

Two-Layer Filters for Sorting Networks

L. Cruz-Filipe^{1,2} M. Codish³ P. Schneider-Kamp¹

¹Dept. Mathematics and Computer Science, Univ. Southern Denmark (Denmark)

²LabMAg (Portugal)

³Ben-Gurion University of the Negev (Israel)

LabMAg Seminar
April 16th, 2014

Outline

- 1 Sorting Networks in a Nutshell
- 2 Reduction Techniques
- 3 A Symbolical Approach
- 4 Conclusions & Future Work


What are sorting networks?

- Oblivious algorithms to sort a given number of outputs
- Easy to implement at the hardware level
- Intrinsically parallel
- Two interesting optimizations:
 - size (production cost)
 - depth (execution time)


What are sorting networks?

- Oblivious algorithms to sort a given number of outputs
- Easy to implement at the hardware level
- Intrinsically parallel
- Two interesting optimizations:
 - size (production cost)
 - depth (execution time)
- See Donald E. Knuth, *The Art of Computer Programming*, vol. 3 for more details

A sorting network




A sorting network



Size and depth

This net has 10 *wires*, 29 *comparators* and 9 *layers*.


A sorting network



Representation

$(5,10);(4,9);(3,8);(2,7);(1,6);(2,5);(7,10);(1,4);(6,9);(1,3);(4,7); \dots$

A sorting network



Representation

```
[(5,10);(4,9);(3,8);(2,7);(1,6)];[(2,5);(7,10);(1,4);(6,9)];[(1,3);(4,7);...];...
```


The optimization problems

The size problem

What is the minimal number of *comparators* on a sorting network on n channels (S_n)?

The depth problem

What is the minimal number of *layers* on a sorting network on n channels (T_n)?

The optimization problems

The size problem

What is the minimal number of *comparators* on a sorting network on n channels (S_n)?

The depth problem

What is the minimal number of *layers* on a sorting network on n channels (T_n)?

n	3	4	5	6	7	8	9	10	11	12
S_n	3	5	9	12	16	19	≤ 25	≤ 29	≤ 35	≤ 39
T_n	3	3	5	5	6	6	7*	7*	8*	8*

Knuth 1973; *Parberry 1991

The optimization problems

The size problem

What is the minimal number of *comparators* on a sorting network on n channels (S_n)?

The depth problem

What is the minimal number of *layers* on a sorting network on n channels (T_n)?

n	13	14	15	16
S_n	≤ 45	≤ 51	≤ 56	≤ 60
T_n	9^\dagger	9^\dagger	9^\dagger	9^\dagger

Knuth 1973; [†]Bundala & Závodný 2014

An exponential explosion

- Upper bounds obtained by concrete examples (1960s)
- Lower bounds obtained by mathematical arguments
- HUGE number of nets

An exponential explosion

- Upper bounds obtained by concrete examples (1960s)
- Lower bounds obtained by mathematical arguments
- HUGE number of nets
- Parberry (1991)
 - exploration of symmetries
 - fixed first layer
 - 200 hours of computation

An exponential explosion

- Upper bounds obtained by concrete examples (1960s)
- Lower bounds obtained by mathematical arguments
- HUGE number of nets
- Parberry (1991)
- Bundala & Závodný (2014)
 - exploration of symmetries
 - reduced set of two-layer prefixes
 - intensive SAT-solving

An exponential explosion

- Upper bounds obtained by concrete examples (1960s)
- Lower bounds obtained by mathematical arguments
- HUGE number of nets
- Parberry (1991)
- Bundala & Závodný (2014)
- These techniques do not scale for T_{17}
 $\approx 211 \times 10^6$ possibilities for each layer when $n = 17$
- These techniques are not directly applicable to the size problem
36 possibilities for each layer when $n = 9$, so
 $36^{24} \approx 2.2 \times 10^{37}$ 24-comparator nets

Outline

- 1 Sorting Networks in a Nutshell
- 2 Reduction Techniques**
- 3 A Symbolical Approach
- 4 Conclusions & Future Work

Comparator networks

- A (*generalized*) *comparator network* C on n wires is a sequence of pairs (i, j) (the comparators) such that $1 \leq i \neq j \leq n$.
- A *standard* comparator network C is a generalized comparator network such that $i < j$ for every comparator $(i, j) \in C$.
- The *output* of comparator (i, j) on $\vec{x} = x_1 \dots x_n$ is \vec{x} , if $x_i \leq x_j$, and \vec{x}' where $x'_i = x_j$, $x'_j = x_i$ and $x'_k = x_k$ for $k \neq i, j$.
- The *output* of C on a sequence $x_1 \dots x_n$, is defined inductively:
 - if C is empty, then $C(x_1 \dots x_n) = x_1 \dots x_n$;
 - if C is $(i, j); C'$ then $C(x_1 \dots x_n) = C'((i, j)(x_1 \dots x_n))$.
- A comparator network C is a *sorting network* if $C(\vec{x})$ is sorted for every input \vec{x} .

Well-known results (Knuth 1973)

0-1 lemma

C is a sorting network on n channels iff C sorts all inputs in $\{0, 1\}^n$.

Proof

The direct implication is straightforward. For the converse, consider the sequences $0^k 1^{n-k}$ and $0^{k+1} 1^{n-k-1}$; if C sorts both, then it must always place the k -th smallest element of its input in the right place. The result follows by induction on k .

Well-known results (Knuth 1973)

0–1 lemma

C is a sorting network on n channels iff C sorts all inputs in $\{0,1\}^n$.

Standardization theorem

If C is a generalized sorting network, then there is a standard sorting network C' with the same size and depth as C .

Proof

Reverse the first comparator (i,j) such that $i > j$ and interchange i and j on every subsequent comparator.
Repeat until the network is standard.

Well-known results (Knuth 1973)

0–1 lemma

C is a sorting network on n channels iff C sorts all inputs in $\{0, 1\}^n$.

Standardization theorem

If C is a generalized sorting network, then there is a standard sorting network C' with the same size and depth as C .

We will only consider binary inputs and use generalized comparator networks whenever needed.

Restricting prefixes (Parberry 1991)

Output lemma

Let C and C' be comparator networks such that $\text{outputs}(C) \subseteq \text{outputs}(C')$. If $C'; N$ is a sorting network, then so is $C; N$.

Restricting prefixes (Parberry 1991)

Output lemma

Let C and C' be comparator networks such that $\text{outputs}(C) \subseteq \text{outputs}(C')$. If $C'; N$ is a sorting network, then so is $C; N$.

Corollary

There is a minimal-depth sorting network on n channels whose first layer contains $\lfloor \frac{n}{2} \rfloor$ comparators.

Restricting prefixes (Parberry 1991)

Output lemma

Let C and C' be comparator networks such that $\text{outputs}(C) \subseteq \text{outputs}(C')$. If $C'; N$ is a sorting network, then so is $C; N$.


Corollary

There is a minimal-depth sorting network on n channels whose first layer contains $\lfloor \frac{n}{2} \rfloor$ comparators.


Normal form for first layer

There is a minimal-depth sorting network on n channels whose first layer F_n contains the comparators $(1, 2)$, $(3, 4)$, $(5, 6)$, &c.


Normalization




Normalization



Normalization



Normalization



Permutations (Bundala & Závodný 2014) [1]

Permutation lemma

- π is a permutation of $1..n$ preserving F_n ;
- the two-layer network $F_n; L$ can be extended to a sorting network C ;
- $\pi(L)$ is standard;

then $F_n; \pi(L)$ can be extended to a standard sorting network of the same size and depth as C .

Proof

Renumber the wires after layer 2 according to π . The resulting network will produce the same output for every input.

Apply the standardization theorem to obtain a standard comparator network producing the same output for every input. But standard networks do not change sorted inputs, so this is a sorting network.

Permutations (Bundala & Závodný 2014) [II]

Permuted output lemma

- C and C' are two-layer standard comparator networks;
- π is a permutation of $1..n$ mapping $\text{outputs}(C)$ into $\text{outputs}(C')$;
- C' can be extended to a sorting network;

then C can also be extended to a standard sorting network of the same depth.

Proof

Consequence of the above results.

Permutations (Bundala & Závodný 2014) [II]

Permuted output lemma

- C and C' are two-layer standard comparator networks;
- π is a permutation of $1..n$ mapping $\text{outputs}(C)$ into $\text{outputs}(C')$;
- C' can be extended to a sorting network;

then C can also be extended to a standard sorting network of the same depth.

Equivalence of networks

Two standard networks C and C' are equivalent if there is a permutation π of $1..n$ such that C' can be obtained from C by renumbering its wires according to π and standardizing the result.

Finding the value of T_{13}

Saturation

Saturation is a syntactic criterion for two-layer networks. Bundala & Závodný prove that it is enough to consider saturated networks.

Finding the value of T_{13}

Saturation

Saturation is a syntactic criterion for two-layer networks. Bundala & Závodný prove that it is enough to consider saturated networks.

Reflection

(Vertical) reflection of a comparator network produces a “dual” net: if input \vec{x} goes to \vec{y} , then $x^{\vec{D}}$ goes to $y^{\vec{D}}$. A two-layer network can be extended to a sorting network of depth d iff the same holds for its reflection.

Finding the value of T_{13}

The strategy

- 1 Generate all saturated two-layer networks with first layer F_{13} .
- 2 Remove equivalent nets.
- 3 Remove nets whose outputs are permuted supersets of others'.
- 4 Remove reflected nets.
- 5 Use a SAT-solver to find out if the remaining nets can be extended to a sorting network.

Finding the value of T_{13}

The strategy

- 1 Generate all saturated two-layer networks with first layer F_{13} .
- 2 Remove equivalent nets.
- 3 Remove nets whose outputs are permuted supersets of others'.
- 4 Remove reflected nets.
- 5 Use a SAT-solver to find out if the remaining nets can be extended to a sorting network.

n	5	6	7	8	9	10	11	12	13
$ G_n $	26	76	232	764	2620	9496	35696	140152	568504
$ S_n $	10	51	74	513	700	6345	8174	93255	113008
$ G_n/\approx $	18	28	74	101	295	350	1134	1236	4288
$ S_n/\approx $	8		29		100		341		1155
red.	6	6	14	15	37	27	88	70	212
$ R_n $	4	5	8	12	22	21	28	50	118

And for higher n ?

This approach does not scale.

Computing equivalence of nets is very expensive (and not working correctly).

Checking output subsumption is even worse (there are 2^n outputs, 2^{2^n} possible sets of outputs, and $n!$ permutations).

Furthermore, $T_{13} = T_{14} = T_{15} = T_{16}$.

To go beyond these values, we need different techniques.

Outline





- 1 Sorting Networks in a Nutshell
- 2 Reduction Techniques
- 3 A Symbolical Approach**
- 4 Conclusions & Future Work

The general idea

Comparator networks can be represented by labeled graphs (Choi & Moon 2002), where each comparator is a node and there is an edge from v_i to v_j labeled min (max) if the minimum (maximum) output of v_i is an input to v_j .

The general idea

Comparator networks can be represented by labeled graphs (Choi & Moon 2002), where each comparator is a node and there is an edge from v_i to v_j labeled min (max) if the minimum (maximum) output of v_i is an input to v_j .



The general idea

Comparator networks can be represented by labeled graphs (Choi & Moon 2002), where each comparator is a node and there is an edge from v_i to v_j labeled min (max) if the minimum (maximum) output of v_i is an input to v_j .

Equivalence theorem

$C \approx C'$ iff $\mathcal{G}_C \approx \mathcal{G}_{C'}$.

The general idea

Comparator networks can be represented by labeled graphs (Choi & Moon 2002), where each comparator is a node and there is an edge from v_i to v_j labeled min (max) if the minimum (maximum) output of v_i is an input to v_j .

Equivalence theorem

$$C \approx C' \text{ iff } \mathcal{G}_C \approx \mathcal{G}_{C'}.$$

But we want to bypass graphs altogether. Can we find a way to generate a set \mathcal{N} of nets such that:

- for every two-layer comparator network C with first layer F_n there is $C' \in \mathcal{N}$ such that $C \approx C'$;
- for every $C, C' \in \mathcal{N}$, $C \not\approx C'$;


in an efficient way?

Word representation for two-layer networks

Idea: represent two layer-networks by words, corresponding to paths in their graphs. (But avoid graphs altogether.)

Word representation for two-layer networks


Idea: represent two layer-networks by words, corresponding to paths in their graphs. (But avoid graphs altogether.)




Head word:
01221

Word representation for two-layer networks

Idea: represent two layer-networks by words, corresponding to paths in their graphs. (But avoid graphs altogether.)




Head word:
01221



Word representation for two-layer networks


Idea: represent two layer-networks by words, corresponding to paths in their graphs. (But avoid graphs altogether.)



Stick word:
21121212
21212112

Word representation for two-layer networks


Idea: represent two layer-networks by words, corresponding to paths in their graphs. (But avoid graphs altogether.)




Stick word:
21121212
21212112

Word representation for two-layer networks

Idea: represent two layer-networks by words, corresponding to paths in their graphs. (But avoid graphs altogether.)




Stick word:
21121212
 21212112



Word representation for two-layer networks

Idea: represent two layer-networks by words, corresponding to paths in their graphs. (But avoid graphs altogether.)



Cycle word:


121221

122121

122112

Word representation for two-layer networks

Idea: represent two layer-networks by words, corresponding to paths in their graphs. (But avoid graphs altogether.)



Cycle word:


121221

122121

122112

Word representation for two-layer networks

Idea: represent two layer-networks by words, corresponding to paths in their graphs. (But avoid graphs altogether.)



Cycle word:


121221

122121

122112

Word representation for two-layer networks

Idea: represent two layer-networks by words, corresponding to paths in their graphs. (But avoid graphs altogether.)




Cycle word:

121221

122121

122112



Word representation for two-layer networks

Idea: represent two layer-networks by words, corresponding to paths in their graphs. (But avoid graphs altogether.)

Every net generates a unique word, and every well-formed word generates a unique net. The functions net-to-word and word-to-net form an adjunction.

A regular language for words

Word ::= Head | Stick | Cycle

Head ::= $0(12 + 21)^*$

Stick ::= $(1 + 2)(12 + 21)^*(1 + 2)$

Cycle ::= $12(12 + 21)^*(1 + 2)$

A regular language for words

Word ::= Head | Stick | Cycle

Head ::= $0(12 + 21)^*$

Stick ::= $(1 + 2)(12 + 21)^*(1 + 2)$

Cycle ::= $12(12 + 21)^*(1 + 2)$

Generating all words and filtering to obtain only the lexicographically smallest is very easy for the relevant values of n .

A regular language for words

Word ::= Head | Stick | Cycle

Head ::= $0(12 + 21)^*$

Stick ::= $(1 + 2)(12 + 21)^*(1 + 2)$

Cycle ::= $12(12 + 21)^*(1 + 2)$

Generating all words and filtering to obtain only the lexicographically smallest is very easy for the relevant values of n . Two-layer comparator networks can be represented by multi-sets of words. By choosing a canonical representation of multi-sets, we can easily generate exactly one representative for all two-layer networks with first layer F_n modulo equivalence.

Saturation, revisited [1]

Saturation

Saturation is a syntactic criterion for two-layer networks. Bundala & Závodný prove that it is enough to consider saturated networks.

We want a semantic characterization of saturation that is optimal.

Saturation, revisited [I]

Saturation

Saturation is a syntactic criterion for two-layer networks. Bundala & Závodný prove that it is enough to consider saturated networks.

We want a semantic characterization of saturation that is optimal.

Saturation (better)

A comparator network C is *redundant* if there exists a network C' obtained from C by removing a comparator such that $\text{outputs}(C') = \text{outputs}(C)$.

A network C is *saturated* if it is non-redundant and every network C' obtained by adding a comparator to C satisfies $\text{outputs}(C') \not\subseteq \text{outputs}(C)$.

Saturation, revisited [II]

Saturation theorem


Let C be a saturated two-layer network. Then C contains none of the following two-layer patterns.



Saturation, revisited [II]

Saturation theorem

Let C be a saturated two-layer network. Then C contains none of the following two-layer patterns.



This is easily encoded in words.

Word ::= Head | Stick | Cycle

Head ::= $0(12 + 21)^*$


Stick ::= $(1 + 2)(12 + 21)^*(1 + 2)$

Cycle ::= $12(12 + 21)^*(1 + 2)$

Saturation, revisited [II]

Saturation theorem

Let C be a saturated two-layer network. Then C contains none of the following two-layer patterns.



This is easily encoded in words.

Word ::= Head | Stick | Cycle

Stick ::= 12 | eStick | oStick

Head ::= 0 | eHead | oHead

oStick ::= $12(12 + 21)^+21$

eHead ::= $0(12 + 21)^*12$

eStick ::= $21(12 + 21)^+12$

oHead ::= $0(12 + 21)^*21$

Cycle ::= $12(12 + 21)^*(1 + 2)$

Saturation, revisited [III]

Multi-sets of words corresponding to saturated nets do not contain eWords and oWords.

We can again generate all saturated networks very efficiently for n up to 40.

Saturation, revisited [III]

Multi-sets of words corresponding to saturated nets do not contain eWords and oWords.

We can again generate all saturated networks very efficiently for n up to 40.

A similar technique eliminates nets that are (equivalent to) reflections of others.

This is more expensive because the test for cycles takes time proportional to the number of wires. However, the smallest asymmetric cycle has 12 wires, so for practical purposes this is not a problem.

Some numerology

n	5	6	7	8	9	10	11	12	13
$ G_n $	26	76	232	764	2,620	9,496	35,696	140,152	568,504
$ S_n $	10	28	70	230	676	2,456	7,916	31,374	109,856
$ R(G_n) $	16	20	52	61	165	152	482	414	1,378
$ R(S_n) $	6	6	14	15	37	27	88	70	212
$ R_n $	4	5	8	12	22	21	28	50	117

n	14	15	16	17	18
$ G_n $	2,390,480	10,349,536	46,206,736	211,799,312	997,313,824
$ S_n $	467,716	1,759,422	7,968,204	31,922,840	152,664,200
$ R(G_n) $	1,024	3,780	2,627	10,187	6,422
$ R(S_n) $	136	494	323	1,149	651
$ R_n $	94	262	211	609	411

Outline

- 1 Sorting Networks in a Nutshell
- 2 Reduction Techniques
- 3 A Symbolical Approach
- 4 Conclusions & Future Work**

Results

- Efficient generation of two-layer prefixes for comparator networks
- Representation can capture different important semantic properties
- Identified relevant sets of networks for open cases
- Bottleneck is now processing each relevant network

Ongoing & Future work

- Evaluation of T_{17} by careful generation of the third layer for each of the 609 possible nets
- Application of similar ideas to the size problem

Thank you!