

## Chapter 36

# Efficiency

In many of my examples, I make a big deal about how to make the program look nice. I don't say much about how to make it more efficient – in other words, making it run faster.

As you write code, there are a lot of little “this way or that way” decisions. For a bunch of different reasons, if you make all these little decisions based on what seems like it would run faster, it doesn't work and turns the code into a mess.

There are things to do to avoid super-slow code, but most of them are when you make the plan – not when you're writing each line.

Some notes:

### 36.1 Most speed-ups don't work

A lot of the things that seem like they'd speed up a program just don't work. They do nothing, or make it run *slower*. Here's a list of why some things don't work:

#### The compiler rearranges things anyway

Writing an equation on one line or on several lines doesn't matter. The computer breaks it into step-by-step anyway. Suppose you have this:

```
x = a*b + c*d;
```

The compiler changes it into:

```
int temp1 = a*b;  
int temp2 = c*d;  
x = temp1 + temp2;
```

By the same rule, declaring lots of extra temporary local variables doesn't hurt you. The computer was going to fake-declare them anyway.

`if` expressions are the same way. For real, the computer can only have one thing in an `if`:

```
// you type this:
if(n>=1 && n<=10)

// compiler turns it into:
if(n>=1)
    if(n<=10)
```

Or's are broken up using a different trick (it's hard to show, but it works).

The compiler might even flip the test. If you write `if(n>0) A else B` the compiler might turn it into `if(n<=0) B else A`.

So go ahead and write `if`'s and conditions how-ever it looks nicest. If there was a faster way you could have re-arranged them, you may as well let the compiler find it.

## The compiler knows the speed-ups

Over the years, we've found a lot of mini-tricks to rearrange code to speed things up. And we've programmed them, and put them into compilers. In other words, if you can look up a speed-up trick, and it works, it's probably already in the compiler.

The computer will pre-do math with constants. For example `n=7*24`; is fine to get hours in a week. The compiler turns it into `168`. When the program runs, the line is just `n=168`;

In general, leave math in whatever form is easiest to read. I like to write `Random.Range(1,6+1)`; to remind me it's going to 6, with the 1-less rule. It's not slower, since the compiler makes it a 7 before the program ever runs.

The compiler can re-arrange to factor out common terms. For example:

```
int n = x+b*3-1;
int m = (y+b*3-1)/2;
```

You might notice both use `b*3-1`. We could compute that once and re-use it:

```
// "improved" version:
int q = b*3-1;
int n=x+q;
int m = (y+q)/2;
```

But a good compiler does this for us. Over 50 years very smart people have written papers about simplifying equations, which made their way into compilers.

If the code would look better, go ahead and precompute. Otherwise let the compiler handle it.

## Speed-up logic is hard

This is an old example of a non-working attempt at a speed up. Suppose resetting a game round resets `health=100;`. That's a waste if our health was already 100. We'll improve things by checking first: `if(health!=100) health=100;`

But `if`'s take time, too. The old way took 1 step. This way takes at least 1 step, usually 2.

But suppose resetting health to 100 also resets the health bar. Would an `if` here save time?:

```
if(health!=100) {
    health=100;
    redrawHealthBar(); // takes 30 steps
}
// else the health bar is already displaying 100 health
```

If still takes one extra step if the player is damaged. If that's the usual case, this way is still slower than without the `if`'s.

## 36.2 Premature Optimization

Let's suppose we can write each part of the program in two ways: runs faster, or easier to read – one or the other. Another way to say the same thing, suppose we can rewrite any part to be a little faster, but more confusing.

When would we decide to do it which way?:

- If we need to test it and are pretty sure there will be some things that need fixing, easy to read is better.
- If we're going to expand it later (for example, a match game using colors and we probably want to add matching shapes) then easy-to-read is better.
- Often part of our code is temporary or experimental – we're 70% sure it will eventually be replaced by something better. There's not much point making it faster.
- If we get a lot of requests from customers for changes in this part, easy to read is better.

- If this part is done, tested and you're confident it won't need to be changed, running faster is an option.

An example, here's my simple "keys move left and right" code:

```
if(Input.getKey(KeyCode.A) pos.x-=0.1f;
if(Input.getKey(KeyCode.D) pos.x+=0.1f;
pos.x=Mathf.Clamp(pos.x, -7, 7); // runs even if we don't move
transform.position=pos; // ditto
```

When we don't move, this code uselessly makes sure `pos.x` is between -7 and 7, and uselessly resets our position. We could speed it up by only running the last two lines after we moved:

```
float oldX=pos.x;
if(Input.getKey(KeyCode.A) pos.x-=0.1f;
if(Input.getKey(KeyCode.D) pos.x+=0.1f;
if(pos.x!=oldX) {
    pos.x=Mathf.Clamp(pos.x, -7, 7);
    transform.position=pos;
}
```

This runs slower! When we're moving, this is two lines *more* than the old code. It runs a little faster when we're not moving – but who wants a game that speeds up when you stop?

But I can keep trying. Here's my next attempt which really is faster:

```
if(Input.getKey(KeyCode.A) {
    pos.x-=0.1f; if(pos.x<-7) pos.x=-7;
    transform.position=pos;
}
else if(Input.getKey(KeyCode.D) {
    pos.x+=0.1f; if(pos.x>7) pos.x=7;
    transform.position=pos;
}
```

This always runs faster, but I think it's harder to read (two lines setting position, extra 7's).

But wait. We weren't done making the game. We want to slide sideways for a few frames after you let go of a key. Then we want wind and bumping to also move the player. Those are going to be total rewrites. Time spent speeding up the version above was just wasted.

### 36.3 Not all code sections are equal

Suppose you have one player and thirty enemies (or anything using a script with Update). For speed, the enemy code is 30 times more important. In other words, if you think of 30 little ways to speed up the player's code, you could have thought of just one way to speed up the enemy code.

Suppose you have a hundred lines in a script, then a nested loop around two lines. Maybe it checks every square in a 100 by 100 grid:

```
do stuff #1
do stuff line #2
...
...
...
do fortieth thing
...
do fiftieth thing
...
do one hundredth thing

for(int i=0;i<100;i++)
  for(int j=0;j<100;j++) {
    loop line #1
    loop line #2
  }
```

The inside of the loop runs 10,000 times. Compared to the hundred lines before it, the loop is 99% of the total time.

Making those hundred lines ten times as fast would be less than a 1% speed-up. Improving the lines in the loop is all that matters.

You can also have a huge loop even if you don't have that many items. Suppose we have an array of 100 foods and compare every one to every other. That's another ten thousand times loop.

Suppose the inside has a function call that also loops over the array. Everything in that function loop runs  $10,000 \times 100 =$  a million times. From just a size 100 array.

Accidentally writing a triple-nested loop isn't all that uncommon. A lot of making code run faster isn't trying to speed up good code – it's looking for places like unnecessary triple loops.

But even more, there are lots of places where you don't care about speed at all: the Start / Pause / GameOver screens. Scrolling text. The code to fix player names longer than 20 characters. Calculating your "rating" after you finish a level (there's a delay and a little spinning animation, anyway).

## 36.4 Is it already fast enough?

There are places where obviously we care about speed, but once we get it fast enough, there's no point getting any faster.

Almost any menu-based web program only has to run faster than human reaction speed, which isn't very fast at all. Humans can't click buttons more than ten times a second. If you can respond in 1/10th of a second (which is forever to a computer), any faster won't matter.

For a game, screens update at 30 or 60 frames/second. We have at least 1/60th of a second to do all the work for the next frame display. If we're making a matching game, that's plenty of time. "Speeding it up" will have no visible effect.

## 36.5 Will it cause more errors?

"Faster" code tends to be more difficult to read and more complex, which means it tends to have more, harder to fix, errors. Beyond that, some speed-up tricks add new possible errors.

A fun one is how the caching trick can add a new out-of-synch error. A typical example of caching is a link to some other Cube, with a shortcut link to it's material:

```
public GameObject otherCube; // this is pre-set
Material otherCMat; // we compute shortcut to Material of otherCube

void Start() {
    otherCMat = otherCube.GetComponent<Renderer>().material;
}

void Update() {
    otherCMat.color=c2; // using the shortcut
    // otherCube.GetComponent<Renderer>().material.color=c2; // the long way
```

The new problem is that we can aim `otherCube` at something else, while forgetting to recompute `otherCMat`. To an unsuspecting tester, that error makes no sense – it shouldn't even be possible. Which means it can be extra-difficult to fix.

## 36.6 Things that work

In general, focus on high-use parts of the code - for example things that run for 50 items, every frame.

A main way of speeding things up is getting rid of accidental nested loops. They can be hidden, like a loop calling a function which runs another loop. A common Unity example is `GameObject.Find`. That's a loop through everything in the game – maybe a few hundred things and not too bad. But if you forget, you might put in inside another loop and it blows up into tens of thousands of steps.

`List`'s are a good way to have accidental nested loops. I mentioned they're really arrays and doing anything not at the back involves a slide-loop. If you don't know this, you might add to the front, `L.Insert(0,w)`; , in a loop. Adding 100 items that way is 5,000 steps, compared to `L.Add(w)` at only 100 steps. Remove is the same way. Removing the last item is 1 step, removing the first is a slide-loop to shift everything down.

Other tricks involve looking at your plan.

For example, code handling monsters often checks line-of-sight to the player, and `Raycast` is doing lots of math. Instead of running it every frame – 60 times/second – we can run it twice a second. That's a 30x speed-up. It will cause monsters to sometimes have up to a 1/2-second reaction delay. That's fine. In fact, it makes them feel more realistic.

Path-finding (computing a path through a maze or building) is crazy slow. You can often run that every few seconds, or only when something changes; and have some monsters just follow others.

Many actions on lists are much faster if the list is sorted. We can often keep lists sorted as we add items. This takes lots of practice and some theory

`LinkedLists` are the other way of storing items in a row. They're very fast to insert and remove from any position, just as fast to loop through end-to-end as arrays, but much slower than arrays if you need to jump around with look-ups. Choosing to use an array or a linked-list can make a big difference, but knowing which takes lots of practice.

Using less space will speed up a program. Memory is divided into high speed, medium speed and regular. The more of your program which fits into the high and medium speed, the faster it runs.

This is partly why Unity uses floats instead of doubles – less space. But it's still a matter of finding the places where it matters. If you have one especially huge array, cutting the size in half might help; but don't waste much time making small arrays smaller.

The last things we can do are making little tweaks to the ultra high-use code parts. If you work like a dog and are lucky, you might get a 2% speed-up.

`v1.magnitude` (length of a vector) vs. `v1.sqrMagnitude` is an example. Finding the distance ends with a square root. The second command leaves it out, to be a bit faster. You could change `if(v1.magnitude<3)` to `if(v1.sqrMagnitude<9)` for a very small speed-up (maybe).