



# Program a Game Engine from Scratch

Mark Claypool

Development Checkpoint #2

Clock & GameManager

This document is part of the book “Dragonfly – Program a Game Engine from Scratch”, (Version 5.0). Information online at: <http://dragonfly.wpi.edu/book/>

Copyright ©2012–2017 Mark Claypool and WPI. All rights reserved.

## 4.4 Game Management

At a high level, “managing the game” is the job description of the entire game engine. Game programmers (and players) often think of this as “running the game”.

### 4.4.1 The Game Loop

The game manager “runs” the game, doing so by repeating the same basic set of actions in a loop (the *game loop*), over and over. A 10,000 foot view of the game loop is presented in Listing 4.13. Each iteration of the game loop is called a “step” (or a “tick”, as in the tick of a loop clock). During one step, the game loop: 1) gets input, say from the keyboard or the mouse (these are player actions in the game); 2) updates the game world state to move objects around, generate needed actions, respond to the input; 3) draws a new image (the current scene) on the graphics buffer; and 4) swaps out the old image for the new image. This process is repeated in the loop until the game is over.

Listing 4.13: The game loop

```

0 while (game not over) do
1   Get input // e.g., from keyboard/mouse
2   Update game world state
3   Draw current scene to back buffer
4   Swap back buffer to current buffer
5 end while

```

Note that the loop in Listing 4.13 runs as fast as it can, updating and drawing the game world as fast as the computer can get through the code. Early game development efforts were often targeted for a machine with a specific speed, where the time to execute a loop was known and objects could be moved an appropriate amount of time each loop. Of course, running the same game code on a faster machine (as would happen when computer speeds improved) meant the game would run faster! Moreover, if a step took more or less time than expected, the update rate of game objects would vary, causing them to move faster or slower.

In order to rectify this problem, the game loop is enhanced with loop timing information, shown in Listing 4.14. In this version of the game loop, one step of the loop is expected to take `TARGET_TIME` (e.g., 33 milliseconds). So, the time to execute the first 4 instructions is carefully measured and, at the end of the loop on line 6, the game is put to sleep (effectively, doing nothing) for whatever is remaining of the `TARGET_TIME`.

Listing 4.14: The game loop with timing

```

0 while (game not over) do
1   Get input // e.g., keyboard/mouse
2   Update game world state
3   Draw current scene to back buffer
4   Swap back buffer to current buffer
5   Measure loop_time // i.e., how long above steps took
6   Sleep for (TARGET_TIME - loop_time)
7 end while

```



An important decision is deciding how long `TARGET_TIME` should be. Setting it too high will result in the game loop progressing slowly, limiting animation rates and game update rates – the game will look less “smooth” and will feel sluggish to the player. Setting it too low will result in the game loop progressing rapidly, giving a smooth, responsive game, but may unnecessarily burden the computer and cause problems, such as visual glitches or unintended game slowdowns, if the game world is too complicated to be fully updated in one step.

Guidelines for setting `TARGET_TIME` can be drawn from video. Video performance is often reported in units of frame rate, the rate at which video images are updated on the screen. The units are typically frames per second (f/s). “Full motion” video, the rate seen in movies or television, is approximately 30 f/s. Frame rates higher than this provides little benefit to visual quality, while frame rates below this look “jerky” for some kinds of video content. Considering the rendered game images as video images, the game loop rate is analogous to video frame rates, provided guidance on the game loop rates. Notably, a reasonable expectation is to update the game screen 30 times per second – equivalently, setting `TARGET_TIME` to 33 milliseconds.

#### 4.4.2 Measuring Computer Time

In order to step through the game loop 30 times per second, the time for one loop iteration must be measured precisely. Modern operating systems provide several different ways (system calls) to measure time. For example, on Unix systems, the `time()` call returns the number of seconds since January 1st, 1970. Subsequent system calls can use that number to extract the hours, minutes and seconds or even the month, day, year. However, the resolution of the `time()` system call is only 1 second, meaning it is too coarse to provide timing on the order of the milliseconds needed for the game loop.

Fortunately, modern computer processors have high-resolution timers provided by hardware registers that count processor cycles. For instance, a 3 GHz processor increments the timer register 3 billion times per second, providing a resolution of 0.3 nanoseconds – plenty of precision for the game loop! The actual system calls to access these high-resolution timers varies with platform. Windows uses `QueryPerformanceCounter()` to get the timer value, and `QueryPerformanceFrequency()` to get the processor cycle rate. Xbox 360 and PS3 game consoles use the `mftb` (which stands for “move from time base”) register to obtain the timer value, with the hardware having a known processor cycle time. Linux uses `clock_gettime()` to get a high-resolution time value (needing to be linked in with the real-time library, `-lrt`).

In order to measure the time the game loop takes (everything between line 1 “Get input” and line 4 “Swap” in Listing 4.14), the method in Listing 4.15 is used. The method starts by recording the time (storing it in a variable). Next, the tasks to be timed are run (for a game loop, this is input, update and so forth). When the tasks are done, the time is again recorded. The elapsed time is obtained by subtracting the before time from the after time.

Listing 4.15: Method to measure elapsed time

```

0 // Record before time
1 // Do processing stuff (get input, update ...)
2 // Record after time

```



```
3 // Compute elapsed time: after - before
```

For a Linux platform, Listing 4.16 provides a code fragment to compute the elapsed time of a block of computation. Note, in this example, the units for elapsed time are in microseconds, which is often used for timing in game engines (as are milliseconds), but it could easily be adjusted to seconds or milliseconds. The system header file `time.h` is needed for the timing routines and compilation needs `-lrt`, to link in the real-time library. The timing function, `clock_gettime()`, fills in the components of a `timespec` structure, which includes fields for seconds (`tv_sec`) and nanoseconds (`tv_nsec`). Computing elapsed time is done by converting the seconds and nanoseconds to microseconds, and subtracting the initial value from the final value.

Listing 4.16: Measuring elapsed time in Linux

```
0 // Compile with -lrt
1 #include <time.h>
2
3 struct timespec before_ts, after_ts;
4 clock_gettime(CLOCK_REALTIME, &before_ts); // Start timing.
5 // Do stuff...
6 clock_gettime(CLOCK_REALTIME, &after_ts); // Stop timing.
7
8 // Compute elapsed time in microseconds.
9 long int before_msec = before_ts.tv_sec*1000000 + before_ts.tv_nsec/1000;
10 long int after_msec = after_ts.tv_sec*1000000 + after_ts.tv_nsec/1000;
11 long int elapsed_time = after_msec - before_msec;
```

For Mac, the system call `clock_gettime()` does not exist (nor does the `rt` library). Instead, the system call `gettimeofday()` (located in `<sys/time.h>`) should be used, as shown in Listing 4.17. A call to `gettimeofday()` fills a `struct timeval` with the number of seconds and microseconds.

Listing 4.17: Measuring elapsed time in Mac OS

```
0 #include <sys/time.h>
1
2 struct timeval before_tv, after_tv;
3
4 gettimeofday(&before_tv, NULL); // Start timing.
5 // Do stuff...
6 gettimeofday(&after_tv, NULL); // Stop timing.
7
8 // Compute elapsed time in microseconds.
9 long int before_msec = before_tv.tv_sec*1000000 + before_tv.tv_usec;
10 long int after_msec = after_tv.tv_sec*1000000 + after_tv.tv_usec;
11 long int elapsed_time = after_msec - before_msec;
```

For Windows, the system call needed is `GetSystemTime()` (located in `<Windows.h>`) is used, as shown in Listing 4.18. A call to `GetSystemTime()` fills a `SYSTEMTIME` structure with the number of minutes, seconds, and milliseconds.

Listing 4.18: Measuring elapsed time in Windows

```
0 #include <Windows.h>
```



```

1
2 SYSTEMTIME before_st, after_st;
3 GetSystemTime(&before_st);
4 // Do stuff...
5 GetSystemTime(&after_st);
6
7 // Compute elapsed time in microseconds.
8 long int before_msec = (before_st.wMinute * 60 * 1000000)
9                       + (before_st.wSecond * 1000000)
10                      + (before_st.wMilliseconds * 1000);
11 long int after_msec = (after_st.wMinute * 60 * 1000000)
12                     + (after_st.wSecond * 1000000)
13                     + (after_st.wMilliseconds * 1000);
14 long int elapsed_time = after_msec - before_msec;

```

### 4.4.3 The Clock Class

It is helpful for both the game engine and the game programmer to have a class that provides convenient access to high-resolution timing – the Clock class. Listing 4.19 provides the header file for the Clock class.<sup>4</sup> The clock functions as a sort of “stopwatch”, so the time is stored in the variable `previous_time`, set when a Clock object is instantiated to the current time. A call to the method `delta()` returns the elapsed time (in microseconds), and resets `previous_time` to the current time. A call to the method `split()` returns the time (in microseconds) since the last `delta()` call, but does not change the value of `previous_time`. The constructor should set `previous_time` to the current time, and both `delta()` and `split()` can be implemented using Listing 4.16, 4.17, or 4.18 (as appropriate), as a reference.

Listing 4.19: Clock.h

```

0 // The clock, for timing (such as in the game loop)
1
2 class Clock {
3
4 private:
5     long int m_previous_time; // Previous time delta() called (in microsec).
6
7 public:
8     // Sets previous_time to current time.
9     Clock();
10
11     // Return time elapsed since delta() was last called, -1 if error.
12     // Resets clock time.
13     // Units are microseconds.
14     long int delta();
15
16     // Return time elapsed since delta() was last called, -1 if error.
17     // Does not reset clock time.
18     // Units are microseconds.
19     long int split() const;
20 };

```

<sup>4</sup>Note, the conditional `#ifdef` directives described in Section 4.3.4 are not shown.



With a Clock class for timing, the last missing piece for providing timing control in the game loop is the ability to sleep (line 6 of Listing 4.14). Linux and Mac provide the `sleep()` system call, but it has only seconds of resolution, meaning it will not allow the game engine to sleep for, say, 20 milliseconds. Since game loop timing needs milliseconds of resolution, so does an appropriate sleep call.

On Linux and Mac, this can be done with `nanosleep()` which sleeps for a given number of nanoseconds.<sup>5</sup> A `#include <time.h>` is needed for `nanosleep()`. The system call `nanosleep()` takes in a pointer to a `struct timespec`, that has the amount of seconds plus nanoseconds to sleep. The example in Listing 4.20 shows a call to `nanosleep()` for 20 milliseconds.

Listing 4.20: `nanosleep()` example for Linux and Mac

```
0 // Sleep for 20 milliseconds.
1 struct timespec sleep_time;
2 sleep_time.tv_sec = 0;
3 sleep_time.tv_nsec = 20000000;
4 nanosleep(&sleep_time, NULL);
```

On Windows, sleeping can be done with `Sleep()` which sleeps for a given number of milliseconds. A `#include <Windows.h>` is needed for `Sleep()`. In order to obtain a millisecond resolution using `Sleep()`, the system call `timeBeginPeriod(1)` needs to be called once, when the game engine starts, to set the timer resolution to the minimum possible. The system call `timeEndPeriod(1)` is called once the game engine exits to clear the initial request for a minimal timer resolution. Both functions return `TIMERR_NOERROR` if successful or `TIMERR_NOCANDO` if the resolution specified is out of range. Note, the best places for these calls are when the GameManager starts up and when the GameManager shuts down, respectively (see Section 4.4.4 on page 69).

Listing 4.21: `Sleep()` example for Windows

```
0 // Sleep for 20 milliseconds.
1 int sleep_time = 20;
2 Sleep(sleep_time);
```

Listing 4.22 provides pseudo-code for how the Clock class and sleep functions can be used together in the game loop. The call to `clock.delta()` at the beginning of the loop starts the timing, while the call to `clock.split()` after most of the loop body provides the elapsed time, measuring how long the game loop took. The game engine then sleeps (via `nanosleep()` for Linux or Mac or `Sleep()` for Windows) for `TARGET_TIME - loop_time`.

Listing 4.22: The game loop with Clock

```
0 Clock clock
1 while (game not over) do
2     clock.delta()
3     // Get input (e.g., from keyboard/mouse)
4     // Update world state
5     // Draw current scene to back buffer
6     // Swap back buffer to current buffer
```

<sup>5</sup>There are 1 billion nanoseconds in 1 second.



```

7  loop_time = clock.split()
8  sleep(TARGET_TIME - loop_time)
9  end while

```

The expectation is that `(TARGET_TIME - loop_time)` is positive, since the `sleep()` call on line 8 of Listing 4.22 expects positive number. But what happens when it is not? First off, consider what it means for `(TARGET_TIME - loop_time)` to be negative. This happens when the time to do the work in the game loop (the input, update, draw and swap) takes longer than the expected time for one iteration of the game loop (longer than `TARGET_TIME`). When this happens, the game engine cannot keep up with the work required to run the game, resulting, at a minimum, in the displayed frame rate that the player sees to decrease. For example, if the `TARGET_TIME` is 33 milliseconds, providing a frame rate of 30 f/s, but the loop time (`loop_time`) takes 50 milliseconds, the frame rate is only be 20 f/s. With longer loop times, the frame rate drops further, decreasing the smoothness of the visual display for the player. The time between getting input from the player also decreases, probably making the game feel less responsive.

If the `loop_time` is greater than the `TARGET_TIME`, do the game objects themselves need to slow down also? Not necessarily. When updating the game world, the engine can be aware of the previous update time, thus knowing how much time has elapsed, and use this to decide how far, say, an object should move. This requires the game engine to pass along timing information to update functions and for those functions to use the information accordingly. However, in *Dragonfly*, the engine does not do this, so if the computer cannot keep up at the expected `TARGET_TIME` pace, the game will look, feel and run slower. Thus, as for programming all games using a game engine, *Dragonfly* game programmers must work within the constraints of the engine to ensure the load their game places on the engine does not cause performance issues.

#### 4.4.3.1 Fine Tuning the Game Loop (optional)

A subtle timing aspect that is important for some games is that when calling operating system sleep functions (e.g., `nanosleep()` or `Sleep()`), the actual amount of sleep time may be longer than requested depending upon other activity in the system and the operating system scheduler. In most cases, this does not matter much, since sleep differences are typically being only a matter of a few milliseconds. However, in some cases, such as trying to synchronize game state on two different machines in a multi-player game, more precision in the total time a game loop takes is required. So, a final adjustment to the loop timing can be made by determining how long the sleep function call actually took. Measurement can be done before and after the sleep call, with any extra time subtracted from the next game loop. Listing 4.23 shows how to put in this adjustment.

Listing 4.23: The game loop with Clock and sleep adjustment

```

0  Clock clock
1  while (game not over) do
2    clock.delta()
3    // Get input (e.g., from keyboard/mouse)
4    // Update world state
5    // Draw current scene to back buffer

```



```

6 // Swap back buffer to current buffer
7 loop_time = clock.split()
8 intended_sleep_time = TARGET_TIME - loop_time - adjust_time
9 clock.delta()
10 sleep(intended_sleep_time)
11 actual_sleep_time = clock.split()
12 adjust_time = actual_sleep_time - intended_sleep_time
13 end while

```

#### 4.4.4 The GameManager

With timing technologies developed for the game loop, implementation of the GameManager can now be started with the class definition provided in Listing 4.24.

The GameManager constructor should set the type of the Manager to “GameManager” (i.e., `setType("GameManager")`) and initialize all attributes.

The method `setGameOver()` lets the game programmer set the game over condition when ready (e.g., the player has indicated they want to quit) and `getGameOver()` returns the game over status.

The `run()` method is used to start the game, effectively running the game loop until the game is over, controlled by the boolean attribute `game_over`.

The GameManager needs start up methods, like all engine managers. The GameManager `startUp()` method instantiates (via `getInstance()`) and starts up (via `startUp()`) all the other game managers and in the right order. For now, the GameManager only starts up the LogManager. The `game_over` variable should be set to `false`. Most games typically use the default of 33 milliseconds (line 3), but games that want to run faster or slower may want to use an alternate frame time.\* If developing for Windows, GameManager `startUp()` should invoke `timeBeginPeriod(1)` (see page 67).

The `shutDown()` method does the reverse, shutting down the LogManager. It calls `setGameOver()` to indicate to any game objects that the game is over, which sets the `game_over` variable to `true`. If developing for Windows, GameManager `shutDown()` should invoke `timeEndPeriod(1)` (see page 67).

Upon success, Manager `startUp()` and Manager `shutDown()` should be called from GameManager `startUp()` and GameManager `shutDown()`, respectively.

Listing 4.24: GameManager.h

```

0 #include "Manager.h"
1
2 // Default frame time (game loop time) in milliseconds (33 ms == 30 f/s).
3 const int FRAME_TIME_DEFAULT = 33;
4
5 class GameManager : public Manager {
6
7 private:
8     GameManager(); // Private since a singleton.

```

\* **Did you know (#2)?** Large dragonflies have an average cruising speed of about 10 mph, with a maximum speed of about 30 mph. – “Frequently Asked Questions about Dragonflies”, *British Dragonfly Society*, 2013.





```

9   GameManager (GameManager const&); // Don't allow copy.
10  void operator=(GameManager const&); // Don't allow assignment.
11  bool game_over; // True, then game loop should stop.
12  int frame_time; // Target time per game loop, in milliseconds.
13
14  public:
15  // Get the singleton instance of the GameManager.
16  static GameManager &getInstance();
17
18  // Startup all GameManager services.
19  int startUp();
20
21  // Shut down GameManager services.
22  void shutDown();
23
24  // Run game loop.
25  void run();
26
27  // Set game over status to indicated value.
28  // If true (default), will stop game loop.
29  void setGameOver (bool new_game_over=true);
30
31  // Get game over status.
32  bool getGameOver () const;
33
34  // Return frame time.
35  // Frame time is target time for game loop, in milliseconds.
36  int getFrameTime () const;
37  };

```

#### 4.4.5 Development Checkpoint #2!

If you have not kept up already, **Dragonfly** development should continue! Steps:

1. Create the Clock class. Create a `Clock.h` header file based on Listing 4.19. Add `Clock.cpp` to the project and stub out each method so it compiles.
2. Implement and test, using a simple program that creates a Clock object, waits for awhile (use an appropriate sleep call), and calls `split()` and/or `delta()`. Verify the times meet expectations. A robust LogManager (developed during Development Checkpoint 4.3.7) can be used for output.
3. Create the GameManager class. Create a `GameManager.h` header file based on Listing 4.24. Add `GameManager.cpp` and stub out each method so it compiles.
4. In the `GameManager.cpp` file, have the `startUp()` method start the LogManager, and the `shutDown()` method stop the LogManager and call `setGameOver()`. Test that `startUp()` and `shutDown()` work as expected before proceeding.
5. Implement the game loop inside the GameManager `run()` method. The body of the loop does not do anything yet (although you can add some “dummy” statements), but the loop should time (via `delta()` and `split()`) and sleep properly. Be sure to



double-check any conversions of units (e.g., milliseconds to microseconds) used. The game loop uses a Clock object. Test thoroughly by timing (with a clock on the wall) that you get the expected number of loop iterations.

6. Add additional functionality to the GameManager, as desired. The frame time option to `startUp()` can be useful.

Since code developed during this Development Checkpoint drives the entire game, it should be tested thoroughly, making sure it is robust and clearly written before proceeding.

