SYMMETRIC CHAIN DECOMPOSITIONS OF THE MINUSCULE LATTICE $M(n)$: TWENTY-TWO INSTANCES AND TWO EXACT COUNTS

KOREA SUPERINTELLIGENCE LABS

Abstract. Let $M(n)$ be the poset of strict integer partitions with largest part at most n, ordered by containment of Young diagrams. Stanley asked in 1980 whether $M(n)$ admits a symmetric chain decomposition; the question is open, and the most recent account, by Dorward, reports that little progress has been made. We exhibit an explicit symmetric chain decomposition of $M(n)$ for every $n \leq 21$; the largest splits the $2\,097\,152$ elements of $M(21)$ into $28\,460$ saturated chains, each symmetric about the middle of the rank range $0, \ldots, 231$. The decompositions are produced by a middle-out construction in which one outward step is a unit-capacity maximum flow. We also compute two counts for which only lower bounds were known: $\#\mathrm{SCD}(M(6)) = 1344$, so that the published bound ≥ 1344 is tight, and $\#\mathrm{SCD}(M(7)) = 295\,069\,056$, about 819 times the published bound $\geq 360\,036$. The second count is out of reach of exhaustive enumeration and is obtained from a dynamic program whose state is the bijection between two opposite ranks realised by the chains still under construction. All statements are machine-checked in Lean 4.

1. Introduction

A finite graded poset P of top rank $\rho(P)$ has a *symmetric chain decomposition* (SCD) if its elements can be partitioned into chains that are *saturated* — each element of a chain is covered by the next — and *symmetric*, meaning that the ranks of the bottom and the top of a chain sum to $\rho(P)$. Such a decomposition is the strongest convenient certificate of good behaviour a ranked poset can carry: it forces the rank numbers to be symmetric and unimodal and it forces the Sperner property, all at once and for elementary reasons. The Boolean lattice has one, by the classical argument of de Bruijn, van Ebbenhorst Tengbergen and Kruyswijk [1].

Stanley [5] asked whether the minuscule lattices $L(m, n)$ and $M(n)$ carry one. The first of these — partitions inside an $m \times n$ box — has been studied intensively. The second is the subject of this paper.

Definition 1.1. $M(n)$ is the set of *strict* integer partitions $\lambda = (\lambda_1 > \lambda_2 > \ldots)$ with $\lambda_1 \leq n$, ordered by containment of Young diagrams: $\lambda \leq \mu$ iff $\lambda_i \leq \mu_i$ for every i, absent parts being read as 0.

Question 1.2 (Stanley [5]; Question 1.1 of [3])**.** Does $M(n)$ admit a symmetric chain decomposition for every n?

Question 1.2 is open. The most recent account, R. Dorward [3], states that "not much progress has been made on SCDs for $M(n)$" and approaches the problem by counting: its Figure 4 tabulates the number $\#\mathrm{SCD}(M(n))$ of symmetric chain decompositions of $M(n)$, exactly for $n \leq 5$ and as lower bounds for $n = 6, 7$, the entries having been produced by a search (AlphaEvolve) run under a two-minute budget, after which the search "hit computational limits". For no n does [3] state that a symmetric chain decomposition was found, and no explicit decomposition of $M(n)$ appears in print for any n beyond the complete lists for $n = 3, 4, 5$ in its Appendix A.

This paper contributes instances and counts, not a general construction.

Results. Section 3 exhibits a symmetric chain decomposition of $M(n)$ for every $n \leq 21$ (Theorems 3.1 and 3.2), together with the construction that produced them: chains are grown outward from the middle rank, and one outward step is exactly a unit-capacity maximum flow whose middle edge forces one chain to receive a new bottom and a new top simultaneously. The construction never failed, at any n it was run on.

Section 4 settles the two entries of Figure 4 of [3] that were left as bounds:

n	1	2	3	4	5	6	7
[3], Figure 4	1	1	2	2	12	≥ 1344	≥ 360036
here	1	1	2	2	12	**1344**	**295069056**

So the lower bound at $n = 6$ is tight (Theorem 4.2, Corollary 4.3), while the lower bound at $n = 7$ is short of the truth by a factor of about 819 (Theorem 4.5, Corollary 4.6). The count at $n = 7$ is beyond enumeration — there are almost three hundred million objects to list — and comes instead from a dynamic program whose state is the bijection between two opposite ranks realised by the chains still under construction (Lemma 4.4).

Section 5 records the checks that tie the formal model to the source's, and Section 6 says what the machine verification covers and what it does not.

We emphasise what is *not* proved. Twenty-two instances are twenty-two instances: Question 1.2 asks for all n and is untouched by anything below. The evidence assembled here does suggest where a proof will not come from: as in the search of [3], no local rule for building the chains emerged, and the construction of Section 3 is global at every step.

2. The poset $M(n)$

A strict partition with parts in $\{1, \dots, n\}$ is the same thing as a subset of $\{1, \dots, n\}$, so

$$|M(n)| = 2^n.$$

The poset is graded by $\rho(\lambda) = |\lambda| = \sum_i \lambda_i$, with unique minimum $\varnothing$ and unique maximum $\{1, 2, \dots, n\}$, so that

$$\rho(M(n)) = \binom{n+1}{2} = \frac{n(n+1)}{2},$$

and the number of elements of rank r is the coefficient of q^r in

$$\prod_{i=1}^{n} (1 + q^i).$$

In particular $M(n)$ is rank-symmetric. Its covers are easy to name: μ covers λ iff μ is obtained from λ either by adjoining the part 1, or by replacing a part v by $v+1$ when $v + 1 \leq n$ is not already a part. We write R_a for the set of elements of rank a and $\rho = \rho(M(n))$ when n is fixed.

Definition 2.1. A chain $C = (c_0 < c_1 < \cdots < c_k)$ in $M(n)$ is *saturated* if c_{j+1} covers c_j for every j, and *symmetric* if $\rho(c_0) + \rho(c_k) = \rho$. A *symmetric chain decomposition* of $M(n)$ is a set of saturated symmetric chains partitioning $M(n)$. We write $\#\mathrm{SCD}(M(n))$ for the number of them.

A symmetric chain spans the rank interval $[a, \rho - a]$ for some $a \leq \rho/2$, so it meets every rank between $\lfloor \rho/2 \rfloor$ and $\lceil \rho/2 \rceil$ exactly once. Hence:

Lemma 2.2. *Every symmetric chain decomposition of $M(n)$ has exactly $|R_{\lfloor \rho/2 \rfloor}|$ chains.*

The chain counts of the decompositions below are therefore forced, and they are listed in Table 1 as a description of the objects rather than as a claim about them.

Finally, complementing the part set inside $\{1, \dots, n\}$ is an order-reversing involution of $M(n)$ sending rank r to rank $\rho - r$; it is the map η_s of [3, §4], coming from the Lusztig involution. It carries symmetric chain decompositions to symmetric chain decompositions, which is how [3, Cor. 4.8] concludes that $\#\mathrm{SCD}(M(n))$ is even for $n > 2$.

3. Symmetric chain decompositions for $n \leq 21$

Theorem 3.1. *$M(19)$, $M(20)$ and $M(21)$ admit symmetric chain decompositions.*

Theorem 3.2. *$M(n)$ admits a symmetric chain decomposition for every $n \leq 21$.*

Both are proved by exhibiting the decomposition and checking it. A candidate is a finite list of finite lists of elements, and the three conditions of Definition 2.1 — every listed chain is a run of covers, every listed chain has $\rho(\text{bottom}) + \rho(\text{top}) = \rho$, and every element of $M(n)$ occurs in exactly one chain — are decided directly from the definitions of the order and the rank, for the particular n at hand. The decompositions themselves are given in the accompanying formal development; they are far too large to print, $M(21)$ alone splitting 2 097 152 elements into 28 460 chains. Table 1 records their shape.

n	1	2	3	4	5	6	7	8	9	10	11
chains	1	1	2	2	3	5	8	14	23	40	70

n	12	13	14	15	16	17	18	19	20	21
chains	124	221	397	722	1314	2410	4441	8220	15272	28460

Table 1. Number of chains in the decomposition constructed for each n, equal by Lemma 2.2 to the middle rank size of $M(n)$.

3.1. **The construction.** The chains are grown outward from the middle rank, which is the shape [3] attributes to the network-flow attempts of its own search. Suppose the chains have been built on the rank interval $[a+1, \rho-a-1]$, so that the *alive* chains — those whose bottom lies in R_{a+1} — are in bijection with R_{a+1}, and by rank symmetry also with $R_{\rho-a-1}$. The next step must extend some subset of the alive chains down onto all of R_a and up onto all of $R_{\rho-a}$, and it must extend the *same* chains in both directions: a chain that gains a new bottom without a new top is no longer symmetric. This is exactly a maximum flow in the unit-capacity network

$$s \longrightarrow R_a \longrightarrow \mathrm{in}(C) \longrightarrow \mathrm{out}(C) \longrightarrow R_{\rho-a} \longrightarrow t,$$

in which each alive chain C appears as a pair of vertices joined by one edge of capacity 1; an element $x \in R_a$ is joined to $\mathrm{in}(C)$ when the bottom of C covers x, and $\mathrm{out}(C)$ is joined to $y \in R_{\rho-a}$ when y covers the top of C. The step succeeds iff the maximum flow has value $|R_a| = |R_{\rho-a}|$, and the single edge $\mathrm{in}(C) \to \mathrm{out}(C)$ is what couples the two directions. Chains carrying no flow terminate; they already span a symmetric interval, so they are complete.

This is a heuristic and not a proof: a maximum flow at one level does not guarantee that later levels remain feasible, so success produces a witness and failure would produce nothing. In practice it never failed. Every n from 1 to 21 succeeded on the first random edge ordering, with no backtracking, the cost being essentially the cost of the flow computations: at most 6.5 seconds of single-core time for $n \leq 16$, then 64, 244, 556, 1090 and 1759 seconds at $n = 17, \ldots, 21$, a growth of roughly 1.6 to 2.3 per step. The search was stopped at $n = 21$ for reasons of budget, not of feasibility; a run at $n = 22$ was abandoned after 500 seconds without having failed.

Remark 3.3. For $n \leq 18$ the objects of Theorem 3.2, though nowhere claimed in print, are within reach of published software: the deterministic script `Mn_single_scd.py` in the code repository [2] accompanying [3] returns a valid symmetric chain decomposition of $M(n)$ for every $n \leq 18$, as we confirmed by re-running it. At $n = 19$ and $n = 20$ its search times out and it falls through to a fallback whose chains are not symmetric; $n = 21$ we did not test there. This is why Theorem 3.1 is stated separately: $n = 19, 20, 21$ are the values for which we have found no procedure, published or otherwise, that produces the object. For $n \leq 18$ what is new is only the formal statement and the machine check.

Remark 3.4. Writing a decomposition of $M(n)$ down at all becomes an issue around $n = 15$: an explicit list literal carries 2^n numerals. The large decompositions are therefore stored in a compact form — per chain, its length, its bottom element, and then one symbol per step naming which of the two kinds of cover was used. Nothing rests on the correctness of this encoding. It is a way of writing a list of lists down, and the three conditions of Definition 2.1 are re-derived from scratch on the decoded object; an error in the encoding could only produce a list that fails the check. It is nevertheless verified against the explicit lists where both exist (Proposition 5.1).

4. Two exact counts

4.1. **Enumeration.** The symmetric chain decompositions of $M(n)$ can be listed without repetition by the following scheme.

Lemma 4.1. *Fix a linear order on $M(n)$ refining the rank. Given a partial decomposition covering a down-set of that order, let x be the first uncovered element. Then in any completion, x is the bottom of its chain, and that chain ends at rank $\rho - \rho(x)$.*

Proof. Every element of rank smaller than $\rho(x)$ is already covered, so the chain through x cannot contain an element below x; hence x is its bottom, and symmetry fixes the rank of its top at $\rho - \rho(x)$. □

Branching over all saturated chains from x to rank $\rho - \rho(x)$ through still-free elements therefore visits every symmetric chain decomposition exactly once, since the decomposition determines the branch taken at every step. Write $N(n)$ for the number of leaves of this search tree, so $N(n) = \#\mathrm{SCD}(M(n))$ by Lemma 4.1.

Theorem 4.2. $\#\mathrm{SCD}(M(6)) = 1344$.

Corollary 4.3. $\#\mathrm{SCD}(M(6)) \geq 1345$ *is false: the lower bound of* [3, Fig. 4] *is tight.*

Theorem 4.2 is the evaluation of $N(6)$, carried out exhaustively: the search above, run to completion over the 64 elements of $M(6)$, has 1344 leaves. The computation was performed inside the proof assistant and takes about three minutes there; it was reproduced by two independently written programs outside it. The enumerator is validated against the source's own exact values and against an independent checker in Proposition 5.1.

4.2. **Counting without enumerating.** At $n = 7$ the same search has 295 069 056 leaves and is hopeless. What makes the count tractable anyway is that the outward sweep of Section 3.1 carries very little information from one rank to the next.

Lemma 4.4. *Let a partial decomposition be built on the rank interval $[a+1, \rho-a-1]$, all of whose chains are saturated and span $[b, \rho-b]$ for some $b \geq a+1$. Whether it can be completed, and in how many ways, depends only on the bijection $R_{a+1} \to R_{\rho-a-1}$ that sends the bottom of each alive chain to its top.*

Proof. A chain that has terminated spans a symmetric interval and is complete; no later step touches it, and the only trace it leaves is that its bottom and top are no longer available, which is already recorded by their absence from the domain and range of the bijection. What the remaining steps may do to an alive chain depends on that chain only through its current bottom and its current top, and the alive chains are otherwise interchangeable. □

Sweeping the ranks outward while carrying a table of states weighted by the number of partial decompositions realising them therefore computes $\#\mathrm{SCD}(M(n))$, one addition per transition. At $n = 7$ the sweep meets 2339 distinct states in total.

Theorem 4.5. $\#\mathrm{SCD}(M(7)) = 295\,069\,056$.

Corollary 4.6. $\#\mathrm{SCD}(M(7)) > 360\,036$*: the lower bound of* [3, Fig. 4] *at* $n = 7$ *is short of the true value by a factor of about* 819*.*

The value is the output of the sweep of Lemma 4.4, computed inside the proof assistant in about 25 seconds. Since the sweep is a different instrument from the enumerator, and since it is Lemma 4.4 rather than a formal proof that identifies its output with $\#\mathrm{SCD}(M(7))$, it is pinned to the enumerator wherever the enumerator can still be run: the two agree for every $n \leq 5$ — the whole range of exact values published in [3, Fig. 4] — and both give 1344 at $n = 6$. Two further programs, written independently and outside the proof assistant, return 295 069 056 as well. As a further consistency check the value is even, as [3, Cor. 4.8] requires.

Remark 4.7. The state of Lemma 4.4 is a bijection between two opposite ranks, and its usefulness runs out quickly: at $n = 8$ the reachable states exhaust more than 12 gigabytes before the sweep finishes, so $\#\mathrm{SCD}(M(8))$ remains unknown. The evident next move is to quotient the state space by the involution η_s of Section 2, which acts on partial decompositions as well as on complete ones and should roughly halve the count of states. We record for the interested reader that the sequence $1, 1, 2, 2, 12, 1344, 295069056$ is, as of this writing, absent from the OEIS [4], as are its initial segments beyond $1, 1, 2, 2, 12$.

5. Validating the model

Formal statements about a poset are only as good as the encoding of the poset. The following checks were run to confirm that the object being decomposed and counted is the one of [3].

Proposition 5.1. (1) *The enumerator of Lemma 4.1 returns* $1, 1, 2, 2, 12$ *for* $n = 1, \ldots, 5$*, which is exactly the published table* [3, Fig. 4]*.*
(2) *For every* $n \leq 5$*, every object the enumerator produces satisfies the independent definition of a symmetric chain decomposition, and the number of objects produced equals the count returned.*
(3) *For* $n = 8, \ldots, 12$ *the compact encoding of Remark 3.4 decodes to exactly the explicit list of chains.*

Part (1) is a comparison against the source; note that the enumerator knows nothing about the decompositions of Section 3 and reaches the published values on its own. Part (2) separates the two roles a definition can play: the search that produces objects and the predicate that accepts them are written independently, and the predicate accepts everything the search produces. The objects produced are also pairwise distinct, so the count counts distinct decompositions.

Alongside these, the encoding of $M(n)$ itself was checked against its defining properties: the top element has rank $\binom{n+1}{2}$ for $n \leq 20$; the rank sizes sum to 2^n and read the same backwards; the rank sizes of $M(6)$ agree with the coefficients of $\prod_{i=1}^{6}(1 + q^i)$; distinct part sets give distinct Young diagrams, so the order is a genuine partial order on the encoding; the shortcut list of upper covers used by the enumerator agrees with covers computed from containment and rank; and complementation of the part set is an order-reversing, rank-complementing involution for $n \leq 8$, as η_s must be.

The definition of a symmetric chain decomposition was also probed from the other side, since a predicate that accepts too much proves nothing. Each of its conditions was shown to be live and independent: a list with a rank gap, a chain listed downwards, a chain that is saturated but not symmetric, a chain list with a repeated element, one that misses an element, and one whose entries leave the ground set are all rejected; so are five surgically damaged versions of the decomposition of $M(8)$ (top removed, one chain reversed, one duplicated, one deleted, one interior element removed), and so is the decomposition of $M(8)$ offered as a decomposition of $M(9)$, and conversely.

6. Verification

Every statement in this paper has been formally verified in Lean 4 (toolchain `v4.32.0`); the development imports no library beyond the Lean core, every definition being executable, and it is free of `sorry`. The order-theoretic content — that containment of Young diagrams, in the encoding used, is reflexive, transitive and antisymmetric, so that a list of consecutive covers really is a chain — is checked by the kernel with no axioms beyond propositional extensionality. What is delegated to compiled evaluation, through Lean's `native_decide`, is exactly the finite computations: the single evaluation that checks all twenty-two decompositions of Theorems 3.1 and 3.2 against the definition (well under two minutes for all of them together, of which about 25 seconds is $n = 19$); the exhaustive enumeration behind Theorem 4.2; the outward sweep behind Theorem 4.5; and each item of Proposition 5.1 and each of the checks listed after it. Corollaries 4.3 and 4.6 are arithmetic consequences of the corresponding theorems.

Two limits of this should be stated plainly. First, the identification of the enumerator's output with $\#\mathrm{SCD}(M(n))$ rests on Lemma 4.1, and the identification of the sweep's output with $\#\mathrm{SCD}(M(n))$ rests on Lemma 4.4; those lemmas are proved here as mathematics, in the ordinary way, and are not themselves formalised. What the machine checks is that the instruments return the stated numbers, that they agree with each other wherever both can be run, and that they agree with the published exact values. Second, the finite computations are run by compiled code rather than by the kernel, which is a stronger trust assumption than the rest; the decomposition checks of Section 3 are the ones where this matters least, since re-checking a given decomposition is a linear scan that any independent implementation can repeat, and the counts of Section 4 were reproduced by independently written programs.

The formal development accompanies this paper; repository reference to be added at submission.

References

[1] N. G. de Bruijn, C. A. van Ebbenhorst Tengbergen and D. Kruyswijk, *On the set of divisors of a number*, Nieuw Arch. Wiskunde (2) **23** (1951), 191–193.

[2] *Symmetric Chain Decompositions of Minuscule Lattices* — the code repository accompanying [3] and cited there as `https://github.com/google-research/symmetric-chain-decomposition`; script `Mn_single_scd.py`. Commit `eb80f808` (18 May 2026), head of the default branch; consulted August 2026.

[3] R. Dorward, *On the number of symmetric chain decompositions of the minuscule lattices $L(m,n)$ and $M(n)$*, arXiv:2606.15891v1, 14 June 2026. Section, figure and appendix numbers above refer to v1, which as of August 2026 is the only version posted; no journal version has appeared.

[4] OEIS Foundation Inc., *The On-Line Encyclopedia of Integer Sequences*, `https://oeis.org`.

[5] R. P. Stanley, *Weyl groups, the hard Lefschetz theorem, and the Sperner property*, SIAM J. Algebraic Discrete Methods **1** (1980), no. 2, 168–184.