RESEARCH CORRESPONDENCE 27 NOVEMBER 2026

Seven extremal one-bit edits at every input width from four onward

Mohamad Al-Zawahreh ORCID 0009-0006-7592-9166

Changing one answer in a Boolean truth table can force the smallest circuit to grow, even if every gate may be redesigned. Krinkin asked whether that growth can be linear at infinitely many widths. In the AND-inverter model the worst-case increase on n inputs is exactly n for $n \geq 2$. When AND, OR and NOT each cost one, it is $n + 1$ for $n \geq 3$. Here we use the restriction argument from Section 4 of the parent preprint to carry all seven extremal edges in the public four-input AIG census to larger widths. The resulting expensive endpoints remain distinct even after input permutation or complementation of inputs and output.

The one-bit upper bounds

Circuits have one output, free fanout and unrestricted gate reuse. In an AND-inverter graph (AIG), binary AND costs one and wire complements are free. Testing equality to an n-bit address costs $n - 1$ ANDs. One more gate adds or removes the true row. Reversing the edit gives $|\operatorname{opt}(f) - \operatorname{opt}(g)| \leq n$. In the counted AND/OR/NOT basis, constants must be constructed. For a mixed address, form the AND U of the one-positions and the OR V of the zero-positions. Equality is $U \wedge \neg V$ and inequality is $V \vee \neg U$. Either test costs at most n. At an endpoint use a single conjunction or disjunction of all inputs and, if needed, one NOT. Combining the test with the old output costs one more gate.

How to lift any attaining pair

Let f, g differ only at address a, with g the more expensive function. On a fresh input z, set $F(x, z) = f(x)$, and set $G(x, 0) = g(x)$, $G(x, 1) = f(x)$. The edit now occurs only at $(a, 0)$. Ignoring z gives $\operatorname{opt}(F) \leq \operatorname{opt}(f)$. Tying z to an old input in any circuit for F gives the reverse inequality.

The fresh bit is essential in G. Its cofactor at zero is the nonconstant g, so the output cannot simply be z or its complement. At least one gate in the output cone takes z or its literal as an input. On restriction to $z = 0$, that gate becomes a wire or a constant. Propagate constants; the restricted output is nonconstant. At least one charged gate disappears, including in a DAG with shared wires. Hence $\operatorname{opt}(G) \geq \operatorname{opt}(g) + 1$. In either listed basis the upper bound on the gap grows by exactly one when n increases by one. If the old pair attains the bound, so does the lift, with the cheaper optimum unchanged. Iterate.

Two base cases

In the AIG basis the two-input pair NAND and XOR differs only at 00 and costs one and three ANDs. XOR has no constant cofactor. A second and final AND using a primary literal would force one; a final AND using only the first output, its complement and constants produces a constant or an already available literal. Three ANDs suffice. In the counted basis, start with $\neg(x_1 \vee x_2)$ and add the row 110 on three inputs. The optima are two and six. Enumerating all attainable wire sets yields $3, 12, 40, 84, 121, 203$ reachable output functions through budgets zero to five, with the target absent. Six gates realize $\neg(x_1 \vee x_2) \vee (x_1 \wedge x_2 \wedge \neg x_3)$. The parent archive supplies the enumeration. These bases and the lift prove the stated maxima.

hexagon:2611.01269v1

Seven distinct families in the free-inversion basis

Krinkin and Busnello's pinned data give exact AIG sizes for all 222 four-input NPN classes. Here NPN equivalence permits input permutation, independent input negations and output complementation. Enumerating the 995 undirected edges in their mutation graph gives seven at distance four in gate count. For each edge we chose the cheaper class's canonical truth table and flipped one physical output bit to reach the more expensive class. Each choice was checked under all input permutations, all input sign masks and both output polarities.

Lower raw/class	Higher raw	Flipped row	Higher class	ANDs before/after	Lower weight
`0001`	`8001`	15	`0180`	$3 \to 7$	1
`0003`	`4003`	14	`0181`	$2 \to 6$	2
`0007`	`1007`	12	`0119`	$3 \to 7$	3
`000f`	`100f`	12	`01aa`	$1 \to 5$	4
`001f`	`801f`	15	`01ea`	$3 \to 7$	5
`0356`	`8356`	15	`19e1`	$5 \to 9$	6
`0357`	`8357`	15	`0776`	$3 \to 7$	7

Raw words are hexadecimal, with output at address i stored in bit i; x_1 is the least significant address bit. Use the first two columns to see the one-row difference. The two independently normalized class names can differ in many bits. For example, the all-width family $f = \neg(x_1 \wedge x_2)$ and $g = f \wedge (x_1 \vee \cdots \vee x_n)$ reads `7777` and `7776` at four inputs. Those words normalize to the fourth class pair in the table.

Let j be the number of true rows in the cheaper seed. The seven values of j are $1, \ldots, 7$, and in every alignment the edited row changes from zero to one. After adding k fresh inputs, the cheaper function has $j2^k$ true rows and the expensive function has $j2^k + 1$. These seven latter counts are distinct and at most half of $16 \cdot 2^k$. Input symmetries preserve weight, and output complementation preserves the smaller of the true- and false-row counts. Therefore at every width $n = 4 + k$ there are *at least seven NPN-inequivalent expensive functions*, each a one-row edit from a function with exactly n fewer AND gates. The argument places no upper limit on the number of other extremal classes at those widths.

What the supplied circuits check

The reference pack instantiates all seven seeds at four through eight inputs. Its 35 pairs are 70 separate binary AIGER files. The low member ignores all new bits. An address test and a final output gate supply the high member with exactly n additional ANDs. Every file was decoded and its truth table computed. The proof gives the exact optima above four inputs. A larger implementation can be functionally correct while missing that optimum.

Sources and credit. K. Krinkin, one-point perturbations (March 2026); K. Krinkin and L. A. Busnello, four-input data and July completion; A. Golovnev et al., prior bounded-difference argument; M. Al-Zawahreh, parent preprint, version 1.1, Section 4. This AI-assisted correspondence note adds the alignment and simultaneous seven-family deduction to the previously published sharp bounds. Text licensed under CC BY 4.0.