

# N2436: Memory region stores flush and reloads force

Document #: N2436  
Date: 2019-09-22  
Project: WG14 Programming Language C  
WG14-WG21 liaison group  
Reply-to: Niall Douglas <[s\\_sourceforge@nedprod.com](mailto:s_sourceforge@nedprod.com)>  
Bob Steagall <[bob.steagall.cpp@gmail.com](mailto:bob.steagall.cpp@gmail.com)>

There is an increasing need in modern code to work efficiently with shared memory where stores have side effects not visible to the current program's abstract machine, and whose contents may change unbeknownst to the current program's abstract machine. `volatile` will do the job, but it does not perform well. This paper proposes a pair of library functions with greatly improved performance, and which additionally provide bare minimum support for NV-DIMM memory.

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Proposed design</b>	<b>2</b>
<b>3</b>	<b>Proposed API</b>	<b>3</b>
<b>4</b>	<b>Implementation cost</b>	<b>4</b>
<b>5</b>	<b>Conclusion</b>	<b>5</b>
<b>6</b>	<b>References</b>	<b>5</b>

## 1 Introduction

`volatile` is the canonical standard way to tell the C compiler to not eliminate loads nor stores to the thus qualified items, and thus not subject stores to dead store elimination, nor loads to reload elision. It additionally prevents the compiler from merging and reordering loads and stores to such regions, which makes qualifying everything with `volatile` an inefficient way of working with shared memory.

What would be ideal is for a mechanism with fewer unwanted side effects for telling the C compiler that stores to the shared memory region must be guaranteed to occur, and loads from the shared memory must always reload rather than reuse previously loaded values. Such stores can be reordered, merged and composed by the compiler so long as their final side effects are rendered;

similarly, reloads from such regions can be reordered and merged by the compiler in any way which is maximally efficient, so long as a complete reload is performed.

Directly mapped non-volatile storage devices are a special kind of shared memory where the `fdatasync()` operation may be efficiently implemented by ensuring that the CPU flushes modified cache lines to ‘main memory’, where the directly mapped storage device has been mapped into the C program. Directly mapped non-volatile storage is expected to become prevalent on high end and low power systems in the medium term, as it offers superior performance and power consumption.

## 2 Proposed design

There are two functions, one to force reloading of a region of memory, and another to flush any pending stores to a region of memory.

There are two strengths of use for each function, with a corresponding impact upon runtime performance:

1. Force the compiler to write out and reload state, with optional CPU memory fence to constrain reordering of loads and stores from the perspective of other CPUs.
2. Force both the compiler and the CPU to write out and reload state, with optional CPU memory fence to constrain reordering of loads and stores from the perspective of other CPUs.

Cache coherency is always preserved. If the CPU is required to reload cache lines from a region of memory by these functions, all modified cache lines for that region in all CPUs are written out beforehand. This would ruin the utility of forcing the CPU to reload main memory which has changed independently of CPU writes, so it is on the user to ensure that no modified cache lines exist for a region reloaded by the CPU i.e. flush all modified lines, tell the device to change the data, reload cache lines.

CPU store flush and reload is not atomic with respect to concurrent CPU store flush and reload upon overlapping bytes. Users should employ additional synchronisation to prevent concurrent CPU store flushes or reloads upon the same cache lines.

No guarantee is given that individual cache line flushes occur in any particular order before the CPU memory release fence i.e. if you flush four modified cache lines, main memory will receive those updates in an unknown order. If you do not wish this, flush each cache line individually with memory fence in between each.

No guarantee is given that individual cache line reloads occur in any particular order after the CPU memory acquire fence. If you reload four cache lines, they may be fetched from main memory in any order. If you do not wish this, reload each cache line individually with memory fence in between each.

### 3 Proposed API

A reference implementation for the proposed library APIs can be found at [https://github.com/ned14/quickcpplib/blob/master/include/mem\\_flush\\_loads\\_stores.hpp](https://github.com/ned14/quickcpplib/blob/master/include/mem_flush_loads_stores.hpp), with support for x86, x64, ARM v7 and AArch64.

It has been in production use in the presented form for about six months, however two previous earlier API designs were in use for some years beforehand. This proposed API design reflects the experience gained in those years.

```
1  enum memory_flush
2  {
3      memory_flush_none,    //!< No main memory flushing.
4
5      memory_flush_retain,  //!< Flush modified cache line in CPU out to main
6                          //!< memory, but retain as unmodified in cache.
7
8      memory_flush_evict    //!< Flush modified cache line in CPU out to main
9                          //!< memory, and evict completely from all caches.
10 };
11
12 /*! \brief Ensures that reload elimination does not happen for a region of
13 memory, optionally synchronising the region with main memory.
14
15 \return The kind of memory flush actually used.
16 \param data The beginning of the byte array to ensure loads from.
17 \param bytes The number of bytes to ensure loads from.
18 \param kind Whether to ensure loads from the region are from main memory.
19 \param order The atomic reordering constraints to apply to this operation.
20
21 \note 'memory_flush_retain' has no effect for reloads from main memory,
22 it is the same as doing nothing. Only 'memory_flush_evict' evicts all the
23 cache lines for the region of memory, thus ensuring that subsequent loads
24 are from main memory. Note that if the cache line being reloaded is modified,
25 it will be flushed to main memory before being reloaded, thus destroying
26 any modified data there. You should therefore ensure that concurrent
27 actors never modify main memory with modified cache lines in your CPU.
28 */
29 memory_flush mem_force_reload_explicit(volatile char *data,
30                                       size_t bytes,
31                                       memory_flush kind,
32                                       memory_order order);
33
34 /*! \brief The same as 'mem_force_reload_explicit()', but with
35 'kind' set to 'memory_flush_none', and 'order' set to 'memory_order_acquire'.
36 This does not reload loads from main memory, and prevents reads and writes
37 to this region subsequent to this operation being reordered to before this
38 operation.
39 */
40 memory_flush mem_force_reload(volatile char *data,
41                              size_t bytes);
42
43 /*! \brief Ensures that dead store elimination does not happen for a region of
44 memory, optionally synchronising the region with main memory.
45
```

```

46 \return The kind of memory flush actually used.
47 \param data The beginning of the byte array to ensure stores to.
48 \param bytes The number of bytes to ensure stores to.
49 \param kind Whether to wait until all stores to the region reach main memory.
50 \param order The atomic reordering constraints to apply to this operation.
51
52 \warning On older Intel CPUs, due to lack of hardware support, we always execute
53 'memory_flush_evict' even if asked for 'memory_flush_retain'. This can produce
54 some very poor performance. Check the value returned to see what kind of flush
55 was actually performed.
56 */
57 memory_flush mem_flush_stores_explicit(volatile const char *data,
58                                       size_t bytes,
59                                       memory_flush kind,
60                                       memory_order order);
61
62 /*! \brief The same as 'mem_flush_stores_explicit()', but with
63 'kind' set to 'memory_flush_none', and 'order' set to 'memory_order_release'.
64 This does not flush stores to main memory, and prevents reads and writes to
65 this region preceding this operation being reordered to after this operation.
66 */
67 memory_flush mem_flush_stores(volatile const char *data,
68                               size_t bytes);

```

## 4 Implementation cost

It is firstly important to note that from the perspective of the compiler, the semantics of the above functions could be exactly the same as for calling an `extern` function – the compiler must write out relevant state before the call, and reload relevant state after the call. The main difference is that compilers *could* treat these functions a bit more optimally, in that `mem_flush_stores()` does not require reloading state afterwards, and `mem_force_reload()` does not require writing out state beforehand. Also, the fact that `mem_flush_stores()` always equals an `atomic_thread_fence(memory_order_release)`, and `mem_force_reload()` always equals an `atomic_thread_fence(memory_order_acquire)`, could be useful to an optimiser.

On compilers without link time optimisation, a conforming implementation of these functions is trivially easy, just being `extern` is enough. The main memory flush support is also trivially easy to implement, on Intel CPUs one loops the `CLWB` opcode over the cache lines, on ARM64 it is the `dc cvac` opcode. Most CPUs have a similar operation – for those which do not, the functions return `memory_flush_none` to indicate that no main memory flushing was performed.

On compilers with link time optimisation, there is a trick which the reference library implementation uses – it calls an externally modifiable function pointer, thus forcing the compiler to flush and reload state even under link time optimisation. Current link time optimisation technology appears to be unable to discern that the function pointer will only ever have one value, and thus deindirect the function call, and thus optimise away the desired semantics.

## 5 Conclusion

Obviously built in support for these operations in the compiler would be much better again still than calling an `extern` function which in the default implementation, simply calls an atomic thread fence and nothing else. This is why these operations are proposed for standardisation: built in (i.e. inlined) understanding would be more efficient than calling a function which calls a single CPU opcode and returns.

One would also be guaranteed that future link time optimisation technology improvements will retain the desired semantics.

Finally there would be the hope that compilers could better optimise how much state to flush and reload based on the byte range supplied, as the current reference implementation is a sledgehammer which flushes and reloads all compiler state.

## 6 References

[Intel CLWB opcode] <https://www.felixcloutier.com/x86/clwb>

[ARM DC opcode] <https://developer.arm.com/docs/ddi0595/latest/aarch64-system-instructions/dc-cvac>