

Matrix-Free Preconditioned Conjugate-Gradient Solver



Preconditioned CG Solver structure


while ($\alpha_0 < \text{tol}$):	
$\vec{v} = A\vec{p}$	applyStencil
$\lambda = \alpha_0 / \langle \vec{v}, \vec{p} \rangle$	dotProduct
$\vec{x} = \vec{x} + \lambda\vec{p}$	axpby
$\vec{r} = \vec{r} - \lambda\vec{v}$	axpby
$\text{res} = \langle \vec{r}, \vec{r} \rangle$	dotProduct
$\vec{z} = P\vec{r}$ (fw)	GSPreCon.1
$\vec{z} = P\vec{r}$ (bw)	GSPreCon.2
$\alpha_1 = \langle \vec{r}, \vec{z} \rangle$	dotProduct
$\vec{p} = \vec{z} + \frac{\alpha_1}{\alpha_0}\vec{p}$	axpby

← SpMV “in disguise”

} ← Gauss-Seidel preconditioner

GS preconditioner: wavefront parallelization

Pipeline parallel processing: OpenMP barrier after each wavefront update (ugh!)



Barrier cost

- Does the OpenMP barrier impact the performance?
- At which problem size can this be expected?

Overhead (1 NUMA LD): thousands of cy

