FLOW-3D*MP//*

Transition from multi-block type decomposition to true domain decomposition

Introduction

FLOW-3D uses multiple mesh blocks to represent the computational domain. The primary reason is to reduce the computational resources required to model complex flows. Generally, the blocks have different cell sizes and require overlays, i.e., interpolation for fluid fraction, density, temperature, etc., as well as velocity and pressure at the inter-block boundaries. The interpolation is time-consuming and sometimes causes a loss in solution accuracy.

FLOW-3D/MP also uses mesh blocks to decompose the computational space for running on a cluster. Users must create blocks (or domains) in order to run the simulation on multiple cores even though the mesh resolution may not change from one block to another. For example, a 32-rank MPI simulation requires 32 blocks (at the minimum). The automatic decomposition tool (ADT) is designed for subdividing the initial user-defined mesh. Even though the cell sizes in the blocks may not have changed, the current implementation requires solution overlays and interpolations. As a result, inaccuracies and even instabilities may occur near the inter-block boundaries, especially in simulations where the solution varies rapidly. This problem may worsen on higher number of MPI ranks (or processes) because of the increase in the number of inter-block boundaries. The interpolation issue has been alleviated in version 4.2 because of the hybrid implementation. The hybrid implementation uses fewer blocks, in turn leading to less inter-block interpolation. However, the interpolation issues still remain even in the hybrid version and the only way to avoid them is to move to the true domain decomposition scheme described in this note.

Implementation

This section describes the new methodology where interpolation is not used in cases when the cell size does not change across inter-block boundaries. In Fig.1 the mesh blocks are shown for the multi-block gravity mold filling example. The 3 blocks (block 1 in blue, block 2 in pink, and block 3 in green) have different mesh resolutions. In this case, interpolation is required in the inter-block boundary cells between blocks 1 and 2 and between blocks 2 and 3.

If this simulation is run using **FLOW-3D/MP** on 4 ranks, ADT creates 6 blocks as shown in Fig. 2. ADT decomposes the original block 3 into 4 blocks, while keeping blocks 1 and 2 unchanged. As the cell size does not change in these 4 new blocks, there is no need for the overlay/interpolation — the data is transferred directly from the real cells of a donor block to the boundary cells of the neighboring acceptor block. Figure 2 shows the data exchange pattern for the new decomposition. The red arrows indicate inter-block boundaries that require overlays and interpolation, while the black arrows show blocks that directly transfer data to boundary cells.

An inter-block boundary contains two layers of cells. In the original multi-block method, the cell size of these 2 cells in the direction normal to the boundary is equal to the size of the nearest real cell. For the new scheme, at boundaries where interpolation is not needed, the boundary cell size is obtained from the real cells in the donor block.
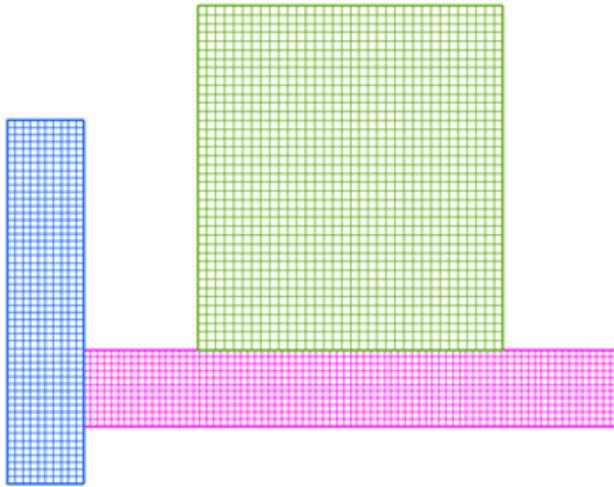
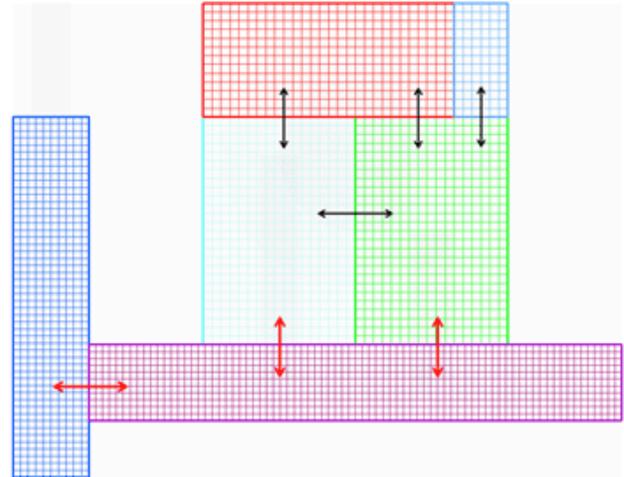Figure 1. Multi-block gravity filling mesh configuration.

Figure 2. New domain decomposition scheme. Red arrows indicate boundaries that require overlays and interpolation. Black arrows indicate direct data transfer to the inter-block boundary cells.

Improvement in Performance

The new decomposition scheme saves computational time whenever interpolation is not required for the solution transfer. Time savings also come from better convergence, resulting in overall performance improvement. The following table shows performance improvements for 2 different simulations with a different number of MPI processes.

Simulation	MPI Processes	Performance Improvement
Lid driven cavity	16	7.3%
Lid driven cavity	32	10.6%
Lid driven cavity	64	11.0%
Sloshing in tank	16	10.2%
Sloshing in tank	32	20.6%
Sloshing in tank	64	25.0%

Table 1. Performance improvements for two different simulations with a different number of MPI processes.

Improvement in Accuracy

The new scheme uses the cell sizes from the neighboring block to form the boundary cells. This results in improved accuracy, as illustrated in the *boxcast* simulation. Figure 3 shows the geometry with the outline of the 4 mesh blocks (from ADT). Figure 4 shows the cell type values (*icstat*) for blocks 2 and 3 and the open active cells are shown in white. The blue cells are the

blocked cells that belong to block 2 and the red cells are the blocked cells that belong to block 3. Figure 5 shows the boundary cells for block 2 to illustrate the difference between the old interpolation scheme and the new scheme. The figure on the left (A) has the boundary cells (x-max) of the same size as the last real cell of block 2. As a result the boundary cells come out as blocked, however, the figure on the right (B) has the boundary cell size of the donor block 3, so the cells are open to flow and heat transfer. This keeps the solution consistent across blocks and results in an improved accuracy.

Figure 3. Mesh-block configuration for the boxcast example.
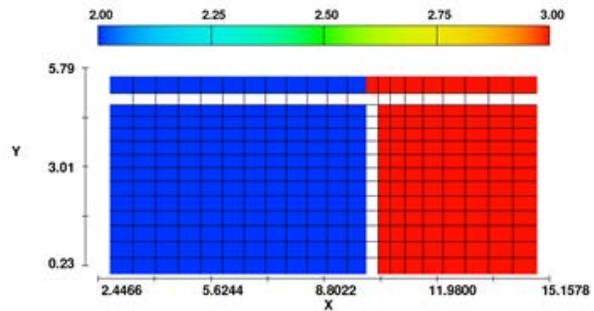

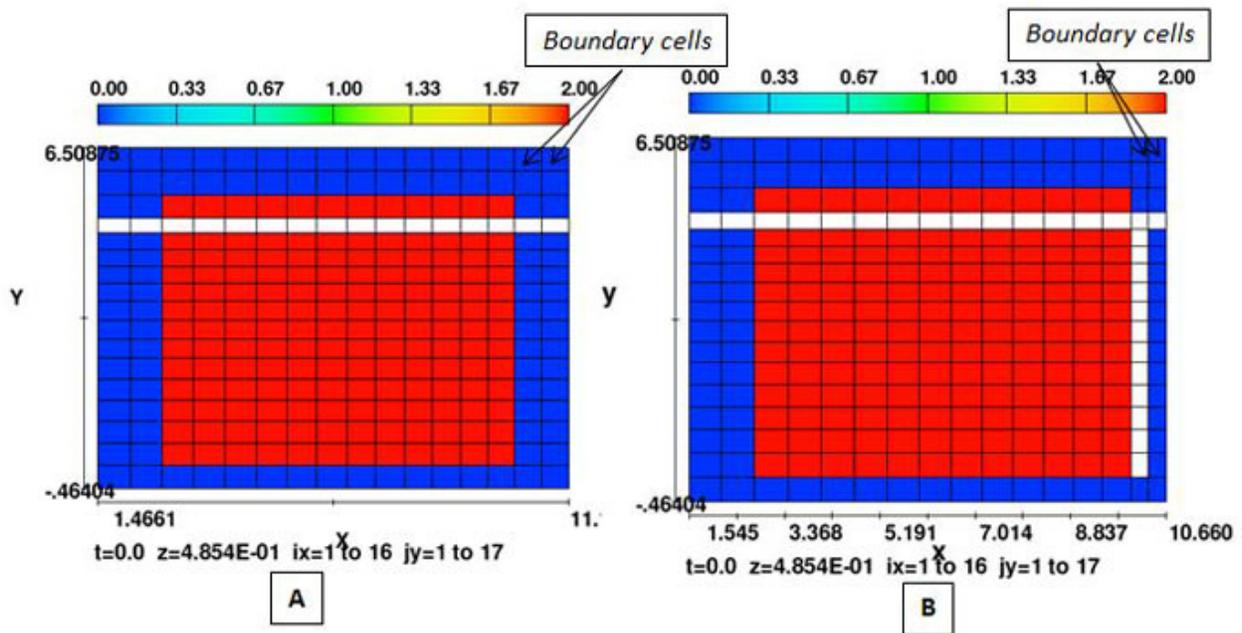Figure 4. The icstat values for mesh blocks 2 and 3 (white color shows open cells).

Figure 5. icstat values for mesh block 2, (A) based on interpolation technique, and, (B) based on the new method.

Conclusion

The true domain decomposition scheme described in this note removes inter-block interpolation where unnecessary. This results in faster runtimes and improves convergence, stability, and accuracy of simulations. For more information about **FLOW-3D/MP**, please contact the Flow Science sales team at sales@flow3d.com.