

INR-Arch: A Dataflow Architecture and Compiler for Arbitrary-Order Gradient Computations in Implicit Neural Representation Processing

Stefan Abi-Karam^{*1,2}, Rishov Sarkar^{*1}, Dejie Xu³, Zhiwen Fan³, Zhangyang Wang³, Cong Hao¹
Georgia Institute of Technology¹, Georgia Tech Research Institute², University of Texas at Austin³
{stefanabikaram, rishov.sarkar, callie.hao}@gatech.edu, {dejia, zhiwenfan, atlaswang}@utexas.edu

Abstract—An increasing number of researchers are finding use for n^{th} -order gradient computations for a wide variety of applications, including graphics, meta-learning (MAML), scientific computing, and most recently, implicit neural representations (INRs). Recent work shows that the gradient of an INR can be used to edit the data it represents directly without needing to convert it back to a discrete representation. However, given a function represented as a computation graph, traditional architectures face challenges in efficiently computing its n^{th} -order gradient due to the higher demand for computing power and higher complexity in data movement. This makes it a promising target for FPGA acceleration. In this work, we introduce INR-Arch, a framework that transforms the computation graph of an n^{th} -order gradient into a hardware-optimized dataflow architecture. We address this problem in two phases. First, we design a dataflow architecture that uses FIFO streams and an optimized computation kernel library, ensuring high memory efficiency and parallel computation. Second, we propose a compiler that extracts and optimizes computation graphs, automatically configures hardware parameters such as latency and stream depths to optimize throughput, while ensuring deadlock-free operation, and outputs High-Level Synthesis (HLS) code for FPGA implementation. We utilize INR editing as our benchmark, presenting results that demonstrate 1.8–4.8 $\times$ and 1.5–3.6 $\times$ speedup compared to CPU and GPU baselines respectively. Furthermore, we obtain 3.1–8.9 $\times$ and 1.7–4.3 $\times$ lower memory usage, and 1.7–11.3 $\times$ and 5.5–32.8 $\times$ lower energy-delay product. Our framework will be made open-source and available on GitHub.**

1. INTRODUCTION

Implicit neural representations (INRs) are enjoying great popularity for a variety of use cases, including 3D neural rendering and stylization in augmented and virtual reality (AR/VR) [1, 2], application-agnostic data representation and compression [3, 4, 5], and super-resolution and inpainting for data across various modalities, such as images and videos [3, 6]. A core strength of INRs lies in their capacity for data compression. As an effective, high-fidelity encoding approach for diverse data types, INRs represent a promising path to efficient data management [4, 5].

However, it is critical to recognize the implications of such compact data encoding for computational hardware requirements. Considering the emerging paradigm where memory costs more than computation, a memory-efficient solution can provide high energy and area efficiency. Therefore, it is crucial to develop methods to perform rapid gradient computations without relying on large, memory-intensive hardware.

Meanwhile, hardware designers are embracing dataflow architectures to achieve low latency through overlapping computation kernels within their designs. Streaming in dataflow designs allows for individual processes to work at a much finer granularity than input and output arrays; instead, they can incrementally produce partial outputs or consume partial inputs through first-in-first-out (FIFO) streams of data. When large numbers of these processes are combined in this way, massive latency savings can be achieved effectively by exploiting the throughput of each process.

^{*}Equal contribution.

^{**}<https://github.com/sharc-lab/inr-arch>

Motivated by the need for efficient INR computation and editing, we propose INR-Arch with emphasis on a dataflow architecture and specialized compiler. Our key contributions are as follows:

- ❶ **Dataflow Architecture:** We propose a dataflow architecture based on FIFO-based array streams and a library of optimized computation kernels that operate on array streams. This approach allows for increased memory efficacy and overlapping computations.
- ❷ **Computation Graph Extraction & Optimization:** We propose an automated method to extract the computation graph of the gradient of a PyTorch tensor, along with several lossless optimization techniques to simplify the resulting graph.
- ❸ **Deadlock Analysis and Optimization:** We propose a novel technique to quickly and accurately determine whether a given set of FIFO depths will cause a dataflow design to deadlock.
- ❹ **FIFO Depth Analysis and Optimization:** We extend the deadlock analysis to compute latency estimates based on a set of FIFO depths, and we propose a procedure to quickly determine a reduced set of FIFO depths that lowers memory usage without impacting performance.
- ❺ **Code Generation:** We propose a compiler that uses the processed computation graph and a set of FIFO depths to generate a dataflow architecture that executes the graph in hardware.
- ❻ **Power, Latency, and Memory Improvements vs. CPU & GPU:** We evaluate INR-Arch applied to INR editing targeting an FPGA platform and compare its latency, memory usage, and energy-delay product against CPU and GPU baselines.

2. BACKGROUND AND MOTIVATION

A. High-Order Gradients

Popular machine learning frameworks such as PyTorch and TensorFlow are capable of automatically computing arbitrary-order gradients of a given function through reverse mode automatic differentiation [7]. This involves first representing the function as a computation graph, where each node represents a primitive operation such as elementwise add, transpose, or matrix multiply. The framework then recursively applies the chain rule of differentiation on the graph to obtain a new computation graph, which represents a function whose output is the gradient of the original function. This automatic differentiation process can be repeated recursively to obtain second- and higher-order gradients of a function.

Higher-order gradients play a crucial role in various fields, such as scientific computing, computer graphics, and deep learning. For instance, in scientific computing, higher-order gradients are essential for accurately modeling complex problems in areas such as fluid dynamics [8].

Similarly, in computer graphics, higher-order gradients are traditionally used to render images with high fidelity and realism. More recently, differentiable rendering techniques [9] have been developed to work along side deep learning to incorporate graphics rendering




Fig. 1: A visual overview of **A)** Implicit Neural Representations (INRs), and **B)** INR Editing using the INSP-Net architecture [12].

in the end-to-end training pipeline. Neural radiance fields (NeRF) [1] are a popular example of differentiable rendering for deep learning.

Moreover, the use of higher-order gradients is becoming increasingly popular in the field of deep learning. Higher-order gradients are used in the traditional training of models as well as for meta-learning such as model-agnostic meta-learning (MAML) [10] and hyperparameter optimization [11]. Additionally, higher-order gradients have been shown to be an effective tool for processing implicit neural representations (INRs) to apply arbitrarily learnable data transformations efficiently.

B. Implicit Neural Representations

Implicit neural representations (INRs) are a way to represent individual data points as entire neural networks, as shown in Fig. 1A. Given a single data sample, such as an image, audio file, 3-D model, etc., and a suitable neural network architecture [3], the data sample can be *encoded* as a set of weights and biases for the neural network and later *decoded*, i.e., reconstructed, from the weights.

Encoding an INR involves training the chosen neural network architecture to predict output coordinates from input coordinates within a single data sample, effectively overfitting the neural network to this one sample. For instance, to encode an image file, we first consider the image to be a “training dataset” for the neural network consisting of input (x, y) 2-D coordinate pairs mapped to 3-D outputs, representing the red, green, and blue (RGB) colors of the pixel at the input (x, y) coordinates within the image. After training the neural network to predict these mappings, the neural network weights and biases can themselves be considered an *implicit* representation of the image in *weight-space*, i.e., an INR.

It follows that decoding an INR involves plugging in discrete input coordinates into a neural network with weights and biases given by the INR to obtain output coordinates; in the case of images, this means plugging in (x, y) coordinate values and obtaining RGB color values in *pixel-space*.

INRs are useful for several purposes. First, since the inputs and outputs of the INR are continuous coordinate values, the INR can be treated as a continuous representation of a discrete input sample, allowing for super-resolution beyond that of the input sample. For instance, given an image encoded as an INR, during decoding, we can plug in non-pixel-aligned input coordinates and obtain output colors corresponding to points between pixels of the original image, effectively providing unlimited image resolution. Second, INRs can be used as an effective compression scheme, as the weights in an INR can require less space than the original data while still maintaining high fidelity [4, 5, 13]. Third, INRs are a universal representation of any type of data that can be represented as a mapping from input coordinates to output coordinates, making them versatile for compression and super-resolution for a wide variety of data formats.

C. INR Editing

A recent work by Xu *et al.* [12] demonstrates that for images encoded as INRs, we can operate directly on the *weight-space representation* of the image to obtain another INR, using a different neural network architecture dubbed **INSP-Net** (shown in Fig. 1B), whose decoding corresponds to a desired signal processing transformation of the original image in *pixel-space*, such as blurring, de-noising, etc. In other words, if we want to edit an image encoded as an INR by, e.g., blurring it, we do not have to decode the INR to pixel-space, apply a blur filter, then re-encode to another INR. By computing the model output and up to the n^{th} -order gradients of the output as input features for a trainable MLP, specific signal processing tasks can be achieved on a distribution of data. However, computing higher-order gradients needed for editing is complex, resulting in exponentially more complex computation graphs as the number of gradients increases. This provides a key motivation for hardware acceleration of the exact gradient computation.

3. PROPOSED METHODOLOGY

With this motivation in mind, we propose **INR-Arch**, a dataflow architecture and compiler for arbitrary-order gradient computations in INR processing. We first discuss the INR-Arch dataflow architecture, followed by our compiler flow that translates gradient computations in PyTorch to a synthesizable and performant HLS design.

A. Dataflow Architecture

1) Challenges

INR editing presents several significant challenges in translating the computation graph to efficient hardware, which we outline below:

- **Many Intermediate Results with Redundant Data Movement:**

The conventional method of buffering intermediate results into scratchpad memory becomes infeasible due to the large computation graph size and the required batch size for effective INR usage. Allocating a buffer for each computation kernel in the graph could dramatically increase memory requirements. Specifically, INR model inference usually involves sampling multiple coordinates simultaneously to reconstruct data points, like pixel locations in an image. Consequently, the models demand an input batch size dimension, which propagates through all computation kernels. If the batch size is substantial, such as 64, it can inflate the memory requirement of scratchpad memories by 64 times.




Fig. 2: An overview of the INR-Arch framework for end-to-end hardware acceleration for INR editing based on the INSP-Net [12] architecture.




Fig. 3: Illustration of the `array_stream` data structure, the library of stream-based kernels, and an example compute graph mapped to a dataflow architecture.

- **Different Computation Kernels with Different Computation Patterns:** Inherent in the INR model is the use of diverse computation kernels, each exhibiting unique computation patterns. The diverse nature of these kernels means they may process data in distinct ways—some might favor sequential processing, while others may benefit from parallelization. This variance introduces

a layer of complexity in optimizing the model as a whole and developing accelerated computation kernels. Effective use of these kernels requires careful coordination and optimal resource allocation to balance the computational load and data movement while minimizing area and latency requirements.

2) Solution

To address these challenges, we propose a **dataflow architecture** for mapping INR editing models to hardware. A visual overview is shown in Fig. 3.

Our dataflow architecture is based on two key components: streaming-based data movement using a proposed *array stream* data structure and a library of computational kernels designed to operate on array streams.

Streams address the issue of “Many Intermediate Results with Redundant Data Movement.” They are conceptualized and physically implemented as fixed-size First-In-First-Out (FIFO) streams with a user-definable depth. Our unique variant, called “array streams,” includes additional metadata about array shape, stream sizes, and block size of represented data, thereby facilitating the structured streaming of intermediate results. Array streams are designed to stream data in row-major order.

This model is advantageous as it allows inputs, outputs, and intermediate activations to be stored as streams rather than relying on buffers such as scratchpad memory. The streams only need to store a fraction of the elements for any given input, output, or intermediate activation, resulting in a memory-efficient implementation compared to traditional buffered computations in CPUs and GPUs. The quantity of data that can be accommodated in the hardware is determined by the FIFO depth. Generally, we find that the FIFO depth can be significantly smaller than the total elements represented by the array stream, leading to substantial memory savings as outlined in Sec. 4.4.

Computational kernels, or compute units designed to interact with data streams, benefit from the array stream’s unified interface. Each kernel is specialized to read and write data in its unique pattern. For instance, some kernels can instantly read and write computed data without buffering (e.g., elementwise add), while others may necessitate buffering (e.g., matrix multiply / MM) or access to array

shape data (e.g., dimension select). Kernels are also categorized by their input-output degree: N:1, 1:1, and 1:N. INR-Arch incorporates a subset of kernels necessary for supporting operations within INR-specific autograd computation graphs (refer to the source code for further exploration of all kernels).

When integrating array streams and computation kernels, the proposed dataflow architecture adheres to the “one-producer, one-consumer” principle. This necessitates that N:1 and 1:1 kernels be capable of mapping their outputs to inputs of downstream computation kernels while following the “one-producer, one-consumer” rule. To achieve this, a special 1:N operation known as “copy stream” is used to multicast a single input stream’s elements to multiple output streams in a round-robin fashion.

The bottom panel in Fig. 3 shows an example of a mapped dataflow architecture for a small computation graph. In general, our work applies this dataflow architecture to larger extracted computation graphs, mapping inputs, outputs, and intermediate activations to array streams and operations to computation kernels.

B. Compiler Methodology

1) Challenges

The dataflow architecture provides a solid foundation for an efficient accelerator, but mapping a gradient computation graph onto this structure presents its own challenges:

- **Complex Computation Graphs with Redundant Operations:** In applying the INSP-Net approach, we build computation graphs by calculating higher-order gradients of the base INR model being edited. This process causes the computation graphs to grow exponentially with each gradient order. Both the base model and the higher-order composition graph share the same computations and redundant sub-graphs. Furthermore, there is an increase in redundant operations within these computation graphs, along with patterns of operations that can cancel each other out, leading to higher redundancy.
- **Susceptibility to Deadlock:** Due to the differing computation patterns of different computation kernels, the generated dataflow architecture is susceptible to deadlock unless FIFO buffers between kernels are carefully provisioned. It is critical to ensure that the generated design will not deadlock, but deadlocks are usually difficult to detect without a full cycle-level simulation, which can take hours or even days for the complex dataflow designs generated by our framework.
- **Latency or Memory Waste from Improper FIFO Buffer Sizing:** Deadlock-free operation is a necessary but not sufficient criterion for an efficient accelerator. Even when there is no deadlock, too-small FIFO depths can degrade performance so the resulting latency is multiple times slower than peak performance. On the other hand, too-large FIFO depths can consume multiple times the memory resources of an equally performant smaller design.
- **Complexity, Correctness, and Runtime Overhead of HLS Code:** Code generation can be an error-prone process. It is important to ensure that the generated code faithfully reproduces the gradient computation carried out by PyTorch while incurring minimal runtime overhead. The generated code should be as simple as possible to aid debugging and minimize the chance of errors.

We propose a four-step compilation process (represented in Fig. 2 as steps 2–5) to address each of these challenges.

2) Computation Graph Extraction & Optimization

The first step of our proposed process is to obtain the computation graph of the higher-order gradient of a desired function expressed as




Fig. 4: Visualization of the computation graph merging optimization. Similar computations are indicated with identical colors to represent their presence both within and across graphs. The merging of these graphs effectively minimizes redundant computations.

a series of PyTorch operations. We take advantage of the computation graph that is automatically built by PyTorch for its automatic differentiation process (“autograd”) as described in Sec. 2.1.

Given a list of PyTorch tensors representing the gradient outputs, we perform a depth-first traversal through the autograd graph of each of the tensors. We construct a combined computation graph from all the output tensors and apply several optimization passes to eliminate redundancy in the graph.

First, since the gradient introduces repeated subsections of the graph due to the chain rule of differentiation, we de-duplicate any common subtrees within the raw graph, indicated by the color-coded sections of Fig. 4. As a result of this de-duplication, the output tensors across multiple gradient orders share most of their computation: for instance, the outputs for the 1st-order gradient are contained entirely within the computation graph of the 2nd-order gradient, with the exception of a few nodes at the end.

Second, the graph can contain “Permute” nodes, which perform an arbitrary permutation of the axes of the input tensor. However, in many cases, these “Permute” nodes simply swap the axes of a two-dimensional input, which is the same as transposing the input. Therefore, when we identify this special case anywhere in the graph, we replace the “Permute” node with a “T” (transpose) node.

Third, since transposing a tensor twice is the same as not modifying it at all, we look for any contiguous sequences of “T” nodes in the graph and remove all matched pairs, leaving zero or one “T” node in place of each sequence.

Finally, when multiple “T” nodes have the same input, we choose one of them to be the canonical node, delete the others, and re-route their outputs to come from the canonical node.

These optimizations massively simplify the graph. De-duplication greatly shrinks the graph size, making it feasible to synthesize accelerators for larger gradient computations. “T” node optimizations help reduce latency significantly, since transposing a tensor requires buffering the entire tensor and thus creates a bottleneck in the dataflow.




Fig. 5: An example of a computation graph that causes deadlock with default FIFO sizing for any non-trivial input. The root cause is the contention between the “Mm” which buffers elements with a delay before writing out data and “Cos” which writes out data every cycle.

3) Deadlock Analysis

Given an optimized computation graph, we must determine suitable buffer sizes for the FIFO streams connecting each kernel to avoid a deadlock in the overall design. To clarify how this issue arises, Fig 5 depicts an example computation graph that is susceptible to deadlock.

Two nodes, Mm and Cos, use the same input and feed the same output, but Cos operates in a fully streaming manner—producing each output element as soon as each input element is available—whereas Mm must fully buffer all the elements from this input before it can produce any output elements. The source node distributes outputs to Mm and Cos in a round-robin fashion, first writing one element to Mm, then the same element to Cos, repeating until all elements are written to both streams. Similarly, the Mul node reads input elements round-robin, reading one element from Mm, then one element from Cos, repeating until both streams are exhausted.

If all FIFOs use their default depth of 2 and there are more than five outputs from the source node, this computation graph is guaranteed to cause a deadlock:

- 1) Mul will first attempt to read an element from the output of Mm.
- 2) However, Mm will not produce an output until it reads all the elements from the source node.
- 3) Meanwhile, Cos will attempt to write its outputs to Mul, which is blocked waiting for Mm’s output; thus, after two output elements, the output stream for Cos will become full, blocking Cos from consuming more elements from its input stream.
- 4) As a result, when the source node attempts to write the fifth output element to the input of Cos, it will stall, thus preventing Mm from receiving any more input.

All four nodes in the computation graph become stalled waiting for each other cyclically, resulting in deadlock.

In this simple example, it is easy to see the cause of the deadlock and to determine a resolution: increase the stream depth of Cos’s input to the total number of elements. However, the computation graphs for higher-order gradients can contain hundreds of nodes, thereby introducing complex dependency chains that cannot be analyzed by hand. Repeated simulation of the dataflow design with different FIFO depths is also infeasible, as the number of FIFOs involved in such a large computation graph leads to a massive design space. Thus we need a systematic approach to detecting and resolving deadlocks.

Our proposed solution is a *dataflow graph* where nodes represent individual FIFO I/O operations (reads and writes) and directed edges represent “happens-before” relations. This graph encodes the entire behavior of a dataflow architecture with a given set of FIFO depths,




Fig. 6: An example showing how the dataflow graph is constructed and used to detect deadlocks by searching for cycles. This example involves two FIFOs, A and B, both with depth 2. Green nodes represent FIFO writes; red nodes represent FIFO reads.

and thus it can be used to determine precisely whether or not there will be a deadlock for some set of FIFO depths.

Fig. 6 shows a simple example of the construction of this graph. In this example dataflow design, a producer process writes to two streams which are then read by a consumer process. Stream A transfers three data elements, A_0 , A_1 , and A_2 , and stream B transfers one, B_0 . Both streams have FIFO depth 2.

We start by determining the ordering of FIFO reads and writes within each process, as shown in Fig. 6(a). To obtain this ordering, we run our dataflow design through LightningSim [14], a trace-based cycle-level simulator for HLS designs. The trace that LightningSim generates internally precisely orders all FIFO operations on a function-by-function basis. FIFO operations that must occur at the same time are grouped into one node, and edges connect nodes in the order defined by the trace. This trace only needs to be generated once for a given design, as the trace order is independent of the FIFO depths.

Then, in Fig. 6(b), we encode read-after-write (RAW) dependencies into the graph by adding edges connecting each write to its corresponding read: read $\#n$ from stream X cannot occur before write $\#n$ to stream X. This establishes the ordering of nodes between dataflow processes. As with Fig. 6(a), this is independent of FIFO depths and only needs to be done once for a given design. The resulting dataflow graph can be interpreted as the dataflow graph for a design where the FIFO depths are “infinite” or unconstrained.

Fig. 6(c) shows the encoding of write-after-read (WAR) dependencies into the graph. WAR dependencies are caused by limited FIFO depths: if a stream X has a depth of d , after d writes to the stream, the stream will be full unless or until at least one read has occurred from the stream. Therefore, write $\#d$ depends on read $\#0$. Following similar logic, it follows that any write $\#n$ where $n \geq d$ depends on read $\#(n - d)$. In Fig. 6(c), with both FIFO depths set to 2, only

write A_2 depends on read A_0 .

Finally, Fig. 6(d) demonstrates the deadlock detection algorithm, which is equivalent to finding cycles in the graph. Since edges represent “happens-before” relations, cycles represent that a node must happen before itself for the computation to proceed, which clearly represents a deadlock. In the figure, the write to A_2 must occur before the write to B_0 (by intra-process order), the write to B_0 must occur before the read from B_0 (by RAW dependency), the read from B_0 must occur before the read from A_0 (by intra-process order), and the read from A_0 must occur before the write to A_2 (by WAR dependency).

To resolve a deadlock, the depths of one or more of the streams with a WAR dependency in the cycle must be increased. In this example, the only WAR dependency in the cycle involves stream A , whose depth must be increased from 2 to 3 to resolve the deadlock.

Different combinations of stream depths can be quickly tested for deadlock by starting from the unconstrained graph, containing only intra-process and RAW dependencies, then adding WAR dependencies according to the stream depths and checking for cycles.

4) FIFO Depth Optimization

Even if we determine a set of FIFO depths that are deadlock-free, it might be far from peak performance, or it might use excessive resources compared to similarly performant designs. We need a procedure to determine the peak performance of the design and find a set of FIFO depths that achieve similar performance without using excessive memory for FIFO buffers.

Luckily, the dataflow graph from Sec. 3.2.3 also allows us to estimate the latency of a dataflow design by assigning a minimum delay to each edge in the graph. We perform a topological sort on the nodes in the graph, then compute each node’s latency as the maximum of its predecessors’ latencies combined with the edge delays. The maximum latency across all nodes in the graph is a very close estimate to the latency of the overall design, excluding stalls incurred by, e.g., off-chip DRAM reads and writes.

Using these latency estimates, we are able to minimize memory usage without impacting performance. We start with the unconstrained graph and compute its latency estimate, which represents the peak performance of the design. Then, one by one, we constrain the depth of each stream to 2—the minimum depth for a FIFO queue—and re-run the latency estimator to see if the constraint changes the overall latency significantly (by more than a threshold α , which is set to 1% in our implementation). If it does, we discard the constraint; otherwise, we accept the new constraint. Once all streams have been evaluated, we run a simulation to determine the actual FIFO depths observed (peak number of FIFO queue slots used at any point in the simulation) under the newly added constraints. We use these observed numbers (with a minimum of 2 for each stream) as our final, optimized set of depths for all FIFOs in the computation graph.

5) Code Generation

The final HLS model is generated (and can be compiled and synthesized) using the code generation component of the presented framework. Code generation is done using a template-based compiler that maps kernels from the INR-Arch hardware library to an HLS implementation of the model using the described dataflow architecture. Most of the implementation is simple initialization for `array_stream` data structures and 1-to-1 mapping of functions in the computation graph to functions in the hardware library.

However, care needs to be taken when mapping hardware kernels to properly insert the hardware kernel calls in the correct topological order, as well as preserve the correct argument order from the

computational graph. Each intermediate activation’s argument order is stored in the associated edge as an edge feature in the processed computations graph, which is then referenced during code generation to generate kernel call argument lists. Care also must be taken to insert `copy_stream` kernels after function calls to effectively “multicast” kernel outputs to the correct downstream kernel inputs. This is done by extracting the edges to successors in the computation graph. These edges then become the edges to which the kernel output is multicast using the `copy_stream` kernel.

The metadata associated with `array_streams` is stored as compile-time information in the `array_stream` struct implementation. The importance of this compile-time information becomes clear when computation kernels access array shape data through `typename` template arguments. This vital information at compile-time during High-Level Synthesis (HLS) can be skillfully utilized within the computation kernels for operations such as unrolling and pipelining of loops, which are dependent upon the array shape and block size specific to an individual `array_stream`, as well as static asserts to check properties about the input arrays (e.g., array sizes for MM). For a more comprehensive overview of using modern C++ features to implement these compile-time design features in HLS, we direct interested readers to [15] as well as our source code.

The Python API for code generation takes in the processed computation graph (Sec. 3.2.2) along with the computed FIFO depths from the deadlock analysis and FIFO depth optimization (Sec. 3.2.3, Sec. 3.2.4). Additionally, the user is able to specify the target FPGA board along with desired fixed-point precision for the implemented HLS model which maps to the Vitis HLS arbitrary-precision fixed-point data structures. The code generation also handles the automated generation, compilation, and execution of a C++ testbench using the fixed-point model. This helps a user experiment and tweak the fixed-point precision used for the best resource usage vs. accuracy trade-off for model inference. Lastly, code generation handles the automated synthesis of the generated HLS model and extraction of synthesis report data for analysis.

4. RESULTS

A. Evaluation Setup

As a case study to evaluate our framework, we measure the performance of two models derived from Xu *et al.* [12], a recent computer vision work that uses high-order gradients of a SIREN model [3] to apply a variety of image transformations, such as blurring or denoising, directly to an image encoded as a SIREN INR.

We evaluate two configurations, namely, the first-order and second-order gradients of the SIREN model as computed in [12], using batch size 64 in both cases. The design for each of these two configurations was generated by the framework and synthesized using Xilinx Vitis HLS for the Xilinx Alveo U50 Data Center Accelerator at 300 MHz. A 32-bit fixed-point format with 10 integer bits was used for the data.

For the first-order model, a hardware parallelism factor of 64 $\times$ was used for all MM operations. However, since the computation graph of the second-order model is so much more complex than that of the first-order model, the second-order model must use a lower parallelism factor of 16 $\times$ for all MMs in order to avoid exceeding available resources on the target device.

FPGA latency results are collected using a highly accurate cycle-level simulator for HLS designs [14], while resource estimates are provided by the HLS tool itself. Baseline results on CPU (Intel Xeon Gold 6226R) and GPU (NVIDIA RTX A6000) were measured directly from the gradient computation code in [12], written using the PyTorch framework.




Fig. 7: Main comparison results for latency, energy-delay product, and memory of 1st-order and 2nd-order INR models between GPU, CPU, and the proposed FPGA implementation. The y-axes are log scales.

TABLE I: Performance comparisons.

Model	Device	Latency (ms)	Memory (MiB)	EDP (J-ms)
1 st Order	CPU	3.34 (1.83x)	7.63 (8.93x)	0.17 (1.67x)
1 st Order	GPU	2.80 (1.53x)	3.64 (4.26x)	0.55 (5.51x)
1 st Order	FPGA	1.83 (1.00x)	0.85 (1.00x)	0.10 (1.00x)
2 nd Order	CPU	12.17 (4.78x)	23.58 (3.07x)	2.20 (11.34x)
2 nd Order	GPU	9.22 (3.62x)	13.08 (1.71x)	6.36 (32.75x)
2 nd Order	FPGA	2.54 (1.00x)	7.67 (1.00x)	0.19 (1.00x)

Comparison factors (parenthesized) are relative to the corresponding FPGA metric.

B. Latency, Energy Efficiency, and Memory

Results are shown in Table I and Table II. Across both first-order and second-order models, the FPGA implementation **beats CPU and GPU baselines in three key metrics: latency, memory usage, and energy-delay product.**

Our first-order gradient model on FPGA achieves a significant speedup over CPU and GPU baselines, but the speedup achieved by our second-order gradient model is even more pronounced, where the generated accelerator achieves nearly 4× speedup over GPU and nearly 5× over CPU.

Notably, Table II shows how, when the same MM parallelism factor is used for different-order gradients, the latencies of the resulting accelerators are very similar. This demonstrates the advantage of our dataflow architecture: because we can overlap most of the kernels in the computation graph, a larger computation graph induced by

TABLE II: Resource usage vs. latency on the Alveo U50.

Model	1 st Order 64x	1 st Order 16x	2 nd Order 16x	
MM Parallelism				
Latency (ms)	1.83	2.55	2.54	Available
BRAM	389 (14%)	233 (9%)	419 (16%)	2,688
DSP	3,343 (56%)	1,039 (17%)	3,889 (65%)	5,952
FF	529k (30%)	277k (16%)	952k (55%)	1,743k
LUT	367k (42%)	234k (27%)	781k (90%)	871k
URAM	—	48 (8%)	192 (30%)	640

TABLE III: Computation graph optimizations.

Optimization	Nodes	Edges	Node Types		
			T	Permute	Other
Original graph	5,531	7,326	438	945	4,148
+ Dedupe common subtrees	459 (−92%)	626 (−91%)	63	5	391
+ Replace “Permute”s → “T”s	459 (±0%)	626 (±0%)	68	0	391
+ Remove “T” pairs	420 (−8%)	587 (−6%)	29	0	391
+ Dedupe common “T”s	396 (−6%)	563 (−4%)	5	0	391

a higher-order gradient does not always mean the latency will be significantly higher. Even when MM parallelism must be reduced for the model to fit within the target device’s resources, it does not result in an increase in latency by the same factor.

(That the 1st-order model with 16× MM parallelism is slightly slower than 2nd-order model with 16× MM parallelism may initially appear erroneous, given that the 1st-order computation graph is a subset of the 2nd-order graph. However, it is a result of our FIFO depth optimization process and will be explained in Sec. 4.4.)

We also see significant memory savings over CPU and GPU baselines, about 9× less memory than CPU and 4× less memory than GPU on the 1st-order model and about 3× and 2× less than CPU and GPU on the 2nd-order model.

Our framework demonstrates its strongest advantage in energy efficiency over CPU and GPU baselines: our model achieves an energy-delay product over 11× lower than CPU and nearly 33× lower than over GPU on the 2nd-order model, thanks to the combination of low latency and low power achieved by our FPGA design.

C. Graph Optimization

We perform an ablation study of our computation graph optimization techniques described in Sec. 3.2.2 and report our findings in Table III. The most significant optimization is the de-duplication of common subtrees in the graph, which accounts for over 90% reduction in both nodes and edges over the unoptimized graph. However, the other optimizations we perform result in significant drops in the number of “Permute” and “T” nodes, collectively dropping their combined total from 68 nodes to just 5. This minimizes bottlenecks in the dataflow computation, as “Permute” and “T” both require buffering the entire input stream before writing outputs.

D. FIFO Depth Optimization

We also evaluate the effectiveness of the FIFO depth optimization scheme described in Sec. 3.2.4 in reducing memory usage. We consider two metrics: the latency of the model and the sum of FIFO depths, which acts as a proxy for the memory consumed by the FIFOs. We evaluate each metric both before and after optimization, where the set of FIFO depths before optimization is determined as the depths actually observed (with a minimum of 2 for each stream) when we run a simulation with all FIFO depths unconstrained (i.e., a simulation of peak performance).

TABLE IV: Before and after FIFO depth optimization.

Model	MM	Before Optimization		After Optimization	
		Latency (ms)	$\sum$ Depths	Latency (ms)	$\sum$ Depths
1 st Order	64×	1.823	125,586	1.828 (+0.3%)	15,579 (-87.6%)
1 st Order	16×	2.538	125,661	2.551 (+0.5%)	15,643 (-87.6%)
2 nd Order	16×	2.545	668,601	2.545 (+0.0%)	96,808 (-85.5%)


MM || = MM parallelism; $\sum$ Depths = Sum of FIFO depthsFig. 8: A trace of FIFO reads for a representative subset of hardware computation kernels in the main dataflow region of the Base INR + 1st Order Gradient INR-DSP model.

Table IV shows our results. In all three cases evaluated, we achieve over 85% reduction in FIFO depths with less than 1% degradation over peak performance.

These results also explain why the 1st-order model with 16×\times MM parallelism runs slightly slower than the 2nd-order model with 16×\times MM parallelism, despite the 1st-order graph being a subset of the 2nd-order graph. At peak performance, the 1st-order model is slightly faster; however, the FIFO depths selected for these two models by the optimization process in Sec. 3.2.4 end up causing the final latency of the 1st-order model to slightly exceed the final latency of the 2nd-order model. This can be avoided by adjusting the acceptable threshold α during depth optimization.

E. Dataflow Trace Visualization

Novel simulation tools [14] are used to dump and inspect simulation traces to analyze FIFO read and FIFO writes and better understand data movement along `array_streams`.

The FIFO reads over time during computationally intensive operations, mainly matrix multiplication (MM), are shown in Fig. 8. Due to the ordering of dependencies in the computation graph, it is clear when some MM operations are computing in parallel, as well as when data is being stalled periodically for downstream computation kernels. Work is ongoing to show other complex simulation behavior of the dataflow to better understand FIFO depths over time for better FIFO sizing and deadlock detection.

5. CONCLUSION

In this paper, we introduced INR-Arch, a framework for dataflow architectures of n^{th} -order gradient computations. This addresses the challenges that traditional architectures encounter when computing

higher-order gradients efficiently. We centered our evaluation application on INR editing and compared our framework against CPU and GPU baselines. We demonstrated significant speed improvements, decreased memory usage, and a lower energy-delay product than both the CPU and GPU baselines.

Future work involves extending our evaluation to include higher-order gradients, examining the applicability of our framework to diverse models, and addressing large, intricate designs like those found in high-performance computing (HPC). These complex designs involve computational kernels or FIFO buffers that may not fit on the board. Furthermore, we plan to continue developing highly optimized and compact model caricatures for additional edge computing applications. By expanding our framework to handle higher-order gradients, we can further illustrate its adaptability and effectiveness across a wider range of applications. Moreover, our goal is to adapt our framework to suit different models, empowering researchers to utilize FPGA acceleration for a multitude of computational tasks beyond the INR editing scenario.

By providing an open-source implementation on GitHub, we invite further exploration, collaboration, customization, and deployment of our framework. This approach can serve the distinct needs of various research domains.

6. ACKNOWLEDGEMENTS

This work and its authors are partially supported by the Center for Research into Novel Computing Hierarchies (CRNCH) at Georgia Tech, the 2022 Qualcomm Innovation Fellowship program, Cisco, and Georgia Tech Research Institute.

REFERENCES

- [1] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, “NeRF: Representing scenes as neural radiance fields for view synthesis.” [Online]. Available: <http://arxiv.org/abs/2003.08934>
- [2] Z. Fan, Y. Jiang, P. Wang, X. Gong, D. Xu, and Z. Wang, “Unified implicit neural stylization.” [Online]. Available: <http://arxiv.org/abs/2204.01943>
- [3] V. Sitzmann, J. N. P. Martel, A. W. Bergman, D. B. Lindell, and G. Wetzstein, “Implicit neural representations with periodic activation functions,” in *Proceedings of the 34th International Conference on Neural Information Processing Systems*, ser. NIPS’20. Red Hook, NY, USA: Curran Associates Inc., Dec. 2020, pp. 7462–7473.
- [4] E. Dupont, A. Golinski, M. Alizadeh, Y. W. Teh, and A. Doucet, “COIN: COMpression with Implicit Neural representations,” in *Neural Compression: From Information Theory to Applications – Workshop @ ICLR 2021*, Apr. 2021.
- [5] E. Dupont, H. Loya, M. Alizadeh, A. Golinski, Y. W. Teh, and A. Doucet, “COIN++: Neural compression across modalities,” *Transactions on Machine Learning Research*, Dec. 2022.
- [6] Z. Chen, Y. Chen, J. Liu, X. Xu, V. Goel, Z. Wang, H. Shi, and X. Wang, “VideoINR: Learning video implicit neural representation for continuous space-time super-resolution.” [Online]. Available: <http://arxiv.org/abs/2206.04647>
- [7] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, “Automatic differentiation in machine learning,” *Journal of Machine Learning Research*, vol. 18, no. 153, pp. 1–43, 2018.
- [8] J. D. Anderson, *Computational fluid dynamics: the basics with applications*, ser. McGraw-Hill series in mechanical engineering. McGraw-Hill, 1995.
- [9] H. Kato, D. Beker, M. Morariu, T. Ando, T. Matsuoka, W. Kehl, and A. Gaidon, “Differentiable rendering: A survey.” [Online]. Available: <http://arxiv.org/abs/2006.12057>
- [10] C. Finn, P. Abbeel, and S. Levine, “Model-agnostic meta-learning for fast adaptation of deep networks.” [Online]. Available: <http://arxiv.org/abs/1703.03400>
- [11] D. Maclaurin, D. Duvenaud, and R. P. Adams, “Gradient-based hyperparameter optimization through reversible learning.” [Online]. Available: <http://arxiv.org/abs/1502.03492>

- [12] D. Xu, P. Wang, Y. Jiang, Z. Fan, and Z. Wang, "Signal processing for implicit neural representations," in *Advances in Neural Information Processing Systems*, Oct. 2022.
- [13] E. Dupont, H. Kim, S. M. A. Eslami, D. J. Rezende, and D. Rosenbaum, "From data to functa: Your data point is a function and you can treat it like one," in *Proceedings of the 39th International Conference on Machine Learning*. PMLR, Jun. 2022, pp. 5694–5725.
- [14] R. Sarkar and C. Hao, "LightningSim: Fast and accurate trace-based simulation for High-Level Synthesis," in *2023 IEEE 31st Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*. Marina Del Rey, CA, USA: IEEE, May 2023.
- [15] S. Lahti, M. Rintala, and T. D. Hämäläinen, "Leveraging modern c++ in high-level synthesis," vol. 42, no. 4, pp. 1123–1132, conference Name: IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems.