

Global MHD simulations of the solar convective zone using a volleyball mesh decomposition. I. Pilot

A. Popovas^{1,2}, Å. Nordlund^{1,3} and M. Szydlarski^{1,2}

¹ Rosseland Centre for Solar Physics, University of Oslo, P.O. Box 1029, Blindern, NO-0315 Oslo, Norway
e-mail: andrius.popovas@astro.uio.no

² Institute of Theoretical Astrophysics, University of Oslo, P.O. Box 1029, Blindern, NO-0315 Oslo, Norway

³ Niels Bohr Institute, University of Copenhagen, Øster Voldgade 5-7, DK-1350 Copenhagen, Denmark

Received; accepted

ABSTRACT

Context. Solar modelling has long been split into “internal” and “surface” modelling, because of the lack of tools to connect the very different scales in space and time, as well as the widely different environments and dominating physical effects involved. Significant efforts have recently been put into resolving this disconnect.

Aims. We address the outstanding bottlenecks in connecting internal convection zone and dynamo simulations to the surface of the Sun, and conduct a proof-of-concept high resolution global simulation of the convection zone of the Sun, using the task-based DISPATCH code framework.

Methods. We present a new ‘volleyball’ mesh decomposition, which has Cartesian patches tessellated on a sphere with no singularities. We use our new entropy based HLLS approximate Riemann solver to model magneto-hydrodynamics in a global simulation, ranging between 0.655–0.995 R_{\odot} , with an initial ambient magnetic field set to 0.1 Gauss.

Results. The simulations develop convective motions with complex, turbulent structures. Small-scale dynamo action twists the ambient magnetic field and locally amplifies magnetic field magnitudes by more than two orders of magnitude within the initial run-time.

Key words. numerical—magnetohydrodynamics (MHD)—Sun: general, interior, magnetic fields—convection

1. Introduction

Throughout most of the history of trying to understand the Sun, solar research has been split into ‘internal’ and ‘surface’ solar physics topics, and an integrated global view of the Sun’s complex plasma dynamics has been lacking. For instance, surface models, however complex they have become (e.g. Carlsson, Mats et al. (2016); Chen et al. (2022); Przybylski et al. (2022)), have to make certain initial and boundary assumptions about the magnetic field. These assumptions are usually specified as experiment input parameters, and to a large extent reflect a need to take larger scales (in space and time) of convection, differential rotation, and magnetic fields into account. Moreover, simulations are usually plane-parallel, which severely restricts the horizontal extent that can be simulated without sacrificing realism because of lacking curvature and large scale differential rotation. In addition, certain features in the atmosphere of the Sun can span large distances, e.g. coronal loops ~ 100 Mm (Peter, H. et al. 2013; Reale 2010), prominences ~ 10 – 100 Mm (Parenti 2014), and supergranulation ~ 30 – 50 Mm (Rieutord & Rincon 2010). Such features are thus difficult to study in isolation. Likewise, dynamo simulations, seeking to reveal the origin of the magnetic field and its cyclic behaviour, usually span from the bottom of the convection zone up to 10–30 Mm below the surface, which gives significant benefits (large time-step, use of incompressible gas approximations, etc.) but at the cost of

having to neglect the surface physics and the existence of sunspots.

To resolve this disconnect many bottlenecks must be addressed: Highly disparate spatial and temporal scales, complex microphysics, the interplay between local and global effects, very wide dynamic range, and required resolution in both space and time to resolve various features at different depths of the atmosphere. Tremendous efforts have been made recently to resolve such disconnects for the Sun and other stars. Hotta & Kusano (2021); Hotta et al. (2022) used global simulations of the solar convection zone to study solar differential rotation. They used spherical geometry and a Yin-Yang grid (Kageyama & Sato 2004) with the sphere spanning 0.71–0.96 R_{\odot} . Guerrero et al. (2022) ran global anelastic convection simulations between 0.6–0.96 R_{\odot} and found that obtaining the correct distribution of angular momentum is not a mere issue of numerical resolution; magnetism and near-surface shear layer may be necessary to simulate the solar interior accurately (Guerrero et al. 2022).

In a larger context, the Sun is only one example of a rather common type of stars, and similar efforts have been made for other stars; e.g. Käpylä (2021) ran 288³ and 576³ grid points star-in-a-box simulations of fully convective stars, showing that the geometry of a star is not the defining criterion for generating differential rotation and large-scale magnetism. These works are just a few examples, using rather different mesh decomposition, physics involved, spatial and time resolution, etc.. What they have in common is an effort to connect the deep interiors with

the atmospheres of stars. Such simulations are extremely computationally expensive, often forcing significant sacrifices. Incompressible or anelastic models are not applicable close to the photosphere of stars, Cartesian boxes are not very suitable for maintaining a spherical hydrostatic equilibrium, and spherical coordinates have singularities at the poles, to mention just a few.

To address outstanding bottlenecks and connect internal convection zone and dynamo simulations to the surface of the Sun, we employ the task-based DISPATCH code framework (Nordlund et al. 2018). This framework has several key features that are critical for this work:

- it uses local time-steps for each patch;
- it is solver agnostic, allowing modular selection between several standard and non-standard methods;
- it uses hybrid MPI/OpenMP parallelism, with only nearest-neighbour MPI communications, which gives theoretically unlimited scaling;
- it can work with curvilinear meshes, or curvilinear arrangements of locally Cartesian meshes
- it can employ static or adaptive mesh refinement;
- it has flexible mechanisms to handle additional physics, such as non-ideal MHD, heat conduction, and radiative energy transfer.

We will touch on each of these features in more detail throughout the following sections of the paper. In sections 2 and 4 we describe the basic equations and set-up. We use a ‘volleyball’ mesh decomposition, which is described in section 3. Simulations are started from a modified hydrostatic equilibrium. Since the setup covers the whole surface of the Sun, we only have to use radial boundary conditions, which, together with the initial conditions, are described in 5. The relaxation of the simulations is described together with the initial results in section 6, followed by discussions and concluding remarks in sections 7 and 8, respectively.

2. Basic equations and set-up

The DISPATCH framework can use several different solvers for magnetohydrodynamics (MHD); an entropy based STAGGER solver, used in Popovas et al. (2018, 2019), an internal energy based Stagger solver taken straight from the Bifrost (Gudiksen et al. 2011) code, and several Godunov-type Riemann solvers based on the HLLC and HLLD solvers from the public domain RAMSES code (Teyssier 2002; Teyssier et al. 2006; Fromang et al. 2006), used in (Pan et al. 2018, 2019). Both of these solver groups have their benefits and limitations. The Stagger group is very good under conditions where magnetic fields are dominating the energy budget, as in the solar corona. However, they are very diffusive and perform poorly under low Mach number condition. The Riemann solver group is much better under conditions where detailed turbulent structures are needed, but are flawed with respect to accurately evolving thermal energy when kinetic or magnetic energies dominate the energy budget—this can give rise to negative pressures in the solar corona, and can cause inaccurate convective flow velocities. This is a well known problem with Godunov and Roe type Riemann solvers based on total energy. Many different *ad hoc* solutions have been suggested (e.g. Ismail & Roe 2009; Winters & Gassner 2016; Gallice et al. 2022), which do not necessarily resolve the problems, but at least make them less visible.

In this work, we use a new, entropy-based HLLD Riemann solver (HLLS; Popovas 2022), which gives the accuracy and performance of a Godunov solver while avoiding the negative pressure problem and providing correct convective flow velocities. Maintaining an accurate hydrostatic equilibrium is also more manageable with entropy as the main energy variable.

We use the MUSCL-Hancock algorithm with constrained transport (CT; Evans & Hawley 1988) for the induction equation, as well as a positivity preserving 3D unsplit TVD slope limiter; see Fromang et al. (2006) for details. The source terms from gravity, Coriolis forces, Newton cooling, etc. are added during both the predictor and the corrector steps. These procedures are explained in detail in Popovas (2022)—here we only briefly summarise the partial differential equations the solver employs.

2.1. Partial Differential Equations

The fluid equations for the total mass density, momentum density, entropy per unit mass and the magnetic field of the system can be written as a system of conservation laws,

$$\frac{\partial}{\partial t} \begin{bmatrix} \rho \\ \rho \mathbf{u} \\ \rho S \\ \mathbf{B} \end{bmatrix} + \nabla \cdot \begin{bmatrix} \rho \mathbf{u} \\ \rho(\mathbf{u} \otimes \mathbf{u}) + P_{tot} - \mathbf{B} \otimes \mathbf{B} \\ \rho \mathbf{u} S \\ \mathbf{B} \otimes \mathbf{u} - \mathbf{u} \otimes \mathbf{B} \end{bmatrix} = \begin{bmatrix} 0 \\ \Phi \\ \mathbf{Q}/T_{gas} + S_{gen} \\ 0 \end{bmatrix}, \quad (1)$$

with

$$\nabla \cdot \mathbf{B} = 0, \quad (2)$$

where ρ , \mathbf{u} , S , P_{tot} , and \mathbf{B} are density, velocity, entropy per unit mass, total pressure and magnetic field respectively; Φ is force per unit volume, which includes force of gravity, Coriolis force, etc; \mathbf{Q} is heating per unit volume, which includes Newton cooling, radiative heat transfer, etc; T_{gas} is gas temperature and S_{gen} is entropy generation from converting kinetic energy to heat. The system is closed by an equation of state (EOS), described in the next subsection.

The system in Eq. 1 can be written in vectorial form, similar to e.g. Londrillo & Del Zanna (2000),

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial \mathbf{F}}{\partial x} + \frac{\partial \mathbf{G}}{\partial y} + \frac{\partial \mathbf{H}}{\partial z} = \Psi, \quad (3)$$

where

$$\mathbf{U} = (\rho, \rho u_x, \rho u_y, \rho u_z, \rho S)^T, \quad (4)$$

$$\Psi = (0, \Phi_x, \Phi_y, \Phi_z, \mathbf{Q}/T_{gas} + S_{gen})^T, \quad (5)$$

and

$$\mathbf{F} = \begin{pmatrix} \rho u_x \\ \rho u_x^2 + P_{tot} - B_x^2 \\ \rho u_x u_y - B_x B_y \\ \rho u_x u_z - B_x B_z \\ \rho u_x S \end{pmatrix} \quad (6)$$

is the flux function. The expressions for the terms \mathbf{G} and \mathbf{H} are analogous.

The divergence-free magnetic field condition (Eq. 2) inside patches is maintained via the very accurate CT algorithm (see e.g. test results in Popovas 2022). As in the Ramses HLLD solver (Fromang et al. 2006), the magnetic field is thus actually evolved with a 2-D Riemann method that corresponds to

$$\frac{\partial \mathbf{B}}{\partial t} = -\nabla \times (\mathbf{B} \times \mathbf{U}), \quad (7)$$

while the \mathbf{B} -part of Eq. 1 is simultaneously used in a 1-D HLLD Riemann solver that produces updates of mass density, momentum, and entropy.

As we interpolate in space and time, the interpolated magnetic field is not necessarily divergence free in the guard zones. To mitigate that, we use a simple, localised divergence-cleaning routine in the guard zones, correcting the perpendicular magnetic field vector component, given the components parallel to each patch face. The procedures will be explained in more detail in Popovas & Nordlund (2022, in prep.b).

2.2. Departure from entropy conservation

In a closed system, energy cannot be created or destroyed; it can only change its form. This is a fundamental requirement and is mathematically fulfilled in energy-based Riemann solvers¹. We have converted the solver into an entropy-based one, and thus we need to account for entropy changes in shocks and magnetosonic waves. Indeed, accounting for this entropy increase is the defining feature of this new solver. The generated entropy S_{gen} depends on the state of the Riemann fan and is characterised by the change of density, normal velocity and total pressure, when crossing a given discontinuity. This generated entropy is then added to ρS , carefully ensuring that it is added to the correct side of the interface between two cells, where an entropy-generating discontinuity occurred. A more detailed procedure description is beyond the scope of this work and is deferred to Popovas (2022).

2.3. Equation of state

We use a tabular EOS (Tomida et al. 2013, updated by Tomida & Hori 2016), which uses logarithmic density and temperature as independent variables. The default table dimensions are 461x761. In order to use the table with our new solver, we transformed the EOS table to use entropy per unit mass instead of temperature as the second variable. The use of the tabular EOS in the HLLS solver, as well as details of the high-performance DISPATCH EOS lookup module will be explained in detail in Popovas & Nordlund (2022, in prep.a), and is only briefly summarised here:

After reading the EOS tables, we add entropy per unit mass to them. This is done by integrating

$$dS = \left[\frac{1}{T} \left(\frac{\partial \varepsilon}{\partial T} \right)_{\rho} \right] dT + \left[\frac{1}{T} \left(\frac{\partial \varepsilon}{\partial \ln \rho} \right)_{T} - \frac{P}{T\rho} \right] d \ln \rho, \quad (8)$$

here ε is internal energy. Next, we re-interpolate the table to a denser grid to have more accurate values in the quick

¹ albeit Riemann solvers in practice rarely include the gravitational energy in their ‘total energy’, and hence require explicit gravitational work terms when forces of gravity are present

look-up routines. Then, we remap all table quantities to a density-entropy table. Lastly, and importantly, we check that our new table satisfies the thermodynamic identity,

$$\left(\frac{\partial \varepsilon}{\partial \rho} \right)_T = \frac{P}{\rho^2} - \frac{T}{\rho^2} \left(\frac{\partial P}{\partial T} \right)_{\rho} = \frac{P}{\rho^2} \left[\left(\frac{\partial \ln P}{\partial \ln T} \right)_{\rho} \right], \quad (9)$$

and gives exact results to numerical precision for requested quantities, when compared to the original table.

This EOS table is later used both to modify the initial hydrostatic equilibrium (see section 5) and to close the PDEs in Eq. 1.

2.4. Code units

There are more suitable choices than CGS or SI units to have the best numerical precision. In our simulation, we use code units with length $l = 10^8$ cm = 1 Mm, time $t = 10^3$ s, and mass density $\rho = 10^{-7}$ g/cm³. All other quantities are then normalised correspondingly, based on their dimensions.

2.5. Experimental setup

We use the ‘volleyball’ domain decomposition described below to conduct the experiment. The radial extent of the simulation domain is from 0.655 to 0.995 R_{\odot} . At the top boundary we are thus less than 3 Mm below the photosphere. To emulate surface cooling we use Newton cooling instead of the short characteristics radiative heat transfer that we plan to use in the photosphere. The smallest cell size near the surface in the simulation described in this paper is 500 km. When relaxation, described below, is complete we will further refine the experiment and go down to a factor of eight smaller cells.

In parallel with the primary simulation, we ran two additional classes of simulations, with increased simplifications. The first one—the ‘small volleyball’ has exactly the same parameters and the same physics, but the resolution is lower, with 1.56 Mm cell size at the finest level. It contains about 30,000 patches.

The second class, denoted ‘sandbox’ runs, is even more simplified:

- the volleyball geometry is replaced with a plane-parallel Cartesian mesh decomposition;
- the vertical extent is identical to the main experiment but the horizontal extent depends on requirement and is periodic;
- a Coriolis force is not present, but
- all other physical effects are identical;
- static mesh refinement is used to reach any required resolution.

These last two simulations accompany the main one as test-beds. They run much faster and are much cheaper. The ‘sandbox’ simulation can run on a local workstation, with just a few cores assigned to it. The ‘small’ simulation can be run on a single compute node. Their main function is to run far ahead of the main simulation to validate the physics modules, the stability of the setup, etc.




Fig. 1. One layer of a volleyball mesh, with the back side removed, for clarity.

3. Domain decomposition

In DISPATCH, we use a set of Cartesian patches arranged in experiment-specific geometries. Each patch contains N^3 cells (in the current work, we generally use $N = 24$, plus 3 guard zone cells on each side, thus 30^3 cells in total). These patches are then placed on a sphere, in a volleyball-like arrangement, cf. Fig. 1. The patches are oriented so their z axes always lie in the radial direction. The tilt between adjacent patches is very small, with exceptions at the ‘seams’ between the six ‘faces’ of the volleyball.

3.1. Volleyball patch arrangement

This mesh decomposition may, from a first glance, appear similar to a cubed-sphere (Ronchi et al. 1996). However, the cubed-sphere is a truly curvi-linear coordinate system in each face region, in which case the solver must be built with the metric coefficients of that coordinate system explicitly known and used. In the volleyball decomposition we have purely Cartesian patches, arranged on a sphere with slight overlaps between them. Note that there is no plane-parallel approximation employed; the force of gravity, for example, is everywhere pointing in the exact radial direction, and as a consequence, hydrostatic solutions inside the patches are small spherical cut-outs, which—despite their small size—deviate ever so slightly from plane-parallel solutions.

We start by tessellating a sphere by generating a set of points. The point object P stores as attributes

1. its position both in spherical and Cartesian coordinates,
2. the size of the patch,
3. its north vector $\underline{\Xi}$, and
4. the coordinate system unit vectors \underline{R} .

An important value here—a horizontal step—is the angular spacing between the patches. It is defined as $\Delta\theta = \theta_0/N_\theta$, where $\theta_0 = \tan^{-1}(2^{-1/2})$ (the angular half-width of a volleyball face) and N_θ is a parameter that can vary from layer to layer (see below). We start with the outermost layer, i.e. $r = r_{max} = 0.995$, with half a patch size L north-east offset from the centre of a front volleyball face (there are 6 faces in total—front, back, left, right, top and bottom—alternatively referred to as front, back, west, east, south and

north). Given the number of intervals N_θ in the θ direction, we generate the point positions and patch sizes L ,

$$L = r \frac{\Delta\theta}{1 - \Delta\theta}. \quad (10)$$

We cover the whole face by stepping $\Delta\theta$ eastwards, i.e. the new ϕ coordinate of the point is then

$$\phi = \phi_{old} + \frac{\Delta\theta}{\cos(\theta)}. \quad (11)$$

This stepping, interleaved with steps northward,

$$\theta = \theta_{old} + \Delta\theta, \quad (12)$$

ends up covering a quarter of a face ($1/24$ or the entire surface). The rest of the surface is covered by executing a sequence of symmetry and flipping operations—we spawn the identical points on the other 6 volleyball faces by simply applying a **symmetry** operation, which involves three **permutation** operations (going from front to right, to top, returning to the front) and two **reflection** operations (going from front to back). These operations have several sub-steps in them; e.g. for **permutation**:

1. permute the Cartesian coordinates;
2. permute the North proxy;
3. convert Cartesian coordinates to spherical;

and for **reflection**:

1. invert Cartesian coordinates;
2. convert Cartesian coordinates to spherical.

After the **symmetry** operation, we go from the east to the west quadrant of the face (an **east-west** operation):

1. $\phi_{new} = -\phi$;
2. convert spherical coordinates to Cartesian;

In the same way, we need to go from the north to the south quadrant (a **north-south** operation):

1. $\theta_{new} = -\theta$;
2. convert spherical coordinates to Cartesian.

We do the last two operations twice more to get back to the original starting point and, of course, do all the **symmetry** operations in-between these steps.

We continue by stepping east, and keep stepping as long as the point is entirely inside the front face, i.e. inside the latitude limit of the front volleyball section, and as long as its lower south west patch corner is outside of the east face. When the patch goes out of bounds in the east direction we go back to the starting point and take a step north. We keep repeating the steps until the whole surface of the particular layer is covered. When this is done we are ready to move one layer down.

In spherical coordinates a volume with a constant angular size $\Delta\theta$ would become physically smaller and smaller with diminishing radius. However, in the interior of the Sun, the size of features and typical timescales are increasing with diminishing radius. Another important fact is the increase of the speed of sound with diminishing radius—together with a decreasing size this can severely restrict the allowed time-step. Instead, we want the cell size to increase with decreasing radius, to adjust to the increasing scales of motion and thus optimise the performance of the




Fig. 2. Radial cut of the volleyball layers.

simulation. The simplest solution is to reduce N_θ , and thus make $\Delta\theta$ larger. Extra care must be taken so that N_θ does not become large enough for gaps to appear in the tessellation. To avoid gaps, we define a smallest allowed N_θ . Figure 2 shows an example radial cut through the layers. As may be seen, patches near the surface are much smaller than patches near the bottom of the model.

3.1.1. Rotation matrices

To rotate the patches into their correct orientation we use rotation matrices $\underline{\underline{R}}$. These matrices have the unit vectors of the local patch coordinate system expressed in the global Eulerian coordinates as columns and the unit vectors of the global coordinate system expressed in local coordinates as rows. The rotation matrix is thus:

$$\underline{\underline{R}} = \begin{bmatrix} \iota & \kappa & \lambda \\ \mu & \nu & \xi \\ \tilde{x} & \tilde{y} & \tilde{z} \end{bmatrix} \quad (13)$$

where $\tilde{P} = [\tilde{x}, \tilde{y}, \tilde{z}] = P/(P \cdot P)^{1/2}$ is the normalised Cartesian position P , representing unit vector ‘up’, $\zeta = \Xi \times \tilde{P}$, $\Psi = [\iota, \kappa, \lambda] = \zeta/(\zeta \cdot \zeta)^{1/2}$ representing a ‘unit vector east’, $\vartheta = \tilde{P} \times \Psi$, $[\mu, \nu, \xi] = \vartheta/(\vartheta \cdot \vartheta)^{1/2}$ representing a ‘unit vector north’. Vectors in the local patch coordinate system can be transformed into vectors in the global coordinate system by noting that each component in the global coordinate system is the scalar product of the local vector with the global unit vectors, which encourages that the local vectors are considered to be column vectors, while global vectors are considered to be row vectors. The transformations may thus be written

$$[g_1 \quad g_2 \quad g_3] = \underline{\underline{R}} \cdot \begin{bmatrix} l_1 \\ l_2 \\ l_3 \end{bmatrix}, \quad (14)$$

while the transformation from global to local coordinates correspondingly is

$$\begin{bmatrix} l_1 \\ l_2 \\ l_3 \end{bmatrix} = [g_1 \quad g_2 \quad g_3] \cdot \underline{\underline{R}}. \quad (15)$$

In Fortran and Python these operations can be performed using the built-in `matmul(a,b)` procedures, which allow both \mathbf{a} and \mathbf{b} to be either vectors or matrices.

If coordinates in the global and local coordinate systems share a common origin these transformations are complete. If they are relative to the centres of the local patches, the distance between the patch centres need to be taken into account, by being added either before (if given in the local coordinate system) or after the matrix multiplication (if given in the local coordinate system).

3.1.2. Relative coordinate transformations

To transform from one local coordinate system to another, for example in the context of guard zone interpolations, one could transform from one local system to a global system and then from there to another local system. It is more convenient, however, to make use of a *relative transformation matrix*, which has the local unit vectors of the ‘source’ patch, expressed in the ‘target’ coordinate system as columns, and correspondingly has the local unit vectors of the ‘target’ patch expressed in the ‘source’ coordinate system as rows. It is clear then that *the ‘source’ coordinates take the role of the ‘local’ coordinates above, while the ‘target’ coordinates corresponds to the ‘global’ coordinates above* (common to a number of different ‘sources’ that contribute to the same ‘target’).

$$\begin{bmatrix} t_1 & t_2 & t_3 \end{bmatrix} = \underline{\underline{R}} \cdot \begin{bmatrix} s_1 \\ s_2 \\ s_3 \end{bmatrix}, \quad (16)$$

and

$$\begin{bmatrix} s_1 \\ s_2 \\ s_3 \end{bmatrix} = \begin{bmatrix} t_1 & t_2 & t_3 \end{bmatrix} \cdot \underline{\underline{R}}. \quad (17)$$

Since the transformation from source to target can also be achieved by first translating the source coordinates to global coordinates and then go from there to target coordinates, it follows that the relative transformation can be obtained as the product of the transposed target rotation matrix and the source rotation matrix:

$$\underline{\underline{R}} = \underline{\underline{R}}_{\text{target}}^T \cdot \underline{\underline{R}}_{\text{source}} \quad (18)$$

3.2. Guard zone interpolations

DISPATCH utilises local time-stepping for individual patches, thus each patch needs guard zone data from neighbour (`nbor`) patches interpolated in both space and time. Since the mass density varies a lot (both vertically and horizontally), we convert to per-unit-mass variables, where applicable, before doing time-interpolations into the `target`’s time. Most patches have only slightly tilted neighbours, making it potentially possible to use standard Cartesian coordinate interpolations, where only the fractional interpolation weight varies from point to point. However, near the ‘seams’ between the six volleyball faces, the patches can have arbitrary overlaps and different north vectors, which require more general interpolation methods. In this context we use the concepts of *regions of authority* (ROA) and *regions of interest* (ROI). The region of authority of a patch is the volume inside which the patch ‘has authority’ and




Fig. 3. A **target** patch (red) and its **source** neighbours (blue). Red dashed lines show the guard zones of the **target**. Note the slight tilt between adjacent patches.

updates the variable values. The region of interest in addition includes the guard zones of the patch, and potentially may include sub-volumes where a neighbour overlapping patch has a higher level of refinement. For simplicity and consistency we currently use these more general point-by-point interpolation methods for all patches. Figure 3 shows a **target** patch with its surrounding **source** neighbours. Note the slight tilt between patches.

For patches with a significant relative tilt between each other, a simple lower-left-corner and upper-right-corner interval range for the ROI is not sufficient. Instead, we take all 8 corner points and find the smallest and largest index in the local frame of reference of the patch we need, to get guard zone data for the **target**.

3.2.1. Interpolation in time

Each patch stores the MHD quantities at 5 different times, in a rotating time-slot buffer. As described in Nordlund et al. (2018), a patch is ready to be updated when all **nbor** patches are ahead in time (including a potential ‘grace’ time interval) of **target**. We perform linear interpolation between the time slots of **nbor** (**source**) to find the time slot indices and weights. These weights are then applied to the selected **source** memory slots (using only the relevant index sub-regions) to get the required quantities for the given **target** time. Higher order interpolation in time is available as an option.

3.2.2. Cell-centred quantities

The rotation matrix \underline{R} is defined at the centre of the patch, thus to get the **target** cell coordinates $p_t = [x_t, y_t, z_t]$ in the **source** system $p_s = [x_s, y_s, z_s]$, we need to take the scalar product of the location in the **target** system with the rotation matrix \underline{R} of the **source** expressed in **target** coordinates. These are available from Eq. 18. The coordinates that go into the operation are the **target** coordinates rel-

Table 1. Mesh decomposition parameters of the simulation.

Bottom radius	0.655 R_\odot
Top radius	0.995 R_\odot
N_θ	36
$N_{\theta, min}$	8
N_θ reduction rate	4
Levels of refinement	3
Initial max resolution	500 km
Final max resolution	62.5 km

ative to the **source** position, i.e. the distance from **target** to **source**, d is subtracted before the multiplication:

$$p_s = [(p_t - d) \cdot \underline{R}] / c + o, \quad (19)$$

where c is the patch’s cell size and o is the patch’s position offset in index space. We go through the **target**’s ROI and map all the **target** points to **source** points in its respective ROA. Of course, not all ROI might be available in a given index space we previously determined, thus we use an array, which tells us if a given data point is **available**.

Lastly, we re-interpolate the MHD quantities from the **source** to the **target**’s frame of reference. For scalars this is just a simple interpolation in space, but for vectors (e.g. cell-centred velocities) we need to transform the vector components to the **target** system by using Eq. 16.

3.2.3. Face-centred quantities

Face-centred vector fields, e.g. the magnetic field, pose an additional complication—different vector components are located on different faces. Equation 19 now needs an additional component for the staggered position of the vector:

$$p_s = [(p_t - d) \cdot \underline{R}] / c + o + h_\mu, \quad (20)$$

where μ is a coordinate axis. This gives us three position data sets for the vector components. We go through the **target**’s ROI and map all the **target** points to **source** points in its respective ROA. For each vector component, we re-stagger the perpendicular vector components into the position of the vector component in question and store them separately. Later these re-staggered components are used to transform the vector to the **target** system of reference by using Eq. 16.

3.3. Simulation setup and static mesh refinement

Table 1 summarises the basic parameters for the mesh decomposition we use. The simulation spans 0.655–0.995 R_\odot . With $N_\theta = 36$, we have 10 layers in the setup. The maximum resolution near the surface is 500km. This is marginally sufficient to resolve the dominant flows at this depth and to relax the initial setup into a steady-state convective Sun in the shortest time possible.

As soon as the simulation reaches a steady-state, differential rotation is fully established, and local dynamo action is actively twisting the magnetic fields, we can increase the resolution by using static mesh refinement. An example of such refinement is shown in 4. Here, only 2 levels of refinement are added at 0.88 and 0.97 R_\odot . In the production




Fig. 4. Static mesh refinement in a ‘sandbox’ version of the simulation. 2 additional levels are added, starting at 0.88 and 0.97 R_{\odot} .

run, after the simulation is relaxed enough, 3 levels of refinement can be added with additional relaxation between each additional level. This gives us a final resolution of 62.5 km. We do a factor of 2 refinement, creating 8 new `leaf` patches on top of a `parent` patch. Newly created patches are co-aligned with the parent patch, so their \underline{R} is the same as `parent`’s.

The top layers are the most numerous and each additional level of refinement means a factor of 2 reduction in time-step and 8 times more cells, so we don’t want to run with maximal resolution the whole time. Such simple mesh refinement can be added during routine restarts of the run, i.e. when node allocation on a compute cluster ends (this is for example 2 days allocation on LUMI and 4 days allocation on Betzy machines), or when the simulation triggers some sort of a condition (elapsed time is a typical one for static mesh refinement).

4. Source terms

Numerical simulations of the Sun can be as complicated as one wants and can provide a wealth of different diagnostics (e.g. Przybylski et al. 2022; Carlsson, Mats et al. 2016; Kohutova & Popovas 2021; Finley et al. 2022; Druett et al. 2022 to name a few), but for now, we limit ourselves to only a few key variables. Additional micro- and macro-physics effects can easily be added later.

4.1. Force of gravity

For smaller 3D Cartesian simulations, it is common to have a globally constant gravitational acceleration g , acting along one of the axes, but since our simulation extends through the whole convective zone we must have a more realistic, radially dependent, $g(r)$. We interpolate the value for g from our 1D model in hydrostatic equilibrium (see subsection 5.1 for details). Since all patches have their z axis pointing radially outwards, in any patch’s local coordinate frame the force per unit mass is simply

$$f_g = -g(r) * \hat{r}, \quad (21)$$

where $\hat{r} = [x, y, z + r_c]/r$, with $[x, y, z]$ being the cell positions in the local frame of reference, relative to the centre of the patch, r_c is the radial distance from the centre of the Sun to the centre of the patch, and $r = \sqrt{x^2 + y^2 + (z + r_c)^2}$ is the radial distance from the centre of the Sun to a given cell.

4.2. Newton cooling

We use Newton cooling in the top few cell layers to drive convection. We interpolate thermal energy per unit mass, ε_0 , from our hydrostatic equilibrium data and obtain its slope, $\Delta\varepsilon$. Heating per unit mass is then

$$Q_N = [(\varepsilon_0 \psi - \Delta\varepsilon s) - \varepsilon_{x,y,z}] \eta, \quad (22)$$

where ψ is a parameter that controls the strength of the Newton cooling; $s = p - r$, with p being the position of a patch’s centre in a global frame of reference; $\eta = \xi/(1 + \xi)\tau_0$, is an exponential decay function, where $\xi = e^{-s}$ and τ_0 is the cooling timescale, which we set to 200 s.

4.3. Coriolis force and rotation

The experiment is set up in a co-rotating frame of reference with respect to the midplane of the Sun, and the whole Sun is initially rotating as a solid body at a constant velocity. This implies that initial longitudinal velocity is zero. The differential rotation develops later with the help of Coriolis and centripetal force. We use $\Omega_{\text{coriolis}} = 2.6 \times 10^{-6} \text{ s}^{-1}$.

5. Initial and boundary conditions

The experiment is initialised in a hydrostatic equilibrium (see below), with small amplitude random perturbations in vertical momentum at radii, above 0.75 R_{\odot} ,

$$\delta(\rho u_z) = \alpha \sin(2\pi\zeta), \quad (23)$$

where $\alpha = 0.01$ is the amplitude of perturbation and $\zeta \in [0, 1]$ is a random number. This small perturbation speeds up the development of convective motions.

5.1. Hydrostatic equilibrium

We start the experiment with a hydrostatic equilibrium, iterated using our EOS tables together with the 1D Model-S from Christensen-Dalsgaard et al. (1996) as a starting structure. We compute the integrated mass, acceleration of gravity per each layer, then we re-integrate the hydrostatic balance, with constant entropy in the convection zone.

5.2. Initial magnetic field

There are many ways to initialise the magnetic field in local, Cartesian simulation domains, e.g. uniform vertical, salt-and-pepper, flux emergence, etc. There are fewer ways to initialise a magnetic field in a global solar simulation. The main requirement—it must be divergence free, the Sun must not become a magnetic monopole (i.e. no purely vertical magnetic field). Thus the magnetic field can be initialised either as a torus, or as a uniform field in the global frame of reference. We initialise a uniform magnetic field, aligned with the vertical axis in the global frame of reference, i.e. $\mathbf{B}_0 = [B_x, B_y, B_z] = [0, 0, 0.1]$ Gauss. This constant magnetic field is then transformed to local frame of reference via Eq. 16. We chose such simple approach because of two reasons:

1. simplicity— $\nabla \cdot \mathbf{B} = 0$ by definition initially;
2. convective motions quickly erase the simplicity of this primordial magnetic field alignment.

5.3. Boundary conditions

The experiment covers the entire surface of the Sun, thus we only need to consider radial boundary conditions (BCs).

5.3.1. Bottom boundary

The bottom boundary is located below the tachocline, where there are no more significant convective motions, while some waves can still propagate. This simplifies things significantly, allowing us to use diminishing boundary condition for the vertical momentum component and symmetric BCs for horizontal velocity components. The vertical magnetic field component has a zero-gradient BC and constrained extrapolation is applied to the horizontal magnetic field components. At the last internal point we enforce a pressure node BC. For inflows (we check vertical velocity one cell above), density and entropy are approaching ρ_0 and S_0 with characteristic time $\tau_{bot} = 1000$ s. For outflows, we adjust towards constant initial pressure P_0 , leaving entropy per unit mass as is. In the guard zone layers both density and entropy are extrapolated logarithmically.

5.3.2. Top boundary

The top boundary is located at 3.5 Mm below the surface. It is cork-like, with an imposed vanishing derivative on the velocity components. The vertical magnetic field component once again has zero gradient BC and constrained extrapolation is applied to the horizontal magnetic field components. The density boundary condition is based on total pressure ($P_{tot} = P_{gas} + P_{mag}$) scale height. The horizontal components of the momenta are assumed to taper off with density, consistent with assuming a no-shear boundary condition for the horizontal velocity components. The vertical velocities are tapered off with a density ratio factor. Entropy per unit mass is kept constant, which implies adiabatic vertical temperature gradients.

6. Relaxation and initial results

Due to finite numerical precision we expect the initial hydrostatic equilibrium not to be perfect and some oscillations might occur while the system is relaxing into an equilibrium state. We dampen these oscillations out, as it would take large amounts of time for them to propagate out of the simulation domain. To do so, we compute a horizontally averaged vertical momentum over an MPI sub-domain at each time-step and for the first 3000 solar seconds we apply the average value to generate friction (as force per unit volume) over a characteristic time-scale $\tau = 100$ s. After 3000 seconds of run time friction is turned off and the experiment is allowed to run unimpeded.

To relax the simulation faster we exaggerate the Newton cooling, i.e. we set $\psi = 0.95$ in Eq. 22. This makes the convective flux about 3 times stronger. As a consequence, the convective cells are also initially larger. After the relaxation is complete we reduce the Newton cooling to values more representative of the solar surface.

We see the first signs of convective motions appearing at around 2 hours of solar time. These have highly repetitive patterns, as initial perturbation is only partially randomised. However, this repetition is very quickly broken down by numerical rounding errors, and at around 6 hours

of solar time the convection pattern is essentially disordered.

Figure 5 shows entropy per unit mass in the experiment. The left hand side panel shows the surface slice, approximately 4.5 Mm below the top boundary. Large convective cells are clearly visible and no repetitive patterns can be identified. The right hand side panel shows a slice through the domain at the prime meridian. Convective flows penetrate deeply into the atmosphere. The tachocline can be identified from the sharp drop in entropy per unit mass.

Figure 6 shows radial velocity in the experiment. Due to enhanced Newton cooling, the radial velocities are correspondingly larger. Velocity vector fields do not show any interpolation glitches between adjacent patches. This also includes seams, where relative angles are large and we need to count on accurate transformations. We took a random sample of individual patches throughout the box and investigated the transition between the guard zones and interiors just after the guard zone interpolation. Transitions in both vector and scalar quantities were smooth, without any discontinuities. This is very encouraging, as smooth transitions between interfaces are rather challenging (e.g. Brchnelova et al. 2022).

Figure 7 shows the radial component of the magnetic field in code units. From both the left and right hand side panels it can be seen that the primordial structure is already mostly hidden by local dynamo effects; in some places the magnetic field magnitude is already two orders of magnitude stronger than the initial mean magnetic field. We can, in fact, see some magnetic vortices forming in some areas. These initial concentration could later give rise to magnetic activity at the surface, e.g. in the forms of network structure, pores, and sunspots. The magnitude of local dynamo effects is greatly exaggerated here, due to much stronger convection, but the topological and morphological effects of the convective flow hierarchy are expected to be largely independent of magnitudes.

Using our new HLLS solver provides us with a unique bug-catching tool. Guard zone interpolation glitches create artificial discontinuities. Discontinuities trigger jump conditions, which in turn increase entropy per unit mass. Even errors in the magnetic field can be misidentified as magnetosonic waves, which can increase entropy. Even if the buildup of entropy per unit mass is marginal, after tens of thousands of updates and continuous buildup we would be able to see glitches in entropy per unit mass, even if other quantities might be able to partially “mask” them. Figure 5 shows no signs of such localised entropy increases.

6.1. Weak and strong scaling

The DISPATCH framework has demonstrated excellent scaling results past 150,000 cores on a fluid dynamics problem (Nordlund et al. 2018). These results stem from the nearest-neighbour-only MPI communications, combined with the near-perfect scaling of OpenMP-parallel task updates on each rank, and thus do not depend on the detailed properties of the solver. Non-Cartesian geometry implies slightly less straightforward MPI decomposition, which in turn might affect load balancing. In volleyball seams patches can have more than 26 neighbours with more awkward overlapping regions of authority. All this, together with additional, experiment-specific physics, such as Newton cooling, gravity, Coriolis force, etc., could affect the

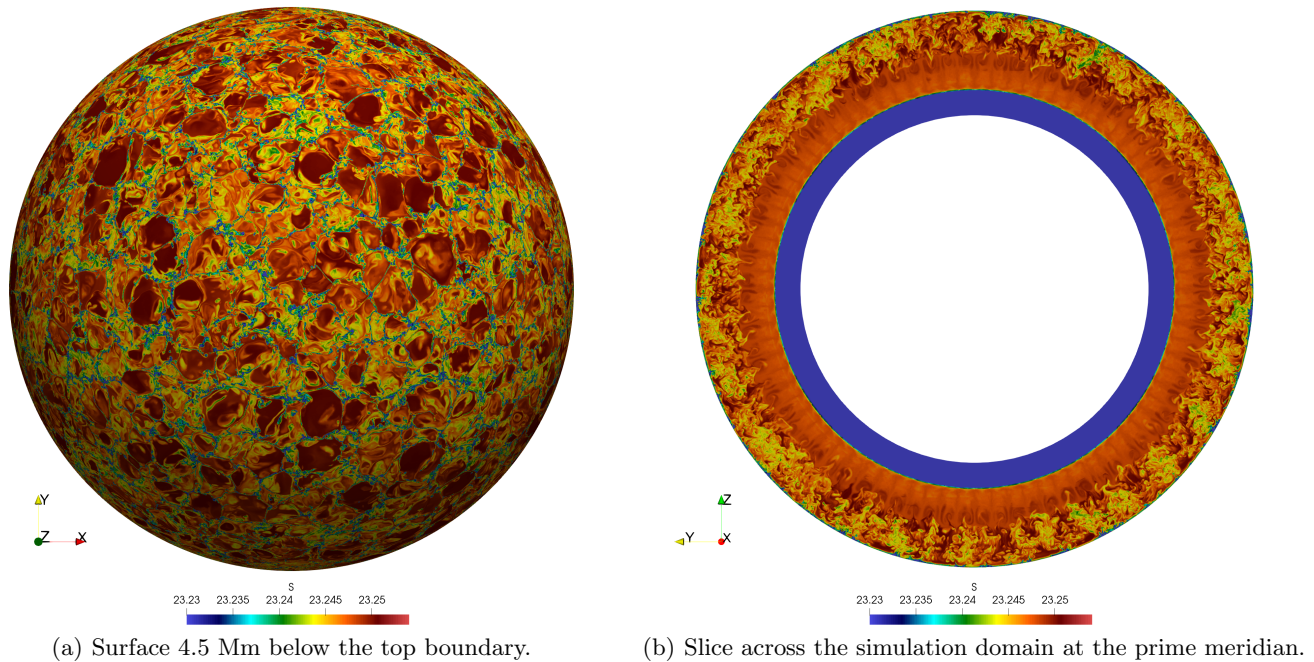


Fig. 5. Entropy per unit mass in the simulation, in code units. Movie available online.

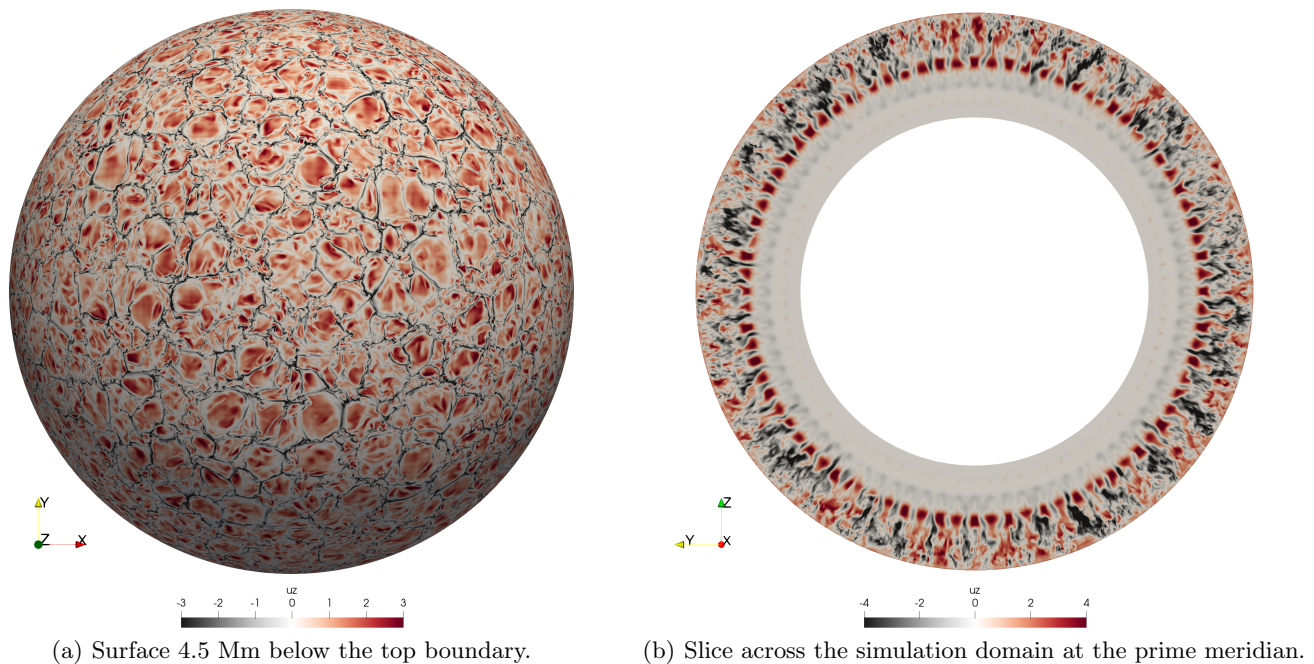



Fig. 6. Radial velocity in the simulation, in code units. Movie available online.

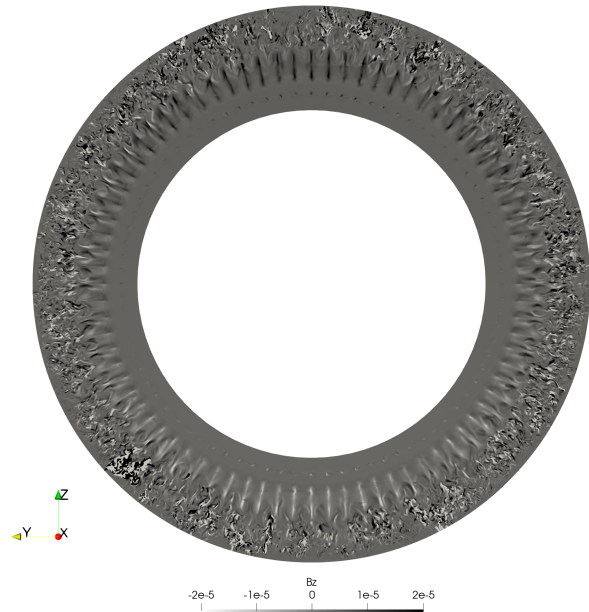
perfect scaling that was previously demonstrated in simpler geometries (Nordlund et al. 2018). Thus we conducted new weak and strong scaling tests with the current experiment, to check that excellent scalability can still be achieved. The tests were conducted on the LUMI and Betzy supercomputer systems. Both of these machines have AMD EPYC CPUs (models 7742 and 7702 respectively), which makes them very similar, to the extent that Betzy can be consid-

ered LUMI’s ‘little sister’. The two machines have different node interconnect and different compilers that we use, though, Intel and GCC respectively. Thus the differences in performance stem from the differences in the system configurations and software.

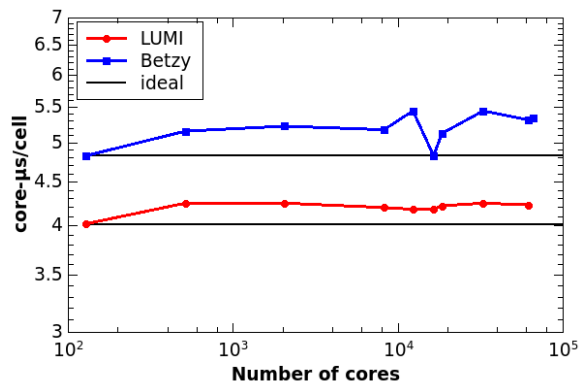
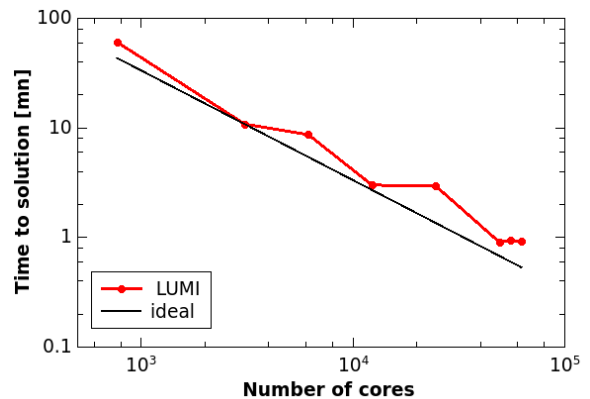
To test the weak scaling, we employ our simplified ‘sandbox’ runs, where we fix the vertical resolution and expand in the horizontal plane with an increasing number of cores.



(a) Surface 4.5 Mm below the top boundary.



(b) Slice across the simulation domain at the prime meridian.

Fig. 7. Radial magnetic field component, in code units.

Fig. 8. Weak scaling results on LUMI (red) and Betzy (blue) supercomputers. For both machines the ideal scaling is shown in black.

Fig. 9. Strong scaling results on LUMI supercomputer.

We fix the patch dimensions to 24^3 cells and keep the number of patches per compute node constant. In all the runs there are 2 MPI ranks per node. In this test we take the average update time (in $\mu\text{s}/\text{cell}$) over an extended period of run-time in the execution phase (we ignore the initialization phase). Figure 8 shows that with our experimental setup we can still achieve nearly perfect weak scaling.

To test the strong scaling we run the complete experimental setup. This setup is kept fixed to check that different MPI decomposition (albeit still carefully chosen to have reasonably good initial load balance) does not negatively affect the run-time and the experiment can be successfully executed on a range of MPI decomposition configurations. In this test we measure the total wall-clock time it takes for an experiment to get from the start of the execution until experiment time reaches 0.2 code units. With a typical time-step of the order of approximately 10^{-3} – 10^{-2} , each

patch is updated between of order 10 to 100 times. Figure 9 shows the result and it is clear that the volleyball mesh decomposition together with the experimental setup has very good strong scaling properties. Note, that at the highest number of cores the efficiency drops, as the number of patches drops to about 2.5 patches per core, which is well below the optimal 20-50 tasks per core. The nearly perfect scaling is re-established if mesh-refinement is used and the workload (the number of patches per core) increases.

7. Discussion

The new volleyball experiment setup provides a unique opportunity to study complex plasma dynamics over a broad range of scales in the Sun, both at depth, in the surface layers, and in the corona.

The pilot relaxation runs described here were carried out on Norway's largest supercomputer, Betzy, and on the

world’s 3rd largest supercomputer to date, LUMI in Finland. On LUMI, we ran the relaxation simulation using 96 nodes (each node contains 128 cores, totalling 12288 cores). With this size, in a 48 hour allocation period, the relaxation simulation evolves 48.3 hours, and thus gives evolution close to real solar-time. The ideal allocation for the refined production run will be about 900 nodes ($\sim 110,000$ cores).

The performance can be further improved by a factor of order 2-3, by re-organizing the solver memory access pattern—this has recently been demonstrated in a test implementation of the ideal gas HLLD solver.

7.1. Challenges

With the number of Cartesian patches ranging from hundreds of thousands to millions, tessellation into a volleyball arrangement gives some challenges, both with respect to performance, and with respect to analysing and visualising the data.

The interpolations between tilted meshes are slightly costlier than in Cartesian mesh decomposition, but this can be significantly mitigated if the relative angles between adjacent patches is small enough to be handled as standard Cartesian coordinate interpolations (with only the fractional interpolation weights changing) can be used to interpolate the guard zone data. This optimisation can of course not be applied at the seams between the six volleyball ‘faces’, but the affected volume fraction is very small, so the global cost would be dominated by the optimised method.

We have in-house Python3 and Julia (Bezanson et al. 2014) language scripts to read in the snapshots and to process the data as a whole and as slices. The processed data can then be visualised with ParaView². Even without mesh refinement it is difficult to visualise the whole surface with a single, albeit powerful workstation³, while with refined meshes one would either have to use the parallel capabilities of ParaView, or else zoom in to smaller regions, using the “lazy reading” capabilities of the low-level DISPATCH data access Python modules. We are also continuously developing new tools and procedures to speed up the data processing.

The current snapshots, without mesh refinement, need 140 gigabytes of memory. Fully refined simulation will require approximately 4 terabytes of memory. Now consider a future simulation with the core and photosphere included (see below), in the context of helioseismology. To do proper wave analysis we would require two weeks of solar time, with a 1 minute cadence for snapshots. That would correspond to 20160 snapshots, with approximately 100 petabytes of data, if the snapshots were stored “as is”. Thus, for such studies we need to develop different data storage and retrieval solutions.

7.2. Future work

A natural next step after fully relaxing the experiment is adding a photosphere. We wish to delay this step until relaxation is complete, as adding a photospheric layer will

make simulations both more complicated and more expensive:

- Newton cooling will be replaced by full multi-bin, short characteristics radiative heat transport;
- Spitzer conductivity becomes relevant in higher layers;
- the resolution needed to sufficiently resolve the surface convective cells is of the order of a few tens of kilometres;
- upper boundary conditions need some modifications to minimise wave reflection

The current experiment includes the bottom of the convection zone. However, to do global helioseismological studies, the convective stable central part of the Sun would need to be included in the simulations. This would require using a different or extended EOS table source, as the tables by Tomida & Hori (2016) can be used only down to approximately $0.5 R_{\odot}$. Including the central region of the Sun also requires switching to a normal Cartesian tessellation at some radius, to avoid the central singularity. As with the seams already present, we do not expect this to lead to significant glitches.

The structural design of the experiment (composed of many small Cartesian grids) facilitates further exploration with zoom-in type of simulations (regridding / refining in space and time). This technique is commonly used in simulations of star formation (e.g. Kuffmeier et al. 2017; Padoan et al. 2020). We will use this kind of setup to be able to focus on individual active regions without having to impose artificial initial and boundary conditions—these will instead be provided from the large scale simulation of the whole Sun.

8. Concluding remarks

This paper describes a method to simulate the solar convective zone with a new ‘volleyball’ mesh decomposition. The method provides a resolution that can be varied with radius, and while covering the full sphere still avoids coordinate singularities.

We show that the guard zone interpolations between patches do not lead to noticeable numerical artefacts, and that the transition between adjacent patches is generally smooth, also at the ‘seams’ between the six volleyball faces.

The proof of concept runs described here also demonstrate how the DISPATCH framework’s flexibility allows us to gradually build up experiments by including more microphysics, and by manipulating the mesh decomposition—adding new layers, refining, or zooming in. This will be indispensable in the near future, when we have completed the relaxation and will be able to use the existing simulations as starting points for further experiments, with potentially different configurations and scientific goals.

Variations on the experimental setup can be used to study, for example:

- local and global dynamo effects
- active regions and flux emergence
- local and global helioseismology
- ...

Studies that will concentrate on specific features in space and/or time can use zoom-in techniques to focus on the areas of interest and enhance the realism (use additional physics modules) all the way to a similar extent as e.g.

² ParaView is an open-source, multi-platform data analysis and visualisation application. <https://www.paraview.org/>

³ 2× AMD EPYC 7742 64-Core Processor + 2× NVidia Tesla V100

Carlsson, Mats et al. (2016), with boundary conditions replaced by immersion in fully spherical initial models.

In a broader context, the volleyball setup can be applied with advantage to any problem with spherical geometry, without the need to handle coordinate singularities, and with local representations that always have one axis pointing in the vertical direction. Contexts where this technique, in combination with the cost-saving advantages of local time-stepping, could significantly reduce simulation costs include for example simulating planet growth by pebble accretion, studying dynamics and escape of early planet atmospheres, supernova explosions, as well as studies of accreting and interacting black holes.

Acknowledgements. This research was supported by the Research Council of Norway through its Centres of Excellence scheme, project number 262622, and through grants of computing time from the Programme for Supercomputing, as well as through the Synergy Grant number 810218 (ERC-2018-SyG) of the European Research Council.

References

- Bezanson, J., Edelman, A., Karpinski, S., & Shah, V. B. 2014, arXiv e-prints, arXiv:1411.1607
- Brchnelova, M., Zhang, F., Leitner, P., et al. 2022, *Journal of Plasma Physics*, 88, 905880205
- Carlsson, Mats, Hansteen, Viggo H., Gudiksen, Boris V., Leenaarts, Jorrit, & De Pontieu, Bart. 2016, *A&A*, 585, A4
- Chen, F., Rempel, M., & Fan, Y. 2022, *ApJ*, 937, 91
- Christensen-Dalsgaard, J., Däppen, W., Ajukov, S. V., et al. 1996, *Science*, 272, 1286
- Druett, M. K., Leenaarts, J., Carlsson, M., & Szydlarski, M. 2022, *A&A*, 665, A6
- Evans, C. R. & Hawley, J. F. 1988, *ApJ*, 332, 659
- Finley, A. J., Brun, A. S., Carlsson, M., et al. 2022, *A&A*, 665, A118
- Fromang, S., Hennebelle, P., & Teyssier, R. 2006, *A&A*, 457, 371
- Gallice, G., Chan, A., Loubère, R., & Maire, P.-H. 2022, *Journal of Computational Physics*, 468, 111493
- Gudiksen, B. V., Carlsson, M., Hansteen, V. H., et al. 2011, *A&A*, 531, A154
- Guerrero, G., Stejko, A. M., Kosovichev, A. G., Smolarkiewicz, P. K., & Strugarek, A. 2022, arXiv e-prints, arXiv:2208.05738
- Hotta, H. & Kusano, K. 2021, *Nature Astronomy*, 5, 1100
- Hotta, H., Kusano, K., & Shimada, R. 2022, *ApJ*, 933, 199
- Ismail, F. & Roe, P. L. 2009, *Journal of Computational Physics*, 228, 5410
- Kageyama, A. & Sato, T. 2004, *Geochemistry, Geophysics, Geosystems*, 5
- Käpylä, P. J. 2021, *A&A*, 651, A66
- Kohutova, P. & Popovas, A. 2021, *A&A*, 647, A81
- Kuffmeier, M., Haugbølle, T., & Nordlund, Å. 2017, *ApJ*, 846, 7
- Londrillo, P. & Del Zanna, L. 2000, *ApJ*, 530, 508
- Nordlund, Å., Ramsey, J. P., Popovas, A., & Küffmeier, M. 2018, *MNRAS*, 477, 624
- Padoan, P., Pan, L., Juvela, M., Haugbølle, T., & Nordlund, Å. 2020, *ApJ*, 900, 82
- Pan, L., Padoan, P., & Nordlund, Å. 2018, *ApJ*, 866, L17
- Pan, L., Padoan, P., & Nordlund, Å. 2019, *ApJ*, 881, 155
- Parenti, S. 2014, *Living Reviews in Solar Physics*, 11, 1
- Peter, H., Bingert, S., Klimchuk, J. A., et al. 2013, *A&A*, 556, A104
- Popovas, A. 2022, arXiv e-prints, arXiv:2211.02438
- Popovas, A. & Nordlund, Å. 2022, in prep.a, *A&A*
- Popovas, A. & Nordlund, Å. 2022, in prep.b, *A&A*
- Popovas, A., Nordlund, Å., & Ramsey, J. P. 2019, *MNRAS*, 482, L107
- Popovas, A., Nordlund, Å., Ramsey, J. P., & Ormel, C. W. 2018, *MNRAS*, 479, 5136
- Przybylski, D., Cameron, R., Solanki, S. K., et al. 2022, *A&A*, 664, A91
- Reale, F. 2010, *Living Reviews in Solar Physics*, 7, 5
- Rieutord, M. & Rincon, F. 2010, *Living Reviews in Solar Physics*, 7, 2
- Ronchi, C., Iacono, R., & Paolucci, P. S. 1996, *Journal of Computational Physics*, 124, 93
- Teyssier, R. 2002, *A&A*, 385, 337
- Teyssier, R., Fromang, S., & Dormy, E. 2006, *Journal of Computational Physics*, 218, 44
- Tomida, K. & Hori, Y. 2016, personal communication
- Tomida, K., Tomisaka, K., Matsumoto, T., et al. 2013, *ApJ*, 763, 6
- Winters, A. R. & Gassner, G. J. 2016, *Journal of Computational Physics*, 304, 72