
StreamFEM-3D

Timothy Knight
Stanford

Timothy Barth
NASA Ames

May 8, 2003

Introduction, Previous Work

- **StreamFEM-2D was presented 6 months ago in the 2002 annual report**
 - Demonstrated good performance (sustained over half of peak)
- **Going from 2D to 3D was an important step**
 - To stress the Merrimac architecture due to the more challenging nature of 3D applications

StreamFEM-3D Overview

- StreamFEM-3D implements the Discontinuous Galerkin finite element method for systems of nonlinear conservation laws in divergence form:

$$u_t + \text{div}(\vec{f}) = 0$$


- Includes models for PDE systems of increasing complexity:
 - Scalar advection equations (1 PDE)
 - Euler equations (5 PDEs)
 - Magnetohydrodynamics (MHD) equations (8 PDEs)
- Includes piecewise polynomial approximations with increasing numbers of degrees of freedom (DOFs):
 - Constant (1), linear (4), quadratic (10), cubic (20)
- Application characteristics
 - Irregular memory accesses on unstructured 3D meshes
 - Variable arithmetic intensity

Implementation Status

- StreamFEM-3D coded in the Brook streaming language and translated to StreamC/KernelC for evaluation in the cycle accurate simulator
- Simulated a single timestep with a 3D tetrahedral element dataset
 - Dataset large enough to make boundary effects negligible

Results from a Brook functional simulation of StreamFEM-2D

→
Mach 3



Woodward-Colella 2D flow over forward-facing step

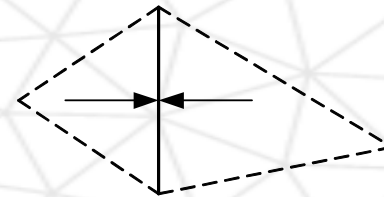
StreamFEM-3D Flow Chart

For each timestep:

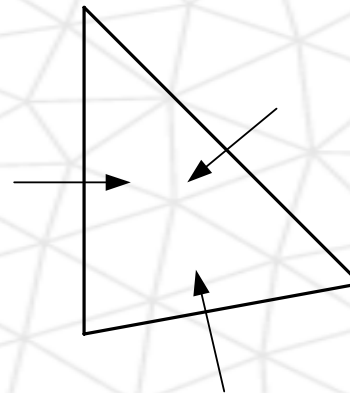
- Loop over element faces:
- Gather 2 element states
 - Compute flux terms
 - Store fluxes to memory

- Loop over elements:
- Gather 4 flux terms
 - Compute interior term and update element
 - Store updated element

- *Figures are drawn in 2D rather than 3D*

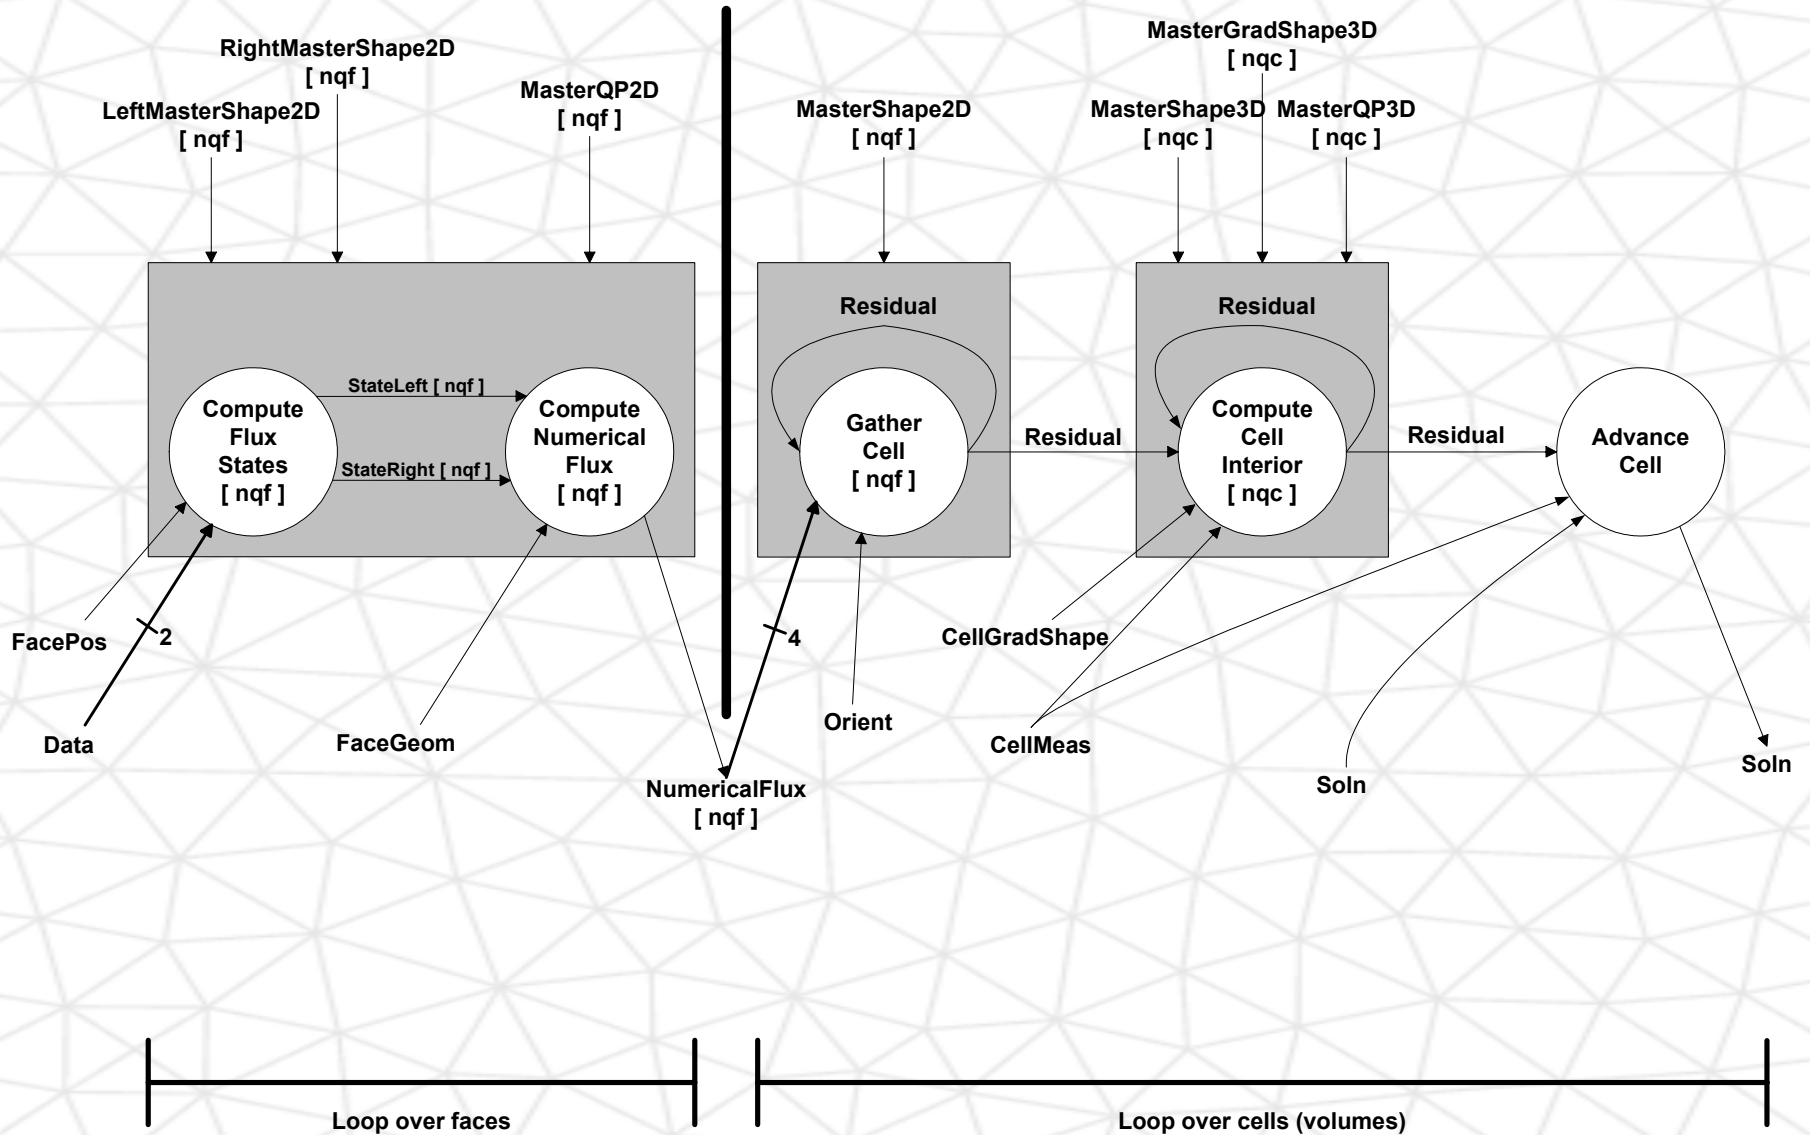


- Loop over **faces** between pairs of tetrahedra



- Loop over **tetrahedra**, gathering flux from each of the 4 faces

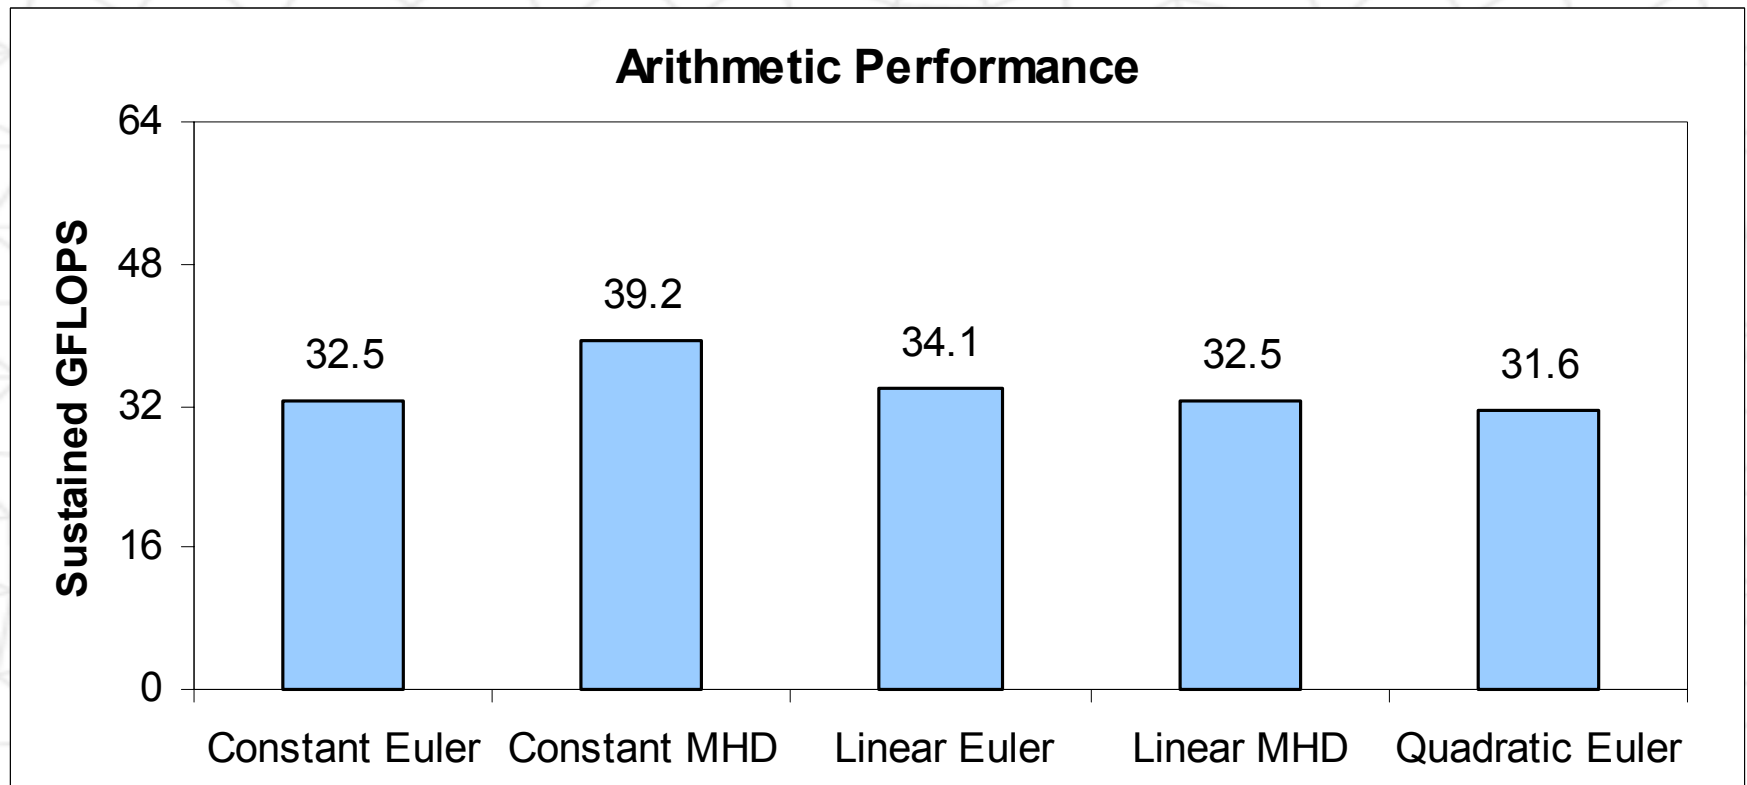
StreamFEM-3D Stream Graph



StreamFEM-3D Performance (1)

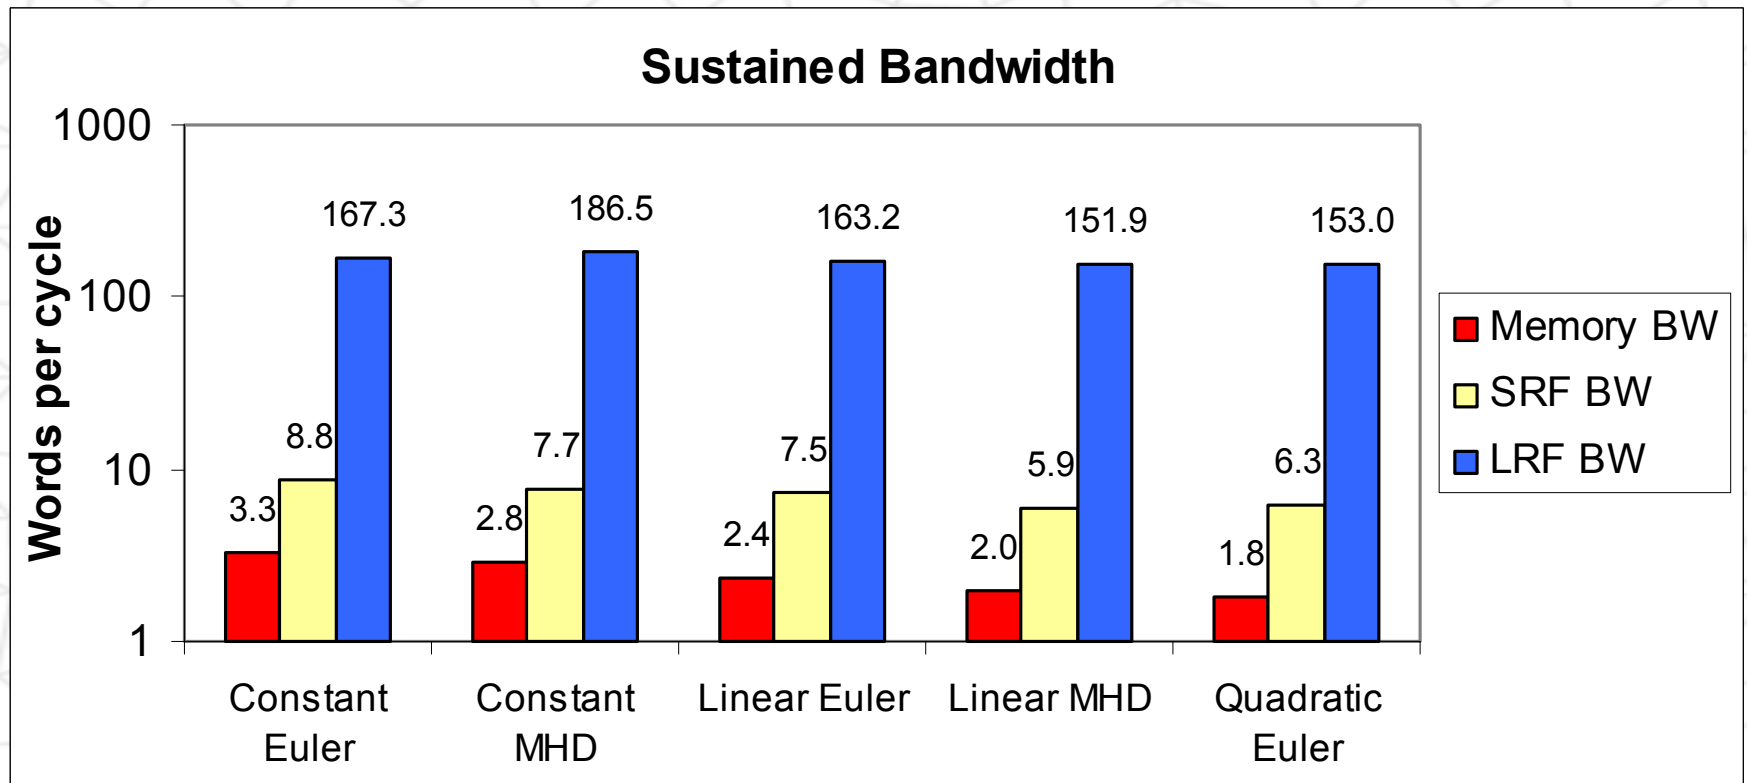
- Performance measured on cycle-accurate Merrimac node simulator
- The current node model in the simulator contains 2-input (MUL or ADD) units rather than the planned 3-input fused MULADD units
 - Peak performance of simulated node is 64 GFLOPS
 - Peak performance of planned node is 128 GFLOPS

StreamFEM-3D Performance (2)



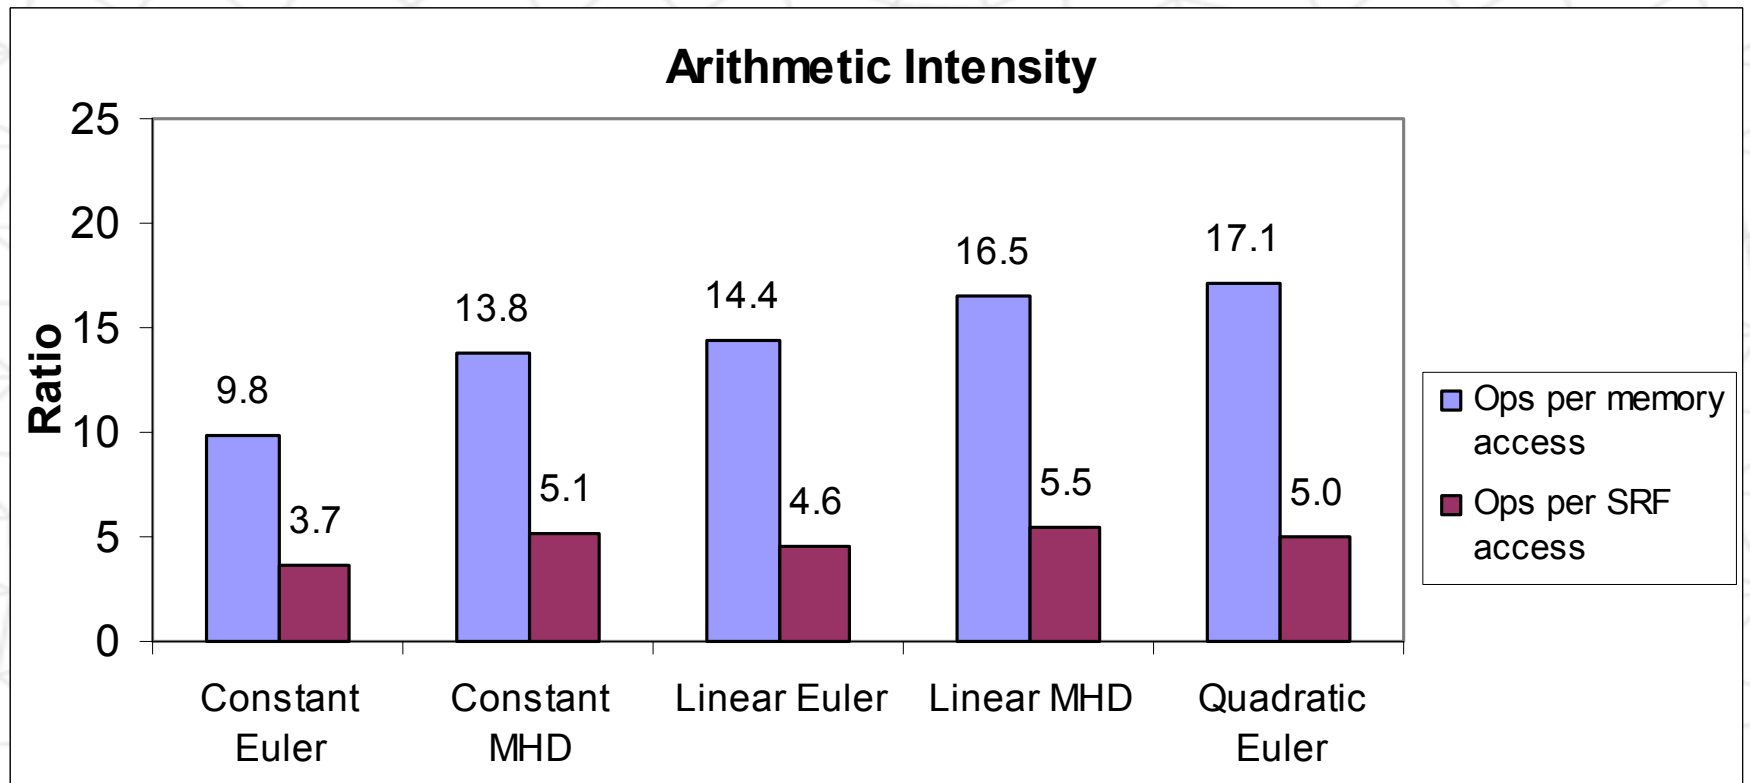
- **Where is performance lost?**
 - The first half of the application (loop over faces) is significantly memory-bandwidth limited

StreamFEM-3D Performance (3)



- 2 orders of magnitude drop in bandwidth demands going from local on-chip storage (LRFs, SRF) to global off-chip storage (memory)
 - Enables the 64 ALUs to be utilized

StreamFEM-3D Performance (4)



Conclusions

- StreamFEM-3D demonstrates that there is sufficient arithmetic intensity in irregular mesh finite element codes to sustain a large fraction of peak on Merrimac
- Going from 2D to 3D significantly increased size of both kernels and data records, but the two versions achieved similar performance (GFLOPS)
 - Shows that Merrimac is able to handle the more challenging 3D applications

Future Work

- **Still work to be done in optimizing StreamFEM-3D implementation**
 - Stream cache may be used to reduce memory bandwidth demands in memory-limited part of application
 - The data elements have 4x reuse
 - Better scheduling and optimization of strip lengths
 - Kernel splitting?
- **Compare performance with C version of application running on conventional superscalar architecture.**
 - Initially, a single-node comparison (P4, MIPS)
 - Later, a multi-node comparison with a C/OpenMP implementation on a superscalar DSM (Origin2000?)
- **Use this and other scientific applications (including StreamMD and StreamFLO) to guide architectural studies**