

Cakewalk® Technology White Paper:
Benefits of Modern CPU Architectures for Digital Audio Applications

By Ron Kuper, CTO, Cakewalk

Benefits of Modern CPU Architectures for Digital Audio Applications

Ron Kuper, CTO Cakewalk
<http://www.cakewalk.com>

Introduction

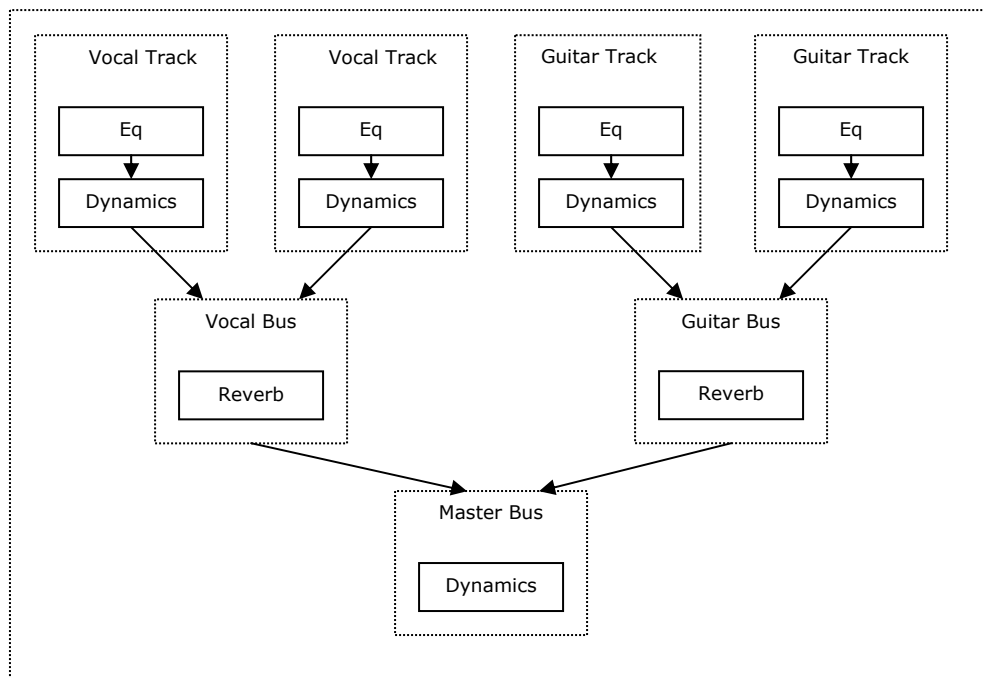
During the last 2 years several major advances have been made in PC processor designs that yield significant benefits to digital audio applications. As usual, clock speeds have been steadily increasing. But even more importantly, many modern PC processors now provide multiple cores, deeper and wider CPU registers, and the ability address a full terabyte of RAM.

Each of these new features has benefits all of its own. But for digital audio applications the whole is greater than the sum of parts, opening up new possibilities for performance, quality and precision in audio applications.

Multicore Computing

Multiprocessing isn't a feature that universally benefits every category of software application. Some types of applications are inherently more parallelizable than others, and require finer or coarser amounts of parallelism. Fortunately for musicians and products, digital audio workstations (DAWs) stand to gain a great deal from parallel processing, both at a coarse and fine level. This has benefited DAWs on multi-processor systems, and also benefits DAWs on multi-core systems.

To reap the full benefits of multiple processors (or multiple cores), a DAW can break down its work load into a dependency graph using parallel subtasks that are processed in data flow order. For example, consider the following signal flow graph of an audio project in a workstation:



Benefits of Modern CPU Architectures for Digital Audio Applications (continued)

In Cakewalk SONAR™ 5 (a multiprocessor optimized DAW), this graph is decomposed into 7 subtasks which can be processed in parallel by separate threads and/or processors. The arrows in the graph represent data flow dependencies. Each node in the graph becomes ready to run when data is available on all of its inputs. Each subtask is assigned to the next available processor as soon as it becomes ready to run. In SONAR, this data flow inspired architecture yields 30-50% performance gains when you go from a single CPU to dual CPU configuration, depending on the project configuration.

Note that parallel processing in a DAW does not come automatically. The DAW must be carefully designed to support parallelism, which means it must carefully break its workload down into smaller parts. Also, it must carefully manage memory and resource contention between different threads and processors. Finally, it should be designed for general scalability across any number of processors, not only for 2 or 4 processors.

With the advent and anticipated mass deployment of multicore processors, nearly everyone will see an immediate and significant performance boost in their favorite DAW – as long as their DAW has been designed to take advantage of parallelism, like Cakewalk SONAR has.

Deeper and Wider Registers

For most programming applications, having more CPU registers translates into better performance. Data that lives in a register is immediately accessible without any of the delay associated with RAM access.

A Pentium® 4 class CPU has 8 32-bit general purpose registers and 8 floating point registers. (An implementation detail for SSE2 means that the floating point registers are addressable either as 80-bit registers on the FPU stack or as 128-bit SSE2 registers).

With the introduction of the x64 CPU architecture, applications now have more registers available to them, and each register is also wider. The x64 architecture provides 16 general purpose registers instead of 8, and each register is 64 bits wide. Furthermore, the x64 architecture provides 16 floating point registers instead of 8.

This doubling of the number of registers can translate into tangible performance gains for digital audio workstations. Benchmarks of Cakewalk SONAR 5 showed the x64 version using equal or less CPU than the Pentium version, as shown below.

64-bit Addressing

A Pentium 4 class CPU can theoretically address a maximum of 4GB of memory. Due to implementation details of the Windows operating system, the practical memory limit for Windows applications is 3 GB.

The x64 architecture extends this limit to 1024 GB or 1 terabyte (TB). This means more of a song's data can be stored in RAM instead of on disk. For example, projects on a DAW will reduce the overhead of disk streaming by keeping audio data in RAM. Also, software synthesizers that employ sampling will be able to load more and larger sample sets.

Having the ability to access up to 1 terabyte of RAM represents a fundamental shift in how sampling based synthesizers can be designed. With sufficient RAM a sampler no longer

needs to stream off disk, and can simply load all sample data into RAM. By not having to stream data from disk, a software sampler would become more responsive to live input.

64-bit and 128-bit floating point processing

All the performance gains described above give rise to a platform that enables new, more powerful, processing systems. In particular, modern CPUs are making 64-bit double precision floating point processing a reality. 128-bit floating point processing is just around the corner.

64-bit double precision processing provides increased dynamic range and sonic clarity over 32-bit floating point systems. An IEEE 64-bit float has 52 bits of mantissa plus a sign bit, giving 53 bits of precision when doing calculations. This amount of precision means audio fidelity is maintained even with dramatic gain staging within processing elements, and also means more accurate processing for recursive DSP such as IIR filtering.

Double precision mixing is especially important when processing 24-bit PCM audio data. Consider the following test program, which simulates mixing one sample in a 3 track project with the same file on each track. Track 1's gain is +6dB. Track 2's gain is -6dB. Track 3's gain is also +6dB and the track's phase is flipped. All 3 tracks are mixed into a bus whose gain is +6dB. The signal produced on the bus output should be the same as the track data but at unity gain, since tracks 1 and 3 would cancel each other out, and the -6dB on track 2 would be offset by the +dB on the mix bus.

```
void test()
{
    const double dScale24 = (double)(1 << 23);
    const double dGain = 2.0;
    const double dAtten = 0.5;

    const int nPCM24Orig = 0x5000FF;
    double dSamp = (double)nPCM24Orig / dScale24;
    float fSamp = (float)dSamp;

    float fSum = 0;
    fSum += (float)( dGain * fSamp );
    fSum += (float)( dAtten * fSamp );
    fSum += (float)( -dGain * fSamp );
    fSum = (float)( dGain * fSum );
    int nPCM24FromFloat = (int)(fSum * dScale24 + 0.5);

    double dSum = 0;
    dSum += (double)( dGain * dSamp );
    dSum += (double)( dAtten * dSamp );
    dSum += (double)( -dGain * dSamp );
    dSum = (double)( dGain * dSum );
    int nPCM24FromDouble = (int)(dSum * dScale24 + 0.5);

    printf("Original: %lX, FromFloat: %lX, FromDouble: %lX\n",
           nPCM24Orig, nPCM24FromFloat, nPCM24FromDouble);
}
```

The output produced by this test program is (surprisingly):

```
Original: 5000FF, FromFloat: 500100, FromDouble: 5000FF
```

Benefits of Modern CPU Architectures for Digital Audio Applications (continued)

What this simple program shows is if X is a 24-bit PCM sample, and the math is done using 32-bit floating point, an inaccuracy is introduced due to summation. In this case the least significant bit is lost. If the gain adjustments are more dramatic, or more gain stages are used, then more bits can be lost.

The SSE processor architecture evolved from 4 x 32-bit float processing, to 2 x 64-bit processing. It is only inevitable that the next step will be 128-bit floating point processing. Such a system will provide truly unparalleled sonic fidelity and dynamic range.

Benchmarks

Intel® Corporation has published benchmarks showing the improved performance of the x64 architecture, using the SPEC benchmark suite. Details can be found at http://www.intel.com/performance/desktop/platform_technologies/em64t.htm.

Cakewalk performed a benchmark using SONAR 5, to compare the performance of x86 and x64 technology. Three project files were used in this benchmark:

- Project-1 consists of 23 audio tracks, 2 send buses, 2 submix buses and 1 master bus. A total of 20 Sonitus audio effects are used: (2) fx:Compressor, (2) fx:Modulator, (13) fx:Surround, (1) fx:Multiband, (1) fx:Reverb and (1) fx:Delay. SONAR's CPU meter was measured during playback.
- Project-2 consists of 19 audio tracks processing live input signal through 4 bands of EQ (a total of 76 bands of EQ). The project also contains 1 instance of Cakewalk Cyclone DXi and 1 instance of Cakewalk TTS-1 DXi. All 4 audio outputs of the TTS-1 are in use. Tracks are bused through 1 send, 2 submix and 1 master bus. A total of 16 Sonitus audio effects are used: (6) fx:Gate, (6) fx:Compressor, (1) fx:Delay, (1) fx:Multiband and (2) fx:Reverb. SONAR's CPU meter was measured while the SONAR engine was processing live data.
- Project-3 extends Project-2 with the addition of 2 instances of the PSYN-2 DXi. Also, 3 of the project's DXi's contain MIDI data on their MIDI tracks. SONAR's CPU meter was measured while playing back live MIDI rendered through the software synthesizers.

The test system was an Intel Pentium 4 with EM64T technology running at 3.6 GHz, with 2GB of RAM. Audio streaming was through the Edirol UA-1000 USB 2.0 audio hardware. To minimize driver overhead testing was done at a relatively high audio buffer size (latency) of 98 msec.

The system was configured to dual boot Windows XP Professional x64 Edition and Windows XP Professional (32-bit version). SONAR 5 was tested on both operating systems using the same Edirol UA-1000 hardware.

Benchmarks were performed using both the SONAR 5 32-bit single precision floating point engine and the 64-bit double precision floating point engine. The results are tabulated below.

Benefits of Modern CPU Architectures for Digital Audio Applications (continued)

Single precision floating point engine

	SONAR 5 x86 WinXP x86	SONAR 5 x86 WinXP x64	SONAR 5 x64 WinXP x64	Approx. Speedup
Project-1	24%	24%	18%	30%
Project-2	25%	25%	22%	14%
Project-3	32%	33%	30%	10%

Double precision floating point engine

	SONAR 5 x86 WinXP x86	SONAR 5 x86 WinXP x64	SONAR 5 x64 WinXP x64	Approx. Speedup
Project-1	26%	26%	20%	30%
Project-2	30%	31%	26%	19%
Project-3	35%	36%	31%	16%

The results of the benchmark show a distinct performance gain for the x64 edition of SONAR 5. Most revealing is that the x64 edition can operate using its double-precision mix engine and still be more efficient than the x86 edition running the single-precision mix engine.

It's also worth noting that the 32-bit version of SONAR ran with little or no performance degradation on Windows XP x64 Edition. This is consistent with other published benchmarks, which seem to demonstrate that the WOW64 emulation layer in Windows XP x64 Edition is nearly transparent.

Conclusion

A digital audio workstation is a complex system with very high processor demands. Cakewalk SONAR 5 x64 demonstrates how a combination of an efficient multithreaded design, and effective utilization of all of the capabilities x64 has to offer, yields a significant performance gain over x86 technology.

Cakewalk is a registered trademark and SONAR, and the Cakewalk logo are trademarks of Twelve Tone Systems, Inc. Other trademarks are held by their respective owners.