

High-Fidelity 4D Flow Assimilation via Hard-Constrained Pseudo-Spectral Meshless Track Assimilation

P. Sperotto^{1,*}, N. Gade-Nielsen¹, D. Hess¹

¹: Dantec Dynamics A/S, Denmark

*Corresponding author: pietro.sperotto@dantecdynamics.com

Keywords: Data Assimilation, LPT, PTV, Pressure from PIV.

ABSTRACT

Meshless Track Assimilation (MTA) provides exceptional accuracy for the 4D Data Assimilation (DA) of Particle Tracking Velocimetry (PTV) data, but its widespread adoption has been bottlenecked by prohibitive computational times, massive memory constraints, and tedious hyperparameter tuning. This work introduces a transformative pseudo-spectral paradigm for MTA by kernelizing compact-support Radial Basis Functions (RBFs), shifting the computational core from expensive sparse matrix products to highly efficient FFT-based convolutions. Furthermore, by embedding the divergence-free condition as a hard spectral constraint and analytically eliminating pressure via variable projection. The overhauled formulation achieves a 20x computational speedup and slashes the DA operator memory footprint from 170 MB to just 2.5 MB while preserving or improving reconstruction accuracy. Crucially, this framework democratizes the method, making fully coupled, billion-variable 4D flow reconstructions feasible in-core on consumer-grade PC hardware with virtually no parameter tuning.

1. Introduction

MTA represents velocity and pressure fields as linear combinations of Radial Basis Functions (RBFs) (Sperotto et al., 2024), with coefficients determined by minimizing a non-linear cost function comprising residuals from experimental data ($\mathbf{R}_{data}$) and from the governing equations (momentum, continuity, temporal continuity, and Poisson equations). While the method proved highly promising in terms of accuracy, it was computationally impractical due to three main limitations: (i) excessive CPU time, (ii) high memory requirements, and (iii) multiple sources of regularization that complicated optimization. The aim of this work is to reformulate the MTA algorithm within a pseudo-spectral framework by kernelizing compact-support RBFs. Improving computational time and memory requirements.

2. Mathematical Framework

The mathematical framework is divided into three parts. Section 2.1 shows the main breakthrough. In sections 2.2 and 2.3 the direct consequences of section 2.1 are shown.

2.1. Kernelization of Compact-Support RBFs

In the original formulation, we employed Wendland functions $\phi \in C^6(\mathbb{R}^3)$ (Wendland, 1995). These are compact-support (CS) RBFs differentiable up to the sixth order in three dimensions. Most of the memory requirements of the MTA implementation originate from the sparse matrices Φ^{DA} , $\partial_{x_j} \Phi^{DA}$, and $\Delta \Phi^{DA} \in \mathbb{R}^{n_b \times n_b}$, with $n_b = n_{bx}n_{by}n_{bz}$ denoting the total number of bases. Here, the subscript $j \in \{x, y, z\}$ denotes spatial derivatives, while the superscript DA indicates that these matrices map the weights to quantities evaluated on the DA grid; for example, $\Phi^{DA} \mathbf{w}_u = \mathbf{u}^{DA}$ with $\mathbf{w}_u$ the RBF weights of a generic quantity, where the governing equations are enforced. In this work, the DA grid is fixed to coincide with the collocation grid.

Under these circumstances, the matrix-vector multiplication can be interpreted as a discrete cross-correlation ($\star$) between the RBF kernel and the weight field:

$$\mathbf{u}^{DA} = \varphi \star \mathbf{w}_u, \quad (1)$$

where φ is a kernelized version of the RBF, i.e. the RBF collocated at the central point of the domain and evaluated on the DA grid. Exploiting the convolution theorem, Eq. (1) can be expressed in the spectral domain as

$$\hat{\mathbf{u}}^{DA} = \overline{\hat{\varphi}} \hat{\mathbf{w}}_u, \quad (2)$$

where the overline denotes complex conjugation, consistent with the cross-correlation definition.

Unlike classical pseudo-spectral methods, we do not assume periodic boundary conditions. Instead, the computational grid is regarded as part of an effectively infinite domain, while only a finite subregion is observed. Because the RBFs possess compact support, only a small number of basis functions outside the observation volume influence the cross-correlation. To ensure accurate gradients near the boundaries, the grid is padded with additional basis functions sufficient to fully cover the kernel support. The influence of these external bases is naturally suppressed through Tikhonov regularization, which stabilizes the inversion and mitigates residual boundary artifacts. In this framework, the Fast Fourier Transform (FFT) serves purely as an efficient computational engine rather than as an assumption of periodicity.

2.2. Divergence-Free Constraint

This formulation unlocks a very efficient way to invert linear, even singular, operators. A typical example of a singular linear operator that must be inverted efficiently is the divergence-free constraint, which enforces incompressibility.

We begin by writing the divergence-free condition in the same spectral form as Eq. (2):

$$\partial_x \bar{\hat{\phi}} \hat{w}_u + \partial_y \bar{\hat{\phi}} \hat{w}_v + \partial_z \bar{\hat{\phi}} \hat{w}_w = 0,$$

$$\partial_x \hat{\phi} \hat{w}_u + \partial_y \hat{\phi} \hat{w}_v + \partial_z \hat{\phi} \hat{w}_w = 0,$$

$$i\mathbf{k}_x \hat{w}_u + i\mathbf{k}_y \hat{w}_v + i\mathbf{k}_z \hat{w}_w = 0,$$

$$\Rightarrow \hat{w}_u = -\frac{\mathbf{k}_y}{\mathbf{k}_x} \hat{w}_v - \frac{\mathbf{k}_z}{\mathbf{k}_x} \hat{w}_w,$$

where $\mathbf{k}_{x,y,z}$ denotes the frequency-domain coordinates.

For compactness, we assume that any frequency bin of $\hat{w}_u$ can be obtained from the previous expression. This assumption fails on the plane $\mathbf{k}_x = 0$, where the relation becomes singular and those values must instead be computed through the standard optimization procedure. Thus, the constraint does not isolate $\hat{w}_u$ alone but enforces a linear combination among all velocity components.

This operation can be performed extremely quickly and with machine precision. However, the gradient must be modified accordingly via the chain rule at the same expense of finding w_u . This projection technique is implemented in both the (non-linear) MTA and its linear version linear-MTA which we consider the evolution of the Physics Informed Interpolation (PII) (Sperotto et al., 2023). The advantages are clear: enforcing the divergence-free condition in this way significantly improves computational performance and removes the need for two hyperparameters previously required in MTA: the weighting of the divergence-free term and of its time derivative. In linear-MTA, this constraint effectively renders the method hyperparameter-free. This technique differs from FlowFit3 Godbersen et al. (2024), where all three velocity components are active parameters and the gradient is projected during optimization to preserve incompressibility, requiring a divergence-free initial guess. In our approach, the parameters inherently satisfy the divergence-free condition: u is formally determined from v and w , meaning that, as a simplification, only the weights of v and w are treated as active in the optimization.

2.3. Variable Projection of the Pressure in Fourier Space

State-of-the-art techniques typically eliminate pressure as an active optimization variable, since pressure often exhibits a different magnitude scale from the velocity components, which can lead

to numerical instability. The standard remedy is to solve a Poisson equation for pressure (Jeon et al., 2022; Godbersen et al., 2024). This, however, requires Dirichlet boundary conditions and multiple Poisson solves to project the gradient at each iteration. Here, we adopt the Variable Projection (VP) approach proposed by Golub & Pereyra (1973) to analytically eliminate the pressure field within the optimization.

Let $\mathbf{F}'_{x_j}$ denote the residuals of the Navier–Stokes equations in direction x_j excluding the pressure gradient, and $\mathbf{F}_{x_j} = \mathbf{F}'_{x_j} + \partial_{x_j} p$ be the complete residuals. The corresponding linear least-squares problem for the pressure reads

$$(\partial_x \Phi^T \partial_x \Phi + \partial_y \Phi^T \partial_y \Phi + \partial_z \Phi^T \partial_z \Phi) \mathbf{w}_p = -(\partial_x \Phi^T \mathbf{F}'_x + \partial_y \Phi^T \mathbf{F}'_y + \partial_z \Phi^T \mathbf{F}'_z). \quad (3)$$

Using the kernelization approach introduced in Section 2.1, this relation can be expressed in the spectral domain as

$$\bar{\hat{\phi}} \hat{\mathbf{w}}_p = -\frac{i\mathbf{k}_x \hat{\mathbf{F}}'_x + i\mathbf{k}_y \hat{\mathbf{F}}'_y + i\mathbf{k}_z \hat{\mathbf{F}}'_z}{\mathbf{k}_x^2 + \mathbf{k}_y^2 + \mathbf{k}_z^2}. \quad (4)$$

This shows that, in the present formulation, the least-squares solver is equivalent to a Poisson solver in Fourier space. The boundary condition reduces to an arbitrary additive bias, which does not influence the Navier–Stokes equations. The key advantage is that the optimal pressure can be evaluated analytically and efficiently at any gradient query, without repeated Poisson solves.

Because the velocity field is already analytically divergence-free, the formulation inherently satisfies both the general and the incompressible forms of the Poisson equation. Consequently, the Poisson penalization term used in previous implementations is no longer required. Except for an optional Tikhonov regularization (not strictly necessary), the only remaining hyperparameter is the weighting of the momentum equations. This considerably simplifies the optimization and enhances robustness. Furthermore, this concept applied with the divergence-free constraint (Section 2.2) implies that MTA optimizes only the weights concerning v and w in a classical non-linear way.

3. Validation

3.1. Transitional Boundary Layer

The smallest dataset of the Transitional Boundary Layer (TBL) used in Sperotto et al. (2024). It is used here as benchmark, this test case is generated from Zaki (2013) DNS data.

In non-dimensional units, we extract a volume of size $30 \times 10 \times 5$ with an inter-frame time of $\Delta t = 0.25$. From the Johns Hopkins Turbulence Database we sample an $89 \times 29 \times 14$ grid and advect 2200 particles for 50 time steps, which are assimilated simultaneously.

We define the Relative Root Mean Square Error (RRMSE) as

$$RRMSE_\varepsilon = \frac{1}{n_t} \sum_{q=0}^{n_t} \sqrt{\frac{1}{N_{grid}} \sum_{i=0}^{N_{grid}} \frac{(\varepsilon^{MTA} - \varepsilon^{GT})^2}{\max_{\mathbf{x} \in \mathbf{X}^{grid}} |\varepsilon^{GT}|}}, \quad (5)$$

where the superscript *grid* indicates quantities sampled on the ground-truth grid used for evaluation.

Table 1 reports the reconstruction accuracy and performance. The new MTA outperforms the original formulation in all quantities except the streamwise velocity component. However, the new MTA was not expected to be "better" than the previous implementation in terms of results but only in terms of computational performance and ease of use. In the original MTA, each velocity component had a separate scaling, which tended to favor the dominant streamwise direction. Despite the slight loss in u , improvements in v , w and p are substantial. It is also worth noting that for the original MTA we have $\mathbf{X}^{grid} \equiv \mathbf{X}^{DA}$, which gives it a slight advantage in this particular benchmark.

MTA	$RRMSE_u$	$RRMSE_v$	$RRMSE_w$	$RRMSE_p$	Memory [GB]	Time [min]
Old	2.8%	12.1%	10.6%	8.4%	2.8	~ 28.8
New	3.0%	10.8%	7.5%	7.5%	1.2	~ 1.5

Table 1. Comparison between old and new MTA on the TBL dataset.

Regarding memory usage, both methods require approximately 500 MB for the LPT matrices. The DA operator, however, requires about 170 MB in the old MTA and only 2.5 MB in the new FFT-based formulation.

The computational savings are dramatic. Even in an unoptimized single-core CPU implementation, the new MTA is about twenty times faster than the original one. These timings exclude initialization and refer only to the reconstruction of a single time step (not including interpolation to the output grid). For larger cases, the computational gap is expected to widen further: the original MTA relies on large sparse CSC multiplications, while the new formulation reduces the core operation to FFT-based convolutions. On GPUs, where FFT operations are heavily optimized and sparse products remain bandwidth-bound, we expect the FFT-based MTA to outperform the classical CSC approach by well over an order of magnitude.

4. Experimental Validation

To validate the method and assess its computational performance, two experimental test cases are considered. All results reported in this section were obtained using the final optimized CPU implementation of the code. Both test cases correspond to turbulent jet flows. The first case is

based on high-quality reconstructed particle tracks, although the particle distribution is relatively coarse; for the purpose of this work, only 37 time steps are evaluated. The second case, instead, is a significantly more challenging dataset, characterized by large positional uncertainty and ghost tracks.

The objective of the first experiment is to validate the reconstruction quality by subsampling the data and comparing the reconstruction obtained from the subsampled dataset with that obtained from the full dataset. The objective of the second experiment is instead to stress-test the method in terms of memory usage, computational cost and robustness.

4.1. Experiment 1

As mentioned above, this experiment is based on a clean particle reconstruction. For this reason, the full-density reconstruction can be used as a proxy for the ground truth. The dataset contains approximately 7.4×10^4 particles per time step in a measurement volume of $170 \times 90 \times 90 \text{ mm}^3$. A second dataset was generated by randomly subsampling the original particle set to 30% of its original density.

To avoid introducing a bias by selecting a single reconstruction method as reference, the full-density dataset was reconstructed using MTA, linear-MTA, and Adaptive Gaussian Window interpolation (AGW). The same three methods were then applied to the subsampled dataset. The error was evaluated for every possible combination between the subsampled reconstruction and the full-density reference reconstruction. In this way, if one method consistently performs best across different reference choices, the result is less dependent on the particular ground-truth proxy.

For MTA, we used 10 bases per particle (bpp) and $\bar{c} = c/h = 4$, where h is the computational grid spacing and c is the inverse shape factor. For linear-MTA, 5 bpp were used. This lower value was selected because it represents the practical limit for linear-MTA in this case, due to the absence of the full Navier–Stokes penalization used in nonlinear MTA. For AGW, radii of 5 mm and 3.5 mm were used for the low- and high-density cases, respectively.

The reference grid size is $137 \times 73 \times 73$ points. The reconstruction error is evaluated using the L_2 norm normalized by the absolute maximum value of the corresponding reference field. The results for u , v , w , and p are reported in Table 2.

The results indicate that MTA is the best-performing method overall among the three considered approaches. We also briefly investigated the effect of increasing the number of bases per particle. For this purpose, the low-density case was reconstructed using MTA with 20 bpp . The resulting errors, computed using the full-density MTA reconstruction as reference, are reported in Table 3.

Figure 1 shows the reconstructed z -component velocity in the central z -plane. The linear-MTA reconstruction, Fig. 1d, is significantly smoother than the corresponding MTA reconstructions,

Error u				Error v			
Method	MTA	lin-MTA	AGW	Method	MTA	lin-MTA	AGW
MTA	0.040	0.034	0.061	MTA	0.074	0.067	0.067
lin-MTA	0.059	0.045	0.066	lin-MTA	0.104	0.087	0.078
AGW	0.081	0.071	0.067	AGW	0.106	0.096	0.068

Error w				Error p		
Method	MTA	lin-MTA	AGW	Method	MTA	lin-MTA
MTA	0.069	0.071	0.068	MTA	0.032	0.052
lin-MTA	0.116	0.093	0.082	lin-MTA	0.114	0.085
AGW	0.111	0.101	0.074			

Table 2. L_2 -normalized error for the subsampled reconstructions. Columns indicate the full-density reconstruction used as reference, while rows indicate the method applied to the subsampled dataset.

u	v	w	p
0.035	0.058	0.057	0.024

Table 3. Error obtained with MTA using 20 *bpp*, computed with respect to the full-density MTA reconstruction.

Fig. 1b and Fig. 1c. This is partly due to the simpler nature of the linear formulation, but also to the smaller number of bases per particle used by linear-MTA. As discussed above, linear-MTA only enforces the divergence-free constraint and is therefore less constrained than nonlinear MTA, which is additionally stabilized by the Navier–Stokes penalization. Comparing the two MTA reconstructions, the 20 *bpp* case appears slightly sharper and more defined than the 10 *bpp* case.

A similar behavior is observed for the reconstructed pressure, shown in Fig. 2. The main difference is that linear-MTA, Fig. 2d, exhibits visible difficulties near the boundaries of the reconstructed volume.

As a final assessment of this experiment, we evaluate the spectral content of the reconstructed fields. This provides a more quantitative measure of the resolution differences observed in the previous figures. In addition, we use the temporal fluctuations as an indicator of temporal consistency. For a generic quantity a , we define its fluctuating component as

$$\tilde{a} = a - \frac{1}{n_t} \sum_{k=0}^{n_t} a_k, \quad (6)$$

where the summation is performed over time. The quantity $\tilde{a}$ can therefore be interpreted as a temporal high-pass component, or equivalently as a proxy for the unsteady fluctuations.

Power Spectral Densities (PSDs) of $\tilde{w}$ and $\tilde{p}$ are computed at the middle time step and averaged over the YZ planes. The results are shown in Fig. 3. A premature detachment of the curves at




Figure 1. Comparison of the reconstructed z -component velocity.

low spatial frequencies can be interpreted as a loss of spatial resolution. Conversely, an overall bias with respect to the reference spectrum indicates a mismatch in the fluctuating content, and therefore a lack of temporal consistency at high temporal frequencies.

In Fig. 3, both AGW and linear-MTA show a clear bias with respect to the reference spectrum. This suggests the presence of excess fluctuation energy or temporally inconsistent high-frequency content in the reconstruction. AGW also appears to smooth the velocity field at significantly lower spatial frequencies. Both MTA reconstructions show better agreement with the reference spectrum, with the 20 *bpp* case being closer to the reference at low and intermediate spatial frequencies. The pressure PSD, Fig. 3b, shows a similar but more pronounced trend. In this case, MTA with 20 *bpp* is able to retain a larger portion of the high-wavenumber energy, although the spectrum still detaches from the reference before the particle-spacing-based cutoff.

The previous results suggest that the additional high-wavenumber energy retained by MTA may not be explained solely by the instantaneous Eulerian particle spacing. Since the reconstruction is performed in a 4D formulation, Lagrangian information carried by the particle tracks may contribute additional space-time constraints, especially in regions where the data are sparse or noisy.

In the present test case, the Eulerian interparticle distance is approximately

$$\Delta x_E \simeq 2.3 \text{ mm}. \quad (7)$$

In contrast, the material displacement between consecutive frames can be estimated as

$$\Delta x_L(\mathbf{x}, t) \simeq \Delta t \|\mathbf{u}(\mathbf{x}, t)\|_2. \quad (8)$$




Figure 2. Comparison of the reconstructed pressure.

This quantity represents the spatial sampling scale along a particle trajectory. In low-speed regions, Δx_L can be significantly smaller than Δx_E , implying that the time-resolved Lagrangian tracks may provide additional information beyond the instantaneous Eulerian sampling.

To test this interpretation, the PSDs of $\tilde{w}$ and $\tilde{p}$ were recomputed only in low-speed regions, defined by $\|\mathbf{u}(\mathbf{x}, t)\|_2 < 0.1$ m/s. The results are shown in Fig. 4. In these regions, the spectral modulation is shifted toward higher spatial frequencies, indicating that the reconstruction retains high-wavenumber content over a larger range than in the global PSD.

To quantify this effect, we report the spatial frequency at which the ratio between the PSD of MTA with 20 *bpp* and the reference PSD reaches 0.9. In the global PSD, this occurs at approximately $f_x = 111 \text{ m}^{-1}$, indicated by the dash-dotted line in Fig. 4. In the low-speed region, the same threshold is reached at approximately $f_x = 134 \text{ m}^{-1}$, indicated by the dotted line. Although this shift may appear modest in one dimension, if interpreted as an isotropic 3D sampling effect it corresponds to an effective sampling-density increase of approximately

$$\left(\frac{134}{111}\right)^3 \simeq 1.75. \quad (9)$$

To exclude the trivial explanation that this effect is caused by a locally higher seeding density, the average interparticle distance was also computed using only particles with velocity below 0.1 m/s. The ratio between the global and low-speed interparticle distance was found to be 0.97, indicating that the low-speed regions are, if anything, slightly more coarsely sampled. Therefore, the observed shift in spectral content cannot be attributed to a locally finer Eulerian particle distribu-




Figure 3. PSDs of the temporal fluctuations of w and p . The grey dashed line indicates the spatial frequency associated with the interparticle distance.

tion. Instead, it supports the interpretation that the 4D formulation is able to exploit Lagrangian space-time information in addition to the instantaneous Eulerian sampling.




Figure 4. PSDs of the temporal fluctuations of w and p computed only in regions where $\|\mathbf{u}(\mathbf{x}, t)\|_2 < 0.1$ m/s. The grey dash-dotted line indicates the modulation frequency identified in the global PSD, while the dotted line indicates the corresponding modulation frequency in the low-speed region.

4.2. Experiment 2

The second test case is obtained from a setup similar to the previous one, but the quality of the particle tracks is significantly lower. In particular, the dataset contains substantial positional uncertainty and a high concentration of ghost tracks caused by the reconstruction procedure. This experiment is therefore not intended as a quantitative accuracy benchmark, since no reliable ground truth is available. Instead, it is used as a large-scale robustness and computational stress test.

The dataset contains approximately 1.74×10^5 particles per time step over 1000 time steps, corresponding to 1.74×10^8 particle samples in total. Storing the particle data alone requires approximately 10 GB in single precision. With 5 bases per particle, the complete 4D reconstruction involves approximately

$$N_b = 1.74 \times 10^5 \times 5 \times 1000 \simeq 8.7 \times 10^8 \quad (10)$$

basis functions. If the velocity and pressure fields were represented explicitly, i.e. by independent weights for u , v , w , and p , the corresponding unconstrained velocity-pressure representation would contain approximately

$$4N_b \simeq 3.48 \times 10^9 \quad (11)$$

scalar coefficients. This number is comparable to the parameter count of a small-to-medium large language model, and is reported only to convey the scale of the representation. Unlike neural-network parameters, these coefficients are deterministic degrees of freedom of a physically constrained 4D flow reconstruction. Moreover, the proposed formulation does not optimize all explicit coefficients independently: incompressibility is imposed as a hard spectral constraint, and pressure is eliminated through variable projection. Nevertheless, the resulting active problem remains of billion-variable scale.

Before discussing the timings, it is important to distinguish between the linear and nonlinear formulations. Both methods require the evaluation of the velocity field at the particle positions through particle-evaluation operators. If these matrices were stored explicitly for all 1000 time steps, the memory requirement would be on the order of several hundred gigabytes, which is incompatible with the objective of running the method on consumer-grade hardware. Therefore, nonlinear MTA is executed in a matrix-free mode, where the particle-evaluation operators are computed on the fly and never stored. This substantially reduces the memory footprint, at the cost of additional computation.

The linear-MTA formulation is less constrained by memory, since it can process only a limited number of time steps simultaneously. In the present implementation, 32 time steps are processed in parallel, corresponding to the number of logical CPU cores. Therefore, the particle-evaluation matrices can be stored explicitly for the standard linear-MTA run. To isolate the computational impact of the matrix-free setting, we also report a matrix-free linear-MTA run.

All timings refer to the complete reconstruction pipeline, including weight computation, inference of velocity, pressure, gradients and vorticity, as well as input/output operations. The DA grid is used as inference grid in order to reduce the cost of the final field evaluation. The computations were performed on a PC equipped with an Intel Xeon w5-3435X CPU, 128 GB of RAM, and an RTX 4500 Ada GPU. The present implementation is CPU-only and does not yet exploit the GPU.

Table 4 reports the peak memory consumption and wall-clock time per time step. The nonlinear MTA reconstruction requires 55 GB of peak memory and 95 s per time step. This is only a single-digit factor slower than the stored-matrix linear-MTA baseline, which requires 14 s per time step, and approximately 3.5 times slower than the matrix-free linear-MTA baseline, which requires 27 s per time step. This is notable because nonlinear MTA solves a substantially richer physics-constrained 4D reconstruction problem, while also operating in a matrix-free setting. The optimized AGW implementation is included only as a fast interpolation reference.

Method	Peak memory use [GB]	Wall time per time step [s]
MTA	55	95
linear-MTA, matrix-free	25	27
linear-MTA	40	14
AGW	15	3

Table 4. Computational performance on the large-scale degraded experimental dataset. Timings include the complete reconstruction pipeline and I/O.

These results show that the proposed pseudo-spectral formulation makes a fully coupled, billion-scale 4D MTA reconstruction feasible in-core on PC type hardware. The nonlinear method remains within the same broad computational regime as the linear reconstruction, despite enforcing the Navier–Stokes constraint and eliminating pressure through the variable-projection procedure. Since the present implementation is still CPU-only, further reductions in wall time are expected from a GPU implementation, for which the algorithm is naturally suited due to its reliance on dense batched operations and FFT-based convolutions.

5. Conclusion

We have presented the pseudo-spectral formulation of MTA, which opens the door to the democratization of high-performance data assimilation algorithms. This is achieved through the kernelization of compact-support RBFs and the resulting efficiency and simplicity of the formulation. Beyond the significant improvements in computational time and memory footprint, an important aspect that was not sufficiently emphasized in Section 3 is the robustness of the method. All results presented in this work were obtained using the same non-dimensional weighting of the Navier–Stokes cost term. In practice, the reconstructions required minimal parameter tuning and were obtained largely on a first attempt. While robustness is inherently difficult to quantify and typically requires extensive long-term validation, the user experience associated with the method provides a strong qualitative indication of its stability and reliability. The target of achieving computational performance comparable to, or within one order of magnitude of, standard interpolation methods has been met. The nonlinear formulation is approximately 30 times slower than

AGW, while the linear formulation is about 4 times slower, placing both approaches within a practical range given their significantly higher physical fidelity. The most important next step is the implementation of a GPU-accelerated version of the algorithm. While a quantitative assessment of the expected speedup is beyond the scope of this work, the structure of the method—based on FFTs and dense operations—suggests that substantial performance gains are achievable.

References

- Godbersen, P., Gesemann, S., Schanz, D., & Schröder, A. (2024). Flowfit3: Efficient data assimilation of lpt measurements. In *21th international symposium on the application of laser and imaging techniques to fluid mechanics*.
- Golub, G. H., & Pereyra, V. (1973). The differentiation of pseudo-inverses and nonlinear least squares problems whose variables separate. *SIAM Journal on numerical analysis*, 10(2), 413–432.
- Jeon, Y. J., Müller, M., & Michaelis, D. (2022, apr). Fine scale reconstruction (VIC#) by implementing additional constraints and coarse-grid approximation into VIC+. *Experiments in Fluids*, 63(4). Retrieved from <https://doi.org/10.1007%2Fs00348-022-03422-9> doi: 10.1007/s00348-022-03422-9
- Sperotto, P., Petersson, P., Watz, B., & Hess, D. (2023). Instantaneous 3d flow field reconstruction with physics informed radial basis functions. In *15th international symposium on particle image velocimetry–ispiv*.
- Sperotto, P., Watz, B., & Hess, D. (2024). Meshless track assimilation (mta) of 3d ptv data. *Measurement Science and Technology*, 35(8), 086005.
- Wendland, H. (1995, dec). Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree. *Advances in Computational Mathematics*, 4(1), 389–396. Retrieved from <https://doi.org/10.1007%2Fbf02123482> doi: 10.1007/bf02123482
- Zaki, T. A. (2013, sep). From streaks to spots and on to turbulence: Exploring the dynamics of boundary layer transition. *Flow, Turbulence and Combustion*, 91(3), 451–473. Retrieved from <https://doi.org/10.1007%2Fs10494-013-9502-8> doi: 10.1007/s10494-013-9502-8