The main Program in Ruby

On this page:

Initializing the Ice Run Time in Ruby

The main entry point to the Ice run time is represented by the local interface `Ice::Communicator`. You must initialize the Ice run time by calling `Ice::initialize` before you can do anything else in your program.

`Ice::initialize` returns a reference to an instance of `Ice:::Communicator`:

Ruby

```ruby
require 'Ice'

status = 0
ic = nil
begin
    ic = Ice::initialize(ARGV)
    # ...
rescue => ex
    puts ex
    status = 1
end

# ...
```

`Ice::initialize` accepts the argument list that is passed to the program by the operating system. The function scans the argument list for any command-line options that are relevant to the Ice run time; any such options are removed from the argument list so, when `Ice::initialize` returns, the only options and arguments remaining are those that concern your application. If anything goes wrong during initialization, `initialize` throws an exception.

Before leaving your program, you *must* call `Communicator.destroy`. The `destroy` operation is responsible for finalizing the Ice run time. In particular, `destroy` ensures that any outstanding threads are joined with and reclaims a number of operating system resources, such as file descriptors and memory. Never allow your program to terminate without calling `destroy` first; doing so has undefined behavior.

The general shape of our program is therefore as follows:

Ruby

```ruby
require 'Ice'

status = 0
ic = nil
begin
    ic = Ice::initialize(ARGV)
    # ...
rescue => ex
    puts ex
    status = 1
end

if ic
    begin
        ic.destroy()
    rescue => ex
        puts ex
        status = 1
    end
end

exit(status)
```

Note that the code places the call to `Ice::initialize` into a `begin` block and takes care to return the correct exit status to the operating system. Also note that an attempt to destroy the communicator is made only if the initialization succeeded.

The `Ice::Application` Class in Ruby

The preceding program structure is so common that Ice offers a class, `Ice::Application`, that encapsulates all the correct initialization and finalization activities. The synopsis of the class is as follows (with some detail omitted for now):

Ruby

```ruby
module Ice
    class Application
        def main(args, configFile=nil, initData=nil)

        def run(args)

        def Application.appName()

        def Application.communicator()
    end
end
```

The intent of this class is that you specialize `Ice::Application` and implement the abstract `run` method in your derived class. Whatever code you would normally place in your main program goes into `run` instead. Using `Ice::Application`, our program looks as follows:

Ruby

```ruby
require 'Ice'

class Client < Ice::Application
    def run(args)
        # Client code here...
        return 0
    end
end

app = Client.new()
status = app.main(ARGV)
exit(status)
```

If you prefer, you can also reopen `Ice::Application` and define `run` directly:

Ruby

```ruby
require 'Ice'

class Ice::Application
    def run(args)
        # Client code here...
        return 0
    end
end

app = Ice::Application.new()
status = app.main(ARGV)
exit(status)
```

You may also call `main` with an optional file name or an `InitializationData` structure. If you pass a configuration file name to `main`, the settings in this file are overridden by settings in a file identified by the `ICE_CONFIG` environment variable (if defined). Property settings supplied on the command line take precedence over all other settings.

The `Application.main` method does the following:

1. It installs an exception handler. If your code fails to handle an exception, `Application.main` prints the exception information before returning with a non-zero return value.
2. It initializes (by calling `Ice::initialize`) and finalizes (by calling `Communicator.destroy`) a communicator. You can get access to the communicator for your program by calling the static `communicator` accessor.
3. It scans the argument list for options that are relevant to the Ice run time and removes any such options. The argument list that is passed to your `run` method therefore is free of Ice-related options and only contains options and arguments that are specific to your application.
4. It provides the name of your application via the static `appName` method. The return value from this call is the first element of the argument vector passed to `Application.main`, so you can get at this name from anywhere in your code by calling `Ice::Application::appName` (which is often necessary for error messages).
5. It installs a signal handler that properly shuts down the communicator.

Using `Ice::Application` ensures that your program properly finalizes the Ice run time, whether your program terminates normally or in response to an exception or signal. We recommend that all your programs use this class; doing so makes your life easier. In addition `Ice::Application` also provides features for signal handling and configuration that you do not have to implement yourself when you use this class.

Catching Signals in Ruby

A program typically needs to perform some cleanup work before terminating, such as flushing database buffers or closing network connections. This is particularly important on receipt of a signal or keyboard interrupt to prevent possible corruption of database files or other persistent data.

To make it easier to deal with signals, `Ice::Application` encapsulates Ruby's signal handling capabilities, allowing you to cleanly shut down on receipt of a signal:

Ruby

```ruby
class Application
    def Application.destroyOnInterrupt()

    def Application.ignoreInterrupt()

    def Application.callbackOnInterrupt()

    def Application.holdInterrupt()

    def Application.releaseInterrupt()

    def Application.interrupted()

    def interruptCallback(sig):
        # Default implementation does nothing.
    end
    # ...
end
```

The methods behave as follows:

- `destroyOnInterrupt`
  This method installs a signal handler that destroys the communicator if it is interrupted. This is the default behavior.

- `ignoreInterrupt`
  This method causes signals to be ignored.

- `callbackOnInterrupt`
  This function configures `Ice::Application` to invoke `interruptCallback` when a signal occurs, thereby giving the subclass responsibility for handling the signal.

- `holdInterrupt`
  This method temporarily blocks signal delivery.

- `releaseInterrupt`
  This method restores signal delivery to the previous disposition. Any signal that arrives after `holdInterrupt` was called is delivered when you call `releaseInterrupt`.

- `interrupted`
  This method returns `True` if a signal caused the communicator to shut down, `False` otherwise. This allows us to distinguish intentional shutdown from a forced shutdown that was caused by a signal. This is useful, for example, for logging purposes.

- `interruptCallback`
  A subclass implements this function to respond to signals. The function may be called concurrently with any other thread and must not raise exceptions.

By default, `Ice::Application` behaves as if `destroyOnInterrupt` was invoked, therefore our program requires no change to ensure that the program terminates cleanly on receipt of a signal. (You can disable the signal-handling functionality of `Ice::Application` by passing the constant `NoSignalHandling` to the constructor. In that case, signals retain their default behavior, that is, terminate the process.) However, we add a diagnostic to report the occurrence of a signal, so our program now looks like:

Ruby

```ruby
require 'Ice'

class MyApplication < Ice::Application
    def run(args)
        # Client code here...

        if Ice::Application::interrupted()
            print Ice::Application::appName() + ": terminating"
        end

        return 0
    end
end

app = MyApplication.new()
status = app.main(ARGV)
exit(status)
```

`Ice::Application` and Properties in Ruby

Apart from the functionality shown in this section, `Ice::Application` also takes care of initializing the Ice run time with property values. Properties allow you to configure the run time in various ways. For example, you can use properties to control things such as the thread pool size or the trace level for diagnostic output. The `main` method of `Ice::Application` accepts an optional second parameter allowing you to specify the name of a configuration file that will be processed during initialization.

Limitations of `Ice::Application` in Ruby

`Ice::Application` is a singleton class that creates a single communicator. If you are using multiple communicators, you cannot use `Ice::Application`. Instead, you must structure your code as we saw in Hello World Application (taking care to always destroy the communicator).

See Also

- Hello World Application
- Properties and Configuration
- Communicator Initialization
- Logger Facility