

# SALSA: Simulated Annealing based Loop-Ordering Scheduler for DNN Accelerators

**Conference Paper****Author(s):**

Jung, Victor J.B.; Symons, Arne; Mei, Linyan; Verhelst, Marian; [Benini, Luca](#) 

**Publication date:**

2023-07-07

**Permanent link:**

<https://doi.org/10.3929/ethz-b-000626725>

**Rights / license:**

[In Copyright - Non-Commercial Use Permitted](#)

**Originally published in:**

2023 IEEE 5th International Conference on Artificial Intelligence Circuits and Systems (AICAS), <https://doi.org/10.1109/aicas57966.2023.10168625>

# SALSA: Simulated Anneling based Loop-Ording Scheduler for DNN Accelerators

Victor J.B. Jung<sup>†</sup>, Arne Symons\*, Linyan Mei\*, Marian Verhelst\*, Luca Benini<sup>†</sup>

<sup>†</sup>Integrated Systems Laboratory, ETH Zürich, Switzerland. \*Department of Electrical Engineering, KU Leuven, Belgium.

**Abstract**—To meet the growing need for computational power for DNNs, multiple specialized hardware architectures have been proposed. Each DNN layer should be mapped onto the hardware with the most efficient schedule, however, SotA schedulers struggle to consistently provide optimum schedules in a reasonable time across all DNN-HW combinations.

This paper proposes SALSA, a fast dual-engine scheduler to generate optimal execution schedules for both even and uneven mapping. We introduce a new strategy, combining exhaustive search with simulated annealing to address the dynamic nature of the loop ordering design space size across layers. SALSA is extensively benchmarked against two SotA schedulers, LOMA [1] and Timeloop [2] on 5 different DNNs, on average SALSA finds schedules with 11.9% and 7.6% lower energy while speeding-up the search by  $1.7\times$  and  $24\times$  compared to LOMA and Timeloop, respectively.

**Index Terms**—DNN, accelerator, scheduling, energy-efficiency, combinatorial optimization, simulated annealing

## I. INTRODUCTION

Convolutional Neural Networks (CNNs) [3] are a very successful class of machine learning (ML) models. This type of Deep Neural Network (DNN) consists of a stack of convolutional layers and reaches state-of-the-art (SotA) performance in the fields of image recognition, classification, segmentation, etc. A wide range of specialized hardware (HW) emerged to accelerate DNN execution [4]. These DNN accelerators vary from datacenter-class [5] to embedded systems. The efficiency of a DNN Accelerator is mainly based on the memory hierarchy, the spatial unrolling, and it heavily relies on efficient schedulers to find optimal temporal mappings [6] of DNN layers onto hardware resources.

As previous work has demonstrated, the scheduling of a NN onto such HW accelerators impacts energy and latency up to orders of magnitude [7]. A subtle change in the characteristics of the NN-HW combination can completely modify the optimal schedule. For example, a change in on-chip memory resources can alter the optimal data allocation scheme and even the most efficient workload execution order to minimize energy or latency.

As a result, many design space exploration (DSE) schedulers [2], [8], [9], [1], have been proposed to automatically find the optimal schedule given a DNN workload and an accelerator HW architecture. However, the above-mentioned schedulers

This work is funded in part by the Convolve project evaluated by the EU Horizon Europe research and innovation program under grant agreement No. 101070374 and has been supported by the Swiss State Secretariat for Education Research and Innovation under contract number 22.00150. E-mail: {jungvi / lbenini}@iis.ee.ethz.ch {arne.symons / linyan.mei / marian.verhelst}@esat.kuleuven.be

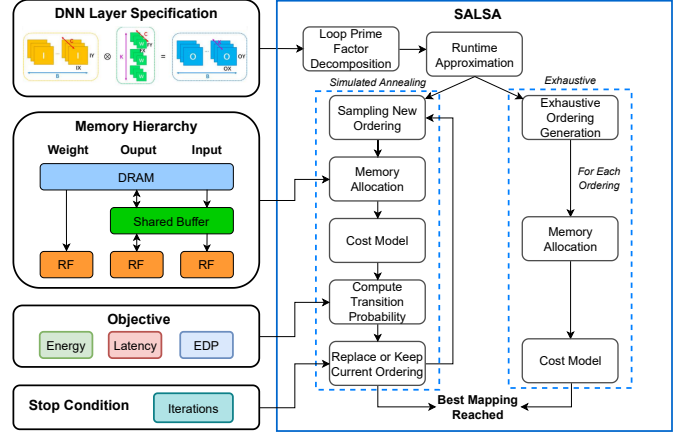


Fig. 1. Overview of the SALSA implementation.

fail to reach near-optimal mappings in a reasonable time. The contributions of this paper are the following:

- 1) We introduce SALSA, a novel scheduler that never shrinks or prunes the schedule search space while having an execution time of a few seconds. Using a dual-engine strategy, SALSA consistently reaches near-optimal schedules with an average error margin of 0.007%.
- 2) To prove its superiority, we extensively compare SALSA with 2 SotA schedulers, LOMA [1] and Timeloop [2]. SALSA always finds schedules with higher or equal quality than Timeloop and LOMA while consequently reducing the search time.

We tested SALSA on 5 commonly used DNNs, benchmarked against Timeloop and LOMA, and evaluated using the SotA cost model ZigZag [10]. **In both benchmarks, SALSA achieves a consequent reduction of the search time, we report  $1.7\times$  and  $24\times$  faster search than LOMA and Timeloop. Most importantly, SALSA reaches superior schedules leading to a reduction of the energy needed to execute the model by 7.6% and 11.9% compared to LOMA and Timeloop, respectively.**

## II. BACKGROUND

### A. DNNs, Accelerators & Schedules

A single convolutional layer consists of 7 nested for-loops, as can be seen in the top-left of Figure 1. The loop dimension sizes determine the tensor size of the three operands; Input (I), Weight (W), and Output (O). Other NN layer topologies (fully connected, pointwise convolutional, matrix-matrix multiplication, etc.) can use the same representation by fixing specific loop dimension sizes to 1. In order to speed up the DNN inference

or increase its energy efficiency, various Application-Specific Integrated Circuit (ASIC) DNN accelerators have been proposed both in academia and by the industry. Such accelerators typically include a spatially unrolled array of Processing Elements (PE) that consist of a Multiply-Accumulate (MAC) unit and local memories to store the operand data. The PEs are connected to larger memories higher up in the memory hierarchy stack through fixed interconnections or a flexible Network-on-Chip (NoC) [4]. These connections allow the multicasting of operand data to multiple PEs, consequently parallelizing the computation. Unrolling a for-loop onto multiple PEs will turn it into a parallel for-loop (parfor-loop). When executing a DNN onto an Accelerator, the set of parfor-loops is named spatial unrolling and indicates the parallelization pattern. Usually, the number of PEs is lower than the dimension of the original for-loops; thus, it is common to split them in order to turn a part of the original for-loop into parfor-loops.

On top of the spatial unrolling, an optimized temporal execution schedule is crucial to extract the full potential of DNN Accelerators. More specifically, a schedule can be decomposed into two elements: 1.) the *loop ordering*, which describes the temporal processing order of the for-loops, and 2.) the *memory allocation*, which assigns the operands of each loop to a specific memory resource. A detailed description of these elements follows later.

### B. Loop Prime Factor Decomposition

The loop ordering of the original nested for-loop representation would not result in an optimal schedule. By decomposing the large loop dimensions into multiple smaller loops, and subsequently re-ordering those smaller loops, better schedules can be found. At the finest level of granularity, each loop is decomposed into the number of prime factors of its loop dimension. The resulting indivisible for-loops are referred to as Loop Prime Factors (LPF). An example of the decomposition of an originally nested for-loop to an LPF ordering is shown in Fig.2 steps A to B.

### C. Loop Ordering Search Space

A loop ordering  $o$  can be seen as a permutation of a finite set of elements, where each element represents a for-loop (Fig.2 step B). The loop ordering search space is thus represented by the Symmetric Group  $S_n$  with  $n$  the number of loops in  $o$ . The order (number of elements) of  $S_n$  is  $n!$  if every element is unique. Due to the LPF decomposition,  $n = 20$  is not uncommon for modern DNN layers. This would require the evaluation of  $O(10^{18})$  orderings. Therefore, exhaustively going through all elements in  $S_n$  is only tractable for small NN layers where  $n < 11$ .

### D. Memory Allocation

Loop ordering has to be combined with the allocation of the data attributed to these loops to specific memory resources in the memory hierarchy (Fig.2 step D). Most mapping representations store the 3 operands (I/W/O) associated with a for-loop at the same memory level. Such mappings are referred to as ‘even memory mappings’. A more complex mapping strategy has

been proposed recently [10], named ‘uneven memory mapping’. This strategy allows to unevenly distribute of operand data of the nested for-loops within shared memories in the hierarchy in order to more efficiently exploit the data reuse at the cost of drastically enlarging the mapping search space.

To reduce the large mapping search space, LOMA [1] proposed a bottom-up memory allocation strategy independent of the loop ordering. This is possible due to the fact that for a single loop ordering  $o$ , the optimal memory allocation  $m$  can be inferred with a one-to-one relationship in a bottom-up fashion.

### E. Cost Model

The energy, latency, or any other performance metric of the inference of a CNN layer on an accelerator depends on four aspects: 1.) the DNN workload  $w$  (size of the 7 loop dimensions); 2.) the accelerator characteristics  $a$  (PE array size, memory organization, memory size, etc.); 3.) the spatial unrolling  $s$  (parallelization strategy across PE array); 4.) the schedule or temporal mapping  $m$ .

This work focuses on temporal DNN mapping optimization, where the inputs  $w$ ,  $a$ , and  $s$  are provided by the user or by an upper-level search engine.

The optimization objective, returned by the cost model, is noted  $V$  and can represent the energy, latency, Energy-Delay Product (EDP), etc.

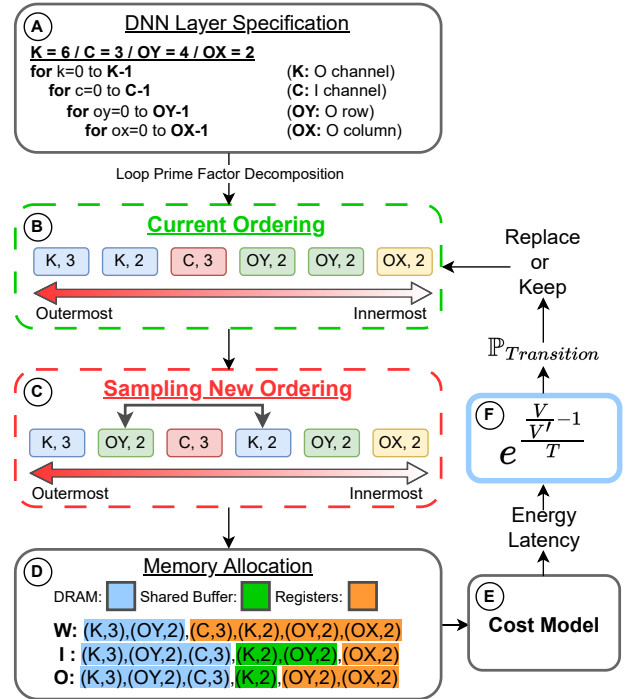


Fig. 2. Detailed example of SALSA's Simulated Annealing path. The workload used in this figure is fictional for the purpose of demonstration, and the Memory Hierarchy is composed of three levels: DRAM, Shared Buffer, and Registers.

## III. RELATED WORK

In recent years, a plethora of tools has been proposed to generate high-quality schedules. Some constrain the search space like CoSA [9], and Pluto [11] to speed up the search. Others, like Interstellar [8] and ZigZag [10] prune some part of the search space during the search through heuristics. LOMA [1] combines

an exhaustive search with optional user-provided constraints, providing both unconstrained and constrained search. Timeloop [2] embeds a random search engine in an unconstrained space, failing to consistently provide near-optimum schedules in fast search time. Alternatively, Mind Mappings [12] trains a DNN to substitute the cost model and make the search space smooth and differentiable in order to apply Stochastic Gradient Descent.

It is important to highlight that, besides the search strategy, the representation of a schedule varies between frameworks. This makes it hard to extensively compare results and performances. All the above-mentioned frameworks implement an even mapping representation (see Section II-D). ZigZag’s and LOMA’s representation also allows uneven mappings. Consequently, its mapping search space becomes more complex.

SALSA overcomes these bottlenecks by implementing a flexible and fast scheduler that allows for both even and uneven mappings generation by separating loop ordering and loop memory allocation in two independent processes. This also allows one to use SALSA with other scheduling representations, e.g., plug in another memory allocation strategy or cost model. As SALSA’s loop ordering algorithm doesn’t use expert knowledge of the cost model or memory allocation, it is robust to drastic changes in the search space.

#### IV. SALSA SCHEDULING APPROACH

To cope with the changing size of the search space from one layer to another, SALSA implements a dual search strategy, as shown in Figure 1. The simulated annealing path is shown in detail in Figure 2.

##### A. Runtime Approximation and Search Method Selection

To decide which of the Exhaustive or Simulated Annealing paths is the fastest (Fig.1), we evaluate and compare their runtime. The Simulated Annealing path’s runtime is constant (depends on a fixed hyperparameter) while the Exhaustive path’s execution time  $T$  is evaluated as follows:

$$T(n, k) = \tau \frac{n!}{\prod_{i=1}^m k_i!} \quad (1)$$

Where  $n$  is the number of elements in the loop ordering,  $k_i$  is the multiplicity of the  $i$ -th element,  $m$  is the number of unique elements in the loop ordering, and  $\tau$  is an HW-dependent constant.

Figure 3 shows how the exhaustive search time exponentially increases with the number of LPFs in a loop ordering while the simulated annealing search time remains constant. We will demonstrate that, even though more LPFs imply a larger permutation space, simulated annealing performs well across all DNN-HW combinations in a constant time.

##### B. Exhaustive Search

The exhaustive search branch is implemented using LOMA’s scheduler [1]. After the exhaustive loop ordering generation, each unique ordering undergoes a bottom-up memory allocation and, finally, a cost model evaluation (both explained next). Most importantly, this exhaustive search engine guarantees to find the global optimum for any preferred optimization criterion at the cost of a potentially infeasible search time.




Fig. 3. Graph illustrating the required search time for different search strategies for varying numbers of LPFs for AlexNet Layer 2. Note the logarithmic y-axis.

##### C. Simulated Annealing Search

In most cases, the exhaustive path would be too time-consuming, and thus the simulated annealing path is taken. Despite its simplicity, simulated annealing [13] and its different variants are widely used and prove to be efficient in combinatorial optimization. Each iteration of the simulated annealing pass will go through the subsequent steps depicted in Figure 2:

1) *Sampling New Ordering*: In order to sample new orderings (Fig.2 step C), we model a neighborhood of nearby states that can serve as the next candidate state [13]. SALSA defines the neighborhood of a loop ordering  $o$  as follows:

$$N_o := \{swap(o, i, j) \mid i \in [0, n), j \in [0, n), i \neq j\} \quad (2)$$

with  $swap(o, i, j)$  the action of swapping the LPFs at indices  $i$  and  $j$  of the ordering  $o$  of size  $n$ . With this neighborhood, any point in the search space can be reached in  $n - 1$  steps.

2) *Memory Allocation & Cost Model Evaluation*: Firstly, we allocate the memory accordingly to the new loop ordering generated by the previous stage (Fig.2 step D). SALSA then uses a cost model to get the performances associated with the candidate state (Fig.2 step E). In this paper, results using the ZigZag as well as the Timeloop cost model will be shown.

3) *Transition Probability Computation & Next Node Selection*: Once the cost  $V'$  of the sampled state  $m'$  is returned by the cost model, SALSA computes the probability of accepting the candidate state  $m'$  using the following formula:

$$\mathbb{P}(m, m') = \exp\left(\frac{V}{V'} - 1\right) \quad (3)$$

where  $V$  and  $V'$  are respectively the optimization objective of the states  $m$  and  $m'$ . The temperature  $T$  is a hyperparameter handling the balance between *intensification* and *diversification* to avoid getting stuck in local optima while focusing the search on promising regions of the search space.

The evolution of  $T$  depends on the number of iterations  $I$  and respect the following geometric progression:  $T_{i+1} = \rho T_i$  where  $\rho = 0.999$




Fig. 4. Mapping energy distribution during a search for layer 2 of AlexNet. using Timeloop and SALSA. Best viewed in color.






Fig. 5. Comparison of SALSA, LOMA 7, and Timeloop for 5 DNN. In this figure, LOMA is configured with an LFP limitation factor of 7. The left part displays the Energy and Search Time for every unique layer of ResNet50, while the right part shows the average Energy Reduction and Speed-up of each DNN. Energy Reduction and Speed-up in the right plots are normalized with Timeloop’s Energy and Time, respectively.

## V. EXPERIMENTAL RESULTS AND BENCHMARKING

### A. Experimental Setup

SALSA is implemented in Python and benchmarked across other schedulers available in the SotA. In our study, we use the 5 following NN: AlexNet, ResNet34 [14], ResNet50 [14], DarkNet19 [15], and MobileNetV2 [16]. The accelerator  $a$  is an Eyeriss-like architecture [4], consisting of a 14 by 12 PE array. Besides a MAC unit, each PE includes a scratchpad for weights, inputs, and outputs. Above the PE array resides a global buffer for storing inputs and outputs, followed by a DRAM that holds all three operands. The spatial dataflow  $s$  is fixed in accordance with the architecture.

The total energy consumption of executing a layer is used as V. Experiments were run on a quad-core CPU @3.6GHz, and with  $I = 1000$ ,  $\rho = 0.999$  and  $T_0 = 0.05$ .

### B. Experimental results

To assess the efficiency of the simulated annealing path of SALSA, we show the energy distribution of mappings using both SALSA and Timeloop (Fig. 4). Note that this energy distribution pattern is consistently found across layers of all studied DNNs. Compared to the random-pruned search of Timeloop, SALSA’s simulated annealing energy distribution is centered on higher-quality states, providing better schedules in a shorter time.

The stochastic nature of SALSA’s simulated annealing motivates an exhaustive search on ResNet34 in order to study the capability of SALSA to consistently reach near-optimal schedules. We used LOMA to exhaustively find the best loop ordering for each unique layer of ResNet34, then we ran SALSA’s simulated annealing engine 500 times per layer. We find that SALSA reaches the global optimum 99.9% of the time. Even when SALSA does not find the global optimum, it still generates high-quality schedules, on average with 0.007% higher energy than the best mapping.

We also compare SALSA against LOMA with various *LFP Limits* (Fig. 6). The *LFP Limit* parameter indicates the maximum size of the orderings considered by LOMA, it limits the number of orderings to evaluate at the cost of the schedule’s energy. We can clearly notice the trade-off between search time and energy between LOMA 6 and SALSA. Since the search for the optimal schedule is done offline, one would always favor lower

energy rather than a reduction of a few seconds in the search time.

Finally, we extensively benchmark SALSA against LOMA and Timeloop (Fig. 5). We choose the LPF limitation factor of LOMA to get a similar search time to SALSA (see Fig. 6). In order to avoid a cost model bias, the schedule found by Timeloop’s engine is evaluated using ZigZag’s cost model. We notice that not all layers benefit from SALSA in the same way: all 3 search engines find similar energy schedules for simple layers (i.e., with fewer loops to permute). However, SALSA significantly outperforms LOMA and Timeloop for more complex layers with a bigger search space, leading to up to 50% of energy reduction. Additionally, SALSA’s search time is drastically lower than Timeloop’s for every layer. Overall, SALSA improves the execution energy by 7.6%, 11.9%, and speed-up the search runtime by 1.7 $\times$ , 24 $\times$ , respectively.




Fig. 6. Performances of SALSA against LOMA X on MobileNetV2 layer 3, X being the LPF limitation factor, shrinking the search space and trading mapping performances for search speed. The configuration of LOMA that does not constrain the search space is noted LOMA Exh (exhaustive).

## VI. CONCLUSION

This paper presented SALSA: a dual-engine, rapid scheduler capable of finding optimal schedules of DNN layers onto an HW accelerator. The simulated annealing-based engine provides an efficient heuristic search guided by any desired performance metric and finds optimal mappings in a short and predictable time. SALSA consistently finds better mappings than current SotA schedulers in a shorter time. It is deployed extensively on 5 DNNs: finding on average 7.6% and 11.9% better energy schedules while speeding up the search by a factor of 1.7 $\times$  and 24 $\times$  compared to LOMA and Timeloop, respectively. By significantly speeding up the process of extracting high-quality temporal mappings, SALSA paves the way for fast spatial unrolling and accelerator architecture search. SALSA is open-sourced and available at [17].

## REFERENCES

- [1] A. Symons, L. Mei, and M. Verhelst, "LOMA: Fast auto-scheduling on dnn accelerators through loop-order-based memory allocation," in *IEEE AICAS*. IEEE, 2021, pp. 1–4.
- [2] A. Parashar, P. Raina, Y. S. Shao, Y.-H. Chen, V. A. Ying, A. Mukkara, R. Venkatesan, B. Khailany, S. W. Keckler, and J. Emer, "Timeloop: A systematic approach to dnn accelerator evaluation," in *IEEE ISPASS*, 2019, pp. 304–315.
- [3] Y. LeCun, Y. Bengio *et al.*, "Convolutional networks for images, speech, and time series," *The handbook of brain theory and neural networks*, vol. 3361, no. 10, p. 1995, 1995.
- [4] Y.-H. Chen, T.-J. Yang, J. Emer, and V. Sze, "Eyeriss v2: A flexible accelerator for emerging deep neural networks on mobile devices," *IEEE JETCAS*, vol. 9, no. 2, pp. 292–308, 2019.
- [5] N. P. Jouppi, C. Young, and e. a. Patil, "In-datacenter performance analysis of a tensor processing unit," *SIGARCH Comput. Archit. News*, vol. 45, no. 2, p. 1–12, jun 2017. [Online]. Available: <https://doi.org/10.1145/3140659.3080246>
- [6] V. Sze, Y.-H. Chen, T.-J. Yang, and J. S. Emer, "Efficient processing of deep neural networks: A tutorial and survey," *Proceedings of the IEEE*, vol. 105, no. 12, pp. 2295–2329, 2017.
- [7] H. Kwon, P. Chatarasi, M. Pellauer, A. Parashar, V. Sarkar, and T. Krishna, "Understanding reuse, performance, and hardware cost of dnn dataflow: A data-centric approach," in *IEEE/ACM MICRO*, 2019.
- [8] X. Yang, M. Gao, Q. Liu, J. Setter, J. Pu, A. Nayak, S. Bell, K. Cao, H. Ha, P. Raina *et al.*, "Interstellar: Using halide's scheduling language to analyze dnn accelerators," in *ACM ASPLOS*, 2020, pp. 369–383.
- [9] Q. Huang, A. Kalaiah, M. Kang, J. Demmel, G. Dinh, J. Wawrzyniek, T. Norell, and Y. S. Shao, "CoSA: Scheduling by constrained optimization for spatial accelerators," in *ACM/IEEE ISCA*. IEEE, 2021, pp. 554–566.
- [10] L. Mei, P. Houshmand, V. Jain, S. Giraldo, and M. Verhelst, "ZigZag: Enlarging Joint Architecture-Mapping Design Space Exploration for DNN Accelerators," *IEEE TC*, 2021.
- [11] U. Bondhugula, A. Acharya, and A. Cohen, "The pluto+ algorithm: A practical approach for parallelization and locality optimization of affine loop nests," *ACM TOPLAS*, vol. 38, no. 3, pp. 1–32, 2016.
- [12] K. Hegde, P.-A. Tsai, S. Huang, V. Chandra, A. Parashar, and C. W. Fletcher, "Mind mappings: enabling efficient algorithm-accelerator mapping space search," in *ACM ASPLOS*, 2021, pp. 943–958.
- [13] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by simulated annealing," *science*, vol. 220, no. 4598, pp. 671–680, 1983.
- [14] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE CVPR*, 2016, pp. 770–778.
- [15] J. Redmon, "Darknet: Open source neural networks in c," 2013.
- [16] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "Mobilenetv2: Inverted residuals and linear bottlenecks," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4510–4520.
- [17] Open-source python implementation of salsa. [Online]. Available: <https://github.com/ZigZag-Project/zigzag>