



HAL
open science

Multi-Level GNN Preconditioner for Solving Large Scale Problems

Matthieu Nastorg, Jean-Marc Gratien, Thibault Faney, Michele Alessandro Bucci, Guillaume Charpiat, Marc Schoenauer

► **To cite this version:**

Matthieu Nastorg, Jean-Marc Gratien, Thibault Faney, Michele Alessandro Bucci, Guillaume Charpiat, et al.. Multi-Level GNN Preconditioner for Solving Large Scale Problems. 2024. hal-04447099

HAL Id: hal-04447099

<https://inria.hal.science/hal-04447099v1>

Preprint submitted on 12 Feb 2024

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 Attribution 4.0 International License

Multi-Level GNN Preconditioner for Solving Large Scale Problems

Matthieu Nastorg
Université Paris-Saclay
CNRS, Inria, LISN
91405 Orsay, France
matthieu.nastorg@inria.fr

Jean-Marc Gratien
Direction Technologie et Numerique
IFP Énergies-Nouvelles
92852 Rueil-Malmaison, France

Thibault Faney
Direction Technologie et Numerique
IFP Énergies-Nouvelles
92852 Rueil-Malmaison, France

Michele Alessandro Bucci
Digital Sciences & Technologies Department
Safran Tech
78114 Châteaufort, France

Guillaume Charpiat
Université Paris-Saclay
CNRS, Inria, LISN
91405 Orsay, France

Marc Schoenauer
Université Paris-Saclay
CNRS, Inria, LISN
91405 Orsay, France

Abstract—Large-scale numerical simulations often come at the expense of daunting computations. High-Performance Computing has enhanced the process, but adapting legacy codes to leverage parallel GPU computations remains challenging. Meanwhile, Machine Learning models can harness GPU computations effectively but often struggle with generalization and accuracy. Graph Neural Networks (GNNs), in particular, are great for learning from unstructured data like meshes but are often limited to small-scale problems. Moreover, the capabilities of the trained model usually restrict the accuracy of the data-driven solution. To benefit from both worlds, this paper introduces a novel preconditioner integrating a GNN model within a multi-level Domain Decomposition framework. The proposed GNN-based preconditioner is used to enhance the efficiency of a Krylov method, resulting in a hybrid solver that can converge with any desired level of accuracy. The efficiency of the Krylov method greatly benefits from the GNN preconditioner, which is adaptable to meshes of any size and shape, is executed on GPUs, and features a multi-level approach to enforce the scalability of the entire process. Several experiments are conducted to validate the numerical behavior of the hybrid solver, and an in-depth analysis of its performance is proposed to assess its competitiveness against a C++ legacy solver.

Index Terms—Graph Neural Networks, Multi-Level Domain Decomposition, Partial Differential Equations, Hybrid Solvers

I. INTRODUCTION

Partial Differential Equations (PDEs) are essential for understanding complex physical or artificial processes in science and engineering. However, solving these equations at scale can be challenging due to the computational cost of resolving the smallest spatio-temporal scales. The steady-state Poisson equation is one of the most common and widely used PDEs. It is ubiquitous [6, 18, 17, 10], and plays a critical role in modern numerical solvers. Despite significant progress in the High-Performance Computing (HPC) community, the solution to the Poisson problem remains a major bottleneck in many numerical simulation processes [27].

Nowadays, data-driven methods like Deep Neural Networks (DNNs) are reshaping the field of numerical simulations.

These deep networks can offer faster predictions, thereby reducing turnaround time for workflows in engineering and science [28, 13]. Early attempts to apply DNNs to solve PDEs involved treating simulation data as images and leveraging Convolutional Neural Networks (CNNs) [20, 9]. CNN models have demonstrated remarkable performance in processing structured, image-like data. However, their effectiveness is limited when dealing with unstructured data, such as meshes encountered in numerical simulations. To overcome this limitation, researchers have turned to Graph Neural Networks (GNNs), which have displayed great potential in solving PDEs on unstructured grids [21, 8]. Using data-driven methods to solve PDEs offers several benefits compared to traditional solvers [23]. Machine Learning (ML) models usually function as “black-box” approaches, approximating the solution without the need for the costly computations of creating and solving the system of equations. Data-driven methods are also well-suited for parallel computations on GPUs, unlike traditional methods that rely on CPU computations with limited parallelization on multiple cores. Conversely, ML methods can suffer from various issues, such as:

Generalization ML models can provide approximate solutions to problems within a specific training distribution, but out-of-distribution samples are often poorly solved. For instance, GNN models for PDEs are usually constrained to small-scale problems with fixed-size meshes, hindering their practical use in industrial contexts.

Level of Accuracy The precision of an ML model is limited. It is as precise as the capabilities of the trained model permit and often deteriorates further when handling out-of-distribution samples. When a very accurate PDE solution is needed, leveraging an ML model can quickly be ill-suited (e.g., accurately solving a Poisson Pressure problem to ensure the consistency of a fractional step method [6]).

In short, traditional solvers guarantee convergence but are

limited to CPU computations. On the other hand, data-driven methods harness the full power of parallel GPU computations but lack generalization and convergence guarantees. To address the generalization capabilities of ML models, a potential solution involves building hybrid solvers that combine Machine Learning with Domain Decomposition Methods (DDMs), which is a recent and promising field [7, 11]. Some research has focused on using ML models to enhance DDMs by approximating optimal interface conditions [12, 26]. Other work involves using different types of DNNs to replace the local sub-domain or coarse solvers. In [16] and [15], the sub-domain solvers in a parallel overlapping Schwarz method are replaced by PINNs [22] and the Deep Ritz method [29], respectively. While these methods yield satisfactory outcomes, they are typically limited to solving PDEs on small-scale Cartesian data. Extending them for large-scale problems with varying shapes and sizes remains elusive, as does fair evaluation of their performance in real industrial contexts. Besides, these approaches still operate as entire data-driven frameworks with no guarantee of converging to any desired precision.

We address these limitations with a novel GNN-based preconditioner, combining GNN solver techniques with multi-level Domain Decomposition. This innovative preconditioner mirrors the structure of an Additive Schwarz Method [3], with GNN models tackling multiple sub-problems. This flexibility allows efficient handling of large-scale meshes with varying sizes and shapes by adapting sub-problem sizes to GNN capabilities. Leveraging GPU parallelism, the method concurrently solves multiple sub-problems in batches. This results in a hybrid solver that can achieve high precision convergence through Krylov methods while significantly boosting their efficiency thanks to the GNN preconditioner, efficiently executed on GPUs. Additionally, the preconditioner can incorporate a two-level method to enhance weak scalability regarding the number of sub-domains.

In this paper, Section II first reviews the main principles of the algorithms used in the fields of Domain Decomposition and statistical GNN-based solvers. Then, Section III introduces the hybrid methodology, which combines algorithms from the two fields to overcome their limitations. Section IV showcases the results validating that our approach is numerically accurate. The performance of our method is benchmarked against that of an optimized C++ legacy linear solver to determine if the methodology is competitive with a state-of-the-art solver used in real CFD software. Finally, we conclude our work and provide perspectives for future research.

II. BACKGROUND

This section provides background materials, reviewing the multi-level Additive Schwarz Preconditioner in Section II-A, and Graph Neural Networks in Section II-B. Let us first state

the problem addressed in the entire paper.

Let $\Omega \subset \mathbb{R}^n$ be a bounded open domain with smooth boundary $\partial\Omega$. Let f be a continuous function defined on Ω , and g a continuous function defined on $\partial\Omega$. In this paper, we focus on solving Poisson problems with Dirichlet boundary conditions, which consists in finding a real-valued function u , defined on Ω , solution of:

$$\begin{cases} -\Delta u = f & \in \Omega \\ u = g & \in \partial\Omega \end{cases} \quad (1)$$

Except in very specific instances, no analytical solution can be derived for the Poisson problem, and its solution must be numerically approximated: the domain Ω is first discretized into an unstructured mesh, denoted Ω_h . The Poisson equation (1) is then spatially discretized using the Finite Element Method. The approximate solution is sought as a vector of values defined on all N degrees of freedom of Ω_h . In this work, we use first-order Lagrange elements, thus N matches the number of nodes in Ω_h . The discretization of the variational formulation of (1) using Galerkin's method results in solving a linear system of the form:

$$A\mathbf{u} = \mathbf{b} \quad (2)$$

where the matrix $A \in \mathbb{R}^{N \times N}$ is sparse and represents the discretization of the continuous Laplace operator, the vector $\mathbf{b} \in \mathbb{R}^N$ comes from the discretization of the forcing term f and of the boundary conditions g , and $\mathbf{u} \in \mathbb{R}^N$ is the solution vector to be sought.

For complex industrial problems, accurate predictions are mandatory, leading to an increasingly large N and, thus, an extensive system (2) to solve. Iterative solvers, particularly Krylov methods like Conjugate Gradient (CG), BiCGStab, and GMRES [23], are favored for their superior convergence compared to stationary algorithms. However, they still face challenges related to efficiency and scalability. Efficiency is tied to the convergence rate, where a higher rate requires fewer iterations to achieve a desired precision, although this rate decreases with larger problem sizes due to matrix conditioning. Scalability concerns how solution time responds to increased computational resources, ideally remaining constant for a fixed problem size per core.

A. Multi-level Additive Schwarz Preconditioner

Preconditioning can significantly improve the efficiency and scalability of Krylov methods. In practice, the rationale is to find a matrix $M \in \mathbb{R}^{N \times N}$ and to use a Krylov method to solve the following preconditioned problem:

$$M^{-1}A\mathbf{u} = M^{-1}\mathbf{b} \quad (3)$$

Preconditioners are generally designed to improve the efficiency of iterative methods, but not all can enhance their scalability. In this context, we review the multi-level Additive Schwarz Method (ASM), which offers both properties. ASM

belongs to Schwarz methods [24] from the broad field of Domain Decomposition Methods (DDMs) [1]. DDMs leverage the principle of “divide and conquer”: the global problem (1) is partitioned into sub-problems of manageable size that can be solved in parallel on multiple processor cores. Initially introduced as stationary iterative methods, Schwarz methods are commonly applied as preconditioners for Krylov methods. Let us consider an overlapping decomposition of Ω into K open sub-domains $(\Omega_i)_{1 \leq i \leq K}$ such that:

$$\Omega = \bigcup_{i=1}^K \Omega_i$$

Let $u^n, n \in \mathbb{N}$ be an approximate solution to the Poisson problem (1). As an iterative solver, the one-level ASM computes an updated approximate solution u^{n+1} by first solving, for each sub-domain $i = 1, \dots, K$, the following local sub-problems:

$$\begin{cases} -\Delta v_i^n = r^n & \text{in } \Omega_i \\ v_i^n = 0 & \text{on } \partial\Omega_i \end{cases} \quad (4)$$

where $r^n = f + \Delta u^n$ is the residual at iteration n . The updated approximation u^{n+1} is then obtained using all sub-solutions v_i^n such that:

$$u^{n+1} = u^n + \sum_{i=1}^K E_i(v_i^n) \quad (5)$$

where E_i is an extension operator that maps functions defined on Ω_i to their extension on Ω that takes value 0 outside Ω_i .

From a numerical perspective, a discrete counterpart can be established to solve the system (2). Let us consider now an overlapping decomposition of the mesh Ω_h into K sub-meshes $(\Omega_{h,i})_{1 \leq i \leq K}$ such that:

$$\Omega_h = \bigcup_{i=1}^K \Omega_{h,i}.$$

We denote by k_i the number of nodes in the sub-mesh i . The restriction of a solution vector $\mathbf{u} \in \mathbb{R}^N$ to a sub-mesh $\Omega_{h,i}$ can be expressed as $R_i \mathbf{u}$, where R_i is a rectangular boolean matrix of size $k_i \times N$. The extension operator is defined as the transpose matrix R_i^T . With these notations, the one-level ASM preconditioner, referred to as $M_{\text{ASM},1}^{-1}$, is defined as:

$$M_{\text{ASM},1}^{-1} = \sum_{i=1}^K R_i^T (R_i A R_i^T)^{-1} R_i \quad (6)$$

The one-level ASM preconditioner (6) is not scalable regarding the number of sub-domains. One mechanism to enhance scalability involves implementing a two-level method with coarse space correction, which connects all sub-domains at each iteration. A potential solution is to leverage a Nicolaides coarse space [19], which defines a matrix R_0 of size $K \times N$. We refer the reader to [1] for additional information. The two-level Additive Schwarz preconditioner, referred to as $M_{\text{ASM},2}^{-1}$, is then defined as:

Algorithm 1 Preconditioned Conjugate Gradient

Compute:

$$\mathbf{r}_0 = \mathbf{b} - A\mathbf{u}_0, \quad \mathbf{z}_0 = M_{\text{ASM}}^{-1}(\mathbf{r}_0), \quad \mathbf{p}_0 = \mathbf{z}_0$$

for $i = 0, 1, \dots$ **do**

$$\rho_i = \langle \mathbf{r}_i, \mathbf{z}_i \rangle, \quad \mathbf{q}_i = A\mathbf{p}_i, \quad \alpha_i = \frac{\rho_i}{\langle \mathbf{p}_i, \mathbf{q}_i \rangle}$$

$$\mathbf{u}_{i+1} = \mathbf{u}_i + \alpha_i \mathbf{p}_i$$

$$\mathbf{r}_{i+1} = \mathbf{r}_i - \alpha_i \mathbf{q}_i$$

if $\|\mathbf{r}_{i+1}\| < \text{tol}$ **then** Break;

end if

$$\mathbf{z}_{i+1} = M_{\text{ASM}}^{-1}(\mathbf{r}_{i+1}), \quad \rho_{i+1} = \langle \mathbf{r}_{i+1}, \mathbf{z}_{i+1} \rangle, \quad \beta_{i+1} = \frac{\rho_{i+1}}{\rho_i}$$

$$\mathbf{p}_{i+1} = \mathbf{z}_{i+1} + \beta_{i+1} \mathbf{p}_i$$

end for

$$M_{\text{ASM},2}^{-1} = \underbrace{R_0^T (R_0 A R_0^T)^{-1} R_0}_{\text{coarse problem}} + \underbrace{\sum_{k=1}^K R_k^T (R_k A R_k^T)^{-1} R_k}_{\text{local problems}} \quad (7)$$

The continuous iterative formulation of ASM, defined by equations (4) and (5), can be expressed in a discrete form by the following preconditioned fixed-point iteration:

$$\mathbf{u}^{n+1} = \mathbf{u}^n + M_{\text{ASM}}^{-1}(\mathbf{b} - A\mathbf{u}^n) \quad (8)$$

$$= \mathbf{u}^n + M_{\text{ASM}}^{-1} \mathbf{r}^n \quad (9)$$

where M_{ASM}^{-1} stands for the one or two level ASM preconditioner. Given its symmetric nature, the multi-level ASM preconditioner M_{ASM}^{-1} performs remarkably well when used as a preconditioner for the Preconditioned Conjugate Gradient (PCG) algorithm [1]. The PCG algorithm is reminded in Algorithm 1, including the application of the multi-level ASM preconditioner to residual vectors highlighted in red.

B. Deep Statistical Solvers

This section introduces Deep Statistical Solvers (DSS) [2], a promising GNN approach used in Machine Learning to solve Poisson problems like (1) on unstructured meshes. In this study, the loss function is a “physics-informed” loss defined as the residual of the discretized Poisson problem. This approach differs from most previous ML methods, which aim to minimize the distance between the output of the model and a “ground-truth” solution. Therefore, a DSS model can be trained without a large set of expensive training sample solutions. Moreover, its physics-informed training process allows for better generalization when approaching problems on varying geometries compared to supervised methods.

The fundamental idea of the DSS approach is, considering a discrete Poisson problem like (2), to build a Machine Learning solver, parametrized by a vector θ and denoted as DSS_θ , which outputs an approximate solution $\tilde{\mathbf{u}}$ such that:

$$\tilde{\mathbf{u}} = \text{DSS}_\theta(A, \mathbf{b}) \quad (10)$$

Because the system (2) was constructed using first-order finite elements, the matrix A can be viewed as the adjacency matrix of its corresponding mesh Ω_h . As a result, a discretized Poisson problem (2) with N degrees of freedom can be interpreted as a graph problem $G = (N, A, \mathbf{b})$, where N is the number of nodes in the graph, $A = (a_{ij})_{(i,j) \in [N]^2}$ is the weighted adjacency matrix that represents the interactions between the nodes, and $\mathbf{b} = (b_i)_{i \in [N]}$ is some local external input applied to all the nodes in the graph. Vector $\mathbf{u} = (u_i)_{i \in [N]}$ represents the state of the graph, with u_i being the state of node i . In addition, we define \mathcal{L}_{res} as the real-valued function that computes the mean squared error of the discretized residual equation:

$$\mathcal{L}_{\text{res}}(\mathbf{u}, G) = \frac{1}{N} \sum_{i \in [N]} \left(-b_i + \sum_{j \in [N]} a_{i,j} u_j \right)^2 \quad (11)$$

Let \mathcal{S} and \mathcal{U} be the set of all such graphs G and all states \mathbf{u} that belong to the same distribution of discretized Poisson problems. In [2], the statistical objective is, given a graph G in \mathcal{S} , to find an optimal state in \mathcal{U} that minimizes (11). Therefore, DSS_θ is trained to predict from G a solution U to solve the following statistical problem:

Given a distribution \mathcal{D} on space \mathcal{S} and the loss \mathcal{L}_{res} , solve:

$$\hat{\theta} = \underset{\theta}{\text{argmin}} \mathbb{E}_{G \sim \mathcal{D}} [\mathcal{L}_{\text{res}}(\text{DSS}_\theta(G), G)] \quad (12)$$

DSS_θ uses an iterative architecture, propagating information through the mesh thanks to a manually set number of different Message Passing Neural Network (MPNN) layers. Interestingly, [2] demonstrates, under mild hypotheses, the consistency of the approach. However, it depends on the number of MPNN layers, which should be directly proportional to the diameter of the meshes at hand. While DSS yields efficient results, it is limited to meshes of the same sizes due to its fixed number of MPNN layers. To tackle meshes at a large scale, the number of iterations must be set accordingly, leading to a significant increase in the model size, eventually becoming intractable.

III. METHODOLOGY

Given the limitations outlined in Section II-B, the Deep Statistical Solvers method cannot effectively handle Poisson problems on a large scale. Instead, we propose using a Conjugate Gradient (CG) method preconditioned with a multi-level GNN preconditioner, referred to as DDM-GNN¹. DDM-GNN adopts an Additive Schwarz approach (as described in Section II-A), where the resolution of the local problems (4) is tackled using a DSS model. This approach aims to leverage both traditional and data-driven methods

to solve Poisson problems like (1) at scale. The proposed hybrid solver combines the convergence guarantee provided by the CG method with enhanced efficiency through the use of the DDM-GNN preconditioner. DDM-GNN, in turn, can effectively handle problems of any size and shape by choosing the size of the sub-problems in line with the capabilities of the DSS model. Additionally, DDM-GNN harnesses GPU parallel computations to further enhance its performance and can be equipped with a two-level method to enable the scalability of the entire process. The resolution of a Poisson problem using the proposed hybrid solver is illustrated in Fig.1, and the code is available here².

Section III-A provides a comprehensive explanation of the DDM-GNN preconditioner, followed by a detailed description of the architecture of the DSS model in Section III-B.

A. DDM-GNN Preconditioner

The DDM-GNN preconditioner is intended to replace the ASM preconditioner in Algorithm 1. Following this setup, DDM-GNN should take a global residual vector $\mathbf{r} \in \mathbb{R}^N$ as input, and output a global correction vector $\mathbf{z} \in \mathbb{R}^N$, ensuring the consistency of the PCG algorithm.

In ASM (Section II-A), the local problems (4) differ from the global one (1). Despite this difference, there are still Poisson problems with Dirichlet boundary conditions, and using a DSS framework (Section II-B) to solve these local problems remains consistent. Notably, the force function f in (1) now corresponds to residuals r . It is noteworthy that the Dirichlet boundary conditions of these local problems are of a homogeneous type instead of being represented by some function g in (1), significantly simplifying the distribution of problems addressed by the DSS model.

As the PCG algorithm aims to minimize the residual with each iteration, the source terms of local Poisson problems (i.e., residual vectors \mathbf{r} in PCG) have increasingly smaller norms. This poses a challenge to our methodology: as the residual norm approaches zero, the DSS model might struggle to generalize properly since \mathbf{r} serves as an input to the model. This could result in a trivial solution, i.e., a solution equal to zero everywhere. Consequently, PCG might stagnate at a certain precision threshold, as DDM-GNN preconditioner would consistently provide a null correction to update the solution. To address this issue, we propose normalizing the source term of the local Poisson problems before providing it as input to the model. Formally, and using the same notations as defined in Section II, DDM-GNN outputs a correction vector \mathbf{z} in three steps, as follows:

1. Resolution of the coarse problem The resolution of the coarse problem is made similarly to Section II-A, using LU decomposition to solve:

¹Domain Decomposition Method - Graph Neural Network

²code will be made public after acceptance.

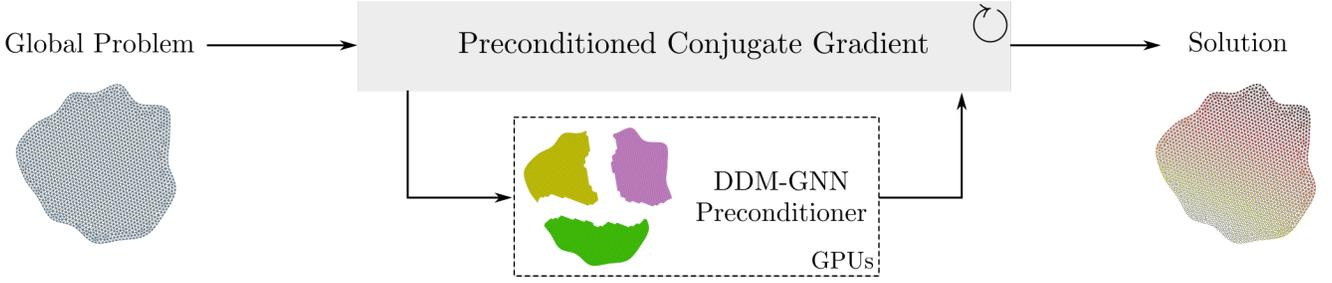


Fig. 1: Illustration of the proposed hybrid solver: A global Poisson problem (left) is solved using a Preconditioned Conjugate Gradient (PCG) method, enhanced by the DDM-GNN preconditioner. DDM-GNN combines Domain Decomposition with a GNN-based solver, harnessing parallel computations on GPUs.

$$\mathbf{r}_c = R_0^T (R_0 A R_0^T)^{-1} R_0 \mathbf{r} \quad (13)$$

2. Resolution of local problems The K local problems are solved concurrently on GPUs using a DSS model such that:

$$[\tilde{\mathbf{r}}_1, \dots, \tilde{\mathbf{r}}_K] = \text{DSS}_\theta ([G_1, \dots, G_K]) \quad (14)$$

where

$$G_i = \left(k_i, R_i A R_i^T, \frac{R_i \mathbf{r}}{\|R_i \mathbf{r}\|} \right) \quad (15)$$

represents the i -th graph associated to the i -th discretized local Poisson problem. Here, (14) suggests that all subdomains are solved simultaneously in one inference of DSS_θ . However, if the number of local problems becomes too large, $[G_1, \dots, G_K]$ can be partitioned into N_b batches, allowing all problems to be solved in N_b inferences of DSS_θ .

3. Gluing everything together The correction vector \mathbf{z} is finally computed by combining the extended coarse problem with the extended local solutions such that:

$$\mathbf{z} = \mathbf{r}_c + \sum_{i=1}^K R_i^T \|R_i \mathbf{r}\| \tilde{\mathbf{r}}_i \quad (16)$$

Fig.2 zooms in on the DDM-GNN preconditioner from Fig.1, illustrating the application of DDM-GNN to a global residual vector \mathbf{r} within the PCG algorithm.

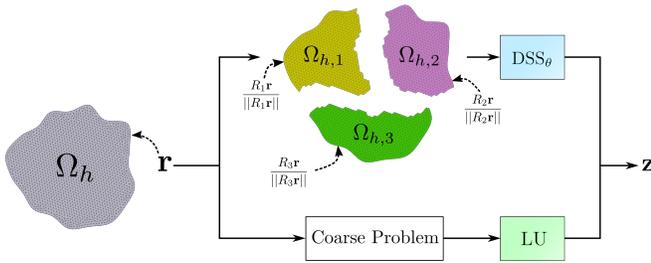


Fig. 2: Illustration of DDM-GNN: It takes a global residual vector \mathbf{r} defined on Ω_h as input and produces a global correction vector \mathbf{z} . This is done by combining the results of solving the coarse problem through LU decomposition (bottom) with those of solving all local problems using DSS_θ (top). Note, we use $\Omega_{h,i}$ instead of $R_i A R_i^T$ for reasons detailed in III-B.

B. Architecture of the DSS_θ model

In addition to our primary contribution, we present a modified version of the original DSS architecture that enables independent inference of solutions regardless of the discretization scheme. In the original DSS approach [2], edge attributes in the MPNNs are derived from the coefficients of A . As discussed in Section II-B, since A serves as the adjacency matrix for Ω_h , it naturally encodes mesh information. To adapt this architecture, we only need to ensure that the graph is undirected everywhere, except for boundary nodes whose edges point toward the interior of the graph. Thus, we propose utilizing the distance between nodes as edge attributes. Although computing the linear system remains necessary for training the model, it is no longer required for inferring solutions, thus saving the computational time of generating the linear system. As a consequence, the graph used to infer solutions, as defined in equation (15), can be equivalently formulated as:

$$G_i = \left(\Omega_{h,i}, \frac{R_i \mathbf{r}}{\|R_i \mathbf{r}\|} \right) \quad (17)$$

The architecture of DSS_θ , displayed in Fig.3, consists of an iterative process that acts on a latent state $H = (\mathbf{h}_j)_{j \in [k_i]} \in \mathcal{H}$ with $\mathbf{h}_j \in \mathbb{R}^d$, $d \geq 1$ for \bar{k} iterations. The entire architecture can be divided into three steps: an Initialization step, a Message Passing step, and a Decoding step, described as follows.

1. Initialization All latent states in H^0 are initialized as null vectors. The latent state H represents an embedding of the physical state \mathbf{r} into a higher-dimensional space. This architectural choice is driven by the need to provide enough space for information propagation throughout the network.

2. Message Passing The Message Passing step, responsible for the flow of information within the graph, performs \bar{k} updates on the latent state variable H using \bar{k} updating blocks of neural networks. To achieve this, at each iteration k , two different messages ϕ_{\rightarrow}^k and ϕ_{\leftarrow}^k are first computed using multi-layer perceptrons (MLPs) $\Phi_{\rightarrow, \theta}^{k+1}$ and $\Phi_{\leftarrow, \theta}^{k+1}$. These messages correspond to outgoing and ingoing links and are defined, node-wise, as follows:

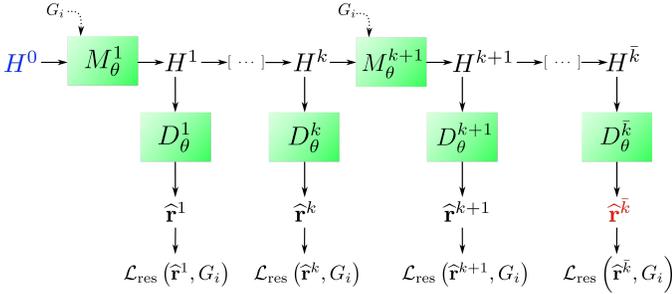


Fig. 3: DSS $_{\theta}$ architecture: The model is initialized with a null latent state H^0 (in blue). An iterative process then propagates information for \bar{k} iterations using $M_{\theta}^{\bar{k}}$ distinct blocks of MPNNs. At each iteration k , a Decoder D_{θ}^k translates the latent state H^k into a physical state $\hat{\mathbf{r}}^k$, with an intermediate loss computed. The green squares denote trainable functions, and their weights are learned during training. The last $\hat{\mathbf{r}}^{\bar{k}}$ (in red) represents the model output.

$$\phi_{\rightarrow,j}^{k+1} = \sum_{l \in \mathcal{N}(j)} \Phi_{\rightarrow,\theta}^{k+1}(\mathbf{h}_j^k, \mathbf{h}_l^k, \mathbf{d}_{jl}, \|\mathbf{d}_{jl}\|) \quad (18)$$

$$\phi_{\leftarrow,j}^{k+1} = \sum_{l \in \mathcal{N}(j)} \Phi_{\leftarrow,\theta}^{k+1}(\mathbf{h}_j^k, \mathbf{h}_l^k, \mathbf{d}_{lj}, \|\mathbf{d}_{lj}\|) \quad (19)$$

where \mathbf{d}_{jl} represents the relative position vector and $\|\mathbf{d}_{jl}\|$ its Euclidean distance. The updated latent state H^{k+1} is then computed using an MLP Ψ_{θ}^{k+1} in a ResNet-like fashion such that, node-wise:

$$\mathbf{h}_j^{k+1} = \mathbf{h}_j^k + \alpha \Psi_{\theta}^{k+1}(\mathbf{h}_j^k, c_j, \phi_{\rightarrow,j}^{k+1}, \phi_{\leftarrow,j}^{k+1}) \quad (20)$$

with $c = \frac{R_i \mathbf{r}}{\|R_i \mathbf{r}\|}$. For purpose of clarity, operations (18), (19) and (20) can be grouped into a single updating block M_{θ}^{k+1} such that the next latent state H^{k+1} is computed as follows:

$$H^{k+1} = M_{\theta}^{k+1}(H^k, G_i) \quad (21)$$

3. Decoder A Decoding step is applied after each iteration to convert the latent state H^{k+1} into a meaningful actual state $\hat{\mathbf{r}}^{k+1}$ by using an MLP D_{θ}^{k+1} such that :

$$\hat{\mathbf{r}}^{k+1} = D_{\theta}^{k+1}(H^{k+1}) \quad (22)$$

The final state $\hat{\mathbf{r}}^{\bar{k}}$ represents the actual output of the algorithm, which corresponds to the sought approximate solution $\hat{\mathbf{r}}_i$ of a local problem i in (14). The training loss is computed as a sum of all intermediate losses, as follows:

$$\text{Training Loss} = \sum_{k=1}^{\bar{k}} \mathcal{L}_{\text{res}}(\hat{\mathbf{r}}^k, G_i) \quad (23)$$

IV. RESULTS

This section analyzes the numerical behavior and performance of the proposed hybrid solver. Section IV-A introduces the dataset and Section IV-B the training configuration of DSS $_{\theta}$. The numerical behavior of the approach is evaluated in Section IV-C, and the impact of DSS $_{\theta}$ hyperparameters on performance is assessed in Section IV-D. Finally, the proposed

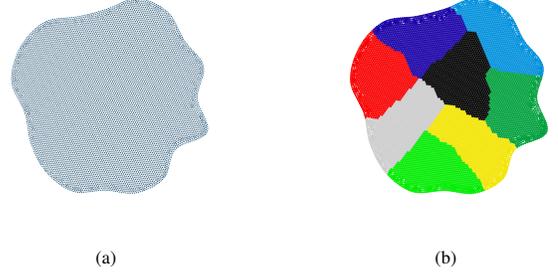


Fig. 4: Example of (4a): one generated global domain with 7420 nodes, and (4b): its partitioning into 8 sub-meshes. Each colored sub-mesh corresponds to a graph sample of the dataset.

method is benchmarked against an optimized C++ legacy linear solver in Section IV-E.

A. Dataset

We generate random 2D domains Ω using 20 points sampled from the unit sphere, connected with Bezier curves to form $\partial\Omega$. Next, GMSH³ is used to discretize Ω into unstructured triangular meshes Ω_h , each containing 6000 to 8000 nodes. For each domain, we solve Poisson problems (1) with forcing functions f and boundary functions g defined as random quadratic polynomials with coefficients $r_{i \in [1, \dots, 9]}$ uniformly sampled in $[-10, 10]$, such that:

$$f(x, y) = r_1(x-1)^2 + r_2y^2 + r_3 \quad (24)$$

$$g(x, y) = r_4x^2 + r_5y^2 + r_6xy + r_7x + r_8y + r_9 \quad (25)$$

Each mesh Ω_h is then partitioned into sub-meshes of approximately 1000 nodes using METIS⁴ with an overlap of 2. We use the PCG solver described in Algorithm 1 preconditioned with a classic two-level ASM method (see Section II-A) to solve the global Poisson problems with a relative residual norm tolerance of 10^{-6} . The dataset consists of discretized local Poisson problems from the two-level ASM preconditioner extracted at each iteration of the PCG algorithm. In that setting, solving 500 global Poisson problems results in 117,138 samples, further split into 70,282 / 23,428 / 23,428 training/validation/test datasets. Figure 4 displays a global problem and its partitioning into several sub-meshes, each of which belongs to the dataset. In subsequent analyses, when considering meshes with varying node counts, this is achieved by increasing the radius of the mesh while maintaining the element size fixed from the original distribution. The force and boundary functions are rescaled accordingly.

B. Model Configuration & Training

DSS $_{\theta}$ is implemented in Pytorch, using the library Pytorch Geometric⁵. The model is trained using the original hyperparameters from [2]; the number of iterations \bar{k} is set to 30, and the latent space dimension d to 10. Each neural network

³<https://gmsh.info/>

⁴<https://github.com/KarypisLab/METIS>

⁵https://github.com/pyg-team/pytorch_geometric

in Eqs. (19), (18), (20), and (22) has one hidden layer of dimension 10 with a ReLU activation function. All model parameters are initialized using Xavier initialization. In Eq. (20), coefficient α is set to 10^{-3} . Training is performed for 400 epochs on 2 P100 GPU using the Adam optimizer with a learning rate of 10^{-2} and a batch size of 100. Gradient clipping is used to prevent exploding gradient issues and is set to 10^{-2} . The *ReduceLrOnPlateau* scheduler from PyTorch is applied, reducing the learning rate by a factor of 0.1 during training. After training, the evaluation of DSS_θ on the test set results in a residual norm of 0.0058 ± 0.002 and a Relative Error with an “exact” solution computed through LU decomposition of 0.13 ± 0.2 .

C. Numerical behavior of the hybrid solver

This section examines the numerical behavior of the hybrid solver across various configurations, including out-of-distribution problems. For each setup, we solve 100 global Poisson problems, sampling parameters as described in Section IV-A. The iteration count required to achieve a relative residual norm of 10^{-6} is evaluated for different problem sizes $N \simeq 2000, 7000$ (training configuration), and 30,000. For each problem size N , we test sub-mesh sizes $N_s = 500, 1000$ (training configuration), and 2000, with an overlap of 2. We also investigate local problems of size 1000 with an overlap of 4. Table I reports mean (\pm standard deviation) values over the 100 Poisson problems. The proposed approach, PCG preconditioned with DDM-GNN, is compared with PCG preconditioned with DDM-LU (i.e. two-level ASM from Section II-A with local problems solved through LU decomposition), and a Conjugate Gradient (CG) without preconditioner. Results show that PCG-DDM-LU always converges in fewer iterations than PCG-DDM-GNN. This entirely expected behavior is due to the fact that DDM-LU uses a direct solver to solve local sub-problems while DDM-GNN uses DSS_θ , which cannot achieve higher precision than a direct LU solver. What is interesting is analyzing the convergence of PCG-DDM-GNN in comparison to PCG-DDM-LU. Firstly, regardless of the problem

TABLE I: Numerical Behaviour

Configuration				Iterations		
N	N_s	K	Overlap	DDM-GNN	DDM-LU	CG
2632 ± 90	1000	3 ± 0	2	22 ± 1	19 ± 1	113 ± 3
–	–	–	4	17 ± 1	14 ± 1	–
–	500	6 ± 0	2	24 ± 1	22 ± 1	–
–	2000	2 ± 0	2	18 ± 1	16 ± 1	–
7148 ± 247	1000	7 ± 1	2	32 ± 2	28 ± 2	183 ± 6
–	–	–	4	27 ± 2	21 ± 2	–
–	500	14 ± 1	2	31 ± 2	29 ± 2	–
–	2000	4 ± 0	2	29 ± 2	24 ± 2	–
33969 ± 1181	1000	35 ± 2	2	49 ± 3	42 ± 3	385 ± 6
–	–	–	4	46 ± 3	36 ± 3	–
–	500	68 ± 2	2	45 ± 3	42 ± 3	–
–	2000	17 ± 1	2	51 ± 3	41 ± 2	–

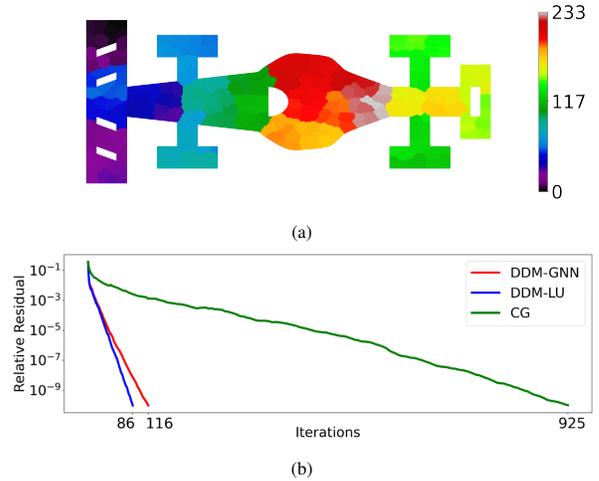


Fig. 5: (5a): partitioning of the mesh into 234 sub-meshes and (5b): the evolution of the relative residual error for the different methods until it reaches an order of 10^{-9} .

or solver setup, PCG-DDM-GNN always converges to the desired precision. Secondly, the difference in iteration counts between PCG-DDM-GNN and PCG-DDM-LU is minimal (always fewer than 10 iterations) across all configurations. It’s worth noting that DSS_θ is trained on datasets with 1000 nodes, yet the entire process consistently converges, even when the sub-mesh size varies (either smaller at 500 or larger at 2000), indicating robustness to out-of-distribution samples. Moreover, the proposed approach maintains the advantageous property of faster convergence with larger overlaps. Overall, the proposed hybrid solver demonstrates consistency; leveraging a GNN-based solver does not significantly degrade results compared to using an “exact solver”. Additionally, the method proves to be scalable in terms of the number of sub-domains and converges much more effectively than CG.

We further conduct a large-scale experiment on a mesh representing a caricatural Formula 1 with 233,246 nodes. This mesh includes “holes” (such as a cockpit and front and rear wing stripes) and is larger than those seen in the training dataset, providing a challenging test of the ability of the hybrid solver to generalize to out-of-distribution (with respect to the geometry of the domain, as well as the mesh size) examples. To solve it, we sample force and boundary functions as described in Section IV-A. The mesh is divided into sub-meshes of about 1000 nodes, resulting in 234 sub-meshes, as shown in Fig.5a. We solve the Poisson problem using PCG preconditioned with DDM-GNN and DDM-LU, as well as CG until the relative residual error reaches 10^{-9} . Figure 5b illustrates the evolution of the relative residual error during the iterations of the various methods. Results show that the proposed hybrid solver still converges when applied to an out-of-distribution and large-scale sample, achieving precision levels significantly lower than those used for training the dataset. Besides, PCG-DDM-GNN is still competitive with respect to PCG-DDM-LU regarding the iteration count and

TABLE II: Metrics for DSS_θ trained with varying \bar{k} and d .

DSS $_\theta$		Metrics		
\bar{k}	d	Residual (10^{-2})	Relative Error	Nb Weights
5	5	4.11 \pm 1.2	0.67 \pm 0.1	1755
5	10	3.84 \pm 1.1	0.66 \pm 0.1	6255
5	20	3.74 \pm 1.1	0.65 \pm 0.1	23505
10	5	2.02 \pm 0.7	0.48 \pm 0.4	3510
10	10	1.58 \pm 0.49	0.44 \pm 0.2	12510
10	20	1.55 \pm 0.49	0.43 \pm 0.1	47010
20	5	0.88 \pm 0.34	0.20 \pm 0.4	7020
20	10	0.68 \pm 0.25	0.20 \pm 0.3	25020
20	20	0.57 \pm 0.21	0.18 \pm 0.1	94020
30	10	0.58 \pm 0.22	0.13 \pm 0.2	37530

converges much more efficiently than a traditional Conjugate Gradient.

D. Impact of DSS_θ hyperparameters on the performance

The size of the DSS_θ network is directly influenced by the latent space dimension d and the number of layers in the network \bar{k} . From a performance perspective, it is essential to consider that a larger network may result in a more precise trained model, but it also incurs increased inference costs. This exploration analyses the impact of the two parameters, d and \bar{k} , on both the numerical behavior (iteration counts at convergence) and the time performance of the hybrid solver. Several DSS_θ models are trained corresponding to parameter sets $\bar{k} \in \{5, 10, 20, 30\}$ and $d \in \{5, 10, 20\}$. Table II summarizes the metrics averaged over the whole test set for these models, including the residual and relative errors, as well as the number of weights. These results show that as \bar{k} and d increase, the metrics improve, but the number of weights also increases.

Next, we apply the PCG-DDM-GNN solver based on each of these trained models to solve 100 Poisson problems of size 10,000. Fig.6a illustrates the inference time (in seconds) to solve a batch of local problems on the left axis and the number of iterations at convergence on the right axis for

each parameter pair (\bar{k}, d) . Fig.6b illustrates the elapsed time (in seconds) of the full resolution on the left axis and the number of iterations at convergence on the right axis for each parameter pair (\bar{k}, d) . Upon analysis of Fig.6a, as expected, it becomes evident that a larger DSS_θ model (represented by the values of \bar{k} and d) results in a greater number of weights, making the preconditioner more accurate at the expense of performance. Fig.6b highlights that the best performance is achieved with the model corresponding to $\bar{k} = 10$ and $d = 10$, even though this model may not be the most accurate.

E. Performance results

We integrated our DDM-GNN preconditioner into the HTSSOLVER package [4, 5], an in-house C++ linear solver. This required extending its multi-level Domain Decomposition framework, reliant on the direct sparse LU solver from the Eigen⁶ library, by incorporating the DSS_θ model. This new solver involves constructing related graph structures for each sub-domain and inferring DSS_θ . Regarding DSS_θ inference, we evaluated several C++ inference engines, including ONNX⁷, TensorRT⁸, and the C++ LibTorch⁹ library. A comparative study of these frameworks is presented in [25]. Our GNN model, implemented using Pytorch-Geometric, requires a 64-bit index type for graph structures, making it incompatible with TensorRT. Ultimately, the C++ LibTorch library was chosen for its superior performance on NVidia GPUs.

The numerical behavior of DDM-GNN implemented in C++ was initially validated by comparing results with the Python framework. We then explore its extensibility concerning global problem sizes and compare its performance with standard optimized preconditioners and the original DDM-LU method, which leverages the sparse LU solver from Eigen, one of the most optimized C++ solver packages. To assess performance in terms of elapsed time, experiments were conducted on Jean Zay, an HPE SGI 8600 multi-node Linux cluster featuring accelerated dual-socket nodes. These nodes are equipped with Intel Cascade Lake 6248 processors at 2.5 GHz, each comprising 20 cores per socket, and either 4 Nvidia Tesla V100 SXM2 GPUs with 16 or 32 GB of memory. Additionally, the cluster includes accelerated dual-socket nodes with AMD Milan EPYC 7543 processors at 2.80 GHz, featuring 32 cores per socket and 8 Nvidia A100 SXM4 GPUs with 80 GB of memory.

To evaluate performance, we analyze the resolution of Poisson problems on increasingly large meshes, resulting in linear systems with sizes approximately $N \simeq 10k, 40k, 100k, 250k, 400k, \text{ and } 600k$. For each system, the graph is partitioned using METIS into K sub-domains of sizes approximately 500, 1000, and 2000 nodes. All experiments

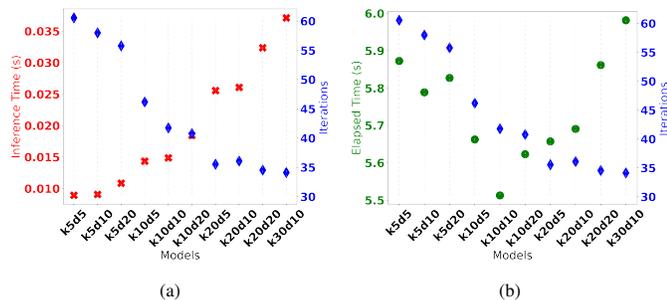


Fig. 6: DSS_θ performance exploration regarding \bar{k} and d .

⁶<https://gitlab.com/libeigen/eigen>

⁷<https://github.com/onnx/onnx>

⁸<https://github.com/NVIDIA/TensorRT>

⁹<https://pytorch.org/cppdocs/>

TABLE III: Benchmark to legacy C++ solver

Configuration		IC(0)		DDM-LU			DDM-GNN		
N	K	N_{iter}	T	N_{iter}	T	T_{lu}	N_{iter}	T	T_{gnn}
10571	5	35	0.0127	17	0.049	0.018	29	0.147	0.116
	11			14	0.041	0.011	25	0.138	0.099
	22			12	0.030	0.008	25	0.142	0.098
41871	21	66	0.0850	18	0.16	0.087	56	0.55	0.44
	42			16	0.18	0.064	33	0.63	0.53
	84			15	0.15	0.062	42	0.80	0.68
100307	50	97	0.298	66	1.97	1.61	35	1.25	1.03
	100			39	1.04	0.76	35	1.30	1.25
	200			54	0.76	0.45	39	1.54	1.43
259604	125	136	1.21	33	1.43	1.03	38	2.77	2.38
	250			25	1.07	0.66	38	2.90	2.39
	500			22	1.08	0.51	36	2.88	2.26
405344	200	168	2.39	35	2.38	1.71	38	3.60	2.83
	400			29	1.92	1.20	38	3.74	2.84
	800			25	1.99	0.90	38	5.64	4.45
609740	300	221	4.9	40	3.13	2.92	38	5.26	4.13
	600			37	3.86	2.33	39	5.67	5.24
	1200			32	2.72	1.73	39	8.88	6.92

solve the linear systems using PCG with a relative residual norm tolerance of 10^{-3} . We compare the performances of the Incomplete Cholesky preconditioner IC(0), DDM-LU, and DDM-GNN. Table III summarizes, for each size N and number of sub-domains K , the total execution time T to solve the entire problem, the inference times T_{lu} and T_{gnn} for applying DDM-LU and DDM-GNN, respectively, and the number of iterations at convergence N_{iter} .

The analysis of the results shows that we recover the numerical behavior in terms of the number of iterations at convergence N_{iter} described previously. The variation of N_{iter} is less sensitive to the global size N of the linear system for the DDM-LU and DDM-GNN preconditioners compared to IC(0). Regarding K , the number of sub-domains, the analysis shows that although the GNN model was trained on meshes of size 1000, the method can generalize even with sub-domains of sizes varying between 500 and 2000 nodes. However, the performance analysis in terms of elapsed time reveals a significant challenge in inferring GNN models in C++ code. The ratios $\frac{T_{lu}}{T}$ and $\frac{T_{gnn}}{T}$ indicate that the majority of the overall resolution time is spent in the GNN and LU resolution. Using the LibTorch inference engine, the GNN resolution remains a substantial challenge, preventing our method from being as competitive as expected compared to state-of-the-art optimized preconditioners. In fact, as described in the literature, for instance, in [14], standard inference engines optimized for Image processing or Large Language Models are not yet mature enough to efficiently support GNN algorithms for HPC software. Future work is necessary to enhance the performance of the GNN resolution by developing an efficient inference engine for GNN structures, contributing to either the ONNX or the TensorRT project.

V. CONCLUSION AND PERSPECTIVES

This paper introduces DDM-GNN, a GNN-based preconditioner that integrates a Graph Neural Network (GNN) model with a Domain Decomposition strategy to enhance the convergence of Krylov methods. This innovative approach allows the GNN model to efficiently handle large-scale meshes in parallel on GPUs. DDM-GNN is also equipped with a two-level method, i.e. a coarse space correction, ensuring the weak scalability of the entire process regarding the number of sub-domains. We validated the numerical behavior of the method on various test cases, demonstrating its applicability and flexibility to a broad range of meshes with varying sizes and shapes. The accuracy of the proposed method ensures the linear solver converges efficiently to any desired level of precision. Implemented within a C++ solver package, the performance of DDM-GNN was evaluated against state-of-the-art optimized preconditioners. However, inferring GNN models in such contexts remains challenging. Existing inference engines like TensorRT or ONNX are not mature enough to handle GNN algorithms efficiently in C++ High-Performance Computing (HPC) codes. Addressing this challenge requires future efforts, including direct contributions to projects like TensorRT or ONNX. To conclude, it is noteworthy that integrating GPU computations into existing legacy codes can be a daunting task. The proposed method holds promise by offering a straightforward methodology that can be easily incorporated into HPC frameworks, harnessing GPU computations.

The long-term objective of this research is to accelerate industrial Computation Fluid Dynamics (CFD) codes on software platforms such as OpenFOAM¹⁰. In upcoming work, we plan to evaluate the efficiency of the proposed approach by implementing it in an industrial CFD code and assessing its impact on performance acceleration.

VI. ACKNOWLEDGEMENTS

This research was supported by DATAIA Convergence Institute as part of the ‘‘Programmed’ Investissement d’Avenir’’, (ANR- 17-CONV-0003) operated by INRIA and IFPEN. This project was provided with computer and storage resources by GENCI at [IDRIS] thanks to the grant 2023-[AD011014247] on the supercomputer [Jean Zay]’s the [V100/A100] partition.

REFERENCES

- [1] Victorita Dolean, Pierre Jolivet, and Frédéric Nataf. *An introduction to domain decomposition methods: algorithms, theory, and parallel implementation*. SIAM, 2015.
- [2] Balthazar Donon et al. ‘‘Deep statistical solvers’’. In: *NeurIPS* (2020).
- [3] Martin J Gander et al. ‘‘Schwarz methods over the course of time’’. In: *Electron. Trans. Numer. Anal* 5 (2008).

¹⁰<https://www.openfoam.com/>

- [4] Jean-Marc Gratien. “A robust and scalable multi-level domain decomposition preconditioner for multi-core architecture with large number of cores”. In: *Journal of Computational and Applied Mathematics* (2019).
- [5] Jean-Marc Gratien. “Introducing multi-level parallelism, at coarse, fine and instruction level to enhance the performance of iterative solvers for large sparse linear systems on Multi-and Many-core architecture”. In: *2020 IEEE/ACM 6thLLVM-HPC Workshop and HiPar Workshop*. IEEE, 2020.
- [6] J-L Guermont and Luigi Quartapelle. “On stability and convergence of projection methods based on pressure Poisson equation”. In: *International Journal for Numerical Methods in Fluids* 9 (1998).
- [7] Alexander Heinlein et al. “Combining machine learning and domain decomposition methods for the solution of partial differential equations—A review”. In: *GAMM-Mitteilungen* 1 (2021).
- [8] Masanobu Horie and Naoto Mitsume. “Physics-Embedded Neural Networks: Graph Neural PDE Solvers with Mixed Boundary Conditions”. In: *NeurIPS* (2022).
- [9] Ekhi Ajuria Illarramendi, Michaël Bauerheim, and Bénédicte Cuenot. “Performance and accuracy assessments of an incompressible fluid solver coupled with a deep convolutional neural network”. In: *Data-Centric Engineering* (2022).
- [10] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. “Poisson surface reconstruction”. In: *Proceedings of the fourth Eurographics symposium on Geometry processing*. 2006.
- [11] Axel Klawonn, Martin Lanser, and Janine Weber. “Machine learning and domain decomposition methods—a survey”. In: *arXiv:2312.14050* (2023).
- [12] Tobias Knoke et al. “Domain Decomposition with Neural Network Interface Approximations for time-harmonic Maxwell’s equations with different wave numbers”. In: *arXiv:2303.02590* (2023).
- [13] Dmitrii Kochkov et al. “Machine learning–accelerated computational fluid dynamics”. In: *Proceedings of the National Academy of Sciences* 21 (2021).
- [14] Alina Lazar et al. “Accelerating the Inference of the Exa. TrkX Pipeline”. In: *Journal of Physics: Conference Series*. 1. IOP Publishing, 2023.
- [15] Ke Li et al. “D3M: A deep domain decomposition method for partial differential equations”. In: *IEEE Access* (2019).
- [16] Wuyang Li, Xueshuang Xiang, and Yingxiang Xu. “Deep domain decomposition method: Elliptic problems”. In: *Mathematical and Scientific Machine Learning*. PMLR, 2020.
- [17] Brock A Luty, Malcolm E Davis, and J Andrew McCammon. “Solving the finite-difference non-linear Poisson–Boltzmann equation”. In: *Journal of computational chemistry* 9 (1992).
- [18] Philip D Mannheim and Demosthenes Kazanas. “Newtonian limit of conformal gravity and the lack of necessity of the second order Poisson equation”. In: *General Relativity and Gravitation* 4 (1994).
- [19] Roy A Nicolaides. “Deflation of conjugate gradients with applications to boundary value problems”. In: *SIAM Journal on Numerical Analysis* 2 (1987).
- [20] Ali Girayhan Özbay et al. “Poisson CNN: Convolutional neural networks for the solution of the Poisson equation on a Cartesian mesh”. In: *Data-Centric Engineering* (2021).
- [21] Tobias Pfaff et al. “Learning mesh-based simulation with graph networks”. In: *arXiv:2010.03409* (2020).
- [22] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations”. In: *Journal of Computational physics* (2019).
- [23] Yousef Saad. *Iterative methods for sparse linear systems*. SIAM, 2003.
- [24] Hermann Amandus Schwarz. *Ueber einen Grenzübergang durch alternirendes Verfahren*. Zürcher u. Furrer, 1870.
- [25] Murat Sever and Sevda Ögüt. “A Performance Study Depending on Execution Times of Various Frameworks in Machine Learning Inference”. In: *2021 15th Turkish National Software Engineering Symposium (UYMS)*. IEEE, 2021.
- [26] Ali Taghibakhshi et al. “Learning interface conditions in domain decomposition solvers”. In: *NeurIPS* (2022).
- [27] Yushan Wang. “Solving incompressible Navier–Stokes equations on heterogeneous parallel architectures”. PhD thesis. Paris 11, 2015.
- [28] Steffen Wiewel, Moritz Becher, and Nils Thuerey. “Latent space physics: Towards learning the temporal evolution of fluid flow”. In: *Computer graphics forum*. 2. Wiley Online Library, 2019.
- [29] Bing Yu et al. “The deep Ritz method: a deep learning-based numerical algorithm for solving variational problems”. In: *Communications in Mathematics and Statistics* 1 (2018).