

TMS320C62x™, TMS320C67x™ DSP Cache Performance on Vocoder Benchmarks

Philip Baltz

C6000™ DSP Applications

ABSTRACT

This application report discusses several multichannel vocoder benchmarks run on the TMS320C6000™ DSP platform devices including:

- TMS320C6201 (C6201)
- TMS320C6202 (C6202)
- TMS320C6203 (C6203)
- TMS320C6204 (C6204)
- TMS320C6205 (C6205)
- TMS320C6211 (C6211)
- TMS320C6701 (C6701)
- TMS320C6711 (C6711)

The C6201, C6202, C6203, C6204, C6205, and C6701 have a Harvard internal memory architecture with an on-chip program cache. The C6211 and C6711 have a two-level cache architecture with dedicated level-one data (L1D) and level-one program (L1P) caches and a unified level-two (L2) cache. The benchmarks include the following vocoders:

- G.729
- G.729a
- G.723
- GSM
- EVRC
- Combinations of the above

The results show that the C6211 performs at 82 – 92%, and the C6201 performs at 87 – 98% cycle efficiency of optimal performance, even when multiple channels or multiple vocoders are all running concurrently. With its large computational power and small device size, the C6203 achieves a very high channel density on these vocoders. The C6204 and C6211 offer good performance on these vocoders at a very low cost per channel.

These ITU vocoder algorithms are compiled from their respective C code. They have varying degrees of optimization at the C and/or assembly level, but are not fully optimized. Better implementations may be available from third party vendors. They are used here to illustrate the efficiency of the internal memory architectures of C6000 devices.

Contents

1	Device Descriptions and Performance Measurements	3
2	Vocoder Descriptions	5
3	Benchmark Descriptions	5
4	Benchmark Memory Usage	6
5	Execution Times by Benchmark Sets	7
5.1	Single Vocoder, Single Channel Execution Time	7
5.2	Single Vocoder, Multiple Channels Execution Time	8
5.3	Multiple Vocoders, Single Channel Execution Time	12
5.4	Multiple Vocoders, Multiple Channels Execution Time	13
Appendix A	Pricing Information	17

List of Figures

Figure 1.	Single Vocoder, Single Channel Execution Time by Device	8
Figure 2.	Measured Maximum Vocoder Channels by Device	9
Figure 3.	Channel Loading by Device	10
Figure 4.	Normalized Cost (\$) per Channel of Vocoders by Device	11
Figure 5.	Multiple Vocoders, Single Channel Execution Time by Device	13
Figure 6.	Measured Maximum Vocoder Channels by Device	14
Figure 7.	Channel Loading by Device	15
Figure 8.	Normalized Cost (\$) per Channel of Vocoders by Device	15
Figure 9.	C6211 and C6201 Efficiency vs. Code Size	16

List of Tables

Table 1.	TMS320C62x/TMS320C67x Device Comparison	4
Table 2.	Benchmark Memory Usage by Device	7
Table 3.	Single Vocoder, Single Channel Execution Cycles by Device	7
Table 4.	Single Vocoder, Single Channel Execution Time (ms) by Device	8
Table 5.	Measured Maximum Vocoder Channels by Device	9
Table 6.	Device Loading at Maximum Vocoder Channels	9
Table 7.	Channel Loading (MHz/Channel) by Device	10
Table 8.	Normalized Cost (\$) per Channel of Vocoders by Device	10
Table 9.	Normalized Efficiency by Device	11
Table 10.	Instruction Size Maximum Vocoder Channels	12
Table 11.	Multiple Vocoders, Single Channel Execution Cycles by Device	12
Table 12.	Multiple Vocoders, Single Channel Execution Time (ms) by Device	12
Table 13.	Measured Maximum Vocoder Channels by Device	13
Table 14.	Device Loading at Maximum Vocoder Channels	14
Table 15.	Channel Loading (MHz/Channel) by Device	14
Table 16.	Normalized Cost (\$) per Channel of Vocoders by Device	15
Table 17.	Normalized Efficiency by Device	16
Table 18.	Instruction Size (Kbytes) of Maximum Vocoder Channels by Device	16
Table A–1.	Cost Calculations	17

1 Device Descriptions and Performance Measurements

The TMS320C6000 DSPs use one of two internal memory architectures:

- Harvard
- Two-level Cache

Table 1 describes the internal memory architecture as well as other features of these devices. The TMS320C620x/TMS320C6701 devices employ a Harvard architecture internally, with dedicated program and data memories. The C6202 and C6203 share an identical internal memory architecture. These devices differ in the type of host port employed and in their core operating voltage. The C6701 shares a very similar architecture to the C6201/C6204/C6205 with the exception that the internal data memory is widened to allow loading of two 64-bit values every cycle. In comparison to the C6201, the C6202 and C6203 increase on-chip memory for both program and data. Cycle results for the C6201 were measured on the C6201 Multichannel Evaluation Module (MCEVM).

The C6202 and C6203 were measured on internal validation boards for these devices. The C6204 and C6205 cycle results are assumed to be identical to those of the C6201. Performance metrics on C6204 and C6205 refer to the C6201 data point in the graphs. C6701 cycle counts can be assumed to be the same as for the those of the C6201/C6204/C6205. However, the total performance is 83%, since the frequency of the C6701 is 167 MHz vs. 200 MHz for the others. This difference in frequency can be directly related to execution time but not to channel performance. For this reason, **note** that the C6701 has been included in the tables that show execution time performance, but *not* in tables that present channel performance.

The C6211 and C6711 processors employ a two-level cache to achieve near optimal performance on a low-cost device with only a limited amount of internal memory. Results for the C6211 were determined on a C6211 DSP Starter Kit (DSK). Results on a C6711 can be assumed to be identical to C6211, and you should refer to the C6211 data point in all graphs.

Table 1. TMS320C62x/TMS320C67x Device Comparison

Device	C6211 C6711	C6201, C6204/C6205, C6701	C6202	C6203
Frequency	150 MHz	200 MHz (C6201) 200 MHz (C6204/C6205) 167 MHz (C6701)	250 MHz	300 MHz
Internal Memory Architecture	Two-level cache	Harvard with program cache		
Total Internal Memory (Bytes)	72K	128K	384K	796K
Dedicated Program (Bytes)	4K	64K	256K	384K
Internal Program Architecture	Cache	1 block: 64K-mapped /cache	2 blocks: 64K-mapped/cache 64K-mapped	2 blocks: 64K-mapped/cache 128K-mapped
Dedicated Data Bytes	4K	64K	128K	512K
Internal Data Architecture	Cache	Mapped		
Internal Unified Memory	64K	Does not apply		
DMA Controller	16-channel Enhanced	4 channels		
Serial Port (McBSPs)	2		3	
32-bit Timers	2			
Host Port	16-bit HPI	16-bit HPI (C6201) 16-bit HPI (C6701) 32-bit XBUS (C6204) PCI (C6205)	32-bit XBUS	
IO Voltage	3.3V			
Core Voltage	1.8V	1.8V (C6201/C6701) 1.5V (C6204/C6205)	1.8 V 1.5V planned	1.5V

2 Vocoder Descriptions

The benchmark vocoders are Code Excited Linear Prediction (CELP) vocoders which process speech signals by dividing the speech signal into segments (or frames) and operating on one frame at a time.

- **GSM:** This is the enhanced full rate (EFR) GSM vocoder. The frame is 20-ms long (or 160 samples with 8kbits/sec sampling rate). The Linear Prediction Coefficients (LPC) parameters are extracted twice per frame. Each frame is equally divided further into four subframes. Pitch value and gain, and stochastic codebook and gain are obtained once per subframe (with the analysis-by-synthesis method). Encoding rate is 13 kbps.
- **EVRC:** This is the vocoder used in IS-95 cellular standard. The frame is 20-ms long. It is a variable rate vocoder with three operational rates: full, half and 1/8, depending on the contents of the current frame. For the full and half rates, LPC and pitch value and gain are obtained once per frame. Then each frame is divided into three subframes with 53, 53, and 54 speech samples. The stochastic codebook and its gain are searched once per subframe. For the 1/8 rate, only the LPC parameters and the frame energy are extracted. Encoding rate is variable.
- **G.729:** This vocoder divides speech into 10-ms frames. Each frame is divided further into two subframes. LPC is updated once per frame, and the rest of the parameters are updated once per subframe. Encoding rate is 8 kbps.
- **G.729A:** This vocoder is the same as G.729, except a simplified version of the stochastic codebook search method. Encoding rate is 8 kbps.
- **G.723:** This vocoder divides speech into 30-ms frames. Each frame is divided further into four subframes. LPC is updated once per frame, and the rest of the parameters are updated once per subframe. The G.723 vocoder can work under two encoding rates: ~5.5 kbps and ~6.5 kbps.

3 Benchmark Descriptions

Four benchmark sets were tested on each device. These benchmark sets are:

- Single Vocoder, Single Channel
- Single Vocoder, Multiple Channels
- Multiple Vocoders, Single Channel
- Multiple Vocoders, Multiple Channels

Sixteen different benchmarks were run on each processor.

1. **Single Vocoder, Single Channel:** A single channel of each vocoder (GSM, EVRC, G.729, G.729a, and G.723) was run on each device. (5 Benchmarks)
2. **Single Vocoder, Multiple Channels:** The maximum number of channels of each vocoder (GSM, EVRC, G.729, G.729a, and G.723) were run on each device. (5 Benchmarks)

3. Multiple Vocoders, Single Channel (3 Benchmarks):

- **Multichannel FrameWork (MCFW)** MCFW uses each of the vocoders (GSM, EVRC, G.729, G.729a, and G.723).
- **Multichannel Voice over Internet Protocol (MCVoIP)** MCVoIP uses the G.729, G.729A, and G.723 vocoders.
- **Multichannel wireless (MCwireless)** MCwireless uses the GSM and EVRC vocoders.

4. Multiple Vocoders, Multiple Channels (3 Benchmarks):

This was done for each of the multiple vocoder sets (MCFW, MCVoIP, MCwireless).

Identical vocoder source code was used on each processor. There were no optimizations or considerations made for cache or peripheral architecture, except to use the DMA to move data on and off chip when appropriate. For this reason, the performance numbers presented might be improved with additional optimizations.

4 Benchmark Memory Usage

Because processors have memory sizes that vary, each benchmark on each processor uses its own set of memory configurations. Table 2 details the usage of internal memory for each benchmark by processor type.

- Because of the small on-chip memory on the low-cost C6211 and C6711, both program and data are loaded in external memory, and the L2 is enabled as a 64K-byte cache.
- Because all vocoder programs are larger than 64K bytes (C6201, C6204, C6205, C6701), the program is loaded in external memory, and the program cache is enabled. On the single channel, single vocoder benchmarks; data is loaded on chip. In every other case, these data cannot fit on chip and are divided between internal and external memory. The DMA pages context data in and out of external memory on a context switch.
- The C6202 has 4x the program memory of C6201 and can typically store up to three vocoders. Thus, for all benchmarks except MCFW, all program is loaded in internal memory. For MCFW, program is loaded in external memory and the instruction cache is enabled. For single channel, single vocoder benchmarks; data is loaded in internal memory. For all other benchmarks, the data is divided between internal and external memory, and the DMA moves data on and off chip as necessary.
- On the C6203, across all benchmarks in this application note, because of its large on-chip memories, all program and data are loaded in internal memory and the instruction cache is disabled.

Table 2. Benchmark Memory Usage by Device

Benchmark Sets	Vocoder Description	Memory	C6201, C6204, C6205, C6701			
			C6211, C6711	C6202	C6203	
Single Vocoder, Single Channel	GSM	Program	Cache	Cache	Internal	Internal
		Data	Cache	Internal	Internal	Internal
Single Vocoder, Multiple Channel	EVRC	Program	Cache	Cache	Internal	Internal
		Data	Cache	DMA	DMA	Internal
Multiple Vocoders, Single Channel Each	G.729	Program	Cache	Cache	Internal	Internal
		Data	Cache	DMA	DMA	Internal
	G.729A	Program	Cache	Cache	Cache	Internal
		Data	Cache	DMA	DMA	Internal
MCwireless MCVoIP	Program	Cache	Cache	Internal	Internal	
	Data	Cache	DMA	DMA	Internal	
MCFW	Program	Cache	Cache	Cache	Internal	
	Data	Cache	DMA	DMA	Internal	
Multiple Vocoder, Multiple Channels Each	MCwireless MCVoIP	Program	Cache	Cache	Internal	Internal
		Data	Cache	DMA	DMA	Internal
	MCFW	Program	Cache	Cache	Cache	Internal
		Data	Cache	DMA	DMA	Internal

5 Execution Times by Benchmark Sets

5.1 Single Vocoder, Single Channel Execution Time

The first benchmark set measures the execution time of a single channel of each individual vocoder for each device. Table 3 lists the number of cycles required to execute 120-ms[†] data for each vocoder by device. Table 4 lists the execution time for each vocoder.

Table 3. Single Vocoder, Single Channel Execution Cycles by Device

Vocoder	C6211, C6711	C6201, C6204, C6205, C6701	C6202	C6203
GSM	1,750,696	1,507,596	1,474,472	1,474,912
G.729a	1,971,212	1,922,752	1,801,772	1,801,772
G.729	4,781,700	4,630,584	4,372,448	4,372,448
G.723	1,639,492	1,456,768	1,400,976	1,400,956
EVRC	3,067,876	2,807,328	2,656,768	2,656,824

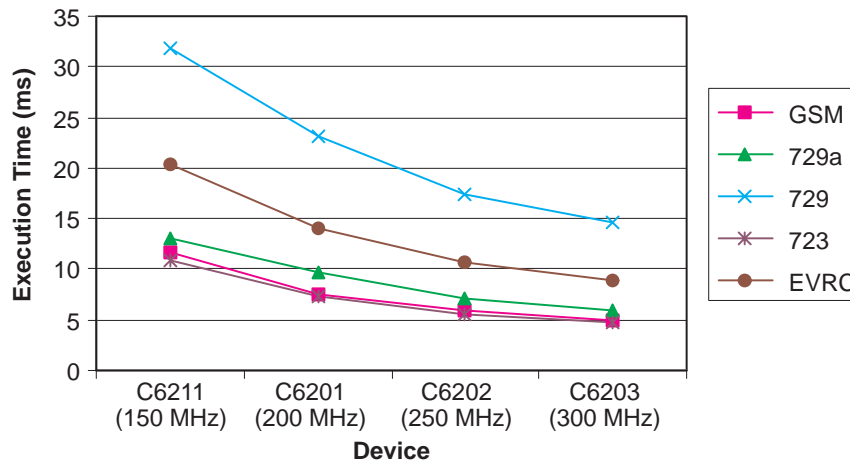
[†] 120 ms was chosen since 120 ms is a multiple of the frame size for each of the vocoders. GSM and EVRC use a 20-ms frame size, G.729 and G.729a use a 10-ms frame size, and G.723 uses a 30-ms frame size.

Table 4. Single Vocoder, Single Channel Execution Time (ms) by Device

Vocoder	C6211, C6711 (150 MHz)	C6201, C6204, C6205† (200 MHz)	C6202 (250 MHz)	C6203 (300 MHz)
GSM	11.7	7.5	5.9	4.9
G.729a	13.1	9.6	7.2	6.0
G.729	31.9	23.2	17.5	14.6
G.723	10.9	7.3	5.6	4.7
EVRC	20.5	14.0	10.6	8.9

† C6701 execution time would scale these results by 5/6.

Figure 1 illustrates execution time decreasing as the speed of the processor increases, as expected. Since the change in the execution time between the C6211 and C6201 is greater than the change between the C6201 and the C6202, or C6202 and C6203; note that there is some performance loss on C6211 because of its smaller memory. However, as shown later in this application report, the small performance loss is much less significant than the cost difference of the C6211 and the C6201 (see Appendix A for pricing information).

**Figure 1. Single Vocoder, Single Channel Execution Time by Device**

The next section shows that performance loss is limited, and the device family provides an attractive set of performance versus cost/performance tradeoffs for the system designer.

5.2 Single Vocoder, Multiple Channels Execution Time

The second benchmark set measures the maximum number of channels of each individual vocoder that can be run on each device. The number of channels is determined by measuring how many channels can be executed in 120 ms (see footnote on page 7). Table 5 lists the number of channels that can execute on each device. The cycle-count measurements are then taken from a multichannel implementation actually running on the device. Thus, all the effects of switching contexts between channels is taken into account. Figure 2 shows that the number of channels executed increases as the speed of the processor increases.

Table 5. Measured Maximum Vocoder Channels by Device

Vocoder	C6211, C6711 (150 MHz)	C6201, C6204, C6205 (200 MHz)	C6202 (250 MHz)	C6203 (300 MHz)
GSM	10	15	19	24
G.729a	8	12	15	19
G.729	3	5	6	8
G.723	11	16	21	25
EVRC	5	8	11	13

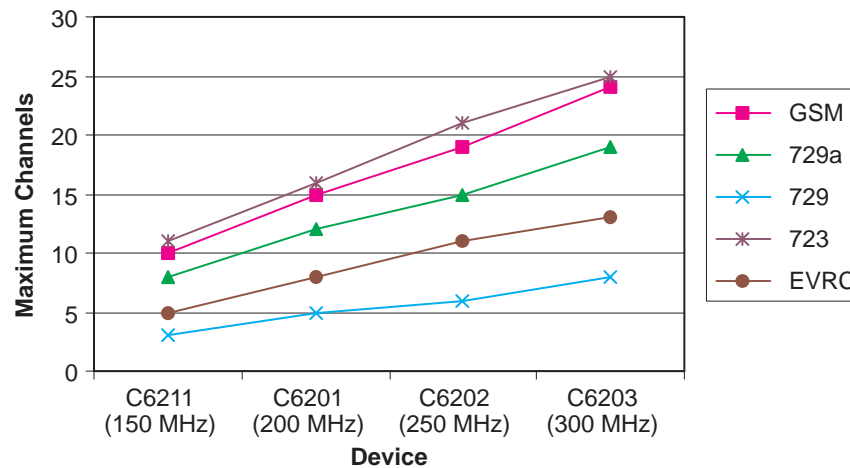


Figure 2. Measured Maximum Vocoder Channels by Device

Table 6 lists the loading for each device, by vocoder, when executing the maximum number of vocoder channels. These values are calculated by dividing the total execution time for the maximum number of channels by 120 ms (see footnote on page 7). A larger device-loading percentage is not necessarily better. For instance, if the device must execute another task in addition to executing a vocoder, then 100% loading by the vocoder means that the number of vocoder channels needs to be reduced to make execution room for the other task. If the device only executes the vocoder, then 100% utilization is optimal.

Table 6. Device Loading at Maximum Vocoder Channels

Vocoder	C6211, C6711 (%)	C6201, C6204, C6205 (%)	C6202 (%)	C6203 (%)
GSM	99	94	96	98
G.729a	94	100	95	98
G.729	80	100	88	98
G.723	98	98	99	97
EVRC	84	94	98	96

Table 7 and Figure 3 show the channel loading for each vocoder by device. Channel loading is calculated by dividing the device speed by the number of channels executed on the device. This results in the amount of device resources required to execute a single channel. This number is then scaled by the device loading shown in Table 6, so that a comparison can be made according to the amount of resources that are actually used to execute the vocoder.

Table 7. Channel Loading (MHz/Channel) by Device

Vocoder	C6211, C6711	C6201, C6204, C6205	C6202	C6203
GSM	14.9	12.5	12.6	12.3
G.729a	17.7	16.7	15.9	15.4
G.729	40.0	40.0	36.8	36.6
G.723	13.4	12.3	11.8	11.6
EVRC	25.3	23.6	22.3	22.1

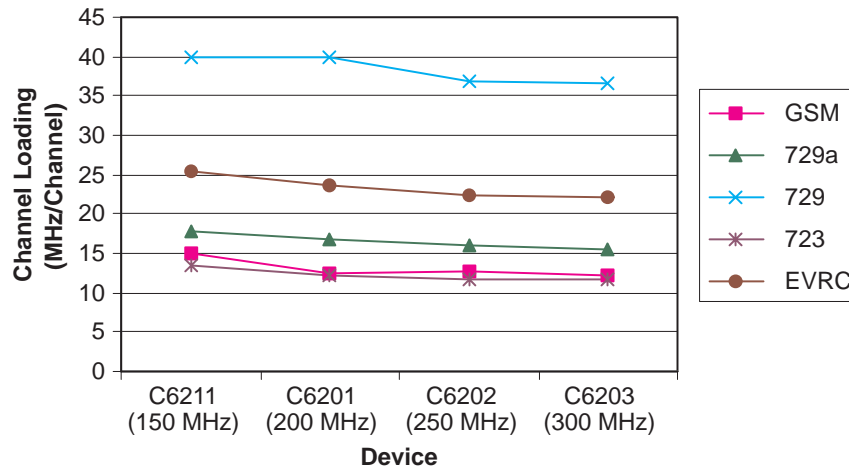


Figure 3. Channel Loading by Device

Table 8 and Figure 4 show the cost per vocoder channel when the device is executing the maximum number of vocoder channels possible (see Appendix A for pricing information). The cost is normalized by device loading to accurately represent the cost of executing the maximum channels on each device.

Table 8. Normalized Cost (\$) per Channel of Vocoders by Device

Vocoder	C6204	C6211	C6201	C6202	C6203
GSM	\$2.93	\$3.97	\$6.29	\$8.16	\$7.35
G.729a	\$3.89	\$4.71	\$8.35	\$10.25	\$9.21
G.729	\$9.33	\$10.67	\$20.04	\$23.83	\$21.88
G.723	\$2.86	\$3.57	\$6.16	\$7.65	\$6.94
EVRC	\$5.49	\$6.74	\$11.80	\$14.47	\$13.23




Figure 4. Normalized Cost (\$) per Channel of Vocoders by Device

Table 9 lists the efficiency for each device on the five vocoders. Since all program and data fit on chip for the C6203, its results are nearest to optimal. Thus, C6203 performance is used as the baseline for the comparison. The efficiency is calculated by scaling the execution time by frequency, then normalizing that by device loading. This normalized, scaled execution performance is then divided by the baseline performance on the C6203 to obtain the efficiency. The resulting numbers show that the two-level cache on the C6211 is very efficient and runs between 82% and 91% of a device which does not stall for external memory accesses. The C6201 performs at 91% to 98% efficiency, thus the C6201 cache is very efficient. The C6202 operates at above 97% efficiency, thus data paging on the C6202 does not significantly affect performance.

Table 9. Normalized Efficiency by Device

	C6211, C6711	C6201, C6204, C6205, C6701	C6202	C6203
Program	Cache	Cache	On-Chip	On-Chip
Data	Cache	Paged by DMA	Paged by DMA	On-Chip
GSM	83%	98%	98%	100%
G.729a	87%	92%	97%	100%
G.729	92%	92%	99%	100%
G.723	87%	94%	98%	100%
EVRC	87%	94%	99%	100%

Table 10 lists the size of the program by vocoder when executing the maximum number of vocoder channels. The program size is similar for each of the devices across all vocoders. This indicates that any differences in performance are not due to a disparate amount of program fetch cycles. Since each device operates on the same data on a particular vocoder, the data sizes are identical and do not cause any differences in device performance.

Table 10. Instruction Size Maximum Vocoder Channels

Vocoder	Instruction Size (Kbytes)
GSM	75
G.729a	95
G.729	95
G.723	79
EVRC	91

5.3 Multiple Vocoder, Single Channel Execution Time

The third benchmark set measures the execution time of a single channel of one of three multiple vocoder tests.

For the multiple vocoders, single channel tests; each multiple vocoder test is run with one channel of each included vocoder. For example, the MCwireless test runs 120 ms (see footnote on page 7) of data on both the GSM and EVRC vocoders. Table 11 lists the number of cycles required to execute 120 ms of data on each multiple vocoder test. Table 12 lists the execution time for each vocoder test. Figure 5 illustrates these results graphically.

Table 11. Multiple Vocoder, Single Channel Execution Cycles by Device

	C6211, C6711	C6201, C6204, C6205, C6701	C6202	C6203
MCFW	13,901,344	13,417,596	12,794,156	11,575,480
MCVoIP	8,755,896	8,481,584	7,580,364	7,551,992
MCwireless	5,051,336	4,626,560	4,106,112	4,076,024

Table 12. Multiple Vocoder, Single Channel Execution Time (ms) by Device

	C6211, C6711 (150 MHz)	C6201, C6204, C6205† (200 MHz)	C6202 (250 MHz)	C6203 (300 MHz)
MCFW	92.7	67.1	51.2	38.6
MCVoIP	58.4	42.4	30.3	25.2
MCwireless	33.7	23.1	16.4	13.6

† C6701 execution time can be assumed to be 5/6 of these values as it runs at 167 MHz.

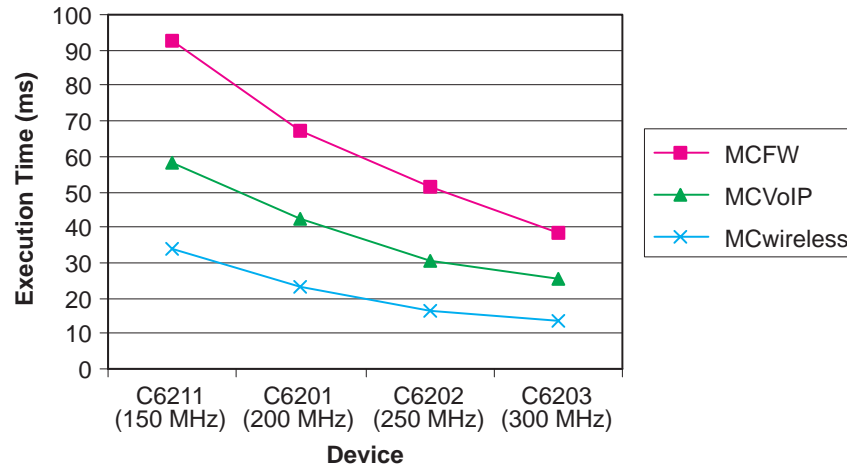


Figure 5. Multiple Vocoders, Single Channel Execution Time by Device

5.4 Multiple Vocoders, Multiple Channels Execution Time

The fourth benchmark set measures the maximum number of channels of each multiple vocoder test that can be executed on each device. The number of channels is determined by measuring how many channels can be executed in 120 ms (see footnote on page 7). The number of channels of each vocoder in each multiple vocoder test is the same. For example, on the C6201, only one channel of G.729, G.729a, and G.723 is used in the MCFW test even though there is enough computational power available to execute multiple G.723 channels concurrent with one G.729 and one G.729a channel. Table 13 lists the maximum number of channels that can execute on each device. Figure 6 illustrates that the maximum number of channels increases with operating frequency.

Table 13. Measured Maximum Vocoder Channels by Device

	C6211, C6711 (150 MHz)	C6201, C6204, C6205 (200 MHz)	C6202 (250 MHz)	C6203 (300 MHz)
MCFW	1	1	2	3
MCVoIP	2	2	3	4
MCwireless	3	5	7	8

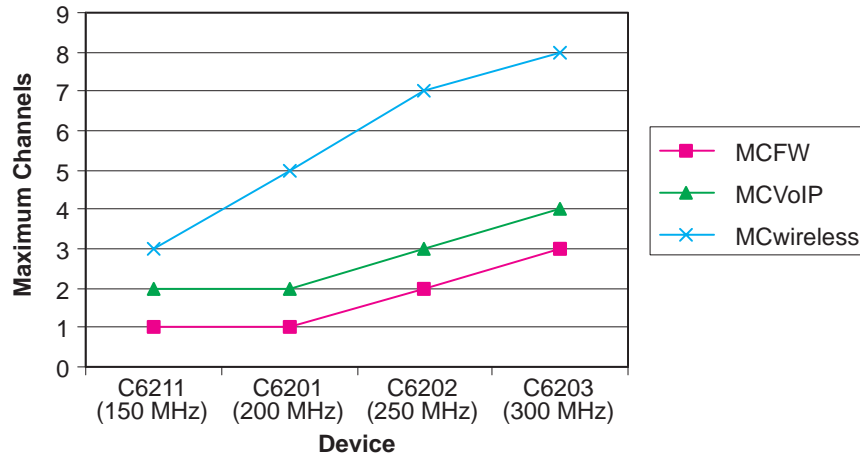


Figure 6. Measured Maximum Vocoder Channels by Device

Table 14 lists the loading for each device by vocoder when executing the maximum number of vocoder channels. These values are calculated by dividing the total execution time for the maximum number of channels by 120 ms (see footnote on page 7).

Table 14. Device Loading at Maximum Vocoder Channels

	C6211, C6711 (150 MHz)	C6201, C6204, C6205 (200 MHz)	C6202 (250 MHz)	C6203 (300 MHz)
MCFW	78%	56%	83%	98%
MCVoIP	96%	70%	77%	84%
MCwireless	82%	92%	98%	90%

Table 15 lists the channel loading for each vocoder by device. Figure 7 shows this data graphically. Channel loading is calculated by dividing the device speed by the number of channels executed on the device, and scaling it by the device-loading percentage shown in Table 14.

Table 15. Channel Loading (MHz/Channel) by Device

	C6211, C6711	C6201, C6204, C6205, C6701	C6202	C6203
MCFW	116.3	111.6	104.1	97.5
MCVoIP	71.9	70.0	63.9	63.2
MCwireless	40.9	36.7	34.8	34.1

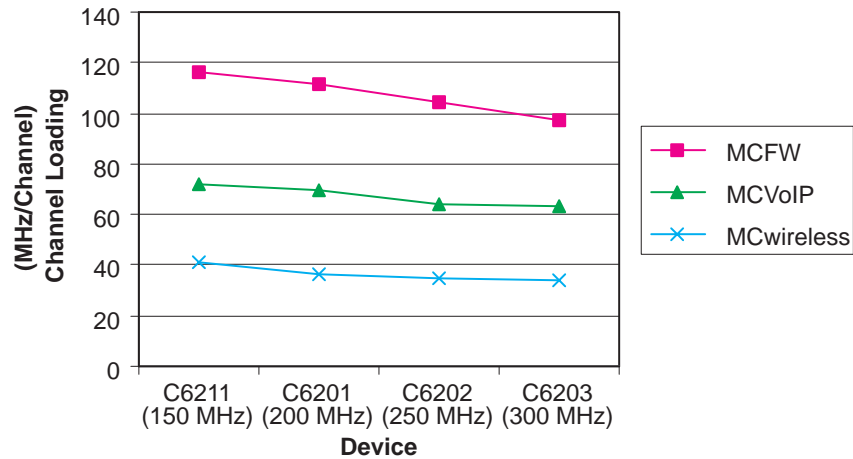


Figure 7. Channel Loading by Device

Table 16 and Figure 8 show the cost per vocoder channel by device when the device is executing the maximum number of multiple vocoder channels possible (see Appendix A for pricing information). This data is normalized by device loading.

Table 16. Normalized Cost (\$) per Channel of Vocoders by Device

	C6204	C6211	C6201	C6202	C6203
MCFW	\$26.02	\$31.00	\$55.92	\$67.44	\$58.35
MCVoIP	\$16.32	\$19.16	\$35.07	\$41.40	\$37.79
MCwireless	\$8.55	\$10.89	\$18.38	\$22.55	\$20.38




Figure 8. Normalized Cost (\$) per Channel of Vocoders by Device

Table 17 lists the efficiency for each device on each of the multi-vocoder tests. These results are calculated by normalizing and scaling execution times for each test. The resulting numbers show that the two-level cache of the C6211 is very efficient and runs between 83% and 88% of a device which does not stall for external memory accesses. The C6201 performs better, and achieves between 87% and 93% efficiency.

Table 17. Normalized Efficiency by Device

	C6211, C6711	C6201, C6204, C6205, C6701	C6202	C6203
Program	Cache	Cache	On-Chip	On-Chip
Data	Cache	Paged by DMA	Paged by DMA	On-Chip
MCFW	84%	87%	94%	100.0%
MCVoIP	88%	90%	99%	100.0%
MCwireless	83%	93%	98%	100.0%

Table 18 lists the size of the instruction for each device by vocoder when executing the maximum number of vocoder channels. The program size is similar for each of the devices across all vocoder tests. Any differences in performance are not due to a disparate amount of program fetch cycles. Since each device operates on the same data on a particular vocoder test, the data sizes are identical and do not cause any differences in device performance.

Table 18. Instruction Size (Kbytes) of Maximum Vocoder Channels by Device

Application	Instruction Size (Kbytes)
MCFW	305
MCVoIP	163
MCwireless	154

Figure 9 illustrates the efficiency numbers presented in Table 16 and Table 17 for the C6211 and C6201. This graph shows that both of these devices attain high levels of cache performance, even on very large data sizes.

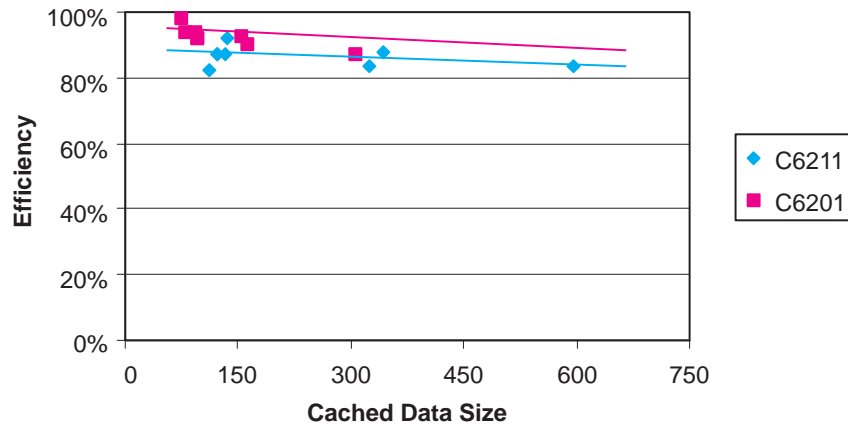


Figure 9. C6211 and C6201 Efficiency vs. Code Size

Appendix A Pricing Information

The following pricing, used for comparisons and channel costs in this Application Report, is for 25K unit purchases. These prices are accurate as of December 1, 1999.

Table A–1. Cost Calculations

Device	C6211	C6201	C6202	C6203	C6204
Price	\$25.00	\$85.21	\$146.92	\$179.53	\$31.63
SDRAM†	\$15.00‡	\$15.00‡	\$15.00‡	\$15.00	\$15.00
Total	\$40.00	\$100.21	\$161.92	\$194.53	\$46.63

† The SDRAM used for comparison is the 16-Mbit, Micron part number MT48LC1M16A1TG-7 S, with a per-unit cost of \$7.50.

‡ The cost calculations for the C6211, C6201, and C6202 devices also include the cost of external SDRAM.

IMPORTANT NOTICE

Texas Instruments and its subsidiaries (TI) reserve the right to make changes to their products or to discontinue any product or service without notice, and advise customers to obtain the latest version of relevant information to verify, before placing orders, that information being relied on is current and complete. All products are sold subject to the terms and conditions of sale supplied at the time of order acknowledgment, including those pertaining to warranty, patent infringement, and limitation of liability.

TI warrants performance of its semiconductor products to the specifications applicable at the time of sale in accordance with TI's standard warranty. Testing and other quality control techniques are utilized to the extent TI deems necessary to support this warranty. Specific testing of all parameters of each device is not necessarily performed, except those mandated by government requirements.

Customers are responsible for their applications using TI components.

In order to minimize risks associated with the customer's applications, adequate design and operating safeguards must be provided by the customer to minimize inherent or procedural hazards.

TI assumes no liability for applications assistance or customer product design. TI does not warrant or represent that any license, either express or implied, is granted under any patent right, copyright, mask work right, or other intellectual property right of TI covering or relating to any combination, machine, or process in which such semiconductor products or services might be or are used. TI's publication of information regarding any third party's products or services does not constitute TI's approval, warranty or endorsement thereof.