Eurographics Conference on Visualization (EuroVis) 2025
W. Aigner, N. Andrienko, and B. Wang
(Guest Editors)

COMPUTER GRAPHICS *forum*
Volume 44 (2025), Number 3

Supplementary materials

1. Workload estimation explained in detail

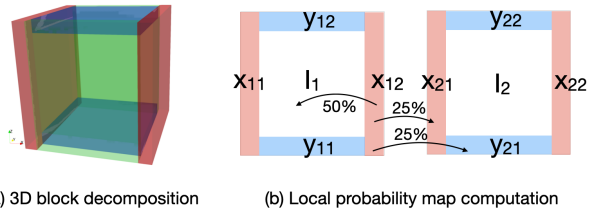

(a) 3D block decomposition (b) Local probability map computation

Figure 1: *Subfigure(a) shows a depiction of our domain decomposition per data block, where each block is divided into seven sub-regions: six boundary regions and one interior region. Testing seeds used for constructing the local flow map are placed in the six boundary regions labeled in different colors. Subfigure(b) shows how probability of local map is computed in 2d domain for region x_{12} as a proof of concept.*

Subfigure 1(b) illustrates details of computing probability for particles located in region x_{12} as proof of concept. For example, region x_{12} may contain 100 test particles, where 50 are advected to the internal block (I_1), 25 to x_{21}, and 25 to y_{21} after the initial round. The corresponding local flow map for x_{12} consists of a vector with three entries. Each entry records the destination identifier, the number of particles advected to that destination, and the total number of steps to that destination for all particles. Additionally, the average number of steps for particles traveling to the destination region can be computed. To ensure that each processing unit maintains a consistent representation of the flow map, we use `AllReduce` operation to gather and broadcast the local flow map to every rank. Associated information is subsequently utilized to determine the probabilistic movement of particles between regions. For instance, a particle initially located in x_{12}has a 50% probability of moving to I_1 in the next round, with the number of steps required being the pre-computed average.

2. Evaluation results of workload estimation accuracy

Figure 2 shows a comparison of block popularity computed from workload estimation (red colored line) and from actual workload execution (blue colored line). Figure 3 illustrates streamline with Tube of Block 21, 22, 25 and 26 for both Supernova and CloverLeaf. Workload estimation errors for these blocks are comparatively high. In order to capture complex structures between these blocks, we need more test seeds, and we may also need to adjust the thickness of boundaries shown in Figure 1(a). In addition, an adaptive setting of the number of testing seeds for each boundary and the number of boundaries along each axis may decrease the error of workload estimation.

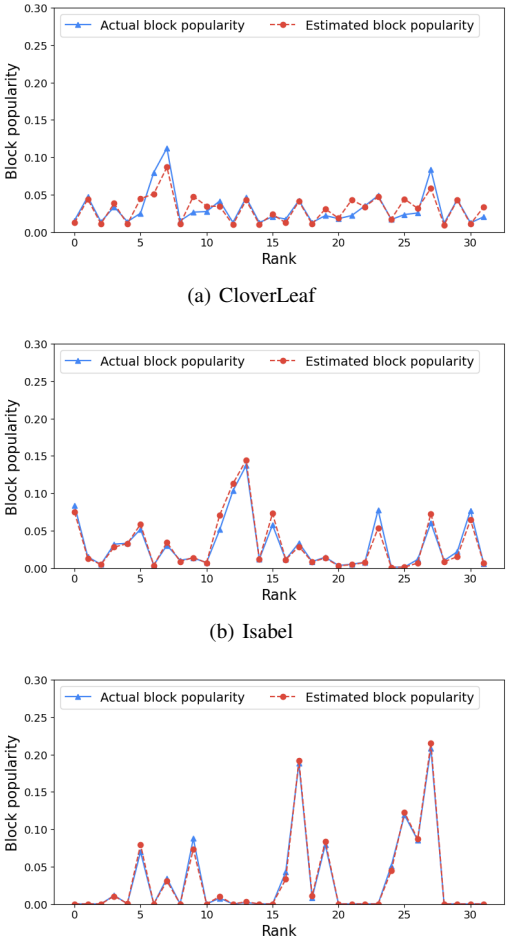

(a) CloverLeaf

(b) Isabel

(c) RedSea

Figure 2: *Detailed results of actual block workload (normalized) and estimated block workload (normalized) for evaluated data sets.*

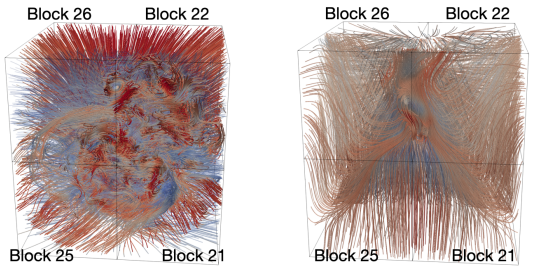

Figure 3: *Illustration of data blocks with relatively higher workload estimation errors in the Supernova and CloverLeaf datasets. Complex vortex structures between these blocks contribute to the estimation inaccuracies.*