

Evaluation and Benchmark Testing: The Freescale Semiconductor i.MX21 Processor

White Paper

Synchromesh Computing conducted benchmark tests on the Freescale Semiconductor i.MX21 processor. Performance of the i.MX21 was compared with the Intel PXA 255 and the Texas Instruments OMAP 1611 through several standard benchmarks. The i.MX21 measures up well to its competition with some distinct, real advantages.

Synchromesh Computing
6300 Bridgepoint Parkway
Bridgepoint Square One Ste. 125
Austin, Texas 78730

Phone: 512.219.0302
Fax: 512.219.0402
E-mail: inquiry@synchromeshcomputing.com
Web: www.synchromeshcomputing.com



Synchromesh Computing



Copyright ©2004, 2006 by Synchromesh Computing. All rights reserved. Duplication or transmission without permission is prohibited.

###

Synchromesh Computing neither warrants nor guarantees that the performance or power consumption measured on processors reflects actual applications. Rather, we stipulate that test results for processors and systems are repeatable in a controlled, normal office environment.

We further assert measurements and scores are obtained fairly, without bias, and no particular platform is shown favoritism. This report is funded in part by Freescale Semiconductor; Synchromesh Computing declares such funding does not influence scoring.

Synchromesh Computing performs its evaluation testing and other work in a modern laboratory using repeatable, documented procedures. While we believe our benchmark scores are accurate, and while great care is taken, benchmarking is an incomplete science with factors occasionally beyond the complete control of the engineer.

Mistakes are possible.

In no event shall Synchromesh or any other company be held responsible for information presented in this report. Specifically, Synchromesh Computing is not liable for damages of any kind, including direct, indirect, special, consequential, or others that might arise from the reading or use of material contained in this document.

Application of technology may change expected results. The term validation means results can be duplicated in a controlled environment and that results are repeatable. Validation does not imply or guarantee performance or fitness for purpose.

All trademarks are the property of their respective owners.



SYNCHROMESH COMPUTING WHITE PAPER

THE FREESCALE SEMICONDUCTOR I.MX21 PROCESSOR

SUMMARY

The Freescale Semiconductor (Freescale) i.MX21 processor was tested by Synchronesh Computing using several standard benchmarks. Synchronesh Computing (we) subjected the Texas Instruments (TI) OMAP 1611 processor and the Intel PXA 255 processor to the same tests.

From those tests, the i.MX21 shows solid performance and energy characteristics, in particular maximum processing speed for minimal power consumption.

INDUSTRY CONTEXT

Standardized microprocessors sold by solid, reliable vendors offer a lot to systems designers. As the Gilder Report notes:

“...a custom chip can forgo all power-using circuitry not needed for its one function. But a general-purpose microprocessor is assured of 10 to 100 times the unit volumes. Thus it can justify 10 to 100 times the design effort of a single-purpose custom chip. Microprocessor designers can spend effort of clock gating, circuit-tuning, and power control that is way beyond what engineers can afford for custom hardware.”¹

Moreover, general-purpose microprocessor vendors are world leaders in offering complex solutions. They offer an abundance of development tools to enable systems designers to quickly reach the market with complete hardware/software platforms.

Having a wide selection of robust compilers, effective debuggers, and a choice of operating systems remains a core strength of the

merchant-chip vendor. The advantage is enhanced if all such features are tested to demonstrate they work well together.

Contents

Summary/Industry Context	1
Test Configuration	3
Benchmark Description.....	4
Test Notes	6
Benchmark Results.....	7
Conclusions	15

No longer can a company simply launch a microprocessor directly to market—leading-edge companies knew this all along. Technology comprises not simply a processor core, rather an entire “program” consisting of the core, peripheral processing blocks, software for those peripherals, firmware, a workable and usable software tool chain, and one or more operating systems.

Consistent profitability and progress in the embedded-processor market requires identifying real niches and filling those niches with appropriate technology... at the right time and at the right price. As an illustration, consider the following two factors:

- ❑ An enormous number of mobile phones is sold each year, estimated at 600 million, with markets in Africa and China growing the most rapidly.
- ❑ Advanced RISC Machine (ARM), Ltd. architecture increasingly dominates the mobile-phone and convergent personal digital assistant (PDA) market segments.

¹ The Gilder Technology Report, September 2004, Vol. IX, No. 9, page 3.



Thus, it makes elegant sense for companies in this market to license ARM and then proceed to add value to their own offerings.

Why is ARM a world leader in microprocessors for handheld devices?

Common opinion holds it was the first to offer very-low power 32-bit processors specifically targeted at small computing devices. Others believe ARM architecture became a critical mass of "installed base," and system developers were loathe to change processor architectures.

*Perhaps it was ARM's early licensing terms, which led to the creation of the intellectual-property microprocessor vendor industry. In any event, ARM and ARM variants form the core of the most popular integrated mobile processors, including those from Intel, Texas Instruments, and Freescale Semiconductor. When Motorola (now Freescale) gave up its own M*core processor, that set the stage for ARM dominance.*

In the future, configurable and reconfigurable processors may become pervasive, along with an open-source operating system such as Linux that might take advantage of this sort of hardware. Now, however, ARM is the worldwide leader.

An excellent means of adding value is to identify important new functionality that should be included as part of standard hardware. We observe this value-addition in many different market segments. For example:

- ❑ Via Technologies has introduced cryptography functionality into its general purpose x86 processors.
- ❑ Transmeta has introduced hardware-level virus protection for the same x86 market.²
- ❑ One of the earliest adopters of this value-adding strategy has been Advanced Micro Devices (AMD), in the form of its Geode line of the x86 processors with integrated 2D graphics.

² Fall Processor Forum 2004, disclosures from Via Technologies and Transmeta, respectively.

We wanted to be consistent with this approach in the context of these benchmarking evaluations. Thus we investigate, for wireless mobile communications and computing processors, whether inclusion of hardware-acceleration blocks is worthwhile.

Freescale, in its i.MX family of processors, focuses on multimedia acceleration. This focus acknowledges that the coming 3G revolution in wireless cellular/mobile phones entails audio and video streaming to handheld devices. Thus, Freescale is adopting the latest marketing strategy of providing general control functions along with "something extra."

The Freescale i.MX21 is based on an ARM926 core, and includes MPEG-4³ Decode and Encode acceleration in hardware, accessible to programmers by a software library. Recognizing the need for efficient performance and excellent power consumption characteristics, Freescale claims it has chosen the best performance per milliwatt envelope. It believes this approach opens up new applications and enables other handheld and portable devices, such as energy management systems and mobile gaming platforms.

Many consumer-level products will be built in the next three years that require general-purpose processors with "something extra." These special features will often take the form of hardware accelerators for specific application functions. Freescale's MPEG-4 hardware block, for both decoding and encoding, provides leading-edge technology in a general-purpose, merchant-market processor. It also supports the fundamentals: mainstream operating systems; high-quality, well-tested C and C++ compilers; debuggers; interface tools; and so on.

Four major operating systems dominate the mobile computing and communications devices (MCCD) world:

³ MPEG-4 is an ISO/IEC multimedia standard developed by MPEG (Moving-Picture Experts Group).



- ❑ Windows CE (and its variants, Smartphone and PocketPC) from Microsoft
- ❑ PalmOS from palmOne
- ❑ Symbian32 from Symbian
- ❑ Embedded Linux

Being able to support these operating systems, as well as offering a complete package to original device manufacturers (ODMs), is vital to participating in today's market. Such lack of support explains why a number of companies with excellent technology, high performance, and thrifty energy consumption haven't broken into the MCCD realm.

TEST CONFIGURATION

Synchronesh Computing is a leader in benchmarking, benchmark certification, performance analysis, and embedded-software tools analysis. For evaluation of the i.MX21 processor, we decided to test its behavior relative to that of two other leading contenders for the emerging MCCD market.

The two other offerings are the TI OMAP 1611 and the Intel PXA 255. In a future report, we plan to compare the new Intel "Bulverde" and the TI OMAP 17xx against the next-generation Freescale i.MX processor.

We also benchmarked the i.MX21's hardware

acceleration for MPEG-4, and devised a scenario to investigate the amount this hardware block offloads the main ARM processor.

Because the processor market is driven both by performance and by low power consumption, these are the areas on which our testing concentrates. In the interests of time, we selected Windows CE 4.x as the common operating system (OS) for the tests.

It is difficult to open up a mobile phone or a PocketPC to take meaningful measurements. Thus, we decided to use what are known as reference platforms that have the processor, memory subsystem, and interfaces laid out conveniently on a special test board. We attempted to configure the same hardware and OS versions for all three processors... which wasn't always easy.

Frequently, a vendor ports Windows CE to a particular reference platform or development system, takes it to market, then proceeds to the Next Big Thing. Sometimes vendors of these systems leave the business. We did our best to eliminate bias, as well as to find systems representative of clock speeds and memory configurations actually used by processor-buying customers.

As you can see in Table 1, we had difficulty obtaining exact "apples to apples" hardware environments, but we came close.

Table 1: Test Object Configurations

	Freescale	Intel	Texas Instruments
Processor	i.MX21	PXA 255	OMAP 1611
Clock speed/bus speed	266 MHz/133 MHz	398 MHz/100 MHz	180 MHz/[?] MHz
RAM	31 MB	31 MB	13 MB
Internal data cache	16K sync (4-way, 32-byte line size)	32K sync (32-way, 32-byte line size)	8K sync (4-way, 32-byte line size)
Internal instruction cache	16K sync write-thru (4-way, 32-byte line size)	32K sync (32-way, 32-byte line size)	16K sync write-thru (4-way, 32-byte line size)
Video system	240x320 16-bit	640x480 32-bit	240x320 16-bit
Development board	Freescale ADS	Accelent	TI Innovator
OS	WinCE	WinCE	WinCE
OS Version	4.2	4.1 (build 908)	4.2 (build 1088)



First, clock speed: the processors have what the vendors actually offer. Second, bus speeds differ; in fact, the bus speed of the OMAP platform could not accurately be determined.

The Intel platform—its video system is both standard VGA and 32-bit at 640x480—requires more work drawing objects on the screen than the other two systems. To be fair to Intel, we selected benchmarks that are not graphically intensive and highlighted processor performance more than the graphics subsystem. We could not completely ignore graphics, however, because all three processors are built to use graphics in a finished system

The OS version issue concerned us. We know from experience that different OS versions can have dramatic effects on performance and power consumption. We obtained the closest practicable match given the nature of this study. We did learn that device driver performance can have a more profound effect on system performance than the OS. We did our best to compensate for that effect.

Memory size for the Freescale and Intel processors is 31 MB, while TI memory size is 13 MB... an oddity perhaps. Still none of the benchmarks (nor the operating system) employ paging, so we determined the memory-size difference had little effect.

To ensure the benchmarks we built from source code were nearly identical, we used Microsoft Embedded Visual C/C++, version 4, with service packs 1, 2, and 3 installed. We built the binaries using the ARM4i target on all processors.

BENCHMARK DESCRIPTION

Benchmarks are software programs and/or workloads designed to help measure performance and other factors important in a design or buying decision.

While benchmarking is certainly used by engineering to help design better—faster,

smaller, more efficient—processors and systems, benchmark results can also play a role in sales and marketing.

Recognizing that reality, we wanted to find benchmarks that would be fair to all parties, at the same time resonating with the fast-paced, quick-adopting MCCD market. The question was whether we could find benchmarks that would highlight performance similarities and differences.

Benchmarks

- ❑ *STREAM*
- ❑ *MPEG-4 Decode/Encode*
 - *With and w/o hardware accel.*
- ❑ *HINT*
- ❑ *ByteMark*
- ❑ *BMQ*

The direction was to use a mixture of popular benchmarks available from the Internet widely used by consumers and companies evaluating technology. These include some benchmarks with solid reputations in computing circles. However, this approach caused us to run into the so-called binary wall. Some of the benchmarks are not available in source-code form, but rather are binaries you run “wholesale,” as is, on the systems.

A note about industry-standard benchmarks: Freescale Semiconductor is a member of the Embedded-Microprocessor Benchmark Consortium (EEMBC) and plans to submit the i.MX21 for benchmark certification by the EEMBC Certification Laboratory (ECL) in the future.

EEMBC is also working on standardizing power-consumption benchmarking, which has taken more than three years to develop but appears close to bearing fruit.



We found that, despite the supposed standard-core architecture for all three processors, some binaries ostensibly built for ARM in fact failed to run on some processors.

For example, SiSoft's SANDRA for Windows CE would not run on the Intel system. Neither would VObench. Yet the equally popular BMQ benchmark tests ran fine on all three platforms.

Even our own SynchroBench, quickly becoming a mainstay in our arsenal of benchmarks for all sorts of Internet-focused platforms, failed to run on all three platforms. This time it was due to incomplete media players. The TI media player had simply not been built into the system image—a common situation in the highly customizable Windows CE world. For MPEG-4, we decided to build command-line benchmarks instead.

VObench, from Virtual Office Systems is a highly popular benchmark suite for PocketPC type devices. So much so that even *PocketPC Magazine* chose it as its own benchmark suite in a roundup of PocketPC platforms in July of 2002.⁴

BMQ benchmarks CPU, memory, and graphics using five tests:

- ❑ Integer
- ❑ Floating Point
- ❑ Drawing
- ❑ Window Operations
- ❑ Memory

In Japan BMQ is very popular, and it comes in source code.

We had good success in building HINT, written by Ames Research Lab, Department of Defense. HINT has gained a reputation for being the most scalable and accurate measure of CPU and memory-subsystem performance.

“Most benchmarks measure either the number of operations that can be performed in a given time

period, or the time required to perform a give fixed calculation.

“HINT does neither; rather, it performs a particular calculation (estimating upper and lower bounds for the definite integral of a monotone-decreasing function) with ever-increasing accuracy. The accuracy of the result at any given time is called the ‘Quality;’ we may measure the improvement in quality at any given time as ‘Quality Improvements per Second,’ or QUIPS.

“As computation progresses and the quality of the result improves, more memory and more operations are required to improve the answer.... Higher is better. HINT curves are a function of raw CPU-processing power, L1 and L2 cache size and speed, and main-memory bandwidth.”

— *Nicholas Coult, PhD, assistant professor of mathematics, Augsburg College.*

STREAM, written by Dr. John McAlpin of Silicon Graphics, is another open-source, industry-standard benchmark suite that does an excellent job of measuring sustainable memory bandwidth. Embedded processors are almost always connected to the Internet, and almost always must process large amounts of data typical of multimedia bit streams.

“The STREAM benchmark is a simple synthetic program that measures sustainable memory bandwidth (in MB/s) and the corresponding computation rate for simple vector kernels. Computer cpus are getting faster much more quickly than computer memory systems. As this progresses, more and more programs will be limited in performance by the memory bandwidth of the system, rather than by the computational performance of the cpu.

“As an extreme example, several current high-end machines run simple arithmetic kernels for out-of-cache operands at 4-5% of their rated peak speeds—that means that they are spending 95-96% of their time idle and waiting for cache misses to be satisfied.

“The STREAM benchmark is specifically designed to work with datasets much larger than the available cache on any given system, so that the

⁴ http://www.pocketpcmag.com/_archives/jul02/



results are (presumably) more indicative of the performance of very large, vector-style applications.”

— John D. McAlpin, creator, *STREAM*

Remember ByteMark? Remember BYTE Magazine? The seminal journal for popular computing (even before PC Magazine), BYTE produced a benchmark suite that can be built from source code, and covers the following basic algorithms:

- ❑ Numeric Sort
- ❑ String Sort
- ❑ Bitfield
- ❑ FP Emulation
- ❑ Fourier
- ❑ Assignment
- ❑ Idea
- ❑ Huffman
- ❑ Neural Net
- ❑ LU Decomposition

ByteMark has the interesting summary metric of comparing processors to a fixed reference (in this case, an original model Pentium at 90 MHz!). But the main reason we picked Byte is illustrated in the following quotation.

“...the benchmarks are designed with a ‘small-code, large-data’ structure. The executable code portion of each test occupies less than 16KB of machine code. It can be argued that this gives only a partial picture of, say, cache performance, because an entire benchmark’s code (and, in some cases, the associated data) will fit inside at least the secondary cache of modern processors.

“Once again, in our defense, this is really the way that a good program should work. Code should exhibit locality—that is, the program should spend the majority of its time executing instructions that are close together (preferably following one another) and relatively little time jumping across large address distances.”⁵

Many embedded applications exhibit the same

characteristics: relatively small code size yet executing real-world algorithms.

We built our own version of MPEG-4 Decode and MPEG-4 Encode from open-source code, to run in a common command-line environment.

TEST NOTES

Power consumption is one of the next frontiers in benchmarking.⁶ Unsurprisingly, curiosities emerged in measuring power consumption on all three platforms. Not all vendors are especially eager to have customers measure power consumption; many do not provide documentation for how to obtain power-consumption measurements. Accessing the various vendor websites doesn’t prove too helpful, either. To its credit, Freescale includes documentation on how to measure core and bus power, a candid move for the benefit of potential customers evaluating its processors and systems.

Fortunately, our lab’s hardware engineers were able to determine power-consumption characteristics. They chose to measure it directly using precise power resistors, rather than relying on the cruder method of determining battery life. (These processors will be deployed predominantly in mobile phones, PDAs and gaming systems using different battery technologies, some of which are rapidly evolving.)

Surprisingly, the method for building Windows CE can determine the amount of power used. If software engineers are not careful, they can build an image that yields much greater power consumption than expected.

It has taken EEMBC, the industry-standard benchmark group, years to develop a standardized methodology for measuring power consumption. When this methodology is approved and the tools for implementing it are in place, Synchronesh Computing will work closely with the ECL to general industry-

⁵ BYTE’s New Benchmarks, March 1995

⁶ Ref. Synchronesh Computing’s White Paper: “Less Power, Engineer Scott!”



standard power consumption figures for its clients. Until then, we believe these measurements are fairly accurate. To get as fair a measurement as possible given the lack of documentation for TI and Intel platforms, we focus on the core power of the part.

Ultimately, what we present are both performance and power consumption benchmarking on MCCD devices, along with their corresponding analyses.

BENCHMARK RESULTS

What follows is a report of performance characteristics and power consumption scores for the three platforms. CPU clock speed, bus clock speed, and memory-access clock speed can all have an effect on performance. If the benchmark is processor-bound, CPU clock speed matters most. If the benchmark exits to main memory, bus speed and memory-access speed are important.

We used the same compiler, with the same switch settings (ARM4, same binary), to eliminate the effects of the compiler as much as possible.

But we believe optimized libraries would quickly jump many of these benchmarks upwards. This is especially true for ByteMark, in which string and memory functions dominate.

The amount of memory, especially in the case of TI's smaller memory size, prohibited some benchmarks from running. Still, since most of these benchmarks run single-threaded to completion, the result is either-or, i.e. either the processor operated properly or it did not.

STREAM

STREAM, a synthetic benchmark, consists of four sub-benchmarks.

These sub-benchmarks are known as kernels, and are as follows:

- COPY
- SCALE
- ADD
- Triad


Where Triad, naturally, is a combination of the three.

Measuring sustainable bandwidth in MB/sec. and the computational load, STREAM's primary focus is on how well a processor integrates with its memory subsystem. Pure computational loads such as ADD and SCALE show the higher clock-speed advantage of Intel, as well as indicating the TI clock-speed disadvantage compared to Freescale and Intel.

However, according to Chart 1, the COPY benchmark test reveals the i.MX21 has substantial abilities in data movement and in memory-to-memory transfers.

These characteristics indicate that processor CPU clock speed is not the only important architectural factor in performance.

Chart 1: STREAM Speed, Raw



	i.MX21	PXA 255	TI 1611
COPY	152.881	47.761	40.506
SCALE	22.069	26.446	14.222
ADD	30.968	34.043	17.455
Triad	16.642	20.601	10.573
GEOMEAN	33.392	30.679	18.057




Adjusting the scores for the clock speeds of the various processors yields an interesting, though not surprising result (ref. Chart 2).

Factoring in power consumption for the processor core, we find the higher clock speed of the Intel platform tends to cause it to burn more milliwatts (ref. Chart 3).

An interesting fact is i.MX21 has a higher difference between its minimum and maximum power consumption than the other processors. The i.MX21 also has the lowest average power consumption, over a geometric mean, than the others. This results in the highest geometric-mean performance per milliwatt on STREAM, as shown in Table 2.


The higher speed-per-MW number is, apparently, what Freescale refers to as Smart Speed™. It is clear the i.MX21 is most efficient at using its energy on this benchmark. According to Chart 3, you get the absolute lowest maximum power consumption with the TI OMAP 1611 and the absolute highest maximum power consumption with the Intel PXA 255. But if efficiency at generating work throughput counts, the i.MX21 wins.

Chart 2: STREAM Speed Adjusted for Clock Speed



	i.MX21	PXA 255	TI 1611
COPY	0.573	0.120	0.225
SCALE	0.083	0.066	0.079
ADD	0.116	0.086	0.097
Triad	0.063	0.052	0.059
GEOMEAN	0.137	0.077	0.100

Chart 3: STREAM Power Consumption



	i.MX21	PXA 255	TI 1611
Min	188.300	345.690	201.600
Max	263.300	368.235	260.400

Table 2: STREAM Power Consumption

	i.MX21	PXA 255	OMAP 1611
Max-Min Power	75.0	22.5	58.8
Avg. Milliwatt	225.8	357.0	231.0
GeoM-Perf/MW	0.161	0.086	0.078

The characteristics of the Intel processor are not surprising when you consider its heritage. The ancestor of the PXA 255 is the Intel StrongARM (previously from Digital Equipment), which was always known for much higher clock speeds than other, competing processors. The StrongARM was



designed by the same team that designed the DEC Alpha, another very-high-performance but very-high-clockspeed processor.

Table 3 shows all the power consumption readings for the three processors in the various benchmarks, except for the i.MX21 running its hardware accelerator on MPEG-4 decoding and encoding.

Note that power does not scale in a linear fashion. It turns out this pattern is repeated consistently with HINT, ByteMark, STREAM, and BMQ.

Table 3: Power Consumption in Milliwatts

	Freescale i.MX21		Intel PXA 255		TI OMAP 1611	
	MIN	MAX	MIN	MAX	MIN	MAX
STREAM execution	188.3	263.3	345.7	368.2	201.6	280.4
HINT execution	292.9	314.5	365.2	386.3	272.2	292.3
ByteMark execution	247.0	293.5	357.7	395.3	240.2	282.4
BMQ execution	183.3	294.2	330.7	398.3	213.4	277.2
Idle at Windows CE screen	105.3	108.1	106.7	109.7	110.9	114.2
Blank screen	105.3	108.1	106.7	109.7	110.9	114.2
Screen draws (fast peak)	105.3	231.3	106.7	396.8	201.6	228.5
Downloading files RS-232 (56K)	183.5	248.5	106.7	335.2	211.7	262.1

We make the following observations:

- ❑ The relative difference between min and max power consumption is greatest in the i.MX21, indicating its ability to throttle down when necessary.
- ❑ However, on screen draws and downloading files through RS-232, the PXA 255 shows superior power management.
- ❑ The highest overall power consumption is Intel, also the highest relative performance.
- ❑ The i.MX21 shows the best performance per milliwatt.
- ❑ The OMAP 1611 has the lowest power consumption, due to lower core frequency.

We really are interested in how the MPEG-4 acceleration block affects power consumption, and we cover that in the section on MPEG-4.

We again stress the importance of measuring power consumption using the real application (which is rarely available) or looking into the upcoming EEMBC power consumption benchmarks. Until these benchmarks emerge, we feel our tests provide an accurate assessment of the processors.

MPEG-4 Decode and Encode

Synchronesh Computing experimented to determine performance “out of the box,” without any acceleration on the i.MX21, and then the effect of MPEG-4 acceleration on decoding and encoding.

MPEG-4 at its heart is similar to MPEG-2 in using a discrete cosign transform (DCT) for encoding. The DCT for encoding is also called quantization, which is a fancy name for “shrinking something while maintaining its essential characteristics, though in an encoded, compressed fashion.” The reverse process, reserved for

MPEG decoding is naturally the inverse DCT (iDCT).

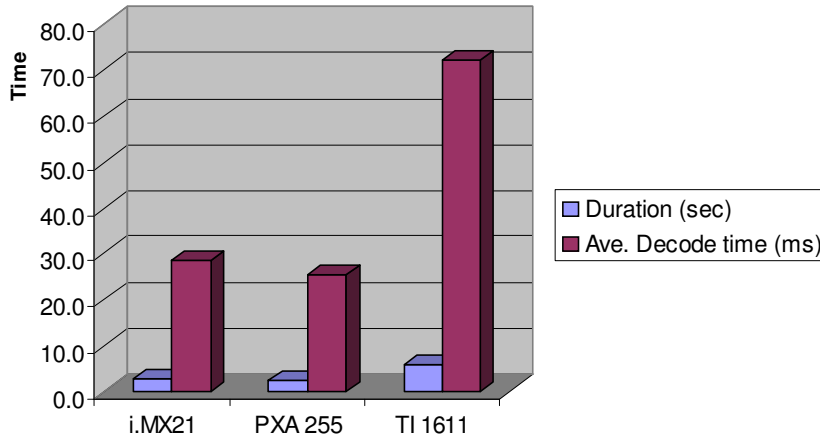
It turns out that encoding is a much more computationally intensive activity, involving a significant amount of mathematics. Smaller caches tend to hurt MPEG-4 encoding more than decoding, since once a frame is decoded it is essentially emitted—either to the graphics processor for display, or to memory for further processing.

MPEG-4 encodings must also deal with motion estimation, motion compensation, and other numerically intensive math functions. It becomes complicated: If you have a hardware-acceleration block that can do some of the work, how you apportion the workload between the main ARM processor core and the accelerator can determine total runtime. Superscalar architectures, which none of our three processors incorporates, have distinct



advantages, as do digital signal processors. Chart 4 shows the MPEG-4 Decode results.

Chart 4: MPEG-4 Decode (Carphone, 93 Frames)



	i.MX21	PXA 255	TI 1611
Duration (sec)	2.64	2.36	5.77
Ave. decode time (ms)	28.41	25.35	72.04

As we can see, the i.MX21's strong data movement abilities become clear. Despite its lower clock speed (266 MHz to the PXA's 398 MHz), it performs nearly as well. The OMAP 1611 suffers from an 8K data cache, which, combined with a probable slow memory bus, severely limits its processing speed. On the MPEG-4 Encode process we find a similar situation, illustrated in Chart 5.

Considering its lower clock speed, the i.MX21 performs well.

MPEG-4 Decode and Encode Using i.MX21 Hardware Acceleration

This subsection describes i.MX21 benchmark results when it incorporates the MPEG-4 hardware acceleration block. Although we do not expect single-cycle Decode or Encode, we do expect substantial improvement in performance—which might justify a price premium or increase the market appeal of the


processor for handheld decoding and encoding applications.

Rather than arrange a complete shootout among the i.MX21, OMAP 1611, and PXA

255, we decided to simply highlight performance improvements from use of the hardware-assisted MPEG-4 block. For example, to be fair to TI we would have to code the MPEG-4 application in Digital Signal Processing (DSP) code, which is outside the scope of this evaluation.

Table 4a shows the MPEG-4 nonoptimized scores for all three platforms. Table 4b factors in hardware assist on the i.MX21, where the lower number is better.

Chart 5: MPEG-4 Encode (Carphone, 96 Frames)



	i.MX21	PXA 255	TI 1611
Duration (sec)	10.52	8.01	19.22
Ave. Encode time (ms)	109.54	83.46	199.72



Table 4a: MPEG-4 Nonoptimized Scores

		i.MX21	PXA 255	OMAP 1611
Nonoptimized MPEG-4 Decode, 93 frames	Duration (sec)	2.64	2.36	5.77
	Avg. Decode time (ms)	28.41	25.39	62.04
Nonoptimized MPEG-4 Encode, 96 frames	Duration (sec)	10.52	8.01	19.12
	Avg. Encode time (ms)	109.54	83.46	199.17

Table 4b: MPEG-4 with Hardware Assist

		Hardware Assist	Nonoptimized
Hardware-assist MPEG-4 Decode, 94 frames	Duration (sec)	0.50	2.64
	Avg. Decode time (ms)	5.30	28.41
Hardware-assist MPEG-4 Encode	Duration (sec)	1.18	10.52
	Avg. Encode time (ms)	12.60	109.54

As you can see, a substantial performance advantage occurs by having hardware-assisted MPEG-4 decoding and encoding. In our judgment, the i.MX21's hardware-assist unit will indeed offload a large amount of

processing from the main CPU. We also wanted to explore the nature of power consumption using the hardware block, as shown in Table 5.

Table 5: Power Consumption

	Measured Voltage		Computed Current (mA)		Computed Core Power (mw)	
	Min	Max	Min	Max	Min	Max
Idle (for comparison)	15.10	15.60	68.64	70.91	113.25	117.00
MPEG-4 Encode	45.80	46.80	208.18	212.72	343.50	351.00
MPEG-4 Decode	41.20	43.20	187.27	196.36	309.00	324.00
Software MPEG-4 Encode/Decode	44.40	46.50	201.82	211.36	333.00	348.75

HINT

HINT, a synthetic benchmark, reveals pure computational processing ability. This benchmark is fairly balanced, performing a lot of computation and then memory movement; despite i.MX21's better data movement, the PXA's cache (double i.MX21's) started to take its toll. We do not believe the "ways" (of the 32 ways to i.MX21's four-way cache) made a difference.

To verify the loss of performance as a function of cache depletion, we arranged to test a Compaq iPAQ with an Intel PXA processor.

Both Intel platforms show a large dropoff when cache is empty—indicating wait states and poorer memory latency performance. The Freescale and TI processors show lesser dropoffs, reflecting a smaller proportional ratio between processor CPU speed and bus speed.

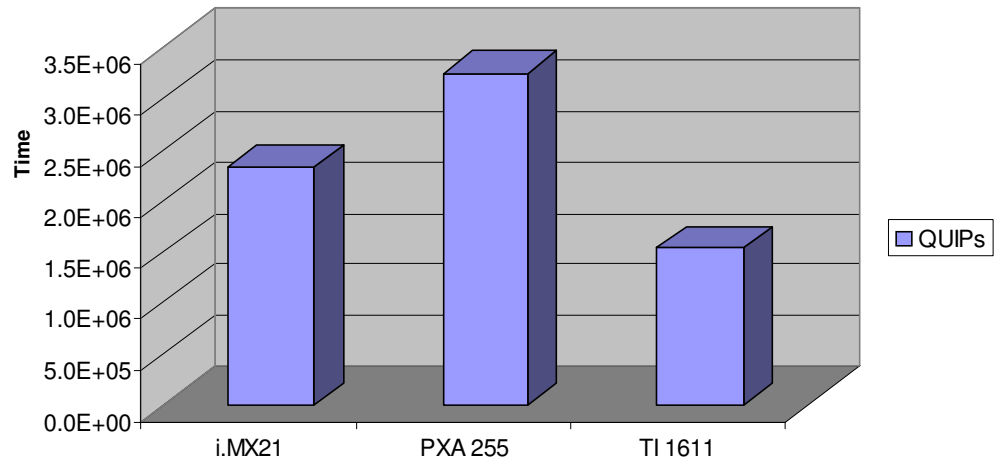
HINT (at <http://hint.byu.edu/>) maintains a large database of other benchmark scores from various systems going back to the early days of computing, and is a benchmark recognized for accuracy and fairness. The HINT site also contains a movie explaining HINT's operation.⁷

⁷ <http://hint.byu.edu/tutorials/hint.mpeg>



Chart 6: HINT Overall QUIPs Scores


What's interesting from Chart 8 is, adjusted for clock speed of the CPU, we find the i.MX21 compares well with the others. The figure exaggerates the differences. Remember the TI 1611 STREAM COPY score? Well, the TI processor's memory subsystem appears to be better than that of the Intel PXA 255 processor.



ByteMark

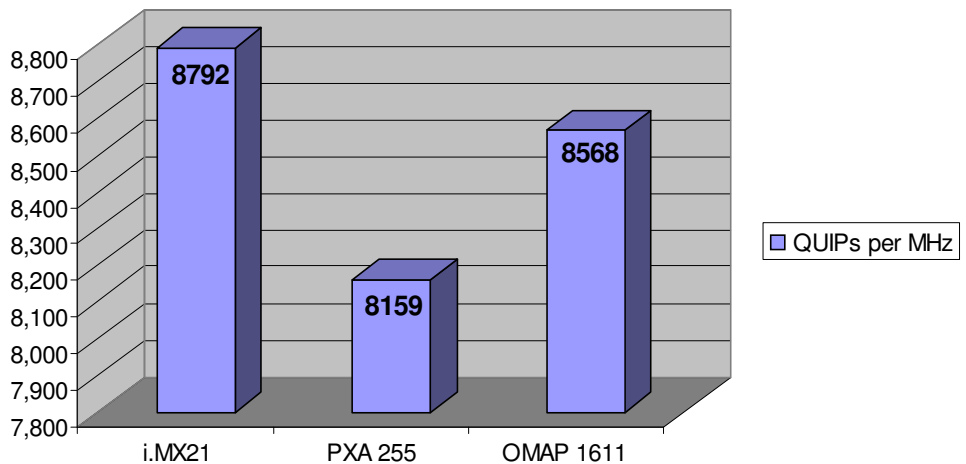
ByteMark, because it uses actual algorithm kernels, is an improvement over Dhrystone. Still, unfortunately, ByteMark tracks two things: CPU clock speed (integer and floating point) and how well a C compiler has implemented its string and memory library functions. None of these companies has special, highly optimized libraries, however, that can be linked in to produce a more efficient binary.

Chart 7: QUIPs Plot Comparison



So why do we like ByteMark? For one thing, so long as libraries are not optimized, it gives a decent representation of a set of common algorithms—EEMBC is much better. Also, it has a vast repository on the Internet of existing benchmark scores. We like that it is written in ANSI C, it exhibits “small-code, real algorithm” characteristics, and it's portable.

Chart 8: QUIPs per Megahertz





ByteMark is not a replacement for EEMBC, but it does give a very good idea of simple workloads. It is also portable to Linux, Windows CE, and practically every other operating system.

We're impressed by the second half of ByteMark for the i.MX21, especially its performance on Neural Net processing, a back-propagation network system. The i.MX21 also showed well on Idea, a block-cipher algorithm; Idea moves through data in 16-bit chunks, which facilitates i.MX21's strong data-movement characteristics.

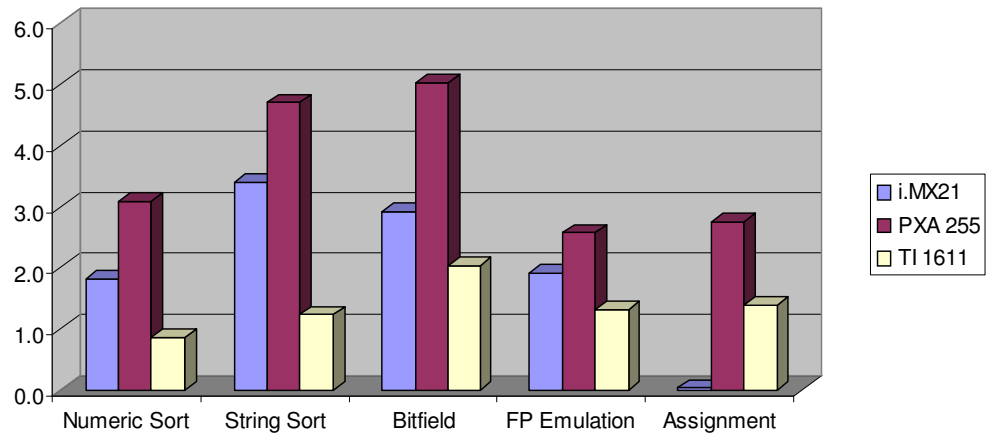
Neural Net moves through very small arrays, and is heavily dependent on the exponential function, as well as, to some extent, floating-point performance. To verify this relationship, in

Table 6 we summarize Integer and Floating Point scores—recall, 1.0 equals a Pentium original-generation processor at 90 MHz.

Table 6: Integer and Floating Point Scores

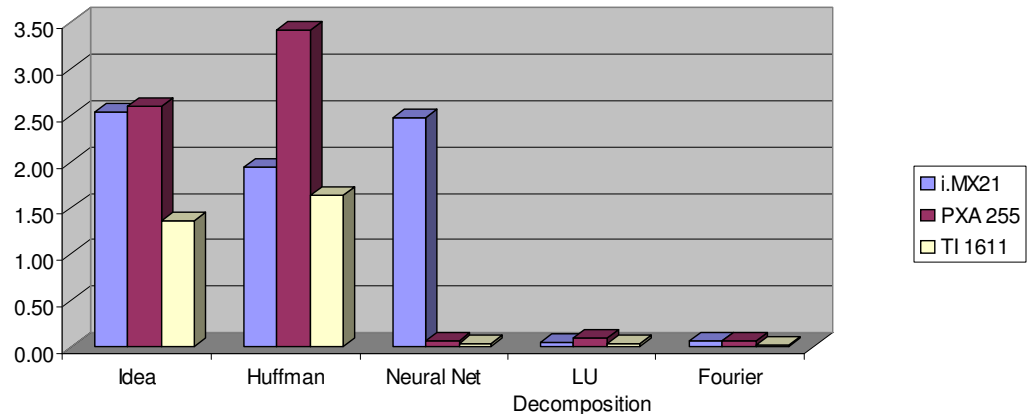
	i.MX21	PXA 255	TI 1611
Integer index	2.370	3.341	1.364
Floating point index	0.050	0.079	0.030

Chart 9: ByteMark Part 1 Results



	Numeric Sort	String Sort	Bitfield	FP Emulation	Assign
i.MX21	1.805	3.399	2.927	1.903	0.040
PXA 255	3.104	4.697	5.027	2.576	2.770
TI 1611	0.867	1.228	2.024	1.318	1.384

Chart 10: ByteMark Part 2 Results

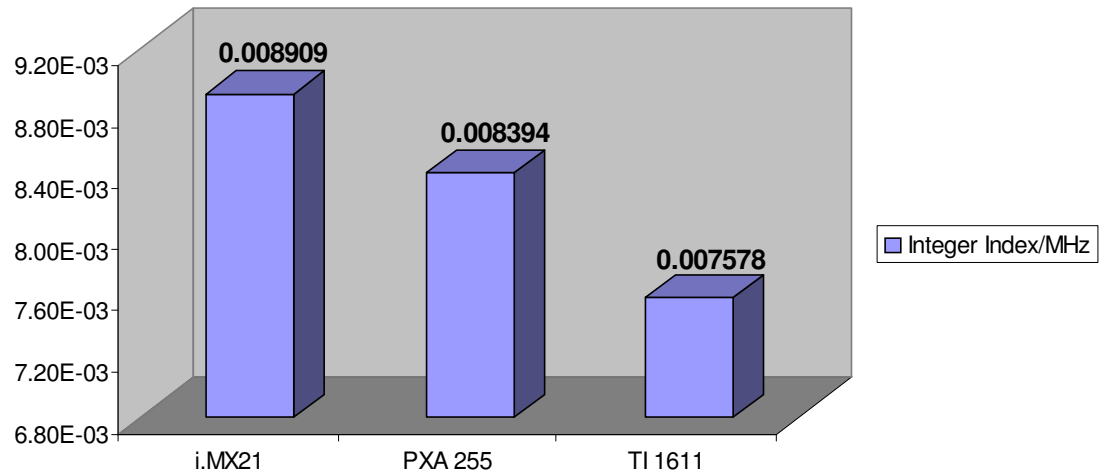


	Idea	Huffman	Neural Net	LU Decomp.	Fourier
i.MX21	2.543	1.946	2.482	0.048	0.065
PXA 255	2.600	3.415	0.073	0.092	0.072
TI 1611	1.355	1.648	0.030	0.038	0.024



Chart 11 shows the integer index factored by clock speeds, which gives a better indication of performance. The integer index is more relevant, because all three processors use software for floating-point emulation.

Chart 11: Byte Integer Index per MHz



BMQ

This series of tests benchmarks CPU, memory, and graphics using five tests:

- Integer
- Floating point
- Drawing
- Window operations
- Memory


The BMQ series is quite popular in Japan for Windows CE platforms, and it comes in source code. Although BMQ does use extremely simple arithmetic—it is really no better than Dhrystone or other simple-loop synthetic benchmarks—it includes some Windows CE “create windows and display them” and “draw on the screen” tests.

What the BMQ tests demonstrate is the Intel PXA 255 system suffers from a deficient video driver.

A Window operations score of 11 is 20-times slower than our research data indicates; a score of 220 is normal for a 400 MHz Xscale type processor.

Chart 12 shows the BMQ-suite results.

Chart 12: BMQ Test Series



	i.MX21	PXA 255	TI 1611
Integer	739	1199	520
Float	58	106	44
Draw	156	117	108
Window	131	11	79
Memory	537	1218	216
Total	324	530	193



CONCLUSIONS

Consideration of the standard benchmarks and the variations we applied yield the following conclusions:

- ❑ Clock speed often hurts “efficiency,” as measured by performance per milliwatt. Efficient data movement is important, and the i.MX21 crossbar switch does its job to keep the processor core fed.
- ❑ Small L1 and L2 caches detract from performance on larger applications (and benchmarks).
- ❑ MPEG-4 hardware assist on the i.MX21 proved worthwhile, dramatically reducing time required to encode YUV (a color standard) files into MPEG files. The strategy of using hardware acceleration to offload the CPU is sound, delivering measurable benefits to users.
- ❑ Optimizing the entire system contributes to performance. All qualities of the system must work well together, such as:
 - Processor core/CPU
 - Memory subsystem
 - Bus speed
 - Data movement
 - Software, such as the device drivers
 - Operating system
 - Libraries
 - The board itself

Failure in any of the above categories can seriously impede the success of a company’s product in the market.

- ❑ The i.MX21 per milliwatt and per clock speed lives up to the Freescale Smart Speed marketing campaign. In fact, it outperformed the TI processor and “pound for pound” ran better than the Intel processor.
- ❑ The Intel PXA 255 has good raw performance based on its higher CPU speed, higher bus speed, higher clock speed, and large caches. We look forward to testing Bulverde, the code name for the follow-on Intel processor; we also look forward to testing the follow-on to the i.MX21, the next-generation i.MX processor.
- ❑ Performance of device drivers affects the performance of Windows CE. This is a conclusion that has little to do with processors per se, rather it points out that a CE-based processor design team needs to have software expertise to achieve expected performance goals.
- ❑ Beware of binary-only benchmarks, as you never really know how they were built—or even if they will run properly.

Looking down the road: The EEMBC scores, especially for Consumer Version 2 (DENmark), should be interesting. They’ll have MPEG-2, MPEG-4, and MP3.



...for performance *matters!*

Synchronesh Computing
6300 Bridgepoint Parkway
Bridgepoint Square One Ste. 125
Austin, Texas 78730
Phone: 512.219.0302
Fax: 512.219.0402
E-mail: inquiry@synchroneshcomputing.com
Web: www.synchroneshcomputing.com

Benchmarking, Development, Business Strategies

Tel: 512.219.0302