The queueing polynomial is convex

P. Graham
https://orcid.org/0009-0007-8493-1505

October 31, 2026

Abstract

We prove that the hypergraph polynomial arising in redundancy scheduling is convex on the nonnegative orthant for every number of vertices, every positive edge size, and every nonnegative degree. This proves Conjecture 1 of Brosch, Laurent and Steenkamp, posed in their work on job occupancy, which asks for convexity on the standard simplex for all parameters at least two. They proved convexity for edge size two and degree three; Polak's Theorem 1.1 extends the edge-size-two case through degree nine. Marshall and Olkin's classical exponential construction gives a moment representation that yields convexity directly. We apply this construction to the conjecture and give a finite proof checked in Lean. A vertex-deletion recurrence expresses the original ordered-tuple polynomial as an average of complete homogeneous symmetric polynomials at nonnegative linear forms. Their convexity follows from McLeod's classical inequality; we give a coefficientwise Hessian proof suited to formal verification. The representation and the full simplex statement are checked in Lean.

This work has not yet been reviewed by experts. The proofs, the Lean formalization and the text were produced by AI models (from OpenAI and Anthropic), working under the author's direction.

Keywords: convex polynomial, redundancy scheduling, complete homogeneous symmetric polynomial, formal proof.

1 Introduction

Let $[n] = \{1, \dots, n\}$ and let $E = \binom{[n]}{L}$ be the set of all L-element subsets. An edge describes a set of servers to which a job may be assigned. For $n = L = 2$, there is a single edge e, and the degree-d polynomial below is $(x_e/2)^d$. In general, for $L \geq 1$ define

$$f_d(x) = \sum_{(e_1,\dots,e_d)\in E^d} \prod_{i=1}^{d} \frac{x_{e_i}}{|e_1 \cup \cdots \cup e_i|}, \qquad x \in \mathbb{R}^E, \quad d \in \mathbb{N}. \tag{1}$$

Tuples are ordered and may repeat edges. Every denominator includes the current edge and is positive. The empty product convention gives $f_0 = 1$. Write $\mathbb{R}^E_+ = \{x : x_e \geq 0 \text{ for every } e\}$ and $\Delta_E = \{x \in \mathbb{R}^E_+ : \sum_{e\in E} x_e = 1\}$.

Theorem 1. *For every $n, d \in \mathbb{N}$ and every integer $L \geq 1$, the polynomial f_d is convex on $\mathbb{R}^E_+$. In particular, for all integers $n, L, d \geq 2$, it is convex on Δ_E.*

The second assertion is exactly Conjecture 1 of Brosch, Laurent and Steenkamp [1, Section 4, p. 2242]. Their paper first appeared in 2020. They prove convexity for $L = 2, d = 3$ in Section 4.2,

Proposition 4.5 and the paragraph preceding it; this gives their Theorem 1.2 on the minimizer. Their related polynomial [1, equation (1.2)] is

$$p_d(x) = \sum_{(e_1,\dots,e_d)\in E^d} \frac{x_{e_1}\cdots x_{e_d}}{|e_1\cup\cdots\cup e_d|} \qquad (d\ge 2).$$

Its convexity is established in the proof of Theorem 1.1, concluding after Corollary 3.8. Since $|e_1| = L$, we have $f_2 = p_2/L$, which gives the degree-two case as noted after their Conjecture 1. Polak [2, Theorem 1.1] proves convexity for $L = 2$ and $d \le 9$. Theorem 1 removes both the edge-size and degree restrictions and enlarges the domain to the nonnegative orthant. If $L > n$, then E is empty and Δ_E is empty, so the simplex assertion includes this degenerate case.

The coefficients in (1) belong to an established product-form queueing framework; see Gardner and Righter [3, Theorem 3.3]. The moment representation from Marshall and Olkin's classical shared-exponential-minimum construction [4, pp. 14–18, especially (4.2)–(4.6)] gives convexity directly; the remark after Lemma 3 states the identity and its precise relation to the finite average. Our contribution is its application to BLS Conjecture 1 with a self-contained finite proof checked in Lean. The convexity ingredient is also classical: McLeod's inequality [5, Section 1, inequality (iii)] makes $h_d^{1/d}$ subadditive, and hence convex by homogeneity; composing with the increasing convex power function gives convexity of h_d on the orthant. The proof below supplies a direct finite argument for both ingredients, so that the Lean theorem starts with the literal sum (1).

2 A finite representation

The idea is to delete the vertices in a uniformly random order. At each stage, sum the coordinates of the edges still present and divide by the number of surviving vertices. The resulting list consists of nonnegative linear functions of x. We will show that f_d is the average of the degree-d complete homogeneous polynomial of these lists. Convexity then passes through linear substitution and finite averaging.

For $S \subseteq [n]$, let $q_d(S;x)$ denote the sum in (1) restricted to tuples whose edges all lie in S. All coordinates remain indexed by the original E. Define

$$r(S;x) = \sum_{\substack{e\in E\\ e\subseteq S}} x_e, \qquad a(S;x) = \frac{r(S;x)}{|S|} \quad (S\neq\varnothing).$$

In particular, $q_d([n];x) = f_d(x)$, $q_0(S;x) = 1$, and $q_d(\varnothing;x) = 0$ for $d > 0$.

Lemma 2. *For $L \ge 1$, every $S \subseteq [n]$, every $d \ge 1$, and every $x \in \mathbb{R}^E$,*

$$|S|q_d(S;x) = r(S;x)q_{d-1}(S;x) + \sum_{v\in S} q_d(S\setminus\{v\};x). \tag{2}$$

Proof. For a tuple $s = (e_1,\dots,e_d)$ supported in S, let $U(s)$ be its union and $w(s)$ its summand in (1). The last denominator of $w(s)$ is $|U(s)|$. Canceling it and summing over the last edge gives

$$\sum_{s\in\binom{S}{L}^d} |U(s)|w(s) = r(S;x)q_{d-1}(S;x).$$

The tuple occurs in $q_d(S\setminus\{v\};x)$ precisely when $v\notin U(s)$, so its multiplicity in the deletion sum is $|S| - |U(s)|$. Adding these two identities proves (2). The cancellation is valid because $d, L \ge 1$; empty sums cause no exception. □

For a list $y = (y_1, \dots, y_m)$, write

$$h_d(y) = \sum_{\substack{\alpha \in \mathbb{N}^m \\ |\alpha| = d}} y^\alpha, \qquad |\alpha| = \sum_i \alpha_i, \qquad y^\alpha = \prod_i y_i^{\alpha_i}.$$

Thus $h_2(u, v) = u^2 + uv + v^2$. For the empty list, $h_0 = 1$ and $h_d = 0$ for $d > 0$. Separating the exponent of the first variable gives

$$h_d(a, y) = \sum_{k=0}^{d} a^k h_{d-k}(y). \tag{3}$$

Choose an ordering of the vertices of S uniformly, and let $S_m = S, S_{m-1}, \dots, S_1$ be the successive nonempty surviving sets, where $m = |S|$. An expectation over such orderings below means a finite average, including the single empty ordering if $S = \varnothing$.

Lemma 3. *For $L \geq 1$, every $S \subseteq [n]$ with $m = |S|$, every $d \in \mathbb{N}$, and every $x \in \mathbb{R}^E$,*

$$q_d(S; x) = \mathbb{E}\, h_d(a(S_m; x), \dots, a(S_1; x)). \tag{4}$$

Proof. Induct on $|S|$. The empty set follows from the conventions above. For nonempty S, put $b_j = |S|^{-1} \sum_{v \in S} q_j(S \setminus \{v\}; x)$ and $a = a(S; x)$. Lemma 2 gives $q_j(S; x) = a q_{j-1}(S; x) + b_j$ for $j > 0$, while $q_0(S; x) = b_0 = 1$. Iterating this scalar recurrence yields

$$q_d(S; x) = \frac{1}{|S|} \sum_{v \in S} \sum_{k=0}^{d} a(S; x)^k q_{d-k}(S \setminus \{v\}; x).$$

Apply the induction hypothesis to each smaller set and use (3). Choosing the first deleted vertex uniformly and then a uniform ordering of the remaining vertices gives a uniform ordering of S, exactly as in (4). □

Relation to the classical exponential construction. Let $(T_v)_{v \in [n]}$ be independent exponential random variables of rate one, and put $Y_e = \min_{v \in e} T_v$. This is the shared-minimum construction in Marshall–Olkin [4, p. 18; see also (4.2)–(4.6)], specialized to vertex shocks. For every $x \in \mathbb{R}^E$ and $d \in \mathbb{N}$ it gives

$$f_d(x) = \frac{1}{d!}\, \mathbb{E}\left[\left(\sum_{e \in E} x_e Y_e \right)^d \right]. \tag{5}$$

Indeed, conditional on the order in which the vertices die, the successive waiting times Δ_k with k vertices surviving are independent exponentials of rates k. Hence $\sum_e x_e Y_e = \sum_{k=1}^{n} r(S_k; x) \Delta_k$. The multinomial expansion and $\mathbb{E}\Delta_k^j = j!/k^j$ show that its conditional dth moment divided by $d!$ equals $h_d(a(S_n; x), \dots, a(S_1; x))$. The death order is uniform, so averaging gives (5) by Lemma 3; the same argument on S gives that lemma's formula. For $x \geq 0$, each realization of $\sum_e x_e Y_e$ is a nonnegative linear form, whose dth power is convex (constant when $d = 0$). Thus the classical construction also proves the orthant convexity directly. The finite argument above and below avoids integration in Lean.

3 Convexity

Here is a direct proof of the classical convexity fact in the form used by the formalization. It avoids introducing probability measures or integrals into the argument.

Lemma 4. *For every $m, d \in \mathbb{N}$, the function h_d is convex on $\mathbb{R}^m_+$.*

Proof. Degrees zero and one are immediate. For $d \geq 2$, fix a direction $z \in \mathbb{R}^m$. The coefficient of y^α, $|\alpha| = d-2$, in the second directional derivative $D_z^2 h_d(y)$ is

$$\sum_i (\alpha_i+2)(\alpha_i+1)z_i^2 + 2\sum_{i<j}(\alpha_i+1)(\alpha_j+1)z_i z_j$$
$$= \left(\sum_i (\alpha_i+1)z_i\right)^2 + \sum_i (\alpha_i+1)z_i^2 \;\geq\; 0.$$

All other coefficients vanish. Thus $D_z^2 h_d(y) \geq 0$ whenever $y \geq 0$. On every line segment in the orthant, the polynomial consequently has nonnegative second derivative and satisfies the convexity inequality. This includes the boundary and the zero-dimensional orthant. □

Proof of Theorem 1. For each fixed deletion ordering, the map $x \mapsto (a(S_n; x), \ldots, a(S_1; x))$ is linear and sends $\mathbb{R}^E_+$ into $\mathbb{R}^n_+$. Lemma 4 makes the corresponding summand of (4) convex on $\mathbb{R}^E_+$. A finite average with nonnegative weights is convex, so Lemma 3 with $S = [n]$ proves the orthant assertion. The simplex is a convex subset of that orthant; restriction gives the second assertion. □

4 Formal verification

The Lean theorem `BLSFinal.source_orthant` states, with natural parameters n, L, d and the hypothesis $1 \leq L$,

$$\text{ConvexOn } \mathbb{R}\ \{x : \texttt{Edge}\ n\ L \to \mathbb{R} \mid \forall e,\ 0 \leq x(e)\}\ (\texttt{sourcePolynomial}\ n\ L\ d).$$

The theorem `BLSFinal.fullEndpoint` states, for every $n, L, d \in \mathbb{N}$, that $2 \leq n$, $2 \leq L$, and $2 \leq d$ imply

$$\text{ConvexOn } \mathbb{R}\ (\texttt{stdSimplex}\ \mathbb{R}\ (\texttt{Edge}\ n\ L))\ (\texttt{sourcePolynomial}\ n\ L\ d).$$

Here `Edge` consists of all L-element subsets of `Fin n`; this replaces the vertex labels $1, \ldots, n$ by $0, \ldots, n-1$. The definition `sourcePolynomial` is exactly (1), with `Fin d`-indexed ordered tuples and inclusive prefix unions. Mathlib's `stdSimplex` imposes coordinate nonnegativity and sum one, and `ConvexOn` is the usual two-point convexity inequality on a convex domain. No representation or convexity hypothesis remains in either exported theorem.

The kernel checks the deletion recurrence, finite representation, Hessian argument, and both endpoint theorems. Their axiom dependencies are `propext`, `Classical.choice`, and `Quot.sound`. Statement checks and small examples are part of the default Lake build: degree zero, the two-variable quadratic h_2, repeated-edge prefixes, and $n = 2, L = 3$ with empty simplex. The acceptance check was same-family: its AI model family had also contributed to the proof.

The accompanying standalone Lean project pins the compiler and Mathlib. Fetching the Mathlib cache and invoking the standard Lake build reproduces the checks; the exact commands and requirements are in the accompanying README. The comparison to the published conjecture and the historical attributions are mathematical reading, rather than kernel assertions.

References

[1] D. Brosch, M. Laurent and A. Steenkamp, Optimizing hypergraph-based polynomials modeling job-occupancy in queuing with redundancy scheduling, *SIAM J. Optim.* **31** (2021), 2227–2254. Published version. doi:10.1137/20M1369592.

[2] S. C. Polak, Symmetry reduction to optimize a graph-based polynomial from queueing theory, *SIAM J. Appl. Algebra Geom.* **6** (2022), 243–266. doi:10.1137/21M1413298; Theorem 1.1 also in arXiv:2104.08264v2.

[3] K. Gardner and R. Righter, Product forms for FCFS queueing models with arbitrary server-job compatibilities: an overview, *Queueing Systems* **96** (2020), 3–51. doi:10.1007/s11134-020-09668-6; Theorem 3.3 in arXiv:2006.05979v1.

[4] A. W. Marshall and I. Olkin, A multivariate exponential distribution, *J. Amer. Statist. Assoc.* **62** (1967), 30–44. doi:10.1080/01621459.1967.10482885. Page and equation references are to the Stanford technical report, 1966.

[5] J. B. McLeod, On four inequalities in symmetric functions, *Proc. Edinburgh Math. Soc.* **11** (1959), 211–219. doi:10.1017/S0013091500021921.