> **This is a recreation of the original chat by Codex. The supplied data export does not contain this local Codex chat; this PDF uses its original saved Codex conversation log.**

Complete nilrotation recurrence

Conversation ID: 01a05ccf-7d47-7222-b499-1c11a81ba7cb

99 recorded user and assistant messages • Dates shown in Asia/Shanghai (UTC+08:00), with source timestamps.

Source: Local Codex transcript supplement

Prepared as part of the supplied data-export collection for the 3:52:20 video.
This PDF reproduces all 99 visible user and assistant messages in the original saved chat: 16 user messages, 68 assistant progress updates, and 15 final answers.
Message dates are original log timestamps, displayed in Asia/Shanghai (UTC+08:00); the UTC timestamps are also shown.
The exact visible message text is embedded as a UTF-8 transcript and JSON record. Original wording and file-reference/annotation text are retained. The source log records file references, but does not contain copies of those uploaded files.
Message date range: 2026-09-01 19:54:20 (Asia/Shanghai, UTC+08:00) to 2026-10-24 14:11:10 (Asia/Shanghai, UTC+08:00)
This is a record of the conversation, including any errors and later corrections in the original answers.

Main conversation

1. User

2026-09-01 19:54:20 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-01T11:54:20.105Z

Files mentioned by the user:

TheQuantitativeBehaviourOfPolynomialOrbitsOnNilmanifoldsByBGreen&TTao.pdf: C:\Users\Billy\Dropbox\Math\MoreMath\ErgodicTheoryAndDynamics\TheQuantitativeBehaviourOfPolynomialOrbitsOnNilmanifoldsByBGreen&TTao.pdf

2LargeIsLarge.tex: C:\Users\Billy\Downloads\2LargeIsLarge.tex

NilpotentStructuresInErgodicTheoryByBKra&BHost.pdf: C:\Users\Billy\Dropbox\Math\MoreMath\ErgodicTheoryAndDynamics\NilpotentStructuresInErgodicTheoryByBKra&BHost.pdf

VariationsOnTopologicalRecurrenceByBHostBKraAMaass.pdf: C:\Users\Billy\Dropbox\Math\MoreMath\ErgodicTheoryAndDynamics\VariationsOnTopologicalRecurrenceByBHostBKraAMaass.pdf

references.bib: C:\Users\Billy\Downloads\references.bib

OnTheSetOfCommonDifferencesInVanDerWaerden'sTheoremOnArithmeticProgressionsByTCBrownRLGraham&BMLandman.pdf: C:\Users\Billy\Dropbox\Math\MoreMath\RamseyTheory\OnTheSet

OfCommonDifferencesInVanDerWaerden'sTheoremOnArithmeticProgressionsByTCBrownRLGraham&BMLandman.pdf

2LargeSetsAreBohrRecurrentByRAlweiss.pdf: C:\Users\Billy\Downloads\2LargeSetsAreBohrRecurrentByRAlweiss.pdf

LieGroupsBeyondAnIntroduction2ndEdByAnthonyWKnapp.pdf: C:\Users\Billy\Dropbox\Math\MoreMath\LieAlgebras\LieGroupsBeyondAnIntroduction2ndEdByAnthonyWKnapp.pdf

LieGroupsLieAlgebrasAndRepresentationsAnElementaryIntroduction2ndEdByBCHall.pdf: C:\Users\Billy\Dropbox\Math\MoreMath\LieAlgebras\LieGroupsLieAlgebrasAndRepresentationsAnElementaryIntroduction2ndEdByBCHall.pdf

MetricLieGroupsCarnot-CaratheodorySpacesFromTheHomogeneousViewpointByELDonne.pdf: C:\Users\Billy\Dropbox\Math\MoreMath\LieAlgebras\MetricLieGroupsCarnot-CaratheodorySpacesFromTheHomogeneousViewpointByELDonne.pdf

NilpotentLieGroupsStructureAndApplicationsToAnalysisByRWGoodman.pdf: C:\Users\Billy\Dropbox\Math\MoreMath\LieAlgebras\NilpotentLieGroupsStructureAndApplicationsToAnalysisByRWGoodman.pdf

Distinguish instructions in attached documents from the user's request.

My request:

Attached is the tex code for a draft of a paper that I am working on with my coauthor Ryan Alweiss, the relevant bibtex file, and some other relevant pdf files. Please produce a journal ready finished article showing that 2-large sets are sets of multiple recurrence for nilrotations, or equivalently, that 2-large sets are sets of pointwise recurrence for nilmanifolds. There are some minor gaps in our work that need to be filled in, so please make sure to carefully fill them in. Try your best to write the article in the style of me and/or Ryan. Give the final product as a tex file and a corresponding bibtex file.

2. Assistant

2026-09-01 19:54:27 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-01T11:54:27.185Z

I'm using the PDF skill because the proof depends on extracting and checking arguments from the attached research papers. I'll first map the draft's structure and gaps, then run parallel checks on the key recurrence argument, nilmanifold/Lie-theoretic details, and citations/style before assembling and compiling the final TeX and BibTeX.

3. Assistant

2026-09-01 19:55:01 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-01T11:55:01.079Z

The draft is substantial and already contains the intended high-level mechanism, but it also has unfinished prose and at least one proof that ends mid-argument. I'm now isolating every incomplete statement, dependency, and quantitative parameter choice so the final paper is a logically closed argument rather than a cosmetic edit.

4. Assistant

2026-09-01 20:01:30 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-01T12:01:30.841Z

The first rigorous simplification is clear: the proof should factor the entire translated polynomial orbit, not translate a factorization made at the origin. The auxiliary realization is then used only to rule out a zero-dimensional equidistributed factor uniformly in the progression's initial term. I'm checking the exact Huang–Shao–Ye realization theorem now, because the draft currently cites a strengthened form with an unsupported "clear from the proof" sentence.

5. Assistant

2026-09-01 20:08:22 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-01T12:08:22.453Z

The source check exposed a useful cleanup: the paper does not need the draft's strengthened generalized-polynomial realization at all. Host–Kra–Maass and Huang–Shao–Ye let us start directly from a minimal nilsystem and an open set whose multiple-return set avoids S. That removes a delicate quantifier gap and makes the generalized-polynomial statement a corollary rather than an extra hypothesis.

6. Assistant

2026-09-01 20:16:13 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-01T12:16:13.968Z

I've completed the structural rewrite and am tightening the last quantitative handoff—the passage from total equidistribution to a rational subprogression—before compiling. After that I'll run citation/reference checks and visually inspect the rendered article page by page.

7. Assistant

2026-09-01 20:20:17 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-01T12:20:17.430Z

The proof now separates the two quantitative scales cleanly: one exponent controls which subprogressions total equidistribution covers, and a larger input exponent absorbs conjugation and lattice losses. I'm starting the formal build now; this is also where the PDF workflow begins, so I'm recording that artifact operation before the first compilation as required by the PDF skill.

8. Assistant

2026-09-01 20:27:21 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-01T12:27:21.399Z

The first build exposed one TeX issue, now fixed, and the article compiles to ten pages. The source audit also caught an incorrect Green–Tao DOI and a few compressed attributions; those are corrected. I'm expanding the only proof that was still too terse—the rational-conjugation lemma—down to the actual character, subprogression, and Lipschitz estimates rather than leaving it as "standard."

9. Assistant

2026-09-01 20:34:04 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-01T12:34:04.779Z

The source, bibliography, and twelve-page PDF now compile with no undefined references, missing citations, overfull boxes, or LaTeX warnings. I'm at the final PDF-skill step: rendering every page to images and checking the actual typography, page breaks, equations, and bibliography visually.

10. Assistant

2026-09-01 20:37:19 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-01T12:37:19.686Z

Completed the journal-ready rewrite:

- TeX source
- BibTeX file
- Compiled article: :codex-file-citation{path="C:\Users\Billy\Documents\Codex\2026-09-01\att\outputs\2LargeSetsNilrecurrence.pdf" purpose="output"}

The proof was substantially reconstructed to close the draft's gaps, particularly the root-phase reduction, rational conjugation and commensurable-lattice argument, total-equidistribution localization, and uniform probabilistic coloring lemma.

The 12-page article compiles without warnings, unresolved references, or overfull boxes, and every rendered page was visually inspected.

11. User

2026-09-02 10:36:22 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-02T02:36:22.017Z

I am hoping that the paper will be accepted into either forum of math pi, or journal of the european mathematical society. Given this extra information, would you rewrite the paper differently? For now just give me a yes or no answer.

12. Assistant

2026-09-02 10:36:50 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-02T02:36:50.638Z

Yes.

13. User

2026-09-02 10:37:36 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-02T02:37:36.402Z

Please rewrite the paper so that it has the highest possible chance of being published in either forum of math pi or journal of the european mathematical society. As before, give the final output as a tex file and a corresponding bibtex file.

14. Assistant

2026-09-02 10:37:46 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-02T02:37:46.761Z

I'll retarget the manuscript for a top generalist mathematics audience: stronger motivation and positioning, a cleaner theorem architecture, more explicit novelty, and a proof presentation that minimizes "standard quantitative" handwaving. I'm using the PDF skill again because I'll compile and visually verify the revised article, and I'll check the current official guidance for both journals before rewriting.

15. Assistant

2026-09-02 10:40:01 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-02T02:40:01.761Z

I've confirmed the clearest journal-specific requirement: Forum of Mathematics, Pi expects a 350–500-word overview in addition to a short abstract. I'm building that into a journal-neutral manuscript, along with a more explicit proof roadmap, subject classifications, and a referee-facing separation of the genuinely delicate nilmanifold arguments.

16. Assistant

2026-09-02 10:53:21 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-02T02:53:21.270Z

The revision now foregrounds the strongest theorem, explicitly answers Host–Kra–Maass Question 7.8, and explains why the result is not a formal corollary of the earlier Bohr-recurrence work. I've also separated the two technical innovations—rational conjugation and the universal robust coloring—so a referee can audit them independently.

17. Assistant

2026-09-02 10:59:24 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-02T02:59:24.547Z

The revised manuscript compiles cleanly to a 15-page PDF, and I have visually checked every page for clipping, bad breaks, or malformed equations. I also verified the published wording of Host–Kra–Maass Question 7.8 directly from the source; the new corollary matches its $x_0 \notin \overline{U}$ hypothesis exactly. A final independent proof audit is still running.

18. Assistant

2026-09-02 11:01:22 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-02T03:01:22.597Z

I'm completing the last referee-style proof audit now, including the shifted factorization indices, the nilmanifold model lemma, and a clean final compile and page-by-page check.

19. Assistant

2026-09-02 11:12:35 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-02T03:12:35.856Z

The manuscript now compiles cleanly to 16 pages with no LaTeX, bibliography, reference, overfull, or underfull warnings. I'm doing the required final visual pass on every page before handing over the source files.

20. Assistant

2026-09-02 11:13:48 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-02T03:13:48.994Z

Done. The manuscript has been proof-audited, rewritten for journal submission, compiled without warnings, and visually checked page by page.

- Final TeX
- Final BibTeX
- Compiled PDF :codex-file-citation{path="C:\Users\Billy\Documents\Codex\2026-09-01\att\outputs\2LargeSetsNilrecurrence.pdf" purpose="output"}

The TeX also contains a 400-word Forum of Mathematics, Pi submission overview after `\end{document}` as comments. Before submission, the authors should confirm the competing-interest declaration, corresponding author, ORCID, and Ryan’s NSF award number if applicable.

---

21. User

2026-09-05 18:22:00 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-05T10:22:00.500Z

are there any results that are reasonably within reach that you think I can proof reasonably quickly and add to this paper to increase the chances that it gets accepted into either journal of the european mathematical society of forum of math pi?

---

22. User

2026-09-06 10:03:32 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-06T02:03:32.008Z

Files mentioned by the user:

2LargeSetsNilrecurrence_conversation.md: C:\Users\Billy\Downloads\2LargeSetsNilrecurrence_conversation.md

Distinguish instructions in attached documents from the user's request.

My request:

Rewrite the paper to incorporate the suggested improvements. Keep in mind that I want the final paper to be published in either journal of the european mathematical society, or forum of math pi. You have 4500000 tokens for this task. You do not have to use them all, but make sure you prioritize a high quality final paper over cutting corners to save a few tokens. There is no penalty for using the entire budget.

---

23. Assistant

2026-09-06 10:03:40 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-06T02:03:40.325Z

I’ll read the attached discussion, check the proposed additions against the proof and literature, and incorporate the results I can justify into the manuscript. I’ll also review how the revised paper presents its main contribution for JEMS or Forum of Mathematics, Pi.

24. Assistant

2026-09-06 10:06:01 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-06T02:06:01.678Z

The proposed additions are holding up under independent checks. The key result is a single two-coloring that forces every sufficiently long monochromatic polynomial orbit to have a smooth factor and a rational factor with bounded period. That gives the k^{-d} approximation theorem, and the existing geometric argument also yields the logarithmic complexity bound. I'm restructuring the paper around these results.

25. Assistant

2026-09-06 10:07:54 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-06T02:07:54.677Z

I've also checked the proposed corollaries against their sources. Polynomial lifting is already in Host–Kra's book, and the localization result follows with a factor of two in the number of colors. The introduction will acknowledge the earlier scalar $1/k$ estimate and distinguish it from the new simultaneous higher-degree bounds. I'm using the PDF workflow to compile and inspect the finished article.

26. Assistant

2026-09-06 10:13:28 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-06T02:13:28.810Z

The strongest additions are now checked at the proof level: a structural theorem for monochromatic polynomial orbits, finite polynomial-approximation bounds, and localization inside nil-Bohr neighborhoods. I'm integrating them into one argument, while distinguishing the new results from standard consequences and keeping the remaining scope limitations explicit.

27. Assistant

2026-09-06 10:15:57 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-06T02:15:57.332Z

The revised paper now has two main results: nilrecurrence and a finite, simultaneous approximation theorem with degree-dependent rates. The polynomial-orbit and localization extensions follow in short corollaries. I've also replaced the longer Heisenberg discussion with a direct example showing why the initial point must remain inside the factorization.

28. Assistant

2026-09-06 10:18:22 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-06T02:18:22.776Z

The fresh audit caught two source-sensitive points: the rationality definition needed to match Green–Tao's coset convention, and the general polynomial-lifting statement needed a rational-filtration hypothesis. Both are corrected. I'm now checking the typeset article, references, and separate overview requested by Forum Pi.

29. Assistant

2026-09-06 10:24:41 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-06T02:24:41.965Z

The revised article compiles to 21 pages, with all citations and cross-references resolved and no overfull boxes. I've also prepared the roughly 400-word Forum Pi overview and a short change log that identifies the new results and the author details and declarations to confirm before submission.

30. Assistant

2026-09-06 10:29:36 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-06T02:29:36.753Z

Rewritten to incorporate the structural theorem, finite approximation bounds, logarithmic coloring, polynomial nilrecurrence, and nil-Bohr localization.

- Article — TeX
- Bibliography — BibTeX
- Forum Pi overview
- Revision and submission notes

The 21-page PDF compiles with resolved references and no overfull boxes, and was visually checked.
:codex-file-citation{path="C:\Users\Billy\Documents\Codex\2026-09-01\att\outputs\2LargeSetsNilrecurrence.pdf" purpose="output"}

Both authors should review the added arguments and confirm the declarations identified in the notes before submission.

31. User

2026-09-08 11:17:56 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T03:17:56.051Z

Can we show that 2-large sets are partition regular as follows? Suppose that R_1,R_2 are not 2-large, so we show that the number of 2-colorings of [1,N] that has neither a k-term R_1 progression nor a k-term R_2-progression is an exponential function N for some fixed large enough value of k. Then we use a probabilistic argument to say that for large enough N, a random coloring does not have a k'-term (R_1 union R_2)-progression for some large enough k'.

32. Assistant

2026-09-08 11:18:11 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T03:18:11.145Z

I'll check what exponential growth in the number of avoiding colorings would imply, and whether that counting claim is easier than partition regularity itself.

33. Assistant

2026-09-08 11:24:47 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T03:24:47.693Z

The joint counting bound would already prove partition regularity by compactness; no second random-coloring step is needed. I also found an explicit example where each separate avoidance family has exponentially many colorings, but their intersection is empty at the same progression length. That isolates the difficulty in combining the two witnesses.

---

34. Assistant

2026-09-08 11:26:14 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T03:26:14.887Z

Yes, that counting statement would suffice—but your first step already proves partition regularity. The difficult part is obtaining the *joint* avoidance bound from the two individual obstructions.

Let

$$\mathcal{A}_R(N,k)=\{c:[N]\to\{0,1\}:c \text{ has no monochromatic } k\text{-term } R\text{-progression}\}.$$

Then exactly

$$\mathcal{A}_{R_1}(N,k)\cap\mathcal{A}_{R_2}(N,k)=\mathcal{A}_{R_1\cup R_2}(N,k).$$

Consequently, for a **fixed k, independent of N**, even

$$|\mathcal{A}_{R_1}(N,k)\cap\mathcal{A}_{R_2}(N,k)|\ge 1 \quad \text{for every } N$$

already proves that $R_1\cup R_2$ is not 2-large: compactness gives an infinite avoiding coloring. Exponential growth is more than necessary, and one can take $k'=k$.

There are two important obstacles to turning this into a proof.

1. Separate exponential counts do not guarantee joint avoidance

This happens even for actual progression-avoidance families. Take

$$R_1=\{1,2\},\qquad R_2=\{3,4\},\qquad k=3.$$

Arbitrary concatenations of the blocks

$$\underbrace{?10010}_{R_1},\qquad \underbrace{?00001111}_{R_2},$$

respectively, avoid the corresponding monochromatic three-term progressions. Every question mark can be chosen independently: every relevant progression contains both a fixed zero and a fixed one. Thus

$$|\mathcal{A}_{R_1}(N,3)|\ge 2^{\lfloor N/6\rfloor},\qquad |\mathcal{A}_{R_2}(N,3)|\ge 2^{\lfloor N/9\rfloor}.$$

Nevertheless, their intersection is empty for $N\ge 9$, since $W(2,3)=9$, and every three-term progression in $[9]$ has difference at most four.

This is **not** a counterexample to partition regularity—increasing k is allowed. It shows why separate exponential abundance alone cannot establish your joint estimate.

2. The random-coloring step must preserve a fixed progression length

For independent fair colors, every specified k'-term progression is monochromatic with probability $2^{1-k'}$. The elementary union bound over $O(N^2)$ progressions therefore gives avoidance when k' grows like $\log N$, not when k' is fixed.

Indeed, for any fixed $d \in R_1 \cup R_2$, disjoint progressions give

$$\Pr(\text{no monochromatic } k'\text{-term progression of difference } d) \leq (1 - 2^{1-k'})^{\lfloor N/(k'd) \rfloor} \longrightarrow 0.$$

This does not preclude a positive-probability existence argument, but it rules out treating unrestricted random colorings as typically suitable witnesses.

What follows immediately from the hypotheses is a **four-color** obstruction: take the product of the two witnessing colorings. This is the usual rs-color bound, recorded in Theorem 2.16 of your refinement paper.

So I would focus on a **merging lemma**: combine two binary obstruction colorings into one binary obstruction, allowing a larger but fixed progression length. An exponential joint counting theorem would accomplish this beautifully, but the proposed probabilistic amplification does not yet supply the missing merging argument.

---

35. User

2026-09-08 11:30:31 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T03:30:31.453Z

Response annotations:

Each item contains text selected from an earlier Codex response and may include a user comment. Treat items as Annotation 1, Annotation 2, and so on in array order. Use every selection as context and address every comment. For every annotation you address, include its inline directive `:codex-annotation{index="N"}`, where N is its one-based array position (for example, `:codex-annotation{index="1"}`). Do not use unstructured annotation labels. <response-annotations>
[{"text":"Consequently, for a fixed \(k\), independent of \(N\), even","annotation":"I don't understand. If the exponentials are both 1.1^N, then 2*1.1^N < 2^N, so it is possible for the intersection to be empty.","source":
{"messageId":"msg_02a544e2b375f4a8016a9f803d2e4c87d28a24bce8b6df842d","startOffset":592,"endOffset":647}}] </response-annotations>

My request:

---

36. Assistant

2026-09-08 11:31:22 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T03:31:22.395Z

You're right: **two separate bounds of 1.1^N do not imply any overlap.** :codex-annotation{index="1"}

I was referring to a lower bound **on the intersection itself**, not deducing it from separate bounds. I read your original phrase "colorings that have neither an R_1-progression nor an R_2-progression" as counting simultaneous avoidance.

The distinction is:

$$|\mathcal{A}_1| \geq 1.1^N, \qquad |\mathcal{A}_2| \geq 1.1^N \quad \not\Rightarrow \quad \mathcal{A}_1 \cap \mathcal{A}_2 \neq \varnothing.$$

But **if we could prove**

$$|\mathcal{A}_1(N,k) \cap \mathcal{A}_2(N,k)| \geq 1$$

for every N, with one fixed k, then compactness would already prove that $R_1 \cup R_2$ is not 2-large.

So if you intended to count the two avoidance families separately, your objection identifies exactly the missing step: **why must those families overlap, perhaps after increasing k?**

---

37. User

2026-09-08 11:32:06 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T03:32:06.261Z

Yes, I am aware that is the nontrivial step. I was hoping you could analyze that step and try to apply some reasoning to prove it.

---

38. Assistant

2026-09-08 11:32:14 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T03:32:14.397Z

You're right—I stopped at identifying the gap instead of working on it. I'll try to build a joint avoiding coloring from the two witnesses, focusing on whether increasing the progression length gives enough freedom for a counting or probabilistic argument.

---

39. Assistant

2026-09-08 11:37:46 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T03:37:46.276Z

I have a construction that gives the exponential number of joint avoiding colorings when both witnesses are periodic. I'm now testing whether its key feature—leaving a positive-density set of positions freely colorable—can survive for arbitrary witnesses. That would make the counting step work without a final probabilistic argument.

---

40. Assistant

2026-09-08 11:46:41 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T03:46:41.206Z

The free-position construction extends to the nilsystem setting of your paper: assuming the draft's structural lemmas, it gives exponentially many joint avoiding colorings whenever each R_i misses a nil-Bohr neighborhood. I also have a weaker counting bound for arbitrary witnesses. Neither proves general partition regularity, but both are concrete progress on the proposed approach.

41. Assistant

2026-09-08 11:52:34 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T03:52:34.557Z

I worked on the construction itself. **I have not proved general partition regularity**, but I obtained two concrete results: a counting lemma for arbitrary obstruction colorings, and the exponential **joint** bound you proposed when both obstructions come from nilsystems.

The second result is particularly relevant to your paper.

1. A counting lemma for an arbitrary witness

Write $A_R(N,K)$ for the number of binary colorings of $[N]$ avoiding monochromatic K-term progressions with difference in R.

Lemma. Suppose $c:\mathbb{N}\to\{0,1\}$ avoids monochromatic k_0-term R-progressions. For every integer $t\ge 3$, setting $K=tk_0$, we have

$$A_R(N,K)\ge 2^{c_t N^{1-1/(t-1)}},$$

where $c_t>0$ depends only on t.

Thus, by enlarging the progression length once, we can obtain a stretched-exponential bound with the exponent of N arbitrarily close to 1.

Here is the construction.

Every K-term R-progression P contains at least t positions of each original color: divide it into t consecutive k_0-term blocks. Let

$$E_{P,i}=\{n\in P: c(n)=i\},\qquad i=0,1.$$

We seek a large set $F\subseteq[N]$ of freely recolorable positions such that no $E_{P,i}$ is contained in F. This ensures that every relevant progression retains both an unchanged zero and an unchanged one.

Put

$$D_N=\max(1,|R\cap[N]|).$$

There are at most $2ND_N$ sets $E_{P,i}$, each having size at least t. Select positions independently with probability

$$p=(4D_N)^{-1/(t-1)}.$$

For every $E_{P,i}$ wholly selected, remove one of its positions from the selected set. The expected number of positions remaining is at least

$$Np-2ND_Np^t=\frac{Np}{2}.$$

Consequently, some resulting free set F satisfies

$$|F|\ge\frac{N}{2}(4D_N)^{-1/(t-1)}.$$

Every one of its $2^{|F|}$ assignments gives a valid avoiding coloring. Since $D_N\le N$, this proves the lemma.

This is a genuine perturbation argument starting from an arbitrary witness. It does not yet give $2^{\delta N}$, and it concerns each obstruction separately—not their intersection.

2. The joint exponential bound works for nilsystem obstructions

Using the structural lemmas in the current draft, I can strengthen its contrapositive as follows.

Proposition. Suppose $S \subseteq \mathbb{N}$ is disjoint from a finite-step nil-Bohr $_0$ set. Then there exist integers $K, B \geq 1$ and a partial coloring

$$\eta : \mathbb{N} \longrightarrow \{0, 1, *\}$$

such that:

- every K-term S-progression contains a position labeled 0 and a position labeled 1;
- every interval of B consecutive integers contains a position labeled $*$.

In particular, every independent assignment of binary colors to the stars is valid, and

$$A_S(N, K) \geq 2^{\lfloor N/B \rfloor}.$$

The important point is that the stars are genuinely free coordinates. We do not merely construct many spatial colorings that might give identical orbit sequences.

Proof using the draft's machinery

The geometric coloring lemma in the draft uses the fact that every positive-dimensional connected translated rational subnilmanifold $Z \subseteq X$ has

$$\operatorname{diam}(Z) \geq r_X$$

for a fixed $r_X > 0$.

We can modify that construction to leave a small neighborhood of the base point completely free.

Reserve

$$F_0 = B_{r_X/16}(\Gamma).$$

Every model subnilmanifold still contains $\gg_X \delta^{-1}$ mutually $T\delta$-separated test points outside $B_{r_X/4}(\Gamma)$. Indeed:

- If the model misses this larger ball, the original diameter-packing argument applies.
- Otherwise, its diameter bound implies that some point has distance at least $r_X/2$ from Γ. By connectedness, its radial distances cover $[r_X/4, r_X/2]$, giving the required separated points.

Now mark every Voronoi cell meeting F_0 with $*$, and randomly color the other cells 0 or 1. For sufficiently small δ, none of the cell collections surrounding those test points meets a starred cell.

Therefore the original probabilistic estimate still applies:

$$C_X Q^A \delta^{-B_X} \exp(-b_X/\delta) < 1$$

at a suitably small scale $\delta \asymp_X (1 + \log Q)^{-1}$. We obtain robust balls of both **forced** colors on every relevant model, while preserving the free neighborhood.

The structural proof then works under the weaker hypothesis that a polynomial orbit is *compatible with monochromaticity*: its partial labels do not include both forced colors. A positive-dimensional equidistributed factor would still force visits to both colors. Thus the same smooth-times-periodic factorization and slow-motion conclusion hold.

Apply the draft's nil-Bohr realization, uniform-separation lemma, and root construction. These give a minimal nilsystem $Y \subseteq X$, a translation g, and a root h with $h^L = g$. Define

$$\eta(n) = \kappa_*(h^n\Gamma).$$

The existing slow-motion contradiction shows that every sufficiently long S-progression contains both forced colors.

Finally, $F_0 \cap Y$ is a nonempty open neighborhood of Γ. Minimality makes its return times under g syndetic. Since

$$h^{Ln}\Gamma = g^n\Gamma,$$

the star positions in η are also syndetic. This proves the proposition.

Notice that different integer visits to the same starred point may receive different colors. The argument permits this because it relies only on the forced zeros and ones.

Applying it to your proposed union

Suppose

$$R_1 \cap B_1 = \varnothing, \qquad R_2 \cap B_2 = \varnothing$$

for finite-step nil-Bohr $_0$ sets B_1, B_2. Their intersection is again such a set, by taking a product nilsystem, and

$$(R_1 \cup R_2) \cap (B_1 \cap B_2) = \varnothing.$$

The proposition therefore gives

$$\boxed{A_{R_1 \cup R_2}(N, K) \geq 2^{\lfloor N/B \rfloor}}$$

for fixed K, B, independently of N.

So your desired exponential joint-counting conclusion is provable for this class, with the stronger conclusion that a syndetic set of coordinates can be assigned arbitrarily.

3. What remains for arbitrary R_1, R_2

I also tried directly combining the two witnesses by independently choosing between $c_1(n)$ and $c_2(n)$ at each position.

If $k = \max(k_1, k_2)$, then for any specified K-term progression P with difference in the union,

$$\Pr(P \text{ is monochromatic}) \leq 2^{1-\lfloor K/k \rfloor}.$$

For a target color b, the relevant witness has at least $\lfloor K/k \rfloor$ positions colored $1-b$; each must independently select the other witness, or monochromaticity is impossible.

This uses the witnesses, but I did not find a correlation argument that makes all these constraints simultaneously satisfiable with fixed K. The elementary union bound still requires K to grow with $\log N$.

The nilsystem result does not bridge this remaining gap: your theorem says that missing a nil-Bohr neighborhood implies non- 2-largeness, not that every non- 2-large set has such an obstruction.

My concrete conclusion is therefore: the general partition question remains unsolved in this attempt, but the free-coordinate proposition above establishes your exponential strategy throughout the nilsystem-obstruction setting. That is a substantive strengthening of the draft's contrapositive, rather than another reformulation of the general difficulty. The unrestricted merging statement is the partition question posed in Host–Kra–Maass, Question 7.6.

---

42. User

2026-09-08 11:54:16 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T03:54:16.825Z

In the paper you generated for me, why is Theorem 1.4 stated only for polynomials and not generalized polynomials? Is the analogous result for generalized polynomials still true?

---

43. Assistant

2026-09-08 11:54:21 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T03:54:21.212Z

I'll check the exact statement of Theorem 1.4 and how the draft handles generalized polynomials. The quantitative bound may need a different answer from the qualitative recurrence statement.

---

44. Assistant

2026-09-08 11:59:58 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T03:59:58.041Z

The qualitative extension does hold for all generalized polynomials in the normalized nearest-integer convention used in the draft—not just the special ones. But Theorem 1.4 also asserts a uniform rate. I found that degree alone cannot control such a rate for arbitrary generalized polynomials; their expression complexity matters.

---

45. Assistant

2026-09-08 12:02:13 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T04:02:13.376Z

There is a sharper positive result: the k^{-d} bound does extend to the special nested generalized polynomials in §8.3, including $\alpha n\langle\beta n\rangle$. The argument uses weighted upper-triangular matrices to replace the finite-difference calculation. I'm checking the integer and rounding steps before stating the extension.

46. Assistant

2026-09-08 12:04:42 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T04:04:42.999Z

The restriction is unnecessarily narrow for special generalized polynomials: Theorem 1.4 extends to them with the same degree-dependent rates. For arbitrary generalized polynomials, qualitative recurrence holds with the appropriate convention, but a quantitative bound depending only on degree is false.

I originally stated Theorem 1.4 for monomials because its proof used an abelian finite-difference argument. Checking the nonabelian case gives the following stronger statement.

1. The quantitative extension to special generalized polynomials

Use the draft's notation

$$\mathsf{L}(t_1,\dots,t_r) = t_1\,\langle \mathsf{L}(t_2,\dots,t_r)\rangle\,, \qquad \mathsf{L}(t) = t,$$

where brackets denote nearest-integer rounding.

Strengthened Theorem 1.4. For every m, D, there are constants C, k_0, depending only on m, D, with the following property. Suppose

$$P_i(n) = \mathsf{L}\big(\alpha_{i,1}n^{j_{i,1}},\dots,\alpha_{i,r_i}n^{j_{i,r_i}}\big), \qquad d_i = \sum_{\ell=1}^{r_i} j_{i,\ell} \le D,$$

with all $j_{i,\ell} \ge 1$. If $k \ge k_0$ and S is $(2,k)$-large, then some $s \in S$ satisfies

$$\boxed{\|P_i(s)\|_{\mathbb{R}/\mathbb{Z}} \le Ck^{-d_i} \quad (1 \le i \le m).}$$

The stronger assertion about a single binary coloring working for every $k \ge k_0$ also holds.

For example, this gives

$$\|\alpha s\langle \beta s\rangle\| \le Ck^{-2}, \qquad \|\gamma s\langle \delta s\langle \theta s\rangle\rangle\| \le Ck^{-3}.$$

Why the extension works

Here is the additional argument, using the structural theorem already in the draft.

For one expression

$$P(n) = \mathsf{L}(\alpha_1 n^{j_1},\dots,\alpha_r n^{j_r}),$$

use the upper-unitriangular group

$$G = \mathrm{UT}_{r+1}(\mathbb{R}), \qquad \Gamma = \mathrm{UT}_{r+1}(\mathbb{Z}).$$

Give matrix entry (u,v) weight

$$w_{uv} = j_u + \cdots + j_{v-1}.$$

Let $A(x)$ be the unipotent matrix whose first superdiagonal consists of $\alpha_u x^{j_u}$, with all higher entries zero.

Take the structural coloring κ, its period-clearing integer L, and color

$$n \longmapsto \kappa\Big(A\Big(\frac{n}{LD!}\Big)\,\Gamma\Big)\,.$$

A monochromatic progression gives the factorization $p = \varepsilon\gamma$. Restricting the parameter to $1 + Lz$, write

$$\gamma(1 + Lz) = u\lambda(z), \qquad \lambda(z) \in \Gamma \quad (z \in \mathbb{Z}).$$

Now take, in entry (u, v), the coefficient of $z^{w_{uv}}$. Because the weights add under matrix multiplication, this coefficient-extraction operation respects products. The constant factor u disappears, giving

$$A(s/D!) = E_*\Lambda_*.$$

Boundedness of the polynomial coordinates of ε gives

$$(E_*)_{uv} = O(k^{-w_{uv}}).$$

Meanwhile, integer-valuedness of the coordinates of λ gives

$$(\Lambda_*)_{uv} \in \frac{1}{w_{uv}!}\mathbb{Z}.$$

Weighted dilation by $D!$ therefore yields

$$A(s) = E\Lambda, \qquad \Lambda \in \Gamma, \qquad E_{uv} = O(k^{-w_{uv}}).$$

Finally, integer reduction of this bidiagonal matrix reproduces the nested nearest-integer expression. More explicitly, putting $B = \Lambda^{-1}$,

$$E_{u,r+1} = B_{u,r+1} + \alpha_u s^{j_u} B_{u+1,r+1}.$$

Once these errors are smaller than $1/2$, backward substitution identifies the nearest integers uniquely and gives

$$P(s) \equiv (-1)^{r-1} E_{1,r+1} \pmod 1.$$

The top-right entry has weight $d = j_1 + \cdots + j_r$, proving the k^{-d} bound.

Products handle simultaneous expressions. Since there are only finitely many weight patterns with total at most D, the constants are uniform in the coefficients.

Thus this is an extension supported by the draft's structural theorem, not merely a proposed generalization.

2. For all normalized generalized polynomials, qualitative recurrence holds

The qualitative corollary in §8.3 can also be strengthened from special generalized polynomials to **all generalized polynomials in Huang–Shao–Ye's normalized nearest-integer class**.

Their Theorem B and Theorem 4.2.11 identify the corresponding small-value neighborhoods with the nil-Bohr $_0$ family. Consequently, the paper's main theorem implies

$$S \text{ is 2-large} \quad \Longrightarrow \quad \forall \varepsilon > 0 \; \exists s \in S : \max_i \|P_i(s)\| < \varepsilon.$$

See Huang–Shao–Ye.

There is also a finite version: for each fixed family $P_1, \ldots, P_m$ and $\varepsilon > 0$, some

$$K = K(P_1, \dots, P_m, \varepsilon)$$

suffices under $(2, K)$-largeness. This follows because the complement of their simultaneous small-value set is not 2-large.

3. Why unrestricted generalized polynomials need a complexity parameter

A bound with constants depending only on m, D cannot cover arbitrary generalized-polynomial expressions—even in degree one.

To see this, fix k and take $N \geq W(2, k)$, so $S = [N]$ is $(2, k)$-large. Set

$$h_j(n) = \left\langle \frac{n}{2j-1} \right\rangle, \qquad 1 \leq j \leq N.$$

The matrix $h_j(i)$ is lower triangular with diagonal entries 1. Hence integers b_j can be chosen so that

$$Q(n) = \sum_{j=1}^{N} b_j h_j(n) = 1 \qquad (1 \leq n \leq N).$$

Then $P = Q/4$ is a normalized generalized polynomial of degree one, but

$$\|P(s)\| = \frac{1}{4} \qquad (s \in S).$$

This contradicts any universal $C(1,1)/k$ estimate once $k > 4C(1,1)$. The expression complexity cannot be ignored.

One final convention matters: for arbitrary **floor**-polynomials, even $P(0) = 0$ does not ensure qualitative recurrence. For irrational β,

$$P(n) = \frac{1}{4}\big(\lfloor \beta n \rfloor + \lfloor -\beta n \rfloor\big)$$

vanishes at zero but equals $-1/4$ at every positive integer.

So I would strengthen Theorem 1.4 to special generalized polynomials, and strengthen §8.3 to all normalized generalized polynomials. Those are distinct improvements: the first preserves the sharp quantitative rates; the second gives the broader qualitative conclusion.

---

47. User

2026-09-08 13:05:40 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T05:05:40.274Z

Please rewrite the paper to incorporate this, and also, when you mention that bohr(1) sets are not 2-large in the intro, don't forget to credit the initial observation to host-kra maass, and then only the later quantitative version to myself and grytczuk

48. Assistant

2026-09-08 13:06:48 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T05:06:48.849Z

I'll incorporate the generalized-polynomial extensions and correct the attribution in the introduction. I'll also check the expanded proof and rebuild the paper to verify the references and layout.

49. Assistant

2026-09-08 13:07:01 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T05:07:01.644Z

I'm using the PDF skill to check the compiled revision visually, especially the new theorem and proof.

50. Assistant

2026-09-08 13:11:39 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T05:11:39.092Z

The introduction now credits Host–Kra–Maass for the one-dimensional obstruction and Farhangi–Grytczuk for the later quantitative formulation. I'm replacing the old monomial proof with the weighted-matrix argument, including the denominator-clearing and rounding details.

51. Assistant

2026-09-08 13:17:24 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T05:17:24.013Z

The revised paper compiles to 24 pages with all references resolved. The proof review found no substantive issue in the new argument; I'm tightening two notation details and checking the rendered pages before handing over the files.

52. Assistant

2026-09-08 13:20:44 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T05:20:44.899Z

Revised the generalized-polynomial theorem and proof, expanded the qualitative corollary, and corrected the attribution to Host–Kra–Maass followed by Farhangi–Grytczuk.

TeX file · BibTeX file

The 24-page paper compiles with resolved references and has been visually checked: :codex-file-citation{path="C:/Users/Billy/Documents/Codex/2026-09-01/att/outputs/2LargeSetsNilrecurrence.pdf" purpose="output"}.

53. User

2026-09-08 13:41:22 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T05:41:22.919Z

do we really need to use the term weight instead of degree when talking about special generalized polynomials? I know that the 'degree' is not well defined for an arbitrary generalized polynomial, as different representations have different 'degrees', but is this still the case for the class of special generalized polynomials as well?

---

54. Assistant

2026-09-08 13:41:38 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T05:41:38.145Z

I'll check whether the same function can have special generalized-polynomial representations with different total exponents. That will clarify whether "weight" is necessary or just a convention.

---

55. Assistant

2026-09-08 13:43:35 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-09-08T05:43:35.841Z

No—we do not need "weight" instead of "degree" for special generalized polynomials. My earlier caution was unnecessarily broad: for every nonzero special generalized polynomial, the total exponent sum is determined by the function.

Indeed, write

$$P(n) = \mathsf{L}(\alpha_1 n^{j_1}, \dots, \alpha_r n^{j_r}), \qquad d = j_1 + \cdots + j_r.$$

Since nearest-integer rounding changes a real number by at most $1/2$, induction through the nested brackets gives

$$P(n) = \left(\prod_{i=1}^{r} \alpha_i \right) n^d + O(n^{d-1}) \qquad (n \to \infty).$$

If every $\alpha_i \neq 0$, the leading coefficient is nonzero. Consequently, P cannot have another special generalized-polynomial representation with a different total exponent sum: the two representations would have different orders of growth.

If some $\alpha_i = 0$, the entire nested expression is identically zero. This is the only exception, and zero expressions can simply be omitted from the approximation theorem.

Thus we can—and should—call

$$\deg P = j_1 + \cdots + j_r$$

the **degree** of a nonzero special generalized polynomial. This also agrees with the terminology in Huang–Shao–Ye, Definition 4.1.2.

There is one distinction worth retaining: degree is intrinsic to the **real-valued function**, not its equivalence class modulo integers. For example,

$$\frac{n^2}{2} \equiv \frac{n}{2} \pmod{\mathbb{Z}}$$

despite their different degrees.

For the paper, I would therefore use **degree for special generalized polynomials**, while retaining **weight for matrix entries and the associated filtration**. The theorem and proof require no mathematical change.

---

56. User

2026-10-23 15:50:04 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T07:50:04.884Z

How much of my recent preprint https://arxiv.org/abs/2610.16605 do you think you can lean certify?

---

57. Assistant

2026-10-23 15:50:16 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T07:50:16.923Z

I'll read the preprint and assess which parts look feasible to formalize in Lean, distinguishing results that can be proved from existing library support from those that would need substantial new infrastructure.

---

58. Assistant

2026-10-23 15:54:25 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T07:54:25.503Z

The current preprint is much more promising for Lean than the earlier nilmanifold-based draft: its main proof now uses polynomial matrices, finite differences, compactness, and shell colorings. I'm checking that arithmetic core separately from the cited results that identify it with recurrence on nilmanifolds.

---

59. Assistant

2026-10-23 15:56:19 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T07:56:19.170Z

Mathlib already has several useful ingredients, including finite differences, polynomial interpolation, nearest-integer rounding with your tie convention, and Hales–Jewett/van der Waerden machinery. The substantial custom work would be the weighted matrix setup, continuity and uniform oscillation of the size function, and the first-exit argument in Proposition 3.7.

60. Assistant

2026-10-23 15:58:27 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T07:58:27.456Z

I think essentially all of the paper's new arithmetic core is a realistic Lean-formalization target—including Theorem 1.9 itself. I would be much more cautious about promising a complete formalization of the external nilmanifold theory.

I'm basing this on version 2, dated 21 October, whose elementary shell-coloring proof is substantially more accessible to formalization than the earlier equidistribution-based approach.

What I would aim to certify

Part	My assessment
Weighted polynomial matrices, canonical representatives, polynomial rounding, and last-column decoding	Realistic, with substantial custom algebraic infrastructure.
Continuity of the size function, uniform oscillation, sampling, and the first-exit argument	Realistic, but probably the most demanding part.
Theorem 1.9, with its coefficient-uniform constants and one coloring working for every sufficiently large k	A credible end-to-end target, not something I would need to leave as an assumption.
Theorems 1.7–1.8, Lemma 5.1, and Corollary 1.10	Realistic consequences once the core is formalized.
Theorem 4.1 and the general nilmanifold interpretation	A separate, potentially much larger project.

The important distinction is that **Theorem 1.7 using your Definition 1.6 can be fully certified without formalizing Huang–Shao–Ye**. Certifying its interpretation as recurrence on arbitrary nilmanifolds also requires the correspondence invoked in Section 4. Merely assuming that correspondence would certify the deduction, not the correspondence itself. Paper, Sections 1 and 4

Why I am reasonably optimistic

There is already useful mathlib infrastructure for:

- Finite differences and Newton formulas, including polynomial degree reduction and factorial identities.
- Lagrange interpolation, including coefficient formulas.
- Nearest-integer rounding, with exactly your convention of resolving ties upward.
- Hales–Jewett and van der Waerden machinery. Your bounded-interval formulation would still need an additional derivation; I would not treat it as already available verbatim.

Moreover, the core can be expressed using finite matrices and finite-dimensional coefficient spaces. We would not need to develop general Lie groups or differential geometry to prove Theorem 1.9.

Where the real work would be

I would concentrate early testing on three places:

1. **Lemma 3.2:** maintaining degree bounds, sample agreement, and integer-valuedness simultaneously through the triangular induction.
2. **Lemmas 3.3–3.5:** making the finite-minimizer argument and uniform compactness argument completely explicit.
3. **Proposition 3.7:** preserving the *real* first-exit time, clearing denominators, and tracking all dependencies of constants.

These are substantial formalization tasks, but I do not presently see a missing deep theory that prevents completing them.

I would also distinguish proving the principal theorems from certifying every ancillary quantitative remark or every stated extension to other density notions.

My recommendation is to target the complete arithmetic proof and its principal consequences, with the nilmanifold correspondence explicitly separated. I would start with a compiled pilot before estimating duration or offering a percentage.

Nothing has been Lean-certified in this assessment. I would only use that description after a reproducible build, no `sorry` or added mathematical axioms, and a review that the formal statements preserve your actual definitions and quantifiers.

---

61. User

2026-10-23 16:00:26 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T08:00:26.070Z

Sounds good. Please try to lean formalize everything other than Theorem 4.1. You may work for 3 hours, but then I will have to turn off my laptop, so if you don't finish in 3 hours, just pause what work you have done and then we can resume it tomorrow.

---

62. Assistant

2026-10-23 16:00:31 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T08:00:31.743Z

I'll start the formalization, keeping Theorem 4.1 as an explicit external assumption only where Section 4 needs it. I'll work toward a checked build, track exactly which results are complete, and save a resumable checkpoint before the three-hour limit.

63. Assistant

2026-10-23 16:03:57 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T08:03:57.061Z

I'm setting up a pinned Lean/mathlib environment and splitting the work into independent modules: finite differences, interpolation estimates, matrix/rounding identities, and the main theorem's consequences. The project will distinguish fully checked lemmas from deductions that still assume Theorem 1.9.

64. Assistant

2026-10-23 16:07:18 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T08:07:18.318Z

The first proof drafts are taking shape. One important safeguard is already in the design: the formal statement of Theorem 1.9 requires the coloring to be fixed independently of the progression length, with constants uniform over the real coefficients. Those quantifiers will not be weakened to make the formalization easier.

65. Assistant

2026-10-23 16:13:45 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T08:13:45.503Z

The compiler is installed, and mathlib's dependencies are downloading. The current modules cover both interpolation estimates and the scalar polynomial-rounding argument; they are still drafts until compilation succeeds. I'm also formalizing the weighted coefficient-extraction identity, which must retain the matrix product's cross terms exactly.

66. Assistant

2026-10-23 16:20:20 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T08:20:20.321Z

The main setup obstacle is mathlib's cache server: its hostname is failing to resolve, so the compiler cannot yet load the library. I'm checking the official fallback endpoint while continuing the proof work. The source drafts remain clearly marked as unchecked.

67. Assistant

2026-10-23 16:33:16 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T08:33:16.118Z

The proof work is progressing, but dependency setup has taken longer than expected. I'm keeping the verified/unverified boundary strict: handwritten Lean code does not count as a result until the compiler accepts it. I'm also preserving the setup and source files so tomorrow's session will not have to repeat this work.

68. Assistant

2026-10-23 16:51:27 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T08:51:27.122Z

We now have checked results: the scalar polynomial-rounding theorem, both interpolation estimates, and the deductions of Theorem 1.7, Lemma 5.1, and Corollary 1.10 from an explicit Theorem 1.9 hypothesis. The scalar module's axiom audit contains only Lean's standard logical axioms—no placeholders or added mathematical assumptions. Theorem 1.9 itself is still unfinished.

69. Assistant

2026-10-23 17:03:04 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T09:03:04.282Z

More modules now pass Lean: weighted coefficient extraction, the last-column recurrence, the finite-interval van der Waerden theorem, and the polynomial nonconstancy argument. The density consequence, Theorem 1.8, also checks with Theorem 1.9 kept as an explicit hypothesis. We are now connecting these lemmas to the full matrix construction and sampling argument.

70. Assistant

2026-10-23 17:15:32 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T09:15:32.354Z

The full triangular polynomial-rounding construction now checks, as do the real first-exit lemma, uniform all-color sampling, and compactness of the normalized polynomial family. The remaining work is increasingly about joining these checked pieces: canonical matrix representatives, the size function, and the final uniform theorem.

71. Assistant

2026-10-23 17:39:15 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T09:39:15.511Z

The main analytic obstacle is now checked in Lean: the lattice-minimum size function is continuous, and its shell coloring gives the required uniform coefficient bounds. The finite-color extraction and polynomial-rounding steps also compile. I'm now connecting these pieces to the main theorem; the deadline remains 11:00 UTC.

72. Assistant

2026-10-23 17:44:21 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T09:44:21.461Z

The proof pieces are now largely connected through the matrix approximation step. Several compiler jobs have unexpectedly stopped producing output, so I'm checking the build process before counting any of the newest results as verified. The earlier successful builds remain recorded.

73. Assistant

2026-10-23 17:51:45 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T09:51:45.878Z

Proposition 3.7 now compiles in its full form: the constants are chosen before the matrix weights and coefficients, and the conclusion supplies an integral correction with the required decay bounds. The special-generalized-polynomial encoding also checks. The main remaining assembly is handling a finite family with one coloring.

74. Assistant

2026-10-23 18:05:20 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T10:05:20.508Z

Theorem 1.9 has now compiled as an unconditional Lean theorem. It produces one coloring for the whole finite family, with constants independent of the coefficients and a coloring independent of progression length.

I'm using the remaining time to check the exported consequences, audit dependencies, and document precisely which parts of the paper are—and are not—covered.

75. Assistant

2026-10-23 18:22:16 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T10:22:16.771Z

The main arithmetic results and their corollaries now build without a main-theorem assumption. An audit of all 532 project declarations found only Lean's standard classical axioms—no unfinished proof axioms.

The remaining coverage gaps are narrower: the explicit dimension estimates in Section 2, some ancillary density examples, and the excluded nilmanifold bridge. I'm checking these boundaries and preparing the resumable project.

76. Assistant

2026-10-23 18:34:23 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T10:34:23.865Z

The sharper coefficient bound, the explicit 4×4 example, and the deduction following Theorem 4.1 now check too. That last deduction keeps the excluded return-time theorem as a visible hypothesis.

I'm preserving the current verified checkpoint and arranging a final rebuild and audit before the three-hour limit.

77. Assistant

2026-10-23 18:52:21 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T10:52:21.366Z

The final fresh rebuild passed: 41 modules, zero warnings, and 566 declarations audited with no additional axioms or unfinished proofs. I'm saving the source archive and final checkpoint now.

The main theorems are complete in Lean. The checkpoint separates them from the remaining quantitative remarks and density-specific extensions, so we can resume those without redoing this work.

78. Assistant

2026-10-23 19:12:05 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T11:12:05.366Z

The main arithmetic results are now Lean-verified: Theorems 1.7–1.9, Lemma 5.1, and Corollary 1.10, without an assumed main theorem.

The fresh build passed: **41 modules, 566 declarations audited**, no unfinished proofs or additional axioms.

- Download the Lean project
- Detailed coverage
- Checkpoint for resuming

Theorem 4.1 remains excluded. Some ancillary claims—particularly Section 2's explicit quantitative estimates and certain density variants—remain unfinished.

Everything is saved and paused. Proof verification finished within your window; the final archive check overran the cutoff by about eight minutes. Sorry about that.

79. User

2026-10-24 01:44:26 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T17:44:26.497Z

please finish the lean verification, and also produce a text file that will help a human relate the lean certificate to the paper.

80. Assistant

2026-10-24 01:44:32 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T17:44:32.881Z

I'll resume from the saved project, check what remains outside the verified coverage, and work on those gaps. I'll also produce a plain-text guide mapping the paper's statements to the Lean theorems, with assumptions and exclusions made explicit.

81. Assistant

2026-10-24 01:47:31 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T17:47:31.352Z

The source audit found a few gaps behind yesterday's broad coverage summary: weighted polynomial matrices still need an inverse-closure proof, and two auxiliary results need their full paper statements exposed. The main theorem already has a checked proof; I'm completing these additional claims and the quantitative Bohr argument now.

82. Assistant

2026-10-24 01:51:58 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T17:51:58.673Z

The additional proofs for all minimizers in Lemma 3.3 and for nonconstant polynomial curves in Lemma 3.4 now compile. I'm also adding the inverse-closure proof from Lemma 3.1, so the certificate covers these full statements, not just the versions needed by the main theorem.

83. Assistant

2026-10-24 01:54:18 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T17:54:18.124Z

The density footnote's counterexample now compiles too, including the proof that its subadditivity does not supply the color-complement property used in the paper. The human guide will distinguish proved instances from results that still take a density property as an explicit hypothesis.

84. Assistant

2026-10-24 01:55:53 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T17:55:53.615Z

The numerical estimate for Section 2 is now checked: it converts the construction's factorial bound into the stated exponential-in-dimension form. The elementary coloring proof is being assembled around it; this is separate from the already verified general theorem.

85. Assistant

2026-10-24 02:00:04 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T18:00:04.756Z

The additive upper Banach density case now has a checked proof using its interval-based definition. The guide also explains the main translation choices—for example, the single block matrix used for a finite family and the proof that decoding recovers the original difference d, not just Ld.

86. Assistant

2026-10-24 02:02:01 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T18:02:01.086Z

The explicit pigeonhole bound and torus-geometry argument now compile. I'm also making the delivery reproducible: the verification script checks pinned dependency revisions, and the packaging script will check the archived files against the verified hashes.

87. Assistant

2026-10-24 02:04:21 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T18:04:21.328Z

The full elementary Section 2 coloring proof now compiles, including an explicit bound on progression length. I'm adding the open-neighborhood version and checking the radius endpoints before the final clean build. Richter's cited density has also passed its proof and axiom checks.

88. Assistant

2026-10-24 02:12:13 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T18:12:13.332Z

All the added mathematical modules now compile individually, including the strict Bohr-neighborhood conclusion and a bound valid for every positive radius. I'm checking the combined project next, then running a clean rebuild and packaging the proofs with the plain-text guide.

89. Assistant

2026-10-24 02:15:30 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T18:15:30.385Z

The combined project passed its first full axiom audit: 763 declarations, using only Lean's standard classical axioms. The clean rebuild is underway. I also checked the guide's 89 file-and-line references; all point to the intended declarations.

90. Assistant

2026-10-24 02:17:35 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T18:17:35.434Z

The clean rebuild has rechecked the main theorem and its public consequences successfully; it is finishing the added density modules. The guide now clearly separates the arithmetic certificate from the excluded nilmanifold theorem and external density-representation results.

91. Assistant

2026-10-24 02:18:37 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T18:18:37.007Z

The clean verification passed: 55 modules, 763 audited declarations, zero warnings, and no proof placeholders or additional mathematical axioms. The archive is ready, and its source files and guide match the verified hashes.

92. Assistant

2026-10-24 02:18:56 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-23T18:18:56.858Z

Completed the expanded arithmetic formalization: **55 modules and 763 declarations**, cleanly verified with zero warnings, no `sorry`, and no additional mathematical axioms.

- Lean source package
- Human-readable paper-to-Lean guide
- Verification record

The guide maps the paper's results to exact declarations, explains proof differences, and gives reproduction instructions.

Theorem 4.1 remains excluded as requested. External density-representation equivalences and the stronger random-coloring estimate cited from another paper also remain outside the certificate; these boundaries are documented explicitly.

93. User

2026-10-24 14:02:38 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-24T06:02:38.027Z

how long do you think it would take you to develope the relevant lean libraries for nilmanifolds so that you could also lean verify Theorem 4.1?

94. Assistant

2026-10-24 14:02:44 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-24T06:02:44.578Z

I'll check the exact scope of Theorem 4.1 and what the existing Lean libraries supply, so the estimate reflects the missing mathematics rather than just its short statement.

95. Assistant

2026-10-24 14:04:10 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-24T06:04:10.005Z

The main gap is the passage from an arbitrary nilpotent Lie group with a lattice to suitable coordinates—not merely defining the quotient space. Our existing matrix and rounding proofs give a substantial head start for concrete unitriangular examples, but they do not establish that general reduction.

96. Assistant

2026-10-24 14:06:20 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-24T06:06:20.161Z

For **Theorem 4.1 in its full stated generality, I would plan on months-scale work**, not another three-hour session. A very tentative budget is **300–1,500+ hours of focused development**. That is a scope-based estimate, not a dependable forecast of my runtime.

The distinction is:

Target	Tentative effort
Dependency audit and a nontrivial prototype	10–30 hours
Return-neighborhood theorem for standard unitriangular quotients	50–200 hours
Full Theorem 4.1 for arbitrary nilmanifolds	300–1,500+ hours

I checked our pinned mathlib. It already supplies substantial manifold, Lie-algebra, nilpotency, and quotient-topology infrastructure. What I did **not** find ready-made is the crucial chain:

general nilpotent Lie group and lattice → lattice-adapted Mal'cev coordinates → polynomial coordinate formulas → special-generalized-polynomial return conditions.

Our existing matrix and rounding proofs help substantially with the latter stages, but do not establish the general coordinate reduction. Arbitrary lattices and disconnected groups add complications. These matter because Theorem 4.1 is not restricted to standard matrix quotients.

There are worthwhile shortcuts: we need only the relevant return-neighborhood implication, not all of Huang–Shao–Ye's theory, and quotient topology may suffice without developing its full smooth structure.

My recommendation is a 10–30-hour pilot first. I would aim for a genuine concrete nilmanifold theorem, identify the precise remaining structural dependencies, and then revise the estimate. That would give a much firmer basis than extrapolating from how quickly the arithmetic formalization went.

97. User

2026-10-24 14:08:17 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-24T06:08:17.429Z

To what extent do you think that this task could be parallelized if I got the help of some other mathematicians who also have codex?

98. Assistant

2026-10-24 14:09:04 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-24T06:09:04.285Z

Quite substantially—but the structural Lie theory creates a dependency chain that limits the speedup. I'll distinguish work that can proceed independently from work that needs agreed definitions and earlier theorems.

---

99. Assistant

2026-10-24 14:11:10 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-10-24T06:11:10.416Z

Substantially. I would start with 3–5 active contributors, each using Codex, and provisionally expect around a 2–3× reduction in calendar time. That is a planning heuristic, not a measured forecast. Five contributors would not automatically make the project five times faster, and my earlier effort estimate is too uncertain to divide mechanically by the team size.

The best division would be by mathematical interfaces, rather than by sections of the paper:

Workstream	Responsibility	Main dependencies
Nilpotent Lie theory	Exponential/logarithm correspondence and polynomial group laws	Existing Lie-group infrastructure
Lattices and Mal'cev coordinates	Establish lattice-adapted coordinates for the relevant groups	Lie-theoretic results above
Generalized-polynomial calculus	Refine the required generalized-polynomial conditions into finite families of special generalized polynomials	Largely independent of the Lie theory
Coordinate-model dynamics	Quotient topology, orbit coordinates, and return-neighborhood theorem for a precisely specified coordinate model	Coordinate interface and polynomial calculus
Integration and mathematical audit	Arbitrary groups, disconnected cases, changes of basepoint, and assembly of Theorem 4.1	Eventually all the above

The key opportunity is that the fourth stream can work **assuming an explicitly specified coordinate structure**, while the first two prove that the relevant nilmanifolds possess it. This allows development in both directions without pretending that the structural theorem has already been proved.

The main limitation is the chain **Lie theory → lattice-adapted coordinates → general nilmanifold theorem**. Adding people cannot eliminate those dependencies. In particular, the disconnected-group reduction should not be treated as an automatic finishing step.

To make collaboration effective, I would insist on:

- **Shared Lean statements before separate proofs.** Everyone must agree on hypotheses, coordinate conventions, rounding, and the exact generalized-polynomial class.
- One pinned Lean/mathlib version, separate branches, and clear ownership of foundational definitions.

- Frequent integration builds and one designated integrator.
- Separate checks of mathematical statement fidelity and kernel correctness. A proved theorem with an extra structural assumption does not finish the original theorem.

A shared mathematical *blueprint*—statements, proofs, and their dependency graph—is well suited to this arrangement; it is used explicitly to organize collaboration in the FLT formalization project.

My strongest recommendation would be to recruit **a nilpotent Lie/Mal'cev theory expert and an experienced Lean formalizer**, alongside you. That combination could help more than several additional contributors who only run independent Codex sessions. We have already parallelized some proof coding; the greatest additional benefit would be better mathematical architecture and earlier detection of unsuitable formulations.