


ASS. 04

- Basic image output program:



- want to split image into horizontal pixel bands so that we can write them independently, IN PARALLEL

- Seems straightforward enough:

- split rows into one
band per CPU core,

- each core writes own
row

- let's say we have n cores

- then split image into


bands: h/n cores

- first obstacle:

- can't write to
file in parallel

- but can write to
memory in parallel,

so long as cores
don't write to the
same pixels

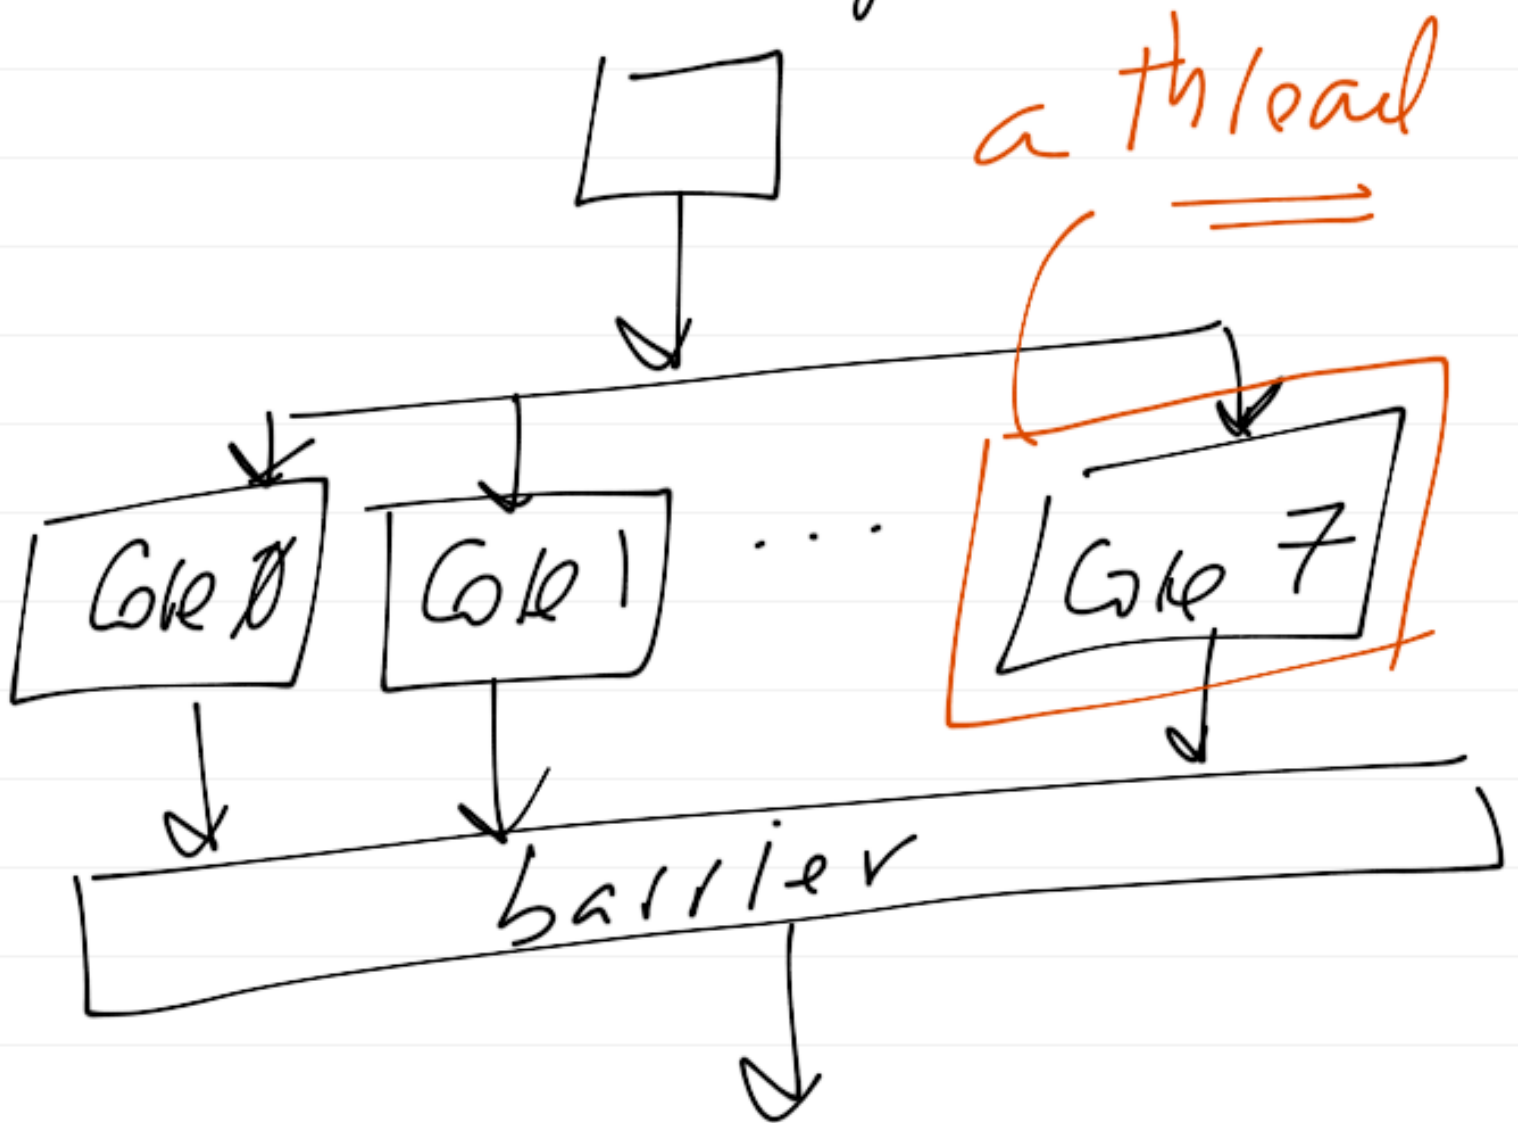


Core \emptyset → $y : \emptyset$ to $\frac{h}{n \text{ cables}}$

Cable 1 → $y : \frac{h}{n \text{ cables}} + 1$ to $\frac{2h}{n \text{ cables}}$

and so on ...

- program flow of control:



- not all cores may finish at the same time

- all can access **shared memory**

- shared memory vs. private memory
- if a core uses index (x, y) , it cannot share this info with any other core
- each core gets a private copy of (x, y)
⇒ some variables must be private

- copying memory among
cores is expensive

- something like our image

(our light \times ing array)

should be shared - no
copies made to cores

- just need to ensure

each core writes to

non-overlapping memory

- parallelizing code is
surprisingly easy, BUT
have to be careful about
memory access - have
to specify what's private
and what's shared

- parallelization on multi-
core architecture done
via OpenMP

- let's do ASG 04 serially first (all that's required)
- parallel part is optional
- OpenMP won't work with our version of Objective-C (gnustep)
- our gnustep install is not thread-safe

- serial version of Code:

```
pixel_t *red =  
    [(pixel_t *) [pixel_t alloc]  
    : init: 1.0: 0.0: 0.0];
```

```
sample *gln (0, 1, 0)
```

```
*bln (0, 0, 1)
```

```
*yln (1, 1, 0)
```

```
*cyn (0, 1, 1)
```

```
*mgn (1, 0, 1)
```

```
*wht (1, 1, 1)
```

```
*blk (0, 0, 0)
```

pixel_t *pix = ... alloc/init
to 0,0,0

int tid, i, x, y;

int w = 640, h = 480;

rgb_t *imgloc, *img;

timer_t *timer =

[(timer_t *) (timer_t alloc)
init];

img = (irgb_t *) malloc
(w * h * sizeof(irgb_t));

[timer_start]; // starts

timer —
just records
current time of day

```
//main loop
```

```
for (y = h-1; y >= 0; y--) {  
  //go back to simulate 8 grees
```

```
  tid = y / (h/8);
```

```
  for (x = 0; x < w; x++) {
```

```
    [pix autorelease];
```

```
    switch (tid) {
```

```
      case 0: pix = [red copy];
```

```
      break;
```

```
      case 1: pix = [green copy];
```

```
      break;
```

```
⋮  
case 7: pix = [blk copy];  
break;  
} // switch
```

```
imgloc = img + y * W + A;
```

```
for (i = 0; i < 3; i++)
```

```
  (*imgloc)[i] =
```

```
  [ [ pix scale: 255.0 ] get : (i) ];
```

```
} // inner for  
} // outer for
```

[timer stop];

printf(stderr, "Gles: 1 \n");

printf(stderr, "%2.2g mrc\n");

[timer elapsed - mrc]);

// write out image

fprintf(stdout, "P6 %d %d 255 \n",
w, h);



```
for (y = h - 1; y >= 0; y--) {  
    for (x = 0; x < w; x++) {  
        imgloc = img + y * w + x;  
        fwrite (imgloc,  
                sizeof (right_t), 1, stdout);  
    }  
}
```



// free memory

[red release];

[green release];

[blue release];

free(img),

[pool drain];

}

- pretty basic program
- want to parallelize
first nested loop
(writing to memory)

Get out the second
(writing to file)

- first, need:

```
#import <omp.h>
```

- then,

```
int chunk, ncores = 1;
```

```
#pragma omp parallel private(tid)
```

```
{  
  tid = omp_get_thread_num();
```

```
  if (tid == 0)
```

```
    ncores =
```

```
    omp_get_wm_threads();  
}
```

chunk = h / ncores; } this
will

split up rows
per core

⋮
[timer start];

```
#pragma omp parallel \
    shared (chunk, w, h, img) \
    private (tid, i, x, y, pix, imgloc)
```

}

```
tid = omp_get_thread_num();
```

```
#pragma omp for  
schedule (static, chunk)
```

```
for (y = b-1; y >= 0; y--) {
```

// pragma parallelizes
just the above for
statement — unrolls
loop & splits up chunk
iterations per core

- rest of code the
same except for
closing } that closes

of #pragma omp parallel

block just above

[timer stop];

- in theory, code runs

more times faster

(in practice there is some overhead)

- what's the timer_t?

```
@interface timer_t : NSObject
```

```
{ double ts; // time start
```

```
double te; // time end
```

```
double tt; // time total
```

```
} (end - start)
```

- init timer to 0.0 for
each of t_s, t_e, t_t

- use this function to timestamp:

- (double) stamp_us

```
{ double s, us, tod;
```

```
  struct timeval tp;
```

```
  gettimeofday(&tp, NULL);
```

```
  s = (double)(tp.tv_sec);
```

```
  us = (double)(tp.tv_usec);
```


$t_{od} = 5 * 1000000.0 + us;$

time of day in microsec

return (t_{od});

}

- need

#include <sys/time.h>

for getting of day

- best of the functions
we use:

- (void) start

```
{ts = [self stamp_us];}
```

- (void) stop

```
{te = [self stamp_us];
```

```
tt = te - ts;
```

```
}
```

- (double) elapsed_us
{ return tt; }

- (double) elapsed_ms
{ return tt / 1000.0; }

- (double) elapsed_s
{ return tt / 1000000.0; }

- for openMP to work,
need to compile
with `-fopenmp` flag

- pretty easy - BUT, for
obj-c to work, each
thread should have its
own autorelease pool

- our version of gsubstep
doesn't support this

- However, Quartz does support some fancy NS code

- let's redo the serial image output program using NS objects:

NSSize

NSBitmapImageRep

NSImage

NSData

- We can use these to write out JPEG images instead of our .ppm ones

- to compile, need:

CFLAGS = `gnustep-config

--objc-flags

↳ backslash greater

LDLIBS = `gnustep-config

--gui-libs

- on Darwin (Mac):

LDLIBS =

- framework Foundation

- framework AppKit

- at top of main.m:

```
#import <Foundation/Foundation.h>
```

```
#import <AppKit/NSImage.h>
```

```
#import <AppKit/NSGraphics.h>
```

```
# ifdef __APPLE__  
# ifdef NSINTEGER  
# define NSUInteger  
    unsigned int
```

```
# endif
```

```
# endif
```


int p, 4, w=640, h=480;

NSize size = NSMakeSize(w, h);

NSBitmapImageRep *bitmap =

[[NSBitmapImageRep alloc]

initWithBitmapDataPlanes: NULL

pixelsWide: size.width

pixelsHigh: size.height

bitsPerSample: 8

samplesPerPixel: 3



has Alpha: NO

is Planar: NO

Color Space Name :

NS Device RGB Color Space

Bitmap Format: 0

bytes per row: size width * 3

bits per Pixel: 24];

NSUInteger red[3] = {255, 0, 0};
NSUInteger green[3] = {0, 255, 0};

⋮

UIImage *img = [UIImage alloc

initWithData:

[bitmap TIFFRepresentation]];

- how the for loops:

```
for (y = h-1; y >= 0; y--) {
```

```
    tid = y / (h/8);
```

```
    for (x = 0; x < w; x++) {
```

```
        switch (tid) {
```

```
            case 0:
```

```
                [bitmap set pixel: red
```

```
                at X: x
```

```
                y: y];
```

```
                break;
```

```
        } }
```

```
        // same for 114, 111, ..
```

NSData + jpegdata =

[hitmap representation Using Type:

NSJPEGFileType

properties: nil];

[jpegdata writeToFile:

@"/image.jpg"

atomically: NO];

[img release];

[bitmap release];

{pool drain};

}

- how to parallelize?
- would need to setup pixel data before feeding it to bitmap.

- so...

- you can use the above
to make your own jpg
images

- a bit more portable
than PPMs (e.g., web
pages)