

# **Active objects in Symbian OS**

Aapo Haapanen

University of Tampere  
Department of Computer Sciences  
Master's Thesis  
April 2008

University of Tampere

Department of Computer Sciences

Aapo Haapanen: Active objects in Symbian OS

Master's Thesis, 51 pages, 17 appendix pages

April 2008

---

This thesis examines a programming construct of Symbian OS called active objects. Active objects encapsulate a request to an asynchronous service and the completion of that request. They can be used to implement cooperative multitasking inside one thread. Active objects are widely used throughout Symbian OS, and the Symbian documentation encourages their usage instead of multithreading.

In this thesis active objects are compared to threads by implementing a solution to classic producer/consumer problem using both programming methods. The performance of the solutions is then compared. The test results show that the active-object based solution performs the operation more quickly and uses significantly less memory than the thread-based solution.

Key words and phrases: Symbian, active objects, cooperative multitasking, asynchronous processing

## Table of Contents

1. Introduction .....	1
2. Background.....	3
2.1 Multitasking.....	3
2.1.1 Preemptive and cooperative scheduling .....	5
2.1.2 Concurrent processes .....	6
2.2 Symbian OS.....	8
2.3 Technical overview of Symbian OS .....	10
3. Threads in Symbian OS.....	13
3.1 RThread .....	14
3.2 Mutexes, semaphores and critical sections.....	15
4. Active objects.....	17
4.1 Implementation of active objects in Symbian OS .....	18
4.1.1 Asynchronous requests .....	18
4.1.2 Handling the completion of asynchronous requests .....	19
4.2 CActive .....	20
4.3 CActiveScheduler .....	24
5. The test programs.....	27
5.1 Implementation of thread-based solution using semaphores .....	28
5.2 Solution using active objects, first version .....	31
5.3 Improved solution using active objects.....	35
6. Tests and analysis of the programs.....	40
6.1 Performance .....	40
6.2 Memory consumption.....	45
6.3 Other considerations .....	46
7. Summary and conclusions.....	48
References.....	50
Appendix A: Full source of the thread-based solution.....	52
Appendix B: Full source of the active-object-based solution, first version.....	57
Appendix C: Full source of improved active-object-based solution .....	62

## 1. Introduction

Almost all computer programs that have a graphical user interface have to be able to perform multiple tasks at the same time. For example, a program has to be able to handle user input while it is performing a time consuming task. Many such programs are also event driven. Most of the time they are doing nothing, except waiting for user input.

Most modern operating systems have the ability to perform multiple tasks at the same time. The user is able to run several programs simultaneously, and within one program several tasks can be performed simultaneously. The most common way to implement multitasking inside one program is to use threads. A thread is one stream of program execution, and the operating system schedules processing time for all threads. This kind of multitasking is called preemptive multitasking.

Symbian operating system (OS) is a C++ based operating system, which is designed for data-enabled mobile devices. It has a fully object oriented design and a compact implementation. Symbian OS has full preemptive multitasking and supports threads, but it also offers an event based, cooperative alternative to threads: active objects. Many programming tasks that are usually implemented with threads can also be implemented with active objects. Active objects are widely used throughout the Symbian OS, thus it is important to understand active objects when developing code to the platform.

The goal of this thesis is to examine the active objects on Symbian OS. It compares the active objects to threads and tries to define some guidelines about when active objects should be used and when using threads would be advantageous. Active objects and threads are analyzed by creating small test programs using both programming techniques, and comparing the

implementations in terms of time and memory consumption. Some less easily measurable considerations are also discussed.

This thesis focuses on Symbian OS version 7.0. That means that the test programs are implemented on that operating system version and the tests are run on Nokia 9300i, which uses that operating system. In Symbian OS v8.0b new kernel architecture was introduced, which included some changes to the handling of threads. Most of the results of this thesis should be valid for also other operating system versions, but especially the test results of the thread-based solution should be taken with a grain of salt when dealing with newer OS versions.

Chapter 2 of this thesis gives background information on cooperative and preemptive multitasking and on Symbian OS. Chapters 3 and 4 describe threads and active objects, respectively. Chapter 5 deals with the test programs and in Chapter 6 those programs are analyzed. In Chapter 7 the conclusions from this study are presented.

## 2. Background

To understand active objects one has to be familiar with issues related to multitasking and Symbian OS. In this chapter those issues are discussed.

### 2.1 Multitasking

The term *multitasking* is used to describe the ability of a computer system to run several programs seemingly simultaneously. Almost all modern operating systems are able to multitask.

However, computers are basically sequential in nature. A computer processor is able to perform only one instruction at the time. A computer that has no multitasking ability operates strictly sequentially. It can only have one thread of control. Such computer can only execute one program at a time, and if the computer has to perform some time-consuming task, all other tasks have to wait until that operation has completed. [Hansen, 1973]

Thus, it is usually operating system's job to offer multitasking ability to a computer system. Operating system is a set of system programs which controls all the computer's resources and provides the base upon which the application programs can be written. [Tanenbaum, 1987]

On single processor systems the operating system gives central processing unit (CPU) time to different tasks based on the tasks' priorities and on the scheduling algorithm used. Such multitasking is usually called *multiprogramming*. On multiprocessor systems the operating system has several CPUs from which it can give processing time, which adds some complexity to the problem. Sharing the workload between several processors is called *multiprocessing*. [Deitel, 1984] However, this thesis concentrates on single processor systems, because currently no multiprocessor systems exist that use Symbian OS.

*Process* is a very central term to the multitasking. A process can be defined as a program in execution. It consists of the executable program, its data and stack, program counter, stack pointer, other registers and other data needed to run the program. The difference between program and process is subtle, but important. Process is an activity, which can be started, suspended etc, while a program is more like a static set of instructions. [Tanenbaum, 1987]

The CPU is able to execute only one process at a time. The process that is currently in execution by the CPU is said to be *running*, or in the *running state*. Other processes that are ready to be executed, but don't right now have the CPU, are *ready* processes, or in the *ready state*. In addition to the processes that are ready to be executed there are processes that are waiting for some event to happen. Those processes are *blocked*, or in the *blocked state*. [Tanenbaum, 1987]

The operating system makes sure that all processes in the ready state have a chance to run. To do this, it alternates the process that is in the running state. The changing of the running process is called *context switching*. During the context switch the operating system takes one of the processes that were in the ready state and switches it to the running state. The process that was previously running is changed to the ready state. The operating system must later be able to activate the process at exactly the same state it was before the change. This means that all information about the process must be saved during the context switch. [Tanenbaum, 1987]

The term *thread* is often used in programming languages to denote some kind of a lightweight process. A common technical distinction between threads and processes is that a process runs in its own address space, while a thread runs inside an address space of a process. In theoretical discussion of concurrency the distinction is often unnecessary and the term process is used exclusively. [Ben-Ari, 2006]

The part of the operating system that gives the CPU to different processes is called *dispatcher*. It decides which process will get the CPU when it next becomes available. This is also called low-level scheduling. [Deitel, 1984]

Through multitasking the CPU usage of a computer system can be maximized. Usually on computer system most of the time is spent waiting for some input or output operation to finish – especially if there is a human

interacting with the computer and doing the input. By running several processes seemingly simultaneously, other processes can still operate normally when one process is waiting for some external event.

### 2.1.1 Preemptive and cooperative scheduling

Scheduling of processes is *nonpreemptive* if, once a process has been given the CPU, it cannot be taken away from that process until the process itself relinquishes the CPU. If the CPU can be taken away from the process, the scheduling is *preemptive*. [Deitel, 1984]

In nonpreemptive systems multitasking can be achieved if the processes cooperate. They can voluntarily yield the CPU to each others. Thus such system is called *cooperative multitasking*. Programs designed to run on cooperative multitasking system must be carefully designed, because if a misbehaving process doesn't relinquish the CPU, other processes won't be able to run at all.

Most current operating systems, including Symbian OS, offer preemptive multitasking. When the dispatcher in preemptive multitasking system gives the CPU to a process, it also sets a timer. If the process has not terminated until the timer runs out, the dispatcher switches the CPU to next process and the previously active thread is moved back into the dispatcher's activation queue. Thus all active processes get to run for a set amount of time in turns.

There are several ways to schedule the execution of the processes. One simple and commonly used scheduling method is *round-robin scheduling* where all the processes that are ready to run are executed in turn for a set amount of time. Each process gets a same share of CPU time. The situation gets more complicated if process priorities are added to the system. A higher priority process may always preempt a lower priority process, or the process with a higher priority may get a larger share of the CPU time than the lower priority process.

Preemptive multitasking is a lot more robust solution than cooperative multitasking. In preemptive multitasking one process can't steal all CPU time to itself. It also guarantees a reasonably quick response time for interactive applications.



However, there is also a drawback to preemptive scheduling. A context switch between two processes involves overhead, and in preemptive systems such a context switches can occur very rapidly. A cooperative multitasking can also make some programming tasks simpler, because the programmer can be certain that the CPU can't be switched to another process in the middle of some critical operation. Therefore, in some circumstances cooperative system may be a better solution.

In Symbian OS active objects are a way to implement cooperative multitasking inside a single thread, while the threads are scheduled preemptively. Active objects will be examined in following chapters.

### 2.1.2 Concurrent processes

Having several processes running simultaneously and perhaps interacting with each other gives rise to some new programming challenges. Two or more processes that exist at the same time are called *concurrent processes*. Although in single-processor systems only one process can be in execution at time, it is useful to think that all the processes that are in the running or ready state are executed concurrently.

On a multitasking system it is important that the data of each process is protected against unintended interference from other processes and that the results of each computation must be independent from the speed at which the computation is performed. [Hansen, 1973]

The simplest situation occurs when two processes are completely independent of each other. Such processes are called *disjoint processes*. Disjoint processes are processes that share no resources between themselves. For instance, they operate on completely different sets of variables. Such processes can't disrupt each other, and the relative order of their execution makes no difference to the results of the operations. [Hansen, 1973]

However, often processes are somehow dependent on each others. If processes operate on common resources, they must not disrupt each other. For instance, when one process is reading data, other processes may not change that data until the reading process is ready. Processes may also be cooperating, and

to facilitate the cooperation they must be able to exchange synchronizing information or results between each other.

The simplest form of interacting between processes is *mutual exclusion*. Mutual exclusion means that some operations in processes A and B should never be executed at the same time. For example, there might be some peripheral device, which only one process can use at a time. Another example is a situation where one process is updating some common variable. No other process should use that variable when one process is updating it. [Hansen, 1973, Deitel, 1984]

An elegant solution to the mutual exclusion problem can be achieved through *critical sections*. Critical section is a part of a program that can only be executed by one process at a time. When a process is executing its critical section, other processes wanting to enter critical sections must wait until there are no other processes in a critical section. If there are several processes waiting to enter a critical section, the method of deciding who is allowed to enter a critical section must be fair. That means that no process may be held off indefinitely from entering a critical section. That can be achieved, for example, by first-in-first-out (FIFO) queue. [Dijkstra, 1965a, Dijkstra, 1965b, Hansen, 1973, Deitel, 1984]

Process interaction through critical sections is quite indirect. For processes to be able to cooperate, they must have a more direct means to communicate with each other. A simple synchronizing signal between processes is called *semaphore*. It is a non-negative variable, which can only be accessed through two specific operations. The names of the operations are usually *V* and *P* or *Signal* and *Wait* in the literature. In Symbian OS implementation of semaphores they are called *Signal* and *Wait*, so those names are used in this thesis. *Signal* increases the value of the semaphore, and *Wait* decreases the value, when it is possible to do so without the value becoming negative. If a *Wait* operation would cause the semaphore to become negative, it will wait for a *Signal* before proceeding. Again, if there are several processes waiting for a *Signal*, the method of deciding which *Wait* operation to complete must be fair. [Dijkstra, 1965a, Deitel, 1984]

The semaphore operations are indivisible. Normally increasing or decreasing the value of a variable is done in at least two separate operations: first reading the old value and then saving the new modified value. With concurrent processes there is a risk that another process accesses the variable between the read and write operations. With semaphore operations there is no such risk. [Dijkstra, 1965a]

A simple example of process synchronization through semaphores is a situation where one process (*processone*) must wait for some condition before it can proceed, and another process (*processtwo*) is able to observe the completion of that condition. In that case a semaphore can be initialized to zero. When *processone* gets to a situation where it must wait for the condition it issues a Wait operation for the semaphore. When *processtwo* observes the fulfillment of the condition it issues a Signal operation for the semaphore, and *processone* may proceed. The mechanism works even if *processtwo* signals the semaphore before *processone* issues the Wait operation. [Deitel, 1984] Note that it is also easy to implement mutual exclusion and critical sections with semaphores.

One common relationship between processes is *producer/consumer relationship*. In producer/consumer relationship one process produces information, while another process processes the information as it becomes available. That means there is a one-way information flow between the processes. Usually there is some kind of a buffer where the producer puts the information it produces, and from where the consumer takes the information it processes. Now there is a need for process synchronization. The producer must not create more information if the buffer is full, and the consumer must not try to process information when the buffer is empty. Additionally, only one process may access the buffer at the same time. [Dijkstra, 1965, Deitel, 1984]

## 2.2 Symbian OS

Symbian OS is an operating system designed for data-enabled mobile phones. The roots of Symbian OS go back to Psion handheld organizers and their operating system, EPOC32. In 1998 Symbian Limited was formed by Psion, Nokia, Motorola and Ericsson, and the operating system was re-branded as

Symbian OS. Later also Siemens, Panasonic and Sony-Ericsson have joined the company, while Psion has sold its share to Nokia. [Symbian]

According to Symbian [Symbian], in 2007 there were 68 phone models using Symbian OS commercially available, and 77.3 million mobile phones using Symbian OS were sold. The market share of phones with Symbian OS was 7% of all mobile phone sales.

Around the time of the formation of Symbian Ltd the mobile phones were getting more and more complicated, and were starting slowly to move from single-use appliances to a platform for many different kinds of applications. Thus, a need for a real operating system emerged. Symbian was formed to develop Psion's EPOC32 operating system into an operating system for data-enabled mobile phones.

According to Mery [2002] the mobile phone market has five key characteristics, which make it unique and require a specifically designed operating system:

- Mobile phones are small and mobile. This limits all resources available for the system. A mobile phone has limited amount of electricity, CPU power, execution memory, and storage size. Yet the operating system should be always available.
- Mobile phones are aimed at the mass market. This means that the operating system should be very reliable. User data should never be lost, and the phone should never lock up or even need to be rebooted.
- Mobile phones are occasionally connected. They can be connected to a mobile phone network or to another device, or they might not be connected to any outside system.
- Manufacturers need to be able to differentiate their products in order to innovate and compete in a fast-evolving market. The operating system needs to adapt to and support different device form-factors and input methods.
- The operating system must allow development of third party applications and services.

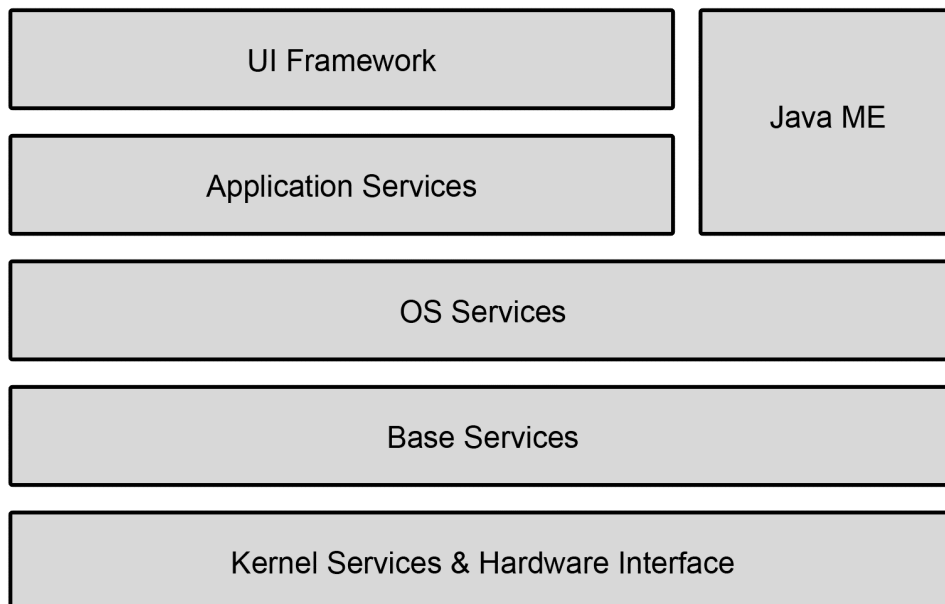
Symbian OS has been developed from these five key points. That has resulted in an operating system that is quite distinct from any

desktop/workstation/server operating system, but which is also quite different from embedded operating systems. [Mery, 2002]

### 2.3 Technical overview of Symbian OS

Symbian OS has a microkernel architecture, where the kernel handles only the lowest level machine resources, CPU cycles and memory. The kernel runs in a supervisor mode, which means that it has a hardware-supported privileged access to hardware resources and only it can execute some CPU operations. Other programs must access hardware through kernel API. Most of the services offered by the operating system are handled by servers that run in user mode. [Tasker, 2000, Morris, 2007]

Symbian OS is fully object-oriented operating system. It is implemented in C++, although some of the design decisions are unique to the Symbian OS. For instance, the exception handling doesn't use the standard C++ model; instead Symbian OS uses its own leave-mechanism.



**Figure 1:** A layered view of the components in Symbian OS [Morris, 2007]

Due to the modular nature of Symbian OS and the various technologies supported, any attempt to draw an architectural image of the system get easily very complicated. Figure 1 gives a one way to look at the different components

in the system, where the components are divided into layers starting from UI layer all the way down into the kernel and hardware interface layer.

UI framework layer provides the basic user interface building blocks for the applications. It acts as an interface between Symbian OS and the variant UI layer. Symbian OS doesn't include an actual user interface. The implementation of the UI is left to the licensees. Two most important user interfaces currently available are S60 by Nokia and UIQ by UIQ Technology AB (fully owned by Sony-Ericsson). The purpose of the UI framework layer is to allow the customization of the UI without fragmentation. [Morris, 2007]

Application services layer provides UI-independent services to the applications. It includes system services used by all applications, such as application framework, technology-specific logic that is used by multiple applications, such as support for messaging and multimedia protocols, and it also includes services supporting specific applications such PIM (personal information management) and office applications. [Morris, 2007]

OS services layer provides services that extend the services provided by the kernel and low-level system libraries into a full operating system. It includes servers, frameworks and libraries that implement support for graphics, communications, connectivity and multimedia. It also includes some generic system components such as C standard library. [Morris, 2007]

Base services layer provides the lowest-level user side services. This layer together with the kernel services layer forms the minimal operating system. Since the actual kernel is a microkernel, everything above the basic operating system privileges is kept outside the kernel and runs in the user mode. [Morris, 2007]

The lowest layer, kernel services and hardware interface layer, contains the operating system kernel itself, and also the components that abstract the interface to the underlying hardware, including device drivers. Everything in this layer runs in the supervisor mode. [Morris, 2007]

As mentioned above, Symbian OS uses client-server model throughout the system to implement services. The kernel itself is a server that offers access to the hardware. File server handles the file system; font and bitmap server offers

access to the physical display; window server handles the windowing system and is the heart of the applications' event handling system. [Morris 2007]

Symbian OS uses asynchronous services widely through the system. That means that the services are implemented using the request—callback model instead of the blocking model. Active objects are used to encapsulate the calling of an asynchronous service and handling the completion of the request. Active objects are examined in more detail in Chapter 4.

Symbian OS is optimized for event handling [Harrison, 2003]. With GUI systems the system spends most of the time waiting for user input, after which it has to quickly respond to that input. User input can be thought of as distinct events that the system handles. When the user presses a key on the keyboard, he creates an event, and when the key is released, another event is created. If the user has a pointing device, that also creates events for the system to react to. Both active objects and servers play a central role in the event handling system. When an application wants to handle certain events, it issues an asynchronous request – using active objects – to the server handling such events. And when the event happens, the request is completed and the application can handle the event and possibly renew the request to receive the events.

Symbian OS version 8.0b introduced a new kernel called EKA2 (EPOC Kernel Architecture 2). The most important changes were the introduction of real-time capabilities, which makes it possible to run GSM protocol stack on the operating system, and the introduction of platform security, which is intended to make the operating system more robust against malicious code. EKA2 also introduced some changes to the handling of threads inside the kernel [Sales, 2005]. This thesis focuses on Symbian OS v7.0, which uses the older EKA1 kernel.

### 3. Threads in Symbian OS

Symbian OS provides preemptively scheduled threads and processes. It also provides standard mechanisms that are needed to synchronize access to resources shared between threads. These methods include semaphores, mutexes and critical sections.

In Symbian OS, a thread is defined as the unit of execution. It is the entity that the kernel schedules, i.e. for which it allocates CPU resources. A process is defined as a collection of threads that share the same address mapping. The process is the fundamental unit of protection in Symbian OS. Each thread within a process can read and write from any others' memory, but they can't directly access the memory of other processes. Each application runs inside its own process. [Sales, 2005, Tasker, 2000]

Threads are preemptively scheduled by the kernel. Each thread has a priority, and the thread that has the highest priority and is not in the blocked state is run by the kernel. If several threads have the same priority, they are time-spliced in a round-robin basis. [Stichbury, 2005] Threads may be blocked because they are waiting for an event to happen, and they resume when the event happens. A context switch between threads inside one process is much cheaper operation than a context switch between threads in different processes. The reason for that is that a context switch between processes requires changes to MMU (memory management unit) settings and various caches need to be flushed. Still, even a context switch between two threads inside the same process is much more expensive operation than, for example, a function call. [Tasker, 2000]

In addition to the user processes, there is also a special kernel process with two threads that run in the supervisor privilege level. One of them is the kernel server thread and it always has the highest priority in the system. The other one



is the idle thread, which always has the lowest priority in the system. The idle thread switches the processor to the idle mode, allowing the system to conserve power. [Stichbury, 2005, Tasker, 2000] The focus of this thesis is user level programs, so by default, user threads are meant when discussing threads.

### 3.1 RThread

The threads in Symbian OS are accessed through *RThread* class. A thread is a kernel object, and *RThread* represents a handle to the object. The full definition of the class is too long to be included here, but Code fragment 1 contains an abbreviated version of the definition, with the most important functions included.

---

```
class RThread : public RHandleBase
{
public:
    inline RThread();
    IMPORT_C TInt Create(const TDesC& aName, TThreadFunction
        aFunction, TInt aStackSize, TAny* aPtr, RLibrary* aLibrary, RHeap*
        aHeap, TInt aHeapMinSize, TInt aHeapMaxSize, TOwnerType aType);
    IMPORT_C TInt Create(const TDesC& aName, TThreadFunction
        aFunction, TInt aStackSize, TInt aHeapMinSize, TInt
        aHeapMaxSize, TAny *aPtr, TOwnerType aType=EOwnerProcess);
    IMPORT_C TInt Create(const TDesC& aName, TThreadFunction
        aFunction, TInt aStackSize, RHeap* aHeap, TAny* aPtr, TOwnerType
        aType=EOwnerProcess);
    IMPORT_C TInt Open(const TDesC& aFullName, TOwnerType
        aType=EOwnerProcess);
    IMPORT_C TInt Open(TThreadId aID, TOwnerType aType=EOwnerProcess);
    IMPORT_C void Resume() const;
    IMPORT_C void Suspend() const;
    IMPORT_C void Kill(TInt aReason);
    IMPORT_C void Terminate(TInt aReason);
    IMPORT_C TThreadPriority Priority() const;
    IMPORT_C void SetPriority(TThreadPriority aPriority) const;
};
```

---

**Code fragment 1:** An abbreviated definition of RThread class from e32std.h

The class is inherited from *RHandleBase* class, which provides some standard handle manipulation functionality, such as *Duplicate()* for duplicating the handle and *Close()* for closing the handle.

The default constructor initializes the class to a pseudo handle set to the constant *KCurrentThreadHandle*. That value is treated as a special case by the kernel and can be used to access the current thread. If a proper handle to the

current thread is needed, it can be created using *Duplicate()* on the pseudo handle [Stichbury, 2005].

A new thread is created with *Create()* function. There are several overloaded versions of the function to allow setting various options associated with the heap of the thread. A thread may have its own heap, or it may use the heap of the thread that created it. The size of the heap can change during the execution of the thread, and the minimum and maximum sizes of the heap can be defined during the thread creation. The size of the thread's stack is fixed to the value defined at the time of the thread creation. The default stack size is 8 KB. Other parameters given to the *Create()* function include the thread's name, a pointer to the function where thread execution starts, and a pointer to data that is passed as a parameter to the thread function. All new threads are created with priority *EPriorityNormal*, and the priority can be changed after the creation with *SetPriority()* function. [Stichbury, 2005]

A handle to an existing thread can be acquired with *Open()* function. The thread to be opened can be indicated by either the name of the thread or the identification number of the thread. [Symbian, 2002]

When a new thread is created, it is initially put into the suspended state. The thread is started with *Resume()* function. The execution can be again suspended with *Suspend()* function. The thread is permanently ended by using functions *Kill()* and *Terminate()*. Both functions take an integer as a parameter. The parameter represents the exit reason of the thread. The functions act in a similar way, but information about which function was used can be retrieved from the ended thread, along with the exit reason. [Stichbury, 2005]

### **3.2 Mutexes, semaphores and critical sections**

For synchronizing the execution of threads, Symbian OS provides mutexes, semaphores and critical sections. Mutexes and semaphores can be either local, which means they are restricted to the current process, or they can be global, in which case they have a name that can be used to find the object in other processes. Critical sections are always local to a single process. [Symbian, 2002]

Mutexes and semaphores are kernel objects, and they are manipulated through handles *RMutex* and *RSemaphore*. Both classes are used in a similar

way. A new mutex or semaphore is created with function *CreateLocal()* or *CreateGlobal()*. The global version of the function takes the name of the object as a parameter. Creating a semaphore also requires the initial value of the semaphore. After the mutex or semaphore is created, it is used with functions *Wait()* and *Signal()*. [Symbian, 2002]

Critical sections are implemented using semaphores, with class *RCriticalSection* inheriting from *RSemaphore*. The critical section handle doesn't expose all of the semaphore's functionality, so only a local critical section can be created. [Symbian, 2002]

#### 4. Active objects

Active objects are a programming concept that is quite unique to Symbian OS. They are a way to implement nonpreemptive multitasking inside a thread. They are also closely related to the event handling system in Symbian OS.

At the core of a thread that uses active objects is an active scheduler. It acts as a kind of a mini-kernel for that thread, while active objects act like nonpreemptively scheduled mini-threads. Each thread may have only one active scheduler, while a thread with an active scheduler has one or more active objects. [Tasker, 1999]

A thread with active objects is basically an event handler. Event handling systems are based around programs requesting some services, which will then complete at a later time. Such completion of a request is called an event. When an outstanding request completes, the requester must handle it. Active objects encapsulate this relation between making a request and handling its completion. Each active object is responsible for making and handling just one outstanding request at a time. [Tasker, 1999]

Often a single thread will issue many outstanding requests. Each request may complete at any time, but each request is guaranteed to complete only once. Active scheduler handles the completion of requests and calls the active object responsible for handling the request. Each active object has a *RunL()* function, which is called when the request of that active object completes. If there are several completed requests to handle, the active scheduler decides the order in which they are handled. [Tasker, 1999]

Since active objects are scheduled nonpreemptively, their *RunL()* functions must return relatively quickly. When an active object is executing its *RunL()* function, no other active object inside the thread can run until the function returns.

In Symbian OS almost all threads have an active scheduler. It follows that almost all code runs inside *RunL()* functions of active objects. However, for most part the programmer doesn't have to think about active objects, because the frameworks of Symbian OS provide them. For instance, a program using GUI (graphical user interface) framework must implement *OfferKeyEventL()* function if it needs to handle keyboard events. That function is called by a *RunL()* function implemented in the framework [Stichbury, 2005]. All the programmer has to know is that the function must return relatively quickly, so that the application remains responsive. Typically applications and servers in Symbian OS contain only one thread, and the asynchronous processing is handled completely by active objects. [Harrison, 2003]

It must be noted that the nonpreemptive multitasking implemented by active objects is a bit different from the traditional cooperative multitasking. In cooperative multitasking tasks must actively tell other tasks when they can be executed, using *Yield()* or similar function. This easily leads to messy programming. In Symbian OS, each event must be handled completely before the other events can be handled. [Harrison, 2003]

Almost everything that can be implemented using threads can also be implemented using active objects. In practice, the use of multithreading in an application is quite rare in Symbian OS. An example of tasks that can't be implemented using active objects is a long-running task that can't be reasonably split into short discrete parts. Another is a task that requires a real-time response because no active object can preempt a running active object. [Tasker, 1999]

## **4.1 Implementation of active objects in Symbian OS**

### **4.1.1 Asynchronous requests**

As mentioned above, active objects are used to encapsulate making an asynchronous request and handling its completion. In Symbian OS any function that takes a *TRequestStatus&* parameter is designed to function asynchronously.

[Tasker, 1999] Classes containing such functions are usually called *asynchronous service providers*. [Harrison, 2003]

*TRequestStatus* is simply a well-encapsulated integer, where the only operations available are assignment and comparison. When an asynchronous request is issued, the asynchronous service provider sets the *TRequestStatus* parameter to *KRequestPending*. When the request completes, the service provider assigns the completion code to *TRequestStatus*. The completion code is usually one of the standard error codes of Symbian OS; anything except *KRequestPending* is permissible. For successful completion of the request the code usually is *KErrNone*. Each request must complete precisely once. [Tasker, 1999, Harrison, 2003]

The issuer of an asynchronous request must call *User::WaitForRequest()* or *User::WaitForAnyRequest()* function to wait for the completion of request, and for each wait, exactly one completion of a request must be handled. When active objects are used to handle the requests, this is done by the active scheduler. [Tasker, 1999] When the asynchronous service provider completes the request, it must call the *User::RequestComplete()* function if the provider is in the same thread as the requestor. If the requestor is in a different thread, it must call *RThread::RequestComplete()* for the requesting thread's handle. [Stichbury, 2005]

Every asynchronous service provider must also provide a cancel function to cancel an outstanding request. When a cancel function is called and the request has not already been completed, the service provider must either immediately complete the request with completion code *KErrCancel*, or if the request can be completed normally in a guaranteed very short time, the request may be completed normally. The cancel function must work synchronously, so that when the function returns, the request has been completed. In any case canceling request may not break the rule that every request must complete exactly once. [Tasker, 1999, Harrison, 2003]

#### **4.1.2 Handling the completion of asynchronous requests**

Code fragment 2 shows how a simple call to an asynchronous function could be implemented. [Tasker, 1999]

---

```

RFile file;
TRequestStatus readStatus;
...
file.Read(buffer, readStatus); // issue request
User::WaitForRequest(readStatus); // wait for completion

```

---

**Code fragment 2:** An example of a call to an asynchronous function

However, with this kind of coding the thread cannot do anything while waiting for the request to complete. Normally a single thread will issue several outstanding requests. Each request may complete at any time, but each request is guaranteed to complete exactly once. The thread must be able to handle this correctly. A wait loop for handling several outstanding requests could look like this [Tasker, 1999]:

1. Call *User::WaitForAnyRequest()*. This will either suspend the thread until a request completes, or if one or more requests have already completed, the function will return immediately.
2. Search through the statuses of requests issued but not yet handled for one that is not *KRequestPending* any more.
3. When such a request is found, handle its completion, and mark the request no longer issued.
4. Go back to 1.

It is important that each pass through this function handles exactly one completion of a request, because there must be as many calls to *User::WaitForAnyRequest()* as there are requests issued. If *User::WaitForAnyRequest()* is called when there are no outstanding requests, the function will never return.

The wait loop above with addition of priorities for requests is encapsulated in *CActiveScheduler*.

## 4.2 CActive

Class *CActive* encapsulates a single request to an asynchronous service provider. It is a pure virtual class, so users of it must inherit from it and implement its pure virtual member functions, such as *RunL*. *CActive* is defined in header *e32base.h* and the definition is shown in Code fragment 3.

---

```

class CActive : public CBase

```

```

    {
public:
enum TPriority
    {
    EPriorityIdle=-100,
    EPriorityLow=-20,
    EPriorityStandard=0,
    EPriorityUserInput=10,
    EPriorityHigh=20,
    };
public:
    IMPORT_C ~CActive();
    IMPORT_C void Cancel();
    IMPORT_C void Deque();
    IMPORT_C void SetPriority(TInt aPriority);
    inline TBool IsActive() const;
    inline TBool IsAdded() const;
    inline TInt Priority() const;
protected:
    IMPORT_C CActive(TInt aPriority);
    IMPORT_C void SetActive();
// Pure virtual
    virtual void DoCancel() =0;
    virtual void RunL() =0;
    IMPORT_C virtual TInt RunError(TInt aError);
public:
    TRequestStatus iStatus;
private:
    TBool iActive;
    TPriQueLink iLink;
    friend class CActiveScheduler;
    friend class CServer;
    };

```

---

**Code fragment 3:** Definition of CActive class

Code fragment 4 shows what a simple class derived from *CActive* might look like.

---

```

class CCompleteAfter5Seconds : public CActive
    {
public:
    static CCompleteAfter5Seconds* NewL();
    ~CCompleteAfter5Seconds();
    void StartTimer();

private:
    // constructors
    CCompleteAfter5Seconds();
    void ConstructL();

    // from CActive
    void RunL();
    void DoCancel();

private:
    // data
    RTimer iTimer;
    };

```

---

**Code fragment 4:** Example definition of a simple active object [Harrison, 2003]



Consider now the implementation of this simple active object. Code fragment 5 shows the implementation of the functions that handle the construction and destruction of the object.

---

```

CCompleteAfter5Seconds* CCompleteAfter5Seconds::NewL()
{
    CCompleteAfter5Seconds* self = new(ELeave) CCompleteAfter5Seconds;
    CleanupStack::Push(self);
    self->ConstructL();
    CleanupStack::Pop(self);
    return self;
}

CCompleteAfter5Seconds::CCompleteAfter5Seconds()
    : CActive(0)
{
    CActiveScheduler::Add(this);
}

void CCompleteAfter5Seconds::ConstructL()
{
    User::LeaveIfError(iTimer.CreateLocal());
}

CCompleteAfter5Seconds::~CCompleteAfter5Seconds()
{
    Cancel();
    iTimer.Close();
}

```

---

**Code fragment 5:** Functions handling the construction and destruction of an example active object [Harrison, 2003]

*NewL()* is a static function that follows the standard two-phased construction pattern of Symbian OS.

Every active object must implement a C++ constructor, which calls the constructor of *CActive* to set the priority of the object. The priority may be either positive or negative, but if there is no good reason to set it otherwise, it should be 0. If the active scheduler has several completed requests to handle at the same time, it uses the priority to decide which request to handle first. The requests with highest priorities are always handled first. After construction the priority can be accessed with *Priority()* and *SetPriority()* functions. The constructor also adds the object to the active scheduler, so that it can be included in the event handling. [Harrison, 2003]

The second phase constructor *ConstructL()* creates a timer object, which is accessed through *RTimer* handle. The timer will be the asynchronous service provider for this active object. [Harrison, 2003]

The destructor cancels any outstanding request by calling the *Cancel()* function defined in base *CActive* class. *Cancel()* checks whether there is an outstanding request, and if there is, it calls *DoCancel()*, which is a pure virtual member function of *CActive* and must be implemented by the derived class. *DoCancel()* must call the cancel function of the asynchronous service provider. The destructor also closes the handle to the timer, which causes the timer object to be destroyed. The destructor of *CActive* base class removes the active object from active scheduler. [Harrison, 2003]

---

```
void CCompleteAfter5Seconds::StartTimer()
{
    __ASSERT_ALWAYS(!IsActive(),
                    User::Panic("CCompleteAfter5Seconds", 1));
    iTimer.After(iStatus, 5000000);
    SetActive();
}

void CCompleteAfter5Seconds::RunL()
{
    // 5 seconds has passed
}

void CCompleteAfter5Seconds::DoCancel()
{
    iTimer.Cancel();
}
```

---

**Code fragment 6:** Implementation of the functions that handle the request

Code fragment 6 shows the functions that handle the asynchronous request. In *StartTimer()* function the asynchronous request is issued. The *iStatus* member function is defined in *CActive* base class and it contains the *TRequestStatus* for the asynchronous function call. But before issuing the asynchronous request, we must check that there isn't already an outstanding request. A single active object can handle only one outstanding request at a time. [Harrison, 2003] Finally, *SetActive()* is called to set the active flag of the active object. The flag tells the active scheduler that this active object has an outstanding request, and needs to be included in the search for the right object when a request completes. [Tasker, 1999]

When the request issued in *StartTimer()* completes the active scheduler will clear the active flag of the active object and call its *RunL()* function. In this example the active object doesn't actually do anything, but the *RunL()* can be very elaborate. After all, most code running in Symbian OS runs inside a *RunL()*

function. *RunL()* can also issue a new request. If in this example *RunL()* called *StartTimer()* again, this active object would stay active until it is cancelled, and the *RunL()* would be called every five seconds. [Tasker, 1999, Harrison, 2003]

Finally, we have *DoCancel()* function, which cancels an outstanding request. This is called by *Cancel()* function, which is defined in *CActive* base class. It checks if there is an outstanding request, so there is no need to check the status in this function. Also if *Cancel()* calls *DoCancel()* to cause the outstanding request to complete, it also calls *User::WaitForRequest()* and clears the active flag so that the active scheduler doesn't handle the completion of this request. That means that the *RunL()* doesn't get called if the request is explicitly cancelled. [Harrison, 2003]

In *CActive* there is also an error handling function called *RunError()*. Our example object is so simple that there is no need to implement *RunError()*, because the default implementation from *CActive* suffices. Actually in this object the *RunError()* will never be called. The active scheduler calls *RunError()* if *RunL()* leaves. If *RunError()* handles the error situation, it must return *KErrNone*. Otherwise it must return some error code, and the error situation is handled in the active scheduler's *Error()* function. The default implementation of *RunError()* just returns the leave code from *RunL()*. [Harrison, 2003]

The final functions in *CActive* that have not yet been mentioned are *IsAdded()* and *Dequeue()*. *IsAdded()* just tells if the active object has been added to the active scheduler. *Dequeue()* removes the active object from active scheduler and it is called by *~CActive()*. There is usually never a reason to call this function elsewhere. [Harrison, 2003]

### 4.3 CActiveScheduler

The active scheduler implements the active object handling loop described in Subsection 4.1.2. It maintains a doubly linked list of all the active objects added to it, ordered by priority. [Stichbury, 2005] The active scheduler is implemented in class *CActiveScheduler*, and it is defined in *e32base.h*. The definition is shown in Code fragment 7.

---

```
class CActiveScheduler : public CBase
{
```

```

public:
    IMPORT_C CActiveScheduler();
    IMPORT_C ~CActiveScheduler();
    IMPORT_C static void Install(CActivescheduler* aScheduler);
    IMPORT_C static CActiveScheduler* Current();
    IMPORT_C static void Add(CActive* anActive);
    IMPORT_C static void Start();
    IMPORT_C static void Stop();
    IMPORT_C virtual void WaitForAnyRequest();
    IMPORT_C virtual void Error (TInt anError) const;
protected:
    inline TInt Level() const;
private:
    TInt iLevel;
    TPriQue<CActive> iActiveQ;
};

```

---

**Code fragment 7:** The definition of the CActiveScheduler class

For all applications using the CONE framework to implement the GUI an active scheduler is already provided by the framework. However, if the thread has no active scheduler, it must be created and started before active objects can be used. *CActiveScheduler* is a concrete class that can be used as is, but it also has two virtual functions that can be redefined in derived classes. [Harrison, 2003]

*Install()* sets a pointer to an active scheduler in a privileged thread-local storage (TLS) location that is very fast to access. After installing the active scheduler it can be accessed with *CActiveScheduler::Current()*. *Add()* adds an active object to the scheduler. [Harrison, 2003]

*Start()* starts the active scheduler. There should be at least one outstanding request on an active object before calling *Start()*, otherwise the thread will hang. *Start()* includes the central wait loop of the active scheduler, and it will not return until *CActiveScheduler::Stop()* is called. *Stop()* causes *Start()* to return after currently active *RunL()* or *Error()* has completed. *Start()* and *Stop()* functions can be nested, and this is used by the UI framework in modal dialogs. However, that can easily lead to messy code, and should be avoided when possible. [Harrison, 2003]

*WaitForAnyRequest()* just calls *User::WaitForAnyRequest()*, but it can be overridden in derived classes to perform some special processing before waiting for a request. The overriding function must call *User::WaitForAnyRequest()*. [Harrison, 2003]

*Error()* handles all errors that have not been handled by active object itself in *CActive::RunError()*. The standard implementation in *CActiveScheduler()* does

nothing, but derived classes can implement some error handling there. The active scheduler in CONE displays a natural-language version of the error code.  
[Harrison, 2003]

## 5. The test programs

To analyze the differences between threads and active objects a small test program is written using both threads and active objects. The program implements a solution to the producer/consumer problem. There are three versions of the solution. First, a thread-based solution is implemented using semaphores as first described by Dijkstra [1965a]. Then two different solutions using active objects are presented.

In the producer/consumer problem two groups of objects use the same shared resource to communicate with each other. The producers create information, which is used by the consumers. The producers create the data one piece at a time. When producer creates a piece of information, it is put into a shared buffer to wait for a consumer to process it. The consumers take the information from the buffer one piece at a time, and process it in some way.

The buffer that is used for communication between the producers and consumers has a limited size. When the buffer is full, the producers must not try to add more pieces of information to it. When the buffer is empty, the consumers must wait for a producer to insert a piece of information to the buffer.

There are many ways to implement such a buffer. One way is to use *cyclic* or *circular buffer*. A cyclic buffer has cursors pointing the location of the first empty space and the location of the next data item to be taken out of the buffer. When data is added to the buffer or removed from it, the corresponding cursor is moved to next location in the buffer, or if the cursor was in the last location, it is moved to the first item. [Bacon and Harris, 2003]


It is important that only one object accesses the shared buffer at a time. A consumer must not try to take a piece of information while a producer is in the middle of adding it to the buffer. Similarly, two producers can't try to add

information at the same time to the buffer, and two consumers must not try to consume the same piece of information.

A simple version of the producer/consumer problem has one producer and one consumer. In the general version of the problem there can be any number of producers and consumers.

### 5.1 Implementation of thread-based solution using semaphores

The thread-based implementation follows the solution given by Bacon and Harris [2003]. The solution is described in Figure 2.



**Figure 2:** Solution to producer/consumer problem with semaphores [Bacon and Harris, 2003]

This solution uses three semaphores to control the execution of the producers and consumers. First we have a mutex (*iGuard*), which allows only one object to access the buffer at a time. Another semaphore (*iItems*) controls the consumers, so that the consumers don't try to process information when the buffer is empty. The value of that semaphore is equal to the value of items in the buffer. Finally, a third semaphore (*iSpaces*) controls the producers, so that they don't try to add anything to the buffer when the buffer is full. The value of that semaphore is equal to the empty spaces in the buffer.

In the Symbian implementation of this solution we have a controller class that creates all the threads and owns the circular buffer that is used for

communication between threads. Code fragment 8 contains the abbreviated class definition of the controller class *CController*.

---

```
class CController
{
public:
    CController(TInt aNumProducers, TInt aNumConsumers);
    ~CController();
    void ConstructL();
    void Start();
    void Stop();
public: // data
    CBuffer *iBuffer;
    RSemaphore iSpaces;
    RSemaphore iItems;
private:
    RArray<RThread> iProducers;
    RArray<RThread> iConsumers;
    TInt iNumProducers;
    TInt iNumConsumers;
};
```

---

**Code fragment 8:** The abbreviated definition of *CController* class

The semaphore *iGuard* is not defined here. Instead, it is owned by class *CBuffer*. Symbian OS provides a circular buffer class *CCirBuf*. However, that class is not thread safe. In other words, if two threads try to add or remove objects from a *CCirBuf* object at the same time, data corruption may occur. *CBuffer* class encapsulated the *CCirBuf* and protects its access function by a mutex. Code fragment 9 contains the definition of *CBuffer* class.

---

```
class CBuffer
{
public:
    CBuffer();
    ~CBuffer();
    void ConstructL();
    void AddInfo(const TData *aData);
    void GetInfo(TData *aData);

private:
    RMutex iGuard;
    CCirBuf<TData>* iBuffer;
};
```

---

**Code fragment 9:** The definition of *CBuffer* class

Code fragment 10 shows the implementation of the access functions. The access to *CCirBuf* is protected by mutex *iGuard*. This means that only one thread can access the buffer at a time. This encapsulation makes the circular buffer thread-safe.



---

```

void CBuffer::AddInfo(const TData *aData)
{
    iGuard.Wait();
    iBuffer->Add(aData);
    iGuard.Signal();
}

void CBuffer::GetInfo(TData *aData)
{
    iGuard.Wait();
    iBuffer->Remove(aData);
    iGuard.Signal();
}

```

---

**Code fragment 10:** The implementation of access functions of CBuffer class

Code fragment 11 contains the producer thread's main function. The thread gets a pointer to the controller class as its start parameter. The main loop of the thread looks very much like the pseudo code in Figure 2. First we create a data item. If the buffer is full, semaphore *iSpaces* has value 0. In that case the thread stops at the *iSpaces.Wait()*, and continues when a consumer thread signals that semaphore. The semaphore is signaled when a consumer has removed an item from buffer, thus creating an empty space for a new item.

---

```

TInt ProducerMain(TAny* aPtr)
{
    CController *controller = (CController*) aPtr;
    FOREVER
    {
        TData data;
        ProduceDataItem(data);
        controller->iSpaces.Wait();
        controller->iBuffer->AddInfo(&data);
        controller->iItems.Signal();
    }
    return KErrNone;
}

```

---

**Code fragment 11:** The implementation of producer thread's main function

Code fragment 12 contains the consumer thread's main function. This also closely resembles the pseudo code in Figure 2. The consumer also gets a pointer to the controller object as a start parameter. If *iItems* has value 0 it means that there are no items in the buffer. In that case the consumer thread waits for a signal from a producer thread. The producer thread signals the *iItems* when it has added an object to the buffer.

---

```

TInt ConsumerMain(TAny* aPtr)
{
    CController *controller = (CController*) aPtr;

```

```

FOREVER
{
    TData data;
    controller->iItems.Wait();
    controller->iBuffer->GetInfo(&data);
    controller->iSpaces.Signal();
    ConsumeDataItem(data);
}
return KErrNone;
}

```

---

**Code fragment 12:** The implementation of consumer thread's main function

The controller class *CController*'s constructor functions *CController()* and *ConstructL()* create the semaphores and the producer and consumer threads. The implementation details of those functions and the rest of the program are not interesting. The full source code of this test program is available in Appendix A.

## 5.2 Solution using active objects, first version

In Chapter 4 we saw that active objects encapsulate an asynchronous request. However, active objects can also be used to implement simple cooperative multitasking without any asynchronous service providers. If we activate an active object and immediately complete the request, *RunL()* will be run as soon as possible by the active scheduler. In this first attempt to solve the producer/consumer problem using active objects we try to create a cooperatively scheduled system of active objects, where an unlimited number of producers and consumers communicate using a limited buffer. In this version of the solution we try to make active objects to behave as much as possible like independently running threads. We will run into some problems with the scheduling of the active objects.

Since all active objects run inside a single thread, we don't need semaphores or mutexes to control the access to a shared object. We know exactly when the control leaves an active object, so we can be sure that all add and remove operation to the buffer are singular. Thus, we don't need to encapsulate the buffer *CBuffer* class. Instead, we can use the *CCirBuf* class directly.

Similarly, we don't need semaphores *iSpaces* and *iItems* to control the access to buffer. Since we know that another object can't access the buffer before

our object is ready, the producer can simply ask the buffer if it is empty or not before adding an item to the buffer. Other objects cannot add items to the buffer between the producer asking if there is space and adding its item to it.

Let's start by examining the code for the producer active object, *CProducer*. Code fragment 13 contains the class definition of *CProducer*. Like the producer thread in the thread-based solution, *CProducer* class also gets a pointer to the controller class when it is created.

---

```
class CProducer : public CActive
{
public:
    CProducer(CController* aController);
    ~CProducer();
    void IssueRequest();
    void CreateNewData();
    void Start();
private: // from CActive
    virtual void DoCancel();
    virtual void RunL();
private:
    CController* iController;
    TData iData;
};
```

---

**Code fragment 13:** The definition of the *CProducer* class

The constructor just adds the newly created active object to the active scheduler. *Start()* creates the first data item and then calls *IssueRequest()*. The implementation of *IssueRequest()* can be seen in Code fragment 14. This function together with *RunL()* form the heart of program logic of the producer class. First the function checks if there is space for a new item in the buffer. If there is space, the data item is added to the buffer and the request is completed with error code *KErrNone*. If the buffer is full, the request is completed with error code *KErrNotReady*.

---

```
void CProducer::IssueRequest()
{
    TInt completion;
    if (iController->iBuffer->Count() < iController->iBuffer->Length())
    {
        iController->iBuffer->Add(&iData);
        completion = KErrNone;
    }
    else
    {
        completion = KErrNotReady;
    }
    SetActive();
    TRequestStatus* status = &iStatus;
```

```
User::RequestComplete(status, completion);
}
```

---

**Code fragment 14:** The implementation of producer active object's `IssueRequest` function

Code fragment 15 shows the implementation of producer's `RunL()` function. From the request's error code we can see if the data was successfully added to the buffer. If so, we create a new data item; otherwise we try to add the same item again. In both cases we immediately create a new request.

---

```
void CProducer::RunL()
{
    if (iStatus == KErrNone)
    {
        CreateNewData();
    }
    IssueRequest();
}
```

---

**Code fragment 15:** The implementation of the producer active object's `RunL` function

The implementation of consumer object is analogous to the producer. Code fragment 16 shows the implementations of `IssueRequest()` and `RunL()` functions. Here we check if there are any items in the buffer. If so, we remove one from the buffer and consume it. In `RunL()` we just create a new request. Again an error code is used to indicate whether an item was successfully consumed, although in this case the code is not used anywhere. Regardless of whether the previous try to consume an item was successful, we simply try again as soon as the request completes.

---

```
void CConsumer::IssueRequest()
{
    TData data;
    TInt completion;
    if (iController->iBuffer->Count() > 0)
    {
        iController->iBuffer->Remove(&data);
        ProcessData();
        completion = KErrNone;
    }
    else
    {
        completion = KErrNotReady;
    }
    SetActive();
    TRequestStatus* status = &iStatus;
    User::RequestComplete(status, completion);
}

void CConsumer::RunL()
{
```

```

    IssueRequest ();
}

```

---

**Code fragment 16:** The implementation of the consumer active object's `IssueRequest` and `RunL` functions

The controller class creates all the active objects and calls `Start()` function in each of them. Also, just like in the thread-based solution, it owns the circular buffer that is used for communication between producers and consumers. Finally, when all the active objects have been started, the active scheduler is started. The call to `CActiveScheduler::Start()` never returns, unless one of the active objects calls `CActiveScheduler::Stop()`.

When we try to execute this program, we quickly notice that the system is not working as intended. Debugging the code reveals that only one active object's `RunL()` function is being called repeatedly. This happens because of the way the active scheduler schedules the completion of the active objects. If there are multiple active objects ready to be called, the active scheduler chooses the one with the highest priority. However, the order of the completion of objects with same priority depends on the implementation details of the active scheduler. The active scheduler maintains a doubly linked list of all its active objects, and the place of the active object is defined when the object is added to the scheduler. [Stichbury, 2005] In our unsophisticated implementation all active objects are always ready to be executed, and all the objects have the same priority. Thus, when the active scheduler starts to look for the next eligible active object to be executed, it searches through the linked list in the priority order, and always finds the same object.

We can circumvent this problem by juggling the priorities of our active objects, thus handling the scheduling of the objects ourselves. If we start each active object with the same priority, and each time an object completes its request we lower its priority by one, we make sure that each object gets a chance to complete its request one after each other. In this way we create a cooperative round robin scheduling for our objects. We also have to make sure that the priority variable doesn't roll over, so we have to add a maintenance function to the `CController`, which raises every object's priority if the priority falls too low.

The full source code of this version of the solution is in Appendix B.

The fact that we have to constantly adjust the priorities of the objects should serve as a strong hint that this is not the proper way to use active objects. This is also an inefficient solution because the active objects are constantly run regardless of whether the producers can add items to the buffer or whether there are any items for the consumer to consume. This is different from the thread-based solution, which uses semaphores to suspend the thread when it can't access the buffer.

However, this kind of solution, where the active object issues a request and immediately completes it by itself, is useful in some situations. One example is a long-running maintenance task that can be completed in small increments. Such a task can be implemented using a low-priority active object. When the maintenance task is started, the active object is activated and the request is immediately completed. The active object's *RunL()* function is called as soon as there are no higher priority active objects to handle. At the *RunL()*, the active object performs one increment of the task and if the task is not completed, issues and completes a new request. This way the task is being run only when the thread would otherwise be idle. [Harrison, 2003]

### 5.3 Improved solution using active objects

The solution to the producer/consumer problem presented in the previous section shows that active objects can be used to create a system that closely resembles cooperatively scheduled threads. However, that solution was clearly not optimal, and we run into some scheduling problems with it.

In this section a different solution is examined. In the previous solution the active objects' asynchronous requests were always completed immediately, regardless of whether the request was successful or not. This time we reintroduce the *CBuffer* class as a wrapper to the circular buffer, but now we make the class into an asynchronous service provider.

The operations to add and remove items from the buffer are made into asynchronous operations. This means that the operation is completed only when the item is successfully added or removed from the buffer, or the active object cancels its request. In the previous solution, when the buffer was full, the producers just kept trying again and again to add a new item to the buffer.

When the add operation is turned into an asynchronous operation, if the producer tries to add an item to the buffer and the buffer is full, the producer active object is suspended until there is a free space in the buffer. Similarly, the consumers that are trying to consume items when the buffer is empty are suspended until there are items to consume.

Let's examine the implementation of the new *CBuffer* class. Code fragment 17 contains the class definition. Here we see the new *Add()* and *Remove()* operations, which take a reference to *TRequestStatus* as a parameter. That indicates that these operations are asynchronous requests. And since the class provides asynchronous services, it also needs to have a cancel function. Since the class handles several requests at the same time, *Cancel()* also takes a reference to *TRequestStatus* as a parameter. This is used to identify the request and the active object that is being cancelled.

The class also needs to keep track of all the outstanding requests. For that reason two new circular buffers are introduced. One of them contains all outstanding producer requests and the other contains all outstanding consumer requests. For each outstanding request a pointer to its request status and a pointer to the data item are saved.

---

```

class CBuffer
{
public:
    CBuffer();
    ~CBuffer();
    void ConstructL(TInt aNumProducers, TInt aNumConsumers);
    void Add(TData* aData, TRequestStatus& aRS);
    void Remove(TData* aData, TRequestStatus& aRS);
    void Cancel(TRequestStatus& aRS);
private:
    void CompleteAddReq(TData* aData, TRequestStatus* aRS);
    void CompleteRemoveReq(TData* aData, TRequestStatus* aRS);

    struct TRequestData
    {
        TRequestStatus* iRS;
        TData* iData;
    };
    CCirBuf<TData>* iBuffer;
    CCirBuf<TRequestData> *iProducerRequests;
    CCirBuf<TRequestData> *iConsumerRequests;
};

```

---

**Code fragment 17:** The definition of *CBuffer* class

Next, let's have a look at the implementation of *Add()*. The implementation is shown in Code fragment 18. First, the status of the request is changed to *KRequestPending*. As long as the status is this, the active object that issued the request stays suspended. If the buffer is not full, the request can be completed immediately. If the buffer is full, the active object's request status and the data item are stored into a circular buffer containing all outstanding producer requests. In that case the request's status stays at *KRequestPending*, and the active object is suspended until the request can be completed.

---

```
void CBuffer::Add(TData* aData, TRequestStatus& aRS)
{
    aRS = KRequestPending;
    if (iBuffer->Count() < iBuffer->Length())
    {
        CompleteAddReq(aData, &aRS);
    }
    else
    {
        TRequestData rd;
        rd.iData = aData;
        rd.iRS = &aRS;
        iProducerRequests->Add(&rd);
    }
}
```

---

**Code fragment 18:** The implementation of *CBuffer::Add* function

The remove operation works similarly to the add operation. If there are items in the buffer, the operation is completed immediately. And if the buffer is empty, a pointer to the consumer object data and the request status are stored into another circular buffer, which contains all outstanding consumer requests.

The completion of add operation can be seen in Code fragment 19. First, the data is added to the circular buffer and the active object's request is completed with status *KErrNone*. Since there now is at least one item in the buffer, we'll check if there are any outstanding consume requests. If there are, one of those is completed.

---

```
void CBuffer::CompleteAddReq(TData* aData, TRequestStatus* aRS)
{
    iBuffer->Add(aData);
    User::RequestComplete(aRS, KErrNone);

    TRequestData rd;
    if (iConsumerRequests->Remove(&rd))
    {
        CompleteRemoveReq(rd.iData, rd.iRS);
    }
}
```

---



}

---

**Code fragment 19:** The implementation of `CBuffer::CompleteAddReq` function

Completing remove requests again works in a similar way to the add requests. First, an item is removed from the buffer, and if there are any outstanding add requests, one of those is completed.

Since *CBuffer* acts as an asynchronous service provider, it must also provide a cancel function for the requests. Code fragment 20 shows the implementation of the cancel function. Since *CBuffer* handles several asynchronous requests simultaneously, the cancel function needs the request status as a parameter. The status is used to identify the request to be cancelled. The function searches through all the outstanding requests for the one to be cancelled, and when it finds the correct request, it calls *User::RequestComplete()* for that request with completion code *KErrCancel*. When the active object is cancelled using its *Cancel()* function, *RunL()* is not called. That means that it is not necessary to handle *KErrCancel* completion code as a special case in the function.

---

```
void CBuffer::Cancel(TRequestStatus& aRS)
{
    TInt reqs = iProducerRequests->Count();
    while (reqs)
    {
        TRequestData rd;
        iProducerRequests->Remove(&rd);
        if (rd.iRS != &aRS)
            iProducerRequests->Add(&rd);
        else
        {
            User::RequestComplete(rd.iRS, KErrCancel);
            return;
        }
        reqs--;
    }

    reqs = iConsumerRequests->Count();
    while (reqs)
    {
        TRequestData rd;
        iConsumerRequests->Remove(&rd);
        if (rd.iRS != &aRS)
            iConsumerRequests->Add(&rd);
        else
        {
            User::RequestComplete(rd.iRS, KErrCancel);
            return;
        }
        reqs--;
    }
}
```

---

 }
 

---

**Code fragment 20:** Implementation of `CBuffer::Cancel` function

In this solution the producer and consumer active objects are simpler than in the previous case. Previously all the program logic was contained in the active object, but this time the *CBuffer* class handles most of the logic. The producer doesn't have to bother if the buffer is full or not. It simply issues the request, and when the request completes, the item has been added to the buffer. The producer object's *IssueRequest()* and *RunL()* functions are shown in Code fragment 21.

The full source code of this solution is available in appendix C.

---

```

void CProducer::IssueRequest ()
{
    iController->iBuffer->Add(&iData, iStatus);
    SetActive();
}

void CProducer::RunL ()
{
    if (iStatus == KErrNone)
    {
        CreateNewData ();
        IssueRequest ();
    }
}

```

---

**Code fragment 21:** The implementation of producer active object's *IssueRequest* and *RunL* functions

At least in theory this solution should be more efficient than the previous one. This time the active object is not completed until the request is successfully completed, and while the request is in pending state, the active object is suspended.

In the next chapter we run some tests for all the three solutions and see how they perform with different amounts of producers and consumers.

## 6. Tests and analysis of the programs

In this chapter the three test programs are compared, and some tests are run for all of them. All the tests are run on Nokia 9300i mobile phone, which runs uses version 7.0 of the Symbian OS. For the timings, the internal clock of the Symbian OS is used. In this version of the operating system the clock offers 64Hz accuracy, which means that the smallest observable increment is 0.0156 seconds [Harrison, 2003].

### 6.1 Performance

There are many variables that can affect the performance of a program on a modern mobile phone. An effort is made to eliminate as many of those as possible. The setup for the tests is as follows:

- The phone is formatted, to make sure there are no installed applications running in the background.
- The phone is booted up, but the mobile phone functionality is not switched on.
- After booting the phone, it is left alone for several minutes to let it complete all the background activities related to booting up.
- The phone battery is fully charged up, but while the tests are run, the charger is not connected.
- The test applications are copied to a MMC card from computer, and the card is inserted to the phone.
- The test applications are run directly from the MMC card by opening File Manager application on the phone and navigating to the MMC directory. The test applications are run by opening the application exes from File Manager.

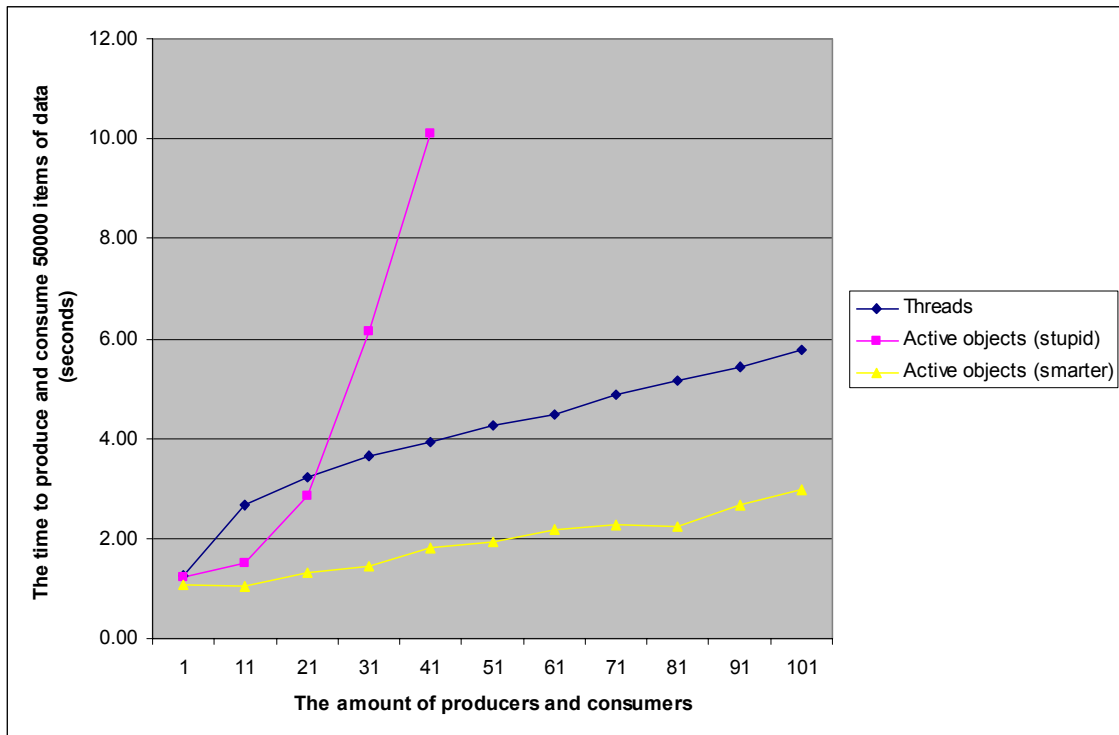
- Since it is impossible to control all the background activity on the test phone, all the tests are run at least five times and the average value of the results is used.

There are also several factors intrinsic to the test programs that might affect the results of the tests. The tests try to analyze the effects of those factors. For instance, how does the amount of the threads or active objects affect the results? Does the increase in the amount of the producers and consumers cause overhead that slows down the application? The first test tries to answer that question.

In the first test the time to produce and consume 50000 items of data is measured. For all three applications the test is run with one producer and one consumer, and then the amount of producers and consumers is increased in steps of ten until we have 101 producers and 101 consumers. The actual data item creation and consuming is very simple in this application, so most of the time should be spent handling the scheduling of threads or active objects and their cooperation. The measured times also include the initialization and destruction of the threads and active objects, and also the creation of the active scheduler for the active-object-based solutions. The size of the buffer is ten items. The results of the test are presented in Figure 3.

From the test results we can clearly see that the first solution using active objects is not performing very well when the amount of active objects increases. It is quite easy to understand why this happens. The thread-based solution uses semaphores to suspend the producers when the buffer is full and the consumers when the buffer is empty. The second active-object-based solution uses the asynchronous nature of active objects for the same purpose. In the first active-object-based solution all the producers and consumers are running all the time, trying again and again to access the buffer until the operation is successful. It is also possible that the constant shuffling of the active objects' priorities causes unnecessary overhead when there are lots of active objects.

Based on these results the first active-object-based solution is discarded from the rest of the tests.



**Figure 3:** The time to produce 50000 items of data with different amount of producers and consumers.


Comparing the results of the thread-based solution and the second active-object based solution (from now on called simply the active-object-based solution), we can see that when there are only one producer and consumer, the results are about the same, but when the amount of threads increases, the thread-based solution takes at least twice as long time to complete the task than the active-object-based solution. For the thread-based solution there is a big leap from one producer and consumer to eleven producers and consumers, while from that point onward the increase in the amount of threads causes only more or less linear increase in the time required to complete the task.

On the other hand, the active-object-based solution doesn't exhibit this kind of jump from one to eleven producers and consumers. There the time to complete the task increases pretty much linearly as the amount of active objects increases.

Next, let's have a closer look at the thread-based and the active-object-based solutions. The first test included both the time to initialize the threads and active objects and the time to actual time to produce and consume the data items. In the next test the time spent in initialization is separated from the rest

of the test. The amount of items to be created is also increased from 50000 to 200000 items.

Figure 4 shows the time taken for the initialization for different amounts of threads and active objects. For the thread-based solution this time includes the creation of the threads, starting the threads, and after the required amount of data has been produced and consumed the killing of all the threads. For the active-object-based solution the time includes the creation of all the active objects, the creation of the active scheduler, activating all the active objects, and when the required amount of data has been consumed, the destroying of the active objects and cancellation of the outstanding asynchronous requests.




**Figure 4:** The time taken to initialize and destroy the threads or active objects

The results show that the thread-based solution takes much more time to initialize its threads than the active-object-based solution takes to initialize itself. The time taken to initialize 202 threads is about 2.5 seconds, while the same amount of active objects is created in about 0.1 seconds. The difference is significant. The time taken for initialization increases linearly as the number of threads or active objects increases, which is logical since there are simply more objects to create.

Figure 5 shows the time taken for the test programs to actually produce and consume 200000 items of data. For the thread-based solution the timing

starts from the moment all threads have been started and ends when the required amount of data has been consumed. For the active-object-based solution the timing starts when the active scheduler has been started and ends when the required amount of data has been consumed.



**Figure 5:** The time taken to produce and consume 200000 items of data, not including the initialization.

From these results we can see that for these amounts of producers and consumers the active-object-based solution is always faster than the thread-based solution. However, the curves are quite different, and it seems likely that if the amount of producers and consumers would still be increased, the thread-based solution would outperform the active-object-based solution. However, those amounts of threads or active objects seem quite unlikely to be reached in a program running inside a mobile phone.

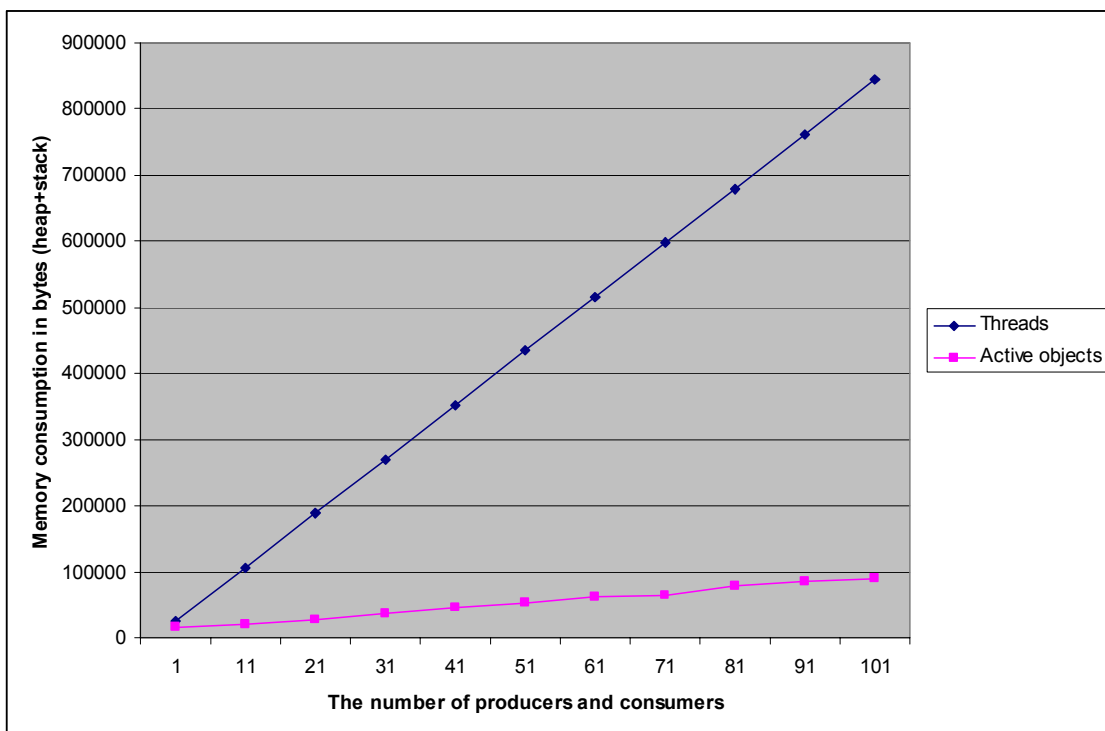
When there are only one producer and one consumer, the performances of both programs are quite similar. When the amount of producers and consumers is increased, the time taken by the thread-based solution increases rapidly. With 11 producers and consumers the thread-based solution takes more than double the time it took with one producer and consumer. After that the performance starts to stabilize, until after about 40 producers and consumers the time taken doesn't increase at all when more threads are added.

The time taken by the active-object-based solution increases more or less linearly as the amount of active objects is increased. The difference in performance between the solutions is at its most striking when there are about 10 producers and consumers each. At that point the thread-based solution takes about 2.5 times as long as the active-object-based solution to consume the required amount of data items.

## 6.2 Memory consumption

The memory consumption of the test programs has been measured using the *RThread::GetRamSizes()* function. For the thread-based solution the function is called for all threads in the program and the values are added together. For the active-object-based solution the memory usage of the single thread is reported.

Figure 6 shows the memory consumption for different amounts of threads or active objects.



**Figure 6:** The memory consumption of the test programs for different amounts of threads or active objects

The results show that the thread-based solution consumes about ten times as much memory as the active-object-based solution. This is not surprising, since each thread needs its own stack. In the test program the stack size per



thread is 4kB. If the stack size is halved from that, the program won't run any more.

### 6.3 Other considerations

The time and memory consumptions of a program are very important factors when analyzing these solutions, especially on mobile phones where these resources are much more limited than on many other platforms. However, there are also many less tangible and less easily measurable considerations that need to be taken into account when analyzing the solutions.

For software projects the ease of programming is often very important. The ease of programming is affected by both how long time it takes to implement the solution, and how error-prone the solution is. It is impossible to give any definite measurements of how easy any given solution is to implement, because it very much depends on the individual developers' experience and skill set. However, we can analyze some of the factors that affect the results.

Concurrent programming is difficult. Especially in preemptive multitasking environment all interactions between different threads of processing must be carefully analyzed. Often the access to data must be protected using synchronizing constructs such as mutexes and semaphores. One must also be careful when using those constructs to avoid the possibility of a deadlock. Since the order of execution of the threads is determined at the runtime by the scheduler, and a context switch may occur at any time, the developer must be able to analyze all possible cases.

The nonpreemptive scheduling of active objects simplifies this somewhat. With active objects the developer knows the exact locations where the execution switches from one active object to the next one. Most of the time this means that there is no need to protect the access to common data, because the developer can be sure that the active object can complete the operation before the next one accesses the same data. The developer must still be careful to avoid deadlocks, but certainly synchronizing active objects is easier than synchronizing threads.

Active objects are very widely used in Symbian OS. All applications that use the CONE GUI framework run mostly inside active objects. Since the

framework provides the active scheduler, adding own active objects to the program is quite simple. However, since those active objects run in the same thread as the code handling user interaction, extra care must be taken to ensure that all active objects complete their functions relatively quickly.

If code is ported to Symbian OS from another operating system, it probably uses threads in many cases where Symbian code would use active objects. Changing the code into using active objects might not be simple. The first attempt to solve the producer/consumer problem shows that simply thinking of active objects as cooperative threads doesn't always work. The active objects might need a different approach to the problem. In such cases it is probably easier to continue using threads also in the code ported to Symbian OS.

## 7. Summary and conclusions

In this thesis I have introduced the concept of active objects in Symbian OS. The practical use of active objects was examined by implementing a solution to the well-known producer/consumer problem using active objects. The solution was examined by comparing it to the more traditional thread-based solution, which uses semaphores for synchronizing the threads.

Active objects are widely used throughout the Symbian OS, and the documentation provided by Symbian strongly encourages their use over multithreading in applications written for Symbian OS. Comparing the thread-based and the active-object-based solutions supported this recommendation. The active-object based solution was more efficient and it consumed significantly less memory than the thread-based solution.

The results also highlighted a potential pitfall in porting thread-based program into using active objects. Using active objects efficiently requires understanding of their nature as an encapsulation of an asynchronous request. Just thinking of active objects as cooperatively scheduled threads may lead to an inefficient code, as was demonstrated by the first attempt to solve the producer/consumer problem with active objects.

There is still a lot of room for further study of active objects in Symbian OS. This thesis only examined one specific problem, and there is no proof that these solutions are the most efficient ones to the given problem. Some of the test results may be influenced by implementation details of the test programs. More tests with different types of programming problems should be run to gain a more complete understanding of active objects and their performance in comparison to threads.

All the tests were run on version 7.0 of Symbian OS. That version of the operating system uses the EKA1 kernel architecture. Symbian OS 8.0b

introduced a new kernel called EKA2, which had some changes to the way threads are handled inside the kernel. Thus, the results are only valid for older Symbian smartphones. To see how the test programs perform on latest phones using Symbian OS 8.0b or a later version, the tests need to be run also on EKA2 kernel.

## References

- [Bacon and Harris, 2003] Jean Bacon and Tim Harris, *Operating Systems: Concurrent and Distributed Software Design*, Addison-Wesley, 2003.
- [Ben-Ari, 2006] M. Ben-Ari, *Principles of Concurrent and Distributed Programming, Second Edition*, Addison-Wesley, 2006.
- [Deitel, 1984] Harvey M. Deitel, *An Introduction to Operating Systems*. Addison-Wesley, 1984.
- [Dijkstra, 1965a] E. W. Dijkstra, *Cooperating Sequential Processes*. Technical Report EWD-123, Technical University, Eindhoven, Netherlands, 1965.
- [Dijkstra, 1965b] E. W. Dijkstra, Solution of a problem in concurrent control. *Communications of ACM* **8**, 9 (Sept. 1965), 569.
- [Hansen, 1973] Peter Brinch Hansen, *Operating System Principles*. Prentice-Hall, 1973.
- [Harrison, 2003] Richard Harrison, *Symbian OS C++ for Mobile Phones*, John Wiley & Sons Ltd, 2003.
- [Mery, 2002] David Mery, Symbian white paper: Why is a different operating system needed?, Symbian Ltd, 2002. Available at <http://www.symbian.com/technology/why-diff-os.html>.
- [Morris, 2007] Ben Morris, *The Symbian OS Architecture Sourcebook – Design and Evolution of a Mobile Phone OS*, John Wiley & Sons Ltd, 2007.
- [Sales, 2005] Jane Sales, *Symbian OS Internals – Real-time Kernel Programming*, John Wiley & Sons Ltd, 2005.

[Stichbury, 2005] Jo Stichbury, *Symbian OS Explained – Effective C++ Programming For Smartphones*, John Wiley & Sons Ltd, 2005.

[Symbian] Symbian website, <http://www.symbian.com>.

[Symbian, 2002] *Symbian OS v7.0s Developer Library*. Available at [http://www.symbian.com/Developer/techlib/v70sdocs/doc\\_source/index.html](http://www.symbian.com/Developer/techlib/v70sdocs/doc_source/index.html).

[Tanenbaum, 1987] Andrew S. Tanenbaum, *Operating Systems Design and Implementation*. Prentice-Hall, 1987.

[Tasker, 1999] Martin Tasker, *Active Objects*, Symbian Ltd, 1999. Available at [http://www.symbian.com/developer/techlib/papers/tp\\_active\\_objects/active.htm](http://www.symbian.com/developer/techlib/papers/tp_active_objects/active.htm).

[Tasker, 2000] Martin Tasker, *Professional Symbian Programming*. Wrox, 2000.

## Appendix A: Full source of the thread-based solution

```

1  #include <e32base.h>
2  #include <e32cons.h>
3  #include <e32svr.h>
4
5  LOCAL_C void doTestsL();
6
7  _LIT(KThreadProducerName, "producer");
8  _LIT(KThreadConsumerName, "consumer");
9
10 const TInt KBufferSize = 10;
11 const TInt KEndCondition = 200000;
12
13 struct TData
14 {
15     public:
16         TInt iDataNumber;
17         TBuf<32> iMessage;
18     };
19
20 class CBuffer
21 {
22     public:
23         CBuffer();
24         ~CBuffer();
25         void ConstructL();
26         void AddInfo(const TData *aData);
27         void GetInfo(TData *aData);
28
29     private:
30         RMutex iGuard;
31         CCirBuf<TData>* iBuffer;
32     };
33
34 class CController
35 {
36     public:
37         CController(TInt aNumProducers, TInt aNumConsumers);
38         ~CController();
39         void ConstructL();
40         void Start();
41         void Stop();
42
43     public: // data
44         CBuffer *iBuffer;
45         RSemaphore iSpaces;
46         RSemaphore iItems;
47         RSemaphore iObjectsConsumed;
48         RSemaphore iEndCondition;
49
50     private:
51         RArray<RThread> iProducers;
52         RArray<RThread> iConsumers;
53         TInt iNumProducers;
54         TInt iNumConsumers;
55     };
56
57 // producer thread
58
59 void ProduceDataItem(TData& aData)
60 {
61     aData.iDataNumber = 1;
62     aData.iMessage = _L("producer is producing!");
63 }
64
65 TInt ProducerMain(TAny* aPtr)
66 {
67     CController *controller = (CController*) aPtr;

```

```

68     FOREVER
69     {
70         TData data;
71         ProduceDataItem(data);
72         controller->iSpaces.Wait();
73         controller->iBuffer->AddInfo(&data);
74         controller->iItems.Signal();
75     }
76     return KErrNone;
77 }
78
79 // consumer thread
80
81 void ConsumeDataItem(TData& /*aData*/)
82 {
83 }
84
85 TInt ConsumerMain(TAny* aPtr)
86 {
87     CController *controller = (CController*) aPtr;
88     FOREVER
89     {
90         TData data;
91         controller->iItems.Wait();
92         controller->iBuffer->GetInfo(&data);
93         controller->iSpaces.Signal();
94         ConsumeDataItem(data);
95         controller->iObjectsConsumed.Signal();
96         if (controller->iObjectsConsumed.Count() >= KEndCondition)
97             controller->iEndCondition.Signal();
98     }
99     return KErrNone;
100 }
101
102 // CBuffer
103
104 CBuffer::CBuffer()
105 {
106     iGuard.CreateLocal();
107 }
108
109 CBuffer::~CBuffer()
110 {
111     delete iBuffer;
112     iGuard.Close();
113 }
114
115 void CBuffer::ConstructL()
116 {
117     iBuffer = new (ELeave) CCirBuf<TData>;
118     iBuffer->SetLengthL(KBufferSize);
119 }
120
121 void CBuffer::AddInfo(const TData *aData)
122 {
123     iGuard.Wait();
124     iBuffer->Add(aData);
125     iGuard.Signal();
126 }
127
128 void CBuffer::GetInfo(TData *aData)
129 {
130     iGuard.Wait();
131     iBuffer->Remove(aData);
132     iGuard.Signal();
133 }
134
135 // CController
136
137 CController::CController(TInt aNumProducers, TInt aNumConsumers)
138 : iNumProducers(aNumProducers), iNumConsumers(aNumConsumers)

```



```

139     {
140     }
141
142 CController::~CController()
143     {
144     // close handles to threads
145     TInt i;
146     for (i = 0; i < iNumProducers; i++)
147     {
148         iProducers[i].Close();
149     }
150
151     for (i = 0; i < iNumConsumers; i++)
152     {
153         iConsumers[i].Close();
154     }
155
156     // delete objects and close all other handles
157     delete iBuffer;
158     iItems.Close();
159     iSpaces.Close();
160     iObjectsConsumed.Close();
161     iEndCondition.Close();
162     iProducers.Close();
163     iConsumers.Close();
164     }
165
166 void CController::ConstructL()
167     {
168     // create semaphores and CBuffer
169     iSpaces.CreateLocal(KBufferSize);
170     iItems.CreateLocal(0);
171     iObjectsConsumed.CreateLocal(0);
172     iEndCondition.CreateLocal(0);
173     iBuffer = new (ELeave) CBuffer;
174     iBuffer->ConstructL();
175
176     // create threads
177     TInt i;
178     for (i = 0; i < iNumProducers; i++)
179     {
180         TName threadName(KThreadProducerName);
181         threadName.AppendFormat(_L("%d"), i+1);
182         RThread t;
183         TInt err = t.Create(threadName, ProducerMain, 2048, NULL, this);
184         if (err) User::Panic(threadName, err);
185         iProducers.Append(t);
186     }
187     for (i = 0; i < iNumConsumers; i++)
188     {
189         TName threadName(KThreadConsumerName);
190         threadName.AppendFormat(_L("%d"), i+1);
191         RThread t;
192         TInt err = t.Create(threadName, ConsumerMain, 2048, NULL, this);
193         if (err) User::Panic(threadName, err);
194         iConsumers.Append(t);
195     }
196     }
197
198 void CController::Start()
199     {
200     TInt i;
201     for (i = 0; i < iNumProducers; i++)
202     {
203         iProducers[i].Resume();
204     }
205     for (i = 0; i < iNumConsumers; i++)
206     {
207         iConsumers[i].Resume();
208     }
209     }

```

```

210
211 void CController::Stop()
212     {
213     TInt i;
214     for (i = 0; i < iNumProducers; i++)
215     {
216         iProducers[i].Kill(KErrNone);
217     }
218     for (i = 0; i < iNumConsumers; i++)
219     {
220         iConsumers[i].Kill(KErrNone);
221     }
222     }
223
224 // run the test
225 LOCAL_C void doTestsL()
226     {
227     CConsoleBase *console =
228         Console::NewL(_L("Tests"), TSize(KConsFullScreen, KConsFullScreen));
229     console->Printf(_L("press any key to start.\n"));
230     console->Getch(); // get and ignore character
231
232     TInt iterations;
233     for (iterations = 1; iterations <= 111; iterations += 10)
234     {
235         TTime startTimeIni;
236         startTimeIni.UniversalTime();
237         TInt numProducers = iterations;
238         TInt numConsumers = iterations;
239         console->Printf(_L("Running with %d prod. and %d cons. threads.\n"),
240             numProducers, numConsumers);
241
242         CController* controller =
243             new (ELeave) CController(numProducers, numConsumers);
244         CleanupStack::PushL(controller);
245         controller->ConstructL();
246         controller->Start();
247
248         TTime startTime;
249         startTime.UniversalTime();
250         controller->iEndCondition.Wait();
251         TTime endTime;
252         endTime.UniversalTime();
253
254         controller->Stop();
255
256         delete controller;
257         CleanupStack::Pop(controller);
258
259         TTime endTimeIni;
260         endTimeIni.UniversalTime();
261
262         TInt fullTime =
263             endTimeIni.MicroSecondsFrom(startTimeIni).Int64().GetTInt();
264         TInt runTime = endTime.MicroSecondsFrom(startTime).Int64().GetTInt();
265         console->Printf(_L("Completed in %d + %d microseconds.\n\n"),
266             fullTime-runTime, runTime);
267     }
268     console->Printf(_L("THE END. press any key to exit.\n"));
269     console->Getch(); // get and ignore character
270     delete console;
271     }
272
273 GLDEF_C TInt E32Main() // main function called by E32
274     {
275     __UHEAP_MARK;
276     CTrapCleanup* cleanup=CTrapCleanup::New();
277     TRAPD(error, doTestsL());
278     __ASSERT_ALWAYS(!error, User::Panic(_L("doTestsL"), error));
279     delete cleanup;
280     __UHEAP_MARKEND;

```

```
281     return 0;  
282 }
```

## Appendix B: Full source of the active-object-based solution, first version

```

1  #include <e32base.h>
2  #include <e32cons.h>
3  #include <e32svr.h>
4
5  LOCAL_C void doTestsL();
6
7  _LIT(KThreadProducerName, "producer");
8  _LIT(KThreadConsumerName, "consumer");
9
10 const TInt KBufferSize = 10;
11 const TInt KEndCondition = 50000;
12
13 class CProducer;
14 class CConsumer;
15
16 struct TData
17 {
18     public:
19         TInt iDataNumber;
20         TBuf<32> iMessage;
21     };
22
23 class CController
24 {
25     public:
26         CController(TInt aNumProducers, TInt aNumConsumers);
27         ~CController();
28         void ConstructL();
29         void ControllerMainL();
30         void Start();
31         void Stop();
32         void ResetPriorities();
33
34     public: // data
35         CCirBuf<TData>* iBuffer;
36         TInt iConsumeCount;
37
38     private:
39         RPointerArray<CProducer> iProducers;
40         RPointerArray<CConsumer> iConsumers;
41         TInt iNumProducers;
42         TInt iNumConsumers;
43     };
44
45 class CProducer : public CActive
46 {
47     public:
48         CProducer(CController* aController, TName aName);
49         void ConstructL();
50         ~CProducer();
51         void IssueRequest();
52         void CreateNewData();
53         void Start();
54     private:
55         virtual void DoCancel();
56         virtual void RunL();
57     private:
58         CController* iController;
59         TData iData;
60         TInt iCount;
61         TName iName;
62     };
63
64 class CConsumer : public CActive
65 {

```

```

66     public:
67         CConsumer(CController* aController, TName aName);
68         void ConstructL();
69         ~CConsumer();
70         void Start();
71         void IssueRequest();
72         void ProcessData();
73     private:
74         virtual void DoCancel();
75         virtual void RunL();
76     private:
77         CController* iController;
78         TName iName;
79     };
80
81 // CProducer
82
83 CProducer::CProducer(CController* aController, TName aName)
84     : CActive(0),
85       iController(aController),
86       iName(aName)
87     {
88     CActiveScheduler::Add(this);
89     }
90
91 void CProducer::ConstructL()
92     {
93     }
94
95 CProducer::~CProducer()
96     {
97     Cancel();
98     }
99
100 void CProducer::Start()
101     {
102     CreateNewData();
103     IssueRequest();
104     }
105
106 void CProducer::CreateNewData()
107     {
108     iData.iDataNumber = 1;
109     iData.iMessage = _L("producer is producing!");
110     }
111
112 void CProducer::IssueRequest()
113     {
114     TInt completion;
115     if (iController->iBuffer->Count() < iController->iBuffer->Length())
116     {
117         iController->iBuffer->Add(&iData);
118         completion = KErrNone;
119     }
120     else
121     {
122         completion = KErrNotReady;
123     }
124     SetActive();
125     TRequestStatus* status = &iStatus;
126     User::RequestComplete(status, completion);
127     }
128
129 void CProducer::RunL()
130     {
131     if (Priority() > KMinTInt)
132         SetPriority(Priority() - 1);
133     else
134         iController->ResetPriorities();
135     if (iStatus == KErrNone)
136     {

```

```

137         CreateNewData();
138     }
139     IssueRequest();
140 }
141
142 void CProducer::DoCancel()
143 {
144 }
145
146 // CConsumer
147
148 CConsumer::CConsumer(CController* aController, TName aName)
149     : CActive(0),
150     iController(aController),
151     iName(aName)
152 {
153     CActiveScheduler::Add(this);
154 }
155
156 void CConsumer::ConstructL()
157 {
158 }
159
160 CConsumer::~CConsumer()
161 {
162     Cancel();
163 }
164
165 void CConsumer::Start()
166 {
167     IssueRequest();
168 }
169
170 void CConsumer::ProcessData()
171 {
172     iController->iConsumeCount++;
173     if (iController->iConsumeCount > KEndCondition)
174         iController->Stop();
175 }
176
177 void CConsumer::IssueRequest()
178 {
179     TData data;
180     TInt completion;
181     if (iController->iBuffer->Count() > 0)
182     {
183         iController->iBuffer->Remove(&data);
184         ProcessData();
185         completion = KErrNone;
186     }
187     else
188     {
189         completion = KErrNotReady;
190     }
191     SetActive();
192     TRequestStatus* status = &iStatus;
193     User::RequestComplete(status, completion);
194 }
195
196 void CConsumer::RunL()
197 {
198     if (Priority() > KMinTInt)
199         SetPriority(Priority() - 1);
200     else
201         iController->ResetPriorities();
202
203     IssueRequest();
204 }
205
206 void CConsumer::DoCancel()
207 {

```

```

208     }
209
210 // CController
211
212 CController::CController(TInt aNumProducers, TInt aNumConsumers)
213 : iConsumeCount(0), iNumProducers(aNumProducers), iNumConsumers(aNumConsumers)
214 {
215     }
216
217 CController::~CController()
218 {
219     TInt i;
220     for (i = 0; i < iNumProducers; i++)
221     {
222         delete iProducers[i];
223     }
224
225     for (i = 0; i < iNumConsumers; i++)
226     {
227         delete iConsumers[i];
228     }
229     iProducers.Close();
230     iConsumers.Close();
231     delete iBuffer;
232 }
233
234 void CController::ConstructL()
235 {
236     iBuffer = new (ELeave) CCirBuf<TData>;
237     iBuffer->SetLengthL(KBufferSize);
238
239     TInt i;
240     for (i = 0; i < iNumProducers; i++)
241     {
242         TName name(KThreadProducerName);
243         name.AppendFormat(_L("%d"), i+1);
244         CProducer* p = new (ELeave) CProducer(this, name);
245         iProducers.Append(p);
246     }
247
248     for (i = 0; i < iNumConsumers; i++)
249     {
250         TName name(KThreadConsumerName);
251         name.AppendFormat(_L("%d"), i+1);
252         CConsumer* c = new (ELeave) CConsumer(this, name);
253         iConsumers.Append(c);
254     }
255 }
256
257 void CController::Start()
258 {
259     TInt i;
260     for (i = 0; i < iNumProducers; i++)
261     {
262         iProducers[i]->Start();
263     }
264
265     for (i = 0; i < iNumConsumers; i++)
266     {
267         iConsumers[i]->Start();
268     }
269 }
270
271 void CController::Stop()
272 {
273     CActiveScheduler::Stop();
274 }
275
276 void CController::ResetPriorities()
277 {
278     TInt i;

```

```

279     for (i = 0; i < iNumProducers; i++)
280     {
281         iProducers[i]->SetPriority(0);
282     }
283
284     for (i = 0; i < iNumConsumers; i++)
285     {
286         iConsumers[i]->SetPriority(0);
287     }
288 }
289
290 // run the test
291 LOCAL_C void doTestsL()
292 {
293     CConsoleBase *console =
294         Console::NewL(_L("Tests"), TSize(KConsFullScreen, KConsFullScreen));
295     console->Printf(_L("press any key to start.\n"));
296     console->Getch(); // get and ignore character
297
298     TInt iterations;
299     for (iterations = 1; iterations <= 111; iterations += 10)
300     {
301         TTime startTime;
302         startTime.UniversalTime();
303         TInt numProducers = iterations;
304         TInt numConsumers = iterations;
305         console->Printf(_L("Running with %d prod. and %d cons. threads.\n"),
306             numProducers, numConsumers);
307
308         CActiveScheduler* activeScheduler = new (ELeave) CActiveScheduler;
309         CleanupStack::PushL(activeScheduler) ;
310         CActiveScheduler::Install(activeScheduler);
311
312         CController* controller =
313             new (ELeave) CController(numProducers, numConsumers);
314         CleanupStack::PushL(controller);
315         controller->ConstructL();
316         controller->Start();
317
318         // Start handling requests
319         CActiveScheduler::Start();
320
321         delete controller;
322         CleanupStack::Pop(controller);
323         CleanupStack::PopAndDestroy(activeScheduler);
324         TTime endTime;
325         endTime.UniversalTime();
326
327         console->Printf(_L("%d objects consumed in %d microseconds.\n\n"),
328             KEndCondition, endTime.MicroSecondsFrom(startTime));
329     }
330     console->Printf(_L("THE END. press any key to exit.\n"));
331     console->Getch(); // get and ignore character
332     delete console;
333 }
334
335 GLDEF_C TInt E32Main() // main function called by E32
336 {
337     __UHEAP_MARK;
338     CTrapCleanup* cleanup=CTrapCleanup::New();
339     TRAPD(error, doTestsL());
340     __ASSERT_ALWAYS(!error, User::Panic(_L("doTestsL"), error));
341     delete cleanup;
342     __UHEAP_MARKEND;
343     return 0;
344 }

```



## Appendix C: Full source of improved active-object-based solution

```

1  #include <e32base.h>
2  #include <e32cons.h>
3  #include <e32svr.h>
4
5  LOCAL_C void doTestsL();
6
7  _LIT(KThreadProducerName, "producer");
8  _LIT(KThreadConsumerName, "consumer");
9
10 const TInt KBufferSize = 10;
11 const TInt KEndCondition = 50000;
12
13 class CProducer;
14 class CConsumer;
15
16 struct TData
17 {
18     public:
19         TInt iDataNumber;
20         TBuf<32> iMessage;
21     };
22
23 class CBuffer
24 {
25     public:
26         CBuffer();
27         ~CBuffer();
28         void ConstructL(TInt aNumProducers, TInt aNumConsumers);
29         void Add(TData* aData, TRequestStatus& aRS);
30         void Remove(TData* aData, TRequestStatus& aRS);
31         void Cancel(TRequestStatus& aRS);
32     private:
33         void CompleteAddReq(TData* aData, TRequestStatus* aRS);
34         void CompleteRemoveReq(TData* aData, TRequestStatus* aRS);
35
36         struct TRequestData
37         {
38             TRequestStatus* iRS;
39             TData* iData;
40         };
41         CCirBuf<TData>* iBuffer;
42         CCirBuf<TRequestData> *iProducerRequests;
43         CCirBuf<TRequestData> *iConsumerRequests;
44     };
45
46 class CController
47 {
48     public:
49         CController(TInt aNumProducers, TInt aNumConsumers);
50         ~CController();
51         void ConstructL();
52         void ControllerMainL();
53         void Start();
54         void Stop();
55
56     public: // data
57         CBuffer* iBuffer;
58         TInt iConsumeCount;
59         TInt iHeap;
60         TInt iStack;
61         TTime iEndTime;
62
63     private:
64         RPointerArray<CProducer> iProducers;
65         RPointerArray<CConsumer> iConsumers;
66         TInt iNumProducers;
67         TInt iNumConsumers;

```

```

68     };
69
70 class CProducer : public CActive
71 {
72     public:
73         CProducer(CController* aController, TName aName);
74         void ConstructL();
75         ~CProducer();
76         void IssueRequest();
77         void CreateNewData();
78         void Start();
79     private:
80         virtual void DoCancel();
81         virtual void RunL();
82     private:
83         CController* iController;
84         TData iData;
85         TInt iCount;
86         TName iName;
87 };
88
89 class CConsumer : public CActive
90 {
91     public:
92         CConsumer(CController* aController, TName aName);
93         void ConstructL();
94         ~CConsumer();
95         void Start();
96         void IssueRequest();
97         TBool ProcessData();
98         TData iData;
99     private:
100        virtual void DoCancel();
101        virtual void RunL();
102    private:
103        CController* iController;
104        TName iName;
105 };
106
107 // CProducer
108
109 CProducer::CProducer(CController* aController, TName aName)
110     : CActive(0),
111       iController(aController),
112       iName(aName)
113 {
114     CActiveScheduler::Add(this);
115 }
116
117 void CProducer::ConstructL()
118 {
119 }
120
121 CProducer::~CProducer()
122 {
123     Cancel();
124 }
125
126 void CProducer::Start()
127 {
128     CreateNewData();
129     IssueRequest();
130 }
131
132 void CProducer::CreateNewData()
133 {
134     iData.iDataNumber = 1;
135     iData.iMessage = _L("producer is producing!");
136 }
137
138 void CProducer::IssueRequest()

```

```

139     {
140     iController->iBuffer->Add(&iData, iStatus);
141     SetActive();
142     }
143
144 void CProducer::RunL()
145     {
146     if (iStatus == KErrNone)
147     {
148         CreateNewData();
149         IssueRequest();
150     }
151     }
152
153 void CProducer::DoCancel()
154     {
155     iController->iBuffer->Cancel(iStatus);
156     }
157
158 // CConsumer
159
160 CConsumer::CConsumer(CController* aController, TName aName)
161     : CActive(0),
162     iController(aController),
163     iName(aName)
164     {
165     CActiveScheduler::Add(this);
166     }
167
168 void CConsumer::ConstructL()
169     {
170     }
171
172 CConsumer::~CConsumer()
173     {
174     Cancel();
175     }
176
177 void CConsumer::Start()
178     {
179     IssueRequest();
180     }
181
182 TBool CConsumer::ProcessData()
183     {
184     iController->iConsumeCount++;
185     if (iController->iConsumeCount > KEndCondition)
186     {
187         iController->Stop();
188         return ETrue;
189     }
190     return EFalse;
191     }
192
193 void CConsumer::IssueRequest()
194     {
195     iController->iBuffer->Remove(&iData, iStatus);
196     SetActive();
197     }
198
199 void CConsumer::RunL()
200     {
201     if (iStatus == KErrNone)
202     {
203         if (!ProcessData())
204             IssueRequest();
205     }
206     }
207
208 void CConsumer::DoCancel()
209     {

```

```

210     iController->iBuffer->Cancel(iStatus);
211     }
212
213 // CCBuffer
214
215 CBuffer::CBuffer() {}
216
217 CBuffer::~CBuffer()
218 {
219     delete iBuffer;
220     delete iProducerRequests;
221     delete iConsumerRequests;
222 }
223
224 void CBuffer::ConstructL(TInt aNumProducers, TInt aNumConsumers)
225 {
226     iBuffer = new (ELeave) CCirBuf<TData>;
227     iBuffer->SetLengthL(KBufferSize);
228
229     iProducerRequests = new (ELeave) CCirBuf<TRequestData>;
230     iProducerRequests->SetLengthL(aNumProducers);
231
232     iConsumerRequests = new (ELeave) CCirBuf<TRequestData>;
233     iConsumerRequests->SetLengthL(aNumConsumers);
234 }
235
236 void CBuffer::Add(TData* aData, TRequestStatus& aRS)
237 {
238     aRS = KRequestPending;
239     if (iBuffer->Count() < iBuffer->Length())
240     {
241         CompleteAddReq(aData, &aRS);
242     }
243     else
244     {
245         TRequestData rd;
246         rd.iData = aData;
247         rd.iRS = &aRS;
248         iProducerRequests->Add(&rd);
249     }
250 }
251
252 void CBuffer::Remove(TData* aData, TRequestStatus& aRS)
253 {
254     aRS = KRequestPending;
255     if (iBuffer->Count() > 0)
256     {
257         CompleteRemoveReq(aData, &aRS);
258     }
259     else
260     {
261         TRequestData rd;
262         rd.iData = aData;
263         rd.iRS = &aRS;
264         iConsumerRequests->Add(&rd);
265     }
266 }
267
268 void CBuffer::CompleteAddReq(TData* aData, TRequestStatus* aRS)
269 {
270     iBuffer->Add(aData);
271     User::RequestComplete(aRS, KErrNone);
272
273     TRequestData rd;
274     if (iConsumerRequests->Remove(&rd))
275     {
276         CompleteRemoveReq(rd.iData, rd.iRS);
277     }
278 }
279
280 void CBuffer::CompleteRemoveReq(TData* aData, TRequestStatus* aRS)

```

```

281     {
282     iBuffer->Remove(aData);
283     User::RequestComplete(aRS, KErrNone);
284
285     TRequestData rd;
286     if (iProducerRequests->Remove(&rd))
287     {
288         CompleteAddReq(rd.iData, rd.iRS);
289     }
290     }
291
292 void CBuffer::Cancel(TRequestStatus& aRS)
293 {
294     TInt reqs = iProducerRequests->Count();
295     while (reqs)
296     {
297         TRequestData rd;
298         iProducerRequests->Remove(&rd);
299         if (rd.iRS != &aRS)
300             iProducerRequests->Add(&rd);
301         else
302         {
303             User::RequestComplete(rd.iRS, KErrCancel);
304             return;
305         }
306         reqs--;
307     }
308     reqs = iConsumerRequests->Count();
309     while (reqs)
310     {
311         TRequestData rd;
312         iConsumerRequests->Remove(&rd);
313         if (rd.iRS != &aRS)
314             iConsumerRequests->Add(&rd);
315         else
316         {
317             User::RequestComplete(rd.iRS, KErrCancel);
318             return;
319         }
320         reqs--;
321     }
322     }
323
324 // CController
325
326 CController::CController(TInt aNumProducers, TInt aNumConsumers)
327 : iConsumeCount(0), iNumProducers(aNumProducers), iNumConsumers(aNumConsumers)
328 {
329     }
330
331 CController::~CController()
332 {
333     iProducers.Close();
334     iConsumers.Close();
335     delete iBuffer;
336     }
337
338 void CController::ConstructL()
339 {
340     iBuffer = new (ELeave) CBuffer;
341     iBuffer->ConstructL(iNumProducers, iNumConsumers);
342
343     TInt i;
344     for (i = 0; i < iNumProducers; i++)
345     {
346         TName name(KThreadProducerName);
347         name.AppendFormat(_L("%d"), i+1);
348         CProducer* p = new (ELeave) CProducer(this, name);
349         iProducers.Append(p);
350     }
351

```

```

352     for (i = 0; i < iNumConsumers; i++)
353     {
354         TName name(KThreadConsumerName);
355         name.AppendFormat(_L("%d"), i+1);
356         CConsumer* c = new (ELeave) CConsumer(this, name);
357         iConsumers.Append(c);
358     }
359 }
360
361 void CController::Start()
362 {
363     TInt i;
364     for (i = 0; i < iNumProducers; i++)
365     {
366         iProducers[i]->Start();
367     }
368
369     for (i = 0; i < iNumConsumers; i++)
370     {
371         iConsumers[i]->Start();
372     }
373 }
374
375 void CController::Stop()
376 {
377     iEndTime.UniversalTime();
378
379     TInt i;
380     for (i = 0; i < iNumProducers; i++)
381     {
382         delete iProducers[i];
383     }
384
385     for (i = 0; i < iNumConsumers; i++)
386     {
387         delete iConsumers[i];
388     }
389
390     CActiveScheduler::Stop();
391 }
392
393 // run the test
394 LOCAL_C void doTestsL()
395 {
396     CConsoleBase *console =
397         Console::NewL(_L("Tests"), TSize(KConsFullScreen, KConsFullScreen));
398     console->Printf(_L("press any key to start.\n"));
399     console->Getch(); // get and ignore character
400
401     TInt iterations;
402     for (iterations = 1; iterations <= 111; iterations += 10)
403     {
404         TTime startTimeIni;
405         startTimeIni.UniversalTime();
406         TInt numProducers = iterations;
407         TInt numConsumers = iterations;
408         console->Printf(_L("Running with %d prod. and %d cons. threads.\n"),
409             numProducers, numConsumers);
410
411         CActiveScheduler* activeScheduler = new (ELeave) CActiveScheduler;
412         CleanupStack::PushL(activeScheduler) ;
413         CActiveScheduler::Install(activeScheduler);
414
415         CController* controller =
416             new (ELeave) CController(numProducers, numConsumers);
417         CleanupStack::PushL(controller);
418         controller->ConstructL();
419         controller->Start();
420
421         // Start handling requests
422         TTime startTime;

```

```

423     startTime.UniversalTime();
424     CActiveScheduler::Start();
425     TTime endTime = controller->iEndTime;
426
427     delete controller;
428     CleanupStack::Pop(controller);
429     CleanupStack::PopAndDestroy(activeScheduler);
430
431     TTime endTimeIni;
432     endTimeIni.UniversalTime();
433
434     TInt fullTime =
435         endTimeIni.MicroSecondsFrom(startTimeIni).Int64().GetTInt();
436     TInt runTime = endTime.MicroSecondsFrom(startTime).Int64().GetTInt();
437     console->Printf(_L("Completed in %d + %d microseconds.\n\n"),
438         fullTime-runTime, runTime);
439 }
440 console->Printf(_L("THE END. press any key to exit.\n"));
441 console->Getch(); // get and ignore character
442 delete console;
443 }
444
445 GLDEF_C TInt E32Main() // main function called by E32
446 {
447     __UHEAP_MARK;
448     CTrapCleanup* cleanup=CTrapCleanup::New();
449     TRAPD(error,doTestsL());
450     __ASSERT_ALWAYS(!error,User::Panic(_L("doTestsL"),error));
451     delete cleanup;
452     __UHEAP_MARKEND;
453     return 0;
454 }
455

```