Technical Brief

NVIDIA nForce IGP
Dynamic **A**daptive **S**peculative **P**re-processor
(DASP)

I. CPU Performance Challenges

The speed at which CPUs are able to process the most complex of calculations is astounding. While it's true that CPU core clock rates (measured in MHz) are advancing at an extremely fast pace thanks to recent, deeply pipelined, superscalar CPU architectures, it's also true that such speed advances invariably introduce other system problems including pipeline thrashing and stalling, severely reducing overall PC system performance.

Through radical latency reduction techniques, NVIDIA®'s nForce integrated, dynamic, adaptive, speculative pre-processor technology (DASP™), boosts CPU performance well beyond its nominal speed grade.

II. Dynamic Adaptive Speculative Pre-Processor (DASP)

DASP is an intelligent agent that monitors CPU requests and looks for access patterns that it can successfully predict. When it recognizes such access patterns, it exploits unused memory bandwidth to load its cache with data the CPU is expected to request later. When the CPU requests the data, it is returned to the CPU immediately rather than after waiting for the memory access. For such requests, latency is reduced from 40% to 60%. DASP is carefully engineered to make efficient use of memory bandwidth while minimizing overall latency.

DASP employs a patent-pending intelligent multi-datapath prediction/prefetching technology. It can track multiple unrelated streams of data requests and employs a proprietary selection process to select candidate cache lines to read into an on-die cache. Its fully pipelined architecture provides concurrent read/write operations from/to the cache respectively. The design is fully adaptive to ensure the lowest latency access for non-speculative CPU traffic.

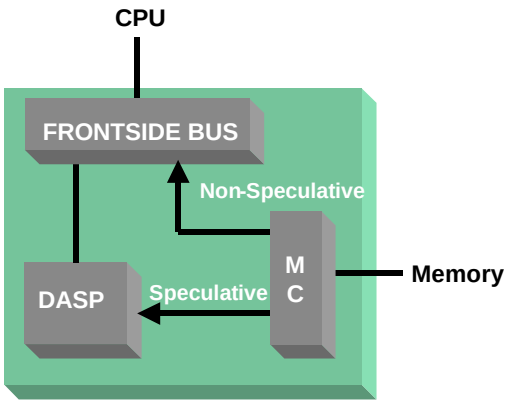

Figure 1: DASP

DASP Performance Enhancement

The following chart shows the relative performance enhancement with DASP disabled versus enabled on the same NVIDIA nForce Platform Processing Architecture. In *StreamD*, a memory-intensive application, DASP delivers a whopping 31% increase in *Copy32* performance. For *SiSoft Sandra 2001SE*, it shows 18% increase for both Integer and Floating-Point benchmarks. In CPU-intensive applications such as *Premiere 5.1* and *Windows Media Encoder 4.0*, it shows an impressive 8.82% and 6.64% performance improvement.

Benchmark	DASP Disabled	DASP Enabled	% Increase
StreamD			
Copy32	**494**	**647**	**30.97%**
Copy64	**545**	**670**	**22.94%**
SiSoft Sandra 2001 SE 1024x768x16 @75Hz			
Integer	**530**	**626**	**18.11%**
Floating Point	**692**	**819**	**18.35%**
BAPCO SysMark2000			
Premiere 5.1	**238**	**259**	**8.82%**
Windows Media Encoder 4.0	**211**	**225**	**6.64%**

All tests were run on a 1.2GHz/133MHz Athlon-based PC, with 256MB of DDR SDRAM under WinMe.

Note: Performance numbers are subject to change and are for illustrative purposes only.

III. Conclusion

With NVIDIA's nForce Platform Processor Architecture's DASP technology, users will see performance benefits in virtually all typical business and scientific applications, as well as improvements in overall system and graphics performance.

Appendix A – Glossary

BIU: Bus Interface Unit to CPU

BTB: Branch Target Buffer

CPUCLK: CPU FSB clock

Core-to-FSB ratio: Core CPU clock to Front-Side-Bus CPU clock ratio

DDR SDRAM: Double Data Rate SDRAM

DIMM: Dual In-Line Memory Module

FSB: Front Side Bus of CPU

GB/sec: Gigabytes per second

GPU: Graphics Processing Unit. The IGP integrates an NVIDIA GeForce2™ GPU on-chip. This white paper will use GPU and graphics processor interchangeably.

MC: Memory Controller

PC133: 133MHz SDRAM 64-bit DIMM system memory

DRDRAM: Rambus Direct RDRAM memory

SDR: Synchronous DRAM

SDRAM: Synchronous DRAM