

Presentation Complexity of Irregular Connections and Stokes Data

Luca Blanchi

June 2026

Abstract. This note applies Presentation Theory to local meromorphic connections with irregular singularities. The classical analytic input is the Hukuhara-Levelt-Turrittin formal decomposition and the Stokes classification. The presentation-theoretic output is a family of reusable lower-bound principles: formal types, Stokes directions, Stokes factors, and turning data become extracted observables whose encoding cost is forced inside any complete presentation system for the corresponding class of connections. Ordinary monodromy is shown to have large residual fibres already for rank-one exponential connections, while local Stokes normal forms give a controlled presentation format. A final section records the corresponding controlled-transfer statement for irregular Riemann-Hilbert presentations on bounded good classes.

1 Introduction

Irregular singularities contain information that is invisible to ordinary monodromy. A meromorphic connection on a punctured disc is not classified by its local system on the punctured disc: formal exponential factors and Stokes matrices are additional data, and they are essential. The classical theory explains precisely which data are needed. This paper asks a different question: once a class of irregular connections is presented by finite descriptions, which parts of the Stokes data must be paid for by any complete presentation?

The point is not to reprove the Stokes classification. The classification is used as input. The contribution is to place it in the language of Presentation Theory: formal type, Stokes directions, Stokes factors, and turning data are treated as observables extracted from a presentation. Whenever such an observable has intrinsic encoding cost, any complete presentation system for the connections must have at least that much presentation complexity, up to the declared overhead of extraction.

The resulting statements have a uniform shape. Let a presentation system describe a class of objects. If a mathematically necessary invariant can be extracted from every presentation with overhead $\backslash(F)$, and if that invariant has cost $\backslash(N)$, then the original presentation has cost at least the inverse $\backslash(F)$ -scale of $\backslash(N)$. In the irregular-connection setting, the invariants include:

- the reduced formal exponential type;
- the cyclic Stokes-direction configuration determined by differences of exponential factors;
- the Stokes factors in the corresponding unipotent Stokes groups;
- global Stokes hypersurfaces and turning data, in classes where good formal structures are available;
- enhanced-sheaf or Stokes-filtered presentations related by irregular Riemann-Hilbert correspondences.

The note is deliberately local first. The local statements are the cleanest and are the natural testing ground for presentation-theoretic lower bounds. The global and Riemann-Hilbert statements at the end are included as controlled transfer principles on bounded good classes, where the required analytic operations are part of the declared package.

2 Presentation Systems and Extracted Observables

Definition 2.1 (Presentation system). Let $\backslash(\mathcal{C})$ be a class of mathematical objects. A presentation system for $\backslash(\mathcal{C})$ is a triple $\backslash[(\mathcal{P}, \rho, \mathbf{C}_{\backslash(\mathcal{P})})]$ where $\backslash(\mathcal{P})$ is a set of finite descriptions, $\backslash(\rho: \mathcal{P} \rightarrow \mathcal{C})$ is a realization map, and $\backslash(\mathbf{C}_{\backslash(\mathcal{P})}: \mathcal{P} \rightarrow \mathbb{R}_{\geq 0})$ is a cost function. The induced presentation cost of $\backslash(X \in \mathcal{C})$ is $\backslash[\mathbf{C}_{\backslash(\mathcal{C})}(X) = \inf\{\mathbf{C}_{\backslash(\mathcal{P})}(p): p \in \mathcal{P}, \rho(p) = X\}]$. The system is complete on a subclass $\backslash(\mathcal{C}_0 \subseteq \mathcal{C})$ if every object of $\backslash(\mathcal{C}_0)$ is isomorphic to

$(\rho(p))$ for some $(p \in \mathcal{P})$.

Definition 2.2 (Extracted observable with overhead). Let $(A: \mathcal{C}_0 \rightarrow \mathcal{A})$ be an isomorphism-invariant observable, and let $(\mathbf{C}_{\mathcal{A}})$ be a cost on the target. We say that (A) is extracted from $(P, \rho, \mathbf{C}_{\mathcal{P}})$ with overhead (F) if $(F: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0})$ is nondecreasing and every presentation $(p \in \mathcal{P})$ with $(\rho(p) \in \mathcal{C}_0)$ determines a presentation of $(A(\rho(p)))$ of cost at most $[F(\mathbf{C}_{\mathcal{P}}(p))]$.

Theorem 2.3 (Extracted-observable lower bound). Let $(A: \mathcal{C}_0 \rightarrow \mathcal{A})$ be extracted with nondecreasing overhead (F) . Then every $(X \in \mathcal{C}_0)$ satisfies $[\mathbf{C}_{\mathcal{A}}(A(X))] \leq F(\mathbf{C}_{\mathcal{C}}(X))$. Equivalently, $[\mathbf{C}_{\mathcal{C}}(X)] \geq \inf\{t \geq 0 : F(t) \geq \mathbf{C}_{\mathcal{A}}(A(X))\}$.

Proof. Choose a presentation (p) of (X) . By the extraction hypothesis, (p) determines a presentation of $(A(X))$ of cost at most $(F(\mathbf{C}_{\mathcal{P}}(p)))$. Therefore $[\mathbf{C}_{\mathcal{A}}(A(X))] \leq F(\mathbf{C}_{\mathcal{P}}(p))$. Taking the infimum over all presentations (p) realizing (X) , and using monotonicity of (F) , gives the first inequality. The second formulation is the same inequality written in inverse-scale form.

Definition 2.4 (Residual fibre). Let $(B: \mathcal{C}_0 \rightarrow \mathcal{B})$ be another observable. The residual fibre of (B) over $(b \in \mathcal{B})$ is $B^{-1}(b) = \{X \in \mathcal{C}_0 : B(X) = b\}$. If (A) is nonconstant on $(B^{-1}(b))$, then (B) does not separate the objects in that fibre.

Proposition 2.5 (Fibre lower bound). Assume that (B) is part of the presentation format, and that a complete presentation of an object in a fibre $(B^{-1}(b))$ must also determine an observable (A) . If (A) is extracted with overhead (F) , then for every $(X \in B^{-1}(b))$ $[\mathbf{C}_{\mathcal{C}}(X)] \geq \inf\{t \geq 0 : F(t) \geq \mathbf{C}_{\mathcal{A}}(A(X))\}$. In particular, a large (A) -profile inside a single (B) -fibre forces presentation cost that is not visible from (B) alone.

Proof. This is the extracted-observable lower bound applied after restricting the presentation system to the fibre $(B^{-1}(b))$. The observable (B) is constant there, so any remaining lower bound comes from the additional observable (A) .

3 Local Irregular Connections

Let $(K = \mathbb{C}((z)))$. A local meromorphic connection may be represented algebraically as a finite-dimensional differential (K) -module, or analytically as a meromorphic connection on a punctured complex disc. After a finite ramification $(z = t^e)$ and a formal gauge transformation, the Hukuhara-Levelt-Turrittin theorem gives a decomposition of the form

$$M \otimes_{\mathbb{C}((z))} \mathbb{C}((t)) = \text{direct-sum}_{\phi \in Q} (E^{\phi} \otimes R_{\phi}^{\text{right}}),$$

where $(\phi \in t^{-1} \mathbb{C}[t^{-1}])$ is an exponential factor, (E^{ϕ}) is the rank-one exponential connection, and (R_{ϕ}) is regular singular. The finite set (Q) , with multiplicities and Galois descent data when ramification is present, is the formal exponential type.

Definition 3.1 (Reduced formal type). For this paper the reduced formal type of (M) , denoted $(Q(M))$, is the finite collection of exponential factors appearing after a chosen minimal ramification, modulo addition of holomorphic terms and modulo the descent equivalence imposed by the ramification. It is understood together with multiplicities and the regular-singular formal blocks attached to the factors when those blocks are part of the chosen presentation problem.

The Stokes directions are determined by differences of exponential factors. If $(\delta = \phi - \psi)$ has leading term

$$\delta(t) = a t^{-m} + \text{lower pole order terms}, \quad a \neq 0,$$

then, on the ramified circle in the (t) -coordinate, the pair (ϕ, ψ) contributes $(2m)$ anti-Stokes directions before possible identifications coming from descent or coincidences with other pairs. Along sectors between consecutive directions, the relative exponential growth order is constant; crossing a Stokes direction permits a Stokes factor.

Definition 3.2 (Local Stokes presentation). Fix a class of local meromorphic connections for which a ramification convention, a formal normal form convention, and a sector convention have been declared. A

local Stokes presentation consists of: $\begin{array}{l} \text{item the reduced formal type } \mathcal{Q}; \\ \text{item the regular singular formal blocks } \mathcal{R}_{\phi}; \\ \text{item the cyclic Stokes-direction arrangement determined by all} \\ \text{differences } \phi - \psi; \\ \text{item one Stokes factor in the appropriate unipotent Stokes group for each Stokes} \\ \text{direction; item the compatibility and descent data required by the chosen ramification and sector conventions.} \end{array}$

4 Formal Type and Directional Cost

Definition 4.1 (Formal-type cost). Let $\mathcal{C}_{\text{form}}(\mathcal{Q})$ be a declared encoding cost for reduced formal types. Examples include a bit-height cost for the coefficients of the exponential factors, a pole-order cost, a ramification cost, or a combined height and combinatorial-size cost. The following results do not depend on a specific choice, only on the fact that the formal type is an extracted observable with a declared overhead.

Theorem 4.2 (Formal-type lower bound). Let $\mathcal{C}_0$ be a class of local meromorphic connections and let $(\mathcal{P}, \rho, \mathcal{C}_{\mathcal{P}})$ be complete on $\mathcal{C}_0$. Suppose the reduced formal type $[\mathcal{M} \mapsto \mathcal{Q}(\mathcal{M})]$ is extracted from presentations with overhead $\mathcal{F}_{\text{form}}$. Then every $(\mathcal{M} \in \mathcal{C}_0)$ satisfies $[\mathcal{C}_{\mathcal{P}}(\mathcal{M})] \geq \inf_{t: \mathcal{F}_{\text{form}}(t) \geq \mathcal{C}_{\text{form}}(\mathcal{Q}(\mathcal{M}))} \mathcal{C}_{\text{form}}(\mathcal{Q}(\mathcal{M}))$.

Proof. This is the extracted-observable lower bound applied to the observable $\mathcal{Q}$. The analytic content is the existence and invariance of the reduced formal type; the cost conclusion is purely presentation-theoretic.

Definition 4.3 (Directional profile). The directional profile $\Theta(\mathcal{M})$ is the cyclic finite set of Stokes directions, with multiplicity data recording which ordered pairs of exponential factors become active at each direction. A cost $\mathcal{C}_{\Theta}$ may charge for the number of directions, their angular precision, the ramification coordinate, and the incidence relation between directions and factor pairs.

Proposition 4.4 (Pole order forces lifted directions). Let ϕ and ψ be two exponential factors after a fixed ramification $(z = t^e)$. If $[\phi(t) - \psi(t) = a t^{-m} + O(t^{-m+1}), \text{quad } a \neq 0]$ then the lifted Stokes-direction set for the pair (ϕ, ψ) on the (t) -circle has cardinality $(2m)$.

Proof. Write $(t = re^{i\theta})$. The leading contribution to the real part of the difference is $[\text{Re}(a t^{-m}) = r^{-m} \text{Re}(a e^{-im\theta})]$. The Stokes directions for the pair are the directions where the leading exponential comparison is purely oscillatory, equivalently where this real part vanishes. As (θ) runs around a circle, the argument of $(a e^{-im\theta})$ winds (m) times. The real part vanishes twice in each turn. Hence there are $(2m)$ lifted directions. Lower-order pole terms do not change the number of leading anti-Stokes directions for sufficiently small (r) ; they only perturb the comparison away from the leading-direction set.

Theorem 4.5 (Directional lower bound). Assume Θ is extracted from a complete local presentation system with overhead $\mathcal{F}_{\Theta}$. Then $[\mathcal{C}_{\mathcal{P}}(\mathcal{M})] \geq \inf_{t: \mathcal{F}_{\Theta}(t) \geq \mathcal{C}_{\Theta}(\Theta(\mathcal{M}))} \mathcal{C}_{\Theta}(\Theta(\mathcal{M}))$. In particular, for any cost $\mathcal{C}_{\Theta}$ that dominates the number of lifted Stokes directions, a pair of exponential factors whose difference has pole order (m) contributes a lower bound of order at least $(2m)$, before identifications imposed by descent or by coincident directions.

Proof. The first inequality is the extracted-observable lower bound. The final assertion follows from the preceding proposition and from the assumed domination of cardinality by $\mathcal{C}_{\Theta}$.

5 Stokes Factors and Ordinary-Monodromy Fibres

Fix a formal type $\mathcal{Q}$ and sector conventions. For each Stokes direction (θ) , the allowed Stokes factor lies in a unipotent algebraic group $\mathcal{U}_{\theta}(\mathcal{Q})$. The full Stokes-factor space is an iterated product

$$\mathcal{U}(\mathcal{Q}) = \prod_{\theta \in \Theta(\mathcal{Q})} \mathcal{U}_{\theta}(\mathcal{Q}),$$

with compatibility and product constraints determined by the formal monodromy convention.

Definition 5.1 (Stokes-factor observable). The Stokes-factor observable sends a local connection $(\mathcal{M})$ with formal type $\mathcal{Q}$ to its tuple $[S(\mathcal{M}) \in \mathcal{U}(\mathcal{Q})]$ in the declared Stokes coordinates. Its cost $\mathcal{C}_{\text{Stokes}}$ may be a coordinate-height cost, a sparsity cost, a dimension-sensitive algebraic cost,

or any declared cost on the corresponding unipotent data.

Theorem 5.2 (Stokes-factor lower bound). Let (S) be extracted from local presentations with overhead (F_{Stokes}) . Then $\mathbf{C}_{\text{Stokes}}(M) \geq \inf_{t: (F_{\text{Stokes}})(t) \geq \mathbf{C}_{\text{Stokes}}(S(M))} t$. Consequently, whenever a family of connections has fixed formal type and varying Stokes factors of increasing coordinate cost, the presentation cost must increase at the corresponding inverse-overhead scale.

Proof. Apply the extracted-observable lower bound to the Stokes-factor observable. The statement with fixed formal type is the same inequality restricted to the fibre of (Q) .

Example 5.3 (Rank-one exponential residual fibre). For $(m \geq 1)$ and $(a \in \mathbb{C}^*)$, consider the rank-one connection $\nabla_{a,m} = d - d(a z^{-m})$ on the punctured disc. A local solution is $(\exp(a z^{-m}))$. Since (z^{-m}) is single-valued around the puncture, the ordinary monodromy of this rank-one local system is trivial. Nevertheless the formal exponential factor $(a z^{-m})$ changes with (a) and (m) .

Proposition 5.4 (Ordinary monodromy does not separate formal type). The ordinary monodromy observable has an infinite residual fibre on the class of rank-one exponential connections $(\nabla_{a,m})$. Inside the fibre of trivial ordinary monodromy, the reduced formal type can have arbitrarily large pole order and arbitrarily large coefficient height.

Proof. The preceding example shows that every $(\nabla_{a,m})$ has trivial ordinary monodromy. If $((a,m) \neq (a',m'))$, then the exponential factors $(a z^{-m})$ and $(a' z^{-m'})$ are distinct modulo holomorphic terms unless the negative Laurent polynomials agree. Hence ordinary monodromy is constant on an infinite set of distinct formal types. The pole order (m) and the height of (a) , when a height model is declared, can be made arbitrarily large.

Corollary 5.5 (Formal cost invisible to ordinary monodromy). Let a presentation system be complete for the rank-one exponential family and suppose the formal type is extracted with overhead (F_{form}) . Then no cost model depending only on ordinary monodromy can bound the presentation cost of this family. More precisely, $\mathbf{C}_{\text{form}}(\nabla_{a,m}) \geq \inf_{t: (F_{\text{form}})(t) \geq \mathbf{C}_{\text{form}}(a z^{-m})} t$.

Proof. Ordinary monodromy is constant on the family, while the formal-type observable is not. The inequality is the formal-type lower bound applied to $(\nabla_{a,m})$.

6 Local Stokes Normal Forms

The previous sections give lower bounds from observables. Conversely, the classical Stokes classification gives a complete local presentation format once conventions have been fixed.

Definition 6.1 (Local normal-form data). For a bounded class of local meromorphic connections, local Stokes normal-form data consist of:

- a ramification degree and descent convention;
- a reduced formal type (Q) ;
- regular-singular formal blocks (R_{ϕ}) ;
- the induced directional profile $(\Theta(Q))$;
- Stokes factors $(S_{\theta} \in U_{\theta}(Q))$;
- the formal monodromy and product compatibility data.

Theorem 6.2 (Completeness of local Stokes normal forms). After fixing the above conventions, local Stokes normal-form data determine the analytic isomorphism class of the corresponding local meromorphic connection. Conversely, every connection in the bounded class determines such data.

Proof. The Hukuhara-Levelt-Turrittin theorem gives the formal decomposition after ramification. On each sector between consecutive Stokes directions, one chooses sectorial normalizations compatible with the formal type. On overlaps of adjacent sectors the transition maps are precisely the Stokes factors, together with the formal monodromy contribution. Thus the listed data reconstruct the sectorial gluing problem. If two connections have the same listed data, their sectorial normal forms can be identified on every sector and the transition functions agree on overlaps, so the sectorial identifications glue to an analytic isomorphism. Conversely, an analytic connection produces the same type of data by formal decomposition and sectorial comparison.

Proposition 6.3 (Quasi-minimality under bounded conventions). Fix a bounded local class in which ramification degree, rank, sector convention, and regular-singular block format are bounded. Suppose a complete presentation system admits extraction of (Q) , (Θ) , and (S) with overheads $(F_{\text{form}}, F_{\Theta}, F_{\text{Stokes}})$. Then every complete presentation of (M) has cost at least the maximum

of the corresponding inverse-overhead lower bounds: $\inf\{t: F_{\text{form}}(t) \geq \text{mathbf C}_{\text{form}}(Q(M))\}, \inf\{t: F_{\text{Theta}}(t) \geq \text{mathbf C}_{\text{Theta}}(\text{Theta}(M))\}, \inf\{t: F_{\text{Stokes}}(t) \geq \text{mathbf C}_{\text{Stokes}}(S(M))\} \text{ Bigr}\}.$ If the declared local Stokes normal-form encoder has total cost bounded above by a function of these three observable costs, then the lower and upper bounds match up to the corresponding overhead functions.

Proof. The lower bound is obtained by applying the extracted-observable theorem to each observable and taking the maximum. The final assertion is an upper-bound statement for the chosen encoder: once the normal-form data have been encoded, the completeness theorem reconstructs the analytic class. Hence the presentation cost is bounded above by the declared cost of that normal-form encoder, while the first part shows that each of its essential components is forced by any complete presentation.

7 Global Stokes Data and Turning Loci

On higher-dimensional bases, good formal structures need not be present everywhere without modification. Turning loci record where the formal decomposition changes or fails to be good in the chosen model. The following statements are therefore formulated for classes where the relevant good-structure operations are part of the declared presentation package.

Definition 7.1 (Global good class). A global good class is a bounded class of meromorphic connections on smooth complex varieties with divisor of poles, together with:

- a bounded format for local charts and normal-crossing data;
- a declared procedure, possibly after admissible modifications, for obtaining good formal structures;
- a finite presentation format for the resulting formal types, Stokes hypersurfaces, gluing data, and descent data.

Definition 7.2 (Turning observable). The turning observable $\mathcal{T}$ records the locus where the original model does not already carry the declared good formal structure, together with the combinatorial or local data needed by the chosen modification procedure. Its cost $\text{mathbf C}_{\mathcal{T}}$ may measure degree, number of components, local multiplicities, or a heightened algebraic presentation of the locus.

Theorem 7.3 (Turning-data lower bound). Let $\mathcal{C}_0$ be a global good class, and suppose the turning observable $\mathcal{T}$ is extracted from presentations with overhead $F_{\mathcal{T}}$. Then $\text{mathbf C}_{\mathcal{C}}(M) \geq \inf\{t: F_{\mathcal{T}}(t) \geq \text{mathbf C}_{\mathcal{T}}(\mathcal{T}(M))\}.$

Proof. This is the extracted-observable lower bound applied to $\mathcal{T}$. The geometric input is that the chosen class has a well-defined turning observable in the declared presentation format.

Definition 7.4 (Good-model resolution cost). Let $\Pi(M)$ be the set of admissible modifications $(\pi: Y \rightarrow X)$ in the declared global class such that the pulled-back connection $(\pi^* M)$ has good formal structure. If $\text{mathbf C}_{\text{mod}}$ is a cost on such modifications, define $\text{mathbf C}_{\text{res}}(M) = \inf_{\pi \in \Pi(M)} \text{mathbf C}_{\text{mod}}(\pi).$ This is not an intrinsic number independent of conventions; it is a presentation cost for resolving the failure of good formal structure in the chosen global model.

Theorem 7.5 (Good-model resolution lower bound). Assume that every complete global Stokes presentation in the chosen class determines an admissible modification $(\pi: Y \rightarrow X)$ for which $(\pi^* M)$ has good formal structure, with extraction overhead F_{res} . Then $\text{mathbf C}_{\mathcal{C}}(M) \geq \inf\{t: F_{\text{res}}(t) \geq \text{mathbf C}_{\text{res}}(M)\}.$

Proof. Let (p) be a complete presentation of (M) . By hypothesis it determines a good-model modification (π_p) whose cost is at most $F_{\text{res}}(\text{mathbf C}_{\mathcal{C}}(p))$. Since $\text{mathbf C}_{\text{res}}(M)$ is the infimum over all such modifications, one has $\text{mathbf C}_{\text{res}}(M) \leq \text{mathbf C}_{\text{mod}}(\pi_p) \leq F_{\text{res}}(\text{mathbf C}_{\mathcal{C}}(p)).$ Taking the infimum over all presentations (p) realizing (M) gives the stated lower bound.

Theorem 7.6 (Solution-side turning lower bound). Work in a class where Teyssier's solution-theoretic criterion for good formal structure applies: the good formal structure locus is detected by the locus where the restrictions of the solution complexes of (M) and $(\text{End} M)$ to the pole divisor are local systems. Suppose this solution-side failure locus is extracted from the chosen presentation format with overhead F_{solturn} . Then $\text{mathbf C}_{\mathcal{C}}(M) \geq \inf\{t: F_{\text{solturn}}(t) \geq \text{mathbf C}_{\text{turn}}(\text{Turn}(M))\}.$

Proof. Under the stated criterion, the solution-side failure locus equals the turning locus in the chosen class. Thus extracting the solution-side local-system failure locus is also an extraction of $(\text{Turn}(M)).$

The result is the turning-data lower bound applied to this particular extraction route. The role of the analytic theorem is to identify the same obstruction on the formal side and on the solution side; the presentation-theoretic conclusion is the cost lower bound forced by that obstruction.

Theorem 7.7 (Global Stokes-data lower bound). In a global good class, suppose the global Stokes presentation $(G(M))$, including local formal types, Stokes hypersurfaces, Stokes factors, and gluing data, is extracted with overhead (F_G) . Then $\| \mathbf{C}_{\mathcal{C}}(M) \| \geq \inf_{t: F_G(t) \geq \mathbf{C}_G(G(M))} \| \mathbf{C}_G(G(M)) \|$. In particular, families with bounded underlying topological local-system data but unbounded global Stokes or turning data cannot have uniformly bounded complete presentations unless the overhead function already absorbs that unbounded cost.

Proof. The inequality is formal. For the final assertion, restrict to a residual fibre of the underlying topological local-system observable. If (G) or (T) is unbounded on that fibre and is required for completeness, the preceding inequality prevents uniformly bounded complete presentations.

8 Controlled Irregular Riemann-Hilbert Transfer

The irregular Riemann-Hilbert correspondence identifies holonomic $(\mathcal{D})$ -module data with enhanced-sheaf data in a way that remembers Stokes phenomena. Presentation Theory separates the categorical equivalence from the cost statement: the equivalence supplies the mathematical bridge, while a controlled transfer package supplies explicit overhead functions for finite presentations.

Definition 8.1 (Controlled irregular Riemann-Hilbert package). On a bounded good class, a controlled irregular Riemann-Hilbert package consists of:

- begin{itemize}[leftmargin=2em] item a presentation system for meromorphic or holonomic $(\mathcal{D})$ -module objects in the class; item a presentation system for the corresponding enhanced-sheaf or Stokes-filtered objects; item forward and inverse realization procedures implementing the correspondence on the class; item nondecreasing overhead functions $(F_{\rightarrow})$ and $(F_{\leftarrow})$ bounding the cost of these procedures; item verification data ensuring that the presentations lie in the bounded good class. end{itemize}

Theorem 8.2 (Controlled transfer of presentation bounds). Assume a controlled irregular Riemann-Hilbert package on a bounded good class. Let (M) be an object on the $(\mathcal{D})$ -module side and let $(\mathrm{Sol}^{\mathrm{enh}}(M))$ denote its enhanced solution object in the package. Then $\| \mathbf{C}_{\mathrm{enh}}(\mathrm{Sol}^{\mathrm{enh}}(M)) \| \leq F_{\rightarrow}(\| \mathbf{C}_{\mathcal{D}}(M) \|)$, and $\| \mathbf{C}_{\mathcal{D}}(M) \| \leq F_{\leftarrow}(\| \mathbf{C}_{\mathrm{enh}}(\mathrm{Sol}^{\mathrm{enh}}(M)) \|)$. Consequently, any lower bound for an extracted observable on one side transfers to the other side after composing with the corresponding overhead function.

Proof. The first two inequalities are exactly the cost guarantees in the controlled package: a presentation of (M) is transformed into a presentation of its enhanced solution object with overhead $(F_{\rightarrow})$, and a presentation of the enhanced object is transformed back with overhead $(F_{\leftarrow})$. For the final statement, suppose an observable (A) on the enhanced side gives $\| \mathbf{C}_{\mathrm{enh}}(\mathrm{Sol}^{\mathrm{enh}}(M)) \| \geq L(M)$. Combining this with the first inequality gives $\| L(M) \| \leq F_{\rightarrow}(\| \mathbf{C}_{\mathcal{D}}(M) \|)$, which is the transferred lower bound on the $(\mathcal{D})$ -module side. The other direction is identical using $(F_{\leftarrow})$.

Corollary 8.3 (Enhanced-sheaf visibility of Stokes cost). In a controlled irregular Riemann-Hilbert package, local formal type, Stokes directions, and Stokes factors give lower bounds for enhanced-sheaf presentations whenever they are extracted from the enhanced side. Conversely, enhanced-sheaf observables that recover Stokes-filtered data give lower bounds for meromorphic-connection presentations.

Proof. Apply the controlled-transfer theorem to the formal-type, directional, and Stokes-factor lower bounds proved above. The enhanced solution object carries the Stokes-filtered asymptotic information; the controlled package is precisely the assertion that this information is available with declared finite overhead on the bounded class.

9 Synthesis

The local classification of irregular connections supplies a natural normal-form presentation, while Presentation Theory explains which parts of that normal form are forced by any complete finite description. Ordinary monodromy can be constant on large residual fibres; formal type and Stokes data separate those

fibres and therefore impose presentation cost. In global settings, turning loci and Stokes hypersurfaces play the same role once a good-structure package has been fixed. Under controlled irregular Riemann-Hilbert transfer, these lower bounds move between meromorphic-connection presentations and enhanced-sheaf presentations.

The main reusable pattern is therefore:

classical irregular invariant $\rightarrow$ extracted observable
 $\rightarrow$ presentation lower bound.

The analytic theory determines the invariant. The presentation-theoretic layer records how much finite description cost is forced by making that invariant available.

References

- [1] P. Boalch. Geometry and braiding of Stokes data; fission and wild character varieties. *Annals of Mathematics* 179 (2014), 301--365.
- [2] A. D'Agnolo and M. Kashiwara. Riemann-Hilbert correspondence for holonomic D-modules. *Publications mathématiques de l'IHES* 123 (2016), 69--197.
- [3] K. S. Kedlaya. Good formal structures for flat meromorphic connections, I: Surfaces. *Duke Mathematical Journal* 154 (2010), 343--418.
- [4] K. S. Kedlaya. Good formal structures for flat meromorphic connections, II: Excellent schemes. *Journal of the American Mathematical Society* 24 (2011), 183--229.
- [5] T. Mochizuki. *Wild Harmonic Bundles and Wild Pure Twistor D-modules*. Asterisque 340, Societe Mathematique de France, 2011.
- [6] C. Sabbah. *Introduction to Stokes Structures*. Lecture Notes in Mathematics 2060, Springer, 2013.
- [7] W. Wasow. *Asymptotic Expansions for Ordinary Differential Equations*. Dover, 1987.