FPSAC 2010, San Francisco, USA DMTCS proc. **AN**, 2010, 191–202

Chamber Structure For Double Hurwitz Numbers

Renzo Cavalieri[1] and Paul Johnson[2†] and Hannah Markwig[3‡]

[1]*Colorado State University, Department of Mathematics, Weber Building, Fort Collins, CO 80523-1874, USA*
[2]*Department of Mathematics, Imperial College London, 180 Queen's Gate, London SW7 2AZ, UK*
[3]*CRC "Higher Order Structures in Mathematics", Georg August Universität Göttingen, Bunsenstr. 3-5, 37073 Göttingen, Germany*

Abstract. Double Hurwitz numbers count covers of the sphere by genus g curves with assigned ramification profiles over 0 and ∞, and simple ramification over a fixed branch divisor. Goulden, Jackson and Vakil (2005) have shown double Hurwitz numbers are piecewise polynomial in the orders of ramification, and Shadrin, Shapiro and Vainshtein (2008) have determined the chamber structure and wall crossing formulas for $g = 0$. We provide new proofs of these results, and extend them in several directions. Most importantly we prove wall crossing formulas for all genera.

The main tool is the authors' previous work expressing double Hurwitz number as a sum over labeled graphs. We identify the labels of the graphs with lattice points in the chambers of certain hyperplane arrangements, which give rise to piecewise polynomial functions. Our understanding of the wall crossing for these functions builds on the work of Varchenko (1987). This approach to wall crossing appears novel, and may be of broader interest.

This extended abstract is based on a new preprint by the authors.

Résumé. Les nombres de Hurwitz doubles dénombrent les revêtements de la sphère par une surface de genre g avec ramifications prescrites en 0 et ∞, et dont les autres valeurs critiques sont non dégénérées et fixées. Goulden, Jackson et Vakil (2005) ont prouvé que les nombres de Hurwitz doubles sont polynomiaux par morceaux en l'ordre des ramifications prescrites, et Shadrin, Shapiro et Vainshtein (2008) ont déterminé la structure des chambres et ont établis des formules pour traverser les murs en genre 0. Nous proposons des nouvelles preuves de ces résultats, et les généralisons dans plusieurs directions. En particulier, nous prouvons des formules pour traverser les murs en tout genre.

L'outil principal est le précédent travail des auteurs exprimant les nombres de Hurwitz doubles comme somme de graphes étiquetés. Nous identifions les étiquetages avec les points entiers à l'intérieur d'une chambre d'un arrangement d'hyperplans, qui sont connu pour donner une fonction polynomiale par morceauz. Notre étude des formules pour traverser les murs de cettes fonctions se base sur un travail antérieur de Varchenko (1987). Cette approche paraît nouvelle, et peut être d'un large intérêt.

Ce résumé élargi se base sur un papier nouveaux des auteurs.

Keywords: Hurwitz Numbers, Lattice Points, Hyperplane arrangements, Graphs

†NSF grants DMS-0602191 and DMS-0902754
‡German Research Foundation (Deutsche Forschungsgemeinschaft (DFG)) through the Institutional Strategy of the University of Göttingen

1 Introduction

Hurwitz theory studies holomorphic maps between Riemann surfaces with specified ramification. Double Hurwitz numbers count covers of $\mathbb{P}^1$ with assigned ramification profiles over 0 and ∞, and simple ramification over a fixed branch divisor.

A systematic study of double Hurwitz numbers in Goulden et al. (2005) shows double Hurwitz numbers are piecewise polynomial in the entries of the partitions defining the special ramification. In Shadrin et al. (2008), this result was investigated further in genus 0; the regions of polynomiality are determined, and a recursive wall crossing formula for how the polynomials change is obtained. This paper gives a unified approach to these results that strengthens them in several ways - the most important being the extension of the results of Shadrin et al. (2008) to positive genus.

This extended abstract is based on Cavalieri et al. (2009).

2 Statement of Results

The double Hurwitz number $H_g(\mathbf{x})$ (where $\mathbf{x} = (x_1, \dots, x_n)$) counts the number of maps $\pi : C \to \mathbb{P}^1$, where C is a connected, genus g curve and π has profiles $\mathbf{x_0} := \{x_i | x_i > 0\}$ (resp. $\mathbf{x}_\infty := \{x_i | x_i < 0\}$) over 0 (resp. ∞), and simple ramification over $r = 2g - 2 + n$ fixed other points. The preimages of 0 and ∞ are marked. Each cover is counted with weight $1/|\mathrm{Aut}(\pi)|$. Since r and g are related by the Riemann-Hurwitz formula, we sometimes use $H^r(\mathbf{x})$ to denote $H_g(\mathbf{x})$ when it makes formulas more attractive.

A ramified cover is essentially equivalent information to a monodromy representation; an equivalent definition of Hurwitz number counts the number of homomorphisms φ from the fundamental group Π_1 of $\mathbb{P}^1 \setminus \{0, \infty, p_1, \dots, p_r\}$ to the symmetric group S_d such that:

- the image of a loop around 0 has cycle type $\mathbf{x_0}$;
- the image of a loop around ∞ has cycle type $\mathbf{x}_\infty$;
- the image of a loop around p_i is a transposition;
- the subgroup $\varphi(\Pi_1)$ acts transitively on the set $\{1, \dots, d\}$.

This number is divided by $|S_d|$, to account both for automorphisms and for different monodromy representations corresponding to the same cover. One organizes this count in terms of graphs as in (Cavalieri et al., Lemma 4.1), a fact which is the starting point of our investigation (see Section 3).

Let $\mathcal{H}$ be the hyperplane $\mathcal{H} = \{\sum_i x_i = 0\} \subset \mathbb{R}^n$. We think of H_g (equiv. H^r) as a map

$$H_g : \mathcal{H} \cap \mathbb{Z}^n \to \mathbb{Q} : \mathbf{x} \mapsto H_g(\mathbf{x}).$$

Our first result is a new proof of the following theorem in Goulden et al. (2005):

Theorem 2.1 (GJV) *The function $H_g(\mathbf{x})$ is a piecewise polynomial function of degree $4g - 3 + n$.*

Our techniques allow us to extend this result and answer a question implicit in the work of Goulden, Jackson and Vakil:

Theorem 2.2 *$H_g(\mathbf{x})$ is either even or odd, depending on the parity of the leading degree $4g - 3 + n$.*

We then extend the results of Shadrin et al. (2008) to all genera. We determine the regions on which $H_g(\mathbf{x})$ is polynomial:

Theorem 2.3 *The chambers of polynomiality of $H_g(\mathbf{x})$ are bounded by **walls** corresponding to the **resonance** hyperplanes W_I, given by the equation $W_I = \left\{\mathbf{x}_I = \sum_{i\in I} x_i = 0\right\}$, for any $I \subset \{1,\dots,n\}$.*

We then describe wall crossing formulas for general genus. Denote the chambers of the resonance arrangement as H-chambers;

Definition 2.4 *Let C_1 and C_2 be two H-chambers adjacent along the wall W_I, with C_1 being the chamber with $x_I < 0$. The Hurwitz number $H^r(\mathbf{x})$ is given by polynomials, say $P_1(\mathbf{x})$ and $P_2(\mathbf{x})$, on these two regions. By a wall crossing formula, we mean a formula for the polynomial*

$$WC_I^r(\mathbf{x}) = P_2(\mathbf{x}) - P_1(\mathbf{x}).$$

With the notation $WC_I^r(\mathbf{x})$ there is no ambiguity about which direction we cross the wall. Since $\mathbf{x}$ lies on the hyperplane $\sum_{i=1}^n x_i = 0$, each wall has two possible labels: W_I and W_{I^c}. We choose the name so that $\mathbf{x}_I$ is increasing.

We use $H^{r\bullet}(\mathbf{x})$ to denote Hurwitz numbers with potentially disconnected covers. Our main theorem is:

Theorem 2.5 (Wall crossing formula)

$$WC_I^r(\mathbf{x}) = \sum_{s+t+u=r} \sum_{|\mathbf{y}|=|\mathbf{z}|=|\mathbf{x}_I|} (-1)^t \binom{r}{s,t,u} \frac{\prod \mathbf{y}_i}{\ell(\mathbf{y})!} \frac{\prod \mathbf{z}_j}{\ell(\mathbf{z})!} H^s(\mathbf{x}_I, \mathbf{y}) H^{t\bullet}(-\mathbf{y}, \mathbf{z}) H^u(\mathbf{x}_{I^c}, -\mathbf{z}) \qquad (1)$$

Here $\mathbf{y}$ is an ordered tuple of $\ell(\mathbf{y})$ positive integers with sum $|\mathbf{y}|$, and similarly with $\mathbf{z}$.

The walls W_I correspond to values of $\mathbf{x}$ where the cover could potentially be disconnected, or where $x_i = 0$. Crossing this second type of wall corresponds to moving a ramification between 0 and ∞. In the traditional view of double Hurwitz numbers, the number of ramification points over 0 and ∞ were fixed separately, rather than just the total number of ramification points. Theorem 2.5 suggests that it is natural to treat them as part of the same problem: the wall crossing formula for $x_i = 0$ is identical to the other wall crossing formulas.

3 Overview of Methods

This paper is an exploration of the consequences of a formula in the author's previous work, Cavalieri et al., which expresses double Hurwitz numbers $H_g(\mathbf{x})$ as a sum over certain directed trivalent graphs Γ with several labelings, which we call monodromy graphs:

Definition 3.1 *For fixed g and $\mathbf{x} = (x_1,\dots,x_n)$, a graph Γ is a* monodromy graph *if:*

- *Γ is a connected, genus g, directed graph.*
- *Γ has n 1-valent vertices called* leaves*; the edges leading to them are* ends*. All ends are directed inward, and are labeled by the weights $x_1,\dots,x_n$. If $x_i > 0$, we say it is an* in-end*, otherwise it is an* out-end.

- *All other vertices of* Γ *are* 3*-valent, and are called* internal vertices. *Edges that are not ends are called* internal edges.
- *After reversing the orientation of the out-ends,* Γ *does not have directed loops, sinks or sources.* [(i)].
- *The internal vertices are ordered compatibly with the partial ordering induced by the directions of the edges.*
- *Every internal edge* e *of the graph is equipped with a* weight $w(e) \in \mathbb{N}$. *The weights satisfy the* balancing condition *at each internal vertex: the sum of all weights of incoming edges equals the sum of the weights of all outgoing edges.*

It follows from (Cavalieri et al., Lemma 4.1) that the Hurwitz number is computed as:

$$H_g(\mathbf{x}) = \sum_{\Gamma} \frac{1}{|\mathrm{Aut}(\Gamma)|} \prod_e w(e), \tag{2}$$

the sum over all monodromy graphs Γ for g and $\mathbf{x}$, and the product over the interior edges of Γ.

In genus zero, the edge labelings $w(e)$ are determined uniquely by $\mathbf{x}$. This makes the genus zero case much easier to treat, and the results for this case were already presented in Cavalieri et al.. In higher genus, if we fix a directed graph and the labels $\mathbf{x}$ for the ends (such data will be called a *directed* $\mathbf{x}$*-graph*), there are many ways to assign edge labels $w(e)$ that satisfy the balancing condition.

The crux of this paper is to understand the space of edge labelings (which we call *flows*) for each directed $\mathbf{x}$-graph. The space of flows consists of the lattice points in a certain bounded polytope which we call an F-chamber. The contribution $s_{\Gamma(d)}(\mathbf{x})$ of a fixed directed $\mathbf{x}$-graph $\Gamma(d)$ to $H_g(\mathbf{x})$ equals

$$s_{\Gamma(d)}(\mathbf{x}) = \frac{1}{|\mathrm{Aut}(\Gamma(d))|} \cdot m(\mathcal{C}) \cdot \sum_{b \in \mathcal{C} \cap \Lambda} \prod_e L_e(\mathbf{x}, b) \tag{3}$$

where $\mathcal{C}$ is the F-chamber associated to $\Gamma(d)$, Λ denotes the lattice and $m(\mathcal{C})$ equals the number of ways to order the vertices of $\Gamma(d)$ as required for a monodromy graph. Here we have written $L_e(\mathbf{x}, b)$ for $w(e)$, as the weight of each edge will be a linear function in $\mathbf{x}$ and the coordinates of Λ.

We illustrate this in an example that we continue to develop throughout. Consider the directed $\mathbf{x}$-graph $\Gamma(\mathbf{x}, d, v)$ on the left hand side in Figure 1. In this example, we use the notation $\Gamma(\mathbf{x}, d, v)$ to indicate that the graph comes with directed edges (d) and with a vertex ordering (v). In the figure, the vertices are labelled to indicate the vertex ordering. We want to understand all monodromy graphs that equal $\Gamma(\mathbf{x}, d, v)$ after forgetting the weights of the internal edges. There are no monodromy graphs that equal $\Gamma(\mathbf{x}, d, v)$ after forgetting the weights if $x_1 + x_3 \leq 0$, so we assume that $x_1 + x_3 > 0$.

We have two degrees of freedom to choose weights for the interior edges such that the balancing condition is satisfied, one for each independent cycle of Γ. Once we label one of the interior edges with the weight i, and another with j, all other weights are determined by the balancing condition, as shown in the right hand side of Figure 1. All possible collections of edge labels are indexed by the lattice points in the polytope defined requiring these labels to be nonnegative:

$$i \geq 0, j \geq 0, j + i - x_2 \geq 0, -x_4 - i - j \geq 0, -x_4 - j \geq 0, j - x_2 \geq 0.$$

(i) We do not consider leaves to be sinks or sources.

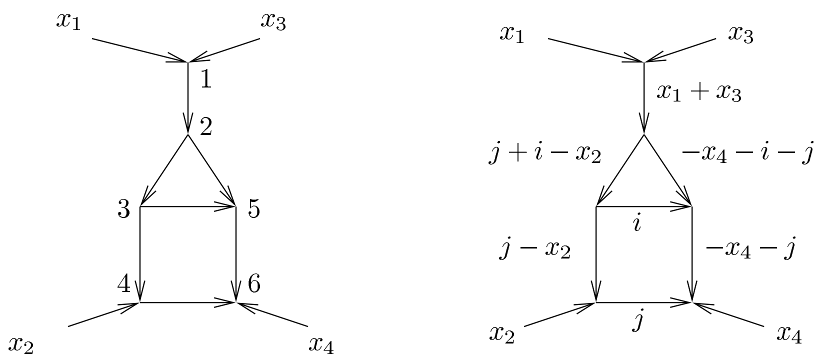

Fig. 1: A directed $\mathbf{x}$-graph and the weights of internal edges determined by the balancing condition

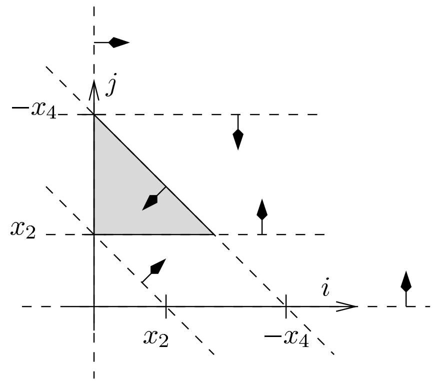

Fig. 2: The F-chamber corresponding to $\Gamma(\mathbf{x}, d, v)$

Figure 2 shows all hyperplanes $w(e) = 0$ with a normal vector indicating on which side of the hyperplane the inequality $w(e) > 0$ is satisfied; this defines the F-chamber corresponding to $\Gamma(\mathbf{x}, d, v)$.

The contribution of $\Gamma(\mathbf{x}, d, v)$ to $H_g(\mathbf{x})$ is given by

$$(x_1 + x_3) \cdot \sum_{i=0}^{-x_4-x_2} \sum_{j=x_2}^{-i-x_4} i \cdot j \cdot (j + i - x_2) \cdot (-x_4 - i - j) \cdot (-x_4 - j) \cdot (j - x_2)$$

where the sum goes over all lattice points (i, j) in the polygon above ($\Gamma(\mathbf{x}, d, v)$ has no automorphisms).

Theorem 2.1 follows from Equation 3 and the general theory of lattice points in polytopes. As we change $\mathbf{x}$, the facets of F-chamber $\mathcal{C}$ translate (their normal directions remain constant). Since for all integral $\mathbf{x}$, the vertices of the F-chamber $\mathcal{C}$ are integers, our sums are piecewise polynomial, and the walls occur when the topology of $\mathcal{C}$ changes.

In the general setup of the theory the resulting polynomials need not be odd or even, so Theorem 2.2 is more subtle: it is related to Ehrhart reciprocity, and depends essentially on the fact that the polynomial

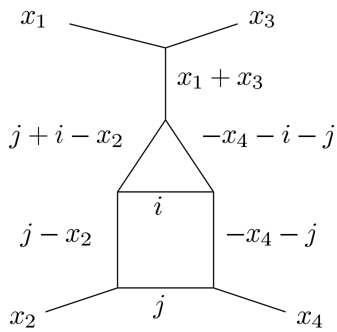

Fig. 3: Labels $w'(e)$ for the undirected graph

we are summing over the polytope vanishes on the boundary of the polytope.

To prove Theorems 2.3 and 2.5 requires understanding for what values of $\mathbf{x}$ the F-chambers change topology, and how they change topology, respectively. To answer these questions, it is helpful to notice that the F-chambers for distinct $\mathbf{x}$-graphs with the same underlying undirected graph Γ fit together as the set $\mathcal{BC}_\Gamma(x)$ of bounded chambers of a natural hyperplane arrangement $\mathcal{A}_\Gamma(\mathbf{x})$ associated to Γ and $\mathbf{x}$.

Returning to our example, we can retain the orientation of the edges in Figure 1 as a reference orientation, and the labels $w'(e)$ for the internal edges obtained from the balancing condition as in Figure 3.

We switch to $w'(e)$ instead of $w(e)$ because we do no longer restrict the edge labels to be positive; instead, any possible value of i and j are allowed. For each edge of Γ, the set of i and j where $w'(e) = 0$ will give a hyperplane, and together these form the hyperplane arrangement $\mathcal{A}_\Gamma(\mathbf{x})$. Inside a chamber of $\mathcal{A}_\Gamma(\mathbf{x})$, a sign for $w'(e)$ is picked for every edge, and thus an orientation for every edge; the chamber will be the F-chamber for that directed $\mathbf{x}$-graph. Figure 4 shows the hyperplane arrangement $\mathcal{A}_\Gamma(\mathbf{x})$, with each F-chamber labeled by the corresponding directed graphs with the induced orientations. Since the orientation of the ends and the edge with label $x_1 + x_3$ does not depend on i and j, we do not include these edges in the pictures.

Only the bounded F-chambers (shaded) correspond to directed $\mathbf{x}$-graphs that contribute to the Hurwitz number. The unbounded F-chambers correspond to graphs with a directed loop, and so the vertices have no compatible total orderings and the multiplicity of these chambers are zero.

For different chambers, the product $\prod w(e)$ differs at most by the sign, since the edge weights $w(e)$ equal plus or minus the edge label $w'(e)$, depending on the side of the hyperplane $w'(e) = 0$ the F-chamber is situated. Thus we can define a sign $\operatorname{sign}(\mathcal{C})$ for each F-chamber $\mathcal{C}$ that is determined by the number of edges that are reversed when compared to the reference orientation.

Summing all the contributions from directed $\mathbf{x}$-graphs $\Gamma(d)$ with the same underlying undirected $\mathbf{x}$-graph Γ, we get the contribution S_Γ of the undirected $\mathbf{x}$-graph Γ to the Hurwitz number as

$$S_\Gamma(\mathbf{x}) = \frac{1}{|\mathrm{Aut}(\Gamma)|} \sum_{\mathcal{C}} \operatorname{sign}(\mathcal{C}) m(\mathcal{C}) \sum_{b \in \mathcal{C} \cap \Lambda} \prod L_e(\mathbf{x}, b),$$

where the sum goes over all bounded chambers $\mathcal{C}$ of $\mathcal{A}_\Gamma(\mathbf{x})$.

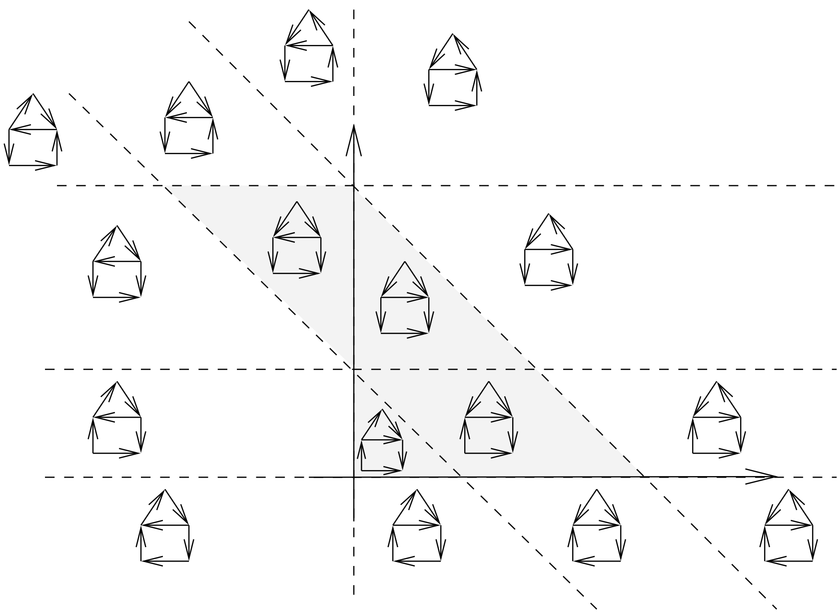

Fig. 4: The parameter space for monodromy graphs corresponding to a given $\mathbf{x}$-graph

Remark 3.2 *In Equation (2)* Γ *is a monodromy graph, while here* Γ *is an* $\mathbf{x}$*-graph, and so the meaning of Aut*(Γ) *is different. An automorphism of a monodromy graph must fix all vertices, while an automorphism of an* $\mathbf{x}$*-graph only needs to fix the ends. These extra automorphisms account for the fact that the same monodromy graph can occur in multiple ways from a single* $\mathbf{x}$*-graph.*

Even for a generic choice of $\mathbf{x}$ the arrangement $\mathcal{A}_\Gamma(\mathbf{x})$ is not simple - that is, there are hyperplanes that do not intersect transversally. This follows from the balancing condition: if two edge labels incident to a vertex are both zero, then the third edge label must be as well. As a consequence, for each vertex we have three hyperplanes intersecting in codimension two. But for generic $\mathbf{x}$, these are the only nontransverse intersections. When we pass through a value of $\mathbf{x}$ with more nontransverse intersections than expected, the topology of the arrangement $\mathcal{A}_\Gamma(\mathbf{x})$ changes, and so do the Hurwitz polynomials. We prove Theorem 2.3 by showing that if $e_1, \ldots, e_k$ are k edges whose hyperplanes intersect in codimension $k-1$ at $\mathbf{x}$, but generically intersect transversally, then these edges disconnect Γ, and each component will contain at least one end. Flows in the intersection of the hyperplanes correspond to flows on the graph where the edges are cut, and so if I is the set of ends on one component, we see that $\mathbf{x}$ must have been a point on the wall W_I.

Our main result is the wall crossing formula (Theorem 2.5). The idea of the proof is simple: matching the contributions to both sides by every directed $\mathbf{x}$-graph. In genus 0, realizing this strategy is straightforward because there is a natural geometric bijection (Cut) between graphs contributing to the wall crossing (LHS) and pairs of graphs contributing to the product of Hurwitz numbers on the RHS (the middle term can easily be seen to equal 1 in genus 0). In higher genus Cut is no longer a function, and a delicate process of inclusion/exclusion is required, leading us to foray into algebraic combinatorics. While to determine the walls it is enough to know where topology of $\mathcal{A}_\Gamma(\mathbf{x})$ changes, to derive the wall crossing formula we must understand how the topology changes; i.e. if $\mathbf{x}_1$ and $\mathbf{x}_2$ lie in two H-chambers C_1 and C_2, adjacent along the wall W_I, how does $\mathcal{A}_\Gamma(\mathbf{x}_1)$ differ from $\mathcal{A}_\Gamma(\mathbf{x}_2)$? This understanding is essential for relating the difference of the contributions $S_\Gamma(\mathbf{x}_1) - S_\Gamma(\mathbf{x}_2)$ to the resp. Hurwitz numbers.

The information how $\mathcal{A}_\Gamma(\mathbf{x}_1)$ differs from $\mathcal{A}_\Gamma(\mathbf{x}_2)$ is conveniently encoded in a linear map

$$\nabla_{12} : \mathbb{R}[\mathcal{BC}_\Gamma(\mathbf{x}_1)] \to \mathbb{R}[\mathcal{BC}_\Gamma(\mathbf{x}_2)]$$

called the Gauss-Manin connection. The basic picture is that as $\mathbf{x}$ passes through a wall, certain F-chambers vanish, and others appear. For any F-chamber, the change in shape as $\mathbf{x}$ crosses a wall can be described in terms of adding or subtracting these appearing F-chambers; ∇_{12} is the map that sends a given F-chamber to this signed sum of F-chambers. It turns out to be easier to declare F-chambers to form an orthonormal basis of $\mathbb{R}[\mathcal{BC}_\Gamma(\mathbf{x}_1)]$, and study the adjoint ∇_{12}^* which records which F-chambers of $\mathcal{A}_\Gamma(\mathbf{x}_1)$ map to a given one in $\mathcal{A}_\Gamma(\mathbf{x}_2)$.

The key point is that integrating a polynomial f over an F-chamber $\mathcal{C}(\mathbf{x})$ gives only a piecewise polynomial function; but if we replace $\mathcal{C}(\mathbf{x})$ by $\nabla_{12}\mathcal{C}(\mathbf{x})$ when we cross a wall, then we get a globally defined polynomial. Results of Varchenko (1987) show that if we replace integration by summing over lattice points, the same result is true if we deal properly with lattice points in the boundary of the polytope. Since our polynomials vanish there, we don't have to worry about this, and so ∇_{12} encodes essentially all the information for Hurwitz wall crossing.

Returning once more to our running example, in Figure 4 showing $\mathcal{A}_\Gamma(\mathbf{x})$, we implicitly assumed that $0 > x_2 + x_4$. The topology of the hyperplane arrangement changes if $0 = x_2 + x_4$. Fix the wall $W_{\{2,4\}}$ and let C_1 and C_2 be two adjacent H-chambers. Assume that in C_1, we have $0 < x_2 + x_4$, and in C_2, we have $x_2 + x_4 < 0$. Figure 5 shows the hyperplane arrangements $\mathcal{A}_\Gamma(\mathbf{x}_1)$ and $\mathcal{A}_\Gamma(\mathbf{x}_2)$ for two points $\mathbf{x}_1 \in C_1$ and $\mathbf{x}_2 \in C_2$. The hyperplanes appear with their defining equations. They are drawn with different line styles in order to emphasize how they move. The bounded F-chambers are labelled with letters. Since the edge with weight $x_1 + x_3$ gives the inequality $x_1 + x_3 > 0$ on the right which is not satisfied on the left, every F-chamber on the right is an appearing chamber, and every F-chamber on the left is vanishing. This can also be seen from the corresponding graphs: since the top most interior edge with weight $x_1 + x_3$ always points down on the right, there is a flow from top to bottom. Figure 6 shows the directed $\mathbf{x}$-graphs corresponding to some of the F-chambers.

To understand the Gauss-Manin connection for this example, we pick an appearing F-chamber on the right, e.g. A, and ask ourselves what F-chambers on the left contain it in their support when carried over the wall, i.e. we determine $\nabla^*_{\Gamma,12}(A)$. To do this, we take chambers on the left, e.g. E, and carry them over, i.e. we first determine $\nabla_{\Gamma,12}(E)$. When we carry E over, we get B and keep the orientation (we switch the summation index twice). In the same way, we get $\nabla_{\Gamma,12}(F) = A$. If we interpret the inequalities of G on the right, we have to switch one summation index, and then we end up with $A + B + C$. Thus $\nabla_{\Gamma,12}(G) = -A - B - C$. Finally, H becomes $D + B + C$. Thus, $\nabla^*_{\Gamma,12}(A) = F - G$, $\nabla^*_{\Gamma,12}(B) = E - G + H$, $\nabla^*_{\Gamma,12}(C) = -G + H$ and $\nabla^*_{\Gamma,12}(D) = H$.

Our key result is that we can express the linear map ∇^*_{12} combinatorially. We define a simple I-cut of a directed $\mathbf{x}$-graph Γ to be a minimal set of edges E so that E disconnects Γ into exactly two components, one containing the ends in I, and the other containing all the ends in I^c. An I-cut in general is a union of simple I-cuts. This could be the empty union, which we call the empty cut. Note that a general I-cut might be expressible as a union of simple I-cuts in many different ways.

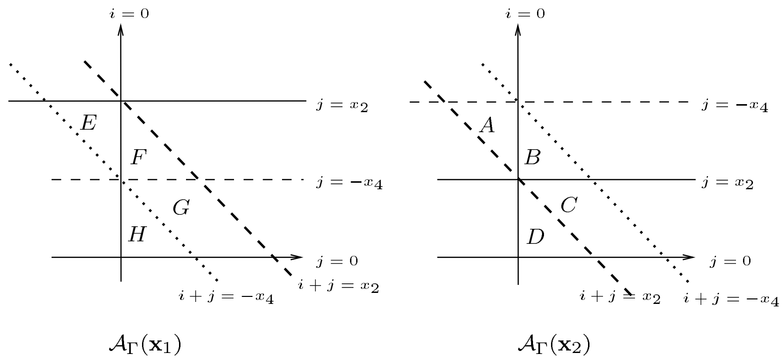

Fig. 5: The hyperplane arrangements $\mathcal{A}_\Gamma(\mathbf{x}_1)$ and $\mathcal{A}_\Gamma(\mathbf{x}_2)$ for two points $\mathbf{x}_1$ and $\mathbf{x}_2$ on opposite sides of a wall.

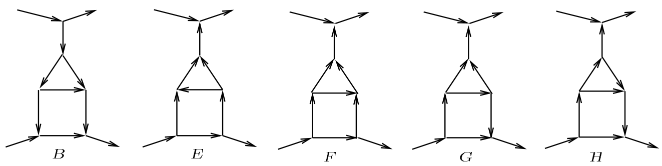

Fig. 6: The directed $\mathbf{x}$-graphs corresponding to the F-chambers B, E, F, G and H of figure 5.

The set of simple I-cuts forms a poset under inclusion, which we denote $C_\Gamma(I)$. As an example, we show the poset of $\{1,3\}$-cuts of the graph Γ_B from above as an example (see Figure 7).

In this case, the $C_\Gamma(I)$ is simply the boolean lattice generated by the simple cuts, although this is not true in general. A key lemma is in our paper is the following:

Lemma 3.3 *$C_\Gamma(I)$ is isomorphic to the face lattice of a certain cone, defined in terms of a different hyperplane arrangement associated to Γ.*

The main importance of Lemma 3.3 is that it shows $C_\Gamma(I)$ is Eulerian, and all Möbius inversion type questions can be translated into questions about Euler characteristics of subsets of the cone. A more immediate consequence is that $C_\Gamma(I)$ is ranked; the empty cut has rank zero.

The key step in our proof of the Main Theorem 2.5 is the following theorem, which expresses ∇_{12}^* in terms of the poset of cuts $C_\Gamma(I)$.

Theorem 3.4 *Let A and B be F-chambers in $\mathcal{BC}_\Gamma(\mathbf{x}_2)$ and $\mathcal{BC}_\Gamma(\mathbf{x}_1)$, respectively. Let Γ_A and Γ_B denote the corresponding orientations of the edges of Γ, and let S be the subset of edges of Γ where these orientations differ. Then*

$$\langle \nabla_{12}^* A, B\rangle = (-1)^{|S|} \sum_{S \subset C \in C_{\Gamma_A}(I)} (-1)^{rk(C)}.$$

Here the notation means we sum only over the I-cuts of Γ_A that contain S.

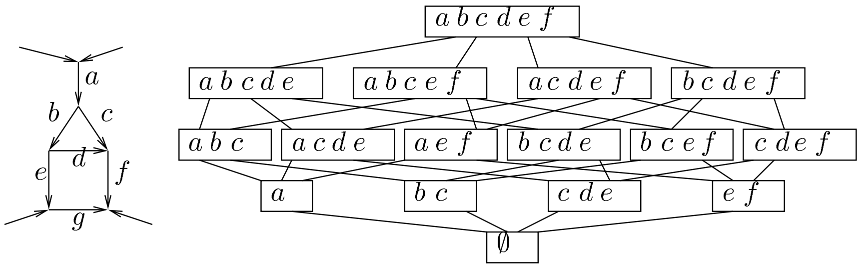

Fig. 7: The directed $\mathbf{x}$-graph Γ_B and its poset of $\{1,3\}$-cuts

Theorem 3.4 does not depend on the graph being trivalent, and could be of independent interest.

The proof of Theorem 3.4 is the technical heart of the paper, and rests upon the following observations from Varchenko (1987): cones are preserved by the Gauss-Manin connection, and every chamber can be written as a signed sum of cones. Thus, it suffices to show ∇^*_{12} preserves cones. We are able to do this by using the understanding of $C_\Gamma(I)$ afforded by Lemma 3.3.

We now illustrate the statement of Theorem 3.4 in the case of our example. Consider the appearing chamber B on the right of the wall. We have seen that $\nabla^*_{\Gamma,12}(B) = E - G + H$, and so we understand the left hand side of Theorem 3.4. We have also determined the poset of I-cuts of the directed graph Γ_B (Figure 7), and so we are able to compute the right hand side as well.

First, let us verify that Theorem 3.4 gives $\langle \nabla^*_{12} B, E\rangle = 1$. We see that to get Γ_E from Γ_B, we must change the orientation of the edges a, b, c, d, e and f, and so $S = \{a,b,c,d,e,f\}$. . There is only one cut that contains S, the maximal cut. Its rank is four (see Figure 7). Since $|S| = 6$, Theorem 3.4 gives $\langle \nabla^*_{12} B, E\rangle = (-1)^6 \cdot (-1)^4 = 1$.

Similarly, we will verify that Theorem 3.4 gives $\langle \nabla^*_{12} B, G\rangle = -1$. In this case, the set S of edges where the orientations of Γ_G and Γ_B differ is $\{a,b,c,e\}$. There are three cuts that cut these edges, namely $abcde$ and $abcef$, both of rank three, and $abcdef$ of rank four. Since $|S| = 4$, Theorem 3.4 gives $\langle \nabla^*_{12} B, G\rangle = (-1)^4 \cdot ((-1)^3 + (-1)^3 + (-1)^4) = -1$.

Additionally, chambers that do not appear in $\nabla^*_{12} B$ should appear with coefficient zero in the right hand side of Theorem 3.4. Let us check that we get $\langle \nabla^*_{12} B, F\rangle = 0$. To get Γ_F from Γ_B, the set S of edges we must reverse is $\{a,b,c,e,f\}$. The cuts that contain S are $abcef$ of rank three and $abcdef$ of rank four, and so we get $(-1)^5 \cdot ((-1)^3 + (-1)^4) = 0$.

A more complicated wall crossing formula than Theorem 2.5 follows rather quickly from Theorem 3.4.

For a cut $C \in C_\Gamma(I)$, removing the edges in C from Γ will cut Γ into multiple components graphs, each of which can can be interpreted as a graph appearing for a simpler Hurwitz number.

As we sum over all Γ, we will sometimes see essentially the same cut C appearing for different Γ - that is, the components of $\Gamma \setminus C$ will be different graphs, but will contribute to the same Hurwitz problem (have the same number of vertices and in and out going ends), and glue together in the same manner. This is the situation illustrated in Figure 8. As a result, we obtain:

Theorem 3.5 (Heavy Formula)

$$WC^r_I(\mathbf{x}) = \sum_{N=0}^{\infty} \sum_{s+(\sum_{j=1}^N t_j)+u=r} \sum_{|\lambda|=|\eta|=d} \sum_{data\ in\ \star} (-1)^N \binom{r}{s, t_1, \ldots, t_N, u} \frac{\prod (\mu^{(i,j)})_k}{\prod \ell(\mu^{(i,j)} j)!}$$

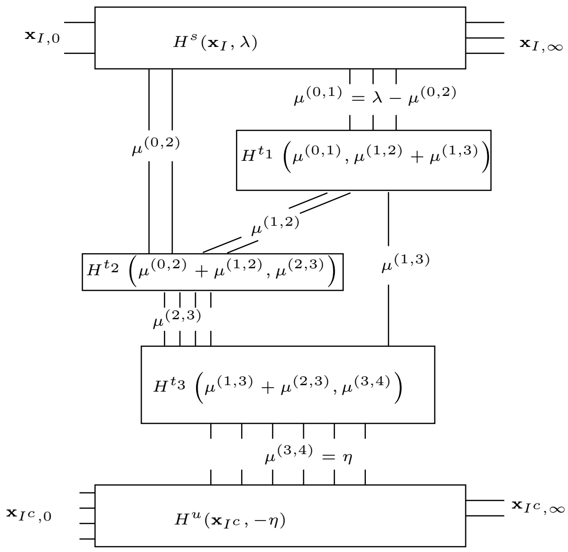

Fig. 8: The data denoted by $\star$ in the heavy formula, Theorem 3.5

$$H^s(\mathbf{x}_I, \lambda)\left(\prod_{j=1}^{N} H^{t_j}(\star)\right) H^u(\mathbf{x}_{I^c}, -\eta)$$

The data denoted by $\star$ is illustrated in Figure 8: it consists in disconnecting a graph with an I-cut in all possible ways with the right numerical invariants. The μ_i^j denote the partitions of weights of the edges connecting the i-th to the j-th connected component, we use $(\mu_i^j)_k$ to denote its parts.

The derivation of Theorem 3.5 from Theorem 2.5 is essentially inclusion-exclusion, and an application of Lemma 3.3.

4 Motivation and Connections to other work

Although our methods are essentially combinatorial, much of the motivation of Goulden et al. (2005), and hence our work, comes from algebraic geometry, in particular the ELSV formula Ekedahl et al. (2001). There, it is shown that similar polynomiality occurs for single Hurwitz numbers (where there is no ramification over ∞), and that the coefficients of these polynomials are the intersection of certain classes in the moduli space of curves $\overline{\mathcal{M}}_{g,n}$. This connection has been vital in understanding these intersections. In Goulden et al. (2005), it is suggested that a similar relationship should hold for one part double Hurwitz, where the map is totally ramified over zero - i.e., where $\mathbf{x}$ has only one positive part. They conjecture that the moduli space of curves should be replaced by some yet to be determined universal Picard space, which would give us a similar understanding of the intersection theory there. One part double Hurwitz numbers are simply one chamber of the Hurwitz problem, however, and it would be wonderful to extend the conjecture of Goulden et al. (2005) to give a formula for double Hurwitz numbers on all chambers, with the wall crossing phenomenon explained in terms of changes in the moduli space. Our work could perhaps be of use in investigating such a conjecture.

Acknowledgements

We would like to thank Federico Ardila and Michael Shapiro for helpful discussions.

References

R. Cavalieri, P. Johnson, and H. Markwig. Tropical Hurwitz numbers. JACO (to appear), arXiv:0804.0579.

R. Cavalieri, P. Johnson, and H. Markwig. Wall crossings for double Hurwitz numbers. Preprint, arXiv:1003.1805, 2009.

T. Ekedahl, S. Lando, M. Shapiro, and A. Vainshtein. Hurwitz numbers and intersections on moduli spaces of curves. *Invent. Math.*, 146(2):297–327, 2001.

I. Goulden, D. Jackson, and R. Vakil. Towards the geometry of double Hurwitz numbers. *Adv. Math.*, 198 (1):43–92, 2005.

S. Shadrin, M. Shapiro, and A. Vainshtein. Chamber behavior of double Hurwitz numbers in genus 0. *Adv. Math.*, 217(1):79–96, 2008.

A. Varchenko. Combinatorics and topology of the arrangement of affine hyperplanes in the real space. *Functional Anal. Appl.*, 21(1):9–19, 1987.