

MODULE – III

Concurrent Computing (Chapter 6)

Throughput computing focuses on delivering high volumes of computation in the form of transactions.

Advances in hardware technologies led to the creation of multi core systems, which have made possible the delivery of high-throughput computations

Multiprocessing is the execution of multiple programs in a single machine, whereas multithreading relates to the possibility of multiple instruction streams within the same program.

This chapter presents the concept of multithreading and describes how it supports the development of high-throughput computing applications.

6.1 Introducing parallelism for single-machine computation

Parallelism has been a technique for improving the performance of computers since the early 1960s.

In particular, multiprocessing, which is the use of multiple processing units within a single machine, has gained a good deal of interest and gave birth to several parallel architectures.

Asymmetric multiprocessing involves the concurrent use of different processing units that are specialized to perform different functions. **Symmetric multiprocessing** features the use of similar or identical processing units to share the computation load.

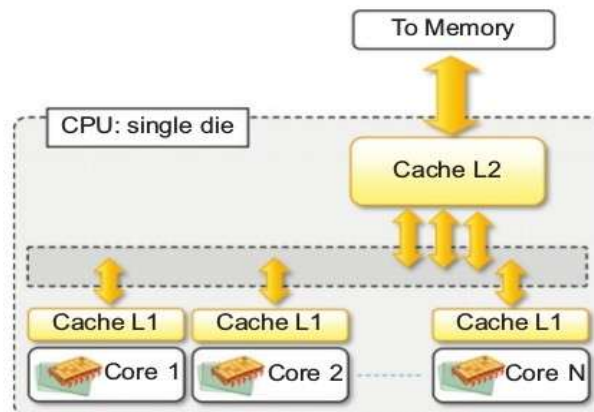


FIGURE 6.1

Multicore processor.

Multicore systems are composed of a single processor that features multiple processing cores that share the memory. Each core has generally its own L1 cache, and the L2 cache is common to all the cores, which connect to it by means of a shared bus, as depicted in Figure 6.1. Dual- and quad-core configurations are quite popular nowadays and constitute the standard hardware configuration for commodity computers. Architectures with multiple cores are also available but are not designed for the commodity market. Multicore technology has been used not only as a support for processor design but also in other devices, such as GPUs and network devices, thus becoming a standard practice for improving performance.

Multiprocessing is just one technique that can be used to achieve parallelism, and it does that by leveraging parallel hardware architectures.

In particular, an important role is played by the operating system, which defines the runtime structure of applications by means of the abstraction of process and thread. A process is the runtime image of an application, or better, a program that is running, while a thread identifies a single flow of the execution within a process. A system that allows the execution of multiple processes at the same time supports multitasking. It supports multithreading when it provides structures for explicitly defining multiple threads within a process.

6.2 Programming applications with threads

Modern applications perform multiple operations at the same time. The use of threads might be implicit or explicit.

Implicit threading happens when the underlying APIs use internal threads to perform specific tasks supporting the execution of applications such as graphical user interface (GUI) rendering, or garbage collection in the case of virtual machine-based languages.

Explicit threading is characterized by the use of threads within a program by application developers, who use this abstraction to introduce parallelism. Common cases in which threads are explicitly used are I/O from devices and network connections, long computations, or the execution of background operations.

6.2.1 What is a thread?

A thread identifies a single control flow, which is a logical sequence of instructions, within a process. By logical sequence of instructions, we mean a sequence of instructions that have been designed to be executed one after the other one.

Operating systems that support multithreading identify threads as the minimal building blocks for expressing running code.

Each process contains at least one thread but, in several cases, is composed of many threads having variable lifetimes. Threads within the same process share the memory space and the execution context.

In a **multitasking** environment the operating system assigns different time slices to each process and interleaves their execution. The process of temporarily stopping the execution of one process, saving all the information in the registers, and replacing it with the information related to another process is known as a **context switch**.

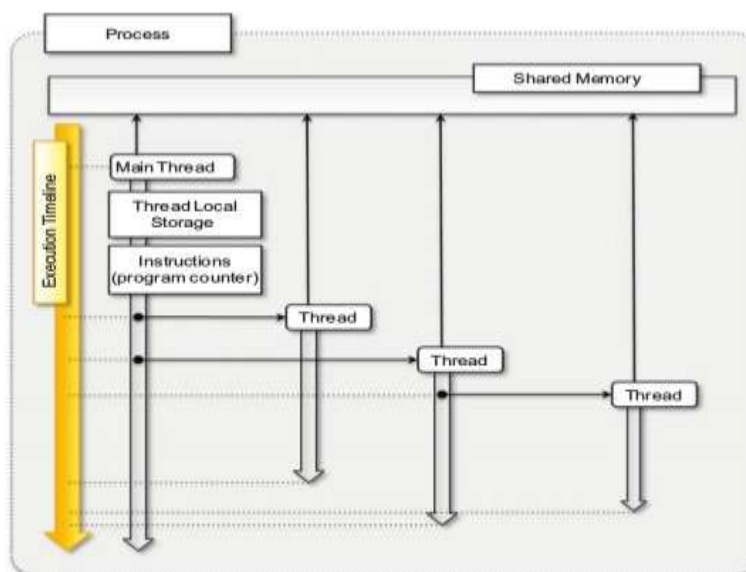


FIGURE 6.2

The relationship between processes and threads.

Figure 6.2 provides an overview of the relation between threads and processes and a simplified representation of the runtime execution of a multithreaded application. A running program is identified by a process, which contains at least one thread, also called the main thread. Such a thread is implicitly created by the compiler or the runtime environment executing the program. This thread is likely to last for the entire lifetime of the process and be the origin of other threads, which in general exhibit a shorter duration. As main threads, these threads can spawn other threads. There is no difference between the main thread and other threads created during the process lifetime. Each of them has its own local storage and a sequence of instructions to execute, and they all share the memory space allocated for the entire process. The execution of the process is considered terminated when all the threads are completed.

6.2.2 Thread APIs

1 POSIX Threads

2 Threading support in java and .NET

1 POSIX Threads

Portable Operating System Interface for Unix (POSIX) is a set of standards related to the application programming interfaces for a portable development of applications over the Unix operating system flavors. Standard POSIX 1.c (IEEE Std 1003.1c-1995) addresses the implementation of threads and the functionalities that should be available for application programmers to develop portable multithreaded applications.

Important to remember from a programming point of view is the following:

- A thread identifies a logical sequence of instructions.
- A thread is mapped to a function that contains the sequence of instructions to execute.
- A thread can be created, terminated, or joined.
- A thread has a state that determines its current condition, whether it is executing, stopped, terminated, waiting for I/O, etc.
- The sequence of states that the thread undergoes is partly determined by the operating system scheduler and partly by the application developers.
- Threads share the memory of the process, and since they are executed concurrently, they need synchronization structures.
- Different synchronization abstractions are provided to solve different synchronization problems.

2 Threading support in java and .NET

Languages such as Java and C# provide a rich set of functionalities for multithreaded programming by using an object-oriented approach. Both Java and .NET execute code on top of a virtual machine, the APIs exposed by the libraries refer to managed or logical threads. These are mapped to physical threads.

Both Java and .NET provide class Thread with the common operations on threads:

start, stop, suspend, resume, abort, sleep, join, and interrupt.

Start and stop/abort are used to control the lifetime of the thread instance.

Suspend and resume are used to programmatically pause and then continue the execution of a thread.

Sleep operation allows pausing the execution of a thread for a predefined period of time.

Join operation that makes one thread wait until another thread is completed.

Waiting states can be interrupted by using the interrupt operation.

6.2.3 Techniques for parallel computation with threads

Developing parallel applications requires an understanding of the problem and its logical structure. Decomposition is a useful technique that aids in understanding whether a problem is divided into components (or tasks) that can be executed concurrently. It allows the breaking down into independent units of work that can be executed concurrently with the support provided by threads.

1 Domain decomposition

2 Functional decomposition

3 Computation vs. Communication

1 Domain decomposition

Domain decomposition is the process of identifying patterns of functionally repetitive, but independent, computation on data. This is the most common type of decomposition in the case of throughput computing, and it relates to the identification of repetitive calculations required for solving a problem. The master-slave model is a quite common organization for these scenarios:

- The system is divided into two major code segments.
- One code segment contains the decomposition and coordination logic.
- Another code segment contains the repetitive computation to perform.

- A master thread executes the first code segment.
- As a result of the master thread execution, as many slave threads as needed are created to execute the repetitive computation.
- The collection of the results from each of the slave threads and an eventual composition of the final result are performed by the master thread.

Embarrassingly parallel problems constitute the easiest case for parallelization because there is no need to synchronize different threads that do not share any data. Embarrassingly parallel problems are quite common, they are based on the strong assumption that at each of the iterations of the decomposition method, it is possible to isolate an independent unit of work. This is what makes it possible to obtain a high computing throughput. If the values of all the iterations are dependent on some of the values obtained in the previous iterations, the problem is said to be **inherently sequential**. **Figure 6.3** provides a schematic representation of the decomposition of embarrassingly parallel and inherently sequential problems.

The matrix product computes each element of the resulting matrix as a linear combination of the corresponding row and column of the first and second input matrices, respectively. The formula that applies for each of the resulting matrix elements is the following:

$$C_{ij} = \sum_{k=0}^{n-1} A_{ik}B_{kj}$$

Two conditions hold in order to perform a matrix product:

- Input matrices must contain values of a comparable nature for which the scalar product is defined.
- The number of columns in the first matrix must match the number of rows of the second matrix.

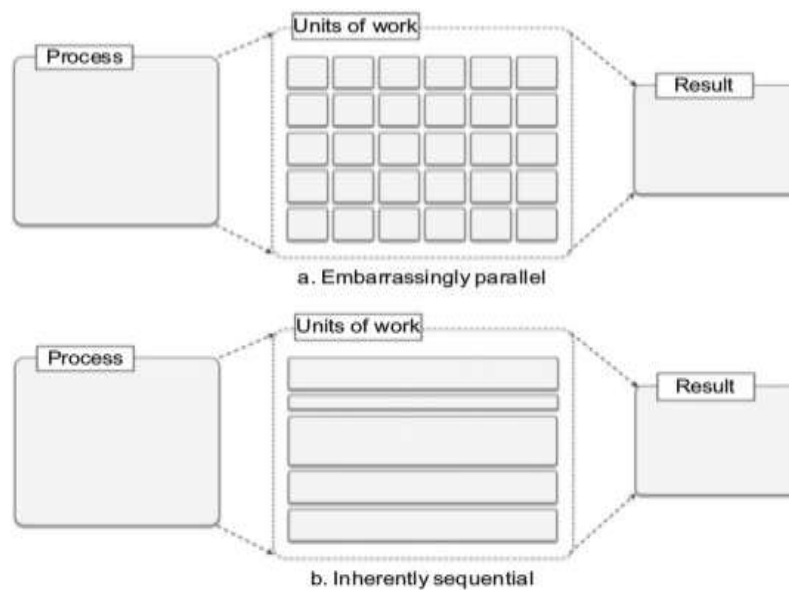



FIGURE 6.3

Domain decomposition techniques.

Figure 6.4 provides an overview of how a matrix product can be performed.

**FIGURE 6.4**

A matrix product.

The problem is embarrassingly parallel, and we can logically organize the multithreaded program in the following steps:

- Define a function that performs the computation of the single element of the resulting matrix by implementing the previous equation.
- Create a double for loop (the first index iterates over the rows of the first matrix and the second over the columns of the second matrix) that spawns a thread to compute the elements of the resulting matrix.
- Join all the threads for completion, and compose the resulting matrix.

The .NET framework provides the `System.Threading.Thread` class that can be configured with a function pointer, also known as a delegate, to execute asynchronously. This class will also define the method for performing the actual computation. **Listing 6.1** shows the class **ScalarProduct**.

```

/// Class ScalarProduct. Computes the scalar product between the row and the column arrays.
public class ScalarProduct
{
    /// Scalar product.
    private double result;
    /// Gets the resulting scalar product.
    public double Result{ get { returnthis.result; } }
    /// Arrays containing the elements of the row and the column to multiply.
    private double[] row, column;
    /// Creates instance of the ScalarProduct class and configures it with the given row & column arrays.
    /// <param name="row">Array with the elements of the row to be multiplied.</param>
    /// <param name="column">Array with the elements of the column to be multiplied.</param>
    public ScalarProduct(double[] row, double[] column)
    {
        this.row = row;
        this.column = column;
    }
    /// Executes the scalar product between the row and the column.
    /// <param name="row">Array with the elements of the row to be multiplied.</param>
    /// <param name="column">Array with the elements of the column to be multiplied.</param>
    public void Multiply()
    {
        this.result = 0;
        for(int i=0; i<this.row.Length; i++)
        {
            this.result += this.row[i] * this.column[i];
        }
    }
}

```

```

    }
}
}

```

LISTING 6.1 ScalarProduct Class.

To control the threads, we need to keep track of them so that we can query their status and obtain the result once they have completed the computation. We can create a simple program that reads the matrices, keeps track of all the threads in an appropriate data structure, and, once the threads have been completed, composes the final result. **Listing 6.2** shows the content of the **MatrixProduct** class using System;

```

using System.Threading;
using System.Collections.Generic;
/// Class MatrixProduct. Performs the matrix product of two matrices.
public class MatrixProduct
{
    /// First and second matrix of the product.
    private static double[,] a, b;
    /// Result matrix.
    private static double[,] c;
    /// Dictionary mapping thread instances to corresponding ScalarProduct instances that are run inside.
    private static IDictionary<Thread, ScalarProduct> workers;
    /// Read the command line parameters and perform the scalar product.
    /// <param name="args">Array strings containing the command line parameters.</param>
    public static void Main(string[] args)
    {
        // reads the input matrices a and b.
        MatrixProduct.ReadMatrices();
        // executes the parallel matrix product.
        MatrixProduct.ExecuteProduct();
        // waits for all the threads to complete and composes the final matrix.
        MatrixProduct.ComposeResult();
    }
    /// Executes parallel matrix product by decomposing problem in independent scalar product between
    /// rows and columns.
    private static void ExecuteThreads()
    {
        MatrixProduct.workers = newList<Thread>();
        int rows = MatrixProduct.a.Length;
        int columns = MatrixProduct.b[0].Length;
        for(int i=0; i<rows; i++)
        {
            for(int j=0; j<columns; j++)
            {
                double[] row = MatrixProduct.a[i];
                // because matrices are stored as arrays of arrays in order to get the columns
                // we need to traverse the array and copy data to another array.
                double[] column = new double[common];
                for(int k=0; k<common; k++)
                {
                    column[j] = MatrixProduct.b[j][i];
                }
                // creates a ScalarProduct instance with the previous rows and
                // columns and starts a thread executing the Multiply method.
            }
        }
    }
}

```

```

        ScalarProduct scalar = newScalarProduct(row, column);
        Thread worker = newThread(newThreadStart(scalar.Multiply));
        worker.Name =string.Format("{0}.{1}",row,column);
        worker.Start();
        // adds the thread to the dictionary so that it can be further retrieved.
        MatrixProduct.workers.Add(worker, scalar);
    }
}
/// Waits for the completion of all the threads and composes the final result matrix.
private static void ComposeResult()
{
    MatrixProduct.c = new double[rows,columns];
    foreach(KeyValuePair<Thread,ScalarProduct>pair in MatrixProduct.workers)
    {
        Thread worker = pair.Key;
        // we have saved the coordinates of each scalar product in the name
        // of the thread now we get them back by parsing the name .
        string[] indices = string.Split(worker.Name, new char[] { '.' });
        int i = int.Parse(indices[0]);
        int j = int.Parse(indices[1]);
        // we wait for the thread to complete
        worker.Join();
        // we set the result computed at the given coordinates.
        MatrixProduct.c[i,j] = pair.Value.Result;
    }
    MatrixProduct.PrintMatrix(MatrixProduct.c);
}
/// Reads the matrices.
private static void ReadMatrices()
{
    // code for reading the matrices a and b
}
/// Prints the given matrix.
///<param name="matrix">Matrix to print.</param>
private static void PrintMatrices(double[,] matrix)
{
    // code for printing the matrix.
}
}
}

```

LISTING 6.2 MatrixProduct Class (Main Program).

2 Functional decomposition

Functional decomposition is the process of identifying functionally distinct but independent computations. The focus here is on the type of computation rather than on the data manipulated by the computation. This kind of decomposition is less common and does not lead to the creation of a large number of threads, since the different computations that are performed by a single program are limited. Functional decomposition leads to a natural decomposition of the problem in separate units of work. **Figure 6.5** provides a pictorial view of how decomposition operates and allows parallelization.

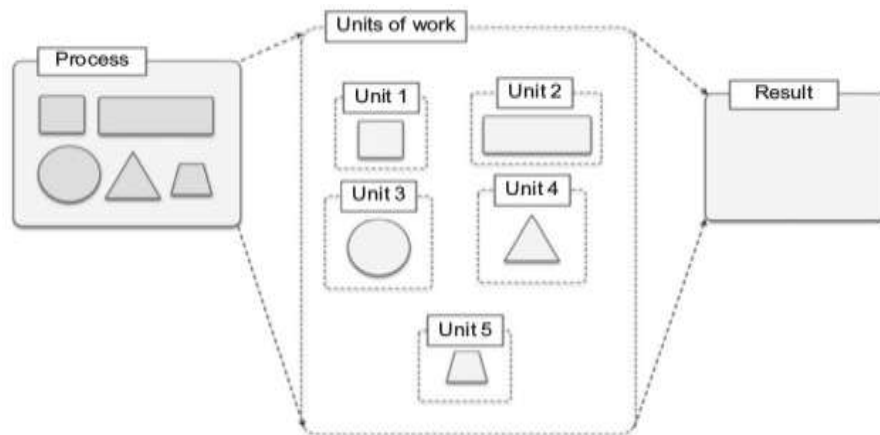


FIGURE 6.5

Functional decomposition.

The problems that are subject to functional decomposition can also require a composition phase in which the outcomes of each of the independent units of work are composed together.

In the following, we show a very simple example of how a mathematical problem can be parallelized using functional decomposition. Suppose, for example, that we need to calculate the value of the following function for a given value of x :

$$f(x) = \sin(x) + \cos(x) + \tan(x)$$

Once the value of x has been set, the three different operations can be performed independently of each other. This is an example of functional decomposition because the entire problem can be separated into three distinct operations. A possible implementation of a parallel version of computation is shown in **Listing 6.3**.

```
using System;
using System.Threading;
using System.Collections.Generic;
/// Delegate UpdateResult. Function pointer that is used to update the final result from the slave threads once
/// the computation is completed.
/// <param name="x">partial value to add.</param>
public delegate void UpdateResult(double x);
/// Class Sine. Computes the sine of a given value.
public class Sine
{
    /// Input value for which the sine function is computed.
    private double x;
    /// Gets the input value of the sine function.
    public double X { get { return this.x; } }
    /// Result value.
    private double y;
    /// Gets the result value of the sine function.
    public double Y { get { return this.y; } }
    /// Function pointer used to update the result.
    private UpdateResult updater;
    /// Creates an instance of the Sine and sets the input to the given angle.
    /// <param name="x">Angle in radians.</param>
    /// <param name="updater">Function pointer used to update the result.</param>
    public Sine(double x, UpdateResult updater)
    {
        this.x = x;
        this.updater = updater;
    }
}
```



```

    }
    /// Executes the sine function.
    public void Apply()
    {
        this.y = Math.Sin(this.x);
        if (this.updater != null)
        {
            this.updater(this.y);
        }
    }
}
/// Class Cosine. Computes the cosine of a given value.
public class Cosine
{
    /// Input value for which the cosine function is computed.
    private double x;
    /// Gets the input value of the cosine function.
    public double X { get { return this.x; } }
    /// Result value.
    private double y;
    /// Gets the result value of the cosine function.
    public double Y { get { return this.y; } }
    /// Function pointer used to update the result.
    private UpdateResultupdater;
    /// Creates an instance of the Cosine and sets the input to the given angle.
    /// <param name="x">Angle in radians.</param>
    /// <param name="updater">Function pointer used to update the result.</param>
    public Cosine(double x, UpdateResultupdater)
    {
        this.x = x;
        this.updater = updater;
    }
    /// Executes the cosine function.
    public void Apply()
    {
        this.y = Math.Cos(this.x);
        if (this.updater != null)
        {
            this.updater(this.y);
        }
    }
}
/// Class Tangent. Computes the tangent of a given value.
public class Tangent
{
    /// Input value for which the tangent function is computed.
    private double x;
    /// Gets the input value of the tangent function.
    public double X { get { return this.x; } }
    /// Result value.
    private double y;
    /// Gets the result value of the tangent function.
    public double Y { get { return this.y; } }
}

```

```

    /// Function pointer used to update the result.
    private UpdateResult updater;
    /// Creates an instance of the Tangent and sets the input to the given angle.
    /// <param name="x">Angle in radians.</param>
    /// <param name="updater">Function pointer used to update the result.</param>
    public Tangent(double x, UpdateResult updater)
    {
        this.x = x;
        this.updater = updater;
    }
    /// Executes the cosine function.
    public void Apply()
    {
        this.y = Math.Tan(this.x);
        if (this.updater != null)
        {
            this.updater(this.y);
        }
    }
}
/// Class Program. Computes the function  $\sin(x) + \cos(x) + \tan(x)$ .
public class Program
{
    /// Variable storing the computed value for the function.
    private static double result;
    /// Synchronization instance used to avoid keeping track of the threads.
    private static object synchRoot = new object();
    /// Read the command line parameters and perform the scalar product.
    /// <param name="args">Array strings containing the command line parameters.</param>
    public static void Main(string[] args)
    {
        // gets a value for x
        double x = 3.4d;
        // creates the function pointer to the update method.
        UpdateResult updater = new UpdateResult(Program.Sum);
        // creates the sine thread.
        Sine sine = new Sine(x, updater);
        Thread tSine = new Thread(new ThreadStart(sine.Apply));
        // creates the cosine thread.
        Cosine cosine = new Cosine(x, updater);
        Thread tCosine = new Thread(new ThreadStart(cosine.Apply));
        // creates the tangent thread.
        Tangent tangent = new Tangent(x, updater);
        Thread tTangent = new Thread(new ThreadStart(tangent.Apply));
        // shuffles the execution order.
        tTangent.Start();
        tSine.Start();
        tCosine.Start();
        // waits for the completion of the threads.
        tCosine.Join();
        tTangent.Join();
        tSine.Join();
        // the result is available, dumps it to console.
    }
}

```

```

        Console.WriteLine("f({0}): {1}", x, Program.result);
    }
    /// Callback that is executed once computation in thread is completed and adds partial value passed
    /// as a parameter to result.
    /// <param name="partial">Partial value to add.</param>
    private static void Sum(double partial)
    {
        lock(Program.synchRoot)
        {
            Program.result += partial;
        }
    }
}

```

LISTING 6.3 Mathematical Function.

3 Computation vs. Communication

It is very important to carefully evaluate the communication patterns among the components that have been identified during problem decomposition. The two decomposition methods presented in this section and the corresponding sample applications are based on the assumption that the computations are independent. This means that:

- The input values required by one computation do not depend on the output values generated by another computation.
- The different units of work generated as a result of the decomposition do not need to interact (i.e., exchange data) with each other.

These two assumptions strongly simplify the implementation and allow achieving a high degree of parallelism and a high throughput.

6.3 Multithreading with Aneka

As applications become increasingly complex, there is greater demand for computational power that can be delivered by a single multicore machine.

Often this demand cannot be addressed with the computing capacity of a single machine. It is then necessary to leverage distributed infrastructures such as clouds.

Decomposition techniques can be applied to partition a given application into several units of work, submitted for execution by leveraging clouds.

Aneka, as middleware for managing clusters, grids, and clouds, provides developers with advanced capabilities for implementing distributed applications. In particular, it takes traditional thread programming a step further. It lets you write multithreaded applications the traditional way, with the added twist that each of these threads can now be executed outside the parent process and on a separate machine. In reality, these “threads” are independent processes executing on different nodes and do not share memory or other resources, but they allow you to write applications using the same thread constructs for concurrency and synchronization as with traditional threads. Aneka threads, as they are called, let you easily port existing multithreaded compute-intensive applications to distributed versions that can run faster by utilizing multiple machines simultaneously, with minimum conversion effort.

6.3.1 Introducing the thread programming model

Aneka offers the capability of implementing multithreaded applications over the cloud by means of the **Thread Programming Model**.

This model introduces the abstraction of distributed thread, also called **Aneka thread**, which mimics the behavior of local threads but executes over a distributed infrastructure.

This model provides the best advantage in the case of **embarrassingly parallel applications**.

The Thread Programming Model exhibits APIs that mimic the ones exposed by .NET base class libraries for threading. In this way developers do not have to completely rewrite applications in order to leverage Aneka by replacing the **System.Threading.Thread** class and introducing the **AnekaApplication** class.

There are three major elements that constitute the object model of applications based on the Thread Programming Model:

Application. This class represents the interface to the Aneka middleware and constitutes a local view of a distributed application. In the Thread Programming Model the single units of work are created by the programmer. Therefore, the specific class used will be `Aneka.Entity.AnekaApplication<T,M>`, with T and M properly selected.

Threads. Threads represent the main abstractions of the model and constitute the building blocks of the distributed application. Aneka provides the `Aneka.Threading.AnekaThread` class, which represents a distributed thread.

Thread Manager. This is an internal component that is used to keep track of the execution of distributed threads and provide feedback to the application. Aneka provides a specific version of the manager for this model, which is implemented in the `Aneka.Threading.ThreadManager` class.

6.3.2 Aneka thread vs. common threads

To efficiently run on a distributed infrastructure, Aneka threads have certain limitations compared to local threads.

These limitations relate to the communication and synchronization strategies.

1 Interface compatibility

2 Thread life cycle

3 Thread synchronization

4 Thread priorities

5 Type serialization

1 Interface compatibility

The `Aneka.Threading.AnekaThread` class exposes almost the same interface as the `System.Threading.Thread` class with the exception of a few operations that are not supported. Table 6.1 compares the operations that are exposed by the two classes.

.Net Threading API	Aneka Threading API
<code>System.Threading</code>	<code>Aneka.Threading</code>
<code>Thread</code>	<code>AnekaThread</code>
<code>Thread.ManagedThreadId (int)</code>	<code>AnekaThread.Id (string)</code>
<code>Thread.Name</code>	<code>AnekaThread.Name</code>
<code>Thread.ThreadState (ThreadState)</code>	<code>AnekaThread.State</code>
<code>Thread.IsAlive</code>	<code>AnekaThread.IsAlive</code>
<code>Thread.IsRunning</code>	<code>AnekaThread.IsRunning</code>
<code>Thread.IsBackground</code>	<code>AnekaThread.IsBackground[false]</code>
<code>Thread.Priority</code>	<code>AnekaThread.Priority[ThreadPriority.Normal]</code>
<code>Thread.IsThreadPoolThread</code>	<code>AnekaThread.IsThreadPoolThread [false]</code>
<code>Thread.Start</code>	<code>AnekaThread.Start</code>
<code>Thread.Abort</code>	<code>AnekaThread.Abort</code>
<code>Thread.Sleep</code>	[Not provided]
<code>Thread.Interrupt</code>	[Not provided]
<code>Thread.Suspend</code>	[Not provided]
<code>Thread.Resume</code>	[Not provided]
<code>Thread.Join</code>	<code>AnekaThread.Join</code>

The basic control operations for local threads such as Start and Abort have a direct mapping, whereas operations that involve the temporary interruption of the thread execution have not been supported.

The reasons for such a design decision are twofold. **First**, the use of the Suspend/Resume operations is generally a deprecated practice, even for local threads, since Suspend abruptly interrupts the execution state of the thread. **Second**, thread suspension in a distributed environment leads to an ineffective use of the infrastructure, where resources are shared among different tenants and applications.

Sleep operation is not supported. Therefore, there is no need to support the Interrupt operation, which forcibly resumes the thread from a waiting or a sleeping state.

2 Thread life cycle

Aneka Thread life cycle is different from the life cycle of local threads. It is not possible to directly map the state values of a local thread to Aneka threads. **Figure 6.6** provides a comparative view of the two life cycles. The white balloons in the figure indicate states that do not have a corresponding mapping on the other life cycle; the shaded balloons indicate the common states.

In local threads most of the state transitions are controlled by the developer, who actually triggers the state transition by invoking methods. Whereas in Aneka threads, many of the state transitions are controlled by the middleware.

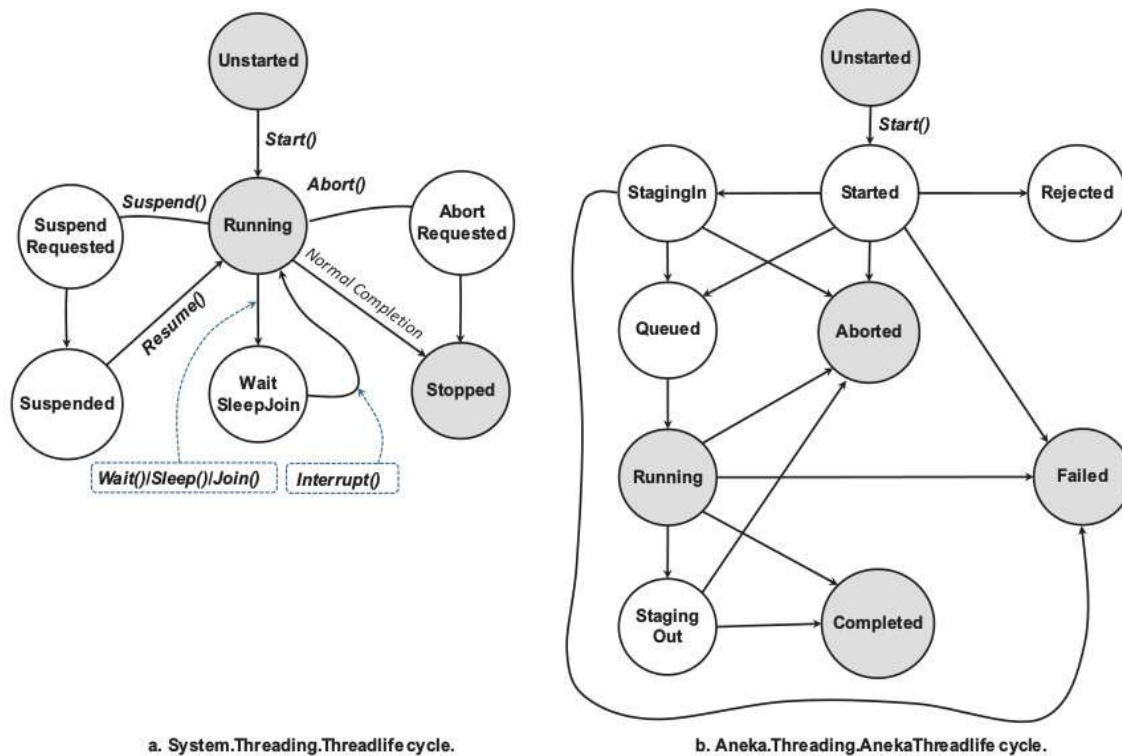


FIGURE 6.6

Thread life-cycle comparison.

Aneka threads support file staging and they are scheduled by the middleware, which can queue them for a considerable amount of time.

An Aneka thread is initially found in the Unstarted state. Once the Start() method is called, the thread transits to the Started state, from which it is possible to move to the StagingIn state if there are files to upload for its execution or directly to the Queued state. If there is any error while uploading files, the thread fails and it ends its execution with the Failed state, which can also be reached for any exception that occurred while invoking Start().

Another outcome might be the Rejected state that occurs if the thread is started with an invalid reservation token. This is a final state and implies execution failure due to lack of rights.

Once the thread is in the queue, if there is a free node where to execute it, the middleware moves all the object data and depending files to the remote node and starts its execution, thus changing the state into Running.

If the thread generates an exception or does not produce the expected output files, the execution is considered failed and the final state of the thread is set to Failed. If the execution is successful, the final state is set to Completed. If there are output files to retrieve, the thread state is set to StagingOut while files are collected and sent to their final destination.

At any point, if the developer stops the execution of the application or directly calls the Abort() method, the thread is aborted and its final state is set to Aborted.

3 Thread synchronization

The .NET base class libraries provide advanced facilities to support thread synchronization by the means of monitors, semaphores, reader-writer locks, and basic synchronization constructs at the language level. Aneka provides minimal support for thread synchronization that is limited to the implementation of the join operation for thread abstraction.

This requirement is less stringent in a distributed environment, where there is no shared memory among the thread instances and therefore it is not necessary.

Providing coordination facilities that introduce a locking strategy in such an environment might lead to distributed deadlocks that are hard to detect. Therefore, by design Aneka threads do not feature any synchronization facility except join operation.

4 Thread priorities

The System.Threading.Thread class supports thread priorities, where the scheduling priority can be one selected from one of the values of the ThreadPriority enumeration: Highest, AboveNormal, Normal, BelowNormal, or Lowest.

Aneka does not support thread priorities, the Aneka.Threading.Thread class exhibits a Priority property whose type is ThreadPriority, but its value is always set to Normal, and changes to it do not produce any effect on thread scheduling by the Aneka middleware.

5 Type serialization

Serialization is the process of converting an object into a stream of bytes to store the object or transmit it to memory, a database, or a file. Main purpose is to save the state of an object to recreate it when needed. The reverse process is called deserialization.

Local threads execute all within the same address space and share memory; therefore, they do not need objects to be copied or transferred into a different address space. Aneka threads are distributed and execute on remote computing nodes, and this implies that the object code related to the method to be executed within a thread needs to be transferred over the network.

A .NET type is considered serializable if it is possible to convert an instance of the type into a binary array containing all the information required to revert it to its original form or into a possibly different execution context. This property is generally given for several types defined in the .NET framework by simply tagging the class definition with the Serializable attribute.

Aneka threads execute methods defined in serializable types, since it is necessary to move the enclosing instance to remote execution method. In most cases, providing serialization is as easy as tagging the class definition with the Serializable attribute; in other cases it might be necessary to implement the ISerializable interface and provide appropriate constructors for the type.

6.4 Programming applications with Aneka threads

To show how it is possible to quickly port multithreaded application to Aneka threads, we provide a distributed implementation of the previously discussed examples for local threads.

6.4.1 Aneka threads application model

The Thread Programming Model is a programming model in which the programmer creates the units of work as Aneka threads. Therefore, it is necessary to utilize the `AnekaApplicationN<W,M>` class, which is the application reference class for all the programming models.

The Aneka APIs support different programming models through template specialization. Hence, to develop distributed applications with Aneka threads, it is necessary to specialize the template type as follows:

AnekaApplication<AnekaThread, ThreadManager>

These two types are defined in the `Aneka.Threading` namespace noted in the Aneka. **Threading.dll** library of the Aneka SDK.

Another important component of the application model is the **Configuration** class, which is defined in the `Aneka.Entity` namespace (**Aneka.dll**). This class contains a set of properties that allow the application class to configure its interaction with the middleware, such as the address of the Aneka index service, which constitutes the main entry point of Aneka Clouds.

Listing 6.4 demonstrates how to create a simple application instance and configure it to connect to an Aneka Cloud whose index service is local.

```
// namespaces containing types of common use
using System;
using System.Collections.Generic;
// common Aneka namespaces.
using Aneka;
using Aneka.Util;
using Aneka.Entity;
// Aneka Thread Programming Model user classes
using Aneka.Threading;
///Creates an instance of the Aneka Application configured to use Thread Programming Model.
/// <returns>Application instance.</returns>
private AnekaApplication<AnekaThread,ThreadManager> CreateApplication();
{
    Configuration conf =new Configuration();
    // this is the common address and port of a local installation of Aneka Cloud.
    conf.SchedulerUri = newUri("tcp://localhost:9090/Aneka");
    conf.Credentials =newUserCredentials("Administrator", string.Empty);
    // we will not need support for file transfer, hence we optimize the
    // application in order to not require any file transfer service.
    conf.UseFileTransfer = false;
    // we do not need any other configuration setting
    // we create the application instance and configure it.
    AnekaApplication<AnekaThread,ThreadManager> app =
    new AnekaApplication<AnekaThread,ThreadManager>(conf);
    return app;
}
```

LISTING 6.4 Application Creation and Configuration.

Listing 6.5 provides a very simple example of how to create Aneka threads.

```
//..... continues from the previous listing
///Thread worker method
private void WorkerMethod()
{
// .....
}
```

```

///Creates a collection of threads that are executed in the context of given application.
/// <param name="app">>Application instance.</param>
private void CreateThreads(AnekaApplication<AnekaThread,ThreadManager> app);
{
    // creates a delegate to the method to execute inside the threads.
    ThreadStart worker = new ThreadStart(this.WorkerMethod);
    // iterates over a loop and creates ten threads.
    for(int i=0; i<10; i++)
    {
        AnekaThread thread = new AnekaThread(worker, app);
        thread.Start();
    }
}

```

LISTING 6.5 Thread Creation and Execution.

6.4.2 Domain decomposition: matrix multiplication

To port to Aneka threads the multithreaded matrix multiplication, we need to apply the considerations made in the previous section. Hence, we start reviewing the code by first making the proper changes to the ScalarProduct class. **Listing 6.6** shows the modified version of ScalarProduct.

The class has been tagged with the Serializable attribute and extended with the methods required to implement custom serialization. Supporting custom serialization implies the following:

1. Including the System.Runtime.Serialization namespace.
2. Implementing the ISerializable interface. This interface has only one method that is void GetObjectData(SerializationInfo, StreamingContext), and it is called when the runtime needs to serialize the instance.
3. Providing a constructor with the following signature: ScalarProduct(SerializationInfo, StreamingContext). This constructor is invoked when the instance is deserialized.

The SerializationInfo class has a central role of providing a repository where all the properties defining the serialized format of a class can be stored and referenced by name.

```

using System.Runtime.Serialization;
/// Class ScalarProduct. Computes the scalar product between the row and the column arrays.
/// The class uses custom serialization. In order to do so it implements ISerializable interface.
[Serializable]
public class ScalarProduct : ISerializable
{
    /// Scalar product.
    private double result;
    /// Gets the resulting scalar product.
    public double Result { get { return this.result; } }
    /// Arrays containing the elements of the row and the column to multiply.
    private double[] row, column;
    /// Creates an instance of ScalarProduct class and configures it with given row and column arrays.
    /// <param name="row">>Array with the elements of the row to be multiplied.</param>
    /// <param name="column">>Array with the elements of the column to be multiplied.
    /// </param>
    public ScalarProduct(double[] row, double[] column)
    {
        this.row = row;
        this.colum = colum;
    }
    /// Deserialization constructor used by .NET runtime to recreate instances of types implementing

```



```

/// custom serialization.
/// <param name="info">Bag containing the serialized data instance.</param>
/// <param name="context">Serialization context (not used).</param>
public ScalarProduct(SerializationInfo info, StreamingContext context)
{
    this.result = info.GetDouble("result");
    this.row = info.GetValue("row", typeof(double[])) as double[];
    this.column = info.GetValue("column", typeof(double[])) as double[];
}
/// Executes the scalar product between the row and the colum.
/// <param name="row">Array with the elements of the row to be multiplied.</param>
/// <param name="column">Array with the elements of the column to be multiplied.</param>
public void Multiply()
{
    this.result = 0;
    for(int i=0; i<this.row.Length; i++)
    {
        this.result += this.row[i] * this.column[i];
    }
}
/// Serialization method used by .NET runtime to serialize instances of types implementing custom
/// serialization.
/// <param name="info">Bag containing the serialized data instance.</param>
/// <param name="context">Serialization context (not used).</param>
public ScalarProduct(SerializationInfo info, StreamingContext context)
{
    this.result = info.GetDouble("result");
    this.row = info.GetValue("row", typeof(double[])) as double[];
    this.column = info.GetValue("column", typeof(double[])) as double[];
}
/// Executes the scalar product between the row and the colum.
/// <param name="row">Array with the elements of the row to be multiplied.</param>
/// <param name="column">Array with the elements of the column to be multiplied.</param>
public void Multiply()
{
    this.result = 0;
    for(int i=0; i<this.row.Length; i++)
    {
        this.result += this.row[i] * this.column[i];
    }
}
/// Serialization method used by .NET runtime to serialize instances of types implementing custom
/// serialization.
/// <param name="info">Bag containing the serialized data instance.</param>
/// <param name="context">Serialization context (not used).</param>
public void GetObjectData(SerializationInfo info, StreamingContext context)
{
    info.AddValue("result",this.result);
    info.AddValue("row", this.row,typeof(double[]));
    info.AddValue("column", this.column,typeof(double[]));
}
}

```

LISTING 6.6 ScalarProduct Class (Modified Version).

The second step is to change the MatrixProduct class to leverage Aneka threads. We need to first create a properly configured application and then substitute the occurrences of the **System.Threading.Thread** class with **Aneka.Threading.Thread** (see **Listing 6.7**).

```
using System;
// we do not anymore need the reference to the threading namespace. using System.Threading;
using System.Collections.Generic;
// reference to the Aneka namespaces of interest. common Aneka namespaces.
using Aneka;
using Aneka.Util;
using Aneka.Entity;
// Aneka Thread Programming Model user classes
using Aneka.Threading;
/// Class MatrixProduct. Performs the matrix product of two matrices.
public class MatrixProduct
{
    /// First and second matrix of the produt.
    private static double[,]a, b;
    /// Result matrix.
    private static double[,] c;
    /// Dictionary mapping the thread instances to the corresponding ScalarProduct instances that are run
    /// inside.The occurrence of the Thread class has been substituted with AnekaThread.
    private static IDictionary<AnekaThread, ScalarProduct> workers.
    /// Reference to the distributed application the threads belong to.
    private static AnekaApplication<AnekaThread, ThreadManager> app;
    /// Read the command line parameters and perform the scalar product.
    /// <param name="args">Array strings containing the command line parameters.</param>
    public static void Main(string[] args)
    {
        try
        {
            // activates the logging facility.
            Logger.Start();
            // creates the Aneka application instance.
            MatrixProduct.app =Program.CreateApplication();
            // reads the input matrices a and b.
            MatrixProduct.ReadMatrices();
            // executes the parallel matrix product.
            MatrixProduct.ExecuteProudct();
            // waits for all the threads to complete and composes the final matrix.
            MatrixProduct.ComposeResult();
        }
        catch(Exception ex)
        {
            IOUtil.DumpErrorReport(ex, "Matrix Multiplication – Error executing " + "the
            application");
        }
        finally
        {
            try
            {
                // checks whether the application instance has been created stops it.
                if (MatrixProduct.app != null)
```

```

        {
            MatrixProduct.app.Stop();
        }
    }
    catch(Exception ex)
    {
        IOUtil.DumpErrorReport(ex, "Matrix Multiplication – Error stopping " + "the application");
    }
    // stops the logging thread.
    Logger.Stop();
}
/// Executes parallel matrix product by decomposing problem in independent scalar product between rows
/// and columns.

private static void ExecuteThreads()
{
    // we replace the Thread class with AnekaThread.
    MatrixProduct.workers = new Dictionary<AnekaThread, ScalarProduct>();
    int rows = MatrixProduct.a.Length;
    // in .NET matrices are arrays of arrays & number of columns is represented by length of 2nd array.
    int columns = MatrixProduct.b[0].Length;
    for(int i=0; i<rows; i++)
    for(int j=0; j<columns; j++)
    {
        double[] row = MatrixProduct.a[i];
        // because matrices are stored as arrays of arrays in order to
        // to get the columns we need to traverse the array and copy the
        // the data to another array.
        double[] column = new double[common];
        for(int k=0; k<common; k++)
        {
            column[j] = MatrixProduct.b[j][i];
        }
        // creates a ScalarProduct instance with the previous rows and
        // columns and starts a thread executing the Multiply method.
        ScalarProduct scalar = new ScalarProduct(row, column);
        // we change the System.Threading.Thread class with the corresponding
        // Aneka.Threading.AnekaThread class and reference the application instance.
        AnekaThread worker = new AnekaThread(new ThreadStart(scalar.Multiply), app);
        worker.Name = string.Format("{0}.{1}", row, column);
        worker.Start();
        // adds the thread to the dictionary so that it can be
        // further retrieved.
        MatrixProduct.workers.Add(worker, scalar);
    }
}
/// Waits for the completion of all the threads and composes the final result matrix.
private static void ComposeResult()
{
    MatrixProduct.c = new double[rows,columns];
    // we replace the Thread class with AnekaThread.
    foreach(KeyValuePair<AnekaThread,ScalarProduct>pair in MatrixProduct.workers)
    {

```

```

        AnekaThread worker = pair.Key;
        // we have saved the coordinates of each scalar product in the name
        // of the thread now we get them back by parsing the name.
        string[] indices = string.Split(worker.Name, new char[] {','});
        int i = int.Parse(indices[0]);
        int j = int.Parse(indices[1]);
        // we wait for the thread to complete
        worker.Join();
        // instead of using the local value of the ScalarProduct instance
        // we use the one that has is stored in the Target property.
        // MatrixProduct.c[i,j] = pair.Value.Result;
        MatrixProduct.c[i,j] = ((ScalarProduct) worker.Target).Result;
    }
    MatrixProduct.PrintMatrix(MatrixProduct.c);
}
/// Reads the matrices.
private static void ReadMatrices()
{
    // code for reading the matrices a and b
}
/// Prints the given matrix.
/// <param name="matrix">Matrix to print.</param>
private static void PrintMatrices(double[,] matrix)
{
    // code for printing the matrix.
}
/// Creates an instance of the Aneka Application configured to use the
/// Thread Programming Model.
/// <returns>Application instance.</returns>
private AnekaApplication<AnekaThread,ThreadManager> CreateApplication();
{
    Configuration conf =new Configuration();
    // this is the common address and port of a local installation of the Aneka Cloud.
    conf.SchedulerUri = newUri("tcp://localhost:9090/Aneka");
    conf.Credentials =newUserCredentials("Administrator", string.Empty);
    // we will not need support for file transfer, hence we optimize the application in order to not require
    // any file transfer service.
    conf.UseFileTransfer = false;
    // we do not need any other configuration setting, we create the application instance and configure it.
    AnekaApplication<AnekaThread,ThreadManager> app =
        new AnekaApplication<AnekaThread,ThreadManager>(conf);
    return app;
}
}
}

```

LISTING 6.7 MatrixProduct Class (Modified Version).

6.4.3 Functional decomposition: Sine, Cosine, and Tangent

There is only one significant difference in this case: Each of the threads has a reference to a delegate that is used to update the global sum at the end of the computation. Since we are operating in a distributed environment, the instance on which the object will operate is not shared among the threads, but each thread instance has its own local copy.

This example also illustrates how to modify the classes Sine, Cosine, and Tangent so that they can leverage the default serialization capabilities of the framework (see **Listing 6.8**).

```

using System;
using System.Collections.Generic;
// reference to the Aneka namespaces of interest. common Aneka namespaces.
using Aneka;
using Aneka.Util;
using Aneka.Entity;
// Aneka Thread Programming Model user classes
using Aneka.Threading;
/// <param name="x">partial value to add.</param>
public delegate void UpdateResult(double x);
/// Class Sine. Computes the sine of a given value.
[Serializable]
public class Sine
{
    /// Input value for which the sine function is computed.
    private double x;
    /// Gets or sets the input value of the sine function.
    public double X { get { return this.x; } set { this.x = value; } }
    /// Result value.
    private double y;
    /// Gets or sets the result value of the sine function.
    public double Y { get { return this.y; } set { this.y = value; } }
    //we need a default constructor, which is automatically provided by compiler if we do not specify
    any constructor.
    /// Creates an instance of the Sine and sets the input to the given angle.
    /// <param name="x">Angle in radiants.</param>
    /// <param name="updater">Function pointer used to update the result.</param>
    public void Apply()
    {
        this.y = Math.Sin(this.x);
        // we cannot use this anymore because there is no shared memory space.
    }
}
/// Class Cosine. Computes the cosine of a given value. The same changes have been
/// applied by removing the code not needed anymore rather than commenting it out.
[Serializable]
public class Cosine
{
    /// Input value for which the cosine function is computed.
    private double x;
    /// Gets or sets the input value of the cosine function.
    public double X { get { return this.x; } set { this.x = value; } }
    /// Result value.
    private double y;
    /// Gets or sets the result value of the cosine function.
    public double Y { get { return this.y; } set { this.y = value; } }
    /// Executes the cosine function.
    public void Apply()
    {
        this.y = Math.Cos(this.x);
    }
}
/// Class Cosine. Computes the cosine of a given value. The same changes have been

```

```

/// applied by removing the code not needed anymore rather than commenting it out.
[Serializable]
public class Tangent
{
    /// Input value for which the tangent function is computed.
    private double x;
    /// Gets or sets the input value of the tangent function.
    public double X { get { return this.x; } set { this.x = value; } }
    /// Result value.
    private double y;
    /// Gets or sets the result value of the tangent function.
    public double Y { get { return this.y; } set { this.y = value; } }
    /// Executes the tangent function.
    public void Apply()
    {
        this.y = Math.Cos(this.x);
    }
}
/// Class Cosine. Computes the cosine of a given value. The same changes have been
/// applied by removing the code not needed anymore rather than commenting it out.
[Serializable]
public class Tangent
{
    /// Input value for which the tangent function is computed.
    private double x;
    /// Gets or sets the input value of the tangent function.
    public double X { get { return this.x; } set { this.x = value; } }
    /// Result value.
    private double y;
    /// Gets or sets the result value of the tangent function.
    public double Y { get { return this.y; } set { this.y = value; } }
    /// Executes the tangent function.
    public void Apply()
    {
        this.y = Math.Tan(this.x);
    }
}
/// Class Program. Computes the function  $\sin(x) + \cos(x) + \tan(x)$ .
public class Program
{
    /// Variable storing the coputed value for the function.
    private static double result;
    ///we do not need synchronization anymore, because update of global sum is done sequentially.
    /// Synchronization instance used to avoid keeping track of the threads.
    private static object synchRoot = new object();
    /// Reference to the distributed application the threads belong to .
    private static AnekaApplication <AnekaThread, ThreadManager> app;
    /// Read the command line parameters and perform the scalar product.
    /// <param name="args">Array strings containing the command line parameters. </param>
    public static void Main(string[], args)
    {
        try
        {

```

```

// activates the logging facility.
Logger.Start();
// creates the Aneka application instance.
app = Program.CreateApplication();
// gets a value for x
double x = 3.4d;
// creates the function pointer to the update method.
UpdateResult updater = new UpdateResult(Program.Sum);
// creates the sine thread.
Sine sine = new Sine(x, updater);
AnekaThread tSine = new AnekaThread(new ThreadStart(sine.Apply), app);
// creates the cosine thread.
Cosine cosine = new Cosine(x, updater);
AnekaThread tCosine = new AnekaThread(new ThreadStart(cosine.Apply), app);
// creates the tangent thread.
Tangent tangent = new Tangent(x, updater);
AnekaThread tTangent = new AnekaThread(new ThreadStart(tangent.Apply), app);
// shuffles the execution order.
tTangent.Start();
tSine.Start();
tCosine.Start();
// waits for the completion of the threads.
tCosine.Join();
tTangent.Join();
tSine.Join();
// once we have joined all the threads the values have been collected back and we use
//Target property in order to obtain the object with updated values.
sine = (Sine) tSine.Target;
cosine = (Cosine) tSine.Target;
tangent = (Tangent) tSine.Target;
Program.result = sine.Target.Y + cosine.Y + tangent.Y;
// the result is available, dumps it to console.
Console.WriteLine("f({0}): {1}", x, Program.result);
}
catch(Exception ex)
{
    IOUtil.DumpErrorReport(ex,"Math Functions – Error executing " + "application" );
}
finally
{
    try
    {
        // checks whether the application instance has been created stops it.
        if (app != null)
        {
            app.Stop();
        }
    }
    catch(Exception ex)
    {
        IOUtil.DumpErrorReport(ex,"Math Functions – Error stopping "+ application" );
    }
}
// stops the logging thread .

```

```

        Logger.Stop();
    }
}
private static void Sum(double partial)
{
    lock(Program.synchRoot)
    {
        Program.result += partial;
    }
}
/// Creates an instance of the Aneka Application configured to use the
/// Thread Programming Model.
/// <returns>Application instance. </returns>
private AnekaApplication <AnekaThread,ThreadManager> CreateApplication();
{
    Configuration conf =new Configuration();
    // this is the common address and port of a local installation of the Aneka Cloud.
    conf.SchedulerUri = newUri("tcp://localhost:9090/Aneka" );
    conf.Credentials = newUserCredentials("Administrator", string.Empty);
    // we will not need support for file transfer, hence we optimize the
    // application in order to not require any file transfer service.
    conf.UseFileTransfer = false;
    // we do not need any other configuration setting
    // we create the application instance and configure it.
    AnekaApplication <AnekaThread,ThreadManager> app =
    new AnekaApplication <AnekaThread,ThreadManager>(conf);
    return app;
}
}

```

LISTING 6.8 Mathematical Function (Modified Version).