



HAL
open science

Systolic Arrays and Structured Pruning Co-design for Efficient Transformers in Edge Systems

Pedro Palacios, Rafael Medina, Jean-Luc Rouas, Giovanni Ansaloni, David Atienza

► To cite this version:

Pedro Palacios, Rafael Medina, Jean-Luc Rouas, Giovanni Ansaloni, David Atienza. Systolic Arrays and Structured Pruning Co-design for Efficient Transformers in Edge Systems. GLSVLSI '25: Great Lakes Symposium on VLSI 2025, Apr 2025, New Orleans LA USA, United States. pp.320-327, <10.1145/3716368.3735158>. <hal-05242392>

HAL Id: hal-05242392

<https://hal.science/hal-05242392v1>

Submitted on 5 Sep 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY-NC-ND 4.0 - Attribution - Non-commercial use - No Derivative Works - International License

Systolic Arrays and Structured Pruning Co-design for Efficient Transformers in Edge Systems

Pedro Palacios
pedro.palaciosalmendros@epfl.ch
Embedded Systems Laboratory (ESL)
École Polytechnique Fédérale de
Lausanne (EPFL)
Lausanne, Switzerland

Rafael Medina
rafael.medinamorillas@epfl.ch
Embedded Systems Laboratory (ESL)
École Polytechnique Fédérale de
Lausanne (EPFL)
Lausanne, Switzerland

Jean-Luc Rouas
jean-luc.rouas@labri.fr
LaBRI, CNRS UMR 5800
Univ. Bordeaux, Bordeaux INP
Talence, France

Giovanni Ansaloni
giovanni.ansaloni@epfl.ch
Embedded Systems Laboratory (ESL)
École Polytechnique Fédérale de
Lausanne (EPFL)
Lausanne, Switzerland

David Atienza
david.atienza@epfl.ch
Embedded Systems Laboratory (ESL)
École Polytechnique Fédérale de
Lausanne (EPFL)
Lausanne, Switzerland

Abstract

Efficient deployment of resource-intensive transformers on edge devices necessitates cross-stack optimization. We thus study the interrelation between structured pruning and systolic acceleration, matching the size of pruned blocks with the systolic array dimensions. In this setting, computations of pruned weight blocks can be skipped, reducing run-time and energy consumption, but potentially impacting quality of service (QoS). To evaluate the trade-offs between systolic array size and sparsity opportunities, we present a novel co-design framework that integrates algorithmic optimization, system simulation, and hardware design. Targeting speech recognition and machine translation using transformers as case study, we analyze how configuration choices across the stack affect performance metrics. Results demonstrate that structured pruning on systems featuring systolic array acceleration can effectively increase performance, while maintaining high QoS levels. Up to 44% system-wide speedups due to structured pruning and quantization were measured, with only 1.4% word error rate degradation on the standard LibriSpeech dataset.

CCS Concepts

• **Hardware** → **Hardware-software codesign**; • **Computer systems organization** → **Systolic arrays**; • **Computing methodologies** → **Neural networks**.

Keywords

Systolic arrays, structured pruning, hardware-software co-design, edge AI.

1 Introduction

Transformers have fostered a revolution in machine learning, with applications ranging from classification [10] to generative models for text and images [7], to speech recognition [31]. However, their complex structure based on multiple attention and feed-forward layers [29] results in unprecedented computational requirements,

posing significant challenges for their deployment. These are particularly acute in edge scenarios, where systems have to operate within constrained energy and performance envelopes.

In this context, a plethora of optimization strategies have been proposed. On the software side [8], commonly used approaches involve reducing the precision of data representations (quantization) and removing parts that contribute the least to inference outcomes (pruning). As for hardware, efforts have mainly focused on the acceleration of the main computational kernel in transformers, i.e. General Matrix Multiplications (GEMMs). Although diverse solutions ranging from analog crossbars [25][15] to near-DRAM computing [11][21] have been investigated to this end, a particularly promising alternative is represented by systolic arrays [34]. These two-dimensional meshes of processing elements can indeed parallelize the computation of a GEMM (or, more precisely, the computation of a GEMM tile), while presenting high parallelism, low resource requirements and only mandating a simple, low-overhead control logic.

Recent works [6, 13, 28, 33, 35] have attempted to co-optimize software algorithms and hardware accelerators dedicated to transformer inference [17]. Such a stance is particularly appealing at the crossroads of model pruning and systolic array acceleration. On the software side, pruning can be performed by eliding weights in regular block patterns (in a “structured” way) rather than as individual elements [32]. While this approach introduces a constraint to pruning, and can hence result in lower overall sparsity rates, it substantially amplifies hardware-side optimization opportunities when matching the sizes of the pruned tile and the accelerator mesh. The exploration of this strategy, which we term Systolic Array Structured Pruning (SASP), is the focus of this work.

SASP opens a complex multidimensional design space which requires careful consideration of metrics spanning from hardware to algorithms. Indeed, while a larger accelerator can expose a higher degree of parallelism, it also requires more resources (area / energy). Moreover, SASP settings with larger tiles may overly penalize the achievable sparsity for a desired quality of service (QoS) or, alternatively, result in high QoS degradation for a fixed pruning rate.

To quantitatively explore these interrelations, we present a novel framework and methodology which, for the first time, integrates methods for a) the structured pruning and quantization of transformer algorithms, b) the system-level modeling of accelerated systems executing them, and c) the hardware synthesis of accelerators. Our environment for SASP exploration builds on state-of-the-art methods for the training of transformers (ESPnet [31]) and for system simulation (gem5 [5]). By employing a novel systolic array architectural template, it supports both floating point and weight-quantized data representations, as supported by ESPnet.

As a test case, we employ our exploration approach to analyze speech recognition (ASR) and machine translation (MT) applications, based on transformers processing the LibriSpeech [22] and MuST-C [9] datasets. We observed that SASP can achieve, for a 32×32 systolic array and a ESPnet ASR model trained on LibriSpeech, up to 44% speedup and 42% energy savings over a non-pruned, non-quantized system when employing a 20% pruning rate, resulting in a marginal Word Error Rate (WER) degradation of 1.4%.

The contributions of this paper are summarized as follows:

- We perform a systematic exploration of Systolic Array Structured Pruning (SASP), a co-design strategy that combines systolic array acceleration and structured pruning with matching accelerator and tile size under quality of service constraints.
- We introduce a cross-stack framework to enable the evaluation of SASP through figures of merit at different abstraction levels, including the assessment of QoS, performance, resource usage, and energy, as well as their trade-offs. We discuss how these insights can be effectively leveraged from the joint perspective of algorithmic optimization, system integration, and systolic array design.
- Across different models and tasks, we show that SASP-based co-optimization of transformers and systolic arrays can offer significant speedup and energy efficiency gains. For example, a joint ASR and MT case study achieves reductions in runtime and energy consumption of up to 51% and 34% while limiting QoS degradation to 4 BLEU points.

2 State of the Art

By providing spatially-distributed computation with low control logic overhead, systolic arrays can effectively parallelize the execution of matrix multiplications, the dominant computing pattern in transformer inference [34]. To evaluate the benefits that systolic arrays can induce, SMAUG [33] and TiC-SAT [2] present system simulation infrastructures able to support complete inferences, including both their hardware-accelerated and their software-executed parts. These works showcase that even small-sized systolic arrays have the potential to reduce run-time by orders of magnitude. Performance is further improved when data is properly laid out in a tiled arrangement in memory according to the accelerator characteristics, in order to maximize spatio-temporal locality. Although such an approach has been adopted [1, 28], no attempt was made therein to prune computations as we do here with Systolic Array Structured Sparsity.

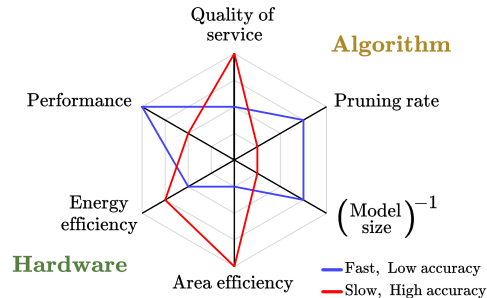


Figure 1: Qualitative radar plot illustrating two SASP solutions with different trade-offs: a slow and accurate one (red) employing a small accelerator and a low pruning rate, and a fast but inaccurate one (blue) using a large accelerator and a high pruning rate. Across all axes, higher is better.

Indeed, the ductility of DNN models makes them highly amenable to pruning. In the context of systolic array acceleration, pruning optimizations are categorized as either fine-grained or structured [17]. In the first case, specialized systolic arrays have been proposed which can leverage the presence of zero values in tiles by either clock gating processing elements [18] or by reordering operands [19, 27]. Nonetheless, fine-grained pruning requires a fair amount of control logic overhead in the accelerator design and impacts the regularity of data layout in memory, which may negate the intended benefits [20, 35].

In this light, structured pruning strategies offer a promising alternative, as tiles of low significance can be entirely skipped before processing them onto accelerators, when the tile size matches the target accelerator parallelism. Hence, speedups can be harnessed without requiring specialized hardware for sparsity management. This position has been adopted by previous works [3, 4, 14, 16, 26]. However, they only provide a partial view of the ensuing design space. In particular, [3, 26] and [14] adopt a system-level stance, exploring the potential for acceleration of co-designed pruning strategies and data-parallel accelerators, but do not investigate the impact on quality of service (QoS, e.g. accuracy, word error rate) of the performed pruning. Conversely, [16] and [4] provide an algorithmic-level assessment of the effect of structured sparsity, but neglect the hardware and architectural implication of adopting a matched accelerator design. To the best of our knowledge, a holistic analysis of the algorithmic-to-hardware space exposed by SASP is hence missing. Our paper aims at filling this gap.

3 Co-designing Accelerators and Sparsity

The dimensions of the design space exposed by Systolic Array Structured Pruning (SASP) solutions are illustrated in Fig. 1. The illustrated figures of merit reside at widely different layers of the hardware/software stack. Hence, to enable co-optimization, we developed the integrated toolflow illustrated in Fig. 2. The input to our framework is a trained transformer model and a target dataset, which, in our implementation, are defined via the ESPnet toolkit for automatic speech processing [31]. Hyper-parameters determine the size of the SASP tile (which sets the granularity of structured sparsity as well as the size of the systolic array) and the target

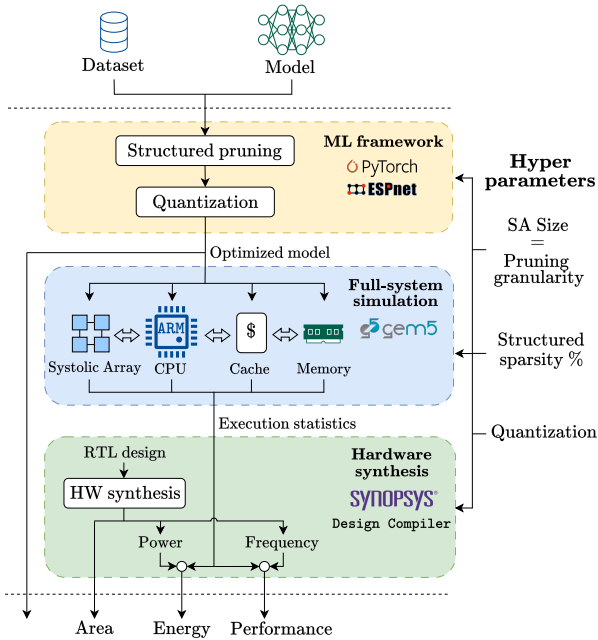


Figure 2: Overview of Hardware-Software co-design framework.




Figure 3: Tiled matrix multiplication with structured pruning.

sparsity rate. Moreover, both floating-point and weight-quantized implementations are supported.

The co-design framework is structured in 3 tiers: in its upper tier, PyTorch APIs are employed to perform pruning and (optionally) quantization. Then, system simulation is employed to gather run-time statistics of the application executing on a virtual system featuring systolic acceleration. Finally, an RTL-level architectural template is used to gather hardware metrics such as energy and area. The implementation of each tier is detailed in the following.

3.1 Structured Pruning and Quantization

Matrices employed by transformer models are commonly much larger than the size of systolic arrays. Hence, operations to perform GEMM must be tiled. As in [2], we herein consider a weight-stationary scenario, in which a tile of parameters is stored in the systolic array, and partial results are computed by streaming inputs/outputs to/from the accelerator. These are eventually aggregated via element-wise addition.

As shown in Fig. 3, in this setting a tile containing only zero values can be completely skipped, saving both the time required to configure the systolic array and the time required to calculate the

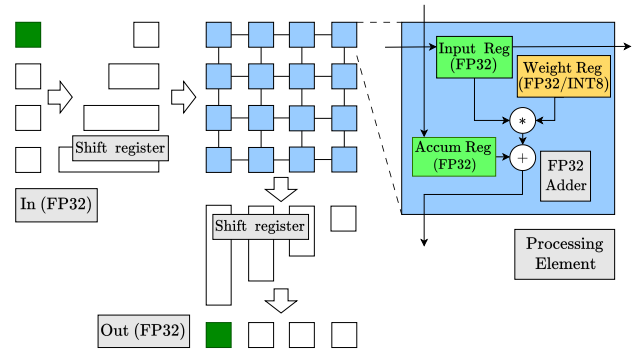


Figure 4: Architectural diagram of the systolic array, supporting FP32 activations and either non-quantized (FP32) or quantized (INT8) weights.

related partial results. In the example in Fig. 3 it can be observed that the sparsity induced by the red weight tile lowers the workload required for the computation of the entire shaded column in the output. We enforce structured sparsity by zeroing a percentage of tiles with the lowest L1-norm (sum of absolute values) across the entire model. This approach allows to heterogeneously prune GEMMs according to their sensitivity. In particular, feed-forward GEMMs are much more amenable to pruning than attention ones, so we focus on these for our exploration in Section 4.3.

After sparsification, post-training quantization can optionally be used to reduce the representation precision of weights from 32-bits floating point (FP32) to 8-bits integer (INT8). Finally, inference is performed on a target dataset, in order to gather QoS metrics such as Word Error Rate (WER) or BLEU.

3.2 Full System Simulation

Run-time statistics on the deployment of the SASP-pruned models are collected in the gem5 [5] simulation environment, which allows specifying complex systems including hardware (processors, memory hierarchy) and software (operating system) components. To this end, we developed a gem5 custom functional unit modeling a tightly coupled systolic array. Similarly to [2], the functional unit employs dedicated instructions, extending the ARM instruction set, to a) program weights, b) perform the systolic array computation, and c) stream inputs/outputs (see Fig. 4). We assume a 32-bit input-output interface, allowing to transfer one input and one output activation per custom instruction. As for weights, either a single FP32 or four INT8 values can be programmed in the array in a quantized or non-quantized setting, respectively.

The implemented instruction set extensions can be employed to accelerate user-level applications via inline assembly pragmas. For convenience, we wrapped these in parametric library functions, allowing to transfer a weight tile or compute a partial GEMM with a single function call. In this way, we gathered the run-time characteristics of executing entire transformer layers under varying architectural and sparsity settings.

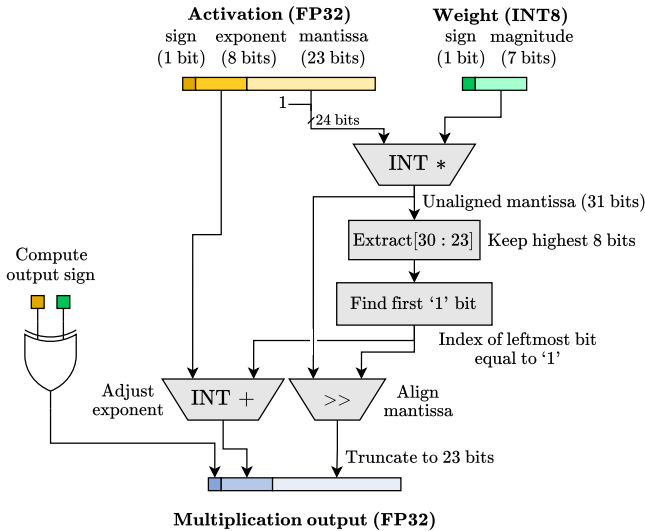


Figure 5: Hardware diagram of the hybrid FP32_INT8 multiplier. This logic is bypassed in case any of the operands is equal to zero.

3.3 Systolic Array Architecture

The systolic array hardware implementation, depicted in Figure (Fig. 4), comprises a mesh of processing elements (PEs) with nearest-neighbor connections. Inside each, an adder and a multiplier implement a MAC operation between input activation, weight and partial result values, the latter being stored in an accumulation register. Notice that inputs are streamed left-to-right, partial results flow from top to bottom, and weights are instead stationary. At the periphery of the array, shift registers of varying depth are employed to skew data along a diagonal, properly aligning inputs and outputs.

Instances of the template can be derived by providing architectural parameters defining its size and the desired data format (either FP32 for weights and activations, or using INT8 weights and FP32 activations). In both versions of the PE, the multiplier and adder are pipelined to meet timing requirements. The pipeline latency is entirely hidden by the activations I/O latency from/to the systolic array. The instances are fully synthesizable using standard digital IC design tools and logic cell libraries.

Both in non-quantized and weight-quantized settings, adders in PEs support FP32 representation both for operands and for the results. Conversely, multipliers can be highly optimized in the weight-quantized case, as they must only support simpler FP32_INT8 arithmetic. A diagram of the hybrid FP32_INT8 multiplier design adopted in our architectural systolic array template is presented in Fig. 5. This implementation correctly computes the multiplication result, except for the case where either of the inputs equal to 0. We handle this as a special case, by employing a dedicated multiplexer. Moreover, to optimize area and energy efficiency, infinities, NaNs, and subnormal numbers are not handled.

In detail, our design assumes that the INT8 weight is represented using a sign-and-magnitude format. Hence, the output sign is computed as the XOR of the activation and weight signs. Then, the FP32 mantissa is expanded by appending the leading ‘1’, which is

implicit in the IEEE format. Furthermore, the expanded mantissa is multiplied by the magnitude of the weight value (INT8). The resulting unaligned output mantissa is right-shifted to align the leading ‘1’ and truncated to 23 bits. Finally, the output exponent is adjusted according to the number of performed mantissa shifts.

Note that the hybrid multiplier design readily generalizes to different floating-point and integer bitwidths beyond the FP32_INT8 considered in this paper, e.g., to support FP16 activations.

4 Experimental Results

4.1 Setup

To evaluate our framework’s capabilities, we focus on automatic speech recognition and machine translation, two challenging and commercially relevant edge AI tasks, using the LibriSpeech [22] and MuST-C [9] corpora. Using ESPnet [31] and PyTorch [23], we implemented transformer models targeting the encoder for SASP optimization, since its execution dominates run-time. The structure and training method of the models correspond to the parameters shown in Table 1. All the quality of service (QoS) results in this paper are reported on the respective corpora test sets and are comparable to state-of-the-art results [12]. System simulations were run in the gem5-X variant [24] of the gem5 simulator [5]. We considered a single-core configuration having a 2-level cache hierarchy and running at 1 GHz, as detailed in Table 2. Hardware syntheses of systolic array instances targeted the same 1 GHz timing constraint, and employed a TSMC 28nm technology node. Floating point arithmetic operators (adders and multipliers) were derived from the FPxx library using SpinalHDL [30], while the hybrid FP32_INT8 multiplier was implemented from scratch according to the design in Section 3.3.

In all experiments, we considered systolic arrays of sizes ranging from 4×4 to 32×32 . We spanned various structured pruning rates and investigated both FP32_FP32 and FP32_INT8 quantization schemes. Experiments at different abstraction tiers collected results on area, energy, run-time, and achieved quality of service (QoS). We detail each in the rest of the section. Then, we provide insights from a cross-tier point of view and summarize our findings in Section 4.6.

4.2 Hardware Exploration

The area and power results for different systolic array sizes and quantization choices are shown in Fig. 6. Since multipliers account for an important part of the area and power budget of the entire systolic array (55.6% and 33.6%, respectively, in the 8×8 , FP32_FP32 implementation), the use of the simpler FP32_INT8 design results in tangible savings. In average, these reductions amount to 35.3% and 19.5% in area and power across different array sizes.

Both area and power grow quadratically with the systolic array dimension (e.g. by ~ 4 times between the 4×4 and the 8×8 instances), as both the number of PEs and the number of elements in input/output shift registers have a quadratic dependency on the number of array rows/columns.

4.3 System Exploration

We first assess the overall impact of SASP across workloads and systolic array sizes on runtime and energy consumption under the target QoS degradations defined in Table 1. As shown in Fig. 7,

Table 1: Parameters and Quality of Service (QoS) of the deployed models, targeting automatic speech recognition (ASR) and machine translation (MT). Training sets are augmented using 3× speed perturbation. QoS of the model on the MuST-C dataset [9] are obtained for their execution in cascade (ASR → MT).

Toolkit	Task	Dataset	Encoder blocks	Decoder blocks	Attention heads	d_{model}	Feed-forward size	Training set	Training epochs	Test set	Model QoS	SASP QoS target
ESPnet	ASR	LibriSpeech	18	6	4	512	2048	960 h	100	5 h	3.5% WER	5% WER
ESPnet2	ASR	LibriSpeech	12	6	8	512	2048	960 h	100	5 h	3.2% WER	5% WER
ESPnet2	ASR + MT	MuST-C	18 / 6	6 / 6	4 / 4	128 / 128	2048 / 1024	560 h	35 / 200	5 h	31 BLEU	27 BLEU

Table 2: Configuration of the simulated system.

Processors	1x in-order ARMv8 core @1.0 GHz
L1-I Cache	32 kB, 2-way, 2 cycle access
L1-D Cache	32 kB, 2-way, 2 cycle access
L2 Cache	1 MB, 2-way, 20 cycle access
Memory	DDR4 2400 MHz, 4 GB
Operating System	Ubuntu LTS 16.04
Systolic array	Tightly coupled, control via custom instructions

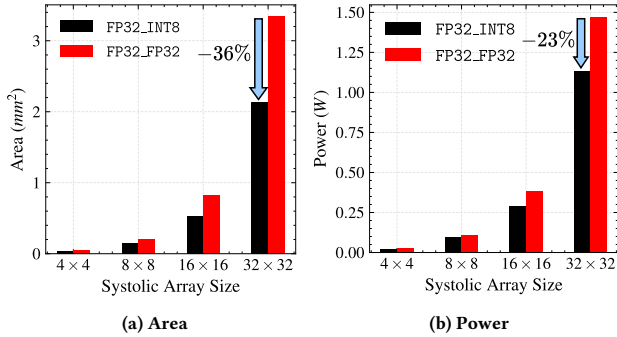


Figure 6: Synthesis results for the systolic array design across configurations of varying size (number of rows of the array) and quantization.

speedup and energy improvement opportunities depend on the workload due to their different sensitivities to pruning. For example, maximum speedup / energy improvements vary from 26% / 21% in ESPnet ASR on LibriSpeech, to 22% / 18% in ESPnet2 ASR on LibriSpeech, to 51% / 34% in ESPnet2 ASR+MT on MuST-C. However, their trends when changing accelerator size are consistent across workloads, e.g., achievable improvements decrease when employing larger systolic arrays. To thoroughly analyze the trade-offs among encoder inference speedup, area-energy product and quality of service, all further explorations will focus on the ESPnet automatic speech recognition model using the LibriSpeech corpus.

Fig. 8 plots the measured per-layer normalized encoder run-time of a systolic-accelerated system performing an inference. Data refers to an 8×8 systolic array, with varying degrees of structured sparsity. Speedup numbers closely follow sparsity levels, as inference run-time is strongly dominated by GEMM computations (exceeding 97% in all cases [2]).

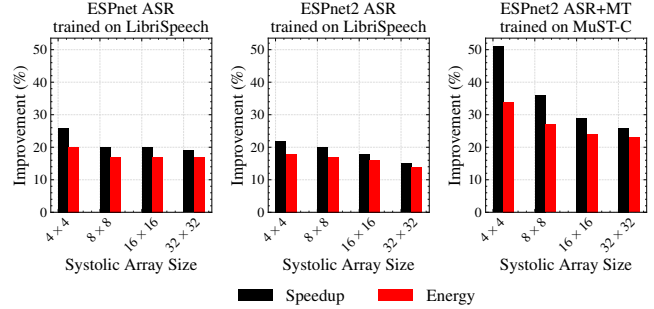


Figure 7: Speedup and energy improvements from Systolic Array Structured Pruning under the quality of service target in Table 1 with respect to non-pruned quantized executions, considering systems with systolic array accelerators of varying size. Results are shown for FP32_INT8 systolic array configurations.

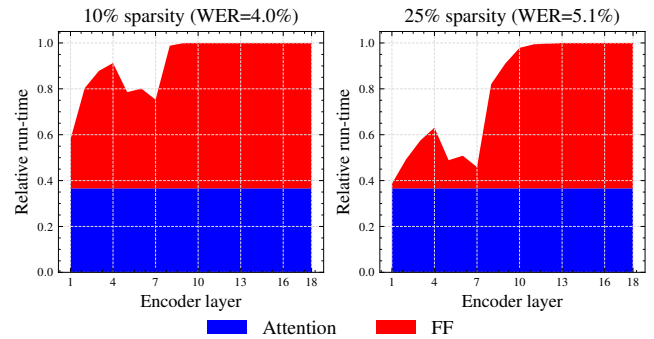


Figure 8: Per-layer normalized run-time of the transformer encoder after applying Systolic Array Structured Pruning at two global sparsity targets. Results are shown for an 8×8 , FP32_INT8 systolic array, with the execution time of each layer normalized to the execution without pruning.

Pruning was performed according to the methodology proposed in Section 3.1 targeting feed-forward layers, which exhibit a much higher degree of resilience and account for the largest part of the workload. Results in Fig. 8 highlight that early feed-forward layers are the most amenable to pruning, while later ones have a higher proportion of tiles with a non-negligible L1-norm, which have a higher impact on inference outcomes.

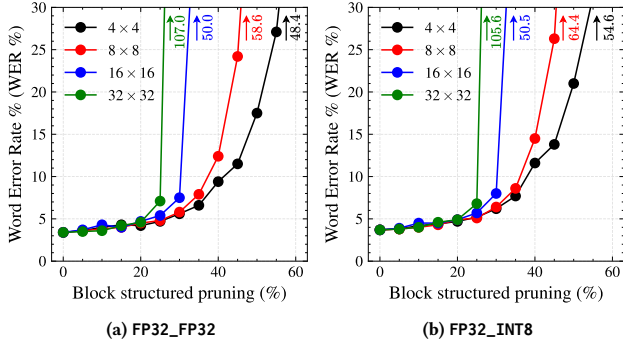


Figure 9: Achieved Word Error Rate when varying the percentage of Systolic Array Structured Pruning.

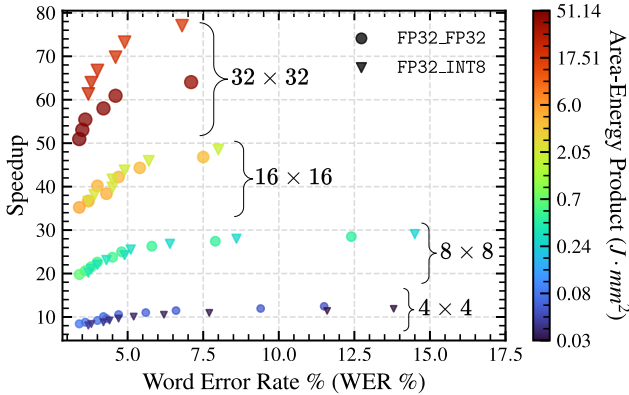


Figure 10: Trade-offs among encoder inference speedup, area-energy product and Word Error Rate across systolic array sizes and structured pruning rates. Speedup is computed with respect to a non-accelerated, non-quantized baseline.

4.4 Impact of SASP on Quality of Service

Fig. 9 shows that Word Error Rate (WER) grows exponentially when increasing the degree of Systolic Array Structured Pruning. Similar trends are present for both the weight-quantized model (indicated as FP32_INT8) and the non-quantized one (FP32_FP32).

As the size of the systolic array grows, trends become steeper, showcasing an abrupt increase in WERs at smaller SASP rates. This effect is caused by the higher brittleness of large-tile structured pruning with respect to small-tile cases. Indeed, while it may be possible to find four prunable 4×4 tiles (containing 64 values in total), selecting a single contiguous 8×8 tile (again, of 64 values) can be considerably more challenging.

4.5 Multidimensional SASP Trade-offs

Focusing again on speech recognition, Fig. 10 shows the variations in performance, QoS and resource usage when changing SASP rate, quantization strategy, and systolic array size. WER and Speedup are plotted on the two axes, the marker shape discriminates between FP32 and weight-quantized implementations, and the marker colors

indicate their resource requirements in terms of Area-Energy product. Data points form four distinct clusters corresponding to each systolic array size. Two curves in each cluster represent the two quantization choices, which have notably different Area-Energy products. The non-quantized FP32_FP32 version achieves lower (better) WER, with the differences becoming more pronounced at higher pruning rates and for larger systolic arrays.

From a run-time perspective, FP32_INT8 configurations allow to reduce the cost of weight transfers by loading four INT8 weights per 32-bit bus access, as opposed to a single FP32 one in the FP32_FP32 configuration. Consequently, FP32_INT8 implementations outperform their FP32_FP32 counterparts for systolic array sizes larger than 4×4 , as the savings in data transfers offset software/system overhead. Nonetheless, while quantization has a large impact on area and energy, its influence on performance is smaller, as the majority of the run-time is not spent in weight data transfers, but instead for streaming inputs / computing outputs, which is equally fast for both quantization schemes.

Within each systolic array size and quantization configuration, the SASP pruning rate guides the trade-off between inference time and QoS. Fig. 10 shows that, up to an inflection point at a WER of $\sim 5\%$, SASP enables to strike effective fine-grained balances between run-time performance and QoS. Beyond this inflection point, further increases in pruning rates cause instead very high WER degradations.

Table 3 illustrates the effect of applying Systolic Array Structured Pruning and weight quantization at the 5% inflection point. In this setting, SASP improves performance and energy consumption up to 26% and 21%, respectively. Furthermore, when combining quantization and structured pruning, performance and energy efficiency improvements reach 44% and 42%, while also decreasing area occupation by 36%. These substantial gains are achieved without increasing the systolic array size, and hence do not incur the hefty associated area and energy costs, as also depicted in Table 3. As an example, scaling from an 8×8 to a 32×32 systolic array does yield a $3.04\times$ speedup for FP32_INT8 quantization, but also requires $15.21\times$ more area and $3.98\times$ more energy.

4.6 Cross-tier Analysis

As discussed above, increasing the systolic array size increases run-time performance by offering higher parallelism, but also restricts the achievable pruning rate for a given QoS. Therefore, when using SASP, increases in the array size (hence the tile size in GEMM computations) result in diminishing gains for a given WER target, as the pruning rate must be lowered to maintain QoS. This trend is illustrated in Fig. 11, which highlights a sublinear relation between systolic array size and speedup, both in the case of FP32 and weight-quantized scenarios. Hardware costs (area / energy), instead, increase quadratically with size, as outlined in Section 4.2, requiring careful co-design considerations, especially when targeting resource-constrained edge systems.

5 Conclusion

Deploying transformers on edge devices requires both software optimization and hardware acceleration to meet strict resource constraints while maintaining performance. In this paper, we have

Table 3: Area, encoder speedup and energy results for different systolic array configurations without SASP ($3.5 \pm 0.2\%$ WER) and with SASP ($5 \pm 0.4\%$ WER). Speedup is computed with respect to a non-quantized baseline executed on CPU.

		Size	4 × 4	8 × 8	16 × 16	32 × 32
FP32_FP32	Area (mm ²)		0.05	0.21	0.83	3.34
	No SASP	Speedup	8.42	19.79	35.22	50.95
		Energy (J)	1.60	3.09	6.37	15.32
	SASP	Pruning (%)	25	25	20	20
		Speedup	10.56	25.01	42.21	60.91
		Energy (J)	1.27	2.43	5.28	12.70
FP32_INT8	Area (mm ²)		0.03	0.14	0.53	2.13
	No SASP	Speedup	8.03	20.18	36.53	61.33
		Energy (J)	1.24	2.67	4.57	10.64
	SASP	Pruning (%)	25	20	20	20
		Speedup	10.08	24.23	43.74	73.25
		Energy (J)	0.99	2.21	3.79	8.82

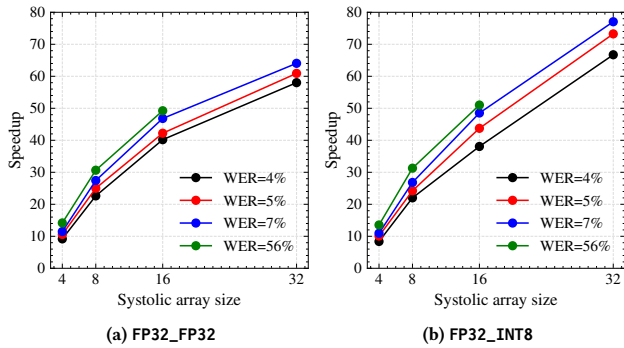


Figure 11: Speedup with respect to software execution of the encoder while varying the systolic array size, for different Word Error Rates.

analyzed their interaction, focusing on structured sparsity and systolic arrays. We explored Systolic Array Structured Pruning (SASP), where the size of pruned blocks is matched to the dimensions of the systolic array, enabling the skipping of entire computation tiles. To assess the benefits and pitfalls of SASP, we presented a cross-stack framework to co-optimize edge AI transformers, which integrates algorithmic optimization, system simulation, and hardware design. Employing it, we performed a comprehensive analysis of how SASP, quantization, and systolic array configurations affect area, energy, performance, and quality of service (QoS) in edge transformers for speech recognition and machine translation.

Our results demonstrate that SASP provides fine-grained control over the trade-off between inference run-time and QoS up to an inflection point, after which increased pruning drastically degrades QoS for small performance gains. Experimental evidence on the ESPnet ASR transformer trained on LibriSpeech has shown that system-wide performance improvements of up to 44% and accelerator energy reductions of up to 42% can be obtained under a 1.4% WER degradation, when employing weight quantization

and a 20% pruning rate. Additionally, we showcased that, although larger systolic arrays do reduce run-time, they also incur substantial energy and area costs, while yielding sublinear speedups for a target QoS. This sublinearity emerges from the reduced structured pruning opportunities, as finding contiguous zero blocks becomes harder with increasing block sizes. For this reason, SASP is particularly well-suited for edge AI accelerators, where stringent resource constraints are present.

Acknowledgments

This work was supported in part by the Swiss State Secretariat for Education, Research, and Innovation (SERI) through the Swiss-Chips Research Project; in part by the Swiss NSF Edge-Companions project (GA No. 10002812); in part by the EC H2020 FVLLMONTI Project under Grant 101016776; and in part by “La Caixa” Foundation (ID 100010434) under Grant LCF/BQ/EU23/12010073. This research was partially conducted by ACCESS – AI Chip Center for Emerging Smart Systems, supported by the InnoHK initiative of the Innovation and Technology Commission of the Hong Kong Special Administrative Region Government.

References

- [1] Alireza Amirshahi, Giovanni Ansaloni, and David Atienza. 2024. Accelerator-Driven Data Arrangement to Minimize Transformers Run-Time on Multi-Core Architectures. In *15th Workshop on Parallel Programming and Run-Time Management Techniques for Many-Core Architectures and 13th Workshop on Design Tools and Architectures for Multicore Embedded Computing Platforms (PARMA-DITAM 2024)*.
- [2] Alireza Amirshahi, Joshua Alexander Harrison Klein, Giovanni Ansaloni, and David Atienza. 2023. TiC-SAT: Tightly-coupled Systolic Accelerator for Transformers. In *2023 28th Asia and South Pacific Design Automation Conference (ASP-DAC)*. 657–663.
- [3] Bahar Asgari, Ramyad Hadidi, Hyesoon Kim, and Sudhakar Yalamanchili. 2019. ERIDANUS: Efficiently Running Inference of DNNs Using Systolic Arrays. *IEEE Micro* 39, 5 (2019), 46–54. doi:10.1109/MM.2019.2930057
- [4] Leila Ben Letaifa and Jean-Luc Rouas. 2023. Variable Scale Pruning for Transformer Model Compression in End-to-End Speech Recognition. *Algorithms* 16, 9 (2023), 398.
- [5] Nathan Binkert, Bradford Beckmann, Gabriel Black, Steven K. Reinhardt, Ali Saidi, Arkaprava Basu, Joel Hestness, Derek R. Hower, Tushar Krishna, Somayeh Sardashti, Rathijit Sen, Korey Sewell, Muhammad Shoaib, Nilay Vaish, Mark D. Hill, and David A. Wood. 2011. The gem5 simulator. *SIGARCH Comput. Archit. News* 39, 2 (Aug. 2011), 1–7. doi:10.1145/2024716.2024718
- [6] Oliver Bringmann, Wolfgang Ecker, Ingo Feldner, Adrian Frischknecht, Christoph Gerum, Timo Hämäläinen, Muhammad Abdullah Hanif, Michael J. Klaiber, Daniel Mueller-Gritschneider, Paul Palomero Bernardo, Sebastian Prebeck, and Muhammad Shafique. 2021. Automated HW/SW Co-design for Edge AI: State, Challenges and Steps Ahead: Special Session Paper. In *2021 International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS)*. 11–20.
- [7] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. In *Proceedings of the 34th International Conference on Neural Information Processing Systems (Vancouver, BC, Canada) (NIPS '20)*. Curran Associates Inc., Red Hook, NY, USA, Article 159, 25 pages.
- [8] Krishna Teja Chitty-Venkata, Sparsh Mittal, Murali Emani, Venkatram Vishwanath, and Arun K Somani. 2023. A survey of techniques for optimizing transformer inference. *Journal of Systems Architecture* (2023), 102990.
- [9] Mattia A. Di Gangi, Roldano Cattoni, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2019. MuST-C: a Multilingual Speech Translation Corpus. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 2012–2017. doi:10.18653/v1/N19-1202

- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xi-aohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*.
- [11] Juan Gómez-Luna, Izzat El Hajj, Ivan Fernandez, Christina Giannoula, Geraldo F Oliveira, and Onur Mutlu. 2021. Benchmarking memory-centric computing systems: Analysis of real processing-in-memory hardware. In *2021 12th International Green and Sustainable Computing Conference (IGSC)*. IEEE, 1–7.
- [12] Pengcheng Guo, Florian Boyer, Xuankai Chang, Tomoki Hayashi, Yosuke Higuchi, Hirofumi Inaguma, Naoyuki Kamo, Chenda Li, Daniel Garcia-Romero, Jiatong Shi, Jing Shi, Shinji Watanabe, Kun Wei, Wangyou Zhang, and Yuekai Zhang. 2021. Recent Developments on Espnet Toolkit Boosted By Conformer. In *IEEE International Conference on Acoustics, Speech and Signal Processing*. 5874–5878. doi:10.1109/ICASSP39728.2021.9414858
- [13] Nithilan Kannappan Jayakodi, Janardhan Rao Doppa, and Partha Pratim Pande. 2021. A General Hardware and Software Co-Design Framework for Energy-Efficient Edge AI. In *2021 IEEE/ACM International Conference On Computer Aided Design (ICCAD)*. 1–7. doi:10.1109/ICCAD51958.2021.9643557
- [14] Hyeon-Ju Kang. 2019. Accelerator-aware pruning for convolutional neural networks. *IEEE Transactions on Circuits and Systems for Video Technology* 30, 7 (2019), 2093–2103.
- [15] Joshua Klein, Irem Boybat, Yasir Mahmood Qureshi, Martino Dazzi, Alexandre Levisse, Giovanni Ansaloni, Marina Zapater, Abu Sebastian, and David Atienza. 2022. Alpine: Analog in-memory acceleration with tight processor integration for deep learning. *IEEE Trans. Comput.* 72, 7 (2022), 1985–1998.
- [16] Bingbing Li, Zhenglun Kong, Tianyun Zhang, Ji Li, Zhengang Li, Hang Liu, and Caiwen Ding. 2020. Efficient transformer-based large scale language representations using hardware-friendly block structured pruning. *arXiv preprint arXiv:2009.08065* (2020).
- [17] Tailin Liang, John Glossner, Lei Wang, Shaobo Shi, and Xiaotong Zhang. 2021. Pruning and quantization for deep neural network acceleration: A survey. *Neurocomputing* 461 (2021), 370–403.
- [18] Zhi-Gang Liu, Paul N Whatmough, and Matthew Mattina. 2020. Systolic tensor array: An efficient structured-sparse GEMM accelerator for mobile CNN inference. *IEEE Computer Architecture Letters* 19, 1 (2020), 34–37.
- [19] Liqiang Lu, Yicheng Jin, Hangrui Bi, Zizhang Luo, Peng Li, Tao Wang, and Yun Liang. 2021. Sanger: A co-design framework for enabling sparse attention using reconfigurable architecture. In *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture*. 977–991.
- [20] Xiaolong Ma, Sheng Lin, Shaokai Ye, Zhezhi He, Linfeng Zhang, Geng Yuan, Sia Huat Tan, Zhengang Li, Deliang Fan, Xuehai Qian, et al. 2021. Non-structured DNN weight pruning—Is it beneficial in any platform? *IEEE transactions on neural networks and learning systems* 33, 9 (2021), 4930–4944.
- [21] Rafael Medina Morillas, Saeideh Alinezhad Chamazcoti, Marina Zapater Sanchó, Giovanni Ansaloni, Timon Evenblij, Alexandre Sébastien Julien Levisse, Dwaipayan Biswas, Francky Catthoor, and David Atienza Alonso. 2024. Bank on Compute-near-Memory: Design Space Exploration of Processing-near-Bank Architectures. In *International Conference on Hardware/Software Codesign and System Synthesis (CODES 2024)*.
- [22] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. Librispeech: An ASR corpus based on public domain audio books. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 5206–5210. doi:10.1109/ICASSP.2015.7178964
- [23] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. *PyTorch: an imperative style, high-performance deep learning library*. Curran Associates Inc., Red Hook, NY, USA.
- [24] Yasir Mahmood Qureshi, William Andrew Simon, Marina Zapater, David Atienza, and Katzalin Olcoz. 2019. Gem5-X: A Gem5-Based System Level Simulation Framework to Optimize Many-Core Platforms. In *2019 Spring Simulation Conference (SpringSim)*. 1–12. doi:10.23919/SpringSim.2019.8732862
- [25] Katie Spoon, Hsinyu Tsai, An Chen, Malte J Rasch, Stefano Ambrogio, Charles Mackin, Andrea Fasoli, Alexander M Friz, Pritish Narayanan, Milos Stanisavljevic, et al. 2021. Toward software-equivalent accuracy on transformer-based deep neural networks with analog memory devices. *Frontiers in Computational Neuroscience* 15 (2021), 675741.
- [26] Wenhao Sun, Deng Liu, Zhiwei Zou, Wendi Sun, Song Chen, and Yi Kang. 2023. Sense: Model-hardware codesign for accelerating sparse CNNs on systolic arrays. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 31, 4 (2023), 470–483.
- [27] Minjin Tang, Mei Wen, Jianchao Yang, Zeyu Xue, and Junzhong Shen. 2024. SP5A: Exploring Sparse-Packing Computation on Systolic Arrays From Scratch. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* (2024).
- [28] Shikhar Tuli, Chia-Hao Li, Ritvik Sharma, and Niraj K. Jha. 2023. CODEBench: A Neural Architecture and Hardware Accelerator Co-Design Framework. *ACM Trans. Embed. Comput. Syst.* 22, 3, Article 51 (April 2023), 30 pages. doi:10.1145/3575798
- [29] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (Long Beach, California, USA) (NIPS'17)*. Curran Associates Inc., Red Hook, NY, USA, 6000–6010.
- [30] Tom Verbeure. [n. d.]. Fpdx Library. <https://github.com/tomverbeure/math>.
- [31] Shinji Watanabe, Takaaki Hori, Shigeki Karita, Tomoki Hayashi, Jiro Nishitoba, Yuya Unno, Nelson Enrique Yalta Soplín, Jahn Heymann, Matthew Wiesner, Nanxin Chen, Adithya Renduchintala, and Tsubasa Ochiai. 2018. ESPnet: End-to-End Speech Processing Toolkit. In *Proceedings of Interspeech*. 2207–2211. doi:10.21437/Interspeech.2018-1456
- [32] Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. 2016. Learning structured sparsity in deep neural networks. In *Proceedings of the 30th International Conference on Neural Information Processing Systems (Barcelona, Spain) (NIPS'16)*. Curran Associates Inc., Red Hook, NY, USA, 2082–2090.
- [33] Sam (Likun) Xi, Yuan Yao, Kshitij Bhardwaj, Paul Whatmough, Gu-Yeon Wei, and David Brooks. 2020. SMAUG: End-to-End Full-Stack Simulation Infrastructure for Deep Learning Workloads. *ACM Trans. Archit. Code Optim.* 17, 4, Article 39 (Nov. 2020), 26 pages. doi:10.1145/3424669
- [34] Rui Xu, Sheng Ma, Yang Guo, and Dongsheng Li. 2023. A Survey of Design and Optimization for Systolic Array-based DNN Accelerators. *ACM Comput. Surv.* 56, 1, Article 20 (Aug. 2023), 37 pages. doi:10.1145/3604802
- [35] Jiecao Yu, Andrew Lukefahr, David Palfaman, Ganesh Dasika, Reetuparna Das, and Scott Mahlke. 2017. Scalpel: Customizing DNN pruning to the underlying hardware parallelism. In *2017 ACM/IEEE 44th Annual International Symposium on Computer Architecture (ISCA)*. 548–560. doi:10.1145/3079856.3080215