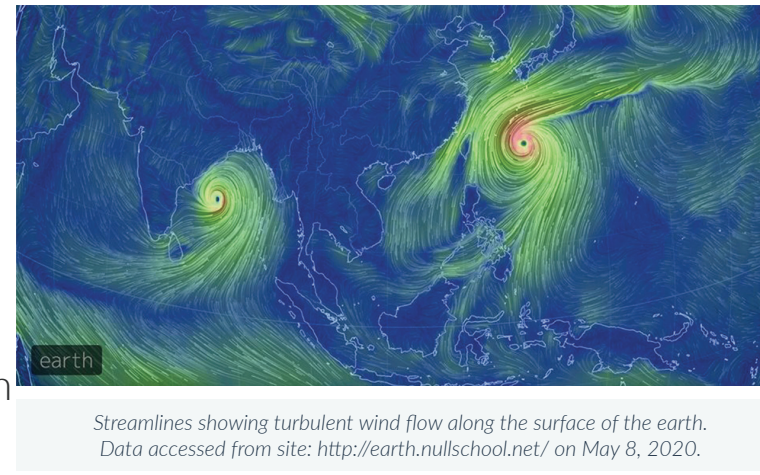


## Large-scale computations of turbulent flows

### Hi-order methods

offer high resolution capabilities that are often required for numerical simulation of complex turbulent flows, which are characterized by chaotic and highly unsteady motion of diverse spatial and temporal scales.



### Recovery-assisted discontinuous Galerkin (RADG) methods

can achieve super convergence  $2p+2$  where  $p$  is degree of the polynomial basis.

Desired accuracy is achieved with lower floating-point operations (FLOP) at high  $p$ .

For the hyperbolic system,

$$\frac{\partial u}{\partial t} + \nabla \cdot \mathcal{F} = 0,$$

$u(x,t)$  is a vector of conserved variables, and  $\mathcal{F}(u)$  are vector-valued flux functions.

DG weak-form is obtained by multiplying the governing equation with the polynomial basis  $\phi_m^k$  of at most degree  $p$ ,


and integrating over each element  $\Omega_m$  with outward normal  $\mathbf{n}$  at the element interface

$$\int_{\Omega_m} \phi_m^k \frac{\partial U_m^h}{\partial t} dx = \int_{\Omega_m} (\nabla \phi_m^k) \cdot \mathcal{F}(U_m^h) dx - \int_{\partial \Omega_m} \phi_m^k (\tilde{\mathcal{F}} \cdot \mathbf{n}^-) ds.$$

Use recovery to evaluate the common flux  $\tilde{\mathcal{F}}$ . The recovered solution  $\mathbf{f}(\mathbf{r})$  over the union of adjacent elements  $\Omega_A$  and  $\Omega_B$  is

$$\mathbf{f}(\mathbf{r}) = \sum_{n=0}^{2K-1} \Psi^n(\mathbf{r}) \hat{\mathbf{f}}^n.$$

The recovery basis is a  $2K$ -dimension polynomial, where  $\mathbf{K} = (p+1)^D$  in  $D$ -dimensions.



The recovered solution is weakly equivalent to numerical solutions  $U_A^h$  and  $U_B^h$  over the union with respect to  $k = 0, 1, \dots, K-1$  DG basis functions of both  $\Omega_A$  and  $\Omega_B$ , i.e.,

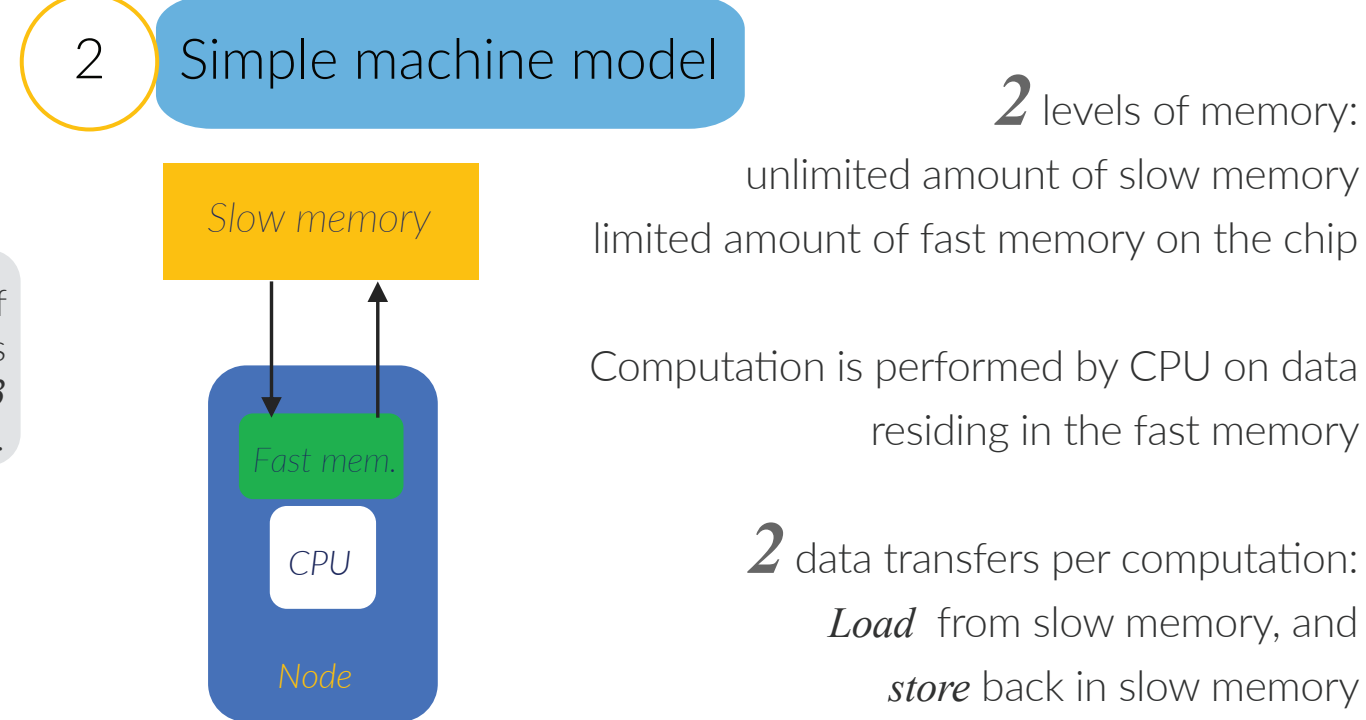
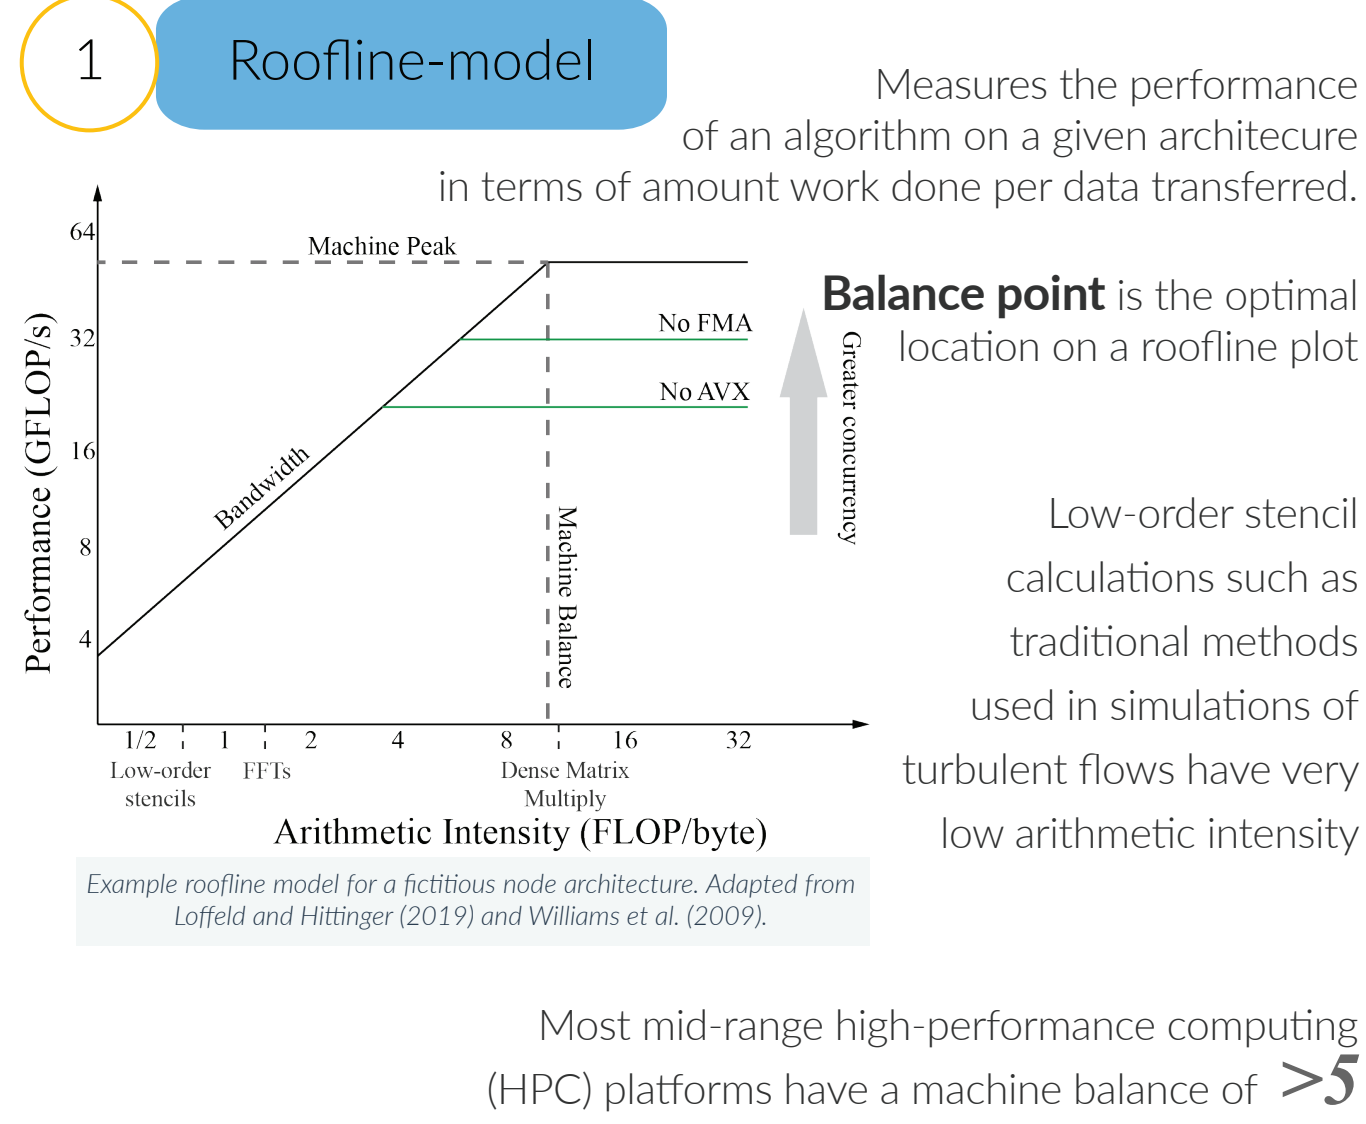
$$\int_{\Omega_A} \phi_A^k \mathbf{f} dx = \int_{\Omega_A} \phi_A^k U_A^h dx, \quad \text{and} \quad \int_{\Omega_B} \phi_B^k \mathbf{f} dx = \int_{\Omega_B} \phi_B^k U_B^h dx.$$

The recovery is exclusively used to calculate the interface quantity  $\mathbf{f}(\mathbf{0})$  located at  $x_p$ , and this entire procedure is encapsulated in a matrix-vector multiplication with recovery operator  $\mathcal{R}$ ,


$$\mathbf{f}(\mathbf{r} = 0) = \mathcal{R} \begin{bmatrix} \hat{\mathbf{U}}_A \\ \hat{\mathbf{U}}_B \end{bmatrix}.$$

Pass the recovered solutions to a Riemann solver to calculate the flux  $\tilde{\mathcal{F}} = \text{Rie}(f_A|_{x_x} + f_B|_{x_x})$

## Methods



**Assumptions** include decomposing vectorize-addition (VXA) and fused-multiply-add (FMA) operations into individual operations. Multithreading has no effect on the number of operations performed.



**Table: Floating-point operations (FLOP) count per time-step for RADG discretization of hyperbolic systems.**  $Q_v = (p+1)^{D-1}$  is the number of quadrature points along an element interface,  $K = Q_v = (p+1)^D$  is the number of degrees of freedom in an element,  $N$  is the number of elements in each direction with  $W=I$ ,  $N_v = D(N+2W+1)(N+2W)^{D-1}$  is the total number of interfaces in  $D$ -dimensions,  $f_r = DQ_v$ , and for Rusanov flux  $f_r = 4DQ_v$ .

#	Steps in the residual evaluation	FLOP count	No-cache case data transfers	Infinite-cache case with intermed. data transfers	Infinite- and finite-cache case without intermed. data transfers
1	$U_i$ and $U_p$ , at interface quadrature points	$8Q_v KN_v$	$2(R+2Q_v+Q_s)N_v$	$Q_v(N+2)^D + (R+2Q_s)N_v$	$Q_v(N^D + (N+2)^D) + 2DR(N+2)^{D-1}$
2	$\tilde{\mathcal{F}}$ , at interface quadrature points	$f_r N_v$	$3Q_s N_v$	$3Q_s N_v$	--
3	$U^h$ , at interior quadrature points	$2Q_v KN^D$	$2Q_v N^D$	$2Q_v N^D$	--
4	$\mathcal{F}(U^h)$ , at interior quadrature points	$f_r N^D$	$(2D+1)Q_v N^D$	$(2D+1)Q_v N^D$	$Q_v \times N^D$

## Arithmetic Intensity

Algorithm design that maintains trade-off between data motion, memory usage, and operations is essential to extract benefits from modern supercomputers with heterogeneous architectures

### 1 Floating-point operation (FLOP) count

For RADG schemes, the residual evaluation given by the right-hand side of DG weak-form as,

$$R_m = \int_{\Omega_m} (\nabla \phi_m^k) \cdot \mathcal{F}(U_m^h) dx - \int_{\partial \Omega_m} \phi_m^k (\tilde{\mathcal{F}} \cdot \mathbf{n}^-) ds,$$

is broken down into steps. The FLOP count of each step is evaluated, which increases with increase in  $p$  and  $N$ .


### 2 No-cache case

Degenerate case that mimics poor cache utilization

Provides a lower limit on the performance

Cache-size is small to hold data for  $I$  operation per step

Large number of data transfers reduce the arithmetic intensity to  $< 0.4$

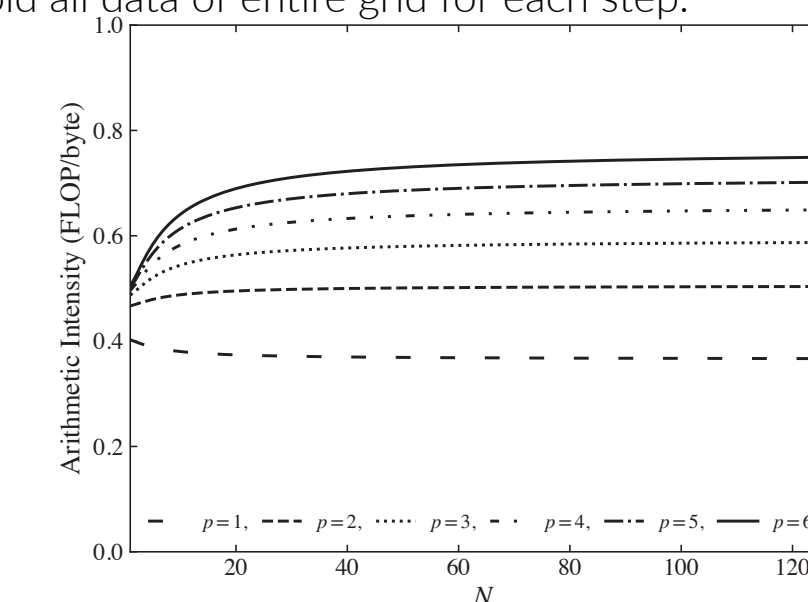


### 3 Infinite-cache case with intermediate data transfers

Cache is large enough to hold all data of entire grid for each step,

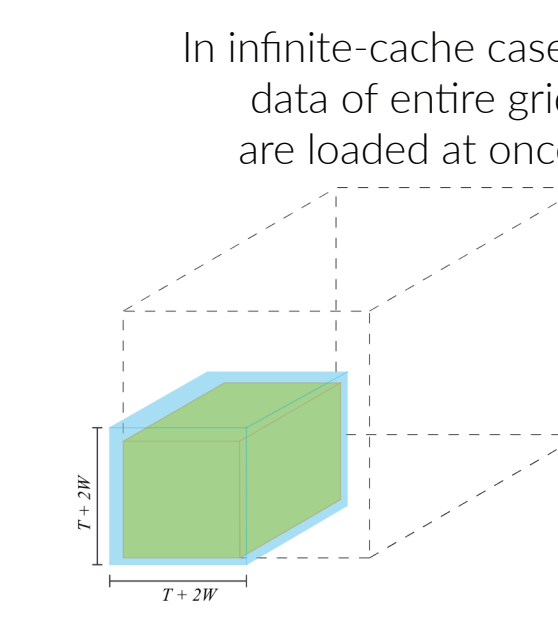
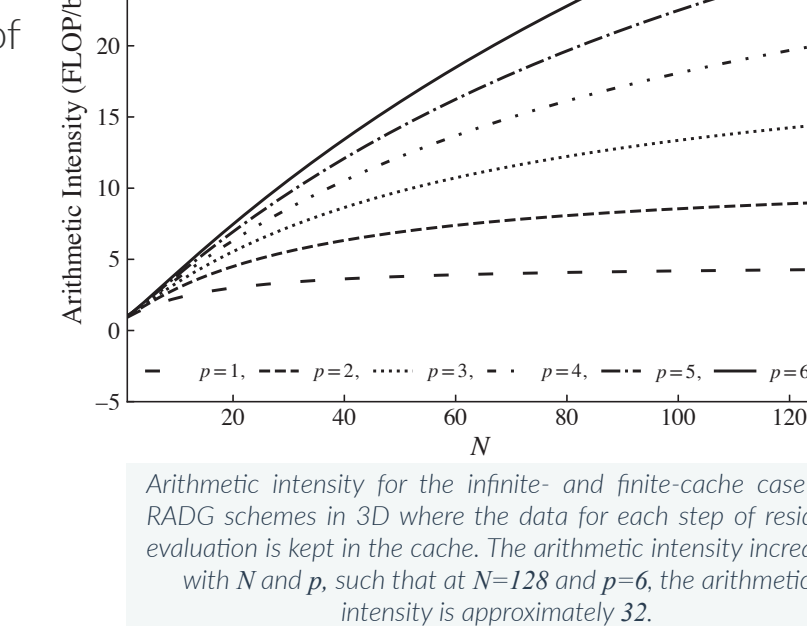
Intermediate data, at each step are transferred back in memory.

The rate of increase in FLOP is overshadowed



### 4 Infinite- & finite-cache case w/o intermediate data transfers

Improve performance by reducing data transfers such that data are loaded for first step, and final output is stored after last step with no other data transfers between steps.



RADG schemes can achieve arithmetic intensity  $> 10$  at large values of  $N$

For  $p > 2$ , RADG schemes are limited by cache size. The recovery matrix  $\mathcal{R}$  requires the most memory. In practice modern HPC platforms have  $\sim 2.5\text{MB}$  of cache per core.

### 5 Improving the cache size requirement for RADG schemes

Choose flattened rectangular tiles of size  $T \times T \times H$  where  $H$  is height of the tile.

Cache space is evaluated as  $Q_v(H+2)(T+2)^2 + Q_v(HT)^2 + Q_s \times N_f$  where  $N_f$  is total number of faces in a rectangular tile.

Cache space reduces to  $2\text{MB}$  for rectangular tile  $8x8x2$  with  $p = 3$  from  $15\text{MB}$  for cubic tile of size  $8^3$ .

### 6 Numerical Experiment

A 3D cube of length  $L = 2\pi$  is initialized with velocity fields  $(u, v, w)$  in the  $x, y$  and  $z$  directions respectively as,

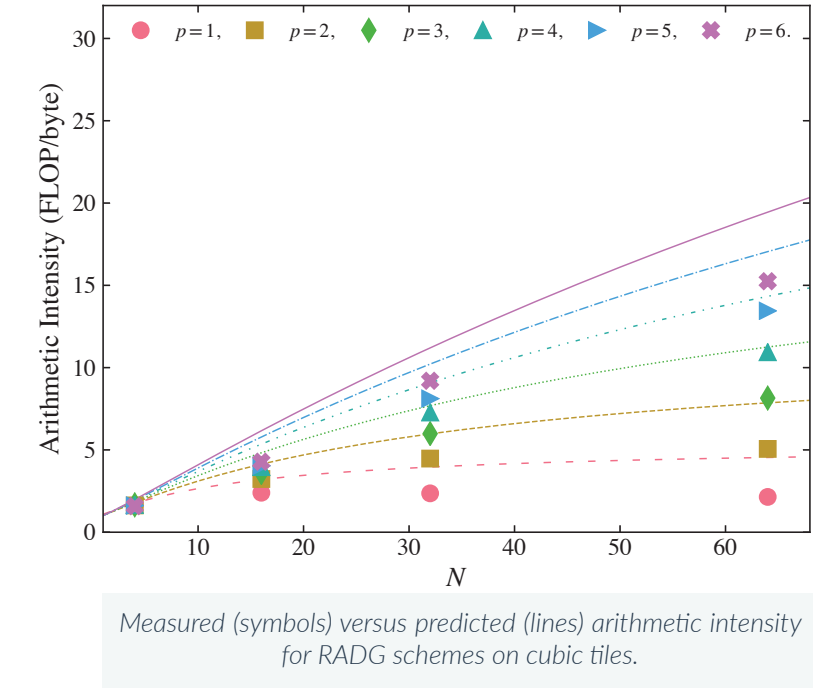
$$u = \sin(x/L) \cos(y/L) \cos(z/L), \\ v = w = 0.$$

A simple linear advection system for single component  $u$  is formulated as

$$\frac{\partial u}{\partial t} + \frac{\partial \mathcal{F}}{\partial x} = 0,$$

where  $\mathcal{F} = a.u$  with advective speed  $a = (\pi, 0, 0)^T$ .

Four cases are considered with  $N = 4, 16, 32$ , and  $64$ .



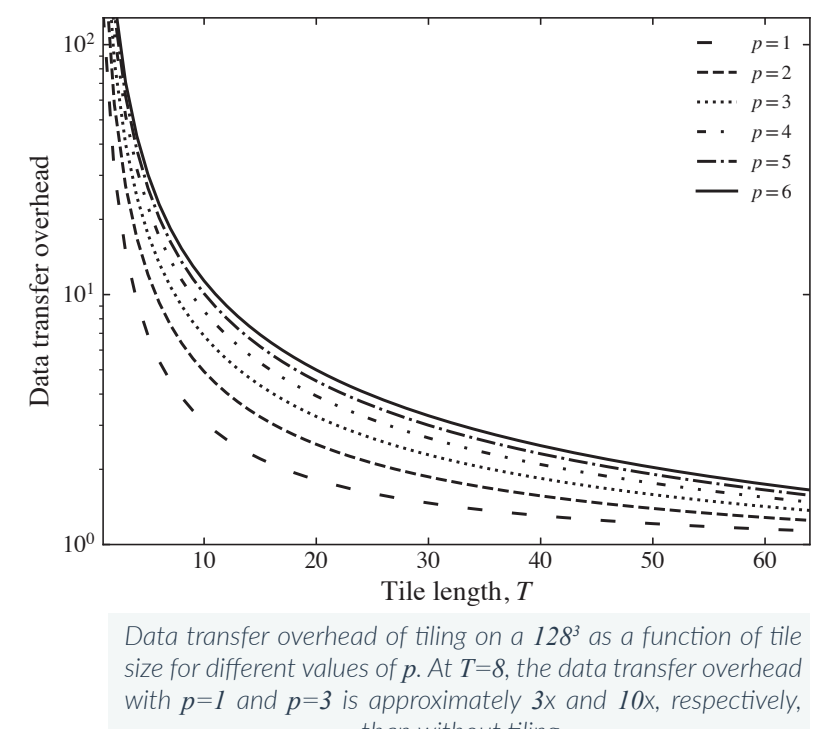
Hardware performance counters on the nodes of the Intel-Knights-Landing-based XSEDE *Stampede2* system measure the FLOP and data transfers.

The counters are read through the Intel Advisor XE. Measurements correspond to the residual evaluation for single time-step with no time-marching.

Measured arithmetic intensity is significantly lower than predicted estimates.

Reasons include, high cache size requirement for  $N > 8$  and  $p > 2$ , which leads to tiling.

For tile lengths  $T < 32$ , the overheads are high that reduces the arithmetic intensity.



## Conclusions

1. It is important to emphasize on-node performance of algorithms along with parallel scalability, and therefore, assessing arithmetic intensity must be considered as an important design criterion.
2. Three data cache strategies are investigated - lower bound no-cache case, ideal infinite-cache and a finite-cache which is more representative of reality.
3. While the RADG schemes are predicted to require high cache size that is mostly dominated by the large recovery operator, the cache size can be relaxed by using a flattened rectangular tiling.
4. RADG schemes are predicted to achieve high arithmetic intensity, which is necessary to better utilize on-node floating-point capabilities of modern HPC platforms.
5. Numerical experiment shows that measured values of arithmetic intensity for RADG are lower than the predicted estimates mostly due to data tiling overheads at large values of  $N$  and  $p$ .
6. Further investigation is needed to modify the implementation of recovery operator that reduces the storage requirement and increases the arithmetic intensity if RADG in practical applications.

## Acknowledgements

This work was supported in part by the Michigan Institute for Computational Discovery and Engineering (MICDE). Additionally, this research used the Extreme Science and Engineering Discovery Environment's Stampede2 system at the Texas Advanced Computing Center (TACC), which is supported by the NSF grant number ASC-130005.

## References

- [1] D. Brown, P. Messina, D. Keyes, et al., "Scientific grand challenges: Crosscutting technologies for computing at the exascale," U.S. Department Of Energy, Pacific Northwest National Laboratory, Tech. Rep., 2010.
- [2] J. Dongarra, J. Hittinger, J. Bell, et al., "Applied Mathematics Research for Exascale Computing," Tech. Rep., 2014.
- [3] B. Cockburn, G. E. Karniadakis, and C.-W. Shu, "The Development of Discontinuous Galerkin Methods," *Discontinuous Galerkin Methods*, pp. 77-88, 2000.
- [4] P. E. Johnson and E. Johnsen, "The Compact Gradient Recovery Discontinuous Galerkin Method for Diffusion Problems," *J. Comp. Phys.*, vol. 398, p. 108872, 2019.
- [5] P. E. Johnson, L. H. Klieu, and E. Johnsen, "Analysis of Recovery-Assisted Discontinuous Galerkin Methods for the Compressible Navier-Stokes Equations," *Journal of Computational Physics*, p. 109813, 2020.
- [6] S. Williams, A. Waterman, and D. Patterson, "Roofline: An Insightful Visual Performance Model for Multicore Architectures," *Communications of the ACM*, vol. 52, no. 4, pp. 65-76, 2009.
- [7] J. Loffeld and J. A. F. Hittinger, "On the Arithmetic Intensity of High-Order Finite-Volume Discretizations for Hyperbolic Systems of Conservation Laws," *International Journal of High Performance Computing Applications*, vol. 33, no. 1, pp. 25-52, 2019.