Impact of mesh partitioning on parallel OPM Flow performance

Andreas Thune

August 2022

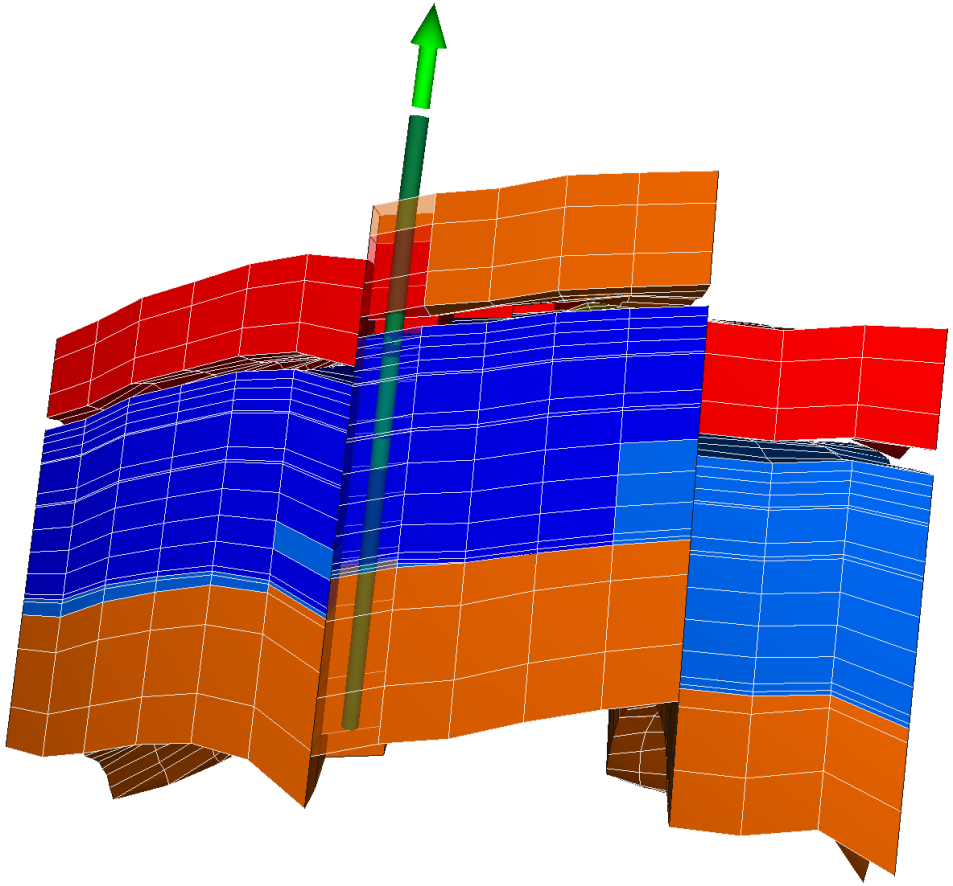

simula

This presentation focuses on distributed wells and improvements in OPM Flows parallel performance.

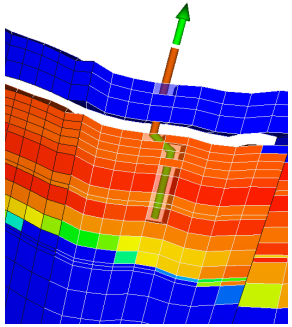

New edge-weighting scheme utilizing distributed well implementation.

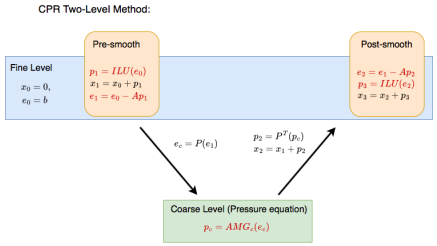

Avoiding non-contributing computations in the CPR preconditioner.

Reservoir mesh is partitioned by first translating it to an edge-weighted graph.

$$\Omega \rightarrow G = (V, E, \omega: E \longrightarrow \mathbb{R})$$

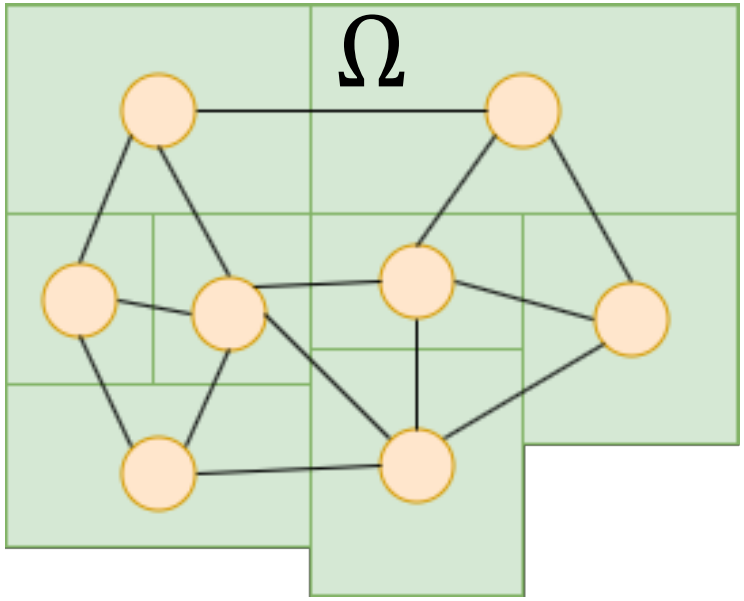

P-way graph partitioning problem

$$min_{\mathcal{C}} \left(J(\mathcal{C}) = \sum_{e \in \mathcal{C}} \omega(e) \right),$$

Subject to: $\frac{P \cdot max_i(|V_i|)}{\sum_i |V_i|} < \epsilon.$

Transmissibility edge-weights are used to yield good linear solver convergence.

Transmissibility and well edge weights:

$$\omega(e) = \begin{cases} T_e, & e \text{ is face edge}, \\ \infty, & e \text{ is well edge}. \end{cases}$$

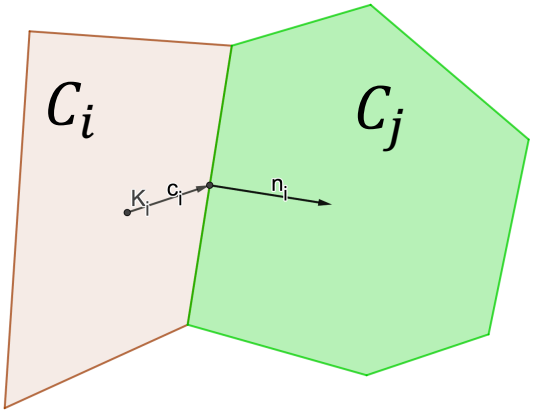

Transmissibility:

$$T_{ij} = \left(\frac{1}{t_i} + \frac{1}{t_j}\right)^{-1} \quad t_i = \frac{\vec{c}_i K_i \vec{n}_i}{||\vec{c}_i||^2}$$

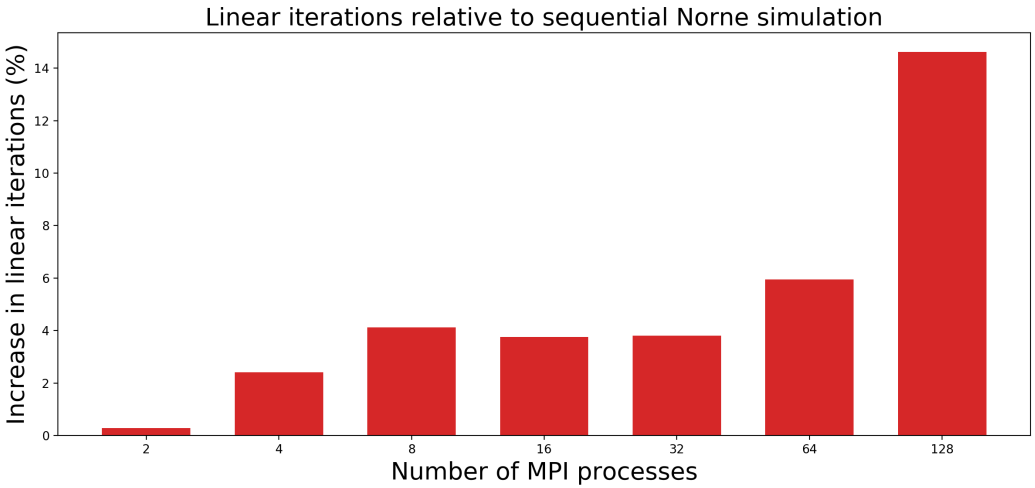

Well edge-weights ensure that wells remain on single subdomain.

Transmissibility and well edge weights:

$$\omega(e) = \begin{cases} T_e, & e \text{ is face edge}, \\ \infty, & e \text{ is well edge}. \end{cases}$$

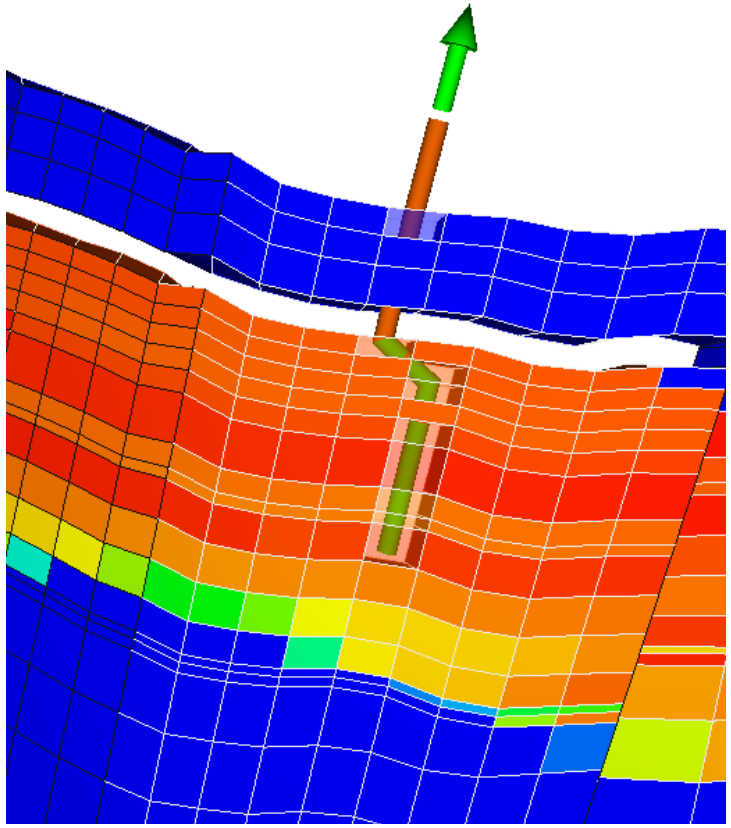

Well edge-weights ensure that wells remain on single subdomain.

Transmissibility and well edge weights:

$$\omega(e) = \begin{cases} T_e, & e \text{ is face edge}, \\ \mathbf{0}, & e \text{ is well edge}. \end{cases}$$

Why 0-weighted well edges:

- Infinity well edge-weights creates poor quality partitions.
- Well contributions are not* added to preconditioner matrix.

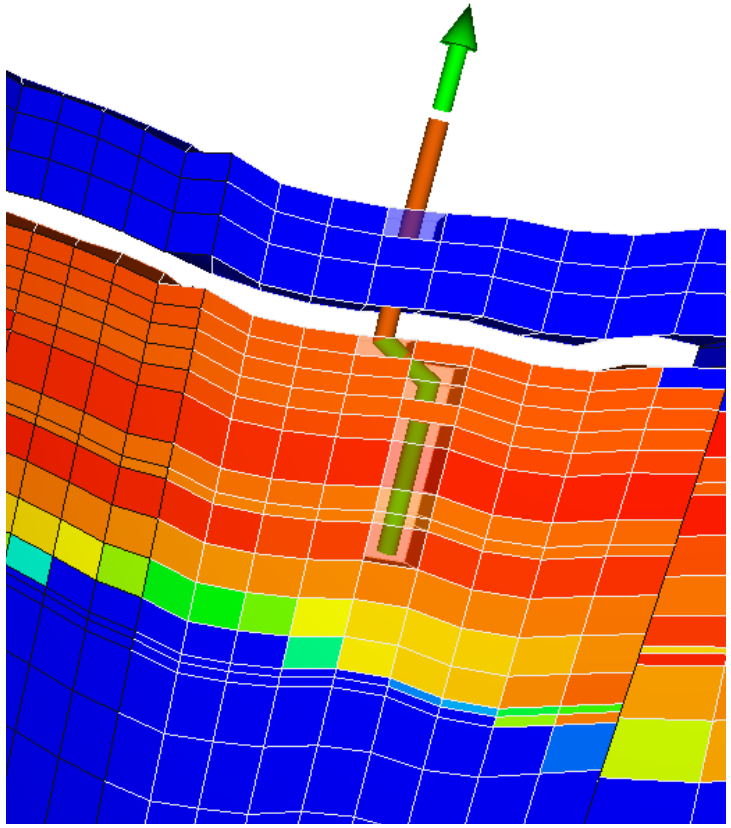

Partitioning without well edge-weights reduces total communication volume.

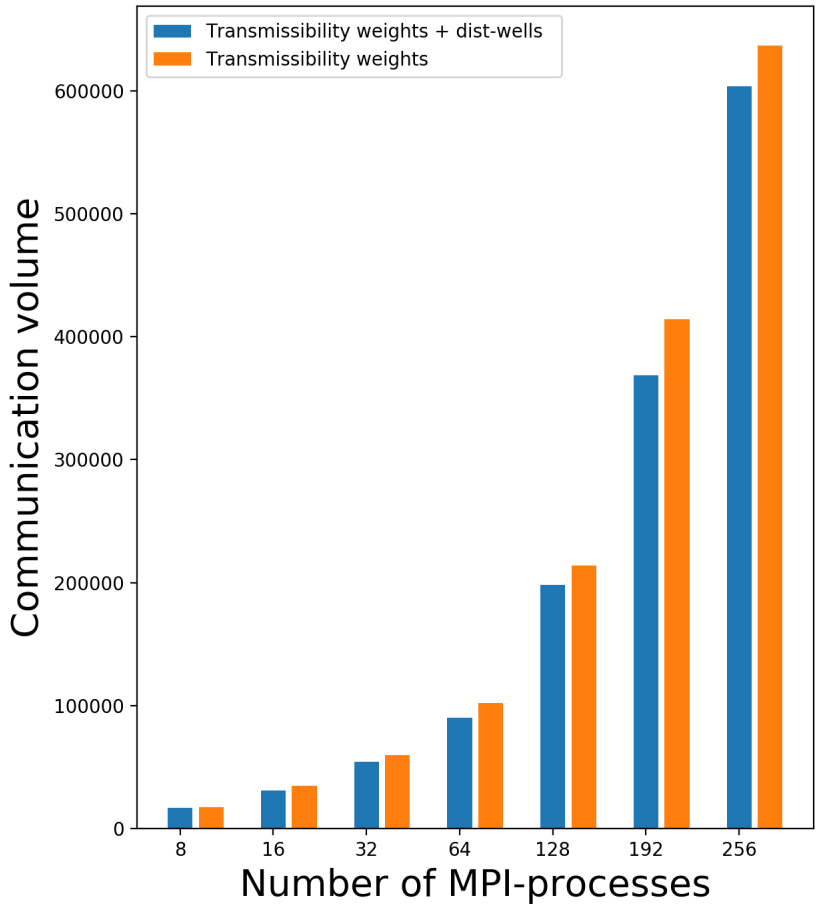

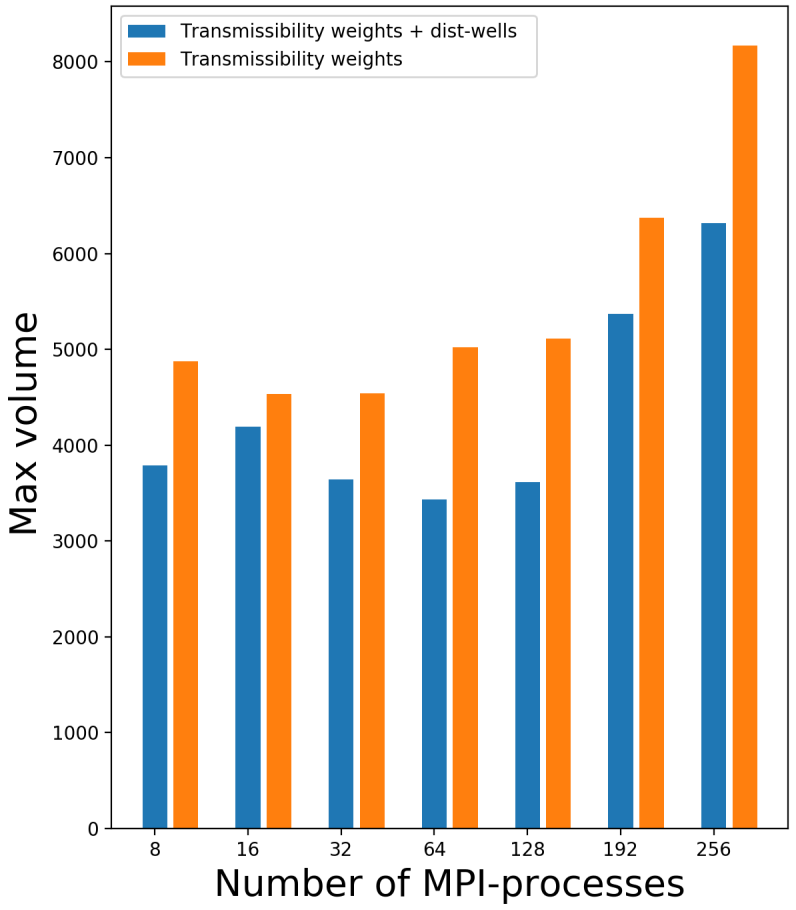

- Case: 1-millon cell black-oil case.

Distributed wells partitioning results in reduced Flow simulation execution time on 1-million cell black-oil case.

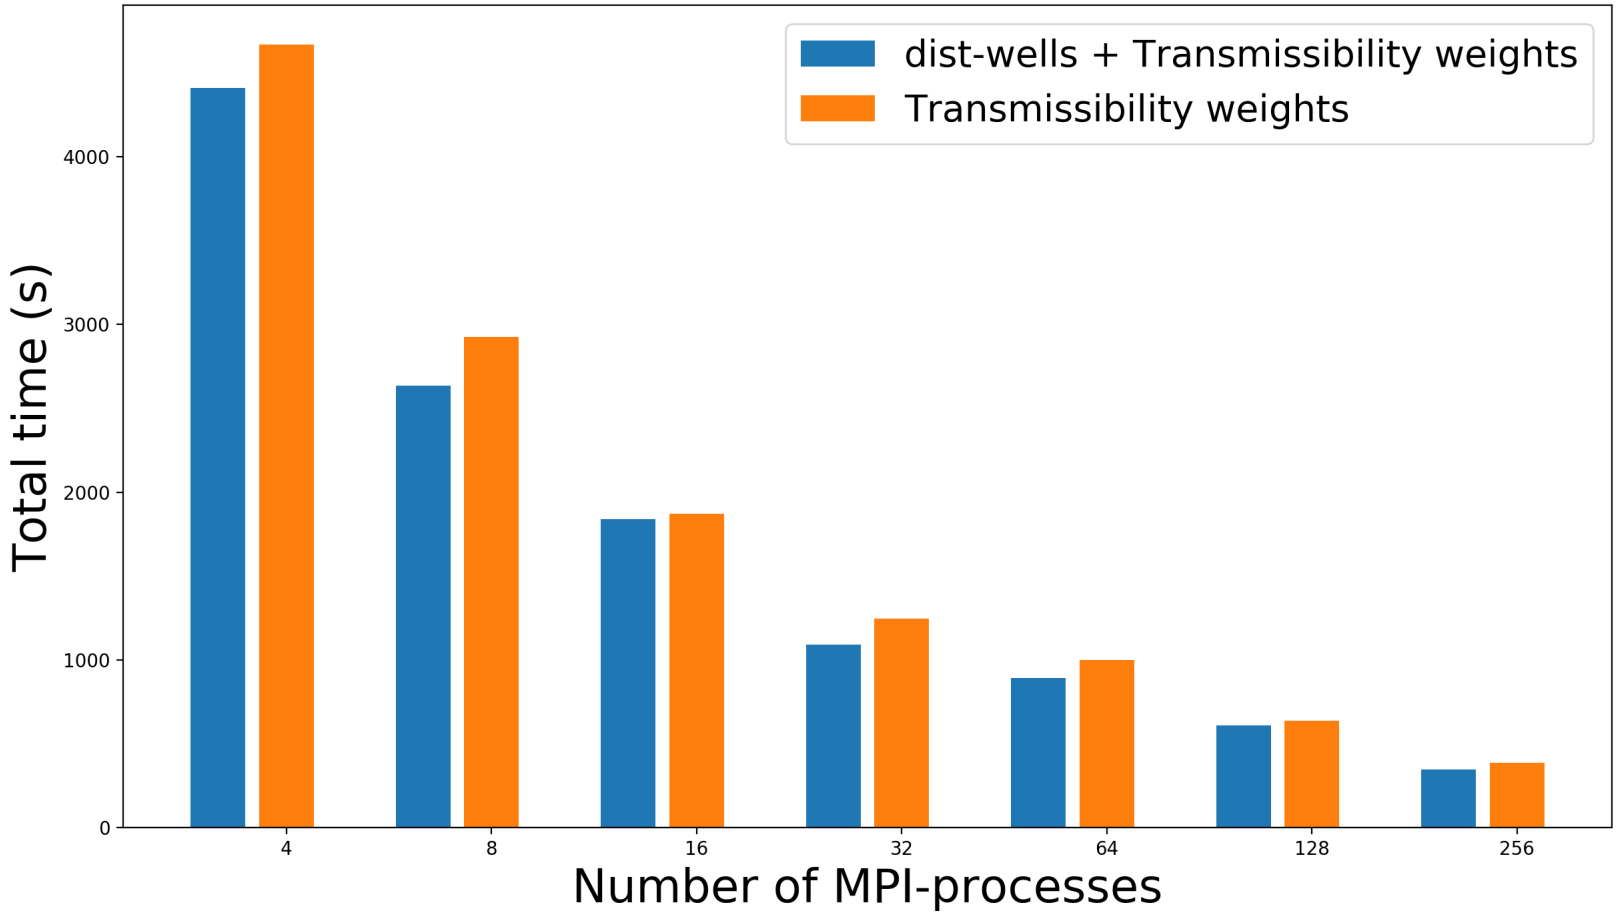

- Case: 1-millon cell black-oil case.
- ILU0 and default settings.
- CPU: Dual-32 core AMD EPYC Naples.

Ghost cells allows for communication free system assembly.

Local grid:

- **Green interior cells.**
- **Red ghost cells.**

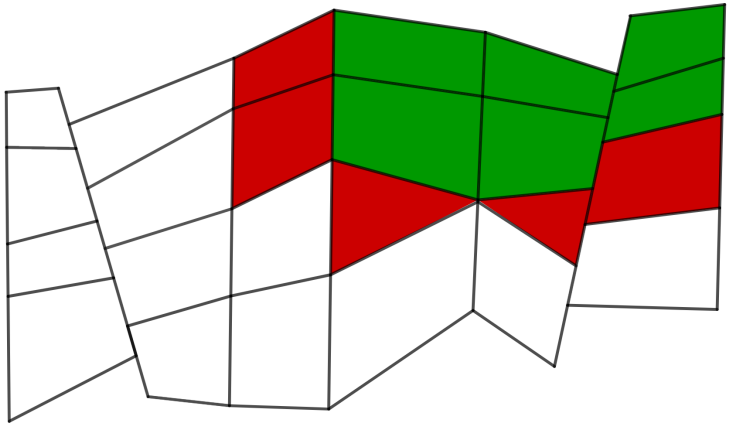

Assembled local system:

$$A_{loc} = \begin{bmatrix} A_{II} & A_{IG} \\ 0 & \mathbb{I}_{GG} \end{bmatrix}$$

Matrix-vector multiplication and smoother operations in CPR and AMG benefit from removing ghost-DoF computations.

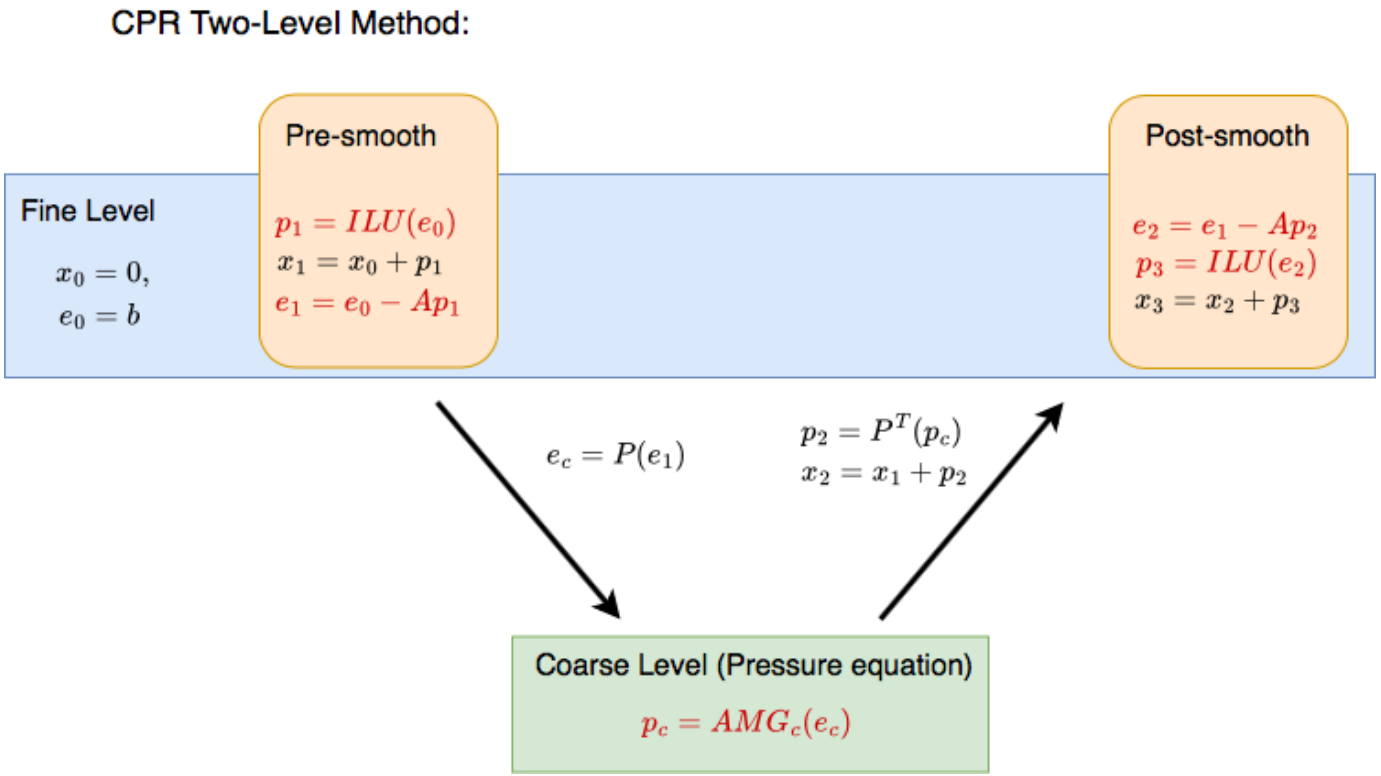

CPR apply time is reduced by 10-20% when ignoring ghost DoFs.

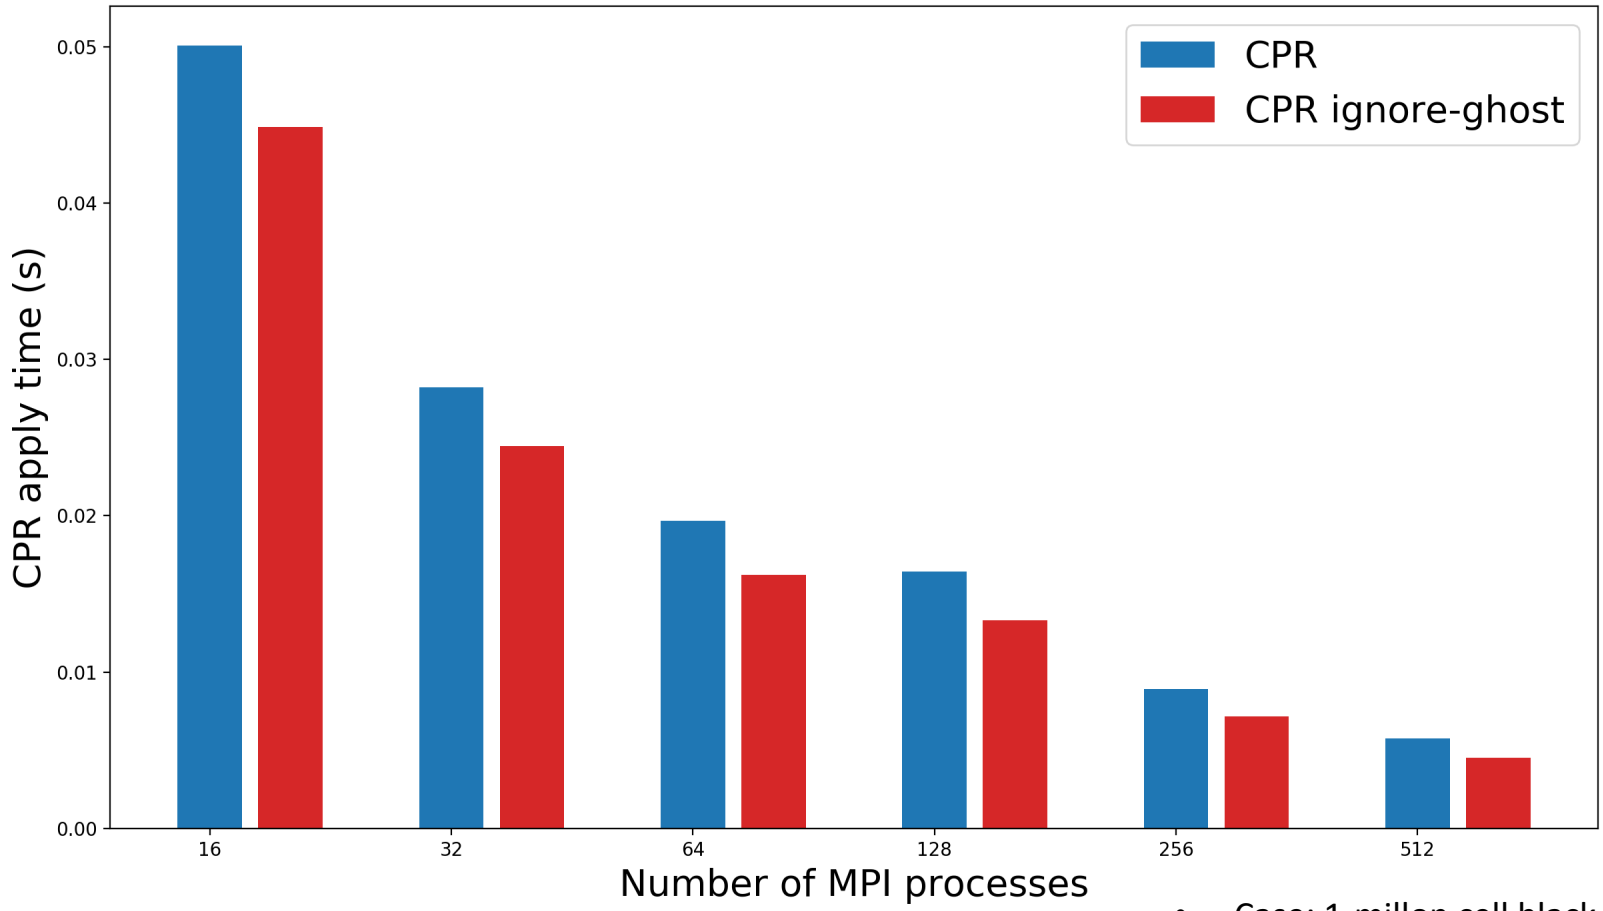

- Case: 1-millon cell black-oil case.
- CPU: Dual-64 core AMD EPYC Milan.

Flow execution time on 1-millon black-oil case is reduced by 5-10% when ignoring ghost DoFs in CPR.

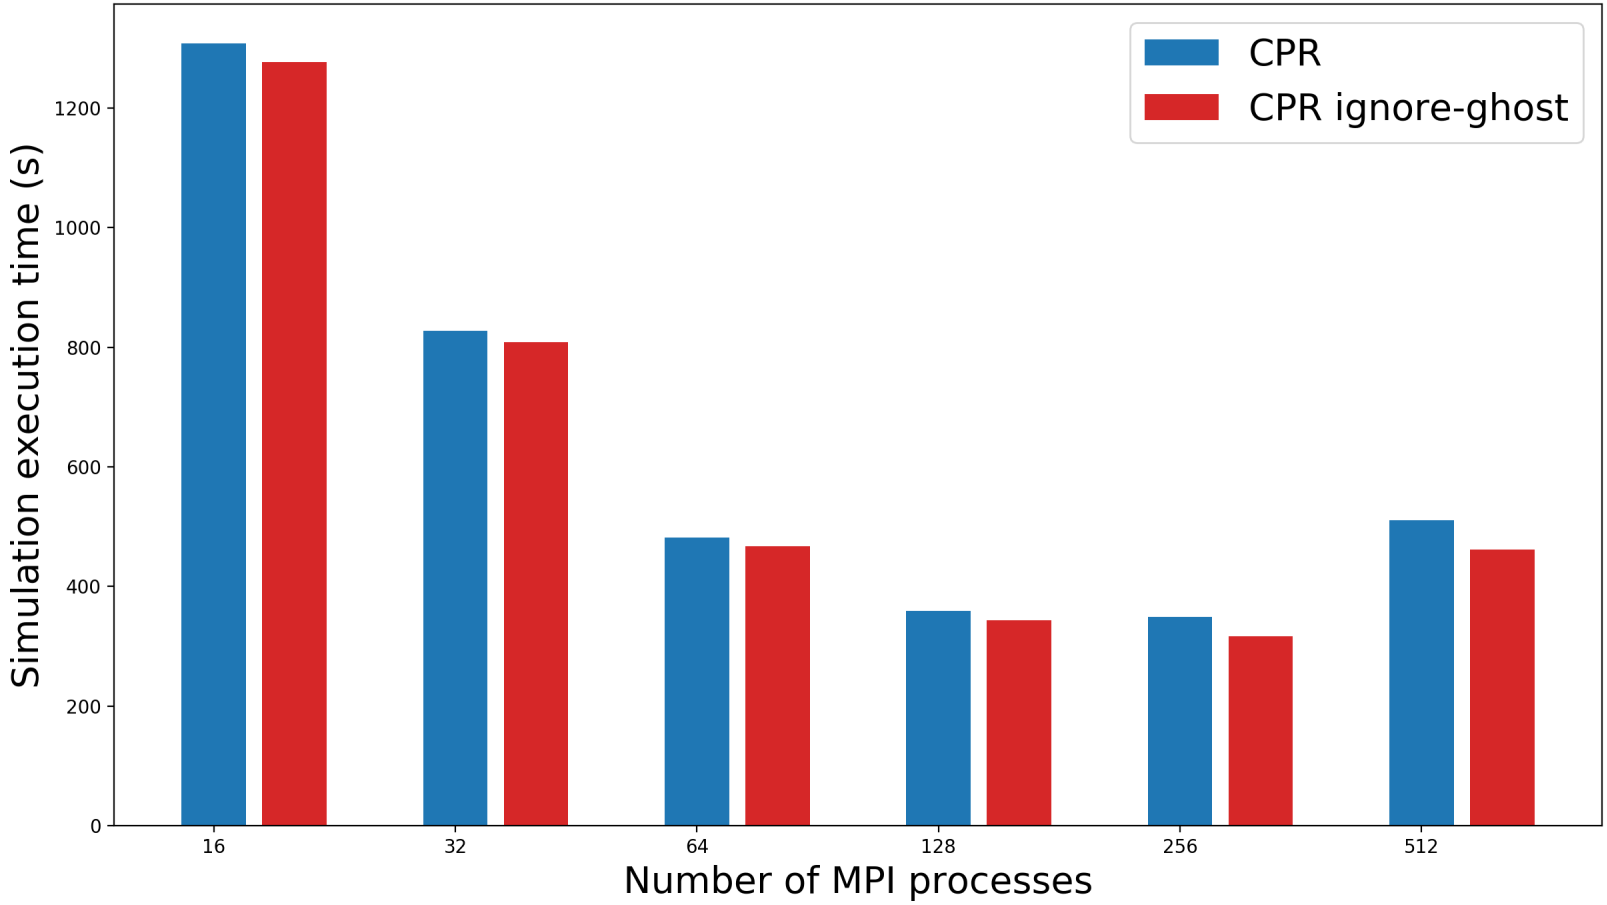

- Case: 1-millon cell black-oil case.
- CPU: Dual-64 core AMD EPYC Milan.

In summary, partitioning edge-weighting strategy impacts OPM Flow performance.

Distributed well implementation can improve partitioning quality and overall performance.

Avoiding non-contributing ghost computations reduce CPR execution time.

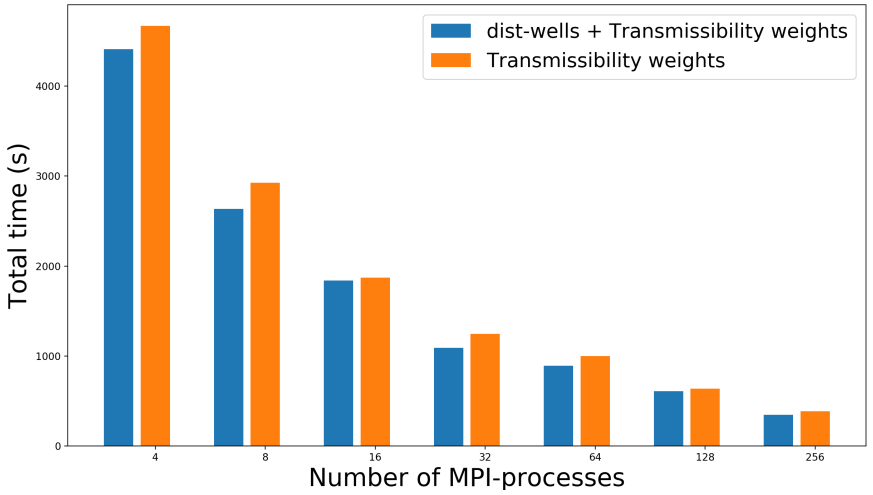

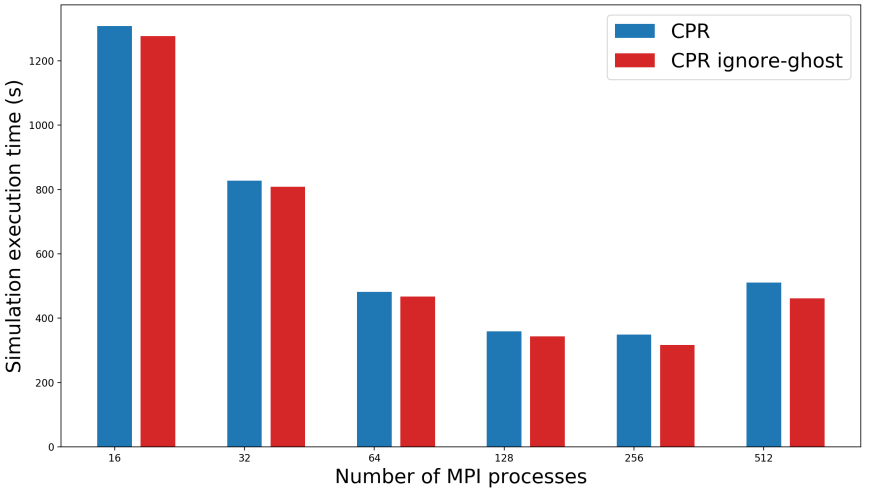

simula

Questions?