



Threads: Principles

ICS332
Operating Systems

Concurrent Programming

- **Concurrency**: the execution of multiple “tasks” at the “same” time
- College students mostly write non-concurrent, or **sequential**, programs
 - At any point, you could stop the program and say exactly which instruction is being executed and what the runtime stack looks like
 - And there is a single answer to all the above for all execution of your program at the same point in its execution
- In a concurrent program, you design the program in terms of **tasks**, where each task has a “life of its own”
 - Each task has a specific job to do
 - Tasks may need to “talk” to each other or “wait” for one another
 - Tasks can be in different regions of the code or in the same region of the code at the same time
- It’s a different way of thinking and of programming

Example #1: Make it Fast

- Consider an input array of 10,000 integers: { 23, 56, 7, 68, 68 ...}
- I want to output a boolean array where each element is true if and only if the corresponding integer in the input array is odd
{ true, false, true, false, false ...}
- Assume it takes one millisecond to test an integer value and update the output array
- **Sequential programming:**
 - Iterate through the array, which would take 10,000 milliseconds.
- **Concurrent programming:**
 - If I create 10 “tasks” that each compute 1,000 output values, i.e., 1/10-th of the work, each task takes 1,000 milliseconds
 - Now if I can execute these 10 tasks independently (on a 10-core CPU), the whole execution could take as few as 1,000 milliseconds, i.e., 10x faster
 - In practice, we can’t go quite 10x faster due to various overheads and bottlenecks (e.g., the memory)
 - But we should go much faster than sequential provided we have multiple cores (which we all do in this day and age!)

Example #2: Make it Responsive

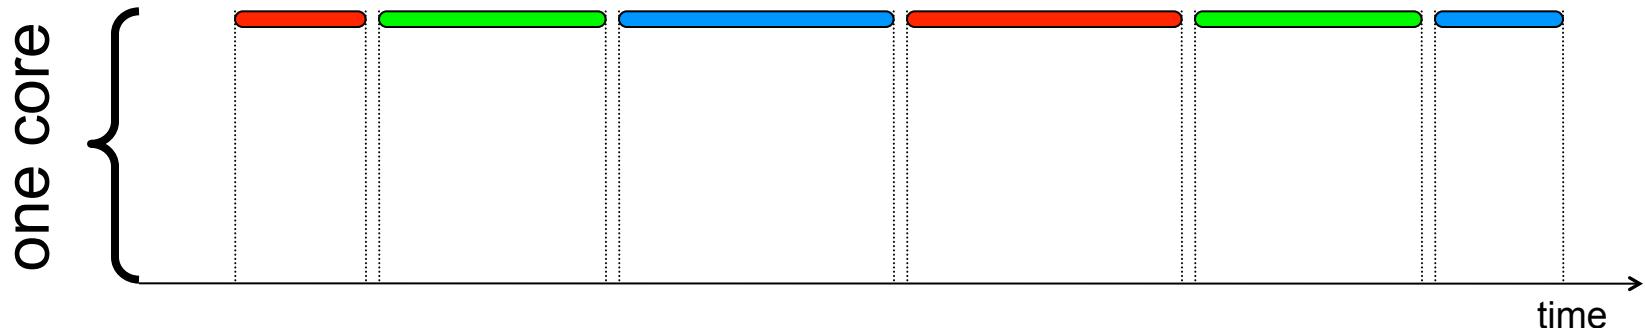
- Consider a Photoshop-like app in which a click of a button launches a transformation filter of all images that a user has selected
 - If many images are selected, this can take minutes
- **Sequential programming:**
 - While the transformation is happening, no other code can run, including the code that reacts to button clicks, meaning that the application is “frozen”, including whatever “Cancel” button one may have tried to implement
 - One solution, which is terrible, is to sprinkle “check whether the button is being clicked” code all over the code that performs the transformation
 - And it may not be feasible if that code is, for instance, a third-party library
- **Concurrent programming:**
 - Create a “task” in charge of watching buttons and reacting to clicks, which runs all the time
 - Whenever the user clicks on some “OK” button to perform the image transformation, create a “task” in charge of the transformation
 - Both tasks then run “at the same time”: while the image transformation is being performed, the user can still interact with the app (e.g., to quit!)

Why Concurrency

- The two previous examples illustrate the two main motivations for concurrency
- Make programs **faster**
 - Because multiple tasks can use different hardware components at the same time
 - e.g., while task #1 uses a core, task #2 uses another core, and task #3 uses the network card
- Make programs **more responsive**
 - While a task is blocked or doing something time-consuming, other tasks can still do their work
 - e.g., while a task is waiting for a network packet to arrive, another can animate a spinning beachball

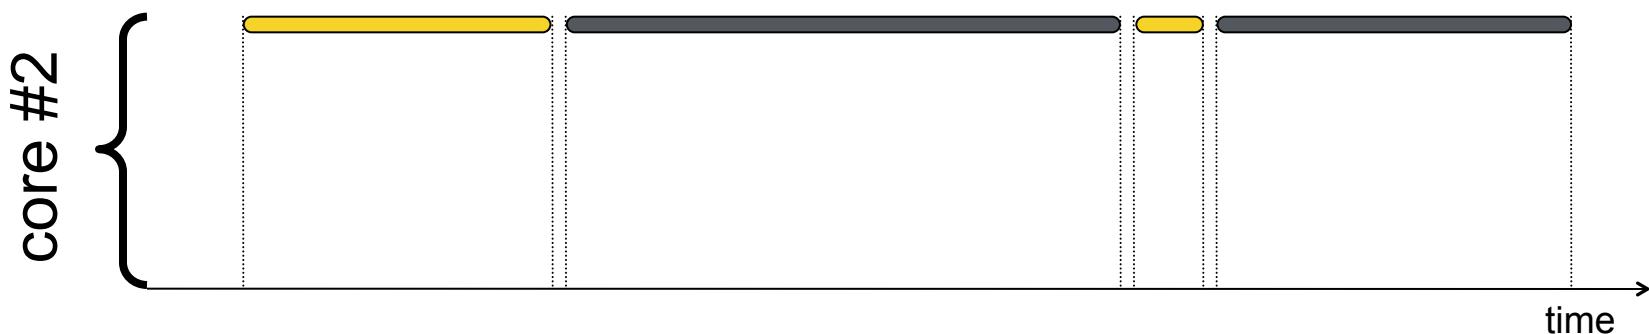
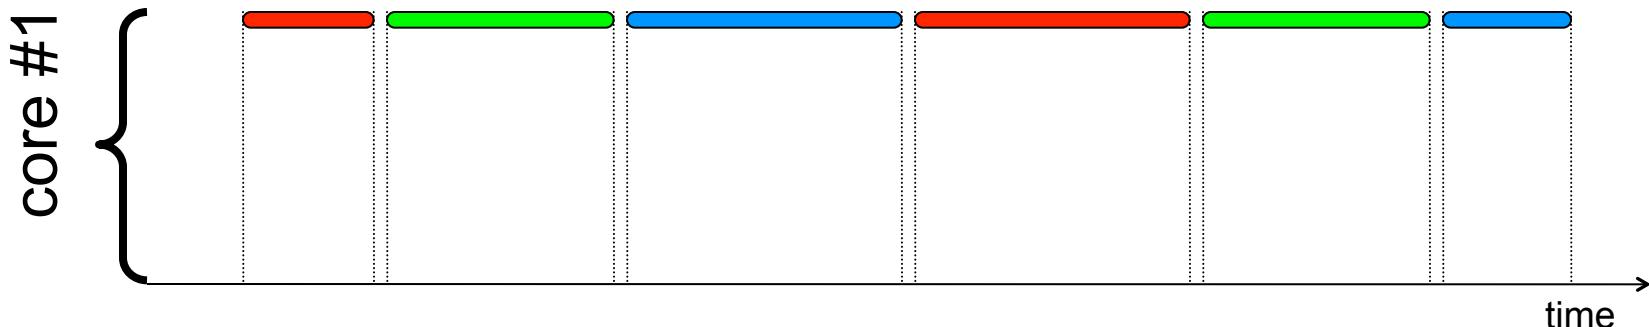
False Concurrency on One Core

- We now know that OSes use context-switching to alternate between processes/threads on a core
- This is known as **False Concurrency**
- Example (gaps = context-switching overhead):



- Provides the illusion of concurrency to a human because time quanta are short
- Increases core utilization because when a process/thread does I/O, the core is used by another process/thread

True Concurrency on Multiple Cores



- False concurrency within each core
- **True concurrency across cores**
 - e.g., the yellow and red threads sometimes experience true concurrency

True/False Concurrency

- The programmer should not have to care/know whether concurrency will be true or false
 - A concurrent program with 10 threads will work on a single-core processor, a quad-core processor, a 32-core processor, etc.
 - Typically you don't know on what kind of computer the program will run anyway
- A multi-threaded program will **reach higher interactivity with True and/or False concurrency**
- A multi-threaded program will **reach higher performance only with True concurrency**
- **Concurrency is not only about cores:** there can be concurrency between any two hardware resources
 - e.g., between the CPU and the Disk (a Web browser can have a thread that reads data from the disk and a thread that renders that data)
- A “let’s just add threads and things will be more interactive and faster” approach often works
- The OS makes it all transparent because it virtualizes the CPU

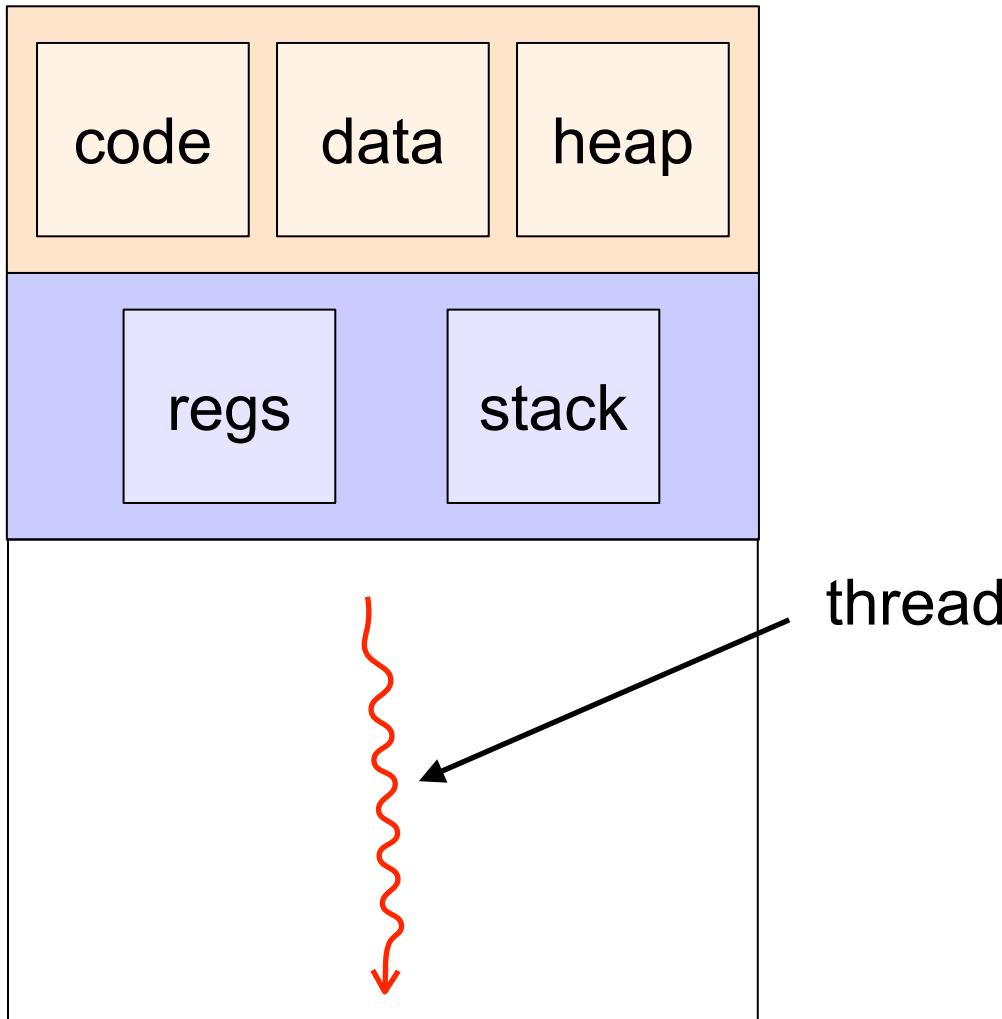
Concurrency with Processes

- We have already talked about concurrency
 - After all it's the 2nd “easy piece” in our textbook
- Processes run concurrently on the computer
 - They were used for concurrent programming a lot until the early 90's
 - And still used today
- But because the OS virtualizes memory, **by default processes don't share memory**
- We have seen that processes can communicate with IPC
 - *Message passing*: often not easy when processes have complicated cooperating behaviors
 - *Shared Memory*: often simpler, but requires many system calls and is cumbersome, up until the arrival of... threads!

Threads

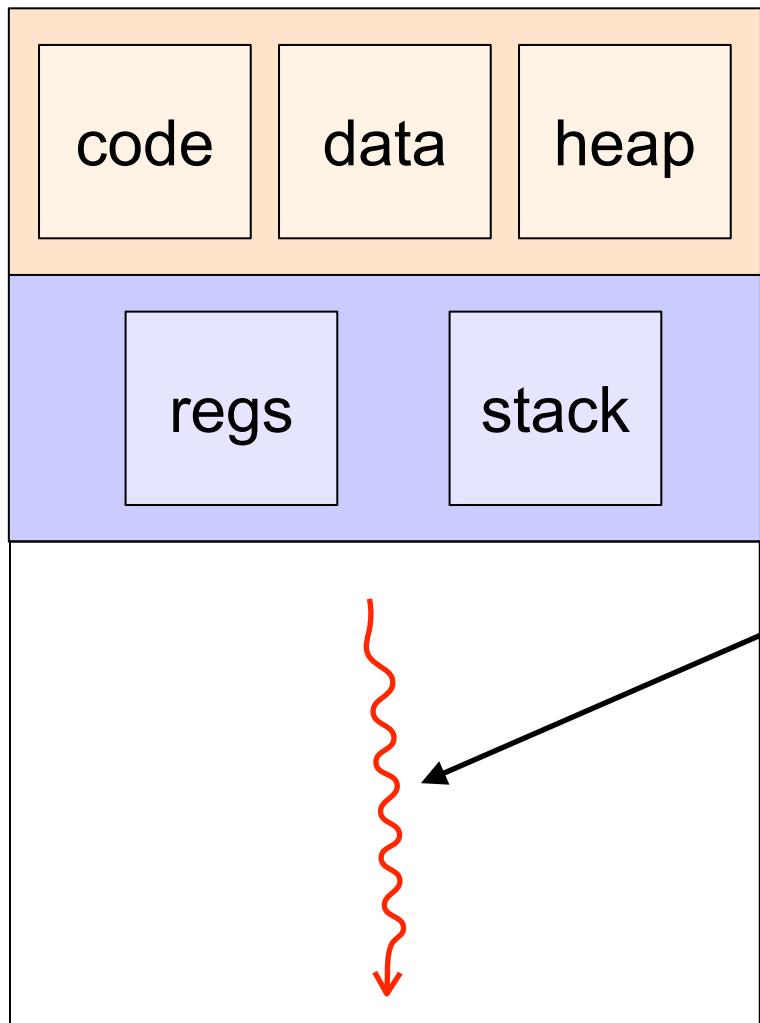
- A **thread** is a basic unit of CPU utilization within a process (i.e., it can be seen as a “task”)
- A **multi-threaded process**: Concurrent executions of different parts of the same running program, where each execution is a thread
- Each thread has its own:
 - Thread ID (assigned by the OS)
 - Program Counter (which instruction the thread currently executes)
 - Registers Set (which values are stored in registers)
 - Stack (bookkeeping of the thread’s function/method invocations)
- The above fully defines “what a thread is doing right now”
- But “within a process” threads share:
 - The code/text section
 - The data segment (global variables)
 - The heap
 - And other things (file descriptors, signal handlers, ...)

Threads: Typical Representation

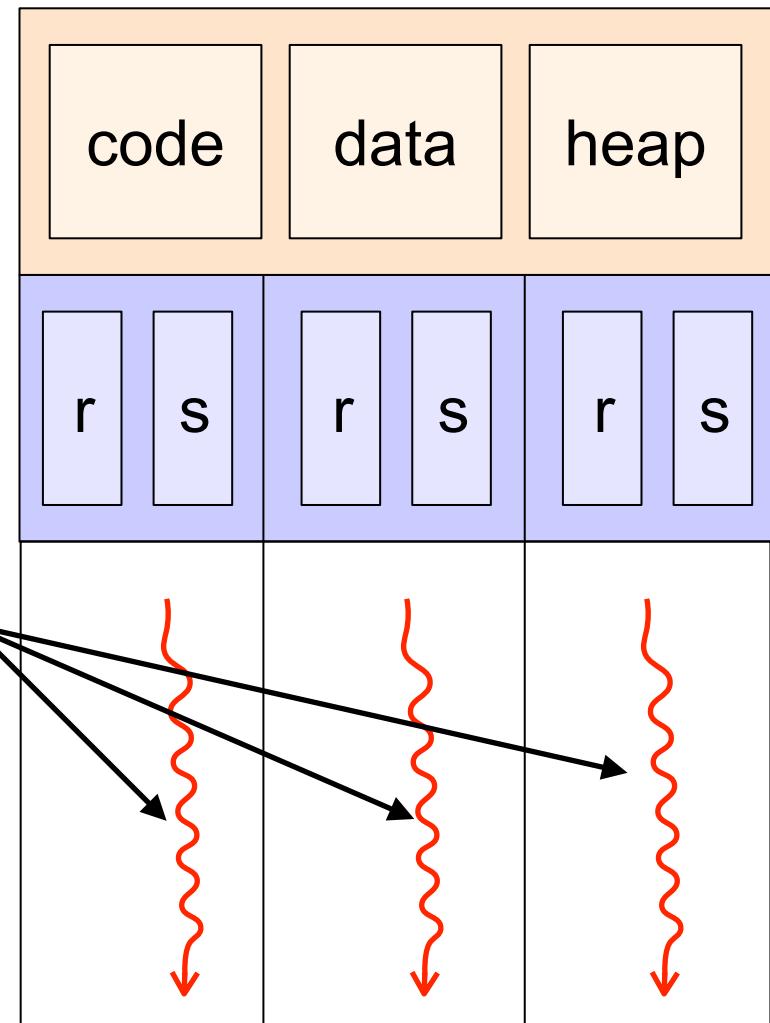


Single-Threaded Process

Threads: Typical Representation



Single-Threaded Process



Multi-Threaded Process

Multi-Threaded Program in Execution

```

Separate algorithms for Pluto and the rest. */
if (np < 1 || np > 9) {
    jstat = -1;
    for (i = 0; i < 5; i++) pvi[i] = 0.0;
    return;
} else {
    ip = np - 1;
}

/* Time since J2000. */
t = (double) (515445.1 / 365250.0);
/* OK start unless remote epoch. */
if (jstat < 0) {
    if (tbs < 2.0 || tbs > 5.0) {
        /* Compute the mean elements. */
        da = (p[0] * (1.0 + e[0] * (1.0 + (d[0] * t) * t) +
                      0.5 * dim(p[0]) + dim(p[1]) * dim(p[2]) * t * t) +
               d[0] * dim(p[0]) + dim(p[1]) * dim(p[2]) * t * t) *
             DASZR;
        de = (p[1] * (1.0 + e[1] * (1.0 + e[2] * (1.0 + (d[1] * t) * t) +
                      0.5 * dim(p[1] * (1.0 + e[0] * (1.0 + (d[0] * t) * t) +
                      dim(p[0]) * dim(p[2]) * t * t) +
               d[1] * dim(p[1]) + dim(p[2]) * t * t) *
             DASZR_D2P);
        d = (p[2] * (1.0 + e[2] * (1.0 + (d[2] * t) * t) +
                      0.5 * dim(p[2] * (1.0 + dim(p[1]) * dim(p[2]) * t * t) +
               d[2] * dim(p[2]) + dim(p[0]) * dim(p[1]) * t * t) *
             DASZR;
        dom = mod((3600.0 * omega(p[0]) + (omegap(p[1]) +
                      omega(p[2]) * t) * t) * DASZR_D2P);
    }
}

/* Apply trigonometric terms. */
dmu = (double) (1.0 / amas(p));
dmu = (double) (1.0 / amas(p));
if (t < 7.0) {
    arga = skd(p[0]) * dmu;
    argi = skd(p[1]) * dmu;
    da += (cal(p[0]) * cos(arga) +
           ip[0] * sin(arga)) * 1e-7;
    d += (cal(p[1]) * cos(argi) +
           ip[1] * sin(argi)) * 1e-7;
}
arga = op(p[0] * dmu;
arga += ((cal(p[0]) * cos(arga) +
           ip[0] * sin(arga)) * 1e-7;
for (j = 1; j < 9; j++) {
    argi = ksd(p[1]) * dmu;
    d += ((cal(p[1]) * cos(argi) +
           ip[1] * sin(argi)) * 1e-7;
}

da += d * D2P;
/* Daily motion. */
dm = G * amas(p);
/* Make first step. */
if (splatp) {
    if (j < 0) {
        /* Pluto. */
        /* ... */
        /* Time. */
        t = (da + dm * t) / amas(p);
        /* OK start. */
        if (jstat < 0) {
            /* Fundan. */
            d = (d + dm * t) * D2P;
            d += (d * d * t) * D2P;
            d += (d * d * t) * D2P;
        }
        /* Initialize coefficients and derivatives. */
        for (i = 1; i < 3; i++) {
            wbr[i] = 0.0;
            wbrd[i] = 0.0;
        }
    }
}

/* Term by term through Meeus Table 36A. */
for (j = 0; j < size(term / size(term[0]); j++) {
    /* Argument and derivative (radians, radians per century). */
    wj = (double) (term[j] * j);
    ws = (double) (term[j].is);
    wp = (double) (term[j].ip);
    /* ... */
}

```

```

/* Time since 2000.0 in Julian millennia. */
t = ( tdb - 515440.0 ) / 3652500.0;

/* ----- Topocentric terms ----- */

/* Convert UT1 to local solar time in radians. */
tsol = dmod( t, 1.0 ) * D2R - w;

/* FUNDAMENTAL ARGUMENTS: Simon et al 1994. */

/* Combine time argument (millenia) with deg/arcsec factor. */
w = t / 3600.0;

/* Sun Mean Longitude. */
elsun = dmod( 280.452665683 + 1296027711.03429 * w, 360.0 ) * DD2R;

/* Sun Mean Anomaly. */
emsun = dmod( 357.52910918 + 1295965810.481 * w, 360.0 ) * DD2R;

/* Mean Elongation of Moon from Sun. */
dmod = dmod( 297.8501954 + 160396899.89453 * w, 360.0 ) * DD2R;

/* Mean Longitude of Jupiter. */
elj = dmod( 34.35151874 + 109306899.89453 * w, 360.0 ) * DD2R;

/* Mean Longitude of Saturn. */
els = dmod( 50.07744430 + 44046398.47038 * w, 360.0 ) * DD2R;

/* TOPOCENTRIC ARGUMENTS: Moyer and Murray 1983. */
w0 = 0.0;
w0 += 0.00100e-10 * u * sin( tsol - 2.0 * emsun )
+ 0.00133e-10 * u * sin( tsol - 1.0 * emsun )
+ 0.00133e-10 * u * sin( tsol + emsun - elj )
+ 0.0230e-10 * u * sin( tsol + 2.0 * elsun + emsun )
+ 0.0230e-10 * u * sin( tsol + 2.0 * elsun + emsun )
+ 0.05312e-10 * u * sin( tsol - emsun )
- 0.13677e-10 * u * sin( tsol + 2.0 * elsun )
- 1.3184e-10 * v * cos( elsun )
+ 3.1769e-10 * u * sin( tsol );

/* ----- Fairhead model ----- */

/* T<sup>0</sup> */
w0 = 0.0;
for ( i = 474; i >= 1; -i ) {
    i3 = i - 3;
    w0 += fairhd[i-3] * sin( fairhd[i-2] * t + fairhd[i-1] );
}

/* T<sup>1</sup> */
w1 = 0.0;
for ( i = 579; i >= 475; -i ) {
    i3 = i - 3;
    w1 += fairhd[i-3] * sin( fairhd[i-2] * t + fairhd[i-1] );
}

/* T<sup>2</sup> */
w2 = 0.0;
for ( i = 764; i >= 680; -i ) {
    i3 = i - 3;
    w2 += fairhd[i-3] * sin( fairhd[i-2] * t + fairhd[i-1] );
}

/* T<sup>3</sup> */
w3 = 0.0;
for ( i = 784; i >= 765; -i ) {
    i3 = i - 3;
    w3 += fairhd[i-3] * sin( fairhd[i-2] * t + fairhd[i-1] );
}

/* T<sup>4</sup> */
w4 = 0.0;
for ( i = 787; i >= 785; -i ) {
    i3 = i - 3;
    w4 += fairhd[i-3] * sin( fairhd[i-2] * t + fairhd[i-1] );
}

/* Multiply by powers of T and combine. */
wf = t * ( t * ( t * ( t * w4 + w3 ) + w2 ) + w1 ) + w0;

/* Adjustments to use JPL planetary masses instead of IAU. */
wj = sin( t * 6697.776754 * 4.021194 ) * 6.5e-10
+ sin( t * 1123.000000 * 1.000000 ) * 1.0e-09
+ sin( t * 6208.294251 * 5.696701 ) * -1.96e-9
+ sin( t * 74.781599 + 4.2359 ) * -1.73e-9
+ 3.638e-8 * t * t;

/* Final result: TDB-TT in seconds. */
return wf + wj;
}

```

Multi-Threaded Program in Execution

```

Separate algorithms for Pluto and the rest. */
if (np < 1) np > 9) {
  /*jstat = -1;
  for (i = 0; i <= 5; i++) p[i] = 0.0;
  return;
} else {
  ip = np - 1;
}
/* Time: Julian millennia since [2000, */
t = (d - 515445.1 / 365250.0;
/* OK start unless remote epoch. */
/*jstat = -1;
if (fabs (t - 1.0) > 0.1) jstat = 0;
/* Comp. the min. arguments. */
da = a[0] + (a[1]*t) + (a[2]*t*t);
da = (0.504 * dim(p[0]) + dim(p[1]*t) + dim(p[2]*t*t));
da += (0.0001 * dim(p[3]) * t*t);
da += (0.0001 * (e[1]*t) + (e[2]*t*t));
dpe = mod (13600.0 * p[0]*p[0] + (p[1]*p[1]) + (p[2]*p[2]) * t*t);
dpe += (0.0001 * dinc(p[0]) + (dinc(p[1]) + dinc(p[2])*t*t));
dpe += (0.0001 * dinc(p[3]) * t*t);
dpm = mod (13600.0 * omega(p[0]) + (omega(p[1]) + omega(p[2])*t*t));
dpm += (0.0001 * omega(p[3]) * t*t);
dASZR = dASZR_D2P1;
dD2P1 = dD2P1_D2P1;
/* Compute the trigonometric terms. */
dmu = 3595362.0;
for (i = 0; i <= 7; i++) {
  arg0 = dphi(p[i]) * dmu;
  arg1 = dphi(p[i+1]) * dmu;
  da += (cal(p[i]) * cos (arg0) +
         (ip[i] * sin (arg0)) * 1e-7;
  d1 += (cal(p[i+1]) * cos (arg1) +
         (ip[i+1] * sin (arg1)) * 1e-7;
  arg0 = (ip[i] * dim(p[0]) + (ip[i+1] * dim(p[1])) +
          (ip[i+2] * dim(p[2])) + (ip[i+3] * dim(p[3])) * 1e-7;
  da += (ip[i] * (ip[i] * cos (arg0) +
                 (ip[i+1] * sin (arg0)) * 1e-7;
  d1 += (ip[i+1] * (ip[i+1] * cos (arg1) +
                    (ip[i+2] * sin (arg1)) * 1e-7;
}
d1 = d1_D2P1;
/* Daily motion. */
dm = G * S * sqrt (1.0 + 1.0 / amas(p)) / (da * da * da);
/* Make prediction. */
splaPlanets (date, 1, date, dl, dom, dpe, da, d1, dm, pv, &j);
/*jstat = -2;
if (i < 0) {
  /* ... */
} else {
  /* ... */
  /* Pluto */
  /* ... */
  /* Time: Julian millennia since [2000, */
  t = (d - 515445.1 / 365250.0;
  /* OK start unless remote epoch. */
  /*jstat = -1;
  /* Fundamentals.
  d1 = (d1 * dmu) / dASZR;
  dpe = (dpe * dmu) / dASZR;
  dm = (dm * dmu) / dASZR;
  /* Initial, coefficients and derivatives. */
  for (i = 1; i < 3; i++) {
    wibr(i) = 0.0;
    wibrd(i) = 0.0;
  }
  /* Term by term through Meeus Table 36A. */
  for (j = 0; j < sizeof term / sizeof term[0]; j++) {
    /* Argument and derivative (radians, radians per century). */
    wj = (double) term[j];
    ws = (double) term[j].is;
    wp = (double) term[j].ip;
    /* ... */
  }
}

```

Multi-Threaded Program in Execution

Multi-Threaded Program in Execution

```

/* Validate the planet number. */
if (np < 1 || np > 9) {
    *jstat = -1;
    for (i = 0; i <= 5; i++) pv[i] = 0.0;
    return;
} else {
    ip = np - 1;
}

/* Separate algorithms for Pluto and the rest. */
if (ip == 9) {
    /* Compute term through Neptune */
    /* Time: Julian millennia since [2000. */
    t = (da - 51544.5) / 365250.0;
    /* OK stat is unless remote epoch */
    if (tbs(t) <= 1.0 || t > 1.0) {
        /* Compute mean elements */
        da = a[0] + (a[1] * t) + a[2] * t * t;
        da += a[3] * sin(a[0]) + a[4] * cos(a[0]);
        da += a[5] * sin(a[1]) * cos(a[2]) * t;
        da += a[6] * sin(a[1]) * sin(a[2]) * t * t;
        da += a[7] * sin(a[2]) * t * t * t;
        /* Compute eccentricity elements */
        de = e[0] + (e[1] * t) + e[2] * t * t;
        de += e[3] * sin(e[0]) + e[4] * cos(e[0]);
        de += e[5] * sin(e[1]) * cos(e[2]) * t;
        de += e[6] * sin(e[1]) * sin(e[2]) * t * t;
        de += e[7] * sin(e[2]) * t * t * t;
        /* Compute inclination elements */
        di = i[0] + (i[1] * t) + i[2] * t * t;
        di += i[3] * sin(i[0]) + i[4] * cos(i[0]);
        di += i[5] * sin(i[1]) * cos(i[2]) * t;
        di += i[6] * sin(i[1]) * sin(i[2]) * t * t;
        di += i[7] * sin(i[2]) * t * t * t;
        /* Compute longitude of perihelion */
        dp = l[0] + (l[1] * t) + l[2] * t * t;
        dp += l[3] * sin(l[0]) + l[4] * cos(l[0]);
        dp += l[5] * sin(l[1]) * cos(l[2]) * t;
        dp += l[6] * sin(l[1]) * sin(l[2]) * t * t;
        dp += l[7] * sin(l[2]) * t * t * t;
        /* Compute argument of perihelion */
        dp = a[8] + (a[9] * t) + a[10] * t * t;
        dp += a[11] * sin(a[8]) + a[12] * cos(a[8]);
        dp += a[13] * sin(a[9]) * cos(a[10]) * t;
        dp += a[14] * sin(a[9]) * sin(a[10]) * t * t;
        dp += a[15] * sin(a[10]) * t * t * t;
        /* Compute mean anomaly */
        dm = G * M * t;
        dm *= (1.0 + 1.0 / asmas(ip)) / (da * da * da);
        /* Make prediction */
        sp = plan9(ip, date, da, dm, dpe, da, de, dp, pv, &j);
        if (j < 0) *jstat = -2;
    } else {
        /* Pluto */
        /* ... */
        /* Time: Julian centuries since [2000. */
        t = (da - 51544.5) / 365250.0;
        /* OK stat is unless remote epoch */
        if (t >= -1.55 && t < -1.0700000000000001; /* 1.0700000000000001; */
            /* Fundamental arguments (radians). */
            d[0] = (d[0] + d[1] * t) / D00R;
            d[2] = (d[2] + d[3] * t) / D00R;
            d[4] = (d[4] + d[5] * t) / D00R;
            /* Initial coefficients and derivatives. */
            for (i = 1; i < 3; i++) {
                wibr[i] = 0.0;
                wibr[i] = 0.0;
            }
        }
    }
}

/* Term by term through Meeus Table 36A. */
for (j = 0; j < sizeof(term / sizeof(term[0])); j++) {
    /* Argument and derivative (radians, radians per century). */
    w[0] = (double) term[j].x();
    ws = (double) term[j].y();
    wp = (double) term[j].z();
    /* Compute mean elements */
    da = a[0] + (a[1] * w[0]) + a[2] * w[0] * w[0];
    da += a[3] * sin(a[0]) + a[4] * cos(a[0]);
    da += a[5] * sin(a[1]) * cos(a[2]) * w[1];
    da += a[6] * sin(a[1]) * sin(a[2]) * w[1] * w[1];
    da += a[7] * sin(a[2]) * w[1] * w[1] * w[1];
    /* Compute eccentricity elements */
    de = e[0] + (e[1] * w[0]) + e[2] * w[0] * w[0];
    de += e[3] * sin(e[0]) + e[4] * cos(e[0]);
    de += e[5] * sin(e[1]) * cos(e[2]) * w[1];
    de += e[6] * sin(e[1]) * sin(e[2]) * w[1] * w[1];
    de += e[7] * sin(e[2]) * w[1] * w[1] * w[1];
    /* Compute inclination elements */
    di = i[0] + (i[1] * w[0]) + i[2] * w[0] * w[0];
    di += i[3] * sin(i[0]) + i[4] * cos(i[0]);
    di += i[5] * sin(i[1]) * cos(i[2]) * w[1];
    di += i[6] * sin(i[1]) * sin(i[2]) * w[1] * w[1];
    di += i[7] * sin(i[2]) * w[1] * w[1] * w[1];
    /* Compute longitude of perihelion */
    dp = l[0] + (l[1] * w[0]) + l[2] * w[0] * w[0];
    dp += l[3] * sin(l[0]) + l[4] * cos(l[0]);
    dp += l[5] * sin(l[1]) * cos(l[2]) * w[1];
    dp += l[6] * sin(l[1]) * sin(l[2]) * w[1] * w[1];
    dp += l[7] * sin(l[2]) * w[1] * w[1] * w[1];
    /* Compute argument of perihelion */
    dp = a[8] + (a[9] * w[0]) + a[10] * w[0] * w[0];
    dp += a[11] * sin(a[8]) + a[12] * cos(a[8]);
    dp += a[13] * sin(a[9]) * cos(a[10]) * w[1];
    dp += a[14] * sin(a[9]) * sin(a[10]) * w[1] * w[1];
    dp += a[15] * sin(a[10]) * w[1] * w[1] * w[1];
    /* Compute mean anomaly */
    dm = G * M * w[0];
    dm *= (1.0 + 1.0 / asmas(ip)) / (da * da * da);
    /* Make prediction */
    sp = plan9(ip, date, da, dm, dpe, da, de, dp, pv, &j);
    if (j < 0) *jstat = -2;
}

```

Multi-Threaded Program in Execution

Or they can be running the same code at the same time (more or less)

```

0.0000 0 in Julian millennia. */
344565250.0;
/* Topocentric terms ..... */
/* local solar time in radians. */
ut1, 1.0) * D2R - wl;
/* ARGUMENTS: Simon et al 1994. */
/* argument (millennia) with deg/arcsec factor. */
/* */
/* Latitude */
(280.46645683 + 1296027711.03429 * w, 360.0) * DD2R;
/* anomaly. */
d (357.59210918 + 1295965810.481 * w, 360.0) * DD2R;
/* Position of Moon from Sun. */
37.85019547 + 16029616012.090 * w, 360.0) * DD2R;
/* Position of Jupiter. */
34.35151874 + 109306899.89453 * w, 360.0) * DD2R;
/* Position of Saturn. */
50.07744430 + 44046398.47038 * w, 360.0) * DD2R;
/* DUTCTR TERMS: Moyer 1981 and Murray 1983. */
wt = 0.00029e-10 * u * sin (tsol + elsun - els);
+ 0.00106e-10 * u * sin (tsol + 2.0 * emsun);
+ 0.00212e-10 * u * sin (tsol + elsun - e);
+ 0.00133e-10 * u * sin (tsol + elsun - elj);
- 0.00229e-10 * u * sin (tsol + 2.0 * elsun + emsun);
- 0.0220e-10 * u * cos (elsun + emsun);
+ 0.00126e-10 * u * cos (elsun + emsun);
+ 0.1367e-10 * u * cos (elsun + elj);
- 0.1384e-10 * u * cos (elsun);
+ 3.1767e-10 * u * sin (elsun);
/* Fairhead model */
/* w0 = 0.0; */
/* for (i = 0; i < 10; i++) */
i3 = i * 3;
w0 = fairhd[i3-3] * sin (fairhd[i-2] * t + fairhd[i-1]);
}
/* w1 = 0.0; */
/* for (i = 67; i >= 4; i--) */
i3 = i * 3;
w1 = fairhd[i3-3] * sin (fairhd[i-2] * t + fairhd[i-1]);
}
/* w2 = 0.0; */
/* for (i = 764; i >= 60; i--) */
i3 = i * 3;
w2 = fairhd[i3-3] * sin (fairhd[i-2] * t + fairhd[i-1]);
}
/* T^2 */
w3 = 0.0;
for (i = 784; i >= 75; i--) {
i3 = i * 3;
w3 = fairhd[i3-3] * sin (fairhd[i-2] * t + fairhd[i-1]);
}
/* T^3 */
w4 = 0.0;
for (i = 787; i >= 75; i--) {
i3 = i * 3;
w4 = fairhd[i3-3] * sin (fairhd[i-2] * t + fairhd[i-1]);
}
/* Multiply by powers of T and combine. */
w = t * (t * (t * (w0 + 4 * w3) + 2 * w1) + w0);
/* An adjustment to the planetary masses instead of IAU. */
w = sin (t * 6069.6754 + 4.02.49) * 1.5e-10
+ sin (t * 213299095 + 5.543132) * 3.3e-10
+ sin (t * 10628251 + 4.86071) * -1.9e-9
+ sin (t * 74.781599 + 2.4559) * -1.7e-9
+ 3.638e-8 * t*t;
}
/* Final result: TDB-TT in seconds. */
return wt + wf + w;
}

```

Multi-Threaded Program in Execution

Or any combination thereof

```

    /* Look for sign */
    switch (idch1) {
        case '+':
        case '-':
            state = seek_sign;
            break;
        case 'N':
        case 'U':
        case 'M':
            state = accept_unt_exp_no_mant;
            break;
        case 'P':
        case 'E':
            state = seek_digit_when_none_before_pt;
            break;
        case '+':
        case '-':
            state = seek_1st_leading_digit;
            break;
        case 'M':
            state = error;
            break;
        case OTHER :
            state = next_field_default;
            break;
        case 'A':
        case 'M':
        case 'D':
            state = null_field;
            break;
        default:
            state = error;
            break;
    }
}

```

Threads vs. Processes

■ 😊 Memory sharing

- Threads naturally share memory among each other
 - Provides a Shared Memory IPC mechanism with no system calls
- Having concurrent activities in the same address space is very powerful
- It makes it possible to implement all kinds of concurrency behaviors/patterns

■ 😡 No memory protection

- This is a “feature” since we *want* threads to share memory
- But this can cause really, really difficult bugs
- More about this in the Synchronization module

■ 😊 Economy

- Creating a thread is cheap
 - Slightly cheaper than creating a process in MacOS/Linux
 - Much cheaper than creating a process in Windows
- Context-switching between threads in a process is cheaper than between different processes
 - Because they share the same address space (TLB... see much later this semester)
- So if you can use threads instead of processes, then you likely can go a bit faster
 - In old OSes (Solaris 4), threads were called “lightweight processes”

Threads vs. Processes

- 😠 **Less fault-tolerance**
 - If a thread fails/crashes, then the whole process fails/crashes, instead of processes, which are independent of each other
 - This motivates developers to use both processes and threads (see next slide)
- 😠 **Possibly more memory-constrained**
 - Since threads execute in the same address space, and an OS can bound the size of a process' address space
 - But that's typically not a big deal (one can configure the OS if need be)
- The advantages here are well worth the drawbacks/limitations
 - The big drawback/feature is “no memory protection”
 - We have developed many, many approaches/solutions to deal with it (see the Synchronization module)
- Natural question: is concurrency with processes obsolete?

Concurrency with Processes?

- Should we still care about concurrency with processes?
- **YES** because many applications consists of multiple processes (each of which are often multi-threaded)
- Well-known examples are some popular Web browser (Chrome, Firefox)!
 - They calls fork() each time you open a tab
 - Each tab is a (possibly heavily) multi-threaded process
 - As a result, the code contains processes that do IPC because they don't "see" the same memory naturally
 - But if a tab crashes (due to running bad JavaScript code, for instance) your browser doesn't crash!
 - Google "firefox chrome processes threads" :)
- In real-world settings you often have to put together different software products to make up a whole system
 - Some may just be executables instead of libraries with nice APIs
 - So you have to create processes
 - You interact with them via stdin/stdout/stderr streams for instance (see our programming assignment) or via any supported IPC mechanisms
- **Bottom-line:** don't drink the "I'll only do threads, not processes" Kool-Aid

User vs. Kernel Threads

- Threads can be supported solely in User Space (**User Threads**)
 - You can write your own thread implementation without help from the OS
 - Often a homework assignment in a graduate OS course
- The main advantage of User Threads is low overhead
 - e.g., because no system calls
- User Threads have several drawbacks:
 - If one thread blocks, all other threads block
 - All threads run on the same core (because the OS doesn't know that there are threads within a process)
- For these reasons User Threads are (no longer) heavily used
 - But Java recently re-introduced its old “Green (user) threads”!!
- All OSes today provide support for threads (**Kernel Threads**) that can run on different cores and be truly independent of each other
- We typically just call them “threads”

Linux/MacOS Threads

- Processes and Threads are implemented as **tasks**
 - Kernel data structure: **task struct**
 - We already looked at it when we talked about processes
- The **clone()** syscall is used to create a task
- It can be invoked with several options, each set or not set
- Each option specifies something the child should share or not share with its parent
- **fork()** just calls **clone()** with a particular set of options
 - Preserved as a system call for backward compatibility to create processes
- From the man page: “if CLONE_VM is set, the calling process and the child process run in the same memory space”
- To create a process, **clone()** is called without the CLONE_VM option
- To create a thread, **clone()** is called with, among other things, the CLONE_VM option

Main Takeaways

- Concurrency is used to make programs faster and/or more responsive
- False and True concurrency
- Threads within a process
 - Share the code and the heap
 - Have their own stack, registers, and program counter
- Concurrency with processes is not obsolete, and many programs use both multiple processes and multiple threads

Conclusion

- Threads have been mainstream for decades: multi-threading today is everywhere, in part due to us having multi-core architectures
- Let's do a `ps auxM` on my MacOS laptop and see how many processes are multi-threaded...
 - When I did this while back writing this slides I got 350 processes and 1157 threads. Almost all processes are multithreaded, with up to 60+ threads for a process
- Let's move on to seeing how to use them in programming languages...
 - Almost all languages have some way to create and manage threads!