

Developing Thread Model Applications

Christian Vecchiola and Xingchen Chu

Abstract

This tutorial describes the Aneka Thread Programming Model and explains how to create distributed applications based on it. It illustrates some examples provided with the Aneka distribution that are built on top the Thread Model. It provides a detailed step by step guide for users on how to create an application by using the Microsoft Visual Studio 2005 Development Environment. After having read this tutorial, the users will be able to develop their own application on top of the Aneka Thread Model.

Document Status

Creation Date:	12/01/08
Version:	0.1
Classification:	User
Authors:	Christian Vecchiola, Xingchen Chu
Last Revision Date:	09/17/09
Status:	Draft

1. Prerequisites

In order to fully understand this tutorial the user should be familiar with the general concepts of Grid and Cloud Computing, Object Oriented programming and generics, distributed systems, and a good understanding of the .NET framework 2.0 and C#.

The practical part of the tutorial requires a working installation of Aneka. It is also suggested to have Microsoft Visual Studio 2005 (any edition) with C# package installed¹ even if not strictly required.

¹ Any default installation of Visual Studio 2005 and Visual Studio 2005 Express comes with all the components required to complete this tutorial installed except of Aneka, which has to be downloaded and installed separately.

2. Introduction

Aneka allows different kind of applications to be executed on the same Grid/Cloud infrastructure. In order to support such flexibility it provides different abstractions through which it is possible to implement distributed applications. These abstractions map to different execution models. Currently Aneka supports three different execution models:

- *Task Execution Model*
- *Thread Execution Model*
- *MapReduce Execution Model*

Each execution model is composed by three different elements: the *WorkUnit*, the *Scheduler*, the *Executor*, and the *Manager*. The *WorkUnit* defines the granularity of the model; in other words, it defines the smallest computational unit that is directly handled by the Aneka infrastructure. Within Aneka, a collection of related work units define an application. The *Scheduler* is responsible for organizing the execution of work units composing the applications, dispatching them to different nodes, getting back the results, and providing them to the end user. The *Executor* is responsible for actually executing one or more work units, while the *Manager* is the client component which interacts with the Aneka system to start an application and collect the results. A view of the system is given in the figure below.

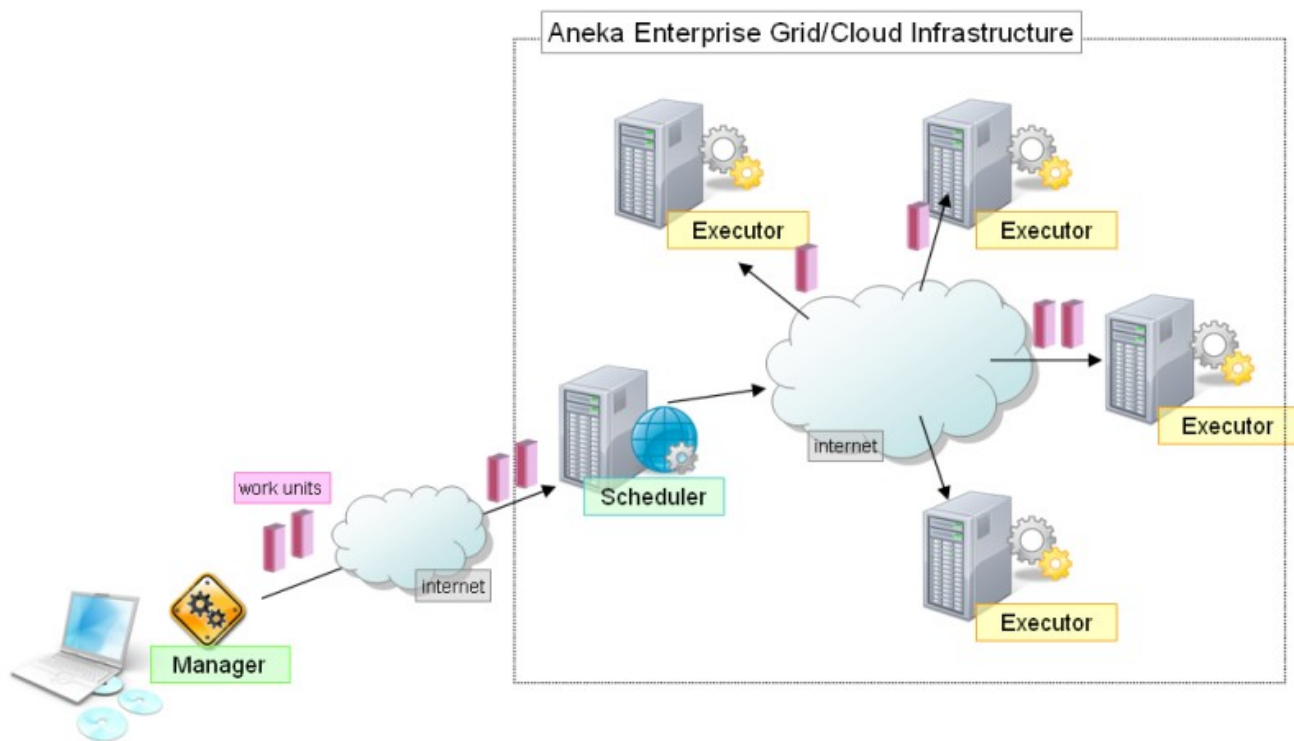


Figure 1. System Components View.

Hence, for the Thread Model there will be a specific *WorkUnit* called *AnekaThread*, a *Thread Scheduler*, a *Thread Executor*, and a *Thread Manager*. In order to develop an application for Aneka the user does not have to know all these components; Aneka handles a lot of the work by itself without the user's contribution. Only few things the users are required to know:

- how to define *AnekaThread* instances specific to the application that is being defined;
- how to create a *AnekaApplication* and starts the execution of threads;
- how to control the *AnekaApplication* and collect the results.

In the remainder of this tutorial will then concentrate on the Thread Model even if many of the concepts described can be applied to other execution models.

3. Thread Model

3.1 Local vs Remote Threads

The modern operating systems provide the abstractions of *Process* and *Thread* for defining the runtime profile of a software application. A *Process* is a software infrastructure that is used by the operating system to control the execution of an application. A *Process* generally contains one or more threads. A *Thread* is a sequence of instructions that can be executed in parallel with other instructions. When an application is running the operating system takes care of alternating their execution on the local machine. It is responsibility of the developer to create a consistent computation as a result of thread execution.

The Thread Model uses the same abstraction of for defining a sequence of instructions that can be remotely executed in parallel with other instructions. Hence, within the Thread Model an application is a collection of remotely executable *threads*. The Thread Model allows developers to virtualize the execution of a local multi-threaded application (developed with the .NET threading APIs) in an almost complete transparent manner. This model represents the right solution when developers want to port the execution of a .NET multi-threaded application on Aneka and still use the same way of controlling the execution of application flow, which is based on synchronization between threads.

Developers that are familiar with multi-threaded applications will find the Thread Model the most natural path to program distributed applications with Aneka. The transition between a .NET thread and an Aneka thread is almost transparent. In the following a sample application will be used to discuss how to use Aneka threads.

3.2 Working with Threads

Within the .NET threading model a thread is represented by the *Thread* sealed class that can be configured with the method to execute through the *ThreadStart* class. The user activates a thread by calling the *Start* method on it and by using the APIs exposed by the *Thread* class it can:

- Check the status of the thread by using the *Thread.State* and the *Thread.IsAlive* properties.
- Control its execution by stopping it (*Thread.Abort*).
- Suspending and resuming their execution (*Thread.Suspend* and *Thread.Resume*).
- Wait for its termination by calling *Thread.Join*.

The *Thread.Join*, *Thread.Suspend*, and *Thread.Resume* are the operations that allow developer to create basic synchronization patterns between threads. The *Thread* class provides additional APIs that cover:

- Thread affinity.
- Volatile read and write.
- Critical region management.
- Asynchronous operations.
- Stack management.

More complex and advanced synchronization pattern can be obtained by using other classes of the .NET threading APIs that does not have any reference to the *Thread* class.

In order to remotely execute a thread, the Thread Model provides a counterpart of the the *Thread* class: *AnekaThread*. The *AnekaThread* class represents the work unit in the Thread Model and exposes a subset of the APIs of *System.Threading.Thread*. It is possible to perform almost all the basic operations described before and the following table identifies the mappings between the two worlds.

Table 1. Local vs Remote Thread

.NET Threading API	Aneka Threading API
System.Threading	Aneka.Threading
Thread	AnekaThread
Thread.ManagedThreadId (int)	AnekaThread.Id (string, from WorkUnit)
Thread.Name	AnekaThread.Name (from WorkUnit)
Thread.ThreadState (ThreadState)	AnekaThread.State (WorkUnitState)
Thread.IsAlive	AnekaThread.IsAlive
Thread.IsRunning	AnekaThread.IsRunning (from

	WorkUnit)
Thread.IsBackground	[Not provided]
Thread.Priority	[Not provided]
Thread.IsThreadPoolThread	[Not provided]
Thread.Start	AnekaThread.Start
Thread.Abort	AnekaThread.Abort
Thread.Sleep	[Not provided]
Thread.Interrupt	[Not provided]
Thread.Suspend	[Not provided]
Thread.Resume	[Not provided]
....	[Not provided]

The *AnekaThread* class implements the basic *Thread* operations but does not give any support for the advanced operations such as: critical region, stack, apartment, culture, and execution context management. Moreover, some basic operations have not been supported, these are:

- Thread Priority Management.
- *Suspend, Resume, Sleep, and Interrupt.*

The reason why these operations have not been supported is because the *AnekaThread* instances are remotely executed on a computation node that generally executes work units coming from different distributed applications. It is not possible to keep the resources of a computation node occupied with a *AnekaThread* instance that is sleeping, or suspended. For what concerns the priority *Aneka* does not provide any facility.

```
namespace Aneka.Threading
{
    /// <summary>
    /// Class AnekaThread. Represents the basic unit of work
    /// in the Thread Model. A AnekaThread instance is configured with a
    /// specific method to execute and its remote execution is
    /// started.
    /// </summary>
    public class AnekaThread : WorkUnit
    {
        /// <summary>
        /// Gets a boolean value indicating whether the AnekaThread instance
        /// instance is alive (not: Unstarted | Completed | Aborted |
        /// Failed)
        /// </summary>
        public bool IsAlive { get; }
    }
}
```

```

    /// <summary>
    /// Gets the reflection and instance information on the method that is
    /// executed on the remote computation node.
    /// </summary>
    public RemoteMethodInfo TargetMethodInfo { get; }
    /// <summary>
    /// Gets the instance representing the target of the method invocation.
    /// </summary>
    public object Target { get; }
    /// <summary>
    /// Creates an instance of the AnekaThread.
    /// </summary>
    /// <param name="start">thread start method information.</param>
    /// <param name="application">grid application.</param>
    public AnekaThread(ThreadStart start,
        AnekaApplication<AnekaThread, ThreadManager> application)
    { ... }
    /// <summary>
    /// Starts the execution of the AnekaThread.
    /// </summary>
    public void Start() { ... }
    /// <summary>
    /// Aborts the execution of the AnekaThread.
    /// </summary>
    public void Abort() { ... }
    /// <summary>
    /// Waits until the execution of AnekaThread is terminated.
    /// </summary>
    public void Join() { ... }
    /// <summary>
    /// Waits until the execution of AnekaThread is terminated.
    /// </summary>
    /// <param name="time">Maximum interval of time to wait.</param>
    public void Join(TimeSpan time) { ... }
}
}

```

Listing 1 - AnekaThread class.

Listing 1 presents the public interface of the *AnekaThread* class. Other than the declaration of the basic operations for controlling the execution of a *AnekaThread*, the class exposes two interesting properties that provide information about the method executed remotely (*TargetMethodInfo*) and target of the invocation (*Target*).

In order to create a *AnekaThread* instance it is necessary to pass to the constructor two parameters: a *ThreadStart* object and a reference to the *AnekaApplication* instance that the threads belongs to. Once the thread has been created its *State* property is set to *WorkUnitState.Unstarted* and it is

possible to access the information about the method that will be executed by the *TargetMethodInfo* property. This property extrapolates all the reflection information used to recreate the execution environment on the remote computation node. This property is of type *RemoteMethodInfo* whose interface is described in Listing 2.

```
namespace Aneka.Threading
{
    /// <summary>
    /// Class RemoteMethodInfo. Wraps all the required information for executing
    /// a method in a remote computation node.
    /// </summary>
    public sealed class RemoteMethodInfo
    {
        /// <summary>
        /// Gets the display name of the assembly containing the definition of
        /// the method to execute.
        /// </summary>
        public string AssemblyName { get; }
        /// <summary>
        /// Serialized information of the target of the method invocation.
        /// </summary>
        public byte[] ObjectInstance { get; }
        /// <summary>
        /// Flags used to invoke the method.
        /// </summary>
        public BindingFlags Flags { get; }
        /// <summary>
        /// Name of the method to invoke.
        /// </summary>
        public string TargetMethod { get; }
        ....
    }
}
```

Listing 2 - Class RemoteMethodInfo.

While the use of *TargetMethodInfo* is mostly internal, the *Target* property is of more interest for the user. This property exposes the updated value of the instance after the execution of the *AnekaThread* and provides a quick way for performing the mapping between the threads and the instances that are object of thread execution. When programming with the .NET threading API developers have to maintain this mapping in a specific data structure (list, hash-table, etc...) and it is their responsibility to keep it updated when the state of the thread changes. By using the Aneka threading APIs no additional code is required.

Static Methods

The .NET Framework allows the execution of static methods as entry points for Thread execution. While this feature makes perfectly sense in a local execution context, it becomes unclear in a distributed environment.

Threads running in the same application domain share a static context. As a result, the side effects of an execution can be captured into static variables and still be accessible to the user. Threads running in different application domains - and this is the case of *AnekaThread* instances - do not share a static context. This makes the use of static methods quite limited.

For this reason the current implementation of the Thread Model does not support the remote execution of static methods.

The second parameter required by the *AnekaThread* constructor is the reference to the *AnekaApplication* class that groups all the instances belonging to the same application. A property common to all the programming models supported by Aneka is the concept of *application*. This can be generally described as a collection of related jobs that constitute a distributed computation. All the programming models must provide a local view of the distributed application through a specific instance of the *AnekaApplication* class. *AnekaApplication* $\langle W, M \rangle$ is a generic class and needs to be specialized with the concrete types that are related to the programming model implemented. In the case of the Thread Model we have that:

- *W*, which must inherit from *WorkUnit*, is *AnekaThread*.
- *M*, which must implement *IApplicationManager*, is *ThreadManager*.

These two generic parameters represent respectively the basic unit of computation of the model and the specific client manager used to handle the interaction with Aneka for the given programming model.

The specific tasks of the *AnekaApplication* are the interaction with Aneka and providing aggregate information on the execution of all the work units that it owns. Whereas in other programming models (see the *Task Model*) the *AnekaApplication* class plays a more concrete role, in the case of the *Thread Model* its role is mostly confined to the setup of the distributed application. The .NET Threading APIs do not have a corresponding application object and the execution flow of the application is mainly controlled by directly operating on the thread instances by calling the methods exposed by the *System.threading.Thread* class. In the case of Aneka the same approach has been maintained and developers can directly operate on *AnekaThread* instances once they have been created and assigned to a *AnekaApplication*

instance.

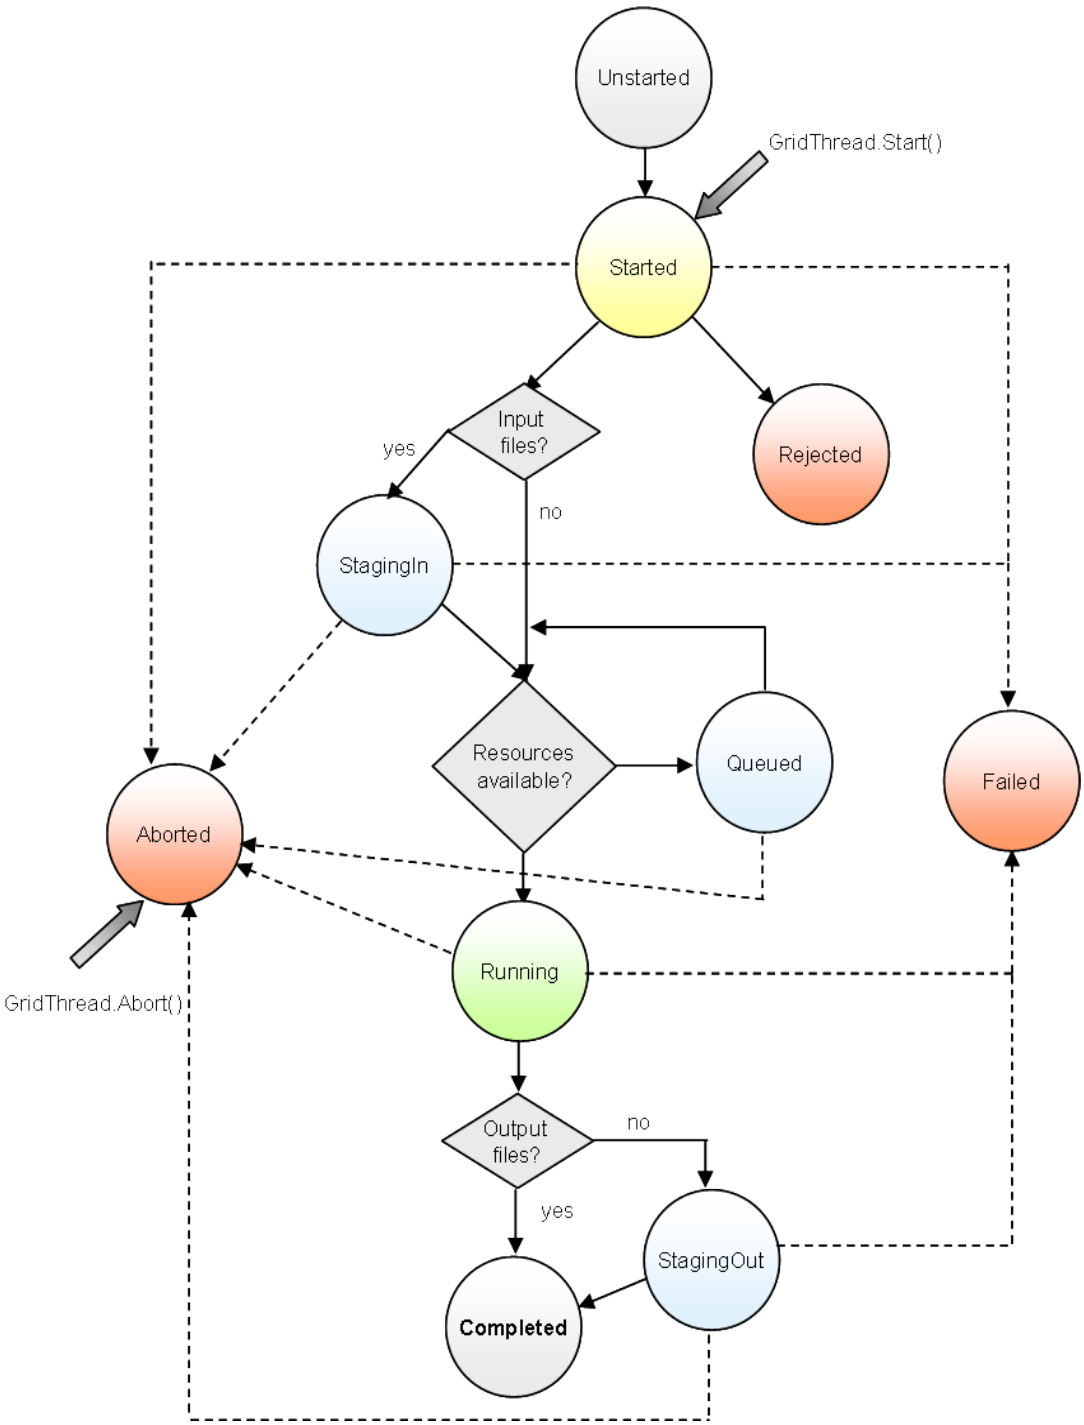


Figure 2. AnekaThread instance state transitions (client view).

The *AnekaThread* class provides all the required facilities to control its life cycle. Figure 2 depicts the life cycle of a *AnekaThread* instance. As soon as the instance is created it is in the *Unstarted* state. A call to *AnekaThread.Start()* makes it move into the *Started* state and causes the

submission of the instance to Aneka. If the *AnekaThread* has some dependent files to be transferred it moves to *StagingIn* state until all the dependent files are transferred. The *AnekaThread* can then move directly to the *Running* state if any computing node is available or being queued, thus moving into the *Queued* state. As soon as execution completes if there are any dependent output files to be downloaded to the client the state is changed to *StagingOut* otherwise it is directly set to *Completed*. At any stage an exception or an error can occur that causes the *AnekaThread* instance to move into the *Failed* state. The user can also actively terminate the execution by calling *AnekaThread.Abort()* and this causes the *AnekaThread* instance to be stopped and its state to be set to *Aborted*.

NOTE: The diagram also shows the Rejected state. This state is related to the negotiation protocol and it is actually not active. *AnekaThread* instances can be allocated to specific slots for their execution that can be reserved exclusively. When starting an *AnekaThread* instance it is possible to associate to it a reservation identifier that will map the instance to the reserved slots. If this reservation identifier is not valid or expired the state of the instance becomes Rejected and it is not allowed to run. The diagram does not show the ReScheduled state. This state is assumed when an *AnekaThread* instance is interrupted during execution and put in the scheduling queue again, This could happen if the execution slot in which the instance was running has been preempted by a reserved *WorkUnit* instance.

3.3 Additional Considerations

3.3.1 Serialization

Since *AnekaThread* instances are moved between different application domains they need to be serialized. The *AnekaThread* instance is declared *serializable*, but this does not guarantee that all *AnekaThread* instances created by the users will be serializable. In particular, since the *AnekaThread* is configured with a *ThreadStart* object referencing the instance that is the target of the method invocation, the type containing the method definition of the method need to be serializable too. The reason for this, is because the infrastructure will serialize the local instance on which the method will be invoked and send it to the remote node.

In case the users provides a method that is not defined in a serializable type, the *AnekaThread* constructor throws an *ArgumentException* alerting the user that the selected method cannot be used to run a *AnekaThread* instance. This prevents the user from creating a work unit that will not run.

3.3.2 Thread Programming Model vs Common APIs

As pointed out in section 3.1 the Thread Model allow developers to completely control the execution of the application by using the operations exposed by

the *AnekaThread* class. Once the *AnekaApplication* instance has been properly set up there is no need to maintain a reference to it. The rationale behind this choice is that developers familiar with the .NET Threading APIs do not have the explicit concept of application but simply coordinate the execution of threads.

Since the Thread Model relies on the common APIs of the infrastructure, it takes advantage of the services these APIs offer and these services can be used by developers too. In this case the *AnekaApplication* class plays an important role in controlling the execution flow, since it allows to:

- Monitor the state of *AnekaThread* instances by using events:
 - *AnekaApplication*<W,M>.WorkUnitFailed
 - *AnekaApplication*<W,M>.WorkUnitFinished
 - *GridApplicaiton*<W,M>.WorkUnitAborted
- Programmatically control the execution of *GirdThread* instances:
 - *AnekaApplication*<W,M>.ExecuteWorkUnit(W)
 - *AnekaApplication*<W,M>.StopWorkUnit(W)
- Terminate the execution of the application:
 - *AnekaApplication*<W,M>.ApplicationFinished
 - *AnekaApplication*<W,M>.StopExecution

These APIs are available to all the models and allows developers to perform the basic operations required to manage the distributed application in a model independent fashion. For what concerns the Thread Model this seems to be unnatural even though can be useful some times. This tutorial will not explore further this option and the reader is suggested to look for the Task Model that naturally uses this APIs.

It is important to notice that the result of using the *AnekaThread* operations or the *AnekaApplication* operations is the same. The reason for this is that both of the two classes relies on the *ThreadManager* class for performing the requested operations.

4. Example: Distributed Warhol Filter.

In this section we will show how to use the *Thread Model* and the *AnekaThread* APIs to create a distributed image filter that performs the *Warhol Effect*. By developing this simple application the user will be able to:

- Create a *AnekaApplication* instance configured for the Thread Model.
- Customize the execution of the a *AnekaApplication* with a configuration file.
- Create and customize *AnekaThread* instances with user specific code.

- Submit and control the execution of AnekaThread instances.

This tutorial is not an exhaustive guide to the APIs provided with the Thread Model but it is a good start for developing applications based on distributed threads with Aneka.

4.1 What is the Warhol Effect?

There is no clear definition of what the Warhol Effect is but the effect transforms a given picture into another that resembles in style the following painting of Marilyn Monroe made by the famous pop artist Andy Warhol (see Figure 3).



Figure 3. Marilyn Monroe prints (Andy Warhol).

Given the fact that the prints are made by a human without any specific algorithm it is quite difficult to automate the process and there are many attempts on the web that are trying to reproduce the same effect by means of a computer algorithm. Any filter available in the web produces a result whose similarity in style with Andy Warhol's paintings varies.

In order to produce a simple computer algorithm that perform the filter we will apply the following restrictions:

- The output image produced by the filter is has a color depth of four colors (many of the Warhol's paintings are basically made by using four colors).
- The colors of the original image are remapped and clustered into the new palette made of three colors according to their brightness.
- The output image will provide an image that is two times the size of the original one and is composed by organizing into a square four different samples of the same image filtered with different palettes.

The outcome of this simple example will be a console application that given an input image will produce a second image as described above. This application will use the Thread Model and Aneka for distributing the execution

of the steps required to perform the filter.

NOTE: in the state transition diagram it also appear a *Rejected* state. This state is related to the reservation infrastructure that is not active at the moment. In simple terms, WorkUnit instances can be executed with a reservation; this means that a specific slot in the system has been reserved for their execution that is identified by an id. The *Rejected* state come into play when a WorkUnit instance provides to the system a reservation identifier that is not valid or expired. The diagram does not show the *ReScheduled* state that appears when a WorkUnit, while running, is terminated by the infrastructure and put into the scheduling queue for being executed again. This could happen because the execution slot assigned to the instance is expired and cannot be extended.

4.2 Application Structure

The source code of this application can be found into the Aneka installation directory under the *Examples/Tutorials/ThreadDemo* directory. There is a convenient Visual Studio 2005 Project that simplifies the build process, but that is not essential for completing the tutorial.

In order to implement the application we will organize the whole application into the three main classes:

- *Aneka.Examples.ThreadDemo.WarholFilter* (see *WarholFilter.cs*): this class performs the filter of the image and given a picture produces another picture that is the same size of the original and remaps its colors.
- *Aneka.Examples.ThreadDemo.WarholApplication* (see *WarholApplication.cs*): this class is responsible for:
 - setting up the *AnekaApplication* instance;
 - configuring it with for the Thread Model;
 - creating the *AnekaThread* instances and starting their execution;
 - waiting for the completion of the *AnekaThread* instances and assembling the four images produced into a single image.
- *Aneka.Examples.ThreadDemo.WarholDriver* (see *Program.cs*): this class constitutes the main entry point of the application and is in charge of parsing the command line parameters, creating and configuring the *WarholApplication* class, and starting its execution. In case the command line parameters are not correct the class displays a simple help that explains to ther user how to launch the application.

A summary view of the operations exposed by the three classes can be seen in Figure 4 (next page).

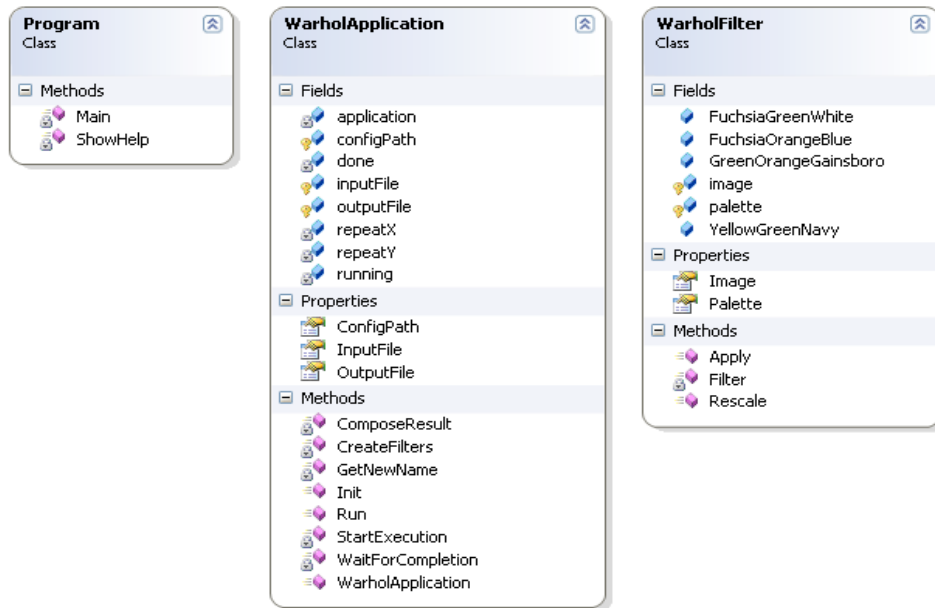


Figure 4. Class diagram.

These three classes are compiled into an executable (wahrolizer.exe) representing the console application that perform the filter. The application can be started by the command line as follows:

wahrolizer input_file [output_file] [conf_file]

where:

input_file: path to the image taken as a input for the filter [mandatory]

output_file: path to the file where the filtered image will be saved to [optional]

conf_file: path to the configuration file for connecting to Aneka [optional]

The only mandatory parameter is *input_file*. If the user does not provide any save path for the output image the application will automatically create a file named *[input-file-name].wahrol.[input-file-ext]* where *[input-file-name]* and *[input-file-ext]* are respectively the name and the extension of the input file. For example:

marilyn.jpg => marilyn.wharol.jpg

If the output file already exists the application will overwrite the file without asking the user permission. For what concerns the settings used to connect to Aneka the user can specify them into an xml file. A sample xml configuration file is already provided (see conf.xml) and it contains the default values that

are used when the user does not specify a configuration file. This file whose content is displayed in Figure 4 can be used a starting point for exploring the configuration settings of Aneka and creates customs configurations.

In the next sections we will describe the implementation of the filter, the main steps carried out by the *WarholApplication* to perform the distributed filtering.

```
<?xml version="1.0" encoding="utf-8" ?>
<Aneka>
  <UseFileTransfer value="false" />
  <Workspace value="." />
  <SingleSubmission value="true" />
  <ResubmitMode value="AUTO" />
  <PollingTime value="1000" />
  <LogMessages value="true" />
  <SchedulerUri value="tcp://localhost:9090/Aneka" />
</Aneka>
```

Figure 5. Aneka Configuration File (conf.xml).

4.3 WarholFilter: Filter Implementation

The *WarholFilter* class implements the basic operation of remapping the colors of an image into a three color palette that can be specified by the user.

```
// File: WarholFilter.cs
using System;
using System.Collections.Generic;
using System.Text;
using System.Drawing;

namespace Aneka.Examples.ThreadDemo
{
    /// <summary>
    /// Class WarholFilter. Performs the color remapping and clustering of an input
    /// image according to a specified palette. This class exposes an Image property
    /// that is set by the user to the input image to filter and that points to the
    /// processed image after the Apply method has been called.
    /// </summary>
    [Serializable]
    public class WarholFilter
    {
        /// <summary>
        /// Sample palette made by Yellow, DarkGreen, and Navy Color constants.
        /// </summary>
        private static readonly Color[] YellowGreenNavy = new Color[3] { ... };
        /// <summary>
        /// Sample palette made by Fuchsia, Orange, and DarkBlue Color constants.
        /// </summary>
    }
}
```

```

private static readonly Color[] FuchsiaOrangeBlue = new Color[3] { ... };
/// <summary>
/// Sample palette made by Green, Orange, and Gainsboro Color constants.
/// </summary>
private static readonly Color[] GreenOrangeGainsboro = new Color[3] { ... };
/// <summary>
/// Sample palette made by Fuchsia, DarkOliveGreen, and WhiteSmoke Color
/// constants.
/// </summary>
private static readonly Color[] FuchsiaGreenWhite = new Color[3] { ... };
/// <summary>
/// Input/Output bitmap.
/// </summary>
protected Bitmap image;
/// <summary>
/// Gets, sets the input image on which the filter is applied. This
/// property stores the filtered bitmap after the Apply() method is
/// called.
/// </summary>
public Bitmap Image
{ get { return this.image; } set { this.image = value; } }
/// <summary>
/// Target color palette.
/// </summary>
protected Color[] palette;
/// <summary>
/// Gets, sets the palette of colors that will be used to remap the image.
/// </summary>
public Color[] Palette
{ get { return this.palette; } set { this.palette = value; } }

/// <summary>
/// Applies the filter.
/// </summary>
public void Apply()
{
    if (this.image == null)
    {
        throw new ArgumentNullException("Image is null!", "image");
    }
    if (this.palette == null)
    {
        throw new ArgumentNullException("Palette is null!", "palette");
    }
    this.image = this.Filter(this.image, this.palette);
}
/// <summary>
/// Remaps the color values of the source image to the color values contained
/// in the given palette by clustering them according to their brightness.

```



```

    /// </summary>
    /// <param name="image">source image</param>
    /// <param name="palette">target palette</param>
    /// <returns>filtered bitmap</returns>
    protected Bitmap Filter(Bitmap source, Color[] palette)
    {
        // Step 1. reorder the palette according to the color brightness
        // Step 2. identify the minimum (min) and the maximum (max) brightness
        //          values for the source image and creates (max - min) / length
        //          clusters where the length is the size of the palette.
        // Step 3. invoke Rescale and compute the minimum brightness threshold
        //          color values for each cluster.
        // Step 4. for each pixel of the image evaluates the brightness and find
        //          find the cluster into which the color will be mapped. Set the
        //          color corresponding to that cluster in the output image.
    }
    /// <summary>
    /// Creates an array of threshold values by recursively dividing the range
    /// identified by max - min and putting the values computed into the given
    /// array.
    /// </summary>
    /// <param name="delta">length of the subarray in values that will be filled
    /// during the call of the method. </param>
    /// <param name="start">index of the first element of the subarray</param>
    /// <param name="midPoint">threshold value for the middle element</param>
    /// <param name="min">minimum threshold value</param>
    /// <param name="max">maximum threshold value</param>
    /// <param name="values">target array</param>
    protected void Rescale(int delta, int start,
                          float midPoint, float min, float max, float[] values)
    { ... }
}
}

```

Listing 3 - Class WarholFilter.

Listing 3 provides a summary view of the class. As it can be noticed by the included namespaces there is nothing in this class that relates to any Aneka library. WarholFilter simply defines the operation that is carried out by the AnekaThread instance when it is executed on the remote computation node. The relevant members of this class are the following:

- *WarholFilter.Image*: this property is used to store either the input or the output bitmap of the filter. More precisely when an instance of WarholFilter is created this property is set to the image that will be filtered. After the apply method is called, this property references the filtered image.
- *WarholFilter.Palette*: this property references the palette that will be used to remap the color values of the image. The WarholFilter class

exposes four ready to use palettes that contain combination of colors that are used in Andy Warhol's paintings:

- *WarholFilter.YellowGreenNavy*
- *WarholFilter.FuchsiaOrangeBlue*
- *WarholFilter.FuchsiaGreenWhite*
- *WarholFilter.GreenOrangeGainsboro*

These palettes will be used by *WarholApplication* to create the four images.

- *WarholFilter.Apply()*: this method performs some argument checking on the values of the two properties and invokes *WarholFilter.Filter* to generate the processed image whose color values are rescaled to the values in *Palette*.

It is not the purpose of this tutorial to dig into the details of the implementation of the filter (a complete explanation of the filter is given into the comments in the code) but it is worth estimating its complexity.

4.3.1 Complexity Analysis

The computational complexity of the filter is located into the *WarholFilter.Filter* method and it depends on two parameters: the size of the palette and the image dimension ($O(L_{pal}, N_{px})$ where L_{pal} is the size of the palette and N_{px} the number of pixel of the image). The steps involved into the filter as pointed out into the listing are the following:

- Step 1: sorting is made by using bubble sort ($O(L_{pal}^2)$)
- Step 2: the entire image is scanned once ($O(N_{px})$)
- Step 3: the threshold computation is $O(L_{pal})$
- Step 4: the entire image is scanned once and for each pixel the palette array is scanned ($O < O(L_{pal} \times N_{px})$)

The complexity of the filter has then an upper bound of $K_1(O(L_{pal}^2) + O(L_{pal}) + O(N_{px}))$. Since $L_{pal} \ll N_{px}$ it is possible to identify an upper bound with $K_2 O(N_{px})$. This upper bound is a rough estimation of the time required to perform one single filter. The application requires four different samples of the filter. Hence we can easily parallelize this step and reduce the overall computation time.

4.4 WarholApplication: Distributed Filtering Coordination

The *WarholFilter* application is responsible of coordinating the distributed execution of the filters and composing their results together. It uses the services of the Thread Model for parallelizing the computation and the services of *WarholFilter* class for performing the single filtered components of the image. In this section we will describe in detail its behavior.

```

// File: WarholApplication.cs
#region Namespaces
using System;
using System.Collections.Generic;    // IList<...> class.
using System.Text;                 // StringBuilder class.
using System.IO;                   // IOException (IO Errors management)
using System.Drawing;              // Image and Bitmap classes.

using Aneka.Entity;                // Aneka Common APIs for all models
using Aneka.Threading;             // Aneka Thread Model
using System.Threading;            // ThreadStart (AnekaThread initialization)
#endregion

namespace Aneka.Examples.ThreadDemo
{
    /// <summary>
    /// Class WarholApplication. This class uses the Thread Model for performing
    /// the distributed filter on a given image and creating the Warhol Effect. It
    /// is responsible of managing the interaction with Aneka and controlling the
    /// execution of the distributed application by managing the AnekaThread
    /// instances required to perform the filter.
    /// </summary>
    public class WarholApplication
    {
        #region Properties
        /// <summary>Path to the input file.</summary>
        protected string inputPath;
        /// <summary>Gets or sets the path to the input file.</summary>
        public string InputPath
        {
            get { return this.inputPath; }
            set
            {
                if ((value == null) || (value == string.Empty))
                {
                    throw new ArgumentException("InputPath is null or empty",
                                                "InputPath");
                }
                this.inputPath = value;
            }
        }
        /// <summary>Path to the output file.</summary>
        protected string outputPath;
        /// <summary>Gets or sets the path to the output file.</summary>
        public string OutputPath
        {
            get { return this.outputPath; }
            set { this.outputPath = value; }
        }
    }
}

```

```

}
/// <summary>Path to the configuration file.</summary>
protected string configPath;
/// <summary>Gets or sets the path to the configuration file.</summary>
public string ConfigPath
{
    get { return this.configPath; }
    set { this.configPath = value; }
}
#endregion

#region Implementation Fields
/// <summary>
/// Application instance that interfaces the client with Aneka.
/// </summary>
protected AnekaApplication<AnekaThread, ThreadManager> application;
/// <summary>
/// List containing the currently running AnekaThread instances.
/// </summary>
protected IList<AnekaThread> running;
/// <summary>
/// List containing the WarholFilter instances that have completed the
/// execution.
/// </summary>
protected IList<WarholFilter> done;
/// <summary>
/// Number of columns that will compose the final image.
/// </summary>
protected int repeatX;
/// <summary>
/// Number of rows that will compose the final image.
/// </summary>
protected int repeatY;
#endregion

#region Public Methods
/// <summary>Creates an instance of WarholApplication.</summary>
public WarholApplication() {}
/// <summary>Applies the filter.</summary>
public void Run() { ... }
#endregion

#region Helper Methods
/// <summary>
/// Reads the configuration and initializes the AnekaApplication instance.
/// </summary>
protected void Init() { ... }
/// <summary>
/// Starts the execution of the distributed application by creating the

```

```

    /// filters wrapping them into AnekaThread instances and starting their
    /// execution.
    /// </summary>
    /// <param name="source">Source bitmap.</param>
    protected void StartExecution(Bitmap source) { ... }
    /// <summary>
    /// Waits for the completion of all the threads and if some thread has
    /// failed its execution restarts it.
    /// </summary>
    protected void WaitForCompletion() { ... }
    /// <summary>
    /// Collects the processed images from each filter and composes them
    /// into a single image.
    /// </summary>
    /// <param name="source">Source bitmap.</param>
    protected void ComposeResult(Bitmap source) { ... }
    /// <summary>
    /// Creates an array of WarholFilter instances each of them configured
    /// with the same input image and a different palette.
    /// </summary>
    /// <param name="source">Source bitmap.</param>
    /// <returns>Array of filters.</returns>
    protected WarholFilter[] CreateFilters(Bitmap source) { ... }
    /// <summary>
    /// Creates a file name by adding the given suffix to the given file name.
    /// </summary>
    /// <param name="name">Source file name.</param>
    /// <param name="suffix">String suffix to append to the name.</param>
    /// <returns>A string containing the new file name.</returns>
    protected string GetNewName(string name, string suffix) { ... }
    #endregion
}
}

```

Listing 4 - Class WarholApplication.

Listing 4 shows the content of the *WarholApplication.cs* file. This file contains the definition of the *WarholApplication* class. We can inspect the code by first dividing the content of the file in five sections delimited by the **#region ... #endregion** preprocessor directives:

- **Namespaces:** this region includes all the required namespaces for developing the application. Among the namespaces included three are of a particular interest:
 - *Aneka.Entity (Aneka.dll):* This namespace contains the definition of the base APIs common to all the models supported by Aneka. Inside this namespace we can find the *AnekaApplication* class, the *Configuration* class, the *WorkUnit* class, and all those types that are

part of the core object model. This namespace needs always to be included while developing applications for Aneka.

- *Aneka.Threading* (*Aneka.Threading.dll*): This namespace contains the definition of the APIs for the Thread Model. In particular it contains the definition of the *AnekaThread* and the *ThreadManager* class. This namespace needs to be included when developing applications based on the Thread Model.

NOTE: The whole set of APIs of a model is generally organized into three namespaces: *Aneka.[model]*, *Aneka.[model].Scheduling*, and *Aneka.[model].Execution*. These namespaces correspond to three different assemblies named with the same convention. When developing applications with a specific model it is generally required to include only the first namespace, since the other two namespaces contain the definition of server side components.

- *System.Threading* (*System.dll*): This namespace contains .NET Threading APIs. When developing applications with the Thread Model it is necessary to initialize *AnekaThread* instances with a *ThreadStart* object whose definition is contained in the *System.Threading* namespace.

The other namespaces that have been included into the file contains the definition of support classes that have been used to program the *WarholApplication* class. For example, since the class deals with images it is necessary to include the *System.Drawing* (*System.Drawing.dll*) namespace.

- Properties: this region contains the three properties that constitute all the parameters required to perform the filter. The path to the input image (*WarholApplication.InputPath*), the path where to save the output image (*WarholApplication.OutputPath*), and to the path configuration file (*WarholApplication.ConfigPath*). The only mandatory parameter is the path to the input image that must point to an existing file. The other two parameters are configured with the default values if set to null.
- Implementation Fields: this region contains the declaration of the protected members that are used to manage the application. There are five fields declared:
 - *AnekaApplication*<*AnekaThread*, *ThreadManager*> *application*;
 - *IList*<*AnekaThread*> *running*;
 - *IList*<*WarholFilter*> *done*;
 - *int repeatX*;
 - *int repeatY*;

The first three fields constitute infrastructure that is generally required while developing applications based on the Thread Model. As for any

other model it is necessary to create a *AnekaApplication* instance that represents the local view of the distributed application. Since *AnekaApplication* is a generic type it needs to be specialized by using the components that identify the model we use. In this case we will use a *AnekaApplication<AnekaThread, ThreadManager>* instance, that specializes the behaviour of *AnekaApplication* for the *Thread Model*. Moreover, two additional data structures are required: a list containing the *running AnekaThread* instances and *done* list containing the filter instances that have completed their execution. It can be noticed that *running* is a list of *AnekaThread* instances while *done* is a list of *WarholFilter* instances. The reason for this is because while the thread is running we need to keep a reference to it in order to join the thread and check its state. Once a thread has completed its execution there is no more need to store a reference to it but we simply keep the *WarholFilter* instance connected to the thread that is accessible by the *AnekaThread.Target* property.

Two more fields have been defined: *repeatX* and *repeatY*. These fields are specific to the application we are developing and maintain the information about the number of columns and rows that compose the final image.

- Public Methods: the public methods of the class are its default constructor and the *Run* method that starts the execution of the filter. The content of this method will be analyzed in detail in the next section.
- Helper Methods: in order to make the more understandable the code, the body of the *Run* method has been divided into logical steps that have been encapsulated into helper methods. In particular we can identify four major steps:
 - *WarholApplication.Init*: initializes the *AnekaApplication* instance with the selected *Configuration* instance.
 - *WarholApplication.StartExecution*: initializes the *running* and *done* lists, creates the *AnekaThread* instances, and starts their execution.
 - *WarholApplication.WaitForCompletion*: waits for the completion of all the *AnekaThread* instances in the *running* list and eventually restarts their execution if they failed. When this method returns, the *running* list is empty and the *done* list contains the reference to the *WarholFilter* instances created.
 - *WarholApplication.ComposeResult*: iterates on the *done* list and compose the final output image by arranging the filtered bitmaps into an image with *repeatX* columns and *repeatY* rows. This method saves the output bitmap into the *WarholApplication.OutputPath* if set, otherwise it generates a new name as described in Section 4.2.

This region contains two more methods that are invoked by the previous one that are in charge of creating the list of different filters that will be applied to the image (*WarholApplication.CreateFilters*) and of

generating a new name for the output file (*WarholApplication.GetNewName*).

In the following we will explore in more details the single steps of the application.

```
namespace Aneka.Examples.ThreadDemo
{
    ...
    public class WarholApplication
    {
        ....

        /// <summary>
        /// Application instance that interfaces the client with Aneka.
        /// </summary>
        protected AnekaApplication<AnekaThread, ThreadManager> application;
        ....

        /// <summary>
        /// Applies the filter.
        /// </summary>
        public void Run()
        {
            if (File.Exists(this.inputPath) == false)
            {
                throw new FileNotFoundException("InputPath does not exist.",
                                                "InputPath");
            }
            try
            {
                // Initializes the AnekaApplication instance.
                this.Init();

                // read the bitmap
                Bitmap source = new Bitmap(this.inputPath);

                // create one filter for each of the four slices that will
                // compose the final image and starts their execution on
                // Aneka by wrapping them into AnekaThread instances...
                this.StartExecution(source);

                // wait for all threads to complete...
                this.WaitForCompletion();

                // collect the processed images and compose them
                // into one single image.
                this.ComposeResult(source);
            }
        }
    }
}
```



```

    }
    finally
    {
        // we ensure that the application closes properly
        // before leaving the method...
        if (this.application != null)
        {
            if (this.application.Finished == false)
            {
                this.application.StopExecution();
            }
        }
    }
}
}
}
}
}
}
}

```

Listing 5 - *Run()* method.

Listing 5 reports the content of the *Run* method. As we can notice body of the method sequentially calls the logical steps identified before and ensures that the application instance is closed if some error occurs. In the finally block we simply check that application instance is not null and whether it has completed its execution by looking at the *AnekaApplication<W,M>.Finished* property. If the application is not finished, the *AnekaApplication<W,M>.StopExecution()* method is invoked to terminate its execution.

NOTE: The lines of code contained in the finally block identify a common programming pattern for all the model supported by Aneka. This pattern ensures the resources allocated by the *AnekaApplication* instance are properly released and no work unit is left running on Aneka.

We can now explore the single steps and see how to set up the *AnekaApplication* instance and configure it for its execution.

```

namespace Aneka.Examples.ThreadDemo
{
    ...
    public class WarholApplication
    {
        ....

        /// <summary>
        /// Application instance that interfaces the client with Aneka.
        /// </summary>

```

```

protected AnekaApplication<AnekaThread, ThreadManager> application;
....

/// <summary>
/// Reads the configuration and initializes the AnekaApplication instance.
/// </summary>
protected void Init()
{
    Configuration configuration = null;
    if (string.IsNullOrEmpty(this.configPath) == true)
    {
        configuration = Configuration.GetConfiguration();
    }
    else
    {
        configuration = Configuration.GetConfiguration(this.configPath);
    }

    // we set this force to false because
    // we want to handle the resubmission
    // of failed threads.
    configuration.SingleSubmission = false;

    // we initialize the AnekaApplication instance with the
    // selected configuration object and the components required
    // for the Thread Model
    this.application =
        new AnekaApplication<AnekaThread, ThreadManager>(configuration);
}
}

```

Listing 6 - Init() method.

Listing 6 reports the body of the *Init* method. The method performs to very basic steps:

- Configuration setup: if the configuration path has been the static method *Configuration.GetConfiguration(string)* is invoked to read to read the information form the given file. If the configuration path has not been set, the static method *Configuration.GetConfiguration()* is called. This method will first look for the default configuration file (in this case: warholizer.exe.config) and if not found it will create a default configuration object. The default values for the *Configuration* class can been seen in the *conf.xml* file (see Figure 5). Once the configuration instance has been created the value of *SingleSubmission* is set to false. The reason for this is because in case some threads are failed they will be restarted again.

NOTE: The Thread Model is naturally based on the coordinated execution of remotely executable threads that can be started at any time. The *SingleSubmission* property forces the behavior of the *AnekaApplication* to a single submission of all the *AnekaThread* instances that have been explicitly added to the *AnekaApplication* before calling the method *AnekaApplication.SubmitExecution*. Since with the Thread Model the *AnekaThread* instances are not explicitly added to the *AnekaApplication* setting *SingleSubmission* to true would lead to an unexpected behaviour and it could cause the premature termination of the application. This problem is even worse when we want to handle the resubmission of the failed threads. It is then a general recommendation to **not to use *SingleSubmission* set to true when using the Thread Model.**

- AnekaApplication initialization: this step is accomplished by simply initializing the *application* field and passing as parameter the Configuration object that has been obtained at the previous step. Since the *AnekaApplication* class is a generic type its initialization implies specifying the actual types used by the application. In the case of the Thread Model we will specialize the *AnekaApplication* class by using *AnekaThread* in place of the *WorkUnit* and *ThreadManager* in place of *IApplicationManager*.

The next step of the process is performed by the *WarholApplication.StartExecution* method that takes as input the Bitmap instance read from the input file and creates the *AnekaThread* instances as described in Listing 7.

```
namespace Aneka.Examples.ThreadDemo
{
    ...
    public class WarholApplication
    {
        ....
        /// <summary>
        /// Application instance that interfaces the client with Aneka.
        /// </summary>
        protected AnekaApplication<AnekaThread, ThreadManager> application;
        ....

        /// <summary>
        /// Starts the execution of the distributed application by creating the
        /// filters wrapping them into AnekaThread instances and starting their
        /// execution.
        /// </summary>
        /// <param name="source">Source bitmap.</param>
    }
}
```

```

protected void StartExecution(Bitmap source)
{
    this.running = new List<AnekaThread>();
    WarholFilter[] filters = this.CreateFilters(source);

    // creates an AnekaThread for each filter
    foreach (WarholFilter filter in filters)
    {
        AnekaThread thread = new AnekaThread(new ThreadStart(filter.Apply),
                                             this.application);

        thread.Start();
        this.running.Add(thread);
    }
}
}
}
}

```

Listing 7 - StartExecution(Bitmap) method.

The interesting bits in this method are concentrated within the foreach loop. For each *WarholFilter* instance that has been created a new instance of *AnekaThread* is initialized and configured to run the *WarholFilter.Apply* method. Each *AnekaThread* also need to have a reference to the *AnekaApplication* it belongs to. The second statement simply starts the execution of the *AnekaThread* instance by calling *AnekaThread.Start()*.

These two statements constitute the common operations required to configure and start a *AnekaThread* instance. We also add this instance to the list of running threads in order to keep track of its reference and being able to get the results once the *AnekaThread* has completed its execution.

```

namespace Aneka.Examples.ThreadDemo
{
    ...
    public class WarholApplication
    {
        ....

        /// <summary>
        /// Starts the execution of the distributed application by creating the
        /// filters wrapping them into AnekaThread instances and starting their
        /// execution.
        /// </summary>
        protected void WaitForCompletion()
        {
            this.done = new List<WarholFilter>();
            bool bSomeToGo = true;

```

```

while (bSomeToGo == true)
{
    foreach (AnekaThread thread in this.running)
    {
        thread.Join();
    }
    for (int i = 0; i < this.running.Count; i++)
    {
        AnekaThread thread = this.running[i];
        if (thread.State == WorkUnitState.Completed)
        {
            this.running.RemoveAt(i);
            i--;
            WarholFilter filter = (WarholFilter) thread.Target;
            this.done.Add(filter);
        }
        else
        {
            // it must be failed...
            thread.Start();
        }
    }
    bSomeToGo = this.running.Count > 0;
}
}
}
}
}

```

Listing 8 - WaitForCompletion() method.

This method exposes another classic synchronization pattern that is used while creating multi-threaded applications: threads synchronization. In this specific case we simply want to implement a barrier for all threads. This can be easily done by calling the *AnekaThread.Join()* method on all the instances that we have started and that are contained in the *running* list.

The method uses a while loop that will terminate once all the results have been collected. Inside the while loop two basic steps are performed:

- Wait for thread completion: the method invoke the *AnekaThread.Join()* method on all the threads contained in the *running* list. This call makes the application to wait until the thread terminate.
- Check thread results: once all the threads have terminated we iterate again the *running* list to check whether some thread has failed its execution or not by looking at the *State* property. If the thread has successfully completed its execution its state is set to *WorkUnitState.Completed*. In this case we simply remove the thread

from the *running* list and add *WarholFilter* instance referenced by the *AnekaThread.Target* property into the *done* list. If the other cases the thread is simply restarted.

The loop terminates when the *running* list is empty. This means that all the threads have successfully completed their execution and all the filters have been collected into the *done* list.

```
namespace Aneka.Examples.ThreadDemo
{
    public class WarholApplication
    {
        ....
        /// <summary>
        /// Starts the execution of the distributed application by creating the
        /// filters wrapping them into AnekaThread instances and starting their
        /// execution.
        /// </summary>
        /// <param name="source">Source bitmap.</param>
        protected void ComposeResult(Bitmap source)
        {
            Bitmap output = new Bitmap(source.Width * this.repeatX,
                                       source.Height * this.repeatY,
                                       source.PixelFormat);

            Graphics graphics = Graphics.FromImage(output);
            int row = 0, col = 0;
            foreach (WarholFilter filter in this.done)
            {
                graphics.DrawImage(filter.Image, row * source.Width,
                                     col * source.Height);

                row++;
                if (row == this.repeatX)
                {
                    row = 0;
                    col++;
                }
            }
            graphics.Dispose();
            if (string.IsNullOrEmpty(this.outputPath) == true)
            {
                this.outputPath = this.GetNewName(this.inputPath, "warhol");
            }
            output.Save(this.outputPath);
        }
    }
}
```

Listing 9 - ComposeResult(Bitmap) method.

The next logical step is constituted by the processing of the result and the creation of the final output image. These tasks are accomplished into the *AnekaThread.ComposeResult* method whose content is reported in Listing 9. The interesting bits in this method are concentrated within the foreach loop. For each *WarholFilter* instance that is contained into the *done* list the processed image exposed by the *WarholFilter.Image* property is drawn into the final output image in the position identified by the *row* and *col* local variables. The member field *repeatX* is used to identify the end of a line and move to the next colum.

After the output bitmap has been composed it is saved to the file pointed by *OutputPath* property or to an automatically generated file name by invoking the *WarholApplication.GetNewName* method.

```
namespace Aneka.Examples.ThreadDemo
{
    ...
    public class WarholApplication
    {
        ....
        /// <summary>
        /// Creates an array of WarholFilter instances each of them configured
        /// with the same input image and a different palette.
        /// </summary>
        /// <param name="source">Source bitmap.</param>
        /// <returns>Array of filters.</returns>
        protected virtual WarholFilter[] CreateFilters(Bitmap source)
        {
            WarholFilter[] filters = new WarholFilter[4];

            WarholFilter one = new WarholFilter();
            one.Image = source;
            one.Palette = WarholFilter.FuchsiaGreenWhite;
            filters[0] = one;

            WarholFilter two = new WarholFilter();
            two.Image = source;
            two.Palette = WarholFilter.YellowGreenNavy;
            filters[1] = two;

            WarholFilter three = new WarholFilter();
            three.Image = source;
            three.Palette = WarholFilter.FuchsiaOrangeBlue;
            filters[2] = three;

            WarholFilter four = new WarholFilter();
```

```

        four.Image = source;
        four.Palette = WarholFilter.GreenOrangeGainsboro;
        filters[3] = four;

        this.repeatX = 2;
        this.repeatY = 2;

        return filters;
    }
}
}

```

Listing 10 - CreateFilters(Bitmap) method.

The last method that we want to explore is the `WarholApplication.CreateFilters` method that is responsible of creating all the filter instances and define the number of rows and columns into which the final image will be organized. By separating the creation of filters into this method we can easily customize the output image by simply overriding this method and, for example, creating a final image that is composed by 9 samples of the original images or simply changing the colors of the palette.

4.5 Program: Putting all together

The program class implements a simple command line parser that reads the arguments of given by the user, checks whether they are correct, configures the `WarholApplication` class and starts the execution of the filter by invoking the `Run()` method. If the user has not provided the right parameters a simple command line help is show.

The class defines only two static methods: one is the entry point of the application (`Program.Main(string[])`) and the other one shows the command line help (`Program.ShowHelp()`).

```

// File: Program.cs
namespace Aneka.Examples.ThreadDemo
{
    /// <summary>
    /// Class Program. Virtualizes the execution of WarholFilter by using the
    /// Thread Model. This class simply parses the command line arguments passed
    /// to the process and sets up the WarholApplication.
    /// </summary>
    public class Program
    {
        /// <summary>
        /// Creates an array of WarholFilter instances each of them configured

```



```

/// with the same input image and a different palette.
/// </summary>
/// <param name="args">Command line arguments.</param>
static void Main(string[] args)
{
    if (args.Length >= 2)
    {
        string inputFile = args[0];
        string outputFile = args[1];
        string confFile = null;
        if (File.Exists(inputFile) == false)
        {
            Console.WriteLine("warholizer: [ERROR] input file [{0}] not" +
                "found. EXIT.", inputFile);

            return;
        }
    }
    else
    {
        // the input file exists...
        // now we check for the configuration file.
        if (args.Length == 3)
        {
            confFile = args[2];
            if (File.Exists(confFile) == false)
            {
                Console.WriteLine("warholizer: [ERROR] configuration" +
                    "file [{0}] not found. EXIT", inputFile);

                return;
            }
        }
        // now we check for the out file to simply issue
        // a warning if the file exists...
        if (File.Exists(outputFile) == true)
        {
            Console.WriteLine("warholizer: [WARNING] output file [{0}]"
                + "already exists and it will be overwritten.", inputFile);
        }
    }
    // ok at this point we have the following conditions
    // 1. inputPath exists
    // 2. confFile exists
    // we can start the application..
    WarholApplication app = new WarholApplication();
    app.InputPath = inputFile;
    app.OutputPath = outputFile;
    app.ConfigPath = confFile;
    try
    {
        app.Run();
    }
}

```

```

        }
        catch (Exception ex)
        {
            Console.WriteLine("warholizer: [ERROR] exception:");
            Console.WriteLine("\tMessage: " + ex.Message);
            Console.WriteLine("\tStacktrace: " + ex.StackTrace);
            Console.WriteLine("EXIT");
        }
    }
    else
    {
        Program.ShowHelp();
    }
}
/// <summary>
/// Shows a command line help about the usage of the application.
/// </summary>
static void ShowHelp() { .... }
}
}

```

Listing 11 - Program class.

4.6 Compiling and building the Application

4.6.1 Building the demo in Visual Studio 2005

It is possible to build and run the application by simply opening the Visual Studio 2005 Project *ThreadDemo.csproj* in the *ThreadDemo* directory and build the project. Visual Studio will create the executable *warholizer.exe* along with all the libraries required to run in the *ThreadDemo\bin\Debug* directory (Configuration: Debug).

Visual Studio 2005 will also copy the *conf.xml* and the *marilyn.jpg* file into the *bin\Debug* directory of the application. These two files can be used to test the execution of the application.

4.6.2 Building the demo from the command line

If you do not have the Visual Studio 2005 installed but you have c# 2.0 compiler (let us assume that the compiler is the one shipped with .NET framework SDK and that is called *cs.exe*) it is possible to compile the application from the command line.

Dependencies

The first step that is required is identifying the dependencies that this application relies on to execute:

1. Most of the support classes that we have used to build the application

are defined in the *System.dll* assembly that is referenced by default and contained in the Global Assembly Cache (GAC).

2. To perform the operations on the images we have used the *Bitmap* class which is defined in the *System.Drawing* namespace (GAC: *System.Drawing.dll*).
3. In order to use the Thread Model we used the types defined in the *Aneka.Threading* namespace that is implemented in the *Aneka.Threading.dll*.
4. Any application that uses the Aneka APIs has an implicit dependency on the following assemblies:
 1. *Aneka.dll* (Namespaces: *Aneka*, *Aneka.Entity*, *Aneka.Security*)
 2. *Aneka.Data.dll* (Namespaces: *Aneka.Data*, *Aneka.Data.Entity*)
 3. *Aneka.Util.dll* (Namespace: *Aneka* (utility classes))

The complete set of dependencies is then given by: *System.dll*, *System.Drawing.dll*, *Aneka.dll*, *Aneka.Threading.dll*, *Aneka.Data.dll*, *Aneka.Util.dll*. We can find the libraries that relate to Aneka into the *[Aneka Installation Directory]\bin* directory. The easiest thing to do is then copy these libraries to the *ThreadDemo* directory. The other two libraries are registered in the GAC, hence we do not need to copy them.

Compilation

Once we have copied the required libraries into the *ThreadDemo* directory we can invoke the C# 2.0 compiler to compile the three files (*WarholFilter.cs*, *WarholApplication.cs*, and *Program.cs*) that compose the application with the following command line:

```
csc /r:System.dll /r:System.Drawing.dll /r:Aneka.dll /r:Aneka.Data.dll  
/r:Aneka.Util.dll /r:Aneka.Threading.dll /t:exe /out:warholizer.exe  
Program.cs WarholFilter.cs WarholApplication.cs
```

The compilation process will create the *warholizer.exe* executable in the *ThreadDemo* directory.

4.6.3 Running the application

In order to run the application it is necessary to connect to have Aneka installed either on the local machine or on a remote machine that can be reached through a TCP connection. We assume, for the sake of simplicity, that Aneka is running on the local machine with the default installation (port: 9090). In this case we can simply run test the application by running the from the command line the following:

```
warholizer.exe marilyn.jpg
```

This application will produce the file *marilyn.warhol.jpg* in the same directory. We can also provide a different name (for example *foo.jpg*) of the output file by executing the following:

```
warholizer.exe marilyn.jpg foo.jpg
```

If we need to customize the way in which the application connects to Aneka. We can simply edit the configuration file *conf.xml* (for example we need to change the address where the application should connect) and run the following:

```
warholizer.exe marilyn.jpg foo.jpg conf.xml
```

Figure 6 shows the input file and a possible outcome of the execution of warholizer on the given input file.



Figure 6. Input (left) and output (right) images (not in original sizes).

5. Conclusions

In this tutorial we have introduced the Thread Model for developing distributed applications based on remotely executable threads with Aneka. The Thread Model allows developers to quickly virtualize multi-threaded applications with Aneka. It introduces the concept of *AnekaThread* that represents a thread that is executed on a remote computing node in the Aneka network. The *AnekaThread* class exposes a subset of the operations offered by the *System.Threading.Thread* class, this makes the transition from a local multi-threaded application to a distributed multi-threaded application straightforward.

As happens for local threads a *AnekaThread* is configured with a *ThreadStart* object that wraps the information required run a method. It is possible to start, join, and abort a thread in the same manner as we do with local threads. Few restrictions apply to the execution of remote threads. As explained in the tutorial, *AnekaThread* instances cannot be paused or run static methods and they do not support all the asynchronous operations of

the *Thread* class. Given these limitations, the Thread Model remains still appealing for developers that want to take advantage of distributed computing systems without learning a new programming model.

In order to explain and illustrate the approach to the development of distributed application by using the Thread Model a simple application has been developed: warholizer. This application is a multi-threaded image filters that reproduces the Warhol Effect on a given picture by leveraging the computation on Aneka. In particular the following aspects have been discussed:

- What is the Thread Model and how it relates with the common .NET threading APIs.
- How to create and configure a *AnekaApplication* for the Thread Model.
- How to create and configure a *AnekaThread*.
- How to control the execution of a Thread Model application by using *AnekaThread* instances.
- How to implement the common synchronization patterns used in multi-threaded applications.
- How to structure the source code of an application that is based on the Thread Model.
- How to compile and build a working example from the command line.

This tutorial has also introduced the general notions concerning a distributed system based on Aneka and the essential information for using the client APIs that are common to all programming models. For a more complete and detailed description of the behavior of these APIs it is possible to explore the APIs documentation.