Practical Profiling

Handout for practicum 1 for INFOMOV - updated for the 2019/2020 edition.



Introduction

A crucial step in the optimization process is profiling. This document takes you through the basics of measuring application performance.

Situation

Our starting point is <u>a working application</u>, and for the purpose of this document, we will assume all intended functionality has been implemented. A minimal performance indicator is shown in the top-left corner. It fluctuates between values close to multiples of 16, which is caused by the use of *GetTickCount* for performance measurements: this timer has a 16ms resolution.

Goal

Our aim is to gain insight in the performance of this application:

- which sections take most time to execute;
- how does performance change over time.

This information will allow us to focus optimization effort on just the expensive parts of the code.

Very Sleepy

We can directly use the profiling tools supplied with Visual Studio. However, their functionality differs greatly between versions, and some important functionality is missing from the free versions of Visual Studio. Luckily, a free alternative exists: *VerySleepy*. Run the tool, and select **Launch** from the file menu:

elect a process to profile:					Select thr	ead(s) to profile: (CTF	L-click for multi	ple)		
Process	Туре	CPU 👻	PID	^	Thread	Location		Thread Usa	ge 👻	
nvcontainer.exe	32-bit	0%	4104						-	
sihost.exe	64-bit	0%	4164							
svchost.exe	64-bit	0%	4172							
taskhostw.exe	64-bit	0%	4244							
mbamtray.exe	Launch an EXE							×		
RuntimeBroker.exe explorer.exe	Enter a command to execute, with any additional arguments.									
ShellExperienceHost.exe SearchUl.exe TSVNCache exe	Working direc	tory.						-		
nvtray.exe										
<						Launch	Cancel			
Refresh 📀 Download	Ŧ		Profile f	or set time	(s) 100		Prof	ïle All →	Profile Selec	ted 🕨
										1

The executable to run is in the project folder of the balls application. This is also the working folder; make sure to specify this as well. Then click Launch. Make sure you profile the release version of the code: the debug version will be much slower, and may have different performance characteristics.

File View Help		_						_		Y	~
Name	Evely	-	Inclusive	% Evelveive	9/ In chusing	Madula	Course File	^	Averages	Call Stacks	F
Name	exciu	*	o 75	/ Exclusive	/o inclusive	WOULLE	Source File		Called From		
memcpy	0.755		0.755	0.80%	0.80%	VCRUNTI	[unknown]		Name		
top4016021	0.015		0.015	1.00%	3.30%	Tenni2017.02	[unknown]				
[00401641]	0.53c		0.53c	4.90%	4.90%	Tmpl2017-02					
[0040712A]	0.335		0.335	3 3296	3 3 2 9%	Tmpl2017-02					
[0040162C]	0.375		0.334	2 97%	2 97%	Tmpl2017-02					
[00401506]	0.245		0.245	2.16%	2.16%	Tmpl2017-02					
[004015E0]	0.235		0.23s	2.07%	2.07%	Tmpl2017-02					
[00403E7D]	0.225		0.22s	2.00%	2.00%	Tmpl2017-02					
[004015E3]	0.21s		0.21s	1.94%	1.94%	Tmpl2017-02					
[004015C8]	0.20s		0.20s	1.82%	1.82%	Tmpl2017-02				_	
[004015E7]	0.20s		0.20s	1.81%	1.81%	Tmpl2017-02			<		
[00401604]	0.19s		0.19s	1.71%	1.71%	Tmpl2017-02			Child Calls		
[00401759]	0.18s		0.18s	1.67%	1.67%	Tmpl2017-02			Name		
[00402136]	0.18s		0.18s	1.65%	1.65%	Tmpl2017-02					
[00401649]	0.16s		0.16s	1.49%	1.49%	Tmpl2017-02					
T0040212C1	0.16s		0.16s	1.49%	1.49%	Tmnl2017-02		Υ.			

An

Source file

To improve this, we need to include *debug information* in the executable. Consult the Lecture 1 slides for details on how to do this. After this change, we get the following output:

Line 1

Q Very Sleepy CS - C:\Users\JBIKKER\AppData\Local\Temp\AA0C.tmp										- 🗆 ×		
File View Help												
Functions										Averages Call Stacks Filters		
Name	Exclusive 🔻	Inclusive	% Exclusive	% Inclusive	Module	Source File	Source Line	Address	^	Main		
Tmpl8::Game::Tick	17.54s	26.15s	62.96%	93.85%	Tmpl2017-02	e:\dropbox\jacc	229	0x401365		Eunction Name		
BuildTree	4.19s	4.19s	15.05%	15.05%	Tmpl2017-02	e:\dropbox\jacc	202	0x401ff1		Madula Temp12017 02		
Tmpl8::Sprite::Draw	2.33s	2.33s	8.37%	8.37%	Tmpl2017-02	e:\dropbox\jacc	358	0x403cf3		Module Impi2017-02		
Tmpl8::operator*	0.34s	0.34s	1.23%	1.23%	Tmpl2017-02	e:\dropbox\jacc	26	0x404a90		Source File		
Tmpl8::Surface::Clear	0.33s	0.33s	1.20%	1.20%	Tmpl2017-02	e:\dropbox\jacc	82	0x402ea0				
radix8	0.05s	0.05s	0.18%	0.18%	Tmpl2017-02	e:\dropbox\jacc	126	0x40237b				
radix1	0.04s	0.04s	0.15%	0.15%	Tmpl2017-02	e:\dropbox\jacc	111	0x402294				
sprintf	0.01s	0.01s	0.04%	0.04%	Tmpl2017-02	c:\program files	1783	0x402475				
Tmpl8::Surface::Line	0.01s	0.01s	0.03%	0.03%	Tmpl2017-02	e:\dropbox\jacc	172	0x403173				
Tmpl8::Surface::Print	0.00s	0.00s	0.01%	0.01%	Tmpl2017-02	e:\dropbox\jacc	93	0x402dab				
operator new	0.00s	0.00s	0.01%	0.01%	Tmpl2017-02	f:\dd\vctools\crt	19	0x405534				
Tmpl8::Surface::CopyTo	0.00s	0.00s	0.01%	0.01%	Tmpl2017-02	e:\dropbox\jacc	221	0x403338				
memchy	0.00¢	0.00¢	0.01%	0.01%	Tmnl2017-02	[unknown]	0	0v40657c	~			
Source Log	Source Log											
· · · · · · · · · · · · · · · · · · ·									^			
<pre>void Surface::Clear(Pixel a_Color) { int s = m_Width * m_Height;).33s for (int i = 0; i < s; i++) m_Buffer[i] = a_Color; }</pre>												
<pre>void Surface::Centre(char* a_String, int y1, Pixel color) { int x = (m_Nidth - (int)strlen(a_String) * 6) / 2; Print(a_String, x, y1, color); }</pre>												
			(tine	1		>				

Note the filter on the right: we are not interested in any code outside our executable. Also note that VerySleepy is not just able to extract function names from the executable, but also the *full source* code. Let this be a warning: never distribute an executable with this information!

A few words on the contents of the columns: an exclusive column (seconds, percentage) shows the time spent in the function. Here, 62.96% of the time was spent in Tmpl8::Game::Tick. The inclusive column includes time spent in functions called by Tick. Only the values in the exclusive columns thus sum to 100%.

The source window provides even more detailed information. Some lines have an indication of the time spent at that line. Not all lines have this information: VerySleepy gathers its information by

periodically checking at which address the CPU is executing code. The time for a line is thus always an estimate, based on the number of times VerySleepy polled that line. Running the application longer will increase the number of lines that have an estimate, but in general, lines without a time estimate may be considered to be not very expensive.

Interpretation

Based on our measurements, we can conclude that the Tick function is by far the most expensive function (even exclusive time spent in called functions). Especially the section under "verlet: satisfy constraints" takes a lot of time. Why a line like const vec3 d = p[idx] - p[i]; takes so much time is a topic for later, but one thing is clear: we should not be focusing on Sort(), nor on BuildTree, at least for the moment. A 35-line section in Tick is all that matters.

Performance over Time

While VerySleepy can report overall application performance and pinpoint bottlenecks, it cannot tell us how application performance is over time. It is for example perfectly possible that the Sort or BuildTree functions dominate frame time for one or two frames, even if they take little time on average. Or perhaps 'satisfying constraints' is far more expensive in some frames.

To get more insight we can instrument the code. The template provides a convenient high-resolution timer, simply called 'timer'. Two methods are of importance to us: reset() and elapsed().

We add a timer to the start of Game::Tick:

timer tm;

Now, we can obtain a split time in milliseconds at any time:

float measurement1 = tm.elapsed();

After gathering a few measurements we can report the timings. It is important *not* to report while timing: reporting may have considerable overhead. So, at the end of Game::Tick, we add:

```
printf( "measurement 1: %f\n", measurement1 );
printf( "measurement 2: %f\n", measurement2 );
```

Some results are shown below to the right. The four columns are for satisfy constraints, sorting, render and finalize.

A few things pop out:

- The first frame is expensive.
- Satisfying constraints quickly goes down: after a second or two it is down to 10ms. Then, after a few more seconds, it goes up again.
- Sorting quickly settles for 0.21ms.
- Rendering starts low, but takes almost 4ms once all the balls have entered the screen.

Watching the text roll by is somewhat inconvenient. We can improve this by drawing a simple graph.

Т	E:\Dropbox\Jacco\lectures\INFOMO	/\examples\balls\	Tmpl2017-02.exe
---	----------------------------------	-------------------	-----------------

applicat:	ion star	ted.		
timings:	23.93	0.40	0.35	0.57
timings:	22.74	0.26	0.28	0.49
timings:	26.24	0.29	0.30	0.69
timings:	21.09	0.25	0.28	0.41
timings:	17.23	0.21	0.14	0.44
timings:	17.98	0.29	0.25	0.51
timings:	15.26	0.24	0.36	0.81
timings:	19.69	0.28	0.34	0.67
timings:	19.61	0.29	0.33	0.58
timings:	17.19	0.27	0.26	0.44
timings:	14.35	0.27	0.20	0.57
timings:	15.77	0.35	0.30	0.50
timings:	15.91	0.26	0.27	0.60
timings:	19.29	0.30	0.51	0.70
timings:	15.63	0.29	0.30	0.63
timings:	16.89	0.27	0.30	0.56
timings:	14.75	0.21	0.15	0.69
timings:	17.93	0.30	0.19	0.76



Here, red is the time for 'satisfy constraints' (10 pixels per millisecond), green the time for sorting, blue the time for render and yellow the time for finalize. The two dips happened when the application was interacted with, causing a 'splash'.

The graph makes performance over time very easy to interpret:

- Start-up is slow, but it gets worse after the first 200 frames;
- Rendering starts fast but converges to a stable time slice;
- Some frames have significant peaks.

Based on this information, it is probably worthwhile to investigate *why* satisfying constraints is cheaper for some frames: perhaps it is possible to find a better algorithm that meets or beats the best time in all situations. This means we have high-level optimization work to do, before we get to low-level optimizations.

TODO

To gain additional insight, we can add additional graphs for the usual suspects:

- count the number of expensive operations (division, square root) to see if there is a relation with raw performance;
- count the number of memory operations (array accesses) to see if the application is held back by those;
- count the number of traversal steps in the binary tree;
- count the number of collisions.

In each case this helps us to identify what consumes time: if e.g. the number of tree traversal steps has no relation to performance, it makes no sense to focus optimization on this part of the code.

Get Active

Take some time to practice the material presented in this document, in case you haven't done so already:

- Replicate the VerySleepy results to make sure the tools are working for you and to get familiar with the workflow.
- Add the high resolution timer. Figure out what the resolution of this timer is, i.e. what is the smallest amount of time for which you get accurate results?
- Implement the graph, and operate the application to see the impact on various parts of the code.
- Add additional graphs to increase your understanding of what affects performance in this particular application.

In case you are not familiar with the template that is being used here:

screen->Line(x1, y1, x2, y2, color) can be used to draw a line for the graph. The screen is 1024x640 pixels, so the last line is at y=639, or simply SCRHEIGHT - 1. Similarly, screen width is stored in SCRWIDTH.

Colors are specified as 32-bit ARGB. Useful colors are (hexadecimal) 0xFF0000 for red, 0x00FF00 for green, 0x0000FF for blue and 0xFFFF00 for yellow.

The End

Questions and comments:

bikker.j@gmail.com or room 4.24.



INFOMOV 2019