A 185/14 Sorter on Our Own Twist

the best known $n = 32$ network, in coordinates where the construction is visible

Original Results by **Mavdil**[*],
with **Heath Hunnicutt**[†]
Ruach Tov Collective
August 10, 2026

The network below sorts 32 inputs in 185 comparators and 14 layers — the best known figures. Its first sixty-four comparators are the four canonical strides. Its next sixteen are the twist our plane rule generates, unaided, from the subspace $\{5, 10, 15\}$. It is not a new bound; it is the known network in coordinates where our construction can be read off the page.

The network

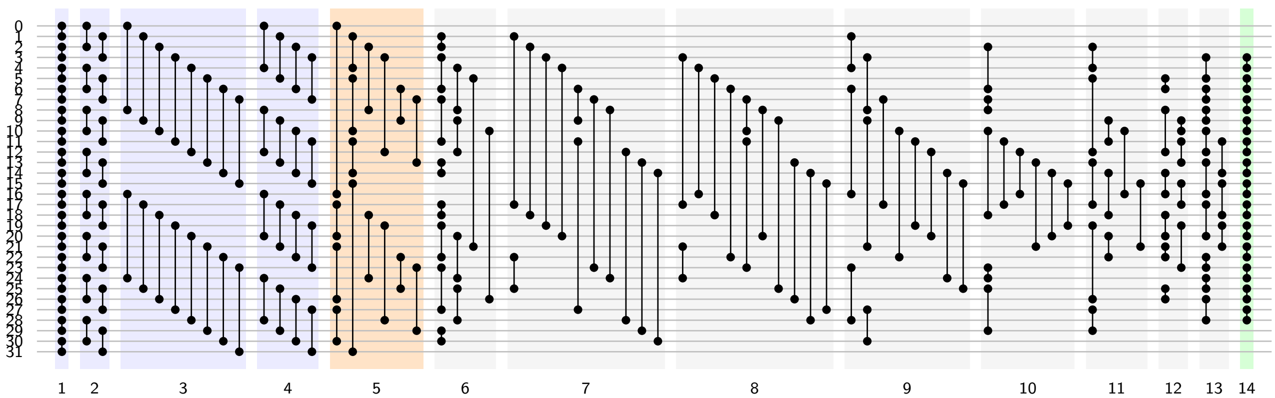

Wire 0 at the top, wire 31 at the bottom; time runs left to right. The blue band is the hypercube prefix, the orange band the twist, the green band the closure.

* Corresponding AI author: `mavdil@ruachtov.ai`
† Corresponding human author: `heath@ruachtov.ai`

The layers

layer	size	directions $d = a \oplus c$
L_1	16	1
L_2	16	2
L_3	16	8
L_4	16	4
L_5	16	$\mathbf{5}$, $\mathbf{10}$, $\mathbf{15}$, 16
L_6	14	$3, 5, 12, 16$
L_7	12	$15, 16$
L_8	12	$13, 16, 18, 20, 23, 28$
L_9	12	$5, 11, 22, 24, 28$
L_{10}	10	$4, 15, 24, 26, 28$
L_{11}	10	$2, 6, 9, 26, 28$
L_{12}	10	$3, 4, 6, 30$
L_{13}	12	$1, 5, 6, 29$
L_{14}	13	$\mathbf{3}$, $\mathbf{7}$, $\mathbf{15}$, $\mathbf{31}$

Two layers deserve their bold type. L_5 carries the plane $\{5, 10, 15\}$ together with the axis 16 on the extremal pairs — the twist, exactly as `plane_gen.py` emits it. L_{14} carries $\{3, 7, 15, 31\}$, the **complete adjacency ladder**: every direction that can appear in a last layer, and no other.

The comparators

L_1	`(0,1)(2,3)(4,5)(6,7)(8,9)(10,11)(12,13)(14,15)(16,17)(18,19)(20,21)(22,23)(24,25)(26,27)(28,29)(30,31)`
L_2	`(0,2)(1,3)(4,6)(5,7)(8,10)(9,11)(12,14)(13,15)(16,18)(17,19)(20,22)(21,23)(24,26)(25,27)(28,30)(29,31)`
L_3	`(0,8)(1,9)(2,10)(3,11)(4,12)(5,13)(6,14)(7,15)(16,24)(17,25)(18,26)(19,27)(20,28)(21,29)(22,30)(23,31)`
L_4	`(0,4)(1,5)(2,6)(3,7)(8,12)(9,13)(10,14)(11,15)(16,20)(17,21)(18,22)(19,23)(24,28)(25,29)(26,30)(27,31)`
L_5	`(0,16)(1,4)(2,8)(3,12)(5,10)(6,9)(7,13)(11,14)(15,31)(17,20)(18,24)(19,28)(21,26)(22,25)(23,29)(27,30)`

The remaining nine layers are in `networks/n32_mavdil_untangled_185.json`.

How our twist compares with Dobbelaere's

The best-known network's own twist comes from the plane $\{6, 9, 15\}$. Ours comes from $\{5, 10, 15\}$. **They agree on half their comparators and disagree on the other half — and the disagreement is a relabelling, not a difference.**

What is the same

Eight of sixteen comparators are **identical**:

$$(0, 16)\ (3, 12)\ (5, 10)\ (6, 9)\ (15, 31)\ (19, 28)\ (21, 26)\ (22, 25).$$

These are the two extremal closers and **all six** pairs the self-mirror 12-class plays. Both planes contain the rung 15, both give it to the self-mirror class, and the class is settled identically.

Everything structural is shared: 16 comparators, forced poset $167/496$, class profile $[6, 6, 4, 4, 4, 4]$, split overlap 78.4%, and — under our own consume — an identical completion at $211/17$.

What is different

The two 8-classes. Ours takes the plane's vectors $\{5, 10\}$; Dobbelaere's takes $\{6, 9\}$. On each group of four wires the two twists are **the same four wires matched the other way**:

wires	ours $\{5, 10\}$	Dobbelaere $\{6, 9\}$
$1, 2, 4, 8$	$(1, 4)$ $(2, 8)$	$(1, 8)$ $(2, 4)$
$7, 11, 13, 14$	$(7, 13)$ $(11, 14)$	$(7, 14)$ $(11, 13)$
$17, 18, 20, 24$	$(17, 20)$ $(18, 24)$	$(17, 24)$ $(18, 20)$
$23, 27, 29, 30$	$(23, 29)$ $(27, 30)$	$(23, 30)$ $(27, 29)$

Why the difference is not a difference

A linear map on $\mathbb{F}_2^5$ carries one twist to the other exactly. The simplest is the axis swap $e_2 \leftrightarrow e_3$; at least three prefix-preserving maps do it. **The two planes are one twist in two coordinate systems**, which is why every measurement taken separately came out the same.

That is also why the transport works. Relabelling Dobbelaere's suffix by the inverse map lands it on our twist. The relabelling is not monotone — it inverts 2 of the 31 adjacent wire pairs — so the transported network is **generalized**: right size, right depth, wrong output order. Knuth's untangling sweep repairs it with size and depth preserved, and the result is the network drawn above.

What this demonstrates, and what it does not

Does. The prefix is derived — four canonical strides. The twist is **calculated**, from a plane chosen by a rule with three independent justifications: no axis, exactly one ladder rung, not the antipodal. The closure is **predicted** — L_{14} is the complete ladder $\{3, 7, 15, 31\}$, which the Adjacency Ladder says is the only admissible content of a final layer. Eighty of the 185 comparators and the shape of the last are ours by construction.

Does not. This is not a new upper bound, and the 105 comparators between the twist and the closure are transported, not derived. Our own consume reaches $208/17$ at best. The middle of the network remains the open phase.

Appendix: the complete network

All 185 comparators, layer by layer. Rules separate the hypercube prefix, the twist, the transported consume, and the closure.

layer	size	comparators
L_1	16	(0,1)(2,3)(4,5)(6,7)(8,9)(10,11)(12,13)(14,15) (16,17)(18,19)(20,21)(22,23)(24,25)(26,27)(28,29)(30,31)
L_2	16	(0,2)(1,3)(4,6)(5,7)(8,10)(9,11)(12,14)(13,15) (16,18)(17,19)(20,22)(21,23)(24,26)(25,27)(28,30)(29,31)
L_3	16	(0,8)(1,9)(2,10)(3,11)(4,12)(5,13)(6,14)(7,15) (16,24)(17,25)(18,26)(19,27)(20,28)(21,29)(22,30)(23,31)
L_4	16	(0,4)(1,5)(2,6)(3,7)(8,12)(9,13)(10,14)(11,15) (16,20)(17,21)(18,22)(19,23)(24,28)(25,29)(26,30)(27,31)
L_5	16	(0,16)(1,4)(2,8)(3,12)(5,10)(6,9)(7,13)(11,14) (15,31)(17,20)(18,24)(19,28)(21,26)(22,25)(23,29)(27,30)
L_6	14	(1,2)(3,6)(4,8)(5,21)(7,11)(9,12)(10,26)(13,14) (17,18)(19,22)(20,24)(23,27)(25,28)(29,30)
L_7	12	(1,17)(2,18)(3,19)(4,20)(6,9)(7,23)(8,24)(11,27) (12,28)(13,29)(14,30)(22,25)
L_8	12	(3,17)(4,16)(5,18)(6,22)(7,10)(8,20)(9,25)(11,23) (13,26)(14,28)(15,27)(21,24)
L_9	12	(1,4)(3,8)(6,16)(7,17)(9,21)(10,22)(11,19)(12,20) (14,24)(15,25)(23,28)(27,30)
L_{10}	10	(2,6)(7,8)(10,18)(11,17)(12,16)(13,21)(14,20)(15,19) (23,24)(25,29)
L_{11}	10	(2,4)(5,12)(9,11)(10,16)(13,17)(14,18)(15,21)(19,26) (20,22)(27,29)
L_{12}	10	(5,6)(8,12)(9,10)(11,13)(14,16)(15,17)(18,20)(19,23) (21,22)(25,26)
L_{13}	12	(3,5)(6,7)(8,9)(10,12)(11,14)(13,16)(15,18)(17,20) (19,21)(22,23)(24,25)(26,28)
L_{14}	13	(3,4)(5,6)(7,8)(9,10)(11,12)(13,14)(15,16)(17,18) (19,20)(21,22)(23,24)(25,26)(27,28)

Appendix II: `plane_gen` at $n = 32$

The generator's output for every one of the 27 serving planes, applied to the four-stride state. Each row is a candidate twist: the plane, its ladder rung, the number of comparators the rule emits, the forced poset afterwards, and the resulting class profile.

plane	rung	size	poset	class profile after
$\{3,5,6\}$	3	10	143/496	$4,4,4,4,4,2,2,2,2$
$\{3,9,10\}$	3	10	143/496	$4,4,4,4,4,2,2,2,2$
$\{3,13,14\}$	3	10	151/496	$4,4,4,4,4,2,2,2,2$
$\{3,17,18\}$	3	10	151/496	$4,4,4,4,4,2,2,2,2$
$\{3,20,23\}$	3	10	151/496	$4,4,4,4,4,2,2,2,2$
$\{3,21,22\}$	3	10	151/496	$4,4,4,4,4,3,3,2,2$
$\{3,24,27\}$	3	10	151/496	$4,4,4,4,4,2,2,2,2$
$\{3,25,26\}$	3	10	151/496	$4,4,4,4,4,3,3,2,2$
$\{3,29,30\}$	3	10	151/496	$4,4,4,4,4,2,2,2,2$
$\{5,10,15\}\star$	15	16	167/496	$6,6,4,4,4,4$
$\{6,9,15\}\star$	15	16	167/496	$6,6,4,4,4,4$
$\{7,9,14\}$	7	10	151/496	$4,4,4,4,4,2,2,2,2$
$\{7,10,13\}$	7	10	151/496	$4,4,4,4,4,2,2,2,2$
$\{7,11,12\}$	7	10	151/496	$4,4,4,4,4,2,2,2,2$
$\{7,17,22\}$	7	10	167/496	$6,6,4,4,4,3,3$
$\{7,18,21\}$	7	10	167/496	$6,6,4,4,4,3,3$
$\{7,19,20\}$	7	10	167/496	$6,6,4,4,4,3,3$
$\{7,25,30\}$	7	10	167/496	$6,6,4,4,4,3,3$
$\{7,26,29\}$	7	10	167/496	$6,6,4,4,4,3,3$
$\{7,27,28\}$	7	10	167/496	$6,6,4,4,4,3,3$
$\{15,17,30\}$	15	8	139/496	$8,8,6,6$
$\{15,18,29\}$	15	8	139/496	$8,8,6,6$
$\{15,19,28\}$	15	16	175/496	$6,6,5,5,4,4$
$\{15,20,27\}$	15	8	139/496	$8,8,6,6$
$\{15,21,26\}$	15	16	175/496	$6,6,5,5,4,4$
$\{15,22,25\}$	15	16	175/496	$6,6,5,5,4,4$
$\{15,23,24\}$	15	8	139/496	$8,8,6,6$

The two starred planes are the pair discussed in the body — one twist in two coordinate systems. They are the only two that emit 16 comparators **and** bisect every class evenly. Three planes in the rung-15 family below them — $\{15,19,28\}$, $\{15,21,26\}$, $\{15,22,25\}$ — also emit 16 and reach a **better** poset, $175/496$, but leave the odd profile $[6,6,5,5,4,4]$. The rung-3 and rung-7 families emit 10, and four rung-15 planes emit only 8.

The pattern by rung is stark. **Rung 15 gives large layers and even bisections; rungs 3 and 7 give small layers and fragmentation.** Of the 27, only the two starred rows produce the profile $[6,6,4,4,4,4]$ that the best known network uses.

Appendix III: the Adjacency Ladder

The definition

Ladder rung. A direction $d \in \mathbb{F}_2^k$ is a **ladder rung** if $d = 2^j - 1$ for some $j \geq 1$ — a block of trailing ones, $\underbrace{1 \cdots 1}_{j}$.

The rungs at $n = 32$ are $1, 3, 7, 15, 31$. The name is justified by two facts, one from the integers and one from the literature.

Adjacency. A direction d carries an adjacent pair $(x, x{+}1)$ if and only if d is a ladder rung, and rung $2^j - 1$ carries exactly 2^{k-j} of them.

This holds because $x \oplus (x{+}1)$ is always a block of trailing ones — incrementing flips a run of ones and the zero above them. **It is a fact about the integers, true at every n, not only at powers of two.**

Codish's last-layer lemma (Codish, Cruz-Filipe and Schneider-Kamp, **Sorting Networks: the End Game**, arXiv:1411.6408, Lemma 3). In a **non-redundant** sorting network, every comparator in the last layer is of the form $(i, i{+}1)$.

Together: **only ladder rungs can appear in a final layer.** At $n = 32$ that is 5 of the 31 nonzero directions; at $n = 64$, 6 of 63.

n	rungs	adjacent pairs each
8	$1, 3, 7$	$4, 2, 1$
16	$1, 3, 7, 15$	$8, 4, 2, 1$
32	$1, 3, 7, 15, 31$	$16, 8, 4, 2, 1$
64	$1, 3, 7, 15, 31, 63$	$32, 16, 8, 4, 2, 1$

The counts halve down the ladder, ending at the antipodal $n{-}1$ with a single adjacent pair: the **centre inversion** $(n/2{-}1,\ n/2)$.

The closure of this network

L_{14} carries 13 comparators, every one adjacent, distributed across the rungs:

rung	taken	comparators
$d{=}1$	0 of 16	—
$d{=}3$	6 of 8	`(5,6)(9,10)(13,14)(17,18)(21,22)(25,26)`
$d{=}7$	4 of 4	`(3,4)(11,12)(19,20)(27,28)`
$d{=}15$	2 of 2	`(7,8)(23,24)`
$d{=}31$	1 of 1	`(15,16)`

Rungs 7, 15 and 31 are taken completely. Rung 3 takes six of its eight — the other two were settled earlier. Rung 1 contributes nothing: by L_{14} its pairs are already comparable.

The closure is therefore a **descent through the top of the ladder**, each rung contributing every pair still live. That is the structure `precis-bundle-basis` derives from remnant geometry, met here in the artefact: at $n = 8$ the closure is $[3, 7]$, at $n = 16$ it is $[3, 7, 15]$, and here $[3, 7, 15, 31]$.

Why the antipodal must bundle

The antipodal $d = n{-}1$ has exactly one adjacent pair. A final layer containing it and nothing else would be a layer for a single comparator, wasting depth. So the antipodal **rides out bundled** with the rungs below it — and Codish's lemma says its partners can only be rungs, since nothing else carries an adjacent pair at all.

Two independent derivations meet here. `precis-antipodal-bundle` reaches the same conclusion from **remnant geometry**: by the time the antipodal is played, transitivity has settled all but its centre pair. Codish reaches it from **non-redundancy**. The ladder is where the two arguments coincide.

The search consequence Codish drew

The lemma was proved to prune a SAT search. Their Theorem 5 counts the possible last layers as $F_{n+1} - 1$:

$$n = 8\colon\ 54 \qquad n = 16\colon\ 2\,583 \qquad n = 32\colon\ 5\,702\,886$$

against $\binom{n}{2}$-many comparators to choose from in the general case. Read in the direction algebra the same lemma says something structural rather than combinatorial: **the tail of a network must descend the suffix-ones ladder**, and the antipodal sits at its foot with one admissible comparator.

Bibliography

The upper bounds we measured against

[Dob] Bert Dobbelaere. **List of sorting networks.**
`https://bertdobbelaere.github.io/sorting_networks_extended.html`
Page updated 7 November 2025; consulted 10 August 2026. The best known (size, depth) pairs for $n \leq 64$, with provenance for each. **The $n = 32$ entry $(185, 14)$ is attributed to "Batcher odd-even merge", not to the author's SorterHunter** — a distinction that corrected our reading of the network's structure. The changelog records "2020-06-27: Small improvement for networks of 29...32 inputs and 14 layers". The same page gives the lower bound 147 for $n = 32$, tightened 2025-04-21 on a suggestion of Jelmer Firet following the principles of [VV72]. **The interval $147 \ldots 185$ is open.**

[SH] Bert Dobbelaere. **SorterHunter** — an evolutionary approach to find small and low latency sorting networks. `https://github.com/bertdobbelaere/SorterHunter`
The program behind most of the entries on that list. It is **not** the source of the $n = 32$ network, which we had wrongly assumed.

[Bat68] K. E. Batcher. Sorting networks and their applications. **AFIPS Spring Joint Computer Conference**, 1968. The odd-even merge. At $n = 32$ it gives $191/15$; we show in the body that it is a **dimensional program** throughout, with ten of its fifteen layers using a single direction.

The renumbering

[Knu98] Donald E. Knuth. **The Art of Computer Programming**, volume 3: **Sorting and Searching**, 2nd edition. Addison-Wesley, 1998. Section 5.3.4, **Exercise 16** — the untangling of a generalized comparator network into a standard one of the same size and depth. This is the procedure that repairs our transported network: relabelling the wires by a linear map on $\mathbb{F}_2^5$ carries one twist to the other but is not monotone, so the result sorts into a permuted order until untangled. Codish et al. [CCS15] cite the same exercise for the same purpose.

The last-layer theory

[CCS15] Michael Codish, Luís Cruz-Filipe and Peter Schneider-Kamp. **Sorting Networks: the End Game.** arXiv:1411.6408; LATA 2015, LNCS 8977. **Lemma 3**: in a non-redundant sorting network every comparator in the last layer is of the form $(i, i+1)$. **Theorem 5**: the number of possible last layers is $F_{n+1} - 1$ — $5\,702\,886$ at $n = 32$. Definition 6 introduces k-blocks and Lemma 9 extends the constraint back one layer. Appendix III reads Lemma 3 in the direction algebra: only the **suffix-ones** directions $2^j - 1$ can end a network.

[CCEMS16] Michael Codish, Luís Cruz-Filipe, Thorsten Ehlers, Mike Müller and Peter Schneider-Kamp. Sorting networks: to the end and back again. **Journal of Computer and System Sciences**, 2016; `doi:10.1016/j.jcss.2016.04.004`. Preprint arXiv:1507.01428. The journal version, sometimes cited as 2017 for its issue date. **There is no separate 2017 arXiv paper in this line** — we checked the arXiv API, which lists seven Codish sorting-network papers, the latest from 2015.

[CCFS16] Michael Codish, Luís Cruz-Filipe, Michael Frank and Peter Schneider-Kamp. Sorting nine inputs requires twenty-five comparisons. **Journal of Computer and System Sciences** 82(3), 2016. Preprint arXiv:1405.5754. Size-optimality at $n = 9$ and $n = 10$, by symmetry-breaking combined with SAT.

Optimality results and bounds

[BZ14] Daniel Bundala and Jakub Závodný. Optimal sorting networks. **LATA 2014**, LNCS 8370, Springer. Depth-optimality for $n \leq 16$.

[Par89] Ian Parberry. A computer-assisted optimal depth lower bound for sorting networks with nine inputs. **Supercomputing '89**.

[Har19] Jannis Harder. **sortnetopt**. `https://github.com/jix/sortnetopt`
Size-optimality at $n = 11$ and $n = 12$. **Exhaustive size-optimality is known only to** $n = 12$; everything above, including all of $n = 32$, is bounds.

[VV72] D. C. Van Voorhis. Toward a lower bound for sorting networks. **IBM Symposium on Complexity of Computer Computations**, 1972. The counting argument behind the lower bounds on [Dob]'s table.

[VM13] Vinod K. Valsalam and Risto Miikkulainen. Using symmetry and evolutionary search to minimize sorting networks. **Journal of Machine Learning Research** 14, 2013.

[TAOCP] Donald E. Knuth. **The Art of Computer Programming**, volume 3, 2nd edition, 1998. Section 5.3.4, **Networks for Sorting** — the classical optima for $n \leq 8$, and the 0-1 principle we verify against.

Held locally

`references/` contains [CCS15], [CCEMS16], [CCFS16], [BZ14], and Goodrich's **Zig-zag Sort** (arXiv:1403.2777). The Dobbelaere list is a web page and is cited by URL and consultation date.