

UNDER RICCI FLOW, A 3-TORUS GOES FLAT

JOHN LOTT

ABSTRACT. We determine the long-time behavior of the canonical Ricci flow on a closed three-manifold of Thurston type $\mathbb{R}^3$, Nil, or Sol, starting from an arbitrary initial metric. In the Euclidean case, the metric $g(t)$ converges exponentially fast to a flat metric. The Gromov–Hausdorff limit of $(M, t^{-1}g(t))$ is a point in the Euclidean and Nil cases, and a metric circle or a compact interval of positive length in the Sol case. In all three cases, the blowdown flows lifted to the universal cover, $s^{-1}\tilde{g}(s\tau)$, converge, in the pointed Cheeger–Hamilton sense, to an explicit homogeneous expanding Ricci soliton solution.

CONTENTS

1. Introduction	3
1.1. Main result and geometric setting	3
1.2. Previous work and the missing diameter estimate	3
1.3. The blockwise argument	4
1.4. The three geometries and the remaining cases	6
1.5. Organization and conventions	6
AI Acknowledgment.	7
2. Long-time inputs from earlier work	7
2.1. Type-III blowdowns	7
2.2. Almost-flatness and stability near flat metrics	8
2.3. Estimates for invariant torus-bundle flows	9
2.4. Bamler’s Type-III estimate and fixed-threshold alternative	9
3. Finite covers and torus-bundle geometry	11
3.1. Finite torus-bundle covers and monodromy	11
3.2. Torus-bundle conventions	13
3.3. Universal-cover noncollapsing	13
3.4. Averaging and elementary collapse	14
4. Finite-time decay of the non-invariant part	15
4.1. Short-orbit averaging	18
4.2. Translation differences and the Ricci–DeTurck system	20
4.3. The lifted estimate and completion of the proof	24

Date: September 13, 2026.

2020 Mathematics Subject Classification. Primary 53E20; Secondary 57K35.

Key words and phrases. Ricci flow on three-manifolds, Type-III Ricci flow, torus bundles, collapsing, Gromov–Hausdorff convergence, homogeneous Ricci solitons.

5.	Comparison of torus fibrations	25
5.1.	Fiber equivalence and diameter	26
5.2.	Fiber-equivalent small-fiber fibrations	27
5.3.	Fiber-inequivalent fibrations and fixed-fiber improvement	28
6.	One DeTurck stage and regular restart	31
6.1.	Rescaled blocks and admissible stage starts	33
6.2.	Averaged endpoints and positive-time regularization	34
6.3.	Stage preparation and stability	36
7.	One-block contraction and stagewise families	39
7.1.	Admissible-start data	40
7.2.	Persistence of the retained fibration	40
7.3.	Contraction on one large-diameter block	41
7.4.	Stagewise families and switching times	45
8.	Quotient-circle energy, length, and componentwise iteration	45
8.1.	One-dimensional estimates on a smooth stage	46
8.2.	Transition jump estimates	46
8.3.	The scalar energy-length recursion	48
8.4.	Iteration on a connected large-diameter interval	49
9.	Descent through finite covers	57
9.1.	Circle limits under finite quotients	57
9.2.	Diameter and smooth descent	59
10.	The Euclidean and Sol cases	60
10.1.	The Euclidean case	62
10.2.	The Sol case	64
11.	The Nil case	65
11.1.	Large-diameter components and almost-flat gaps	65
11.2.	The global diameter bound and Nil blowdown	66
12.	Proof of the main theorem	67
	Appendix A. Uniform lifted parabolic estimates	68
	Estimates used in the body	68
A.1.	Parabolic Hölder spaces and lifted bounded geometry	69
A.2.	From slice geometry to parabolic coordinate systems	71
A.3.	Scalar-principal systems and coefficient regularity	75
A.4.	Uniform Schauder estimates on lifted cylinders	77
A.5.	Quasilinear derivative gain	82
A.6.	Continuation, smoothing, and regular restart	83
	Appendix B. Source guide for the external inputs	89
	Appendix C. Dependency trees	92
	Appendix D. Parameters and persistent notation	96
	References	101

1. INTRODUCTION

1.1. Main result and geometric setting. Given an initial Riemannian metric on a closed three-manifold, there is a canonical Ricci flow through singularities [7, 15]. One can ask whether its long-time behavior realizes Thurston's geometric decomposition. This is still not completely understood.

After some finite time, the canonical Ricci flow is known to be either empty or smooth; if the original manifold is aspherical, then the manifold that remains is diffeomorphic to it [2, 7]. Hence, to study the long-time behavior in the cases considered here, it is enough to consider immortal Ricci flows, meaning smooth solutions defined for all $t \in [0, \infty)$. The title emphasizes the simplest consequence of the main theorem below: every immortal Ricci flow on a three-dimensional torus converges to a flat metric as $t \rightarrow \infty$.

Theorem 1.1 (Main theorem). *Let M be a closed connected three-manifold which admits a locally homogeneous metric modeled on one of $\mathbb{R}^3$, Nil, or Sol, and let $g(t)$, $0 \leq t < \infty$, be any immortal Ricci flow on M . In all three cases, for every basepoint $\tilde{m} \in \widetilde{M}$ in the universal cover, the lifted blowdown flows $(\widetilde{M}, \tilde{m}, s^{-1}\tilde{g}(s\tau))$, $\tau > 0$, converge as $s \rightarrow \infty$ in the pointed Cheeger–Hamilton sense to a homogeneous expanding Ricci soliton solution. The compact and lifted limits are as follows.*

- (1) *If M has type $\mathbb{R}^3$, then, for every $m \geq 0$, $g(t)$ converges exponentially fast in the C^m -norm of any fixed smooth background metric, as $t \rightarrow \infty$, to a flat metric on M . The Gromov–Hausdorff limit of $(M, t^{-1}g(t))$, as $t \rightarrow \infty$, is a point. The limit of the lifted blowdown flows is the flat expanding soliton solution*

$$(\mathbb{R}^3, 0, g_{\text{Euc}}(\tau)), \quad g_{\text{Euc}}(\tau) = dx^2 + dy^2 + dz^2, \quad \tau > 0.$$

- (2) *If M has type Nil, then the Gromov–Hausdorff limit of $(M, t^{-1}g(t))$, as $t \rightarrow \infty$, is a point. The limit of the lifted blowdown flows is the Nil expanding soliton solution*

$$\left(\mathbb{R}^3, 0, \frac{1}{3\tau^{1/3}} \left(dx + \frac{1}{2}y dz - \frac{1}{2}z dy \right)^2 + \tau^{1/3}(dy^2 + dz^2) \right).$$

- (3) *If M has type Sol, then the Gromov–Hausdorff limit of $(M, t^{-1}g(t))$, as $t \rightarrow \infty$, is a metric circle or a compact interval of positive length. The limit of the lifted blowdown flows is the Sol expanding soliton solution*

$$(\mathbb{R}^3, 0, e^{-2z}dx^2 + e^{2z}dy^2 + 4\tau dz^2).$$

Throughout the universal-cover conclusions, convergence as $s \rightarrow \infty$ is understood in the sense used in [17]: for every sequence $s_j \rightarrow \infty$, the flows $(\widetilde{M}, \tilde{m}, s_j^{-1}\tilde{g}(s_j\tau))$ converge as $j \rightarrow \infty$, in the pointed Cheeger–Hamilton sense, to the displayed pointed Ricci flow.

1.2. Previous work and the missing diameter estimate. The analogous result for Ricci flow on $\mathbb{T}^2$ was proved by Hamilton [14]. In that case, since Ricci flow preserves the conformal structure and the area, one knows exactly what flat metric one gets in the limit.

In contrast, in part (1) of Theorem 1.1, there is no *a priori* way to know what the limiting flat metric will be.

Each manifold M in Theorem 1.1 has a finite regular orientable cover that fibers over a circle with torus fiber and whose monodromy lies in $\mathrm{SL}(2, \mathbb{Z})$. The cover may be chosen to be $\mathbb{T}^3$ with identity monodromy in the Euclidean case, a Nil torus bundle with nontrivial unipotent monodromy in the Nil case, and a Sol torus bundle with hyperbolic monodromy having positive eigenvalues in the Sol case. We may therefore reduce the analytic part of the proof to these three cases and then descend the conclusions through the fixed finite cover.

There are three main geometric inputs. First, the compact and universal-cover blowdown conclusions used here, together with analogous results for the other aspherical Thurston geometries, were proved by the author [17] under the Type-III curvature estimate $|\sec_{g(t)}| = O(t^{-1})$ and the scale-invariant diameter bound $\mathrm{diam}(M, g(t)) = O(\sqrt{t})$. Second, the needed estimates for flows with local $\mathbb{T}^2$ -symmetry were established by the author and Nataša Šešum [18]. Third, Bamler proved the Type-III curvature estimate and a large-time description by increasingly accurate locally $\mathbb{T}^2$ -invariant time-slice models [2, 3, 4, 5, 6].

More precisely, fix K so that $|\sec_{g(t)}| \leq K/t$, choose a geometric cutoff $c_{\mathrm{af}} = c_{\mathrm{af}}(K) > 0$, and divide each sufficiently late time slice into two cases:

$$\begin{aligned} \textit{almost-flat case:} \quad & \mathrm{diam}(M, g(t)) < c_{\mathrm{af}}\sqrt{t}, \\ \textit{torus-bundle case:} \quad & \mathrm{diam}(M, g(t)) \geq c_{\mathrm{af}}\sqrt{t}. \end{aligned}$$

The cutoff c_{af} is chosen only so that the first inequality makes

$$\left(\sup_M |\sec_{g(t)}| \right) \mathrm{diam}(M, g(t))^2$$

smaller than the three-dimensional almost-flat threshold. In the second case, $t^{-1}g(t)$ is increasingly close in high C^k -norms to an invariant metric for a possibly time-dependent torus-bundle structure, and the torus fibers have diameter $o(1)$ with respect to $t^{-1}g(t)$.

The principal missing geometric estimate is

$$\mathrm{diam}(M, g(t)) = O(\sqrt{t}).$$

The almost-flat case already gives this estimate. If the flow had an exact local $\mathbb{T}^2$ -symmetry in the other case, then the diameter bound would follow from [18]. The difficulty is that the error in Bamler's approximation may tend to zero arbitrarily slowly. When the comparison is restarted on successive time intervals, these errors can accumulate and swamp the desired diameter bound. Moreover, Bamler's fibration may change with time. Thus one cannot simply start a locally $\mathbb{T}^2$ -invariant Ricci flow at each time and compare it with the original flow for an indefinitely long interval.

1.3. The blockwise argument. We use three terms consistently. A *block* is a physical-time interval $[S, 100S]$ on which $\mathrm{diam}(M, g(t)) \geq \delta_0\sqrt{t}$ for a fixed threshold $\delta_0 > 0$; parabolic rescaling by S^{-1} identifies it with the normalized interval $[1, 100]$. A *stage* is the Ricci-DeTurck comparison on this normalized interval between the rescaled original flow and an

invariant Ricci-flow background. A *restart* is the factor-100 endpoint rescaling and averaging that produces the initial metrics for the next stage.

We first prove a finite-time decay theorem for the non-invariant part of a metric. Let $h(t)$, $1 \leq t \leq 100$, be an invariant Ricci flow with bounded curvature and a quotient circle of uniformly positive length. If g_{RDT} is a Ricci–DeTurck flow relative to h , sufficiently close to h , and the local $\mathbb{T}^2$ -orbits are sufficiently small, then we show that the size of the non-invariant part of g_{RDT} on $[10, 100]$ is an arbitrarily small multiple of its size at time 1. Intuitively, at the linearized level, the non-invariant fluctuations correspond to torus modes whose Lichnerowicz eigenvalues become very negative if the torus fibers are small. These negative eigenvalues then cause a fast decay for the non-invariant part of the metric. The actual proof uses an averaging inequality on short torus orbits and parabolic Schauder estimates. The Schauder estimates are all carried out on the universal cover, which is uniformly noncollapsed under the geometric hypotheses. In this way, the decay estimates are uniform with respect to collapsing fibers. Using the Ricci–DeTurck gauge is essential for the Schauder estimates.

We iterate this estimate on intervals $[T_k, T_{k+1}]$, where $T_{k+1} = 100T_k$, using one fixed fibration. In order to apply the finite-time decay theorem, we need to know that the torus fibers remain uniformly small on blocks. The argument for this uses the time-dependent fibration supplied by Bamler.

This is incorporated into the central block-contraction theorem, which proves the conceptual implication

$$\text{admissible block start} \quad \longmapsto \quad \text{admissible restart}, \quad b_{\text{new}} \leq \frac{1}{2}b_{\text{old}},$$

where b is the non-invariant error. To apply it at the first stage, we choose a sufficiently late time and a fibration supplied by Bamler. We parabolically rescale the original flow, average the initial metric with respect to the chosen fibration, evolve the average by Ricci flow, and put the original rescaled flow in Ricci–DeTurck gauge relative to that invariant solution. At the end of the stage, we rescale, average again, and continue with the same fixed fibration. In physical time this produces a continuous, piecewise smooth pullback $g'(t)$ of $g(t)$ and a piecewise Ricci-flow comparison family $h(t)$, which may jump at switching times; see Figure 2.

For the invariant comparison flow $h(t)$, the estimates of [18] control the increase of the quotient-circle length on each smooth stage. The jumps of h between stages produce a summable sequence of length errors. A scalar recursion then shows that the length of the quotient circle is $O(\sqrt{t})$. On a persistent large-diameter tail, the quantitative lower bound for the quotient length forces the normalized stage-start lengths to converge and the energies to tend to the self-similar value 2. A two-sided estimate at each restart and the smooth-stage inequalities propagate this convergence throughout every block; comparison with Bamler’s shrinking fibers then gives a full metric-circle limit on the working cover. The fiber-diameter bounds yield the same diameter estimate for $h(t)$, and hence for $g(t)$.

1.4. The three geometries and the remaining cases. The block construction and quotient-length recursion are applied separately on each maximal interval of sufficiently late times for which the normalized diameter $t^{-1/2} \text{diam}(M, g(t))$ stays above the fixed threshold; the argument also controls a possible terminal partial block. The same block and quotient-length machinery is used in the Euclidean, Nil, and Sol cases. The Thurston type matters only in determining whether the normalized diameter eventually stays above the threshold or sometimes drops below it. For a Sol manifold, the latter alternative would produce an almost-flat slice and is topologically impossible. Hence the iteration applies for all sufficiently large times and gives both the required diameter estimate and the full metric-circle limit on the working cover. The Sol universal-cover blowdown comes from [17]; finite-cover descent then gives part (3) of Theorem 1.1.

In the Euclidean case, if one sufficiently late time slice is sufficiently almost flat, then the stability results of [13, 18] give exponential convergence to a flat metric. Otherwise, the torus-bundle case persists on a tail. The resulting diameter bound, combined with [17], then forces the flow into the almost-flat case. This proves part (1).

For a Nil manifold, we apply the same iteration separately on each connected component of

$$\{t : t^{-1/2} \text{diam}(M, g(t)) \geq c_{\text{af}}\}.$$

The complementary times are already almost flat and satisfy the required diameter bound, while Type-III metric distortion controls the partial blocks at the ends of each component. Once $\text{diam}(M, g(t)) = O(\sqrt{t})$ is known, [17] proves part (2).

In a broader context, one may speculate that an immortal Ricci flow on a compact three-manifold M has diameter growth $O(\sqrt{t})$ if and only if M admits a locally homogeneous Riemannian metric [17, Remarks 1.4–1.5]. The “only if” implication follows by combining Bamler’s Type-III estimate with [17, Proposition 3.5]. The “if” implication is known for H^3 from Perelman’s work, and for $\mathbb{R}^3$, Nil, and Sol from the present paper. The remaining geometries are $H^2 \times \mathbb{R}$ and $\widetilde{\text{SL}(2, \mathbb{R})}$. The methods here do not extend directly to these cases. Bamler’s results do not provide an approximation by an S^1 -invariant metric on a circle bundle over a higher-genus surface [2, Theorem 1.4(e)]. Moreover, even for an S^1 -invariant Ricci flow on such a circle bundle, the long-time behavior is not known in general. Section 2 of [18] treats the special case of a warped-product metric, but curvature of the circle bundle introduces additional complications.

1.5. Organization and conventions. The structure of the paper is the following. Section 2 states the long-time inputs from earlier work and the fixed-threshold formulation used in the block argument. Section 3 selects the working torus-bundle cover and develops the local bundle geometry. Section 4 proves finite-time decay of the non-invariant part. Section 5 compares torus fibrations to bound the diameters of fibers. Section 6 constructs one Ricci–DeTurck stage and its regular restart. Section 7 proves the one-block contraction. Section 8 combines that contraction with the quotient-circle energy and length estimates to obtain a single iteration theorem for any connected large-diameter interval. Section 9 collects the finite-cover descent facts used in the final arguments. Sections 10 and 11 prove

the Euclidean, Sol, and Nil conclusions on the working cover, and Section 12 descends them and completes the proof of Theorem 1.1. Appendix A contains the uniform lifted parabolic estimates, and Appendix B records the detailed guide from the external formulations in Section 2 to the cited sources.

For the reader's convenience, Appendix C gives dependency trees of the principal statements in the paper. Each arrow runs from a dependency to a result that uses it, and selected edge labels record the main parameters or quantities carried along that implication. Appendix D lists the persistent parameters and notation alphabetically, first by Greek letter and then by Roman letter.

Unnamed constants denoted by C may change from line to line. Named constants retain the values assigned to them. All constants depend only on the parameters shown in the relevant statements, not on the particular collapsed manifold once the collapse parameter is sufficiently small.

AI Acknowledgment. The author developed the proof structure, intermediate results, and proof sketches. ChatGPT-5.6 Sol was used to elaborate detailed arguments from this material and to assist in preparing the figures and the dependency and notation appendices. The author assumes responsibility for all mathematical claims and source attributions in the manuscript.

2. LONG-TIME INPUTS FROM EARLIER WORK

The proof uses several long-time results from the literature as external inputs. This section states them in forms that are uniform under passage to covers and compatible with the later rescaling-and-restart argument. It also fixes the thresholds and quantitative formulations that will be used without further alteration in the later analytic and geometric constructions. A detailed guide from these formulations to the cited statements is given in Appendix B.

2.1. Type-III blowdowns. The following proposition is the part of [17, Theorem 1.2] used below.

Proposition 2.1 (Blowdowns under Type-III curvature and diameter control). *Let M be a closed connected orientable three-manifold modeled on $\mathbb{R}^3$, Nil, or Sol, and let $g(t)$ be an immortal Ricci flow. Suppose that, for all sufficiently large t ,*

$$(2.1) \quad |\mathrm{Rm}_{g(t)}| \leq C/t, \quad \mathrm{diam}(M, g(t)) \leq C\sqrt{t}.$$

Then the compact and universal-cover blowdown conclusions of [17, Theorem 1.2] hold in the following form. The compact normalized metrics converge to a point in the Euclidean and Nil cases. In the Sol case, every Gromov–Hausdorff subsequential limit is a metric circle or a compact interval. The lifted blowdown flows converge in the pointed Cheeger–Hamilton sense to the corresponding homogeneous expanding soliton solutions. In the Euclidean case, the limit is the static Euclidean flow $(\mathbb{R}^3, g_{\mathrm{Euc}}(\tau))$, where $g_{\mathrm{Euc}}(\tau) = dx^2 + dy^2 + dz^2$ for $\tau > 0$.

The new geometric estimate needed in this paper is the second inequality in (2.1).

2.2. Almost-flatness and stability near flat metrics. Fix a dimensional constant $\varepsilon_{\text{af}} > 0$ below the three-dimensional Gromov–Ruh threshold.

We use the following almost-flat theorem of Gromov and Ruh [12, 19].

Proposition 2.2 (Almost-flat topological theorem). *If a metric g on a closed three-manifold satisfies*

$$(2.2) \quad \left(\sup_M |\sec_g| \right) \text{diam}(M, g)^2 < \varepsilon_{\text{af}},$$

then the manifold is an infranilmanifold; in particular, its fundamental group is virtually nilpotent.

The following uniform version of the flat stability theorem of Guenther–Isenberg–Knopf [13] is the external analytic input used in the Euclidean argument.

Proposition 2.3 (Stability near a compact family of flat metrics). *Let M be a fixed torus and let $\mathcal{F}$ be a compact set of flat metrics on M (not modulo diffeomorphisms). For a fixed $\rho \in (0, 1)$, there is a uniform $h^{2+\rho}$ -neighborhood of $\mathcal{F}$, measured using one fixed smooth reference metric on M , such that the Ricci flow starting from any smooth Riemannian metric in this neighborhood exists for all future time and, without applying a time-dependent pullback, converges exponentially as a tensor field on M to a smooth flat metric. Here $h^{2+\rho}$ is the little Hölder space, namely the closure of smooth metrics in the $C^{2,\rho}$ -norm. The decay rate and the constants in the fixed-background stability estimates may be chosen uniformly on $\mathcal{F}$. Convergence holds in every fixed C^m -norm, with higher-order constants that may also depend on higher norms of the smooth initial metric.*

Proof. For one flat background, [13, Theorem 3.7] gives exponential stability for the Ricci flow. In its proof, the corresponding Ricci–DeTurck convergence is established first, and [13, Proposition 3.6] is then used to transfer this convergence to the ungauged Ricci flow. Fix a smooth reference metric on the underlying torus and use its little Hölder norm to metrize the space of $h^{2+\rho}$ -metrics. For every $h \in \mathcal{F}$, choose an open stability neighborhood U_h and the exponential estimate supplied by [13]. There are $h_1, \dots, h_J \in \mathcal{F}$ such that $\mathcal{F} \subset U := \bigcup_{j=1}^J U_{h_j}$. Since $\mathcal{F}$ is compact and U is open, there is $\delta > 0$ such that the δ -neighborhood of $\mathcal{F}$, in this one fixed $h^{2+\rho}$ -metric, is contained in U . Thus every initial metric in that uniform neighborhood lies in at least one of the finitely many U_{h_j} . The minimum of the corresponding exponential rates and the maximum of the constants, together with the equivalence of the finitely many little Hölder norms occurring in these estimates, give uniform stability data. For smooth initial data, positive-time parabolic estimates bootstrap the exponentially decaying Ricci–DeTurck perturbation in every fixed C^m -norm. The resulting estimates for the DeTurck vector field give convergence of the associated diffeomorphisms in the corresponding orders, transferring the exponential convergence to the ungauged flow. The decay rate remains uniform on $\mathcal{F}$, while the higher-order constants may also depend on higher norms of the smooth initial metric. $\square$

The diameter-normalized near-flat compactness statement that places a metric in this fixed-coordinate stability neighborhood is specific to the Euclidean argument. It is formulated and proved in Subsection 10.1, immediately before it is used.

2.3. Estimates for invariant torus-bundle flows. The next proposition states the two invariant-flow estimates from [18, Lemmas 3.3 and 3.7] used in the iteration.

Proposition 2.4 (Invariant energy and quotient-length inequalities). *Let $h(t)$ be an invariant Ricci flow on a torus bundle, written in the notation of [18] and Subsection 8.1, and set*

$$E(t) = \max_{S^1} h^{yy} \operatorname{Tr}((G^{-1}G_y)^2), \quad L(t) = L(S^1, h_{yy}(t)).$$

Then [18, Lemmas 3.3 and 3.7] give

$$(2.3) \quad \frac{d}{dt} \log L(t) \leq \frac{1}{4} E(t).$$

If $t_1 \geq t_0$ and $E(t_0) > 0$, they also give

$$(2.4) \quad \frac{1}{E(t_1)} - \frac{1}{E(t_0)} \geq \frac{1}{2}(t_1 - t_0).$$

If $E(t_0) = 0$, the maximum-principle inequality underlying the energy estimate gives $E(t) = 0$ for every $t \geq t_0$; in that case the reciprocal inequality is not invoked and (2.3) has zero right-hand side. These inequalities are used only on the smooth invariant stages constructed in Section 7.

2.4. Bamler's Type-III estimate and fixed-threshold alternative. We use Bamler's Type-III estimate in the following form [2, Corollary 1.2].

Proposition 2.5 (Global scale-invariant curvature bound). *Let M be a closed connected three-manifold of type $\mathbb{R}^3$, Nil, or Sol, and let $g(t)$ be an immortal Ricci flow. There is a constant $K < \infty$ such that, for every $t > 0$,*

$$(2.5) \quad |\sec_{g(t)}| \leq K/t.$$

The next proposition is a fixed-threshold, tail-uniform reformulation of the topology-specific alternatives in [2, Theorem 1.4(b)–(d)]. The detailed construction of the invariant comparison metrics, together with the derivative and fiber estimates used here, is given in [6, Section 4.4, especially Proposition 4.9], using [3, 4, 5]. The initial normalization in the cited theorem is removed by a constant parabolic rescaling, as explained in Appendix B; no normalization of $g(0)$ is assumed below.

Proposition 2.6 (Fixed-threshold almost-flat/torus-bundle alternative). *Let M be $\mathbb{T}^3$, a compact Nil torus bundle with nontrivial unipotent monodromy, or a compact Sol torus bundle with monodromy $H \in \operatorname{SL}(2, \mathbb{Z})$ satisfying $\operatorname{Tr} H > 2$, and let $g(t)$ be an immortal Ricci flow on M . Let K be a Type-III constant supplied by Proposition 2.5. Fix any $c_0 \in (0, 1)$ such that*

$$(2.6) \quad Kc_0^2 < \varepsilon_{\text{af}}.$$

The parameter c_0 is used only as the diameter threshold separating the two geometric alternatives below; it is not an analytic smallness parameter and will not be decreased later to meet compactness or stability tolerances.

There exist a finite time $T_B(c_0) < \infty$ and a nonincreasing function

$$\varepsilon_B^{c_0} : [T_B(c_0), \infty) \longrightarrow (0, \tfrac{1}{2}], \quad \varepsilon_B^{c_0}(t) \longrightarrow 0,$$

with

$$(2.7) \quad \varepsilon_B^{c_0}(t) < c_0, \quad t \geq T_B(c_0),$$

such that every $t \geq T_B(c_0)$ belongs to exactly one of the following diameter-defined alternatives.

(1) If

$$(2.8) \quad \text{diam}(M, g(t)) < c_0 \sqrt{t},$$

then

$$(2.9) \quad \left(\sup_M |\sec_{g(t)}| \right) \text{diam}(M, g(t))^2 < K c_0^2 < \varepsilon_{\text{af}}.$$

Thus the slice is quantitatively almost flat.

(2) If

$$(2.10) \quad \text{diam}(M, g(t)) \geq c_0 \sqrt{t},$$

then there exist a torus fibration $\pi_t^B : M \rightarrow S_t^1$ and an invariant metric g_t^B , in the sense described in Subsection 3.2, such that

$$(2.11) \quad (1 + \varepsilon_B^{c_0}(t))^{-1} g_t^B \leq g(t) \leq (1 + \varepsilon_B^{c_0}(t)) g_t^B.$$

Moreover, for every integer m with $0 \leq m < (\varepsilon_B^{c_0}(t))^{-1}$, the metric $t^{-1} g_t^B$ is $\varepsilon_B^{c_0}(t)$ -close to $t^{-1} g(t)$ in C^m , relative to $t^{-1} g(t)$, and every fiber has $g(t)$ -diameter at most $\varepsilon_B^{c_0}(t) \sqrt{t}$.

Labels (1) and (2) refer only to the diameter inequalities. A torus comparison may also exist in the first alternative, but it is not used there. In the second alternative the fibration, quotient circle, and comparison metric may depend on t , with no asserted compatibility or uniqueness.

Fixed global application. For the global almost-flat/torus-bundle decomposition, fix once and for all a cutoff $c_{\text{af}} = c_{\text{af}}(K) \in (0, 1)$ satisfying

$$K c_{\text{af}}^2 < \varepsilon_{\text{af}},$$

and write $T_B(c_{\text{af}})$ and $\varepsilon_B^{c_{\text{af}}}$ for the corresponding tail data. Any additional application of Proposition 2.6 is labeled by its own cutoff and has its own tail data. The Euclidean argument will make one such additional application.

Working-cover application. All later applications of Proposition 2.6 are made directly on the single working cover fixed after Proposition 3.1; no additional finite cover is taken.

3. FINITE COVERS AND TORUS-BUNDLE GEOMETRY

This section establishes preliminary results about the topology, local torus action, lifted noncollapsing, and averaging formalism. Its conclusions allow the later parabolic estimates to be uniform despite collapse of the compact quotient.

The section is organized as follows. The first subsection, *Finite torus-bundle covers and monodromy*, passes to a finite regular orientable torus-bundle cover that, according to the geometry, is $\mathbb{T}^3$ with identity monodromy, a Nil bundle with nontrivial unipotent monodromy, or a Sol bundle with hyperbolic monodromy having positive eigenvalues. The second subsection, *Torus-bundle conventions*, interprets the bundle as a twisted principal $\mathbb{T}^2$ -bundle and fixes the local translation actions and the meaning of local $\mathbb{T}^2$ -invariance. The third subsection, *Universal-cover noncollapsing*, proves a collapse-independent injectivity-radius lower bound on the universal cover from an invariant bilipschitz comparison, a uniform volume lower bound, and the Cheeger–Gromov–Taylor estimate. The final subsection, *Averaging and elementary collapse*, defines the global fiberwise averaging operator and the non-invariant part, proves fixed-time contraction of averaging in invariant Hölder norms, and quantifies collapse to the metric quotient circle when the invariant fibers are short.

We use distinct notation for the three covers that occur in the proof. A fixed finite regular cover of the original manifold is denoted $\pi_{\text{fin}} : M_{\text{fin}} \rightarrow M$, with pulled-back metric g_{fin} . If M is a torus bundle, the infinite cyclic cover obtained by unwrapping its base circle is denoted $\pi_{\text{cyc}} : M_{\text{cyc}} \rightarrow M$; a pulled-back tensor is written with the subscript cyc. The universal cover is denoted $\widetilde{M}$, and universal-cover lifts carry tildes. In particular, the compact finite cover and the noncompact cyclic cover are never represented by the same symbol.

3.1. Finite torus-bundle covers and monodromy. For a matrix $H \in \text{SL}(2, \mathbb{Z})$, we use the standard terminology that H is *elliptic* if it has finite order, *parabolic* if it has infinite order and $|\text{Tr } H| = 2$, and *hyperbolic* if $|\text{Tr } H| > 2$. These are the mutually exclusive possibilities in $\text{SL}(2, \mathbb{Z})$, and replacing H by a positive power preserves its type.

Proposition 3.1 (Finite torus-bundle covers of the three geometries). *Let M be a closed connected three-manifold modeled on $\mathbb{R}^3$, Nil, or Sol. Then M has a finite regular orientable cover M_{fin} which is the total space of a $\mathbb{T}^2$ -bundle over S^1 , with monodromy $H \in \text{SL}(2, \mathbb{Z})$. The cover may be chosen so that the following hold:*

- (1) *if M is modeled on $\mathbb{R}^3$, then $M_{\text{fin}} \cong \mathbb{T}^3$ and $H = I$;*
- (2) *if M is modeled on Nil, then H is nontrivial unipotent;*
- (3) *if M is modeled on Sol, then H is hyperbolic with positive eigenvalues, equivalently $\text{Tr } H > 2$.*

In particular, the monodromy is elliptic, parabolic, or hyperbolic according as the original geometry is $\mathbb{R}^3$, Nil, or Sol.

Proof. For a compact Euclidean three-manifold, the translation subgroup gives a finite cover by $\mathbb{T}^3$. A compact Nil manifold has a finite cover which is a torus bundle over a circle with nontrivial unipotent monodromy, while a compact Sol manifold has a cover of degree at most

four which is a torus bundle with hyperbolic monodromy. These facts, together with the monodromy classification, are recorded in [20, Sections 4–5, especially Theorems 4.16, 5.3, and 5.5]. Pulling back first by the degree-two cover of the base changes the monodromy from H_0 to $H = H_0^2$ and places it in $\mathrm{SL}(2, \mathbb{Z})$ without changing the associated geometry. In the Sol case this also makes both eigenvalues positive and gives $\mathrm{Tr} H = (\mathrm{Tr} H_0)^2 - 2 \det H_0 > 2$. Finally, we take the cover corresponding to the core of this subgroup in $\pi_1(M)$, meaning the intersection of all its conjugates, or equivalently the largest normal subgroup of $\pi_1(M)$ contained in it. This is a finite regular refinement. To see directly that the torus-bundle structure and its type persist under the refinement, write the bundle group as

$$\Gamma_H = \mathbb{Z}^2 \rtimes_H \mathbb{Z}.$$

For any finite-index subgroup $\Gamma' \leq \Gamma_H$, the intersection $L = \Gamma' \cap \mathbb{Z}^2$ is a finite-index lattice in $\mathbb{Z}^2$, and the image of Γ' under the projection to $\mathbb{Z}$ is $m\mathbb{Z}$ for some $m \geq 1$. Choosing an element of Γ' over m identifies Γ' with $L \rtimes_{H^m} \mathbb{Z}$; conjugation preserves L and is the restriction of H^m . Hence the corresponding connected finite cover is again a toroidal mapping torus. The stronger geometry-specific conclusion is also preserved. If $H = I$, then the refined monodromy is again the identity and the total space is a three-torus. If H is nontrivial unipotent, then $H^m - I$ is nonzero on $\mathbb{Q}^2$, so its restriction to the finite-index lattice L is nonzero; hence the refined monodromy remains nontrivial unipotent. In the Sol case, the eigenvalues of H^m remain positive and off the unit circle, so the refined monodromy still has trace greater than 2. In terms of the original monodromy, this power is H_0^{2m} , regardless of the parity of m . This applies in particular to the core refinement. $\square$

Working-cover convention. Apply Proposition 3.1 once to the original manifold, fix the resulting regular orientable torus-bundle cover, and lift the flow to it. Take the identity cover only when all the geometry-specific conditions of that proposition already hold; in the Sol case this requires $\mathrm{Tr} H > 2$. The phrase *working cover* always refers to this one fixed cover and its lifted flow. It is $\mathbb{T}^3$ with identity monodromy in the Euclidean case, a Nil torus bundle with nontrivial unipotent monodromy in the Nil case, and a Sol torus bundle with $\mathrm{Tr} H > 2$ in the Sol case. Thus it lies exactly in the scope of Proposition 2.6. The base double cover and regular refinement are part of this initial choice; no further cover is needed to apply Bamler’s theorem.

Apply that proposition with the fixed global choice $c_0 = c_{\mathrm{af}}$ directly to the flow on this cover, and carry out all block constructions on the same cover. The curvature constant is unchanged by pullback, while $T_B(c_{\mathrm{af}})$ and $\varepsilon_B^{c_{\mathrm{af}}}$ are the tail data obtained for the lifted flow. Any later auxiliary cutoff is applied to this same flow and receives separately labeled tail data. In the local constructions below, the working cover is denoted simply by M ; the notation M_{fin} is restored only when it must be distinguished from the original manifold in the final assembly. The cyclic cover used for torus translations is always M_{cyc} , and its universal cover is $\widetilde{M}$. All descent statements are collected in Section 9.

3.2. Torus-bundle conventions. Under the [working-cover convention](#), fix the bundle projection $\pi : M \rightarrow S^1$ on the already selected torus-bundle cover. Its monodromy is represented by a fixed matrix $H \in \mathrm{SL}(2, \mathbb{Z})$, with the terminology fixed in Subsection 3.1. No further preliminary passage to a cover is made in the local constructions below.

We view such a bundle as a *twisted principal $\mathbb{T}^2$ -bundle*. Choose fiber orientations and an atlas over intervals in the base in which the fibers are principal $\mathbb{T}^2$ -spaces. On overlaps, the fiber translations are conjugated by constant elements of $\mathrm{SL}(2, \mathbb{Z})$ coming from the monodromy. Equivalently, after choosing a coordinate on the base, the bundle is the mapping torus of an affine automorphism of $\mathbb{T}^2$ whose linear part is H . The local principal translations give local $\mathbb{T}^2$ -actions. A tensor field, and in particular a metric, is called *invariant* if, in each such chart, it is invariant under these local fiber translations. This condition is independent of the chart, because the transition maps identify the local translation actions by torus automorphisms.

3.3. Universal-cover noncollapsing. The next lemma is tailored to the torus-bundle structures that occur throughout the paper.

Lemma 3.2 (Universal-cover noncollapsing for torus bundles). *Fix $B \geq 1$. There is a number $i_B > 0$ with the following property. Let M be the total space of a $\mathbb{T}^2$ -bundle over S^1 , let h be an invariant metric for that bundle structure, and let g be another metric satisfying*

$$B^{-1}h \leq g \leq Bh.$$

If $|\sec_g| \leq \Lambda$ for some $\Lambda > 0$, then the universal cover satisfies

$$\inf_{\tilde{x} \in \tilde{M}} \mathrm{inj}_{\tilde{g}}(\tilde{x}) \geq i_B \Lambda^{-1/2}.$$

Proof. By scaling, it is enough to prove the result for $\Lambda = 1$. Lift the bundle and both metrics to the universal cover. The base becomes $\mathbb{R}$, the local torus translations become a free isometric $\mathbb{R}^2$ -action for $\tilde{h}$, and horizontal transport gives a global identification

$$\tilde{M} \cong \mathbb{R}^2 \times \mathbb{R}.$$

Use $\tilde{h}$ -arclength s on the base and horizontally transported fiber coordinates. Then

$$\tilde{h} = ds^2 + G_{ij}(s) dx^i dx^j,$$

with $G(s)$ positive definite. This formula is global; the monodromy affects only the identification after one period downstairs.

Consider $\tilde{x} = (0, 0)$, and for $r > 0$ define

$$\Omega_r = \{(x, s) : |s| < r, x^T G(s) x < r^2\}.$$

To reach a point of Ω_r , first move horizontally from $(0, 0)$ to $(0, s)$, then move along a straight segment in the $\mathbb{R}^2$ -orbit. Both pieces have length less than r , so $\Omega_r \subset B_{\tilde{h}}(\tilde{x}, 2r)$. For each fixed s , the fiber section is a radius- r disk for the metric $G(s)$, hence has area πr^2 . Therefore

$$\mathrm{vol}_{\tilde{h}}(\Omega_r) = 2\pi r^3.$$

The bilipschitz comparison gives $d_{\tilde{g}} \leq B^{1/2}d_{\tilde{h}}$ and the volume-measure inequality $d\mu_{\tilde{g}} \geq B^{-3/2}d\mu_{\tilde{h}}$. Taking $r = (2\sqrt{B})^{-1}$ yields

$$\begin{aligned} \text{vol}_{\tilde{g}} B_{\tilde{g}}(\tilde{x}, 1) &\geq B^{-3/2} \text{vol}_{\tilde{h}}(\Omega_r) \\ &= B^{-3/2} 2\pi(2\sqrt{B})^{-3} = \frac{\pi}{4B^3}. \end{aligned}$$

The same lower bound holds at every basepoint. Since $|\sec_{\tilde{g}}| \leq 1$, the Cheeger–Gromov–Taylor estimate [8, Section 4] gives

$$\inf_{\tilde{x} \in \tilde{M}} \text{inj}_{\tilde{g}}(\tilde{x}) \geq i_B > 0.$$

Scaling back gives the stated lower bound $i_B \Lambda^{-1/2}$. $\square$

3.4. Averaging and elementary collapse. We next establish the averaging estimate and the elementary consequence that small invariant fibers make the metric space close to its quotient circle.

For a tensor field U on M , write $\bar{U}$ for its average over the chosen local $\mathbb{T}^2$ -actions and write $U^\perp = U - \bar{U}$. Normalized Haar measure is preserved by every automorphism of $\mathbb{T}^2$, in particular by the $\text{SL}(2, \mathbb{Z})$ changes of fiber coordinates on overlaps. Hence the local averages agree on overlaps and define a global tensor. In particular, for a Riemannian metric g we write $g^\perp = g - \bar{g}$. If a tensor U is pulled back to the cyclic cover M_{cyc} obtained by unwrapping the base circle, we denote the pullback by U_{cyc} . Its further pullback to the universal cover is denoted by $\tilde{U}$.

Lemma 3.3 (Averaging in invariant Hölder norms). *Let h be an invariant metric for the fixed torus structure, and let $\mathcal{A}$ denote the corresponding averaging operator. For every integer $k \geq 0$, every $\sigma \in (0, 1)$, and every tensor field U ,*

$$(3.1) \quad \|\mathcal{A}U\|_{C^{k,\sigma}(M,h)} \leq \|U\|_{C^{k,\sigma}(M,h)}.$$

Averaging also commutes with constant metric scaling.

Proof. Pass to M_{cyc} , on which the fiber translations are the global maps R_c , $c \in \mathbb{T}^2$. Each R_c is an isometry of h_{cyc} . Pullback by R_c therefore commutes with the Levi-Civita connection, preserves pointwise tensor norms and spatial distances, and intertwines the parallel transports used in the intrinsic Hölder seminorms. Minkowski's inequality and integration against normalized Haar measure give (3.1) on the cyclic cover. The integral is invariant under the deck transformation, because the monodromy acts on $\mathbb{T}^2$ by a Haar-measure-preserving automorphism, and hence the estimate descends to M . Constant scaling changes neither the maps R_c nor normalized Haar measure. $\square$

For an invariant metric h , let

$$d_{\mathbb{T}^2}(h) = \sup_{y \in S^1} \text{diam}_h(F_y), \quad F_y = \pi^{-1}(y),$$

where the diameter is computed in the ambient metric space (M, h) . We also write $d_{\mathbb{T}^2}^{\text{int}}(h)$ for the corresponding supremum of the intrinsic diameters of the flat torus fibers with their induced metrics. Then $d_{\mathbb{T}^2}(h) \leq d_{\mathbb{T}^2}^{\text{int}}(h)$.

Lemma 3.4 (Collapse to the quotient circle). *Let $\pi : (M, h) \rightarrow (S^1, h_{S^1})$ be the metric quotient of an invariant torus-bundle metric, and suppose*

$$\sup_{y \in S^1} \text{diam}_h(\pi^{-1}(y)) \leq \delta,$$

where the fiber diameters are ambient diameters in (M, h) . Then, for all $x, x' \in M$,

$$(3.2) \quad 0 \leq d_h(x, x') - d_{h_{S^1}}(\pi(x), \pi(x')) \leq 2\delta.$$

Consequently, π has distortion at most 2δ and

$$d_{\text{GH}}((M, h), (S^1, h_{S^1})) \leq \delta.$$

Proof. The quotient map is distance nonincreasing, which gives the first inequality in (3.2). Given $\eta > 0$, choose $y \in \pi^{-1}(\pi(x))$ and $y' \in \pi^{-1}(\pi(x'))$ so that

$$d_h(y, y') \leq d_{h_{S^1}}(\pi(x), \pi(x')) + \eta.$$

The ambient fiber-diameter bound gives $d_h(x, y) \leq \delta$ and $d_h(x', y') \leq \delta$. The triangle inequality, followed by $\eta \downarrow 0$, proves the second inequality. Since π is onto, a surjective map of distortion at most 2δ gives the stated Gromov–Hausdorff bound. $\square$

4. FINITE-TIME DECAY OF THE NON-INVARIANT PART

This section develops the finite-time decay estimate behind the iteration. It shows that, on a fixed rescaled time interval, the non-invariant part of a sufficiently small Ricci–DeTurck perturbation of a locally $\mathbb{T}^2$ -invariant Ricci flow background contracts to any prescribed fraction of its initial size, provided the local torus orbits are short and the quotient circle does not collapse. The estimate is proved on lifted coordinate systems, so its constants are independent of collapse on the compact manifold and can be reused at every stage.

Here is the proof mechanism. Set

$$a_0 = \|g_{\text{RDT}}^\perp(1)\|_{C^{N,\sigma}(M, h(1))}.$$

The case $a_0 = 0$ follows at once from invariance and uniqueness. When $a_0 > 0$, pass to the cyclic cover and write $u = \pi_{\text{cyc}}^* g_{\text{RDT}}$. For every torus translation R_a , the normalized difference

$$y_a = a_0^{-1}(u - R_a^* u)$$

satisfies a homogeneous linear parabolic system with scalar principal part. A maximum principle argument gives a collapse-independent C^0 -bound for y_a . Lifting to the universal cover and applying interior Schauder estimates then gives a uniform C^1 -bound on the time interval $[10, 100]$. Averaging the translation differences recovers $a_0^{-1} g_{\text{RDT}}^\perp$. Since this tensor has zero orbit average, the short-orbit estimate makes its C^0 -norm small, proving the C^0

contraction. Higher derivative contraction follows from interpolation between the small C^0 -norm and higher derivative bounds.

The section is organized as follows. First, we define the lifted bounded-geometry conditions BG and PBG, give the parabolic averaging estimate, fix the Ricci–DeTurck equation and the intrinsic parabolic spaces X_s^m and Y_s^m , and then state the finite-time decay theorem. The first subsection, *Short-orbit averaging*, controls the geometry of the torus orbits, compares their ambient and intrinsic diameters, and proves that a zero-average tensor has small C^0 -norm when the orbits are short. The second subsection, *Translation differences and the Ricci–DeTurck system*, normalizes the differences between a solution and its torus translates, derives a linear tensor system with scalar principal part, obtains a uniform C^0 -bound on a compact associated bundle, and upgrades it by lifted Schauder estimates. The final subsection, *The lifted estimate and completion of the proof*, averages the translation differences to recover the non-invariant part and combines the short-orbit C^0 -estimate with parabolic interpolation to obtain any prescribed contraction on $[10, 100]$. The detailed collapse-independent Schauder estimates used in this argument are proved in Appendix A. *Bounded lifted geometry notation.* Let $m \geq 2$ and $\Lambda \geq 1$. For a metric s on a compact manifold, the notation

$$\text{BG}_m(s; \Lambda)$$

means that, on the universal cover,

$$(4.1) \quad \text{inj}_{\tilde{s}} \geq \Lambda^{-1}, \quad |\tilde{\nabla}^j \text{Rm}_{\tilde{s}}| \leq \Lambda \quad (0 \leq j \leq m+3).$$

In applications, the condition is imposed after the parabolic rescaling that makes the stage start time equal to one. For a Ricci flow $s(t)$, $a \leq t \leq b$, the notation

$$\text{PBG}_m(s; \Lambda; [a, b])$$

means that, on the universal cover,

$$(4.2) \quad \text{inj}_{\tilde{s}(t)} \geq \Lambda^{-1}, \quad |\tilde{\nabla}^j \text{Rm}_{\tilde{s}(t)}| \leq \Lambda \quad (0 \leq j \leq m+3), \quad a \leq t \leq b,$$

and that the fixed-coordinate, uniformly locally finite parabolic atlas constructed in Lemma A.1 may be chosen with chart radius at least Λ^{-1} and all coefficient, overlap, and cutoff bounds at most Λ . Enlarging Λ by a function of the curvature and injectivity bounds provides this atlas automatically. Thus BG encodes the finite-order geometry at one stage start, while PBG encodes the corresponding geometry and coordinate control on a whole parabolic interval. The tensor Hölder norms below are the intrinsic lifted norms associated with the reference metric, as specified in Paragraph A.1; in particular, they are defined for every smooth reference flow, without a PBG hypothesis. Whenever PBG_m holds, these intrinsic norms are uniformly equivalent to the fixed-coordinate uniformly local norms furnished by Lemma A.1. Thus the coordinate estimates used below have constants independent of collapse of the compact quotient.

Let $h(t)$ be an invariant Ricci flow. A family $g_{\text{RDT}}(t)$ satisfies the Ricci–DeTurck equation relative to $h(t)$ if

$$(4.3) \quad \frac{\partial}{\partial t} g_{\text{RDT}}(t) = -2 \text{Ric}(g_{\text{RDT}}(t)) + \mathcal{L}_{W(t)} g_{\text{RDT}}(t),$$

where

$$(4.4) \quad W^k(t) = g_{\text{RDT}}(t)^{ij} (\Gamma_{ij}^k(g_{\text{RDT}}(t)) - \Gamma_{ij}^k(h(t))).$$

The background flow $h(t)$ itself satisfies (4.3) relative to $h(t)$.

For a reference Ricci flow $s(t)$ on $[a, b]$, an integer $m \geq 2$, and $\sigma \in (0, 1)$, set

$$(4.5) \quad \begin{aligned} X_s^m[a, b] &= C^{m+\sigma, (m+\sigma)/2}(M \times [a, b], s), \\ Y_s^m[a, b] &= C^{m-2+\sigma, (m-2+\sigma)/2}(M \times [a, b], s), \end{aligned}$$

where these symbols denote the intrinsic lifted parabolic norms of Paragraph A.1.

Lemma 4.1 (Averaging in parabolic Hölder norms). *Let $h(t)$, $a \leq t \leq b$, be a smooth family of invariant metrics for the fixed torus structure, and let $\mathcal{A}$ be the corresponding averaging operator. For every integer $m \geq 2$ and every time-dependent tensor field U ,*

$$(4.6) \quad \|\mathcal{A}U\|_{X_h^m[a, b]} \leq \|U\|_{X_h^m[a, b]}, \quad \|\mathcal{A}U\|_{Y_h^m[a, b]} \leq \|U\|_{Y_h^m[a, b]}.$$

If $\text{PBG}_m(h; \Lambda; [a, b])$ holds, the analogous estimates in the fixed-chart uniformly local norms hold with a constant depending only on m, σ, Λ , and $b - a$; in those norms the constant need not equal one. Averaging commutes with constant metric scaling.

Proof. On M_{cyc} , the translations R_c are time-independent and satisfy $R_c^* h_{\text{cyc}}(t) = h_{\text{cyc}}(t)$. Hence their pullbacks commute with the spatial Levi-Civita connection and the metric-compatible time derivative $\mathcal{D}_t^{h_{\text{cyc}}}$, and preserve the parallel transports in the intrinsic Hölder seminorms of Paragraph A.1. Minkowski's inequality and integration against normalized Haar measure therefore give (4.6). Deck invariance lets the estimates descend to M . Under PBG_m , the fixed-chart estimates follow from Lemma A.2. The scaling assertion is immediate. $\square$

We will prove the following theorem in this section. For the Ricci–DeTurck metric appearing in its statement, set

$$g_{\text{RDT}}^\perp = g_{\text{RDT}} - \overline{g_{\text{RDT}}},$$

where the average is taken in the fixed local $\mathbb{T}^2$ -structure.

Informal content. On the normalized interval $[1, 100]$, sufficiently short torus orbits suppress the non-invariant part by any prescribed factor, uniformly in the amount of collapse of the compact quotient.

Theorem 4.2 (Finite-time decay of the non-invariant part). *Given $0 < \alpha < 1$, $K \geq 0$, $\lambda > 0$, an integer $N \geq 2$, and $\sigma \in (0, 1)$, there are constants $\beta_{\text{orb}} > 0$ and $\beta_{\text{pert}} > 0$ with the following property. Let $h(t)$ be an invariant Ricci-flow solution on M , defined for $t \in [1, 100]$, such that $|\sec_{h(t)}| \leq K$ and such that, for every $t \in [1, 100]$, the quotient circle*

for the chosen local $\mathbb{T}^2$ -structure has length at least λ , while every local $\mathbb{T}^2$ -orbit has $h(t)$ -diameter less than β_{orb} . Let $g_{\text{RDT}}(t)$ be a solution of the Ricci–DeTurck equation relative to $h(t)$, defined for $t \in [1, 100]$, and assume

$$\|g_{\text{RDT}} - h\|_{X_h^N[1,100]} < \beta_{\text{pert}}.$$

Then

$$(4.7) \quad \|g_{\text{RDT}}^\perp\|_{X_h^N[10,100]} \leq \alpha \|g_{\text{RDT}}^\perp(1)\|_{C^{N,\sigma}(M,h(1))}.$$

The orbit-diameter hypothesis refers to the same fixed local $\mathbb{T}^2$ -structure used in the averaging; Lemma 5.5 provides it in the geometric applications.

4.1. Short-orbit averaging. Once K is fixed, put

$$i_K = i_1 \max\{K, 1\}^{-1/2},$$

where i_1 is the constant from Lemma 3.2 for $B = 1$. In the applications below, Lemma 3.2, Shi’s estimates, and Lemma A.1 produce the lifted charts used in Appendix A. The cyclic cover used in the proof of the decay theorem serves only to make the local torus translations global.

The next fixed-time estimate says that a tensor with zero orbit average is controlled by its derivatives when the orbits are short.

Lemma 4.3 (Small-orbit averaging estimate). *Fix $K < \infty$ and $\lambda > 0$. There are constants $\beta_0 > 0$ and $C < \infty$, depending only on K and λ , with the following property. Let h be an invariant metric satisfying $|\sec_h| \leq K$, whose quotient circle has length at least λ , and whose local $\mathbb{T}^2$ -orbits have diameter at most $d \leq \beta_0$. If U is a smooth symmetric two-tensor with zero average over the local $\mathbb{T}^2$ -orbits, then*

$$(4.8) \quad \|U\|_{C^0(M,h)} \leq Cd \|U\|_{C^1(M,h)}.$$

Proof. Step 1. Control the variation of the orbit metrics.

Unwrap the base circle, obtaining $\pi_{\text{cyc}} : M_{\text{cyc}} \rightarrow M$. The local torus translations become a global isometric $\mathbb{T}^2$ -action R_a , $a \in \mathbb{T}^2$. In a global horizontal gauge, with arclength s on the lifted base,

$$h_{\text{cyc}} = ds^2 + G_{ij}(s) dx^i dx^j.$$

The orbit shape operator is $A = \frac{1}{2}G^{-1}G_s$. Along a horizontal geodesic it satisfies

$$A_s + A^2 + R_{\partial_s} = 0, \quad R_{\partial_s}(V) = R(V, \partial_s)\partial_s.$$

Let $\lambda_{\min}$ be the smallest eigenvalue of A . At points of differentiability, and in the barrier sense elsewhere,

$$\lambda'_{\min} \leq K - \lambda_{\min}^2.$$

If $\lambda_{\min}(s_0) < -\sqrt{K}$, comparison with $y' = K - y^2$ forces blow-up to $-\infty$ in finite forward time. That is impossible because the vertical Jacobi tensor is nonsingular, equivalently

because $G(s)$ stays smooth and positive definite for all $s \in \mathbb{R}$. Thus $\lambda_{\min} \geq -\sqrt{K}$. Reversing s replaces A by $-A$, so $\lambda_{\max} \leq \sqrt{K}$. Hence

$$(4.9) \quad \|A(s)\| \leq C_K \quad \text{for every } s \in \mathbb{R}.$$

Consequently, horizontal transport through base distance r changes vertical lengths by at most the factor $e^{C_K r}$.

Step 2. Compare intrinsic and ambient orbit diameter.

Decrease β_0 so that $\beta_0 < \lambda/10$. Since $d \leq \beta_0$, we then have $d < \lambda/10$. Take two points in one orbit and choose $0 < \varepsilon < \lambda - d$. Join them by an ambient path of length at most $d + \varepsilon$. Lift the path to the cyclic cover after choosing a lift of its initial point. Its endpoint must lie in the same lifted orbit as its initial point. Indeed, an endpoint in a nontrivial deck translate of that orbit would make the projected path wind nontrivially around the quotient circle, whose shortest noncontractible loop has length at least λ ; this contradicts $d + \varepsilon < \lambda$. The base coordinate of the lifted path therefore stays within $d + \varepsilon$ of that lifted orbit. Project the path horizontally back to this orbit. By (4.9), the projected path has length at most $e^{C_K(d+\varepsilon)}(d + \varepsilon)$. Letting $\varepsilon \downarrow 0$ gives

$$(4.10) \quad d_{\mathbb{T}^2}^{\text{int}}(h) \leq C_K d.$$

Step 3. Control the differential of a translation.

Fix $a \in \mathbb{T}^2$ and $x \in M_{\text{cyc}}$, and join x to $R_a x$ by an intrinsic minimizing geodesic in the flat orbit. Its length is at most $C_K d$. Translation-invariant orthonormal vertical frames are parallel for the intrinsic flat metric. Their ambient covariant derivatives and the derivative of the invariant horizontal unit normal are bounded by the second fundamental form, hence by (4.9). Let E_1, E_2 be an invariant orthonormal vertical frame, and let N be the invariant horizontal unit normal along the orbit geodesic. The differential dR_a carries the frame (E_1, E_2, N) at x to the same invariant frame at $R_a x$. Integrating the bounds for $\nabla_{\dot{\gamma}} E_i$ and $\nabla_{\dot{\gamma}} N$ along the geodesic and comparing with ambient parallel transport $\mathcal{P}_{x, R_a x}$ gives

$$(4.11) \quad \|dR_a - \mathcal{P}_{x, R_a x}\| \leq C_K d.$$

Here and below, U denotes the pullback to M_{cyc} . Comparing tensor values first by parallel transport and then by (4.11), we obtain

$$|U(x) - R_a^* U(x)| \leq C_K d (\|\nabla U\|_{C^0} + \|U\|_{C^0}).$$

Step 4. Average.

Because the orbit average of U is zero,

$$U = \int_{\mathbb{T}^2} (U - R_a^* U) da.$$

Integrating the previous estimate with normalized Haar measure and taking the supremum gives

$$\|U\|_{C^0} \leq C_K d (\|\nabla U\|_{C^0} + \|U\|_{C^0}).$$

Decrease β_0 once more so that $C_K\beta_0 \leq 1/2$. Since $d \leq \beta_0$, we have $C_K d \leq 1/2$, and the last term can be absorbed. This is (4.8). $\square$

4.2. Translation differences and the Ricci–DeTurck system. We give a preview of the rest of the argument. The two covers in the proof have distinct roles:

$$\widetilde{M} \xrightarrow{\text{universal cover}} M_{\text{cyc}} \xrightarrow{\text{cyclic cover}} M.$$

- M_{cyc} Global torus translations and the compact associated-bundle maximum-principle argument.
- $\widetilde{M}$ Bounded geometry, fixed-radius harmonic-coordinate cylinders, and all parabolic Schauder estimates.

Outline of the proof. The following table summarizes the translation-difference argument and the final proof of Theorem 4.2.

Step	Hypothesis or operation	Conclusion
1	Compare u with each translate R_a^*u .	A homogeneous tensor system with scalar principal part.
2	Apply the maximum principle on the compact associated bundle.	A collapse-independent C^0 -bound for the normalized differences.
3	Lift to $\widetilde{M}$ and regularize the coefficients.	Uniform finite-order coefficient bounds in fixed-radius charts.
4	Apply the lifted Schauder estimate.	A uniform $C^{N+1+\sigma, (N+1+\sigma)/2}$ -bound on $[9, 100]$.
5	Average over translations and interpolate on $[10, 100]$.	The prescribed contraction of the non-invariant part.

The proposition below proves Steps 1–4; Step 5 is the concluding proof of the decay theorem.

Proposition 4.4 (Uniform lifted estimate for translation differences). *Under the hypotheses of Theorem 4.2, put*

$$a_0 = \|g_{\text{RDT}}^\perp(1)\|_{C^{N,\sigma}(M,h(1))},$$

and assume that $a_0 > 0$. Let $\pi_{\text{cyc}} : M_{\text{cyc}} \rightarrow M$ be the cyclic cover and let R_a , $a \in \mathbb{T}^2$, denote its global torus action. Write h_{cyc} and $(g_{\text{RDT}})_{\text{cyc}}$ for the lifted metrics, and set

$$u = (g_{\text{RDT}})_{\text{cyc}}, \quad u_a = R_a^*(g_{\text{RDT}})_{\text{cyc}}, \quad z_a = u - u_a, \quad y_a = a_0^{-1}z_a.$$

If tildes denote the further lifts to the universal cover, then

$$\sup_{a \in \mathbb{T}^2} \|\widetilde{y}_a\|_{C_{\text{uloc}}^{N+1+\sigma, (N+1+\sigma)/2}(\widetilde{M} \times [9, 100], \widetilde{h})} \leq C.$$

Here the subscript uloc denotes the uniformly local fixed-chart norm defined in Paragraph A.1. The constant C depends only on the fixed data in Theorem 4.2, and not on the collapsed quotient.

Proof. Step 1. Apply the general Ricci–DeTurck difference lemma.

Let

$$u_s = u_a + s(u - u_a), \quad 0 \leq s \leq 1.$$

Choose β_{pert} below a constant depending only on the fixed data so that each u_s is a metric and $2^{-1}h_{\text{cyc}} \leq u_s \leq 2h_{\text{cyc}}$. Apply Lemma A.4 to u_a and u , with background h_{cyc} , and denote its coefficients by

$$P_a^{pq} = \int_0^1 u_s^{pq} ds, \quad B_a^p = \int_0^1 \mathcal{B}_s^p ds, \quad C_a = \int_0^1 \mathcal{C}_s ds.$$

Dividing the resulting equation for z_a by a_0 gives the homogeneous local system

$$(4.12) \quad \partial_t y_a - P_a^{pq} \nabla_p \nabla_q y_a = B_a^p \nabla_p y_a + C_a y_a.$$

The second-order part acts separately on every tensor component by the same operator $P_a^{pq} \nabla_p \nabla_q$; B_a^p and C_a are endomorphism-valued and contain all coupling between components. Uniformly in a ,

$$c h_{\text{cyc}}^{pq} \xi_p \xi_q \leq P_a^{pq} \xi_p \xi_q \leq C h_{\text{cyc}}^{pq} \xi_p \xi_q, \quad |P_a - h_{\text{cyc}}^{-1}| \leq C \beta_{\text{pert}}.$$

Since $N \geq 2$, (A.17) also gives, on the full cylinder,

$$(4.13) \quad |B_a| \leq C \beta_{\text{pert}}, \quad |C_a| \leq C(K + \beta_{\text{pert}}).$$

No derivatives of $\text{Rm}(h_{\text{cyc}})$ are needed for these bounds. Indeed, writing $u_s = h_{\text{cyc}} + e_s$, we have $\nabla^{h_{\text{cyc}}} u_s = \nabla^{h_{\text{cyc}}} e_s$ and $(\nabla^{h_{\text{cyc}}})^2 u_s = (\nabla^{h_{\text{cyc}}})^2 e_s$, because $\nabla^{h_{\text{cyc}}} h_{\text{cyc}} = 0$. Because $R_a^* h_{\text{cyc}} = h_{\text{cyc}}$, the tensor e_s inherits the assumed X_h^N -bound uniformly in a and s . Thus the derivative terms in B_a and C_a are controlled by that bound, while the remaining background term is algebraic in $\text{Rm}(h_{\text{cyc}})$ and is controlled by K .

The cyclic cover may be collapsed, and no bounded-geometry chart on M_{cyc} is used at this stage. Equation (4.12), the pointwise ellipticity bound, and (4.13) are the only coefficient information needed before we lift the system to the universal cover in Step 3. All higher coefficient estimates will be proved there.

At $t = 1$, the invariant parts cancel:

$$z_a(1) = (g_{\text{RDT}})_{\text{cyc}}^\perp(1) - R_a^*(g_{\text{RDT}})_{\text{cyc}}^\perp(1).$$

Therefore

$$(4.14) \quad \|y_a(1)\|_{C^{N,\sigma}(M_{\text{cyc}}, h_{\text{cyc}}(1))} \leq 2.$$

By the norm convention stated immediately after the parabolic Hölder norm definition in Subsection A.1, the norm in (4.14) is computed after lifting to the universal cover; it does not presuppose a fixed-radius atlas on M_{cyc} .

Step 2. Obtain a uniform C^0 -bound.

Let $\phi_a = |y_a|_{h_{\text{cyc}}(t)}^2$. Since $\partial_t h_{\text{cyc}} = -2 \text{Ric}(h_{\text{cyc}})$, differentiating the norm and using (4.12) gives

$$\begin{aligned} (\partial_t - P_a^{pq} \nabla_p \nabla_q) \phi_a &\leq -c |\nabla y_a|_{h_{\text{cyc}}}^2 + C |y_a|_{h_{\text{cyc}}} |\nabla y_a|_{h_{\text{cyc}}} + C |y_a|_{h_{\text{cyc}}}^2 \\ &\leq C \phi_a. \end{aligned}$$

The cyclic cover is noncompact, but the maximum principle can be applied on a compact associated bundle. Let γ be its deck generator. If H is the monodromy, then $\gamma R_a \gamma^{-1} = R_{Ha}$, and the lifted metrics are deck-invariant. Hence

$$\gamma^* y_{Ha} = y_a, \quad \phi_{Ha}(\gamma x_{\text{cyc}}, t) = \phi_a(x_{\text{cyc}}, t).$$

Thus $(x_{\text{cyc}}, a) \mapsto \phi_a(x_{\text{cyc}}, t)$ descends to

$$\mathcal{Z} = (M_{\text{cyc}} \times \mathbb{T}^2) / ((x_{\text{cyc}}, a) \sim (\gamma x_{\text{cyc}}, Ha)),$$

which is compact. It is the $\mathbb{T}^2$ -bundle over $M_{\text{cyc}} / \langle \gamma \rangle = M$ associated to the monodromy action $a \mapsto Ha$. At a spatial maximum on $\mathcal{Z}$, with a fixed, the M_{cyc} -gradient vanishes and the P_a -trace of the M_{cyc} -Hessian is nonpositive. Applying the scalar maximum principle to $e^{-Ct} \phi$ and using (4.14) gives

$$(4.15) \quad \sup_{a \in \mathbb{T}^2} \|y_a\|_{C^0(M_{\text{cyc}} \times [1, 100], h_{\text{cyc}})} \leq C.$$

Step 3. Lift to the universal cover and prove the finite-order Schauder estimate.

The goal is the uniform estimate (4.20) for the translation differences on $\widetilde{M} \times [9, 100]$. There are two parts. First we obtain exactly the finite amount of interior regularity of $\widetilde{u}$ needed to control the coefficients of the equation for $\widetilde{y}_a$. We then estimate that linear tensor system through order $N + 1 + \sigma$.

The scalar estimates and their reduction to finite tensor systems with scalar principal part are recorded in Remark A.5 and Proposition A.7. We apply that proposition after obtaining the finite coefficient regularity required for the translation-difference equation.

Let

$$\pi_\infty : \widetilde{M} \longrightarrow M_{\text{cyc}}$$

be the universal covering map. We decorate all lifted quantities by tildes. In particular, $\widetilde{y}_a$ satisfies

$$(4.16) \quad \partial_t \widetilde{y}_a - \widetilde{P}_a^{pq} \widetilde{\nabla}_p \widetilde{\nabla}_q \widetilde{y}_a = \widetilde{B}_a^p \widetilde{\nabla}_p \widetilde{y}_a + \widetilde{C}_a \widetilde{y}_a$$

on $\widetilde{M} \times [1, 100]$, where $\widetilde{\nabla}$ is the connection of $\widetilde{h}(t)$. The C^0 -estimate (4.15) lifts unchanged to this equation.

Lemma 3.2, with $B = 1$ and $\Lambda = \max\{K, 1\}$, gives a time-uniform injectivity-radius lower bound for $\widetilde{h}(t)$. No curvature-derivative bound at the initial time is assumed here. Since

$\tilde{h}(t)$ is a complete Ricci flow with uniformly bounded curvature on $[1, 100]$, Shi's derivative estimates [21], applied with the fixed positive time gap from 1 to 7, give

$$\sup_{\tilde{M} \times [7, 100]} |\tilde{\nabla}^j \text{Rm}_{\tilde{h}}| \leq C_j(K, N), \quad 0 \leq j \leq N + 4.$$

Together with the preceding injectivity-radius estimate and Lemma A.1, these bounds give, for a constant Λ_h depending only on the fixed data,

$$\text{PBG}_{N+1}(h; \Lambda_h; [7, 100]).$$

By Lemma A.2, the assumed intrinsic $X_h^N[1, 100]$ -bound controls, with a constant depending only on the fixed data, the order- N fixed-chart norm on $[7, 100]$ used below. Let r_0 be the corresponding atlas radius from Lemma A.1, reduced so that $(4r_0)^2 \leq 1$. For a center $(\tilde{x}_0, t_0) \in \tilde{M} \times [8, 100]$, write

$$Q_\rho(\tilde{x}_0, t_0) = B_{\tilde{h}(t_0)}(\tilde{x}_0, \rho) \times [t_0 - \rho^2, t_0]$$

for the corresponding one-sided backward cylinder in the fixed spatial coordinates of that lemma. The choice $(4r_0)^2 \leq 1$ ensures that every $Q_{4r_0}(\tilde{x}_0, t_0)$ with $t_0 \in [8, 100]$ is contained in $\tilde{M} \times [7, 100]$.

Claim 1 (finite coefficient regularity). Uniformly in a , the coefficients of (4.16) satisfy

$$(4.17) \quad \begin{aligned} & \|\tilde{P}_a\|_{C^{N+1+\sigma, (N+1+\sigma)/2}(\tilde{M} \times [8, 100], \tilde{h})} \\ & + \|\tilde{B}_a\|_{C^{N+\sigma, (N+\sigma)/2}(\tilde{M} \times [8, 100], \tilde{h})} \\ & + \|\tilde{C}_a\|_{C^{N-1+\sigma, (N-1+\sigma)/2}(\tilde{M} \times [8, 100], \tilde{h})} \leq C. \end{aligned}$$

To prove the claim, use a fixed local frame for symmetric two-tensors in one of the preceding charts. The lifted Ricci–DeTurck equation has the form

$$(4.18) \quad \partial_t \tilde{u}^A - \tilde{u}^{pq} \partial_p \partial_q \tilde{u}^A = F^A(x, t, \tilde{u}, \tilde{u}^{-1}, \partial \tilde{u}; \tilde{h}, \tilde{h}^{-1}, \partial \tilde{h}, \partial^2 \tilde{h}),$$

where the right-hand side contains no second derivatives of $\tilde{u}$. The same equation holds for $\tilde{u}_a$. The assumed $X_h^N[1, 100]$ -closeness, lifted to the universal cover, gives a uniform $C^{N+\sigma, (N+\sigma)/2}$ -bound for both metrics and makes their principal matrices uniformly elliptic. The PBG_{N+1} -bound gives the required $C^{N+1+\sigma, (N+1+\sigma)/2}$ -control of the reference metric in the fixed coordinates. The same atlas, supplied by Lemma A.1, bounds $\|\tilde{h}^{-1}\|_{C^0}$ uniformly there. Thus the reference nondegeneracy hypothesis in the derivative-gain lemma is also satisfied.

Apply Lemma A.9 with $k = N$, outer cylinder Q_{4r_0} , and inner cylinder Q_{2r_0} . It gives

$$(4.19) \quad \|\tilde{u}\|_{C^{N+1+\sigma, (N+1+\sigma)/2}(Q_{2r_0})} \leq C.$$

The same bound holds for $\tilde{u}_a$, with the same constant, because $R_a^* h_{\text{cyc}} = h_{\text{cyc}}$. The derivative count in that lemma is exactly the one needed here: differentiating $N-1$ times uses $\tilde{u}$ through order N and $\tilde{h}$ through order $N+1$. The formulas (A.17) now give (4.17): P_a is algebraic in the metrics and their inverses, B_a uses one derivative, and C_a uses at most two derivatives.

The norms on the noncompact universal cover are the suprema of the corresponding local norms in this fixed-radius atlas.

Claim 2 (translation-difference estimate). By Claim 1, the lifted system (4.16) satisfies the hypotheses of Proposition A.7 with $m = N + 1$ on $[8, 100]$, uniformly in a . Its forcing term is zero. The uniformly local interior estimate (A.22), with one unit of room at the initial side, and the lift of the C^0 -bound (4.15) therefore give

$$(4.20) \quad \sup_{a \in \mathbb{T}^2} \|\tilde{g}_a\|_{C_{\text{uloc}}^{N+1+\sigma, (N+1+\sigma)/2}(\widetilde{M} \times [9, 100], \tilde{h})} \leq C.$$

The estimate is one-sided backward at the top time and is uniform in the translation parameter. This is the required universal-cover Hölder bound. $\square$

4.3. The lifted estimate and completion of the proof.

Proof of Theorem 4.2. Set

$$a_0 = \|g_{\text{RDT}}^\perp(1)\|_{C^{N, \sigma}(M, h(1))}.$$

If $a_0 = 0$, the initial metric is invariant. On the cyclic cover that unwraps the base, the local translations form a global $\mathbb{T}^2$ -action, and the Ricci–DeTurck equation with invariant background is equivariant under that action. Uniqueness, or equivalently the zero-data maximum principle used below, preserves the symmetry. Hence $g_{\text{RDT}}^\perp \equiv 0$, and the estimate is immediate. Assume from now on that $a_0 > 0$. Use the notation of Proposition 4.4. That proposition gives the uniform estimate (4.20).

Averaging the translation differences recovers the non-invariant part:

$$\begin{aligned} a_0^{-1}(g_{\text{RDT}})_{\text{cyc}}^\perp &= a_0^{-1} \left((g_{\text{RDT}})_{\text{cyc}} - \int_{\mathbb{T}^2} R_a^*(g_{\text{RDT}})_{\text{cyc}} da \right) \\ &= \int_{\mathbb{T}^2} y_a da. \end{aligned}$$

Because γ conjugates R_a to R_{Ha} , and the automorphism $a \mapsto Ha$ preserves normalized Haar measure, this integral is deck-invariant. It descends to

$$w = a_0^{-1} g_{\text{RDT}}^\perp.$$

After lifting this identity to $\widetilde{M}$, we have

$$\tilde{w} = \int_{\mathbb{T}^2} \tilde{y}_a da.$$

By the same Minkowski argument as in Lemma 4.1, integration is bounded in each local Hölder norm of the fixed universal-cover atlas. Hence (4.20) gives

$$(4.21) \quad \|w\|_{X_h^{N+1}[9, 100]} \leq C.$$

The tensor w has zero orbit average.

Let

$$d = \sup_{t \in [1, 100]} d_{\mathbb{T}^2}(h(t)) \leq \beta_{\text{orb}}.$$

Applying Lemma 4.3 to $w(t)$ for each $t \in [9, 100]$ and then using (4.21) gives

$$(4.22) \quad \|w\|_{C^0(M \times [9, 100], h)} \leq Cd \leq C\beta_{\text{orb}}.$$

Apply the parabolic interpolation inequality [16, Proposition 4.2, p. 49] in the fixed-radius universal-cover charts, take the uniformly local supremum, and then use the intrinsic/fixed-chart norm equivalence. With one unit of room at the initial side and backward cylinders at the final side, this gives some $\mu = \mu(N, \sigma) \in (0, 1)$ such that

$$(4.23) \quad \begin{aligned} \|w\|_{X_h^N[10, 100]} &\leq C \|w\|_{C^0(M \times [9, 100], h)}^\mu \|w\|_{X_h^{N+1}[9, 100]}^{1-\mu} \\ &\leq C\beta_{\text{orb}}^\mu. \end{aligned}$$

First choose β_{pert} below all perturbative thresholds used in Proposition 4.4. With that choice fixed, choose β_{orb} below the threshold in Lemma 4.3 and so that $C\beta_{\text{orb}}^\mu \leq \alpha$. Since $g_{\text{RDT}}^\perp = a_0 w$, this is precisely (4.7). $\square$

Theorem 4.2 is the analytic contraction estimate used below: once a fixed fibration has uniformly short fibers and its quotient circle has a uniform positive lower bound, the non-invariant part can be made to decay by any prescribed factor on the latter part of a finite rescaled stage.

5. COMPARISON OF TORUS FIBRATIONS

The preceding decay estimate for the non-invariant part is predicated on the assumption that the fibers of the background invariant metric are short. In order to apply this in the subsequent iterative process, we will need to know that this property is preserved. We do this by comparing the fixed torus fibration retained by the iteration with the possibly different fibration supplied by Bamler at each time. This section proves the geometric dichotomy that achieves the comparison: equivalent short-fiber fibrations transfer smallness to one another, whereas two inequivalent short-fiber fibrations force the entire manifold to have small diameter. On a large-diameter block, the second alternative is excluded, so the retained fibration persists and its quotient circle remains quantitatively nondegenerate.

The section is organized as follows. The first subsection, *Fiber equivalence and diameter*, bounds the total diameter in terms of a quotient-circle length and a fiber diameter, and defines fiber equivalence through the integral cohomology class, or equivalently the common kernel of the induced maps to $\mathbb{Z}$. The second subsection, *Fiber-equivalent small-fiber fibrations*, uses that common kernel to show that horizontal projection of a short comparison fiber onto a fixed fiber has degree ± 1 , thereby transferring the small-fiber estimate to the fixed fibration. The final subsection, *Fiber-inequivalent fibrations and fixed-fiber improvement*, first proves that two inequivalent short-fiber fibrations force the whole manifold to have small diameter, and then combines the two alternatives into a pointwise bootstrap: at definite total diameter, the comparison fibration is fiber-equivalent to the retained one, the fixed fibers become strictly shorter, and the fixed quotient-circle length acquires a positive lower bound. These conclusions are later inserted into a first-hitting argument on each large block.

5.1. Fiber equivalence and diameter. We begin with the fibration-compatibility facts used later in the DeTurck iteration. These statements are purely topological and metric; they do not use the Ricci flow. Throughout this section, when a torus fibration $\pi : M \rightarrow S^1$ has been fixed, a *fixed fiber* means a fiber of this chosen map π , as opposed to a different time-slice fibration. If r is a locally $\mathbb{T}^2$ -invariant metric for this fixed fibration, its *fixed quotient length* means the length of the metric quotient circle of π with respect to r .

Lemma 5.1 (Base length and fiber diameter control diameter). *Let h be an invariant metric for a torus fibration $\pi : M \rightarrow S^1$, with quotient metric h_{S^1} . Let g be another metric satisfying $g \leq Bh$; no invariance is assumed for g . If every π -fiber has g -diameter at most d_f , then*

$$\text{diam}(M, g) \leq B^{1/2}L(S^1, h_{S^1})/2 + d_f.$$

In particular, if $d_{\mathbb{T}^2}(h) \leq d$, then

$$\text{diam}(M, g) \leq B^{1/2}(L(S^1, h_{S^1})/2 + d).$$

Proof. Take two points $x, x' \in M$. Join their projections by a shortest base arc and lift that arc horizontally for h , starting at x . Because $g \leq Bh$, the lift has g -length at most $B^{1/2}L(S^1, h_{S^1})/2$. Its endpoint z lies in the fiber containing x' , so the ambient fiber-diameter bound gives $d_g(z, x') \leq d_f$. The triangle inequality therefore gives

$$d_g(x, x') \leq d_g(x, z) + d_g(z, x') \leq B^{1/2}L(S^1, h_{S^1})/2 + d_f.$$

The path used to estimate $d_g(z, x')$ need not stay in the fiber. Taking the supremum over x, x' proves the first estimate. If every fiber has h -diameter at most d , then it has g -diameter at most $B^{1/2}d$; substituting this into the first estimate gives the last one. $\square$

We next fix the convention for comparing torus fibrations and prove the two compatibility lemmas used in the fixed-fiber bootstrap. The first says that a fiber-equivalent small-fiber fibration makes the fixed fibers small. The second says that a fiber-inequivalent small-fiber fibration forces the whole diameter to be small.

We shall compare fibrations by their induced cohomology classes rather than by identifying their target circles.

Definition 5.2 (Fiber-equivalent torus fibrations). Let

$$\pi_i : M \rightarrow S_i^1, \quad i = 0, 1,$$

be torus fibrations. The target circles are not part of the fixed data and may be different. Choose a basepoint $x \in M$ and, for each $i = 0, 1$, an identification $\pi_1(S_i^1, \pi_i(x)) \cong \mathbb{Z}$, and let

$$\alpha_{\pi_i, x} : \pi_1(M, x) \longrightarrow \mathbb{Z},$$

be the surjective homomorphism induced by π_i . Changing the identification changes $\alpha_{\pi_i, x}$ by a sign. Equivalently, π_i determines a cohomology class $[\pi_i] \in H^1(M; \mathbb{Z})$, well-defined up to sign. We say that π_0 and π_1 are *fiber-equivalent* if

$$[\pi_0] = \pm[\pi_1] \in H^1(M; \mathbb{Z}).$$

Equivalently, after choosing a common basepoint,

$$\ker \alpha_{\pi_0, x} = \ker \alpha_{\pi_1, x} \subset \pi_1(M, x).$$

For a different basepoint, the two kernels are changed by the same conjugacy convention. Thus, in fiber-subgroup language, the definition says that the images of the two fiber fundamental groups agree, up to this usual conjugacy ambiguity.

If a bundle equivalence is given by a diffeomorphism of M isotopic to the identity and a diffeomorphism of the target circles, then the fibrations are fiber-equivalent; the sign records the orientation of the base map. An arbitrary self-diffeomorphism of M need not preserve the cohomology class. Conversely, fiber-equivalence specifies only the underlying cohomology class, or equivalently the common fiber kernel, and not a chosen bundle equivalence. Because maps to S^1 classify integral cohomology classes in degree one, the relation

$$[\pi_0] = \pm[\pi_1] \in H^1(M; \mathbb{Z}),$$

means that, after choosing a degree ± 1 map $\varphi : S_0^1 \rightarrow S_1^1$, the two fibrations represent the same homotopy class of maps from M to a circle. In other words, there is a homotopy-commutative diagram of spaces

$$\begin{array}{ccc} M & \xrightarrow{\pi_0} & S_0^1 \\ \text{Id}_M \downarrow & & \downarrow \varphi \\ M & \xrightarrow{\pi_1} & S_1^1. \end{array}$$

Passing to fundamental groups gives the agreement of the kernels of the two maps to $\mathbb{Z}$, and that common kernel is the only topological input used in the next lemma.

5.2. Fiber-equivalent small-fiber fibrations.

Lemma 5.3 (Fiber-equivalent small fibers force small fixed fibers). *Fix $\lambda > 0$. For every $\theta > 0$ there are $\eta > 0$ and $\gamma > 0$ with the following property. Let r be an invariant metric for a torus fibration $\pi : M \rightarrow S^1$, and suppose that the quotient circle has length at least λ . Let q be an invariant metric for a possibly different torus fibration $\pi_B : M \rightarrow S_B^1$. Assume that*

$$\|q - r\|_{C^1(M, r)} < \gamma,$$

that π_B is fiber-equivalent to π , and that every π_B -fiber has q -diameter at most η . Then

$$d_{\mathbb{T}^2}(r) < \theta.$$

Proof. Step 1. Place a comparison fiber over a short base arc.

See Figure 1. Choose η and γ so small that every π_B -fiber has r -diameter less than 2η and

$$4\eta < \min\{\lambda/10, \theta/10\}.$$

This is possible because q and r are bilipschitz when γ is small. Since $\pi : (M, r) \rightarrow S^1$ is distance nonincreasing, the image of a connected π_B -fiber has diameter less than 2η . The

quotient circle has length at least λ , so that image lies in a base arc of length less than $4\eta < \min\{\lambda/10, \theta/10\}$.

Fix a fiber $F = \pi^{-1}(y_0)$, choose $x \in F$, and let $F_B = \pi_B^{-1}(\pi_B(x))$. Choose an arc $J \subset S^1$ containing $\pi(F_B)$, with

$$\text{length}(J) < \min\{\lambda/10, \theta/10\}.$$

Because $x \in F_B$, the point $y_0 = \pi(x)$ lies in J . Horizontal transport over J identifies $\pi^{-1}(J)$ with $F \times J$. Let $P : \pi^{-1}(J) \rightarrow F$ be the resulting projection; it fixes x .

Step 2. Prove that $P|_{F_B}$ is onto.

Use x as basepoint. Since π_B and π are fiber-equivalent, their kernels in $\pi_1(M, x)$ agree by Definition 5.2. The inclusions of F and F_B identify their fundamental groups with this common kernel. Because J is an arc, P is the endpoint of a deformation retraction of $\pi^{-1}(J)$ onto F that fixes x . Therefore, with $i_F : F \hookrightarrow M$ and $i_B : F_B \hookrightarrow M$,

$$(i_F)_* \circ (P|_{F_B})_* = (i_B)_* : \pi_1(F_B, x) \longrightarrow \pi_1(M, x).$$

The map $(i_F)_*$ is injective, and the two inclusion maps have the same image. Hence $(P|_{F_B})_*$ is an isomorphism. Thus $P|_{F_B} : F_B \rightarrow F$ has degree ± 1 , so it is onto.

Step 3. Bound the fixed fiber.

Given $p, q \in F$, choose $p_B, q_B \in F_B$ with $P(p_B) = p$ and $P(q_B) = q$. The horizontal segments from p_B, q_B to p, q each have r -length at most $\text{length}(J)$, while $d_r(p_B, q_B) \leq 2\eta$. Hence

$$d_r(p, q) \leq 2 \text{length}(J) + 2\eta.$$

After decreasing η and γ , the right side is less than θ . Taking the supremum over p, q and then over all fibers F proves the lemma. $\square$

The complementary lemma says that two fiber-inequivalent small-fiber fibrations can occur only when the whole manifold has small diameter.

5.3. Fiber-inequivalent fibrations and fixed-fiber improvement.

Lemma 5.4 (Fiber-inequivalent small fibers force small diameter). *Fix $B \geq 1$. There is a constant $C_{\text{fib}} = C_{\text{fib}}(B)$ with the following property. Let h be a metric on a closed three-manifold M . Suppose that M has two torus fibrations*

$$\pi_i : M \rightarrow S_i^1, \quad i = 0, 1,$$

and invariant comparison metrics h_i , with quotient metrics $h_{S^1, i}$, such that

$$B^{-1}h_i \leq h \leq Bh_i, \quad i = 0, 1.$$

Suppose also that, for some $d_f > 0$, every fiber of each π_i has h -diameter at most d_f . If π_0 and π_1 are fiber-inequivalent, then

$$\text{diam}(M, h) \leq C_{\text{fib}} d_f.$$




FIGURE 1. Fiber-equivalent short fibers. The comparison fiber lies over the short fixed-base arc J . Horizontal projection in $\pi^{-1}(J) \cong F \times J$ maps it onto the fixed fiber with degree ± 1 , transferring the comparison-fiber and base-arc bounds to the target estimate $\text{diam}_r(F) < \theta$.

Proof. Fix $x \in M$, choose representatives of the two induced maps to $\mathbb{Z}$, and set $K_i = \ker \alpha_{\pi_i, x} \subset \pi_1(M, x)$. Let F_1 be the π_1 -fiber through x , with inclusion $i_1 : F_1 \hookrightarrow M$. Suppose first that

$$\alpha_{\pi_0, x} \circ (i_1)_* : \pi_1(F_1, x) \longrightarrow \mathbb{Z},$$

were zero. Then $K_1 = (i_1)_* \pi_1(F_1, x) \subset K_0$, and this inclusion would induce a surjection

$$\pi_1(M, x)/K_1 \cong \mathbb{Z} \longrightarrow \pi_1(M, x)/K_0 \cong \mathbb{Z}.$$

After identifying the two quotient groups with $\mathbb{Z}$, this is a surjective homomorphism $\mathbb{Z} \rightarrow \mathbb{Z}$, and hence an isomorphism. Since its kernel is K_0/K_1 , we would have $K_0 = K_1$, contrary to the assumption that the fibrations are fiber-inequivalent. Therefore $\alpha_{\pi_0, x} \circ (i_1)_* \neq 0$. Equivalently, $\pi_0|_{F_1} : F_1 \rightarrow S_0^1$ induces a nonzero map on H_1 , and is thus onto.

The quotient map $\pi_0 : (M, h) \rightarrow (S_0^1, h_{S^1,0})$ is $B^{1/2}$ -Lipschitz: it is distance nonincreasing for h_0 , and $h_0 \leq Bh$. Since $\pi_0(F_1) = S_0^1$,

$$\text{diam}(S_0^1, h_{S^1,0}) \leq B^{1/2} \text{diam}_h(F_1) \leq B^{1/2} d_f.$$

Hence $L(S_0^1, h_{S^1,0}) \leq 2B^{1/2} d_f$. Applying Lemma 5.1 to π_0 gives

$$\text{diam}(M, h) \leq B^{1/2} L(S_0^1, h_{S^1,0})/2 + d_f \leq (B+1)d_f.$$

Thus one may take $C_{\text{fib}} = B+1$. $\square$

The next lemma isolates the pointwise geometric input used in the block argument. If a metric of definite diameter is close to invariant metrics for two torus fibrations and both fibrations have sufficiently small fibers, then the fibrations are fiber-equivalent. The lemma also improves the fixed-fiber bound and gives a definite lower bound for the fixed quotient circle.

Lemma 5.5 (Pointwise fixed-fiber improvement). *Fix $\delta_0 > 0$, $0 < \lambda_p < \delta_0/10$, and $\theta_{\max} > 0$. There are constants*

$$0 < \theta < \theta_{\max}, \quad \beta_{\text{fib}} > 0, \quad \varepsilon_{\text{fib}} > 0,$$

with the following property.

Let q be a metric on a closed three-manifold M such that

$$(5.1) \quad \text{diam}(M, q) \geq \delta_0.$$

Let $\pi : M \rightarrow S^1$ be a torus fibration, and let r be an invariant metric for π . Suppose that

$$(5.2) \quad \|q - r\|_{C^1(M,r)} < \beta_{\text{fib}}, \quad d_{\mathbb{T}^2}(r) \leq \theta.$$

Suppose also that M has a second torus fibration $\pi_B : M \rightarrow S_B^1$ and an invariant metric h_B for π_B such that

$$(5.3) \quad \|h_B - q\|_{C^1(M,q)} < \varepsilon_{\text{fib}}, \quad \text{diam}_q(F_B) \leq \varepsilon_{\text{fib}},$$

for every π_B -fiber F_B . Then π_B is fiber-equivalent to π , and

$$(5.4) \quad d_{\mathbb{T}^2}(r) < \frac{\theta}{2}, \quad L(S^1, r_{S^1}) \geq \lambda_p.$$

Proof. Choose preliminary upper bounds for β_{fib} and ε_{fib} so small that q , r , and h_B are pairwise 2-bilipschitz. On this neighborhood there is a universal comparison constant C_{cmp} such that

$$(5.5) \quad \begin{aligned} \|h_B - r\|_{C^1(M,r)} &\leq C_{\text{cmp}} (\|h_B - q\|_{C^1(M,q)} + \|q - r\|_{C^1(M,r)}), \\ \text{diam}_{h_B}(F_B) &\leq 2^{1/2} \text{diam}_q(F_B). \end{aligned}$$

The first estimate follows from the standard formula for the difference of two Levi-Civita connections; the second is the bilipschitz comparison.

Let $C_{\text{fib}} = C_{\text{fib}}(2)$ be the constant in Lemma 5.4. Choose $0 < \theta < \theta_{\max}$ so that

$$(5.6) \quad 2^{1/2} \left(\frac{\lambda_p}{2} + \theta \right) < \delta_0, \quad C_{\text{fib}} 2^{1/2} \theta < \frac{\delta_0}{2}.$$

Apply Lemma 5.3, with base-length lower bound λ_p and target fiber diameter $\theta/2$, to obtain constants $\eta > 0$ and $\gamma > 0$. Now choose ε_{fib} and β_{fib} , still inside the preliminary bilipschitz neighborhood, so that

$$(5.7) \quad \begin{aligned} C_{\text{fib}} \varepsilon_{\text{fib}} &< \frac{\delta_0}{2}, & 2^{1/2} \varepsilon_{\text{fib}} &< \eta, \\ C_{\text{cmp}}(\varepsilon_{\text{fib}} + \beta_{\text{fib}}) &< \gamma. \end{aligned}$$

We first prove the quotient-length bound. If $L(S^1, r_{S^1}) < \lambda_p$, then Lemma 5.1, applied with $h = r$, $g = q$, and $B = 2$, gives

$$\text{diam}(M, q) \leq 2^{1/2} \left(\frac{\lambda_p}{2} + \theta \right) < \delta_0,$$

contrary to (5.1). Hence

$$(5.8) \quad L(S^1, r_{S^1}) \geq \lambda_p.$$

The π -fibers have q -diameter at most $2^{1/2}\theta$, while the π_B -fibers have q -diameter at most ε_{fib} . If the two fibrations were fiber-inequivalent, set

$$d_* := \max\{2^{1/2}\theta, \varepsilon_{\text{fib}}\}.$$

Lemma 5.4, applied with the common fiber-diameter bound $d_f = d_*$, would give

$$\text{diam}(M, q) \leq C_{\text{fib}} d_* < \delta_0,$$

by (5.6) and (5.7), again contradicting (5.1). Thus π_B is fiber-equivalent to π .

Finally, (5.5) and (5.7) give

$$\|h_B - r\|_{C^1(M, r)} < \gamma, \quad \text{diam}_{h_B}(F_B) \leq 2^{1/2} \varepsilon_{\text{fib}} < \eta.$$

Together with (5.8), Lemma 5.3 yields $d_{\mathbb{T}^2}(r) < \theta/2$, completing the proof. $\square$

6. ONE DETURCK STAGE AND REGULAR RESTART

This section turns the geometric data at the beginning of a large block into a controlled Ricci–DeTurck evolution and supplies the regularity needed to recover the analytic stage-start conditions at the next scale. Starting from a metric close to its fiberwise average, it constructs a locally $\mathbb{T}^2$ -invariant Ricci-flow background, compares the original flow to it throughout one stage, and uses positive-time smoothing, averaging, and parabolic rescaling at the endpoint. The resulting stage estimate can be iterated with constants independent of collapse.

The section is organized as follows. It begins by fixing the metric and gauge notation and, for Sections 6–8, the derivative order N and Hölder exponent σ . The first subsection, *Rescaled blocks and admissible stage starts*, defines a parabolically rescaled large block, states the analytic and geometric conditions at an admissible stage start, and gives Bamler’s time-slice comparison after pulling it back by the same diffeomorphism used to define the rescaled block. The second subsection, *Averaged endpoints and positive-time regularization*, explains the finite regularity needed for a restart, proves that high-order bounded

geometry survives averaging, and states the positive-time smoothing result that reproduces the bounded-geometry part of the analytic start conditions at the rescaled endpoint. The final subsection, *Stage preparation and stability*, evolves the averaged initial metric by an invariant Ricci flow, places the rescaled original flow in Ricci–DeTurck gauge relative to that background, proves uniform continuous dependence through the whole stage, and verifies that rescaling and averaging at time 100 preserve bounded geometry and control the end-point error. Finite-time decay in Section 7 contracts that error, while fibration comparison reproduces the short-fiber condition. Thus this section proves the analytic estimate used on every block of the iteration.

Standing regularity convention. Throughout Sections 6–8, fix once and for all an integer $N \geq 2$ and a Hölder exponent $\sigma \in (0, 1)$. Constants in these sections may depend on N and σ unless stated otherwise.

Standing fixed-threshold application. Throughout Sections 6–8, fix a particular application of Proposition 2.6 to the working-cover flow. Its cutoff is denoted by c_0 , its tail time by $T_B(c_0)$, and its nonincreasing comparison error by $\varepsilon_B^{c_0}$. Thus

$$Kc_0^2 < \varepsilon_{\text{af}}, \quad \varepsilon_B^{c_0}(t) \longrightarrow 0.$$

Every block, stage, and component statement in these sections is relative to this one fixed application, and every large-diameter threshold δ_0 satisfies $\delta_0 \geq c_0$. The analytic constants depend only on the parameters displayed in their statements; only the late start times may additionally depend on $T_B(c_0)$ and $\varepsilon_B^{c_0}$. Later arguments use this machinery with $c_0 = c_{\text{af}}$, or, in the auxiliary Euclidean argument, with $c_0 = \zeta$.

Metric and gauge notation. The following notation is used throughout the block construction.

Level	Notation	Role
physical time	$g(t), g'(t), h(t)$	the original Ricci flow, its continuous stagewise DeTurck pullback, and the right-continuous piecewise invariant comparison family;
one rescaled stage	$p(\tau), q(\tau), r(\tau)$	the parabolically rescaled original flow, its Ricci–DeTurck representative, and the invariant Ricci-flow background;
auxiliary construction	$r_D(\tau)$	the Ricci–DeTurck solution with background p and initial data $r(1)$, used only to construct and control the invariant Ricci flow r ;
time-slice auxiliary	$h_B(\tau), \pi_B(\tau)$	Bamler’s invariant comparison metric and its possibly time-dependent torus fibration, pulled back to the coordinates of the current rescaled stage;
restart	$\rho^+, q^+(1), r^+(1)$	the rescaled invariant endpoint, the rescaled DeTurck endpoint, and the averaged metric used to restart the next stage;
fixed structure	π	the torus fibration retained throughout the block iteration.

Unsubscripted p, q, r denote a generic stage, while p_k, q_k, r_k denote the corresponding objects on stage k .

Figure 2 shows the relation among g, g' , and h after the stages are assembled.




FIGURE 2. The original flow g , its continuous stagewise DeTurck pullback g' , and the invariant comparison family h . The family h may jump at a switching time because the parabolically rescaled endpoint is averaged before the next stage. The drawing is schematic and not to scale; vertical placement only separates the three families and has no metric meaning.

6.1. Rescaled blocks and admissible stage starts. All unqualified Hölder norms in this section are the intrinsic lifted norms defined in Paragraph A.1, with the displayed reference metric. Fixed-coordinate uniformly local norms are used only inside the Schauder and bootstrapping arguments, after a PBG bound has been established; Lemma A.2 then converts between the two conventions with constants independent of collapse. When the reference metric is not invariant, Lemma 3.2 provides the lifted injectivity-radius bound through the invariant bilipschitz comparison metric in the hypotheses. Let $\{g(t)\}_{t \geq 0}$ be the immortal Ricci flow under consideration.

Definition 6.1 (Parabolically rescaled large block). Fix $S > 0$, a diffeomorphism Φ , and $\delta_0 \geq c_0$, where c_0 is the cutoff in the standing fixed-threshold application. Set

$$(6.1) \quad p(\tau) = \Phi^*(S^{-1}g(S\tau)), \quad \frac{1}{2} \leq \tau \leq 101.$$

We call $p(\tau)$ a *parabolically rescaled δ_0 -large block starting at S* if $S/2 \geq T_B(c_0)$ and

$$\text{diam}(M, g(S\tau)) \geq \delta_0 \sqrt{S\tau}, \quad 1 \leq \tau \leq 100.$$

Definition 6.2 (Admissible stage start). Fix a torus structure and constants $\Lambda, \eta, \theta > 0$. We call the pair (q_0, r_0) , where $r_0 = \bar{q}_0$, an *admissible stage start with constants Λ, η, θ* if r_0 is Riemannian and

$$(6.2) \quad \text{BG}_N(r_0; \Lambda), \quad \|q_0 - r_0\|_{C^{N,\sigma}(M, r_0)} < \eta, \quad d_{\mathbb{T}^2}(r_0) < \theta.$$

The Riemannian condition and the first two displayed conditions are the *analytic start conditions*; the last displayed condition is the *geometric start condition*. Stage preparation and positive-time smoothing reproduce bounded geometry and control the endpoint error; finite-time decay and averaging contract that error, while fibration comparison reproduces the geometric condition. No quotient-length condition is part of an admissible stage start; the required lower bound is reconstructed from the large-diameter hypothesis inside each block.

On a block in the sense of Definition 6.1,

$$(6.3) \quad |\sec_{p(\tau)}| \leq K/\tau, \quad \frac{1}{2} \leq \tau \leq 101.$$

For each $\tau \in [1, 100]$, Bamler's large-diameter case provides, after pullback by a diffeomorphism, a torus fibration $\pi_B(\tau) : M \rightarrow S_B^1(\tau)$ and an invariant comparison metric $h_B(\tau)$ satisfying

$$(1 + \varepsilon_B^{c_0}(S\tau))^{-1} h_B(\tau) \leq p(\tau) \leq (1 + \varepsilon_B^{c_0}(S\tau)) h_B(\tau),$$

whose fibers have $p(\tau)$ -diameter at most $\varepsilon_B^{c_0}(S\tau)\sqrt{\tau}$. After increasing the starting time, the same comparison is as small as required in every fixed C^k -norm used below. The pair $(\pi_B(\tau), h_B(\tau))$ may vary arbitrarily with τ , and all these statements remain true after any further time-dependent pullback of the block.

6.2. Averaged endpoints and positive-time regularization. The next lemma uses the following finite-regularity scheme for the high-order restart:

$$(6.4) \quad \begin{array}{c} \text{BG}_N(q; \Lambda), \quad m_0 \text{ invariant}, \quad \|q - m_0\|_{C^{N+5}, \sigma(M, q)} < \varepsilon_{\text{av}}, \\ \Downarrow \\ |\tilde{\nabla}^j \text{Rm}_{\tilde{m}_0}|, |\tilde{\nabla}^j \text{Rm}_{\tilde{q}}| \leq C \quad (0 \leq j \leq N+3), \\ \Downarrow \\ \text{BG}_N(m_0; \Lambda'), \quad \text{BG}_N(\bar{q}; \Lambda'), \\ \Downarrow \end{array}$$

the bounded-geometry input for the next initial-face estimate is recovered.

Here ε_{av} is from Lemma 6.3 with $m = N$. The loss of two derivatives from the metric to curvature explains the passage from $N+5$ metric derivatives to curvature derivatives through order $N+3$. The $C^{N+5, \sigma}$ -control is not assumed at the old initial face; it is recovered at positive time by the smoothing proposition below.

Lemma 6.3 (High-order regularity of an averaged slice). *Fix an integer $m \geq 2$, $\sigma \in (0, 1)$, and $\Lambda_0 < \infty$. There are constants $\varepsilon_{\text{av}} > 0$, $C_{\text{av}} < \infty$, and $\Lambda_1 < \infty$, depending only on these data, with the following property. Let q and m_0 be metrics on a torus bundle, suppose that m_0 is invariant, and let $\bar{q}$ denote the average of q with respect to this torus structure. If $\text{BG}_m(q; \Lambda_0)$ holds and*

$$(6.5) \quad \|q - m_0\|_{C^{m+5}, \sigma(M, q)} < \varepsilon_{\text{av}},$$

then m_0 and $\bar{q}$ are Riemannian and satisfy $\text{BG}_m(m_0; \Lambda_1)$ and $\text{BG}_m(\bar{q}; \Lambda_1)$. Moreover,

$$(6.6) \quad \|\bar{q} - m_0\|_{C^{m+5,\sigma}(M,m_0)} \leq C_{\text{av}} \|q - m_0\|_{C^{m+5,\sigma}(M,q)}.$$

Proof. For sufficiently small ε_{av} , the two metrics are uniformly equivalent and their finite-order tensor Hölder norms are uniformly equivalent. The standard formulas comparing two Levi-Civita connections, followed by the formulas for curvature and its covariant derivatives, transfer the bounds for $\tilde{\nabla}^j \text{Rm}_{\bar{q}}$, $0 \leq j \leq m+3$, to $\tilde{m}_0$. Since m_0 is invariant, Lemma 3.2 then gives a uniform injectivity-radius lower bound for $\tilde{m}_0$.

Because m_0 is fixed by averaging,

$$\bar{q} - m_0 = \overline{q - m_0}.$$

Lemma 3.3, together with equivalence of the q - and m_0 -based norms, proves (6.6). Decreasing ε_{av} preserves positive definiteness. The same connection and curvature comparison formulas give the required curvature-derivative bounds for $\bar{q}$; this metric is invariant, so Lemma 3.2 gives its lifted injectivity-radius bound. Enlarging a single constant gives Λ_1 . $\square$

The next proposition states the positive-time smoothing used at every switching time. Informally, its hypotheses, smoothing step, and conclusion are as follows:

Hypotheses	A curvature-controlled invariant flow r and a Ricci–DeTurck flow q that is close to r in $X_r^N[1, 100]$.
Smoothing	Interior regularization raises the difference estimate from order N to order $N+5$ near $\tau = 100$.
Conclusion	Rescale by 100^{-1} , average, and obtain a $C^{N+5,\sigma}$ -small restart with a uniform lifted-geometry bound.

This argument propagates the high-order initial-face regularity needed for the next Ricci–DeTurck stage.

Proposition 6.4 (Positive-time smoothing and regular restart). *Fix $K_0 < \infty$. There are constants*

$$\beta_{\text{reg}} > 0, \quad C_{\text{reg}} < \infty, \quad \Lambda_{\text{reg}} < \infty,$$

depending only on K_0, N, σ , with the following property.

Hypotheses. *Let $r(\tau)$, $1 \leq \tau \leq 100$, be an invariant Ricci flow satisfying*

$$(6.7) \quad |\sec_{r(\tau)}| \leq K_0/\tau.$$

Let $q(\tau)$ be a Ricci–DeTurck flow relative to $r(\tau)$ and assume

$$(6.8) \quad \|q - r\|_{X_r^N[1,100]} \leq \beta_{\text{reg}}.$$

Rescale and average. *Put*

$$q^+(1) = 100^{-1}q(100), \quad \rho^+ = 100^{-1}r(100), \quad r^+(1) = \overline{q^+(1)}.$$

Conclusion. *Then $r^+(1)$ is Riemannian,*

$$\text{BG}_N(\rho^+; \Lambda_{\text{reg}}), \quad \text{BG}_N(r^+(1); \Lambda_{\text{reg}}),$$

and

$$(6.9) \quad \begin{aligned} & \|q^+(1) - \rho^+\|_{C^{N+5,\sigma}(M,\rho^+)} + \|r^+(1) - \rho^+\|_{C^{N+5,\sigma}(M,\rho^+)} \\ & \leq C_{\text{reg}} \|q - r\|_{C^0(M \times [1,100],r)}. \end{aligned}$$

The proof of Proposition 6.4 is given in Subsection A.6.

6.3. Stage preparation and stability. The next lemma prepares one analytic stage from the analytic part of an admissible stage start. It averages the initial metric, evolves the average by Ricci flow, and places the rescaled original flow in Ricci–DeTurck gauge relative to that invariant flow. Its additional high-order input is reproduced at the rescaled endpoint by Proposition 6.4; no short-fiber hypothesis is used in this analytic lemma.

Lemma 6.5 (Stage preparation, stability, and regular restart). *There are constants*

$$K_{\text{st}} < \infty, \quad \Lambda_* < \infty, \quad \beta_{\text{reg}} > 0, \quad C_{\text{st}} < \infty,$$

depending only on K, N, σ , with the following property. For every $0 < \beta_ \leq \beta_{\text{reg}}$, there is a constant $\eta_{\text{st}} > 0$, depending only on K, N, σ, β_* , such that the following holds. Let $p(\tau)$ be any parabolically rescaled δ_0 -large block, for some $\delta_0 \geq c_0$, under the standing fixed-threshold application. Fix the torus structure used during this stage and put*

$$q(1) = p(1), \quad r(1) = \overline{q(1)}.$$

Put

$$b := \|q(1) - r(1)\|_{C^{N,\sigma}(M,r(1))}.$$

Assume that $r(1)$ is Riemannian, $\text{BG}_N(r(1); \Lambda_)$ holds, and $b < \eta_{\text{st}}$. These are exactly the analytic start conditions in Definition 6.2, with $\Lambda = \Lambda_*$ and $\eta = \eta_{\text{st}}$. Then the invariant Ricci flow $r(\tau)$ and the Ricci–DeTurck flow $q(\tau)$ relative to it exist on $[1, 100]$, and*

$$(6.10) \quad |\sec_{r(\tau)}| \leq K_{\text{st}}/\tau, \quad 1 \leq \tau \leq 100,$$

$$(6.11) \quad \|q - r\|_{X_r^N[1,100]} < \beta_*,$$

$$(6.12) \quad \|q - r\|_{X_r^N[1,100]} \leq C_{\text{st}} b.$$

If

$$q^+(1) = 100^{-1}q(100), \quad \rho^+ = 100^{-1}r(100), \quad r^+(1) = \overline{q^+(1)},$$

then $r^+(1)$ is Riemannian, $\text{BG}_N(r^+(1); \Lambda_)$ holds, $|\sec_{r^+(1)}| \leq K_{\text{st}}$, and*

$$(6.13) \quad \begin{aligned} & \|q^+(1) - \rho^+\|_{C^{N+5,\sigma}(M,\rho^+)} + \|r^+(1) - \rho^+\|_{C^{N+5,\sigma}(M,\rho^+)} \\ & \leq C_{\text{st}} b. \end{aligned}$$

In particular,

$$(6.14) \quad \|r^+(1) - \rho^+\|_{C^{N,\sigma}(M,\rho^+)} \leq C_{\text{st}} b.$$

Before the proof, it is useful to display the two distinct gauges used in the construction:

$$(6.15) \quad \begin{array}{ccccc} r(1) & \xrightarrow[\text{background } p]{\text{Ricci-DeTurck}} & r_D & \xrightarrow{\text{DeTurck diffeomorphisms}} & r, \\ p(1) & \xrightarrow[\text{background } r]{\text{Ricci-DeTurck}} & q & \xrightarrow{\text{DeTurck diffeomorphisms}} & p. \end{array}$$

The first line of (6.15) constructs the comparison Ricci flow r , proves its curvature control, and verifies that it remains invariant. Once that background is available, the second line gauges the original rescaled flow relative to r and produces the stage metric q . The auxiliary metric r_D is not part of the inductive state.

Proof. We first choose constants that do not depend on the requested stage closeness β_* . The geometry of every rescaled block p is bounded uniformly through order N by the argument in Step 1 below. Let $\eta_1^{(0)}, C_1^{(0)}$ be the uniform constants in Proposition A.11 for the reference flow p , relabeled from η_{cd} and C_{cd} . The standard comparison formulas for nearby metrics and their Levi-Civita connections give uniform constants $b_{\text{eq}} > 0$ and $C_{\text{eq}} \geq 1$ such that $b < b_{\text{eq}}$ implies

$$\|r(1) - p(1)\|_{C^{N,\sigma}(M,p(1))} \leq C_{\text{eq}} b.$$

Set

$$\eta_1 := \min\{b_{\text{eq}}, \eta_1^{(0)}/C_{\text{eq}}\}, \quad C_1 := C_1^{(0)} C_{\text{eq}}.$$

Restrict the curvature comparison to the neighborhood

$$\|w - p(\tau)\|_{C^2(M,p(\tau))} \leq \frac{1}{2}.$$

The C^0 -part then gives $\frac{1}{2}p(\tau) \leq w \leq \frac{3}{2}p(\tau)$, so the inverse of w is uniformly controlled relative to $p(\tau)$. The curvature comparison formula and the uniformly bounded reference geometry give a constant $C_{\text{curv}}^{(0)}$ such that, on this neighborhood,

$$\sup_M |\sec_w| \leq \sup_M |\sec_{p(\tau)}| + C_{\text{curv}}^{(0)} \|w - p(\tau)\|_{C^2(M,p(\tau))}.$$

Set $C_{\text{curv}} = C_{\text{curv}}^{(0)} C_1$. Thus $\|w - p(\tau)\|_{C^2(M,p(\tau))} \leq C_1 b \leq \frac{1}{2}$ implies

$$\sup_M |\sec_w| \leq \sup_M |\sec_{p(\tau)}| + C_{\text{curv}} b.$$

These constants depend only on the fixed reference bounds, not on the stage threshold to be chosen below. Put

$$K_0 := K + 100C_{\text{curv}}, \quad K_{\text{st}} := 2K_0 + 1.$$

Apply the positive-time smoothing result, Proposition 6.4, with curvature constant K_{st} , and denote its constants by $\beta_{\text{reg}}, C_{\text{reg}}, \Lambda_{\text{reg}}$.

We next fix the geometry constant needed at the first stage. There is a constant $\Lambda_{\text{III}} = \Lambda_{\text{III}}(K, N, \sigma)$ such that every sufficiently late rescaled slice $q_0 = T^{-1}g(T)$, written in a Bamler torus-bundle gauge, satisfies $\text{BG}_N(q_0; \Lambda_{\text{III}})$. Indeed, the Type-III estimate on $[T/2, T]$ and Shi's estimates give the required curvature-derivative bounds after rescaling. The Bamler

metric is invariant and uniformly bilipschitz to q_0 , so Lemma 3.2 gives the lifted injectivity-radius bound. Apply Lemma 6.3 with $\Lambda_0 = \Lambda_{\text{III}}$, and call its output geometry constant Λ_{init} . Set

$$\Lambda_* := \max\{\Lambda_{\text{reg}}, \Lambda_{\text{init}}\}.$$

Let $s(\tau)$, $1 \leq \tau \leq 100$, be any reference Ricci flow whose initial metric satisfies $\text{BG}_N(s(1); \Lambda_*)$ and whose sectional curvature is bounded by K_{st} . Applying Lemma A.3 with $m = N$, $\Lambda_0 = \Lambda_*$, $K_0 = K_{\text{st}}$, and $L = 99$ gives a constant $\Lambda_{\text{par}} = \Lambda_{\text{par}}(K_{\text{st}}, \Lambda_*, N, \sigma)$ such that

$$\text{PBG}_N(s; \Lambda_{\text{par}}; [1, 100]).$$

In particular, the initial curvature-derivative bounds, their propagation near $\tau = 1$, the later Shi estimates, the lifted injectivity-radius bound, and the resulting parabolic coordinate control are all uniform at every stage. Let η_2, C_2 be the corresponding uniform constants in Proposition A.11. Choose

$$C_{\text{st}} \geq \max\{C_2, C_{\text{reg}}C_2, 1\}.$$

After fixing $0 < \beta_* \leq \beta_{\text{reg}}$, choose $\eta_{\text{st}} > 0$ no larger than η_1, η_2 , and sufficiently small that

$$(6.16) \quad C_{\text{st}}\eta_{\text{st}} < \beta_*, \quad C_1\eta_{\text{st}} \leq \frac{1}{2}, \quad \eta_{\text{st}} < 1,$$

and that the endpoint positivity and curvature perturbations used in Step 3 below are valid. All constants in the statement have now been fixed.

Step 1. First gauge: construct and control the invariant flow.

For every $\tau \in [1, 100]$, Lemma 3.2 applies to $g = p(\tau)$ and $h = h_B(\tau)$, with $B = 2$ and $\Lambda = \max\{K, 1\}/\tau$. Hence

$$\text{inj}_{\tilde{p}(\tau)} \geq i_2 \sqrt{\frac{\tau}{\max\{K, 1\}}} \geq \frac{i_2}{\sqrt{\max\{K, 1\}}}.$$

The curvature bound on $[1/2, 101]$ and Shi's estimates [21] give the curvature-derivative bounds through order $N + 3$ on $[1, 100]$. Thus $\text{PBG}_N(p; \Lambda_p; [1, 100])$ holds for a uniform constant Λ_p .

Let r_{D} be the Ricci-DeTurck solution with background p and initial metric $r(1)$. By the definitions of η_1 and C_1 , the bound $b < \eta_{\text{st}} \leq \eta_1$ implies the initial-data hypothesis of Proposition A.11, with the norm measured relative to $p(1)$, and gives

$$(6.17) \quad \|r_{\text{D}} - p\|_{X_p^N[1, 100]} \leq C_1 b.$$

The DeTurck diffeomorphisms convert r_{D} into the ordinary Ricci flow r starting from $r(1)$. By (6.16), $C_1 b < C_1 \eta_{\text{st}} \leq \frac{1}{2}$. Thus the C^2 -part of (6.17) places $r_{\text{D}}(\tau)$ in the curvature-comparison neighborhood above, and the curvature perturbation formula gives

$$\begin{aligned} \sup_M |\sec_{r(\tau)}| &= \sup_M |\sec_{r_{\text{D}}(\tau)}| \\ &\leq \frac{K}{\tau} + C_{\text{curv}} b \leq \frac{K_0}{\tau}, \quad 1 \leq \tau \leq 100. \end{aligned}$$

Lift r to the cyclic cover M_{cyc} . There the local fiber translations form a global $\mathbb{T}^2$ -action, say R_a , and $R_a^* r_{\text{cyc}}(1) = r_{\text{cyc}}(1)$. For each $a \in \mathbb{T}^2$, $R_a^* r_{\text{cyc}}(\tau)$ and $r_{\text{cyc}}(\tau)$ are complete bounded-curvature Ricci flows with the same initial metric; completeness and the curvature bound follow by lifting the smooth flow on the compact quotient. The uniqueness theorem for complete noncompact bounded-curvature Ricci flows [9, Theorem 1.1] therefore gives $R_a^* r_{\text{cyc}}(\tau) = r_{\text{cyc}}(\tau)$ for $1 \leq \tau \leq 100$. Hence $r(\tau)$ remains invariant.

Step 2. Second gauge: compare the original block with the invariant flow.

The analytic start conditions include $\text{BG}_N(r(1); \Lambda_*)$. Together with the curvature estimate just proved, Lemma A.3 gives

$$\text{PBG}_N(r; \Lambda_{\text{par}}; [1, 100]).$$

This is the high-order inductive input that was not supplied by the low-order start error alone.

Apply Proposition A.11 with reference flow $s = r$, and let q be the resulting Ricci–DeTurck flow with $q(1) = p(1)$. Since $b < \eta_{\text{st}} \leq \eta_2$,

$$(6.18) \quad \|q - r\|_{X_r^N[1, 100]} \leq C_2 b \leq C_{\text{st}} b.$$

By (6.16), this is smaller than β_* , proving (6.11) and (6.12). The ordinary Ricci flow obtained from q by the DeTurck diffeomorphisms starts from $p(1)$, so uniqueness identifies it with p .

Step 3. Smooth, rescale, and average the endpoint.

Since $\|q - r\|_{X_r^N} \leq \beta_* \leq \beta_{\text{reg}}$, Proposition 6.4 applies to q and r . Combining (6.9) with (6.18) gives (6.13). The same proposition shows that $r^+(1)$ is Riemannian and satisfies $\text{BG}_N(r^+(1); \Lambda_{\text{reg}})$, hence also $\text{BG}_N(r^+(1); \Lambda_*)$. Equation (6.14) is the lower-order part of (6.13).

Finally,

$$|\sec_{\rho^+}| = 100 |\sec_{r(100)}| \leq K_0.$$

The C^2 -part of (6.13) changes sectional curvature by at most $CC_{\text{st}}b$. The last smallness requirement in the choice of η_{st} , together with $K_{\text{st}} > 2K_0$, gives $|\sec_{r^+(1)}| \leq K_{\text{st}}$. Thus the endpoint reproduces the high-order initial-face hypothesis needed for the next stage. $\square$

Thus one analytic stage can be constructed with constants that are uniform under collapse: averaging produces an invariant reference flow, the original block remains controlled in Ricci–DeTurck gauge, and the averaged rescaled endpoint stays in the same stability neighborhood.

7. ONE-BLOCK CONTRACTION AND STAGEWISE FAMILIES

This is the central inductive section of the paper. It combines the analytic stage estimate with the fixed-fibration geometry to produce a block map that reproduces its hypotheses while reducing the non-invariant error by a definite factor. The quotient-circle lower bound is not carried as an inductive coordinate: the large-diameter hypothesis and fibration comparison reconstruct it inside every block.

The first subsection, *Admissible-start data*, records the unified start conditions from Definition 6.2 and the two numerical state variables used in the block map. The second, *Persistence of the retained fibration*, uses the pointwise fixed-fiber improvement in a first-hitting argument to keep the retained fibers uniformly short and to recover a positive quotient-length lower bound throughout the block. The third, *Contraction on one large-diameter block*, combines that persistence with stage stability and finite-time decay, first in a concise block contraction theorem and then in a quantitative proposition. The final subsection, *Stagewise families and switching times*, records the stagewise metric families and switching-time notation used for iteration on a connected large-diameter interval in Section 8.

7.1. Admissible-start data. For a stage start $(q_k(1), r_k(1))$, with $r_k(1) = \overline{q_k(1)}$, write

$$(7.1) \quad b_k = \|q_k(1) - r_k(1)\|_{C^{N,\sigma}(M, r_k(1))}, \quad d_k = d_{\mathbb{T}^2}(r_k(1)).$$

Thus the numerical start data consist only of the non-invariant error and the retained-fiber diameter, while lifted bounded geometry is recorded separately. By Definition 6.2, the pair is an admissible stage start with constants Λ_* , η_{blk} , θ_{blk} precisely when

$$b_k < \eta_{\text{blk}}, \quad d_k < \theta_{\text{blk}}, \quad \text{BG}_N(r_k(1); \Lambda_*).$$

The block theorem below proves the direct implication

$$(7.2) \quad \left. \begin{array}{l} b_k < \eta_{\text{blk}}, \quad d_k < \theta_{\text{blk}}, \\ \text{BG}_N(r_k(1); \Lambda_*) \end{array} \right\} \implies \left\{ \begin{array}{l} b_{k+1} \leq \frac{1}{2} b_k, \quad d_{k+1} < \theta_{\text{blk}}, \\ \text{BG}_N(r_{k+1}(1); \Lambda_*). \end{array} \right.$$

No quotient-length condition is required at any stage start.

7.2. Persistence of the retained fibration.

Lemma 7.1 (Persistence of the retained fibration on a large block). *Fix $\delta_0 > 0$, $0 < \lambda_p < \delta_0/10$, and $\theta_{\text{max}} > 0$. Let*

$$\theta, \quad \beta_{\text{fib}}, \quad \varepsilon_{\text{fib}}$$

be any fixed choice of constants for which Lemma 5.5 holds with these parameters.

Let $p(\tau)$, $1 \leq \tau \leq 100$, be a parabolically rescaled δ_0 -large block, and let $q(\tau) = \Psi_\tau^ p(\tau)$ for a smooth family of diffeomorphisms. Let $r(\tau)$ be a smooth family of invariant metrics for one fixed fibration $\pi : M \rightarrow S^1$. Suppose*

$$(7.3) \quad d_{\mathbb{T}^2}(r(1)) < \theta$$

and, for every $\tau \in [1, 100]$,

$$(7.4) \quad \|q(\tau) - r(\tau)\|_{C^1(M, r(\tau))} < \beta_{\text{fib}}.$$

Suppose also that at every such τ there exist an invariant metric $h_B(\tau)$ and a torus fibration $\pi_B(\tau) : M \rightarrow S_B^1(\tau)$ such that, for every $\pi_B(\tau)$ -fiber F_B ,

$$(7.5) \quad \|h_B(\tau) - q(\tau)\|_{C^1(M, q(\tau))} < \varepsilon_{\text{fib}}, \quad \text{diam}_{q(\tau)} F_B \leq \varepsilon_{\text{fib}}.$$

Then, throughout the block,

$$(7.6) \quad d_{\mathbb{T}^2}(r(\tau)) < \theta,$$

$$(7.7) \quad L(S^1, r_{S^1}(\tau)) \geq \lambda_p.$$

Moreover, every $\pi_B(\tau)$ is fiber-equivalent to π .

Proof. The function $\tau \mapsto d_{\mathbb{T}^2}(r(\tau))$ is continuous. If it reached θ , let τ_* be the first such parameter. Since $q(\tau) = \Psi_\tau^* p(\tau)$ and the block is δ_0 -large,

$$\text{diam}(M, q(\tau_*)) = \text{diam}(M, p(\tau_*)) \geq \delta_0 \sqrt{\tau_*} \geq \delta_0.$$

At τ_* , equations (7.4) and (7.5), together with $d_{\mathbb{T}^2}(r(\tau_*)) = \theta$, verify all hypotheses of Lemma 5.5. That lemma gives

$$d_{\mathbb{T}^2}(r(\tau_*)) < \theta/2,$$

a contradiction. This proves (7.6).

Now fix any $\tau \in [1, 100]$. The same diameter estimate, (7.6), and the two comparison assumptions again allow Lemma 5.5 to be applied at time τ . Its quotient conclusion gives (7.7), and its fibration conclusion says that $\pi_B(\tau)$ is fiber-equivalent to π . $\square$

7.3. Contraction on one large-diameter block.

Theorem 7.2 (Block contraction theorem). *Under the standing fixed-threshold application, fix $\delta_0 \geq c_0$. There are constants*

$$T_{\text{blk}} < \infty, \quad \eta_{\text{blk}} > 0, \quad \theta_{\text{blk}} > 0, \quad \Lambda_* < \infty$$

with the following property. The constants $\eta_{\text{blk}}, \theta_{\text{blk}}, \Lambda_$ depend only on K, δ_0, N, σ . The start-time threshold T_{blk} may additionally depend on the fixed Bamler tail data $T_B(c_0)$ and $\varepsilon_B^{c_0}$ for the given flow.*

Let $p(\tau)$ be a parabolically rescaled δ_0 -large block starting at $S \geq T_{\text{blk}}$, fix a torus fibration $\pi : M \rightarrow S^1$, and put

$$q(1) = p(1), \quad r(1) = \overline{q(1)}.$$

Suppose that $r(1)$ is Riemannian and, with

$$b = \|q(1) - r(1)\|_{C^{N,\sigma}(M, r(1))}, \quad d = d_{\mathbb{T}^2}(r(1)),$$

one has

$$b < \eta_{\text{blk}}, \quad d < \theta_{\text{blk}}, \quad \text{BG}_N(r(1); \Lambda_*).$$

Equivalently, $(q(1), r(1))$ is an admissible stage start in the sense of Definition 6.2, with constants $\Lambda_, \eta_{\text{blk}}, \theta_{\text{blk}}$. Then the invariant Ricci flow r and the Ricci–DeTurck flow q relative to it exist on $[1, 100]$. At the rescaled endpoint, put*

$$q^+(1) = 100^{-1}q(100), \quad \rho^+ = 100^{-1}r(100), \quad r^+(1) = \overline{q^+(1)},$$

and set

$$b^+ = \|q^+(1) - r^+(1)\|_{C^{N,\sigma}(M, r^+(1))}.$$




FIGURE 3. One fixed-fibration block. Stability and fibration comparison make the finite-time decay estimate applicable; rescaling and averaging then produce a contracted restart with the same hypotheses.

Then $r^+(1)$ is Riemannian and

$$(7.8) \quad b^+ \leq \frac{1}{2}b, \quad d_{\mathbb{T}^2}(r^+(1)) < \theta_{\text{blk}}, \quad \text{BG}_N(r^+(1); \Lambda_*).$$

Thus the same admissible-start conditions hold at the restart. The ordinary Ricci flow corresponding to q is p .

Proposition 7.3 (Quantitative estimates for one block). *The constants in Theorem 7.2 can be chosen together with $C_{\text{st}} < \infty$, depending only on K, δ_0, N, σ , so that, under the hypotheses and in the notation of that theorem,*

$$(7.9) \quad \|q - r\|_{X_r^N[1,100]} \leq C_{\text{st}}b,$$

$$(7.10) \quad d_{\mathbb{T}^2}(r(\tau)) < \theta_{\text{blk}},$$

$$(7.11) \quad L(S^1, r_{S^1}(\tau)) \geq \delta_0/100, \quad 1 \leq \tau \leq 100.$$

At the restart, $r^+(1)$ is Riemannian, $\text{BG}_N(r^+(1); \Lambda_*)$ holds, and

$$(7.12) \quad b^+ \leq \frac{1}{2}b,$$

$$(7.13) \quad \|r^+(1) - \rho^+\|_{C^{N,\sigma}(M, \rho^+)} \leq C_{\text{st}}b,$$

$$(7.14) \quad d_{\mathbb{T}^2}(r^+(1)) < \theta_{\text{blk}}.$$

In particular, these estimates imply the qualitative conclusion (7.8).

Figure 3 summarizes the logical order of the one-block construction used in the proof.

Proof of Theorem 7.2 and Proposition 7.3. The proof has two steps. We first choose the constants in their dependency order. We then apply stage stability, the retained-fibration lemma, finite-time decay, and the restart estimate in that order.

Step 1. Choose the constants.

The constants are selected in the order

$$\begin{aligned} (K, \delta_0, N, \sigma) &\longrightarrow (K_{\text{st}}, \Lambda_*, \beta_{\text{reg}}, C_{\text{sc}}) \longrightarrow \alpha_{\text{blk}} \longrightarrow (\beta_{\text{orb}}^{\text{dec}}, \beta_{\text{pert}}^{\text{dec}}) \\ &\longrightarrow (\theta_{\text{blk}}, \beta_{\text{fib}}, \varepsilon_{\text{fib}}) \longrightarrow \beta_* \longrightarrow (\eta_{\text{blk}}, T_{\text{blk}}). \end{aligned}$$

The stage-stability constant C_{st} , which depends only on K, N, σ , is fixed together with K_{st}, Λ_* , and β_{reg} ; it is omitted from the chain because it is not itself a later smallness target. The principal dependency diagrams appear in Appendix C.

Take $K_{\text{st}}, \Lambda_*, \beta_{\text{reg}}$, and C_{st} from Lemma 6.5. Choose C_{sc} so that, whenever q is sufficiently close to an invariant metric r ,

$$(7.15) \quad \|100^{-1}u\|_{C^{N,\sigma}(M, \overline{100^{-1}q})} \leq C_{\text{sc}} \|u\|_{C^{N,\sigma}(M,r)},$$

for every symmetric two-tensor u . This is the scaling rule for tensor Hölder norms followed by equivalence of nearby reference metrics. Choose

$$0 < \alpha_{\text{blk}} < (2C_{\text{sc}})^{-1},$$

and let

$$\beta_{\text{orb}}^{\text{dec}}, \quad \beta_{\text{pert}}^{\text{dec}}$$

be the two thresholds supplied by Theorem 4.2 for $(\alpha_{\text{blk}}, K_{\text{st}}, \delta_0/100, N, \sigma)$.

Apply Lemma 5.5 with diameter lower bound δ_0 and

$$\lambda_{\text{p}} = \delta_0/100, \quad \theta_{\text{max}} = \beta_{\text{orb}}^{\text{dec}}/4.$$

Denote its outputs by $\theta_{\text{blk}}, \beta_{\text{fib}}$, and ε_{fib} , and choose

$$0 < \beta_* \leq \min\{\beta_{\text{pert}}^{\text{dec}}, \beta_{\text{fib}}, \beta_{\text{reg}}\}.$$

Let η_{st} be the stage threshold corresponding to this β_* in Lemma 6.5. Choose $\eta_{\text{blk}} \leq \eta_{\text{st}}$ below the finitely many smallness thresholds needed for the endpoint norm comparison and restart: (7.15) must apply when the endpoint perturbation has $C^{N,\sigma}$ -size at most $C_{\text{st}}\eta_{\text{blk}}$, and a perturbation of this size must carry fibers of diameter $< \theta_{\text{blk}}/10$ to fibers of diameter $< \theta_{\text{blk}}$.

Finally, since $\varepsilon_B^{c_0}(t) \rightarrow 0$ and Bamler's comparison converges in every fixed order used below, choose $T_{\text{blk}} \geq 2T_B(c_0)$ so large that

$$10\varepsilon_B^{c_0}(T_{\text{blk}}) < \varepsilon_{\text{fib}},$$

and so that on every block starting at $S \geq T_{\text{blk}}$, after any simultaneous stage pullback, the Bamler metric comparison is within ε_{fib} in C^1 and lies in all the already fixed Hölder and bilipschitz neighborhoods.

Step 2. Verify the stage hypotheses, retain the fibration, and contract.

The parabolically rescaled block hypotheses supply the Type-III curvature bound and an invariant small-fiber comparison at every slice. Hence Lemma 6.5 applies and yields q , r , (7.9), and

$$(7.16) \quad \|q(\tau) - r(\tau)\|_{C^1(M, r(\tau))} < \beta_*, \quad 1 \leq \tau \leq 100.$$

The DeTurck diffeomorphisms make $q(\tau)$ a time-dependent pullback of $p(\tau)$. Pull back Bamler's comparison metric and fibration by the same diffeomorphisms, and denote the pulled-back objects again by $h_B(\tau)$, $\pi_B(\tau)$, and $F_B(\tau)$. By the choice of T_{blk} , for every $\tau \in [1, 100]$ and every $\pi_B(\tau)$ -fiber $F_B(\tau)$,

$$(7.17) \quad \|h_B(\tau) - q(\tau)\|_{C^1(M, q(\tau))} < \varepsilon_{\text{fib}}, \quad \text{diam}_{q(\tau)}(F_B(\tau)) \leq \varepsilon_{\text{fib}}.$$

Indeed, since $\varepsilon_B^{c_0}$ is nonincreasing and $1 \leq \tau \leq 100$, Bamler's parabolically rescaled fiber estimate gives

$$\text{diam}_{q(\tau)}(F_B(\tau)) \leq \varepsilon_B^{c_0}(S\tau)\sqrt{\tau} \leq 10\varepsilon_B^{c_0}(T_{\text{blk}}) < \varepsilon_{\text{fib}}.$$

Apply Lemma 7.1 with

$$\lambda_p = \delta_0/100, \quad \theta_{\max} = \beta_{\text{orb}}^{\text{dec}}/4.$$

Its output θ is θ_{blk} by the choice in Step 1. Its initial-fiber hypothesis is the assumed inequality $d_{\mathbb{T}^2}(r(1)) < \theta_{\text{blk}}$; its stage closeness follows from (7.16) and $\beta_* \leq \beta_{\text{fib}}$; and its comparison hypotheses are (7.17). It therefore gives (7.10) and (7.11) throughout the block. The curvature bound and the quotient-length estimate (7.11), together with (7.9) and $\beta_* \leq \beta_{\text{pert}}^{\text{dec}}$, give the perturbative hypothesis of Theorem 4.2. Moreover, (7.10) and $\theta_{\text{blk}} < \beta_{\text{orb}}^{\text{dec}}/4$ give its short-orbit hypothesis. Hence

$$(7.18) \quad \|q^\perp\|_{X_r^N[10, 100]} \leq \alpha_{\text{blk}} b.$$

Averaging commutes with constant scaling, so

$$q^+(1) - r^+(1) = 100^{-1} q(100)^\perp.$$

Using (7.15) and (7.18) at $\tau = 100$, we obtain

$$b^+ \leq C_{\text{sc}} \alpha_{\text{blk}} b \leq \frac{1}{2} b,$$

which is (7.12). The endpoint clause of Lemma 6.5 gives (7.13) and $\text{BG}_N(r^+(1); \Lambda_*)$. By (7.10), the fibers of ρ^+ have diameter $< \theta_{\text{blk}}/10$; the choice of η_{blk} and (7.13) give (7.14). Since $b^+ \leq b/2 < \eta_{\text{blk}}$, (7.12), (7.14), and $\text{BG}_N(r^+(1); \Lambda_*)$ reproduce the three admissible-start conditions at the restart. The last assertion follows from the DeTurck correspondence and uniqueness for Ricci flow. $\square$

Thus every large block reconstructs the quotient-length lower bound needed for finite-time decay, while only the error, the retained-fiber diameter, and the lifted regularity are passed to the next block.

7.4. Stagewise families and switching times. The block contraction theorem is now assembled over the geometric sequence of switching times. This subsection constructs the global gauge, the continuous physical-scale family g' , and the piecewise invariant comparison family h .

We now iterate the block contraction theorem. Once T_0 and the fixed fibration are chosen, put

$$T_k = 100^k T_0, \quad k = 0, 1, 2, \dots$$

On stage k , write

$$q_k(1) = T_k^{-1} g'(T_k), \quad r_k(1) = \overline{q_k(1)},$$

let $q_k(\tau), r_k(\tau)$ be the corresponding flows on $[1, 100]$, and define b_k, d_k by (7.1). Define

$$g'(t) = T_k q_k(t/T_k), \quad T_k \leq t \leq T_{k+1},$$

and use the right-continuous convention

$$h(t) = T_k r_k(t/T_k), \quad T_k \leq t < T_{k+1}.$$

At a switching time,

$$(7.19) \quad h(T_{k+1}-) = T_k r_k(100), \quad h(T_{k+1}) = T_{k+1} r_{k+1}(1).$$

Thus g' is continuous and piecewise smooth, while h is an invariant Ricci flow on each stage and may jump when the endpoint is averaged.

8. QUOTIENT-CIRCLE ENERGY, LENGTH, AND COMPONENTWISE ITERATION

This section converts the blockwise invariant comparison into the geometric diameter estimate. For a locally $\mathbb{T}^2$ -invariant torus-bundle flow, the relevant base geometry is governed by a one-dimensional vertical energy and the length of the quotient circle; smooth evolution, fixed rescaling, and the small restart jumps therefore reduce to a scalar recursion. The resulting bounds are inserted into one iteration theorem for an arbitrary connected large-diameter time interval.

The first subsection, *One-dimensional estimates on a smooth stage*, recalls the one-dimensional estimates on a smooth invariant stage. The second, *Transition jump estimates*, controls the scalar changes at a restart, and the third, *The scalar energy-length recursion*, solves the resulting energy-length recursion. The final subsection, *Iteration on a connected large-diameter interval*, combines these estimates with the block contraction theorem, including the initial full block and a possible terminal partial block, and obtains quotient-length and total-diameter control on the connected interval. On a persistent tail, the concluding corollary upgrades the scalar bounds to convergence of the normalized quotient length and then to the full metric-circle limit.

8.1. One-dimensional estimates on a smooth stage. We recall the one-dimensional estimates for a smooth invariant Ricci flow. In an invariant local trivialization, write the metric in Kaluza–Klein form

$$\begin{aligned} g(t) &= h_{yy}(y, t) dy^2 \\ &\quad + G_{ij}(y, t)(dx^i + A^i(y, t)dy)(dx^j + A^j(y, t)dy), \\ G(y+1, t) &= H^T G(y, t) H, \end{aligned}$$

where $H \in \mathrm{SL}(2, \mathbb{Z})$ by the convention of Subsection 3.2. Here $h_{yy}(y, t) dy^2$ is the quotient metric on the base circle and $G_{ij}(y, t)$ is the metric on the torus fibers; the connection term $A^i dy$ does not enter the scalar quantities below. Since the base is one-dimensional, the local evolution formulas for the quotient coefficient and vertical block may be computed in horizontal gauges with no mixed term. Changes of such gauges only reparametrize the base and apply fixed linear changes to the vertical coordinates, so the scalar quantities E and L are unaffected. Put

$$\mathcal{E} = h^{yy} \mathrm{Tr} \left((G^{-1} G_y)^2 \right), \quad E(t) = \mathcal{E}_{\max}(t),$$

and let

$$L(t) = L(S^1, h_{yy}(t)) = \int_{S^1} \sqrt{h_{yy}(y, t)} dy,$$

be the length of the quotient circle. The inequalities in Proposition 2.4 are written in [18] in the modified one-dimensional gauge obtained from the ordinary Ricci flow by a horizontal Lie derivative. Since $E(t)$ and $L(t)$ are invariant under reparametrizing the base circle, the same estimates hold for the ordinary invariant Ricci flow. Thus, whenever the denominators are nonzero,

$$(8.1) \quad \frac{1}{E(t_1)} - \frac{1}{E(t_0)} \geq \frac{1}{2}(t_1 - t_0), \quad t_1 \geq t_0,$$

and

$$(8.2) \quad \frac{d}{dt} \log L(t) \leq \frac{1}{4} E(t).$$

If $E(t_0) = 0$, then the maximum-principle inequality underlying (8.1) gives $E(t) = 0$ for every $t \geq t_0$. In this case, the reciprocal estimate is not invoked, and (8.2) has zero right-hand side.

8.2. Transition jump estimates.

Lemma 8.1 (Transition jump estimate). *Fix $0 \leq E_0 < \infty$. There are constants*

$$\delta_J > 0, \quad C_J < \infty,$$

depending only on E_0 , with the following property. Let ρ and $\tilde{\rho}$ be smooth invariant metrics for any fixed torus structure. Suppose that

$$E(\rho) \leq E_0, \quad \delta := \|\tilde{\rho} - \rho\|_{C^1(M, \rho)} < \delta_J.$$

Then

$$(8.3) \quad E(\tilde{\rho}) \leq E(\rho) + C_J \delta, \quad \left| \log \frac{L(\tilde{\rho})}{L(\rho)} \right| \leq C_J \delta.$$

Proof. Step 1. Choose invariant coordinates and control the coefficients.

Choose ρ -arclength s on a base arc and fiber coordinates for which the connection one-form of ρ vanishes. In this horizontal gauge for ρ ,

$$\rho = ds^2 + G_{ij}(s) dx^i dx^j.$$

Any invariant metric γ , written in the same coordinates, has the form

$$\gamma = a_\gamma(y) dy^2 + G_{\gamma,ij}(y)(dx^i + A_\gamma^i(y)dy)(dx^j + A_\gamma^j(y)dy),$$

where $a_\gamma(y)dy^2$ is the quotient metric and G_γ is the metric on the torus fibers. Therefore

$$E(\gamma) = \max_{S^1} a_\gamma^{-1} \operatorname{Tr} ((G_\gamma^{-1} \partial_y G_\gamma)^2), \quad L(\gamma) = \int_{S^1} \sqrt{a_\gamma} dy.$$

The connection form A_γ does not enter either expression.

Write

$$a = a_\rho, \quad G = G_\rho, \quad \tilde{a} = a_{\tilde{\rho}}, \quad \tilde{G} = G_{\tilde{\rho}}.$$

The energy bound implies

$$(8.4) \quad |G^{-1/2} G_s G^{-1/2}| \leq C(E_0).$$

The vertical block is the restriction of the metric to the fiber directions. For $x \in \pi^{-1}(y)$, let $\mathcal{V}_x = \ker d\pi_x$ be the vertical subspace. The quotient coefficient can be characterized by

$$a_\gamma(y) = \min_{V \in \mathcal{V}_x} \gamma_x(\partial_y + V, \partial_y + V).$$

Local $\mathbb{T}^2$ -invariance makes the minimum independent of the choice of x in the fiber. It also shows directly that a_γ depends smoothly on the metric coefficients and on the inverse of the vertical block. At a fixed basepoint, make a constant linear change of vertical frame so that $G = I$. In this ρ -orthonormal adapted frame, the relative $C^1(M, \rho)$ -bound and (8.4) give

$$(8.5) \quad |\hat{a} - 1| + |\tilde{G} - I| + |\partial_s \tilde{G} - \partial_s G| \leq C(E_0) \delta,$$

where $\hat{a} = \tilde{a}/a$ is the quotient coefficient of $\tilde{\rho}$ in the reference arclength coordinate. For the derivative term, the first covariant derivative of $\tilde{\rho} - \rho$ controls the ordinary s -derivative of the vertical block; the conversion terms involve only reference connection coefficients, bounded by (8.4). Because the frame was normalized before applying matrix estimates, these constants do not depend on the condition number of the original G , the invariant atlas, the monodromy, or the fibration class.

Step 2. Compare quotient lengths.

After decreasing δ_J , the first term in (8.5) gives

$$\left| \log \sqrt{\hat{a}} \right| \leq C(E_0) \delta.$$

Since $dL(\rho) = ds$ and $dL(\tilde{\rho}) = \sqrt{\hat{a}} ds$, integration around the quotient circle yields

$$e^{-C(E_0)\delta} L(\rho) \leq L(\tilde{\rho}) \leq e^{C(E_0)\delta} L(\rho),$$

which is the length estimate in (8.3).

Step 3. Compare energies.

At the chosen basepoint set

$$B_\rho = \partial_s G, \quad \tilde{B} = \hat{a}^{-1/2} \tilde{G}^{-1/2} (\partial_s \tilde{G}) \tilde{G}^{-1/2}.$$

These are the normalized vertical derivatives in the reference-orthonormal frame. The traces of their squares are the pointwise energy densities of ρ and $\tilde{\rho}$. For δ small, $\tilde{G}$ stays in a fixed neighborhood of the identity, where $S \mapsto S^{-1}$ and $S \mapsto S^{-1/2}$ have uniformly bounded derivatives. Thus (8.5) and $|B_\rho| \leq C(E_0)$ imply

$$|\tilde{B} - B_\rho| \leq C(E_0)\delta, \quad |\tilde{B}| \leq C(E_0).$$

Consequently,

$$\begin{aligned} |\operatorname{Tr}(\tilde{B}^2) - \operatorname{Tr}(B_\rho^2)| &= |\operatorname{Tr}((\tilde{B} - B_\rho)\tilde{B}) + \operatorname{Tr}(B_\rho(\tilde{B} - B_\rho))| \\ &\leq C(E_0)\delta. \end{aligned}$$

The basepoint was arbitrary. Taking the maximum over S^1 and enlarging C_J to cover both estimates proves (8.3). $\square$

8.3. The scalar energy-length recursion.

Lemma 8.2 (Scalar quotient recursion and convergence). *Fix $C_0 \geq 0$, and let $b_k \geq 0$, $k \geq 0$, be a summable sequence. Suppose that the sequences $E_k \geq 0$ and $\ell_k > 0$, $k \geq 0$, satisfy*

$$(8.6) \quad E_{k+1} \leq \frac{100E_k}{1 + 99E_k/2} + C_0 b_k, \quad \ell_{k+1} \leq (1 + C_0 b_k) \left(\frac{1 + 99E_k/2}{100} \right)^{1/2} \ell_k.$$

Then $\sup_k E_k < \infty$ and $\sup_k \ell_k < \infty$, with a bound depending only on E_0 , ℓ_0 , C_0 , and $\sum_k b_k$. If, in addition,

$$(8.7) \quad \inf_k \ell_k > 0,$$

then there is a number $\ell_\infty > 0$ such that

$$(8.8) \quad \ell_k \longrightarrow \ell_\infty, \quad E_k \longrightarrow 2.$$

Proof. For a real number x , write $x_+ = \max\{x, 0\}$. Set $y_k = (E_k - 2)_+$. Since

$$\frac{100x}{1 + 99x/2} - 2 = \frac{x - 2}{1 + 99x/2},$$

and $(x + y - 2)_+ \leq (x - 2)_+ + y$ for $y \geq 0$, the energy recurrence gives

$$y_{k+1} \leq \frac{1}{100} y_k + C_0 b_k.$$

Iterating this inequality and using $\sum_k b_k < \infty$ shows that

$$(8.9) \quad \sum_k y_k < \infty.$$

In particular, (E_k) is uniformly bounded.

On that bounded range, put

$$Q(x) = \left(\frac{1 + 99x/2}{100} \right)^{1/2}.$$

Then $Q(x) \leq 1$ for $x \leq 2$, while $\log Q(x) \leq C(x - 2)_+$ for $x \geq 2$. Taking logarithms in the length recurrence and iterating gives

$$\log \frac{\ell_k}{\ell_0} \leq C \sum_{j < k} b_j + C \sum_{j < k} y_j.$$

Both sums are uniformly bounded, so (ℓ_k) is uniformly bounded as well.

Assume now (8.7). Define

$$a_k = \log(1 + C_0 b_k) + \log Q(E_k).$$

The positive parts are summable:

$$(8.10) \quad \sum_k (a_k)_+ \leq C_0 \sum_k b_k + C \sum_k (E_k - 2)_+ < \infty.$$

Writing $u_k = \log \ell_k$, the second inequality in (8.6) gives $u_{k+1} \leq u_k + a_k$. Hence

$$z_k = u_k + \sum_{j=k}^{\infty} (a_j)_+$$

is nonincreasing. It is bounded below by $\log(\inf_j \ell_j)$, so it converges. The tail in (8.10) tends to zero, and therefore u_k , and hence ℓ_k , converges to a positive limit ℓ_∞ .

It remains to identify the limiting energy. Since $\ell_{k+1}/\ell_k \rightarrow 1$ and $b_k \rightarrow 0$, the second inequality in (8.6) rules out $E_k \leq 2 - \varepsilon$ along any subsequence: on such a subsequence its right-hand multiplicative factor is eventually bounded above by a number strictly smaller than one. Thus $\liminf_k E_k \geq 2$. On the other hand, (8.9) gives $(E_k - 2)_+ \rightarrow 0$, and hence $\limsup_k E_k \leq 2$. This proves (8.8). $\square$

8.4. Iteration on a connected large-diameter interval.

Theorem 8.3 (Iteration on a connected large-diameter interval). *Under the standing fixed-threshold application with cutoff c_0 , fix $\delta_0 \geq c_0$, an integer $N \geq 2$, and $\sigma \in (0, 1)$. There are constants $\eta_0 > 0$ and $T_p < \infty$ with the following properties. The constant η_0 depends only on K, δ_0, N, σ ; the threshold T_p may additionally depend on the fixed Bamler tail data $T_B(c_0)$ and $\varepsilon_B^{c_0}$. Neither depends on the number D_1 below.*

Let I be a connected time interval on which

$$(8.11) \quad \text{diam}(M, g(t)) \geq \delta_0 \sqrt{t}.$$

Suppose that $S \in I$, $S \geq T_p$, and

$$(8.12) \quad \text{diam}(M, g(S)) \leq D_1 \sqrt{S}$$

for some $0 < D_1 < \infty$. At time S , choose the torus fibration from case 2 of Proposition 2.6, put

$$q_0(1) = S^{-1}g(S), \quad r_0(1) = \overline{q_0(1)},$$

where the average is taken in that fibration, and suppose that $r_0(1)$ is Riemannian and

$$(8.13) \quad b_0 = \|q_0(1) - r_0(1)\|_{C^{N,\sigma}(M, r_0(1))} < \eta_0.$$

Set $T_k = 100^k S$. Call the k th block controlled if $[T_k, T_{k+1}] \subset I$. On every controlled block, the same retained torus fibration supports stagewise families $q_k(\tau)$, $r_k(\tau)$, $1 \leq \tau \leq 100$, and the estimates

$$(8.14) \quad b_{k+1} \leq \frac{1}{2} b_k,$$

$$(8.15) \quad \|q_k - r_k\|_{X_{r_k}^N[1,100]} \leq C_{\text{st}} b_k,$$

$$(8.16) \quad \|r_{k+1}(1) - 100^{-1}r_k(100)\|_{C^{N,\sigma}} \leq C_{\text{st}} b_k.$$

Here and below the norm in the last line is taken relative to $100^{-1}r_k(100)$. On the union of the controlled blocks these stages assemble into a continuous piecewise smooth pullback $g'(t) = \Psi(t)^*g(t)$ and a right-continuous piecewise invariant family $h(t)$, as in Section 7. There is a constant $C < \infty$, depending only on K, δ_0, N, σ and D_1 , such that

$$(8.17) \quad L(S^1, h_{yy}(t)) \leq C\sqrt{t}$$

on every controlled block, for either one-sided value at a switching time, and

$$(8.18) \quad \text{diam}(M, g(t)) \leq C\sqrt{t}$$

for every $t \in I \cap [S, \infty)$. Thus the same theorem covers both an unbounded persistent tail and a finite component with a terminal partial block.

Figure 4 summarizes the energy–length recursion used in the proof below.

Quotient-Circle Length Bound (Stagewise Recursion)




FIGURE 4. Energy–length recursion. The smooth-stage inequalities, the fixed rescaling laws, and the summable jump errors reduce the geometric iteration to a scalar contraction about the energy level 2. In general this gives the uniform length bound shown in the figure; on a persistent tail, the positive lower length bound upgrades the same recursion to convergence.

Proof. **Step 1. Choose the constants and prepare the initial slice.**

Apply Theorem 7.2 and Proposition 7.3 for the present δ_0, N, σ . Denote their constants by

$$T_{\text{blk}}, \quad \eta_{\text{blk}}, \quad \theta_{\text{blk}}, \quad C_{\text{st}}, \quad \Lambda_*.$$

Let δ_J and C_J be the constants in Lemma 8.1 for reference energy bound 3. Choose

$$(8.19) \quad 0 < \eta_0 \leq \eta_{\text{blk}}, \quad C_{\text{st}}\eta_0 < \delta_J, \quad C_J C_{\text{st}}\eta_0 \leq 1.$$

Choose $T_p \geq T_{\text{blk}}$ so large that whenever $t \geq T_p$ and $\text{diam}(M, g(t)) \geq \delta_0\sqrt{t}$, the Type-III estimate on $[t/2, t]$, Shi's estimates, and Lemma 3.2 give the uniform lifted geometry used in Lemma 6.3. Require also that Bamler's comparison at time t be sufficiently small in $C^{N+5, \sigma}$

for that lemma; C^{N+6} -smallness suffices. If r_t is the average of $t^{-1}g(t)$ in Bamler's fibration, then

$$(8.20) \quad \text{BG}_N(r_t; \Lambda_*), \quad d_{\mathbb{T}^2}(r_t) < \theta_{\text{blk}}.$$

These choices have the dependence asserted in the theorem and are independent of D_1 . At $t = S$, equations (8.13) and (8.20) show that $(q_0(1), r_0(1))$ is an admissible stage start in the sense of Definition 6.2, with constants $\Lambda_*, \eta_{\text{blk}}, \theta_{\text{blk}}$.

The Type-III estimate also gives a constant C_{III} , depending only on K , such that

$$(8.21) \quad C_{\text{III}}^{-1}g(t_1) \leq g(t_2) \leq C_{\text{III}}g(t_1) \quad (T_p \leq t_1 \leq t_2 \leq 100t_1).$$

Indeed, $|\partial_t g|_g = 2|\text{Ric}|_g \leq C/t$, and integration gives the comparison.

Step 2. Iterate all controlled blocks and assemble the gauge.

If $[S, 100S] \not\subset I$, connectedness gives $\sup I \leq 100S$. Equations (8.12) and (8.21) then imply

$$\text{diam}(M, g(t)) \leq CD_1\sqrt{S} \leq CD_1\sqrt{t}, \quad t \in I \cap [S, \infty),$$

so there is nothing more to prove. Assume henceforth that the first block is controlled.

Set $\Phi_0 = \text{Id}$. Inductively suppose that the k th block is controlled and that, with

$$p_k(\tau) = \Phi_k^*(T_k^{-1}g(T_k\tau)), \quad q_k(1) = p_k(1), \quad r_k(1) = \overline{q_k(1)},$$

the pair $(q_k(1), r_k(1))$ is an admissible stage start in the sense of Definition 6.2. Since (8.11) holds on the whole block, Theorem 7.2 and Proposition 7.3 construct the stage and give (8.14)–(8.16), together with

$$d_{\mathbb{T}^2}(r_{k+1}(1)) < \theta_{\text{blk}}, \quad \text{BG}_N(r_{k+1}(1); \Lambda_*).$$

Thus the next start is admissible whenever the next block is controlled.

Let $\psi_k(\tau)$, $\psi_k(1) = \text{Id}$, be the stage DeTurck diffeomorphisms. Then

$$q_k(\tau) = \psi_k(\tau)^* p_k(\tau).$$

At the endpoint,

$$q_{k+1}(1) = 100^{-1}q_k(100) = (\Phi_k \circ \psi_k(100))^*(T_{k+1}^{-1}g(T_{k+1})).$$

Hence the induction continues with $\Phi_{k+1} = \Phi_k \circ \psi_k(100)$. On every controlled block define

$$g'(T_k\tau) = T_k q_k(\tau), \quad h(T_k\tau) = T_k r_k(\tau) \quad (1 \leq \tau < 100),$$

using the right-continuous restart convention for h , and set

$$\Psi(T_k\tau) = \Phi_k \circ \psi_k(\tau).$$

The endpoint identities make g' and Ψ continuous across controlled switching times, and

$$(8.22) \quad g'(t) = \Psi(t)^* g(t).$$

Step 3. Control the first full block directly.

From (8.12) and (8.21),

$$(8.23) \quad \text{diam}(M, g(S\tau)) \leq CD_1\sqrt{S} \leq CD_1\sqrt{S\tau}, \quad 1 \leq \tau \leq 100.$$

The stage estimate (8.15) makes $r_0(\tau)$ uniformly bilipschitz to the pullback $q_0(\tau)$. Since the quotient map for an invariant metric is distance nonincreasing and a metric circle has length twice its diameter,

$$(8.24) \quad L(S^1, h_{yy}(S\tau)) = \sqrt{S} L(S^1, (r_0)_{yy}(\tau)) \leq CD_1\sqrt{S} \leq CD_1\sqrt{S\tau}.$$

Thus both desired geometric bounds hold on the first controlled block without any hypothesis on its initial vertical energy or quotient length.

Step 4. Create uniform scalar data at the first restart.

Let

$$e_0(\tau) = \mathcal{E}_{\max}(r_0(\tau)), \quad \rho_0^+ = 100^{-1}r_0(100).$$

The energy decay on the first smooth stage gives

$$E(\rho_0^+) = 100e_0(100) \leq \frac{100e_0(1)}{1 + 99e_0(1)/2} < \frac{200}{99} < 3.$$

By (8.16) and (8.19),

$$\|r_1(1) - \rho_0^+\|_{C^1(M, \rho_0^+)} \leq C_{\text{st}}b_0 < \delta_J.$$

The transition-jump lemma therefore gives

$$(8.25) \quad E_1 := E(r_1(1)) \leq 4.$$

The stage and restart estimates also make $r_1(1)$ uniformly bilipschitz to $100^{-1}q_0(100)$. By (8.23), the latter has diameter at most CD_1 . Hence

$$(8.26) \quad \ell_1 := L(S^1, (r_1)_{yy}(1)) \leq CD_1.$$

Step 5. Propagate the scalar energy and length bounds.

For every controlled stage $k \geq 1$, define

$$e_k(\tau) = \mathcal{E}_{\max}(r_k(\tau)), \quad E_k = e_k(1),$$

and

$$\ell_k(\tau) = L(S^1, (r_k)_{yy}(\tau)), \quad \ell_k = \ell_k(1).$$

Constant scaling gives

$$(8.27) \quad L(S^1, h_{yy}(T_k\tau)) = \sqrt{T_k} \ell_k(\tau), \quad 1 \leq \tau < 100,$$

with the same formula for the left-hand value at $\tau = 100$. Integrating (8.1) and (8.2) gives

$$(8.28) \quad e_k(\tau) \leq \frac{E_k}{1 + \frac{1}{2}(\tau - 1)E_k}, \quad \ell_k(\tau) \leq \ell_k \left(1 + \frac{1}{2}(\tau - 1)E_k \right)^{1/2}.$$

In particular,

$$(8.29) \quad e_k(100) \leq \frac{E_k}{1 + \frac{99}{2}E_k}, \quad \ell_k(100) \leq \ell_k \left(1 + \frac{99}{2}E_k\right)^{1/2}.$$

Put $\rho_k^+ = 100^{-1}r_k(100)$. By (8.14),

$$(8.30) \quad b_k \leq 2^{-k}b_0, \quad C_{\text{st}}b_k < \delta_J.$$

The restart estimate gives

$$(8.31) \quad \|r_{k+1}(1) - \rho_k^+\|_{C^1(M, \rho_k^+)} \leq C_{\text{st}}b_k.$$

Constant scaling and (8.29) imply

$$E(\rho_k^+) = 100e_k(100) < 3, \quad L(\rho_k^+) = 10^{-1}\ell_k(100).$$

Applying Lemma 8.1 gives

$$(8.32) \quad E_{k+1} \leq 100e_k(100) + C_J C_{\text{st}}b_k, \quad \left| \log \frac{\ell_{k+1}}{10^{-1}\ell_k(100)} \right| \leq C_J C_{\text{st}}b_k.$$

Since $b_k \leq b_0$ and b_0 was chosen small, exponentiating the second estimate and enlarging one constant to C_2 gives

$$(8.33) \quad E_{k+1} \leq \frac{100E_k}{1 + 99E_k/2} + C_2b_k, \quad \ell_{k+1} \leq (1 + C_2b_k) \left(\frac{1 + 99E_k/2}{100} \right)^{1/2} \ell_k.$$

The sequence (b_k) is summable. Applying Lemma 8.2 to the shifted sequences beginning at $k = 1$, with initial data (8.25) and (8.26), gives

$$(8.34) \quad \sup_{k \geq 1} E_k < \infty, \quad \sup_{k \geq 1} \ell_k < \infty.$$

The bounds depend on D_1 only through (8.26).

Step 6. Return to physical time on all controlled blocks.

For $k \geq 1$, equations (8.28) and (8.34) give $\ell_k(\tau) \leq C\sqrt{\tau}$. Together with (8.27), this proves (8.17) on every such block; the first block was handled in (8.24). The same argument applies to either one-sided value at a switching time.

Proposition 7.3 gives $d_{\mathbb{T}^2}(r_k(\tau)) < \theta_{\text{blk}}$ on every controlled stage, so the physical retained fibers of $h(t)$ have diameter $O(\sqrt{t})$. The stage estimate (8.15) makes $g'(t)$ uniformly bilipschitz to $h(t)$. Applying Lemma 5.1 to the retained fibration, using the quotient-length and fiber bounds, gives

$$\text{diam}(M, g'(t)) \leq C\sqrt{t}$$

on every controlled block. Equation (8.22) and invariance of diameter under pullback give the same estimate for $g(t)$.

Step 7. Control a terminal partial block.

Let k_* be the first uncontrolled index, if one exists. The case $k_* = 0$ was handled at the beginning of Step 2. If $k_* > 0$, then $T_{k_*} \in I$, the preceding full block controls the metric at

T_{k*} , and connectedness gives $\sup I \leq T_{k*+1} = 100T_{k*}$. Applying (8.21) from T_{k*} extends the same $C\sqrt{t}$ diameter bound to $I \cap [T_{k*}, \infty)$. If no uncontrolled index exists, the controlled blocks cover $I \cap [S, \infty)$. This proves (8.18) and completes the proof. $\square$

Corollary 8.4 (Diameter and circle convergence on a persistent large-diameter tail). *Let M be $\mathbb{T}^3$, a compact Nil torus bundle with nontrivial unipotent monodromy, or a compact Sol torus bundle with monodromy $H \in \mathrm{SL}(2, \mathbb{Z})$ satisfying $\mathrm{Tr} H > 2$, and let $g(t)$, $t \geq 0$, be an immortal Ricci flow on M . Fix a particular application of Proposition 2.6, with cutoff c_0 and tail data $T_B(c_0), \varepsilon_B^{c_0}$, and suppose that*

$$\mathrm{diam}(M, g(t)) \geq c_0\sqrt{t}$$

for every sufficiently large t . Then there are constants $C < \infty$ and $\ell_\infty > 0$ such that

$$\mathrm{diam}(M, g(t)) \leq C\sqrt{t}$$

for all sufficiently large t , and

$$(8.35) \quad (M, t^{-1}g(t)) \longrightarrow C_{\ell_\infty}$$

in the Gromov–Hausdorff sense, where C_{ℓ_∞} is the metric circle of length ℓ_∞ .

Proof. Fix an integer $N \geq 2$ and $\sigma \in (0, 1)$. Choose S after the beginning of the persistent large-diameter tail and large enough for the averaged Bamler slice at S to satisfy the start condition (8.13) with $\delta_0 = c_0$; the high-order comparison also makes the average Riemannian. Apply Theorem 8.3 after taking this fixed-threshold application as the standing one, with

$$I = [S, \infty), \quad D_1 = S^{-1/2} \mathrm{diam}(M, g(S)).$$

Every block is controlled, and (8.18) gives the diameter conclusion.

We use the stage notation from the proof of Theorem 8.3. Thus $T_k = 100^k S$, the non-invariant errors satisfy $b_k \leq 2^{-k} b_0$, and, for $k \geq 1$,

$$E_k = \mathcal{E}_{\max}(r_k(1)), \quad \ell_k(\tau) = L(S^1, (r_k)_{yy}(\tau)), \quad \ell_k = \ell_k(1).$$

The quantitative one-block estimate (7.11), applied on every stage, gives

$$(8.36) \quad \ell_k(\tau) \geq c_0/100, \quad k \geq 1, \quad 1 \leq \tau \leq 100.$$

The sequences satisfy (8.33). Hence Lemma 8.2, together with (8.36), gives

$$(8.37) \quad \ell_k \longrightarrow \ell_\infty > 0, \quad E_k \longrightarrow 2.$$

The two-sided restart estimate (8.32) and $b_k \rightarrow 0$ then imply

$$(8.38) \quad 10^{-1} \ell_k(100) \longrightarrow \ell_\infty.$$

We next propagate the convergence through each smooth stage. Put

$$A_k(\tau) = 1 + \frac{1}{2}(\tau - 1)E_k, \quad B_k = 1 + \frac{99}{2}E_k.$$

The length estimate in (8.28) gives

$$\ell_k(\tau) \leq \ell_k A_k(\tau)^{1/2}.$$

Applying the same invariant energy and length inequalities with initial time τ gives

$$\ell_k(100) \leq \ell_k(\tau) \left(1 + \frac{1}{2}(100 - \tau)e_k(\tau)\right)^{1/2} \leq \ell_k(\tau) \left(\frac{B_k}{A_k(\tau)}\right)^{1/2},$$

where the last inequality uses the energy bound in (8.28). Consequently,

$$\ell_k(100) \left(\frac{A_k(\tau)}{B_k}\right)^{1/2} \leq \ell_k(\tau) \leq \ell_k A_k(\tau)^{1/2}.$$

Equations (8.37) and (8.38) squeeze the two sides uniformly for $1 \leq \tau \leq 100$, and give

$$(8.39) \quad \sup_{1 \leq \tau \leq 100} \left| \frac{\ell_k(\tau)}{\sqrt{\tau}} - \ell_\infty \right| \longrightarrow 0.$$

It remains to show that the retained fibers vanish in the same normalization. Let $t = T_k \tau$, and pull Bamler's comparison at time t into the assembled stage gauge by setting

$$\widehat{g}_{k,\tau}^B = T_k^{-1} \Psi(t)^* g_t^B, \quad \widehat{\pi}_{k,\tau}^B = \pi_t^B \circ \Psi(t).$$

The C^1 -part of Bamler's comparison, the stage estimate (8.15), and the monotonicity of $\varepsilon_B^{c_0}$ show, uniformly for $1 \leq \tau \leq 100$, that

$$(8.40) \quad \|\widehat{g}_{k,\tau}^B - r_k(\tau)\|_{C^1(M, r_k(\tau))} \longrightarrow 0.$$

The bilipschitz comparison and Bamler's fiber estimate also give

$$(8.41) \quad \sup_{1 \leq \tau \leq 100} \sup_{F \text{ a } \widehat{\pi}_{k,\tau}^B\text{-fiber}} \text{diam}_{\widehat{g}_{k,\tau}^B} F \longrightarrow 0.$$

Indeed, after scaling by T_k^{-1} , the right side of the original fiber estimate is at most a fixed multiple of $\varepsilon_B^{c_0}(T_k)\sqrt{\tau}$. The persistence conclusion in Lemma 7.1 says that every $\widehat{\pi}_{k,\tau}^B$ is fiber-equivalent to the retained fibration. Fixing an arbitrary $\theta > 0$, apply Lemma 5.3 with $\lambda = c_0/100$, using (8.36)–(8.41). For all sufficiently large k , it gives $d_{\mathbb{T}^2}(r_k(\tau)) < \theta$ simultaneously for all $\tau \in [1, 100]$. Thus

$$(8.42) \quad \sup_{1 \leq \tau \leq 100} d_{\mathbb{T}^2}(r_k(\tau)) \longrightarrow 0.$$

Finally, the identity

$$t^{-1} \Psi(t)^* g(t) = \tau^{-1} q_k(\tau)$$

shows that $(M, t^{-1}g(t))$ is isometric to $(M, \tau^{-1}q_k(\tau))$. The C^0 -part of (8.15), together with the uniform normalized diameter bound, implies

$$d_{\text{GH}}((M, \tau^{-1}q_k(\tau)), (M, \tau^{-1}r_k(\tau))) \longrightarrow 0$$

uniformly in τ . Lemma 3.4 gives

$$d_{\text{GH}}((M, \tau^{-1}r_k(\tau)), C_{\ell_k(\tau)/\sqrt{\tau}}) \leq \frac{d_{\mathbb{T}^2}(r_k(\tau))}{\sqrt{\tau}}.$$

Together with (8.39) and (8.42), this proves (8.35) uniformly on every late block, and hence for the full limit as $t \rightarrow \infty$. $\square$

9. DESCENT THROUGH FINITE COVERS

The main construction is performed on a fixed finite orientable torus-bundle cover, while the theorem concerns the original manifold. This section shows how the diameter estimates, collapse statements, smooth limits, and pointed universal-cover limits obtained on the working cover transfer to the original manifold through the fixed finite quotient.

The section is organized as follows. The first subsection, *Circle limits under finite quotients*, studies a continuous family of finite quotients whose covering spaces converge to a metric circle: equivariant compactness produces a limiting cyclic or dihedral action, the possible quotients are respectively circles or intervals of explicitly related length, and continuity rules out switching among different quotient types on the late-time tail. The second subsection, *Diameter and smooth descent*, first compares the diameters of a manifold and a fixed finite Riemannian cover, and then states the descent of normalized diameter bounds, point collapse, circle limits, pointed universal-cover limits, and fixed-coordinate smooth or exponential convergence. These results are collected here so that the Euclidean, Nil, and Sol case arguments can be carried out entirely on one working cover and descended only once in the final proof.

9.1. Circle limits under finite quotients. The following compactness lemma is used when a circle limit on a fixed finite cover is descended to the original manifold.

Lemma 9.1 (Finite quotients of a circle collapse). *Let Γ be a finite group acting on a compact space $\widehat{X}$, and let $\widehat{d}_t$, $t \geq T$, be a family of Γ -invariant length metrics that is continuous in the uniform norm on $\widehat{X} \times \widehat{X}$. Put $X_t = (\widehat{X}, \widehat{d}_t)/\Gamma$. Suppose that*

$$(\widehat{X}, \widehat{d}_t) \longrightarrow C_L$$

in the Gromov–Hausdorff sense, where C_L is a metric circle of length $L > 0$. Then X_t has a full Gromov–Hausdorff limit. The limit is either a metric circle of length L/m or a closed interval of length $L/(2m)$, for some positive integer m dividing $|\Gamma|$. If every limiting action of Γ on C_L is orientation preserving, then the limit is a circle.

Proof. We first identify the possible subsequential limits and include the finite-group compactness argument. Let $t_j \rightarrow \infty$, and choose Gromov–Hausdorff ε_j -approximations

$$f_j : (\widehat{X}, \widehat{d}_{t_j}) \longrightarrow C_L, \quad \varepsilon_j \longrightarrow 0.$$

Fix a countable dense set $A = \{a_1, a_2, \dots\} \subset C_L$. For every n choose $x_{j,n} \in \widehat{X}$ with $d_{C_L}(f_j(x_{j,n}), a_n) < 2\varepsilon_j$. Since C_L is compact and $\Gamma \times \mathbb{N}$ is countable, a diagonal subsequence may be chosen so that

$$f_j(\gamma x_{j,n})$$

converges for every $\gamma \in \Gamma$ and every n . Define its limit to be $\rho(\gamma)a_n$. The approximation property of f_j and the fact that γ is an isometry of $(\widehat{X}, \widehat{d}_{t_j})$ give

$$d_{C_L}(\rho(\gamma)a_n, \rho(\gamma)a_m) = d_{C_L}(a_n, a_m).$$

Thus $\rho(\gamma)$ extends uniquely to a distance-preserving map of C_L .

We record the convergence property needed to pass the group law and the quotient metric to the limit. If $z_j \in \widehat{X}$ and $f_j(z_j) \rightarrow z \in C_L$, then

$$(9.1) \quad f_j(\gamma z_j) \longrightarrow \rho(\gamma)z.$$

Indeed, approximate z by some a_n . The approximation property of f_j , the isometry of γ , and the already established convergence of $f_j(\gamma x_{j,n})$ show that the two sides of (9.1) are as close as desired. Applying this property twice and using the exact group relation on $\widehat{X}$ gives

$$\rho(\gamma\gamma') = \rho(\gamma)\rho(\gamma') \quad (\gamma, \gamma' \in \Gamma).$$

In particular, $\rho(\gamma^{-1})$ is the inverse of $\rho(\gamma)$, so each extension is an isometry and

$$\rho : \Gamma \longrightarrow \text{Isom}(C_L)$$

is a homomorphism.

If $f_j(x_j) \rightarrow x$ and $f_j(y_j) \rightarrow y$, then finiteness of Γ , (9.1), and the approximation property give

$$\begin{aligned} d_{X_{t_j}}([x_j], [y_j]) &= \min_{\gamma \in \Gamma} \widehat{d}_{t_j}(x_j, \gamma y_j) \\ &\longrightarrow \min_{\gamma \in \Gamma} d_{C_L}(x, \rho(\gamma)y) = d_{C_L/\rho(\Gamma)}([x], [y]). \end{aligned}$$

The convergence is uniform over pairs: otherwise a violating sequence of pairs would have a subsequence whose f_j -images converge, contradicting the displayed limit. Since the maps f_j are asymptotically onto, the induced correspondence between the two orbit spaces is asymptotically onto as well. Hence

$$(9.2) \quad X_{t_j} \longrightarrow C_L/\rho(\Gamma).$$

Every finite subgroup of $\text{Isom}(C_L) \cong O(2)$ is either a rotation group of order m or a dihedral group of order $2m$; here the case $m = 1$ in the second alternative is the group generated by one reflection. In the rotation case, the quotient in (9.2) is a circle of length L/m . In the dihedral case, it is a closed interval of length $L/(2m)$. In either case, m divides $|\Gamma|$. Hence only finitely many isometry classes can occur as subsequential limits; denote their set by $\mathcal{S}$.

It remains to rule out switching among these finitely many possibilities. On the fixed orbit set $\widehat{X}/\Gamma$, the quotient distance is

$$d_t^\Gamma([x], [y]) = \min_{\gamma \in \Gamma} \widehat{d}_t(x, \gamma y).$$

Consequently,

$$\|d_t^\Gamma - d_s^\Gamma\|_{C^0} \leq \|\widehat{d}_t - \widehat{d}_s\|_{C^0},$$

so $t \mapsto X_t$ is continuous in the Gromov–Hausdorff metric. The subsequential compactness just proved implies

$$\text{dist}_{\text{GH}}(X_t, \mathcal{S}) \longrightarrow 0;$$

otherwise a sequence staying a fixed positive distance from $\mathcal{S}$ would have a subsequence satisfying (9.2). If $\mathcal{S}$ has only one member, the full-limit assertion follows immediately. Otherwise, put

$$\Delta = \min\{d_{\text{GH}}(Y, Z) : Y, Z \in \mathcal{S}, Y \not\cong Z\} > 0$$

and choose $0 < r < \Delta/3$. The Gromov–Hausdorff balls $B_{\text{GH}}(Y, r)$, $Y \in \mathcal{S}$, are pairwise disjoint. For all sufficiently large t , the space X_t lies in their union. Since the interval of sufficiently large times is connected and $t \mapsto X_t$ is continuous, its image lies in one of these balls, say $B_{\text{GH}}(Y_*, r)$. The choice $r < \Delta/3$, together with $\text{dist}_{\text{GH}}(X_t, \mathcal{S}) \rightarrow 0$, then gives $d_{\text{GH}}(X_t, Y_*) \rightarrow 0$. This proves the full-limit assertion. Finally, if every limiting action is orientation preserving, its image is cyclic, so the selected limit is a circle. $\square$

9.2. Diameter and smooth descent. We first prove the elementary comparison used whenever the proof passes to a fixed finite cover.

Lemma 9.2 (Diameter under a finite Riemannian cover). *Let $\pi_{\text{fin}} : (M_{\text{fin}}, g_{\text{fin}}) \rightarrow (M, g)$ be a connected, d -sheeted Riemannian covering of closed connected manifolds. Then*

$$\text{diam}(M, g) \leq \text{diam}(M_{\text{fin}}, g_{\text{fin}}) \leq 2d \text{diam}(M, g).$$

Proof. The first inequality follows because the quotient map is distance nonincreasing. For the second, put $D = \text{diam}(M, g)$, fix $p \in M$, and enumerate the fiber

$$\pi_{\text{fin}}^{-1}(p) = \{p_1^{\text{fin}}, \dots, p_d^{\text{fin}}\}.$$

Every point $x^{\text{fin}} \in M_{\text{fin}}$ lies within distance D of one of the p_i^{fin} : project x^{fin} to M , choose a minimizing path from its projection to p , and lift this path starting at x^{fin} .

For $\varepsilon > 0$, the path-connected open balls $B(p_i^{\text{fin}}, D + \varepsilon)$ therefore cover the connected space M_{fin} . Their intersection graph is connected. Hence any two centers can be joined by a chain of at most $d - 1$ intersecting balls, and successive centers in the chain are at distance less than $2(D + \varepsilon)$. Two arbitrary points of M_{fin} are each within D of a center, so

$$\text{diam}(M_{\text{fin}}, g_{\text{fin}}) \leq 2D + 2(d - 1)(D + \varepsilon).$$

Letting $\varepsilon \downarrow 0$ proves the second inequality. $\square$

Proposition 9.3 (Finite-cover descent). *Let*

$$\pi_{\text{fin}} : M_{\text{fin}} \longrightarrow M$$

be a connected, d -sheeted finite regular cover of closed connected manifolds, with deck group Γ . Let $g(t)$, $t \geq T$, be a C^0 -continuous family of Riemannian metrics on M , let $g_{\text{fin}}(t) = \pi_{\text{fin}}^ g(t)$, and let $a(t) > 0$ be continuous. Put*

$$\widehat{g}(t) = a(t)g(t), \quad \widehat{g}_{\text{fin}}(t) = a(t)g_{\text{fin}}(t).$$

Then the following conclusions hold.

(1) *At every time,*

$$\text{diam}(M, \widehat{g}(t)) \leq \text{diam}(M_{\text{fin}}, \widehat{g}_{\text{fin}}(t)) \leq 2d \text{diam}(M, \widehat{g}(t)).$$

Consequently, any $O(1)$, $o(1)$, or two-sided positive diameter bound for either normalized family transfers to the other, with constants depending only on d .

(2) *If*

$$(M_{\text{fin}}, \widehat{g}_{\text{fin}}(t)) \longrightarrow \text{pt}$$

in the Gromov–Hausdorff sense, then $(M, \widehat{g}(t)) \rightarrow \text{pt}$.

(3) *Suppose that*

$$(M_{\text{fin}}, \widehat{g}_{\text{fin}}(t)) \longrightarrow C_L$$

for a metric circle of length $L > 0$. Then $(M, \widehat{g}(t))$ has a full Gromov–Hausdorff limit. It is either a metric circle of length L/m or a closed interval of length $L/(2m)$, for some positive integer m dividing $|\Gamma|$. The limit is a circle if every limiting action of Γ on C_L is orientation preserving; it is a closed interval if a limiting action contains an orientation-reversing isometry.

(4) *The manifolds M_{fin} and M have the same universal cover, and the lifts of $g_{\text{fin}}(t)$ and $g(t)$ to it agree. Hence every pointed Cheeger–Hamilton limit of one lifted family is the corresponding limit of the other.*

(5) *Suppose, for an integer $m \geq 0$, that $\widehat{g}_{\text{fin}}(t)$ converges in a fixed-coordinate C^m -norm to a Riemannian metric $g_{\text{fin}, \infty}$. Then the limit is Γ -invariant, descends to a metric g_∞ on M , and $\widehat{g}(t) \rightarrow g_\infty$ in C^m . A uniform exponential rate on the cover descends with the same exponent.*

Proof. Part 1 is Lemma 9.2, applied after the common constant scaling by $a(t)$. Convergence of a compact metric space to a point is equivalent to convergence of its diameter to zero, so part 2 follows from the first inequality.

For part 3, let $\widehat{d}_t$ be the distance function of $\widehat{g}_{\text{fin}}(t)$. The deck group acts isometrically for every t . Continuity of the metric family on the compact manifold implies continuity of $t \mapsto \widehat{d}_t$ in the uniform norm: nearby metrics are uniformly bilipschitz, and their distance functions are therefore uniformly close. Lemma 9.1 applies. If a limiting action contains an orientation-reversing isometry, its finite image in $\text{Isom}(C_L)$ is dihedral and its quotient is a closed interval; the full-limit assertion in that lemma rules out a different quotient type along another sequence.

Part 4 is immediate from covering-space theory and the identity between the two lifted metrics. Finally, every deck transformation preserves $\widehat{g}_{\text{fin}}(t)$. Passing to the fixed-coordinate limit shows that it preserves $g_{\text{fin}, \infty}$, which therefore descends. Pullback by a fixed finite covering identifies the local C^m -norms upstairs and downstairs, proving part 5 and its statement about rates. $\square$

10. THE EUCLIDEAN AND SOL CASES

This section completes the long-time analysis in the Euclidean and Sol cases. On the working cover, the Euclidean case has identity monodromy and is eventually forced into the

uniform flat-stability regime, while the Sol case has hyperbolic monodromy with positive eigenvalues, which rules out the almost-flat alternative and leaves a persistent large-diameter tail to which the block and quotient-length estimates apply. The resulting cover-level conclusions are exponential convergence to a flat metric in the Euclidean case, and $O(\sqrt{t})$ diameter together with convergence of the compact normalized flow to a metric circle in the Sol case.

The section is organized as follows. After stating the two case conclusions, the first subsection, *The Euclidean case*, proves that almost-flat slices occur arbitrarily far out, uses diameter normalization and near-flat compactness to enter a uniform flat-stability neighborhood on a finite torus cover, and obtains exponential convergence there. This descends directly to the working cover. The second subsection, *The Sol case*, excludes the almost-flat alternative by the non-virtual-nilpotence of the bundle group, applies the persistent-tail corollary to obtain both the matching $O(\sqrt{t})$ upper bound and the full metric-circle limit, and invokes the Type-III blowdown theorem only for the universal-cover Sol limit. The case conclusions are established on the working cover; the original manifold is reintroduced only in the proof of the main theorem.

Working cover. Throughout this section and Section 11, we use the single [working-cover convention](#): M denotes that fixed cover and $g(t)$ its lifted flow. The original manifold is reintroduced only in Section 12, where the conclusions are descended at once using Proposition 9.3. The only later finite cover is the stability cover supplied by Proposition 10.2 in the Euclidean proof. It is automatically regular because the working manifold there is a three-torus with abelian fundamental group, and it is used only to enter the flat-stability theorem and descend the resulting flat limit, not to alter the working monodromy. We use throughout $|\text{Rm}_{g(t)}| \leq C/t$, which follows from (2.5) in dimension three.

Theorem 10.1 (Euclidean and Sol long-time limits). *Let $g(t)$, $t \geq 0$, be an immortal Ricci flow on the orientable total space M of a $\mathbb{T}^2$ -bundle over S^1 , under the [working-cover convention](#). Let $H \in \text{SL}(2, \mathbb{Z})$ be its monodromy, as in Subsection 3.2. Assume either that $H = I$, so that $M \cong \mathbb{T}^3$ is in the Euclidean case, or that $\text{Tr } H > 2$, so that M is in the Sol case with positive-eigenvalue monodromy. Then the following hold.*

- (1) *If M has Euclidean type (equivalently, $H = I$), then for every $m \geq 0$, $g(t)$ converges exponentially fast in the C^m -norm of any fixed smooth background metric, as $t \rightarrow \infty$, to a flat metric on M .*
- (2) *If M has Sol type (equivalently, $\text{Tr } H > 2$ under the working-cover convention), then there are constants $0 < c_- < C < \infty$ such that*

$$c_- \sqrt{t} \leq \text{diam}(M, g(t)) \leq C \sqrt{t},$$

for all sufficiently large t . In addition, the Gromov–Hausdorff limit of $(M, t^{-1}g(t))$ is a circle. For every basepoint $\tilde{m} \in \tilde{M}$, the lifted blowdown flows

$$(\tilde{M}, \tilde{m}, s^{-1}\tilde{g}(s\tau)), \quad \tau > 0,$$

converge, as $s \rightarrow \infty$, in the pointed Cheeger–Hamilton sense to a pointed flow isometric to the Sol expanding soliton solution

$$(\mathbb{R}^3, 0, e^{-2z}dx^2 + e^{2z}dy^2 + 4\tau dz^2), \quad \tau > 0.$$

10.1. The Euclidean case. The following proposition is specific to the Euclidean argument. It packages the near-flat compactness step and the passage from compactness modulo diffeomorphisms to the fixed-coordinate stability neighborhood used below.

Proposition 10.2 (Near-flat entry into a uniform stability neighborhood). *Fix $\rho \in (0, 1)$. There exist an integer $m_{\text{flat}} \geq 2$, a number $\varepsilon_{\text{flat}} > 0$, a fixed model torus $\mathbb{T}_*^3$, a compact set $\mathcal{F}$ of flat metrics on $\mathbb{T}_*^3$, and a uniform $h^{2+\rho}$ -stability neighborhood $\mathcal{U}_{\mathcal{F}}$ of $\mathcal{F}$, as in Proposition 2.3, with the following property.*

Let q be a smooth metric on $\mathbb{T}^3$, put $D = \text{diam}(\mathbb{T}^3, q) > 0$, and suppose that

$$(10.1) \quad D^{j+2} \sup_{\mathbb{T}^3} |\nabla^j \text{Rm}_q|_q < \varepsilon_{\text{flat}}, \quad 0 \leq j \leq m_{\text{flat}}.$$

Then there is a connected finite covering map

$$p_{\text{st}} : \mathbb{T}_*^3 \longrightarrow \mathbb{T}^3$$

such that

$$(10.2) \quad p_{\text{st}}^*(D^{-2}q) \in \mathcal{U}_{\mathcal{F}}.$$

In particular, the Ricci flow starting from the metric in (10.2) converges exponentially in fixed coordinates to a flat metric.

Proof. Claim 2.16 of [18] gives the following diameter-normalized compactness statement. If q_i are metrics on $\mathbb{T}^3$ with diameter one and $\|\text{Rm}_{q_i}\|_{C^0} \rightarrow 0$, then, after passing to connected finite covers and applying diffeomorphisms of the covering tori, the lifted metrics are C^1 -close to a compact subset of the moduli space of flat three-tori. The paragraph immediately following that claim observes that, when scale-normalized curvature-derivative bounds are also available, the same argument gives closeness in any prescribed finite C^α -topology.

Choose an integer $\alpha > 2 + \rho$, and choose m_{flat} large enough for this C^α -compactness argument. To choose representatives, consider the finite-dimensional space of marked flat metrics on a fixed model torus $\mathbb{T}_*^3$. The mapping-class group $\text{GL}(3, \mathbb{Z})$ acts properly discontinuously on this space, and the quotient map to flat moduli is open. For each point of the compact subset of moduli obtained above, choose a lift and a relatively compact neighborhood of that lift. Finitely many of the quotient images cover the compact subset, so the union of the closures of the corresponding neighborhoods is a compact set $\mathcal{F}$ of marked flat metrics that maps onto it. Thus $\mathcal{F}$ is a compact set of flat metrics on $\mathbb{T}_*^3$ representing every limiting moduli class. Proposition 2.3 gives a uniform open stability neighborhood $\mathcal{U}_{\mathcal{F}}$ in the fixed $h^{2+\rho}$ -topology.

If no $\varepsilon_{\text{flat}}$ had the stated property, there would be a sequence of diameter-one metrics for which the left side of (10.1) tends to zero for every $0 \leq j \leq m_{\text{flat}}$, but for which no finite-cover pullback, after any identification of the covering torus with $\mathbb{T}_*^3$, belongs to $\mathcal{U}_{\mathcal{F}}$.

The preceding C^α -compactness gives finite covers and diffeomorphisms whose pulled-back metrics converge in C^α , hence in $h^{2+\rho}$, to $\mathcal{F}$. For large indices they lie in $\mathcal{U}_{\mathcal{F}}$, a contradiction.

For a general q , rescaling by D^{-2} makes the diameter one and turns (10.1) into the preceding scale-one hypotheses. The compactness argument initially produces a covering map $p_0 : \mathbb{T}_*^3 \rightarrow \mathbb{T}^3$ and a diffeomorphism $\chi : \mathbb{T}_*^3 \rightarrow \mathbb{T}_*^3$ such that $\chi^* p_0^*(D^{-2}q) \in \mathcal{U}_{\mathcal{F}}$. Replacing p_0 by $p_0 \circ \chi$ absorbs the diffeomorphism into the covering map and gives (10.2). The final assertion is Proposition 2.3. $\square$

The constants in Proposition 10.2 are independent of the cutoff in any application of Proposition 2.6. In the proof that follows, the analytic tolerance $\varepsilon_{\text{flat}}$ determines the separate normalized-diameter parameter δ_{flat} ; neither the auxiliary cutoff ζ nor the fixed global cutoff c_{af} is changed.

Proof of the Euclidean case. We first prove the scale-invariant statement

$$(10.3) \quad \liminf_{t \rightarrow \infty} t^{-1/2} \text{diam}(M, g(t)) = 0.$$

We now make a fresh application of Proposition 2.6, independent of the fixed global application with $c_0 = c_{\text{af}}$. This new application does not alter c_{af} , $T_B(c_{\text{af}})$, or $\varepsilon_B^{c_{\text{af}}}$. Fix an auxiliary cutoff $\zeta \in (0, 1)$ with $K\zeta^2 < \varepsilon_{\text{af}}$, apply the block machinery with $c_0 = \zeta$, and denote the corresponding tail data by $T_B(\zeta)$ and ε_B^ζ . If the ζ -small alternative did not occur at arbitrarily large times, then on some tail

$$\text{diam}(M, g(t)) \geq \zeta \sqrt{t}.$$

Corollary 8.4, applied with $c_0 = \zeta$, $T_B(\zeta)$, and ε_B^ζ , would give $\text{diam}(M, g(t)) = O(\sqrt{t})$. The Euclidean case of Proposition 2.1 would then imply

$$(M, t^{-1}g(t)) \longrightarrow \text{pt},$$

and hence $t^{-1/2} \text{diam}(M, g(t)) \rightarrow 0$, contradicting the displayed lower bound. Thus the ζ -small alternative occurs arbitrarily late for every such ζ , which proves (10.3).

Fix $\rho \in (0, 1)$, and let

$$m_{\text{flat}}, \quad \varepsilon_{\text{flat}}, \quad \mathbb{T}_*^3, \quad \mathcal{F}, \quad \mathcal{U}_{\mathcal{F}}$$

be the data in Proposition 10.2. By the [working-cover convention](#), $M \cong \mathbb{T}^3$ and $H = I$ already, so no preliminary cover is needed to reduce the monodromy to the identity.

The Type-III estimate on $[t_0/2, t_0]$ and Shi's derivative estimates give constants $C_j(K)$, independent of t_0 , such that

$$(10.4) \quad \sup_M |\nabla^j \text{Rm}|_{g(t_0)} \leq C_j(K) t_0^{-1-j/2}, \quad 0 \leq j \leq m_{\text{flat}},$$

whenever t_0 is sufficiently large. Choose $\delta_{\text{flat}} > 0$ so small that

$$(10.5) \quad C_j(K) \delta_{\text{flat}}^{j+2} < \varepsilon_{\text{flat}}, \quad 0 \leq j \leq m_{\text{flat}}.$$

By (10.3), choose a sufficiently large t_0 for which, with $D = \text{diam}(M, g(t_0))$,

$$(10.6) \quad \frac{D}{\sqrt{t_0}} < \delta_{\text{flat}}.$$

This choice is where the analytic tolerance $\varepsilon_{\text{flat}}$ enters; it is independent of both the auxiliary geometric cutoff ζ and the fixed global cutoff c_{af} .

Combining (10.4), (10.5), and (10.6) gives

$$D^{j+2} \sup_M |\nabla^j \text{Rm}|_{g(t_0)} < \varepsilon_{\text{flat}}, \quad 0 \leq j \leq m_{\text{flat}}.$$

Proposition 10.2, applied directly to $q = g(t_0)$, therefore gives a connected finite covering map

$$p_{\text{st}} : \mathbb{T}_*^3 \longrightarrow M$$

such that

$$(10.7) \quad p_{\text{st}}^*(D^{-2}g(t_0)) \in \mathcal{U}_{\mathcal{F}}.$$

The diffeomorphism arising in compactness modulo diffeomorphisms has already been absorbed into p_{st} in that proposition. Set

$$M_{\text{st}} = \mathbb{T}_*^3, \quad \pi_{\text{st}} = p_{\text{st}}, \quad g_{\text{st}}(t) = \pi_{\text{st}}^*g(t).$$

Define the restarted normalized flow

$$(10.8) \quad \widehat{g}_{\text{st}}(\tau) = D^{-2}g_{\text{st}}(t_0 + D^2\tau), \quad \tau \geq 0.$$

By uniqueness, this is the actual Ricci flow starting from the metric in (10.7). By Proposition 2.3, $\widehat{g}_{\text{st}}(\tau)$ converges exponentially in fixed coordinates to a flat metric $\widehat{g}_{\text{st},\infty}$. The fixed-background input there is [13, Theorem 3.7], which is stated for the actual Ricci flow. In its proof, Ricci–DeTurck convergence is established first, and [13, Proposition 3.6] transfers it to the ungauged flow, without a time-dependent pullback in the conclusion. Scaling back gives exponential convergence of $g_{\text{st}}(t)$ to $D^2\widehat{g}_{\text{st},\infty}$ in every fixed C^m -norm.

Since the fundamental group of M is abelian, the connected finite cover $\pi_{\text{st}} : M_{\text{st}} \rightarrow M$ is already regular. Part 5 of Proposition 9.3 therefore applies directly: the flat limit $D^2\widehat{g}_{\text{st},\infty}$ descends to a flat metric on M , and $g(t)$ converges to it exponentially in every fixed C^m -norm. $\square$

10.2. The Sol case.

Proof of the Sol case. Use the fixed global application of Proposition 2.6 with $c_0 = c_{\text{af}}$, tail time $T_B(c_{\text{af}})$, and error $\varepsilon_B^{c_{\text{af}}}$. Its almost-flat alternative cannot occur at any sufficiently large time. Indeed, if it occurred, then (2.9) and the almost-flat topological theorem, Proposition 2.2, would imply that $\pi_1(M)$ is virtually nilpotent. This is impossible in the Sol case. Indeed, $\pi_1(M)$ has a finite-index subgroup that is a cocompact lattice in the Lie group Sol, which has exponential volume growth. By the volume-growth comparison for compact quotients [22] and invariance of growth type under passage to finite-index subgroups, $\pi_1(M)$ therefore has exponential growth. On the other hand, $\pi_1(M) = \mathbb{Z}^2 \rtimes_H \mathbb{Z}$ is polycyclic, so if

it were virtually nilpotent then it would have polynomial growth by [23, Theorem 4.3(1)]. Hence it is not virtually nilpotent.

Consequently, the torus-bundle case occurs for every sufficiently large time, and

$$\text{diam}(M, g(t)) \geq c_{\text{af}} \sqrt{t}.$$

Corollary 8.4, applied with $c_0 = c_{\text{af}}$, gives the matching upper bound and a number $\ell_\infty > 0$ such that

$$(M, t^{-1}g(t)) \longrightarrow C_{\ell_\infty}.$$

Since the working cover has Thurston type Sol, Proposition 2.1 gives the pointed universal-cover blowdown to the Sol expanding soliton. Its weaker classification of subsequential compact limits in the Sol case is not used to obtain the full circle limit. $\square$

11. THE NIL CASE

The Nil case, with nontrivial unipotent monodromy, is more delicate because the flow can move repeatedly between large-diameter and almost-flat regimes. The late-time axis is therefore decomposed into connected components of the normalized-diameter superlevel set. Theorem 8.3 was formulated to handle exactly such a component: it starts from one sufficiently late Bamler slice, retains that fibration on all complete blocks in the component, creates uniform scalar data after the first restart, propagates the quotient-length bound, and controls a terminal partial block by the global scale-invariant curvature estimate. Thus no separate Nil-specific iteration is needed.

11.1. Large-diameter components and almost-flat gaps. For a finite large-diameter component, choose a starting time S just to the right of its left endpoint. The short interval before S is controlled directly by continuity of the normalized diameter. The complete geometric blocks and the terminal partial block are then controlled by Theorem 8.3; the complementary gaps lie in the almost-flat alternative. The schematic decomposition is shown in Figure 5. If the component has no first uncontrolled block, the controlled blocks continue indefinitely.




FIGURE 5. Schematic decomposition of a finite large-diameter component in the Nil case. Theorem 8.3 begins at $S = T_0$ and runs on every full block contained in the component; its diameter conclusion also covers the terminal partial block. If T_{k*} is the first uncontrolled start time, connectedness gives $t_I^+ \leq T_{k*+1} = 100T_{k*}$, so Type-III distortion controls the terminal interval.

11.2. The global diameter bound and Nil blowdown.

Theorem 11.1 (Nil diameter bound and blowdown). *Let $g(t)$, $t \geq 0$, be an immortal Ricci flow on the orientable total space M of a $\mathbb{T}^2$ -bundle over S^1 with nontrivial unipotent monodromy, under the [working-cover convention](#). Then there is a constant $C < \infty$ such that*

$$\text{diam}(M, g(t)) \leq C\sqrt{t},$$

for all sufficiently large t . In addition,

$$\lim_{t \rightarrow \infty} (M, t^{-1}g(t)) = \text{pt}$$

in the Gromov–Hausdorff sense. On every compact τ -interval in $(0, \infty)$, the lifted blowdown flows on the universal cover converge as $s \rightarrow \infty$ in the pointed Cheeger–Hamilton sense to a pointed flow isometric to

$$\left(\mathbb{R}^3, 0, \frac{1}{3\tau^{1/3}} \left(dx + \frac{1}{2}y dz - \frac{1}{2}z dy \right)^2 + \tau^{1/3}(dy^2 + dz^2) \right).$$

Proof. Fix $N \geq 2$ and $\sigma \in (0, 1)$, and set

$$\mathcal{D}(t) = t^{-1/2} \text{diam}(M, g(t)).$$

Let η_0 and T_p be the constants in Theorem 8.3, applied with the fixed global application $c_0 = c_{\text{af}}$ and $\delta_0 = c_{\text{af}}$. Choose $T_* \geq T_p$ so large that, at every torus-bundle slice $t \geq T_*$, averaging the parabolically rescaled metric in Bamler’s fibration produces a Riemannian metric whose $C^{N, \sigma}$ distance from the rescaled metric is less than η_0 , as required by (8.13). Define

$$\mathcal{L} = \{t \geq T_* : \mathcal{D}(t) \geq c_{\text{af}}\}.$$

On $\mathcal{L}^c$, the desired estimate holds with constant c_{af} , and the slices are in the almost-flat alternative in Proposition 2.6. At every time in $\mathcal{L}$, the torus-bundle comparison is available by the other case of the same dichotomy. It remains to control each connected component of $\mathcal{L}$.

Let I be a nontrivial component with finite left endpoint $t_I^- > T_*$. By continuity, $\mathcal{D}(t_I^-) = c_{\text{af}}$. Choose $S \in I$ close enough to t_I^- that

$$\mathcal{D}(t) \leq 2c_{\text{af}} \quad \text{for } t \in I \cap [t_I^-, S].$$

At time S , use the torus-bundle fibration and average in that same fibration. By the choice of T_* , the resulting start metric satisfies the smallness hypothesis of Theorem 8.3. Apply that theorem with the fixed global application $c_0 = c_{\text{af}}$, with $\delta_0 = c_{\text{af}}$ and $D_1 = 2c_{\text{af}}$. Its constants are uniform over all such components. The short interval $[t_I^-, S] \cap I$ is already controlled by $\mathcal{D} \leq 2c_{\text{af}}$. If I has a finite right endpoint t_I^+ , then $\mathcal{D}(t_I^+) = c_{\text{af}}$; the component theorem already includes that endpoint. Singleton components also satisfy $\mathcal{D} = c_{\text{af}}$.

There may be one component with left endpoint T_* . If it is bounded, it does not affect the large-time estimate. If it is unbounded, choose one sufficiently large S in it and set $D_1 = \mathcal{D}(S) < \infty$. Applying Theorem 8.3 with the fixed global application $c_0 = c_{\text{af}}$ at that S

controls the component from S onward; the constant may depend on this single finite value of D_1 . The bounded interval before S is irrelevant.

Combining the estimates on $\mathcal{L}$ and $\mathcal{L}^c$ gives

$$\text{diam}(M, g(t)) = O(\sqrt{t}).$$

Together with the Type-III curvature estimate, this is precisely the hypothesis of Proposition 2.1. Its Nil case gives the Gromov–Hausdorff collapse of $(M, t^{-1}g(t))$ to a point and the pointed Cheeger–Hamilton blowdown to the Nil expanding soliton displayed in the statement. $\square$

12. PROOF OF THE MAIN THEOREM

The purpose of the section is to return from the fixed working cover to the original manifold and assemble the Euclidean, Nil, and Sol conclusions into the three parts of the main theorem.

Proof of Theorem 1.1. For this proof only, write M_0 for the manifold denoted by M in the statement. Apply Proposition 3.1 once and fix the finite regular orientable torus-bundle cover

$$\pi_{\text{fin}} : M_{\text{fin}} \longrightarrow M_0$$

supplied by its stronger conclusion. Let $g_{\text{fin}}(t) = \pi_{\text{fin}}^* g(t)$. By construction, $(M_{\text{fin}}, g_{\text{fin}}(t))$ satisfies the [working-cover convention](#). Apply Proposition 2.6 directly to this lifted flow with the fixed global choice $c_0 = c_{\text{af}}$. This cover is fixed for the remainder of the proof, and each case conclusion is descended only after the corresponding cover-level theorem has been applied.

The Euclidean case. The Euclidean conclusion of Theorem 10.1 gives exponential convergence of $g_{\text{fin}}(t)$ to a flat metric. Part 5 of Proposition 9.3 descends the limit and all fixed-order exponential estimates to M_0 .

The C^2 -convergence estimate gives $|\text{Rm}_{g_{\text{fin}}(t)}| = O(e^{-\lambda t})$ for some $\lambda > 0$, while the diameter remains bounded. Proposition 2.1 therefore gives point collapse of $(M_{\text{fin}}, t^{-1}g_{\text{fin}}(t))$ and the static Euclidean blowdown on the universal cover. Parts 2 and 4 of Proposition 9.3 give the corresponding conclusions on M_0 . This proves part (1).

The Nil case. Applying Theorem 11.1 to $g_{\text{fin}}(t)$ gives point collapse on the compact cover and the Nil expanding-soliton blowdown on the common universal cover. Parts 2 and 4 of Proposition 9.3 descend these conclusions. This proves part (2).

The Sol case. The Sol conclusion of Theorem 10.1 gives a metric-circle limit on the compact cover and the Sol blowdown on the common universal cover. Parts 3 and 4 of Proposition 9.3 show that the limit on M_0 is a metric circle or a compact interval of positive length; they also preserve the universal-cover limit. This proves part (3). $\square$

APPENDIX A. UNIFORM LIFTED PARABOLIC ESTIMATES

This appendix provides the analytic foundation that makes the main iteration uniform in the collapsing regime. All collapse-uniform estimates are formulated on the universal cover with constants controlled by finite-order lifted geometry, rather than by the injectivity radius of the compact quotient. It develops the linear and quasilinear Schauder theory and establishes the continuous-dependence, positive-time smoothing, and regular-restart results invoked in Sections 4 and 6.

Standard inputs and deductions made here. The imported analytic ingredients are the harmonic-radius estimate of Anderson–Cheeger [1, Theorem 0.3 and Section 1], the harmonic-coordinate elliptic identity and regularity theorem of DeTurck–Kazdan [11, Lemma 4.1 and Theorem 4.5(b)], the scalar parabolic estimates from Lieberman itemized in Remark A.5, and DeTurck’s compact-manifold local existence argument [10, Theorem and proof, pp. 158–161]. The fixed time-slab atlas, its uniform overlap and cutoff bounds, the passage to scalar-principal tensor systems, the lifted patching estimates, the finite-order induction, and the collapse-independent restart bounds are the deductions carried out in this appendix.

The appendix is organized as follows. The opening unnumbered subsection, *Estimates used in the body*, displays the precise estimates imported by the body and the chain from a bounded initial slice to uniform parabolic geometry, linear control, stage stability, and regular restart. The subsection *Parabolic Hölder spaces and lifted bounded geometry* fixes the Euclidean and intrinsic parabolic norms, the metric-compatible time derivative, and the uniformly local convention on the universal cover. The subsection *From slice geometry to parabolic coordinate systems* constructs a fixed-coordinate harmonic atlas with uniform overlap and cutoff bounds, proves equivalence of intrinsic and chart norms, and propagates a BG_m initial-face bound to a full PBG_m bound on a finite stage. The subsection *Scalar-principal systems and coefficient regularity* derives the difference equation for two Ricci–DeTurck solutions and gives the finite derivative counts for its principal, first-order, and zeroth-order coefficients. The subsection *Uniform Schauder estimates on lifted cylinders* proves the required finite-order Cauchy–Dirichlet estimate, including the initial and artificial boundary faces, and patches the local estimates to global and interior collapse-independent estimates on the universal cover. The subsection *Quasilinear derivative gain* obtains one positive-time derivative at a time for Ricci–DeTurck systems by differentiating the equation and applying the linear theory on nested cylinders. The final subsection, *Continuation, smoothing, and regular restart*, closes the uniform continuous-dependence bootstrap, uses a fixed-quotient continuation criterion only after all collapse-independent bounds have been established, gains the five endpoint derivatives needed by the body, and verifies the rescaled averaged restart geometry. The spaces X_s^m , Y_s^m , BG_m , and PBG_m were introduced in Section 4 and are used here with the precise conventions given below.

Estimates used in the body. The body uses the following estimates from the appendix. All intrinsic norms are evaluated after lifting to the universal cover. Whenever a coordinate estimate is used, PBG provides fixed lifted charts whose norms are uniformly equivalent

to the intrinsic ones; the resulting constants are therefore independent of collapse of the compact quotient.

Result	Hypotheses	Conclusion and use in the body
Lemma A.3	A stage-start bound $\text{BG}_m(s(a); \Lambda_0)$, a sectional-curvature bound, and a bounded time interval.	A full parabolic bound $\text{PBG}_m(s; \Lambda_1; [a, b])$; this provides the fixed lifted charts for every stage.
Proposition A.7	PBG_m , a finite tensor system with scalar principal part, and controlled coefficients.	Collapse-independent global and interior Schauder estimates; used for translation differences and endpoint smoothing.
Proposition A.11	PBG_N and a sufficiently small $C^{N,\sigma}$ initial perturbation.	Existence on the whole stage and a linear X_s^N stability estimate; used in stage preparation.
Proposition 6.4	Stage closeness in X_r^N and curvature control for the locally $\mathbb{T}^2$ -invariant background.	A $C^{N+5,\sigma}$ -small rescaled endpoint, an averaged restart, and uniform BG_N ; used at every switch.

The estimates are related by

$$\text{BG}_m \text{ at the initial face} \xrightarrow{\text{Lemma A.3}} \text{PBG}_m \text{ on the stage,}$$

$$\text{PBG}_m + \text{scalar principal part} \xrightarrow{\text{Proposition A.7}} \text{uniform linear control,}$$

uniform linear control + quasilinear bootstrap $\implies$ stage stability and regular restart.

Lemma A.1 constructs the fixed-coordinate atlas, and Lemma A.9 provides the finite number of positive-time derivative gains. The remainder of the appendix proves each implication displayed above.

A.1. Parabolic Hölder spaces and lifted bounded geometry. We first fix the analytic conventions used throughout the paper. Spatial Hölder spaces on a time slice are denoted by $C^{m,\sigma}$. On a Euclidean parabolic cylinder Q , put

$$d_p((x, t), (y, s)) = \max\{|x - y|, |t - s|^{1/2}\}.$$

For a scalar function, or for a component of a tensor in a fixed frame, set

$$[v]_{\sigma;Q}^p = \sup_{P \neq P' \in Q} \frac{|v(P) - v(P')|}{d_p(P, P')^\sigma},$$

$$\langle v \rangle_{\theta,t;Q} = \sup_{\substack{(x,t),(x,s) \in Q \\ t \neq s}} \frac{|v(x, t) - v(x, s)|}{|t - s|^{\theta/2}}.$$

For an integer $m \geq 0$ and $\sigma \in (0, 1)$, define

$$\begin{aligned}
 \|v\|_{C^{m+\sigma, (m+\sigma)/2}(Q)} &= \sum_{|\nu|+2j \leq m} \|\partial_x^\nu \partial_t^j v\|_{C^0(Q)} \\
 &+ \sum_{|\nu|+2j=m} [\partial_x^\nu \partial_t^j v]_{\sigma; Q}^p \\
 &+ \sum_{|\nu|+2j=m-1} \langle \partial_x^\nu \partial_t^j v \rangle_{1+\sigma, t; Q},
 \end{aligned}
 \tag{A.1}$$

where the last sum is empty when $m = 0$. For a finite tensor system, the norm is the sum of the component norms in the chosen fixed frame. This is Lieberman's $H_{m+\sigma}$ -norm from [16, Chapter IV, Section 1], written in the $C^{m+\sigma, (m+\sigma)/2}$ notation used here. Thus one spatial derivative has parabolic weight one and one time derivative has weight two.

Norm convention. The spaces X_s^m and Y_s^m are defined intrinsically, before any bounded-geometry hypothesis is imposed. Lift $s(t)$ and a tensor U to the universal cover. For a time-dependent vector field V , define the metric-compatible covariant time derivative

$$\mathcal{D}_t^s V = \partial_t V + \frac{1}{2} s(t)^{-1} (\partial_t s(t)) V,
 \tag{A.2}$$

and extend $\mathcal{D}_t^s$ to arbitrary tensor fields by duality and the Leibniz rule. Its time-parallel transport $\mathsf{T}_{u \rightarrow t}^s(x)$ is an isometry from the tensor norm induced by $s(u)$ at x to that induced by $s(t)$.

Set

$$r_s = \frac{1}{8} \min \left\{ 1, \inf_{(x,t) \in \bar{M} \times [a,b]} \text{inj}_{\bar{s}(t)}(x) \right\}.
 \tag{A.3}$$

For the fixed smooth family on the compact time interval, the infimum is positive. This canonical choice also has a uniform positive lower bound whenever a PBG bound is imposed. If V is a time-dependent tensor, let

$$[V]_{\alpha; s}^x = \sup_t \sup_{0 < d_{s(t)}(x, y) < r_s} \frac{|V(x, t) - \mathsf{P}_{y \rightarrow x}^{s(t)} V(y, t)|_{s(t)}}{d_{s(t)}(x, y)^\alpha},
 \tag{A.4}$$

$$\langle V \rangle_{\alpha; s}^t = \sup_x \sup_{u \neq t} \frac{|V(x, t) - \mathsf{T}_{u \rightarrow t}^s(x) V(x, u)|_{s(t)}}{|t - u|^{\alpha/2}},
 \tag{A.5}$$

$$[V]_{\alpha; s}^p = [V]_{\alpha; s}^x + \langle V \rangle_{\alpha; s}^t.
 \tag{A.6}$$

Here $\mathsf{P}_{y \rightarrow x}^{s(t)}$ denotes parallel transport along the unique short $s(t)$ -geodesic. Replacing the numerical factor $1/8$ in (A.3) by any fixed smaller positive factor gives an equivalent norm.

For an integer $m \geq 0$, define

$$\begin{aligned}
 \|U\|_{\mathcal{C}_s^{m+\sigma, (m+\sigma)/2}[a,b]} &= \sum_{k+2j \leq m} \|(\nabla^{s(t)})^k (\mathcal{D}_t^s)^j U\|_{C_s^0} \\
 &+ \sum_{k+2j=m} [(\nabla^{s(t)})^k (\mathcal{D}_t^s)^j U]_{\sigma;s}^p \\
 &+ \sum_{k+2j=m-1} \langle (\nabla^{s(t)})^k (\mathcal{D}_t^s)^j U \rangle_{1+\sigma;s}^t,
 \end{aligned}
 \tag{A.7}$$

where the last sum is empty when $m = 0$, and C_s^0 denotes the space-time supremum measured with $s(t)$. Taking the supremum on the universal cover defines the deck-invariant norm. The spaces in (4.5) are

$$\|U\|_{X_s^m[a,b]} = \|U\|_{\mathcal{C}_s^{m+\sigma, (m+\sigma)/2}[a,b]}, \quad \|U\|_{Y_s^m[a,b]} = \|U\|_{\mathcal{C}_s^{m-2+\sigma, (m-2+\sigma)/2}[a,b]}.$$

Thus they are well defined for every smooth reference flow, even before a uniform atlas is specified.

For a tensor on the universal cover that need not be deck-invariant, we retain

$$\|U\|_{C_{\text{uloc}}^{m+\sigma, (m+\sigma)/2}}$$

for the supremum of the Euclidean component norms (A.1) in a fixed lifted parabolic atlas. This notation is used for translation differences only after the required PBG bound has been established.

A.2. From slice geometry to parabolic coordinate systems.

Lemma A.1 (Uniform parabolic atlas from bounded geometry). *Fix $n, m \geq 2, \sigma \in (0, 1), \Lambda \geq 1$, and $L > 0$. Let $s(t)$, $a \leq t \leq b$, be a Ricci flow with $b - a \leq L$, and suppose that the inequalities in (4.2) hold with this Λ . Then there are $r_* > 0$ and $A_* < \infty$, depending only on n, m, σ, Λ, L , with the following properties. For every $(\tilde{x}_0, t_0)$, there are harmonic coordinates for $\tilde{s}(t_0)$ on $B_{\tilde{s}(t_0)}(\tilde{x}_0, 4r_*)$. Keeping the spatial coordinates fixed, they remain valid on*

$$B_{\tilde{s}(t_0)}(\tilde{x}_0, 4r_*) \times ([t_0 - (4r_*)^2, t_0 + (4r_*)^2] \cap [a, b]).$$

Parabolic estimates may therefore be applied on backward, forward, or endpoint-truncated subcylinders contained in this set. In these coordinates the coefficients of $\tilde{s}$, $\tilde{s}^{-1}$, and the Christoffel symbols have the finite parabolic Hölder bounds through order m needed below, all bounded by A_ . On overlaps, the transition maps and their inverses have uniform finite-order parabolic Hölder bounds through the same order. Radius- $r_*/4$ one-sided or truncated subcylinders can be chosen to cover space-time, while the radius- $2r_*$ cylinders have uniformly bounded overlap and admit spatial cutoffs with uniform derivative bounds.*

Proof. Choose $p > 2n$ so large that $1 - 2n/p > \sigma$. The local $L^{1,p}$ harmonic-radius estimate of Anderson–Cheeger [1, Theorem 0.3 and Section 1, pp. 268–274], applied on the universal cover after rescaling on the injectivity-radius scale, gives a radius $r_h > 0$ on which harmonic

coordinates have uniform ellipticity and uniform $W^{1,p}$ -bounds. Qualitative existence and regularity of the harmonic coordinate functions are given by [11, Lemma 1.2 and Corollary 1.4]; the radius and uniform constants used here come from the Anderson–Cheeger estimate and the quantitative bootstrap below. In harmonic coordinates the metric satisfies the uniformly elliptic system

$$s^{ab}\partial_a\partial_b s_{ij} = -2\text{Ric}_{ij}(s) + \mathcal{Q}_{ij}(s^{-1}, \partial s).$$

The first interior $W^{2,p/2}$ -estimate gives uniform $C^{1,\alpha}$ -control for some $\alpha \in (\sigma, 1)$. The bounds for $\nabla^j \text{Rm}$, converted inductively to coordinate bounds using the already controlled Christoffel symbols, and repeated interior elliptic estimates for this system then bootstrap the center-time coefficients through the finite order needed here. This is the quantitative harmonic-coordinate regularity argument of [11, Lemma 4.1 and Theorem 4.5(b), pp. 254–256]. All of these estimates are on the universal cover and depend only on n, m, σ, Λ .

Since $|\text{Ric}(s)|_s \leq n|\text{Rm}(s)|_s$ and $\partial_t s = -2\text{Ric}(s)$, the curvature bound gives $|\partial_t s|_s \leq 2n\Lambda$. Choose $r_* < r_h/16$ and put $\delta_* = (4r_*)^2$, decreasing r_* so that $e^{2n\Lambda\delta_*} \leq 4/3$. It follows that

$$e^{-2n\Lambda|t-t_0|}s(t_0) \leq s(t) \leq e^{2n\Lambda|t-t_0|}s(t_0).$$

Thus a harmonic chart chosen once at time t_0 on the $s(t_0)$ -ball of radius $8r_*$ is kept fixed as a spatial chart throughout $[t_0 - \delta_*, t_0 + \delta_*] \cap [a, b]$; the metric comparison ensures that all cylinders in the statement remain inside its original spatial domain. No time-dependent coordinate change is made.

For completeness, put $\ell_* = (r_*/4)^2$ and divide $[a, b]$ into at most $2 + L/\ell_*$ closed time intervals J_ν , each of length at most ℓ_* , with center time t_ν . At each level choose a maximal $r_*/8$ -separated set in the universal cover and harmonic charts of radius $8r_*$ at its points. The radius- $r_*/8$ balls at the center time cover space. Keep every chart fixed on $I_\nu = [t_\nu - \delta_*, t_\nu + \delta_*] \cap [a, b]$. Since $J_\nu \subset I_\nu$, the products of the covering balls with J_ν are covered by a fixed number of radius- $r_*/4$ backward, forward, or endpoint-truncated parabolic cylinders in these fixed charts. They therefore cover space-time; call them the selected subcylinders.

Curvature bounds, the injectivity-radius lower bound, volume comparison, and the disjointness of the radius- $r_*/16$ balls about the chosen centers give a uniform spatial overlap bound for the radius- $2r_*$ cylinders. Only a uniformly bounded number of the intervals I_ν meet any given time, so the space-time overlap is uniform as well. The spatial families may be infinite on the universal cover, but they are uniformly locally finite; none of these constants involves the compact quotient.

It remains to establish the coefficient bounds in these fixed coordinates. At t_ν they follow from the preceding harmonic-coordinate bootstrap. On I_ν , differentiate $\partial_t s = -2\text{Ric}(s)$ in the fixed coordinates. More explicitly, the Ricci-flow evolution and commutation formulas used in Shi’s derivative estimates [21] express, for $j \geq 1$, $\partial_x^\mu \partial_t^j s$ in terms of lower-weight coordinate derivatives of s and $\nabla^q \text{Rm}(s)$ with $q \leq |\mu| + 2j - 2$. The center-time harmonic-coordinate bounds, induction on $|\mu| + 2j$, and Grönwall for $j = 0$ therefore give

$$\sup |\partial_x^\mu \partial_t^j s_{kl}| \leq C \quad \text{whenever } |\mu| + 2j \leq m + 2.$$

The same bounds hold for s^{-1} , and the formula for the Christoffel symbols gives the required bounds for them. The bounds through parabolic weight $m + 2$ provide the extra spatial and time derivatives needed in (A.1); Euclidean mean-value inequalities then give its Hölder seminorms.

It remains to justify the corresponding bounds for transition maps. The time projection of each selected subcylinder lies in one of the intervals J_ν . If two selected subcylinders overlap, their associated intervals are equal or adjacent, so their center times differ by at most $2\ell_* < \delta_*$; each of the two fixed charts is therefore valid at the center time of the other. Let x be a chart chosen at time t_ν and y one chosen at time t_μ . On a uniform neighborhood of such an overlap, each component of $y \circ x^{-1}$ satisfies, in the x -coordinates, the scalar equation

$$\Delta_{\tilde{s}(t_\mu)} y^\alpha = 0,$$

because the y -coordinates are harmonic for $\tilde{s}(t_\mu)$. The coefficients of this equation in the x -chart are uniformly elliptic and have the finite-order bounds already proved. The radius- $8r_*$ chart domains and the radius- $2r_*$ enlargements of the selected subcylinders leave a uniform interior margin, and the harmonic-coordinate normalization gives a uniform C^0 -bound for the coordinate functions. Interior elliptic Schauder estimates, iterated a finite number of times, therefore give uniform spatial Hölder bounds for $y \circ x^{-1}$ through the required order on the overlap. Interchanging x and y gives the same bounds for the inverse transition map. Since both coordinate systems are kept fixed in time, these transition maps have zero time derivatives, and the spatial estimates are precisely the required parabolic Hölder bounds. Thus the overlap-map bounds are not an additional compactness theorem: they follow from scalar interior elliptic regularity on the uniform interior margins just constructed.

Finally, no quotient-dependent partition-of-unity estimate is used. Choose one fixed Euclidean cutoff on the model chart, rescale it by the already fixed radius r_* , and pull it back through each harmonic chart. The coefficient bounds above then give uniform spatial derivative bounds for these cutoffs. This proves every assertion of the lemma, with constants independent of collapse downstairs. $\square$

Lemma A.2 (Equivalence of intrinsic and fixed-chart norms). *Fix $n, m \geq 2, \sigma \in (0, 1), \Lambda \geq 1$, and $L > 0$. Suppose that $\text{PBG}_m(s; \Lambda; [a, b])$ holds and $b - a \leq L$. Let $\|\cdot\|_{X_{s,\text{fc}}^m[a,b]}$ and $\|\cdot\|_{Y_{s,\text{fc}}^m[a,b]}$ denote the suprema of the Euclidean component norms in the fixed lifted atlas supplied by Lemma A.1. Then there is $C = C(n, m, \sigma, \Lambda, L)$ such that every deck-invariant tensor U satisfies*

$$(A.8) \quad C^{-1} \|U\|_{X_s^m[a,b]} \leq \|U\|_{X_{s,\text{fc}}^m[a,b]} \leq C \|U\|_{X_s^m[a,b]},$$

$$(A.9) \quad C^{-1} \|U\|_{Y_s^m[a,b]} \leq \|U\|_{Y_{s,\text{fc}}^m[a,b]} \leq C \|U\|_{Y_s^m[a,b]}.$$

The analogous comparison holds for every spatial $C^{k,\sigma}$ -norm with $0 \leq k \leq m$. For tensors on the universal cover that are not deck-invariant, the same local comparison holds with the fixed-chart side interpreted as the uniformly local norm.

Proof. In the atlas of Lemma A.1, the metric and inverse metric, the Christoffel symbols, the coefficients of the time connection $\mathcal{D}_t^s$, the transition maps, and the cutoffs have uniform finite-order parabolic Hölder bounds. Covariant space-time derivatives are therefore linear combinations of coordinate derivatives of the same or lower parabolic order, with uniformly controlled coefficients; the converse follows by solving the same triangular relations for the coordinate derivatives. Parallel transport along the short spatial geodesics and the time transport of (A.2) differ from componentwise identification by uniformly controlled lower-order terms. The spatial and temporal Hölder seminorms are consequently equivalent in both directions. Taking suprema over the uniformly locally finite atlas proves (A.8)–(A.9); the proof is local and is therefore unchanged for a tensor that is not deck-invariant. $\square$

Lemma A.3 (Slice-to-parabolic bounds). *Fix $n, m, \sigma \in (0, 1)$, and $\Lambda_0, K_0, L > 0$, with $m \geq 2$ and $\Lambda_0 \geq 1$. There is $\Lambda_1 < \infty$, depending only on these data, with the following property. Let $s(t)$, $a \leq t \leq b$, be a Ricci flow on a compact n -manifold, with $b - a \leq L$, such that*

$$\text{BG}_m(s(a); \Lambda_0), \quad \sup_{M \times [a, b]} |\text{sec}_{s(t)}| \leq K_0.$$

Then

$$\text{PBG}_m(s; \Lambda_1; [a, b]).$$

The constant is independent of the injectivity radius of the compact quotient.

Proof. All curvature estimates may be lifted to the universal cover; since the quantities being estimated are deck-invariant, the maximum-principle arguments can be carried out on the compact quotient. The curvature bound gives $|\text{Ric}| \leq c_n K_0$, and hence

$$(A.10) \quad e^{-c_n K_0 L} s(a) \leq s(t) \leq e^{c_n K_0 L} s(a) \quad (a \leq t \leq b).$$

We first establish the curvature-derivative bounds at the lower time face. For $j \geq 1$, the standard differentiated curvature evolution equation has the schematic maximum-principle form

$$\begin{aligned} (\partial_t - \Delta) |\nabla^j \text{Rm}|^2 &\leq - |\nabla^{j+1} \text{Rm}|^2 + C_j |\text{Rm}| |\nabla^j \text{Rm}|^2 \\ &\quad + C_j \sum_{\substack{p+q=j \\ p, q < j}} |\nabla^p \text{Rm}| |\nabla^q \text{Rm}| |\nabla^j \text{Rm}|. \end{aligned}$$

Starting with the bounds in $\text{BG}_m(s(a); \Lambda_0)$, induction on j and the maximum principle give uniform estimates for $|\nabla^j \text{Rm}|$, $0 \leq j \leq m + 3$, on a fixed initial interval whose length depends only on n, m, Λ_0, K_0 . On the part of $[a, b]$ separated from a by half that interval, Shi's derivative estimates [21], applied to the curvature bound on $[a, b]$, give the same finite collection of estimates. If $b - a$ is shorter than the fixed initial interval, the first argument already covers the whole flow. Thus

$$(A.11) \quad \sup_{M \times [a, b]} |\nabla^j \text{Rm}_{s(t)}| \leq A_1, \quad 0 \leq j \leq m + 3,$$

for a constant A_1 depending only on the stated data.

It remains to make the lifted injectivity estimate uniform in time. At time a , the injectivity and curvature bounds in $\text{BG}_m(s(a); \Lambda_0)$ give numbers $\rho, v_0 > 0$, depending only on n, Λ_0 , such that

$$\text{vol}_{\tilde{s}(a)} B_{\tilde{s}(a)}(\tilde{x}, r) \geq v_0 r^n \quad (0 < r \leq \rho).$$

For example, one may take a fixed fraction of Λ_0^{-1} for ρ and use normal-coordinate volume comparison. Put $A = e^{c_n K_0 L}$. The metric comparison (A.10) gives

$$B_{\tilde{s}(a)}(\tilde{x}, A^{-1/2} \rho) \subset B_{\tilde{s}(t)}(\tilde{x}, \rho), \quad d\mu_{\tilde{s}(t)} \geq A^{-n/2} d\mu_{\tilde{s}(a)}.$$

Thus the $\tilde{s}(t)$ -volume of the radius- ρ ball is at least $v_0 A^{-n} \rho^n$, uniformly for $t \in [a, b]$. Together with the sectional-curvature bound, the Cheeger–Gromov–Taylor estimate [8, Section 4] gives

$$(A.12) \quad \inf_{\tilde{M} \times [a, b]} \text{inj}_{\tilde{s}(t)} \geq i_1 > 0,$$

where i_1 again depends only on the fixed data.

Equations (A.11) and (A.12) give the curvature and injectivity-radius part of PBG_m . Applying Lemma A.1 and increasing one constant to include its radius, coefficient, and time-Hölder bounds gives $\text{PBG}_m(s; \Lambda_1; [a, b])$. This also makes explicit that the coordinate control at the initial time comes from the propagated initial curvature derivatives, whereas the later-time control comes from Shi's estimates. $\square$

A.3. Scalar-principal systems and coefficient regularity. We use the Ricci–DeTurck equation (4.3) and the parabolic spaces in (4.5), both introduced in Section 4.

Lemma A.4 (Difference equation for two Ricci–DeTurck solutions). *Let $s(t)$ be a Ricci-flow background and let $v_0(t)$ and $v_1(t)$ be two positive definite solutions of the Ricci–DeTurck equation relative to $s(t)$ on the same time interval. Put*

$$U = v_1 - v_0, \quad v_\vartheta = v_0 + \vartheta(v_1 - v_0), \quad 0 \leq \vartheta \leq 1,$$

and assume that every v_ϑ is positive definite. With ∇ the connection of $s(t)$, the difference satisfies a homogeneous linear system

$$(A.13) \quad \partial_t U - P^{pq} \nabla_p \nabla_q U = B^p \nabla_p U + CU,$$

where the second-order part is diagonal in the tensor components: in a fixed frame, every component U^A is acted on by the same operator $P^{pq} \nabla_p \nabla_q$, and no second derivative of a different component occurs. More precisely, for

$$\mathcal{Q}_s(v) = -2 \text{Ric}(v) + \mathcal{L}_{W(v,s)} v,$$

one has the structural formula

$$(A.14) \quad \begin{aligned} \mathcal{Q}_s(v) = & v^{pq} \nabla_p \nabla_q v + v^{-1} * v^{-1} * \nabla v * \nabla v \\ & + \mathcal{R}(\text{Rm}(s), v, v^{-1}), \end{aligned}$$

and the mean-value identity

$$(A.15) \quad \begin{aligned} \partial_t U &= \mathcal{Q}_s(v_1) - \mathcal{Q}_s(v_0) \\ &= \int_0^1 D\mathcal{Q}_s|_{v_\vartheta}[U] d\vartheta. \end{aligned}$$

Differentiating (A.14) gives

$$(A.16) \quad D\mathcal{Q}_s|_{v_\vartheta}[Z] = v_\vartheta^{pq} \nabla_p \nabla_q Z + \mathcal{B}_\vartheta^p \nabla_p Z + \mathcal{C}_\vartheta Z,$$

with

$$(A.17) \quad \begin{aligned} \mathcal{B}_\vartheta &= v_\vartheta^{-1} * v_\vartheta^{-1} * \nabla v_\vartheta, \\ \mathcal{C}_\vartheta &= v_\vartheta^{-1} * v_\vartheta^{-1} * \nabla^2 v_\vartheta + v_\vartheta^{-1} * v_\vartheta^{-1} * v_\vartheta^{-1} * \nabla v_\vartheta * \nabla v_\vartheta \\ &\quad + \mathcal{R}_1(\text{Rm}(s), v_\vartheta, v_\vartheta^{-1}). \end{aligned}$$

Here $*$ denotes universal tensor contractions, and $\mathcal{R}$ and $\mathcal{R}_1$ are algebraic in the displayed arguments; in particular, $\mathcal{R}_1$ contains no derivatives of v_ϑ . The coefficients in (A.13) are

$$P^{pq} = \int_0^1 v_\vartheta^{pq} d\vartheta, \quad B^p = \int_0^1 \mathcal{B}_\vartheta^p d\vartheta, \quad C = \int_0^1 \mathcal{C}_\vartheta d\vartheta.$$

Consequently, uniform metric equivalence and finite-order bounds for v_0, v_1 , and s give the corresponding ellipticity and Hölder bounds for P, B, C .

Proof. Formula (A.14) is the usual covariant DeTurck calculation [10]; all second derivatives of the unknown metric occur in its scalar principal term. The fundamental theorem of calculus gives (A.15). Differentiating the principal term gives

$$D(v^{pq} \nabla_p \nabla_q v)[Z] = v^{pq} \nabla_p \nabla_q Z - (v^{-1} Z v^{-1})^{pq} \nabla_p \nabla_q v,$$

while

$$\begin{aligned} D(v^{-1} * v^{-1} * \nabla v * \nabla v)[Z] &= v^{-1} * v^{-1} * \nabla v * \nabla Z \\ &\quad + v^{-1} * v^{-1} * v^{-1} * Z * \nabla v * \nabla v. \end{aligned}$$

The curvature term varies only algebraically in Z . This proves (A.16)–(A.17); integrating in ϑ proves (A.13) and the last assertion. $\square$

Remark A.5 (What is imported from Lieberman). The scalar analytic inputs are the interpolation inequality [16, Proposition 4.2, p. 49], the interior Schauder estimate [16, Theorem 4.9, p. 59], and the second-order Cauchy–Dirichlet estimate [16, Theorem 4.28, p. 77]. The initial–lateral corner compatibility condition for a cylindrical domain is stated explicitly in [16, p. 78]. The finite-order version for the scalar-principal systems used here is proved in Lemma A.6, following the spatial differentiation procedure indicated in [16, Exercise 4.5(b), p. 84]. For a localized cylinder $Q^+ = \Omega \times [a, a + \theta]$, the *artificial spatial boundary* means $\partial\Omega \times [a, a + \theta]$; it is introduced by localization and is not a boundary of M . Theorem 4.28 is used below only after localization to such a fixed smooth spatial cylinder.

A localization point used throughout is worth recording here. Each spatial cutoff is chosen identically zero in a collar of the artificial spatial boundary. After subtraction of an extension of the initial trace, the localized unknown has zero data on both the initial and lateral faces, while the cutoff-generated forcing vanishes to all orders at their intersection. Hence the standard Cauchy–Dirichlet corner compatibility conditions are automatic; no additional compatibility condition is imposed on the original Cauchy data away from this artificial corner.

The passage from the scalar estimates to the finite tensor systems used below is carried out in this paper. In a fixed frame, *scalar principal part* means that the same uniformly parabolic operator $\partial_t - a^{ij}\partial_i\partial_j$ acts componentwise, with the coefficient matrix a^{ij} carrying no tensor-component indices.

A.4. Uniform Schauder estimates on lifted cylinders.

Lemma A.6 (Finite-order Cauchy–Dirichlet estimate for scalar-principal systems). *Fix $m \geq 2$ and $\sigma \in (0, 1)$. Let $Q_i^+ = \Omega_i \times [a, a + \theta]$, where $\Omega_1 \Subset \Omega_2$ are smooth Euclidean domains. Suppose that a smooth finite-dimensional vector-valued function Z satisfies on Q_2^+*

$$\partial_t Z - a^{ij}\partial_i\partial_j Z = \mathbf{B}^i\partial_i Z + \mathbf{C}Z + G,$$

where $a^{ij} = a^{ji}$ is uniformly parabolic and $\partial_t - a^{ij}\partial_i\partial_j$ acts componentwise with the same coefficient matrix on every component; only $\mathbf{B}^i$ and $\mathbf{C}$ may couple components. The coefficients $a^{ij}, \mathbf{B}^i, \mathbf{C}$ have uniformly bounded $C^{m-2+\sigma, (m-2+\sigma)/2}$ -norms. Assume that $Z = 0$ on the initial face, and that Z and G vanish in a fixed spatial collar of $\partial\Omega_2 \times [a, a + \theta]$. Then

$$\|Z\|_{C^{m+\sigma, (m+\sigma)/2}(Q_1^+)} \leq C \left(\|G\|_{C^{m-2+\sigma, (m-2+\sigma)/2}(Q_2^+)} + \|Z\|_{C^0(Q_2^+)} \right),$$

where C depends only on the displayed coefficient and parabolicity bounds, m, σ, θ , the rank, and the fixed domains $\Omega_1 \Subset \Omega_2$.

The componentwise second-order estimate and interpolation inequality are imported from Lieberman; the nested-domain absorption and the finite-order spatial-differentiation induction for the coupled lower-order terms are the arguments supplied here.

Proof. Write $I = [a, a + \theta]$ and $L_0 = \partial_t - a^{ij}\partial_i\partial_j$. We first record the localized second-order estimate that will be used at every induction step. Suppose that a finite system V , with zero initial trace, satisfies

$$L_0 V = \mathbf{A}^i\partial_i V + \mathbf{D}V + H$$

on $\Omega'' \times I$, where $\Omega' \Subset \Omega'' \Subset \Omega_2$ and the displayed coefficients have uniform $C^{\sigma, \sigma/2}$ -bounds. Multiply by a spatial cutoff which is one on Ω' and zero in a collar of $\partial\Omega''$. The localized unknown has zero data on the initial face and on the artificial spatial boundary. The scalar estimate [16, Theorem 4.28, p. 77], applied componentwise, together with [16, Proposition 4.2, p. 49], gives, for every $\varepsilon > 0$,

$$(A.18) \quad \begin{aligned} \|V\|_{C^{2+\sigma, 1+\sigma/2}(\Omega' \times I)} &\leq C_\varepsilon \left(\|H\|_{C^{\sigma, \sigma/2}(\Omega'' \times I)} + \|V\|_{C^0(\Omega'' \times I)} \right) \\ &\quad + \varepsilon \|V\|_{C^{2+\sigma, 1+\sigma/2}(\Omega'' \times I)}. \end{aligned}$$

Here the terms involving derivatives of the cutoff, as well as the first- and zeroth-order couplings, are estimated by interpolation; C_ε also depends on the fixed separation of Ω' from $\partial\Omega''$.

For clarity, the last term in (A.18) can be removed without assuming an outer a priori estimate. Choose a smoothly nested family $\Omega(r)$, $r_0 \leq r < r_*$, with $\Omega_1 \Subset \Omega(r_0)$ and $\overline{\Omega(r_*)} \Subset \Omega_2$, and set

$$\Phi(r) = \|V\|_{C^{2+\sigma, 1+\sigma/2}(\Omega(r) \times I)}, \quad A = \|H\|_{C^{\sigma, \sigma/2}(Q_2^+)} + \|V\|_{C^0(Q_2^+)}.$$

Fix $\varepsilon > 0$. Keeping track of cutoff derivatives in (A.18) gives, for $r < R < r_*$,

$$\Phi(r) \leq C_\varepsilon (R - r)^{-p} A + \varepsilon \Phi(R)$$

for a fixed p . With $r_j = r_* - (r_* - r_0)\tau^j$, choose $\varepsilon\tau^{-p} < 1/2$ and iterate. Then

$$\Phi(r_0) \leq CA \sum_{j=0}^{J-1} (\varepsilon\tau^{-p})^j + \varepsilon^J \Phi(r_J).$$

The last term tends to zero as $J \rightarrow \infty$, since V is smooth on the fixed compact cylinder $\overline{\Omega(r_*)} \times I$. Thus (A.18) yields the second-order estimate with no outer top-order term.

We now perform the finite-order induction. Choose smooth domains

$$\Omega_1 \Subset D_{m-2} \Subset D_{m-3} \Subset \cdots \Subset D_0 \Subset D_{-1} \Subset \Omega_2,$$

with the evident interpretation when $m = 2$. For $0 \leq \ell \leq m - 2$, let V_ℓ be the finite vector consisting of all $\partial^\alpha Z$ with $|\alpha| = \ell$. Differentiating the equation spatially and collecting equal-order derivatives gives

$$(A.19) \quad L_0 V_\ell = \mathbf{B}_\ell^i \partial_i V_\ell + \mathbf{C}_\ell V_\ell + H_\ell.$$

The first-order coefficient $\mathbf{B}_\ell$ is assembled from $\mathbf{B}$ and ∂a . The zeroth-order coefficient $\mathbf{C}_\ell$ is assembled from $\mathbf{C}$, $\partial \mathbf{B}$, and, when $\ell \geq 2$, $\partial^2 a$. All remaining commutators are in H_ℓ . Schematically,

$$(A.20) \quad \begin{aligned} H_\ell = & \partial^\ell G + \sum_{3 \leq |\beta| \leq \ell} \partial^\beta a * \partial^{\ell-|\beta|+2} Z \\ & + \sum_{2 \leq |\beta| \leq \ell} \partial^\beta \mathbf{B} * \partial^{\ell-|\beta|+1} Z + \sum_{1 \leq |\beta| \leq \ell} \partial^\beta \mathbf{C} * \partial^{\ell-|\beta|} Z, \end{aligned}$$

where each sum includes the usual multiindex splittings and is omitted when its lower limit exceeds ℓ . In particular, every derivative of Z in H_ℓ has order at most $\ell - 1$. The assumed $C^{m-2+\sigma, (m-2+\sigma)/2}$ -bounds therefore give uniform $C^{\sigma, \sigma/2}$ -bounds for $\mathbf{B}_\ell$ and $\mathbf{C}_\ell$.

For $\ell = 0$, the localized second-order estimate on $D_0 \Subset D_{-1}$ gives

$$\|V_0\|_{C^{2+\sigma, 1+\sigma/2}(D_0 \times I)} \leq C \left(\|G\|_{C^{\sigma, \sigma/2}(Q_2^+)} + \|Z\|_{C^0(Q_2^+)} \right).$$

Suppose inductively that, for every $0 \leq j \leq \ell - 1$, we have proved

$$\|V_j\|_{C^{2+\sigma, 1+\sigma/2}(D_j \times I)} \leq C \left(\|G\|_{C^{j+\sigma, (j+\sigma)/2}(Q_2^+)} + \|Z\|_{C^0(Q_2^+)} \right).$$

Formula (A.20), the parabolic product estimates, and the induction hypothesis then give

$$\|H_\ell\|_{C^{\sigma,\sigma/2}(D_{\ell-1} \times I)} + \|V_\ell\|_{C^0(D_{\ell-1} \times I)} \leq C \left(\|G\|_{C^{\ell+\sigma,(\ell+\sigma)/2}(Q_2^+)} + \|Z\|_{C^0(Q_2^+)} \right).$$

Applying the localized second-order estimate to (A.19) on $D_\ell \Subset D_{\ell-1}$ advances the induction. At every step the localized unknown has zero initial trace and vanishes on the artificial spatial boundary. The full induction uses only the time-jet identities through $r_m = \lfloor m/2 \rfloor$; after ℓ spatial differentiations, only those with $\ell + 2j \leq m$ occur. Since the cutoff is zero in a spatial collar, the localized unknown, forcing, and every jet in this finite recursion vanish at the artificial corner. Hence the required compatibility identities are automatic. The original hypotheses that Z and G vanish near $\partial\Omega_2$ give the same conclusion if the outer boundary itself is used.

At $\ell = m - 2$ we have controlled all spatial derivatives of Z through order m , with the required parabolic Hölder seminorms. The remaining mixed derivatives are recovered recursively from

$$\partial_t Z = a^{ij} \partial_i \partial_j Z + B^i \partial_i Z + CZ + G.$$

Indeed, if $|\nu| + 2j \leq m$, apply $\partial_x^\nu \partial_t^{j-1}$ to this identity. Every coefficient derivative that occurs has parabolic weight at most $m - 2$, while every derivative of Z on the right either has fewer time derivatives or has already been controlled by the spatial induction. Induction on j , together with the parabolic product estimates, gives the sup norms, the σ -seminorms at weight m , and the time seminorms at weight $m - 1$ appearing in (A.1). Restriction from D_{m-2} to Ω_1 proves the stated estimate. This is the finite-system version of the spatial differentiation procedure indicated in [16, Exercise 4.5(b), p. 84]. $\square$

Proposition A.7 (Uniform lifted linear Schauder estimate). *Fix $m \geq 2$ and $\sigma \in (0, 1)$. Let $s(t)$, $a \leq t \leq b$, be a Ricci flow satisfying $\text{PBG}_m(s; \Lambda; [a, b])$ for some $\Lambda \geq 1$. Let E be a fixed finite-rank tensor bundle and consider, on the universal cover and in the charts of Lemma A.1, a system*

$$(A.21) \quad \mathcal{P}U = \partial_t U - a^{ij} \nabla_i \nabla_j U - B^i \nabla_i U - CU = F.$$

Assume that the second-order part is diagonal in the bundle components and uses the same coefficient matrix a^{ij} on every component. For fixed constants $0 < \lambda_- \leq \lambda_+ < \infty$, suppose that

$$\lambda_- s^{ij} \xi_i \xi_j \leq a^{ij} \xi_i \xi_j \leq \lambda_+ s^{ij} \xi_i \xi_j,$$

and that the local $C^{m-2+\sigma, (m-2+\sigma)/2}$ -norms of a, B, C are bounded by A_0 .

Every smooth bounded solution on the universal cover satisfies, for $\theta > 0$ with $a + \theta \leq b$,

$$(A.22) \quad \|U\|_{C_{\text{uloc}}^{m+\sigma, (m+\sigma)/2}(\widetilde{M} \times [a+\theta, b], s)} \leq C_\theta \left(\|U\|_{C^0(\widetilde{M} \times [a, b], s)} + \|F\|_{C_{\text{uloc}}^{m-2+\sigma, (m-2+\sigma)/2}} \right).$$

For the global estimate, assume in addition that $\mathcal{P}$ is deck-invariant, equivalently that the coefficient fields a, B, C descend to M . If U and F are deck-invariant and descend to tensors

u and f on M , with $u(a) \in C^{m,\sigma}$ and $f \in Y_s^m[a, b]$, then

$$(A.23) \quad \|u\|_{X_s^m[a,b]} \leq C \left(\|u(a)\|_{C^{m,\sigma}(M,s(a))} + \|f\|_{Y_s^m[a,b]} \right).$$

The constants depend only on $m, \sigma, \Lambda, \lambda_-, \lambda_+, A_0$, on θ in the interior estimate, and on an upper bound for $b - a$. The interior and initial-face estimates are based, respectively, on Lieberman's interior Schauder theorem [16, Theorem 4.9, p. 59] and his Cauchy–Dirichlet Schauder theorem [16, Theorem 4.28, p. 77]. All atlas-radius, overlap, and cutoff constants, fixed-coordinate ellipticity constants, localized initial-face boundary-chart constants, and coefficient Hölder constants are uniform over the admissible family specified by PBG_m and the displayed bounds. In particular, the constants are independent of collapse of the compact quotient.

Proof. Fix one lifted parabolic chart, choose concentric cylinders $Q_0 \subseteq Q_1 \subseteq Q_2$ whose radii are fixed fractions of the atlas radius, and choose a spatial cutoff χ supported in Q_2 with $\chi = 1$ on Q_1 . In the fixed tensor frame, (A.21) reads

$$\partial_t U^A - a^{ij} \partial_i \partial_j U^A = B^{iA}{}_D \partial_i U^D + C^A{}_D U^D + F^A,$$

where $A, D \in \{1, \dots, \text{rank } E\}$, repeated D -indices are summed, and the terms from the reference connection have been included in B and C . The base-order interior estimate follows by the same localization, interpolation, and radius-iteration argument used to remove the outer top-order term in the proof of Lemma A.6: apply the scalar interior estimate [16, Theorem 4.9, p. 59] componentwise to χU^A , estimate the first- and zeroth-order couplings and the cutoff terms with [16, Proposition 4.2, p. 49], and iterate over concentric cylinders. This gives

$$\|U\|_{C^{2+\sigma, 1+\sigma/2}(Q_0)} \leq C \left(\|F\|_{C^{\sigma, \sigma/2}(Q_2)} + \|U\|_{C^0(Q_2)} \right).$$

For the higher-order estimate, repeat the spatial-differentiation induction in the proof of Lemma A.6, with the interior estimate above in place of the localized Cauchy–Dirichlet estimate. On a fixed chain of nested cylinders, collect the components $\partial^\alpha U^A$, $|\alpha| = \ell$, into one finite system. Differentiation preserves the scalar principal part; the terms with one derivative on a^{ij} become first-order couplings, and the terms linear in $\partial^\alpha U$ become zeroth-order couplings. The remaining commutators contain derivatives of F of order at most ℓ and derivatives of U of order at most $\ell + 1$. The coefficient bounds, the induction hypothesis, and the same interpolation and radius-iteration argument therefore advance the induction through $\ell = m - 2$. The last paragraph of the proof of Lemma A.6, applied to the present equation, then recovers the mixed time derivatives of parabolic weight at most m . This proves the full $C^{m+\sigma, (m+\sigma)/2}$ -bound on Q_0 . The argument up to this point uses only fixed-coordinate uniform parabolicity, the stated coefficient bounds, and the cylinder separations, and therefore proves the local estimate (A.25). The PBG_m hypothesis is used next only to make these local data uniform over the lifted atlas.

Taking the supremum over the atlas of Lemma A.1 proves (A.22). This part of the argument does not require U to descend to the compact quotient.

Suppose now that $\mathcal{P}$, U , and F are deck-invariant. The equation for $|u|_s^2$, the coefficient bounds, the scalar maximum principle on the compact quotient, and Grönwall's inequality give

$$\|u\|_{C^0(M \times [a,b],s)} \leq C \left(\|u(a)\|_{C^0(M,s(a))} + \|f\|_{C^0} \right).$$

For cylinders meeting the initial time slice $t = a$, write $Q_i^+ = Q_i \cap \{t \geq a\}$, choose the spatial sections Ω_i to be concentric Euclidean balls, choose the same type of spatial cutoff χ , and put $L_0 = \partial_t - a^{ij} \partial_i \partial_j$ in the fixed frame. The artificial spatial boundary is the one defined in Remark A.5. If $u_0^A(x) = U^A(x, a)$, define the time-independent extension

$$E^A(x, t) = \chi(x) u_0^A(x), \quad Z^A = \chi U^A - E^A.$$

Then $Z = 0$ on the initial face $t = a$ and on the artificial spatial boundary, and its equation is

$$(A.24) \quad L_0 Z^A = \chi F^A + \chi B^{iA}{}_D \partial_i U^D + \chi C^A{}_D U^D + [L_0, \chi] U^A + a^{ij} \partial_i \partial_j E^A.$$

Using $\chi U = Z + E$, rewrite the two coupling terms as

$$\chi B^i \partial_i U = B^i \partial_i Z + B^i \partial_i E - B^i (\partial_i \chi) U, \quad \chi C U = C Z + C E.$$

Thus the terms involving Z have the first- and zeroth-order form in Lemma A.6. Denote by $\mathcal{G}$ the sum of the remaining terms, which form its forcing. The remainder of the initial-face argument is the finite-order induction from the proof of Lemma A.6, with two points requiring comment. First, every term involving E is bounded by $\|u(a)\|_{C^{m,\sigma}}$. Second, derivatives of the cutoff terms in $\mathcal{G}$ contain derivatives of U of order at most $\ell + 1$ at the ℓ -th spatial-differentiation step. This is one order below the $C^{\ell+2+\sigma, (\ell+2+\sigma)/2}$ -norm being estimated. Lieberman's interpolation inequality [16, Proposition 4.2, p. 49] therefore bounds these terms by an arbitrarily small multiple of the corresponding norm on the next larger cylinder, plus a constant times $\|U\|_{C^0}$; the nested-radius iteration used in the proof of Lemma A.6 absorbs the former term.

The compatibility check is also the same as in that proof. With $r_m = \lfloor m/2 \rfloor$, the required identities are $(\partial_t^j Z)(\cdot, a)|_{\partial\Omega_2} = 0$ for $0 \leq j \leq r_m$, where the higher time jets are generated recursively from (A.24); after ℓ spatial differentiations, only those with $\ell + 2j \leq m$ are used. Since χ , and hence E , Z , and $\mathcal{G}$, vanish in a full collar of the artificial boundary, every jet in this recursion vanishes there. Thus all required corner identities are automatic. Repeating the induction in the proof of Lemma A.6, using the componentwise Cauchy–Dirichlet estimate [16, Theorem 4.28, p. 77] for its localized second-order step and then recovering the mixed time derivatives from the equation as in the lemma's final paragraph gives

$$\begin{aligned} \|U\|_{C^{m+\sigma, (m+\sigma)/2}(Q_1^+)} &\leq C \left(\|F\|_{C^{m-2+\sigma, (m-2+\sigma)/2}(Q_2^+)} + \|u(a)\|_{C^{m,\sigma}} \right. \\ &\quad \left. + \|U\|_{C^0(Q_2^+)} \right). \end{aligned}$$

Taking the supremum of these initial-face and interior estimates and using the preceding maximum-principle bound for the C^0 -term proves (A.23). No sum over quotient charts occurs, which is why the constants are uniform under collapse. $\square$

Remark A.8. The proof also gives the following purely local estimate. Let $Q_0 \Subset Q_2$ be fixed nested coordinate cylinders in one chart, with a fixed positive spatial separation and lower-time gap; their upper time faces may agree. Suppose that, in a fixed tensor frame on Q_2 , the system is

$$\partial_t U^A - a^{ij} \partial_i \partial_j U^A = B^{iA}{}_D \partial_i U^D + C^A{}_D U^D + F^A,$$

where a^{ij} is common to all components and is uniformly parabolic with fixed-coordinate constants λ_- and λ_+ , and the $C^{m-2+\sigma, (m-2+\sigma)/2}$ -norms of a, B, C are at most A_1 . Then

$$(A.25) \quad \|U\|_{C^{m+\sigma, (m+\sigma)/2}(Q_0)} \leq C \left(\|U\|_{C^0(Q_2)} + \|F\|_{C^{m-2+\sigma, (m-2+\sigma)/2}(Q_2)} \right).$$

Here C depends only on $m, \sigma, \lambda_-, \lambda_+, A_1$, the rank and dimension, and the fixed cylinder-separation data. This local estimate requires neither PBG_m , a covering space, nor deck invariance.

A.5. Quasilinear derivative gain.

Lemma A.9 (Local one-derivative gain for Ricci–DeTurck systems). *Fix an integer $k \geq 2$ and $\sigma \in (0, 1)$. Let Q and Q' be concentric coordinate cylinders in a fixed spatial chart, with Q' having smaller spatial radius and a positive time gap from the lower time face of Q ; the two cylinders may share their upper time face. Let $r(t)$ be a smooth reference metric on Q , and let $q(t)$ be a smooth positive definite solution of the Ricci–DeTurck equation relative to $r(t)$. In a fixed tensor frame this equation has the scalar-principal form*

$$(A.26) \quad \partial_t q^A - q^{ij} \partial_i \partial_j q^A = \mathcal{F}^A(x, t, q, q^{-1}, \partial q; r, r^{-1}, \partial r, \partial^2 r),$$

where $\mathcal{F}$ is a universal smooth expression and contains no second spatial derivatives of q . Suppose, in the fixed coordinates, that

$$\lambda I \leq (q_{ij}) \leq \Lambda I, \quad \|r^{-1}\|_{C^0(Q)} \leq A, \\ \|q\|_{C^{k+\sigma, (k+\sigma)/2}(Q)} + \|r\|_{C^{k+1+\sigma, (k+1+\sigma)/2}(Q)} \leq A.$$

Then

$$(A.27) \quad \|q\|_{C^{k+1+\sigma, (k+1+\sigma)/2}(Q')} \leq C,$$

where C depends only on the displayed bounds, the rank and dimension, and the fixed spatial and lower-time separations between Q' and Q . Consequently, the estimate is uniform in any bounded-geometry atlas with fixed nested-cylinder ratios.

Proof. The bound on r^{-1} , together with the assumed Hölder bound on r , controls r^{-1} in $C^{k+1+\sigma, (k+1+\sigma)/2}(Q)$. Indeed, differentiate $rr^{-1} = I$ and apply the parabolic Hölder product estimates. In particular, the reference Christoffel symbols and their derivatives through parabolic order k are uniformly bounded. The ellipticity bound on q similarly controls q^{-1} to the order supplied by the hypotheses.

Let $|\alpha| = k - 1$, and collect the tensors $V_\alpha^A = \partial^\alpha q^A$ into one finite system. Spatially differentiating (A.26) and moving the term in which all derivatives remain on q^A to the left

gives

$$(A.28) \quad (\partial_t - q^{ij}\partial_i\partial_j)V_\alpha^A = B_{\alpha\beta D}^{iA}\partial_i V_\beta^D + C_{\alpha\beta D}^A V_\beta^D + G_\alpha^A.$$

Here β ranges over the multiindices of length $k-1$. The first-order terms include those in which one derivative falls on q^{ij} and the highest derivatives arising from the dependence of $\mathcal{F}$ on ∂q . The terms linear in an undifferentiated V_β are placed in C . After all terms linear in V_β or ∂V_β have been placed in C or B , every factor in G_α contains a derivative of q of order at most $k-1$, or a derivative of r of order at most $k+1$. The hypotheses, the preceding control of both inverse metrics, and the parabolic Hölder product estimates therefore give

$$\|B\|_{C^{\sigma,\sigma/2}(Q_1)} + \|C\|_{C^{\sigma,\sigma/2}(Q_1)} + \|G\|_{C^{\sigma,\sigma/2}(Q_1)} \leq C$$

on any fixed intermediate cylinder $Q' \subset Q_1 \subset Q$.

Equation (A.28) is a finite system with the same scalar principal operator on every component. Applying the purely local estimate (A.25) from Remark A.8 gives $V \in C^{2+\sigma,1+\sigma/2}(Q')$, with a bound depending only on the stated data; its C^0 -term is already controlled by the assumed $C^{k+\sigma,(k+\sigma)/2}$ -norm of q . Thus all spatial derivatives of q through order $k+1$ are bounded with the required Hölder control. The equation and its differentiated forms then recover the mixed time derivatives of parabolic weight at most $k+1$. This proves (A.27). $\square$

A.6. Continuation, smoothing, and regular restart.

Lemma A.10 (Fixed-quotient Ricci–DeTurck continuation). *Let M be compact, let g_0 be a fixed smooth metric, and let $s(t)$ be a fixed smooth background family on $[a, b]$. Given $c > 0$, $C < \infty$, and $\sigma \in (0, 1)$, there is $\tau_M > 0$ with the following property. For every $t_0 \in [a, b]$ and every smooth metric v_0 satisfying*

$$cg_0 \leq v_0 \leq Cg_0, \quad \|v_0\|_{C^{2,\sigma}(M,g_0)} \leq C,$$

the Ricci–DeTurck equation relative to $s(t)$, with initial value $v(t_0) = v_0$, has a unique solution on $[t_0, t_0 + \tau_M] \cap [a, b]$. The number τ_M may depend on the fixed quotient, atlas, background family, and displayed bounds; no collapse-uniformity is asserted.

Proof. In a fixed finite atlas, the equation has scalar strongly parabolic principal part

$$\partial_t v_{ij} = v^{pq}\partial_p\partial_q v_{ij} + F_{ij}(x, t, v, \partial v; s, \partial s, \partial^2 s).$$

The two-sided metric bound gives a common parabolicity constant, while the $C^{2,\sigma}$ -bound and the fixed smooth background place the coefficients and nonlinearities in a bounded set with uniform Hölder and local Lipschitz bounds. The contraction argument in DeTurck’s local existence proof [10, Theorem and proof, pp. 158–161] therefore gives a lifespan bounded below by one number τ_M for every such restart slice. The same argument gives uniqueness on overlaps. $\square$

Proposition A.11 (Uniform lifted Ricci–DeTurck continuous dependence). *Fix $N \geq 2$ and $\sigma \in (0, 1)$. Let $s(t)$, $a \leq t \leq b$, be a Ricci flow satisfying $\text{PBG}_N(s; \Lambda; [a, b])$ for some $\Lambda \geq 1$. There are constants $\eta_{\text{cd}} > 0$ and $C_{\text{cd}} < \infty$, depending only on these data and an upper bound for $b - a$, with the following property.*

Hypotheses. If $v(a)$ is a smooth Riemannian metric and

$$\|v(a) - s(a)\|_{C^{N,\sigma}(M,s(a))} < \eta_{\text{cd}},$$

Conclusion. Then the Ricci–DeTurck flow $v(t)$ relative to $s(t)$ exists on $[a, b]$ and

$$(A.29) \quad \|v - s\|_{X_s^N[a,b]} \leq C_{\text{cd}} \|v(a) - s(a)\|_{C^{N,\sigma}(M,s(a))}.$$

The constants are uniform under collapse of the quotient.

Proof. **Step 1. Derive the quadratic perturbation estimate.**

Put $u = v - s$ and write $\mathcal{L}_s = \partial_t - \Delta_{L,s}$, where the Lichnerowicz Laplacian on symmetric two-tensors is

$$(\Delta_{L,s}u)_{ij} = s^{pq}\nabla_p\nabla_q u_{ij} + 2R_{ipqj}(s)u^{pq} - R_i^p(s)u_{pj} - R_j^p(s)u_{ip}.$$

All contractions and covariant derivatives in this formula use s . For the spatial Ricci–DeTurck operator

$$\mathcal{Q}_s(w) = -2\text{Ric}(w) + \mathcal{L}_{W(w,s)}w,$$

Taylor’s formula at s gives

$$(A.30) \quad \mathcal{L}_s u = \mathcal{N}_s(u), \quad \mathcal{N}_s(u) = \int_0^1 (1 - \vartheta) D^2 \mathcal{Q}_s|_{s+\vartheta u}[u, u] d\vartheta.$$

Every term in $\mathcal{N}_s(u)$ is at least quadratic and has schematic form

$$u * \nabla^2 u + \nabla u * \nabla u + \text{Rm}(s) * u * u.$$

The parabolic Hölder spaces in (4.5) are Banach algebras. Consequently, on every interval $[a, t]$ contained in the existence interval and on which $\|u\|_{X_s^N[a,t]}$ is sufficiently small,

$$(A.31) \quad \|\mathcal{N}_s(u)\|_{Y_s^N[a,t]} \leq C \|u\|_{X_s^N[a,t]}^2.$$

Proposition A.7, applied to $\mathcal{L}_s$ on the same interval, yields

$$(A.32) \quad \|u\|_{X_s^N[a,t]} \leq C_{\text{lin}} \|u(a)\|_{C^{N,\sigma}} + C \|u\|_{X_s^N[a,t]}^2.$$

We also record the initial time derivatives needed to start the bootstrap. Put $\varepsilon_0 = \|u(a)\|_{C^{N,\sigma}(M,s(a))}$. For ε_0 below a fixed uniform threshold, the equation recursively gives

$$\sum_{j=0}^{\lfloor N/2 \rfloor} \|(\mathcal{D}_t^s)^j u|_{t=a}\|_{C^{N-2j,\sigma}(M,s(a))} \leq C_{\text{init}} \varepsilon_0.$$

To see this, evaluate $\mathcal{L}_s u = \mathcal{N}_s(u)$ at $t = a$, then differentiate it in time and substitute the already determined lower time derivatives. At parabolic weight at most N , this recursion uses spatial derivatives of $u(a)$ through order N only. The background coefficients are controlled by PBG_N , and every resulting term vanishes when $u(a) = 0$. The Hölder product estimates therefore give the displayed linear bound on a sufficiently small fixed neighborhood. Passing from ordinary to metric-compatible time derivatives uses only the same controlled

background coefficients. Enlarge C_{init} to include the fixed-chart norm-equivalence constants, and increase C_{lin} so that $C_{\text{lin}} \geq \max\{C_{\text{init}}, 1\}$.

Choose $\delta > 0$ below this initial-data threshold and so that every tensor u with $\|u\|_{X_s^N[a,t]} \leq \delta$ makes $s + u$ positive definite and keeps the Ricci–DeTurck system uniformly parabolic, and also so that $C\delta \leq 1/4$. This δ is the *bootstrap radius*. On the fixed compact quotient, Lemma A.10 gives a unique short-time solution. The prescribed smooth time-dependent background enters only as a known coefficient in that fixed-atlas quasilinear system.

Let $[a, t_*]$ be any interval on which $\|u\|_{X_s^N[a,t_*]} \leq \delta$. If $\|u(a)\|_{C^{N,\sigma}} \leq \delta/(4C_{\text{lin}})$, then the preceding inequality and $\|u\|^2 \leq \delta\|u\|$ give

$$\|u\|_{X_s^N[a,t_*]} \leq \frac{\delta}{4} + \frac{1}{4}\|u\|_{X_s^N[a,t_*]}, \quad \text{hence} \quad \|u\|_{X_s^N[a,t_*]} \leq \frac{\delta}{3}.$$

We call this strict improvement from δ to $\delta/3$ the *bootstrap estimate*.

Step 2. Close the bootstrap and continue to the final time.

We now make the continuation argument quantitative. Let $[a, T_{\text{max}})$, with $T_{\text{max}} \leq b$, be the maximal existence interval. For the smooth local solution, the purely temporal Hölder seminorms in $X_s^N[a, t]$ tend to zero as $t \downarrow a$: their time exponents are $\sigma/2$ and $(1 + \sigma)/2$, both strictly less than one. The sup norms and spatial Hölder parts converge to those of the initial time derivatives. Consequently, the initial-jet estimate in Step 1 gives

$$\limsup_{t \downarrow a} \|u\|_{X_s^N[a,t]} \leq C_{\text{init}}\varepsilon_0 \leq \delta/4.$$

Thus the bootstrap condition holds on a short initial interval. The length of this interval may depend on higher derivatives of the fixed smooth initial metric; no collapse-uniform lower bound for it is needed. Let T_* be the supremum of the times up to which the norm is at most δ . The estimate just proved gives the strict bound $\delta/3$ on $[a, T_*]$; continuity of the parabolic Hölder norm on smooth compact subintervals then extends the bootstrap condition past T_* . Hence $T_* = T_{\text{max}}$. Consequently,

$$(A.33) \quad \|u\|_{X_s^N[a,t]} \leq \delta/3, \quad a < t < T_{\text{max}}.$$

In particular, every slice $v(t) = s(t) + u(t)$ is uniformly positive definite, and its lifted $C^{N,\sigma}$ -norm is bounded by a constant depending only on the data in the proposition.

The remaining continuation argument is made on the fixed compact quotient; no collapse-uniform restart time is needed. Choose a finite smooth coordinate atlas on this quotient and a fixed smooth background metric g_0 . The intrinsic lifted estimate (A.33), restricted to finitely many lifted coordinate neighborhoods and compared with this fixed atlas, gives constants $c_M > 0$ and $C_M < \infty$ such that

$$(A.34) \quad c_M g_0 \leq v(t) \leq C_M g_0, \quad \|v(t)\|_{C^{N,\sigma}(M,g_0)} \leq C_M, \quad a < t < T_{\text{max}}.$$

Here the constants may depend on the particular compact quotient and its chosen atlas; this is harmless, since all collapse-independent estimates have already been obtained in (A.33). The coefficients of the known background $s(t)$ have bounded fixed-atlas norms on $[a, b]$.

Therefore, the Ricci–DeTurck equation is a uniformly strongly parabolic quasilinear system on this fixed compact manifold, with a uniform $C^{2,\sigma}$ -bound on its solution slices.

Lemma A.10, applied with the bounds in (A.34), gives a number $\tau_M > 0$, independent of the restart time $t_0 < T_{\max}$, for which the solution with initial value $v(t_0)$ exists on $[t_0, t_0 + \tau_M] \cap [a, b]$. Only the $C^{2,\sigma}$ -bound and uniform parabolicity enter this fixed-quotient lifespan; the higher finite-order bounds are then propagated by differentiating the system and applying the local parabolic estimates already established above.

If $T_{\max} < b$, choose $t_0 < T_{\max}$ so close to $T_{\max}$ that $t_0 + \tau_M > T_{\max}$. The restarted solution extends past $T_{\max}$, and uniqueness for the strongly parabolic Ricci–DeTurck system identifies it with the original solution on their overlap, contradicting maximality. Thus $T_{\max} = b$. Choosing $t_0 \in (b - \tau_M, b)$ and restarting on $[t_0, b]$ gives the endpoint $t = b$. Again uniqueness identifies the two solutions on their overlap, so there is no gluing loss at the limiting time.

Step 3. Obtain the linear global stability estimate.

Finally, on $[a, b]$, combine (A.33) with (A.31) and the linear estimate. Since $C\delta \leq 1/4$, the quadratic term can be absorbed, giving

$$\|u\|_{X_s^N[a,b]} \leq 2C_{\text{lin}} \|u(a)\|_{C^{N,\sigma}}.$$

Taking $\eta_{\text{cd}} = \delta/(4C_{\text{lin}})$ and $C_{\text{cd}} = 2C_{\text{lin}}$ proves the proposition. The usual DeTurck ODE then identifies this solution with a pullback of the corresponding Ricci flow. $\square$

The preceding results give the continuous-dependence estimate needed in the body. We finish by proving the positive-time smoothing and restart statement used at each switching time.

Proof of Proposition 6.4. Step 1. Create high-order parabolic geometry for the background.

We work on the universal cover. From (6.7), local $\mathbb{T}^2$ -invariance, and Lemma 3.2, the metrics $\tilde{r}(\tau)$ have a uniform injectivity-radius lower bound on $[10, 100]$. Shi’s estimates give bounds for all curvature derivatives through the finite order needed below on $[10, 100]$. Hence, after increasing a constant Λ_{sm} depending only on the fixed data,

$$\text{PBG}_{N+5}(r; \Lambda_{\text{sm}}; [10, 100])$$

holds. The norm equivalence in Lemma A.2 therefore converts (6.8) into the uniform order- N fixed-chart bound on $[10, 100]$ used in the bootstrap below.

Step 2. Gain five derivatives for the Ricci–DeTurck metric.

Choose β_{reg} initially so small that (6.8) makes q uniformly bilipschitz to r and keeps the Ricci–DeTurck equation uniformly parabolic. In the universal-cover charts for r , the equation for q has the quasilinear scalar-principal form (A.26). We now give the finite bootstrap that will be used below.

Choose once and for all six nested spatial cylinders in each fixed chart, with radii decreasing by a uniform amount. For $j = 0, \dots, 5$, put $I_j = [10 + 2j, 100]$. We claim inductively

that

$$(A.35) \quad \|q\|_{C_{\text{uloc}}^{N+j+\sigma, (N+j+\sigma)/2}(\widetilde{M} \times I_j, r)} \leq C_j.$$

For $j = 0$, this follows from (6.8) and the PBG_{N+5} -bound for r . Suppose it holds for some $j < 5$. Apply Lemma A.9 with $k = N + j$, using the j -th spatial cylinder as the outer cylinder and the next one as the inner cylinder. The two units of temporal room between I_j and I_{j+1} provide the required separation from the lower time face; backward cylinders allow the estimate to include $\tau = 100$.

The hypotheses of the lemma are the induction hypothesis, uniform ellipticity of q , and the PBG_{N+5} -bound for r . In particular, Lemma A.1 gives a uniform coordinate C^0 -bound for r^{-1} on every outer cylinder, so the reference nondegeneracy assumption is satisfied at all five steps. More explicitly, differentiating the equation $N + j - 1$ times uses derivatives of q of order at most $N + j$ and derivatives of r of order at most $N + j + 1 \leq N + 5$. The lemma therefore gives the asserted $C^{N+j+1+\sigma, (N+j+1+\sigma)/2}$ -bound on I_{j+1} . This proves the induction. After the five steps,

$$\|q\|_{C_{\text{uloc}}^{N+5+\sigma, (N+5+\sigma)/2}(\widetilde{M} \times [20, 100], r)} \leq C_5.$$

The argument is carried out entirely in the uniform universal-cover atlas, so its constants do not involve the injectivity radius of the collapsed quotient. For related local derivative estimates, see [7, Appendix A.2, “Local derivative estimates”]. Lemma A.9 provides the precise finite-order estimate used here.

Step 3. Smooth the difference by the homogeneous linear system.

Set $u = q - r$. Since q and r solve the Ricci–DeTurck equation with the same background r , Lemma A.4, applied on the universal cover, gives a homogeneous system

$$\partial_\tau u - P^{ij} \nabla_i \nabla_j u = B^i \nabla_i u + Cu$$

with scalar principal part. The preceding regularity of q and r gives the coefficient bounds required by Proposition A.7 with $m = N + 5$. Its interior estimate, applied with room at the lower time face, yields

$$(A.36) \quad \|q - r\|_{X_r^{N+5}[30, 100]} \leq C \|q - r\|_{C^0(M \times [1, 100], r)}.$$

Step 4. Rescale, average, and verify the restart geometry.

The fixed factor-100 rescaling changes the spatial Hölder norm by a factor depending only on N and σ . Evaluating (A.36) at $\tau = 100$ gives the first term on the left-hand side of (6.9). Since ρ^+ is invariant, Lemma 3.3 and

$$r^+(1) - \rho^+ = \overline{q^+(1) - \rho^+}$$

give the second term.

It remains to verify the uniform geometry of the restart. Shi’s estimates at $\tau = 100$, followed by the fixed rescaling, bound $|\widetilde{\nabla}^j \text{Rm}_{\widetilde{\rho}^+}|$ for $0 \leq j \leq N + 3$. Lemma 3.2, applied to the invariant metric $r(100)$ and then rescaled, gives a uniform lower bound for $\text{inj}_{\widetilde{\rho}^+}$. Thus $\text{BG}_N(\rho^+; \Lambda')$ holds for a uniform Λ' . By decreasing β_{reg} once more, (6.9) makes

$r^+(1)$ a sufficiently small $C^{N+5,\sigma}$ -perturbation of ρ^+ . The formulas for the curvature tensor and its covariant derivatives then give uniform bounds through order $N + 3$, while positivity follows from the C^0 -part. Finally, $r^+(1)$ is invariant, so Lemma 3.2 gives its uniform lifted injectivity-radius lower bound. Enlarging one constant gives $\text{BG}_N(\rho^+; \Lambda_{\text{reg}})$ and $\text{BG}_N(r^+(1); \Lambda_{\text{reg}})$, proving the proposition. $\square$

This completes the uniform analytic estimates used in the body of the paper.

APPENDIX B. SOURCE GUIDE FOR THE EXTERNAL INPUTS

Section 2 states the external results in the forms used by the block argument. This appendix records precisely what is quoted, what is a harmless reformulation, and what uniformization is supplied in the present paper.

The Type-III blowdown theorem. Proposition 2.1 is the portion of [17, Theorem 1.2] used here, restricted to the three geometries under consideration and written with the two hypotheses displayed in (2.1). It asserts only compact blowdown information and the lifted limits. In the Euclidean and Nil cases, the compact conclusion is full convergence to a point. In the Sol case, the quoted compact conclusion only classifies any Gromov–Hausdorff subsequential limit as a circle or interval; no quantitative lower length bound, precompactness, or uniqueness of the compact limit is imported. The full metric-circle limit on the working cover is instead proved in Corollary 8.4: the scalar recursion and the positive quotient-length lower bound give convergence of ℓ_k and $E_k \rightarrow 2$, the two-sided restart estimate propagates the length limit through each block, and the comparison with Bamler’s fibers makes the retained fibers vanish. Proposition 9.3 then gives a full circle or interval limit of positive length on the original manifold. Exponential convergence of the unrescaled flow in the Euclidean case is a later consequence of near-flat compactness and Proposition 2.3, not part of this quoted input.

The almost-flat theorem. Proposition 2.2 is the three-dimensional topological consequence of the Gromov–Ruh almost-flat theorem [12, 19]. No analytic conclusion is added: it is used only to turn the scale-invariant small-curvature condition into virtual nilpotence of the fundamental group.

Flat stability and gauge transfer. For one flat background, [13, Theorem 3.7] gives exponential stability for the Ricci flow. In its proof, the corresponding Ricci–DeTurck convergence is established first, and [13, Proposition 3.6] is then used to transfer this convergence to the ungauged Ricci flow, including convergence of the relevant diffeomorphisms. The compact-family version in Proposition 2.3 is not a verbatim quotation: it follows by taking a finite subcover of a compact set of actual flat representatives and then taking the minimum decay rate and maximum constant among the finitely many fixed-background estimates. The proposition is stated for smooth initial metrics in the little-Hölder neighborhood. Its all-orders conclusion is obtained by positive-time parabolic estimates in DeTurck gauge and convergence of the associated diffeomorphisms; higher-order constants may also depend on higher norms of the smooth initial metric, while the decay rate is uniform on the compact family. When near-flat compactness is initially stated modulo diffeomorphisms, the relevant diffeomorphism must first be applied to obtain actual representatives on the fixed torus; Proposition 10.2 absorbs it into the finite covering map.

Invariant energy and quotient length. The length inequality in Proposition 2.4 is [18, Lemma 3.7], and the reciprocal energy inequality is the integrated form of [18, Lemma 3.3].

The sentence treating $E(t_0) = 0$ only records the zero solution of the same maximum-principle inequality, so no reciprocal expression is used at vanishing energy.

The global scale-invariant curvature bound. Bamler's [2, Corollary 1.2] gives the estimate in Proposition 2.5 for all sufficiently large times. Increasing K , if necessary, over the compact initial time interval gives the stated bound for every $t > 0$. Pulling the flow back to a finite cover leaves the sectional-curvature estimate and its constant unchanged.

The fixed-threshold Bamler alternative. The initial normalization required by [2, Theorems 1.1 and 1.4] is obtained by a constant parabolic rescaling. Choose $a > 0$ sufficiently large that $\widehat{g}(0) = ag(0)$ satisfies the normalized initial conditions, and set

$$(B.1) \quad \widehat{g}(u) = ag(u/a), \quad u = at > 0, \quad u^{-1}\widehat{g}(u) = t^{-1}g(t).$$

Regard $\widehat{g}$ as a flow with no surgeries, apply the cited theorem, and scale back. The normalized metric, its comparison norms, and the normalized fiber diameters are unchanged; a tail time $\widehat{T}$ becomes $\widehat{T}/a$, and an error function $\widehat{\varepsilon}(u)$ becomes $\widehat{\varepsilon}(at)$. This justifies using the theorem for an arbitrary immortal flow on the working cover.

Proposition 2.6 is a fixed-threshold reformulation, not a verbatim quotation. Theorem 1.4(b)–(d) of [2] supplies the topology-specific global alternatives and a single time-dependent error function $\varepsilon(t) \rightarrow 0$. The detailed construction of the invariant comparison metrics, together with the high-order closeness and fiber-diameter estimates, is given in [6, Section 4.4, especially Proposition 4.9], using [3, 4, 5]. In the whole-manifold torus-bundle cases, these fiberwise local actions constitute the twisted torus-bundle invariance fixed in Subsection 3.2; no globally defined untwisted $\mathbb{T}^2$ -action is being asserted when the monodromy is nontrivial. In the Sol case, Proposition 3.1 and the working-cover convention ensure $H \in \mathrm{SL}(2, \mathbb{Z})$ and $\mathrm{Tr} H > 2$. Its positive eigenvalues give a real logarithm $A = \log H$ with trace zero; $\mathbb{R}^2 \rtimes_{e^{tA}} \mathbb{R}$ is isomorphic to the connected Lie group Sol, and $\mathbb{Z}^2 \rtimes_H \mathbb{Z}$ is a cocompact lattice in it. The working torus bundle is therefore a Solv manifold in the direct scope of [2, Theorem 1.4(d)]. The base double cover is taken before the core refinement, whose monodromy is a power of H_0^2 ; no later cover or descent of Bamler's comparison metric is invoked. Apply the cited results directly to the flow on this fixed working cover. After fixing a cutoff $c_0 \in (0, 1)$ with $Kc_0^2 < \varepsilon_{\mathrm{af}}$, choose the associated tail time $T_B(c_0)$ so that

$$\varepsilon(u) \leq c_0/2 \quad \text{for every } u \geq T_B(c_0).$$

In the flat $\mathbb{T}^3$ and small-Nil branches, the normalized diameter is then eventually below this fixed threshold; in the remaining branches, the cited alternatives give the invariant torus-bundle comparison whenever the diameter is at least $c_0\sqrt{t}$. Replacing the error by its nonincreasing envelope

$$\varepsilon_B^{c_0}(t) = \sup_{u \geq t} \varepsilon(u)$$

only weakens the metric and fiber estimates. Moreover, since $\varepsilon_B^{c_0}(t) \geq \varepsilon(t)$, every integer m with $0 \leq m < (\varepsilon_B^{c_0}(t))^{-1}$ also satisfies $m < \varepsilon(t)^{-1}$, so the available derivative order still suffices. This gives the diameter-defined, tail-uniform formulation in Proposition 2.6.


In every application, the working cover is selected once according to the [working-cover convention](#); the fixed-threshold alternative is invoked only on that cover and never on an arbitrary torus bundle representing the Euclidean, Nil, or Sol geometry. In particular, a negative-trace hyperbolic bundle is handled through its initial working cover, not by applying Proposition [2.6](#) directly to that bundle. The fixed global application uses $c_0 = c_{\text{af}}$, with tail data $T_B(c_{\text{af}}), \varepsilon_B^{c_{\text{af}}}$; the Euclidean proof makes one separate application with $c_0 = \zeta$, $T_B(\zeta)$, and ε_B^ζ , without altering the fixed global application.

APPENDIX C. DEPENDENCY TREES


This appendix is a navigational aid rather than an additional part of the proof. Tree I traces the lifted geometric and analytic inputs through finite-time decay, fibration comparison, and stage preparation to the quantitative one-block estimates and hence to the qualitative block contraction theorem. Tree II combines the block results with the locally $\mathbb{T}^2$ -invariant energy-length estimates to obtain iteration on a connected large-diameter interval. On a persistent tail it also records the quantitative quotient-length lower bound, retained-fibration equivalence, transition control, and elementary circle collapse used to prove convergence of the normalized quotient length and the full circle limit. Tree III takes those outputs as inputs rather than re-expanding their internal block dependencies; it records the application of Bamler's alternative on the fixed torus-bundle cover, then combines cover topology and finite-cover descent with the Euclidean, Nil, and Sol arguments to obtain the main theorem.

Arrows run from a dependency to a result that uses it. When present, an edge label lists only the named parameters passed from the source vertex to the target vertex; an arrow carrying no such parameter is left unlabeled. White boxes denote results whose dependencies are expanded in the displayed tree. Light-gray boxes denote inputs carried in from earlier sections or earlier trees, including external results; the main theorem is shown in darker gray. The trees show only the principal nonredundant dependencies; definitions and remarks are omitted, and a numbered result appears at most once in each diagram. Blue headings indicate that the entire box is an internal hyperlink to the corresponding numbered statement.


Tree I: Analytic estimates, fibration comparison, and quantitative one-block control



Tree II: Componentwise iteration and the persistent-tail circle limit



Tree III: Long-time limits and assembly of the main theorem



APPENDIX D. PARAMETERS AND PERSISTENT NOTATION

This appendix is a reference guide for the quantitative data carried through the proof. It collects only symbols that persist across major statements or section boundaries, making the hypotheses and conclusions of the stage and block constructions easier to track; notation local to a single proof is deliberately excluded. The appendix is not used as an additional mathematical input.

The table lists Greek-letter symbols first and Roman-letter symbols second, in alphabetical order within each group. Upper- and lower-case variants are grouped by their base letter; numerical or symbolic subscripts precede alphabetic subscripts, which are ordered without regard to case. Related symbols share an entry under the first symbol shown. For each entry, the second column links to the statement, equation, subsection, or anchored paragraph where the symbol is introduced or fixed, and the third column describes the role that the symbol continues to play later in the paper. The table is restricted to structural constants, maps, metrics, state variables, and scales that cross statement or section boundaries; local coordinate variables, tensor indices, cutoff functions, and every symbol used only inside one proof are omitted. It is intended both as a reading aid for the constant dependencies and as a compact reference for the persistent data in the stage and block constructions.

Symbol	Defined or fixed in	Meaning
<i>Greek-letter parameters and maps</i>		
α	Theorem 4.2	Prescribed factor $0 < \alpha < 1$ contracting the non-invariant part on $[10, 100]$ relative to its initial $C^{N,\sigma}$ -norm.
β_*	Lemma 6.5 and the joint proof of Theorem 7.2 and Proposition 7.3	Stage-closeness target in the preparation lemma, chosen in the block proof below the decay, fibration-comparison, and restart thresholds.
β_{fib}	Lemmas 5.5 and 7.1	Closeness threshold produced by the pointwise comparison and propagated as the stage-closeness input in the retained-fibration lemma.
β_{orb}	Theorem 4.2	Short-orbit diameter threshold used in the zero-average C^0 -estimate.
β_{pert}	Theorem 4.2	Independent smallness threshold for the Ricci–DeTurck perturbation in the lifted parabolic estimates.
β_{reg}	Proposition 6.4 and Lemma 6.5	Smallness threshold for positive-time regular restart, carried into stage preparation and the block parameter choice.
γ	Lemma 5.3	Relative C^1 -closeness threshold in fiber-equivalent comparison; chosen from the quotient-length lower bound and the target fixed-fiber diameter, and reused in Lemma 5.5 and Corollary 8.4.
Γ	Lemma 9.1 and Proposition 9.3	Fixed finite group acting isometrically on the covering metrics; in finite-cover descent it is the deck group, with $ \Gamma = d$. Its limiting action on the circle need not be faithful.

Symbol	Defined or fixed in	Meaning
δ_0	Definition 6.1 and Theorem 8.3	Lower bound for the normalized diameter on a rescaled large block or connected large-diameter interval; it is at least the cutoff in the fixed-threshold application being used.
δ_{flat}	Equations (10.5)–(10.6)	Normalized-diameter tolerance chosen from $\varepsilon_{\text{flat}}$; it is independent of both the auxiliary cutoff ζ and the fixed global cutoff c_{af} .
δ_J	Lemma 8.1	Upper threshold for the relative C^1 -jump, depending only on the chosen reference energy bound E_0 . The component iteration uses $E_0 = 3$.
ε_{af}	Almost-flat threshold paragraph	Dimensional almost-flat threshold used to choose every admissible cutoff c_0 , including the fixed global cutoff $c_{\text{af}}(K)$.
ε_{av}	Lemma 6.3	High-order smallness threshold under which averaging preserves bounded lifted geometry.
$\varepsilon_B^{c_0}(t)$	Proposition 2.6 and the standing fixed-threshold application	Nonincreasing comparison error attached to the application with cutoff c_0 ; different cutoffs have separately labeled error functions.
ε_{fib}	Lemmas 5.5 and 7.1	Comparison-metric and comparison-fiber threshold produced by the pointwise lemma and used uniformly throughout a retained-fibration block.
$\varepsilon_{\text{flat}}$	Proposition 10.2	Analytic smallness threshold for entering the uniform fixed-coordinate flat-stability neighborhood; it is independent of the geometric cutoff used in Bamler’s alternative.
$\zeta, T_B(\zeta), \varepsilon_B^\zeta$	Proof of the Euclidean case in Theorem 10.1	Auxiliary cutoff, with $K\zeta^2 < \varepsilon_{\text{af}}$, and its separately labeled tail data. The liminf argument allows every such ζ , without changing the fixed global application.
η	Lemma 5.3 and Definition 6.2	In fiber comparison, a threshold for comparison-fiber diameter, chosen with γ . In the admissible-start definition, the same letter denotes a generic non-invariant error threshold; these are distinct uses.
η_0	Theorem 8.3	Initial non-invariant error threshold for iteration on a connected large-diameter interval; it depends only on K, δ_0, N, σ and is independent of D_1 .
η_{blk}	Theorem 7.2	Start-and-restart error threshold for the block contraction.
η_{cd}	Proposition A.11	Initial perturbation threshold for uniform Ricci–DeTurck continuous dependence.
η_{st}	Lemma 6.5 and Equation (6.16)	Stage-start error threshold for the prescribed target β_* , also satisfying $C_1\eta_{\text{st}} \leq \frac{1}{2}$; the block threshold η_{blk} is chosen below it.
θ	Lemmas 5.5 and 7.1	Preliminary fixed-fiber target produced by the pointwise comparison and retained strictly throughout one large block.
θ_{blk}	Theorem 7.2 and Proposition 7.3	Fixed-fiber target for the block iteration, reproduced by every completed block.
θ_{max}	Lemmas 5.5 and 7.1	Prescribed upper target for the fixed-fiber diameter; in the block proof it is set to $\beta_{\text{orb}}^{\text{dec}}/4$.

Symbol	Defined or fixed in	Meaning
λ	Theorem 4.2	Positive lower bound for the quotient-circle length in the finite-time decay theorem.
Λ_*	Lemma 6.5 and Theorem 7.2	Uniform BG_N -bound, chosen as $\max\{\Lambda_{\text{reg}}, \Lambda_{\text{init}}\}$, depending only on K, N, σ . It controls the first averaged slice and is reproduced at every restart.
λ_p	Lemmas 5.5 and 7.1	Quotient-length lower threshold in the pointwise comparison and the retained-fibration conclusion; the block construction sets $\lambda_p = \delta_0/100$.
Λ_{reg}	Proposition 6.4	Uniform BG_N -bound for the rescaled locally $\mathbb{T}^2$ -invariant endpoint and the averaged restart.
$\pi, \pi_B, \pi_{\text{cyc}}, \pi_{\text{fin}}$	Proposition 2.6, Section 3, and the metric-and-gauge notation	The retained fibration; the pulled-back Bamler fibration $\pi_B(\tau)$, denoted π_t^B in physical time; the infinite cyclic covering map; and the fixed finite regular covering map, respectively.
ρ	Propositions 2.3 and 10.2	Exponent $0 < \rho < 1$ in the little-Hölder space $h^{2+\rho}$ used for flat stability. It is unrelated to the invariant metric called ρ in Lemma 8.1.
ρ^+, ρ_k^+	Proposition 6.4 and Theorem 8.3	Rescaled invariant endpoint: $\rho^+ = 100^{-1}r(100)$ generically and $\rho_k^+ = 100^{-1}r_k(100)$ on stage k .
σ	Standing regularity convention and Appendix A	Hölder exponent $0 < \sigma < 1$, fixed once for the stage and block constructions and used in the spatial and parabolic norms.
τ	Definition 6.1	Parabolically rescaled time: the ambient block p is defined on $[1/2, 101]$, while the stage families q and r are used on $[1, 100]$.
Φ_k	Theorem 8.3	Stage-start pullback diffeomorphism, with $p_k(\tau) = \Phi_k^*(T_k^{-1}g(T_k\tau))$ and $\Phi_{k+1} = \Phi_k \circ \psi_k(100)$.
ψ_k, Ψ	Theorem 8.3	Stage diffeomorphisms satisfying $q_k = \psi_k^*p_k$ and $\psi_k(1) = \text{Id}$, and the assembled map $\Psi(T_k\tau) = \Phi_k \circ \psi_k(\tau)$; hence $g' = \Psi^*g$.
<i>Roman-letter parameters and notation</i>		
$\mathcal{A}, \bar{U}, U^\perp$	Subsection 3.4 and Lemmas 3.3 and 4.1	Averaging projection for the fixed local $\mathbb{T}^2$ -structure, the average $\bar{U} = \mathcal{A}U$, and the non-invariant part $U^\perp = U - \mathcal{A}U$; the two lemmas give the fixed-time and parabolic Hölder bounds.
b, b_k	Equation (7.1), Proposition 7.3, and Theorem 8.3	First numerical admissibility coordinate: the non-invariant start error on one stage and at the start of stage k ; the block estimate gives $b_{k+1} \leq b_k/2$.
$BG_m(s; \Lambda)$	Section 4	Scale-one bounded geometry on the universal cover: $\text{inj}_{\bar{s}} \geq \Lambda^{-1}$ and $ \tilde{\nabla}^j \text{Rm}_{\bar{s}} \leq \Lambda$ for $0 \leq j \leq m+3$, where $m \geq 2, \Lambda \geq 1$. This is a separate admissibility condition.
c_0	Proposition 2.6 and the standing fixed-threshold application	Generic cutoff of the particular Bamler application used throughout one block or component iteration.
$c_{\text{af}} = c_{\text{af}}(K)$	The fixed global application following Proposition 2.6	Fixed global cutoff used for the global almost-flat/torus-bundle decomposition; in the generic block machinery it is the specialization $c_0 = c_{\text{af}}$.
C_{cd}	Proposition A.11	Linear continuous-dependence constant for one Ricci–DeTurck comparison.

Symbol	Defined or fixed in	Meaning
C_{fib}	Lemma 5.4	Constant, equal to $B + 1$ in the proof, controlling total diameter when two short-fiber fibrations are fiber-inequivalent.
C_J	Lemma 8.1	Constant depending only on the reference energy bound E_0 , controlling a one-sided energy increase and the absolute logarithmic length change. The component iteration uses $E_0 = 3$.
C_L	Lemma 9.1 and Corollary 8.4	Metric circle of circumference $L > 0$. In the working-cover limit $L = \ell_\infty$; its diameter is $L/2$.
C_{reg}	Proposition 6.4	Constant in the high-order positive-time smoothing estimate.
C_{st}	Lemma 6.5	Stagewise linear stability constant, also controlling the high-order endpoint jump.
d, d_k	Equation (7.1), Theorem 7.2, and Theorem 8.3	Second numerical admissibility coordinate: the retained-fiber ambient diameter at a generic stage start and at the start of stage k . This use of d is unrelated to the covering degree below.
d (cover degree)	Lemma 9.2 and Proposition 9.3	Positive integer degree of the fixed finite cover; for a regular cover, $d = \Gamma $. This is not the stage-start fiber diameter denoted by d .
D_1	Theorem 8.3	Initial upper bound for the normalized diameter at the chosen start of a connected large-diameter interval.
$\mathcal{D}_t^s$	Equation (A.2)	Metric-compatible covariant time derivative used in the intrinsic parabolic Hölder norms.
$d_{\mathbb{T}^2}(r)$	Subsection 3.4	Supremum of the ambient diameters of the fixed torus fibers in the metric space (M, r) .
$E(t), e_k(\tau), E_k$	Proposition 2.4, Lemma 8.2, and Theorem 8.3	Maximum vertical energy of an invariant flow, its stagewise profile $e_k(\tau)$, and the start values $E_k = e_k(1)$ created after the first restart and propagated by the scalar recursion; on a persistent tail, $E_k \rightarrow 2$.
$g(t)$	Theorem 1.1 and the working-cover convention	Immortal Ricci flow in physical time. The same notation is used for its lift while the block construction is carried out on the working cover, denoted simply by M .
$g'(t)$	Theorem 8.3	Continuous, piecewise smooth DeTurck-gauge pullback of the original physical-time flow on the controlled blocks.
$g_{\text{fin}}(t)$	Sections 3, 9, and 12	Pullback of the original Ricci flow to the fixed finite regular working cover.
g_{RDT}	Theorem 4.2	Ricci–DeTurck solution relative to an invariant background in the finite-time decay theorem.
g_t^B	Proposition 2.6	Bamler’s locally $\mathbb{T}^2$ -invariant comparison metric at physical time t .
H	Working-cover convention and Subsection 3.2	Monodromy matrix on the fixed working cover: I in the Euclidean case, nontrivial unipotent in the Nil case, and hyperbolic with $\text{Tr } H > 2$ in the Sol case.
$h(t)$	Theorem 8.3	Right-continuous invariant comparison family: $h(T_k\tau) = T_k r_k(\tau)$ for $1 \leq \tau < 100$, with the averaged restart at T_{k+1} . It is a Ricci flow on each smooth piece, but may jump at switching times.

Symbol	Defined or fixed in	Meaning
$h_B(\tau)$	Metric-and-gauge notation and the rescaled-block comparison paragraph	Bamler comparison metric on a rescaled block, possibly for a time-dependent comparison fibration.
I	Theorem 8.3	Connected time interval on which the normalized diameter stays above δ_0 ; it may be an unbounded tail or a finite component with a terminal partial block.
i_1, i_2, i_B	Lemma 3.2	Universal-cover injectivity-radius constants for $B^{-1}h \leq g \leq Bh$, with h invariant. In Tree I, i_1 and i_2 are the cases $B = 1$ and $B = 2$, respectively: they supply the invariant-metric input to Proposition 6.4 and the comparison-metric input to Lemma 6.5. For $ \sec_g \leq \Lambda$, the lower bound is $i_B \Lambda^{-1/2}$. The i_1 used locally in the proof of Lemma A.3 is a separate constant with the dependences stated there.
K	Proposition 2.5	Global scale-invariant curvature constant obtained by enlarging Bamler's eventual constant over the compact initial interval.
K_{st}	Lemma 6.5	Uniform sectional-curvature constant for the invariant reference flow on one stage.
$L(S^1, r_{S^1})$	Subsection 8.1	Length of the quotient circle for an invariant metric.
ℓ_∞	Lemma 8.2 and Corollary 8.4	Positive limiting normalized quotient length on a persistent tail: $\ell_k(\tau)/\sqrt{\tau} \rightarrow \ell_\infty$ uniformly on $[1, 100]$. It is the circumference of the working-cover circle limit.
$\ell_k(\tau), \ell_k$	Theorem 8.3 and Lemma 8.2	Quotient-circle length along stage k and its start value $\ell_k = \ell_k(1)$, introduced for $k \geq 1$ after the first restart and then propagated by the scalar recursion; it is not an admissibility coordinate, and on a persistent tail ℓ_k converges.
m (circle quotient)	Lemma 9.1 and Proposition 9.3	Order of the rotation subgroup of the limiting circle action. The quotient has circumference L/m in the cyclic case and interval length $L/(2m)$ in the dihedral case; m divides $ \Gamma $.
$M_{\text{cyc}}, M_{\text{fin}}, \widetilde{M}$	Sections 3, 4, and 9	Infinite cyclic cover, fixed finite regular working cover, and universal cover; tensors on the cyclic cover carry the subscript <i>cyc</i> , while universal-cover lifts carry tildes.
m_{flat}	Proposition 10.2	Finite curvature-derivative order used in the scale-invariant near-flat entry criterion.
N	Standing regularity convention	Integer spatial regularity order $N \geq 2$, fixed once for the stage and block constructions.
$p(\tau)$	Definition 6.1	Pulled-back parabolic rescaling $\Phi^*(S^{-1}g(S\tau))$ of the original flow.
p_{st}	Proposition 10.2	Connected finite covering map that absorbs the compactness diffeomorphism and places the diameter-normalized metric in $\mathcal{U}_{\mathcal{F}}$.
$\text{PBG}_m(s; \Lambda; [a, b])$	Section 4 and Equation (4.2)	The BG_m injectivity and curvature-derivative bounds at every time, together with a fixed-coordinate uniformly locally finite parabolic atlas whose radius, coefficients, overlaps, and cutoffs are uniformly controlled.
$q(\tau), q_k(\tau)$	Lemma 6.5 and Theorem 8.3	Ricci-DeTurck metric on one stage and on stage k .

Symbol	Defined or fixed in	Meaning
$q^+(1)$	Proposition 6.4	Rescaled Ricci–DeTurck endpoint $q^+(1) = 100^{-1}q(100)$, which becomes the next stage’s initial metric.
$r(\tau), r_k(\tau)$	Lemma 6.5 and Theorem 8.3	Invariant Ricci flow starting from the averaged stage-start metric.
$r^+(1)$	Proposition 6.4	Averaged restart metric $r^+(1) = \overline{q^+(1)}$, used as the next invariant stage start.
$r_D(\tau)$	Equation (6.15) and Lemma 6.5	Auxiliary Ricci–DeTurck flow with background p and initial metric $r(1)$, used only to construct and control the invariant Ricci flow r .
S, T_k	Definition 6.1 and Theorem 8.3	Physical start time $T_0 = S$ of the first rescaled block and the geometric sequence $T_k = 100^k S$ of later block starts.
$\mathbb{T}_*^3, \mathcal{F}, \mathcal{U}_{\mathcal{F}}$	Propositions 2.3 and 10.2	Fixed model torus, compact family of actual flat representatives on it, and the associated uniform fixed-coordinate stability neighborhood for smooth initial metrics in the Euclidean near-flat entry argument.
$T_B(c_0)$	Proposition 2.6 and the standing fixed-threshold application	Tail threshold attached to the application with cutoff c_0 ; for example, the fixed global and Euclidean applications use $T_B(c_{\text{af}})$ and $T_B(\zeta)$.
T_{blk}	Theorem 7.2	Late start threshold depending on K, δ_0, N, σ and, possibly, the chosen tail data $T_B(c_0), \varepsilon_B^{c_0}$. The geometric and analytic block thresholds do not acquire this tail-data dependence.
T_p	Theorem 8.3	Late start threshold for component iteration, depending on K, δ_0, N, σ and, possibly, the selected Bamler tail data, but not on D_1 .
$X_s^N[a, b]$	Equation (4.5), Paragraph A.1, and Lemma 4.1	Intrinsic lifted solution norm $\mathcal{C}_s^{N+\sigma, (N+\sigma)/2}$. Averaging is contractive when $s(t)$ is invariant under the fixed torus structure.
$X_{s, \text{fc}}^N[a, b]$	Lemmas A.2 and 4.1	Fixed-chart uniformly local solution norm, uniformly equivalent to X_s^N under PBG. For $s(t)$ invariant under the fixed torus structure and satisfying that bound, averaging is uniformly bounded.
$Y_s^N[a, b]$	Equation (4.5), Paragraph A.1, and Lemma 4.1	Intrinsic lifted forcing norm $\mathcal{C}_s^{N-2+\sigma, (N-2+\sigma)/2}$. Averaging is contractive when $s(t)$ is invariant under the fixed torus structure.
$Y_{s, \text{fc}}^N[a, b]$	Lemmas A.2 and 4.1	Fixed-chart uniformly local forcing norm, uniformly equivalent to Y_s^N under PBG. For $s(t)$ invariant under the fixed torus structure and satisfying that bound, averaging is uniformly bounded.

REFERENCES

- [1] M. T. Anderson and J. Cheeger, *C^α -compactness for manifolds with Ricci curvature and injectivity radius bounded below*, J. Differential Geom. **35** (1992), no. 2, 265–281.
- [2] R. Bamler, *Long-time behavior of 3-dimensional Ricci flow. Introduction*, Geom. Topol. **22** (2018), no. 2, 757–774.
- [3] R. Bamler, *Long-time behavior of 3-dimensional Ricci flow. A: Generalizations of Perelman’s long-time estimates*, Geom. Topol. **22** (2018), no. 2, 775–844.

- [4] R. Bamler, *Long-time behavior of 3-dimensional Ricci flow. B: Evolution of the minimal area of simplicial complexes under Ricci flow*, Geom. Topol. **22** (2018), no. 2, 845–892.
- [5] R. Bamler, *Long-time behavior of 3-dimensional Ricci flow. C: 3-manifold topology and combinatorics of simplicial complexes in 3-manifolds*, Geom. Topol. **22** (2018), no. 2, 893–948.
- [6] R. Bamler, *Long-time behavior of 3-dimensional Ricci flow. D: Proof of the main results*, Geom. Topol. **22** (2018), no. 2, 949–1068.
- [7] R. Bamler and B. Kleiner, *Uniqueness and stability of Ricci flow through singularities*, Acta Math. **228** (2022), 1–215.
- [8] J. Cheeger, M. Gromov and M. Taylor, *Finite propagation speed, kernel estimates for functions of the Laplace operator, and the geometry of complete Riemannian manifolds*, J. Differential Geom. **17** (1982), no. 1, 15–53.
- [9] B.-L. Chen and X.-P. Zhu, *Uniqueness of the Ricci flow on complete noncompact manifolds*, J. Differential Geom. **74** (2006), no. 1, 119–154.
- [10] D. M. DeTurck, *Deforming metrics in the direction of their Ricci tensors*, J. Differential Geom. **18** (1983), no. 1, 157–162.
- [11] D. M. DeTurck and J. L. Kazdan, *Some regularity theorems in Riemannian geometry*, Ann. Sci. École Norm. Sup. (4) **14** (1981), no. 3, 249–260.
- [12] M. Gromov, *Almost flat manifolds*, J. Differential Geom. **13** (1978), no. 2, 231–241.
- [13] C. Guenther, J. Isenberg and D. Knopf, *Stability of the Ricci flow at Ricci-flat metrics*, Comm. Anal. Geom. **10** (2002), no. 4, 741–777.
- [14] R. Hamilton, *The Ricci flow on surfaces*, in *Mathematics and General Relativity*, Contemporary Mathematics **71** (1988), 237–262.
- [15] B. Kleiner and J. Lott, *Singular Ricci flows I*, Acta Math. **219** (2017), 65–134.
- [16] G. M. Lieberman, *Second order parabolic differential equations*, World Scientific Publishing Co., Inc., River Edge, NJ, 1996.
- [17] J. Lott, *Dimensional reduction and the long-time behavior of Ricci flow*, Comment. Math. Helv. **85** (2010), no. 3, 485–534.
- [18] J. Lott and N. Šešum, *Ricci flow on three-dimensional manifolds with symmetry*, Comment. Math. Helv. **89** (2014), no. 1, 1–32.
- [19] E. A. Ruh, *Almost flat manifolds*, J. Differential Geom. **17** (1982), no. 1, 1–14.
- [20] P. Scott, *The geometries of 3-manifolds*, Bull. London Math. Soc. **15** (1983), no. 5, 401–487.
- [21] W.-X. Shi, *Deforming the metric on complete Riemannian manifolds*, J. Differential Geom. **30** (1989), no. 1, 223–301.
- [22] A. S. Švarc, *A volume invariant of coverings*, Dokl. Akad. Nauk SSSR (N.S.) **105** (1955), 32–34 (Russian).
- [23] J. A. Wolf, *Growth of finitely generated solvable groups and curvature of Riemannian manifolds*, J. Differential Geom. **2** (1968), 421–446.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, BERKELEY, BERKELEY, CA 94720-3840, USA

Email address: lott@berkeley.edu