
Implications of memory interface configurations on low-power STM32 microcontrollers

Introduction

The low-power STM32 microcontrollers have a rich variety of configuration options regarding the Flash memory interface.

This application note showcases the different settings under various test conditions, providing guidelines for the optimization of the application power consumption.

Reference documents

The reference documents are available on STMicroelectronics web site www.st.com:

- *Ultra-low-power STM32L0x3 advanced ARM[®]-based 32-bit MCUs* Reference Manual (RM0367)
- *STM32L100xx, STM32L151xx, STM32L152xx and STM32L162xx advanced ARM[®]-based 32-bit MCUs* Reference Manual (RM0038)
- *STM32L4x6 advanced ARM[®]-based 32-bit MCUs* Reference Manual (RM0411)

This application note applies to the products listed in [Table 1](#).

Table 1. Applicable products

Type	Product series
Microcontrollers	STM32L0 Series, STM32L1 Series, STM32L4 Series

Contents

- 1 Definitions 5**
- 2 System architecture 5**
- 3 Operation modes 6**
 - 3.1 STM32L1 Series device options 6
 - 3.2 STM32L0 Series device options 7
 - 3.3 STM32L4 Series device options 7
 - 3.4 Execution from a volatile memory 8
- 4 Power consumption and performance comparison
using STM32L1 Series devices 9**
 - 4.1 Dhrystone benchmark 9
 - 4.2 32-bit instruction code 12
 - 4.3 Memory read stress test 14
- 5 Power consumption and performance comparison
using STM32L0 Series devices 16**
 - 5.1 Dhrystone benchmark 16
 - 5.2 Memory read stress test 18
- 6 Power consumption and performance comparison
using STM32L4 Series devices 21**
 - 6.1 Influence of prefetch and cache with zero Flash latency 21
 - 6.2 Dhrystone benchmark 22
 - 6.3 Memory read stress test 24
- 7 Conclusion 25**
- 8 Revision history 26**

List of tables

Table 1.	Applicable products	1
Table 2.	List of acronyms	5
Table 3.	Configurations available on STM32L1 Series devices with regulator range 1	6
Table 4.	Configurations available on STM32L0 Series devices with regulator range 1	7
Table 5.	Device option summary	7
Table 6.	Dhrystone results with no background transfer	9
Table 7.	Dhrystone results with DMA simultaneously reading data from the Flash memory	10
Table 8.	32-bit code result with no background transfer.	12
Table 9.	32-bit code result with DMA simultaneously reading data from the Flash memory	13
Table 10.	Literal pool with no additional data read from the Flash memory	14
Table 11.	Literal pool reading with DMA simultaneously reading the Flash memory	15
Table 12.	Dhrystone with no additional data read from the Flash memory.	16
Table 13.	Dhrystone with DMA simultaneously reading data from the Flash memory	17
Table 14.	Literal pool with no additional data read from the Flash memory	19
Table 15.	Literal pool with DMA simultaneously reading data from the Flash memory	20
Table 16.	Dhrystone test using core voltage range 1 and HSI clock	22
Table 17.	Literal measurements	24
Table 18.	Document revision history	26

List of figures

Figure 1. Dhrystone results with no background transfer 10
Figure 2. Dhrystone results with DMA simultaneously reading data from the Flash memory 11
Figure 3. 32-bit code result with no background transfer. 12
Figure 4. 32-bit code result with DMA simultaneously reading data from the Flash memory 13
Figure 5. Literal pool reading with no additional data read from the Flash memory 14
Figure 6. Literal pool reading with DMA simultaneously reading data from the Flash memory 15
Figure 7. Dhrystone with no additional data read from the Flash memory. 17
Figure 8. Dhrystone with DMA simultaneously reading data from the Flash memory 18
Figure 9. Literal pool with no additional data read from the Flash memory 19
Figure 10. Literal pool with DMA simultaneously reading data from the Flash memory 20
Figure 11. Dhrystone test plot of energy needed for execution 22
Figure 12. Energy cost of the dhrystone test loop 23
Figure 13. Literal pool chart plot of energy efficiency 24

1 Definitions

Table 2. List of acronyms

Term	Description
NV	Non-volatile (memory), also referred as Flash memory
HSI	High-speed internal clock
SPI	Serial peripheral Interface bus
MCU	Microcontroller
CPU	Central processing unit (part of the MCU)
NVIC	Nested vector interrupt controller
DMA	Direct memory access
RM	Reference manual
SWD	Single wire debug interface

2 System architecture

The memory interface manages the read and write accesses from the core/bus matrix towards the non-volatile memory. This holds for both the instruction and data access.

For configuring the non-volatile memory read access during the program execution, the configuration flags are accessible in the access control register.

The latency serves the purpose of reducing the rate at which the NVM is read. An extra wait cycle must be enabled for a system clock higher than 16 MHz for the highest voltage regulator range. For lower core voltages this threshold frequency goes lower.

To compensate this bandwidth deficiency, a prefetch can be configured. The memory controller will then attempt to have the next instruction ready before the core requests it.

The STM32L1 memory interface can also use a 64-bit read access internally to be able to serve the core with data and instruction close to its own space. The extra 32 bits can be used by the prefetch to load the next instruction and provide it to the core immediately when needed.

The STM32L0 memory interface does not have the 64-bit wide bus, but the memory controller is capable of data pre-read. This simple buffer is similar to the prefetch, but works also for data.

The STM32L4 memory interface has a full 64-bit wide (plus 8-bit ECC) connection to the bus matrix, shared between data and instruction. The Flash interface incorporates an ART accelerator, a prefetch mechanism and a cache designed to minimize the effect of memory latency. The Flash interface is then capable of transferring data and instruction simultaneously, under the condition that they are ready in the cache.

All the performance improvements resulting from the memory interface settings, come at a cost of an increased power consumption. A 32-bit access with no latency, no pre-read and

no prefetch is considered as a low-power mode. The following section sheds light on the kind of tradeoffs they represent.

3 Operation modes

The following operation modes are used to assess the impact of the memory interface settings on the performance and power consumption. All the measurements have been done using $V_{CC} = 3.3\text{ V}$ and the voltage regulator range 1. The speed and consumption would be lower using lower regulator levels, but linearly lower relative to the range 1 measurements. For example with the voltage regulator range 3 and the system clock speed at 2 MHz (from MSI) the power consumption would be roughly 10 times lower for all the measurements and the performance roughly 10 times lower for all the measured configurations. There is no point in repeating the measurement for all the configuration combinations.

3.1 STM32L1 Series device options

[Table 3](#) lists a short summary of the device options. For a detailed description refer to *read interface* section of the RM0038 reference manual.

Table 3. Configurations available on STM32L1 Series devices with regulator range 1

Frequency	<16 MHz					>16 MHz	
	0	0	0	1	1	1	1
Latency	0	0	0	1	1	1	1
64-bit	0	1	1	1	1	1	1
Prefetch	0	0	1	0	1	0	1

The table of valid configurations is clearly demonstrating the following simple rules:

- The latency cannot be turned off with clock speeds exceeding 16 MHz.
- When the latency is set to 1, the 64-bit access is mandatory.
- The prefetch is impossible without the 64-bit access.

3.2 STM32L0 Series device options

[Table 4](#) lists a short summary of the device options. For a detailed description refer to *reading the NVM* section of the RM0367 reference manual.

Table 4. Configurations available on STM32L0 Series devices with regulator range 1

Frequency	<16 MHz										>16 MHz				
Latency	0	1	0	1	0	1	0	1	0	1	1	1	1	1	1
Pre-read	0	0	1	1	1	1	0	0	X	X	0	1	1	0	X
Prefetch	0	0	0	0	1	1	1	1	X	X	0	0	1	1	X
Buffer disable	0	0	0	0	0	0	0	0	1	1	0	0	0	0	1

The table of valid configurations is clearly demonstrating the following simple rules:

- The latency cannot be turned off with clock speeds exceeding 16 MHz.
- When the buffer is disabled, it cannot be configured.
- Prefetch and pre-read configure the usage of the 6 words in the internal buffer, not their total amount.

3.3 STM32L4 Series device options

[Table 5](#) lists a short summary of the device options. For a detailed description refer to *reading the NVM* section of the RM0411 reference manual.

Table 5. Device option summary

Frequency	<16 (at V _{CORE} range 1)								>16 (at V _{CORE} range 2)							
Latency	0								>1							
Data cache	0	0	0	1	0	1	1	1	0	0	0	1	0	1	1	1
Instruction cache	0	0	1	0	1	1	0	1	0	0	1	0	1	1	0	1
Prefetch	0	1	0	0	1	0	1	1	0	1	0	0	1	0	1	1

The table is clearly demonstrating that the prefetch, the data cache and the instruction cache settings are mutually independent. Any of these three features can be activated or deactivated regardless of the frequency or any other setting.

The settings are only simple when the voltage regulator settings are disregarded. But the read access latency strongly depends on the voltage regulator settings. For example at a 16 MHz speed, while with range 1 the latency on a Flash read is 1 CPU cycle, with range 2 the latency on the same core frequency increases to 3 CPU cycles.

For more details see the *Read access latency* section in the RM0411 reference manual.

3.4 Execution from a volatile memory

The intuitive way to avoid the Flash memory speed issues would be to use the RAM for selected portions of code. There are several reasons not to do that.

1. The RAM is a scarce resource on small devices.
2. Most of data are likely to be placed in the RAM, accessing the code in the RAM eliminates the advantage of Harvard architecture approach in the STM32L1 Series.
3. To switch off the Flash memory and conserve more energy, also the interrupt table and interrupt handlers need to be in the RAM.

In case of a typical microcontroller application, the overall energy budget of the RAM execution is roughly the same as the execution on the 32 MHz system clock with the Flash memory latency set. Which means that if the Flash memory can run without the latency enabled, it will be a better option most of the time. In other words the RAM execution tends to be about 30% slower than the execution of the same code from the Flash memory and the current consumption should not decrease more than the same 30% range.

4 Power consumption and performance comparison using STM32L1 Series devices

To assess the performance of the MCU with different memory controller settings, several benchmark tests have been used. All the tests have been executed on a STM32L152RC Discovery board using all the available memory interface settings, listed in [Section 3.1](#). All the tests have been executed both in standalone and in parallel with a DMA transfer, constantly reading from the program NV memory. The DMA channel was directed to the SPI output configured to the highest available speed ($f_{PCLK}/2$) and low priority.

Three clock configurations have been used in the measurements. One with the plain 16 MHz HSI clock as the system clock and no latency set, another with the same clock but the Flash latency configured (Flash memory running effectively on lower clock) and the third with the PLL set to produce the 32 MHz system clock.

All the measurements are taken on a single sample of Discovery board at ambient temperature. The values provided are an arithmetic mean from several measurements.

4.1 Dhrystone benchmark

Although the Dhrystone benchmark is often deemed outdated, it is still somewhat representative of many microcontroller applications.

Table 6. Dhrystone results with no background transfer

Frequency	16 MHz					32 MHz	
	0	1	2	3	4	5	6
Latency	0	0	0	1	1	1	1
64-bit	0	1	1	1	1	1	1
Prefetch	0	0	1	0	1	0	1
Timing for 50000 cycles [s]	2.57	2.57	2.57	3.05	2.86	1.52	1.46
Average current [mA]	5.75	5.78	6.11	5.13	5.62	10.42	11.08
Energy [mJ]	48.77	49.02	51.82	51.63	53.04	52.27	53.38

Figure 1. Dhrystone results with no background transfer

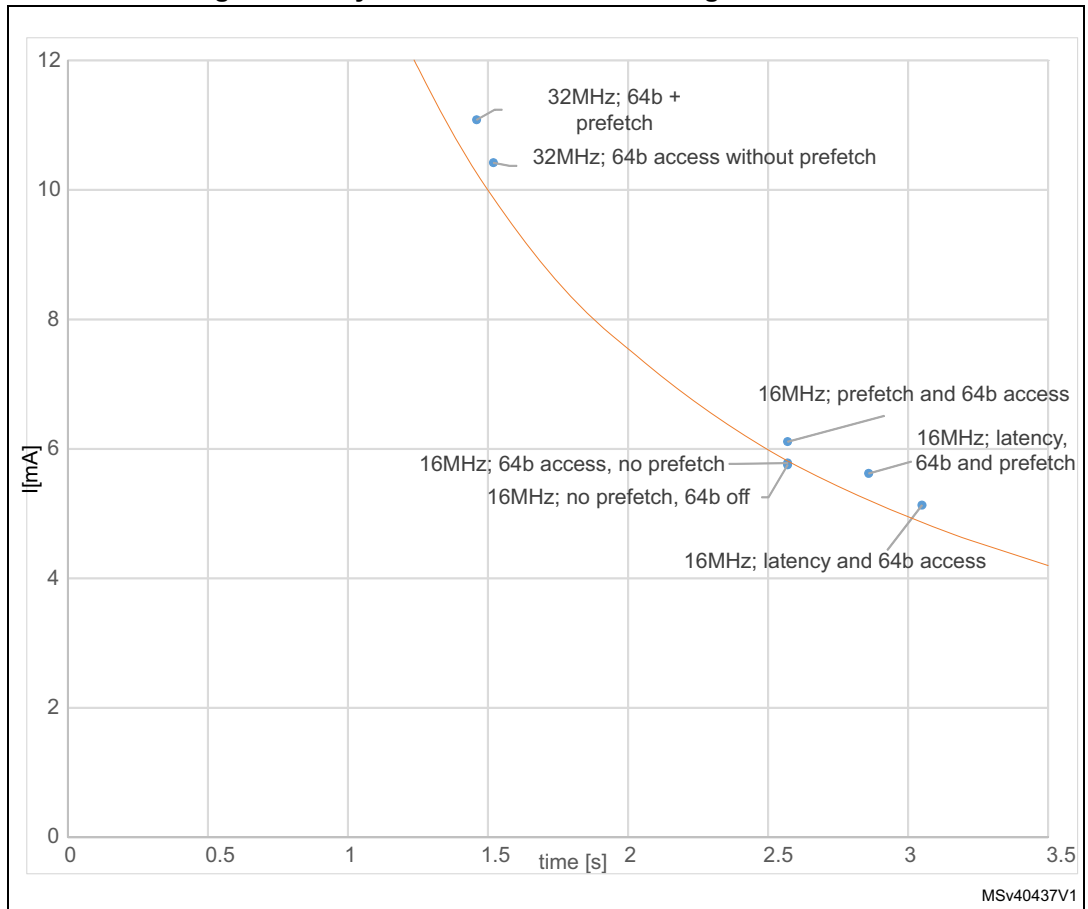
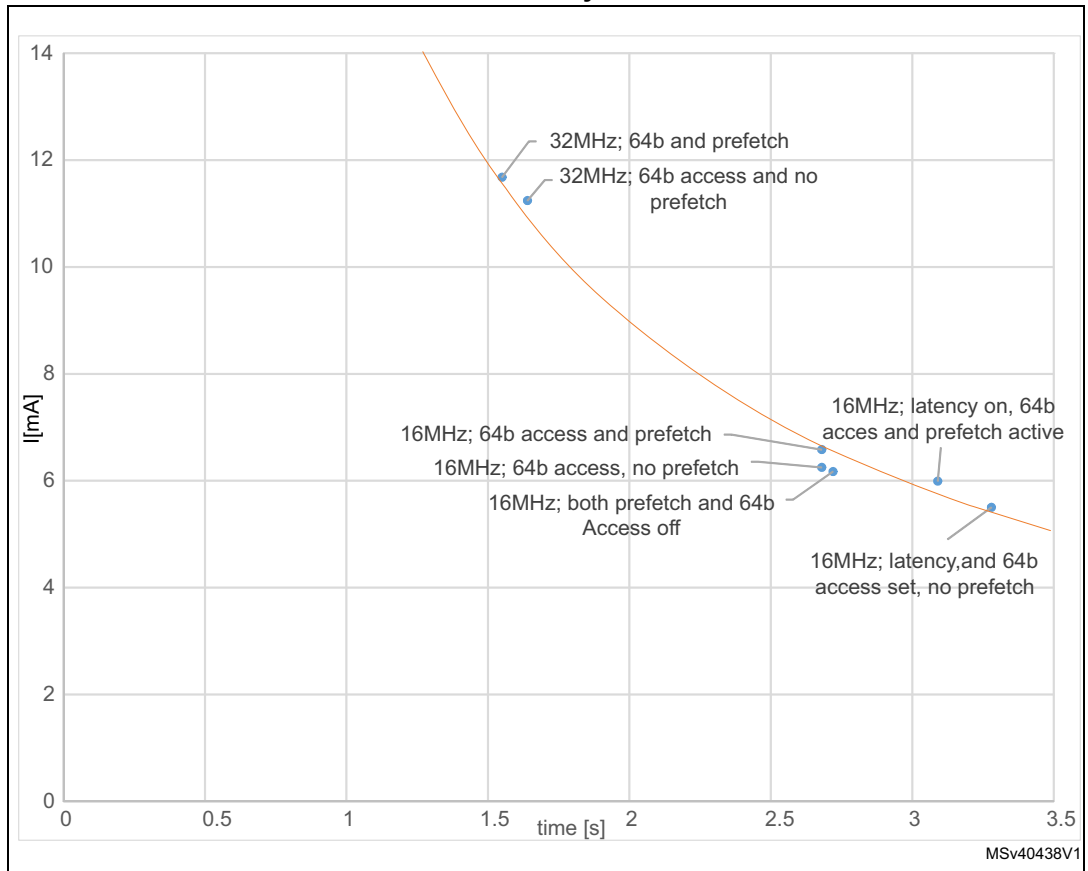


Table 7. Dhrystone results with DMA simultaneously reading data from the Flash memory

Frequency	16 MHz					32 MHz	
	0	1	2	3	4	5	6
Latency	0	0	0	1	1	1	1
64-bit	0	1	1	1	1	1	1
Prefetch	0	0	1	0	1	0	1
Timing for 50000 cycles [s]	2.72	2.68	2.68	3.28	3.09	1.64	1.55
Average current [mA]	6.17	6.25	6.58	5.50	5.99	11.24	11.68
Energy [mJ]	55.38	55.28	58.19	59.53	61.08	60.83	59.74

Figure 2. Dhrystone results with DMA simultaneously reading data from the Flash memory



Configuring a 64-bit access or a prefetch makes a very small difference on a low clock speed where the latency can be avoided. On the contrary, setting the latency may lead to a lower power consumption in situations where the speed is not critical. At higher speeds the efficiency of the prefetch is situational, leading to ultimate performance but the gain in speed may be lower than the consumption increase.

4.2 32-bit instruction code

A stress test consists of executing 12 aligned 32-bit instructions manipulating data in registers in a loop of 500000 cycles. The code with a higher ratio of 32-bit instructions is more likely to find a bottleneck in the memory interface than a typical Thumb code with prevalent 16-bit instructions.

Table 8. 32-bit code result with no background transfer

Frequency	16 MHz					32 MHz	
	0	0	0	1	1	1	1
Latency	0	0	0	1	1	1	1
64-bit	0	1	1	1	1	1	1
Prefetch	0	0	1	0	1	0	1
Timing for 500000 cycles [s]	0.9	0.9	0.9	1.06	0.964	0.59	0.497
Average current [mA]	5.25	5.41	5.63	4.82	5.11	9.09	9.78
Energy [mJ]	15.59	16.07	16.72	16.86	16.26	17.70	16.04

Figure 3. 32-bit code result with no background transfer


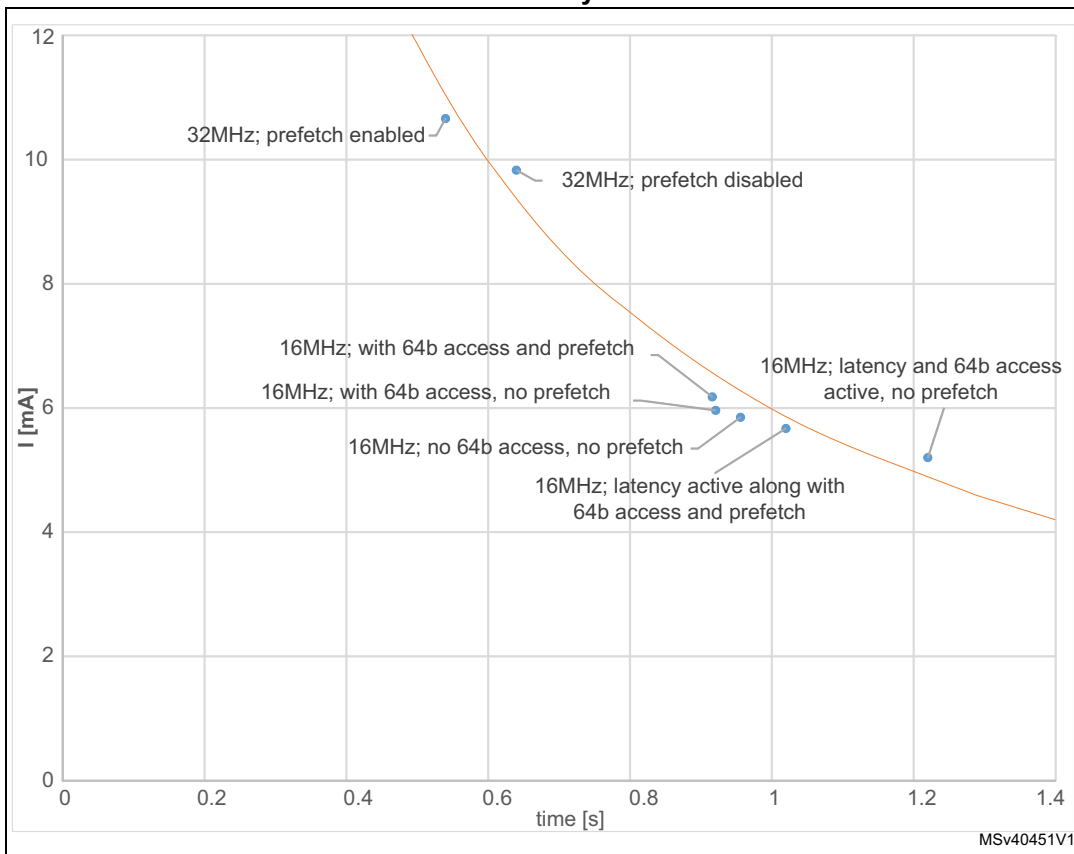


Table 9. 32-bit code result with DMA simultaneously reading data from the Flash memory

Frequency	16 MHz					32 MHz	
	0	0	0	1	1	1	1
Latency	0	0	0	1	1	1	1
64bit	0	1	1	1	1	1	1
Prefetch	0	0	1	0	1	0	1
Timing for 500000 cycles [s]	0.956	0.921	0.916	1.22	1.02	0.64	0.54
Average current [mA]	5.85	5.96	6.18	5.20	5.67	9.83	10.66
Energy [mJ]	18.46	18.11	18.68	20.94	19.09	20.76	19.00

Figure 4. 32-bit code result with DMA simultaneously reading data from the Flash memory



The findings are in line with the expectations: a code with high share of 32-bit instructions benefits a lot from the prefetch once the memory latency is in place. But with zero latency the extra bandwidth is likely to be useless.

4.3 Memory read stress test

A stress test consists of executing 20 LDR instructions fetching data from the program NV memory to the CPU core registers in a loop of 500000 cycles. This way, not only the instructions are fetched from the memory but another read access is generated during the instruction execution, again creating a choke point at the memory interface. Fetching of subsequent instruction is then likely to be delayed. The code simulates a case when a heavy load of literal pools (string constants) like for example predefined messages, is read from a non-volatile memory very often.

Note: The memory reading by LDM instructions is not used as it is not demonstrating limits of the memory interface, only the memory itself.

Table 10. Literal pool with no additional data read from the Flash memory

Frequency	16 MHz					32 MHz	
Latency	0	0	0	1	1	1	1
64-bit	0	1	1	1	1	1	1
Prefetch	0	0	1	0	1	0	1
Timing for 500000 cycles [s]	3.66	2.73	2.72	3.38	3.32	1.69	1.66
Average current [mA]	5.44	5.58	6.12	4.85	5.33	9.78	10.73
Energy [mJ]	65.70	50.27	54.93	54.10	58.40	54.54	58.78

Figure 5. Literal pool reading with no additional data read from the Flash memory

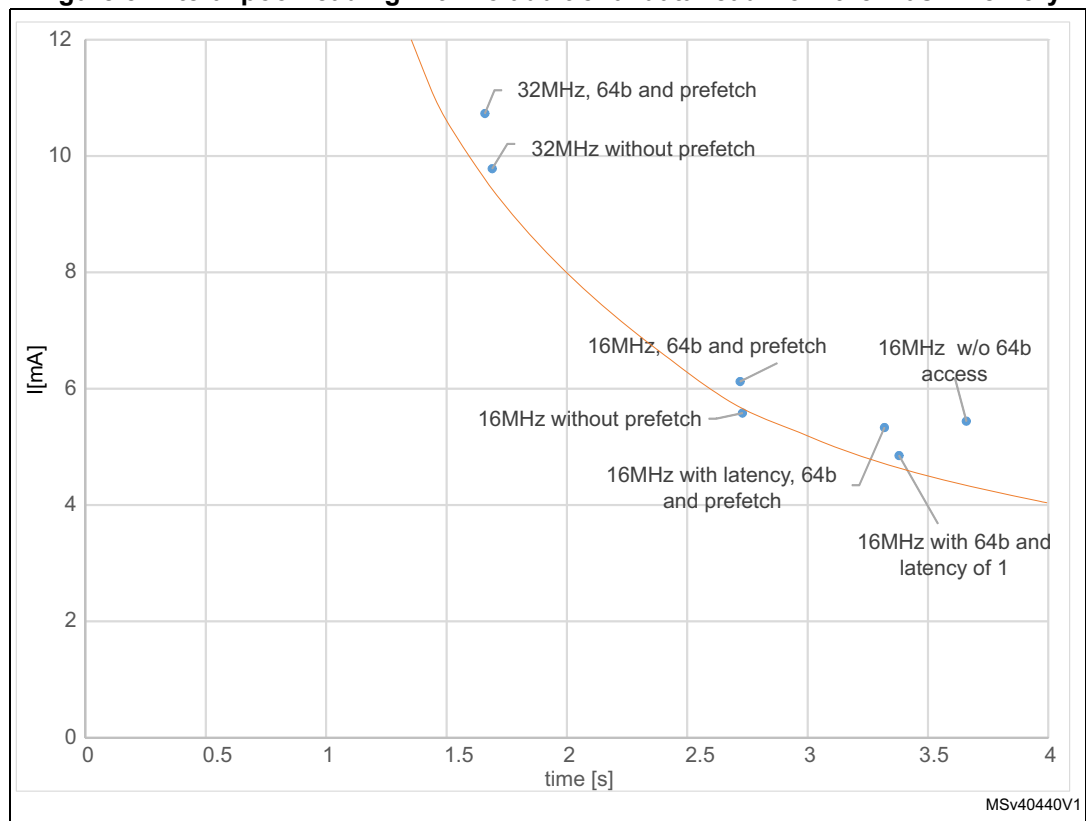



Table 11. Literal pool reading with DMA simultaneously reading the Flash memory

Frequency	16 MHz					32 MHz	
	0	0	0	1	1	1	1
Latency	0	0	0	1	1	1	1
64-bit	0	1	1	1	1	1	1
Prefetch	0	0	1	0	1	0	1
Timing for 500000 cycles [s]	3.98	2.94	2.94	3.92	3.88	1.97	1.96
Average current [mA]	6.04	6.26	6.73	5.40	5.72	10.62	11.59
Energy [mJ]	79.33	60.73	65.29	69.85	73.24	69.04	74.96

Figure 6. Literal pool reading with DMA simultaneously reading data from the Flash memory



As expected, mostly in case of a data read transfer the effect of the prefetch is lower, but a 64-bit memory access makes a significant difference even with zero memory latency.

5 Power consumption and performance comparison using STM32L0 Series devices

The Cortex-M0+ core is much simpler compared to the Cortex-M3 used in the STM32L1 Series. The 32-bit instruction benchmark is dropped as the Thumb-2 instruction set support in the M0+ core is very limited and an extensive usage of 32-bit code is not realistic with a code compiled for the STM32L0 Series.

The remaining tests have been executed on a STM32L073-EVAL board using all the available memory interface settings, listed in [Section 3.2](#). All the tests have been executed both standalone and in parallel with a DMA transfer constantly reading from the program NV memory. The DMA channel was directed to the SPI output configured to the highest available speed ($f_{PCLK}/2$), but low priority.

Two clock configurations have been used in the measurements. One with the plain 16 MHz HSI clock as the system clock and no latency set, the other with the PLL set to produce the 32 MHz system clock and of course the Flash memory latency set to 1.

All the measurements are taken on a single sample of Nucleo board at ambient temperature. The values provided are an arithmetic mean from several measurements.

5.1 Dhrystone benchmark

The Dhrystone code is executed and the task consists of processing 50000 cycles of the test code.

Table 12. Dhrystone with no additional data read from the Flash memory

Frequency	16 MHz					32 MHz				
Latency	0	0	0	0	0	1	1	1	1	1
Prefetch	1	0	0	1	0	1	0	0	1	0
Pre-read	1	1	0	0	0	1	1	0	0	0
Disabled buffer	0	0	1	0	0	0	0	1	0	0
Time [ms]	3769	3766	3771	3769	3769	2139	2667	2720	2130	2667
Average current [mA]	4.32	4.42	4.54	4.40	4.39	8.14	7.52	7.52	8.04	7.43
Energy [mJ]	53.73	54.93	56.49	54.72	54.60	57.46	66.20	67.49	56.51	65.40

Figure 7. Dhystone with no additional data read from the Flash memory

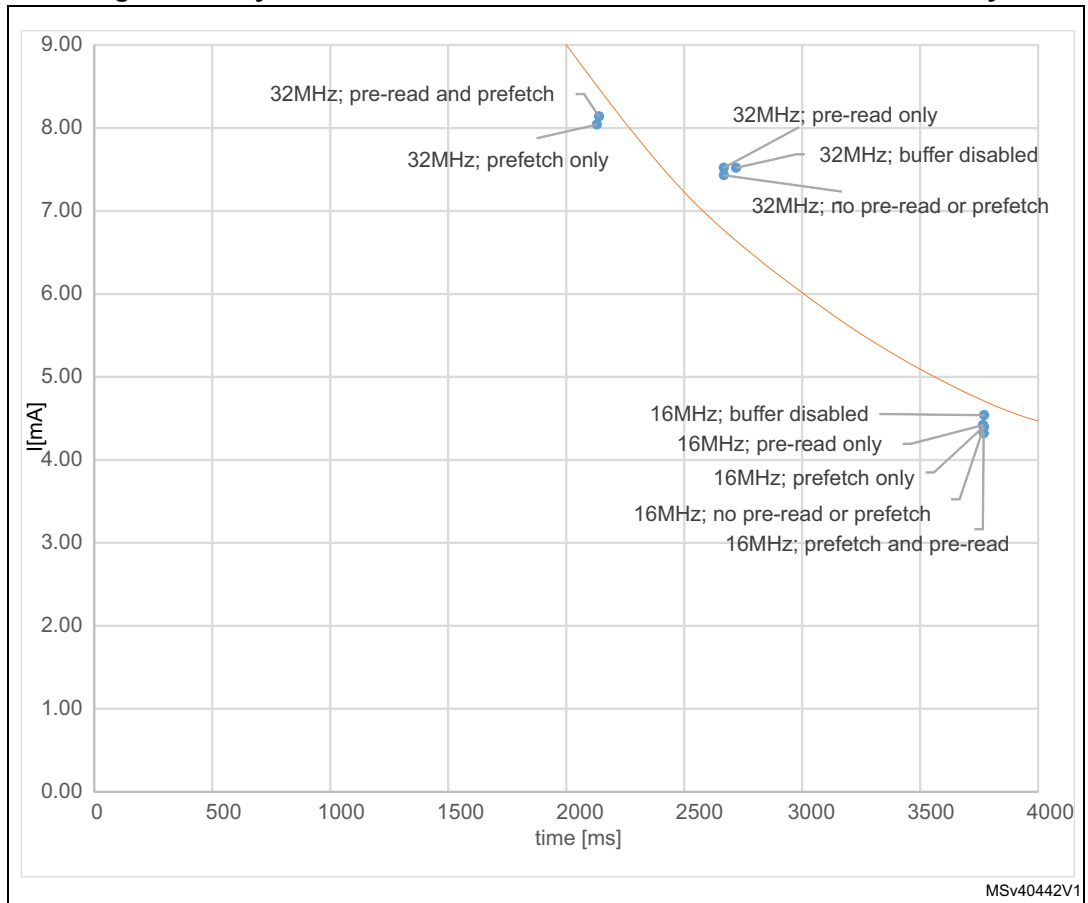



Table 13. Dhystone with DMA simultaneously reading data from the Flash memory

Frequency	16 MHz					32 MHz				
	0	1	2	3	4	5	6	7	8	9
Latency	0	0	0	0	0	1	1	1	1	1
Prefetch	1	0	0	1	0	1	0	0	1	0
Pre-read	1	1	0	0	0	1	1	0	0	0
Disabled buffer	0	0	1	0	0	0	0	1	0	0
Time [ms]	3903	3901	3906	3906	3904	2377	2853	2956	2334	2843
Average current [mA]	4.69	4.77	4.87	4.68	4.59	8.58	8.21	8.15	8.66	7.80
Energy [mJ]	69.40	61.41	62.77	60.32	59.13	67.29	77.31	79.31	66.70	73.17

Figure 8. Dhrystone with DMA simultaneously reading data from the Flash memory



This example clearly shows that the internal 6 word buffer improves the energy efficiency even if it is not well utilized, like in case of zero latency. The best option is to keep it on, but to disable the prefetch and pre-read.

In case of the configuration with the latency is enabled, the prefetch is probably worth using. The pre-read is obviously not used by the DMA channel and does not represent an improvement.

5.2 Memory read stress test

A stress test consists of executing 20 LDR instructions fetching data from program NV memory to CPU core registers in a loop of 500000 cycles. This way, not only the instructions are fetched from the memory but another read access is generated during the instruction execution, again creating a choke point at the memory interface. Fetching of subsequent instruction is then likely to be delayed. The code simulates a case when a heavy load of literal pools, like for example predefined messages, is read from a non-volatile memory very often.

Note: The memory reading by LDM instructions are not used as it is not demonstrating limits of the memory interface, only the memory itself.

Table 14. Literal pool with no additional data read from the Flash memory

Frequency	16 MHz					32 MHz				
	Latency	0	0	0	0	0	1	1	1	1
Prefetch	1	0	0	1	0	1	0	0	1	0
Pre-read	1	1	0	0	0	1	1	0	0	0
Disabled buffer	0	0	1	0	0	0	0	1	0	0
Time [ms]	2402.5	2401.5	2403	2403	2399.5	2009	2058.5	2091	1817	1819
Average current [mA]	3.4	3.42	3.36	3.14	3.19	6.03	6.05	5.94	5.83	5.73
Energy [mJ]	26.95	27.10	26.64	24.89	25.25	39.97	41.09	40.98	34.95	34.39

Figure 9. Literal pool with no additional data read from the Flash memory





Table 15. Literal pool with DMA simultaneously reading data from the Flash memory

Frequency	16 MHz					32 MHz				
	0	0	0	0	0	1	1	1	1	1
Latency	0	0	0	0	0	1	1	1	1	1
Prefetch	1	0	0	1	0	1	0	0	1	0
Pre-read	1	1	0	0	0	1	1	0	0	0
Disabled buffer	0	0	1	0	0	0	0	1	0	0
Time [ms]	2533.5	2533.5	4854.5	4587	4591	2292.5	2301	2420	2299	2302.5
Average current [mA]	3.86	3.86	3.38	3.32	3.29	7.42	7.39	7.34	7.25	7.18
Energy [mJ]	32.27	32.27	54.15	50.26	49.84	56.13	56.11	58.62	55.00	54.56

Figure 10. Literal pool with DMA simultaneously reading data from the Flash memory



This example finally demonstrates the advantage of the pre-read setting. It can greatly improve the efficiency when more than one stream of data is read from the Flash memory and there is no latency. The prefetch is not useful when dealing mostly with data, that is no surprise. Again it is a good idea to keep the buffer enabled. The only reason to disable the buffer is if the timing needs to be more deterministic, whatever the efficiency cost may be.

6 Power consumption and performance comparison using STM32L4 Series devices

The STM32L4 Series devices are based on the Cortex-M4 core connected to the 32-bit multilayer AHB bus matrix that connects up to six master and eight slave devices supporting concurrent operations as long as the bus masters are accessing different bus slaves.

The tests have been executed on a STM32L476G-Discovery board using all the available memory interface settings, listed in [Table 5](#). The results of execution with a concurrent DMA transfer are not included for the STM32L4 Series. The impact of the DMA on timing is minimal and the added current consumption is approximately the same regardless of the Flash interface configuration, making the results not interesting.

One set of tests has been executed only with V_{CORE} range1 to provide a comparison with other series featured in this overview and to assess the impact of the prefetch and caches.

Other set of measurements has been executed using different latency, frequency and voltage regulator settings to assess the energy needed for different operations in case of a battery powered application.

All the measurements are taken on a single sample of Discovery board at ambient temperature. The values provided are an arithmetic mean from several measurements.

6.1 Influence of prefetch and cache with zero Flash latency

One fact must be clarified before more measurement results presentation. Neither the prefetch or caches have any influence on the execution speed when the Flash memory is available with zero latency. But the impact on the power consumption may be significant.

The prefetch actively tries to read the following instruction from the Flash memory and the energy used to read the instruction may be wasted in case of branch. In case of a correct instruction prefetch there is no timing advantage, as the instruction is also ready within one clock cycle from the Flash memory. It is recommended to disable the prefetch when the latency is zero. The measured input current difference is 10% in case of dhrystone.

On the contrary the caches tend to conserve the energy when they are activated. Both the instruction and data cache are likely to replace an access to the Flash memory with an access to the cache, which needs significantly less current. The test have proven that enabling the caches lowers the power consumption by 20%.

With both contributors combined, the STM32L476G in a worst configuration of the Flash interface, runs at significantly higher current consumption than that using optimal settings (both at 16 MHz, latency 0, V_{CORE} range 1).


6.2 Dhrystone benchmark

The Dhrystone code is executed only on optimal settings with a zero latency, where the timing is still the same. The task consists of processing 50000 cycles of the test loop, using the HSI or the HSI sourced PLL as the clock source.

Table 16. Dhrystone test using core voltage range 1 and HSI clock

Frequency	16 MHz	32 MHz							
Latency	0	1							
D-cache	1	0	0	0	1	1	1	1	0
I-cache	1	0	0	1	1	1	0	0	1
Prefetch	0	0	1	1	1	0	0	1	0
Time [ms]	2561	1552	1473	1313	1281	1283	1498	1430	1310
Average current [mA]	3.12	6.55	6.61	5.87	5.9	5.65	6.56	6.6	5.71
Energy [mJ]	24.80	31.51	30.19	23.89	23.42	22.48	30.45	29.25	23.19

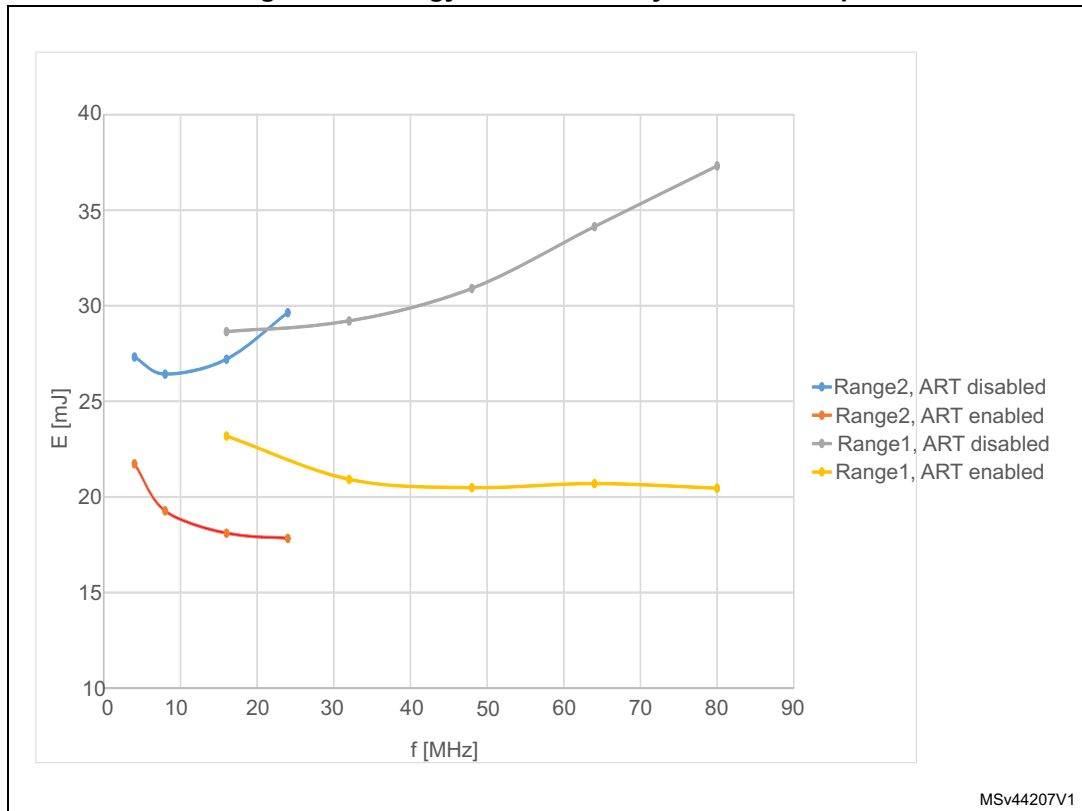
Figure 11. Dhrystone test plot of energy needed for execution



This example clearly demonstrates that while the prefetch can lead to an improved performance, especially if the instruction cache is enabled, it does not bring a significant additional advantage in case of the dhrystone test code. The prefetch complements the caches and helps in the code sections with minimum loops, where the caches cannot help.

The optimal configuration of the Flash interface being identified, how the cache behaves using different core clock speeds. A higher clock speed leads to a higher latency, forcing the core to wait for a read access to the Flash memory if the instruction and data are not available in the ART cache. The core waiting for the memory still needs energy, reducing the overall efficiency.

Figure 12. Energy cost of the dhrystone test loop



In [Figure 12](#) the same test loop of 50000 dhrystone tests is executed with different clock settings using either the MSI, or in case of a 64 MHz and a 80 MHz PLL, a module with the MSI as the source clock. The additional power consumption of the PLL causes a slight drop in the efficiency visible on the chart.

Otherwise the chart shows us that at least in case of a dhrystone test, which includes lot of loops, the ART accelerator cache is able of improving the MCU execution efficiency by increasing the core clock. This is a remarkable feature.

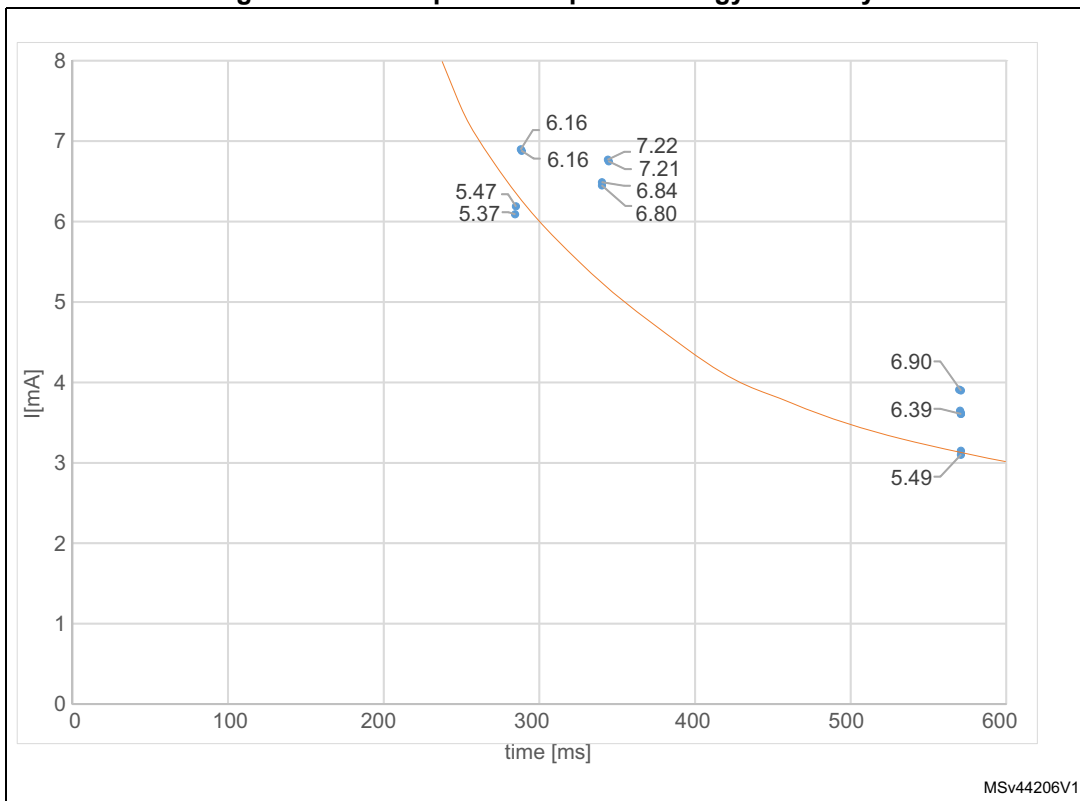
6.3 Memory read stress test

A stress test consists of executing 20 LDR instructions fetching data from program NV memory to CPU core registers in a loop of 100000 cycles. This test demonstrates mainly the power of the data cache in such situations.

Table 17. Literal measurements

Frequency	16 MHz	32 MHz							
Latency	0	1							
D-cache	1	0	0	0	1	1	1	1	0
I-cache	1	0	0	1	1	1	0	0	1
Prefetch	0	0	1	1	1	0	0	1	0
Time [ms]	570	344.5	344	340.2	284.9	284.3	288.1	288.7	340.2
Average current [mA]	3.10	6.75	6.77	6.49	6.19	6.09	6.9	6.88	6.45
Energy [mJ]	5.49	7.21	7.22	6.84	5.47	5.37	6.16	6.16	6.80

Figure 13. Literal pool chart plot of energy efficiency



In case of data literal pool loop the data cache tends to improve significantly the execution speed, while the instruction cache tends to rather contribute to the power consumption. What is not visible from the plot is that the efficiency improvement tends to grow slowly with several hundred iterations before reaching a maximum.

7 Conclusion

Each low-power STM32 microcontroller Series requires a slightly different approach to optimize the energy efficiency.

The measured results provide the guidance for decision whether or not to enable the different memory interface settings. The features like the prefetch, improving the benchmark result, also lead to a higher power consumption and the overall efficiency is dependent on the task processed by the microcontroller.

There is no significant benefit in tweaking the settings when the Flash memory latency is not in place. This makes sense only if the Flash memory contains frequently used literal pools (predefined data constants) or if the cache access leads to lower energy consumption.

With the Flash memory latency in place, the Flash interface should be setup carefully, as the performance difference between the optimal and default configuration may be significant. It is definitely possible to activate some Flash interface settings only temporarily for particular operations and disable them afterwards.

8 Revision history

Table 18. Document revision history

Date	Revision	Changes
19-Jan-2016	1	Initial release.
24-Oct-2016	2	Updated cover adding STM32L4 Series in Table 1: Applicable products . Updated Section 2: System architecture adding STM32L4 memory interface description. Added Section 3.3: STM32L4 Series device options . Added Section 6: Power consumption and performance comparison using STM32L4 Series devices . Updated Section 7: Conclusion .

IMPORTANT NOTICE – PLEASE READ CAREFULLY

STMicroelectronics NV and its subsidiaries ("ST") reserve the right to make changes, corrections, enhancements, modifications, and improvements to ST products and/or to this document at any time without notice. Purchasers should obtain the latest relevant information on ST products before placing orders. ST products are sold pursuant to ST's terms and conditions of sale in place at the time of order acknowledgement.

Purchasers are solely responsible for the choice, selection, and use of ST products and ST assumes no liability for application assistance or the design of Purchasers' products.

No license, express or implied, to any intellectual property right is granted by ST herein.

Resale of ST products with provisions different from the information set forth herein shall void any warranty granted by ST for such product.

ST and the ST logo are trademarks of ST. All other product or service names are the property of their respective owners.

Information in this document supersedes and replaces information previously supplied in any prior versions of this document.

© 2016 STMicroelectronics – All rights reserved

