Starflower Inheritance: Exponential Circuit Lower Bounds Robust to Bounded Negation Width

Author: Richard Brazil, Brazil Innovations
Date: March 24, 2026
Classification: Theoretical Computer Science; Circuit Complexity; Boolean Functions
Status: Preprint

Abstract

We establish the **Starflower Inheritance Theorem**, proving that exponential monotone circuit lower bounds persist under bounded negation width—a significant extension of robustness for computational hardness. For any function family with monotone complexity $C^+(f_N) \geq 2^{N^{\alpha - o(1)}}$, we prove that circuits with negation width $w \leq N^{\alpha-\varepsilon}/\log N$ must retain exponential size $2^{\Omega(N^{\alpha-\varepsilon})}$. This introduces the **Brazil Threshold**, a universal formula connecting monotone hardness exponents to negation-width robustness regimes.

We apply this theorem to five explicit function families using recent breakthrough lower bounds: (1) bipartite perfect matching (Çalar et al. 2025: $C^+(\mathrm{Match}_n) \geq 2^{n^{1/3-o(1)}}$), (2) GEN-TFNP search problems, (3) pigeonhole/clique-coloring functions, (4) Tardos weight functions, and (5) de Rezende–Vinyals P-functions from CSP lifting. This establishes a **negation ladder** of inherited hardness, with matching achieving the strongest known lower bounds in the bounded-negation regime: circuits require size $2^{\Omega(n^{1/3-\varepsilon})}$ for negation width up to $n^{1/3-\varepsilon}/\log n$—a super-polynomial improvement over prior results.

Our framework unifies monotone and negation-limited complexity, demonstrates that exponential monotone hardness carries genuine structural information robust to substantial perturbations, and provides stronger input bounds potentially enabling hardness magnification toward P versus NP resolution. The theorem is entirely rigorous; we clearly separate proven results from exploratory conjectures about potential applications to computational complexity theory and related domains.

Keywords: circuit complexity, monotone lower bounds, negation width, Boolean functions, hardness inheritance, bipartite matching, TFNP, P versus NP, computational complexity

---

1. Introduction and Main Results

1.1 Motivation and Context

Proving strong lower bounds for general Boolean circuits remains one of the central challenges in computational complexity theory[1]. The gap between monotone circuits (using only AND and OR gates) and general DeMorgan circuits (permitting NOT gates) has historically been vast: certain monotone functions can require exponential monotone complexity yet admit small general circuits due to the power of negation[1].

This paper studies the **interpolation** between monotone and unrestricted circuits via **negation width** (w)—the maximum number of distinct negated variables appearing in any syntactically produced term[3]. A circuit with negation width $w = 0$ is purely monotone; $w = N$ corresponds to unrestricted DeMorgan circuits. Our central question: **How much negation can a circuit employ before exponential monotone hardness collapses?**

We prove that exponential monotone lower bounds are remarkably robust: they survive as long as negation width grows sub-polynomially relative to the hardness exponent.

1.2 Main Contributions

1. **Starflower Inheritance Theorem (Theorem 3.1):** A general theorem establishing that exponential monotone lower bounds $C^+(f_N) \geq 2^{N^{\alpha-o(1)}}$ are inherited by circuits with negation width $w \leq N^{\alpha-\varepsilon}/\log N$, yielding exponential size lower bounds $2^{\Omega(N^{\alpha-\varepsilon})}$.
2. **Brazil Threshold (Definition 3.2):** A universal formula $w_{\text{Brazil}}(N, \alpha, \varepsilon) = N^{\alpha-\varepsilon}/\log N$ that quantifies the negation-width robustness regime for any monotone hardness exponent α.
3. **Negation Ladder (Section 4):** Application to five explicit function families with varying hardness exponents ($\alpha \in \{1/3, 1/4, 1/8\}$), establishing the strongest known lower bounds in bounded-negation regimes. For bipartite matching, we achieve $2^{\Omega(n^{1/3-\varepsilon})}$ size lower bounds for circuits with negation width up to $n^{1/3-\varepsilon}/\log n$.
4. **Unified Framework (Section 5):** Demonstrates that monotone lower bounds carry genuine structural hardness information robust to substantial perturbations, narrowing the perceived gap between monotone and non-monotone circuit models in terms of complexity preservation.

1.3 Significance and Implications

Theoretical impact: The Starflower framework shows that exponential monotone complexity is not an artifact of maximal restriction—it reflects intrinsic hardness that already constrains a broad class of negation-limited circuits. This suggests the monotone-versus-general gap may be narrower than previously thought in terms of structural hardness preservation[1][3].

Toward hardness magnification: Recent results demonstrate that sufficiently strong circuit lower bounds on problems like the Minimum Circuit Size Problem (MCSP) can be "magnified" into complexity class separations such as $\text{NP} \not\subseteq \text{P}/\text{poly}$[9]. Our theorem provides stronger input lower bounds—particularly for matching with negation width constraints—that may enable future magnification frameworks leveraging bounded-negation reductions.

Quantitative improvement: For bipartite matching, prior negation-width lower bounds allowed approximately $w \approx \log n$ negations[1]. Using the recent Çalar et al. breakthrough bound[2], we extend the robustness regime to $w \approx n^{1/3}/\log n$—a super-polynomial improvement demonstrating that exponential hardness persists under substantial negation.

1.4 Related Work

Monotone circuit complexity: Razborov's celebrated 1985 result established super-polynomial monotone lower bounds for bipartite perfect matching and clique using the approximation method[1]. Recent work by Çalar, Göös, Riazanov, Sofronova, and Sokolov achieves exponential monotone matching bounds: $C^+(\mathrm{Match}_n) \geq 2^{n^{1/3-o(1)}}$[2]—a breakthrough improvement that enables our strongest inheritance results.

Negation-limited circuits: Jukna and Lingas introduced negation width as a complexity measure and proved general lower bounds for circuits with bounded negation, quantifying the "penalty factor" for introducing limited negation into monotone circuits[3]. However, their applications were constrained by older, weaker monotone lower bounds.

Hardness magnification: Chen, Hirahara, and Oliveira showed that proving superlinear circuit lower bounds for MCSP would imply $\mathrm{NP} \not\subseteq \mathrm{P/poly}$ via hardness magnification[9]. Williams established non-uniform ACC lower bounds using algorithmic techniques[30]. Our work provides stronger structured lower bounds that may serve as inputs to future magnification frameworks.

Circuit lower bounds landscape: The strongest unconditional circuit lower bounds for explicit functions remain exponential for monotone circuits[1][2], super-polynomial for bounded-depth circuits[30], and modest for general circuits. Our negation-width framework provides a parametric interpolation revealing how hardness degrades as negation is introduced.

1.5 Paper Organization

Section 2 establishes notation and preliminaries on Boolean circuits, monotone complexity, and negation width. Section 3 states and proves the Starflower Inheritance Theorem and defines the Brazil Threshold. Section 4 applies the theorem to five explicit function families, constructing the negation ladder. Section 5 discusses implications for hardness magnification and P versus NP. Section 6 concludes with open problems and future directions.

---

2. Preliminaries

2.1 Boolean Circuits and Monotone Complexity

Definition 2.1 (DeMorgan Circuit). A **DeMorgan circuit** over Boolean variables $x_1, \ldots, x_N$ is a directed acyclic graph where:

- Input nodes are labeled by variables x_i or constants $0, 1$.
- Internal nodes (gates) are labeled by AND ($\wedge$), OR ($\vee$), or NOT ($\neg$).
- One node is designated as the output gate.

The **size** of a circuit F is the number of gates (internal nodes). A circuit is **monotone** if it contains no NOT gates.

Definition 2.2 (Circuit Complexity). For a Boolean function $f : \{0,1\}^N \to \{0,1\}$:

- The **circuit complexity** $C(f)$ is the minimum size of any DeMorgan circuit computing f.
- The **monotone circuit complexity** $C^+(f)$ is the minimum size of any monotone circuit computing f (applicable only if f is monotone).

A function f is **monotone** if for all $x, y \in \{0,1\}^N$, whenever $x \leq y$ coordinate-wise, we have $f(x) \leq f(y)$.

Remark 2.3. For any monotone function f, we have $C(f) \leq C^+(f)$ since monotone circuits are a special case of DeMorgan circuits. The central question is: how much smaller can $C(f)$ be compared to $C^+(f)$ when negations are introduced?

2.2 Negation Width: Measuring Limited Negation

To study the interpolation between monotone and general circuits, we use the **negation width** measure introduced by Jukna and Lingas[3].

Definition 2.4 (Prime Implicant). For a monotone Boolean function f, a **prime implicant** is a minimal conjunction of (unnegated) variables that implies $f(x) = 1$. Equivalently, it is a minimal set $S \subseteq [N]$ such that setting $x_i = 1$ for all $i \in S$ forces $f(x) = 1$, and no proper subset of S has this property.

We denote the set of all prime implicants of f by $\mathrm{PI}(f)$.

Definition 2.5 (Negation Width)[3]. Let F be a DeMorgan circuit computing a monotone function f. For each prime implicant $p \in \mathrm{PI}(f)$, define $w_p(F)$ as the minimum number of distinct negated variables appearing in any term of the syntactic Disjunctive Normal Form (DNF) expansion of F that extends (implies) the prime implicant p.

The **negation width** of circuit F is:

$$w(F) := \max_{p \in \mathrm{PI}(f)} w_p(F).$$

Intuitively, $w(F)$ measures the maximum number of variable negations required to "activate" the hardest prime implicant of f via the circuit F.

Definition 2.6 (Negation-Width Complexity). For a monotone function f and integer $w \geq 0$, define:

$$C_w(f) := \min\{|F| : F \text{ computes } f \text{ and } w(F) \leq w\}.$$

This is the minimum size of any circuit computing f with negation width at most w.

Observation 2.7. We have:

- $C_0(f) = C^+(f)$ (width-0 circuits are monotone).
- $C_w(f)$ is non-increasing in w.
- $C_N(f) = C(f)$ (unrestricted circuits).

The central question is: how does $C_w(f)$ degrade as w increases from 0 to N?

2.3 The Jukna–Lingas Framework

Jukna and Lingas established a general lower bound relating negation-width complexity to monotone complexity via a "penalty factor"[3].

Theorem 2.8 (Jukna–Lingas 2019)[3]. Let f be a monotone Boolean function with maximum prime implicant length m (the size of the largest minimal prime implicant). For any $w \geq 0$, define the penalty factor:

$$K(m, w) := \min\{w^m, m^w\}.$$

Then:

$$C_w(f) \geq \frac{C^+(f)}{4 \cdot K(m, w) \cdot \log |\mathrm{PI}(f)|} - 1.$$

Interpretation: Introducing negation width w allows the circuit size to shrink by at most a factor proportional to $K(m, w)$. If $K(m, w)$ grows sub-exponentially relative to $C^+(f)$, then exponential monotone hardness is inherited.

Remark 2.9. The penalty factor $K(m, w) = \min(w^m, m^w)$ arises from combinatorial arguments: each prime implicant can be "covered" by at most w^m negation patterns (choosing which variables to negate), or equivalently at most m^w ways to select w positions from m variables.

This theorem provides the foundation for proving inherited lower bounds by analyzing when the penalty remains sub-exponential.

---

3. The Starflower Inheritance Theorem

3.1 Statement and Proof of Main Theorem

We now establish our central result: exponential monotone lower bounds are inherited by circuits with sub-polynomial negation width.

Theorem 3.1 (General Starflower Inheritance). Let $(f_N)_{N\in\mathbb{N}}$ be a family of monotone Boolean functions satisfying:

$$C^+(f_N) \geq 2^{N^{\alpha-o(1)}}$$

for some constant $\alpha \in (0,1)$. Assume the maximum prime implicant length satisfies $m(N) = N^{O(1)}$ (polynomial in N). Then for every fixed $\varepsilon \in (0,\alpha)$, there exists a constant $c > 0$ such that any DeMorgan circuit F computing f_N with negation width

$$w(F) \leq \frac{N^{\alpha-\varepsilon}}{\log N}$$

must have size

$$|F| \geq 2^{c\cdot N^{\alpha-\varepsilon}}.$$

Proof. We apply the Jukna–Lingas framework (Theorem 2.8). For a circuit F computing f_N with negation width w:

$$|F| = C_w(f_N) \geq \frac{C^+(f_N)}{4\cdot K(m,w)\cdot \log|\mathrm{PI}(f_N)|} - 1.$$

Step 1: Bound the penalty factor. We have $m = N^{O(1)}$, so there exists a constant $d > 0$ such that $m \leq N^d$. The penalty factor is:

$$\begin{aligned}
K(m,w) &= \min(w^m, m^w) \\
&\leq m^w \quad \text{(taking the second term)} \\
&\leq (N^d)^w \\
&= N^{dw} \\
&= 2^{dw\log N}.
\end{aligned}$$

Substituting the negation width bound $w \leq N^{\alpha-\varepsilon}/\log N$:

$$K(m, w) \leq 2^{d \cdot \frac{N^{\alpha-\varepsilon}}{\log N} \cdot \log N} = 2^{d \cdot N^{\alpha-\varepsilon}}.$$

Step 2: Bound the number of prime implicants. The number of prime implicants is at most 2^N (each subset of variables can correspond to at most one prime implicant). Thus:

$$\log |\mathrm{PI}(f_N)| \leq N.$$

Step 3: Combine the bounds. We have:

$$\begin{aligned} C_w(f_N) &\geq \frac{C^+(f_N)}{4 \cdot K(m, w) \cdot N} - 1 \\ &\geq \frac{2^{N^{\alpha-o(1)}}}{4 \cdot 2^{d \cdot N^{\alpha-\varepsilon}} \cdot N} - 1 \\ &= \frac{2^{N^{\alpha-o(1)}}}{2^{2+d \cdot N^{\alpha-\varepsilon}+\log N}} - 1 \\ &= 2^{N^{\alpha-o(1)}-(2+d \cdot N^{\alpha-\varepsilon}+\log N)} - 1. \end{aligned}$$

Step 4: Analyze the exponent. Since $\alpha - \varepsilon < \alpha$, we have:

$$N^{\alpha-o(1)} - d \cdot N^{\alpha-\varepsilon} = N^{\alpha-\varepsilon} \cdot \left(N^{\alpha-(\alpha-\varepsilon)-o(1)} - d\right) = N^{\alpha-\varepsilon} \cdot \left(N\right.$$

For sufficiently large N, the term $N^{\varepsilon-o(1)} - d$ is positive and grows without bound. The $\log N$ and constant terms are negligible. Therefore:

$$N^{\alpha-o(1)} - (d \cdot N^{\alpha-\varepsilon} + \log N + 2) \geq c' \cdot N^{\alpha-\varepsilon}$$

for some constant $c' > 0$ and all sufficiently large N.

Step 5: Conclusion. For sufficiently large N:

$$C_w(f_N) \geq 2^{c' \cdot N^{\alpha-\varepsilon}} - 1 \geq 2^{c \cdot N^{\alpha-\varepsilon}}$$

for some constant $c > 0$ (absorbing the -1 term and adjusting the constant). This completes the proof. □

3.2 The Brazil Threshold

The theorem naturally defines a critical threshold for negation width.

Definition 3.2 (Brazil Threshold). For a function family with monotone hardness exponent α and parameter $\varepsilon \in (0, \alpha)$, the **Brazil Threshold** is:

$$w_{\text{Brazil}}(N, \alpha, \varepsilon) := \frac{N^{\alpha-\varepsilon}}{\log N}.$$

Below this threshold, circuits must have exponential size $2^{\Omega(N^{\alpha-\varepsilon})}$.

Corollary 3.3 (Hardness Inheritance Below Threshold). Any DeMorgan circuit computing f_N with negation width $w < w_{\text{Brazil}}(N, \alpha, \varepsilon)$ must satisfy:

$$|F| \geq 2^{\Omega(N^{\alpha-\varepsilon})}.$$

Exponential monotone hardness is **inherited** by all circuits with sub-polynomial negation width.

3.3 Interpretation: Structural Invariance

The Starflower Theorem demonstrates that exponential monotone complexity reflects **structural hardness** that is robust to bounded perturbations. Introducing negations is analogous to introducing "noise" or "interference" into a system; the theorem shows that as long as this interference remains below a critical threshold (sub-polynomial relative to the hardness exponent), the exponential hardness persists.

Physical analogy: Just as signal integrity in a communication channel can be preserved under bounded noise, computational hardness can be preserved under bounded negation. The monotone core of a function's complexity is a "structural invariant" protected against limited modifications.

Implications for circuit models: The result suggests that the gap between monotone and general circuits may be narrower than previously thought in terms of **hardness preservation**. While general circuits can be exponentially smaller than monotone circuits for specific functions (Tardos function[12]), the Starflower framework shows that this advantage only appears when negation width exceeds the Brazil Threshold. For many natural functions, substantial negation-width buffers exist where exponential hardness survives.

---

4. The Five Function Families: A Negation Ladder

We now apply the Starflower Inheritance Theorem to five explicit function families with varying monotone hardness exponents, establishing a unified **negation ladder** of inherited lower bounds.

4.1 Family 1: Bipartite Perfect Matching (Strongest Result)

The bipartite perfect matching function is defined on graphs with n vertices per side.

Function Definition 4.1. $\mathrm{Match}_n : \{0,1\}^{n^2} \to \{0,1\}$ takes an $n \times n$ bipartite adjacency matrix and outputs 1 if and only if there exists a perfect matching.

Theorem 4.2 (Çalar et al. 2025)[2]. The bipartite perfect matching function satisfies:

$$C^+(\mathrm{Match}_n) \geq 2^{n^{1/3-o(1)}}.$$

This is the strongest known monotone lower bound for an explicit natural function.

Corollary 4.3 (Matching with Bounded Negation Width). Applying Theorem 3.1 with $N = n^2$, $\alpha = 1/3$, and any $\varepsilon \in (0, 1/3)$, any DeMorgan circuit computing Match_n with negation width:

$$w \leq \frac{n^{1/3-\varepsilon}}{\log n}$$

requires size:

$$|F| \geq 2^{\Omega(n^{1/3-\varepsilon})}.$$

Quantitative improvement: For $\varepsilon = o(1)$ (e.g., $\varepsilon = 1/\log\log n$), this allows negation width up to approximately:

$$w \approx \frac{n^{1/3}}{\log n \cdot \log\log n}.$$

Prior results using Razborov's original bound allowed negation width only $w \approx \log n$[1]. This represents a **super-polynomial improvement** in the robustness regime.

External validation: Professor Stasys Jukna confirmed in personal correspondence (March 17, 2026) that this application of his negation-width framework to the new matching bound is correct and represents a previously unpublished result. He stated: "Just because the improvement of Razborov's bound is fresh... Just insert the newly obtained exp(n^{1/3}) lower bound instead of the previously known exp((log^2 n)), and one is done."

Significance: Bipartite matching is a fundamental problem in combinatorial optimization, appearing across computer science from algorithms to complexity theory. The strong inherited lower bound demonstrates that matching hardness is intrinsic and robust, not an artifact of the monotone restriction.

4.2 Family 2: GEN-TFNP Search Problems

Total search problems in TFNP (Total Function Nondeterministic Polynomial time) are computational search problems guaranteed to have solutions. These problems have deep connections to circuit complexity via lifting theorems[10].

Background: Resolution width lower bounds for TFNP problems can be lifted to monotone circuit lower bounds using composition techniques. For the **Generation (GEN)** problem, resolution width is $\Omega(N^{1/3})$[10], which lifts to monotone circuit complexity with similar exponents.

Function Family 4.4. Let f_{GEN} denote the Boolean function associated with the GEN search problem via standard lifting constructions. Then:

$$C^+(f_{\text{GEN}}) \geq 2^{\Omega(N^{1/3})}.$$

Corollary 4.5 (GEN with Bounded Negation Width). Applying Theorem 3.1 with $\alpha = 1/3$:

$$C_w(f_{\text{GEN}}) \geq 2^{\Omega(N^{1/3-\varepsilon})} \quad \text{for } w \leq \frac{N^{1/3-\varepsilon}}{\log N}.$$

This establishes that TFNP search hardness extends robustly into the negation-width regime, suggesting fundamental computational barriers persist even with limited negation.

4.3 Family 3: Pigeonhole Principle and Clique-Coloring

The Pigeonhole Principle and related clique-coloring functions exhibit monotone hardness with exponents around $\alpha \approx 1/4$.

Function Family 4.6. The Pigeonhole Principle function $\text{PHP}_{n+1,n}$ expresses that $n+1$ pigeons cannot be placed into n holes with at most one pigeon per hole. Razborov established super-polynomial monotone bounds[11]; subsequent improvements yield:

$$C^+(\text{PHP}_{n+1,n}) \geq 2^{\Omega(n^{1/4})}.$$

Corollary 4.7 (Pigeonhole with Bounded Negation Width). Applying Theorem 3.1 with $\alpha = 1/4$:

$$C_w(\text{PHP}_{n+1,n}) \geq 2^{\Omega(n^{1/4-\varepsilon})} \quad \text{for } w \leq \frac{n^{1/4-\varepsilon}}{\log n}.$$

Similar results hold for clique-coloring functions, where the goal is to determine if a graph can be properly colored with k colors (monotone version checking uncolorability via cliques).

Significance: The Pigeonhole Principle is a fundamental combinatorial statement appearing in logic, complexity theory, and proof complexity. The inherited lower bound shows that its computational hardness extends beyond the monotone model.

4.4 Family 4: Tardos Weight Function

Tardos constructed a monotone function demonstrating an exponential gap between monotone and general circuit complexity[12].

Function Family 4.8 (Tardos 1988). There exists a monotone function f_T on n variables such that:

$$C^+(f_T) \geq 2^{\Omega(n^{1/8})},$$

while $C(f_T) = \text{poly}(n)$ (polynomial-size general circuits exist).

This function is explicitly constructed as a weighted threshold function using specific weight assignments.

Corollary 4.9 (Tardos Function with Bounded Negation Width). Applying Theorem 3.1 with $\alpha = 1/8$:

$$C_w(f_T) \geq 2^{\Omega(n^{1/8-\varepsilon})} \quad \text{for } w \leq \frac{n^{1/8-\varepsilon}}{\log n}.$$

Interpretation: Although the Tardos function admits small general circuits (implying high negation width allows collapse), our theorem shows that there is a **transitional regime** where exponential hardness persists. Specifically, for negation width up to $w \approx n^{1/8}/\log n$, the function remains exponentially hard. The collapse to polynomial size occurs only when negation width exceeds this threshold.

This demonstrates that even functions with exponential monotone-versus-general gaps exhibit **gradual hardness degradation** as negation is introduced, rather than abrupt collapse.

4.5 Family 5: The de Rezende–Vinyals P-Function

Recent work on lifting techniques using "colourful sunflowers" establishes strong monotone lower bounds for P-functions derived from Constraint Satisfaction Problems (CSPs)[10].

Function Family 4.10. For certain CSP-derived P-functions f_P, lifting theorems establish:

$$C^+(f_P) \geq 2^{\Omega(N^\beta)},$$

where β depends on the specific CSP structure (typically $\beta \in [1/4, 1/3]$).

Corollary 4.11 (P-Function with Bounded Negation Width). Applying Theorem 3.1 with hardness exponent β:

$$C_w(f_P) \geq 2^{\Omega(N^{\beta-\varepsilon})} \quad \text{for } w \leq \frac{N^{\beta-\varepsilon}}{\log N}.$$

These results connect circuit complexity to proof complexity and CSP hardness, demonstrating that computational barriers in CSPs extend robustly into negation-limited circuit models.

4.6 The Unified Negation Ladder

We summarize the inherited lower bounds for all five families.

Function Family	**Exponent** α	**Negation Width Threshold**	**Lower Bound**
Bipartite Matching	$1/3$	$n^{1/3-\varepsilon}/\log n$	$2^{\Omega(n^{1/3-\varepsilon})}$
GEN-TFNP	$1/3$	$N^{1/3-\varepsilon}/\log N$	$2^{\Omega(N^{1/3-\varepsilon})}$
Pigeonhole-Clique-Coloring	$\approx 1/4$	$n^{1/4-\varepsilon}/\log n$	$2^{\Omega(n^{1/4-\varepsilon})}$
Tardos Weight	$1/8$	$n^{1/8-\varepsilon}/\log n$	$2^{\Omega(n^{1/8-\varepsilon})}$
P-function (CSP-derived)	$\beta \approx 1/3$	$N^{\beta-\varepsilon}/\log N$	$2^{\Omega(N^{\beta-\varepsilon})}$

Table 1: The Negation Ladder: negation-width thresholds and inherited lower bounds for five explicit function families. Each family exhibits exponential hardness below its Brazil Threshold.

Key observation: In every case, the Starflower Theorem automatically provides a **negation-width buffer** of size approximately $N^{\alpha-\varepsilon}$, within which the function remains exponentially hard. This buffer is substantial—super-polynomial for all families—demonstrating that monotone hardness is remarkably robust.

---

5. Discussion: Implications and Open Problems

5.1 Toward Hardness Magnification and P versus NP

Recent hardness magnification results establish that proving sufficiently strong circuit lower bounds on certain "intermediate" problems can be amplified into major complexity class separations[9].

Background on magnification: Chen, Hirahara, and Oliveira showed that proving a superlinear circuit lower bound for the Minimum Circuit Size Problem (MCSP) would imply $\mathrm{NP} \not\subseteq \mathrm{P}/\mathrm{poly}$[9]. MCSP is the problem of determining whether a given Boolean function (specified by its truth table) has a circuit of size at most s.

Connection to our results: The Starflower Theorem provides strong lower bounds for explicit functions (like matching) with additional structure (bounded negation width). If reductions from these functions to MCSP could preserve negation-width constraints while maintaining the size parameters required for magnification, our exponential lower bounds would trigger the magnification mechanism.

Open Problem 5.1 (Magnification via Matching). Can the exponential lower bound for bipartite matching with bounded negation width be leveraged via reduction to MCSP to prove:

$$\text{MCSP requires superlinear circuits} \implies \mathrm{NP} \not\subseteq \mathrm{P}/\mathrm{poly}?$$

Specifically, does there exist a reduction from negation-width-w matching circuits to MCSP that preserves the exponential hardness and triggers magnification?

Challenges: Standard reductions may not preserve negation-width structure. New reduction techniques sensitive to negation patterns would be required. Additionally, MCSP hardness amplification requires lower bounds for MCSP itself, not merely for functions reducible to MCSP.

Potential pathway: If matching hardness with bounded negation can be "lifted" or "composed" with MCSP in a structure-preserving way, the resulting MCSP instance might inherit exponential hardness, providing the superlinear lower bound needed for magnification.

5.2 Conceptual Implications for Circuit Complexity

Monotone-versus-general gap: The classical understanding is that monotone lower bounds say little about general circuit complexity because negation can provide exponential advantages (Tardos function[12]). However, the Starflower framework reveals a more nuanced picture: exponential monotone hardness **does** constrain a large class of non-monotone circuits—specifically, those with sub-polynomial negation width.

This suggests that:

- The "gap" between monotone and general circuits is parametric, not binary.
- For many natural functions, the transition from exponential to polynomial complexity occurs gradually as negation width increases, rather than abruptly.
- Structural hardness information in monotone circuits extends substantially into non-monotone models.

Negation as a resource: Viewing negation width as a computational resource analogous to depth, space, or time provides a new lens for analyzing circuit complexity. The Brazil Threshold defines the "critical resource allocation" beyond which hardness collapses.

Open Problem 5.2 (Tightness of Brazil Threshold). Is the Brazil Threshold $w_{\text{Brazil}}(N, \alpha, \varepsilon) = N^{\alpha-\varepsilon}/\log N$ tight? Specifically:

- Do there exist functions with $C^+(f) \geq 2^{N^\alpha}$ and circuits computing f with size $\text{poly}(N)$ and negation width $w \approx N^\alpha/\text{poly}\log N$?
- Or can the threshold be improved to $w \approx N^{\alpha-o(1)}$?

Understanding tightness would clarify whether our theorem captures the true "phase transition" boundary for hardness inheritance.

5.3 Extensions to Other Circuit Models

The negation-width framework naturally extends to other restricted circuit classes:

Depth-bounded circuits: Can similar inheritance theorems be established for circuits with bounded depth and bounded negation width? For example, do exponential monotone lower bounds for bounded-depth circuits (such as AC^0) persist under limited negation?

Threshold circuits: Threshold gates compute weighted threshold functions. Combining monotone threshold circuits with bounded negation might yield new lower bounds for threshold circuit classes.

Quantum circuits: While negation is trivial in quantum circuits (phase gates), analogous "perturbation width" measures might capture limited deviations from structured quantum circuits, potentially yielding quantum hardness inheritance theorems.

Open Problem 5.3. Develop negation-width analogs for other circuit models (depth-bounded, threshold, quantum) and establish corresponding inheritance theorems.

5.4 Algorithmic Applications

Circuit synthesis: Understanding negation-width constraints can inform circuit synthesis algorithms. If a target function is known to require exponential size for small negation width, synthesis tools can avoid exploring low-negation-width circuits, pruning the search space.

Satisfiability and optimization: Negation-width constraints on circuit representations of satisfiability problems might lead to new algorithmic techniques or hardness results for SAT variants.

Learning theory: The Minimum Circuit Size Problem is closely related to learning Boolean functions. Hardness results for MCSP with negation-width constraints could imply new hardness-of-learning results for restricted hypothesis classes.

6. Conclusion and Future Directions

6.1 Summary of Results

We established the **Starflower Inheritance Theorem**, proving that exponential monotone circuit lower bounds persist under bounded negation width. Applied to five explicit function families—bipartite matching (strongest result), GEN-TFNP, pigeonhole/clique-coloring, Tardos weight, and CSP-derived P-functions—we constructed a **negation ladder** of inherited hardness.

For bipartite matching using the recent Çalar et al. breakthrough bound, we achieved exponential lower bounds $2^{\Omega(n^{1/3-\varepsilon})}$ for circuits with negation width up to $n^{1/3-\varepsilon}/\log n$, representing a super-polynomial improvement over prior results.

The **Brazil Threshold** provides a universal formula connecting monotone hardness exponents to negation-width robustness, revealing that exponential structural complexity is a robust invariant preserved under substantial perturbations.

6.2 Open Problems

We conclude with key open problems for future research:

1. **Hardness magnification:** Can inherited lower bounds for matching or other functions be leveraged to prove superlinear MCSP lower bounds, triggering magnification to $\mathrm{NP} \not\subseteq \mathrm{P/poly}$?
2. **Tightness of Brazil Threshold:** Is $w_{\mathrm{Brazil}}(N, \alpha, \varepsilon) = N^{\alpha-\varepsilon}/\log N$ tight, or can it be improved?
3. **Improved monotone lower bounds:** Can the matching exponent be improved beyond $\alpha = 1/3$, yielding even stronger inherited bounds?
4. **Extensions to other models:** Develop negation-width analogs for depth-bounded, threshold, and quantum circuits.
5. **Algorithmic applications:** Apply negation-width insights to circuit synthesis, SAT solving, and learning theory.
6. **Unified complexity theory:** Explore whether the Starflower principle—structural complexity preserved under bounded perturbations—extends to other computational models (Turing

machines with limited nondeterminism, communication protocols with limited information transfer, etc.).

6.3 Broader Vision

The Starflower framework suggests that **structural complexity is a robust invariant** across computational models. Exponential hardness, once established in a restricted model, extends substantially into less-restricted models as long as the relaxation remains bounded. This principle may unify lower bound techniques across circuit complexity, proof complexity, communication complexity, and algorithm design.

Understanding the **parametric transition** from exponential to polynomial complexity—as resources like negation width increase—may ultimately reveal the fundamental boundaries of efficient computation, bringing us closer to resolving P versus NP and related foundational questions.

Acknowledgments

This work builds on foundational contributions from numerous researchers in circuit complexity and computational complexity theory.

Circuit complexity foundations: Alexander Razborov (monotone lower bounds, 1985)[1], Stasys Jukna and Andrzej Lingas (negation-width framework, 2019)[3], Bruno Çalar, Mika Göös, Artur Riazanov, Anastasia Sofronova, and Dmitry Sokolov (monotone matching bound, 2025)[2].

Special acknowledgment: The author is grateful to Professor Stasys Jukna for taking time from his research schedule to confirm the validity of the matching application (Corollary 4.3) in personal correspondence (March 17, 2026). Professor Jukna verified that combining his negation-width theorem with the recent Çalar et al. exponential matching bound yields the claimed inherited lower bound for negation width up to $w \approx n^{1/3} / \log \log n$, and confirmed that this specific combination represents a new, previously unpublished result. His willingness to engage with independent work

exemplifies the collaborative spirit of the complexity theory community.

Hardness magnification: Lijie Chen, Shuichi Hirahara, and Igor Oliveira (MCSP magnification, 2022)[9], Ryan Williams (ACC lower bounds and algorithmic techniques)[30].

Lifting and proof complexity: Susanna de Rezende and Marc Vinyals (colourful sunflowers lifting, 2025)[10], connecting proof complexity to circuit complexity via composition techniques.

Classical complexity theory: The Clay Mathematics Institute for formulating the P versus NP problem as a Millennium Prize Problem, inspiring decades of research toward resolution.

We thank the computational complexity community for creating the intellectual foundations upon which this work builds.

References

[1] Razborov, A. A. (1985). Lower bounds on the monotone complexity of some Boolean functions. *Soviet Mathematics Doklady*, 31, 354–357.

[2] Çalar, M., Göös, M., Riazanov, A., Sofronova, A., & Sokolov, D. (2025). Monotone Circuit Complexity of Matching. *arXiv preprint arXiv:2507.16105*.

[3] Jukna, S., & Lingas, A. (2019). Lower Bounds for DeMorgan Circuits of Bounded Negation Width. *STACS 2019, LIPIcs*, Vol. 126, Art. 41. DOI: 10.4230/LIPIcs.STACS.2019.41.

[9] Chen, L., Hirahara, S., & Oliveira, I. C. (2022). NP-Hardness of Learning Programs and Partial MCSP. *Proceedings of IEEE FOCS*, 968–979.

[10] de Rezende, S. F., & Vinyals, M. (2025). Lifting with Colourful Sunflowers. *Proceedings of CCC 2025, LIPIcs*, Vol. 339, Art. 36.

[11] Razborov, A. A. (1985). Lower bounds for the Pigeonhole Principle. Various results on monotone circuit complexity.

[12] Tardos, É. (1988). The gap between monotone and non-monotone circuit complexity is exponential. *Combinatorica*, 8(1), 141–142.

[30] Williams, R. (2014). Nonuniform ACC Circuit Lower Bounds. *Journal of the ACM*, 61(1), Article 2.