



Modeling and tuning of SpMV and a lattice QCD kernel on the A64FX

Christie Alappat, Jan Laukemann, Thomas Gruber, Georg Hager, Gerhard Wellein, Nils Meyer, Tilo Wettig

Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany Erlangen National High Performance Computing Center (NHR@FAU) University of Regensburg, Germany



Paper

C. Alappat, N. Meyer, J. Laukemann, T. Gruber,
G. Hager, G. Wellein, and T. Wettig: *ECM modeling and performance tuning of SpMV and Lattice QCD on A64FX.*Concurrency and Computation: Practice and
Experience, e6512 (2021).
Available with Open Access.

DOI: <u>10.1002/cpe.6512</u>

SPECIAL ISSUE PAPER

ECM modeling and performance tuning of SpMV and Lattice QCD on A64FX

 $\label{eq:christie} Christie Alappat^{*1} \mid Nils \ Meyer^2 \mid Jan \ Laukemann^1 \mid Thomas \ Gruber^1 \mid Georg \ Hager^1 \mid Gerhard \ Wellein^1 \mid Tilo \ Wettig^2$

 ¹Erlangen National High Performance Computing Center, Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen, Germany
 ²Department of Physics, University of Regensburg, Regensburg, Germany

Correspondence

*Christie Alappat, Erlangen National High Performance Computing Center, Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen, Germany. Email: christie.alappat@fau.de

Summary

The A64FX CPU is arguably the most powerful Arm-based processor design to date. Although it is a traditional cache-based multicore processor, its peak performance and memory bandwidth rival accelerator devices. A good understanding of its performance features is of paramount importance for developers who wish to leverage its full potential. We present an architectural analysis of the A64FX used in the Fujitsu FX1000 supercomputer at a level of detail that allows for the construction of Execution-Cache-Memory (ECM) performance models for steady-state loops. In the process we identify architectural peculiarities that point to viable generic optimization strategies. After validating the model using simple streaming loops we apply the insight gained to sparse matrix-vector multiplication (SpMV) and the domain wall (DW) kernel from quantum chromodynamics (QCD). For SpMV we show why the CRS matrix storage format is not a good practical choice on this architecture and how the SELL-C- σ format can achieve bandwidth saturation. For the DW kernel we provide a cache-reuse analysis and show how an appropriate choice of data layout for complex arrays can realize memory-bandwidth saturation in this case as well. A comparison with state-of-the-art high-end Intel Cascade Lake AP and Nvidia V100 systems puts the capabilities of the A64FX into perspective. We also explore the potential for power optimizations using the tuning knobs provided by the Fugaku system, achieving energy savings of about 31% for SpMV and 18% for DW.

KEYWORDS:

ECM model, A64FX, sparse matrix-vector multiplication, lattice quantum chromodynamics





Single-core analysis

ECM model





Single-core ECM model



Hofmann et.al.: Bridging The Architecture Gap: Abstracting Performance-Relevant Properties Of Modern Server Processors, https://doi.org/<u>10.14529/jsfi200204</u>

In-core prediction

Application knowledge

STREAM TRIAD a[i] = b[i] + s * c[i]

.L18:

ld1d z4.d, p5/z, [x21, x9, lsl 3]
ld1d z5.d, p5/z, [x20, x9, lsl 3]
fmad z5.d, p5/m, z2.d, z4.d
st1d z5.d, p5, [x19, x9, lsl 3]
add x8, x9, 8
whilelo p5.d, w8, w7
b.any .L18

Reservation Stations RSE0 RSA0 RSA1 **RSBR** RSE1 20 entries 20 entries 19 entries 10 entries 10 entries FLB EAGA EAGB FLA PR EXA EXB BR μ<mark>Ο</mark>Ρ μΟΡ μΟΡ μΟΡ μ<mark>Ο</mark>Ρ μ<mark>Ο</mark>Ρ μΟΡ μ<mark>Ο</mark>Ρ BR Predicate manipul. Int ALU Int ALU Int ALU Int ALU LD LD FP arith. MUL DIV AGU AGU FP arith. Int ST FMA FMA FP DIV VEC addr. calc **Execution Units** FP ST





In-core prediction

Application knowledge

STREAM TRIAD a[i] = b[i] + s * c[i]

L18:

ld1d z4.d, p5/z, [x21, x9, lsl 3] ld1d z5.d, p5/z, [x20, x9, lsl 3] fmad z5.d, p5/m, z2.d, z4.d st1d z5.d, p5, [x19, x9, lsl 3] add x8, x9, 8 whilelo p5.d, w8, w7 b.any .L18



Machine knowledge



https://github.com/RRZE-HPC/OSACA

In-core prediction

Instruction	Reciprocal Throughput [cy]	Latency [cy]
ld1d	0.5	11
st1d	1.0	_
fadd	0.5	9
fmad	0.5	9
faddv	11.5	49

Reservation Stations					
RSE0 20 entries	RSE1 20 entries	RSA0 10 entries	RSA1 10 entries	RSBR 19 entries	
FLAPRFLBμOPμOPμOP	ΕΧΑ ΕΧΒ μΟΡ μΟΡ	EAGA μΟΡ	<mark>EAGB</mark> μΟΡ	BR µOP	
Int ALUPredicate manipul.Int ALUFP arith.FP arith.FP arith.	Int ALU Int ALU MUL DIV	LD	LD	BR	
FP DIV VEC addr. calc					
FP ST			Exe	cution Units	

OS ACA Static analysis and prediction of in-core contribution

https://github.com/RRZE-HPC/OSACA

Data transfer for STREAM triad



Overlap hypotheses for A64FX



Model validation (FX1000, large pages)



Multicore (in-memory data set)



Sufficient unrolling is crucial (but sometimes it's not enough)





SpMV

Sparse Matrix-Vector Multiplication





SpMV



Sparse Matrix-Vector Multiplication (SpMV): b=Ax

In Compressed Row Storage (CRS) format

```
for i = 0:nrows-1 //Long outer loop
for j = row_ptr[i]:row_ptr[i+1]-1 // Short inner loop
b[i] = b[i] + A[j] * x[col_idx[j]]
```

Minimum code balance: $B_c^{min} = 6 \frac{\text{byte}}{\text{flop}}$

SpMV – dRECT vs. HPCG-128³



- dRECT: 4000-column tall & skinny dense matrix ($N_{nzr} = 4000$)
- HPCG: matrix from HPCG benchmark ($N_{nzr} = 27$), 128³ rows

SpMV

Assembly of the short inner-loop



The problem with SpMV on A64FX

We need both:

- SIMD vectorization
- Modulo Variable Expansion (MVE)

With CRS, both must be implemented in the inner loop. The partial sums accumulation adds to the overhead.

Can we get rid of the partial sums accumulation and separate SIMD from MVE?





$CRS \rightarrow SELL-C-\sigma$

Change data storage format



SELL-C- σ

Idea

- Sort rows according to length within sorting scope σ
- Store nonzeros column-major in zero-padded blocks of height C



Example C = 4 without further unrolling \rightarrow longer inner loop, but still an LCD

```
for(i = 0; i < N/4; ++i)
ſ
  for(j = 0; j < cl[i]; ++j)</pre>
  ſ
    y[i*4+0] += val[cs[i]+j*4+0] *
              x[col[cs[i]+j*4+0]];
    y[i*4+1] += val[cs[i]+j*4+1] *
              x[col[cs[i]+j*4+1]];
                                        C = 4
    y[i*4+2] += val[cs[i]+j*4+2] *
              x[col[cs[i]+j*4+2]];
    y[i*4+3] += val[cs[i]+j*4+3] *
              x[col[cs[i]+j*4+3]];
```

How to choose the parameters?

- C
 - $n \times SIMD$ width to allow good utilization of SIMD units
 - n > 1 useful for hiding ADD pipeline latency

• *o*

- As small as possible, as large as necessary
- Large σ reduces zero padding (brings β closer to 1)
- Sorting alters RHS access pattern $\rightarrow \alpha$ depends on σ

M. Kreutzer, G. Hager, G. Wellein, H. Fehske, and A. R. Bishop: A *unified sparse matrix data format for efficient general sparse matrix-vector multiplication on modern processors with wide SIMD units*. SIAM Journal on Scientific Computing **36**(5), C401–C423 (2014). DOI: 10.1137/130930352,



SpMV performance with SELL-C- σ (1 CMG)

 SELL-C-σ separates SIMD from sum reduction

reduction of fmla

latency impact

C>8 allows for

Performance [Gflop/s]



SpMV performance with SELL-C- σ (full chip)







Domain Wall (DW) kernel

from Quantum Chromodynamics (QCD)



Context

- Lattice QCD simulates the strong interaction
- Iterative multigrid techniques on regular (4D or 5D) lattices
- Core component: Apply Dirac operator D to quark-field vector Ψ
- Domain Wall (DW) formulation: quark field lives on 4D boundary of a 5D space-time volume $V_4 \times L_s$

 $(D\psi)(n,s)_{\alpha a} =$

$$\sum_{\mu=1}^{n} \sum_{\beta=1}^{n} \sum_{b=1}^{n} \left\{ U_{\mu}(n)_{ab} (1+\gamma_{\mu})_{\alpha\beta} \psi(n+\hat{\mu},s)_{\beta b} + U_{\mu}^{\dagger} (n-\hat{\mu})_{ab} (1-\gamma_{\mu})_{\alpha\beta} \psi(n-\hat{\mu},s)_{\beta b} \right\}$$

DW stencil kernel (simplified)

<pre>#define x_p 1 // x-plus direction #define x_m 2 // x-minus direction #define y_p 3 // y-plus direction</pre>	 "Grid" lattice QCD framework Uses SVE intrinsics Data type: double complex
<pre>#pragma omp parallel for schedule(static)</pre>	
for {t, z, y, x} = 1:{ $L_t - 2, L_z - 2, L_y - 2, L_x - 2$ } // collaps	ed loop over 4d space-time
{	
for (int $s=0$; $s; ++s) //loop over 5th dimen$	sion
{	
$O[t][z][y][x][s] = R(x_p) \cdot U[x_p][t][z][y][x$	$] \cdot P(x_p) \cdot I[t][z][y][x+1][s] +$
$\mathbf{R}(\mathbf{x}_{m}) \cdot \mathbf{U}[\mathbf{x}_{m}][t][z][y][x]$	$] \cdot P(x_m) \cdot I[t][z][y][x-1][s] +$
$R(y_p) \cdot U[y_p][t][z][y][x]$	$] \cdot P(y_p) \cdot I[t][z][y+1][x][s] +$
$\mathbf{R}(\mathbf{y}_{m}) \cdot \mathbf{U}[\mathbf{y}_{m}][\mathbf{t}][\mathbf{z}][\mathbf{y}][\mathbf{x}]$	$] \cdot P(y_m) \cdot I[t][z][y-1][x][s] +$
$R(z_p) \cdot U[z_p][t][z][y][x$	$] \cdot P(z_p) \cdot I[t][z+1][y][x][s] +$
$\mathbf{R}(\mathbf{z}_{m}) \cdot \mathbf{U}[\mathbf{z}_{m}][\mathbf{t}][\mathbf{z}][\mathbf{y}][\mathbf{x}]$	$] \cdot P(z_m) \cdot I[t][z-1][y][x][s] +$
$R(t_p) \cdot U[t_p][t][z][y][x]$	$] \cdot P(t_p) \cdot I[t+1][z][y][x][s] +$
$\mathbf{R}(\mathbf{t}_{\mathbf{m}}) \cdot \mathbf{U}[\mathbf{t}_{\mathbf{m}}][\mathbf{t}][\mathbf{z}][\mathbf{y}][\mathbf{x}]$	$] \cdot P(t_m) \cdot I[t-1][z][y][x][s];$

Complex numbers data layout choice



RRII R R R R R R I I I I I I I I I I I I R R R ···

Observed performance

- Starting point: RIRI layout, ACLE intrinsics, GCC/FCC
- 1320 flops/LUP (theoretical)
- Measured code balance: 1500 byte/LUP

• A64FX (FX1000): $B_m = 0.25 \frac{\text{byte}}{\text{flop}} \rightarrow \text{expect memory bound}$



 $B_c \approx 1.14 \frac{\text{byte}}{\text{flop}}$

Layer Conditions (LC) analysis

- LC determine traffic to/from different caches (i) of size s_i
- w = 2: write-allocate factor (would be 1 if WA evasion applies)
- RIRI (RRII) layout has d = 1 (d = 2)

Name	V _j in byte/LUP	$s_i >$	
		Scalar code	512-bit vectorized code
no reuse	$(8 \cdot 9 + 8 \cdot 12 + w \cdot 12) \cdot 16$	0	0
LC_s	$(8 \cdot 9/L_s + 8 \cdot 12 + w \cdot 12) \cdot 16$	2880	11520
LC_x	$(8 \cdot 9/L_s + 7 \cdot 12 + w \cdot 12) \cdot 16$	$2 \cdot L_s(8 \cdot 9/L_s + 8 \cdot 12 + 12) \cdot 16$	$d \cdot 8 \cdot L_s(8 \cdot 9/L_s + 8 \cdot 12 + 12) \cdot 16$
LC_y	$(8 \cdot 9/L_s + 5 \cdot 12 + w \cdot 12) \cdot 16$	$2 \cdot L_s \cdot L_x(8 \cdot 9/L_s + 7 \cdot 12 + 12) \cdot 16$	$d \cdot 8 \cdot L_s \cdot L_x (8 \cdot 9/L_s + 7 \cdot 12 + 12) \cdot 16$
LC_z	$(8 \cdot 9/L_s + 3 \cdot 12 + w \cdot 12) \cdot 16$	$2 \cdot L_s \cdot L_x \cdot L_y (8 \cdot 9/L_s + 5 \cdot 12 + 12) \cdot 16$	$8 \cdot L_s \cdot L_x \cdot L_y (8 \cdot 9/L_s + 5 \cdot 12 + 12) \cdot 16$
LC_t	$(8 \cdot 9/L_s + 1 \cdot 12 + w \cdot 12) \cdot 16$	$2 \cdot L_s \cdot L_x \cdot L_y \cdot L_z (8 \cdot 9/L_s + 3 \cdot 12 + 12) \cdot 16$	$4 \cdot L_s \cdot L_x \cdot L_y \cdot L_z (8 \cdot 9/L_s + 3 \cdot 12 + 12) \cdot 16$

SIMD-friendly data layout

- One partition per SIMD lane
- Partition size:

$$L_x \times L_y \times \frac{L_z}{2} \times \frac{L_t}{2} \times L_s$$

 → SIMD makes LCs more stringent (need more cache to fulfill)



LC effect



Summary of optimizations for DW

- Software prefetching decreases L2 data volume
- -O1 makes compiler obey the ordering hints in the computational kernel (more efficient OoO execution)
- RRII data layout
 - Prevents use of complex arithmetic instructions fcmla/fcadd
 - Removes imbalance between FLA and FLB ports in the core
 - Some register spills occur, but still better than RIRI
 - Measurement falls short of ECM prediction by 2.3x (GCC) or 3.1x (FCC)

DW kernel optimizations and ECM model



CMG performance RIRI vs. RRII

- RRII saturates already at 8 cores
- Sharing across cores in L2 gives slight increase @ 12 cores



Comparison with other architectures



- ECM model constructed for single-core performance of A64FX
- Partially overlapping memory hierarchy → high single-core memory bandwidth (even more so with large pages)
- If performance is bad, the single-core performance is usually the culprit
- SpMV requires proper data format for efficient single-core execution
- DW kernel benefits from prefetching and OoO improvements
- Performance modeling is invaluable for navigating optimization efforts





Thank You.







