



Introduction



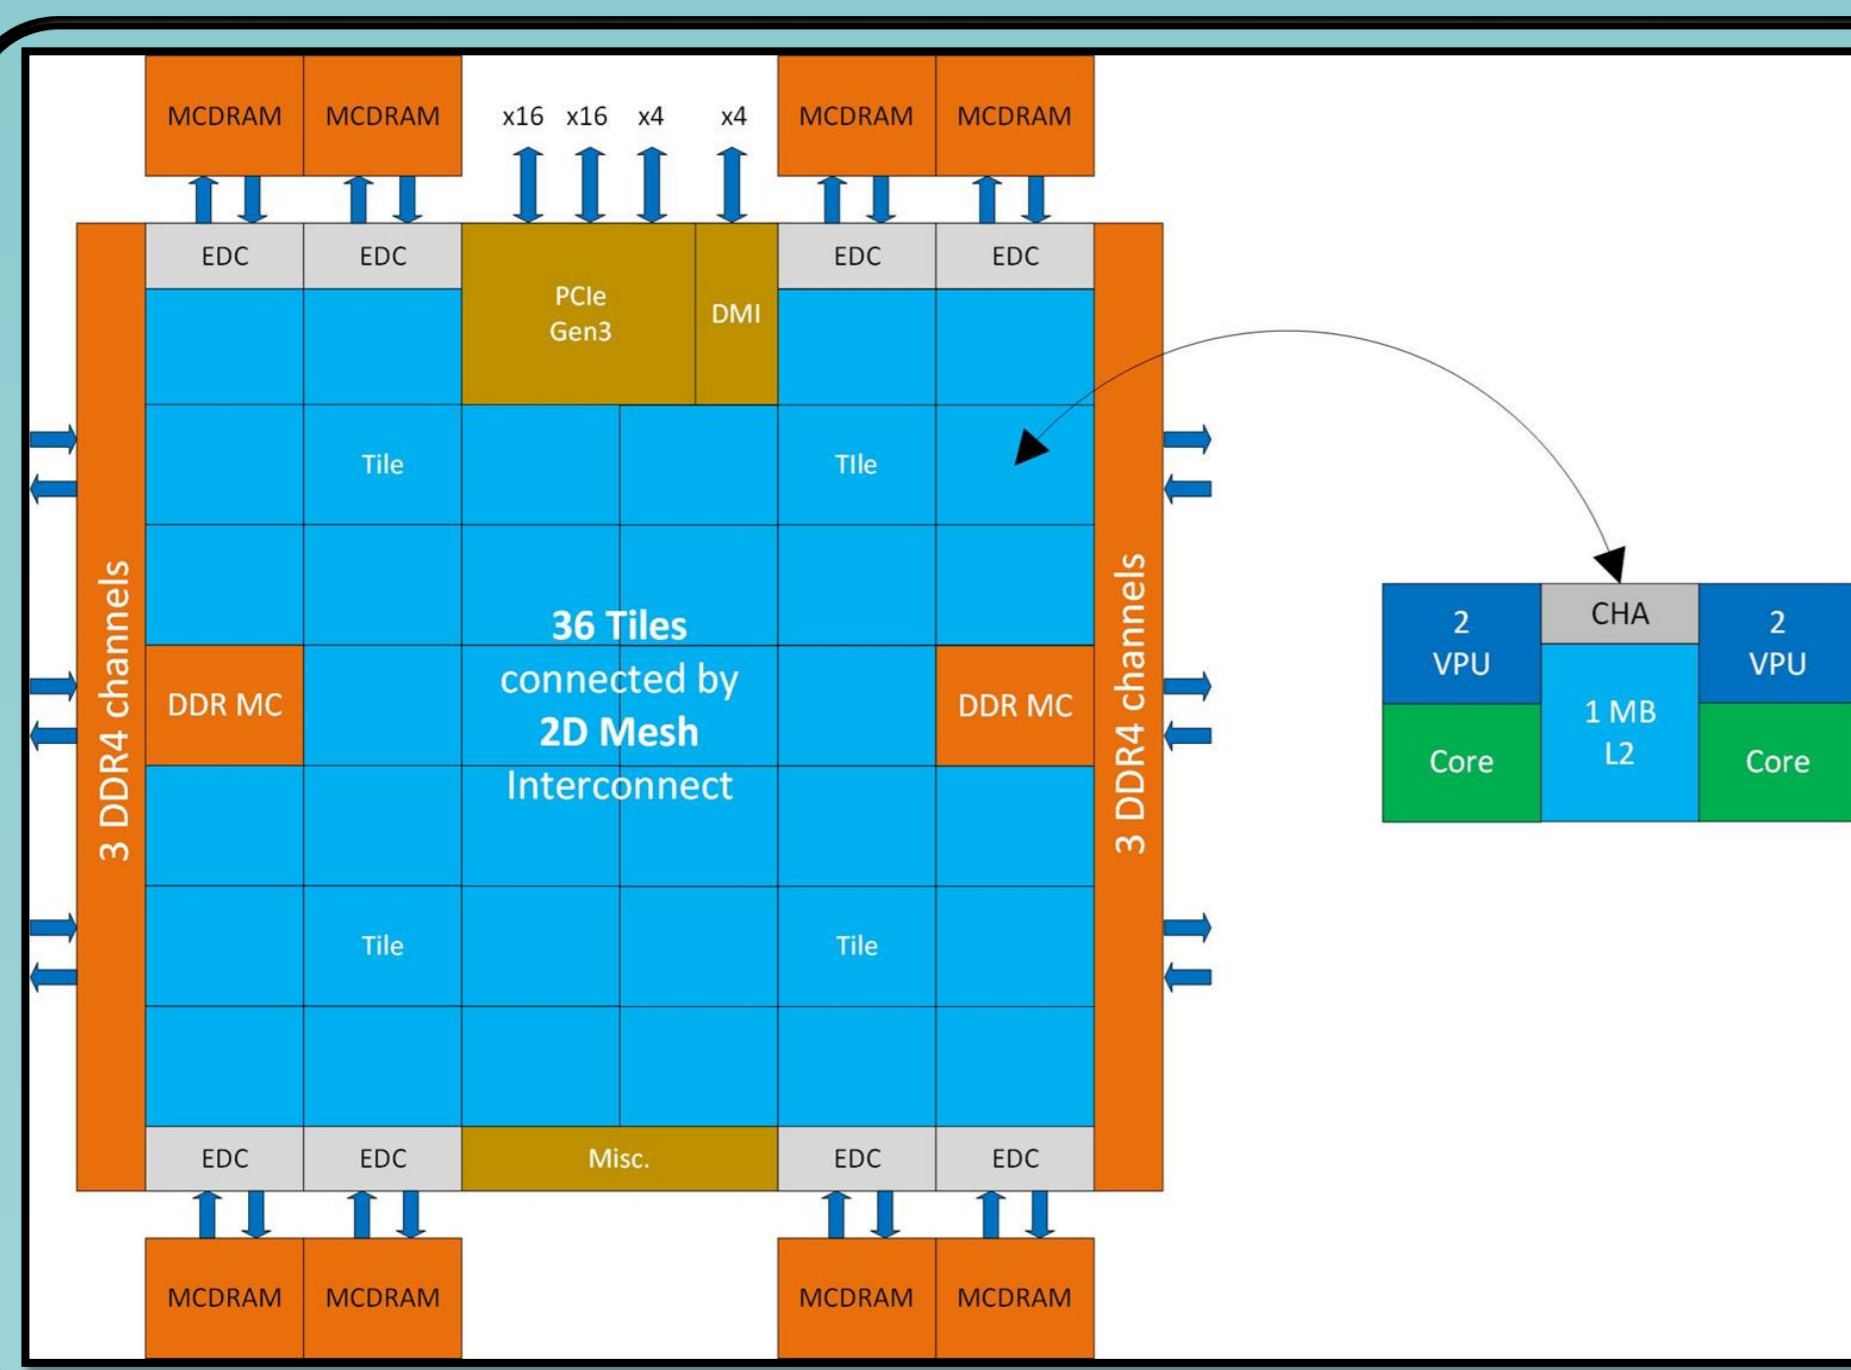
Time series motif discovery has a huge interest in many fields:

- Bioinformatics, seismology
- Entomology, energy conservation
- Traffic prediction, voice recognition
- Climate, robotics, health care

Matrix Profile (SCRIMP) is method used for detecting motifs between subsequences of a time series.

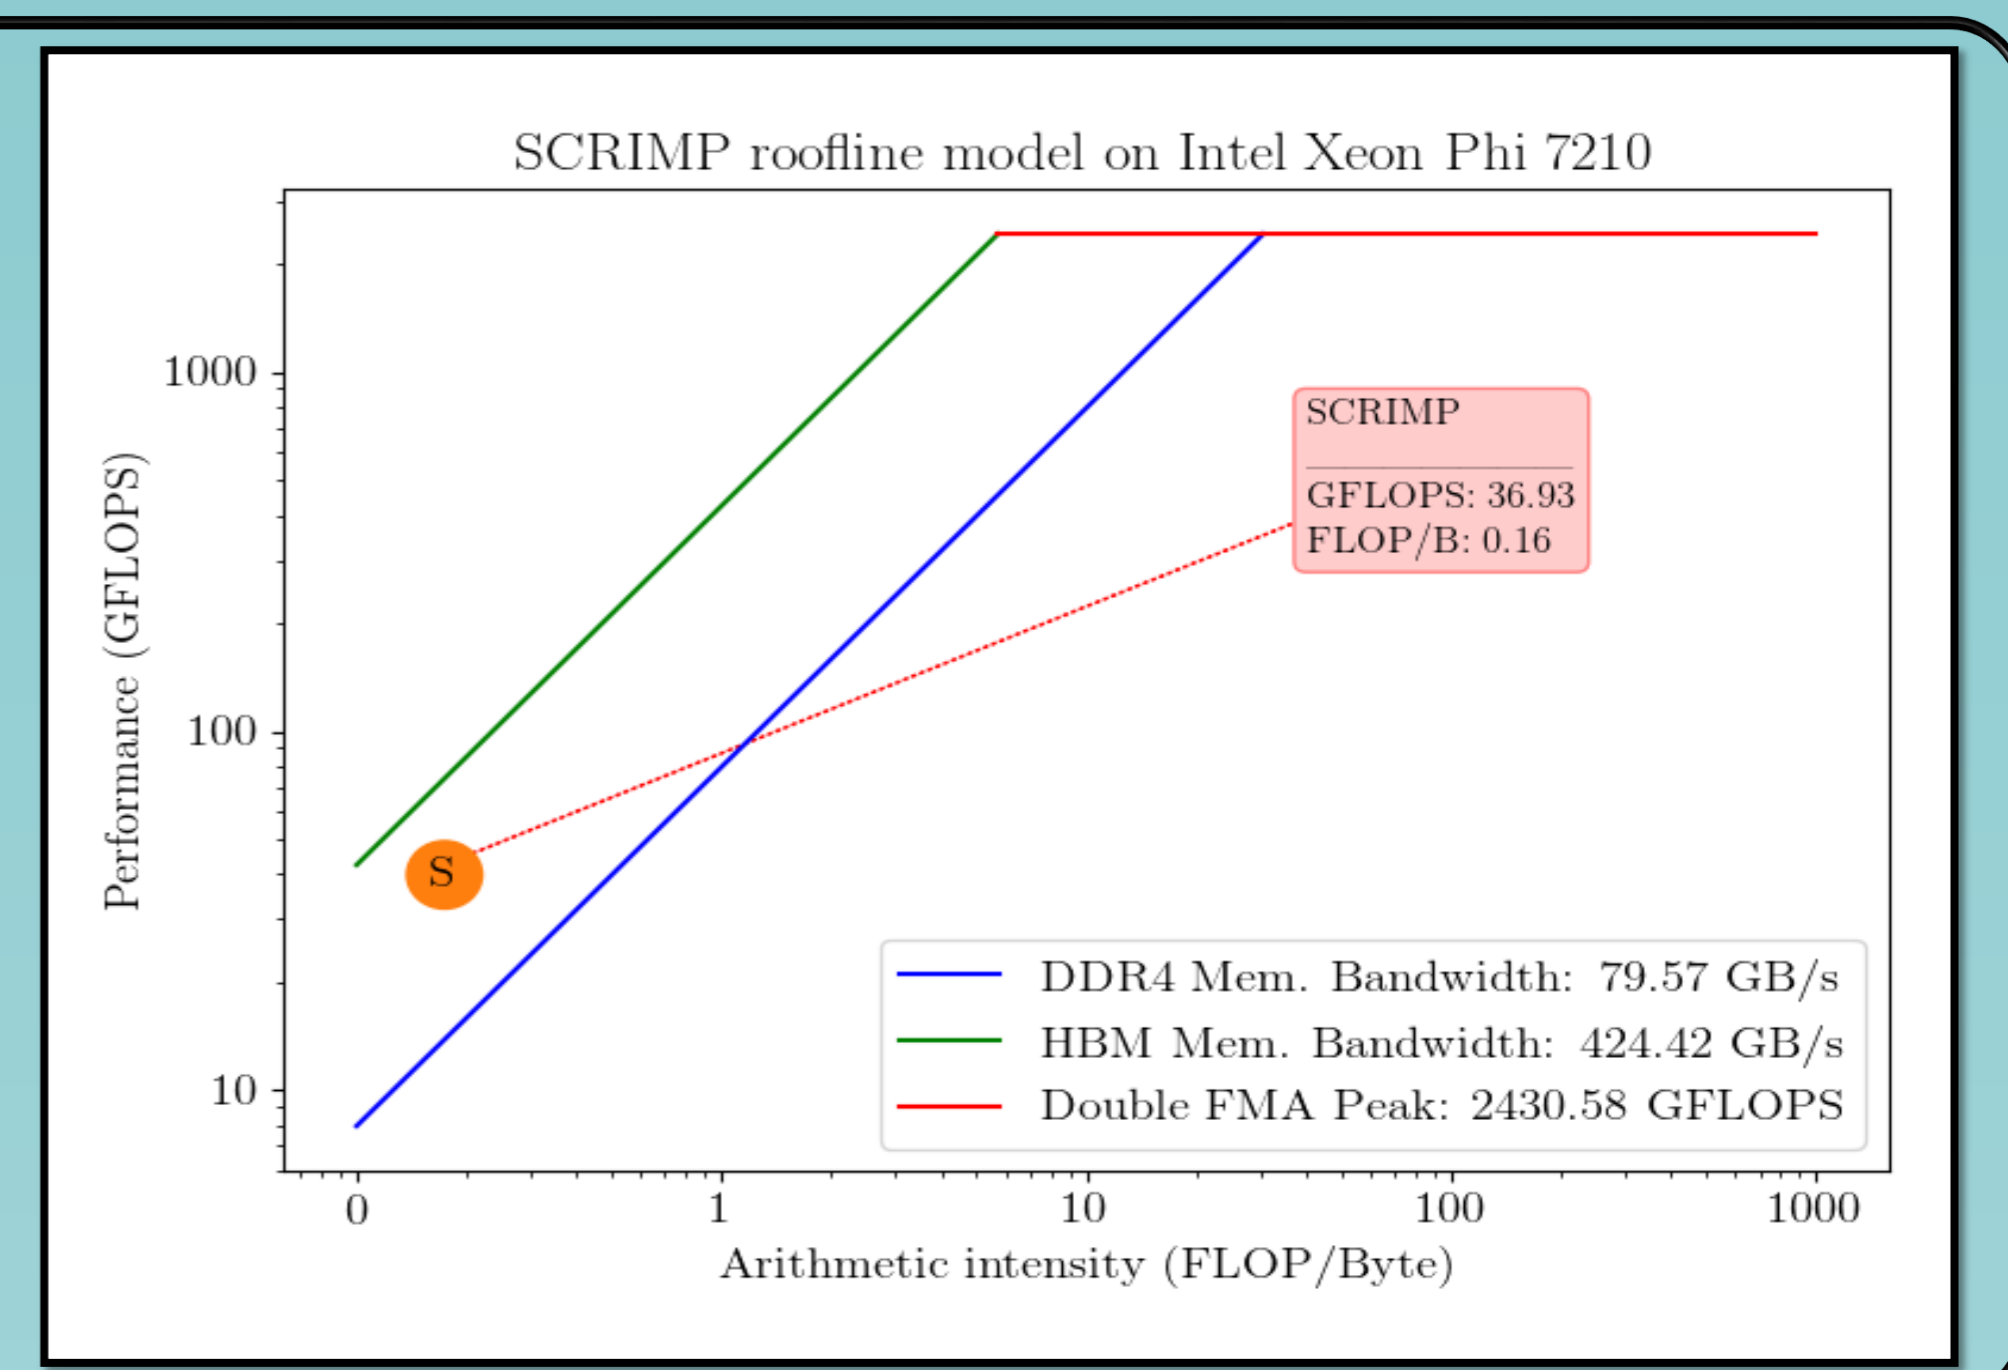


Problem and Motivation




SCRIMP is memory bound when executed in a many core machine as the Intel Xeon Phi Knights Landing.

Arithmetic intensity is very low, and cores spend most of the computation time waiting for the memory requests.




Matrix Profile




We utilize Matrix Profile as it is an exact method that allows to detect anomalies and similarities using Euclidean distances.

$$d_{i,j} = \sqrt{2m \left(1 - \frac{Q_{i,j} - m\mu_i\mu_j}{m\sigma_i\sigma_j} \right)}$$

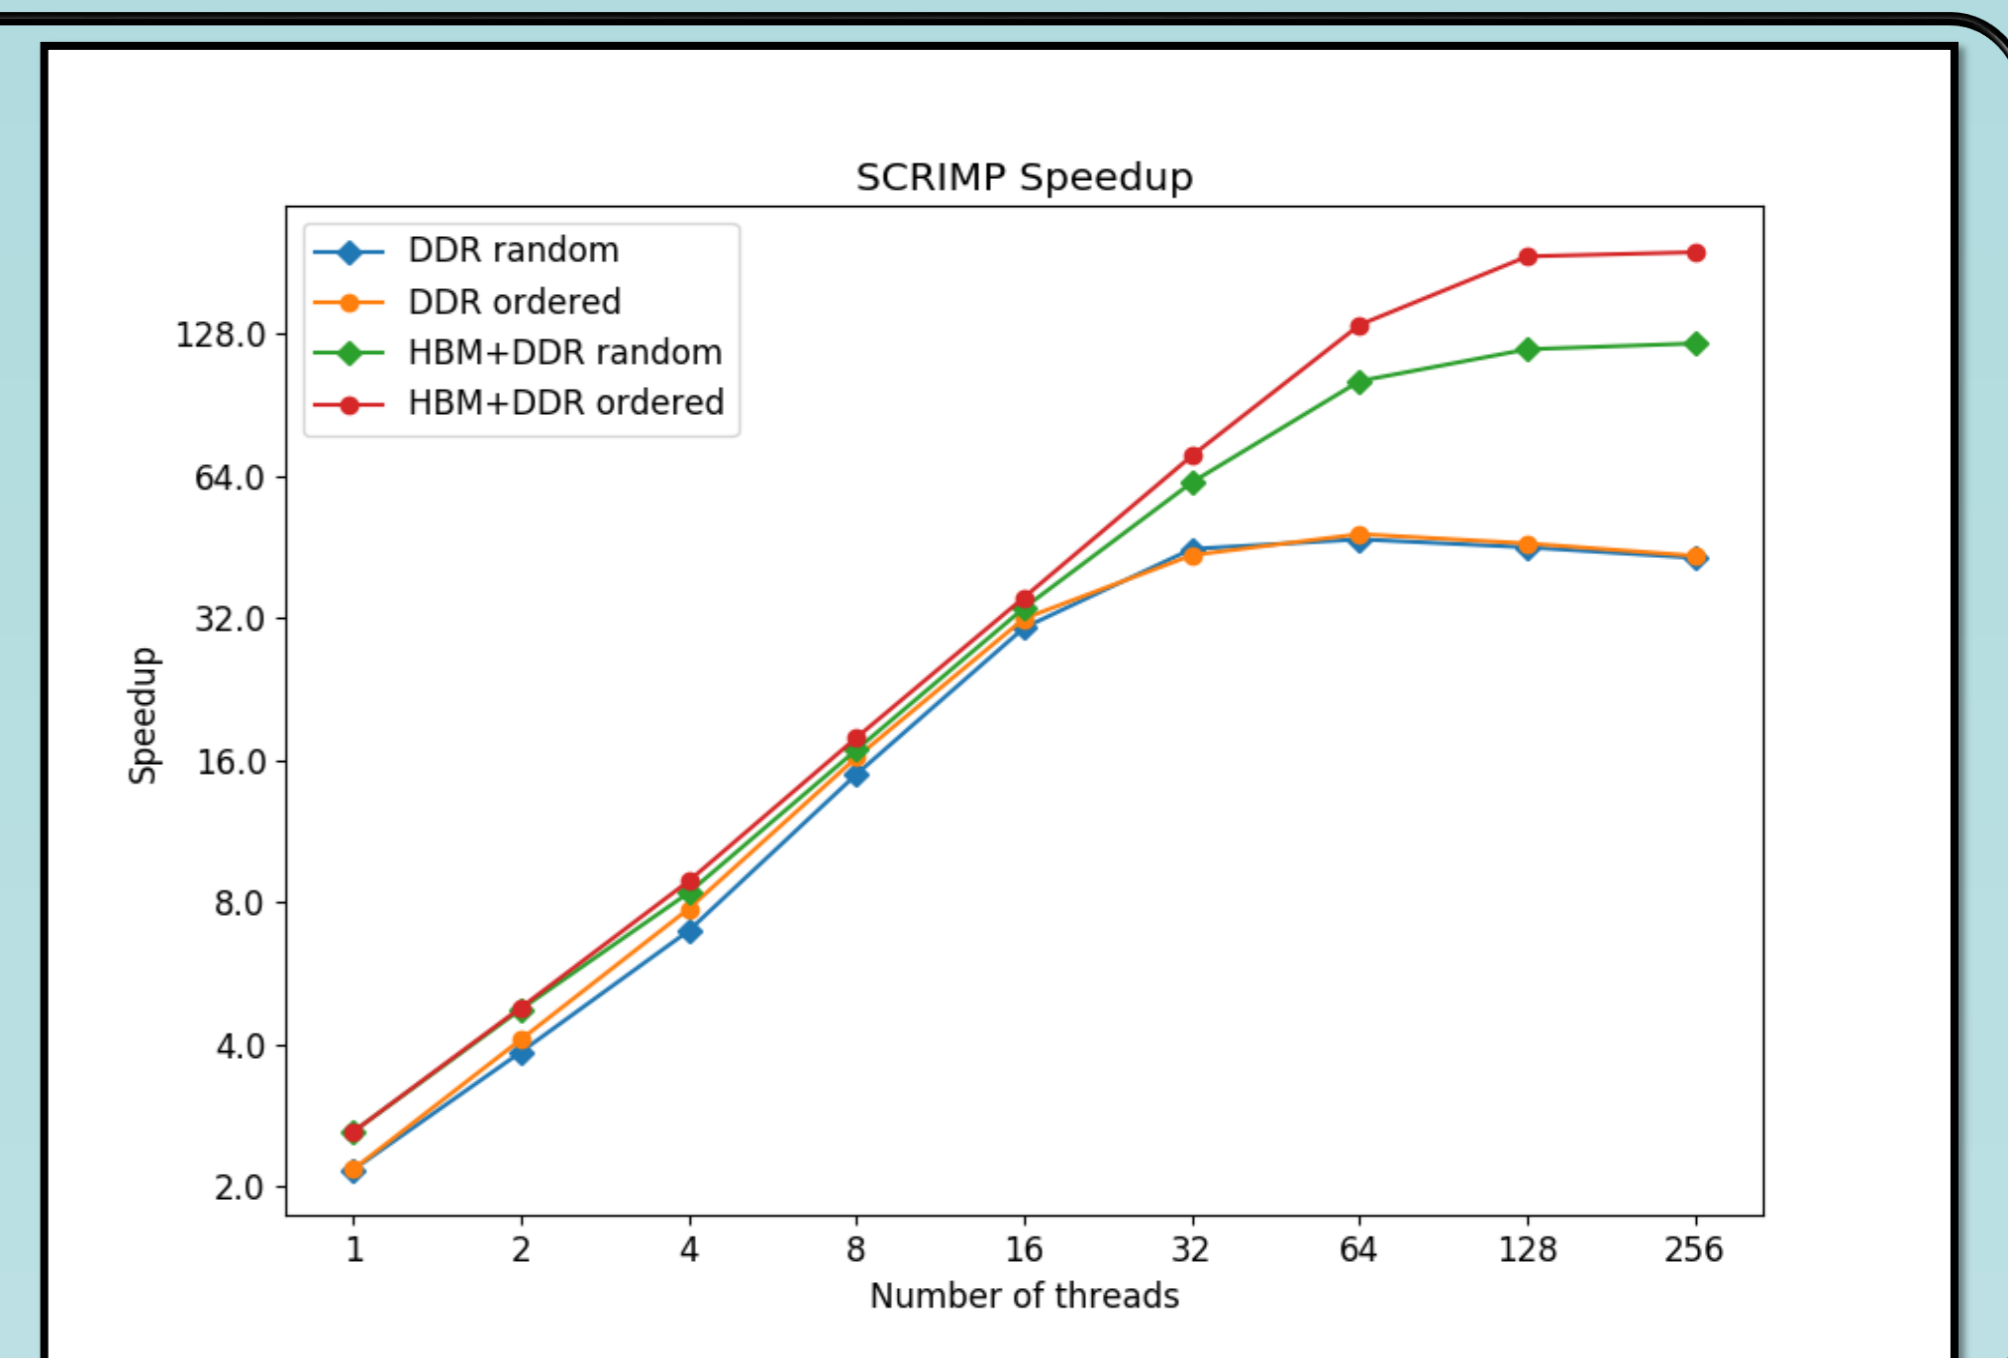


Results



We parallelize and vectorize SCRIMP implementation of Matrix Profile, that allows for sequential or random order computation.

Results show high memory bandwidth utilization, so good scalability is reached when using High Bandwidth Memory.



Conclusions

We introduce a novel implementation of the SCRIMP Matrix Profile algorithm tuned for an Intel Xeon Phi KNL architecture, provided with 3D-stacked high-bandwidth memory. Performance is improved up to 190x with respect to sequential execution (128 th. + vect). Using HBM outperforms by 5x the DDR4-only solution, proving the benefits of HBM for memory bound problems.

