

In order to compare ConcRT, OpenMP and TBB technologies, we implemented a few algorithms from different areas of numeric computation and compared their performance with C++ equivalents. Four functions were implemented.

The first two functions were taken from linear algebra. Function 1 computes determinant of a matrix with non-zero leading principal minors. Algorithm of this function is based on LU decomposition of the input matrix. The second function computes a conjugated transposed matrix from a complex input. Function 3 is taken from signal processing. It computes the convolution of two matrices. The algorithm explicitly implements the definition of 2D convolution. Finally, function four is a solution to a system of ordinary differential equations via implementation of the well-known N-body problem. The Runge – Kutta algorithm of 4-th order was used to solve the system of differential equations. We have the following equation:

$Y'(t)=F(t,Y)$; $Y(t_0)=Y_0$, then the value of the next point will be computed as:
 $Y_{n+1}=Y_n+h/6 \cdot (K_1+2 \cdot K_2+2 \cdot K_3+K_4)$, where $K_1=F(t,Y)$,
 $K_2=F(t+h/2, Y+h/2 \cdot K_1)$, $K_3=F(t+h/2, Y+h/2 \cdot K_2)$, $K_4=F(t+h, Y+h \cdot K_3)$.

All four algorithms were initially implemented in C++ and then converted into parallel form using the OpenMP, ConcRT and TBB technologies. Tests were run at both: Windows 7 and Windows HPC Server 2008. The Windows 7 machine has next characteristics: Quad-Core Intel(R) Xeon(R) E5440 2.83GHz Processor, 64-bit OS with 4G RAM. The Windows HPC Server 2008 machine has next characteristics: Quad-Core Intel Xeon E5110 1.60GHz Processor, 64-bit OS with 4G RAM.

All routines were built with VS2010 compiler, OpenMP 2.0 and TBB 3.0 libraries were used to implement parallel versions of programs.

All tests were run 10+ times. Times of execution were measured using GetTickCount() routine. Average execution time was used in this report to compute speed-up. All comparative charts are placed in Appendices 1-3.

Brief review of technologies

Concurrency Runtime (ConcRT) technology includes 2 libraries: Parallel Patterns Library (PPL) and Asynchronous Agents Library.

The PPL provides general-purpose containers and algorithms for performing fine-grained parallelism. The PPL enables imperative data parallelism by providing parallel algorithms that distribute computations on collections or on sets of data across computing resources. It also enables task parallelism by providing task objects that distribute multiple independent operations across computing resources.

The Asynchronous Agents Library (or just Agents Library) provides both an actor-based programming model and message passing interfaces for coarse-grained dataflow and pipelining tasks. Asynchronous agents enable you to make productive use of latency by performing work as other components wait for data.

Parallel patterns and Agents libraries present in Visual Studio 2010 (VS 2010).

OpenMP uses the fork-join model of parallel execution. Although this fork-join model can be useful for solving a variety of problems, it is somewhat tailored for large array-based applications. OpenMP is intended to support programs that will execute correctly both as parallel programs (multiple threads of execution and a full OpenMP support library) and as sequential programs (directives ignored and a simple OpenMP stubs library).

MS OpenMP v2.5 is built-in in VS 2010. OpenMP is free to be used. Also, there is a Fortran version of OpenMP.

Intel Threading Building Blocks is a runtime-based parallel programming model for C++ code that uses threads. It consists of a template-based runtime library to help you harness the latent performance of multicore processors. Intel Threading Building Blocks is used to write scalable applications that:

- Specify logical parallel structure instead of threads
- Emphasize data parallel programming
- Take advantage of concurrent collections and parallel algorithms

Both, commercial aligned open source and commercial versions of TBB are available.

Summarizing

For most mathematical functions parallelism is achieved by executing loops in parallel or by executing independent tasks simultaneously. These mechanisms are available in all examined parallel technologies. After analyzing performance results, we achieved next results.

OpenMP technology showed the biggest speed-up on both Windows 7 and Windows HPC Server 2008. Also, this technology showed the best scalability by data size, i.e. the same code works equivalently good on data of different sizes. The code changes to convert serial code into parallel were minimal comparing to ConcRT and TBB. OpenMP can be used in both, C and C++ projects. If OpenMP program executes on 1 core machine, then the slowdown is small.


TBB library demonstrated almost the same speed-up as OpenMP, the scalability by data size also looks good. But code changes to convert serial code into parallel are significantly bigger than changes for OpenMP - when parallelizing for-loop, the inner action must be wrapped like a function object, or as lambda function (this feature is present in TBB v3.0). TBB can be used only in C++ projects. The slowdown is small, if TBB application is executed on 1 core machine.

ConcRT was the slowest of all tested technologies. The code changes to convert serial code into parallel are only slightly easier than for TBB technology and much bigger than for OpenMP. ConcRT technology can be used only in C++ projects. But, while implementing parallel version of chosen mathematical algorithms, only PPL was used. So, ConcRT has much wider range of use then OpenMP and TBB, which can't be demonstrated by chosen mathematical problems.


Anyway, all these technologies can be used simultaneously, so in one project one can mix TBB, OpenMP and ConcRT. My opinion is that for mathematical algorithms, OpenMP is the best technology.

Appendix 1


Below are results of the execution time measurements of ConcRT, OpenMP and TBB programs at Windows 7 machine using 1, 2, 3, and 4 cores.



Speed-up of ConcRT technology comparing with C++, Windows 7




Speed-up of TBB technology comparing with C++, Windows 7




Speed-up of OpenMP technology comparing with C++, Windows 7

Appendix 2


Below are results of the execution time measurements of ConcRT, OpenMP and TBB programs at Windows HPC Server 2008 machine using 1, 2, 3, and 4 cores.



Speed-up of ConcRT technology comparing with C++, Windows HPS Server 2008




Speed-up of TBB technology comparing with C++, Windows HPS Server 2008



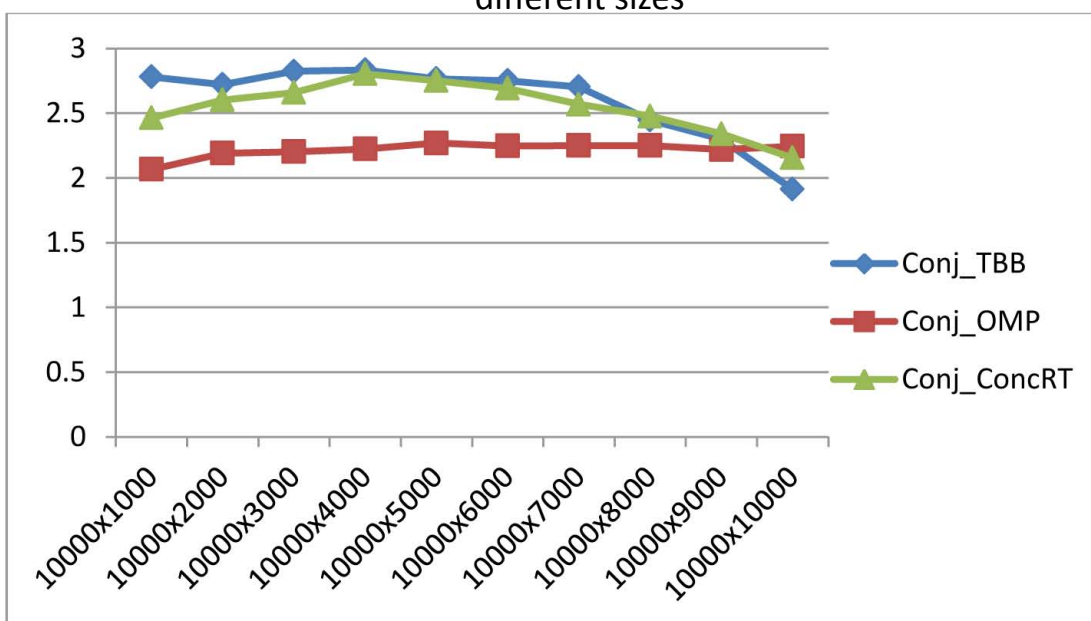
Speed-up of OpenMP technology comparing with C++, Windows HPS Server 2008

Appendix 3


Below are results of the speed-up of parallel applications comparing to serial equivalents for different sizes of input parameters. Applications were tested on Windows 7 machine using 4 cores.




Speed-up of parallel technologies for computing determinant of matrices of different sizes



Speed-up of parallel technologies for conjugated transpose of matrices of different sizes



Speed-up of parallel technologies for convolution of matrices of different sizes



Speed-up of parallel technologies for solving N-body problem for different amount of bodies