


Application Performance Characterization Project (APEX)


- What are the dominant code-features, which determine performance?
- Do codes with the same features have similar performance behavior?
- Can we write a flexible, general synthetic benchmark with parameters representing these features?
- How well can we approximate code performance with this synthetic benchmark?
- How few features do we need?
- How general is this approach?

Memory Access Probe: Apex-MAP Version 1

- Focus on **spatial** and **temporal** data locality.
- Parameters:
 - Data set size: M
 - Spatial Locality (L):
 - Blocked access to L contiguous data elements.
 - Temporal Locality (α):
 - More frequent access to memory locations by using **non-uniform random** addresses of blocks
 - Starting addresses distributed according to a **power law**.
 - Characterize temporal locality with the exponent α in $[0,1]$
 - $\alpha=1$: Uniform random access.
 - $\alpha=0$: Access to a single vector only.



Parallel Performance Surfaces 256 Processors – MPI




Using Performance Models to Analyze System Responses

Linear timing model for two memory levels

$$T = f \cdot P(c/m) \cdot (a+b \cdot (L-1)) + (1-P(c/m)) \cdot (c+d \cdot (L-1))$$

J/L

Residual errors show particularities of systems very well!



APEX-Map Parameters to Characterize Kernels

	HPCC Stream	HPCC GUPS	NAS CG	NBODY	MM Stride	MM Vector	MM-DGEMM	Transpose
Pattern	Random	Random	Random	Random	Stride	Random	Random	Stride
Temp Locality	1	1	0.01-0.03	1		1.0	0.001	
Spatial Locality	N	1	1	4		K	64	
Mem	N	Table	Matrix+Vector	Mem	Matrix b	Matrix b	Matrix b	Matrix
Reg Pressure	1	1	1	1	1	1	1	1
Comp. Intensity	1	1	1	16	1	2	12	1
No. of Data Access	1	Nupdate	Matrix / Vector	1		M	M	
LHS	Array	Array	Scalar	Scalar		Array	Scalar	
Loop Structure	Nested	Nested	Fused	Nested		Nested	Nested	
Stride					N			N

APEX-Map Version 3

- Insert a parameterized compute kernel
- Two new computational parameters:
 - Computational Intensity CI
 - Register Pressure R
- Left Hand Side can have writes:
 - Scalar (as Version 1 and 2)
 - Array (but only stride 1, direct)
- List of all performance parameters and feature:
 - PARALLEL: NY (V1 only)
 - PARALLEL LANGUAGE: MPI / SHMEM / UPC / CAF (V1 only)
 - ACCESS PATTERN: RANDOM / STRIDE
 - SPATIAL LOCALITY (L): [1, M] Default: {1, 4, 16, ..., 65536}
 - INDEPENDENCE LENGTH (I): [1, X] Default: 1024
 - TEMPORAL LOCALITY (a): [0,1] Default: {1.0 0.5 0.25 0.1 0.05 ... 0.0025 0.001}
 - MEMORY SIZE (M): Default: 67,108,864 Words = 512MB / process
 - REGISTER PRESSURE (R): [1, X] Default: 1
 - COMPUTATIONAL INTENSITY (CI): [1, X] Default: 1
 - ACCESS MODE: FUSED / NESTED
 - RESULTS: SCALAR / ARRAY (left hand side of statement)

```

•Code Example for: Random, Nested, Scalar, R=3, CI = 3
for (i = 0; i < times; i++) {
  index-length = B / L;
  initIndexArray(index-length);
  CLOCK(time1);
  for (j = 0; j < index-length; j++) {
    for (k = 0; k < L; k++) {
      W0 = W0+c0*(data0[ind[j]+k]+c0*(data1[ind[j]+k]+c0*(data2[ind[j]+k])));
      W1 = W1+c1*(data1[ind[j]+k]+c1*(data2[ind[j]+k]+c1*(data0[ind[j]+k))));
      W2 = W2+c2*(data2[ind[j]+k]+c2*(data0[ind[j]+k]+c2*(data1[ind[j]+k))));
    }
  }
  CLOCK(time2); }
    
```



This work was sponsored by the Department of Energy Office of Science.
Contact: Erich Strohmaier, estrohmaier@lbl.gov

NAS CG Approximation

```

DO j = 1, lastrow-firstrow+1
sum = 0.d0
DO k = rowstr(j), rowstr(j+1)-1
sum = sum + a(k)*p(colidx(k))
ENDDO
w(j) = sum
ENDDO
    
```


Apex-Map	Stream
Pattern	Random
Temp Locality	???
Spatial Locality	1
Mem	E + N
Reg. Pressure	1
Comp. Intensity	1
No. of Data Access	E/N (C: 240)
LHS	SCALAR
Loop Structure	FUSED



The results indicate that the performance of CG for different data sets can be simulated with Apex-Map using one stream with temporal locality ranging from 0.03 - 0.01


GUPS Approximation

Apex-Map	Stream
Pattern	Random
Temp Locality	1
Spatial Locality	1
Mem	TableSize
Reg. Pressure	1
Comp. Intensity	1
Concurrency	NUPLICATE
Results	ARRAY
Access Mode	NESTED




DGEMM Approximation

Apex-Map	Stream
Pattern	Random
Temp Locality	0.001
Spatial Locality	64
Mem	Matrix b
Reg. Pressure	1
Comp. Intensity	12
Concurrency	M
Results	SCALAR
Access Mode	NESTED



nBody Approximation

Apex-Map	Stream
Pattern	Random
Temp Locality	1
Spatial Locality	4
Mem	Total MEM
Reg. Pressure	1
Comp. Intensity	15
Concurrency	1
Results	SCALAR
Access Mode	NESTED



•FDIV, and FSQR! require calibration for CI values
• Use No. Loops/second executed as performance metric