Parallelization in PUMA

- Implements implicit time integration schemes on distributed memory computers.
- Each compute node reads its own portion of the grid file at startup.
- Cells are divided among the active compute nodes at runtime based on cell ID (shown in figure) and only faces associated with local cells are read.
- Faces on the interface surface between adjacent computational domains are duplicated in both domains. Fluxes through these faces are computed in both domains.
- Solution variables are communicated between domains at every timestep which ensures that the computed solution is independent of the number of compute nodes.
- Communication of the solution across domains is all that is required for first-order spatial accuracy, since Q_L and Q_R are simply cell averages to the first order.
- If the left and right states are computed to higher-order, then Q_L and Q_R are shared explicitly with all adjacent domains. The fluxes through each face are then computed in each domain to obtain the residual for each local cell.