

Object Lifecycle Manager

A Complementary Pattern for Controlling Object Creation and Destruction

David L. Levine, Christopher D. Gill, and Douglas C. Schmidt

{levine,cdgill,schmidt}@cs.wustl.edu

Department of Computer Science, Washington University

St. Louis, MO 63130, USA *

July 21, 1999

To appear at the Pattern Languages of Programming Conference, Allerton Park, Illinois, USA, 15 – 18 August 1999.

Abstract

Creational patterns, such as Singleton and Factory Method [1], address object construction and initialization, but do not consider object destruction. In some applications, however, object destruction is as important as object construction. The *Object Lifecycle Manager* pattern addresses issues associated with object destruction. Object Lifecycle Manager is also an example of a *complementary pattern*, which completes or extends other patterns. In particular, the Object Lifecycle Manager pattern completes creational patterns by considering the entire lifecycle of objects.

1 Intent

The *Object Lifecycle Manager* pattern can be used to govern the entire lifecycle of objects, from creating them prior to their first use to ensuring they are destroyed properly at program termination. In addition, this pattern can be used to replace static object creation (and destruction) with dynamic object preallocation (and deallocation) that occurs automatically during application initialization (and termination).

2 Example

Singleton [1] is a common creational pattern that provides a global point of access to a unique class instance and defers

creation of the instance until it is first accessed. If a singleton is not needed during the lifetime of a program, it will not be created. The Singleton pattern does not address the issue of when its instance is destroyed, however, which is problematic for certain types of applications and operating systems.

To illustrate why it is important to address destruction semantics, consider the following logging component that provides a client programming API to a distributed logging service [2]. Applications use the logging component as a front-end to the distributed logging service to report errors and generate debugging traces.

```
class Logger
{
public:
    // Global access point to Logger singleton.
    static Logger *instance (void) {
        if (instance_ == 0)
            instance_ = new Logger;
        return instance_;
    }

    // Write some information to the log.
    int log (const char *format, ...);

protected:
    // Default constructor (protected to
    // ensure Singleton pattern usage).
    Logger (void);

    static Logger *instance_;
    // Contained Logger singleton instance.

    // ... other resources that are
    // held by the instance ...
};

// Initialize the instance pointer.
Logger *Logger::instance_ = 0;
```

The `Logger` constructor, which is omitted for brevity, allocates various OS endsystem resources, such as socket handles, shared memory segments, and/or system-wide semaphores, that are used to implement the logging service.

*This work was supported in part by Boeing, NSF grant NCR-9628218, DARPA contract 9701516, and Nortel.
Copyright ©1999, David L. Levine, Christopher D. Gill, and Douglas C. Schmidt. Permission is granted to copy for the PLoP 1999 conference. All other rights reserved.

To reduce the size and improve the readability of its logging records, an application may choose to log certain data, such as timing statistics, in “batch” form rather than individually. For instance, the following statistics class batches timing data for individual identifiers:

```
class Stats
{
public:
    // Global access point to the
    // statistics singleton.
    static Stats *instance (void) {
        if (instance_ == 0)
            instance_ = new Stats;
        return instance_;
    }

    // Record a timing data point.
    int record (int id,
                const timeval &tv);

    // Report recorded statistics
    // to the log.
    void report (int id) {
        Logger::instance ()->
            log ("Avg timing %d: "
                "%ld sec %ld usec\n",
                id,
                average_i (id).tv_sec,
                average_i (id).tv_usec);
    }

protected:
    // Default constructor.
    Stats (void);

    // Internal accessor for an average.
    const timeval &average_i (void);

    // Contained Stats singleton instance.
    static Stats *instance_;

    // ... other resources that are
    // held by the instance ...
};

// Initialize the instance pointer.
Stats *Stats::instance_ = 0;
```

When a user-defined threshold is reached, it uses the `Logger` singleton to report average timing statistics for an identifier.

Both the `Logger` and `Stats` classes provide distinct services to the application: the `Logger` class provides general logging capabilities, whereas the `Stats` class provides specialized batching and logging of time statistics. These classes are designed using the Singleton pattern, so that a single instance of each is used in an application process.

The following example illustrates how an application might use the `Logger` and `Stats` singletons.

```
int main (int argc, char *argv[])
{
```

```
    // Interval timestamps.
    timeval start_tv, stop_tv;

    // Logger, Stats singletons
    // do not yet exist.

    // Logger and Stats singletons created
    // during the first iteration.
    for (int i = 0; i < argc; ++i) {
        ::gettimeofday (&start_tv);
        // do some work between timestamps ...
        ::gettimeofday (&stop_tv);
        // then record the stats ...
        stop_tv.sec -= start_tv.sec;
        stop_tv.usec -= start_tv.usec;
        Stats::instance ()->record (i, stop_tv);

        // ... and log some output.
        Logger::instance ()->
            log ("Arg %d [%s]\n", i, argv[i]);
    }

    // Logger and Stats singletons are not
    // cleaned up when main returns.
    return 0;
}
```

Note that the `Logger` and `Stats` singletons are not constructed or destroyed explicitly by the application, *i.e.*, their lifecycle management is decoupled from the application logic. It is common practice to not destroy singletons at program exit [3].

Several drawbacks arise, however, from the fact that the Singleton pattern only addresses the *creation* of singleton instances and does not deal with their destruction. In particular, when the main program above terminates, neither the `Logger` nor the `Stats` singletons are cleaned up. At best, this can lead to false reports of leaked memory. At worst, important system resources may not be released and destroyed properly.

For instance, problems can arise if the `Logger` and/or `Stats` singletons hold OS resources, such as system-scope semaphores, I/O buffers, or other allocated OS resources. Failure to clean up these resources gracefully during program shutdown can cause deadlocks and other synchronization hazards. To alleviate this problem, each singleton’s destructor should be called.

One way of implementing the Singleton pattern that attempts to ensure singleton destruction is to declare a static instance of the class at file scope [3]. For example, the following `Singleton_Destroyer` template provides a destructor that deletes the singleton.

```
template <class T>
Singleton_Destroyer
{
public:
    Singleton_Destroyer (void): t_ (0) {}
    void register (T *) { t_ = t; }
    ~Singleton_Destroyer (void) { delete t_; }
```

```
private:
    T *t_; // Holds the instance.
};
```

To use this class, all that's necessary is to modify the `Logger` and `Stats` classes by defining a static instance of the `Singleton_Destroyer`, such as the following example for `Logger`:

```
static Singleton_Destroyer<Logger>
    logger_destroyer;

// Global access point to the
// Logger singleton.
static Logger *instance (void) {
    if (instance_ == 0) {
        instance_ = new Logger;
        // Register the singleton so it will be
        // destroyed when the destructor of
        // <logger_destroyer> is run.
        logger_destroyer.register (instance_);
    }
    return instance_;
}

// ... similar changes to Stats class ...
```

Note how `logger_destroyer` holds the singleton and deletes it when the program exits. A similar `Singleton_Destroyer` could be used by the `Stats` singleton, as well.

Unfortunately, there are several problem with explicitly instantiating static `Singleton_Destroyer` instances. In C++, for example, each `Singleton_Destroyer` could be defined in a different compilation unit. In this case, there is no guaranteed order in which their destructors will be called, which can lead to undefined program behavior. In particular, if singletons in different compilation units share resources, such as socket handles, shared memory segments, and/or system-wide semaphores, the program may fail to exit cleanly. The undefined order of singleton destruction in C++ makes it hard to ensure these resources are released by the OS before (1) the last singleton using the resource is completely destroyed, but not before (2) a singleton that has not exited uses the resource(s).

In summary, the key forces that are not resolved in these examples above are: (1) ensuring resources allocated by a singleton are ultimately released when a program exits, (2) managing the order of creation and destruction of static instances, and (3) providing a framework that encapsulates these details within a well-defined interface.

3 Context

An application or system where full control over the lifecycle of the objects it creates is necessary for correct operation.

4 Problem

Many applications do not handle the entire lifecycle of their objects properly. In particular, applications that use creational patterns, such as `Singleton`, often fail to address object destruction. Similarly, applications that use static objects to provide destruction often suffer from inconsistent initialization and termination behavior. Both of these problems are outlined below.

Problems with singleton destruction: `Singleton` instances may be created dynamically.¹ A dynamically allocated singleton instance becomes a resource leak if it is not destroyed, however. Often, these types of leaks are ignored because (1) they aren't significant in many applications and (2) on most multi-user general-purpose operating systems, such as UNIX or Windows NT, they are cleaned up when a process terminates.

Unfortunately, resource leaks can be troublesome in the following contexts:

- **When graceful shutdown is required [3]:** A singleton may be responsible for system resources, such as system-wide locks, open network connections, and shared memory segments. Explicit destruction of these singletons may be desirable to ensure these resources are released at a well defined point during program termination. For instance, if the `Logger` class in Section 2 requires system-wide locks or shared memory, the `Logger` must release these resources as soon as the `Logger` is not needed.

- **When singletons maintain references to other singletons:** Explicitly managing the order of destruction of singletons may be necessary to avoid problems due to dangling references during program termination. For example, if the `Stats` class in the example above uses the `Logger` instance in its `report` method, this method could be invoked during the `Stats` instance destruction, which renders the behavior of the program undefined. Likewise, to support useful behaviors, such as logging previously unreported values during program shutdown, the termination ordering of these singletons must be controlled.

- **When checking for memory leaks:** Memory leak detection tools, such as NuMega BoundsCheck, ParaSoft Insure++, and Rational Purify, are useful for languages, such as C and C++, that require explicit allocation and deallocation of dynamic memory. Such tools will identify singleton instances as leaked memory, which make it hard to identify the relevant memory leaks.

¹Singleton is used as an example in much of this pattern description because (1) it is a popular creational pattern and (2) it highlights challenging object destruction issues nicely. However, Object Lifecycle Manager can complement other creational patterns, such as `Factory Method`, and does not assume that its managed objects are singletons or homogeneous types.

For instance, if a large number of identifiers is used to record `Stats` data in our running example, it may appear that a sizeable amount of memory is leaking during program operation. In large-scale applications, these “leaks” can result in numerous erroneous warnings, thereby obscuring the actual memory leaks and hindering system debugging.

- **Dynamic memory allocation may be from a global pool:** Some real-time operating systems, such as VxWorks [4] or pSOS [5], have only a single, global heap for all applications. Therefore, application tasks must release dynamically allocated memory upon task termination; otherwise, it cannot be reallocated to other applications until the OS is rebooted. For instance, failure to explicitly release memory that was dynamically allocated by the `Logger` and `Stats` singletons in our running example represents actual resource leaks on such OS platforms.

Problems with static object lifecycle: Some objects must be created prior to any use. In C++, such instances traditionally have been created as *static objects*, which are intended to be constructed prior to invocation of the main program entry point and destroyed at program termination. However, there are several important drawbacks to static objects:

- **Unspecified order of construction/destruction.** C++ only specifies the order of construction/destruction of static objects *within* a compilation unit (file); the construction order matches the declaration order and destruction is the reverse order. However, there is no constraint specified on the order of construction/destruction *between* static objects in different files. Therefore, construction/destruction ordering is implementation dependent.

For example, the versions of the `Logger` and `Stats` classes that use the `Singleton Destroyer` in Section 2 illustrate the problems that arise due to the undefined order of destruction of the `Stats` and `Logger` singleton instances.

It is difficult to write portable C++ code that uses static objects possessing initialization dependencies. Often, it is simpler to avoid using static objects altogether, rather than trying to analyze for, and protect against, such dependencies. This approach is particularly appropriate for reusable components and frameworks, which should avoid unnecessary constraints on how they are used and/or initialized.

On some platforms, explicit singleton management is necessary for correct program operation because some platforms destroy singletons prematurely. For instance, the garbage collector in older Java Development Kits (JDKs) may destroy an object when there are no longer any references to it, even if the object was intended to be a singleton [6]. Though this deficiency has been fixed in later JDKs, the Object Lifecycle Manager could solve it, under application control, by maintaining singleton references.

Another problem in Java applications is the sharing of namespaces, which allow sharing (intended or otherwise) between singletons in separate applets [7]. Again, the Object Lifecycle Manager could be used to register singleton instances. Applets would then access their singletons from this registry.

- **Poor support by embedded systems.** Embedded systems have historically used C. Therefore, they do not always provide seamless support for OO programming languages features. For instance, the construction/destruction of static objects in C++ is one such feature that often complicates embedded systems programming. The embedded OS may have support for explicit invocation of static constructor/destructor calls, but this is not optimal from the programmer’s perspective.

Some embedded operating systems do not support the notion of a *program* that has a unique entry point. For example, VxWorks supports multiple *tasks*, which are similar to threads because they all share the address space. However, there is no designated *main* task for each application. Therefore, these embedded systems platforms can be configured to call static constructors/destructors at module (object file) load/unload time, respectively. On such platforms, it is not otherwise necessary to unload and load between repeated executions. To properly destroy and construct static objects, however, the static object destructors/constructors must either be called manually, or the module unloaded and loaded again, which hinders repeated testing.

In addition, placement of data in read-only memory (ROM) complicates the use of static objects [8]. The data must be placed in ROM prior to run-time; however, static constructors are called at run-time. Therefore, embedded systems sometimes do not support calls of static constructors and destructors. Moreover, if they are supported it may be under explicit application control, instead of by implicit arrangement of the compiler and run-time system.

One or more of these drawbacks of static objects typically provides sufficient motivation for removing them from a program. Often, it is better not use them in the first place, but to apply the following solution instead.

5 Solution

Define an *Object Lifecycle Manager*, which is a singleton that contains a collection of *Preallocated Objects* and *Managed Objects*. The Object Lifecycle Manager is responsible for constructing and destroying the Preallocated Objects at program initialization and termination, respectively. It is also responsible for ensuring all of its Managed Objects are destroyed properly at program termination.

6 Applicability

Use Object Lifecycle Manager when:

Singletons and other dynamically created objects must be removed without application intervention at program termination: Singleton and other creational patterns do not typically address the question of when the objects they create should be removed, or who should remove them. Object Lifecycle Manager provides a convenient, global object that deletes dynamically created objects. Creational pattern objects can then register with the Object Lifecycle Manager for deletion, which usually occurs at program termination.

Static objects must be removed from the application: As described in Section 4, static objects can be troublesome, especially in some languages and on some platforms. Object Lifecycle Manager provides a mechanism to replace static objects with Preallocated Objects. Preallocated Objects are dynamically allocated before the application uses them, and deallocated at program termination.

The platform does not support static object construction/destruction: Some platforms, such as VxWorks and pSOS, do not always construct static objects at program start and destroy them at program termination.² In general, it is best to remove all static objects, *e.g.*, to support repeated testing of a program. Another situation where static objects can cause difficulty is when they are placed in ROM. Objects in ROM cannot be initialized at run-time, because they cannot be modified at all.

The underlying platform does not provide a notion of a program, though the application needs it: The root cause for lack of support for static object construction/destruction on some platforms is their lack of support for the notion of a program, as discussed in Section 4. The Object Lifecycle Manager pattern can be used to emulate programs by partitioning this address space. The scope of each Object Lifecycle Manager delineates a program, from the application perspective.

Destruction order must be specified by the application: Dynamically created objects can be *registered* with the Object Lifecycle Manager for destruction. The Object Lifecycle Manager can be implemented to destroy objects in any desired order.

The application requires explicit singleton management: As described in Section 4, singletons may be destroyed prematurely, for example, on earlier Java platforms. The Object

²On certain platforms, such as VxWorks and pSOS, static objects can be constructed when the module is loaded and destroyed when it is unloaded. After loading, an entry point can be called more than once before unloading. Therefore, a *program* can be run more than once after constructing static objects, without ever destroying them. Conversely, static object constructors and destructors can be invoked explicitly.

Lifecycle Manager delays singleton destruction until program termination.

7 Structure and Participants

The structure and participants of the Object Lifecycle Manager pattern are shown using UML in Figure 1 and described below.

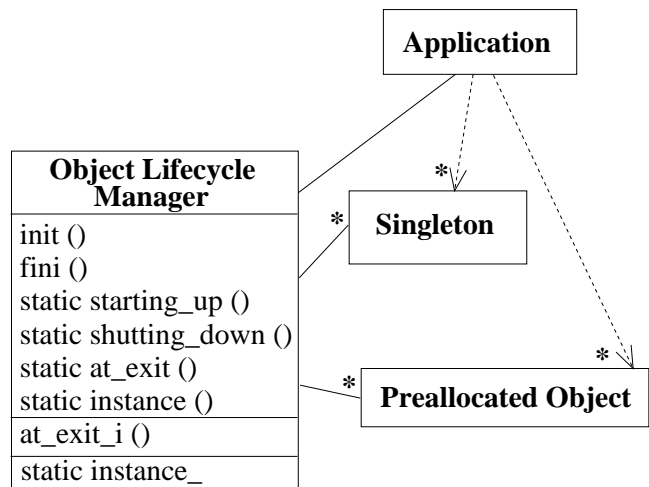


Figure 1: Object Lifecycle Manager Pattern Structure

Object Lifecycle Manager: Each Object Lifecycle Manager is a singleton that contains collections of Managed Objects and Preallocated Objects.

Managed Objects: Any object may be *registered* with an Object Lifecycle Manager, which is responsible for destroying the object. Object destruction occurs when the Object Lifecycle Manager itself is destroyed, typically at program termination.

Preallocated Objects: An object may be hard-coded for construction and destruction by an Object Lifecycle Manager. Preallocated objects have the same lifetime as the Object Lifecycle Manager, *i.e.*, the lifetime of the process that executes the program.

Application: The Application initializes and destroys its Object Lifecycle Managers, either implicitly or explicitly. In addition, the application registers its Managed Objects with an Object Lifecycle Manager, which may contain Preallocated Objects.

8 Dynamics

The dynamic collaborations among participants in the Object Lifecycle Manager pattern are shown in Figure 2. The diagram depicts four separate activities:

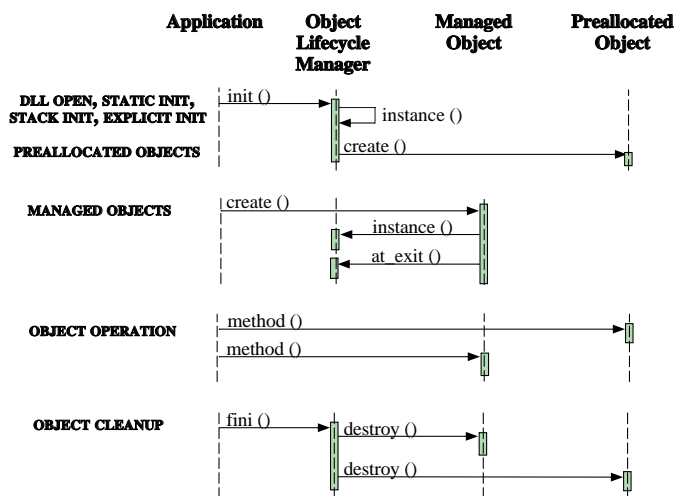


Figure 2: Object Lifecycle Manager Pattern Collaborations

1. Object Lifecycle Manager creation and initialization, which in turn creates the Preallocated Objects;
2. Managed Object creation by the Application, and registration with the Object Lifecycle Manager;
3. Use of Preallocated and Managed Objects by the Application; and
4. Destruction of the Object Lifecycle Manager, which includes destruction of all Managed and Preallocated Objects.

Within each activity, time increases down the vertical axis.

9 Implementation

The Object Lifecycle Manager pattern can be implemented using the steps presented below. This implementation is based on the Object Manager provided in the ACE framework [9], which motivates many of the interesting implementation issues in this section. ACE is written in C++, so some of the steps discussed below are language specific. Appendix A illustrates even more concretely how the Object Lifecycle Manager pattern has been implemented in ACE.

1. Define the Object Lifecycle Manager. This component provides applications with an interface with which to register objects whose lifecycle must be managed to ensure proper destruction upon program termination. In addition, this component defines a repository that ensures proper destruction of its managed objects. The Object Lifecycle Manager is a container for the Preallocated Objects and for the Managed Objects registered for destruction at termination.

The following substeps can be used to implement the Object Lifecycle Manager.

- **Migrate common interfaces and implementation details into a base class:** Factoring common internal details into an Object Lifecycle Manager Base class can make the Object Lifecycle Manager implementation simpler and more robust. Defining an Object Lifecycle Manager Base class also supports the creation of multiple Object Lifecycle Managers, each of a separate type. To simplify our discussion, we only touch on the use of multiple Object Lifecycle Managers briefly. They do not add consequences to the pattern, but are useful for partitioning libraries and applications.

- **Define a cleanup function interface:** The Object Lifecycle Manager should allow applications to register arbitrary types of objects. When a program is shut down, the Object Lifecycle Manager cleans up these objects automatically.

The following example illustrates a specialized `CLEANUP_FUNC` used in ACE to register an object or array for cleanup:

```
typedef void (*CLEANUP_FUNC)(void *object,
                             void *param);

class Object_Lifecycle_Manager
{
public:
    // ...
    static int at_exit (void *object,
                       CLEANUP_FUNC cleanup_hook,
                       void *param);
    // ...
};
```

The static `at_exit` method registers an object or array of objects for cleanup at process termination. The `cleanup_hook` argument points to a global function or static member function that is called at cleanup time to destroy the object or array. At destruction time, the Object Lifecycle Manager passes the `object` and `param` arguments to the `cleanup_hook` function. The `param` argument contains any additional information needed by the `cleanup_hook` function, such as the number of objects in the array.

Another alternative to registering a `CLEANUP_FUNC` with the object manager would be to use the C-library `atexit` function to invoke the termination functions at program exit. However, not all platforms support `atexit`. Furthermore, `atexit` implementations usually have a limit of 32 registered termination functions.

- **Define a cleanup base class interface:** This allows applications to register for destruction with the Object Lifecycle Manager any object whose class derives from

a Cleanup base class. The Cleanup base class should have a virtual destructor and a cleanup method that simply calls `delete this;`, which in turn invokes all derived class destructors.

The following code fragment illustrates the interface used for such registrations in ACE:

```
class Object_Lifecycle_Manager
{
public:
    // ...
    static int at_exit (Cleanup *object,
                       void *param = 0);
    // ...
};
```

This static `at_exit` method registers a Cleanup object for cleanup at process termination. At destruction time, the Object Lifecycle Manager calls the Cleanup object's cleanup function, passing in the `param` argument. The `param` argument contains any additional information needed by the cleanup function.

- **Define a singleton adapter:** Although it is possible to explicitly code singletons to use the Object Lifecycle Manager this is tedious and error-prone. Therefore, it is useful to define a Singleton adapter class template that encapsulates the details of creating singleton objects and registering them with the Object Lifecycle Manager. In addition, the Singleton adapter can ensure the thread-safe Double-Checked Locking pattern [10] is used to construct and access an instance of the type-specific Singleton.

The following code fragment illustrates how a singleton adapter is implemented in ACE:

```
template <class TYPE>
class Singleton : public Cleanup
{
public:
    // Global access point to the
    // wrapped singleton.
    static TYPE *instance (void) {
        // Details of Double Checked
        // Locking omitted ...
        if (singleton_ == 0)
            singleton_ = new Singleton<TYPE>;

        // Register with the Object Lifecycle
        // Manager for destruction.
        Object_Lifecycle_Manager::
            at_exit (singleton_);

        return &singleton_->instance_;
    }

protected:
    // Default constructor.
    Singleton (void);

    // Contained instance.
    TYPE instance_;

    // Instance of the singleton adapter.
```

```
static Singleton<TYPE> *singleton_;
};
```

The Singleton class template is derived from the Cleanup class. This allows the Singleton instance to register itself with the Object Lifecycle Manager. The Object Lifecycle Manager then assumes responsibility for dynamically deallocating the Singleton instance and with it the adapted TYPE instance.

- **Define a termination function interface:** Lifecycle management functionality has been discussed so far in terms of destruction of objects at program termination. However, the Object Lifecycle Manager can provide a more general capability – the ability to call a function at program termination – using the same internal implementation mechanism. For example, to ensure proper cleanup of open Win32 WinSock sockets at program termination, the `WSACleanup` function must be called. This capability is described in more detail in Appendix A.

- **Define a preallocation mechanism:** Some application resources, such as synchronization locks, must be created before their first use in order to avoid race conditions. Encapsulating their creation within the Object Lifecycle Manager's own initialization phase ensures this will occur, without adding complexity to application code. The Object Lifecycle Manager should be able to preallocate objects or arrays. The Object Lifecycle Manager can either perform these preallocations statically in global data or dynamically on the heap.

- **Determine the destruction order of registered objects:** As noted in Section 6, the Object Lifecycle Manager can be implemented to destroy registered objects in any desired order. For example, priority levels could be created, and destruction could proceed in decreasing order of priority. An interface could be provided for objects to set and change their destruction priority.

We have found that destruction in reverse order of registration has been a sufficient policy. An application can, in effect, specify destruction order by controlling the order in which it registers objects with the Object Lifecycle Manager.

2. Determine the set of Preallocated Objects. Preallocated Objects must always be created before the application's main processing begins. Thus, these objects must be hard-coded into each Object Lifecycle Manager class. An efficient implementation is to store each Preallocated Object in an array. Certain languages, such as C++, do not support arrays of heterogeneous objects, however. Therefore, in these languages pointers must be stored instead of the objects themselves. The actual objects are dynamically allocated by the Object Lifecycle Manager when it is instantiated,

and destroyed by the Object Lifecycle Manager destructor.

The following substeps should be used to implement preallocated objects:

- **Limit exposure:** To minimize the exposure of header files, identify the Preallocated Objects by enumerated literals, *e.g.*,

```
enum Preallocated_Object_ID
{
    ACE_FILECACHE_LOCK,
    ACE_STATIC_OBJECT_LOCK,
    ACE_LOG_MSG_INSTANCE_LOCK,
    ACE_DUMP_LOCK,
    ACE_SIG_HANDLER_LOCK,
    ACE_SINGLETON_NULL_LOCK,
    ACE_SINGLETON_RECURSIVE_THREAD_LOCK,
    ACE_THREAD_EXIT_LOCK,
};
```

- **Use cleanup adapters:** Next, use template functions or macros for allocation/deallocation, *e.g.*,

```
#define PREALLOCATE_OBJECT(TYPE, ID) {\
    Cleanup_Adapter<TYPE> *obj_p;\
    obj_p = new Cleanup_Adapter<TYPE>;\
    preallocated_object[ID] = obj_p;\
}

#define DELETE_PREALLOCATED_OBJECT(TYPE, ID)\
cleanup_destroyer (\
    (Cleanup_Adapter<TYPE> *)\
    preallocated_object[ID], 0);\
preallocated_object[ID] = 0;
```

The Cleanup Adapter adapts any object to use the simpler Object Lifecycle Manager registration interface, discussed in Appendix A. Similarly, Cleanup Destroyer uses the Cleanup Adapter to destroy the object.

An analogous array, enum, and macro pair can be supplied for preallocated arrays, if necessary.

- **Define an accessor interface to the Preallocated Objects:** Applications need a convenient and typesafe interface to access Preallocated Objects. The following code fragment illustrates how this interface is provided via a class template adapter in ACE:

```
template <class TYPE>
class Managed_Object
{
public:
    static TYPE *
    get_preallocated_object
        (Object_Lifecycle_Manager::
         Preallocated_Object_ID id)
    {
        // Cast the return type of the object
        // pointer based on the type of the
        // function template parameter.
```

```
        return &((Cleanup_Adapter<TYPE> *)
                 Object_Lifecycle_Manager::
                 preallocated_object[id])->object ();
    }
    // ... other methods omitted.
};
```

- 3. **Determine how to manage the lifecycle of the Object Lifecycle Manager itself.** The Object Lifecycle Manager is responsible for initializing other global and static objects in a program. However, that begs the important bootstrapping question of how this singleton initializes and destroys itself. The following are the alternatives for initializing the Object Lifecycle Manager singleton instance:

- **Static initialization:** If an application has no static objects with constraints on their order of construction or destruction, it's possible to create the Object Lifecycle Manager as a static object. For example, ACE's Object Lifecycle Manager can be created as a static object. The ACE library has no other static objects that have constraints on order of construction or destruction.

- **Stack initialization:** In this approach, create the Object Lifecycle Manager on the stack of the main program thread. This approach to initializing the Object Lifecycle Manager assumes that there is one unique main thread per program. This thread defines the program: it is *running* if, and only if, the main thread is alive. This approach has a compelling advantage: the Object Lifecycle Manager instance is automatically destroyed via any path out of *main*.³

Stack initialization is implemented transparently in ACE via a preprocessor macro named *main*. The macro renames the main program entry point to another, configurable name, such as *main.i*. It provides a main function that creates the Object Lifecycle Manager instance as a local object (on the stack) and then calls the renamed application entry point.

There are two drawbacks to the Stack initialization approach:

1. *main* (*int*, *char* *[]) must be declared with arguments, even if they're not used. All of ACE is converted to this, so only applications must be concerned with it.
2. If there are any static objects that depend on those that are destroyed by the Object Lifecycle Manager, their destructors might attempt to access the destroyed objects. Therefore, the application developer is responsible for ensuring that no static objects depend on those destroyed by the Object Lifecycle Manager.

³On platforms such as VxWorks and pSoS that have no designated main function, we simulate the main thread by instantiating the Object Lifecycle Manager on the stack of one thread, which is denoted by convention as the *main* thread.

- **Explicit initialization:** In this approach, create the Object Lifecycle Manager explicitly under application control. The Object Lifecycle Manager `init` and `fini` methods allow the application to create and destroy the Object Lifecycle Manager when desired. This option was provided primarily to alleviate complications when using dynamically loaded libraries. However, it should not be used frequently, given the availability of the dynamically loaded library initialization mechanism, below.

- **Dynamically loaded library initialization:** In this approach, create and destroy the Object Lifecycle Manager when its dynamic library is loaded and unloaded, respectively. Most dynamic library facilities include the ability to call an initialization function when the library is loaded, and a termination function when the library is unloaded.

10 Example Resolved

The discussion in Section 9 demonstrates how the Object Lifecycle Manager pattern can be used to satisfy a variety of design forces related to managing object lifecycles. In particular, the unresolved forces described in Section 2 can be adequately addressed by applying the Object Lifecycle Manager pattern.

The following example shows how the Object Lifecycle Manager pattern can be applied to the original `Logger` and `Stats` examples from Section 2. Using the `Singleton` adapter template described in Section 9 greatly simplifies managed object implementations by encapsulating key implementation details, such as registration with the Object Lifecycle Manager. For instance, the original `Stats` class can be replaced by a managed `Stats` class, as follows.

```
class Stats
{
public:
    friend class Singleton<Stats>;

    // Destructor: frees resources.
    ~Stats (void);

    // Record a timing data point.
    int record (int id,
                const timeval &tv);

    // Report recorded statistics
    // to the log.
    void report (int id) {
        Singleton<Logger>::instance ()->
            log ("Avg timing %d: "
                "%ld sec %ld usec\n",
                id,
                average_i (id).tv_sec,
                average_i (id).tv_usec);
    }

protected:
    // Default constructor.
```

```
Stats (void)
{
    // Ensure the <Logger> instance
    // is registered first, and will be
    // cleaned up after, the <Stats>
    // instance.
    Singleton<Logger>::instance ();
}

// Internal accessor for an average.
const timeval &average_i (void);

// ... other resources that are
// held by the instance ...
};
```

Notice that the singleton aspects have been factored out of the original `Stats` class and are now provided by the `Singleton` adapter template. Similar modifications can be made to the original `Logger` class so that it uses the `Singleton` adapter template.

Finally, the following example shows how an application might use the `Logger` and `Stats` classes.

```
int main (int argc, char *argv[])
{
    // Interval timestamps.
    timeval start_tv, stop_tv;

    // <Logger> and <Stats> singletons
    // do not yet exist.

    // <Logger> and then <Stats> singletons
    // are created and registered on the first
    // iteration.
    for (int i = 0; i < argc; ++i) {

        ::gettimeofday (&start_tv);
        // do some work between timestamps ...
        ::gettimeofday (&stop_tv);

        // then record the stats ...
        stop_tv.sec -= start_tv.sec;
        stop_tv.usec -= start_tv.usec;
        Singleton<Stats>::instance ()->
            record (i, stop_tv);

        // ... and log some output.
        Singleton<Logger>::instance ()->
            log ("Arg %d [%s]\n", i, argv[i]);
    }

    // <Logger> and <Stats> singletons are
    // cleaned up by Object Lifecycle Manager
    // upon program exit.
    return 0;
}
```

The following key forces are resolved in this example: (1) ensuring resources allocated by an instance are subsequently released, (2) managing the order of creation and destruction of singletons, and (3) providing a framework that encapsulates these details within a well-defined interface.

11 Known Uses

Object Lifecycle Manager is used in the Adaptive Communication Environment (ACE) [9, 11] to ensure destruction of singletons at program termination and to replace static objects with dynamically allocated, managed objects. ACE is used on many different OS platforms, some of which do not support static object construction/destruction for every program invocation. ACE supports placement of objects in ROM. Therefore, it cannot contain any objects that must be initialized at run-time.

Gabrilovich [12] augmented the Singleton pattern to permit applications to specify destruction order. A local static `auto_ptr`⁴ is responsible for the destruction of each singleton instance. Destruction of singleton instances proceeds by phase; an application may optionally register its singleton instances for destruction in a specific phase.

An interesting example of a “small” Object Lifecycle Manager is the *strong pointer* [13, 14].⁵ A strong pointer manages just one object; it destroys the object when its scope is exited, either normally or via an exception. There can be many strong pointers in a program, behaving as Function-as-Owner (or Block-as-Owner). Moreover, the strong pointers themselves have transient lifetimes, *i.e.*, that of their enclosing blocks.

In contrast, there is typically just one Object Lifecycle Manager per-program (or per-large-scale-component). Likewise, Object Lifecycle Managers live for the duration of the program invocation. This reflects the specific intent of the Object Lifecycle Manager to destroy objects at program termination, but not sooner. Such objects may be used after the current block or function has been exited, and destruction/creation cycles are not possible or desired.

12 See Also

The Object Lifecycle Manager pattern is related to the Manager [15] pattern. In both patterns, a client application uses a collection of objects, relying upon a manager to encapsulate the details of how the objects themselves are managed. This separation of concerns makes the application more robust, as potentially error-prone management aspects are hidden behind a type-safe interface. Using these managers also makes the application more extensible, as certain details of the managed objects can be varied independent of the manager implementation. For example, a manager for a certain class can be used to manage objects of classes derived from that base class.

⁴The `auto_ptr` is a local static object in the singleton instance accessor method.

⁵A C++ `auto_ptr` is an implementation of a strong pointer.

The Object Lifecycle Manager pattern differs from the Manager pattern in the types of the managed objects. Where the Manager pattern requires that the managed objects have a common base type, the Object Lifecycle Manager pattern allows objects of unrelated types to be managed. The Manager pattern relies on inheritance for variations in the manager and managed object classes. In contrast, the Object Lifecycle Manager relies on object composition and type parameterization to achieve greater decoupling of the manager from the managed objects.

The Object Lifecycle Manager pattern also differs from the Manager pattern in the management services it provides the application. The Manager pattern provides search, iteration, and deletion services for managed objects, which the Object Lifecycle Manager pattern does not provide. The application can *register* a pre-existing object with the Object Lifecycle Manager, which then assumes responsibility for the remaining lifecycle of the managed object. The Manager pattern only allows on-demand creation of objects, so that the lifecycle of objects is managed in an “all-or-none” manner. Fundamentally, the Manager pattern focuses on the repository aspects of object management, while the Object Lifecycle Manager pattern emphasizes the lifecycle aspects instead.

Object Lifecycle Manager complements creational patterns, such as Singleton, by managing object instance destruction. Singleton addresses only part of the object lifecycle because it just manages instance creation. However, destruction is usually not an important issue with Single because it does not retain *ownership* of created objects [3]. Ownership conveys the responsibility for managing the object, including its destruction. Singleton is the prototypical example of a creational pattern that does not explicitly transfer ownership, yet does not explicitly arrange for object destruction. Object Lifecycle Manager complements Singleton by managing the destruction portion of the object lifecycle.

Object Lifecycle Manager can complement other creational patterns, such as Abstract Factory and Factory Method. Implementations of these patterns could register dynamically allocated objects for deletion at program termination. Alternatively (or additionally), they could provide interfaces for object destruction, corresponding to those for object creation.

Cargill presented a taxonomy of the dynamic C++ object lifetime [16]. The Localized Ownership pattern language includes patterns, such as Creator-as-Owner, Sequence-of-Owners, and Shared Ownership, which primarily address object ownership. Ownership conveys the responsibility for destruction.

Creator-as-Owner is further subdivided into Function-as-Owner, Object-as-Owner, and Class-as-Owner. The Singleton destruction capability of Object Lifecycle Manager may be viewed as new category of Creator-as-Owner: Program-as-Owner. It is distinct from Function-as-Owner, because static

objects outlive the program entry point (`main`). Object Lifecycle Manager's Preallocated Objects similarly can be viewed logically, at least, as outliving the main program function.

When Singleton is used on multi-threaded platforms, a mutex should be used to serialize access to the instance pointer. Double-Checked Locking [10] greatly reduces the use of this mutex by only requiring it prior to creation of the singleton. However, the mutex is still required, and it must be initialized.⁶ Object Lifecycle Manager solves the chicken-and-egg problem of initialization of the mutex by preallocating one for each singleton.

Object Lifecycle Manager uses Adapter for type-safe storage of objects of any class. By using inline functions, the Managed Object Adapter should have no size/performance overhead. We confirmed this with the GreenHills 1.8.9 compiler for VxWorks on Intel targets. However, some compilers do not inline template member functions. Fortunately, the size overhead of Managed Object is very small, *i.e.*, we measured 40 to 48 bytes with `g++` on Linux and LynxOS. The ACE Cleanup Adapter template class has a bit higher size overhead, about 160 bytes per instantiation.

13 Consequences

The consequences of using Object Lifecycle Manager derive from its complete control over the entire object lifecycle. Object Lifecycle Manager may retain this control, *i.e.*, for Preallocated Objects. Or, it may provide control only for the destruction portion of the lifecycle. In the latter case, it works in concert with other creational patterns, such as Singleton or Factory Method.

The benefits of using Object Lifecycle Manager include:

Destruction of Singletons and other Managed Objects at program termination: The Object Lifecycle Manager allows the program to shut down cleanly, releasing memory for these objects along with the resources they hold at program termination. All heap-allocated memory can be released by the application. This supports repeated testing on platforms where heap allocations outlive the program.⁷ It also eliminates reports of memory in use for singletons at program end by memory access checkers.

Specification of destruction order: The order of destruction of objects can be specified. The order specification mechanism can be as simple or as complex as desired. As noted in

⁶POSIX 1003.1c [17] mutexes can be initialized without calling a static constructor. However, they are not available on all platforms.

⁷On some operating systems, notably some real-time operating systems, there is no concept of a *program*. There are tasks, *i.e.*, threads, but no one task has any special, main identity. Therefore, there is no cleanup of dynamically allocated memory, open files, *etc.*, at task termination.

Section 9, simple mechanisms are generally sufficient in practice.

Removal of static objects from libraries and applications:

Static objects can be replaced by Preallocated Objects. This prevents applications from relying on the order in which static objects are constructed/destroyed. Moreover, it allows code to target embedded systems, which sometimes have little or no support for constructing/destroying static objects.

The following liabilities must be considered when using the Object Lifecycle Manager pattern:

Lifecycle of the manager itself: The application must ensure that it respects the lifecycle of the Object Lifecycle Manager, and does not attempt to use its services outside that lifecycle. For example, the application must not attempt to access Preallocated Objects prior to the complete initialization of the Object Lifecycle Manager. Similarly, the application must not destroy the Object Lifecycle Manager prior to the application's last use of a managed or preallocated object. Finally, the implementation of the Object Lifecycle Manager is simplified if it can assume that it will be initialized by only one thread. This preclude the need for a static lock to guard its initialization.

Use with shared libraries: On platforms that support loading and unloading shared libraries at run-time, the application must be *very* careful of platform-specific issues that impact the lifecycle of the Object Lifecycle Manager itself. For example, on Windows NT, the Object Lifecycle Manager should be initialized by the application or by a DLL that contains it. This avoids potential a deadlock situation due to serialization, within the OS, of DLL loading.

A related issue arises with singletons that are created in DLLs, but managed by an Object Lifecycle Manager in the main application code. If the DLL is unloaded before program termination, the Object Lifecycle Manager would try to destroy it using code that is no longer linked into the application. For this reason, we have added an unmanaged Singleton class to ACE. An unmanaged Singleton is of the conventional design, *i.e.*, it does not provide implicit destruction. ACE uses a managed Singleton by default because we found the need for unmanaged Singletons to be very unusual.

14 Concluding Remarks

Many creational patterns specifically address only object *creation*. They do not consider when or how to *destroy* objects that are no longer needed. Object Lifecycle Manager provides mechanisms for object destruction at program termination. Thus, it complements many creational patterns by covering the entire object lifecycle.

The Singleton pattern provides a notable example where coordinated object lifecycle management is important. In particular, deletion at program termination ensures that programs have no memory leaks of singletons. Moreover, applications that employ the Object Lifecycle Manager pattern do not require use of static object constructors and destructors, which is important for embedded systems. In addition, the Object Lifecycle Manager pattern supports replacement of static objects with dynamically Preallocated Objects, which is useful on embedded platforms and with OO languages, such as C++.

One of Object Lifecycle Manager's more interesting aspects is that it addresses weaknesses of another pattern, at least in some contexts. Our initial motivation was to remedy these weaknesses by registering and deleting singletons at program termination. The utility, applicability, novelty, and complexity of Object Lifecycle Manager seemed to be on par with those of Singleton, so we felt that it deserved consideration as a pattern. Because it can address just part of the object lifetime, however, we consider Object Lifecycle Manager to be a *complementary* pattern.

Another interesting question was: "how do we categorize Object Lifecycle Manager?" It was not (originally) a Creational pattern, because it handled only object destruction, not creation. Again, it seemed appropriate to refer to Object Lifecycle Manager as complementing Singleton.

During the shepherding process, it became apparent that Object Lifecycle Manager had another use that was very related to destroying singletons. Static objects create problems similar to those of singletons, *i.e.*, destruction, especially on operating systems that have no notion of a program, and order of construction/destruction. Object Lifecycle Manager's Preallocated Objects were added to support removal of static objects. Our first Preallocated Object was the mutex used for Double-Checked Locking [10] in the ACE implementation of . . . Singleton.

Current development efforts include breaking the one instance into multiple Object Lifecycle Managers, to support subsetting. Each layer of ACE, *e.g.*, OS, logging, threads, connection management, sockets, interprocess communication, service configuration, streams, memory management, and utilities, will have its own Object Lifecycle Manager. When any Object Lifecycle Manager is instantiated, it will instantiate each dependent Object Lifecycle Manager, if not already done. And similar, configured-in cooperation will provide graceful termination.

A highly portable implementation of the Object Lifecycle Manager pattern and the Singleton adapter template is freely available and can be downloaded from www.cs.wustl.edu/~schmidt/ACE-obtain.html.

Acknowledgments

Thanks to Matthias Kerkhoff, Per Andersson, Steve Huston, Elias Sreih, and Liang Chen for many helpful discussions on the design and implementation of ACE's Object Lifecycle Manager. And thanks to Brad Appleton, our PLoP '99 shepherd, and Evgeniy Gabrilovich for many helpful suggestions on the content and presentation of this pattern.

References

- [1] E. Gamma, R. Helm, R. Johnson, and J. Vlissides, *Design Patterns: Elements of Reusable Object-Oriented Software*. Reading, MA: Addison-Wesley, 1995.
- [2] D. C. Schmidt, "Reactor: An Object Behavioral Pattern for Concurrent Event Demultiplexing and Event Handler Dispatching," in *Pattern Languages of Program Design* (J. O. Coplien and D. C. Schmidt, eds.), pp. 529–545, Reading, MA: Addison-Wesley, 1995.
- [3] J. Vlissides, "To Kill a Singleton," *C++ Report*, vol. 8, pp. 10–19, June 1996.
- [4] Wind River Systems, "VxWorks 5.3." <http://www.wrs.com/products/html/vxworks.html>.
- [5] Integrated Systems, Inc., "pSOSystem." <http://www.isi.com/products/psosystem/>.
- [6] "Java Singleton." <http://c2.com/cgi/wiki?JavaSingleton>, May 1999.
- [7] "Another Java Singleton Problem." <http://c2.com/cgi/wiki?AnotherJavaSingletonProblem>, Jan. 1999.
- [8] D. Saks, "Ensuring Static Initialization in C++," *Embedded Systems Programming*, vol. 12, pp. 109–111, Mar. 1999.
- [9] D. C. Schmidt, "ACE: an Object-Oriented Framework for Developing Distributed Applications," in *Proceedings of the 6th USENIX C++ Technical Conference*, (Cambridge, Massachusetts), USENIX Association, April 1994.
- [10] D. C. Schmidt and T. Harrison, "Double-Checked Locking – An Object Behavioral Pattern for Initializing and Accessing Thread-safe Objects Efficiently," in *Pattern Languages of Program Design* (R. Martin, F. Buschmann, and D. Riehle, eds.), Reading, MA: Addison-Wesley, 1997.
- [11] D. C. Schmidt, "The ADAPTIVE Communication Environment (ACE)." www.cs.wustl.edu/~schmidt/ACE.html, 1997.
- [12] E. Gabrilovich, "Controlling the Destruction Order of Singleton Objects," *C/C++ Users Journal*, Oct. 1999. To appear.
- [13] B. Milewski, "Strong Pointers and Resource Management in C++," *C++ Report*, vol. 10, pp. 23–27, Sept. 1998.
- [14] B. Milewski, "Strong Pointers and Resource Management in C++, Part 2," *C++ Report*, vol. 11, pp. 36–39,50, Feb. 1999.
- [15] P. Sommerland and F. Buschmann, "The Manager Design Pattern," in *Proceedings of the 3rd Pattern Languages of Programming Conference*, September 1996.
- [16] T. Cargill, "Localized Ownership: Managing Dynamic Objects in C++," in *Pattern Languages of Program Design 2* (J. M. Vlissides, J. O. Coplien, and N. L. Kerth, eds.), Reading, MA: Addison-Wesley, 1996.

- [17] “Information Technology – Portable Operating System Interface (POSIX) – Part 1: System Application: Program Interface (API) [C Language],” 1995.

A Detailed Implementation

This section provides an example of a concrete Object Lifecycle Manager implementation in C++ described in Section 9. It is based on the The ACE [9, 11] Object Lifecycle Manager implementation. The ACE implementation reveals some interesting design issues. Its most visible purpose is to manage cleanup of singletons at program termination, and create/destroy Preallocated Objects. In addition, it performs other cleanup actions, such as shutdown of services provided by the ACE library, at program termination.

The Object Lifecycle Manager Base abstract base class, shown in Figure 3, provides the initialization and finalization mechanisms for an Object Lifecycle Manager. Subclasses must specialize and provide implementations, described below.

In addition, Object Lifecycle Manager Base supports chaining of Object Lifecycle Managers. Object Lifecycle Managers are Singletons, each with its own locus of interest. An application may have a need for more than one Object Lifecycle Manager, *e.g.*, one per major component. Chaining permits ordered shutdown of the separate components.

Figure 4 shows an example Object Lifecycle Manager class. It is a Singleton, so it provides a static instance accessor. In addition, it provides static `starting_up` and `shutting_down` state accessors. An enumeration lists identifiers for the Preallocated Objects that it owns.

An interesting detail is the (boolean) reference count logic provided by the derived class `init` and `fini` methods. There are several alternatives for constructing an Object Lifecycle Manager, discussed in Section 9. The reference count ensures that an Object Lifecycle Manager is only constructed once, and destroyed once.

The implementations of the `instance`, `init`, and `fini` methods are shown in Figure 5 and Figure 6. The `instance` method is typical of Singleton instance accessors, but includes logic to support static placement instead of dynamic allocation. In addition, it is not thread safe, requiring construction before the program spawns any threads. This avoids the need for a lock to guard the allocation.

The `init` and `fini` methods show creation and destruction of Preallocated Objects, respectively. They show application-specific startup and shutdown code. Finally, they show maintenance of the Object Lifecycle Manager state.⁸

⁸This should be moved up to the base class.

```
class Object_Lifecycle_Manager_Base
{
public:
    virtual int init (void) = 0;
    // Explicitly initialize. Returns 0 on success,
    // -1 on failure due to dynamic allocation
    // failure (in which case errno is set to
    // ENOMEM), or 1 if it had already been called.

    virtual int fini (void) = 0;
    // Explicitly destroy. Returns 0 on success,
    // -1 on failure because the number of fini ()
    // calls hasn't reached the number of init ()
    // calls, or 1 if it had already been called.

    enum Object_Lifecycle_Manager_State {
        OBJ_MAN_UNINITIALIZED,
        OBJ_MAN_INITIALIZING,
        OBJ_MAN_INITIALIZED,
        OBJ_MAN_SHUTTING_DOWN,
        OBJ_MAN_SHUT_DOWN
    };

protected:
    Object_Lifecycle_Manager_Base (void) :
        object_manager_state_ (OBJ_MAN_UNINITIALIZED),
        dynamically_allocated_ (0),
        next_ (0) {}

    virtual ~Object_Lifecycle_Manager_Base (void) {
        // Clear the flag so that <fini>
        // doesn't delete again.
        dynamically_allocated_ = 0;
    }

    int starting_up_i (void) {
        return object_manager_state_ <
            OBJ_MAN_INITIALIZED;
    }
    // Returns 1 before Object_Lifecycle_Manager_Base
    // has been constructed. This flag can be used
    // to determine if the program is constructing
    // static objects. If no static object spawns
    // any threads, the program will be
    // single-threaded when this flag returns 1.

    int shutting_down_i (void) {
        return object_manager_state_ >
            OBJ_MAN_INITIALIZED;
    }
    // Returns 1 after Object_Lifecycle_Manager_Base
    // has been destroyed.

    Object_Lifecycle_Manager_State object_manager_state_;
    // State of the Object_Lifecycle_Manager;

    u_int dynamically_allocated_;
    // Flag indicating whether the
    // Object_Lifecycle_Manager instance was
    // dynamically allocated by the library.
    // (If it was dynamically allocated by the
    // application, then the application is
    // responsible for deleting it.)

    Object_Lifecycle_Manager_Base *next_;
    // Link to next Object_Lifecycle_Manager,
    // for chaining.
};
```

Figure 3: Object Lifecycle Manager Base Class

```

class Object_Lifecycle_Manager :
    public Object_Lifecycle_Manager_Base
{
public:
    virtual int init (void);

    virtual int fini (void);

    static int starting_up (void) {
        return instance_ ?
            instance_>starting_up_i () : 1;
    }

    static int shutting_down (void) {
        return instance_ ?
            instance_>shutting_down_i () : 1;
    }

    enum Preallocated_Object
    {
# if defined (MT_SAFE) && (MT_SAFE != 0)
        OS_MONITOR_LOCK,
        TSS_CLEANUP_LOCK,
# else
        // Without MT_SAFE, There are no
        // preallocated objects. Make
        // sure that the preallocated_array
        // size is at least one by declaring
        // this dummy.
        EMPTY_PREALLOCATED_OBJECT,
# endif /* MT_SAFE */
        // This enum value must be last!
        PREALLOCATED_OBJECTS
    };
    // Unique identifiers for Preallocated Objects.

    static Object_Lifecycle_Manager *instance (void);
    // Accessor to singleton instance.

public:
    // Application code should not use these
    // explicitly, so they're hidden here. They're
    // public so that the Object_Lifecycle_Manager
    // can be onstructed/destroyed in <main>, on
    // the stack.
    Object_Lifecycle_Manager (void) {
        // Make sure that no further instances are
        // created via <instance>.
        if (instance_ == 0)
            instance_ = this;
        init ();
    }

    ~Object_Lifecycle_Manager (void) {
        // Don't delete this again in <fini>
        dynamically_allocated_ = 0;
        fini ();
    }

private:
    static Object_Lifecycle_Manager *instance_;
    // Singleton instance pointer.

    static void *
        preallocated_object[PREALLOCATED_OBJECTS];
    // Array of Preallocated Objects.
};

```

Figure 4: Object Lifecycle Manager Class

```

Object_Lifecycle_Manager *
Object_Lifecycle_Manager::instance_ = 0;
// Singleton instance pointer.

Object_Lifecycle_Manager *
Object_Lifecycle_Manager::instance (void)
{
    // This function should be called during
    // construction of static instances, or
    // before any other threads have been created
    // in the process. So, it's not thread safe.
    if (instance_ == 0) {
        Object_Lifecycle_Manager *instance_pointer =
            new Object_Lifecycle_Manager;

        // instance_ gets set as a side effect of the
        // Object_Lifecycle_Manager allocation, by
        // the default constructor. Verify that . . .
        assert (instance_pointer == instance_);

        instance_pointer->dynamically_allocated_ = 1;
    }
    return instance_;
}

int
Object_Lifecycle_Manager::init (void)
{
    if (starting_up_i ()) {
        // First, indicate that this
        // Object_Lifecycle_Manager instance
        // is being initialized.
        object_manager_state_ = OBJ_MAN_INITIALIZING;

        if (this == instance_) {
# if defined (MT_SAFE) && (MT_SAFE != 0)
            PREALLOCATE_OBJECT (mutex_t,
                OS_MONITOR_LOCK)
            // Mutex initialization omitted.

            PREALLOCATE_OBJECT (recursive_mutex_t,
                TSS_CLEANUP_LOCK)
            // Recursive mutex initialization omitted.
# endif /* MT_SAFE */

            // Open Winsock (no-op on other
            // platforms).
            socket_init (/* WINSOCK_VERSION */);

            // Other startup code omitted.
        }

        // Finally, indicate that the
        // Object_Lifecycle_Manager instance
        // has been initialized.
        object_manager_state_ = OBJ_MAN_INITIALIZED;
        return 0;
    } else {
        // Had already initialized.
        return 1;
    }
}

```

Figure 5: Object Lifecycle Method Implementations

```

int
Object_Lifecycle_Manager::fini (void)
{
    if (shutting_down_i ())
        // Too late. Or, maybe too early. Either
        // fini () has already been called, or
        // init () was never called.
        return object_manager_state_ ==
            OBJ_MAN_SHUT_DOWN ? 1 : -1;

    // Indicate that the Object_Lifecycle_Manager
    // instance is being shut down.
    // This object manager should be the last one
    // to be shut down.
    object_manager_state_ = OBJ_MAN_SHUTTING_DOWN;

    // If another Object_Lifecycle_Manager has
    // registered for termination, do it.
    if (next_) {
        next_>fini ();
        // Protect against recursive calls.
        next_ = 0;
    }

    // Only clean up Preallocated Objects when
    // the singleton Instance is being destroyed.
    if (this == instance_) {
        // Close down Winsock (no-op on other
        // platforms).
        socket_fini ();

        // Cleanup the dynamically preallocated
        // objects.
        # if defined (MT_SAFE) && (MT_SAFE != 0)
            // Mutex destroy not shown . . .
            DELETE_PREALLOCATED_OBJECT (mutex_t,
                MONITOR_LOCK)

            // Recursive mutex destroy not shown . . .
            DELETE_PREALLOCATED_OBJECT (
                recursive_mutex_t,
                TSS_CLEANUP_LOCK)
        # endif /* MT_SAFE */
    }

    // Indicate that this Object_Lifecycle_Manager
    // instance has been shut down.
    object_manager_state_ = OBJ_MAN_SHUT_DOWN;

    if (dynamically_allocated_)
        delete this;

    if (this == instance_)
        instance_ = 0;

    return 0;
}

```

Figure 6: Object Lifecycle Method Implementations, cont'd.