

# Searching for optimal Boolean chains

Adam P. Goucher

February 28, 2023

# Boolean chains

A **Boolean chain** is a sequence of 2-input Boolean gates.

# Boolean chains

A **Boolean chain** is a sequence of 2-input Boolean gates.

For example, the full adder has  $n = 5$  gates and  $k = 3$  inputs:


1  $y_1 = x_1 \oplus x_2;$

2  $y_2 = x_1 \wedge x_2;$

3  $y_3 = y_1 \oplus x_3;$

4  $y_4 = y_1 \wedge x_3;$

5  $y_5 = y_2 \vee y_4.$



Each gate can only depend on inputs or previously-computed values.

# Five normal gates

Without loss of generality we can assume that all gates  $\circ$  are:

- **Nontrivial:**  $a \circ b$  depends on both  $a$  and  $b$ ;
- **Zero-preserving:**  $0 \circ 0 = 0$ .

Knuth (2011) calls these **normal** chains.

# Five normal gates

Without loss of generality we can assume that all gates  $\circ$  are:

- **Nontrivial:**  $a \circ b$  depends on both  $a$  and  $b$ ;
- **Zero-preserving:**  $0 \circ 0 = 0$ .

Knuth (2011) calls these **normal** chains.

Out of the  $2^{2^2} = 16$  functions, 8 are zero-preserving, of which 5 are nontrivial:

$$\mathcal{O} = \{\oplus, \wedge, \vee, <, >\}$$

# Five normal gates

Without loss of generality we can assume that all gates  $\circ$  are:

- **Nontrivial:**  $a \circ b$  depends on both  $a$  and  $b$ ;
- **Zero-preserving:**  $0 \circ 0 = 0$ .

Knuth (2011) calls these **normal** chains.


Out of the  $2^{2^2} = 16$  functions, 8 are zero-preserving, of which 5 are nontrivial:

$$\mathcal{O} = \{\oplus, \wedge, \vee, <, >\}$$

These correspond to AVX instructions **vpxor**, **vpand**, **vpor**, **vpandn**.

# Rewriting

Many logic synthesis tools (e.g. Berkeley's ABC) work by **local rewriting**: replacing small subcircuits with more efficient equivalents.



## Prior work

**Berkeley's ABC** finds 4-input cuts and optimally rewrites those.



## Prior work


**Berkeley's ABC** finds 4-input cuts and optimally rewrites those.

**Nan Li and Elena Dubrova (2011)** found significant benefits (5% cost reduction) by using a library of **1200** 5-input functions.

## Prior work

**Berkeley's ABC** finds 4-input cuts and optimally rewrites those.

**Nan Li and Elena Dubrova (2011)** found significant benefits (5% cost reduction) by using a library of **1200** 5-input functions.



We shall take this to its ultimate logical conclusion: finding all optimal chains for **616125** of the 616126 equivalence classes of 5-input functions.

# Equivalence classes

We can transform a  $k$ -input  $\ell$ -output function into an equivalent function by:

- Permuting inputs ( $k!$  possibilities);
- Negating inputs ( $2^k$  possibilities);
- Permuting outputs ( $\ell!$  possibilities);
- Negating outputs ( $2^\ell$  possibilities);

which generate the group  $(S_2 \wr S_k) \times (S_2 \wr S_\ell)$  of order  $2^{k+\ell}(k!)(\ell!)$ .

# Equivalence classes

We can transform a  $k$ -input  $\ell$ -output function into an equivalent function by:


- Permuting inputs ( $k!$  possibilities);
- Negating inputs ( $2^k$  possibilities);
- Permuting outputs ( $\ell!$  possibilities);
- Negating outputs ( $2^\ell$  possibilities);

which generate the group  $(S_2 \wr S_k) \times (S_2 \wr S_\ell)$  of order  $2^{k+\ell}(k!)(\ell!)$ .

$$0001 < 0010 < 0100 < 0111 < 1000 < 1011 < 1101 < 1110$$

We describe the lexicographically first truth table in an equivalence class as **canonical**.

## Counting equivalence classes



We can count the equivalence classes using **Burnside's lemma**; the dominant term will be:

$$|C| \approx \frac{2^{2^{k\ell}}}{2^{k+\ell}(k!)(\ell!)}$$

## Number of classes of functions of each cost

$n$	5-input 1-output	4-input 2-output
0	2	4
1	2	8
2	5	38
3	20	193
4	93	916
5	389	4869
6	1988	27219
7	11382	135402
8	60713	475926
9	221541	713796
10	293455	117828
11	26535	19
12	1	0
<b>Total</b>	<b>616126</b>	<b>1476218</b>

# Size of search space

With  $k$  inputs and  $n$  gates, the total number of Boolean chains is:

$$\prod_{i=0}^{n-1} 5 \binom{i+k}{2}$$

## Size of search space

With  $k$  inputs and  $n$  gates, the total number of Boolean chains is:

$$\prod_{i=0}^{n-1} 5 \binom{i+k}{2}$$

In particular, for 5 inputs and 11 gates, the number of chains is:

$$18874939423183593750000000 \approx 1.89 \times 10^{25}$$



## Size of search space

With  $k$  inputs and  $n$  gates, the total number of Boolean chains is:

$$\prod_{i=0}^{n-1} 5 \binom{i+k}{2}$$


In particular, for 5 inputs and 11 gates, the number of chains is:

$$18874939423183593750000000 \approx 1.89 \times 10^{25}$$

This is about **60x more** than the number of floating-point operations used to train GPT-3 ( $3.14 \times 10^{23}$ ).

## Room for improvement I: canonical ordering


In many cases, a single DAG can correspond to multiple Boolean chains:




For example, the order of computing variables 5 and 6 could be swapped.

## Room for improvement II: triangle-free


Also, triangles can be removed from our DAG without loss of generality:



# Key insight



## Key insight



Simply represent this as the set  $\{0x0ff0, 0x6666, 0x6996\}$ .

# Reconstruction of chains from tt-sets

By performing a **backtracking search**, we can reconstruct all possible chains corresponding to a tt-set.

- Beginning with the set of truth tables corresponding to inputs, try applying all possible gates which result in a truth table belonging to the tt-set.

# Reconstruction of chains from tt-sets

By performing a **backtracking search**, we can reconstruct all possible chains corresponding to a tt-set.

- Beginning with the set of truth tables corresponding to inputs, try applying all possible gates which result in a truth table belonging to the tt-set.
- We enforce the 'canonical ordering' and 'triangle free' properties at all times.

# Reconstruction of chains from tt-sets

By performing a **backtracking search**, we can reconstruct all possible chains corresponding to a tt-set.

- Beginning with the set of truth tables corresponding to inputs, try applying all possible gates which result in a truth table belonging to the tt-set.
- We enforce the 'canonical ordering' and 'triangle free' properties at all times.
- The total runtime is  $O(Sn^3)$  where  $S$  is the number of chains outputted by the algorithm.




## Ends of tt-sets

A tt-set  $T$  is *attainable* if it corresponds to at least one chain.

## Ends of tt-sets


A tt-set  $T$  is *attainable* if it corresponds to at least one chain.



If  $T \setminus \{e\}$  is still attainable, then we say that  $e$  is an *end* of  $T$ .


## Appending a gate

Appending a gate cannot cause the number of ends to decrease by more than 1.



# Symmetries

Recall that we have a large symmetry group at play.



The search becomes much faster if we only store **canonical** tt-sets.

# Breadth-first search

Define  $N(T)$  to be the set of possible 'next functions':

$$N(T) = \{x \circ y : x, y \in T \cup I \text{ and } \circ \in \mathcal{O}\} \setminus (T \cup I \cup \{0\})$$

# Breadth-first search

Define  $N(T)$  to be the set of possible 'next functions':

$$N(T) = \{x \circ y : x, y \in T \cup I \text{ and } \circ \in \mathcal{O}\} \setminus (T \cup I \cup \{0\})$$

We construct all canonical attainable tt-sets by induction on size:

- Define  $X_0 = \{\{\}\}$  to be the set containing the empty tt-set.
- Define
$$X_{n+1} = \{\text{canonicalise}(T \cup \{e\}) : T \in X_n \text{ and } e \in N(T)\}$$

# Breadth-first search

Define  $N(T)$  to be the set of possible 'next functions':

$$N(T) = \{x \circ y : x, y \in T \cup I \text{ and } \circ \in \mathcal{O}\} \setminus (T \cup I \cup \{0\})$$

We construct all canonical attainable tt-sets by induction on size:

- Define  $X_0 = \{\{\}\}$  to be the set containing the empty tt-set.
- Define
$$X_{n+1} = \{\text{canonicalise}(T \cup \{e\}) : T \in X_n \text{ and } e \in N(T)\}$$

For a given search depth, we can also discard tt-sets with too many ends.

# How does previous work differ?

Previous approaches directly searched the (much larger) space of DAGs rather than canonical attainable tt-sets.



# How does previous work differ?

Previous approaches directly searched the (much larger) space of DAGs rather than canonical attainable tt-sets.

- **Donald Knuth (2011)** used 'top-down' and 'bottom-up' reductions together with brute-force searches for 'special' (irreducible) DAGs.
- Knuth was interested in determining the minimum cost for each function, rather than finding all optimal chains.

# How does previous work differ?

Previous approaches directly searched the (much larger) space of DAGs rather than canonical attainable tt-sets.


- **Donald Knuth (2011)** used 'top-down' and 'bottom-up' reductions together with brute-force searches for 'special' (irreducible) DAGs.
- Knuth was interested in determining the minimum cost for each function, rather than finding all optimal chains.

**Haaswijk, Soeken, Mishchenko, and De Micheli (2018)** used SAT solvers to search for DAGs of a particular topology implementing a given function.

# Reducing constant factors


Canonicalising a tt-set is expensive. To do an 11-gate search, we do the following:

- Use the algorithm to compute all tt-sets in  $X_9$  with  $\leq 3$  ends.
- Brute-force all possibilities for the last two gates.




## Checking optimality

We want to check optimality, i.e. determine whether the purple vertex cannot be attained with fewer gates.



## Checking optimality

We want to check optimality, i.e. determine whether the purple vertex cannot be attained with fewer gates.



We use a 256 MB array (one bit for each of the  $2^{31}$  normal functions) to indicate whether they've previously been constructed with lower cost.

# Practical problems

Even a depth-9 search does not fit in memory (11 terabytes uncompressed).

$n$	tt-sets	uncompressed size	compressed size
6	3 million	73 MB	5 MB
7	130 million	3.6 GB	159 MB
8	6.2 billion	197 GB	6.4 GB
9	(310 billion)	(11 TB)	

## Practical problems

Even a depth-9 search does not fit in memory (11 terabytes uncompressed).

$n$	tt-sets	uncompressed size	compressed size
6	3 million	73 MB	5 MB
7	130 million	3.6 GB	159 MB
8	6.2 billion	197 GB	6.4 GB
9	(310 billion)	(11 TB)	

We need to somehow **partition the search space** into manageable chunks, but this is not particularly easy.

# Invariants and signatures

An **invariant** is an easily-computable function  $f$  from truth tables to an arbitrary finite set which is constant on equivalence classes.



# Invariants and signatures

An **invariant** is an easily-computable function  $f$  from truth tables to an arbitrary finite set which is constant on equivalence classes.

For example, the **bias**, or absolute difference between the number of '1's and '0's in the truth table, is an invariant.



# Invariants and signatures

An **invariant** is an easily-computable function  $f$  from truth tables to an arbitrary finite set which is constant on equivalence classes.

For example, the **bias**, or absolute difference between the number of '1's and '0's in the truth table, is an invariant.



Given an invariant  $f$ , the **signature** of a tt-set  $T$  is the **multiset**  $\{f(x) : x \in T\}$ .

# Distributing the workload

Proceed as before, but partition the creation of each  $X_n$  into separate tasks, one for each size- $n$  signature.

# Distributing the workload

Proceed as before, but partition the creation of each  $X_n$  into separate tasks, one for each size- $n$  signature.

The task corresponding to a size- $n$  signature  $S$  reads the files saved behind by the tasks corresponding to size- $(n - 1)$  signatures  $S \setminus \{s\}$  for each  $s \in S$ .

# Distributing the workload

Proceed as before, but partition the creation of each  $X_n$  into separate tasks, one for each size- $n$  signature.

The task corresponding to a size- $n$  signature  $S$  reads the files saved behind by the tasks corresponding to size- $(n - 1)$  signatures  $S \setminus \{s\}$  for each  $s \in S$ .

We avoid saving tt-sets to disk in the last level of the search tree ( $n = 9$ ), because there are no downstream tasks to consume them.

# Results

We applied this depth-11 exhaustive search procedure to two search problems:

- 5-input 1-output functions (616126 equivalence classes);
- 4-input 2-output functions (1476218 equivalence classes).

Each of the two searches took a few days on an AWS **r5a.24xlarge** instance (96 virtual cores + 768 GB memory), costing  $< \$1000$ .

# Results

We applied this depth-11 exhaustive search procedure to two search problems:

- 5-input 1-output functions (616126 equivalence classes);
- 4-input 2-output functions (1476218 equivalence classes).

Each of the two searches took a few days on an AWS **r5a.24xlarge** instance (96 virtual cores + 768 GB memory), costing < \$1000.

But how do we use these results?

# Building a database

For each equivalence class of functions, we took all Boolean chains with **minimum cost and delay on the Pareto frontier**.

00f5c400	21 30 43 70 81 84 96 a7	b5 dc 00 b0 6f 66 e7 01	!0Cp.....011..
00f5c4c0	21 30 43 70 81 84 96 a3	b5 dc 00 b0 6f 66 e7 01	!0Cp.....of...
00f5c4d0	21 30 43 70 81 84 96 a0	b5 dc 00 b0 6f 7e e7 01	!0Cp.....o~...
00f5c4e0	21 30 43 70 82 84 96 a7	b5 dc 00 b0 6f 6e e7 01	!0Cp.....on...
00f5c4f0	21 30 43 70 82 84 96 a3	b5 dc 00 b0 6f 66 e7 01	!0Cp.....of...
00f5c500	21 30 43 70 82 84 96 a0	b5 dc 00 b0 6f 7e e7 01	!0Cp.....o~...
00f5c510	00 00 00 00 00 00 00 00	00 00 00 00 00 00 00 00	.....
*			
00f5c540	c8 f3 19 00 58 00 0b 00	04 30 44 33 01 50 33 45	....X...0D3.P3E
00f5c550	01 50 55 52 03 40 66 62	02 80 62 78 00 00 00 00	..PUR.@fb..bx...
00f5c560	21 21 50 54 63 97 98 ba	00 00 00 e0 af f4 0c 00	!!PTc.....
00f5c570	10 21 21 64 73 95 98 ba	00 00 00 30 bf e4 0c 00	!!ds.....0....
00f5c580	21 21 50 54 63 87 98 ba	00 00 00 e0 af fc 08 00	!!PTc.....
00f5c590	10 21 21 64 73 85 98 ba	00 00 00 30 bf ec 0c 00	!!ds.....0....
00f5c5a0	30 43 54 61 72 85 91 ba	00 00 00 e0 ad fe 07 00	@CTar.....
00f5c5b0	21 40 50 54 61 97 a3 b8	00 00 00 e0 bb 04 07 00	!@PTa.....
00f5c5c0	21 41 65 70 74 81 a3 b9	00 00 00 70 ca 1d 07 00	!Aept.....p....
00f5c5d0	21 41 54 65 80 91 a3 b7	00 00 00 e0 16 1e 07 00	!Ate.....
00f5c5e0	21 41 65 70 74 81 a3 b9	00 00 00 e0 c2 1d 07 00	!Aept.....
00f5c5f0	30 54 62 74 86 93 a1 b8	00 00 00 e0 71 1d 07 00	@Tbt.....q....
00f5c600	30 54 62 75 86 93 a1 b8	00 00 00 e0 79 1d 07 00	@Tbu.....y....
00f5c610	00 00 00 00 00 00 00 00	00 00 00 00 00 00 00 00	.....
*			
00f5c640	e6 d5 1b 00 09 01 24 00	02 40 55 24 01 50 54 34	.....\$.@US.PT4
00f5c650	02 30 55 34 01 50 44 35	02 60 63 35 09 50 66 42	..@U4.PD5..c5.PfB
00f5c660	01 40 43 44 02 30 44 44	01 20 55 44 02 50 44 52	..@CD.0DD. UD.PDR
00f5c670	01 40 66 52 06 30 55 53	03 50 34 54 01 50 24 66	..@fR.0US.P4T.P5f
00f5c680	01 50 33 66 01 50 52 66	00 00 00 00 00 00 00 00	..P3f.PfR.....
00f5c690	21 21 31 53 60 97 98 b4	ca 00 00 b0 93 9e 3f 00	!!!S`.....?.
00f5c6a0	21 21 31 53 60 97 98 b4	ca 00 00 c0 b3 9e 3f 00	!!!S`.....?.
00f5c6b0	10 20 43 63 65 81 97 a4	cb 00 00 40 7f 28 3b 00	..Cce.....@.(;
00f5c6c0	21 21 31 53 60 74 84 b9	ca 00 00 b0 93 de 39 00	!!!S`t.....9..
00f5c6d0	21 21 31 53 60 74 84 b9	ca 00 00 c0 b3 de 39 00	!!!S`t.....9..
00f5c6e0	20 30 43 51 61 84 95 b7	ca 00 00 e0 9d ed 38 00	..@CQa.....8..
00f5c6f0	20 43 53 65 70 74 91 ba	c8 00 00 e0 ee 7b 3b 00	..CSept.....{ ..
00f5c700	20 53 60 64 64 71 85 a9	cb 00 00 e0 cd ef 3b 00	..S`ddq.....;..
00f5c710	20 21 60 74 74 85 91 b3	ca 00 00 e0 b9 f6 38 00	..!`tt.....8..
00f5c720	10 21 60 74 74 85 91 b3	ca 00 00 e0 b9 f6 38 00	..!`tt.....8..



# The k-perfect hashtable

We hash a canonical 5-input 1-output truth table using the following pair of functions:

```
auto h1 = (x ^ (x >> 7) ^ (x >> 4)) & 0x1ffff;  
auto h2 = (x ^ (x >> 7) ^ (x >> 24)) & 0x1ff;
```

# The $k$ -perfect hashtable

We hash a canonical 5-input 1-output truth table using the following pair of functions:

```
auto h1 = (x ^ (x >> 7) ^ (x >> 4)) & 0x1ffff;  
auto h2 = (x ^ (x >> 7) ^ (x >> 24)) & 0x1ff;
```

The function  $h_1$  is 15-perfect, mapping at most 15 of the 616124 canonical nontrivial truth tables to the same bucket.

The function  $h_2$  is 1-perfect (injective) within each bucket.

## Looking up a canonical truth table

By squeezing each bucket into 64 bytes, it is possible to lookup the cost and location of any canonical function by retrieving a single cache line:

```
uint32_t get_cl_and_cost(const uint32_t* db, uint32_t x, int n_bits) {  
    uint32_t clc = 0;  
    uint32_t hash1 = hh::boolchains::hthash1(x, n_bits);  
    uint32_t hash2 = hh::boolchains::hthash2(x);  
  
    uint32_t cache_line[16];  
  
    memcpy(cache_line, &(db[hash1 << 4]), 64);  
  
    for (size_t i = 0; i < 15; i++) {  
        uint32_t element = cache_line[i];  
        if ((element & 0x1fff) == hash2) {  
            uint32_t cl = ((element >> 9) & 0x7fff) + cache_line[15];  
            uint32_t cost = (element >> 24);  
            clc = (cl << 4) | cost;  
            break;  
        }  
    }  
  
    return clc;  
}
```

# Overview

To find all chains for a (not necessarily canonical) function:

- **Canonicalize it**: this would naively take  $10 \mu\text{s}$ , but **we** obsessively reduced it down to **200 ns**.

# Overview

To find all chains for a (not necessarily canonical) function:

- **Canonicalize it:** this would naively take  $10\ \mu\text{s}$ , but **we obsessively reduced it** down to **200 ns**.
- **Lookup the cost and location:** this is a single cache line retrieval plus a handful of instructions, so takes around **100 ns**.

# Overview

To find all chains for a (not necessarily canonical) function:

- **Canonicalize it:** this would naively take  $10 \mu\text{s}$ , but **we obsessively reduced it** down to **200 ns**.
- **Lookup the cost and location:** this is a single cache line retrieval plus a handful of instructions, so takes around **100 ns**.
- **Iterate over optimal chains:** this step takes variable time, depending on the number of chains, but they're contiguous in memory so this is again reasonably fast.