

CRUSOE™ PROCESSOR BENCHMARK REPORT

MOBILE PLATFORM BENCHMARK RESULTS

TRANSMETA CORPORATION

FEBRUARY 3, 2000

Property of:
Transmeta Corporation
3940 Freedom Circle
Santa Clara, CA 95054 USA
(408) 919-3000
<http://www.transmeta.com>

The information contained in this document is provided solely for use in connection with Transmeta products, and Transmeta reserves all rights in and to such information and the products discussed herein. This document should not be construed as transferring or granting a license to any intellectual property rights, whether express, implied, arising through estoppel or otherwise. Except as may be agreed in writing by Transmeta, all Transmeta products are provided "as is" and without a warranty of any kind, and Transmeta hereby disclaims all warranties, express or implied, relating to Transmeta's products, including, but not limited to, the implied warranties of merchantability, fitness for a particular purpose and non-infringement of third party intellectual property. Transmeta products may contain design defects or errors which may cause the products to deviate from published specifications, and Transmeta documents may contain inaccurate information. Transmeta makes no representations or warranties with respect to the accuracy or completeness of the information contained in this document, and Transmeta reserves the right to change product descriptions and product specifications at any time, without notice.

Transmeta products have not been designed, tested, or manufactured for use in any application where failure, malfunction, or inaccuracy carries a risk of death, bodily injury, or damage to tangible property, including, but not limited to, use in factory control systems, medical devices or facilities, nuclear facilities, aircraft, watercraft or automobile navigation or communication, emergency systems, or other applications with a similar degree of potential hazard.

Transmeta reserves the right to discontinue any product or product document at any time without notice, or to change any feature or function of any Transmeta product or product document at any time without notice.

Trademarks: Transmeta, the Transmeta logo, Crusoe, the Crusoe logo, Code Morphing, LongRun and combinations thereof are trademarks of Transmeta Corporation in the USA and other countries. Other product names and brands used in this document are for identification purposes only, and are the property of their respective owners.

This document contains confidential and proprietary information of Transmeta Corporation. It is not to be disclosed or used except in accordance with applicable agreements. This copyright notice does not evidence any actual or intended publication of such document.

Copyright © 2000, Transmeta Corporation. All rights reserved.

OVERVIEW

Today's truly mobile computing devices operate in an untethered environment characterized by wireless Internet connectivity and battery-powered operation. This mobile computing paradigm is significantly different from the conventional desktop computing model, and introduces a new set of evaluation criteria for platform characterization and comparison. Truly mobile platform user's requirements and expectations include mobile application compatibility, performance sufficiency, and long battery life, in addition to the traditional peak performance requirement of desktop computing.

Computer industry benchmarks for system characterization evolved out of the desktop-workstation-server-mainframe environment that demands continuous peak performance. These traditional benchmarks are inadequate to address the requirements of the new mobile computing paradigm. Truly mobile platforms should be gauged with the attributes important to mobile users. The benchmarks appropriate for truly mobile computer products are different than benchmarks appropriate for desktop and desktop-replacement portable systems.

The benchmark methodology described in the Transmeta white paper "Mobile Platform Benchmarks" [1] was used to develop the benchmark tests and instrumentation for this report. This new benchmark technique measures performance in combination with the energy consumption and battery life penalty for performance, as well as the more traditional peak performance measurement. This mobile benchmark characterizes performance and energy consumption using real mobile application workloads in mobile systems configured and operating the way users configure and operate them, utilizing battery-powered operating modes. The systems used to generate these results were pre-production engineering prototype mobile units and reference design platforms.

This paper reports the results of a new mobile platform benchmark specifically designed to address the platform evaluation needs of the truly mobile computing industry. Crusoe processors were characterized across a range of mobile application workloads and usage profiles. The results reported use the mobile platform benchmark metrics of peak performance (WCR) and energy efficiency (WCE). Future revisions of this report will extend the range of applications workload and usage profiles, and will include a wider variety of mobile operating system environments. As commercially available end-user systems using Crusoe processors appear in the market, they will be characterized using the mobile platform benchmark and the results reported here. Representative competitive systems will also be characterized and reported for comparative purposes.

INTRODUCTION

As discussed in "Mobile Platform Benchmarks" [1], there are a variety of compelling reasons for developing more appropriate benchmark tests and metrics for better characterizing the important

attributes of truly mobile computing platforms. The methodology described in the white paper was used to develop a new set of benchmark tests and platform instrumentation for collecting the test results reported here. The focus in this initial report is processor-centric, in that complete characterization of all the system-level attributes of the test platforms was not attempted. Rather, the benchmark was structured to analyze the Crusoe processor device in relation to typical mobile application peak performance and energy efficiency. These metrics were measured across a range of common mobile application workloads.

PC INDUSTRY BENCHMARKS

Standard PC industry benchmarks were not used for this report for a variety of reasons. Fundamentally, most of these benchmarks are highly flawed with respect to mobile platform and processor characterization. The majority are synthetic benchmarks with poor correlation with user perceivable attributes. The popular application-level benchmarks have also been determined to be fundamentally flawed, and unsuitable for Crusoe characterization. These benchmarks are poor candidates for the mobile platform benchmark methodology described in [1].

The most popular PC industry battery life benchmark is a synthetic battery run-down test with a single and somewhat controversial usage profile. This benchmark uses no real end-user application code, and runs with the arbitrary fixed usage profile of 75% idle. This configuration is so narrow in scope it is highly non-representative of the vast variety of actual mobile platform usage scenarios and real workload requirements. Thus, this benchmark is of limited use for system-level mobile Windows® PC battery life testing, and shows poor correlation with many users' actual battery life observations. This benchmark is of no use whatsoever for component-level mobile device characterization, the focus of this report.

The most common PC industry benchmarks are widely used by the PC product testing organizations within various publications, MIS departments, small enterprises, PC product vendors, OEMs, web-based product review sites, and even end users. These benchmarks are generally applied to desktop and mobile system alike, with the exception of the battery life benchmarks reserved exclusively for mobile products.

The standard procedure used in many testing organizations is to run application performance benchmarks on mobile systems plugged into the wall and not operating off battery power. The system power management is also turned off to force the machine into a configuration favoring peak performance, but highly unrepresentative of system operation on battery power with full power management enabled. Since truly mobile platforms are intended to run primarily, if not exclusively, under battery power, this methodology is flawed.

Standing back from the particular issues with PC industry benchmarks, it is clear that truly mobile computing platforms require a new approach to system characterization. The Crusoe processor was developed to address the needs of mobile computer users. It should be measured using the attributes important to mobile users. Mobile computer benchmarks should measure performance in combination

with the energy consumption penalty for that performance, and they should address the issue of performance sufficiency as well as peak performance. Mobile benchmarks should measure performance and energy consumption using real mobile application workloads in mobile systems configured and operating the way users configure and operate them, under battery power. Although this may seem obvious, the industry has to date not provided an adequate set of benchmarks to address mobile computer products.

In the absence of adequate industry-standard mobile benchmarks, and in conjunction with the development of a new type of microprocessor addressing the needs of the mobile computer user, Transmeta has also developed a new benchmark methodology for evaluating mobile computer products.

MOBILE PLATFORM BENCHMARKS

The engineers at Transmeta have spent years developing a new microprocessor technology. A fundamental part of that work involves analyzing the trade-offs between software and hardware microprocessor feature implementations. This requires an in-depth analysis of the operating characteristics of PCs, operating systems, applications software, and the usage profiles of real PC users. The same understanding and knowledge gained for the optimal implementation of the Crusoe processor and Code Morphing™ software also led to a thorough rethinking of benchmarking, and the unique requirements for mobile platform benchmarks.

Benchmarks, in general, are supposed to allow comparison of alternatives and yield results that correlate with the real-world experience of users. Mobile platform benchmarks should allow comparison of various component and system-level mobile platform solutions, and the results should correlate with real-world and real-user mobile platform operating experience.

The methodology of the mobile platform benchmark is to run real-world mobile x86 application software (workloads) in usage profiles that correspond with the way users typically interact with the software and the mobile platform. A set of benchmark metrics was selected that incorporate energy usage and efficiency, peak performance (throughput), and performance sufficiency in a way that is plausible, measurable, and correlates with real-world experience.

WORKLOAD COMPLETION RATE (PEAK PERFORMANCE) METRIC

The first mobile platform benchmark metric is called Workload Completion Rate (WCR). WCR is a classic application-level peak performance metric. Another common name for this type of metric is throughput. Workload completion rate is defined as:

$$\text{WCR} = (\text{Mobile Workload Completed}) / (\text{Time to Complete Workload})$$

The measurement units for WCR is workload units per hour.

WORKLOAD COMPLETION EFFICIENCY (ENERGY EFFICIENCY) METRIC

The second mobile platform benchmark metric is called Workload Completion Efficiency (WCE). WCE describes an energy consumed-work completed relationship, and may be thought of as the energy efficiency for a given workload. Workload completion efficiency is defined as:

$$\text{WCE} = (\text{Mobile Workload Completed}) / (\text{Energy Consumed to Complete Workload})$$

The measurement units for WCE is workload units per Watt-hour.

WCE provide a much better measure of mobile systems operation then WCR (peak performance) alone, because mobile users experience the energy penalty of accelerated system activity as shortened operating life under battery operation. WCR is important, however, to establish thresholds for application performance sufficiency. Some mobile applications will be unusable if WCR (peak performance) is not sufficiently large.

WORKLOADS

The mobile platform benchmark methodology is focused on mobile platform components, systems, and applications, and is designed to reflect real-world mobile platform usage. The fundamental characteristic of mobile platform usage is untethered operation, which implies battery operation and either wireless or no connectivity. The applications selected for this mobile benchmark are typical for the mobile platforms tested, and are run in battery-powered operating mode. For Windows-based mobile PC platforms, these applications include Microsoft® Word®, Excel®, PowerPoint®, Outlook®, and Internet Explorer®. Other applications used for Windows-based platforms are Acrobat™, Quake™, MP3, and soft DVD, as well as Windows operating system functions.

MEASUREMENTS

Energy efficiency benchmark results are captured by running mobile application program workloads on instrumented test platforms and measuring the energy consumption of components (e.g. CPU), subsystems, and/or the total system. The results reported here are restricted to processor and chipset core logic (north bridge) measurements. Performance related benchmarks also capture time for completion of application workloads. Application workloads that are user rate-limited are scripted at typical user interaction rates. Systems are setup to operate in untethered mobile configurations (i.e. power management enabled and operational) whenever possible and appropriate.

REPORTING

To cover the space of mobile applications and usage profiles, reporting of metrics is not aggregated into composite benchmark results. Rather, a large family of results is reported, which in total provides a much more detailed picture of what the particular operating characteristics are for the mobile platform components and systems evaluated.

RESULTS

Table 1. Load Operating System (LOS) Workload

Windows 98 SE	Mobile Pentium® III 500 MHz	TM5400 266-533 MHz LongRun	TM3200 400 MHz 1.50V	Units
Workload Completion Time	0.0178	0.0194	0.0211	Hr
Processor + Northbridge Average Power	10.3	2.76	3.37	W
Processor + Northbridge Energy Consumed	0.183	0.0536	0.0712	WHr
Workload Units Completed	1.07	1.07	1.07	U _{LOS}
Workload Completion Rate (WCR)	60.0	55.2	50.7	U _{LOS} /Hr
Workload Completion Efficiency (WCE)	5.85	20.0	15.0	U _{LOS} /WHr

Table 2. Windows Desktop Idle (WDI) Workload

Windows 98 SE	Mobile Pentium® III 500 MHz	TM5400 266-533 MHz LongRun	TM3200 400 MHz 1.50V	Units
Workload Completion Time	0.0836	0.0833	0.0836	Hr
Processor + Northbridge Average Power	5.04	0.795	1.16	W
Processor + Northbridge Energy Consumed	0.421	0.0662	0.0970	WHr
Workload Units Completed	5.02	5.00	5.02	U _{WDI}
Workload Completion Rate (WCR)	60.0	60.0	60.0	U _{WDI} /Hr
Workload Completion Efficiency (WCE)	11.9	75.5	51.8	U _{WDI} /WHr

Table 3. Office 2000® (O2K) Workload

Windows 98 SE	Mobile Pentium® III 500 MHz	TM5400 266-533 MHz LongRun	TM3200 400 MHz 1.50V	Units
Workload Completion Time	0.133	0.167	0.178	Hr
Processor + Northbridge Average Power	7.10	2.00	2.74	W
Processor + Northbridge Energy Consumed	0.947	0.334	0.488	WHr
Workload Units Completed	8.00	8.00	8.00	U _{O2K}
Workload Completion Rate (WCR)	60.0	48.0	44.9	U _{O2K} /Hr
Workload Completion Efficiency (WCE)	8.45	24.0	16.4	U _{O2K} /WHr

Table 4. Web Browser (WEB) Workload

Windows 98 SE	Mobile Pentium® III 500 MHz	TM5400 266-533 MHz LongRun	TM3200 400 MHz 1.50V	Units
Workload Completion Time	0.164	0.164	0.166	Hr
Processor + Northbridge Average Power	6.10	1.49	2.19	W
Processor + Northbridge Energy Consumed	1.00	0.244	0.363	WHr
Workload Units Completed	9.84	9.84	9.84	U _{WEB}
Workload Completion Rate (WCR)	60.0	60.0	59.3	U _{WEB} /Hr
Workload Completion Efficiency (WCE)	9.84	40.3	27.1	U _{WEB} /WHr

Table 5. MP3 Playback (MP3) Workload

Windows 98 SE	Mobile Pentium® III 500 MHz	TM5400 266-533 MHz LongRun	TM3200 400 MHz 1.50V	Units
Workload Completion Time	0.123	0.123	0.123	Hr
Processor + Northbridge Average Power	5.46	1.15	1.74	W
Processor + Northbridge Energy Consumed	0.672	0.142	0.214	WHr
Workload Units Completed	7.38	7.38	7.38	U _{MP3}
Workload Completion Rate (WCR)	60.0	60.0	60.0	U _{MP3} /Hr
Workload Completion Efficiency (WCE)	11.0	52.0	34.5	U _{MP3} /WHr

Table 6. DVD Playback (DVD) Workload

Windows 98 SE	Mobile Pentium® III 500 MHz	TM5400 266-533 MHz LongRun	TM3200 400 MHz 1.50V	Units
Workload Completion Time	0.167	0.167	0.167	Hr
Processor + Northbridge Average Power	6.78	2.17	2.87	W
Processor + Northbridge Energy Consumed	1.13	0.362	0.479	WHr
Workload Units Completed	10.0	10.0	10.0	U _{DVD}
Workload Completion Rate (WCR)	60.0	60.0	60.0	U _{DVD} /Hr
Workload Completion Efficiency (WCE)	8.85	27.6	20.9	U _{DVD} /WHr

Table 7. Workload Completion Rate (WCR) Summary, Crusoe TM5400 vs. Mobile Pentium III

Windows 98 SE Workload	Mobile Pentium III 500 MHz	TM5400 266-533 MHz LongRun	TM5400 Mobile Pentium III Ratio
Load Operating System (LOS)	60.0	55.2	0.92
Windows Desktop Idle (WDI)	60.0	60.0	1.00
Office 2000 (O2K)	60.0	48.0	0.80
Web Browser (WEB)	60.0	60.0	1.00
MP3 Playback (MP3)	60.0	60.0	1.00
DVD Playback (DVD)	60.0	60.0	1.00

Table 8. Workload Completion Efficiency (WCE) Summary, Crusoe TM5400 vs. Mobile Pentium III

Windows 98 SE Workload	Mobile Pentium III 500 MHz	TM5400 266-533 MHz LongRun	TM5400 Mobile Pentium III Ratio
Load Operating System (LOS)	5.85	20.0	3.42
Windows Desktop Idle (WDI)	11.9	75.5	6.34
Office 2000 (O2K)	8.45	24.0	2.84
Web Browser (WEB)	9.84	40.3	4.10
MP3 Playback (MP3)	11.0	52.0	4.73
DVD Playback (DVD)	8.85	27.6	3.12

MOBILE BENCHMARK TEST CONFIGURATIONS:

Pentium III data from a commercially available major OEM mobile PC, 64M PC-100 SDRAM, 0.18u integrated L2 Coppermine @ 1.54V core on 440BX mobile module, ACPI power management at maximum power savings.

TM5400 data measured on TM FTM platform, 64M PC-133 SDRAM, 266 MHz @ 1.225V - 533 MHz @ 1.6V, ACPI at maximum power savings, LongRun™ power management enabled.

TM3200 data measured on TM WTM platform, 64M PC-133 SDRAM, ACPI power management at maximum savings.

REFERENCES

[1] Daniel McKenna, “Mobile Platform Benchmarks”, *Transmeta Corporation White Paper, February 3, 2000*.

[2] Alexander Klaiber, “The Technology Behind Crusoe™ Processors”, *Transmeta Corporation White Paper, January 2000*.

To learn more about mobile platform benchmarks, visit <http://www.transmeta.com>.

Microsoft, Windows, Office 2000, Word, Excel, PowerPoint, Outlook, and Internet Explorer are registered trademarks of Microsoft Corporation. Pentium and Pentium III are registered trademarks of Intel Corporation. Acrobat is a trademark of Adobe Systems Incorporated. Quake is a trademark of id Software, Inc.