic algebraic structure of a *partial binary operation* (or “bin”) and considering how to match different principles of associativity with different well-oriented bridgeless planar 3-valent maps. (Since bins are partial, even the unzip move isn’t available, and the result is an infinite hierarchy of independent, higher associative laws.)

5 Polarized flows and bidirectional typing

We close by briefly sketching another perspective on imploid-valued flows that makes explicit their connection to linear logic proof-nets [14], while also being implicitly tied to the important type-theoretic concept of *bidirectional typing* [32, 11] (and related ideas such as *polarized subtyping* [10]).

One reason the notion of imploid-valued flow is more subtle than the classical notion of abelian group-valued flow is that the defining relations (Figure 2) intertwine the preorder with the implication operation. In the corresponding type system for linear lambda terms (Prop. 3.11), *subtyping* is built into the rules for typing application and abstraction, so that typing a term reduces to checking a big collection of constraints of the form $a \multimap b \leq c$ or $c \leq a \multimap b$.

In contrast, the definition of a *polarized flow* (see Figure 4) employs a more rigid separation between \multimap and \leq , relying on the presence of both 3-valent and 2-valent vertices (for simplicity, we leave out 1-valent vertices from this discussion: they can be dealt with similarly to 3-valent vertices). Formally, now we are working with maps which are not merely oriented but also *signed* (cf. [26]), that is, equipped with a function $\pi : E \rightarrow \{+1, -1\}$ assigning each edge a positive (red) or negative (blue) polarity. We assume that π is *proper* in the sense that the sum of polarities around each 3-valent vertex is either $+1$ or -1 , and the sum around each 2-valent vertex is 0 . We likewise assume that π is compatible with the

underlying orientation in the sense that the right half of Figure 4 can be overlaid onto Figure 2. (This means that the orientation markers are for the most part redundant, although it is still necessary to distinguish vertices with one negative input and one positive output, which we color in white, from vertices with one positive input and one negative output, which we color in black.)

Definition 5.1. Let T be a well-oriented 3-valent map, and let π be a well-polarized, well-oriented essentially 3-valent map. We say that π is a **polarization of T** , written $\pi \sqsubset T$, if the underlying oriented map of π is an edge subdivision of T .

Definition-Proposition 5.2. Any well-oriented 3-valent map T has a **minimal polarization** $\pi_T \sqsubset T$, which can be constructed by applying the two fixed patterns of polarization for 3-valent vertices, and then subdividing any edge whose ends have opposite polarity by a 2-valent vertex (either white or black).

Proposition 5.3. Black vertices of π_T correspond to β -redices of T .

Clearly, any polarized flow is a flow, by forgetting polarities. Conversely, any flow on a 3-valent map can be turned into a polarized flow after sufficient edge subdivision, although the resulting polarization may not necessarily be the minimal one.

Our interest in polarization is that it gives another way of decomposing flows inductively, different from and complementary to the inductive definition of the topological orientation.

Definition 5.4. Let π be a well-polarized, well-oriented essentially 3-valent map. The **w–b orientation** of π is given by reversing the orientation of negative edges (hence white vertices become sources, black vertices become sinks).

Proposition 5.5. If the underlying orientation of π is topological then its w–b orientation is acyclic.

By Prop. 5.5, any topological π can be decomposed into a forest of rooted trees, with their leaves glued together by white vertices and their roots glued together by black vertices. (In the literature on proof-nets, these trees are called *formulas*, white vertices are called *axioms*, and black vertices are called *cuts*.) This makes it possible to reduce the problem of building a polarized flow on π to the following recipe:

1. At each white vertex w , assign its negative and positive ends some values w^- and w^+ such that $w^- \leq w^+$.
2. Apply the rules on the right side of Figure 4 to propagate values to the remaining edges.
3. At each black vertex β whose ends have now been assigned values a^+ and b^- , check that $a^+ \leq b^-$.

In more abstract terms following the discussion of Section 3.6, any topological π has a **universal polarized flow** valued in the imploid freely generated from its white vertices modulo the relations induced on its black vertices.

The universal polarized flow on (the minimal polarization of) a linear lambda term is analogous to its *principal type-scheme* [17, 18], but with the difference that it includes explicit subtyping constraints corresponding to β -redices. A consequence is that the universal polarized flow can be computed very efficiently in a single traversal of the term, without performing any β -normalization either explicitly or implicitly (getting around Mairson’s P-completeness result [29]). Of course, complexity may come back into the picture if we want to *instantiate* the universal polarized flow to obtain (say) a flow valued in a free imploid (or a nowhere-unit flow valued in a finite imploid), since this involves the discharging of these subtyping constraints.

An extended example (paying tribute to one of Tutte’s earliest contributions to map coloring [43]) may be found in Appendix B.

References

- [1] K. Appel and W. Haken. Every planar map is four colorable. parts i and ii. *Illinois Journal of Mathematics*, 21:429–567, 1977.
- [2] Dror Bar-Natan and Zsuzsanna Dancso. Homomorphic expansions for knotted trivalent graphs. *J. Knot Theory and Its Ramifications*, 22, 2013.

- [3] O. Bodini, D. Gardy, and A. Jacquot. Asymptotics and random sampling for BCI and BCK lambda terms. *Theoretical Computer Science*, 502:227–238, 2013.
- [4] John Bourke and Stephen Lack. Braided skew monoidal categories. December 2017. arXiv:1712.0827.
- [5] Marius Buliga. Graphic lambda calculus. *Complex Systems*, 22(4):311–360, 2013.
- [6] Alonzo Church. An unsolvable problem of elementary number theory. *American Journal of Mathematics*, 58(2):345–363, 1936.
- [7] Julien Courtiel, Karen Yeats, and Noam Zeilberger. Connected chord diagrams and bridgeless maps. arXiv:1611.04611, October 2017.
- [8] B. J. Day and M. L. Laplaza. On embedding closed categories. *Bull. Austral. Math. Soc.*, 18:357–371, 1978.
- [9] Philippe de Groote. An algebraic correctness criterion for intuitionistic multiplicative proof-nets. *Theoretical Computer Science*, 224(1–2):115–134, 1999.
- [10] Stephen Dolan. *Algebraic Subtyping*. Phd thesis, University of Cambridge, 2016.
- [11] Joshua Dunfield and Neelakantan R. Krishnaswami. Complete and easy bidirectional typechecking for higher-rank polymorphism. *SIGPLAN Notices*, 48(9):429–442, September 2013.
- [12] Samuel Eilenberg and G. Max Kelly. Closed categories. In *Proceedings of the Conference on Categorical Algebra*, pages 421–562. Springer-Verlag, 1966.
- [13] Bertrand Eynard. *Counting Surfaces*. Number 70 in Progress in Mathematical Physics. Birkhäuser, 2016.
- [14] Jean-Yves Girard. Linear logic. *Theoretical Computer Science*, 50:1–102, 1987.
- [15] Georges Gonthier. Formal proof—the Four Color Theorem. *Notices of the AMS*, 55(11):1382–1393, 2008.
- [16] Andrew Goodall, Thomas Krajewski, Guus Regts, and Lluís Vena. A tutte polynomial for maps. *Combinatorics, Probability and Computing*, page to appear, 2018.
- [17] J. Roger Hindley. The principal type-scheme of an object in combinatory logic. *Transactions of the American Mathematical Society*, 146:29–60, 1969.
- [18] J. Roger Hindley. Bck-combinators and linear lambda-terms have types. *Theoretical Computer Science*, 64(1):97–105, 1989.
- [19] J. Roger Hindley. *Basic Simple Type Theory*. CUP, 1997.
- [20] Ian Holyer. The np-completeness of edge-coloring. *SIAM Journal on Computing*, 10(4):718–720, 1981.
- [21] Martin Hyland. Classical lambda calculus in modern dress. *Mathematical Structures in Computer Science*, 27(5):762–781, 2017.
- [22] F. Jaeger. Flows and generalized coloring theorems in graphs. *J. Combinatorial Theory Series B*, 26:205–216, 1979.
- [23] Gareth A. Jones and David Singerman. Theory of maps on orientable surfaces. *Proceedings of the London Mathematical Society*, 37:273–307, 1978.
- [24] Gareth A. Jones and David Singerman. Maps, hypermaps, and triangle groups. In L. Schneps, editor, *The Grothendieck Theory of Dessins d’Enfants*. CUP, 1994.
- [25] David Joyce. A classifying invariant of knots: the knot quandle. *J. Pure Applied Algebra*, 23:37–65, 1982.
- [26] Louis H. Kauffman. A tutte polynomial for signed graphs. *Discrete Applied Mathematics*, 25:105–127, 1989.
- [27] Sergei K. Lando and Alexander K. Zvonkin. *Graphs on Surfaces and Their Applications*. Number 141 in Encyclopaedia of Mathematical Sciences. Springer, 2004.
- [28] Victoria Lebed. Qualgebras and knotted 3-valent graphs. *Fundamenta Mathematicae*, 230(2):167–204, 2015.
- [29] Harry G. Mairson. Linear lambda calculus and ptime-completeness. *J. Functional Programming*, 14(6):623–633, November 2004.
- [30] Martin Markl, Steve Schneider, and Jim Stasheff. *Operads in Algebra, Topology and Physics*, volume 96 of *Mathematical Surveys and Monographs*. AMS, 2002.
- [31] Paul-André Melliès. Ribbon tensorial logic, 2018. This volume.
- [32] Benjamin C. Pierce and David N. Turner. Local type inference. *ACM Transactions on Programming Languages and Systems*, 22(1):1–44, January 2000.
- [33] Jean-Pierre Serre. *Trees*. Springer-Verlag, 1980. Translated from the French by John Stillwell.
- [34] N. J. A. Sloane. The On-Line Encyclopedia of Integer Sequences, 2018. Published electronically at <https://oeis.org>.
- [35] Richard Statman. *Structural complexity of proofs*. Phd thesis, Stanford University, 1974.

- [36] Ross Street. Skew-closed categories. *J. Pure Applied Algebra*, 217(6):973–988, 2013.
- [37] Kornel Szlachányi. Skew-monoidal categories and bialgebroids. *Advances in Mathematics*, 231(3–4):1694–1730, 2012.
- [38] P. G. Tait. Remarks on the colouring of maps. *Proc. Royal Soc. Edinburgh*, 10(4):501–503, 1880.
- [39] Dov Tamari. *Monômes préordonnés et chaînes de Malcev*. Thèse, Université de Paris, 1951. Partially published in *Bull. Soc. Math. France* 82 (1954), 53–96.
- [40] Dov Tamari. A graphic theory of associativity and word-chain patterns. In A. Dold and B. Eckmann, editors, *Combinatorial Theory*, volume 969 of *Lecture Notes in Mathematics*, pages 302–320. Springer, 1982.
- [41] Robin Thomas. An update on the four-color theorem. *Notices of the American Mathematical Society*, 45(7):848–859, 1998.
- [42] Dylan P. Thurston. The algebra of knotted trivalent graphs and turaev’s shadow world. In *Geometry & Topology Monographs*, volume 4 of *Invariants of knots and 3-manifolds (Kyoto 2001)*, pages 337–362. 2004.
- [43] W. T. Tutte. On hamiltonian circuits. *Journal of the London Mathematical Society*, 21:98–101, 1946.
- [44] W. T. Tutte. A contribution to the theory of chromatic polynomials. *Can. J. Math.*, 6:80–91, 1954.
- [45] W. T. Tutte. A census of hamiltonian polygons. *Can. J. Math.*, 14:402–417, 1962.
- [46] W. T. Tutte. A census of planar triangulations. *Can. J. Math.*, 14:21–38, 1962.
- [47] W. T. Tutte. A census of planar maps. *Can. J. Math.*, 15:249–271, 1963.
- [48] W. T. Tutte. On the algebraic theory of graph colourings. *Journal of Combinatorial Theory*, 1:15–50, 1966.
- [49] W. T. Tutte. *Graph Theory*, volume 21 of *Encyclopedia of Mathematics and its Applications*. Addison-Wesley, 1984.
- [50] W. T. Tutte. *Graph Theory as I Have Known it*. Oxford, 1998.
- [51] Samuel Vidal. *Groupe Modulaire et Cartes Combinatoires: Génération et Comptage*. Phd thesis, Université Lille I, France, July 2010.
- [52] David Yetter. Quantales and (non-commutative) linear logic. *J. Symbolic Logic*, 55:41–64, 1990.
- [53] Noam Zeilberger. Counting isomorphism classes of β -normal linear lambda terms. arXiv:1509.07596, 2015.
- [54] Noam Zeilberger. Linear lambda terms as invariants of rooted trivalent maps. *J. Functional Programming*, 26, 2016.
- [55] Noam Zeilberger. A sequent calculus for a semi-associative law. In *Formal Structures for Computation and Deduction (FSCD 2017)*, pages 33:1–33:16, 2017.
- [56] Noam Zeilberger and Alain Giorgetti. A correspondence between rooted planar maps and normal planar lambda terms. *Logical Methods in Computer Science*, 11(3:22):1–39, 2015.

A Proofs of results

Section 2

Proof of Prop. 2.6. To derive (exch) from (dni), we first derive an alternate form of the composition law:

$$a \multimap b \leq (b \multimap c) \multimap (a \multimap c) \quad (\text{comp}')$$

Derivation of (comp') from (dni):

$$\begin{aligned} a \multimap b &\leq ((a \multimap b) \multimap (a \multimap c)) \multimap (a \multimap c) && (\text{dni}) \\ &\leq (b \multimap c) \multimap (a \multimap c) && (\text{comp in negative pos.}) \end{aligned}$$

Derivation of (exch) from (comp') and (dni):

$$\begin{aligned} a \multimap (b \multimap c) &\leq ((b \multimap c) \multimap c) \multimap (a \multimap c) && (\text{comp}') \\ &\leq b \multimap (a \multimap c) && (\text{dni in negative pos.}) \end{aligned}$$

In the other direction, we derive (dni) from (exch) under assumption of left normality:

$$\begin{aligned} I &\leq (a \multimap b) \multimap (a \multimap b) && (\text{id}) \\ &\leq a \multimap ((a \multimap b) \multimap b) && (\text{exch}) \\ a &\leq (a \multimap b) \multimap b && (\text{left normality}) \end{aligned}$$

□

Proof of Prop. 2.10. Unwinding definitions, it is easy (but instructive) to check that the reverse inclusions

$$\begin{aligned} (R \bullet S) \bullet T &\supseteq R \bullet (S \bullet T) \\ I \bullet R &\supseteq R \\ R &\supseteq R \bullet I \end{aligned}$$

follow from the axioms (comp), (id), and (unit), respectively, for any upsets $R, S, T \sqsubset^\uparrow P$ of an imploid. (We leave this as a fun warmup exercise for the reader!) □

Proof of Observation 2.11. 1. Suppose P satisfies (dni), and let $p \in S \bullet R$. By definition, there exists r such that $r \multimap p \in S$ and $r \in R$. But then $(r \multimap p) \multimap p \in R$ by (dni) and upwards closure, hence $p \in R \bullet S$. Conversely, suppose that $R \bullet S \supseteq S \bullet R$ for all $R, S \sqsubset^\uparrow P$, and consider $R = a^\uparrow$, $S = (a \multimap b)^\uparrow$ where $a, b \in P$ are arbitrary. It is easy to check that $b \in S \bullet R$, but then if $b \in R \bullet S$ there must exist p such that $a \leq p \multimap b$ and $a \multimap b \leq p$, implying $a \leq (a \multimap b) \multimap b$.

2. Suppose P satisfies (exch), and let $p \in (R \bullet T) \bullet S$. By definition, there exists s such that $s \multimap p \in R \bullet T$ and $s \in S$, and t such that $t \multimap (s \multimap p) \in R$ and $t \in T$. But then $s \multimap (t \multimap p) \in R$ by (exch) and upwards closure, from which $p \in (R \bullet S) \bullet T$. Conversely, suppose that $(R \bullet S) \bullet T \supseteq (R \bullet T) \bullet S$ for all $R, S, T \sqsubset^\uparrow P$, and consider $R = a \multimap (b \multimap c)^\uparrow$, $T = b^\uparrow$, $S = a^\uparrow$ where $a, b, c \in P$ are arbitrary. It is easy to check that $c \in (R \bullet T) \bullet S$, but then if $c \in (R \bullet S) \bullet T$ there must exist p and q such that $a \multimap (b \multimap c) \leq p \multimap (q \multimap c)$ and $b \leq p$ and $a \leq q$, implying $a \multimap (b \multimap c) \leq b \multimap (a \multimap c)$. □

Proof of Prop. 2.12. Again, it is easy to check that for any downsets $J, K, L \sqsubset^\downarrow M$ of a skew monoid, the inclusions

$$\begin{aligned} K \multimap L &\subseteq (J \multimap K) \multimap (J \multimap L) \\ I &\subseteq K \multimap K \\ I \multimap K &\subseteq K \end{aligned}$$

follow from axioms (assocr), (lunit), and (runit), respectively. □

Proof of Prop. 2.14. We just have to check $(a \multimap b)^{\uparrow\downarrow} \equiv a^{\uparrow\downarrow} \multimap b^{\uparrow\downarrow}$, which reduces to showing that

$$R \ni a \multimap b \iff \forall S. S \ni a \Rightarrow \exists c. R \ni c \multimap b \wedge S \ni c.$$

The implication from left to right is immediate taking $c = a$, while the implication from right to left is immediate taking $S = a^{\uparrow}$. \square

Proof of Corollary 2.18. This is just an adaptation of the standard free monoid construction (appropriately dualized) to the skew monoidal setting: by mechanically unrolling the definition of $!R$ and applying the skew monoid laws, we can show that $!R$ is the greatest comonoid in \check{P} under R . \square

Proof of Prop. 2.20. That $a \leq b$ entails $I \leq a \multimap b$ immediately implies $\#_R$ is an extension of \leq (and reflexive), since $R \supseteq I^{\uparrow}$. For transitivity, suppose that $a \multimap b \in R$ and $b \multimap c \in R$. Applying (comp) and upwards closure to the latter assumption we obtain $(a \multimap b) \multimap (a \multimap c) \in R$, and hence $a \multimap c \in R$ by deductive closure. \square

Proof of Prop. 2.22. The operation

$$DNI = S \mapsto S \wedge \{c \mid a \in S, (a \multimap b) \multimap b \leq c\}$$

defines a monotone operator $DNI : \check{P} \rightarrow \check{P}$ on the complete lattice \check{P} , and so by the Knaster-Tarski theorem we can compute the dni-closure of R as the greatest fixed point of DNI containing R . (NB: since the ordering in \check{P} is reverse inclusion, $\wedge = \cup$, and “greatest” = “is contained in any other fixed point containing R ”.) \square

Proof of Prop. 2.23. The equivalence (1) \Leftrightarrow (2) is similar to the proof of Observation 2.11(1). For (1) \Rightarrow (3): (3a) Let $a \in R$. Then $(a \multimap a) \multimap a \in R$ by dni-closure, hence $I \multimap a \in R$ by (id) in negative position and upwards closure. (3b) Let $a \multimap b \in R$, and $c \in P$. Then $((a \multimap b) \multimap (a \multimap c)) \multimap (a \multimap c) \in R$ by dni-closure, hence $(b \multimap c) \multimap (a \multimap c) \in R$ by (comp) in negative position and upwards closure.

For (3) \Rightarrow (1): Let $a \in R$ and $b \in P$. Then $I \multimap a \in R$ by (3a), and $(a \multimap b) \multimap (I \multimap b) \in R$ by (3b). But then $(a \multimap b) \multimap b \in R$ by (unit) and upwards closure.

Finally, the equivalence (3) \Leftrightarrow (4) is immediate after noting that the closure condition $[b \multimap c \in R \Rightarrow (a \multimap b) \multimap (a \multimap c) \in R]$ is always valid for any upset R . \square

Proof of Corollary 2.24. The first part follows easily from case (2) of Prop. 2.23 and some applications of (assocr), while the second follows from the first and the formula for the deductive closure $!R$ in Corollary 2.18. \square

Proof of Prop. 2.25. We’ve already verified that P/R is an imploid by Prop. 2.20 and Prop. 2.23(4), while left normality ($I \#_R a \multimap b \Rightarrow a \#_R b$) follows from the (unit) axiom and upwards closure. \square

Proofs of Propositions 2.26, 2.27, 2.29 and 2.30. Immediate. \square

Proof of Prop. 2.31. The function $[-]$ is automatically a strong homomorphism, and the claim that $R = \ker[-]$ amounts to $a \in R$ iff $I \#_R a$, which follows from Prop. 2.23(4). Now suppose that $f : P \rightarrow Q$ is a (potentially lax) homomorphism into some left normal imploid, and that $R \subseteq \ker f$. We define $\bar{f} : P/R \rightarrow Q$ to be identical to f on elements of P , and just have to check $a \#_R b$ implies $f(a) \leq_Q f(b)$. Well, if $a \multimap b \in R$ then $I_Q \leq f(a \multimap b)$ by assumption that $R \subseteq \ker f$. But then $f(a) \leq_Q f(b)$ follows from the assumptions that f is a homomorphism and that Q is left normal. \square

Section 3

Proof of Prop. 3.11. This is immediate from the definition of flow and our operadic description of topological orientations in Section 3.3, although a small amount of care should be taken in interpreting the rule of variable exchange. Formally, a 3-valent map-with-boundary corresponds to a linear *term-in-context*, that is, a lambda term T equipped with a specific ordering x_1, \dots, x_n of its free variables. So even though the variable exchange rule is written (in the traditional way) with the same term T appearing in both the premise and the conclusion, in fact these correspond to different maps-with-boundary when interpreted as terms-in-context. See [54] for a more detailed discussion. \square

Proof of Lemma 3.12. By Prop. 3.11, to prove that every topological flow is global we have to show that

$$x_1 : a_1, \dots, x_n : a_n \vdash T : b \text{ implies } I \leq (a_1, \dots, a_n) \multimap b.$$

When T is planar, this follows by induction on exchange-free derivations, appealing in the application case

$$\frac{\Gamma \vdash T_1 : c \quad \Delta \vdash T_2 : a}{\Gamma, \Delta \vdash T_1(T_2) : b} \quad (c \leq a \multimap b)$$


to the following lemma, proved by induction on Δ (here we take Γ and Δ to range over lists of elements of P):

Lemma A.1. *If $I \leq \Gamma \multimap (a \multimap b)$ and $I \leq \Delta \multimap a$ then $I \leq (\Gamma, \Delta) \multimap b$.*


In the non-planar case we also have to deal with the exchange rule, but then we simply appeal to (exch). \square

Proof of Corollary 3.13. By Lemma 3.12 and the characterization of bridgeless maps as unitless terms [54, Proposition 7.3]. \square


Proof of Lemma 3.14. We first note that any well-oriented 3-valent map admits two different constant $\hat{2}$ -flows in which either every edge is assigned the value 1 or every edge is assigned the value 2 (these are trivially flows since $1 \multimap 1 = 1$ and $2 \multimap 2 = 2$). We now proceed by case analysis of the possible orientations of T , and induction on the number of vertices:




By assumption, either T_1 or T_2 must be non-topologically oriented. In the first case, by the induction hypothesis T_1 has a flow in which all its non-root boundary arcs are assigned 1 or 2 and its root is assigned 0. Then after taking the constant-2 (or constant-1) flow on T_2 , we assign 0 to T 's root to obtain our desired flow (using that $0 = y \multimap 0$ for all $y > 0$):





Symmetrically, in the second case we apply the induction hypothesis to T_2 and take the constant-2 (or constant-1) flow on T_1 (using that $y \leq 0 \multimap 0 = 2$):








In the first case we take the constant-1 flow on T_1 and the constant-2 flow on T_2 (using that $2 \rightarrow 1 = 0$), and in the second we take the constant-2 flow on T_1 and either the constant-2 or constant-1 flow on T_2 (using that $0 \rightarrow x = 2$ for all x):



In the first case, by assumption, T_1 must be non-topologically oriented. We apply the induction hypothesis (again using that $y \rightarrow 0 = 0$ for all $y > 0$):




In the second case we simply take the constant-2 flow on T_2 and assign 0 to the root (using that $0 \rightarrow 2 = 2$):




Since T_1 has two outputs, it is not topological. We apply the induction hypothesis taking the root

of T_1 to be the arc counterclockwise from the root of T (again using that $y \leq 0 \rightarrow 0 = 2$):



(Notice this case accounts for why we can't restrict to globally well-oriented maps in the induction.)



Impossible by assumption that T is equipped with a non-topological orientation. □


Proof of Theorem 3.15. Implied by Lemmas 3.12 and 3.14. □

Section 4

Proof of Prop. 4.2. The existence of boundary-preserving homomorphisms corresponding to the first four moves reduces to the definition of a preorder (i.e., reflexivity + transitivity) and of the local flow relations for 3-valent vertices (Figure 2) combined with monotonicity of implication (imp). As already mentioned, the justification of β and η amounts to totality and uniqueness, respectively, of the implication operation, combined with (imp). Similarly, the τ , ι , and δ transformations may be justified by appeal to the imploid axioms (comp), (id), and (unit), respectively, with the right-to-left direction of ι corresponding to left normality. Finally, all of the remaining moves may be derived from the above. (For example, ρ can be derived using β and $\bar{\tau}$, which can in turn be derived from τ using β and η .) We leave the details as an exercise for the reader. □

Proof of Prop. 4.3. The χ move reduces directly to the (exch) axiom, while C can be easily derived from χ in combination with ι and η . The γ move reduces directly to (dni), which is valid in any symmetric left normal imploid (Prop. 2.6). □

Proof of Prop. 4.4. To be completely unambiguous, we should first clarify that the definition of topological orientation (Def.-Prop. 3.10) extends to rooted maps with vertices of degree 2 or 1 by the addition of the following cases:



Since the moves listed in Propositions 4.2 and 4.3 only include vertices of degree < 3 with their unique possible topological orientation, all that needs to be checked is the treatment of trivalent vertices. We already explained in Section 4.1 that the graphical β and η moves preserve topological orientation, precisely because they restrict to the standard rewriting rules from lambda calculus when interpreted on linear terms. One way of seeing that τ preserves topological orientation is to verify that it corresponds to the following transformation, which takes linear terms to linear terms:


$$\lambda x.T \xrightarrow{\tau} \lambda x.\lambda y.T[x(y)/x].$$

Similarly, γ and χ correspond to transformations





$$T_1(T_2) \xrightarrow{\gamma} T_2(T_1) \quad \text{and} \quad (T_1T_2)T_3 \xrightarrow{\chi} (T_1T_3)T_2.$$

The remaining moves can be considered analogously, or by deriving them from the above. □

Proof of Theorem 4.5. We give a graphical proof by induction on topological orientations, exhibiting a sequence of moves $V'_n \Rightarrow T$ to realize any possible rooted essentially 3-valent map T with n non-root boundary arcs. (It is important to note that most maps can be constructed in multiple different ways, and the proof we give here only corresponds to one particular encoding.) For simplicity we ignore the presence of 2-valent vertices (i.e., work modulo edge subdivision), which otherwise only requires a bit of extra book-keeping.



To be explicit, the first (= rightmost) step of the derivation is a δ move, the second corresponds to a sequence of ρ moves, and the last to two parallel applications of the induction hypothesis.



Immediate since $T = V_0$.

□

Section 5

Proof of Prop. 5.3. By inspection of Figure 4, black vertices can only arise in the minimal polarization in configurations of the form




which correspond to β -redices of the unpolarized map. □

Proof of Prop. 5.5. Immediate by induction on topological orientations. (The converse is easily seen to be false: for instance consider the minimal polarizations of the non-topological maps in Example 3.16, which do not have cycles in their w-b orientations.) □


B An extended example

A bridgeless planar 3-valent map (the "Tutte graph")




(From W. T. Tutte, "On Hamiltonian Circuits", *Journal of the London Mathematical Society* 21 (1946), 98–101.)

A rooting of the above together with its topological orientation




The corresponding linear lambda term (with variables and unique β -redex indicated)




$\lambda a \lambda b \lambda c \lambda d \lambda e \lambda f \lambda g \lambda h \lambda i . a (\lambda j \lambda k . ((\lambda l \lambda m \lambda n . b (\lambda o . c (\lambda p . d (l (m ((n o) p))))))) (\lambda q \lambda r \lambda s . e (\lambda t . f (\lambda u . g (q (r ((s t) u))))))) (\lambda v \lambda w . h (\lambda x . i (j ((k v) (w x))))))$

The minimal polarization of the above



The corresponding universal polarized flow




The universal polarized flow is valued in the imploid freely generated over the intervals $a^- \leq a^+, \dots, y^- \leq y^+$ modulo the following relation (writing $[\sigma\tau]$ for $\sigma \rightarrow \tau$):

$$\beta : [[v^+u^-][[w^+v^-][[y^+[x^+w^-]]r^+]]] \leq [[[[n^+m^-][[o^+n^-][[q^+[p^+o^-]]j^+]]][[i^-][[h^+g^-]c^+]]b^-]]$$

The type assigned to the root is:

$$[[[[[f^+e^-][[i^+[g^+f^-]]b^+]]a^-][[[y^-s^+]]r^-][[[x^-t^+]]s^-][[u^+t^-][[q^-k^+]]j^-][[p^-l^+]]k^-][[m^+l^-][[[h^-d^+]]c^-][[e^+d^-]]a^+]]]]]]]]$$

A \mathbb{V} -flow realized as an instance of the universal polarized flow



$$\begin{aligned} a = d = g = m = o = r = u = w = y &= R \\ b = f = i = j = k = l = s = t = v &= G \\ c = e = h = n = p = q = x &= B \\ \beta : G &= G \end{aligned}$$