Architectural Implications of a Family of Irregular Applications

David O'Hallaron
Jonathan Shewchuk
Thomas Gross

School of Computer Science
Carnegie Mellon Univ.

droh@cs.cmu.edu
www.cs.cmu.edu/~droh/

The San Fernando Meshes

	sf10	sf5	sf2	sf1
nodes	7,294	30,169	378,747	2,461,694
elements	35,025	151,239	2,067,739	13,980,162
edges	44,922	190,377	2,509,064	16,684,112

High-level SMVP performance model

$$T_{SMVP} = T_{comp} + T_{comm}$$

$$(1) \quad T_c = \frac{F}{C_{\max}}\left(\frac{1-E}{E}\right)T_f$$

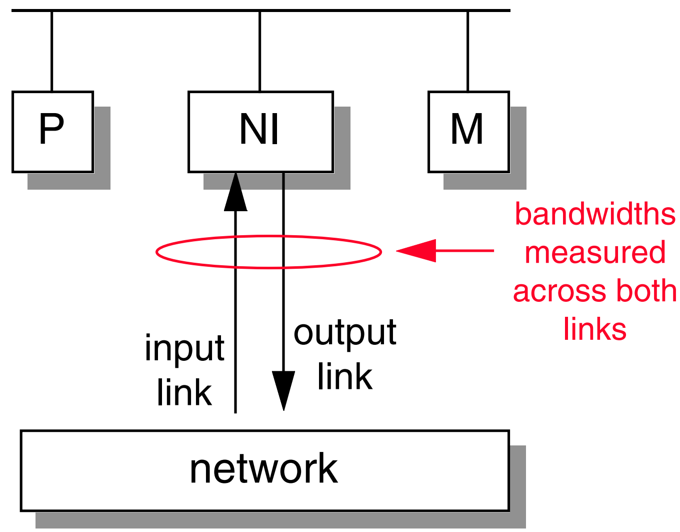

T_c : time per communicated word (inverse of sustained bandwidth)

T_f : time per flop (inverse of sustained MFLOPS)

F : number of flops per PE

$C_{\max}$: maximum number of communicated words per PE

E : efficiency $\left(E = FT_f / \left(FT_f + C_{\max}T_c\right)\right)$

Low-level SMVP performance model

$$T_{SMVP} = T_{comp} + T_{comm} \le FT_f + \left(B_{\max} T_l + C_{\max} T_w\right)$$

$$\boxed{(2) \quad T_c = \frac{T_{comm}}{C_{\max}} \le \frac{B_{\max}}{C_{\max}} T_l + T_w}$$

T_f : avg time per flop during comp phase (inv of sustained MFLOPS)

T_c : avg time per communicated word during comm phase (inv of sustained bandwidth)

T_l : block transfer latency

T_w : avg time per additional block word (inv of burst bandwidth)

$B_{\max}$: maximum blocks per PE

$C_{\max}$: maximum number of communicated words per PE

Sustained bisection bandwidth required for the sf2 SMVP

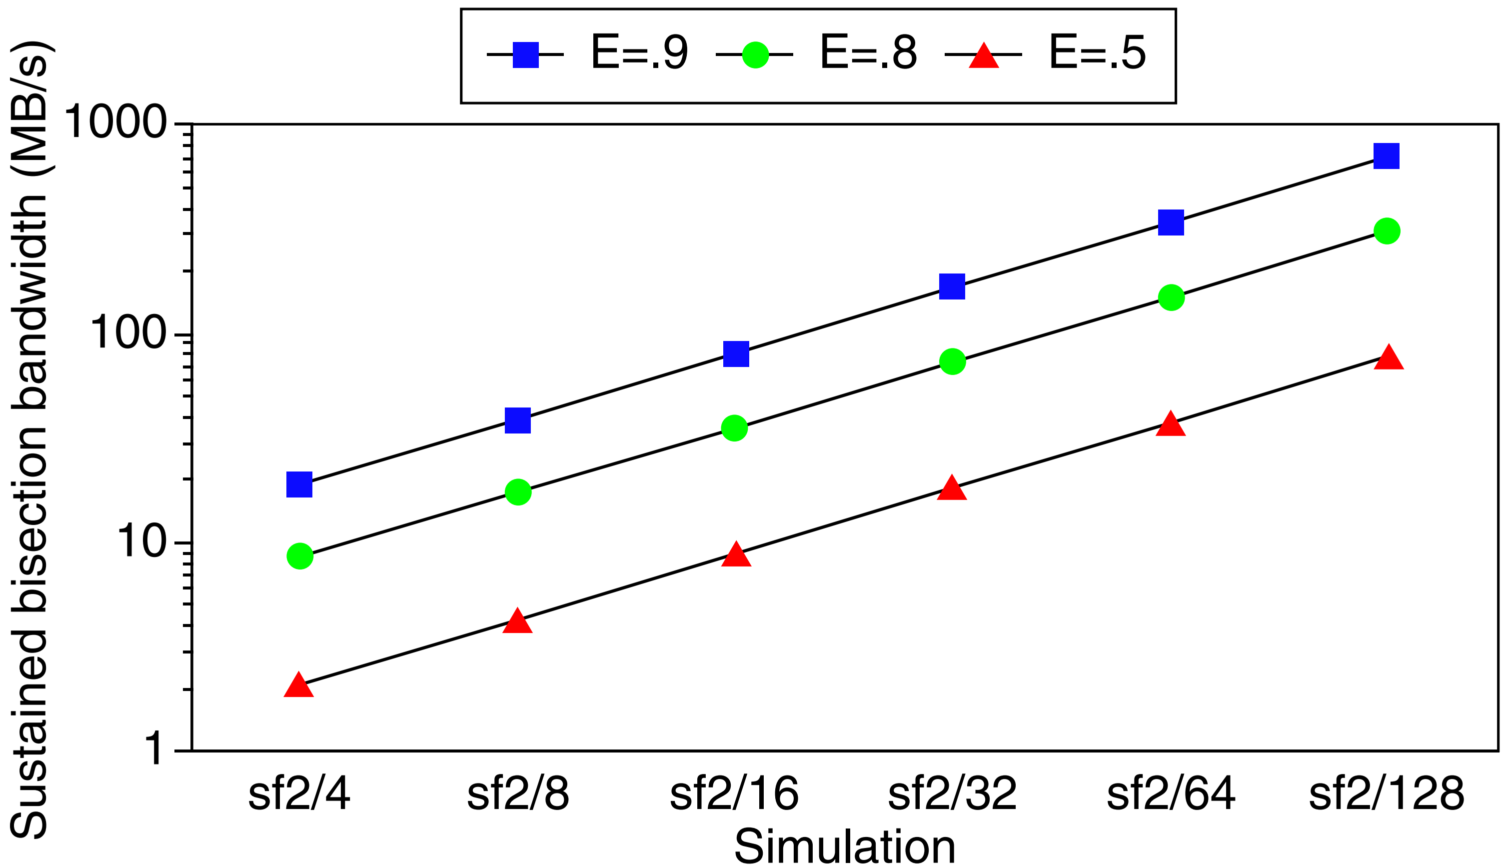

T_f = 5 ns (200 MFLOPS)

Sustained per-PE bandwidth required for the sf2 SMVP

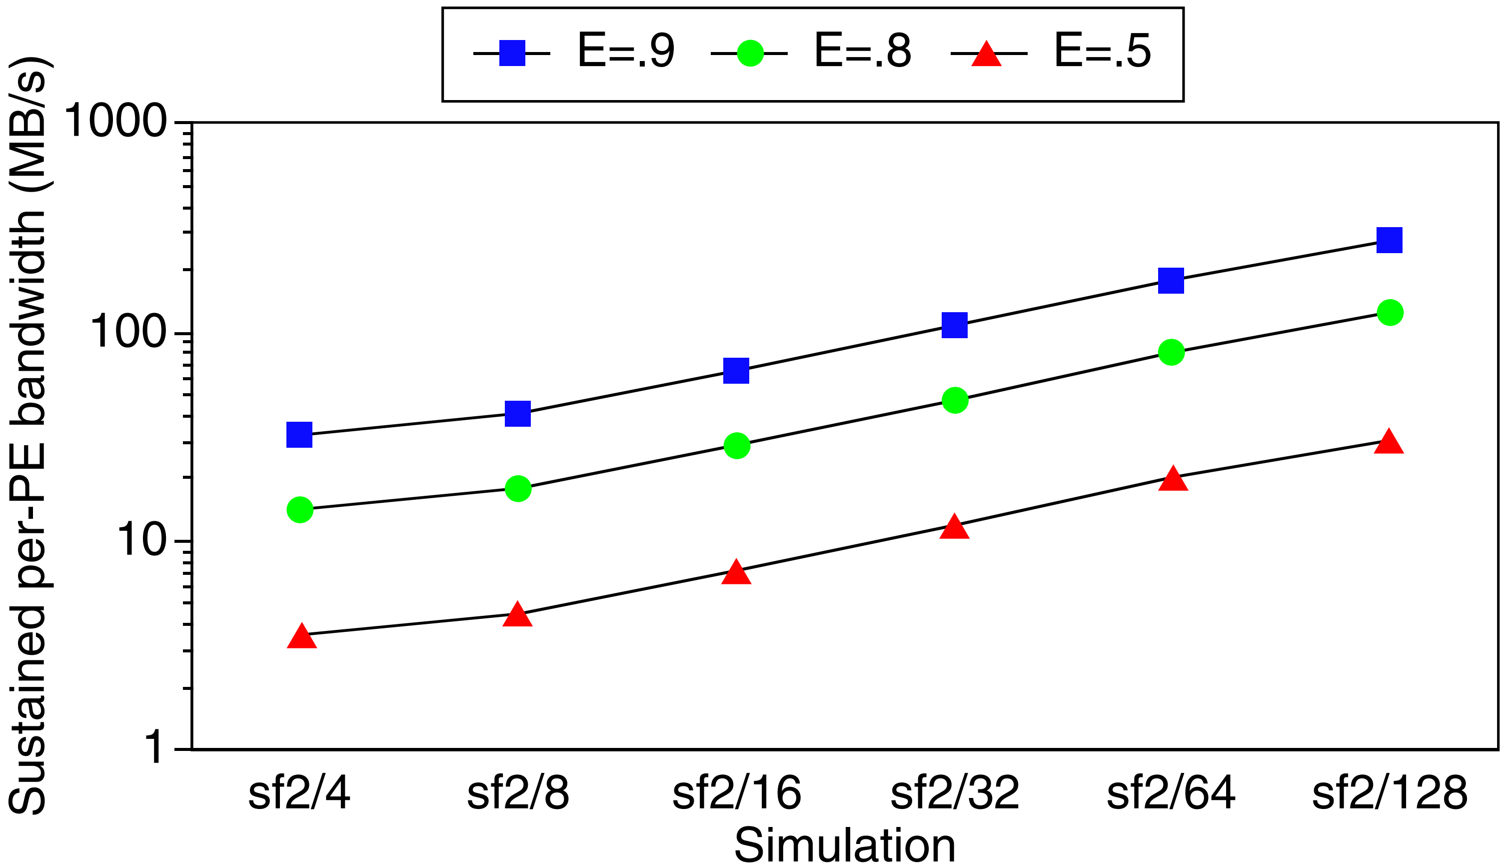

T_f = 5 ns (200 MFLOPS)

Burst bandwidth and latency tradeoffs for sf2/128

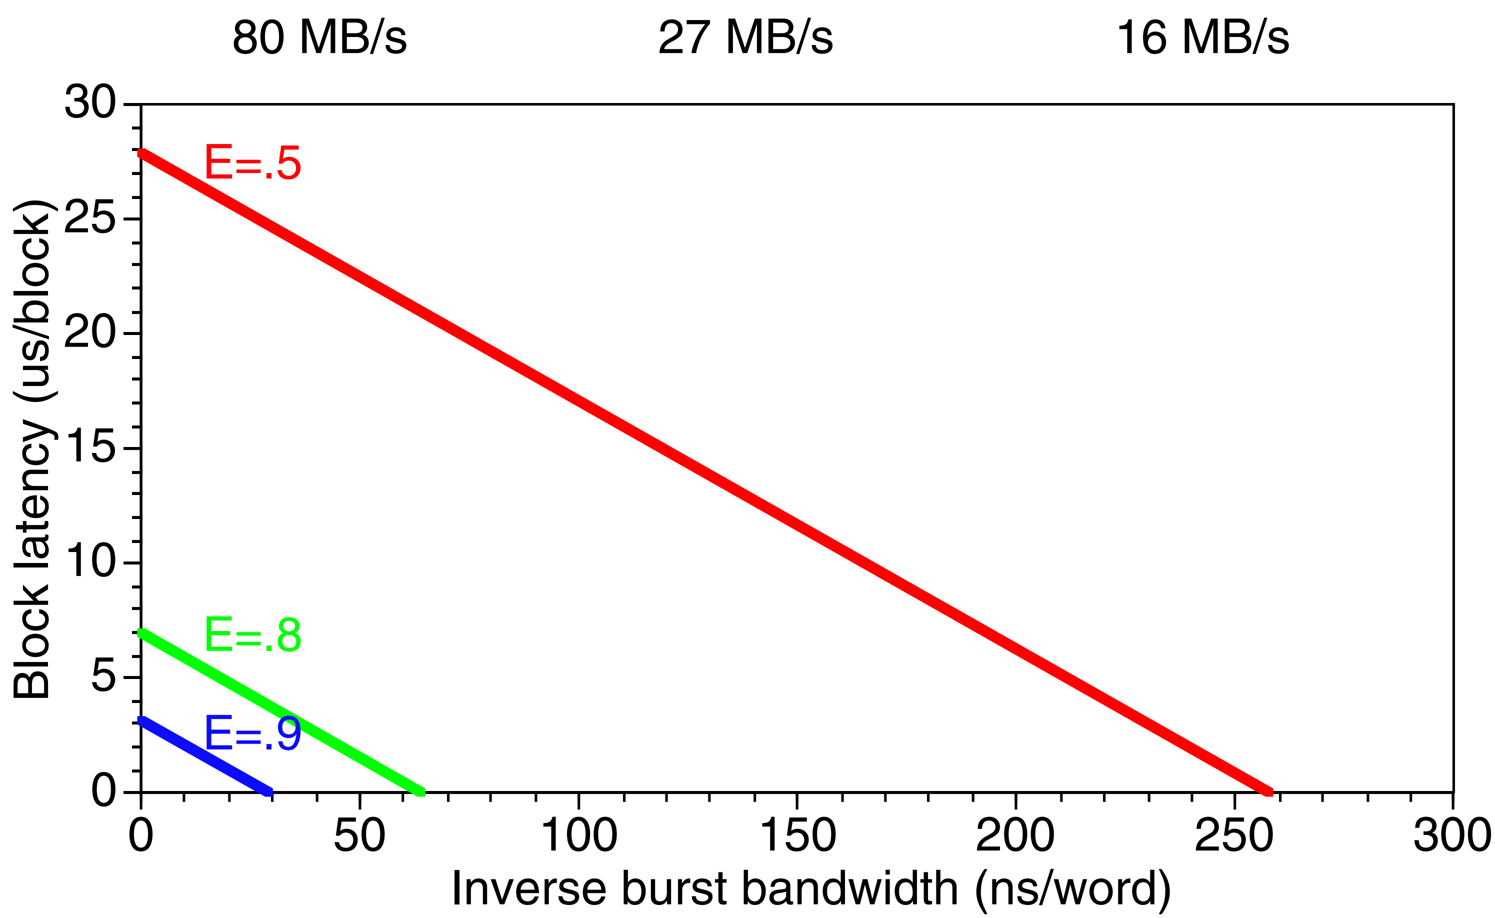

maximally aggregated blocks, Tf = 5ns (200 MFLOPS)

Burst bandwidth and latency tradeoffs for the sf2/128 SMVP

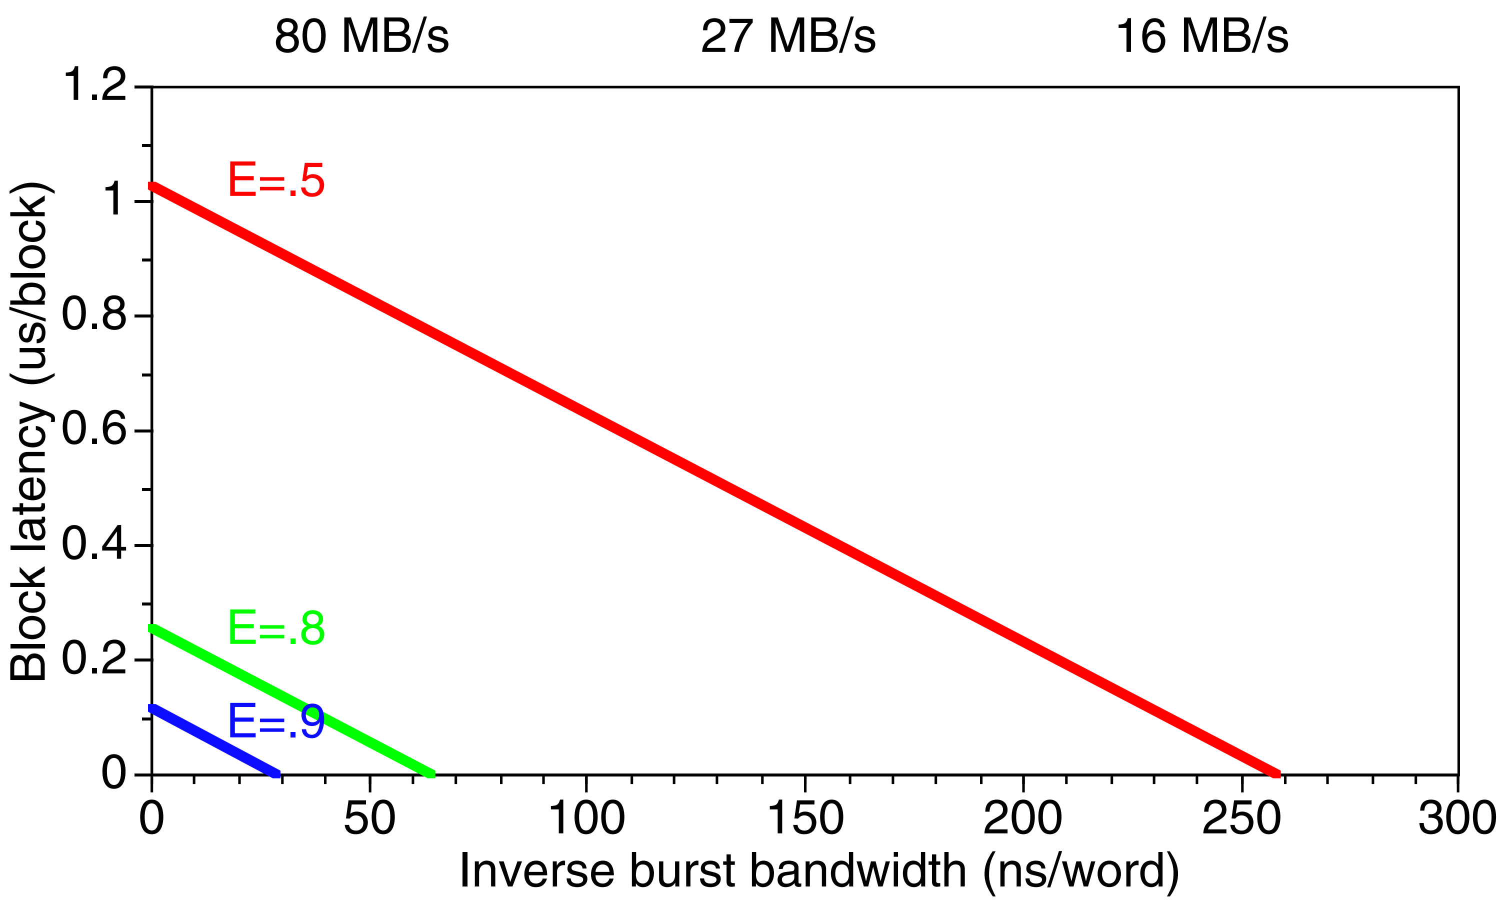

4-word block sizes, Tf = 5ns (200 MFLOPS)

Half-bandwidth and latencies for the sf2 SMVP