Fast Trajectory Generation in High Fidelity Geopotentials using Finite Elements, Mascons, and Parallelism

AF STTR Phase I: AF09-BT02

Final Report v2
(closeout date 2-14-2011)

submitted to:
Analytical Mechanics Associates, Inc. (AMA)

by:

Ryan P. Russell
Nitin Arora

Georgia Institute of Technology
Atlanta Georgia, 30332

ryan.russell@gatech.edu

SUMMARY

The project objectives are to provide fast and accurate alternatives for calculating the geopotential and propagating the trajectory of a spacecraft therein. This report summarizes the activities performed by the Georgia Tech team, whose primary responsibilities are the geopotential calculation. Two fundamentally different solution methods for calculating the geopotential are considered: Local Weighted Interpolation (LWI) model (and Global Point Mascon (PMC) model. Here, we include high level details on the current status of both directions of research. In each case we have conducted detailed literature surveys and error analyses, explored multiple implementation strategies, chose preliminary paths to move forward, and have working prototype codes. The results are very promising according to the primary metrics that include: compute speed, accuracy, and memory. In experiments with both methods we demonstrate accuracies more than an order of magnitude lower than the published values for the spherical harmonics (SH) model. For the PMC we demonstrate approximately the same compute speed when performed serially on a CPU and one order of magnitude speedups when implemented on a GPU. For the LWI method, we demonstrate a global 200x200 field with ~400-800 times speedups requiring ~1.3 GB of memory.

INTRODUCTION

Trajectory propagation using conventional spherical harmonics (SH) in a high fidelity geopotential model is computationally burdensome, both in terms of wall clock time and floating point operations.[1] Furthermore, the fast implementations are recursive and therefore difficult if not impossible to make parallel. For real- or near real-time applications, such as the regular tracking and prediction of the space catalogue, it is simply not feasible given the current resources to account for complete high fidelity geopotential models in the special perturbation propagation techniques.[2,3] As the number of Earth orbiting objects continues to grow alongside with competing directives for the tracking resources, there is a clear need to develop faster techniques for computing high fidelity geopotentials.

Literature on gravity field formulation, estimation, and implementation (Geodesy) has a long rich history rooted in Earth and space sciences (see for example Refs. [4, 5, 6]). With the expansion of robotic spacecraft exploration missions to include irregular shaped celestial bodies such as comets and asteroids, there is renewed interest in alternative techniques for representing and calculating gravity fields.[7, 8, 9, 10, 11] A robust and elegant solution for irregular small bodies uses a polyhedral model although the extra computational requirements are cumbersome. Unlike conventional SH, the polyhedral models converge anywhere in the exterior of the surface, including inside the Brillouin (circumscribing) sphere. Similarly volumetric models composed of cube or sphere elements are suitable everywhere in the feasible domain. While the cube and sphere models lead to simpler computation requirements for each element, 3D models require a volume integral as opposed to the surface integral required for the polyhedral models. Furthermore, volume models suffer from large errors for evaluations near the surface. [7, 8] Neither the volume model nor the polyhedral models are justified at the Earth because spacecraft never pierce the Earth's Brillouin sphere. The extra computation required by the polyhedral methods is too expensive, and the volume models require too many elements to achieve sufficient resolutions.

An alternative method for computing gravity potentials involves large scale 3D interpolation models. Such methods are applicable for both irregular and near-spherical shaped bodies, and expedite computations by effectively trading computer memory for run-time speed. Essentially first proposed[12] by Junkins in 1976, the interpolation methods have been bolstered recently by the extraordinary memory resources of common computers. Depending on the interpolation method, a variety of techniques and basis functions are employed including weighting functions[12,13], wavelets[14], splines[14,15], octrees [16] and psuedocenters[17]. Each interpolation method balances accuracy with the desires to maximize runtime and minimize memory footprint while achieving

exactness, continuity and smoothness as appropriate. Despite their impressive speed gains, the main drawback of interpolation schemes is their intensive memory requirements as well as the implementation complexity. The second portion of this report will focus on a new modernized interpolation method. Here, we propose a new global adaptive finite element based method which relies primarily on a weighted interpolation scheme and uses a modified version of Junkin's finite element based gravity model. This new modified method includes scaled error control and an adaptive coefficient generation. The coefficients for the complete gravity model currently extend to the moon and are generated using a parallel version of the algorithm. At the expense of 1.3GB memory, we demonstrate more two orders of magnitude speed improvements for accurate representation of a 200x200 spherical harmonics field.

Finally, mascon models offer another alternative to SH and examples from the in the literature illustrate a diverse set of implementations and applications. In this study, a global point mascon (PMC) model is revisited for accurate and fast representation of high fidelity geopotential. Achieving the same or better levels of accuracy in comparison to the published SH fields, for the PMC model in this study, we demonstrate approximately the same compute speed when performed serially on a CPU and one order of magnitude speedups when implemented on a GPU.

SPHERICAL HARMONICS TRUTH MODEL

The base high fidelity geopotential model is chosen to be the GGM02C 200x200 SH field derived from GRACE spacecraft data augmented by terrestrial data for the higher frequency terms.[18] For purposes of this study, several truncations of the GGM02C are used as truth models for estimating PMC and LWI representations of varying fidelity. The SH code implemented for measurements and performance metrics is based on the Pines[19,20,21] singular-free formulation using the efficient non-singular recursion formulas from Ref. [[22]]. Consistency checks and a second speed benchmark are performed using an independent normalized version of the classic Legendre formulation.[23] Speedups for the LWI method are calculated with the Legendre formulation while speedups in the PMC approach are calculated with the Pines singular-free formulation. In our implementation the Pines is ~5-25% faster that the Legendre method depending on compiler settings and processor.

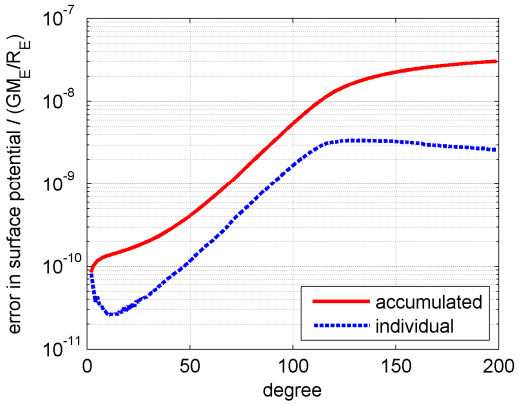

Figure 1: Expected uncertainty profile for GGM02C solution

Estimated accuracies of the GGM02C solution are given in the release notes from Ref. [[24]]. The accumulated error as a function of SH degree is replicated in Figure 1. For example, up to degree and order 70, the accumulated error for the geoid height is 7 mm or ~1×10^{-9} in normalized units. Up to degree and order 200, the normalized error is 3×10^{-8}. Therefore the confidence of the 200×200 and 70×70 field is approximately 8 and 9 digits of accuracy respectively. Furthermore, we note that a global analysis of both fields shows that they only agree to ~6 digits at worst and ~7-8

digits on average. To be discussed later, the accumulated error curve in Figure 1 is important as it will serve as a target for the residual error level in the PMC or LWI estimation problem. Such a curve effectively provides a calibration tool for selecting the appropriate number of mascons given a target SH field resolution.

CONTINUITY AND EXACTNESS

Most previous potential field interpolation methods have ignored continuity (even to first order in some cases) and simply used the forces (first derivatives of the potential) as three separate functions to fit. Generally, they argue that the discontinuities occur at a digit beyond the error tolerance, and are therefore insignificant. Alternatively, we seek to fit only the potential and directly differentiate our resulting interpolation function to obtain force. Aside from a minor potential speed benefit, this 'exactness' property is more elegant, and we speculate it will lead to better behaved (e.g. more consistent) numerics in terms of spacecraft trajectory propagation. Although the first order derivatives are of most importance to us, high order continuity is deemed important as well because many astrodynamics applications (e.g. orbit determination, optimization, targeting) require the second and third order derivatives of the geopotential.

POINT MASCON MODEL APPROACH

INTRODUCTION

The volumetric cube and sphere models used for small body gravity fields are two well known mascon models.[7, 8] The idea to use a collection of localized mass elements to augment geopotential computation is not new.[25] In fact, many of the original satellite geodesy applications in the early 1970's used spacecraft tracking data directly to fit mascon models consisting of point-masses or finite surfaces of constant density.[26, 27, 28, 29, 30] Their results demonstrated excellent agreement with state-of-the-art pure SH models for both the earth and the moon at the time (on the order of 10x10 fields). In particular, mascon models were attractive in the early days of satellite geodesy due to the limited amount of global data. Both terrestrial gravity anomalies and spacecraft tracking arcs were localized and combinations of the global SH with the localized mascon models provided higher resolution in regions with better data.[29] In the case of the moon, continuous tracking existed for coverage of the near side leading to strong gravity signatures, however the global short wavelength terms of SH were difficult to resolve due to lack of direct coverage on the far side.[27] Therefore a supplementary mascon model is naturally attractive at the moon and other planetary bodies where data are limited.[31, 32] Other regional applications of point mass models include local vertical modeling of the Earth's gravity field for geology and other Earth sciences.[33, 34, 35]

A known drawback of a mascon model composed of point masses is the singularity that exists at each of the elements. The simplest method to combat the singularity issue is to simply bury the point masses below the surface.[25] An alternative solution proposed by Koch[26,28,29], Morrison[30], Wong[27], and others is to replace the point-masses with finite masses with constant density profiles. Koch[26] and Koch and Morrison[29] introduced a tessellation of the earth surface into rectangles that reduces the geopotential computation to a summation of flat plate surface integrals (similar in spirit to the polyhedral method utilized for comets and asteroids[7]). Still evaluating over a connected surface layer composed of 2D rectangles, Morrison later shows that approximations of surface integrals using simple cubature rules are equivalent to point mass models with a higher resolution, and instead introduces more realistic cubature formulae.[26,30] Finally, Wong et. al.[27] sacrifice a connected surface layer in order to avoid approximations in the surface integral. They

distribute across the surface discrete masses with finite area, each represented by the potential of a circular disk. The exact potential of the disk avoids the singularity, but (similar to point mass or point sphere models) gaps or overlapping regions are unavoidable for any surface tessellation. Furthermore, while the disk enjoys a relatively simple expression for the potential compared to other finite shapes[27], the evaluation requires inverse trig functions and is much slower to compute than a simple point mass.

In this study, a global point mascon (PMC) model is revisited for accurate and fast representation of high fidelity geopotential. The gravity field inversion problem, particularly when related to PMC distributions, is well known to be ill-conditioned due to the unavoidable nature of the problem (poor observability).[7, 8, 34, 36, 37] Here, experiments are conducted with new model configurations in order to minimize the impact of the ill-conditioned inversion and simultaneously reduce the number of required mascons. Although it is emphasized that future work includes model fits using raw spacecraft measurement data, the scope of the current study is limited to fitting models to existing high fidelity SH models. Similar to most prior mascon and interpolation models, the proposed global PMC model fits the geopotential terms beyond the two-body plus J_2.

In the context of modern multi-core and multi-thread computer processors, the PMC model is attractive due to its simplicity and naturally embedded parallelism. Furthermore, the proposed model is memory light and is extremely simple to implement at any derivative level. Point masses are chosen over the finite shapes and surface layers for simplicity and to accommodate the high resolutions afforded by parallel computation on Graphics Processing Units (GPUs) where speedups are most favorable for problems exhibiting many thousands of independent, parallel calculations. Similar to the surface layer models, the proposed PMCs are distributed amidst a single two dimensional surface. To avoid the singularity and associated resolution problems at the surface, the mascons are buried. Unlike the mascon models proposed during the early age of spacecraft geodesy, the normal equations for reduction of a medium fidelity point mass model (e.g. 5,000 elements) can be readily solved in a few minutes on a common personal computer. Detailed numerical experiments are now feasible, thus enabling a fine tuning of important model parameters such as mascon distribution.

After a variety of experiments, the chosen solution method fixes the PMC locations and estimates only the associated gravitational parameters. Accordingly, the measurement model is linear and the mass estimation is reduced to an iteration-free linear least squares problem. The normal equations are solved using orthogonal projection methods that are known to be the most precision conserving and accurate solution methods available. The ill-conditioning is minimized through the judicious choice of the mascon distribution. The tessellation pattern (in latitude and longitude) of the point masses is found via solution to the Thomson problem from classical mechanics that seeks the minimum energy configuration of N electrons distributed on a sphere.[38,39] The point masses have a common global radius (or equivalently bury distance) that is selected according to an optimization on the Root Mean Square (RMS) of the several least squares solutions. The algorithm is run on a variety of cases that varies 1) number of mascons from 120 to 30,720 and 2) SH model sizes from 6x6 to 156x156. Currently the number of mascons is limited based on the amount of memory available on a single desktop server. The normal equation solution for the 30,720 mascon model requires ~16 GB or Random Access Memory (RAM). Future work includes parallel computation of the normal equation solutions using distributed RAM. The resulting model with 30,720 mascons matches the potential and acceleration evaluated at 100 km compared to the 156x156 SH model with an RMS of $5\text{x}10^{-11}$ and $5\text{x}10^{-9}$ normalized units respectively. The agreement in potentials at the surface is almost two orders of magnitude below the published accuracy of the SH model ($2\text{x}10^{-8}$). A single CPU implementation is found to be approximately equal compute speed compared to SH. Parallel GPU implementations lead to 13x speedups.

POINT MASCON MODEL FORMULATION

The PMC model is implemented using a linear least squares formulation. In the general

problem, an arbitrary number of mascons are placed at arbitrary locations within or on the Earth's surface. Measurements of the fitting function (the GGM02C solution truncated to specific degree and order) are taken at surface locations distributed approximately equally across the globe. If the locations of the mascons are fixed and the number of measurements exceeds the number of mascons, then minimizing the sum of the residuals between the mascon model and the measurements reduces to the classic weighted linear least squares problem with the following performance index J:

$$J = \sum_{i=1}^{m} w_i \varepsilon_i^{\,2} = \sum_{i=1}^{m} w_i \left[U_{PM}\left(\boldsymbol{\eta}_i\right) - U_{SH}\left(\boldsymbol{\eta}_i\right) \right]^2 \tag{1}$$

where there are m measurements, ε_i and $\boldsymbol{\eta}_i$ are the residual and location of the i^{th} measurement, U_{SH} is the measurement and is the potential evaluated according to SH, U_{PM} is the measurement model and is the potential according to the mascon model:

$$U_{PM}\left(\mathbf{r}\right) = U_{2B+J_2} + \sum_{j=1}^{n} GM_j \Big/ \left|\mathbf{r} - \boldsymbol{\rho}_j\right| \tag{2}$$

where there are n mascons, GM_j is the gravitational parameter for the j^{th} mascon, $\boldsymbol{\rho}_j$ is the location of the j^{th} mascon, and U_{2B+J2} is the potential due to the two-body plus J_2 terms:

$$U_{2B+J_2} = \frac{GM_E}{r}\left[1 - J_2\left(\frac{R_E}{r}\right)^2\left(\frac{3z^2}{2r^2} - \frac{1}{2}\right)\right] \tag{3}$$

where z is the third component of $\mathbf{r}$ and the GM_E, R_E and J_2 values come from the GGM02C SH solution and are the gravitational parameter, radius and un-normalized oblateness parameter for the Earth respectively (R_E=6378.1363 km, GM_E=398600.4415 km^3/s^2, J_2=0.0010826356665511). For this study, normalized units are chosen so that GM_E and R_E are unity (1 LU=6378.1363 km, 1 TU=806.810991306733 s).

As most previous authors of both interpolation and mascon models have observed, there is a tremendous benefit if the two body and J_2 contributions are removed from the fitting function. In normalized units the two body and J_2 contributions are ~1 and 1 × 10^{-3} units respectively. When removing both, the sum of all the other terms adds up to ~1 × 10^{-5} units. As an example, if the goal is to target ~10 absolute digits of accuracy, fitting only the higher order terms leaves just ~5 relative digits to match. Note that experiments were performed with removing a few extra terms beyond J_2. After sorting the remaining SH terms in descending magnitude, four specific terms (C_{22}, S_{22}, C_{31}, S_{33}) each are found to be more than 50% larger than the others. However, in numerical experiments, residuals did not appreciably improve when removing these extra terms beyond J_2. Therefore, the simpler J_2 only reduction is kept according to Eq. (2) and (3). Like the SH formulation (and unlike most interpolation schemes), it is important to note that the resulting PMC geopotential model is continuous to any order and 'exact' in the sense that the accelerations are gradients of a conservative potential.

The least squares minimization problem stated in (1) reduces to the classical normal equations:

$$\left(\mathbf{H}^T\mathbf{W}\mathbf{H}\right)\mathbf{x} = \mathbf{H}^T\mathbf{W}\mathbf{y} \tag{4}$$

where $\mathbf{x}$ and $\mathbf{y}$ are the unknowns and measurement vectors respectively:

$$\mathbf{x}_{n\times 1} = \begin{bmatrix} GM_1 \\ GM_2 \\ \cdots \\ GM_n \end{bmatrix}, \quad \mathbf{y}_{m\times 1} = \begin{bmatrix} U_{SH}\left(\boldsymbol{\eta}_1\right) \\ U_{SH}\left(\boldsymbol{\eta}_2\right) \\ \cdots \\ U_{SH}\left(\boldsymbol{\eta}_m\right) \end{bmatrix} = \mathbf{H}\mathbf{x} + \boldsymbol{\varepsilon} \tag{5}$$

where $\mathbf{g}$ is the measurement model vector and $\mathbf{H}$ and $\mathbf{W}$ are the sensitivity and weight matrices respectively:

$$\mathbf{g}_{m\times 1}=\begin{bmatrix} U_{PM}\left(\boldsymbol{\eta}_1\right) \\ U_{PM}\left(\boldsymbol{\eta}_2\right) \\ \cdots \\ U_{PM}\left(\boldsymbol{\eta}_m\right) \end{bmatrix},\quad \mathbf{H}_{m\times n}=\frac{\partial \mathbf{g}}{\partial \mathbf{x}}=\begin{bmatrix} \frac{1}{\left|\boldsymbol{\eta}_1-\boldsymbol{\rho}_1\right|} & \frac{1}{\left|\boldsymbol{\eta}_1-\boldsymbol{\rho}_2\right|} & \cdots & \frac{1}{\left|\boldsymbol{\eta}_1-\boldsymbol{\rho}_n\right|} \\ \frac{1}{\left|\boldsymbol{\eta}_2-\boldsymbol{\rho}_1\right|} & \cdots & \cdots & \cdots \\ \cdots & \cdots & \cdots & \cdots \\ \frac{1}{\left|\boldsymbol{\eta}_m-\boldsymbol{\rho}_1\right|} & \cdots & \cdots & \frac{1}{\left|\boldsymbol{\eta}_m-\boldsymbol{\rho}_n\right|} \end{bmatrix},\quad \mathbf{W}_{m\times m}=\begin{bmatrix} w_1 & 0 & \cdots & 0 \\ 0 & w_2 & \cdots & 0 \\ \cdots & \cdots & \cdots & 0 \\ 0 & 0 & 0 & w_n \end{bmatrix} \tag{6}$$

It is well known that the inverse gravity problem is ill-conditioned. Numerically the problem is combated using householder rotations, an orthogonal solution method, to solve directly the minimum norm problem in Eq. (1). The orthogonal solution is implemented with the standard dgelsy routine from LAPACK. Note that the dgelsy routine requires the entire dense **H** matrix as input. For a model with 30,000 mascons, and assuming a 2:1 ratio of measurements to unknowns, the **H** matrix would include 1.8 billion entries occupying 13.7 GB of memory when represented in double precision. In order to save memory, experiments were performed to accumulate sequentially the symmetric $\mathbf{H}^T\mathbf{W}\mathbf{H}$ matrix that includes only $n(n+1)/2$ entries, then solve the linear system in Eq. (4). Important information in the trailing digits of the $\mathbf{H}^T\mathbf{W}\mathbf{H}$ matrix is irrecoverably lost in the process of accumulating $\mathbf{H}^T\mathbf{H}$, and therefore the memory improvements are not deemed worthy of the performance hit in terms of residuals. As an example, the same level of residuals can be achieved with ~40% fewer mascons (representative improvement) when solving the equations via the orthogonal methods. Therefore, the improved numerical treatment due to the orthogonal rotations justifies the high memory requirements. Also, for obvious reasons, it is preferred to augment the i^{th} row of **H** with a multiplier of $w_i^{1/2}$ instead of dealing directly with the large sparse **W** matrix.

It is also worth mentioning that experiments were performed using quad arithmetic, and similar performance gains are possible. However, due to the due to precision loss associated with the dynamic range of the terms in the mascon summation in Eq. (2), the improved performance afforded by the quad precision solution to the normal equations is only maintained if the runtime summation is also computed in quad precision. Therefore, the remainder of the study will proceed with double precision in efforts to reduce compute runtimes and retain compatibility with GPUs.

Radial Mascon Distribution

We find from numerical experiments and based on intuition that the ill-conditioned nature of the problem is largely a function of the geometry (observability) between the locations of the measurements and the mascons. If the mascons are too close to one another, their individual gravity signature degrades. If the mascons are too far from one another the resolution of the resulting composite function degrades. This competition between resolution and ill-conditioning points toward an optimization problem that allows the mascon geometry to adjust. Ideally, each of the three coordinates of each mascon in addition to its mass value should be a free parameter in the least squares problem. However, this approach leads to a nonlinear measurement model and the number of unknowns increases to $4n$. The nonlinear least squares optimization problem may not converge at all or may require hundreds or more iterations, noting that a single solution to the normal equations with 10^4 unknowns is on the order of 1 CPU day.

Instead, to remain practical, the least squares problem is kept linear by fixing the mascon distribution and solving the linear problem inside each iteration of an optimization loop. The global mascon radius acts as the single optimization variable because it is particularly sensitive to the resolution and conditioning. Each linear problem has a global solution in just one iteration, and the one-dimensional optimization problem is solved via a simple, gradient-free, quadratic polynomial interpolation method, and generally converges in five to ten iterations. The independent variable is the global radius (equivalently bury distance) for the mascon locations while the performance index is the RMS of the residuals of the linear least squares problem. Figure 2 gives a diagram overview

of the proposed global mascon model.

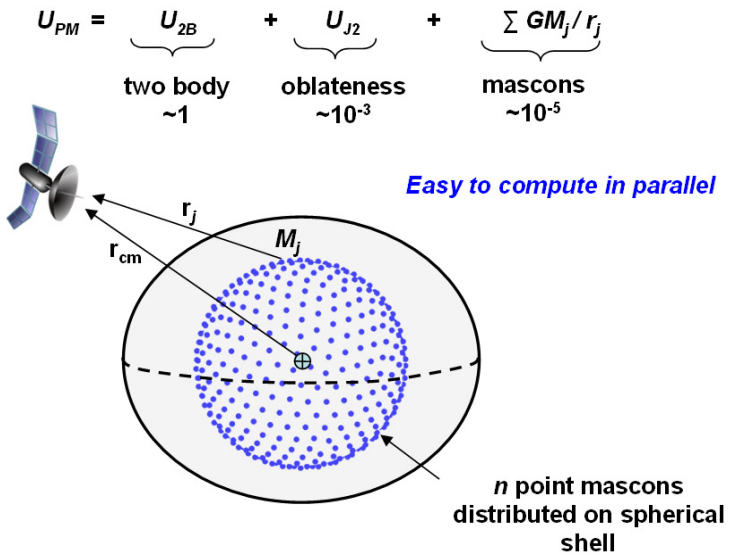

Figure 2: Point mascon model

The Thomson Problem and Lateral Mascon Distribution

To further improve conditioning, the lateral distribution of the mascons is chosen to approximate a solution to the classic "Thomson Problem," which poses the equivalent question of how to equally space *N* points across the surface of a sphere.[38,39,40] Approximate solutions are obtained with simple spiral algorithms that divide lines of latitude into parallel bands of equal area and place nodes along the spiral at a spacing consistent with the distance between coils. In this study a variant of the spiral algorithm from Ref. [40] is used as an initial guess to a rudimentary steepest descent algorithm that adjusts the node locations in the direction of a minimum energy configuration. The full problem is equivalent to an all-in-all *n*-body gravitational simulation and the compute times increase exponentially with the number of nodes. The solution is considered converged when the total system energy changes less than 1 part in 10^{12} from one iteration to the next. Depending on the quality of the initial guess and size of *n*, solutions can take on the order of 100 to 100,000 iterations to converge. For this study, solutions are found and archived for $n = 30\times2^q$ for $q = 0\rightarrow12$. This specific sequence of numbers allows for up to the order of 100,000 mascons, and is specially chosen to be efficient in the case of parallel computations on a GPU where computing threads are available in large blocks. The *n*=122,880 solution required on the order of 1 CPU day to compute (using an Intel Xeon 3.2 GHz processor). While other more efficient algorithms are possible, the details of the Thomson problem and solutions are not the focus of the current work. Figure 3 illustrates an example initial guess and converged solution for the case of $n = 960$. While the spiral algorithm provides a good initial guess as seen in the top row, visually it is apparent in the bottom row that the solution achieves a more equal distribution of spacing between the nodes. When the Thompson solution spacing is used for the mascon distribution, numerical experiments reveal that the RMS of the residual errors can be improved by a representative 50% when compared to using the spiral algorithm spacing.

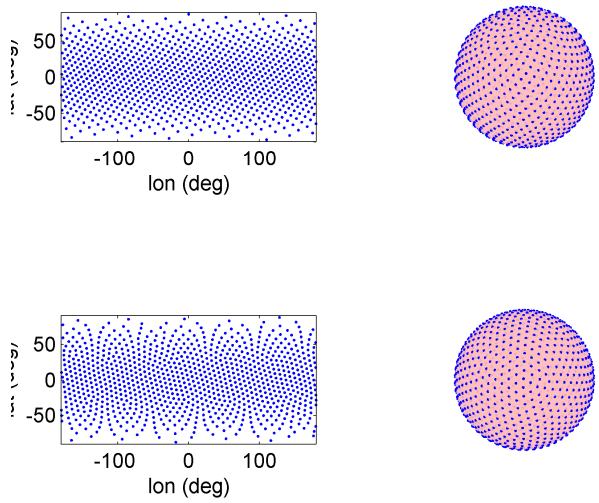

Figure 3: Thomson problem, initial guess from spiral algorithm (top row) and converged solution (bottom row) for* n=*960.

Note, that the Thomson problem spacing is ideal in the current application where the fitting function is SH and therefore global with similar frequencies across the whole domain. This global nature highlights a main drawback of SH, namely that it is impossible to increase the resolution locally using the SH basis function. Alternatively mascons are well suited for multi-resolution applications, like those of the early days of geodesy or in the case of planetary moons or other small celestial bodies where good data are not available globally. For the Earth, high quality global measurements from satellites (and a pair of satellites in the case of GRACE) are now available making SH well suited for geopotential models. It is speculated that this consistency between quality global data and the SH global basis function has contributed to the lack of general interest in mascon models in the past few decades. Although the proposed mascon models in this study are fixed to a global resolution due to the SH fitting function (and an intended global domain), it is trivial to add or remove local resolution by adding or removing mascons in regions of interests. For example, consider the problem where a reference path for a guided missile or other vehicle is known a priori. One could include a fine resolution of mascons near the ground track of the vehicle, and a course resolution otherwise. The resulting potential evaluations for the vehicle would be highly accurate, and require fewer computations, yet would only be valid in the domain near the ground track. Such a scenario is impossible using SH.

Constraints on the total mass and dipole moment

It is sensible that the sum of the masses in the current problem formulation will result to approximately zero because the true surface integral of the SH fitting function evaluated over the full surface is zero. As a high order precaution, extra care is taken to ensure that the sum of the mascon masses is indeed zero in order to maintain a consistent mean motion for a spacecraft as compared to the two-body only model. Therefore, the desired constraint is:

$$\sum_{j=1}^{n} GM_j = 0 \tag{7}$$

The least squares solution indirectly solves the problem to first order, noting from experience that in double precision the normalized mass summations of the solutions are typically on the order of 10^{-12} or smaller. In order to constrain that the total summation is exactly zero (to the numerical extent possible) according to Eq. (7), the Thomson problem and the least squares problem is solved instead with n-1 mascons, and the final n^{th} mascon is chosen to exactly negate the summation of the first n-1 terms:

$$GM_n = -\sum_{j=1}^{n'} GM_j \qquad n' = n-1 \tag{8}$$

As will be discussed later, significant errors accumulate in the summation term in Eq. (8) due to a variety of reasons including order of operations and large magnitude variations. To avoid such considerations at this stage, the summation term in Eq. (8) is performed in extended arithmetic utilizing quad precision (16 bytes), and the resulting negated sum is truncated to double precision and stored in GM_n.

The position of the n^{th} mascon is located, if possible, such that the dipole moment of the mascon distribution is the zero vector. It is tempting to require that the center of mass is located at the origin, however the center of mass definition includes the mass summation in the denominator and therefore can be singular when negative masses are allowed. Instead the numerator only is considered and is defined as the dipole moment (from electrostatics where charges can be negative). Accordingly Eq. (9) requires that the total dipole moment is zero, and has the effect of requiring the gravity forces to be exactly radial when evaluating the field far from the body.

$$\sum_{j=1}^{n} GM_j \boldsymbol{\rho}_j = \mathbf{0} \tag{9}$$

Then, analogous to Eq. (8), the position of the n^{th} mascon is computed (also in quad precision for the same reasons) from Eq. (9):

$$\boldsymbol{\rho}_n = -\frac{1}{GM_n}\left(\sum_{j=1}^{n'} GM_j \boldsymbol{\eta}_j\right) \tag{10}$$

If the final mass location is too close to the surface we enforce the safeguard in Eq. (11) and live with the resulting small violation to Eq. (9). Note that the extra mass will cause a slight decrease in the quality of the solution. Post-fit checks are performed to ensure that the contribution of the additional term is well below the RMS levels.

$$\boldsymbol{\rho}_n = \min\left(0.3R_E, \rho_n\right) \times \boldsymbol{\rho}_n / \rho_n \tag{11}$$

An alternative method to enforce the constraints in Eqs. (9) and (10) is to impose formal constraints on the least squares problem in the form of extra 'measurements' or penalty constraints in the performance index. According the original performance index from Eq. (1) can be augmented as:

$$J' = J + \left[w_s \left(\varepsilon_{m+1}\right)^2 + \sum_{i=1}^{3} w_{di} \left(\varepsilon_{m+1+i}\right)^2 \right] \tag{12}$$

and the terms in Eqs. (5) and (6) are appended to become:

$$\mathbf{y}_{m+1:m+4} = \begin{bmatrix} y_{m+1} \\ y_{m+2} \\ y_{m+3} \\ y_{m+4} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \quad \mathbf{g}_{m+1:m+4} = \begin{bmatrix} g_{m+1} \\ g_{m+2} \\ g_{m+3} \\ g_{m+4} \end{bmatrix} = \sum_{j=1}^{n-1} \begin{bmatrix} GM_j \\ GM_j \left(\rho_x\right)_j \\ GM_j \left(\rho_y\right)_j \\ GM_j \left(\rho_z\right)_j \end{bmatrix} \tag{13}$$

$$\mathbf{H}_{m+1:m+4,1:n-1} = \begin{bmatrix} 1 & 1 & \cdots & 1 \\ \left(\rho_x\right)_1 & \left(\rho_x\right)_2 & \cdots & \left(\rho_x\right)_{n-1} \\ \left(\rho_y\right)_1 & \left(\rho_y\right)_2 & \cdots & \left(\rho_y\right)_{n-1} \\ \left(\rho_z\right)_1 & \left(\rho_z\right)_2 & \cdots & \left(\rho_z\right)_{n-1} \end{bmatrix}, \quad \mathbf{W}_{m+1:m+4,m+1:m+4} = m \begin{bmatrix} w_s & 0 & 0 & 0 \\ 0 & w_{d1} & 0 & 0 \\ 0 & 0 & w_{d2} & 0 \\ 0 & 0 & 0 & w_{d3} \end{bmatrix} \tag{14}$$

where w_s, w_{d1}, w_{d2} and w_{d3} are weights that can be adjusted to influence the constraints. The quality of the solution residuals is not highly sensitive to these weights (for low values), noting that the constraints indeed can be driven lower at a very modest expense in the RMS of the residuals. However, numerical experiments reveal that weights too large can degrade accuracy for function evaluations at altitudes above where measurements were taken. Weights less than or $O(10^2)$ give reasonable performance. For results presented in this study, both approaches (explicit constraints and the additional mass) are utilized. Each component of the dipole sum must be smaller than the

mass sum in order to keep the final mass location in Eq. (10) close to Earth center and avoid the Eq. (11) safeguard. Therefore, it is generally suggested that w_{di} is larger than w_s.

Lastly, it is acknowledged that the summation constraint could also be solved explicitly (by removing the last mass from the unknown vector but keeping it in the measurement model) instead of including it in the performance index. However, an equivalent treatment to the dipole constraint makes the measurement model nonlinear and therefore not appealing.

Target Levels for Residual Errors

The curve in Figure 1 provides a calibration for residual error level targets in the mascon estimation problem. To remain conservative, it is desired to achieve RMS residual errors of approximately one order of magnitude lower than the published accuracy of the original SH GGM02C solution. The moving target changes depending on the fidelity of the field that is being fit. The highest fidelity field used in the current study is 156×156 and corresponds to a normalized error of 2×10-8 or ~8 digits of confidence. Therefore the residual RMS error target for the associated mascon solution is O(10-9). The mascon solutions that fit lower fidelity fields, such as a 10×10 truncation, target residual errors O(10-11). Accordingly, the residuals for the mascon estimation solutions will be masked by the accuracy of the original SH GGM02C solution. These residual targets provide a calibration for the selection of number of mascons given a target degree SH field, or inversely, the selection of the degree SH field given a target number of mascons. Note that the accelerations derived from the SH model and the mascon model will have larger discrepancies due to the differentiation. However, the original SH solution is fit using potentials (not accelerations) and errors compared to the unknown truth are also introduced when differentiating the SH potential.

PROCEDURAL DETAILS

In this section details are presented about the procedure and its implementation.

Summation procedure

The choice of summation method can have a significant effect on the result of multi-argument summation,[41] especially in cases such as the mascon application where terms include both signs and vary in size by several orders of magnitudes. As an example a 30,720 mascon solution modeling the 156×156 SH field includes normalized mass term magnitudes as small as ~1×10^{-5} and as large as ~10. Therefore, the mass magnitude range is up to ~6 digits, and this ill-conditioning is a main culprit in precision loss. To efficiency combat this precision loss to the extent possible, two well-known summation techniques are suggested.[41] The first is a divide and conquer algorithm known as 'pairwise' or 'cascade' summation and requires the same number of additions as a naive sum but includes minor overhead for extra loops (and function calls in the case of a recursive implementation). A second slightly more accurate (but slower) method known as 'Kahan' or 'compensated' summation that uses an additional term that accumulates small errors.

From experiments with the mascon models (and consistent with results in the literature), it is found that pairwise summation implemented with ~6 recursive calls (see Algorithm 1)) provides an extra ~1-2 digits of precision in comparison to a naive summation, and comes at a modest speed penalty of a few percent. The Kahan approach only provides an extra ~half digit of precision over the pairwise approach, and therefore doesn't justify the extra compute time for most applications. Note also that the sort order of the summation vector can also play a large role in precision loss or preservation. For this application, an ascending order sort of the mascon magnitudes is found to be marginally preferred when compared to a variety of other sort orders. It is emphasized that the suggested sort order and pairwise summation method results only from cursory experiments. Users of the published PMC models are encouraged to test other methods for potential improvement. Readers are referred to Ref. [41] and the references therein for details

on summation algorithms. Finally, it is also noted that a parallel summation algorithm designed for speed will use a divide and conquer method similar to that from the pairwise algorithm. Therefore, it is expected that parallel results from the GPU will achieve similar precision levels to that of the CPU using pairwise summation.

Algorithm 1: Pseudo code for recursive pair-wise summation. For this study, p=int(n/50).

RECURSIVE:

```
sum=pairwise_Sum(x,n,p)
```

- input: **x** vector of *n* elements, *p* tuning parameter controlling number of recursions
- output: *sum*

```
IF (n < p)
        sum = naive_Sum(x)
ELSE
        i = n/2
        sum = pairwise_Sum(x(1:i),i,p) + pairwise_Sum(x(i+1:n),n-i,p)
```

Scaling parameters

The runtime required to solve the linear least squares problem increases system increases with $O(n\times m) \approx O(\gamma\times n^2)$, where a new multiplier variable is defined such that

$$m = \mathrm{int}\left(\gamma \times n'\right) \tag{15}$$

where γ relates the number of mascons to the number of measurements. A reasonable value for γ (and the value chosen for the results of this study) is 1.6 noting that larger values cause unnecessarily long runtimes for the linear least squares problem while smaller values may cause poor observability and result in residual statistics that are not representative of the full domain.
The number of features in the SH function topology is of $O(d^2)$, where *d* is the degree and order of the SH fitting function. Therefore, a second scaling parameter is introduced such that:

$$n' = \mathrm{int}\left(\alpha \times d^2\right) \tag{16}$$

Or similarly

$$d = \mathrm{int}\left(\sqrt{n'/\alpha}\right) \tag{17}$$

where *α* relates the number of mascons to the resolution of the SH fitting function. The specification of *α* provides control over the resolution of the mascon model, and naturally accounts for the squaring effect of the field size. In practice, it is found that $\alpha \approx 1.2$ is the lower limit and approximately represents the Nyquist limit for capturing the highest frequencies of the SH function. Scaling the mascon numbers according to Eq. (16) or (17) is preferred because a single value for *α* results in similar magnitudes of residuals independent of the number of mascons or size of the SH field. To be discussed further in later sections, values of in the range $1.2 < \alpha < 2.6$ lead to approximate global fit errors of $\sim 1\times10^{-8} > \mathrm{RMS}(\boldsymbol{\varepsilon}) > \sim 1\times10^{-13}$ respectively.

Considering the new parameter *α* the runtime complexity is $O(n\times m) \approx O(\gamma\times\alpha\times d^4)$. Therefore, the complexity grows with fourth power of the size of the SH fitting function. As discussed already, the maximum value of *d* in the current study is 156 based primarily on memory limitations but also due to runtime considerations using a single desktop server. While the largest solution run in this study $\{n' = 30719, \alpha = 1.25, \gamma = 1.6\} \rightarrow \{d = 156, m = 49{,}150\}$ requires on the order of 1 CPU day to complete, a similar run with *d* cut in half (while keeping *α* and γ) leads to a $2^4 = 16$ fold reduction in runtime. Note that the memory requirements are of the same order as the runtime complexity.

Optimization loop

In the optimization loop, the RMS of the residuals of the linear least squares solution is minimized over the single global mascon radius or equivalently the bury distance, b. This optimization problem is performed using a non-gradient based quadratic polynomial interpolation scheme. The method samples the solution space initially at a course resolution and when a minimum is detected between three neighboring points, a quadratic polynomial is interpolated and

the next iteration is initialized using a finer resolution and a sampling centered at the interpolated minimum. The algorithm iterates until the RMS errors of the three points are of the same order of magnitude. Any one dimensional optimization method will work, however, the described method is chosen as a simple solution that balances complexity with required function calls.

The optimization loop generally converges in 5-15 iterations depending on the quality of the initial guess. For small field solutions, the runtime for 15 linear least squares calls is trivial. To reduce the number of ite-rations for the large field solutions where the runtimes are very significant, groups of solutions are sought for a given γ and α, and a sequence of n values. A near-linear relationship is found between n and the optimized bury distance, b* when n is plotted on a log scale. Accordingly the sequence for n is chosen such that n ap-proximately doubles in successive problems (also leading to efficient block sizes for use on the GPU as dis-cussed before). Then, to improve the initial guess for b in the case of runs with large values of n, a simple linear extrapolation is performed based on the b* values resulting from the previous two solutions.

Algorithm Summary

In summary, Algorithm 2 is a pseudo algorithm for computing a sequence of different sized PMC solutions each with approximately similar residual error statistics. By changing the input value of α, the algorithm can be repeated as desired for different error targets. Meaningful error targets are provided in Figure 1 according to the expected errors of the geoid height for the SH fitting function.

Algorithm 2: Point mascon solution generation

Main Inputs: $\{\alpha, n_{max}\}$ where:

- α is primary mechanism to control the residual error levels (reasonable values $1.2<\alpha<2.5$)
- n_{max} is the largest number of mascons desired (the upper limit for current implementation using a single processor with 16GB RAM is ~32,720).

A. Set general tuning parameters $\{\gamma, w_j, w_s, w_{d1}, w_{d2}, w_{d3}\}$ where

- γ adjusts the ratio of measurements to mascons (reasonable value ~2.0).
- w_j is the weight for each of the m measurements (reasonable value ~1)
- w_s and $[w_{d1}, w_{d2}, w_{d3}]$ are weights to optionally control mass and dipole constraints (reasonable values < $\sim 10^2$)

B. Initialize counters: $q=0$, $n=60$

C. DO WHILE ($n \le n_{max}$)

1. Increment loop: $q=q+1$, $n=n\times 2$

2. Set resolution of mascon field: $n'=n-1$

3. Set resolution of SH field d from Eq. (17)

4. Set number of measurements m from Eq. (15)

5. Set the latitude and longitude coordinates for m measurement locations via the spiral algorithm or Thomson problem solution if available. Set radius of each location to 1 mean Earth radius.

6. Set the latitude and longitude coordinates for n' mascons via Thomson problem solution

7. Choose the global mascon bury distance b_* and solve the associated least squares problem

CASE A: Solve the 1D optimization problem for RMS(ε) as function of b:

- Initial b guess: IF $q>2$ THEN extrapolate $b(n, n_{q-1}, n_{q-2}, b_{*q-1}, b_{*q-2})$ ELSE $b=0.5$
- Minimize the RMS of the residuals $\boldsymbol{\varepsilon}$ resulting from the fixed b linear least squares problem of Eq. (5) and described orthogonal solution method with definitions for **y**, **W**, and **H** from Eqs. (6), (7), (13), and (14). Computation of **y** is independent of b and thus only necessary on first iteration.
- Record the resulting b_* and associated solutions GM_j for $j=1\rightarrow n'$ mascons

CASE B: Choose predetermined b value (based on previous solution optimization for example) and solve least squares problem

8. Archive the mascon solutions (GM_j and $\boldsymbol{\rho}_j$ for $j=1\rightarrow n'$ mascons) in ascending order of $|GM_j|$

9. Solve for GM and $\boldsymbol{\rho}$ of the n^{th} mascon from Eqs. (8), (10) and (11) to enforce the mass and dipole constraints

Main Outputs: GM_k and $\boldsymbol{\rho}_k$ for $k=1\rightarrow n$ mascons for each of the q solutions

POINT MASCON MODEL SOLUTIONS

To demonstrate representative PMC models, the algorithm is applied for a medium size maximum resolution of n_{max}=7,680 mascons and is repeated over a broad range of α values spanning 0.7 to 2.55 at increments of 0.05. In this preliminary run γ is set to 1.6, the measurement weights are set to 1, w_s=10, w_{d1}= w_{d2} =10,000 and w_{d3}=10,000.

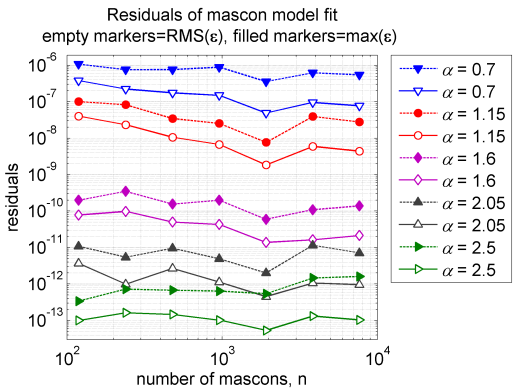

Figure 4: Broad survey of PMC solutions: residuals vs. number of mascons

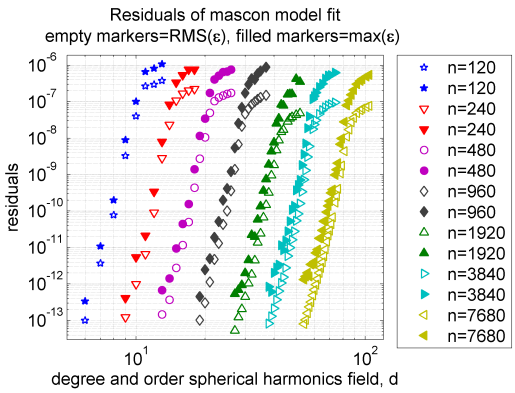

Figure 5: Broad survey of PMC solutions: residuals vs. size of SH field

The results in terms of residual statistics, resolution of the SH field, and number of mascons are illustrated in Figure 4-Figure 8. Figure 4 shows a subset of results for sampling of runs with α increments of 0.45. Each run is represented by two nearly horizontal sets of points where dashed lines and empty markers give RMS of the residuals and the full lines and filled markers give max residuals. The five runs demonstrate the general trends that 1) order of magnitude of residuals is approximately fixed for a given α, and 2) the residuals improve with increasing α. Figure 5 shows the data for all runs (every α) with the resolution of the SH field on the abscissa. Lines are not connecting the points resulting from a single run because the natural ordering occurs for constant n as opposed to the constant α (per the particular design of the algorithm). In this representation, each family of constant n reveals a hook toward the top of the plot in the 10^{-8} to 10^{-6} low resolution range. Below the hook, each family exhibits a near linear relationship between residuals and degree of the SH fitting function in this log-log scale. The linear region below the hook exist for approximately $\alpha \geq 1.2$ for all cases and indicates the region where reducing d (i.e. increasing α) leads to efficient improvements in accuracy. This α~1.2 boundary represents the Nyquist limit for the minimum number of mascons to accurately capture the high frequency terms in the SH fitting function. Figure 6 illustrates the RMS data from Figure 5 except in terms of α. Because d as computed from Eq. (17) is not single valued in terms of α, redundant calls occur across runs with small changes in α when n' is small. To better illustrate the resolution characteristics as a function of α, these redundancies in Figure 6 are removed. The efficiency boundary at α ~1.2 is evident in Figure 6 where the curvature switches sign. Note, also that this boundary is evident in Figure 4 from the non-linear spacing of the families.

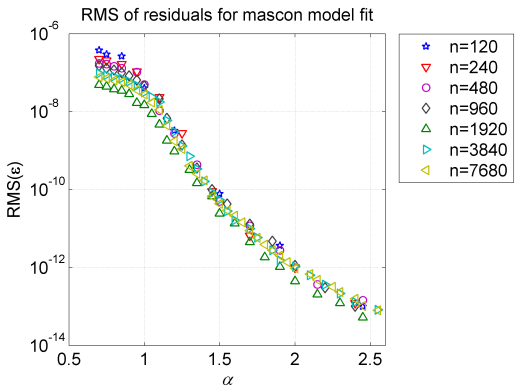

Figure 6: Broad survey of PMC solutions: residuals vs. α. Redundant runs removed.

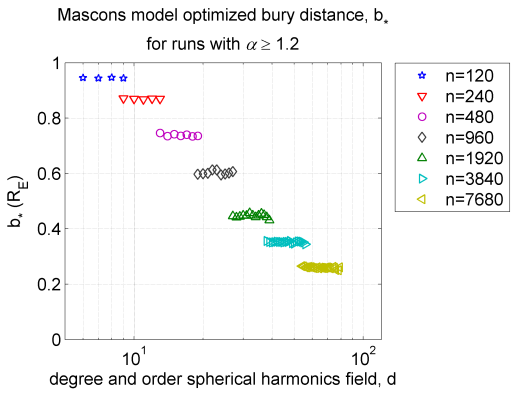

Figure 7: Broad survey of PMC solutions: bury distance vs. size of SH field

Considering only runs with α above the 1.2 boundary, Figure 7 demonstrates the noteworthy result that the optimized bury distance is (approximately) independent of d for a given n (noting the small variations may be due to the loose convergence tolerances in the RMS optimizations). Accordingly, this result suggests that the algorithm as presented only needs to be run once with a high value of n_{max}. Then, on successive runs with different α values, the optimization portion of the algorithm can be initiated or simply restricted to one iteration using the stored b_* values from the original α run (see step 7 CASE B in Algorithm 2). Furthermore, because the normal equations need only to be solved once, multiple measurement vectors **y** representing different truncations of the fitting function can be handled simultaneously. The circle markers in Figure 8 illustrate the mean of each group of solutions with the same n from Figure 7 as a function of n. The near-linear trend over short intervals in Figure 8 is useful for both interpolation and extrapolation for future runs with a different number of mascons.

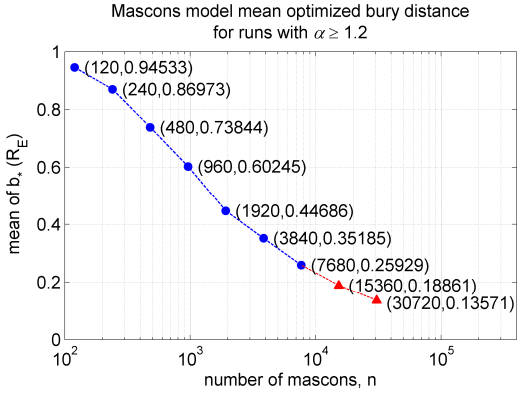

***Figure 8: Mean optimized bury distance. Circles include data from preliminary run (Figure 7) while triangles include data from extending runs with α=[1.25, 1.5, 1.75] to large* n.**

From Figure 7, the largest degree and order SH field is 79 and results when n is highest (7680) and α is lowest (1.2) considering only the efficient range of $\alpha \geq 1.2$. To accommodate higher order SH fields, the algorithm is rerun for larger n_{max} with just a few targeted α values based on the error curves in Figure 1 and the residuals achieved in Figure 6. Included in Figure 8 are data from these new runs (α = 1.25, 1.5 and 1.75) that extend to the 30,720 mascon case. Long runtimes prohibited all the α cases from being extended to large n. As discussed before, a two-fold increase on d leads to a four-fold increase in n and a 16-fold increase in runtime and memory requirement.

A practical limit of 16 GB RAM is imposed. This constraint leads to an upper bound on n_{max} at 32,720 mascons (when γ = 2) and requires on the order of 1 CPU day to solve the normal equations considering a non-parallel implementation on a single CPU (Xeon 3.2 GHz). For α = 1.25, the n=32,720 solution leads to a 156×156 SH field.

When evaluating potential and acceleration comparisons to the SH function at altitudes above the surface, it is observed that some cases of this preliminary run suffer from degradation in accuracy in the form of unusually large localized biases near the north pole. Ensuing experiments showed that increasing γ to ~2 (using an archived Thomson solution spacing with 2n-1 measurements) and reducing the dipole and mass weights (to w_s=0, w_{di}= 10) improved performance at these higher altitudes (noting that surface residuals are not highly sensitive to these changes in γ or w). Accordingly, for the final solution files, the cases of α=1.25, 1.5, and 1.75 are rerun using the simplified form of the algorithm (CASE B in Algorithm 2 step 7) using the previously optimized b_* values from Figure 8. Finally, four of the resulting solutions are chosen for further characterization: PMC11n240, PMC33n1920, PMC71n7680 and PMC156n32720 in order of increasing fidelity (where the nomenclature of the labels indicate the size of the fitting field and the number of mascons). The particular solutions are chosen such that the RMS errors of the least squares solution are at least one order of magnitude smaller than the associated accumulated error levels of the SH fitting function from Figure 1. Table 1 shows this conservatism ratio as well as other details on each of the chosen solutions.

Table 1: Archived point mascon models

Model descriptor	Size of SH fitting field, d×d	Number of mascons, n	α	Mascon bury distance, b_* (R_E)	Accumulated normalized error of SH field	RMS of residuals for mascon fit	Conservatism ratio, (SH err)/(PMC resids)
PMC11n240	11×11	240	1.75	0.8697	1.4×10^{-10}	5.5×10^{-12}	25
PMC33n1920	33×33	1,920	1.75	0.4469	2.3×10^{-10}	1.5×10^{-11}	15
PMC71n7680	71×71	7,680	1.5	0.2593	1.1×10^{-9}	4.2×10^{-11}	26
PMC156n32720	156×156	30,720	1.25	0.1357	2.4×10^{-8}	5.4×10^{-10}	44

POINT MASCON MODEL PERFORMANCE

In this section, the PMC solutions from Table 1 are evaluated for performance in terms of runtime and accuracy in comparison to computations using their associated SH fitting functions. The speed test is performed using random evaluations where the SH model is evaluated using the singular-free and Legendre methods described previously. The PMC model is implemented in two different manners: 1) on the same CPU as the SH model and 2) on a GPU to investigate potential for parallelism. Two accuracy tests are performed: The first is a global grid evaluation at various heights comparing the potentials and accelerations using a regular grid of resolution 10d × 5d in the longitude-latitude space (to ensure the capture of the highest frequency SH terms). The second test evaluates the PMC model performance in representative trajectory propagations.

Accuracy: Global grid evaluations

Using the high resolution grid, Figure 9 and Figure 10 show global contour maps of the PMC156n30720 geopotential and difference with the 156×156 SH function evaluated at the surface, respectively. Note that the conventional measure of geoid height[42] is recoverable from Figure 9 by treating the units as non-dimensional and then multiplying by R_E. The spacing of the mascons and the bumpy mascon signature is apparent in Figure 10. Consistent with the RMS of the least squares solution residuals, the RMS of the regular grid evaluation is 5×10^{-10} . The mean is three orders of magnitude smaller indicating a well behaved fit.

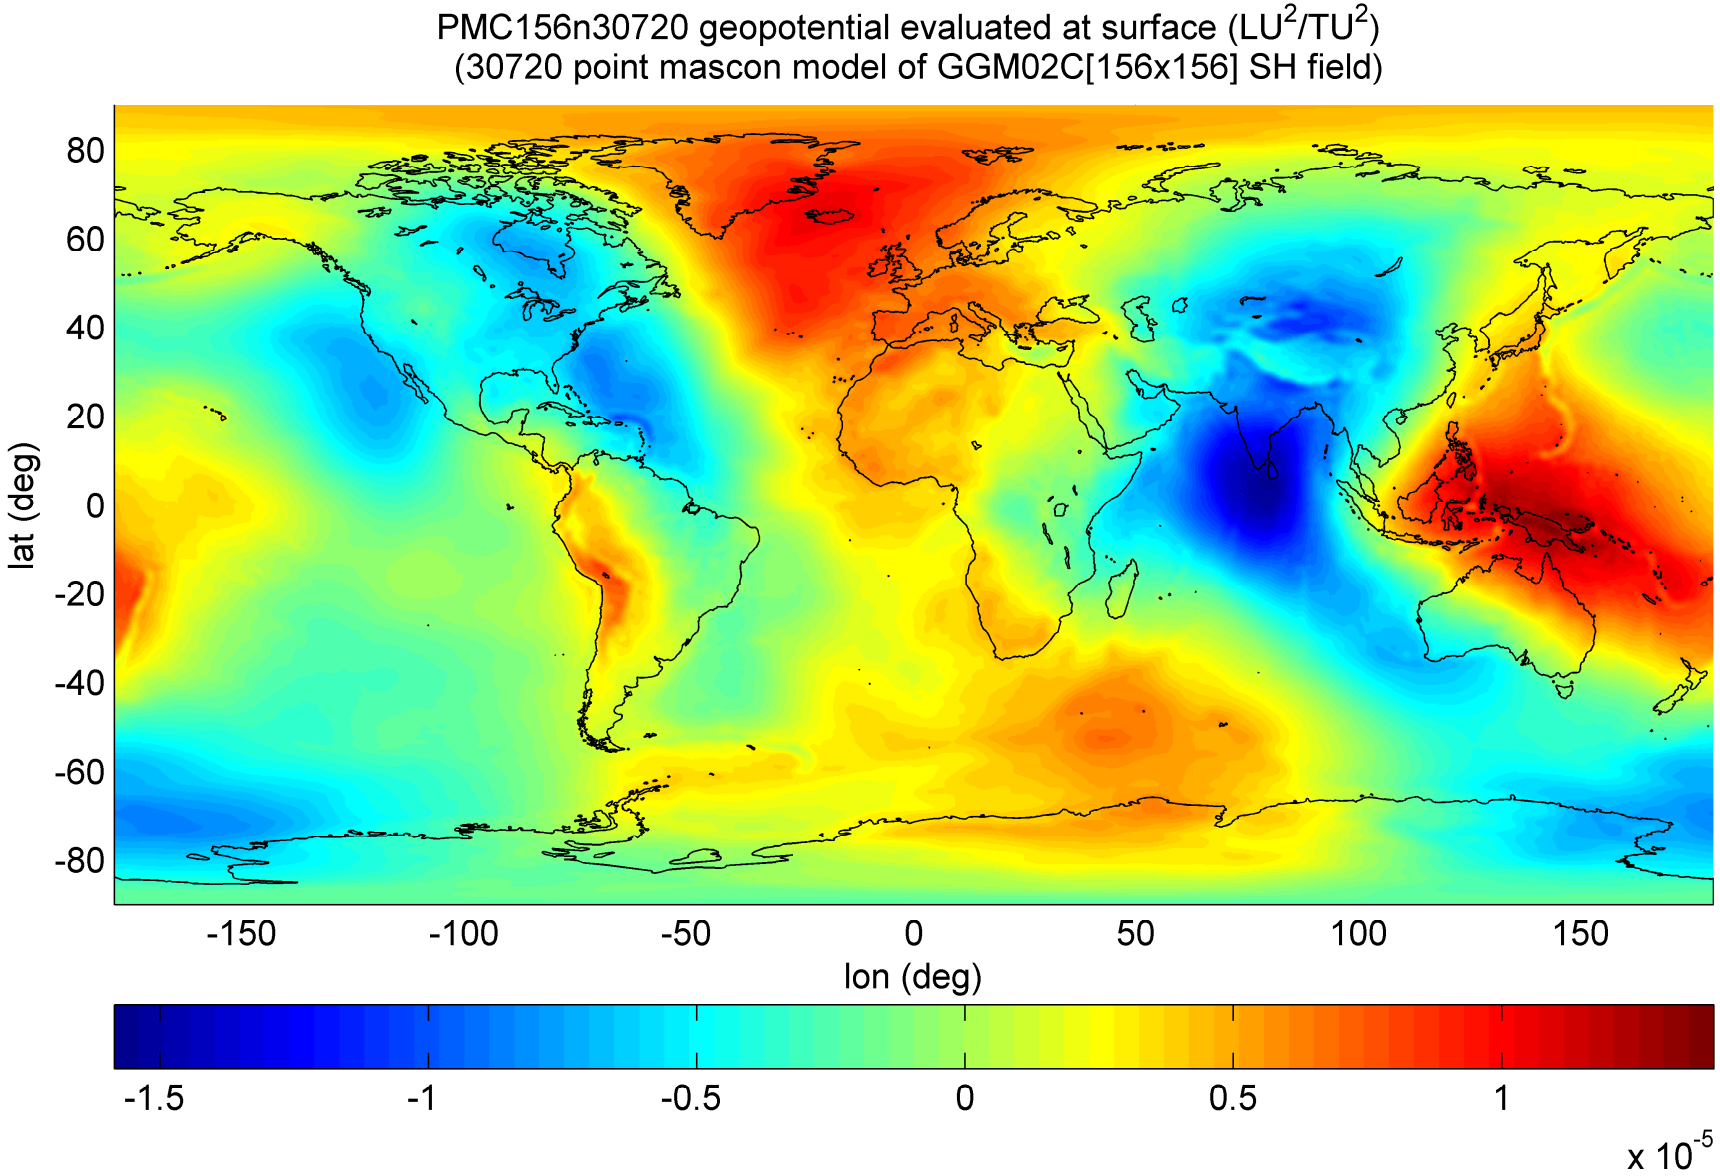

Figure 9: Contour of geopotential (with two-body and J_2 terms removed) evaluated at surface using high fidelity point mascon model: PMC156n30720.

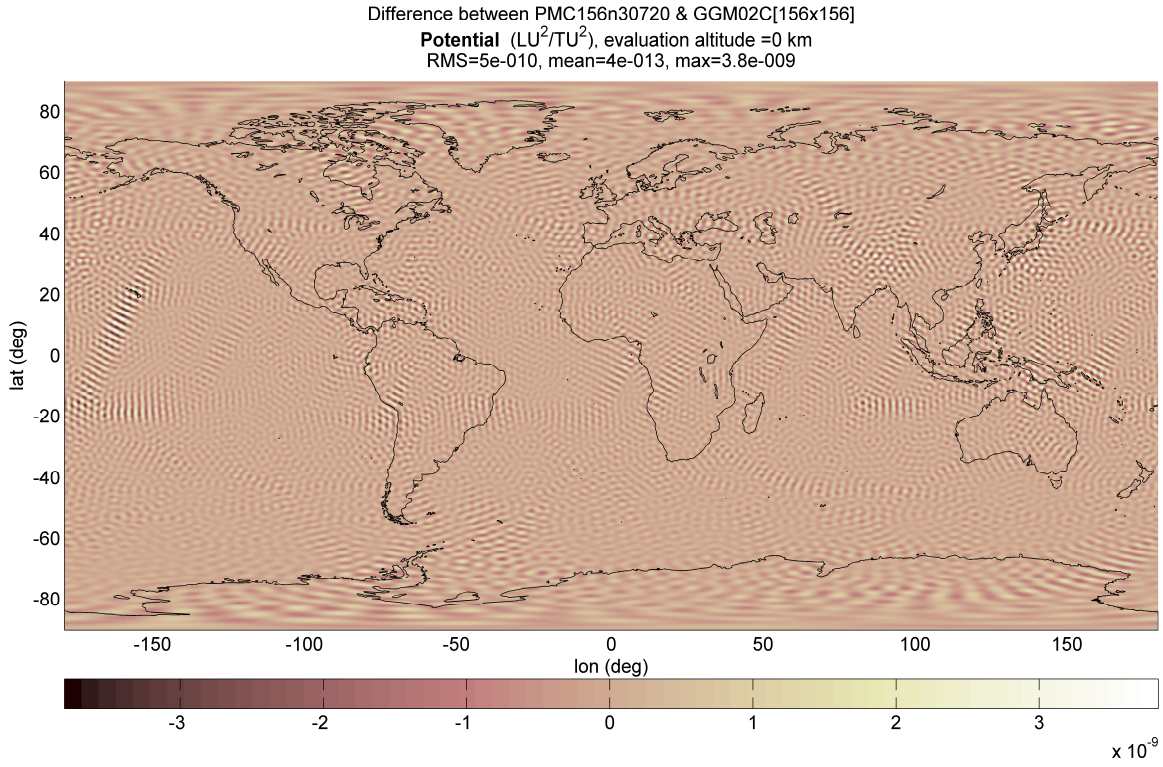

Figure 10: Difference between surface potential evaluated with 30,720 point mascons and 156×156 spherical harmonics fitting function

To assess the performance in the spacecraft operational domain (i.e. not on the surface), each of the PMC models from Table 1 are compared to the SH models at each grid node at altitudes of 50×2^i km for $i=0\rightarrow7$ spanning 50 km to ~1 R_E. Figure 11 shows a sample of the results illustrating statistics for differences in potentials and norms of the accelerations. Despite that the bumps are accentuated in the acceleration plots, the maximum difference at 200 km is small at $\sim3\times10^{-9}$ LU/TU^2 or $\sim 3\times10^{-5}$ mm/s^2. Furthermore, regions of higher residuals are limited in magnitude and reasonably distributed across the domain, although some localized biasing is evident. As expected the fit is best in terms of distribution of residuals at the surface where the measurements are taken, while the biasing increases with altitude despite the steady decrease in RMS.

The statistics for all of the runs for each of the PMC models are summarized in Figure 12. In general, the curves do reflect that PMC models have the natural (and physically accurate) property that residuals reduce with increasing evaluation altitude. The PMC model has known singularities at each of the mascons; therefore the higher fidelity models with shallower mascon bury distances suffer more dramatic curvature near the surface (left side) of Figure 12. As expected for a given field the differences in acceleration are higher near the surface due to the differentiation of both models. It is re-emphasized however, that SH coefficients are also originally fit using a potential model, and there is a natural loss in accuracy associated with differentiating the SH model.

Speed: Random evaluations

For each of the PMC models from Table 1, runtimes are compared for potential and acceleration computations using both the PMC and SH methods evaluated at randomized locations. The CPU for this test is a Xeon processor E5520 (2.27Ghz) and the source code is written in Fortran 95 and compiled with Intel Fortran 12.0 using the -O2 optimization settings. The absolute compute times are given in Figure 13.

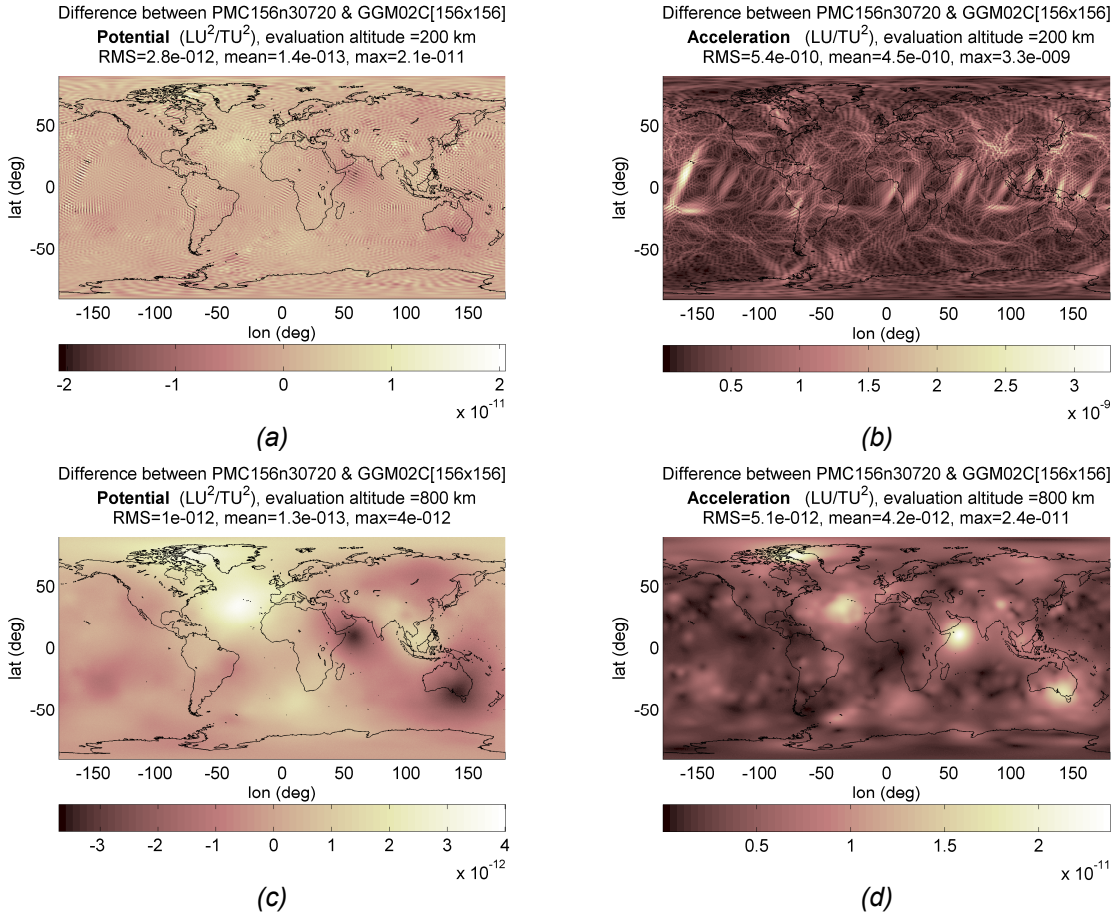

Figure 11: Contour map of difference between high resolution mascon and SH models. (a) and (b) evaluated at 200 km, (c) and (d) evaluated at 800 km; (a) and (c) potentials, (b) and (d) accelerations

The two implementations of SH vary by ~25% in runtime where the Lagrange method is slower than the singular-free approach (contrary to the results of Ref. 1, although it is emphasized that presently no higher order derivatives are computed and other compiler optimization settings or general implementations may lead to different results). The parallel implementation is performed on two different GPU cards: the Tesla C2050 and the GeForce GTX 465. Due to the nature of the simple computations being performed in parallel for the mascon application, the procedure does not benefit much from the high end compute power available in the C2050 card (currently retails for ~$US 2500). To the contrary, the PMC problem leads to a classic GPU application that is limited by available bandwidth connecting the data transfer between the CPU and GPU. The more affordable GeForce GTX 465 card (currently retails for ~$US 250) enjoys similar bandwidth to the C2050 and thus achieves nearly similar speedups.

Figure 13 shows that the serial CPU implementation of the PMC model is on the same order of magnitude in speed compared to the SH implementations, with minor differences depending on the SH method, the size of the field (and, not shown here, the choice of processor, operating system, compiler and compiler settings). The GPU computation gains efficiency with increasing number of mascons. For the highest fidelity (156×156) field, the GPU implementation demonstrates a 13× speedup in comparison to the singular-free SH method, and a ~16× speedup compared to the Lagrange method. The breakeven point where a GPU implementation begins to favor that of a CPU is approximately a ~25×25 field. Future PMC models that fit larger fields such as the full GGM02C 200×200 field (or the extended 360×360 version using the EGM-96 high order terms) should expect higher speedup values. Preliminary experiments

demonstrate an expected doubling of speedup performance (25× to 30×) for the case of a PMC field with 122,880 mascons which could fit a 313×313 SH field using the resolution parameter of α=1.25. Such a solution would require 240 GB of RAM to solve the normal equations and therefore requires a parallel implementation.

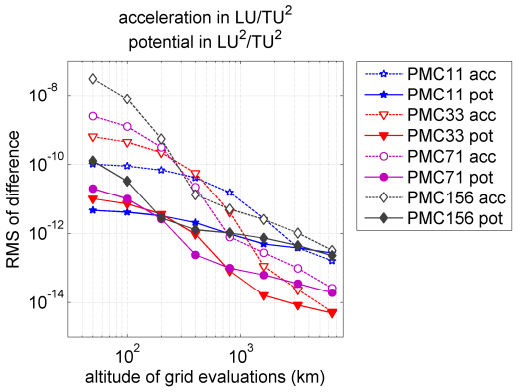

Figure 12: Grid evaluation results on accuracy compared to sing free SH.

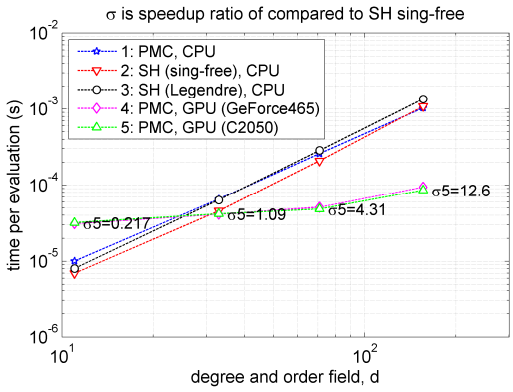

Figure 13: Compute times for randomized calls for PMC and SH models from Table 1. GPUs demonstrate ~13× to ~16× speedups.

Accuracy: Spacecraft trajectory computations

Lastly, five representative spacecraft trajectories (see Table 2) are simulated in order to gauge speed and accuracy performance in a common astrodynamics application. Each of the simulations is repeated for each of the PMC models from Table 1. For all cases the geopotential model is computed in the same 5 methods as were used in the speed tests in Figure 13. The equations of motion are in the body fixed rotating frame and given in Eq. (18) where U is the total potential and Ω is the system rotation and equals 0.00007292115 rad/s.

$$\ddot{x} = 2\Omega\dot{y} + \partial W/\partial x, \quad \ddot{y} = -2\Omega\dot{x} + \partial W/\partial y, \quad \ddot{z} = \partial W/\partial z$$
$$W = \left(\Omega^2/2\right)\left(x^2 + y^2\right) - U \tag{18}$$

This Hamiltonian system admits the well known Jacobi integral of motion, C , that is useful for calibrating accuracy of numeric integration routines:

$$C = 2W - \left(\dot{x}^2 + \dot{y}^2 + \dot{z}^2\right) \tag{19}$$

The integrations are performed using a variable step-size Runge-Kutta 7(8) algorithm with step size error tolerance set to 1×10^{-14}. This step size tolerance translates to a preservation of ~11 digits of the Jacobi constant over a single rev in the case of Simulation 1 and the d=156 case (for both SH and PMC). It is important to mention here that fixed step integrations using very small step sizes are found to preserve C to as many as ~15 digits for SH implementations, while PMC implementations can only achieve ~13 digits (and only ~12 digits in the case of a naive summation). This discrepancy highlights a main drawback of the PMC models in that they inherently loose a few digits of precision (in the form of continuity in the last digits) due to the summation (as discussed earlier). For most practical purposes, this drawback may not be a serious concern as indicated by the fact that the very tight step size tolerance of 10^{-14} leads to nearly identical C evolution for both PMC and SH methods. Tighter tolerances (< ~10^{-14}) are not recommended for use with PMC models.

Table 2: Spacecraft trajectory simulations. Time of flight for each is 3 days.

Simulation	Description	Perigee altitude (km)	Apogee altitude (km)	Inclination (deg)	number of revs revs
1	very low altitude, circular, near polar	150	150	85	49.4
2	low altitude, circular, near polar	450	450	85	46.2
3	medium altitude, circular, near polar	1,350	1,350	85	38.3
4	high altitude, circular, near polar	4,050	4,050	85	24.5
5	low perigee, highly eccentric, sun-sync	150	$3R_E$	63.5	12.8

The results from these preliminary tests (see Figure 14) show that the PMC models perform similarly to SH in the context of typical Earth orbiting spacecraft trajectories. The maximum difference between the SH and PMC models across all of the simulations and models is found to be 0.6 meters after 3 days of flight time while most of the differences are in the cm or mm level. For the PMC case, both GPU cards lead to identical results and are generally very close to the CPU implementation. Speedups from the GPU implementations increase for the higher fidelity PMC models where the force computation dominates the computation effort. For the 156×156 case, the speedup values from Figure 13 are almost conserved achieving 11.5× and 15.5× speedups (approximately the same for all simulations) when compared to the singular-free and Legendre SH cases respectively.

Finally, in preliminary experiments with a Matlab implementation, the large matrix nature of the mascon approach lends very well performing only ~4 times slower than the Fortran implementation of the PMC156n30720 case of simulation 1, while the two-body only matlab version of the same simulation is ~140 times slower. These fast Matlab results are only achievable using a naive summation and is likely due to Matlab use of compiled and efficient code for large matrix manipulations. Therefore we expect the mascon models to have broad appeal in the Matlab community. Future work is needed to fully examine the mascon performance in Matlab.

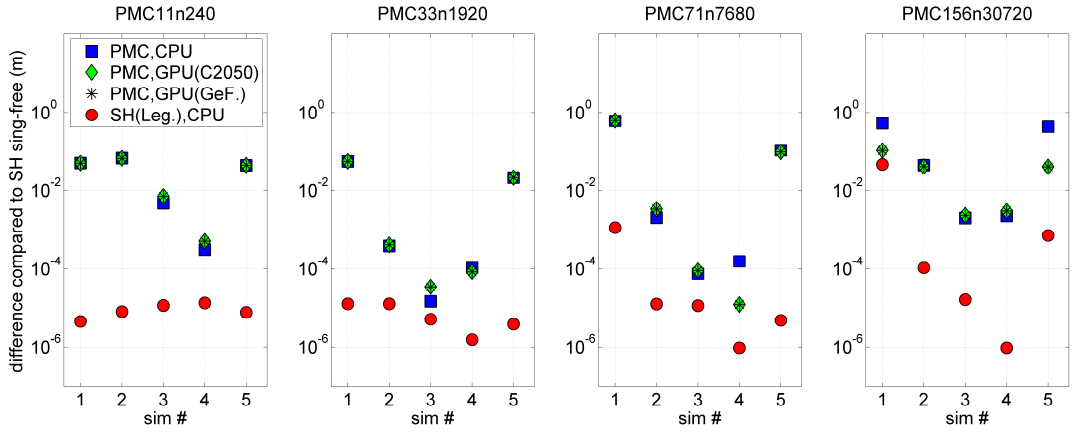

Figure 14: Comparisons of position differences after 3 day orbit simulations. Points are compared to singular-free implementation of associated SH field.

LOCAL WEIGHTED INTERPOLATION MODEL APPROACH

Localized representation of the gravity field follows naturally from the fact that there are uneven gravity undulations over the whole planet's surface. Hence, a localized expression should be more compact and computationally more efficient. In order to separate the global effects from the local undulations we propose the conventional localized gravity model as given by Eq. (20).

$$U = U_{J_2} + U_{local} \tag{20}$$

The J_2 in Eq. (20) (which signifies oblateness) is ~2-3 orders of magnitude more significant than the terms following later in the series. Calculation of the J_2 term is given by (21). Splitting up of J_2 and other higher order terms buys us 4-5 digits of accuracy in the potential at an almost negligible computational cost.

$$U_{J_2} = \frac{Gm}{r}[J_2(\frac{R_e}{r})^2 P_2(sin(\phi)] \tag{21}$$

The U_{local} term in Eq. (20) represents the potential term calculated via a localized polynomial whose coefficients are generated from a least square fit of the gravity potential values over a local region. Following the weight function approach (discussed later) the local potential is given by:

$$U_{local} = \sum_{i=0}^{1}\sum_{j=0}^{1}\sum_{k=0}^{1} w_{ijk} U_{ijk}^{cell} \tag{22}$$

Here, w_{ijk} are Hermite weight functions normalized in (0, 1) space and U_{ijk}^{cell} represents localized per-cell potential normalized in (-1, 1) space. Unlike the quadrature performed by the Junkins approach, here we obtain the term U_{ijk}^{cell} via a least square fit of the local potential data.

Furthermore, the derivatives (accelerations and higher order derivatives) are calculated by taking the gradient of the local approximation function, so they are exact. This is an attractive feature of the local finite gravity model in which both continuous and exact derivatives are naturally computed. Before going into further details of these terms we give a brief overview of the global discretization scheme used for generating the finite elements.

Discretization scheme

Figure 15 shows the global discretization scheme used for creating the 3 dimensional finite element grid. The weight function technique requires a uniform lat-lon grid over the complete discretized domain. Hence, each finite element cell has the same breadth and width for any given altitude (*h*). The distance between two shells determines the height of each finite element cell between them. Optimal shell spacing is determined so as to minimize the number of coefficients required to fit the global field. This optimal shell spacing also helps in minimizing the memory storage requirements at runtime.

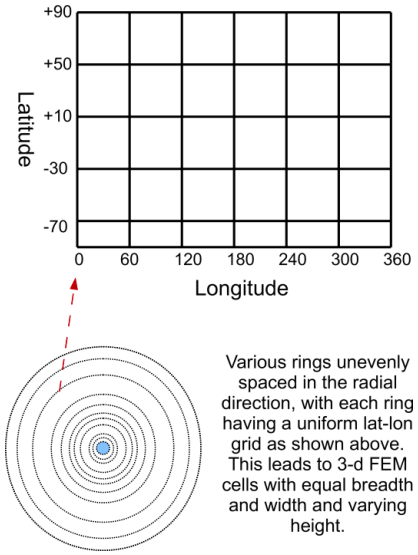

Figure 15: Global Discretization Scheme

The scheme comprises of a two level grid structure; a global grid consisting of vertices from each

cell and a local approximation grid for each cell. The local grid is normalized between -1 to 1. This choice has a big benefit which will be discussed later.

Continuity

Continuity in potential and higher order derivatives is a desirable property while performing trajectory determination / optimization. Many classes of numerical methods fail if this continuity criterion is not met stringently. Keeping this in mind, we adopt the weight function method developed by Junkins which leads to a continuous approximation across all finite element cells irrespective of their local fitting polynomial. This further helps us in selecting the minimum order polynomial required for each cell relative to a user prescribed error tolerance.

Figure 16 shows the two dimensional weight function approach, which is implemented by the two level grid formulation discussed earlier. For the two dimensional case four local grids give rise to an "Overlapped region" as shown. The potential and higher order derivatives are then calculated as a weighted sum of the four adjacent local approximations on this overlapped region. Positive Hermite weight functions are used satisfying the constraint (for the 2-dimensional case) given as:

$$\sum_{i=0}^{1}\sum_{j=0}^{1} w_{ij}(x1, x2) = 1 \tag{23}$$

Here, x1, x2 are the localized coordinates of the evaluation point. We note that the overlapped region is normalized from 0 to 1 while the adjacent local approximation cells are normalized between -1 to 1.

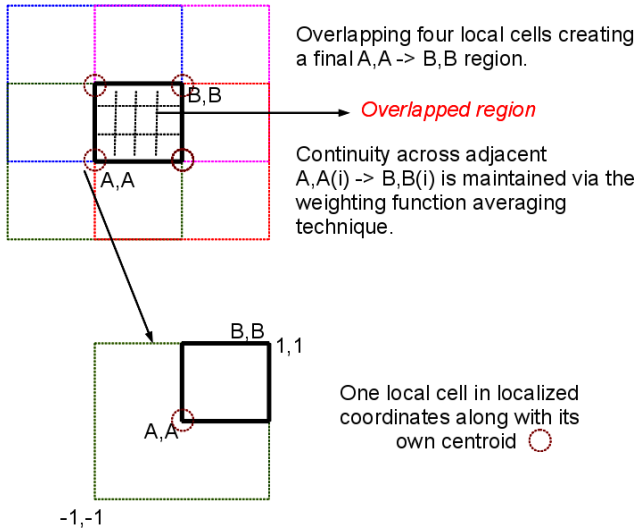

Figure 16: Global local approximation

Each local cell has a polynomial approximation associated with it. Continuity is achieved by averaging out the local approximation with Hermite weighting functions. The fact that continuity across various local approximation cells is independent of the local approximation function provides robustness and a smooth error distribution over the whole solution space. The main downside for using this approach is the 8:1 weight function averaging (for the 3D case) computed during runtime for each potential and higher order derivative calculation.

Radial grid spacing

Even though we require a regular lat-lon grid, we are free to choose grid spacing in the radial direction. Having closely packed shells leads to smaller radial step size. This further, leads to fewer coefficients per cell. While on the other hand choosing a very large radial step size will increase the

number of coefficients per cell making the runtime potential calculation slower. For the present study we heuristically determine the optimal cell spacing using a log based multiplier. We find that for a gravity model extending all the way to the moon, we need approximately 52 unevenly spaced shells.

Localized Gravity Field Approximation

The local approximation within each cell is obtained from a least square fit of the measured data within each cell. We can divide the process of coefficient generation in three basic steps as follows:

1. Analytic coefficient generation
2. Scalable degree selection
3. Adaptive coefficient selection

The general least square process used for generating coefficients is given by:.

$$\vec{c} = \mathbf{B}\vec{u} \tag{24}$$

Here, vector *c* is a column vector with coefficient estimates, vector *u* represents the measurement vector and **B** denotes the least square inverse (LSI) matrix which is further given by:

$$\mathbf{B} = (\mathbf{H}^T\mathbf{W}\mathbf{H})^{-1}\mathbf{H}^T\mathbf{W} \tag{25}$$

Here, *m* represents the number of measurements and *n* represents the total number of coefficients. The matrix **H** (size *m* by *n*) represents the function evaluation matrix generated for various degrees of polynomials along each direction.

Coefficient generation

We remind ourselves that each localized cell is defined in normalized -1 to 1 space. If the number of observations (*m*) and their respective positions within each cell do not change, we then need to compute the LSI matrix only once (analytically) and use it over the whole grid. Apart from computational benefits, this trick also allows us to compute the LSI matrix using a symbolic manipulator (like Maple) without resorting to numeric procedures (and their associated errors). We analytically compute the LSI matrix for all possible combinations of polynomials along each direction within each cell and store the results in separate files which are read at runtime.

For this study we have created a Maple worksheet capable of creating LSI matrices corresponding to all possible combinations of polynomials within each cell. The maximum polynomial degree (p_n) in each direction is fixed at 6. We note that for a 3D polynomial with 6 degrees in each direction, the resulting LSI is of dimension 7^3x7^3. Inverting such a matrix analytically is a computationally expensive. The worksheet takes p_n and *m* as inputs and writes out LSI matrices using Maple's code generation option. These matrices are read once during the coefficient generation process and reused later when needed. The number of observations per cell & per direction is decided based on the local cell size. We must ensure that number of observations in each direction is at minimum p_n +1. Analytic LSI matrices ensure an accurate and consistent coefficient generation process. To obtain the coefficients at runtime only a matrix multiply is needed for each local cell, where otherwise we would require the solution to a large linear system of equations. This analytic inversion process is an important contribution allowing for an adaptive error control scheme.

Scalable degree selection

Once the LSI matrix has been generated all that is left is to generate the measurement vector *u* for each cell to obtain the coefficients. One way to do this is directly to use the spherical harmonics code for a fixed degree (200) and generate the measurement vector. Unfortunately, we find that it takes a long time (on the order of weeks) to generate all the coefficients for a global gravity model extended up to $60*R_e$ (all the way to the moon). Almost all the computational burden is due to the measurement vector obtained using spherical harmonics gravity model.

To overcome this problem we propose a scalable spherical harmonics degree selection method. The method is based on the pruning the contribution of each term in the series evaluation of the potential. We calculate the contribution of each term in the spherical harmonics series as we move from the surface up to where the absolute magnitude of each term becomes less than 1E-15. The *maximum* degree with absolute magnitude greater than 1E-15 for a given a radial distance is noted. We then create a "Degree vs. Radial distance" plot as shown in Figure 17.

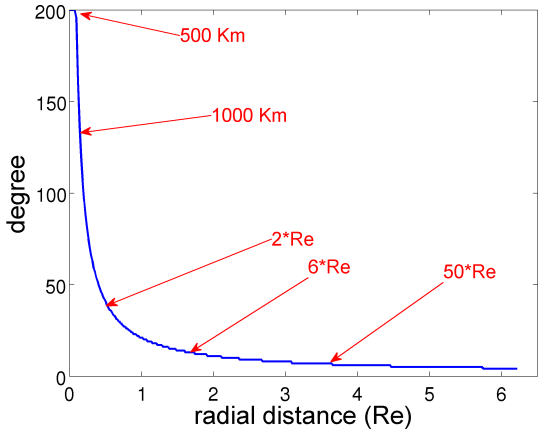

Figure 17: Degree selection curve

After performing a non-linear least square rational polynomial fit on the curve we obtain an approximate equation (Eq 8) for this plot. Equation (26) takes in radial distance (*r*) as input and returns the degree value which we can use to compute the spherical harmonics gravity field at that radial distance.

$$d = \text{int}\left(\frac{20.105042266558979}{\log(r) + 0.007876575109652}\right) \tag{26}$$

where *r* is the radial distance in normalized units from the surface of the earth. It is evident from Figure 17 that as we reach the radial distance of $1*R_e$ from the surface of the earth most of the high order spherical harmonics terms are not significant even in double precision. As we move radially outwards the measurement computation time decreases very rapidly thereby significantly speeding up coefficient generation. This degree selection technique can also be used in other calculations where high altitude, high fidelity trajectory integrations are required.

Adaptive coefficient selection

The weight function approach gives us the freedom to choose any order polynomial for the local approximation function within each cell. We can utilize this fact to make our least square polynomial fitting process adaptive by automatically choosing a different degree polynomial for each direction within each cell (see Figure 18). The polynomial having the minimum number of coefficients while still satisfying a runtime error tolerance is selected and its coefficients are stored.

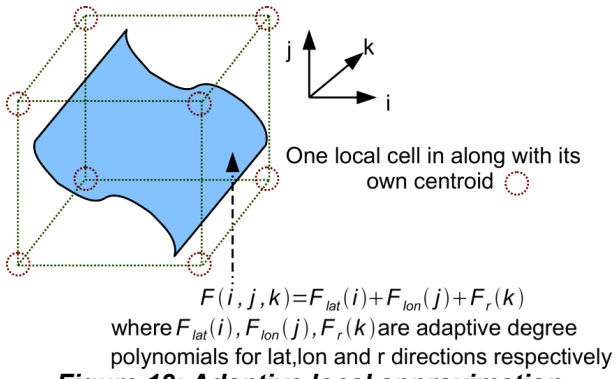

Figure 18: Adaptive local approximation

The error tolerance is determined at runtime depending on the radial distance of the current cell from the surface of the earth. For this purpose the "*actual"* error of the 200 degree spherical harmonics field as given by the GGMO2C gravity model at the surface is computed. At runtime, this error is then mapped to the radial distance where the lowest vertex of the current cell is located. Hence, we obtain an error scaling graph specifying the cutoff tolerances as a function of radial distance (Figure 19).

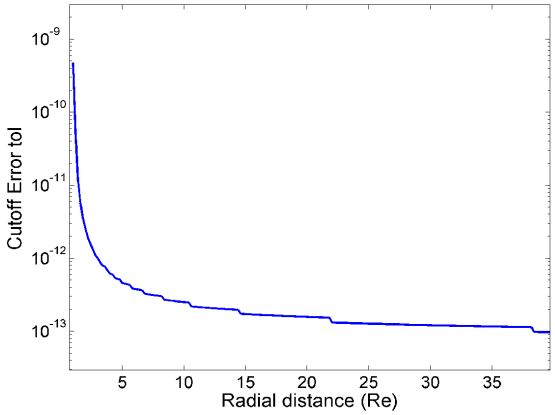

Figure 19: Error scaling graph

Once we have decided upon the degree (from the previous step) to be used at a given radial distance, we calculate the corresponding max error from this error scaling graph. We divide this error by a user defined number (currently 5). This final error value acts as our cutoff tolerance. This adaptive coefficient selection approach minimizes computation cost at runtime and decreases memory requirements for coefficient storage. And importantly, the error profile assumes a realistic form that falls of naturally with increasing radius.

After selecting the polynomials for a cell the final local approximation is given by Eq. (27) or (28).

$$U_{ijk}^{cell} = \sum_{i=0}^{N1max} \sum_{j=0}^{N2max} \sum_{k=0}^{N3max} c_{ijk} f_{ijk}(X1, X2, X3) \tag{27}$$

$$U_{ijk}^{cell} = \sum_{n1=0}^{N1max} \sum_{i=0}^{n1} \sum_{j=0}^{n1-i} c_{ijk} f_{ijk}(X1, X2, X3) \qquad \mathrm{k = n\text{-}i\text{-}j} \tag{28}$$

where c_{ijk} represents the coefficients corresponding to the polynomial f_{ijk}(*X1,X2,X3*), and *X1, X2, X3* represent the coordinates of the evolution point expressed in -1 to 1 normalized coordinates. Equation (28) represents a reduced set case of coefficients which was proposed by Junkins and has been

incorporated in our polynomial selection routines. Depending on the number of coefficients required to meet the error tolerance, the code automatically decides whether to use Junkins reduced polynomial set or not.

Parallel coefficient generation

In spite of various algorithmic optimizations as stated in the previous sections, the coefficient generation process for a complete global model would require approx 20 to 30 days on a single modern CPU. As we only need to generate the coefficients once, we can parallelize this process using Message Passing Interface or MPI. Our target is to bring the coefficient generation time to within a day on supercomputing cluster with ~100 CPUs.

To achieve this, we evenly split the whole solution domain amongst all the processors. This is achieved by discretizing each processor on the basis of latitude band; hence each processor works on a thin vertical slice of the whole domain extending all the way up to the moon. This domain decomposition is ideally suited for explicit parallelism as there is little communication among processors. The parallel version of the code is implemented in FORTRAN 2003 and can be compiled using either Intel MPI compiler of the OPEN MPI compiler. Now we briefly overview the runtime routines needed for running the code, followed by the results.

RUNTIME GRAVITY POTENTIAL AND ACCELERATION COMPUTATION

There are only 3 runtime routines which the user needs to call for using the new finite element based gravity model. These routines are listed below here:

Table 3: Runtime routines

Routine name	Description
ALFem_init()	Initializes the gravity model and reads in the coefficients from a user supplied database
get_local_u	Calculates the gravity potential
get_local_ud	Calculates the gravity potential plus acceleration both in spherical and Cartesian coordinates

The accelerations are computed by taking the gradient of the potential with respect to radial, latitude and longitude directions. The acceleration vector thus obtained is in spherical coordinates and needs to be transformed back to Cartesian coordinates using coordinate transformations. A FORTRAN optimized code has been generated for calculating these transformations. The derivates are calculated in an efficient manner with focus on maximum reuse of already computed data.

RESULTS

For performance comparison purposes we use the Legendre spherical harmonics implementation. For the local finite element based method, the radial direction is made adaptive and controlled by a user input step array. The target abs error tolerance is scaled for an equivalent 200 by 200 gravity model error, as discussed in previous sections. The system configuration used for benchmarking is an Intel Xeon processor E5520 (2.27Ghz) with 16 GB of RAM. The benchmarking is done in two ways: 1) a random accuracy test on various radial bands of the complete finite element domain; and 2) a speed test for both potential only and potential plus acceleration computation. The code is compiled and linked on the Intel FORTRAN Compiler version 12.0 on Linux and has been subjected to ``-O2'' compiler optimizations.

Generation of Coefficients

Table 4 lists all the user options which were used for the generation the full localized finite element based gravity model.

Table 4: User options

Range (latitude)	0.05 to 179.95 degree

Range (longitude)	0 to 360 degree
Range (radial)	1 to 60*R_e
Cell size (lat , lon)	0.6 degree
Number of lat-lon shells	52
Maximum polynomial degree along each direction	6
Measurements along each direction	7
Number of CPUs used for coefficient generation	148
Sample size for benchmarking	10000
Total number of finite elements	9173632

The latitude and longitude size of each cell directly affects the memory requirement and runtime performance of the method. After, a fair bit of experience with the problem and using engineering judgment a value of 0.6 degrees is used for fitting a 200 degree spherical harmonics gravity field. A third order continuous weight function was selected for this study, which provides a continuous Hessian if calculated. The weight function is given by Eq. (29).

$$w = \prod_{i=1}^{3} X_i^4(35 - 84X_i + 70X_i^2 - 20X_i^3) \tag{29}$$

The coefficient generation took 11 hours on the Garuda cluster using 150 processors at the Georgia Institute of Technology. The memory requirement for storing all the coefficients at runtime is 1.3 GB.

Performance

Performance benchmarking was done both on the potential and potential plus acceleration routines. Table 5 gives a summary of the error and speed performance test. As expected, the root mean square (RMS) values for both the potential and acceleration errors are under the required tolerances. The max error for acceleration was found to be up to two orders of magnitude more than max error in potential. This behavior was noted mostly when the evaluation point was located near the poles. Future, studies will investigate this matter further.

Table 5: Results Summary

Region (altitude)	Potential Error RMS/MAX	Accel Error RMS/MAX	Potential + Accel speedup	Potential Only speedup
100 to 500 km	5E-11 / 2.7e-10	1e-9 / 7e-8	407	454
500 to 2000 km	3e-11 / 2.5E-10	7E-10 / 6E-8	410	460
2000 km to 6 Re	1E-11 / 1E-10	5E-10 / 2E-8	440	492
6 to 10 Re	2e-11 / 7e-11	1e-10 / 3e-09	540	590
10 to 20 Re	1E-11 / 5E-11	9E-11 / 1E-09	560	620
20 to 35 Re	4E-11 / 2E-11	7E-11 / 2E-10	570	640
35 to 50 Re	3e-13 / 3e-12	2E-12 / 5E-10	585	672

Figure 20 gives the speedup values for both the potential only and potential plus acceleration routines over a 200 degree spherical harmonics evaluation. We can see from Table 5 and Figure 20 that we are able to achieve speedup values of over 400 times over the spherical harmonics code even for as low as 100 km altitude. The maximum speedup obtained for the potential only routine is almost 700 times and that of potential plus acceleration routine is around 590 times. *Note that scaling these results to comparisons with the singularity-free implementation from the PMC section would lower the speedup*

values by ~25%.

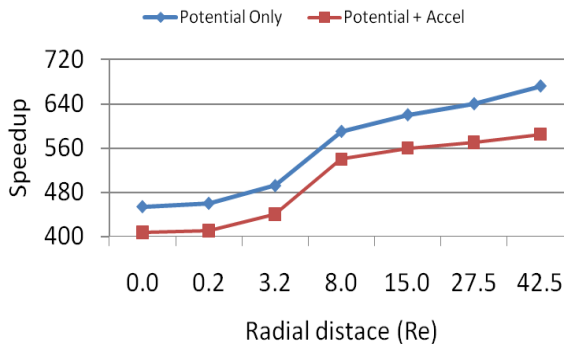

Figure 20: Speedup over Spherical Harmonics

CONCLUSIONS

In summary, we have developed two new methods with working prototype software to achieve multiple order of magnitude runtime improvements in geopotential computation. The benefit of the point mascon formulation is its extreme simplicity, very low memory requirement, and potential for parallelism. The benefit of the interpolation approach is its extreme speed, noting that our results are up to an order of magnitude faster than competing interpolation techniques. Unlike competing approaches, we also emphasize that both the mascon and interpolation approach preserve exactness (i.e. acceleration is exact gradient of force etc.) and continuity to at least 3 orders.

In the prototype global point mascon models we find representative speedups (using a single Graphics Processing Unit or GPU) of ~10 times for high resolution gravity fields (degree and order 156x156) while requiring a memory footprint on the same order as that of spherical harmonics (<1 MB). A remarkable result is that without the GPU, a CPU only implementation is of the same order of magnitude (slightly faster) in speed as the spherical harmonics formulation. The mascon model could therefore find application in a broad community.

In the interpolation approach we have presented an adaptive, local finite element based gravity model which trades memory for speed. The finite element gravity model preserves continuity across various cells using a weight function averaging technique. We have also worked towards a new way of understanding the spherical harmonics problem in which the degree of the field is dependent on the distance from the surface of the earth and its degree error profile generated from the state of the art GGMO2C gravity model. This in turn benefits the coefficient generation process for the global finite element gravity model. A novel adaptive polynomial degree selection ensures a optimal degree of polynomial is selected for the adaptive error tolerance. A parallel coefficient generation algorithm has also been implemented making global model coefficient computation practical. Our results show that the proposed method is multiple orders of magnitude faster than tuned spherical harmonics codes, while still preserving accuracy. Furthermore, we obtain exact derivatives by taking the gradient of the local potential approximations directly. This in turn can benefit high fidelity optimization problems where exact derivatives are important.

In addition to minor technical improvements, in future works we intend to formalize the procedures and software into usable products for the sponsor and community at large. We would also focus on decreasing the memory storage requirements by investigating alternate storage options and memory reduction techniques. Higher order exact derivative calculations will also be implemented. In addition, significant work remains for both methods including systematic design trades, testing, further exploration of approach modifications, and coefficient generation improvement. Currently we use only a desktop server for the mascon coefficient generation (future work includes parallel versions) yet we have a preliminary parallel implementation to achieve the global interpolation model.

Both new method and the corresponding tools are capable of providing large scale performance gains for a variety of astrodynamics problems, including but not limited to high fidelity trajectory design and navigation, orbit determination, and optimization problems.

REFERENCES

[1] Casotto, S., Fantino, E., "Evaluation of methods for spherical harmonic synthesis of the gravitational potential and its gradients", *Advances in Space Research*, Vo. 40, 2007, pp. 69–75.

[2] Hoots, F.R., Schumacher, P.W. Jr., Glover, R.A., "History of Analytical Orbit Modeling in the U.S. Space Surveillance System," *Journal of Guidance, Control, and Dyn*amics, Vol. 27, No. 2, 2004, pp. 174-185.

[3] Miller, J. G., "A New Sensor Allocation Algorithm for the Space Surveillance Network," *Military Operations Research*, v. 12, n. 1, pp. 57 – 70, 2007.

[4] Lambeck. *Geophysical Geodesy*, Oxford University Press, 1988.

[5] Kaula, *Theory of Satellite Geodesy*, Dover, New York, 2000.

[6] Tapley, B. D., Schutz, B. E., and Born, G. H., *Statistical Orbit Determination*, Elsevier Academic Press, Burlington, MA, 2004.

[7] Werner, R. and Scheeres, D., "Exterior Gravitation of a Polyhedron Derived and Compared with Harmonic and Mascon Gravitation Representations of Asteroid 4769 Castalia," *Celestial Mechanics and Dynamical Astronomy*, Vol. 65, 1997, pp. 313–344.

[8] Park, R.S., Werner, R.A., Bhaskaran, S. "Estimating Small-Body Gravity Field from Shape Model and Navigation Data", *Journal of Guidance, Control, and Dynamics*, Vo. 33, No. 1, 2010, pp. 212-221.

[9] Casotto, S. and Musotto, S., "Methods for Computing the Potential of an Irregular, Homogeneous, Solid Body and its Gradient," Paper AIAA-2000-4023

[10] Scheeres, D., Khushalani, B., and Werner, R., "Estimating Asteroid Density Distributions from Shape and Gravity Information," *Planetary and Space Science*, Vol. 48, 2000, pp. 965–971.

[11] Siouris, G.M., "Gravity Modeling in Aerospace Applications", *Aerospace Science and Technology (2009)*, doi:10.1016/j.ast.2009.05.005

[12] Junkins, J.L., "Investigation of Finite-Element Representations of the Geopotential", *AIAA Journal*, Vo. 14, No. 6, 1976 pp. 803-808.

[13] Engels, R.C., Junkins, J.L., "Local Representation of the Geopotential by Weighted Orthonormal Polynomials", *Journal of Guidance and Control*, Vo. 3, No. 1, 1980, pp. 55-61.

[14] Beylkin, G., Cramer, R., "Toward multiresolution estimation and efficient representation of gravitational fields", *Celestial Mechanics and Dynamical Astronomy*,2002, Vol. 84 Iss. 1, pp. 87–104,.

[15] Lekien, F., Marsden, J., "Tricubic Interpolation in Three Dimensions," *International Journal for Numerical Methods in Engineering*, Vo. 63, 2005, pp. 455–471.

[16] Colombi, A., Hirani, A. H., Villac, B. F., "Adaptive Gravitational Force Representation for Fast Trajectory Propagation Near Small Bodies," *Journal of Guidance, Control, and Dynamics*, Vo. 31, No. 4, 2008, pp. 1041-1051.

[17] Hujsak, R.S., "Gravity Acceleration Approximation Functions" Paper AAS-96-123.

[18] Tapley, B., et. al.,"GGM02 - "An improved Earth gravity field model from GRACE", *Journal of Geodesy*, Nov 2005, Vol. 79, Is. 8, pp. 467-478.

[19] Pines, S., "Uniform Representation of the Gravitational Potential and its Derivatives," *AIAA Journal*, Vol. 11, Nov. 1973, pp. 1508-1511.

[20] Casotto, S., Fantino, E., "Evaluation of methods for spherical harmonic synthesis of the gravitational potential and its gradients", *Advances in Space Research*, Vo. 40, 2007, pp. 69–75.

[21] Fantino, E., Casotto, S., " Methods of harmonic synthesis for global geopotential models and their first-, second- and third-order gradients", *Journal of Geodesy*, Vo. 83, 2009, pp. 595–619.

[22] Lundberg, J.B., Schutz, B.E., "Recursion Formulas of Legendre Functions for Use with Nonsingular Geopotential Models", *Journal of Guidance, Control, and Dynamics*, Vol. 11, No. 1, 1988, pp. 32-38.

[23] Spier, G.W., "Design and Implementation of Models for the Double Precision Trajectory Program (DPTRAJ)," Technical Memorandum 33-451, NASA Jet Propulsion Laboratory, Pasadena, California, Apr. 1971.

[24] http://www.csr.utexas.edu/grace/gravity/ggm02/GGM02_Notes.pdf

[25] Weightman, J. A., "Gravity, Geodesy and Artificial Satellites. A Unified Analytical Approach," The Use of Artificial Satellites for Geodesy, Volume 2. Proceedings of the International Symposium, Athens, Greece, April 27-May 1, 1965. Edited by George Veis. Athens: National Technical University of Athens, 1967., p.467.

[26] Koch, K-R. "Simple Layer Model of the Geopotential in Satellite Geodesy," *The Use of Artificial Satellites for Geodesy, Geophysics Monograph Series*, Vol. 15., Edited by Henriksen, S.W., Mancini A., Chovitz, B.H., Washington, DC, American Geophysical Union, 1972. pp. 107-109.

[27] Wong, L., G. Buechler, W. Downs, W. Sjogren, P. Muller, and P. Gottlieb (1971), A Surface-Layer Representation of the Lunar Gravitational Field, J. Geophys. Res., 76(26), 6220-6236.
[28] Koch, K. R., and B. U. Witte, "Earth's Gravity Field Represented by a Simple Layer Potential from Doppler tracking of Satellites", Journal of Geophysical Research, Vol. 76, No. 35, 1971, pp. 8471-8479.
[29] Koch, K. R., Morrison, F., "A Simple Layer Model of the Geopotential from a combination of Satellite and Gravity Data", Journal of Geophysical Research, Vol. 75, No. 8, 1970, pp. 1483-1492.
[30] Morrison, F., "Algorithms for Computing the Geopotential using a Simple Density Layer," Journal of Geophysical Research, Vol. 81, No. 26, 1976, pp. 4933-4936.
[31] Melosh, H. J., "Mascons and the moon's orientation", *Earth and Planetary Science Letters*, Vol. 25, No. 3, Apr. 1975, pp. 322-326.
[32] Palguta, J., Anderson, J.D., Schubert, G., Moore, W.B. "Mass Anomalies on Ganymede," *Icarus*, Vo. 180, no. 2, 2006, pp. 428-441.
[33] Negi, J., "Inversion of regional gravity anomalies and main features of the deep crustal geology of India", *Tectonophysics*, Vo. 165, No. 1-4, 1989, pp. 155-158.
[34] Bowin, C., "Mass Anomalies and the Structure of the Earth", *Physics and Chemistry of the Earth Part A*, Vo. 25, No. 4, 2000, pp. 343-353, DOI: 10.1016/S1464-1895(00)00056-9.
[35] Antunes, C., Pail, R., Catalão, J., "Point Mass Method Applied to the Regional Gravimetric Determination of the Geoid, " *Studia Geophysica et Geodaetica*, Vol. 47, No. 3, pp. 495-509, DOI: 10.1023/A:1024836032617
[36] Baur, O., "Tailored least-squares solvers implementation for high-performance gravity field research," *Computers & Geosciences*, Vo. 35, No. 3, March 2009, pp. 548-556
[37] Cicci, D., "Improving gravity field determination in ill-conditioned inverse problems", *Computers & Geosciences*, Vo. 18, No 5, pp. 509-516.
[38] Ashby, N., Brittin, W. E., "Thomson's problem," *American Journal of Physics*, Vol. 54, No. 9, 1986, pp. 776-777.
[39] http://thomson.phy.syr.edu/thomsonapplet.htm
[40] Saff, E. B., Kuijlaars, A.B.J., "Distributing many points on a sphere," *The Mathematical Intelligencer*, Vol. 19, No. 1, pp. 5-11, DOI: 10.1007/BF03024331.
[41] Higham, N.J., "The Accuracy of Floating Point Summation, " *SIAM Journal of Scientific Computing*, Vol 14, No. 4, pp. 783-799.
[42] Lambeck, *Geophysical Geodesy*, Oxford University Press, Oxford, 1988, pg. 20.