

MACH-B: FAST MULTIPOLE METHOD APPROACHES IN PARTICLE ACCELERATOR SIMULATIONS FOR THE COMPUTATIONAL AND INTENSITY FRONTIERS*

M. Harper Langston^{†,2}, Richard Lethin¹, Pierre-David Letourneau¹, Matthew J. Morse², Julia Wei¹
¹Reservoir Labs Inc., New York, NY 10012, USA

²Courant Institute of Mathematical Sciences, New York University, New York, NY 10003, USA

Abstract

The **MACH-B (Multipole Accelerator Codes for Hadron Beams)** project is developing a Fast Multipole Method [1–7] (FMM)-based tool for higher fidelity modeling of particle accelerators for high-energy physics within the next generation of Fermilab’s Synergia simulation package [8]. MACH-B incorporates (1) highly-scalable, high-performance and generally-applicable FMM-based algorithms [5–7, 9] to accurately model space-charge effects in high-intensity hadron beams and (2) boundary integral approaches [10–12] to handle singular effects near the beam pipe using advanced quadratures. MACH-B will allow for more complex beam dynamics simulations that more accurately capture bunch effects and predict beam loss. Further, by introducing an abstraction layer to hide FMM implementation and parallelization complexities, MACH-B removes one of the key impediments to the adoption of FMMs by the accelerator physics community.

INTRODUCTION

The majority of numerical approaches for accelerator multiparticle-tracking solve the macroscale problem by employing Particle-In-Cell (PIC) methods [8, 13–17]. These methods incorporate an Eulerian method for solving the necessary equations and Lagrangian techniques to advect particles through the domain (e.g., see Figure 1). The specific differences in PIC methods are in how mesh values and particle values are mapped back and forth.

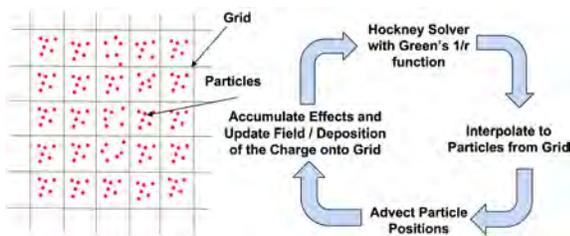


Figure 1: PIC-based Hockney solver. Given a cloud of charged particles, iterate (1) Grid charge deposition; (2) Compute potential; (3) Compute forces at grid points; (4) Compute forces at particle locations.

Since space-charge modeling in high-intensity hadron beams for the accelerator physics community requires scal-

* This work was supported by the U.S. Department of Energy as part of the SBIR Phase I Project DE-SC0020934.

[†] langston@reservoir.com

able and high-fidelity algorithmic approaches, all new computational approaches must satisfy the following:

1. **Are inherently multiscale:** Algorithms that can handle a variety of particle, field, and material distributions. For heterogeneous particle bunches (in terms of composition, number or size), PIC approaches are insufficient.
2. **Exploit locality:** While PIC approaches handle near-field effects appropriately, modeling these effects is often as much art as it is science. Algorithms should provide greater control over near-field interactions.
3. **Reduce expense of non-locality while handling accuracy:** When lower numerical accuracy is acceptable, new techniques should offer increased computational performance while reducing far-field interaction costs.
4. **Guarantee high accuracy when needed:** Prior to smoothing, macroparticles, or self-interactions, approaches must be inherently high-accuracy, solving the underlying partial-differential equations. PIC codes demand a very high-level of understanding to accurately manipulate, often requiring decades of expertise. Alternative approaches should have a lower barrier-to-entry, providing greater accessibility to adaptation by experts.
5. **Handle a variety of complex geometries:** Using a similar framework as the domain solver, new approaches need to solve boundary-dependent problems.

MACH-B TECHNOLOGY

MACH-B addresses the above five key elements, maintaining the strengths of PIC codes and approaches while further improving upon some of their weaknesses, allowing domain experts to evaluate and optimize various scenarios for complex high-energy physics experiments. The MACH-B technology is based on existing and new mathematical frameworks, providing new scalable, high-performance algorithms that will assist in accurately and rapidly computing a variety of complex particle accelerator simulations; specifically (1) **Fast Multipole Methods (FMM)** and (2) **Boundary Integral Solvers (BIS)**.

Combining these two technologies, MACH-B constructs an Embedded Boundary Solver (EBS) made up of (1) a multiscale, adaptive FMM solver for computing near and far-field interactions within the volume and (2) a boundary integral solver (BIS) based on Quadrature by Expansion

(QBX) [12, 18–20] for fast evaluation of layer potentials near and on a complex geometry. As an example, consider Figure 2: a complex shape with interior domain ω , boundary $\delta\omega$ and an inhomogeneous interior PDE of the form: $L(u)(x) = f(x) \in \omega$, $u(x) = g(x) \in \delta\omega$, where $L(u)$ represents the integral operator for the fundamental solution for the underlying PDE (in this case $L(u) = \Delta u$).

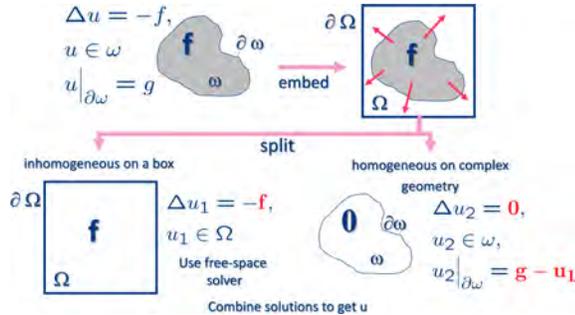


Figure 2: EBS approach for splitting a PDE into (1) an inhomogeneous problem in a simple domain and (2) a homogeneous problem on a complex domain.

The EBS approach performs the following steps in Algorithm 1, using FMMs for both major computational steps.

Algorithm 1: Embedded Boundary Solver

1. Embed the interior domain ω in a regular (e.g., a rectangular prism) domain Ω with boundary $\delta\Omega$ and solve $L(u_1)(x) = f(x)$ in Ω , where f is given at interior locations.
 2. Solve a boundary integral problem, where border data is modified from step 1:
 $L(u_2)(x) = 0, u_2 = g - u_1 \in \delta\omega$.
 3. Combine u_1 and u_2 to obtain the final solution u .
-

Fast Multipole Methods (FMM)

Originally designed for Coulomb interactions and belonging to the class of tree codes, FMMs achieve linear scaling by separating near- and far-field interactions (e.g., see Figure 3) on a spatial hierarchy using tree data structures. As they achieve arbitrary precision at modest cost with straightforward error estimates [1, 2, 4–6, 9, 10, 22–25], FMMs are well-suited for problems requiring high accuracy at large scales, such as in particle accelerator simulations.

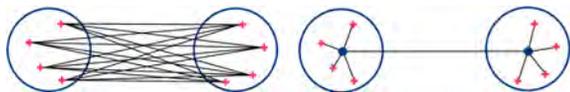


Figure 3: (Left): A naive $O(N^2)$ approach for computing the interactions between well-separated sources and targets. (Right): Using multipole and local expansions to reduce far-field costs, based on refinement.

FMMs are **inherently multiscale**, separating a regular domain into disjoint near- and far-field sets, using a tree

structure to **exploit locality** as well as **reduce the expense of non-locality** through low-rank approximation multipole expansions [1, 4]. FMMs compute the total field at target domain B as the sum of (a) the field due to the sources contained in its near field \mathcal{N}^B and (b) its far field \mathcal{F}^B . Contributions from \mathcal{N}^B are computed directly using dense summations, while contributions from \mathcal{F}^B are obtained by evaluating approximating expansion coefficients. These coefficients are constructed to achieve far-field low-rank approximations at pre-specified levels of accuracy for computationally-efficient and provably-accurate methods. Through the use of two simple parameters ((1) for points per smallest grid in the hierarchy and (2) for number of coefficients in the expansions), **high accuracy is guaranteed** [2, 3, 21].

The FMM uses *upward* and *downward passes* on a hierarchical tree structure, employing multiple operators for converting expansion (multipole and local) coefficients to **achieve optimal $O(n)$ complexity** (See Figure 4 from [4]).

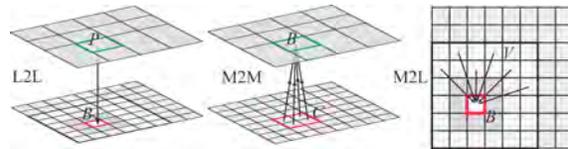


Figure 4: Local to Local (L2L), Multipole to Multipole (M2M) and Multipole to Local (M2L) operators translate coefficients efficiently throughout the FMM algorithm.

Boundary Integral Solvers (BIS)

For smooth/piecewise-smooth boundaries, such as those seen near particle accelerator pipe walls, boundary integral equation approaches (1) require no need for complex mesh generation for calculating potentials, (2) satisfy far-field boundary conditions, and (3) result in higher degrees of accuracy. At the beam pipe, a BIS can be specifically designed to couple with MACH-B’s proposed domain-based FMM solver to produce an embedded boundary solver (EBS). In cases where periodic or mixed boundary conditions may be required, FMMs can be tailored to handle these with minimal complexity [11, 12, 17, 19–21, 24]. Combining these two technologies, MACH-B constructs an Embedded Boundary Solver (EBS) consisting of (1) a multiscale, adaptive FMM solver for computing near- and far-field interactions within the volume (based on [6, 9]) and (2) a Quadrature by Expansion (QBX)-based boundary integral solver (BIS) package, called hedgehog [12] for fast evaluation of layer potentials near and on a source geometry.

NUMERICAL RESULTS

To study the accuracy of the FMM-based solver within Synergia, we performed a comparison between the accuracy of Synergia’s PIC-based Hockney solver (including interpolation from the grid to particles) methods and the FMM in a general case: a cloud of charged particles. For this purpose, we generated a 3D ensemble of charges (positive/protons)

containing up to 32768 particles. Then, we used both Synergia’s PIC-based solver and Synergia’s new FMM-based solver to compute all interactions, comparing to the naive $O(n^2)$ solutions. Results are shown in Table 1.

Parameter (M/p)	Relative error (PIC/FMM)
32/8	$1.41 \cdot 10^{-1} / 2.66 \cdot 10^{-7}$
64/8	$1.23 \cdot 10^{-1} / 2.66 \cdot 10^{-7}$
128/8	$9.92 \cdot 10^{-2} / 2.66 \cdot 10^{-7}$
256/8	$6.56 \cdot 10^{-2} / 2.66 \cdot 10^{-7}$

Table 1: Study of Synergia’s PIC-based solver vs. FMM solver for $\epsilon = 1e - 8$ FMM precision with 32768 particles per M -sized grid cell.

Comparing MACH-B’s FMM implementation, 3D Hockney (Synergia) and naive force computations on arbitrary point clouds, the FMM preserves its accuracy regardless of particle distribution, whereas the PIC-based methods, when macroparticles are absent, suffer from significant numerical errors associated with interpolation and finite differences.

MACH-B incorporates FMM-based boundary integral solvers based on the hedgehog software package [12], studying relative accuracy when evaluating potential and gradient solutions to Laplace’s equation with Dirichlet boundary conditions (*Top*: potential; *Bottom*: gradient) on a simple cylinder mesh as seen in Figure 5.



Figure 5: Magnitude of potential (top) and gradient (bottom) on cylinder for $\delta = 0.05$ and 168 mesh patches. Note that $+x$ axis faces outward, and $+z$ axis faces to the right.

We studied the behavior of the boundary solver for a cylinder with radius $0.5m$ and length $10m$. For a fixed distance δ and a point \mathbf{y} on the discretized boundary, we have the interior point $s(\mathbf{y})$, $s(\mathbf{y}) = \mathbf{y} - \delta n(\mathbf{y})$. For a source point $s = [1, 0, 0]$, we considered an FMM accuracy of $1e - 6$ and increasing levels of surface discretization (*Number of patches*) and evaluation distances of increasing value near the boundary. As can be seen in Table 2, we achieve desired levels of accuracy for both kernels even in regimes very close to the surface as well as further away.

For studying the full EBS solver, we modeled the potential of a single-bunch, Gaussian distribution of charges within a cylindrical conducting pipe. Approximated open boundary conditions at the ends (when a source distribution is concentrated locally) are achieved by extending the

Kernel	δ	Number of patches		
		168	672	2688
Pot.	0.05	$6.27 \cdot 10^{-5}$	$4.68 \cdot 10^{-6}$	$8.87 \cdot 10^{-7}$
	0.1	$1.70 \cdot 10^{-5}$	$3.56 \cdot 10^{-6}$	$6.54 \cdot 10^{-7}$
	0.2	$9.07 \cdot 10^{-6}$	$1.79 \cdot 10^{-6}$	$3.20 \cdot 10^{-7}$
Grad.	0.05	$8.43 \cdot 10^{-3}$	$3.17 \cdot 10^{-5}$	$5.84 \cdot 10^{-6}$
	0.1	$8.49 \cdot 10^{-5}$	$1.96 \cdot 10^{-5}$	$3.86 \cdot 10^{-6}$
	0.2	$5.53 \cdot 10^{-5}$	$1.20 \cdot 10^{-5}$	$2.25 \cdot 10^{-6}$

Table 2: Relative errors for potential and gradient kernels for $\epsilon = 1e - 6$ requested precision as we increase δ and the mesh discretization for a cylinder mesh (See Figures 6 and 7).

length of the pipe. We set the interior charge values to $p(\mathbf{x}) = \frac{1}{\sigma^3(2\pi)^{3/2}} \exp(-\frac{\|\mathbf{x}\|^2}{2\sigma^2})$ ($\sigma = 0.1$) for $N = 10^4$ source points in the interior along the Gaussian. Results for the x - y and x - z planes exhibit expected decay in Figures 6 and 7.

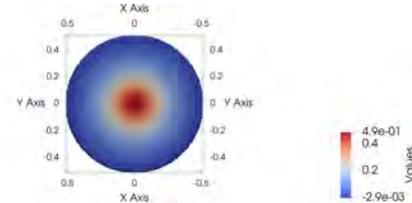


Figure 6: Cylinder cross-section at $z = 0$ for potential.

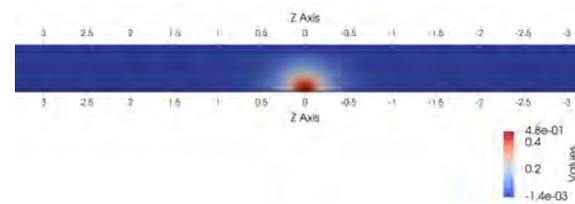


Figure 7: Cylinder cross-section at $y = 0$ for potential, showing only the $+x$ half-plane.

CONCLUSION

Fast Multipole Methods and Integral Equation Solvers for complex boundaries provide alternative approaches to standard PIC-based solvers for the particle accelerator community. Through MACH-B, we have shown the applicability of these methods within Synergia, and ongoing work continues to improve these methods and develop modules for other simulation frameworks. We envision that FMM-based tools will become a strong verification and analysis tool for domain experts as well as open opportunities for new types of simulations involving more complex bunch dynamics.

ACKNOWLEDGMENT

Thank you to Drs. Eric Stern and Qiming Lu at Fermilab for helpful conversations and guiding us through Synergia as well as Drs. Robert Ryne and Jean-Luc Vay at LBNL for many very fruitful discussions.

REFERENCES

- [1] L. Greengard and V. Rokhlin: *A Fast Algorithm for Particle Simulations*, Journal of Computational Physics, 73:2, 325–348, 1987.
- [2] L. Greengard: *fast Algorithms for Classical Physics*, Science, 265:5174, 909–914, 1994.
- [3] L. Ying, G. Biros and D. Zorin: *A Kernel-Independent Adaptive Fast Multipole Algorithm in Two and Three Dimensions*, Journal of Computational Physics, 196:2, 591–626, 2004.
- [4] M.H. Langston, L. Greengard, D. Zorin: *A Free-Space Adaptive FMM-Based PDE Solver in Three Dimensions*, Comm. in Appl. Math. and Comp. Sci., 6:1, 79–122, 2011.
- [5] I. Lashuk, A. Chandramowlishwaran, M.H. Langston, T. Nguyen, R.S. Sampath, A. Shringarpure, R.W. Vuduc, G. Biros: *A Massively Parallel Adaptive Method on Heterogeneous Architectures*, Proc. of ACM/IEEE SC’09.
- [6] D. Malhotra and G. Biros: *PVFMM: A Parallel Kernel Independent FMM for Particle and Volume Potentials*, Comm. in Comp. Phys., 18:3, 808–830, 2015.
- [7] W. Yan and R. Blackwell: *Kernel Aggregated Fast Multipole Method: Efficient Summation of Laplace and Stokes Kernel Functions*, arXiv preprint arXiv:2010.15155, 2020.
- [8] J. Amundson, P. Spentzouris, J. Qiang and R. Ryne: *Synergia: An Accelerator Modeling Tool with 3-D Space Charge*, Journal of Computational Physics, 211:1, 229–248, 2006.
- [9] W. Yan and M. Shelley: *Flexibly Imposing Periodicity in Kernel Independent FMM: A Multipole-to-Local Operator Approach*, Journal of Computational Physics, 355:15, 214–232, 2018.
- [10] L. Ying: *An Efficient and High-Order Accurate Boundary Integral Solver for the Stokes Equations in Three Dimensional Complex Geometries*, Ph.D. Dissertation, Courant Institute of Mathematical Sciences, New York University, 2004.
- [11] M.H. Langston: *An Efficient and High-Order Accurate Boundary Integral Solver for the Stokes Equations in Three Dimensional Complex Geometries*, Ph.D. Dissertation, Courant Institute of Mathematical Sciences, New York University, 2004.
- [12] M.J. Morse, A. Rahimian and D. Zorin: *A Robust Solver for Elliptic PDEs in 3D Complex Geometries*, arXiv:2002.04143, 2020.
- [13] A. Friedman, D.P. Grote, and H. Irving: *Three-Dimensional Particle Simulation of Heavy-Ion Fusion Beams*, Phys. Fluids B, 4, 2203–2210, 1992.
- [14] R. Hockney and J. Eastwood: *Computer Simulation Using Particles*, CRC Press, 1988.
- [15] J. Qiang, R.D. Ryne, Y.S. Habib and V. Decyk: *An Object-Oriented Parallel Particle-In-Cell Code for Beam Dynamics Simulation in Linear Accelerators*, Journal of Computational Physics, 163:2, 434–451, 2000.
- [16] P. Gibbon and G. Sutmann: *Long-Range Interactions in Many-Particle Simulation, Quantum Simulations of Complex Many-Body Systems: From Theory to Algorithms*, In Lecture Notes, J. Grotendorst, D. Marx, A. Muramatsu (Eds.), John von Neumann Institute for Computing, Jülich, NIC Series, 467–506, 2002.
- [17] J. Qiang: *Fast 3D Poisson Solvers in Elliptical Conducting Pipe for Space-Charge Simulation*, Phys. Rev. Accel. Beams 22:10, 104601, 2019.
- [18] A. Klöckner, A. Barnett, L. Greengard, and M. O’Neil: *Quadrature by Expansion: A New Method for the Evaluation of Layer Potentials*, Journal of Computational Physics, 252:1 332–349, 2013.
- [19] M. Wala and A. Klöckner: *A Fast Algorithm with Error Bounds for Quadrature by Expansion*, Journal of Computational Physics, 374:1, 135–162, 2018.
- [20] M. Wala and A. Klöckner: *A Fast Algorithm for Quadrature by Expansion in Three Dimensions*, Journal of Computational Physics, 388:1, 655–689, 2019.
- [21] R. Kress: *Linear Integral Equations*, Springer, 82, 1989.
- [22] A-K Tornberg and L. Greengard: *A Fast Multipole Method for the Three-Dimensional Stokes Equations*, Journal of Computational Physics, 227:3, 1613–1619, 2008.
- [23] H. Zhang and M. Berz: *The Fast Multipole Method in the Differential Algebra Framework*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment (The Eighth International Conference on Charged Particle Optics), 645:1, 338–344, 2011.
- [24] L. Ying, G. Biros, D. Zorin and M.H. Langston: *A New Parallel Kernel-Independent Fast Multipole Method*, Proc. of ACM/IEEE SC’03.
- [25] M.H. Langston, M.M. Baskaran, B. Meister, N. Vasilache and R. Lethin: *Re-Introduction of Communication-Avoiding FMM-Accelerated FFTs with GPU Acceleration*, IEEE Conference on High Performance Extreme Computing (HPEC’13), 2013.