

# N-body Simulation

# Physics

## Gravitational N-body dynamics:


Newton's law of universal gravitation:

$$M_i \vec{R}_i''(t) = G \sum_j \frac{M_i M_j}{|\vec{R}_i - \vec{R}_j|^3} (\vec{R}_j - \vec{R}_i)$$

where:

$$|\vec{R}_i - \vec{R}_j| = \sqrt{(R_{i,x} - R_{j,x})^2 + (R_{i,y} - R_{j,y})^2 + (R_{i,z} - R_{j,z})^2}$$

particles are attracted to each other  
with **the gravitational force**



# Application

- 1 Astrophysics:
  - ▶ planetary systems
  - ▶ galaxies
  - ▶ cosmological structures
- 2 Electrostatic systems:
  - ▶ molecules
  - ▶ crystals

This work: “toy model” with all-to-all  $O(n^2)$  algorithm. Practical N-body simulations may use tree algorithms with  $O(n \log n)$  complexity.




Source: [APOD](#), credit: Debra Meloy Elmegreen (Vassar College) et al., & the Hubble Heritage Team (AURA/ STScI/ NASA)

# Comparative Benchmarks and System Configuration


Colfax ProEdge SXP8600p  
rack-mountable workstations (cluster of 4)



Mellanox Connect-IB  
InfiniBand HCA (FDR)



Intel Xeon Phi 7120P coprocessors  
(4 per system)



Dual-socket Intel Xeon E5-2697 v2  
processor




<http://xeonphi.com/workstations>

# Initial Implementation of the N-Body Simulation

# Illustration of “Toy Model” Calculation Pattern

- All-to-all interaction
- $O(n^2)$  complexity
- All particles fit in memory of each compute node
- No multipole approximation, tree algorithms, Debye screening, etc.
- Basis for more efficient real-life models
- Good educational example



# All-to-All Approach ( $O(n^2)$ Complexity Scaling)

Each particle is stored as a structure:

```
1 struct ParticleType {  
2     float x, y, z;  
3     float vx, vy, vz;  
4 };
```

main() allocates an array of ParticleType:

```
1 ParticleType* particle = new ParticleType[nParticles];
```

Particle propagation step is timed:

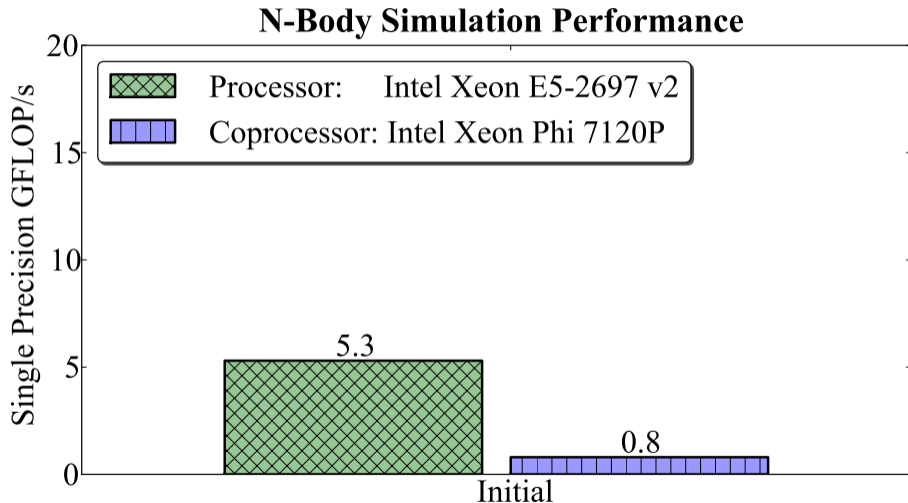
```
1 const double tStart = omp_get_wtime(); // Start timing  
2 MoveParticles(nParticles, particle, dt);  
3 const double tEnd = omp_get_wtime(); // End timing
```

# Particle Update Engine

```
1 void MoveParticles(int nParticles, ParticleType* particle, float dt) {
2     for (int i = 0; i < nParticles; i++) { // Particles that experience force
3         float Fx = 0, Fy = 0, Fz = 0; // Gravity force on particle i
4         for (int j = 0; j < nParticles; j++) { // Particles that exert force
5             // Newton's law of universal gravity
6             const float dx = particle[j].x - particle[i].x;
7             const float dy = particle[j].y - particle[i].y;
8             const float dz = particle[j].z - particle[i].z;
9             const float drSquared = dx*dx + dy*dy + dz*dz + 1e-20;
10            const float drPower32 = pow(drSquared, 3.0/2.0);
11            // Calculate the net force
12            Fx += dx/drPower32; Fy += dy/drPower32; Fz += dz/drPower32;
13        }
14        // Accelerate particles in response to the gravitational force
15        particle[i].vx+=dt*Fx; particle[i].vy+=dt*Fy; particle[i].vz+=dt*Fz;
16    }
17    ...
}
```



# Performance of Initial Implementation



# Optimization: Thread Parallelism

# Incorporating Thread Parallelism


Before:

```
1  for (int i = 0; i < nParticles; i++) { // Particles that experience force
2      float Fx = 0, Fy = 0, Fz = 0; // Gravity force on particle i
3      for (int j = 0; j < nParticles; j++) { // Particles that exert force
4          // Newton's law of universal gravity
5          ...
```

After:

```
1  #pragma omp parallel for
2      for (int i = 0; i < nParticles; i++) { // Particles that experience force
3          float Fx = 0, Fy = 0, Fz = 0; // Gravity force on particle i
4          for (int j = 0; j < nParticles; j++) { // Particles that exert force
5              // Newton's law of universal gravity
6              ...
```

# Performance with Thread Parallelism



# Optimization: Vectorization

# Vectorizing with Unit-Stride Memory Access

Before:


```
1 struct ParticleType {  
2     float x, y, z, vx, vy, vz;  
3 }; // ...  
4     const float dx = particle[j].x - particle[i].x;  
5     const float dy = particle[j].y - particle[i].y;  
6     const float dz = particle[j].z - particle[i].z;
```

After:


```
1 struct ParticleSet {  
2     float *x, *y, *z, *vx, *vy, *vz;  
3 }; // ...  
4     const float dx = particle.x[j] - particle.x[i];  
5     const float dy = particle.y[j] - particle.y[i];  
6     const float dz = particle.z[j] - particle.z[i];
```

# Why AoS to SoA Conversion Helps: Unit Stride


Array of Structures  
(sub-optimal)



Structure of Arrays  
(optimal)



# Performance with Improved Vectorization





# Optimization: Scalar Tuning

# Improving Scalar Expressions

Before:

```
1  const float drSquared = dx*dx + dy*dy + dz*dz + 1e-20;  
2  const float drPower32 = pow(drSquared, 3.0/2.0);  
3  // Calculate the net force  
4  Fx += dx/drPower32;  Fy += dy/drPower32;  Fz += dz/drPower32;
```

After:

```
1  const float drRecip    = 1.0f/sqrtf(dx*dx + dy*dy + dz*dz + 1e-20f);  
2  const float drPowerN32 = drRecip*drRecip*drRecip;  
3  // Calculate the net force  
4  Fx += dx*drPowerN32;  Fy += dy*drPowerN32;  Fz += dz*drPowerN32;
```

- Strength reduction (division → multiplication by reciprocal)
- Precision control (suffix -f on single-precision constants and functions)
- Reliance on hardware-supported reciprocal square root

# Compilation with Relaxed Precision


For the CPU architecture (Intel Xeon E5-2697 v2 processor):

```
vega@lyra% # Compile with relaxed precision: (-fp-model fast=2)
vega@lyra% icpc -o nbody-CPU -qopenmp -fp-model fast=2 nbody.cc
vega@lyra% export KMP_AFFINITY=compact
vega@lyra% ./nbody-CPU
```

For the MIC architecture (Intel Xeon Phi 7120P coprocessor):

```
vega@lyra% # Compile for Xeon Phi with relaxed precision: (-fp-model fast=2)
vega@lyra% icpc -o nbody-MIC -mmic -qopenmp -fp-model fast=2 nbody.cc
vega@lyra% export KMP_AFFINITY=compact
vega@lyra% export SINK_LD_LIBRARY_PATH=$MIC_LD_LIBRARY_PATH
vega@lyra% micnativeloadex ./nbody-MIC
```

# Performance after Scalar Tuning



# Optimization: Memory Traffic

# Improving Cache Traffic

Before:

```
1  for (int i = 0; i < nParticles; i++) { // Particles that experience force
2      float Fx = 0, Fy = 0, Fz = 0; // Gravity force on particle i
3      for (int j = 0; j < nParticles; j++) { // Particles that exert force
4          // ...
5          Fx += dx*drPowerN32; Fy += dy*drPowerN32; Fz += dz*drPowerN32;
```

After: (tileSize = 16)

```
1  for (int ii = 0; ii < nParticles; ii += tileSize) { // Particle blocks
2      float Fx[tileSize], Fy[tileSize], Fz[tileSize]; // Force on particle block
3      Fx[:] = Fy[:] = Fz[:] = 0;
4      #pragma unroll(tileSize)
5      for (int j = 0; j < nParticles; j++) { // Particles that exert force
6          for (int i = ii; i < ii + tileSize; i++) { // Traverse the block
7              // ...
8              Fx[i-ii] += dx*drPowerN32; Fy[i-ii] += dy*drPowerN32; Fz[i-ii] += dz*drPowerN32;
```

# Performance with Cache Optimization (Loop Tiling)

