FROM SCALAR & SERIAL TO VECTOR & PARALLEL

Hands-on Lab

Part 1 of 3

GOALS

- **Part 1: see multi-threading on Intel architecture in action**
- Part 2: learn to vectorize & future-proof vectorization
- Part 3: experiment with what it takes to make the memory subsystem happy

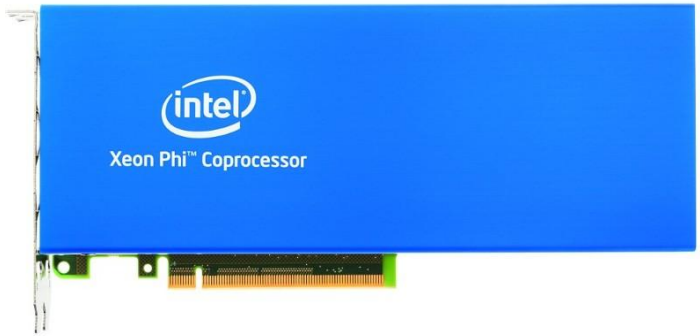

EXAMPLE PROBLEM: BINNING

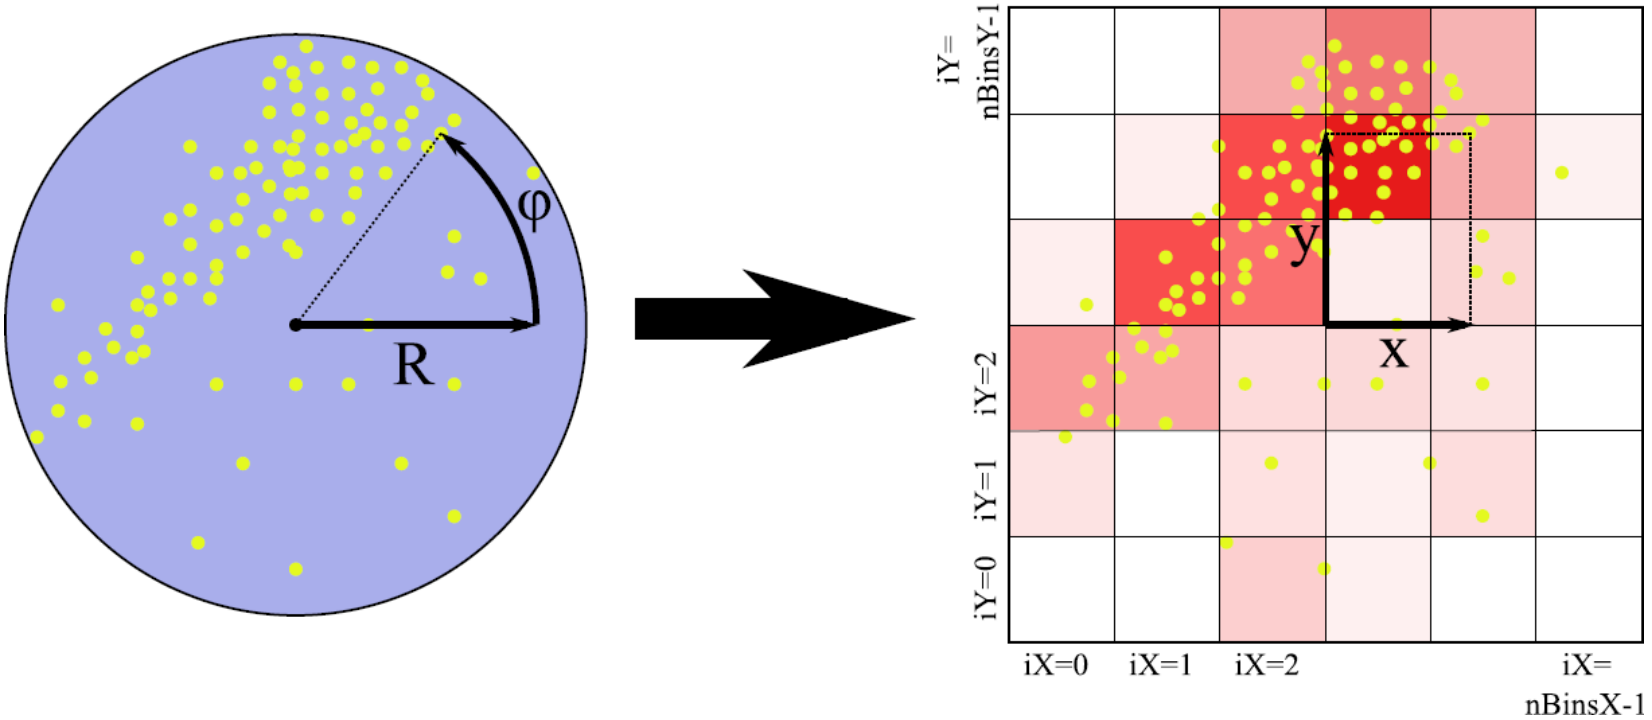

SERVER FOR EXERCISES

- Instructions at **uni.colfax-intl.com/cdt**
- Find code of lab + exercises from **xeonphi.com/book**

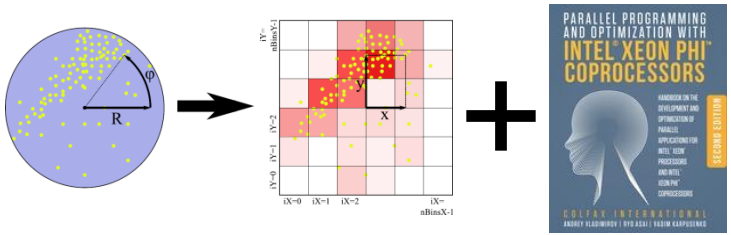

- Enjoy remote access until midnight

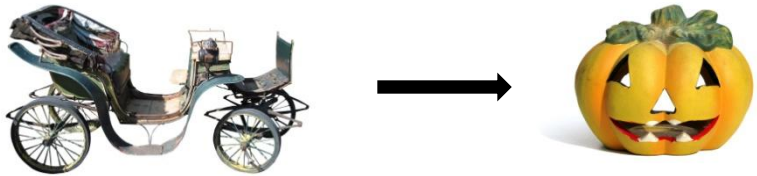

INITIAL APPROACH

```
 // Reference implementation: scalar, serial code without optimization
void BinParticlesReference(
        const InputDataType  & inputData, BinsType & outputBins) {
  // Loop through all particle coordinates
  for (int i = 0; i < inputData.numDataPoints; i++) {
    // Transforming from cylindrical to Cartesian coordinates:
    const FTYPE x = inputData.r[i]*COS(inputData.phi[i]);
    const FTYPE y = inputData.r[i]*SIN(inputData.phi[i]);

    // Calculating the bin numbers for these coordinates:
    const int iX = int((x - xMin)*binsPerUnitX);
    const int iY = int((y - yMin)*binsPerUnitY);

    // Incrementing the appropriate bin in the counter:
    outputBins[iX][iY]++;
  }
}
```

OPENMP THREADS

```
 // Parallel, but INCORRECT implementation
void BinParticlesReference(
         const InputDataType  & inputData,
         BinsType & outputBins) {
  // Distribute loop iterations across threads
#pragma omp parallel for
  for (int i = 0; i < inputData.numDataPoints; i++)
    {

      // …
      // Incrementing the appropriate
      // bin in the counter:
      outputBins[iX][iY]++;
    }
}
```

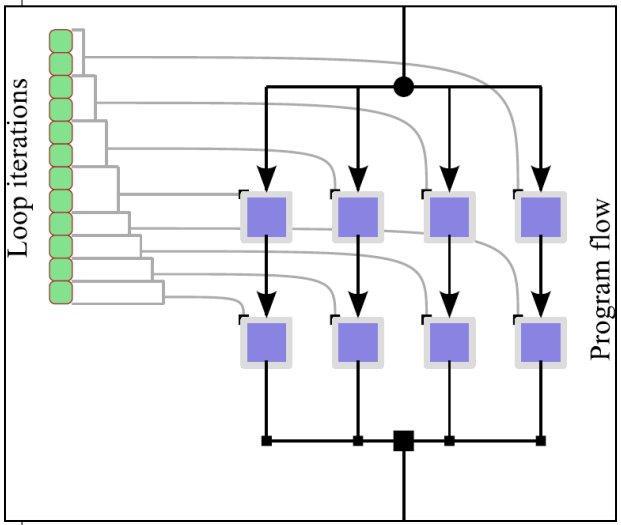

Software threads utilize logical processors in an Intel Xeon or an Intel Xeon Phi processor

DATA RACES

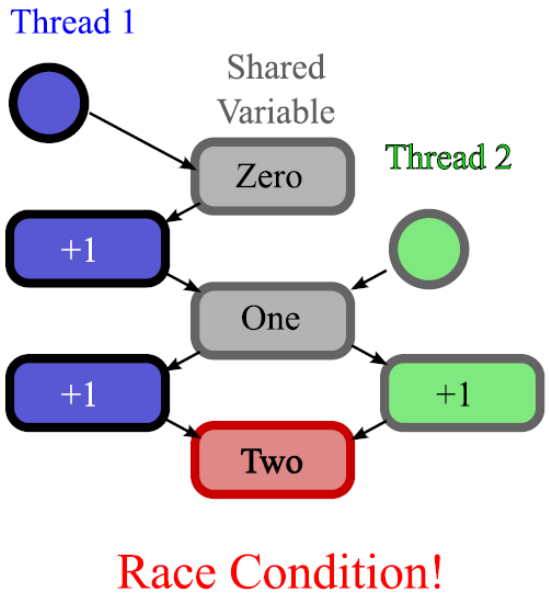
Thread 1
Shared Variable
Zero
Thread 2
+1
One
+1
+1
Two
Race Condition!

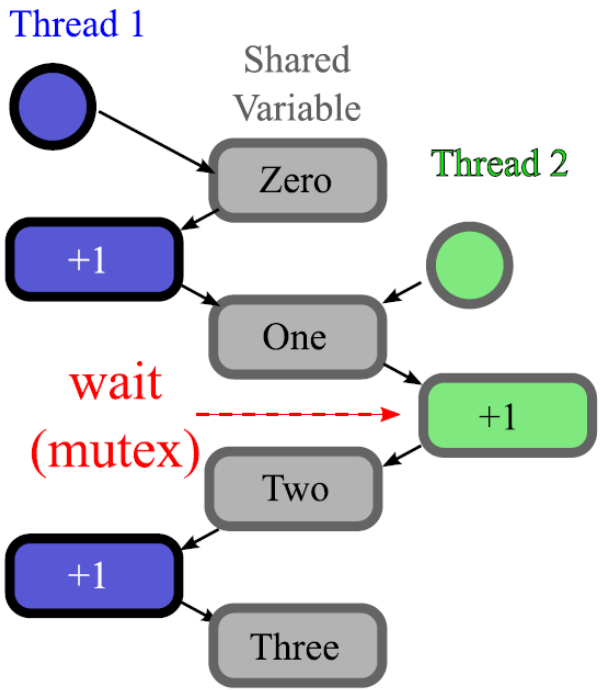
Thread 1
Shared Variable
Zero
Thread 2
+1
One
wait
(mutex)
+1
Two
+1
Three

PARALLEL REDUCTION

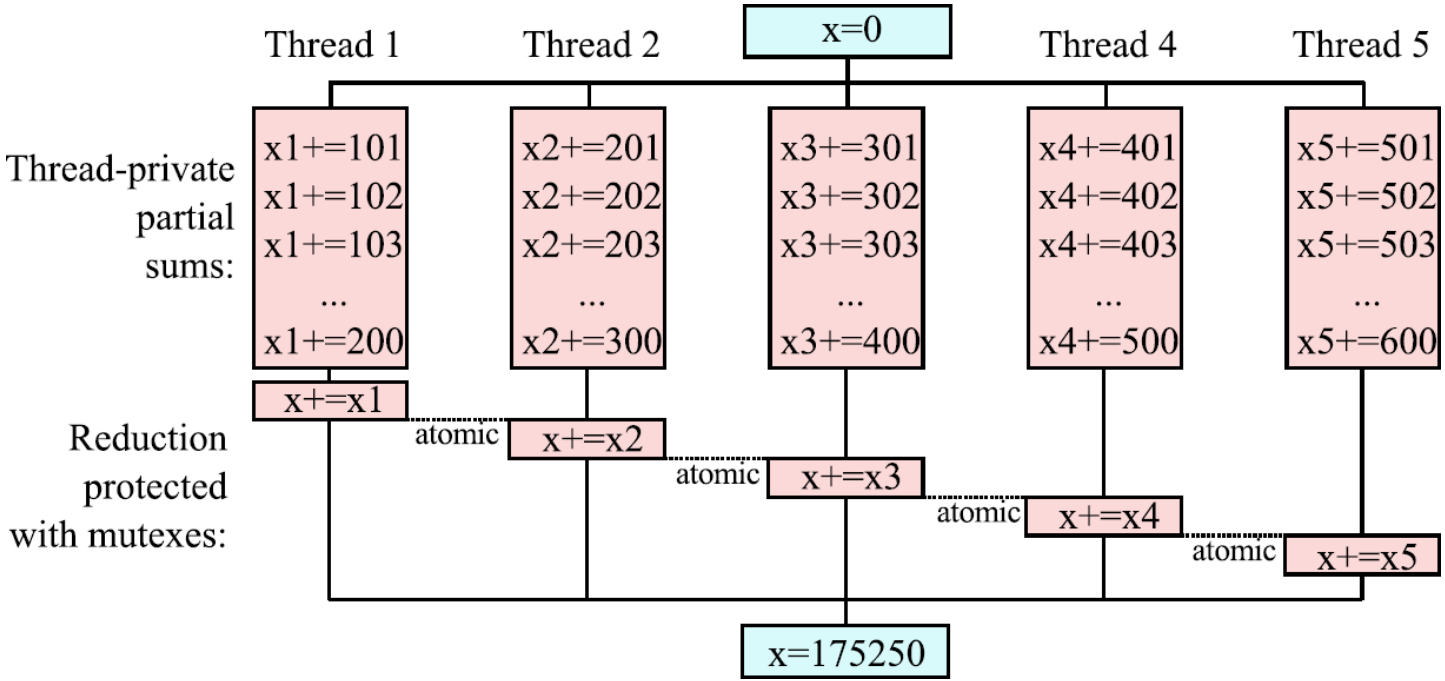

Key: using thread-private containers for partial sums; mutexes only after the parallel loop

OPTIMIZED VERSION

```
void BinParticles_2(const InputDataType& inputData, BinsType& outputBins){
#pragma omp parallel
  { BinsType threadPrivateBins; // Declare thread-private containers
    for (int i = 0; i < nBinsX; i++)
      for (int j = 0; j < nBinsY; j++)
        threadPrivateBins[i][j] = 0;
#pragma omp for
    for (int i = 0; i < inputData.numDataPoints; i++) {
      // …transforming from cylindrical to Cartesian coordinates:
      // …calculating the bin numbers for these coordinates:
      // Incrementing the appropriate bin in the thread-private counter:
      threadPrivateBins[iX][iY]++;
    }
    for(int i = 0; i < nBinsX; i++) // Reduction outside the parallel loop
      for(int j = 0; j < nBinsY; j++)
#pragma omp atomic
        outputBins[i][j] += threadPrivateBins[i][j];
  } }
```

THREAD AFFINITY

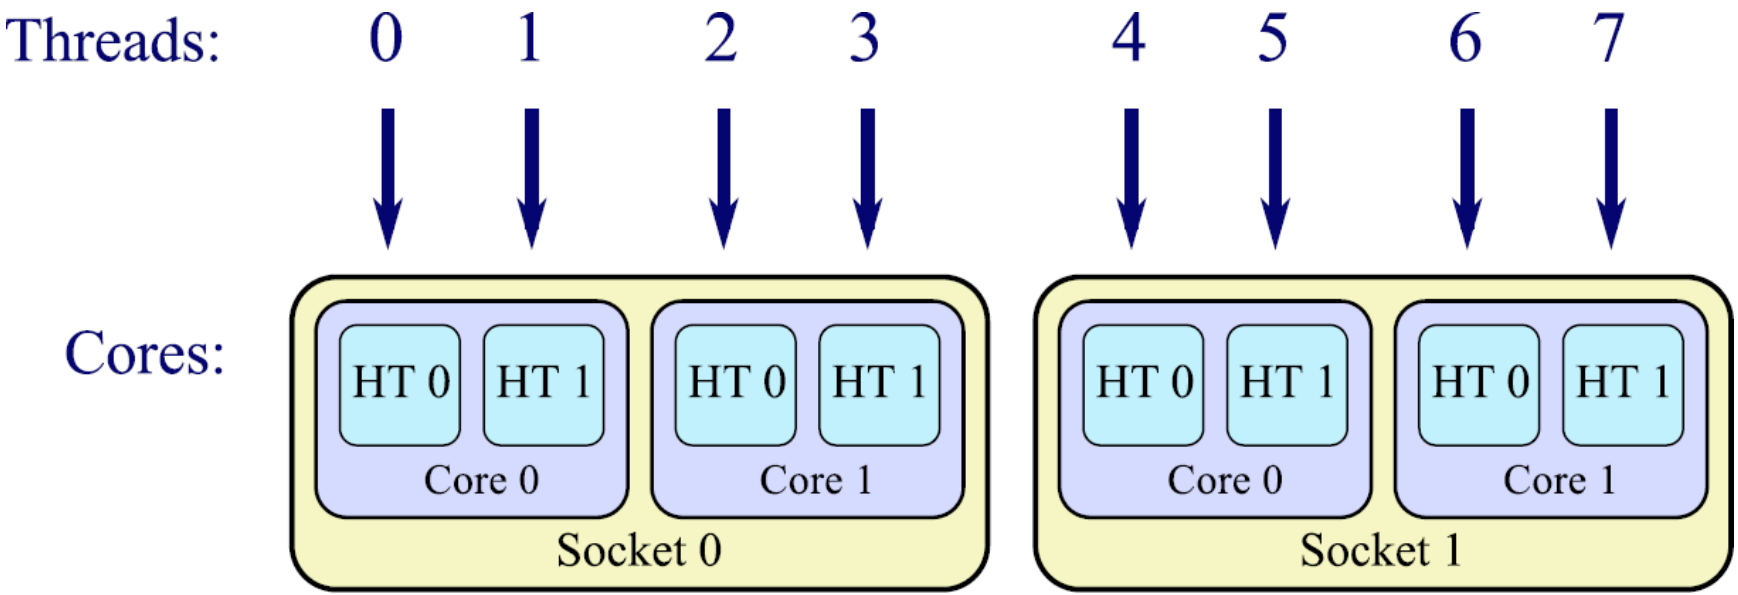

PERFORMANCE RESULTS

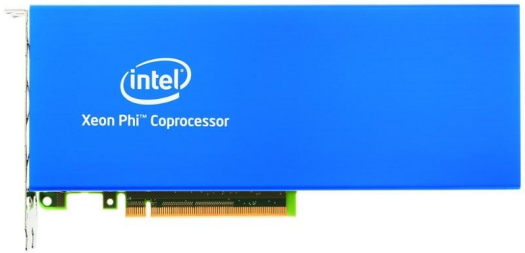

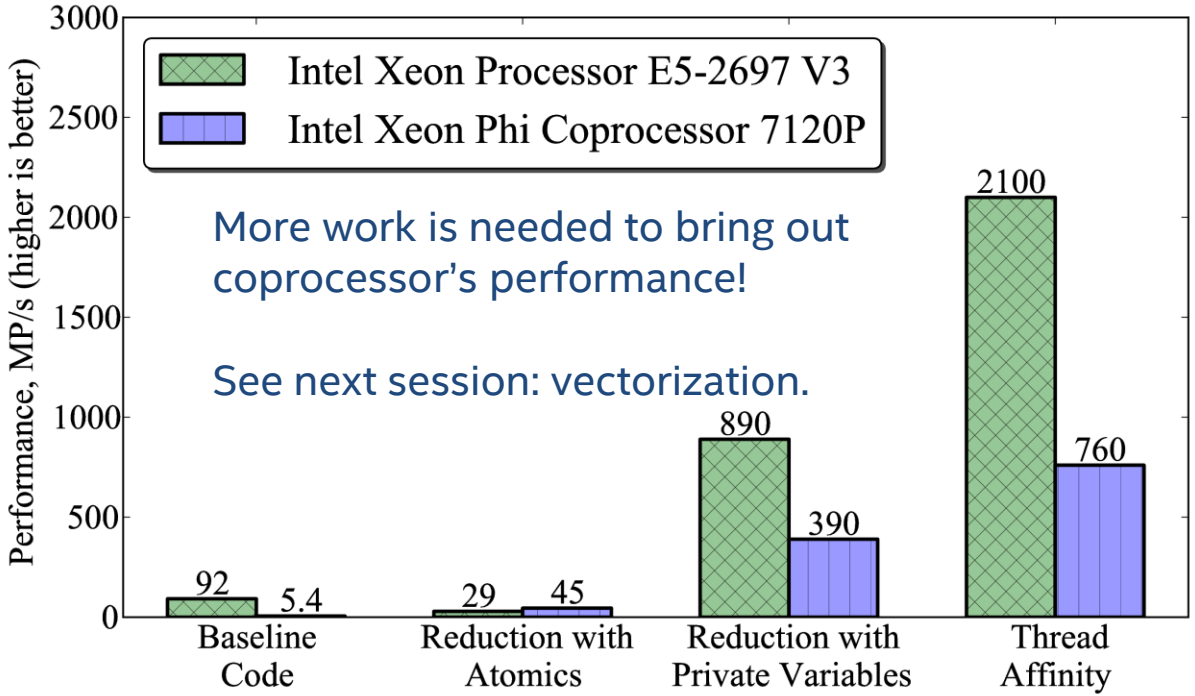

NEXT SESSION

- Part 2: vectorization – another level of parallelism
- Part 3: making the memory subsystem happy

By the way:
download this tutorial
colfaxresearch.com