

Apple Silicon CPU Optimization Guide: 3.0

2024-03-21

Author
Apple Inc.

Copyright © 2024 Apple Inc. All rights reserved.

This material may contain confidential and proprietary information of Apple Inc., which may include intellectual property of Apple Inc., whether registered or otherwise. Any use, reproduction, disclosure, or distribution of this material is subject to the restrictions imposed as a condition for receiving this document, including but not limited to those restrictions provided in the Apple Developer Program License Agreement, the Apple Developer Agreement, and the Limited License to Apple Silicon CPU Optimization Guide. Use of copyright notice is precautionary and does not imply publication or disclosure. Apple and the Apple logo are trademarks of Apple Inc., registered in the U.S. and other countries.



Chapter 2. ISA Optimization: Overview & Integer Unit

The following chapter describes aspects of the processor's Instruction Set Architecture. However, before coding and tuning custom software, explore the optimized universally-installed frameworks and libraries provided for common data types, functions, and algorithms.

2.1 Apple Platform Technologies

Xcode is Apple's integrated development environment (IDE). It offers developers access to a large number of optimized universally-installed platform [technologies](#) that include data types, library classes and functions, and services for many common tasks. Prior to writing custom code, explore the provided tuned [technologies](#) to speed development.

- **Accelerate:** Of particular interest, the [Accelerate](#) framework consists of routines for neural networks (BNNS), image and video processing (vImage), digital signal processing (vDSP), transcendentals (vForce), and linear algebra (Sparse Solvers, BLAS, and LAPACK). Outside of Accelerate, Xcode offers multithreaded compression functions (Apple Archive), compression algorithms for LZFS, LZ4, LZMA, and ZLIB (Compression), and small vector and matrix operations (simd).
- **simd:** The [simd](#) module provides basic data types and simple functions related to small vectors and matrices.
- **Grand Central Dispatch and Background Task:** The [Grand Central Dispatch](#) framework and [Background Task](#) framework provide support for scheduling tasks. Also see [Chapter 5, Asymmetric Multiprocessor \(AMP\) Optimization](#)
- **Many Other Technologies:** For example: Core Image, Core Media, Create ML, Image I/O, ML Compute. Use the Metal framework for rendering advanced 3D graphics and performing data-parallel computations using graphics processors.
- **memcpy():** Use highly tuned `memcpy()` for copying blocks of data.

Recommendation: Explore Platform Technologies and Libraries for Optimized Common Functions and Algorithms:

[Magnitude: High | Applicability: High] Xcode provides developers access to tuned libraries of mathematical data types, functions, and algorithms. These algorithms are optimized for the available hardware and offer the most portable solution between Apple products and generations. Before developing custom solutions, explore the Accelerate and other platform [technologies](#). They may speed both development time as well as application performance without the need for extensive custom code.

2.2 Arm AARCH64 ISA

Apple silicon CPUs support the Arm A-Profile 64b (AARCH64) ISA. The cores do not support AARCH32 nor ISA prior to Arm v8.0. When used, the term "Arm ISA" refers to Arm AARCH64 v8.

All Apple silicon CPUs documented in this guide (see [Section 1.1, "Generation and Series Naming Conventions"](#))x support floating point, floating point half precision converts, Advanced SIMD, and the CRC32 instruction.

[Table 2.1](#) lists the base Arm ISA version and supported ELO (unprivileged, colloquially "user code") features, base and optional, for each CPU.

Software can dynamically check for the existence of the features via the `sysctl` interface. See [Appendix B, Dynamic Determination of Chip-Specific Capabilities](#) for details.

Table 2.1. Supported Arm Base ISA and ELO Features

Feature	Description [Clang ISA Extension Name]	CPU Support		
		M1 Gen. A14 Bionic	M2 Gen. A15 Bionic	M3 Gen. A16 Bionic
	Base ARM ISA	v8.5 (excluding FEAT_BTI)	v8.6	v8.6
FEAT_AES	Advanced SIMD AES instructions	Yes	Yes	Yes
FEAT_AFP	Alternate floating-point behavior	No	No	Yes
FEAT_BF16	Storage and arithmetic instructions of the Brain Floating Point (BF16) data type [bf16]	No	Yes	Yes
FEAT_BTI	Instructions to guard against the execution of instructions that aren't the intended target of a branch	No	Yes	Yes
FEAT_DPB	Clean data cache by address to Point of Persistence DC CVAP instruction	Yes	Yes	Yes
FEAT_DPB2	Clean data cache by address to Point of Deep Persistence DC CVADP instruction	Yes	Yes	Yes
FEAT_DotProd	Advanced SIMD Int8 dot product instructions	Yes	Yes	Yes
FEAT_ECV	Support for Enhanced Counter Virtualization, which enhances the Generic Timer architecture	No	Yes	Yes
FEAT_FCMA	Floating-point complex number instructions	Yes	Yes	Yes
FEAT_FHM	Floating-point half-precision multiplication instructions	Yes	Yes	Yes
FEAT_FlagM	Condition flag manipulation instructions	Yes	Yes	Yes
FEAT_FlagM2	Enhancements to condition flag manipulation instructions	Yes	Yes	Yes
FEAT_FP16	General half-precision floating-point data processing instructions	Yes	Yes	Yes

coding scheme where the registers are named generically, such as v_n , and the (element, lanes) specification is attached to the instruction mnemonic instead. For example, the following are equivalent:

```
ADD    V2.4S, V3.4S, V4.4S
ADD.4S V2, V3, V4
```

Note: This simplified scheme of attaching the vector register organization to mnemonic instead of each operand is not part of the Arm notation, but some examples use it to improve readability.

In cases of mismatched (element size, lanes) specifications between sources and destinations, Apple assemblers will apply the specification attached to the instruction mnemonic of the destination.

Process State (`PSTATE`): The register that holds the contents of the arithmetic flags: Negative Condition (`N`), Zero Condition (`Z`), Carry Condition (`C`) and Overflow Condition (`V`). Note though that the Carry Condition (`C`) in the Arm ISA has opposite semantics during subtraction compared with other ISAs.

2.7 Separate Source and Destination Registers

Arm instructions typically specify a destination register separate from the source register(s). Other architectures have a combined source/destination register. This is often referred to as "destructive", as the source value is "destroyed" when the instruction writes it's result.

In the Arm ISA, source registers are generally not automatically overwritten unless the destination register explicitly matches the source register. As a side benefit, few move operations are required to preserve registers values before using them as source/destination.

For example, a basic "Add (register)" takes a destination register `xd` as well as two source registers `xn` and `xm`:

```
ADD Xd, Xn, Xm    // ([Xd] = [Xn] + [Xm])
```

Some ASIMD&FP operations even take 3 source registers along with a destination register, such as "floating point fused Multiply-Add to accumulator":

```
FMADD Dd, Dn, Dm, Da    // ([Dd] = [Da] + ([Dn] * [Dm]))
```

Note that some operations do exist in the ISA where a source/destination register is specified, such as "Signed Absolute difference and Accumulate" that accumulate into a register:

```
SABA Vd.T, Vn.T, Vm.T    // ([Vd] = [Vd] + ABS([Vn] - [Vm]))
                        // where T is the element organization
```

The Arm ISA generally denotes these source/destination registers just as destination registers `<*d>` but the values will also be read as sources according to the instruction

Table 2.3. Condition Codes

Mnemonic extension	Meaning (integer)	Meaning (floating point)	Conditional flags
EQ	Equal	Equal	Z==1
NE	Not equal	Not equal, or unordered [†]	Z==0
CS/ HS	Carry set/Unsigned higher or same	Greater than, equal, or unordered [†]	C==1
CC / LO	Carry clear/Unsigned lower	Less than	C==0
MI	Minus, negative	Less than	N==1
PL	Plus, positive or zero	Greater than, equal, or unordered [†]	N==0
VS	Overflow	Unordered [†]	V==1
VC	No overflow	Not unordered [†]	V==0
HI	Unsigned higher	Greater than, or unordered [†]	C==1 and Z==0
LS	Unsigned lower or same	Less than or equal	C==0 or Z==1
GE	Signed greater than or equal	Greater than or equal	N==V
LT	Signed less than	Less than, or unordered [†]	N!=V
GT	Signed greater than	Greater than	Z==0 and N==V
LE	Signed less than or equal	Less than, equal, or unordered [†]	Z==1 or N!=V
AL / NV ^{††}	Always (unconditional)	Always (unconditional)	Any

[†]Unordered means at least one NaN operand.

^{††}The NV condition code does not appear in the Arm Architecture Reference Manual, however some assemblers accept it. While the NV condition code implies "Never" and is encoded to suggest inverted AL "Always", the behavior of NV is equivalent as AL which is "always taken". See `shared/functions/system/ConditionHolds()` shared pseudocode in the Arm reference manual.

2.10 Conditional Instructions

The Arm instruction set offers a number of instructions that modify register values depending on a condition. See [Section 4.4.5, "Conditional Branch Mispredicts and Conditional Instructions"](#) for a list of these instructions and recommendations on when to employ them.

2.11 Fused Op with Add/Accumulate Operations (Integer Unit)

The Arm ISA offers a limited set of "fused" instructions that combine an operation with an add or accumulate. For example, MADD (Multiply-Add) multiplies two register values, adds a third register value, and writes the result to the destination register.

The benefits are similar to the Advanced SIMD Fused Op with Add/Accumulate Instructions and both are detailed jointly in [Section 3.5.4, "Fused Op with Add/Accumulate Instructions \(Integer and ASIMD&FP Types\)"](#).

Recommendation: Avoid Frequent JIT-Generated Code and Carefully Follow Prescribed Steps:

Self-modifying or JIT-generated code requires executing several steps to ensure the new code is visible to the instruction fetch unit in the processor. One step includes invalidating the instruction cache, which will evict other useful cached instructions. Avoid frequent code deployment, and hence frequent instruction cache invalidations, as much as possible.

[Magnitude: Low | Applicability: Low]

2.14 ISA Characterization

Broad bucket characterization of ISA usage is available through the following metrics.

Table 2.4. Common Instruction Mix Metrics

Name and Formula (Event Definitions: Section 6.2, "Performance Monitoring Events")	Description
Branch Density[†]: <Ev INST_BRANCH> / <Ev INST_ALL>	Proportion retired branch instructions (including calls and returns) of all retired instructions
Integer Operation Density: <Ev INST_INT_ALU> / <Ev INST_ALL>	Proportion retired integer execution instructions of all retired instructions, excluding branches and memory operations
Advanced SIMD and FP Operation Density: <Ev INST_SIMD_ALU> / <Ev INST_ALL>	Proportion retired Advanced SIMD and FP execution instructions of all retired instructions, excluding memory operations
Advanced SIMD Operation Density: <Ev INST_SIMD_ALU_VECTOR> / <Ev INST_ALL>	Proportion retired Advanced SIMD execution instructions (including integer and floating point data types) of all retired instructions, excluding memory operations. Note: Available on M2 Generation and following, and A15 Bionic and following
Load and Store Density: <Ev INST_LDST> / <Ev INST_ALL>	Proportion retired load and store instructions of all retired instructions
Integer Load Density: <Ev INST_INT_LD> / <Ev INST_ALL>	Proportion retired integer load instructions of all retired instructions
Integer Store Density: <Ev INST_INT_ST> / <Ev INST_ALL>	Proportion retired integer store instructions of all retired instructions
Advanced SIMD and FP Load Density: <Ev INST_SIMD_LD> / <Ev INST_ALL>	Proportion retired Advanced SIMD and FP load instructions of all retired instructions
Advanced SIMD and FP Store Density: <Ev INST_SIMD_ST> / <Ev INST_ALL>	Proportion retired Advanced SIMD and FP store instructions of all retired instructions
Data Barrier Density: <Ev INST_BARRIER> / <Ev INST_ALL>	Proportion retired data barrier (DSB, DMB) instructions of all retired instructions

will be evident. The `NM` variants use IEEE NaN handling instead, always selecting any actual numerical value over NaN value in any particular lane. In this mode, NaN values are propagated to the result *only* if *both* inputs are NaN values, and can disappear otherwise.

- `FMUL` instructions have a special `x` variant that produces a value of `2.0` for any occurrences of zero times infinity, instead of just zero. The result is negative if only one of the values is negative, otherwise the result is positive. This represents a common limit case for some very specific mathematical operations. Unrelated, the "x" suffix on the `FRECPX` instruction represents "exponent".

These three suffixes are seldom used in general-purpose computations.

3.3.14 FP Rounding Mode: A / I / M / N / P / X / Z

Both `CVT` and `INT` instructions use a variety of suffixes to select different rounding modes for FP-to-integer or FP-to-fixed-point conversions, which often have a fractional part that will need to be rounded off. Variations include round towards: nearest, zero, $+\infty$, $-\infty$, and others. These variants are usually used to determine the range of possible rounding error that may be introduced by conversion operation. See Arm Architecture Reference Manual for full descriptions and options.

There are also two special `INT` instruction rounding modes:

- **I**: The instruction uses the current default `FPSCR` rounding mode. This is also the default option for `CVT` operations that have no explicit rounding mode suffix.
- **X**: The instruction uses the "eXact" mode that triggers an exception when any rounding occurs. Use this mode to guarantee that only integers are being converted. A signal handler must be registered prior to the use of these instructions to catch the resulting exceptions.

Do not confuse the rounding `N` suffix with the narrowing `N` suffix, because both can be used with different kinds of `CVT` instructions. Rounding `N` always includes an output type specifier from [Section 3.3.15, "Different Output Type: \(none\) / F / S / U"](#) just after the `N` suffix.

3.3.15 Different Output Type: (none) / F / S / U

Most instructions output the same type as the input, from the type prefix described in [Section 3.3.1, "Overall instruction data type: \(none\) / BF / F / S / U"](#). However, most `CVT` instructions use explicit type suffixes to indicate the different destination type for the conversion (e.g. `F` for floating point, `s` for signed integer, or `U` for unsigned integer). There are also a few signed-to-unsigned integer/fixed-point `SHR` and `XT` instructions that can produce an unsigned output from a signed input, stripping the sign bit off of the output in order to get an extra bit of range. These are indicated with a `U` suffix.

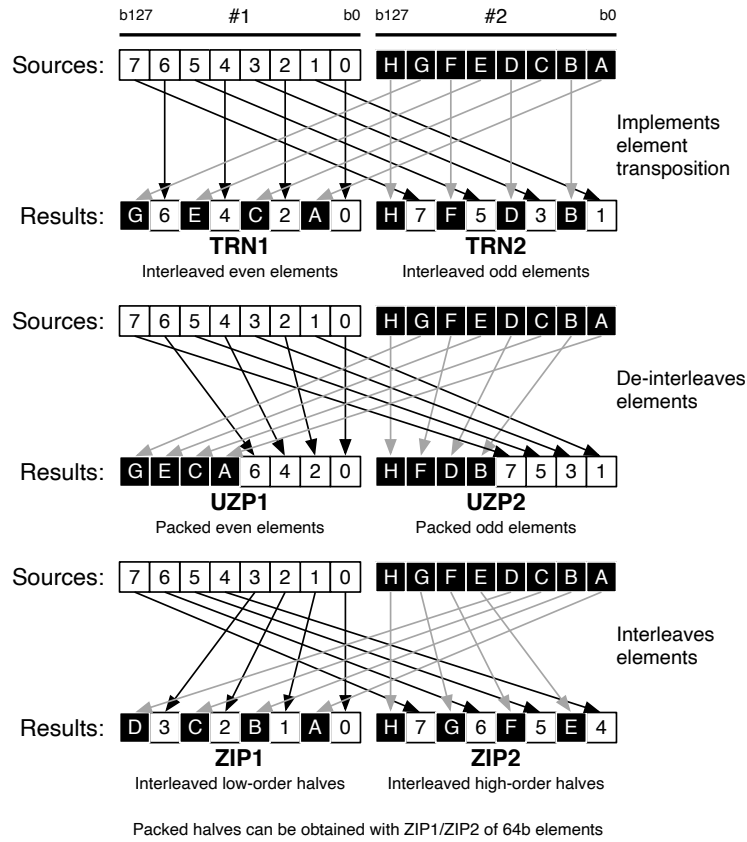
3.3.16 Narrow, Long, Wide: N / L / W

By default, Advanced SIMD operations have the same data type for the input and output. These three suffixes are used to select different bit widths for the input and output:

- **Narrow (n)**: Destination is half the size of the source type(s): `64→32 bit (D→S)`, `32→16 bit (S→H)`, and `16→8 bit (H→B)` conversions

idea further and shows how different instruction variants can be used to generate a wide variety of transformations. Note that while these permutation instructions are often used in 1 / 2 pairs, there are many situations when using only one or the other might still be useful.

Figure 3.4. Advanced SIMD Byte Transformations



3.3.19 Cross-Lane Paired or Horizontal: P / V

Advanced SIMD has two different types of “horizontal” reduction operations, which combine results in different vector lanes together. The P-suffix *paired* versions combine adjacent vector lanes across two vectors together to produce a single vector as a result. Figure 3.5: “Advanced SIMD Cross Lane Paired and Horizontal Instructions” (B) illustrates a paired instruction. Trees of repeated *paired* operations will gradually combine all lanes of input down to result element #0. These versions are available for any data type. Meanwhile, the powerful v-suffix *across vector* versions compute a sum, minimum, or maximum across *all* elements of a vector with a single instruction, as is depicted in Figure 3.5: “Advanced SIMD Cross Lane Paired and Horizontal Instructions” (C). For complexity reasons, these are not available for all data types, particularly FP types. They can be synthesized with repeated P-suffix operations if necessary, however, as is described in Section 3.5.3, “Horizontal Arithmetic and Test Instructions”. For comparison, Figure 3.5: “Advanced SIMD Cross Lane Paired and Horizontal Instructions” (A) depicts the data flow through normal operations *within* vector lanes.

These suffixes can also be combined with the L (long) suffix from Section 3.3.16, “Narrow, Long, Wide: N / L / W” to create paired or cross-vector operations that produce

Each of these intrinsics just extracts the requested half of the 128-bit vector as a 64-bit vector, using the type on the same row of [Table 3.11: “Advanced SIMD Intrinsic Data Types”](#).

3.4.4.6 Byte Rotations Operations

The “extract” operation allows a vector to be extracted from any bitwise rotation of a pair of vectors. This instruction is most often used for realigning vectors that have been read in an unaligned fashion, but can also be used for a wide variety of vector rotations. The format of the operation follows standard guidelines:

```
<vector_dest> = vext[q]<type>(<vector_src_low>, <vector_src_high>, <shift_amt>);
```

The result will be the high-order elements of `<vector_src_low>`, rotated down to the low-order end of the destination vector, and the low-order elements of `<vector_src_high>`, rotated in to the high-order end of the destination. The overall rotation is controlled by `<shift_amt>`. These operations all use the EXT Advanced SIMD instructions, which only come in byte-typed versions with a shift amount of 0-15 bytes. The *intrinsics* automatically scale this amount to match the type:

- **Half precision:** shift amounts of 0-7 halfwords for 128-bit vectors (0-3 for 64-bit vectors)
- **Single precision:** shift amounts of 0-3 words for 128-bit vectors (0-1 for 64-bit vectors)
- **Double precision:** shift amounts of 0-1 doublewords for 128-bit vectors (doubleword operations with 64-bit vectors do not perform any operation).

Note that due to the allowed range of immediate values, the result can be the entire `<vector_src_low>` vector (effectively making this a NOP when the shift amount is 0), but can never be the entire `<vector_src_high>` vector.

3.4.4.7 Data Transposition Operations

Three main data transposition instructions are available in the Advanced SIMD instruction set:

- **Transpose (TRN):** Selects and packs even (or odd) numbered elements from two vectors. These are designed to facilitate matrix transposition (see [Section 3.5.7.1, “Byte Shuffle Example: Transposing Matrices”](#)).
- **Unzip (UZP):** De-interleaves two vectors of interleaved data.
- **Zip (ZIP):** Interleaves two vectors of de-interleaved data.

These operations are usually used in “primary” low half (‘1’) and “secondary” high half (‘2’) pairs to mix two vectors together. However, there are also other ways to use these instructions for more arbitrary data transformations (see [Section 3.3.18, “Lower-Half / Upper-Half: \(none\) / 1 / 2 ”](#) and [Section 3.5.7, “Shuffle and Permute Instructions”](#)). The intrinsics or these instructions is straightforward:

```
<vec_dest_low> = v[trn|uzp|zip]1[q]<type>(<vec_src_low>, <vec_src_high>);
<vec_dest_high> = v[trn|uzp|zip]2[q]<type>(<vec_src_low>, <vec_src_high>);
```

```

UMINV.16B B1, V0
FMOV      W0, S1
CBNZ     W0, target

```

Similarly efficient code can be constructed for reductions and other cross-vector operations.

The integer and brain floating point (`bf16`) dot product instructions also sum across the results of the lane multiplies. See [Section 3.5.6, "Dot Product and Matrix Multiply Instructions"](#).

Recommendation: Use Horizontal Instructions for Common Vector Tests:

Common any- and all-lane tests can be accomplished through the `UMINV` and `UMAXV` instructions. The result can be moved to the integer registers where it can be used as a branch condition.

[Magnitude: Medium | Applicability: Medium]

3.5.4 Fused Op with Add/Accumulate Instructions (Integer and ASIMD&FP Types)

The Arm ISA offers a limited set of "fused" instructions that combine an operation with an add or accumulate — for example, `FMLA`, is a floating point fused multiply-add to accumulator. The processor decodes, issues, executes, and retires these instructions as single μ ops for increased computation efficiency.

Some fused two-operation instructions use three source registers and a separate destination register. This style of instruction is used for fused instructions that execute on the Integer Execution Unit. For example:

- **MADD:** Multiply-Add multiplies two integer register values, adds a third integer register value, and writes the result to the integer destination register.

Others, as noted in [Section 2.7, "Separate Source and Destination Registers"](#), use 2 source registers and a third combined source+destination register. This style of instruction is used for Advanced SIMD and FP instructions, including those that operate on integer typed-elements. For example:

- **MLA:** Multiply-Add to Accumulator multiplies corresponding integer elements in the vectors of the two source ASIMD&FP registers, and accumulates the results with the vector elements of the destination ASIMD&FP register.

In this section, "op-add" refers generically to all relevant instructions, whether they add to a separate destination or into an accumulator.

There are several benefits to op-add instructions:

- **Increased decode and issue throughput:** All op-add instructions effectively combine two operations. As long as both operations are useful, combining them in the same instruction effectively doubles throughput of basic operations throughout the processor.
- **Lower overall latency:** In many cases, the op-add instruction executes with the same latency as an instruction that only executes the base op. However, the latency

to select for each output byte position. For any scenario where a number of different arrangements will need to be used regularly in quick succession, several registers may have to be dedicated to holding these mapping control vectors.

As a result, it is a better to use one of the “dedicated” byte-rearranging instructions (`DUP`, `EXT`, `INS`, `REV`, `TRN`, `UZP`, and `ZIP`) if possible. [Table 3.13: “Advanced SIMD Shuffling Operations”](#) lists these various instructions, along with brief descriptions and a summary of their primary intended uses.

Table 3.13. Advanced SIMD Shuffling Operations

Instruction	Description	Primary Use
<code>DUP</code>	Copies one vector lane of input to all lanes of output	Copying any single value across all vector lanes
<code>EXT #n</code>	Appends 2 full vectors together, rotates n bytes, and extracts 1 vector from the center	Aligning unaligned input; shifting upper-lane elements down to lower lanes
<code>INS</code>	Inserts a value from one vector lane of input into one vector lane of output	Moving of one value around within a vector
<code>REV16</code>	Reverses operand order, within halfwords	Switching endian-ness of half word data
<code>REV32</code>	Reverses operand order, within words	Switching endian-ness of word data
<code>REV64</code>	Reverses operand order, within doublewords	Switching endian-ness of doubleword data
<code>TBL</code>	Selects bytes from 1–4 input vectors to assemble an output vector, based on a selection mask vector	Assembling a vector, byte-by-byte
<code>TBX</code>	Selects bytes from 1–4 input vectors to insert into an output vector, based on a selection mask vector	Inserting bytes into an existing vector
<code>TRN</code>	Interleaves the even or odd numbered elements from two input vectors	Transposing blocks of vectors
<code>UZP</code>	De-interleaves the even or odd numbered elements from two input vectors	De-interleaving packed objects like pixels into separate buffers
<code>ZIP</code>	Interleaves the low or high half elements from two input vectors	Interleaving separated values together, like merging per-color buffers into pixels

The various possible output permutations of the input bytes from each of the Advanced SIMD byte-shuffling instructions are listed in the following tables. Input #1 bytes are indicated using hexadecimal (with lower case letters), while input #2 bytes are indicated using alphabetic capital letters and are shaded gray. Note that `TRN`/`UZP`/`ZIP` all behave identically to each other when the inputs are each comprised of two 64-bit operands and can be used interchangeably.

In many cases, algorithms require specific well-structured byte rearrangement. It may be possible to achieve the desired rearrangement through a series of steps through these dedicated shuffling instructions. Two or four neighboring elements may be moved together by specifying a larger element size than the elements themselves. See the [Section 3.5.7.1, “Byte Shuffle Example: Transposing Matrices”](#) for an example.

specifiers. These loops, which are indicated with comments, are nevertheless written as such to avoid repetition of code that only differs by a single constant value per line.

3.6.1 Matrix Operations

Advanced SIMD is often used to accelerate dense matrix computations. The structure and regularity of these algorithms makes them well suited for the parallel operations offered by the Advanced SIMD instructions. Nevertheless, there are a number of subtle details about Advanced SIMD functionality that can be learned by looking at these examples. The remainder of this section discusses how 32-bit floating point matrix-matrix and matrix-vector multiplies can be constructed.

3.6.1.1 Matrix-Matrix Multiply

Dense matrix-matrix multiplies are commonly performed using vector instructions. The following block of code shows how a simple matrix-matrix multiply can be performed with $M \times P \times P \times N = M \times N$ row-major matrices:

```
// Simple macros to read/write vectors from/to scalar arrays
#define LD_F32V(float32Ref)  (*(float32x4_t *)&(float32Ref))
#define ST_F32V(float32Ref)  (*(float32x4_t *)&(float32Ref))
#define VEC_SIZE_F32V 4
int i, j, k;           // Induction variables
float32x4_t a, c;     // Temporaries
float32_t *A;        // MxP source matrix #1
float32_t *B;        // PxN source matrix #2
float32_t *C;        // MxN destination matrix
float32x4_t zeroVec = vmovq_n_f32(0.0);

// Loop over columns of destination matrix C (rows of A)
for (i=0; i < M; i++) {
    // Loop across rows of destination matrix C (columns of B)
    // -- Calculate VEC_SIZE_F32V columns at once
    for (j=0; j < N; j+=VEC_SIZE_F32V) {
        // Loop over the "inner" P cols/rows of A/B matrices
        c = zeroVec;
        for (k=0; k < P; k+=VEC_SIZE_F32V) {
            // Read a vector full of A elements
            a = LD_F32V(A[i*P + k]);
            // Use multiply-by-element to scan down the "a" vector
            // -- First column of A (element #0) x first row of B
            c = vfmaq_laneq_f32(c, LD_F32V(B[(k+0)*N + j]), a, 0);
            // -- Second column of A (element #1) x second row of B
            c = vfmaq_laneq_f32(c, LD_F32V(B[(k+1)*N + j]), a, 1);
            // -- Third column of A (element #2) x third row of B
            c = vfmaq_laneq_f32(c, LD_F32V(B[(k+2)*N + j]), a, 2);
            // -- Fourth column of A (element #3) x fourth row of B
            c = vfmaq_laneq_f32(c, LD_F32V(B[(k+3)*N + j]), a, 3);
        }
        ST_F32V(C[i*N + j]) = c;
    }
}
}
```

The resulting assembly code from Clang intermingles the intrinsic instructions (5 LDR and 4 FMLA) with loop control and pointer management C code. Only the innermost loop is shown. For instance, the compiler simplifies the " $B[(k+\{0..3\})N + j]$ " indices in the



Chapter 4. Core Microarchitecture Optimization

The details covered in [Chapter 4, Core Microarchitecture Optimization](#) and [Appendix A, Instruction Latency and Bandwidth](#) represent the most important information to consider when optimizing software for the microarchitecture. The CPUs employ additional techniques to improve performance beyond what is described here. These may change from generation to generation, and it may be difficult or counter-productive to try to transform software to match the heuristics. The guidelines in this chapter represent the optimization opportunities with the most significant impact across P cores and E cores, and across CPU series and generations.

4.1 Apple Silicon CPU and Chip Overview

Apple silicon M Series and recent A Series chips feature two types of general purpose CPU cores, performance cores (P cores) and efficiency cores (E cores):

- P cores: Designed to achieve maximum performance. They are aggressive, out-of-order, superscalar, pipelined microprocessors that feature advanced forms of dynamic branch prediction, register renaming, out-of-order speculative execution, memory dependence prediction, and many other state-of-the-art features.
- E cores: Designed to achieve maximum efficiency, thus saving power and increasing battery life while reducing heat generation and fan noise (where applicable). They are based on a similar microarchitecture to the P cores, and still provide compelling high performance. E cores are the most efficient place to run lighter-weight and everyday tasks, allowing the performance cores to be used for the most demanding workflows. Use the E cores to meaningfully increase throughput in heavily threaded applications.

Microarchitectural recommendations broadly apply to both core types, but with a focus on the P cores. In some cases, a recommendation may apply to only one CPU core type, but it will not adversely affect performance of the other type.

On Apple silicon, the system decides whether to schedule work on the P cores or E cores based on many factors. The system does not offer any direct controls to software or to end users regarding task placement. However, software can provide hints to the system about where to schedule tasks via the Quality of Service controls. See [Section 5.1, "Prioritizing Work"](#).

Several cores of the same type are grouped together and share a second level cache and are collectively known as a cluster.

The overall parametrized topology is shown in [Figure 4.1: "General Purpose Compute Complex and Related Components"](#). For instance, the M1 chip consists of two clusters. The first cluster contains 4 P cores that share a second level cache. The second cluster contains 4 E cores that share their own second level cache. Both clusters are connected to each other and to other system components via the fabric. Also connected to the fabric is a high performance memory controller featuring a memory cache. Many other

the instruction window. When the branches actually execute, the predicted direction or target is compared against the result. Occasionally the actual result of the branch will not match the prediction, and a misprediction occurs. Thus, the instructions in the instruction window that are younger than the branch are not the proper instructions. At this point, the processor flushes the younger instructions from the instruction window, and new instructions are fetched from memory at the target of the branch.

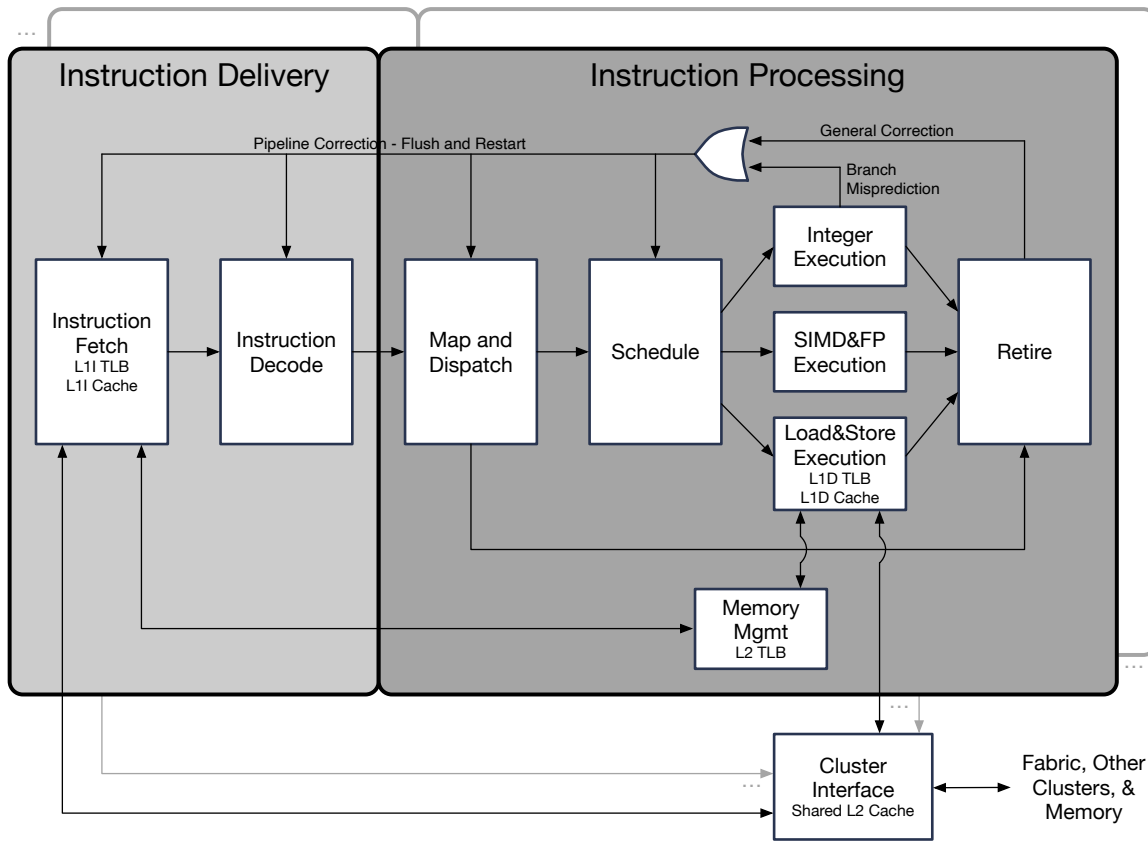
Similarly, the processor may speculate past memory dependences. That is, the processor may guess which, if any, older stores in the instruction window may write values to memory that need to be read by younger loads in the window. In cases where it guesses wrong and misses a dependence, the load will have read incorrect data and potentially fed that to subsequent younger instructions. In such cases, the processor must flush those instructions and re-execute them.

- **Superscalar:** To further increase throughput, the processor employs duplicated, or parallel, specialized hardware components to operate on more than one instruction per cycle. For example, the M1 P cores are capable of decoding 8 instructions each clock cycle, and performing 6 integer add operations per cycle (plus SIMD and memory operations as well).

At the highest level, the processor consists of two main components, Instruction Delivery and Instruction Processing. Instruction Delivery fetches and decodes instructions while Instruction Processing manages the instruction window and executes the instructions. Instruction Delivery predicts branches and fetches instructions as fast as possible to keep the instruction window in Instruction Processing as full as possible. Instruction Processing searches through all of the instructions in the window identifying and executing independent work, while retiring the oldest instructions. The Cluster Interface Unit is outside of the core boundary and serves both Instruction Delivery and Instruction Processing.

Instruction Delivery and Instruction Processing each consist of several units and connect to the Cluster Interface Unit.

Figure 4.2. Abstract View of the Processor Pipeline.



Conceptually, the processor operates on ISA instructions, as previously described. More technically, however, the processor decodes (translates) ISA instruction bytes into executable units called microoperations (μ ops). These μ ops are highly tailored to the microarchitecture. Most instructions are translated into a single μ op, but some complex instructions require multiple μ ops. In many cases, the distinction is not particularly important, but the term "instruction" will be used when specifying a particular instruction at the source code level and " μ op" when referring to specific hardware.

Instruction Delivery consists of

- **Instruction Fetch Unit (Fetch):** Reads instruction bytes from the memory system based on various prediction mechanisms. Contains the Instruction Translation Lookaside Buffer (L1 TLB), which supports [virtual memory](#) by obtaining physical addresses from virtual addresses. Contains the L1 Instruction Cache to improve performance of reads of instruction bytes from memory.
- **Instruction Decode Unit (Decode):** Translates instructions bytes into executable microoperations (μ ops).

Instruction Processing consists of

- **Map and Dispatch Unit (Map):** Inserts new μ ops into the instruction window. It obtains the necessary resources from the Execution and Retirement portions of the processor to implement out-of-order execution of the μ ops.
- **Schedule Unit (Schedule):** Selects ready μ ops and issues them to the appropriate execution units.

4.4.1 L1 Instruction (L1I) TLB

The Instruction Translation Lookaside Buffer caches virtual-to-physical address mappings. Application pages are 16KiB.

The L1I TLB is organized as follows. The L1D TLB and the L1I TLB are backed by a much larger Shared L2 TLB as well as a Page Table Walker. See discussion in [Section 4.6.2, “L1 Data \(L1D\) TLB and Shared L2 TLB”](#).

Table 4.2. L1I 16KiB-Page TLB Organization

Chip	Entries (Coverage)	
	Performance Cores	Efficiency Cores
M1 Generation and A14 Bionic	192 entries (3MiB)	128 entries (2MiB)
M2 Generation and A15 Bionic		192 entries (3MiB)
M3 Generation and A16 Bionic		

Many applications and libraries contain a moderately sized set of commonly used functions and a larger set of less commonly used functions. To minimize the number of TLB entries required, and thus the number of TLB misses encountered, place the commonly used functions into a small set of pages. Alternatively, restructure the algorithm to improve the temporal locality of function calls, where possible. That is, organize work to heavily use a function in a phase of execution, and then move on to another phase with different functions, rather than interleaving many functions.

Fetches that miss the L1I TLB are delayed by several cycles or more. These delays can result in no μ ops being delivered to Map until the translation is available.

Table 4.3. Common L1 Instruction TLB Metrics

Name and Formula (Event Definitions: Section 6.2, “Performance Monitoring Events”)	Description
L1I TLB Miss Density [†] : <Ev L1I_TLB_MISS_DEMAND> / <Ev INST_ALL>	Frequency of speculative L1 instruction TLB misses due to demand fetches compared to the count of all retired instructions
L1I TLB Fill Density [†] : <Ev L1I_TLB_FILL> / <Ev INST_ALL>	Frequency of speculative L1 instruction TLB fills for any reason compared to the count of all retired instructions

[†]Additional TLB and Address Translation Metrics are available in [Table 4.17: “Common L1D TLB, Shared L2 TLB, and Address Translation Metrics”](#)

Recommendation: Collect Commonly Used (Hot) Functions into a Small Set of Pages or Improve Temporal Access Locality to Reduce Virtual Address Translation Overhead:

[Magnitude: Medium | Applicability: Medium] Separate functions into commonly used (hot) and uncommonly used (cold). Place the commonly used functions into a small set of pages. Or, where possible structure algorithms to heavily use a small set of functions and then move on to another set.

function to be called from multiple call points while predicting correct return addresses. Like other indirect branch predictions, the return target is verified during execution, and may result in a misprediction (pipeline flush). Most well-formed code should not need modification to optimize for Return Address Stack behavior. However, follow these guidelines:

- **Use recognized call and return instructions.** Use the expected instructions so that the processor properly manages the Return Address Stack. Avoid custom code sequences, such as regular branches, to implement returns.
 - For calls: `BL` or `BLR`. Branch with Link branches to a PC-relative offset or to an address in a register, setting the register `x30` to `PC+4`. It provides a hint that this is a subroutine call.
 - For returns: `RET`, which is a specific flavor branch through register. It provides a hint that this is a subroutine return.
- **Use properly paired calls and returns.** Avoid using a single return instruction to simultaneously return through multiple stack levels. This restriction mostly affects unusual code such as the C language `longjmp` standard library routine, which can corrupt the RAS if it jumps back to the location of the last `setjmp` call without simulating any returns that would have otherwise occurred. Counterintuitively, a “fast” single `RET` that skips over several intervening returns may end up being slower overall than a “slow” return sequence that performs all of the `RET` instructions. By skipping returns, the Return Address Stack may become misaligned with actual execution, resulting in many mispredictions afterwards.
- **Restructure algorithms to avoid deep call stacks.** Inline short functions where possible to reduce call stack depth, especially those that primarily just call another function. Use recursion sparingly, if at all. Inlining also integrates the function body into the caller, eliminating any computational redundancy and data movement overhead, often resulting in further performance improvement.

Use the following metrics to measure return mispredict rates. Occasional return mispredictions will occur even if all of the above guidelines are followed. For instance, the processor and operating system do not save and restore the contents of the RAS during process switches. Therefore returns executed immediately after the process is switched back into the processor may encounter an empty RAS.

Table 4.9. Return Mispredict Metrics

Name and Formula (Event Definitions: Section 6.2, “Performance Monitoring Events”)	Description
Mispredicted Return (of instructions): <Ev BRANCH_RET_INDIR_MISPRED_NONSPEC> / <Ev INST_ALL>	Proportion retired mispredicted function returns of all retired instructions
Mispredicted Return (of returns): <Ev BRANCH_RET_INDIR_MISPRED_NONSPEC> / <Ev INST_BRANCH_RET>	Proportion retired mispredicted function returns of all retired returns

Recommendation: Leverage the Return Address Stack By Using Well Formed Call-Return Sequences and Limiting Call Depth:

[Magnitude: Medium | Applicability: Low] When calling and returning from functions, use properly paired recognized call and return instructions. Examine the call stacks for any regions of execution that commonly exhibit high return mispredicts. Correct any improperly paired or non-standard calls and returns. If significant mispredicts remain, limit call stack depth through algorithmic changes and function inlining.

4.4.8 Multi- μ op Instructions

Some complex instructions require multiple μ ops to implement the instruction's functionality. These include instructions that read more than three source registers, write multiple destination registers, and use functionality from multiple pipelines. However, use of these instructions improves computational density by packing a lot of functionality into a single 32b instruction.

Cracked instructions require sequences of 2 or 3 μ ops. The processor inserts the μ op sequence seamlessly between μ ops from the prior and subsequent instructions.

Microcoded instructions require sequences of 4 or more μ ops. In addition, the processor may not be able to seamlessly integrate these μ ops into the μ op stream and thus will incur further overhead.

As a general guideline, when using the full capability of these instructions, the added microoperations are more than worth their cost, compared with implementing the same functionality with a collection of available individual instructions. However, if the full functionality is not needed, consider alternatives.

Multi- μ op instructions include:

- **Load and store operations with address writeback:** These common cracked instructions require an extra μ op to perform the address update. Avoid chains of these instructions because of the added μ op bandwidth consumed as well as the added dependencies. Use one adjustment per block of loads and/or stores. (See [Section 2.8.1, "Avoid Chains of Pre- and Post-Indexed Operations"](#)). Formats:

```
LDx/STx Rt, [Xn, #imm]!           // Pre-index addressing mode
LDx/STx Rt, [Xn], #imm           // Post-index addressing mode
```

- **Paired load operations:** These common cracked instructions have two destination registers and are cracked before renaming. However, unless the operands are Q-sized, the processor will re-fuse them back into a single μ op before sending them to the Load and Store Execution Units. Use these instructions wherever possible. Example instructions:

```
LDP{SW} Rt1, Rt2, ...
LD{N}P Dt1, Dt2, ...
```

- **Paired ASIMD&FP store operations:** These common cracked instructions require an additional μ op. Use them wherever possible to improve code density. Example instructions:

just moves the FP register contents directly over to the integer register file without any type of conversion, while `UMOV` and `SMOV` perform a move along with a sign or zero extend of the integer value being moved, requiring extra latency. This may be common especially for byte- and halfword-sized integers. However, for cases when sign or zero extension are not needed, use `FMOV`, even if the value being moved is an integer. These data movement instructions execute in 3 or 4 cycles of latency.

Recommendation: Use `FMOV` Instruction When Moving Data from Vector Registers to GPRs when No Conversion is Needed:

[Magnitude: Low | Applicability: Medium] The `FMOV` instruction does not perform any conversion when moving data from the vector registers to the general purpose registers, whereas `UMOV` and `SMOV` do. Use the lower bandwidth and lower latency `FMOV` wherever possible.

The Arm ISA offers a number of related instructions to convert and move floating point values in vector registers H/S/D to integer (or fixed point) values in the general purpose registers: `FCVT(AMNPZ)(SU)(scalar)`. These require an operation to perform the conversion as well as an operation to move the data, resulting in a latency of 6 or 7 cycles.

4.5.3 Preferred Instructions for Common Operations

Use the following preferred instructions for common operations.

4.5.3.1 Register Data Copy

Use the following preferred instructions for creating copies of values in registers:

Table 4.12. Register Data Copy Instructions

Register Type	Register Data Copy Instructions
Integer	<code>MOV Xd, Xn</code> <code>ORR Xd, XZR, Xn // OR a 0 with Xn into Xn</code> <code>ADD Xd, Xn, #0 // Add a 0 to Xn into Xd</code>
ASIMD&FP	<code>FMOV Dd, Dn</code> <code>MOV.2D Qd, Qn</code> <code>ORR.16B Vd, Vn, Vm when (Vm==Vn) // Or Vm with itself into Vd</code>

When creating multiple copies of a value in registers, avoid chaining, and instead create multiple copies from the original:

```
MOV X6, X5
MOV X7, X6 // Avoid chaining copies

MOV X6, X5
MOV X7, X5 // Prefer creating multiple copies from the original
```

Recommendation: Use Preferred Instructions for Creating Copies of a Value in Registers:

[Magnitude: Medium | Applicability: Medium] Use the preferred instructions and avoid chaining.

Table 4.15. L1D 16KiB-Page TLB Organization (cont.)

Chip	Entries (Coverage)	
	Performance Cores	Efficiency Cores
A16 Bionic	160 entries (2.5MiB) (reduction)	

If the processor cannot find a translation in the L1D TLB, it looks up the virtual address in the much larger Shared L2 TLB. This L2 TLB also backs the L1I TLB.

Table 4.16. Shared L2 16KiB-Page TLB Organization

Chip	Entries (Coverage)	
	Performance Cores	Efficiency Cores
M1 Generation and A14 Bionic	3072 entries (48MiB)	1024 entries (16MiB)
M2 Generation and A15 Bionic		2048 entries (32MiB)
M3 Generation and A16 Bionic		

If the processor cannot find a translation in any of the TLBs, then it automatically initiates a page table walk for the virtual address via Memory Management Unit. A full page table walk is a multi-step walk of a tree structure starting at the base node and ending at a leaf node that contains the page's physical address. Steps require a memory operation to access the appropriate node in the page table structure. The MMU issues the memory request to the Shared L2 Cache, which may have the data cached. Otherwise, the L2 Cache will send the request over the fabric to the memory system.

By leveraging multilevel TLBs and the Shared L2 Cache, the overhead due to virtual memory translation gradually worsens as the working set size increases. Nonetheless, Shared L2 TLB and MMU bandwidth is limited.

Use the following metrics to evaluate TLB performance.

Table 4.17. Common L1D TLB, Shared L2 TLB, and Address Translation Metrics

Name and Formula (Event Definitions: Section 6.2, "Performance Monitoring Events")	Description
L1D TLB Miss Retired Density (of loads and stores): <Ev L1D_TLB_MISS_NONSPEC> / <Ev INST_LDST>	Proportion retired load and store instructions that missed L1D TLB of all retired load and store instructions
L1D TLB Miss Density (of L1D TLB accesses): <Ev L1D_TLB_MISS> / <Ev L1D_TLB_ACCESS>	Proportion speculative L1D TLB misses of L1D TLB accesses including loads, stores, prefetches, etc.
L1D TLB Fill Density (of L1D TLB accesses): <Ev L1D_TLB_FILL> / <Ev L1D_TLB_ACCESS>	Proportion speculative L1D TLB fills for any reason of L1D TLB accesses including loads, stores, prefetches, etc.
Instruction L2 TLB Miss Density (of L1I TLB misses): <Ev L2_TLB_MISS_INSTRUCTION> / <Ev L1I_TLB_FILL>	Proportion speculative L2 TLB instruction misses of unique L1I TLB misses (counted by fills into the L1I TLB)

Table 4.17. Common L1D TLB, Shared L2 TLB, and Address Translation Metrics (cont.)

Name and Formula (Event Definitions: Section 6.2, "Performance Monitoring Events")	Description
Data L2 TLB Miss Density (of L1D TLB misses): <Ev L2_TLB_MISS_DATA> / <Ev L1D_TLB_FILL>	Proportion speculative L2 TLB instruction misses of unique L1D TLB Misses (counted by fills into the L1D TLB)
Instruction Table Walk Request Density (of data L2 TLB misses): <Ev MMU_TABLE_WALK_INSTRUCTION> / <Ev L2_TLB_MISS_INSTRUCTION>	Mean speculative instruction table walk read requests per instruction L2 TLB miss
Data Table Walk Request Density (of data L2 TLB misses): <Ev MMU_TABLE_WALK_DATA> / <Ev L2_TLB_MISS_DATA>	Mean speculative data table walk requests per data L2 TLB miss

Recommendation: Compact Data or Improve Temporal Access Locality to Reduce Virtual Address Translation Overhead:

[Magnitude: Medium | Applicability: Medium] Avoid frequent long-latency page table walks due to sporadically accessing sparse data. Utilize data structures that keep the working page set within the TLB capacity limits. Or, organize accesses to data to improve page temporal locality.

4.6.3 L1 Data (L1D) Cache

Each core uses its own private L1 Data Cache. It usually operates as a write-back write-allocate cache. The write allocation policy keeps the recently written data in the L1 cache, where it is conveniently located for re-reading. Load μops that hit in the L1D Cache execute with a latency of 4 cycles. See [Section A.3.1, "Load Latency"](#) for latency ranges of more complex load instructions.

While the system operates on 128B cachelines, the L1D Cache operates on 64B cachelines. See [Section 4.6.6, "Improving Cache Hierarchy Performance"](#) for recommendations on improving data cache hierarchy performance.

Table 4.18. L1D Cache Organization

Chip	Capacity, Associativity, and Line Size	
	Performance Cores	Efficiency Cores
M1 Generation and A14 Bionic	128KiB, 8-way, 64B lines	64KiB, 8-way, 64B lines
M2 Generation and A15 Bionic		
M3 Generation and A16 Bionic		

Note: Capacity and associativity specifications are provided for reference. Software should check the appropriate `sysctl` parameter to dynamically adjust to CPU and chip configurations. See [Appendix B, Dynamic Determination of Chip-Specific Capabilities](#) for more information.

Use the following metrics to evaluate L1D Cache performance.

Table 4.19. Common Data Cache Metrics

Name and Formula (Event Definitions: Section 6.2, "Performance Monitoring Events")	Description
Load L1D Cache Miss Retired Density (of instructions): $\langle \text{Ev L1D_CACHE_MISS_LD_NONSPEC} \rangle / \langle \text{Ev INST_ALL} \rangle$	Proportion retired L1D Cache loads that miss the cache of all instructions.
Store L1D Cache Miss Retired Density (of instructions): $\langle \text{Ev L1D_CACHE_MISS_ST_NONSPEC} \rangle / \langle \text{Ev INST_ALL} \rangle$	Proportion retired L1D Cache stores that miss the cache of all instructions.
Load L1D Cache Miss Retired Density (of retired loads): $\langle \text{Ev L1D_CACHE_MISS_LD_NONSPEC} \rangle / (\langle \text{Ev INST_INT_LD} \rangle + \langle \text{Ev INST_SIMD_LD} \rangle)$	Proportion retired L1D Cache loads that miss the cache.
Store L1D Cache Miss Retired Density (of retired stores): $\langle \text{Ev L1D_CACHE_MISS_ST_NONSPEC} \rangle / (\langle \text{Ev INST_INT_ST} \rangle + \langle \text{Ev INST_SIMD_ST} \rangle)$	Proportion retired L1D Cache stores that miss the cache.
Load L1D Cache Miss Density (of load pipeline accesses): $\langle \text{Ev L1D_CACHE_MISS_LD} \rangle / \langle \text{Ev LD_UNIT_UOP} \rangle$	Proportion speculative load pipeline accesses that miss the L1D Cache. Accesses may count more than once if the load is replayed within the Load and Store Execution Unit or is split across a cacheline, and includes prefetches and other operations that may use the pipeline.
Store L1D Cache Miss Density (of store pipeline accesses): $\langle \text{Ev L1D_CACHE_MISS_ST} \rangle / \langle \text{Ev ST_UNIT_UOP} \rangle$	Proportion speculative store pipeline accesses that miss the L1D Cache. Accesses may count more than once if the store is replayed within the Load and Store Execution Unit or is split across a cacheline, and includes prefetches and other operations that may use the pipeline.
L1D Dirty Writeback Density (of load and store pipeline accesses): $\langle \text{Ev L1D_CACHE_WRITEBACK} \rangle / (\langle \text{Ev LD_UNIT_UOP} \rangle + \langle \text{Ev ST_UNIT_UOP} \rangle)$	Proportion speculative load and store pipeline accesses that result in writeback of dirty data out of the L1D Cache toward the L2 Cache.

4.6.4 Shared L2 Cache

Each cluster of cores uses a Shared L2 Cache which operates on 128B cachelines. It is shared between instructions and data, and between all cores in the cluster. Load mops that hit in the Shared L2 Cache execute with an average latency of 15 cycles under unloaded conditions. Some loads may be a few cycles faster and some slower, depending on microarchitectural conditions.

Table 4.20. Shared L2 Cache Organization

Chip	Capacity, Associativity, and Line Size	
	Performance Cluster	Efficiency Cluster
M1 Generation	12MiB, 12-way, 128B lines	4MiB, 16-way, 128B lines
M2 Generation	16MiB, 16-way, 128B lines	
M3 Generation		
A14 Bionic	8MiB, 16-way, 128B lines	
A15 Bionic	12MiB, 12-way, 128B lines	
A16 Bionic	16MiB, 16-way, 128B lines	

Note: Capacity and associativity specifications are provided for reference. Software should check the appropriate `sysctl` parameter to dynamically adjust to chip configuration. See [Appendix B, Dynamic Determination of Chip-Specific Capabilities](#) for more information.

Shared L2 Misses may be serviced by either the memory system, including the Memory Cache (See [Section 4.6.5, "Memory Cache"](#)), or another cluster's caching hierarchy depending on data ownership. Load μ ops that hit in another cluster's caching hierarchy execute with a latency in the range of 50ns, which is closer to that of the Memory Cache than the local cluster's L2. Heavy system traffic may increase these latencies.

4.6.5 Memory Cache

Attached to the DRAM interface is a Memory Cache (M Cache) operating on 128B cachelines. It is shared by all agents on the chip that can access DRAM directly, including the GPU and other co-processors and fixed function hardware. Load μ ops that hit in the M Cache execute with a latency in the range of 35ns, while load μ ops that access DRAM will execute with latency in the range of 95ns. These latencies may increase if the system is experiencing heavy traffic.

Table 4.21. Memory Cache Organization

Chip	Capacity, Associativity, and Line Size
M1	8MiB, 16-way, 128B lines
M1 Pro	24MiB, 12-way, 128B lines
M1 Max	48MiB, 12-way, 128B lines
M1 Ultra	96MiB, 12-way, 128B lines
M2	8MiB, 16-way, 128B lines
M2 Pro	24MiB, 12-way, 128B lines
M2 Max	48MiB, 12-way, 128B lines
M2 Ultra	96MiB, 12-way, 128B lines
M3	8MiB, 16-way, 128B lines
M3 Pro	12MiB, 16-way, 128B lines (reduced)
M3 Max	48MiB, 12-way, 128B lines
A14 Bionic	16MiB, 16-way, 128B lines
A15 Bionic	32MiB, 16-way, 128B lines

Table 4.21. Memory Cache Organization (cont.)

Chip	Capacity, Associativity, and Line Size
A16 Bionic	24MiB, 12-way, 128B lines (reduced)

Recommendation: Perform Cache Blocking Optimization Based on L1D and L2 Shared Cache Sizes:

[Magnitude: Medium | Applicability: Medium] Some algorithms benefit from dividing up the data set into blocks that will fit into the cache, then operating on the blocks one-by-one. Do not block algorithms based on the M Cache capacity. The M Cache is shared amongst several agents and large portions of the cache may be consumed by those agents.

4.6.6 Improving Cache Hierarchy Performance

Consider these data layout optimizations to improve cache hierarchy performance:

- **Eliminate false sharing:** A performance bottleneck can occur when two or more independent variables are allocated to the same 128B cacheline and are modified by threads running on different core clusters. When a core in one cluster writes to one of the independent variables, it will cache the line in its own Shared L2 Cache and L1D Cache. This forces the line to be evicted from all other caches in the system. Meanwhile, a core in the other cluster can attempt to write to the second variable, request the cache line, and the caching subsystem will move the entire line containing both variables over to its Shared L2 Cache and L1D Cache. If the first core continues to modify its variable, the line will then move back to the first core, and so on as the cycle repeats.

This unnecessary cache line *thrashing*, or rapid and repeated movement of a cache line back and forth between the clusters, can result in poor multithreaded performance. To avoid this problem, simply add padding as needed between any independent variables that may be used simultaneously by more than one core, thereby ensuring that they are always allocated to different 128B cachelines. In properly tuned code, any remaining cache line thrashing should only occur when “true” sharing occurs, when *individual* variables are actively being shared by multiple cores.

- **Use cache line sharing constructively:** Having multiple variables packed together in the same cacheline can be beneficial. If two or more variables are commonly used together by one core, but seldom used by different cores at the same time, pack them into a single cache line to improve performance. Pack variables or data structure elements that are accessed together into one 64B cacheline to match the L1D Cache line size. When one datum is requested, the related data will naturally arrive with it in the same cacheline, avoiding further cache misses. Data packed into the same 128B cacheline will also benefit from constructive sharing when this larger line is cached in the Shared L2 Cache, although two L1D misses may be incurred to bring both constituent 64B lines into the L1D Cache from the Shared L2 Cache.

More broadly, identify the most commonly accessed (“hot”) variables in structures and place them together so they fall in the same cache line(s). Afterwards place infrequently used (“cold”) variables where space allows. Splitting “hot” from “cold” in

this manner may reduce the working set size, and thus improve the effectiveness of the cache, because only the “hot” portion of the structures will occupy cache space in the common case. For example, group all of the pointers interlinking data nodes together at the beginning of nodes for data structures such as linked lists or trees, so that only one cache line per node loads during a typical list or tree traversal.

- Order data fields to minimize padding:** Some high-level languages do not allow the compiler to reorder individual data elements from that specified in the source, such as fields in a structure in C. Further, these fields are also often aligned according to their size (1B elements to 1B boundaries, 4B elements to 4B boundaries, etc.). To achieve this alignment, padding bytes may be present in between fields. Order data elements such that few gaps are present from the end of one element to the start of the next, according to their natural alignment. In this example, placing the 1 byte `bool` elements next to each other reduces padding bytes and overall structure size:

```
typedef struct _mystruct {           // sizeof() = 12
    bool valid;                     // byte 0
    int data;                        // [3 byte gap]; bytes 4-7
    bool ready;                      // byte 7
} mystruct;

typedef struct _mystruct2 {         // sizeof() = 8
    bool valid;                     // byte 0
    bool ready;                     // byte 1
    int data;                       // [2 byte gap]; bytes 4-7
} mystruct2;
```

False sharing often causes dramatic performance loss, and fixing such issues can lead to significant performance improvement. Typically, identify and fix these issues first. Then identify hot variables, placing them together in a compact order according to their natural alignment. Pack cold variables into gaps to eliminate padding. Finally, pack the remaining cold variables.

For information on the sizes of various data types, see [Table 3.11: “Advanced SIMD Intrinsic Data Types”](#) and [Writing Arm64 Core for Apple Platforms](#).

Recommendation: Avoid False Sharing by Allocating Independent Shared Variables to Different 128B Cachelines:

[Magnitude: High | Applicability: Low] Avoid false sharing bottlenecks by allocating each independent shared variable to a different 128B cacheline. This avoids situations where a cached shared variable, possibly related a software semaphore, is not inadvertently invalidated from the cache due another core using a different shared variable. Thrashing in this manner may significantly slow multi-threaded performance.

Recommendation: Pack Hot Variables into the Smallest Set of Cachelines for Improved Cache Hierarchy Performance, Concentrating Data Commonly Used Together into the same 64B Cachelines.:

[Magnitude: High | Applicability: Medium] Identify and place the commonly used variables into cachelines in a compact order to reduce padding according to natural alignments. Place cold variables into gaps, and pack remaining cold variables adjacent to the hot variables.

- Worst Performance:** When a younger load is predicted to not depend on an older store and is allowed to speculatively execute. As noted above, when the store's address is computed and the dependence violation is detected, the processor flushes the pipeline and restarts the load and subsequent instructions. This memory order violation leads to wasted execution resources both to perform the flush, as well as to re-execute of all μ ops that had speculatively executed.
- Poor Performance:** When a younger load is predicted to depend on an older store, but actually does not. These "phantom" dependences unnecessarily delay load execution.

In some cases, the predictor must decide whether or not to insert a dependency before data addresses can be computed for both the load and store in a particular store-to-load pair. As a result, the processor must make an educated guess whether or not a dependence may occur based on the past behavior for that pair of instructions. Unfortunately, it may predict too conservatively when a store-to-load pair is only occasionally dependent. Pipeline flushes are quite expensive, so it may favor the modest expense of enforcing dependences. This might occur when address pointers for the load and store are briefly sweeping past each other or when buffers only partially overlap, but usually the instructions are not dependent. Thrashing may occur if the dependent and not-dependent cases both regularly occur for the same store-load pairs. This may result in patterns where dependences are enforced when none are needed and never enforced when they are needed, leading to significant slowdown.

Table 4.24. Store-to-Load Dependence Prediction Outcome

		Predicted Outcome	
		Not Dependent	Dependent
Actual Outcome	Load Not dependent on Store (Non-Overlapping Memory Range)	High Performance: Loads and stores may execute at any time without conflict	Poor Performance: While a load was dependent on a store at some point in the past, this overly conservative prediction unnecessarily forces further instances of that load to stall and wait for store completion
	Load Dependent on Store (At Least Partially Overlapping Memory Range)	Worst Performance: Memory Order Violation: If the older store's address is unknown when the younger load is ready to execute, the younger load may prematurely execute and require a pipeline flush with restart. OR Modest Performance: If the older store's address is known, the processor may stall the load until the store is known.**	Modest Performance: Loads are delayed until just after stores complete, properly honoring dependencies but not stalling excessively. **

case, these legitimate dependencies tend to train the predictor conservatively, so that store-to-load dependence enforcement still occurs during later FFT passes. But during these later passes, the working set is larger, and so the reuse of the same data elements are spaced out much farther apart in time and seldom cause true dependencies.

- **Small, in-place, two-dimensional buffers:** In operations involving two-dimensional buffers, nearest neighbor elements are often read as input to help update the value of element (m, n) . If the elements are modified in-place, then element $(m, n-1)$ will be read back in to help compute (m, n) shortly after $(m, n-1)$ was written. Depending on the operation and data set size, the instructions to compute an entire row may fit into a core's active instruction window. When this happens, store-to-load forwarding scenarios can occur as the loads from (m, n) start executing before the stores from $(m, n-1)$ have completed. Similar scenarios can be constructed for in-place buffers with larger numbers of dimensions.
- **Passing Pointers:** If a buffer is being accessed and updated by multiple pointers that are incremented by different amounts during each loop iteration, they may cross at some point during the loop. A one-time flush will occur after the pointer collision, which will then be followed by a number of iterations of overly conservative dependence enforcement.
- **List of Pointers to a Small Object Pool:** Perhaps the most pernicious access pattern can be triggered when loop iterations must perform read-modify-write updates to a small pool of data structures based on a long list of pointers. This pattern is quite common in graphics or physics algorithms that maintain a pool of objects and then a larger list of "links" between objects that are close to each other in space and need to be considered during each simulation step. When the same object just happens to be used in calculations during adjacent loop iterations, then dependence collisions and flushes are quite likely. This sort of hard-to-predict dependence behavior can easily cause thrashing between overly-conservative serialization and too little serialization.
- **Frequently modified top of stack:** In code that regularly pushes data to a stack data structure and then quickly pops it back off again, load-store dependence violations are likely. Moreover, if these locations on the stack are then used to hold different data variables for other functions soon thereafter, then the stale store-to-load predictions trained into the predictor for the earlier stack contents may trigger excessive serialization of loads and stores for these later uses of the stack. On the other hand, address calculations for stack accesses are often quite simple, so store and load addresses can often be calculated quickly enough to avoid relying on the predictor.

While these scenarios are not particularly common, speculative store-to-load prediction thrashing can seriously degrade performance. Production code examples have been seen that reduce performance by 50%.

Consider these strategies to avoid dependence prediction problems:

- **Avoid Read-Modify-Write Data Accesses:** The simplest way to avoid speculative store-to-load thrashing is to avoid looping algorithms that require read-modify-write access to data structures that are too large to be kept in registers. Algorithms where the input and output buffers for each loop are separate naturally avoid memory dependence violations. In contrast, violations may occur when one data structure

iterations access the same objects and shuffling them farther apart in the operation sequence.

- **Unroll Loops:** In some algorithms, a load from every n th iteration of the loop consumes data from a prior store. Unrolling the loop n times creates n copies of the loads and stores and stabilizes the dependence pattern. 1 copy of the load always obtains data from a prior store, while $n-1$ copies always obtain data from memory. Unrolling may generally help even when such a pattern is not obvious. To an extent, more copies of the loads and stores create more opportunities for the predictor to find stable dependence patterns between them.

These improvements or others like them will not always be possible, either because the original algorithm is not amenable to change or because the changes cause performance degradation for other reasons, such as lower cache hit rates. However, if they are possible, then performance of the algorithm will often increase in *any* out-of-order core. Application of these techniques can result in performance increases of 50-150% on programs that were originally limited by speculative store-to-load prediction thrashing.

Recommendation: Reduce Hard-to-Predict Store-to-Load Forwarding:

[Magnitude: High | Applicability: Medium] Avoid algorithms where loads occasionally read recently written data. Consider restructuring data accesses to increase the number of dynamic instructions between the write and the subsequent read, so the read will find the data already written into the cache. Alternatively, consider restructuring algorithms to avoid sporadic reads after writes, potentially leveraging registers to communicate the data.

When multiple in-flight stores simultaneously forward bytes to a single load, the predictor may be unable to properly enforce dependence ordering between loads and individual stores. For instance, a string buffer may be written using `STRB` instructions and subsequently loaded as a 16-byte aligned vector, instead. In these cases, the predictor may force the load to wait until all older stores have issued instead of being dependent on a specific store. Nevertheless, performance is still usually better than if the load receive incorrect data, and the processor is forced to flush the pipeline and re-execute the load. Because of the large instruction window, stores several hundred instructions older than the load may not yet have written the cache when the load is read to execute.

Recommendation: Reduce Scenarios Where Multiple Smaller Stores Simultaneously Forward to a Single Larger Load:

[Magnitude: Medium | Applicability: Medium] Minimize cases where memory is accessed as two differently-sized data types, especially if those accesses may occur within a several hundred instructions of each other.

Table 4.25. Common Memory Dependence Metrics

Name and Formula (Event Definitions: Section 6.2, "Performance Monitoring Events")	Description
Memory Order Violation Density: $\frac{\langle \text{Ev ST_MEMORY_ORDER_VIOLATION_NONSPEC} \rangle}{\langle \text{Ev ST_UNIT_UOP} \rangle}$	Proportion retired stores that caused a pipeline flush and replay of loads and subsequent instructions of all stores.

4.6.10.1 Memory Dependence Prediction Example

Consider the following code example. The "Add AND assignment" operation (+=) on `autoc[coeff]` performs a load of the existing value, adds to it, and stores out the new value. As long as the value of `d` remains ≥ 0 , the load will need the value from the previous store. But, in some iterations when $d < 0$, `coeff` index will be incremented such that the load will need the value from memory as opposed to the previous store.

```

void filter(float *autoc, float *data, uint32_t data_len, uint32_t lag)
{
    float d;
    uint32_t sample, coeff = 0;
    const uint32_t limit = data_len - lag;

    for(coeff = 0; coeff < lag; coeff++)
        autoc[coeff] = 0.0;

    for(sample = 0; sample <= limit; sample++) {
        d = data[sample];
        if (d < 0) {
            coeff++;
            if (coeff == lag)
                break;
        }
        autoc[coeff] += d * data[sample+coeff];
    }
}

```

Knowledge of the data and algorithm may be useful when rewriting code to eliminate the memory dependences. In this example, all updates to a particular `autoc[coeff]` happen sequentially. The algorithm can be altered to load the initial value of `autoc[coeff]` into a local variable `sum`, accumulate all updates into that local variable, then write out that local variable into `autoc[coeff]` when the algorithm moves onto the next `coeff`. The compiler is likely to register allocate the local variable `sum` eliminating the per-iteration memory loads and stores altogether.

```

void filter_opt(float *autoc, float *data, uint32_t data_len, uint32_t lag)
{
    float d;
    uint32_t sample, coeff = 0;
    float sum = 0.0f;
    const uint32_t limit = data_len - lag;

```

```
for(coeff = 0; coeff < lag; coeff++)
    autoc[coeff] = 0.0;

for(sample = 0; sample <= limit; sample++) {
    d = data[sample];
    if (d < 0) {
        autoc[coeff] = sum;
        coeff++;
        sum = 0.0f;
        if (coeff == lag)
            break;
    }
    sum += d * data[sample+coeff];
}

if (coeff < lag)
{
    autoc[coeff] = sum;
}
}
```

4.6.11 Store-to-Load Forwarding

When the processor predicts that a load will read data written all or in part by an older store, the load will wait to execute until the data portion of the store is available. In such cases, the load may read its data from internal buffers prior to the store actually writing the cache. This is commonly referred to as Store-to-Load Forwarding. The P cores feature an aggressive algorithm that allows loads to read data from multiple sources generally without additional delay, including reading some bytes from various younger stores (that have not yet written to the cache) and bytes from the cache itself.

While data alignment for store-to-load-forwarding isn't often a limiter, cache line and page spanning loads may experience delays. (See [Section 4.6.9, "Unaligned Accesses"](#)). Similarly, when multiple stores simultaneously feed a single load, the dependence predictor may enforce conservative dependencies, requiring the store's addresses to be known prior to forwarding. (See [Section 4.6.10, "Memory Dependence Prediction"](#)).

In some instances, the processor may be able to further optimize Store-to-Load Forwarding when resources are available:

- Single element 4B or 8B integer loads and stores only (without sign extension).
- Simple base + offset addressing only

Recommendation: Leverage Flexible Store-to-Load Forwarding When Register Communication is Complex:

[Magnitude: Low | Applicability: Low] For algorithms with complex memory access patterns, forwarding data via registers might require complex data selection and shift/merge sequences. Rely on the aggressive store-to-load forwarding network to complete such data movement. Further, rely on optimized store-to-load forwarding of whole integer registers with simple addressing in situations such as register spills/fills.

4.6.12 Prefetching

For memory-intensive applications, prefetching data into the caches from main memory prior to use can result in dramatic performance increases. To accommodate this, the cores include mechanisms for both hardware-driven and software-driven prefetching. This section gives a brief description of how to avoid pitfalls that might cause significant performance problems.

4.6.12.1 Hardware Prefetcher

The hardware prefetcher is a unit attached to each core that analyzes the cache misses generated by that core in an effort to find predictable streams of misses. When its heuristic identifies a predictable stream, it attempts to launch main memory accesses in advance in order to avoid future cache misses to the same stream of references. In general, any large data structures accessed in a linear fashion should be successfully prefetched, and even somewhat irregular patterns are also often quite prefetchable. In general, it is successful on all but the most irregular and hard-to-predict scenarios. Nevertheless, the hardware prefetcher does have a few limitations, such as the maximum stride that can be detected successfully.

Benefits to relying on the hardware prefetcher:

- **No developer intervention:** This mechanism is automatic and does not require any code modifications.
- **Automatically tuned for each core:** The hardware prefetching aggressiveness is tuned for each core's performance characteristic and each chip's latencies, so that data arrives in the core just before it is used. In contrast, the count of instructions between software prefetches and the subsequent demand accesses must be manually tuned for each core to ensure that they are early enough, but not too early. Because the performance cores usually require prefetching 4 to 30 lines ahead to allow the data to arrive in time (depending on the specific software algorithm), this can be quite difficult to tune manually. In contrast, the efficiency cores usually require less aggressive prefetching. When the same code is executed on the efficiency cores, the prefetcher's aggressiveness is retuned for those characteristics. Such adjustment is not possible utilizing software prefetches.
- **Reduced resource contention with loads and stores:** Software prefetch instructions occupy load units in the pipeline just like any other load. Reducing the number of instructions that must go through those units can help avoid turning them into a critical bottleneck. In contrast, the hardware prefetcher usually only inserts prefetches into the load or store pipelines on spare idle cycles. Hardware prefetches may sometimes still compete for memory system resources with demand accesses, but the hardware prefetcher algorithm adjusts for that too. Software prefetches are *always* mandatory and must execute even in the face of heavy memory system contention. In contrast, hardware prefetches are always optional and can be dropped if necessary.

4.6.12.2 Software Prefetch Instructions

The cores support the `PRFM` and `PRFUM` software prefetch instructions, which differ only in their available selection of addressing modes. The former have the same addressing

modes as `LDR` or `STR`, while the latter handles modes available to `LDUR` or `STUR`. These prefetch instructions include prefetch options that include operation type {prefetch load, prefetch store}, data destination {cache levels 1 or 2}, and temporal behavior {temporal, non-temporal}. When using software prefetch instructions, issue only 1 instruction per cacheline, striding the address based on the particular cache's line size. The address specified within a particular cacheline does not matter.

Table 4.26. Recommended Software Prefetch Instructions

Instruction	Operation	Destination	Temporal Behavior
<code>PRF (U)M PLDL1KEEP</code>	Load / Read	L1D Cache	Temporal
<code>PRF (U)M PLDL1STRM</code>			Non-Temporal
<code>PRF (U)M PLDL2KEEP</code>		Shared L2 Cache	Temporal
<code>PRF (U)M PLDL2STRM</code>			Non-Temporal
<code>PRF (U)M PSTL1KEEP</code>	Store / Write	L1D Cache	Temporal
<code>PRF (U)M PSTL1STRM</code>			Non-Temporal
<code>PRF (U)M PSTL2KEEP</code>		Shared L2 Cache	Temporal
<code>PRF (U)M PSTL2STRM</code>			Non-Temporal

Unlike normal memory accesses, these instructions cannot trigger exceptions (e.g., unmapped virtual to physical addresses, illegal page access request types, etc.). As a result, software prefetch instructions may safely fetch past the ends of arrays and access `NULL` pointers, without slowing down the system with spurious faults. Because prefetch instructions often access data a few iterations ahead of the actual demand accesses in the loop, preventing "bad" accesses would require extra range checks and branches that would complicate and slow the loop. Unfortunately, this feature also prevents prefetch instructions from triggering true page faults early; only actual demand misses can force the core to handle page faults.

When the hardware prefetcher can successfully predict memory access patterns, inclusion of software prefetch instructions may cause a decrease in performance because the software prefetch instructions may inhibit the abilities of the hardware prefetcher. The prefetch instructions don't trigger the hardware prefetcher themselves, so their use can mask any natural miss patterns in the code that would normally trigger the detection of hardware prefetchable patterns. In addition, the prefetch instructions tie up resources within the cores, which may increase competition for critical pipeline resources. In particular, they occupy decode slots, take up space in the reorder buffer, and can compete with actual loads and stores for issue slots in the Load and Store Units.

Furthermore, you may find it difficult to properly place software prefetches into your code. To be fully effective, software prefetches must launch memory requests a hundred or more cycles ahead of the natural use in the code. Thus developers need to insert software prefetch instructions potentially hundreds of dynamic instructions ahead of the use. This distance is likely to increase on future faster processors.

Recommendation: Springly Use Software Prefetch Instructions:

[Magnitude: Medium | Applicability: Low] Consider adding software prefetch instructions to improve load and store latency, but only after analyzing data miss patterns and exploring algorithmic improvements to better facilitate hardware prefetching. Select the appropriate software prefetch instruction based on expected demand operation type and temporal locality.

Before inserting software prefetches, evaluate cache performance using the metrics listed in [Table 4.19: “Common Data Cache Metrics”](#). Identify loads and/or stores that commonly miss and analyze the access patterns. The follow scenarios may warrant use of software prefetching. After attempting insertion of software prefetches, it is best to analyze performance both with and without the software prefetches to see whether or not they are properly placed and verify that they provide an actual performance improvement.

- **Any access pattern:** The biggest advantage for software prefetches is that they can prefetch data using any memory access pattern, while the hardware prefetcher is limited to a finite number of repetitive patterns. Hence, prefetching in complex data structures such as trees is really only possible using software prefetching. That being said, when accessing these sorts of data structures it is often too difficult or sometimes even impossible to calculate the necessary addresses sufficiently early. For example, pointer-chasing delays usually cannot be avoided by prefetching, because the program doesn't know the addresses to prefetch in advance.
- **Multiple streams too close together:** The hardware prefetcher is designed to detect streams while accounting for out-of-order execution of accesses. When multiple separate streams are located close together, typically on the same page, and are accessed simultaneously in an interleaved manner, the prefetcher may sometimes identify them as a single irregular stream. The prefetcher may then initiate either erratic prefetching or, in the worst case, no prefetching at all. For example, significant performance degradation can occur when a single loop reads values from adjacent rows in 2D data buffers (matrices, photos, and video frames, for example). The full buffers in these scenarios are usually large enough to easily exceed the sizes of the caches, so prefetching is important for good performance. However, rows of these large buffers can still be short enough so that streams of accesses to two adjacent rows of the buffer at once can be misinterpreted as a single irregular stream. There are several ways to address this problem, so it is covered in more detail later in this section.
 - **Large strides:** Software prefetching can also help if large strides are needed, typically those nearing the size of memory page or larger. Any stride larger than about a memory page may potentially be missed and tracked as separate streams, instead. For example, a loop accessing a large (i.e. thousands x thousands of elements) row-major array in a column-wise manner will often read the array using strides that greatly exceed the maximum stride that the hardware prefetcher can detect.

4.6.12.3 Prefetch Performance Tips

While rare, the most difficult-to-debug problems with the hardware prefetcher tend to occur when software simultaneously accesses multiple streams that are too close together in memory. To avoid these problems, don't mix two or more simultaneous

streams of cache misses within the same memory page. Consider these transformations to improve the ability of the hardware prefetcher to lock on to streams:

- **Avoid adjacent rows:** If possible, change the algorithm so that it does not need to access adjacent rows of a 2D buffer simultaneously. For example, consider an algorithm that operates row-by-row over a large 2D buffer. You might be tempted to unroll the outer loop such that the algorithm operates on two rows (m and $m+1$) simultaneously, exposing parallelism, like the following:

```
// Fused outer loop, loops over even/odd pairs of rows together
FOR (m = 0; m < M; m += 2) {
    // Inner loop over columns of the buffer
    // -- Cache misses on both rows m & m+1, interleaved
    FOR (n = 0; n < N; n++) {
        Compute row pair output elements (m,n) and (m+1,n)
        // -- Can interleave instructions from both rows together now
    }
}
```

Unfortunately, this loop accesses two rows of the 2D buffer simultaneously. If both rows are located in the same memory page, the hardware prefetcher may be unable to clearly identify the streams.

However, because the algorithm processes each row independently, pick two non-adjacent rows to process simultaneously. For example, instead of adjacent rows, interleave a row from the top half of the buffer (m) and a row from the bottom half ($m+(M/2)$):

```
// Fused outer loop, loops over pairs of rows together
FOR (m = 0; m < M/2; m++) {
    // Inner loop over columns of the buffer
    // -- Cache misses on both rows m & m+M/2, interleaved
    FOR (n = 0; n < N; n++) {
        Compute row pair output elements (m,n) and (m+(M/2),n)
        // -- Can interleave instructions from both rows together now
    }
}
```

Simple changes like this can ensure that any two (or more) rows being read simultaneously are always far apart instead of just a single row apart.

- **Rearrange access patterns:** The next possible solution is to rearrange the accesses to memory so the cache misses all occur in a single stream that the hardware prefetcher can easily understand. Instead of interleaving two streams in confusing patterns like $x, X+M, x+1, X+M+1, x+2, X+M+2, \dots$, for rows of length M , one can instead rearrange the references into streams like $x, x+1, x+2, \dots, x+M-1, X+M, X+M+1, X+M+2, \dots$. The simplest way to do this is to add a small “read ahead” loop that triggers the cache misses ahead of time for all rows or all but the last row. Here is a simple example for when two rows are read together that triggers the misses to the first of the two rows early:

```
// Original outer loop, loops over even/odd pairs of rows
FOR (m = 0; m < M; m += 2) {
```

```

// Added "read ahead" loop
FOR (n = 0; n < N; n += L1_CACHE_LINE_SIZE) {
    Load at (m,n) to trigger cache misses from row m in correct order
}
// Original inner loop, loops over columns of the buffer
// -- Originally missed on both rows m & m+1, now just misses on m+1
FOR (n = 0; n < N; n++) {
    Compute row pair output elements (m,n) and (m+1,n)
}

```

This small addition reorders all of the cache misses so they will occur in the “correct” order and thereby generally present the hardware prefetcher with a single, simple stream of misses. Depending upon the nature of the inner loop, it may sometimes be better to load ahead for both rows m and $m+1$ in a read ahead loop like the following:

```

// Added "read ahead" loop
FOR (n = 0; n < 2*N; n += L1_CACHE_LINE_SIZE) {
    Load at (m,n) to trigger cache misses from row m and then row m+1
}

```

After the processor executes the “read ahead” loop, the computation loop won’t generate any cache misses at all, because all of its data will have been fetched into the cache by the “read ahead” loop. Experiment with different solutions to see which performs better in a given situation.

- Partial software prefetch:** Typically, read ahead loops use normal “demand” accesses which leverage all of the hardware prefetcher’s advantages, such as having the prefetch aggressiveness automatically scaled properly for the underlying core. However, read ahead loops using demand accesses may stall the processor waiting for load data that will immediately be discarded. In select cases, use software prefetch instructions in the “read ahead” loop to avoid these stalls. In this case, the read ahead loop looks slightly different:
-

```

// Added "read ahead" loop
FOR (n = 0; n < N; n += L1_CACHE_LINE_SIZE) {
    PRFM at (m+2,n) to prefetch from row m+2
}

```

When coded with demand accesses, the read ahead loop accesses data used in the subsequent loop iteration. The hardware prefetcher identifies the pattern and prefetches out ahead, prefetching data for future iterations. However, software prefetches do not train the hardware prefetcher. A read ahead loop consisting of software prefetches must manually prefetch out ahead for future iterations. This can be seen in the above code example where the `PRFM` instruction accesses $(m+2, n)$.

This is because of a fundamental difference between hardware and software prefetching. The hardware prefetcher spots patterns in a stream of accesses and then automatically tries to get far enough ahead of them to avoid pipeline stalls, but software prefetch instructions only initiate prefetches at the time that they are executed. Because they do not attempt to “look ahead” in any way, code that uses software prefetches must always manually insert them far enough in advance so that any needed cache misses can complete before the eventual demand loads occur. In

this case, one must prefetch for the computation loop of the next outer loop iteration instead of the one that will execute immediately after this “read ahead” loop in order to provide enough lookahead.

While using software prefetch to access the “even” (m) rows, this hybrid scheme continues to access the “odd” ($m+1$) rows using standard demand misses during the computation loop. As a result, the code is actually still reliant on the hardware prefetcher, but it now only has to deal with prefetching a single stream of references (for row $m+1$) during each inner loop and hence should work much better.

- **All software prefetch:** The final step is to just software prefetch everything and remove the hardware prefetcher from the equation altogether. This can be accomplished with a straightforward variation on the last read ahead loop:

```
// Added "read ahead" loop
FOR (n = 0; n < N; n += L1_CACHE_LINE_SIZE) {
    PRFM at (m+2,n) to prefetch from row m+2
    PRFM at (m+3,n) to prefetch from row m+3
}
```

Much like the second read ahead loop variation from the “rearrange access patterns” case, this runs ahead and fetches two rows at once. Unlike that scenario, however, one can fetch the two rows using an interleaved access pattern. Because the code now bypasses the hardware prefetcher entirely, there is no reason to keep the accesses laid out in a strictly linear fashion. (It may still be a good idea to use the linear access pattern from the “rearrange access patterns” case, because it may improve DRAM performance. However, it is not necessary to guide prefetching, so this alternate code works too.) In fact, it may even be desirable to fuse the “read ahead” loop and the computation loop together so that the software prefetches are fed into the Load and Store Units gradually over the course of the inner loop instead of in a big chunk at between inner loops, more like the following:

```
// Original outer loop, loops over even/odd pairs of rows
FOR (m = 0; m < M; m += 2) {
    // Fused read ahead+inner loops, loops over columns of the buffer // --
    Compute rows m & m+1 + prefetch rows m+2 & m+3
    FOR (ns = 0; ns < N; ns += L1_CACHE_LINE_SIZE) {
        // Read ahead section
        PRFM at (m+2,ns) to prefetch from row m+2
        PRFM at (m+3,ns) to prefetch from row m+3
        // Compute section
        FOR (n = ns; n < ns + L1_CACHE_LINE_SIZE; n++) {
            Compute row pair output elements (m,n) and (m+1,n)
        }
    }
}
```

This kind of pattern can be continued to prefetch even farther ahead, but be careful not to prefetch ahead so far that the data starts falling out of the L1D Cache before it can be used by the compute loop. If L1D Cache capacity is a problem, then it would probably be better to go back to one of the solutions that uses the hardware prefetcher. Finally, note that it is not possible to fuse the read ahead loop like this

unless the read ahead loop is only using software prefetches, because otherwise the code will still have multiple load streams and lead to prefetcher confusion.

Few algorithms access multiple rows of 2D buffers at once in ways that will cause cache misses on multiple rows at once like this, but for those that do these techniques can result in very large performance improvements.

Last, neither software nor hardware prefetching is likely to improve the performance of pointer chasing code. That is, it is hard to predict the need for cachelines far enough in advance for algorithms that predominantly hop from one linked node to the next with minimal additional computation. Many of these data structures feature dynamically allocated nodes, for which the allocation and link ordering form no discernible memory address pattern. Performance improvement may still be possible through intelligent ordering of structure fields to minimize the number of cachelines that are needed for typical access to a node. Specifically, put the key node data items and "next" pointer into a single cacheline. See the discussion on "hot/cold" fields and variables in [Section 4.6.3, "L1 Data \(L1D\) Cache"](#).

Recommendation: Recode Algorithms to Simplify Access Patterns to Aid the Hardware Prefetcher:

[Magnitude: Medium | Applicability: Medium] While software prefetch instructions may be able to prefetch complex patterns, it is often difficult to insert them far enough ahead, but not too far ahead, of the demand use. The hardware prefetcher, however, dynamically adjusts according to the circumstances. Before attempting to insert software prefetches, identify regions of high cache miss behavior, and explore pattern simplifications that may allow the hardware prefetcher to identify and prefetch the stream.



Chapter 5. Asymmetric Multiprocessor (AMP) Optimization

As mentioned in the introduction, Apple silicon chips feature two types of general-purpose CPU cores, performance cores (P cores) and efficiency cores (E cores). Both types of cores use the same instruction set architecture and coherently access the same memory, so threads may freely move back-and-forth between the two types of cores. The E cores are physically smaller cores with shorter wires and smaller transistors, but are based on a similar microarchitecture as the P cores, thus workloads optimized for one are expected to perform well on the other. Having both types allows the system to optimize for performance when performance is a priority, and optimize for efficiency (for example, improving battery life) when it isn't. Because E cores offer compelling high performance on their own, certain tasks may perform sufficiently well on the E cores which can free up P cores for other more demanding tasks. Lastly, E cores can be used in multithreaded workloads along with the P cores to add significant additional performance.

Optimizing multithreaded applications requires careful consideration. In addition to the discussion below, see [Tuning Your Code's Performance for Apple Silicon](#) and [Addressing Architectural Differences in Your macOS Code](#) on the Apple developer website.

To programmatically determine the CPU and chip configurations, see [Appendix B, Dynamic Determination of Chip-Specific Capabilities](#).

5.1 Prioritizing Work

Developers can prioritize work by providing hints to the system as to which work items are the most performance-critical, and hence should be favored for execution on the P cores, using the quality of service (QoS) flags. Likewise, developers can indicate which work items are less performance-critical ("background") and thus should be favored for execution on the E cores. Nevertheless, the system may execute performance-critical threads on E cores when there are more such threads than available P cores. Especially when the desired performance level is unspecified, the system monitors applications dynamically and automatically prioritizes CPU-intensive threads for the P cores for maximum performance. Similarly, to save power, the system may assign less CPU-intensive threads to the E cores. Documentation on the use of the QoS classes can be found on the [Prioritize Work at the Task Level](#) page of the [Energy Efficiency Guide for Mac Apps](#) guide.

Recommendation: Use Quality of Service (QoS) APIs to Prioritize Work:

[Magnitude: Medium | Applicability: Low] Software cannot directly control the type of core on which any task executes. However, use Quality of Service classes to provide hints to the system.

5.2 Partitioning Work

When developing applications with large numbers of worker threads that all need maximum performance, consider the implications of asymmetric multiprocessing. At a high level, there are two general approaches used by developers to divide up work on symmetric MP systems, but only one works well on an asymmetric system:

- **Static Partitioning:** The partitioning algorithm divides up work into pre-determined, equal-sized chunks. Software assigns each worker thread a particular set of chunks to complete. Synchronization techniques such as barriers occasionally synchronize all threads, potentially after a phase of execution. While this works well in a symmetric multi-processor system it often doesn't work well in an asymmetric system because the underlying assumption doesn't apply: in a symmetric system, software can divide a task into in equal size chunks because the underlying processing elements are all equally capable. But this is not the case in an asymmetric system.
- **Dynamic Partitioning:** The partitioning algorithm divides up work into tasks, potentially of various sizes, stashing them in a task queue. Each worker thread retrieves a new task from the queue when it finishes the previous task. While this strategy may incur slightly higher task-management overhead, the benefits of dynamically assigning chunks of the work to available threads typically allows the work to be completed faster.

[Section 5.2.1, "Avoid Static Partitioning"](#) describes some of the pitfalls associated with static partitioning, while [Section 5.2.2, "Use Dynamic Partitioning"](#) describes the benefits of dynamic partitioning.

5.2.1 Avoid Static Partitioning

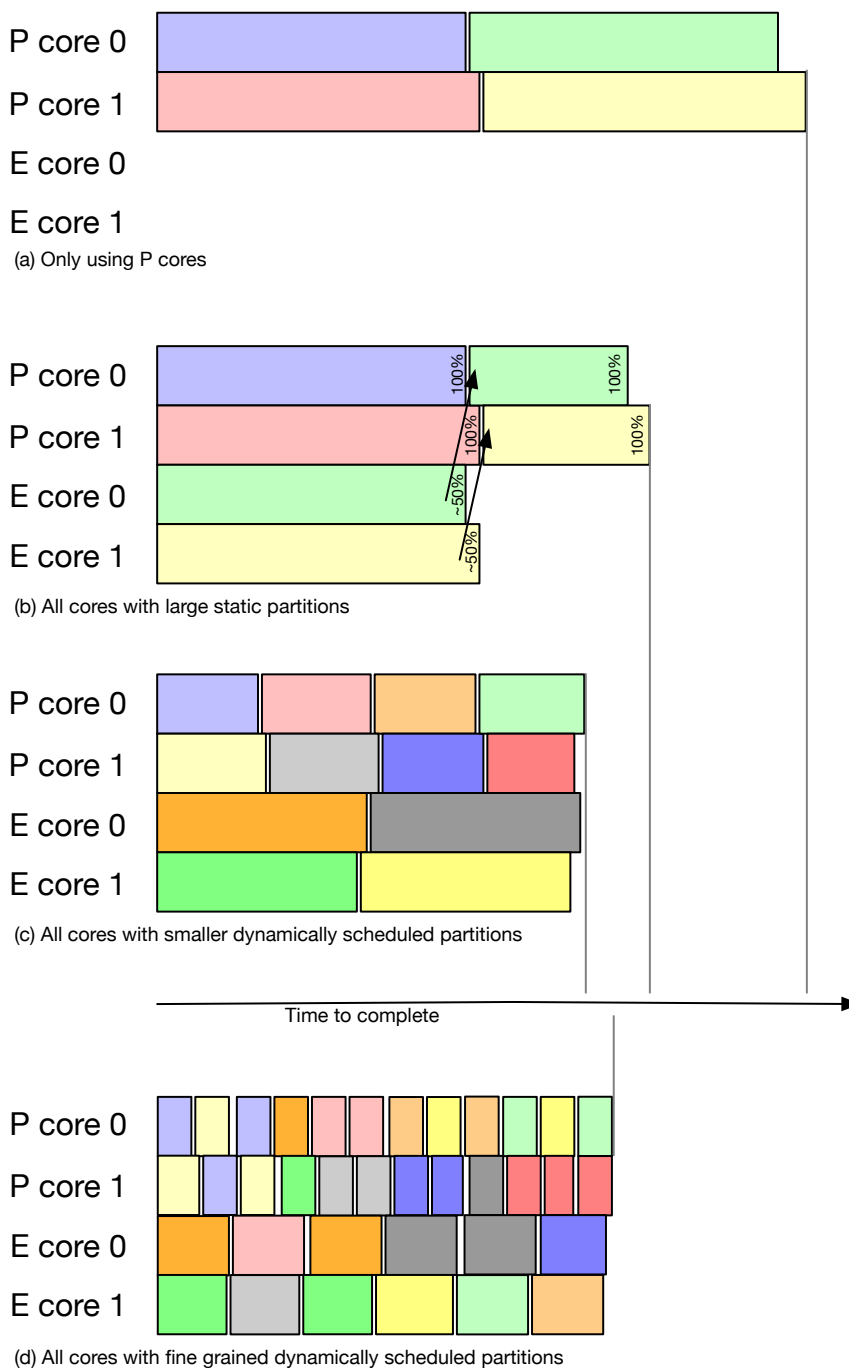
Static partitioning can lead to longer latency execution on asymmetric MP systems. While these equal-sized chunks may all complete in about the same amount of time on a symmetric MP system, asymmetric cores will execute them at different speeds and thereby introduce dynamic load imbalance. Some threads will end up running faster (on P cores) and some slower (on E cores), sometimes switching cores in the middle of execution, usually in a manner that the program itself cannot control.

Because of the asymmetry, threads running primarily on P cores will make more progress than those running primarily on E cores. While the system will likely move those primarily E core threads to the P cores when P cores become available (through rebalancing), those primarily E core threads may already be behind. Note that OS will often rotate threads amongst the P and E cores to load balance when running for multiple OS time quanta.

Similarly, static partitioning strategies tend to assume that little or no other time-consuming tasks are running on the system, and that the static partitioning algorithm can statically and evenly partition work across all of the compute resources. When other unrelated tasks are running on a core, a static slice of the MP application may be delayed, which will delay the overall application in the same manner as static slice running on a slower core. Therefore, even on symmetric multiprocessors, dynamic partitioning is often a better strategy.

If static thread scheduling is the only viable option, use `sysctl` queries to determine the core count. See [Appendix B, *Dynamic Determination of Chip-Specific Capabilities*](#).

Figure 5.1. Partitioning Strategies: An Illustrative Example



As illustrated in [Figure 5.1: "Partitioning Strategies: An Illustrative Example"](#), a dynamic work partitioning strategy will most often lead to the best performance because of its flexible work assignment. Threads synchronized using static techniques like barriers will not work well when some cores can execute code significantly faster than others. Management of multiprocessing at a low level can be accomplished through C++ threads or [pthreads](#). These allow for direct creation and control of multiple POSIX-compatible threads managed by the application, usually with at least one thread per

available core. Distribute work amongst these threads by building dynamic task queues on top of these APIs using basic primitives such as locks and condition variables. However, because this code is complex and often requires extensive tuning, consider a pre-built task management API like [Grand Central Dispatch](#). This framework provides mechanisms to manage and synchronize large numbers of concurrent parallel tasks and automatically handles low-level thread management.

Recommendation: Use Grand Central Dispatch for Dynamic Work and Thread Management:

[Magnitude: Medium | Applicability: Low] Avoid static work partitioning techniques that assign equal portions of the work to each thread, with one thread per core. Instead, divide the work into more (smaller) pieces, 3x the number of cores as a starting point. Use dynamic techniques that allow faster threads to pull more work from a shared work queue. Consider using Grand Central Dispatch which is tuned for each platform to offer the best performance and fairness. C++ threads and pthreads are also available as cross-platform alternatives.

5.3 Avoid Spin-Wait

Regardless of work partitioning strategy, avoid keeping threads active but effectively idle in spin-wait loops. In the rare circumstance the spin-wait is known to be short lived, where the wait time is on the same order as the thread scheduling overhead, spin-wait may be a workable choice. However, most wait times are unknown and often longer than the scheduling overhead. By occupying cores, especially P cores, with spin-wait loops, work that might have fallen behind on E cores may be forced by the system to continue on the E cores. Likewise, the system often has other unrelated work to perform. Spin-waiting prevents the system from using that effectively idle time for other system work, and may force a context switch of application work at less optimal times. See "Don't Keep Threads Active and Idle" at [Tuning Your Code's Performance for Apple Silicon](#). –

Recommendation: Block Threads When Idle and Avoid Spin-Wait Loops:

[Magnitude: Medium | Applicability: Low] When the wait is expected to be very short, on par with the cost of a thread switch, use spin-wait loops to cause a thread to wait for next stage. In all other cases, block the thread such that the operating system can schedule other work on the core.

5.4 APIs for Synchronization and Thread Communication

Experienced MP developers may be tempted to try to achieve maximum performance by building their own primitives (e.g. spin locks) from scratch using Arm ISA Load/Store exclusive or atomic instructions. However, these primitives require complex algorithms (e.g., queued locks) to ensure fairness between all threads, avoiding bias (or even monopolization) to cores within a cluster. Also, topologies, protocols, and latencies are likely to change between generations and even between different products within a particular generation. Custom synchronization primitives tuned for one product may not function well in another.

Last, avoid false sharing, where two or more independent variables occupy the same 128B cacheline. This organization may lead to thrashing of the cacheline between cores. See [Section 4.6.6, “Improving Cache Hierarchy Performance”](#).

Recommendation: Minimize Data Sharing When Increasing Thread Counts:

[Magnitude: Medium | Applicability: Low] Sharing modified data between cores in a cluster is relatively fast due to the Shared L2 Cache. When expanding a workload beyond the number of threads found in a single cluster, consider how much data is written by one core and consumed by another. Excessive sharing may lead to lower-than-expected performance especially with increased thread counts. Reduce the amount of data shared where possible, either algorithmically or via increased data packing within cachelines without introducing false sharing.

5.6 Xcode Sanitizer Tools

Xcode offers a number of tools to aid in multithreaded code development. The [Thread Sanitizer](#) tool inserts diagnostics into the code to record each memory read or write operation. These diagnostics generate a timestamp for each operation, as well as its location in memory. The tool then reports any operations that occur at the same location at approximately the same time. The tool also detects other thread-related bugs, such as uninitialized mutexes and thread leaks.

Recommendation: Use Thread Sanitizer and Other Xcode Tools to Aid Multithreaded Code Development:

[Magnitude: Medium | Applicability: Low] Correct multithreaded applications require the developer to pay careful attention to synchronization, data ordering, and leaks. Use available tools to verify correctness.



Chapter 6. Performance Monitoring

6.1 Xcode Instruments Tool

Xcode features an [Integrated Development Environment \(IDE\)](#) tool called **Instruments**. This tool can help you profile your apps in order to better understand and optimize their behavior and performance.

For more information on how to use Instruments, see [Getting Started with Instruments \(WWDC2019 Video\)](#).

6.2 Performance Monitoring Events

Event names are denoted `<EV event_name>` where the `event_name` follow this convention:

- **INST_***: ISA instruction type profiling event. These events count when a matching instruction retires. By definition, these are non-speculative.
- ***_NONSPEC**: Non-speculative microarchitectural event. These events count when an instruction that caused or experienced the microarchitectural event retires.
- **(other)**: Microarchitectural events. These events count when the event occurs. The observed PC at the time of the associated counter increment is not necessarily indicative of an instruction that caused or experienced the event. Often referred to as "speculative" since they count before it is known whether the related instructions retire or are cleared due to a misprediction.

Event Name	Brief Description
ATOMIC_OR_EXCLUSIVE_FAIL	Atomic or exclusive instruction failed (due to contention)
ATOMIC_OR_EXCLUSIVE_SUCC	Atomic or exclusive instruction successfully completed
BRANCH_CALL_INDIR_MISPRED_NONSPEC	Retired indirect call instructions mispredicted
BRANCH_COND_MISPRED_NONSPEC	Retired conditional branch instructions that mispredicted
BRANCH_INDIR_MISPRED_NONSPEC	Retired indirect branch instructions including calls and returns that mispredicted
BRANCH_MISPRED_NONSPEC	Retired branch instructions including calls and returns that mispredicted
BRANCH_RET_INDIR_MISPRED_NONSPEC	Retired return instructions that mispredicted
CORE_ACTIVE_CYCLE	Cycles while the core was active
FETCH_RESTART	Fetch Unit internal restarts for any reason. Does not include branch mispredicts
FLUSH_RESTART_OTHER_NONSPEC	Pipeline flush and restarts that were not due to branch mispredictions or memory order violations
INST_ALL	All retired instructions
INST_BARRIER	Retired barrier instructions
INST_BRANCH	Retired branch instructions including calls and returns

Event Name	Brief Description
INST_BRANCH_CALL	Retired subroutine call instructions
INST_BRANCH_INDIR	Retired indirect branch instructions including indirect calls
INST_BRANCH_RET	Retired subroutine return instructions
INST_BRANCH_TAKEN	Retired taken branch instructions
INST_INT_ALU	Retired non-branch and non-load/store Integer Unit instructions
INST_INT_LD	Retired load Integer Unit instructions
INST_INT_ST	Retired store Integer Unit instructions
INST_LDST	Retired load and store instructions
INST_SIMD_ALU	Retired non-load/store Advanced SIMD and FP Unit instructions
INST_SIMD_ALU_VECTOR	Retired non-load/store Advanced SIMD instructions (including integer and floating point data types) Note: Available on M2 Generation and following, and A15 Bionic and following
INST_SIMD_LD	Retired load Advanced SIMD and FP Unit instructions
INST_SIMD_ST	Retired store Advanced SIMD and FP Unit instructions
INTERRUPT_PENDING	Cycles while an interrupt was pending because it was masked
L1D_CACHE_MISS_LD	Loads that missed the L1 Data Cache
L1D_CACHE_MISS_LD_NONSPEC	Retired loads that missed in the L1 Data Cache
L1D_CACHE_MISS_ST	Stores that missed the L1 Data Cache
L1D_CACHE_MISS_ST_NONSPEC	Retired stores that missed in the L1 Data Cache
L1D_CACHE_WRITEBACK	Dirty cache lines written back from the L1D Cache toward the Shared L2 Cache
L1D_TLB_ACCESS	Load and store accesses to the L1 Data TLB
L1D_TLB_FILL	Translations filled into the L1 Data TLB
L1D_TLB_MISS	Load and store accesses that missed the L1 Data TLB
L1D_TLB_MISS_NONSPEC	Retired loads and stores that missed in the L1 Data TLB
L1I_CACHE_MISS_DEMAND	Demand instruction fetches that missed in the L1 Instruction Cache
L1I_TLB_FILL	Translations filled into the L1 Instruction TLB
L1I_TLB_MISS_DEMAND	Demand instruction fetches that missed in the L1 Instruction TLB
L2_TLB_MISS_DATA	Loads and stores that missed in the L2 TLB
L2_TLB_MISS_INSTRUCTION	Instruction fetches that missed in the L2 TLB
LD_UNIT_UOP	Uops that flowed through the Load Unit
LD_NT_UOP	Load uops that executed with non-temporal hint
LDST_X64_UOP	Load and store uops that crossed a 64B boundary
LDST_XPG_UOP	Load and store uops that crossed a 16KiB page boundary
MAP_DISPATCH_BUBBLE	Cycles while the Map Unit was not stalled and Decode Unit did not send any uops
MAP_INT_UOP	Mapped Integer Unit uops

Unavailable features will have a parameter value of 0. Unimplemented features will return an error (ENOENT).

For the full list of implemented ISA feature names, see [Section 2.2, "Arm AARCH64 ISA"](#).

Some Arm features also have legacy parameter names created prior to standardization. Legacy parameters will continue to exist to support existing software. However, use the new standardized names whenever possible.

Table B.3. Legacy Feature sysctl Parameter Names

Legacy Parameter	New Parameter
hw.optional.neon (The term "NEON" is no longer broadly used by Arm. See Chapter 3, ISA Optimization: Advanced SIMD and FP Unit.)	hw.optional.AdvSIMD
hw.optional.neon_hpfp	hw.optional.AdvSIMD_HPFPcvt
hw.optional.neon_fp16	hw.optional.arm.FEAT_FP16
hw.optional.armv8_1_atomics	hw.optional.arm.FEAT_LSE
hw.optional.armv8_2_fhm	hw.optional.arm.FEAT_FHM
hw.optional.armv8_2_sha512	hw.optional.arm.FEAT_SHA512
hw.optional.armv8_2_sha3	hw.optional.arm.FEAT_SHA3
hw.optional.armv8_3_compnum	hw.optional.arm.FEAT_FCMA

These sysctl feature parameters are also documented at https://developer.apple.com/documentation/kernel/1387446-sysctlbyname/determining_instruction_set_characteristics.

B.2 Cache and Topology Characteristics

The sysctl interface contains parameters that describe cache organization and topology. Software can query the interface to determine the configuration and adapt the workload accordingly, such as to block a data set for a particular cache size. The cache topology parameters do not include the M Cache, and as noted, software should not optimize for its capacity.

[Table B.4: "Per Performance Level Sysctl Parameters"](#) lists parameters that clearly define the system topology and include an example of M1 Ultra. The parameters include a hierarchy level called `perfllevel{N}` that specify parameters by core type. Higher performing cores have lower perfllevels. On M1 Ultra, the P core parameters are located under `perfllevel0` and the E core parameters are located under `perfllevel1`. These parameters are available in macOS 12, iOS 15, iPadOS 15, and later.

Table B.4. Per Performance Level Sysctl Parameters

New Per Performance Level Parameters	Definition	M1 Ultra Values	
		perfllevel0 P core	perfllevel1 E core
hw.nperfllevels	Number of types of general purpose cores in the chip. The	2	

