**Proceedings of the 9th International and 49th National Conference on Fluid Mechanics and Fluid Power (FMFP)**
**December 14-16, 2022, IIT Roorkee, Roorkee-247667, Uttarakhand, India**
**FMFP2022-6791**

# Multi-domain reduced-order model for store trajectory prediction

**Navdeep Pandey[1] and Aniruddha Sinha[1]**

[1]Department of Aerospace Engineering, IIT Bombay, Mumbai 400076, India

## ABSTRACT

Two-body aerodynamic analysis is crucial whenever there is a store that separates from its parent body, i.e., the aircraft. Whenever a new aircraft is developed or an existing one undergoes some modification with its associated store, it has to undergo a meticulous analysis to predict the path of the separated store across a range of operating parameters (freestream conditions). The present work demonstrates an efficient albeit approximate technique that employs a reduced order model based on proper orthogonal decomposition in conjunction with a multi-domain-decomposition approach for predicting the flow field around a store-aircraft dyad. Encouraging preliminary results are obtained in the validation that is pursued on a two-dimensional problem for simplicity. The approach can be readily extended to three-dimensional problems as well.

**Keywords**: MDDROM, POD, Store-trajectory.

## 1. INTRODUCTION

The aircraft-store separation analysis is crucial, whenever a new aircraft-store duo is developed or an existing one undergoes some design changes. In this analysis, trajectories followed by the store after its release from the parent body (i.e., the aircraft) are computed under different operating conditions (i.e., freestream parameters) to predict the safe-separation flight envelope. Although initial attempts employed expensive and risky flight tests, and subsequent efforts relied on wind tunnel model testing using captive trajectory system, almost all current store trajectory predictions are pursued using computational fluid dynamics (CFD).

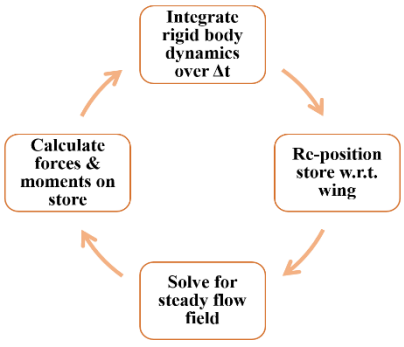


**Figure 1: Flow chart showing the iterative steps involved in store trajectory prediction using quasi-steady CFD.**

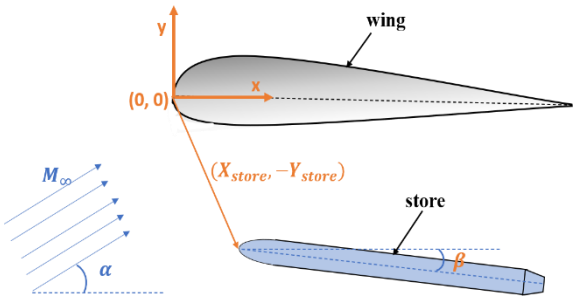


**Figure 2: Setup of the two dimensional 2-body store separation problem.**

The CFD-based approach to store-trajectory simulation is summarized in fig. 1. One starts with the initial configuration where the store is on the verge of separating from the aircraft. Invariably, a quasi-steady approach is employed, wherein CFD simulations of the flow field for any time instant assume that the flow is steady. Once the instantaneous forces and moments are obtained on the store from such a calculation, they are supplied to a 6 degrees of freedom (DOF) rigid body dynamics solver to determine the position of the store at the next time step. Steady CFD calculations are again conducted for this new aircraft-store configuration, and the simulation proceeds iteratively in this manner. The calculations are ended once the store is out of the influence zone of the aircraft. Even with the exponential development of computational power and memory, the quasi-steady CFD calculations are very resource intensive and take time. Hence, there is an opportunity for developing a method that can predict the approximate trajectory of the store within limits of practical applicability with minimal computational cost and time. The reduced-order model (ROM) approach discussed herein is one such empirical technique. ROMs are well known for their ability to predict flow fields efficiently with a small turnaround time [1–6], which makes them suitable for applications where rapid design decisions have to be made – e.g., multi-disciplinary analysis and optimization. In the context of ROMs, the well-resolved CFD approach is called full-order model (FOM); the latter is taken as the 'truth' solution against which the performance of the ROM is evaluated.

The overall objective of this work is to establish and validate a ROM to efficiently predict a store's trajectory under different operating conditions of the parent aircraft – viz. its Mach number $M_\infty$, angle-of-attack $\alpha$, etc. For simplicity, we have

considered a two-dimensional (2D) problem to demonstrate the proposed method. The problem setup shown in fig. 2 consists of two bodies – viz. a wing (i.e., an airfoil in 2D) and another body that represents the store whose aerodynamics are influenced by the wing. We attempt to predict the forces and moments acting on the store when it is positioned relative to the aircraft flying at a given ($M_\infty$, $\alpha$) condition. The actual computation of trajectory can be performed by integrating a 3 DOF rigid body solver with the present setup.

In the proposed ROM approach, we first parameterize the flow field using a minimal basis derived from proper orthogonal decomposition (POD) [7, 8]. This leverages the observed fact that apparently high-dimensional flow fields can often be approximated very well using a much lower dimensional embedding. The reduced-order POD basis is identified empirically, which means that it is based on a 'learning' database generated by a FOM. Subsequently, the coefficients of the basis modes towards the solution for a new case are determined in an optimization step that attempts to minimize the residual of the governing equations while satisfying the boundary conditions. The latter approach is the same as in any CFD, with the sole difference being the severely curtailed DOF to optimize due to the reduced POD basis being employed. POD requires that the learning database, as well as the new cases to be predicted, be for the same geometry and mesh. This is impossible in our store trajectory prediction problem since the store continuously changes its position and orientation relative to the aircraft. Hence, the usual POD-ROM approaches for single-body aerodynamics [1–6] cannot be employed for the entire flow domain. To circumvent this limitation, Ref. [9] proposed a POD-based domain decomposition ROM (DDROM) approach, wherein the overall flow domain was decomposed into one sub-domain that enveloped the aircraft and extended to the far-field boundary but had a 'hole' or 'dropbox' in it located below the aircraft, and the other sub-domain that accounted for the dropbox containing the moving store. The former 'staticzone' subdomain was the same across the entire learning database so that POD-ROM could be applied. The resulting solution was iteratively matched with FOM calculations on the dropbox. Since the majority of the mesh cells were in the staticzone, the use of POD-ROM promised efficiency gains. The present work is an extension of the DDROM approach, wherein further efficiencies are realized by decomposing the dropbox sub-domain itself into two subdomains, and applying POD-ROM to one of them again. It is called multi-domain-decomposed ROM (MDDROM). We demonstrate here that the MDDROM delivers reasonable accuracy vis-a-vis the FOM 'truth'.

## 2. METHODOLOGY

### 2.1 POD

Let us denote the flow vector field by $\boldsymbol{q}(\boldsymbol{x};\boldsymbol{\mu})$, where $\boldsymbol{x} := (x, y)$ is the 2D Cartesian coordinate and $\boldsymbol{\mu}$ is the parameter vector. For example, in a 2D problem governed by Euler equations, $\boldsymbol{q} = [\rho, \rho u, \rho v, p]^{\mathrm{T}}$, where $\rho$ is density, $u$ and $v$ are $x$- and $y$-components of velocity, and $p$ is pressure. In POD, we assume that the flow vector field can be approximated as

$$\boldsymbol{q}(\boldsymbol{x};\boldsymbol{\mu}) \approx \overline{\boldsymbol{q}}(\boldsymbol{x}) + \sum_{n=1}^{N_p} \eta^n(\boldsymbol{\mu})\widetilde{\boldsymbol{q}^n}(\boldsymbol{x}) \qquad (1)$$

where, $\overline{\boldsymbol{q}}(\boldsymbol{x})$ is the mean flow vector field (typically averaged over the flow solutions in the learning database) and the remaining 'fluctuations' are approximated as linear combinations of spatial basis functions $\{\widetilde{\boldsymbol{q}^n}(\boldsymbol{x})\}_{n=1}^{N_p}$ called POD modes, weighted by POD coefficients $\{\eta^n(\boldsymbol{\mu})\}_{n=1}^{N_p}$. For later reference, $\boldsymbol{\eta} := (\eta^1, \dots \dots \eta^{N_p})^T$. The actual determination of the POD modes follows the established 'snapshot' POD approach [6, 8, 10], and is not repeated here.

### 2.2 Reduced-order model

The ROM predicts the flow field for a new parameter vector $\boldsymbol{\mu_0}$ by invoking the governing equations. This is a more robust and accurate approach than the more straightforward interpolation in the parameter space. The POD-based ROM technique employed here was originally developed for single-body steady aerodynamics [1–5]; particular details relevant to this work may be found in Ref. [6]. Here we give only a brief overview. Let the vector of governing (unsteady) conservation equations and boundary conditions be represented as

$$\frac{\partial(\boldsymbol{C}(\boldsymbol{q}))}{\partial t} = \boldsymbol{R}(\boldsymbol{q}),\ \boldsymbol{x} \in \Omega,\ s.t.\ \boldsymbol{B}(\boldsymbol{q}) = \boldsymbol{0},\ \boldsymbol{x} \in \partial\Omega. \qquad (2)$$

Here, $\boldsymbol{C}$ is the operator that maps $\boldsymbol{q}$ to the vector of conserved flow variables, $\boldsymbol{R}(\boldsymbol{q})$ is a shorthand notation for the terms other than the local time derivative in the vector governing equations, $\Omega$ represents the flow domain, and $\boldsymbol{B}(\boldsymbol{q})$ = 0 codifies the conditions imposed on the boundary $\partial\Omega$. Since the solution $\boldsymbol{q}(\boldsymbol{x};\boldsymbol{\mu_0})$ must be steady, we should ideally have $\boldsymbol{R}(\boldsymbol{q}(\boldsymbol{x};\boldsymbol{\mu_0})) = 0$ along with $\boldsymbol{B}(\boldsymbol{q}(\boldsymbol{x};\boldsymbol{\mu_0})) = 0$. Just as in the FOM, we cannot hope for the ROM to find such a solution that exactly satisfies these conditions at all interior and boundary points in the flow; of course, the match is expected to be worse for the ROM. Once the POD modes are determined from the learning database, only the coefficient vector $\boldsymbol{\eta}$ is unknown in the approximate expansion of eqn. (1). Thus, we can write $\boldsymbol{R}\big(\boldsymbol{q}(\boldsymbol{x};\boldsymbol{\mu_0})\big) \approx : \widetilde{\boldsymbol{R}}(\boldsymbol{x};\boldsymbol{\eta}(\boldsymbol{\mu_0}))$ and $\boldsymbol{B}\big(\boldsymbol{q}(\boldsymbol{x};\boldsymbol{\mu_0})\big) \approx : \widetilde{\boldsymbol{B}}(\boldsymbol{x};\boldsymbol{\eta}(\boldsymbol{\mu_0}))$ This notation reinforces the fact that the residual at any point only requires knowledge of the POD coefficients, as does the boundary condition function. For a given $\boldsymbol{\mu_0}$, the optimization problem is posed as:

$$\min_{\boldsymbol{\eta}} \left\|\widetilde{\boldsymbol{R}}(\cdot;\boldsymbol{\eta}(\boldsymbol{\mu_0})\right\|_{\Omega,\mathcal{L}_p}^{p} \quad s.t. \quad \left\|\widetilde{\boldsymbol{B}}(\cdot;\boldsymbol{\eta}(\boldsymbol{\mu_0})\right\|_{\Omega,\mathcal{L}_p}^{r} < \epsilon. \qquad (3)$$

Here, $\mathcal{L}_p$ is the $p$-norm, and the domain of evaluation is indicated in the subscript of the norm too; $\epsilon$ is a suitably chosen threshold. In the present work, we have used the $\mathcal{L}_1$ norm of the residual. Moreover, there are so-called hyper-reduction techniques that drastically reduce the number of control volumes in which the residual has to be evaluated [6].

### 2.3 Multi-domain-decomposed ROM

Consider the two-body problem in fig. 3, where the store is under the influence of the wing and can take any arbitrary position relative to the wing. Due to this continuous change in the placement of the store, we cannot implement the ROM based on POD for the full domain as discussed in section 1. Following Ref. [9], we propose a methodology termed multi-

domain decomposition ROM (MDDROM) to circumvent this issue. As shown in fig. 3, we decompose the overall domain into multiple sub-domains as follows:

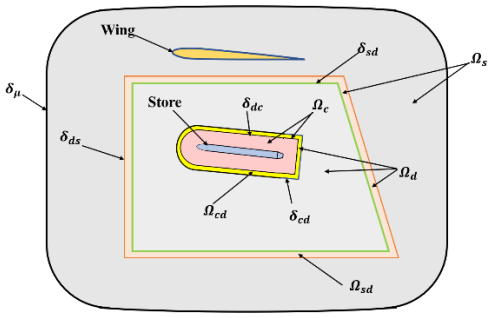


**Figure 3: Three domain-decomposed ROM approach for the two-body problem of fig. 2.**

• *Capsule*, $\Omega_c$ – region immediately surrounding the store, which comprises a grid that does not change as the store moves.
• *Dropbox*, $\Omega_d$ – maximal region enveloping the capsule where it may be expected to reach in its separation trajectory and still remain under the influence of the wing.
• *Staticzone*, $\Omega_s$ – remaining flow domain surrounding the wing and dropbox and extending to the far-field boundary.

The staticzone and dropbox sub-domains are designed to have a small overlap – the staticzone-dropbox overlap $\Omega_{sd}$ in fig. 3. Similarly, the overlap of the capsule with the dropbox yields $\Omega_{cd}$. With careful attention, the overall mesh may be designed so as to remain unchanged in the capsule and staticzone sub-domains. The dropbox is the only subdomain where the mesh needs to change in the course of the store separation. Evidently, POD-based ROM may be applied to $\Omega_c$ and $\Omega_s$, and it is only in $\Omega_d$ that one has to look for a FOM solution. But this region, being away from solid boundaries, should also have the least number of mesh cells, thereby potentiating significant computational savings. In the work of Ref. [9], the dropbox and capsule were a single sub-domain. By identifying the capsule subdomain containing a finely-resolved mesh around the store as a POD-ROM domain, we seek even greater efficiencies now.

The MDDROM approach is an iterative procedure as shown in fig. 4. The iterative process starts with computing initial solutions on $\Omega_s$ and $\Omega_c$ by interpolating POD coefficients in parameter space from the learning database. From these solutions, we extract the flow variable on interior interfaces $\delta_{ds}$ and $\delta_{dc}$, respectively (see fig. 3). These together serve as the (non-uniform) Dirichlet boundary conditions for the FOM solved in $\Omega_d$. From this FOM solution, we extract the flow variables on the interior interfaces $\delta_{sd}$ and $\delta_{cd}$ (see fig. 3 again).

The flow information on $\delta_{sd}$, along with the far-field boundary condition, and wall boundary condition on the wing, serve as constraints for the POD based ROM solution pursued in the staticzone sub-domain $\Omega_s$. Similarly, the flow information on $\delta_{cd}$ and the wall boundary condition on the store, fully specify the POD-based ROM to be solved for the capsule sub-domain $\Omega_c$. Now that we have the ROM solution in $\Omega_s$ and FOM solution in $\Omega_d$, the solutions are compared on the overlap regions $\Omega_{sd}$; a similar check is done for the other overlap region $\Omega_{cd}$.

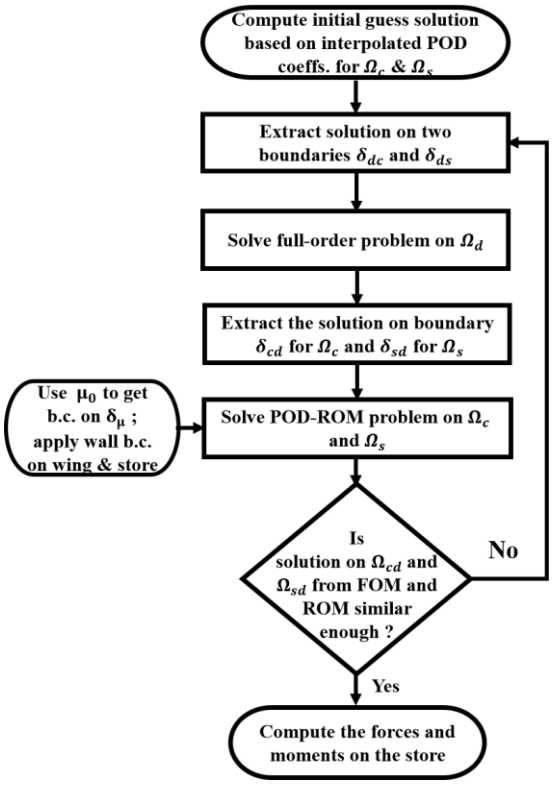


**Figure 4: Flow chart of steps involved in three domain-decomposed ROM for 2-body flow analysis.**

If the both comparisons are satisfactory, then we have a self-consistent solution over the entire flow domain. Else, we go for another iteration.

The learning database for the present work is generated by solving Euler equations; so, the ROM and FOM involved in the MDDROM will also use the Euler equations. It is shown in many of the earlier works on ROM for single body aerodynamics referred earlier that, even if the learning database uses RANS and the ROM minimizes Euler residuals, the predicted flow fields automatically satisfy the no-slip wall condition. Furthermore, the dropbox sub-domain being away from boundaries, one may solve the Euler equations therein too without significant penalty.

## 3. RESULTS AND DISCUSSION

We have evaluated results based on the proposed multidomain-decomposed ROM method for its validation. A two-dimensional domain having two bodies, a wing-like body (airfoil in 2D) and a store-like body under the influence of the wing is considered for the analysis fig. 2.

.

### 3.1 Geometry and mesh generation

A circular domain having two bodies, i.e., airfoil and store is constructed. The airfoil is a standard RAE-2822 profile, with chord $c$. The store is a rectangular slender body, with a nose and a boat-shaped tail. The length of the store is 0.5$c$ and its

slenderness ratio is 10, thereby making its thickness $d = 0.05c$. The nose is a circular-blunted tangent ogive, with nose radius and nose length as $0.05d$ and $1.5d$. The boat-tail has a half wedge angle of $20^0$ and its length equals $d$. A circle centered at the leading edge of the airfoil and having a radius $50c$ constitutes the far-field boundary. Gmsh v4.8.4, an open-source software, is employed for mesh generation [11]. A structured mesh is constructed around the airfoil and the store as well as for the overlap regions for ease of handling different domains, and an unstructured mesh is generated for the remaining domain. Figure 5, depicts the mesh topology for the domain, having quadrilateral cells for structured mesh and triangular cells for unstructured mesh. The overall mesh generation process is automated using Python scripts in integration with Gmsh. This allows us to ensure the same mesh (with the same node and cell numbering) exists in the capsule and staticzone subdomains for all positions of the store.

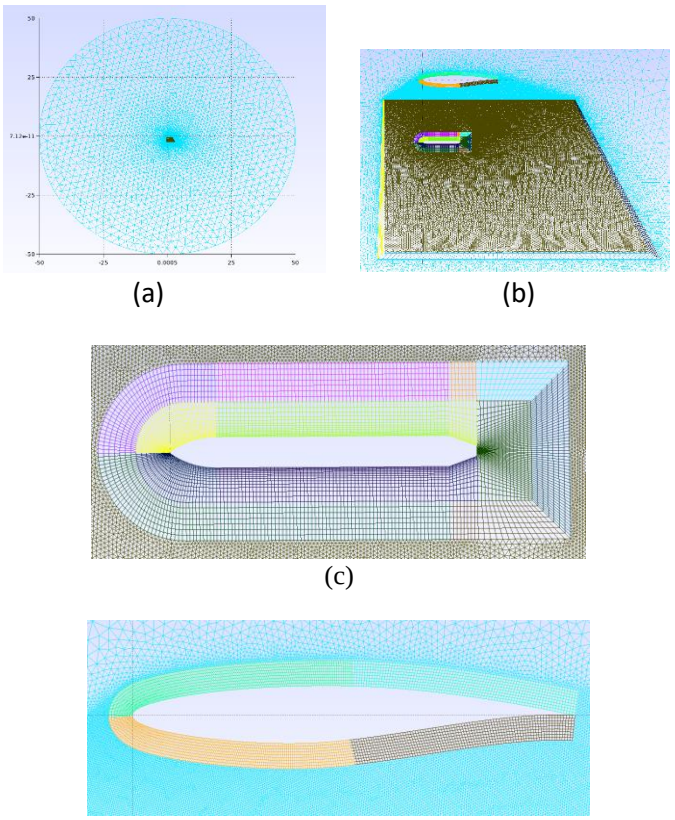


**Figure 5: Mesh topology for wing-store in 2D. (a) Overall domain, Zoomed views of (b) dropbox sub-domain, (c) store with body-fitted structured mesh as capsule sub-domain, and (d) wing with body-fitted structured mesh.**

### 3.2 Database generation and POD

The final learning database should comprise of snapshots of the entire store-separation trajectory for various choices of wing $M_\infty - \alpha$ pairs. However, for the present validation exercise, we treat these wing operating conditions and the three position DOFs of the store as independent parameters of a much smaller learning database. Thus, our parameter vector is $\boldsymbol{\mu} = (M_\infty,\ \alpha,\ X_{store},\ Y_{store},\ \beta)$ (see fig. 2). The sets of values used for these five parameters are given in table 1; all possible combinations of them yield 162 snapshots. As mentioned earlier, the present work uses Euler calculations, for which SU2 v7.1.1 [12] is used. The farfield boundary is provided with the free-stream condition, and the wing and store surfaces are given a no-throughflow condition. The database generation is automated using Python scripts that integrate Gmsh and SU2.

**Table 1: Operating parameters for learning database.**

| Parameter, (μ) | Values |
|---|---|
| Angle of attack ($\alpha$) | $0^0$, $1^0$, $2^0$ |
| Free stream Mach no. ($M_\infty$) | 0.45, 0.47, 0.49 |
| x-position of store ($X_{store}$) | 0.0, 0.2, 0.4 |
| y-position of store ($Y_{store}$) | -0.9, -1.0, -1.2 |
| Rotation angle of store ($\beta$) | $0^0$, $-1^0$ |

The conserved variables ($\rho$, $\rho u$, $\rho v$, $E$) ($E$ being the energy) obtained from the SU2 solution are converted into the chosen set of POD variables mentioned in section 2.1 (i.e., pressure is used instead of energy). Ref. [6] gives the rationale for the choice of this particular set of variables for POD. The flow variables are normalized such that they are of order unity. From the generated database, the flow solutions for staticzone and capsule sub-domains are extracted, and vector POD is performed on the snapshots in each subdomain separately for these two sub-domains.

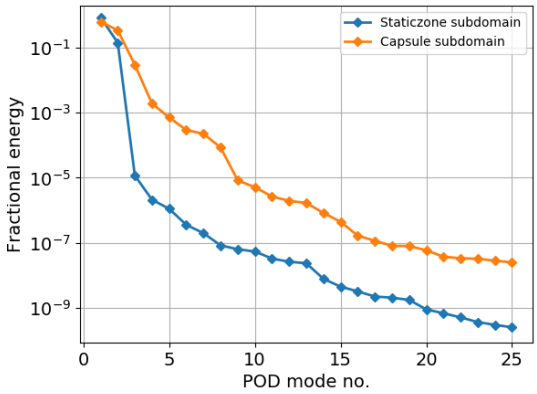


**Figure 6: Vector POD eigenspectra for the staticzone and capsule sub-domains in terms of the fractional energy content – viz. $\lambda_i / \sum_j \lambda_j$, where $\lambda_i$ is the $i^{th}$ POD eigenvalue.**

The fractional eigenspectra of POD for the two subdomains are shown in fig. 6. There is a rapid decrease (4 orders of magnitude) in the 'energy' content (or relative importance) of the POD modes from mode 2 to 3 in the staticzone sub-domain. The reason for this rapid decrease is that the first few modes capture the flow characteristics for the overall domain, whereas subsequent higher modes capture the localized flow features near the wing. On the contrary, the same is not true for the capsule sub-domain, because the capsule sub-domain consists of a very narrow region near the store, and all the parameters' variations affect the flow in this region. For the staticzone, POD mode 10 accounts for about $10^{-6}$ less 'energy' vis-a-vis POD mode 1; for the capsule, it is POD mode 15 that has about this relation with its POD mode 1. Therefore, for reconstructing the

flow field using ROM, we use 10 and 15 POD modes for the staticzone and capsule sub-domains, respectively.

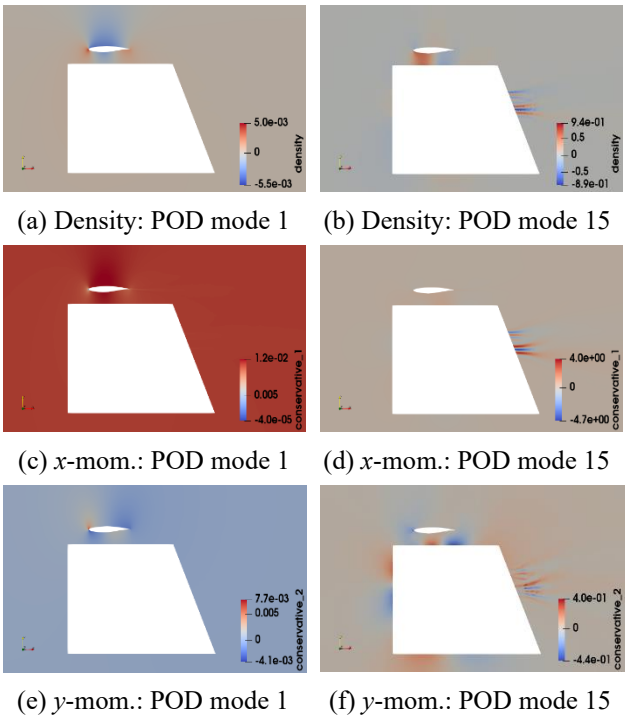


(a) Density: POD mode 1 (b) Density: POD mode 15

(c) *x*-mom.: POD mode 1 (d) *x*-mom.: POD mode 15

(e) *y*-mom.: POD mode 1 (f) *y*-mom.: POD mode 15

**Figure 7: POD mode 1 (left column) and mode 15 (right column) of the staticzone sub-domain.**

The shapes of POD modes 1 and 15 are presented in Figures 7 and 8 for the staticzone and capsule subdomains, respectively. Since the flows considered are weakly compressible, pressure follows density, and is hence omitted from the plots. The main motivation in presenting these extreme POD modes is to highlight the difference in the flow features captured by lower and higher POD modes. As stated previously, POD mode 1 of staticzone captures the flow characteristics for the overall domain, whereas its mode 15 counterpart displays localized flow features – viz. wakes of the store corresponding to its three *y*-positions included in the learning database. In the case of the capsule, POD mode 1 already displays some fine-scale structure; however, the structures in mode 15 are much finer.

### 3.3 Multi-domain-decomposed ROM validation

The MDDROM approach necessitates an iterative solution of the POD-ROM in the staticzone and capsule subdomains and the FOM in the intermediate dropbox sub-domain. As mentioned earlier, Euler residuals are minimized in the POD-ROM. The sequential least-square quadratic programming (SLSQP) method available in the scipy library of Python is employed for this optimization. The FOM solves the Euler equation using SU2, which has been augmented to implement the non-uniform boundary conditions arising at the interface between the sub-domains. The POD-ROM and SU2-FOM are driven iteratively from a Python script.

The MDDROM approach is validated using five new cases as delineated in table 2. Full-order solutions for these cases are computed using SU2's Euler solver and are used as truth solutions in the validation exercise. For illustration, representative flow variables from the case I computed using the MDDROM approach are compared with their 'truth counterparts in fig. 9. The visual agreement is very reasonable.

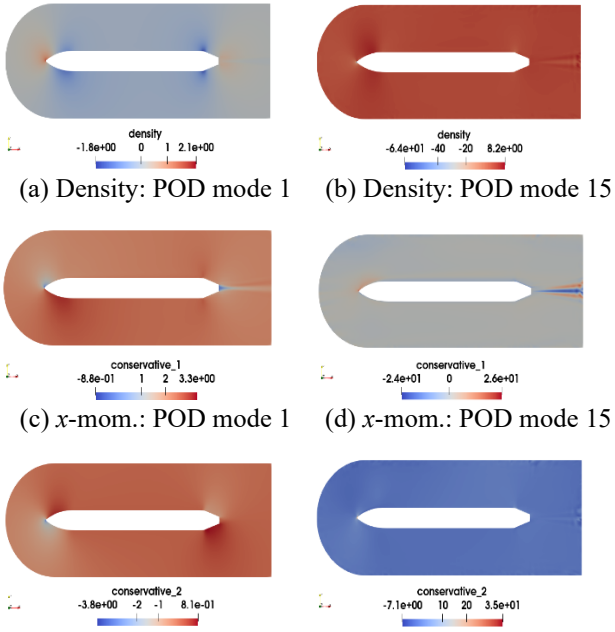


(a) Density: POD mode 1 (b) Density: POD mode 15

(c) *x*-mom.: POD mode 1 (d) *x*-mom.: POD mode 15

(e) *y*-mom.: POD mode 1 (f) *y*-mom.: POD mode 15

**Figure 8: POD mode 1 (left column) and mode 15 (right column) of the capsule sub-domain.**

The primary objective of the proposed methodology is to predict the store-separation trajectory by computing the aerodynamic forces acting on the store in an efficient manner. Hence, the store's lift coefficient is computed from the true solution and the MDDROM, and they are compared in the bottom half of table 2. The percentage absolute error is below 2%, which may be well within the range of practical applicability of the results.

**Table 2: Top half: parameters of validation cases for the MDDROM approach. Bottom half: comparison of the store's lift coefficient $c_l$ in the respective cases.**

| Case no. | I | II | III | IV | V |
|---|---|---|---|---|---|
| $\alpha$ | 0.50 | 0.50 | 0.95 | 1.2 | 1.2 |
| $M_\infty$ | 0.46 | 0.46 | 0.47 | 0.48 | 0.46 |
| $X_{store}$ | 0.20 | 0.20 | 0.20 | 0.40 | 0.40 |
| $Y_{store}$ | -1.00 | -0.90 | -0.10 | -1.20 | -1.20 |
| $\beta$ | $-1.0^0$ | $0.0^0$ | $-1.0^0$ | $0.0^0$ | $0.0^0$ |
| Truth $c_l$ | 0.110 | 0.041 | 0.134 | 0.073 | 0.072 |
| MDDROM $c_l$ | 0.108 | 0.041 | 0.133 | 0.073 | 0.071 |
| % error in $c_l$ | 1.80 | 0.0 | 0.70 | 0.0 | 1.40 |

## 4. CONCLUSIONS

This paper presents a POD-based ROM approach in conjunction with multi-domain decomposition methodology (MDDROM) for two-body aerodynamics problems. The overall motivation is to reduce the computational expense of store-separation trajectories prediction, as required in the

certification of the safe-separation flight envelope of aircraft-store dyads. The proposed approach replaces Euler simulations on the flow domain with a relatively faster multi-variable optimization problem for the major portion of the domain. It resorts to Euler calculations for a very small number of mesh cells that necessarily change when the store moves relative to the aircraft. The approach is validated on a two-dimensional two-body problem, where Euler calculations are pursued in the

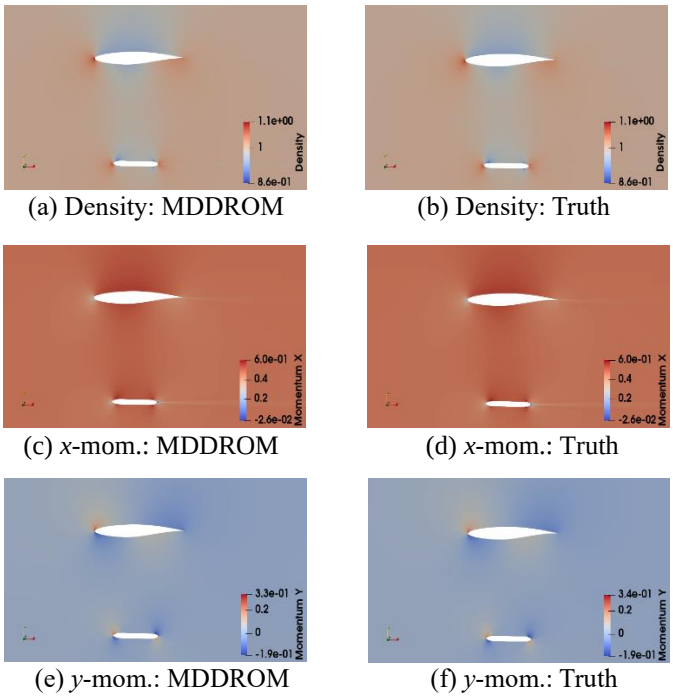


(a) Density: MDDROM (b) Density: Truth

(c) *x*-mom.: MDDROM (d) *x*-mom.: Truth

(e) *y*-mom.: MDDROM (f) *y*-mom.: Truth

**Figure 9: MDDROM solution (left column) and the 'truth' result (right column) for case I of table 2.**

full domain for the generation of the underlying empirical database. Comparison of MDDROM results with full-order Euler simulations reveals very encouraging agreement, both in terms of flow structures as well as the lift coefficient of the store.

For preliminary demonstration and validation of the proposed method, the low-fidelity Euler solver is implemented for FOM calculations at present. We propose to use RANS calculations in the future, and investigate if the MDDROM method continues to deliver reasonably accurate results. Moreover, the actual store-separation problem will also be addressed with the MDDROM.

## REFERENCES

[1] P. LeGresley and J. Alonso. Investigation of non-linear projection for POD based reduced order models for aerodynamics. In 39th Aerospace Sciences Meeting and Exhibit, AIAA paper 926, 2001.

[2] D. Alonso, J. M. Vega, and A. Velazquez. Reduced-order model for viscous aerodynamic flow past an airfoil. AIAA journal, 48(9):1946–1958, 2010.

[3] D. Alonso, A. Velazquez, and J. M. Vega. A method to generate computationally efficient reduced order models. Computer Methods in Applied Mechanics and Engineering, 198(33-36):2683–2691, 2009.

[4] D. Alonso, J. Vega, A. Velazquez, and V. Pablo. Reduced-order modeling of three-dimensional external aerodynamic flows. Journal of Aerospace Engineering, 25:588–599, 10 2012.

[5] R. Zimmermann and S. G ortz. Non-linear reduced order models for steady aerodynamics. Procedia Computer Science, 1:165–174, 05 2010.

[6] A. Sinha, R. Kumar, and J. Umakant. Reduced-order model for efficient generation of a subsonic missile's aerodynamic database. The Aeronautical Journal,126(303):1546–1567, 2020.

[7] J. L. Lumley. The structure of inhomogeneous turbulent flows. In A. M. Yaglom and V. I. Tatarsky, editors, Atm. Turb. and Radio Wave Prop., pages 166–178. Nauka, Moscow, 1967.

[8] P. Holmes, J. L. Lumley, G. Berkooz, and C. W. Rowley. Turbulence, coherent structures, dynamical systems and symmetry. Cambridge University Press, 2012.

[9] A. Sinha and S. Garg. Reduced-order modeling of steady aerodynamics for 2D store separation analysis. In 2018 Applied Aerodynamics Conference, AIAA paper 3168, 2018.

[10] Lawrence Sirovich. Turbulence and the dynamics of coherent structures. i - coherent structures. ii - symmetries and transformations. iii - dynamics and scaling. Quarterly of Applied Mathematics, 45, 10 1987.

[11] C. Geuzaine and J.-F. Remacle. Gmsh: A 3-d finite element mesh generator with built-in pre-and post- processing facilities. International Journal for Numer-ical Methods in Engineering, 79(11):1309–1331, 2009.

[12] F. Palacios, T. D. Economon, A. C. Aranake, S. R. Copeland, A. K. Lonkar, T. W. Lukaczyk, D. E. Manos- alvas, K. R. Naik, A. S. Padr´on, B. Tracey, A. Variyar, and J. J. Alonso. Stanford university unstructured (su2): Open-source analysis and design technology for turbulent flows. In 52nd AIAA Aerospace Sciences Meeting, AIAA paper 0243, 2014.