

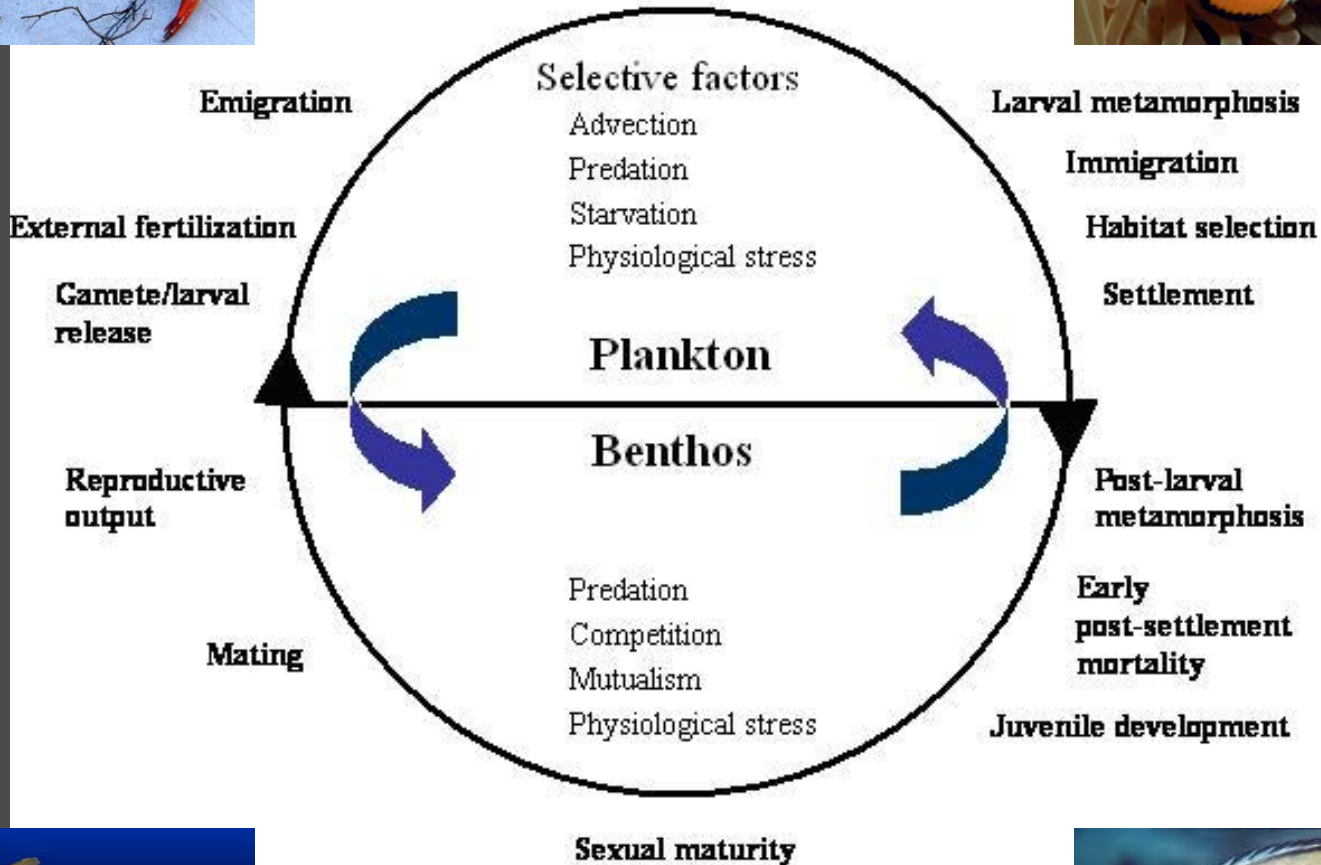


Benjamin Jones


# Parallelization of a particle-tracking model



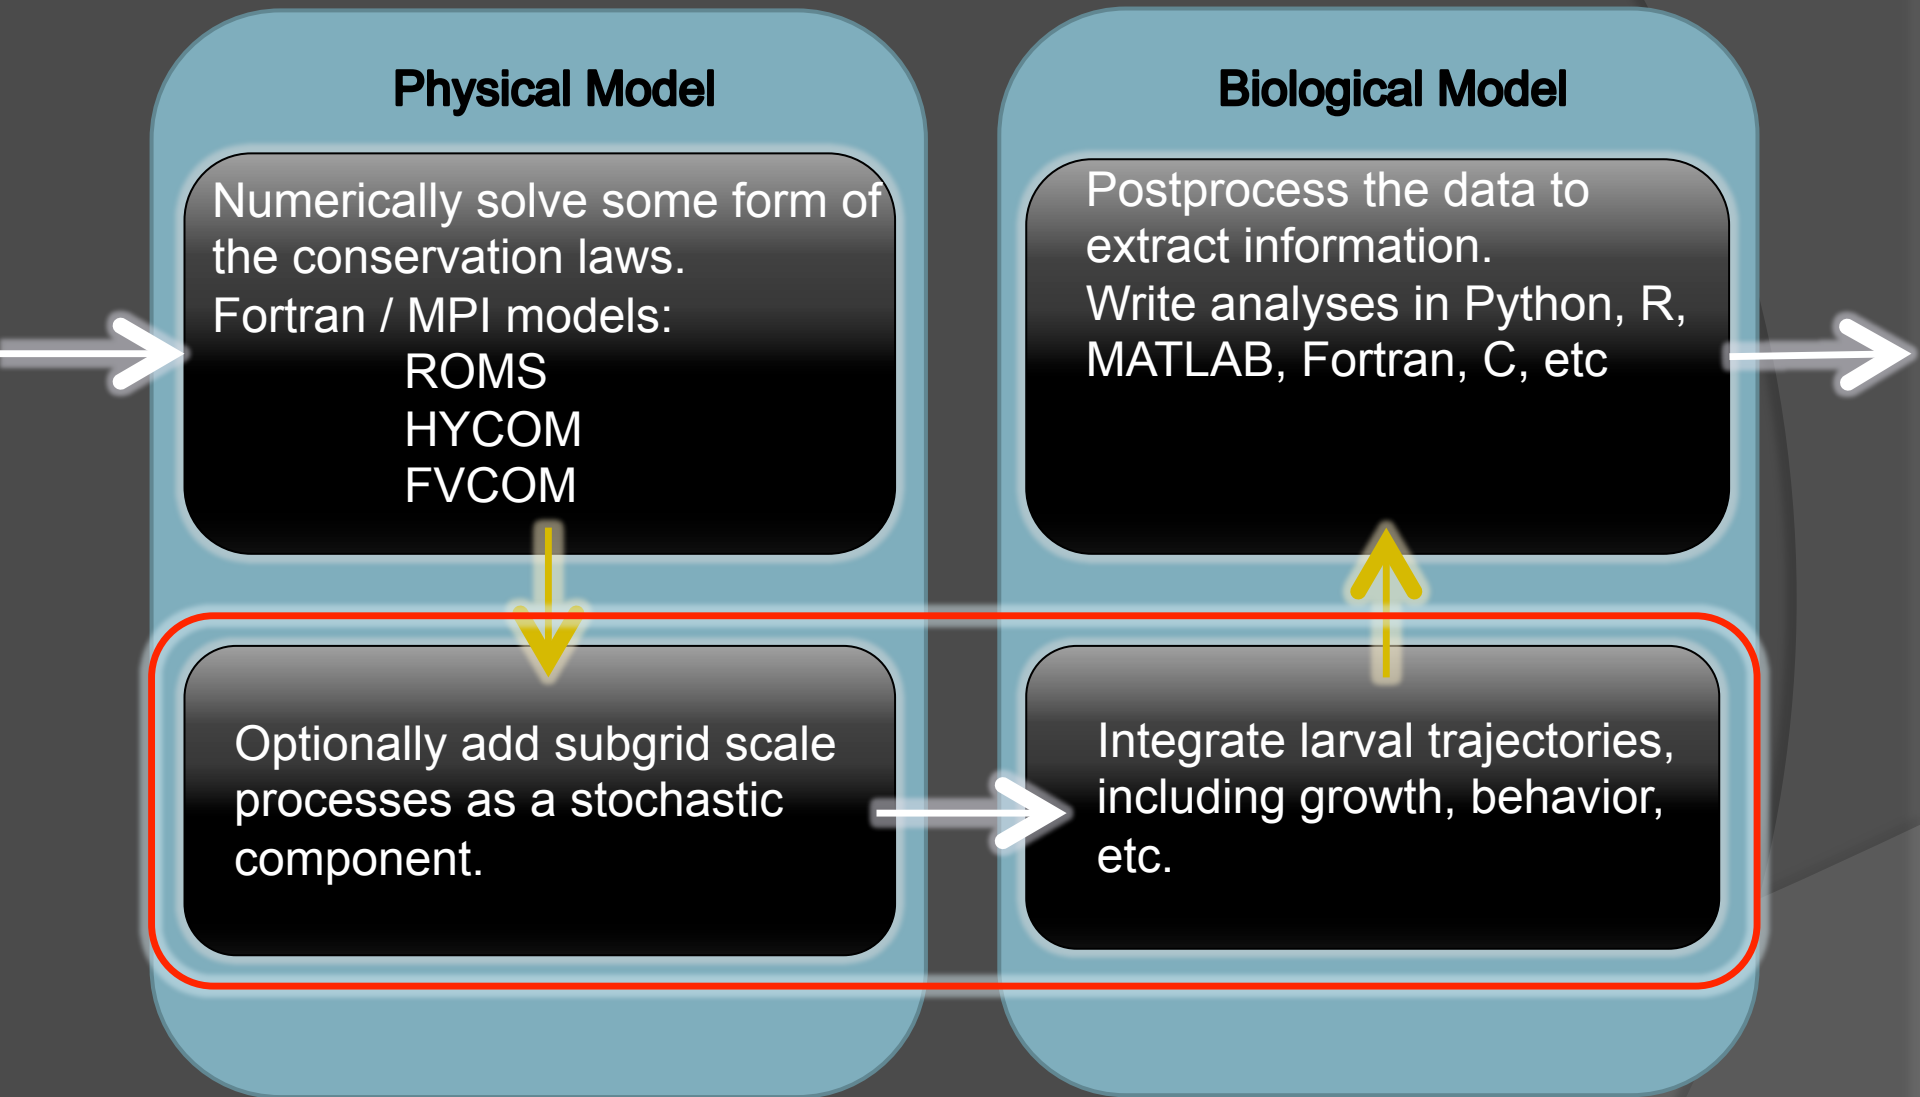
**Larval development and transport**

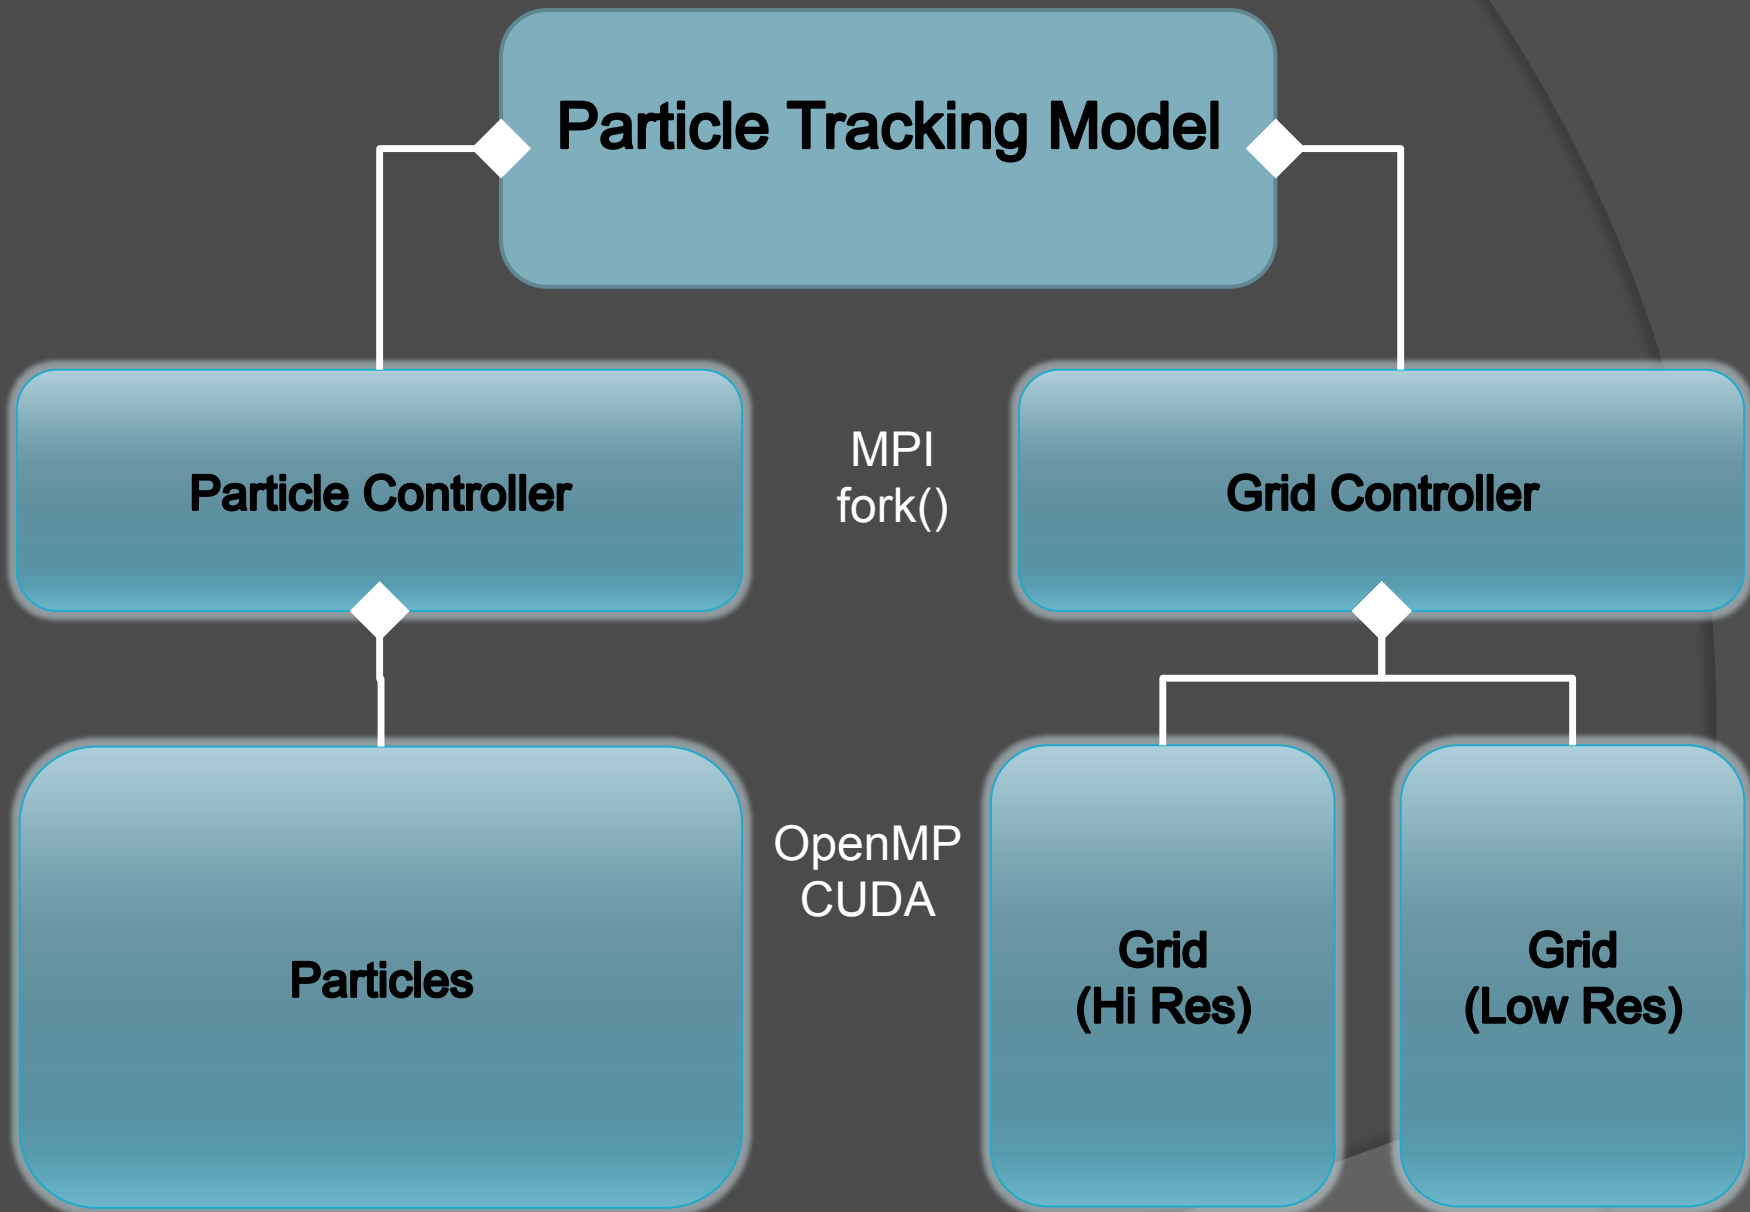


are in the life cycle of marine invertebrates (Adapted from)



# Biophysical Model Overview





# OpenMP Parallelization

Particle Tracking Model

Particle Controller

Grid Controller

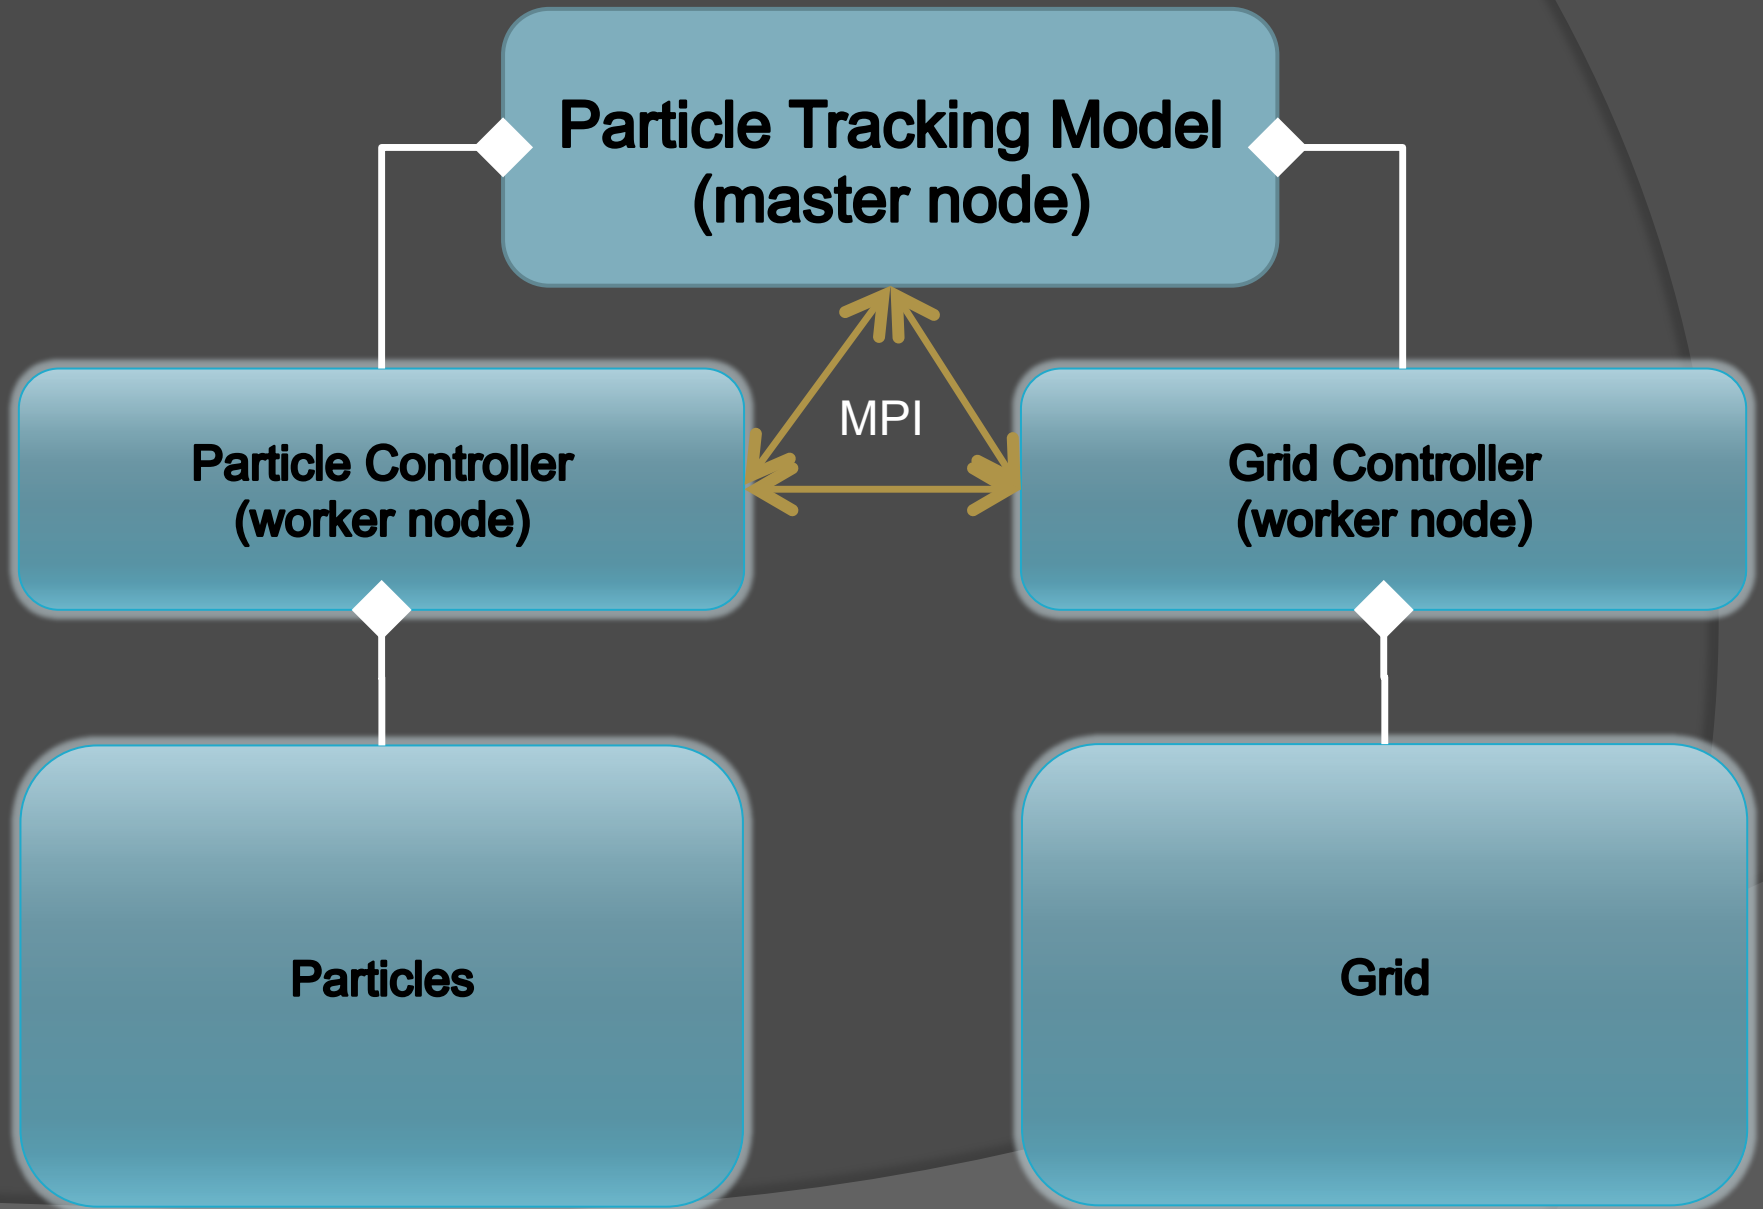
Particles

```
#pragma omp parallel for  
for(int i=0; i < n; ++i)  
  particles[i].move()
```

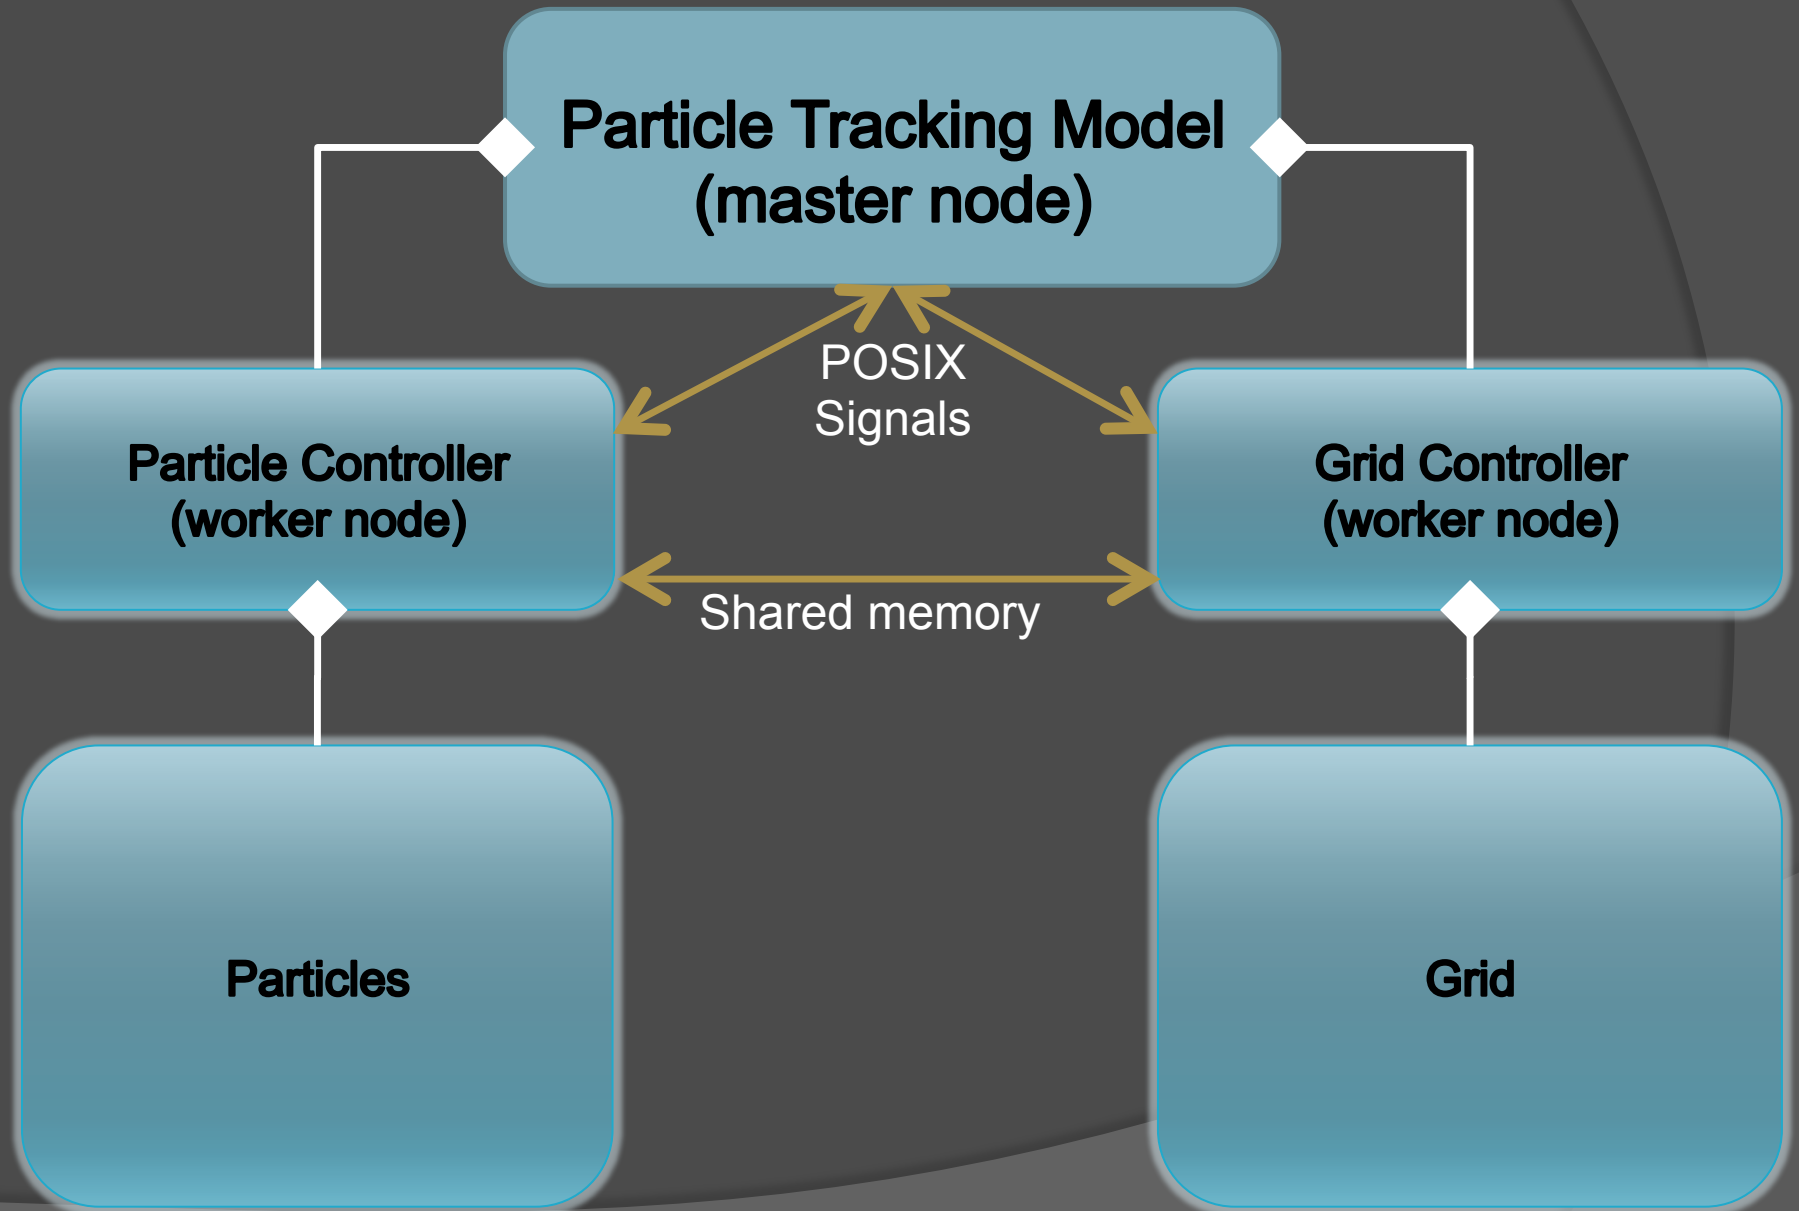
Grid

```
#pragma omp parallel for  
for(int i=0; i< npts; i++) {  
  u[i] = interp(prev_v[i],  
               next_v[i], k)  
  v[i] = interp(prev_v[i],  
               next_v[i], k)  
}
```

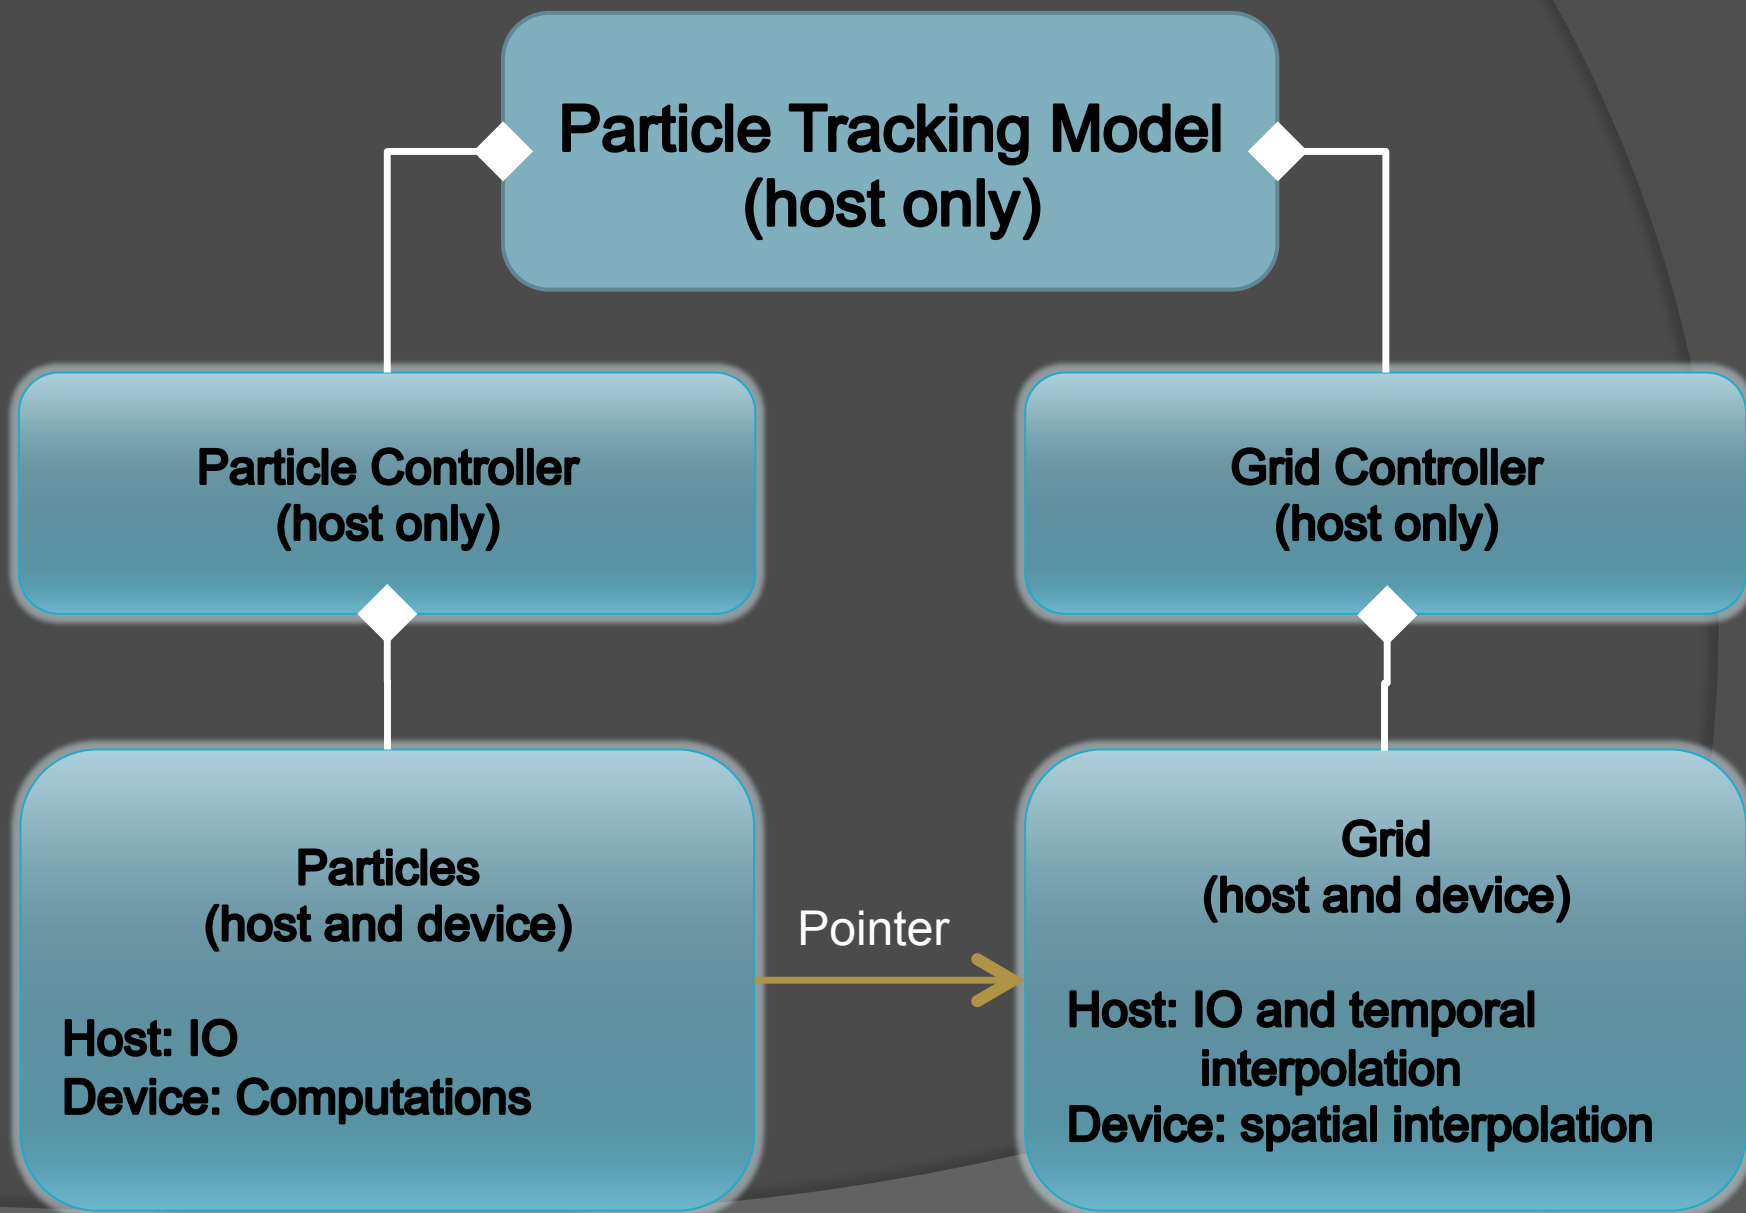
# MPI Parallelization



# Multi-process parallelization




# CUDA








# Flow field and parameterization

- 2D steadily translated Rankine vortex
- 500km X 1000km grid
- 1.5km and 4km lattice









# Model Type

Serial

CUDA

OpenMP

MPI

fork()


Operation

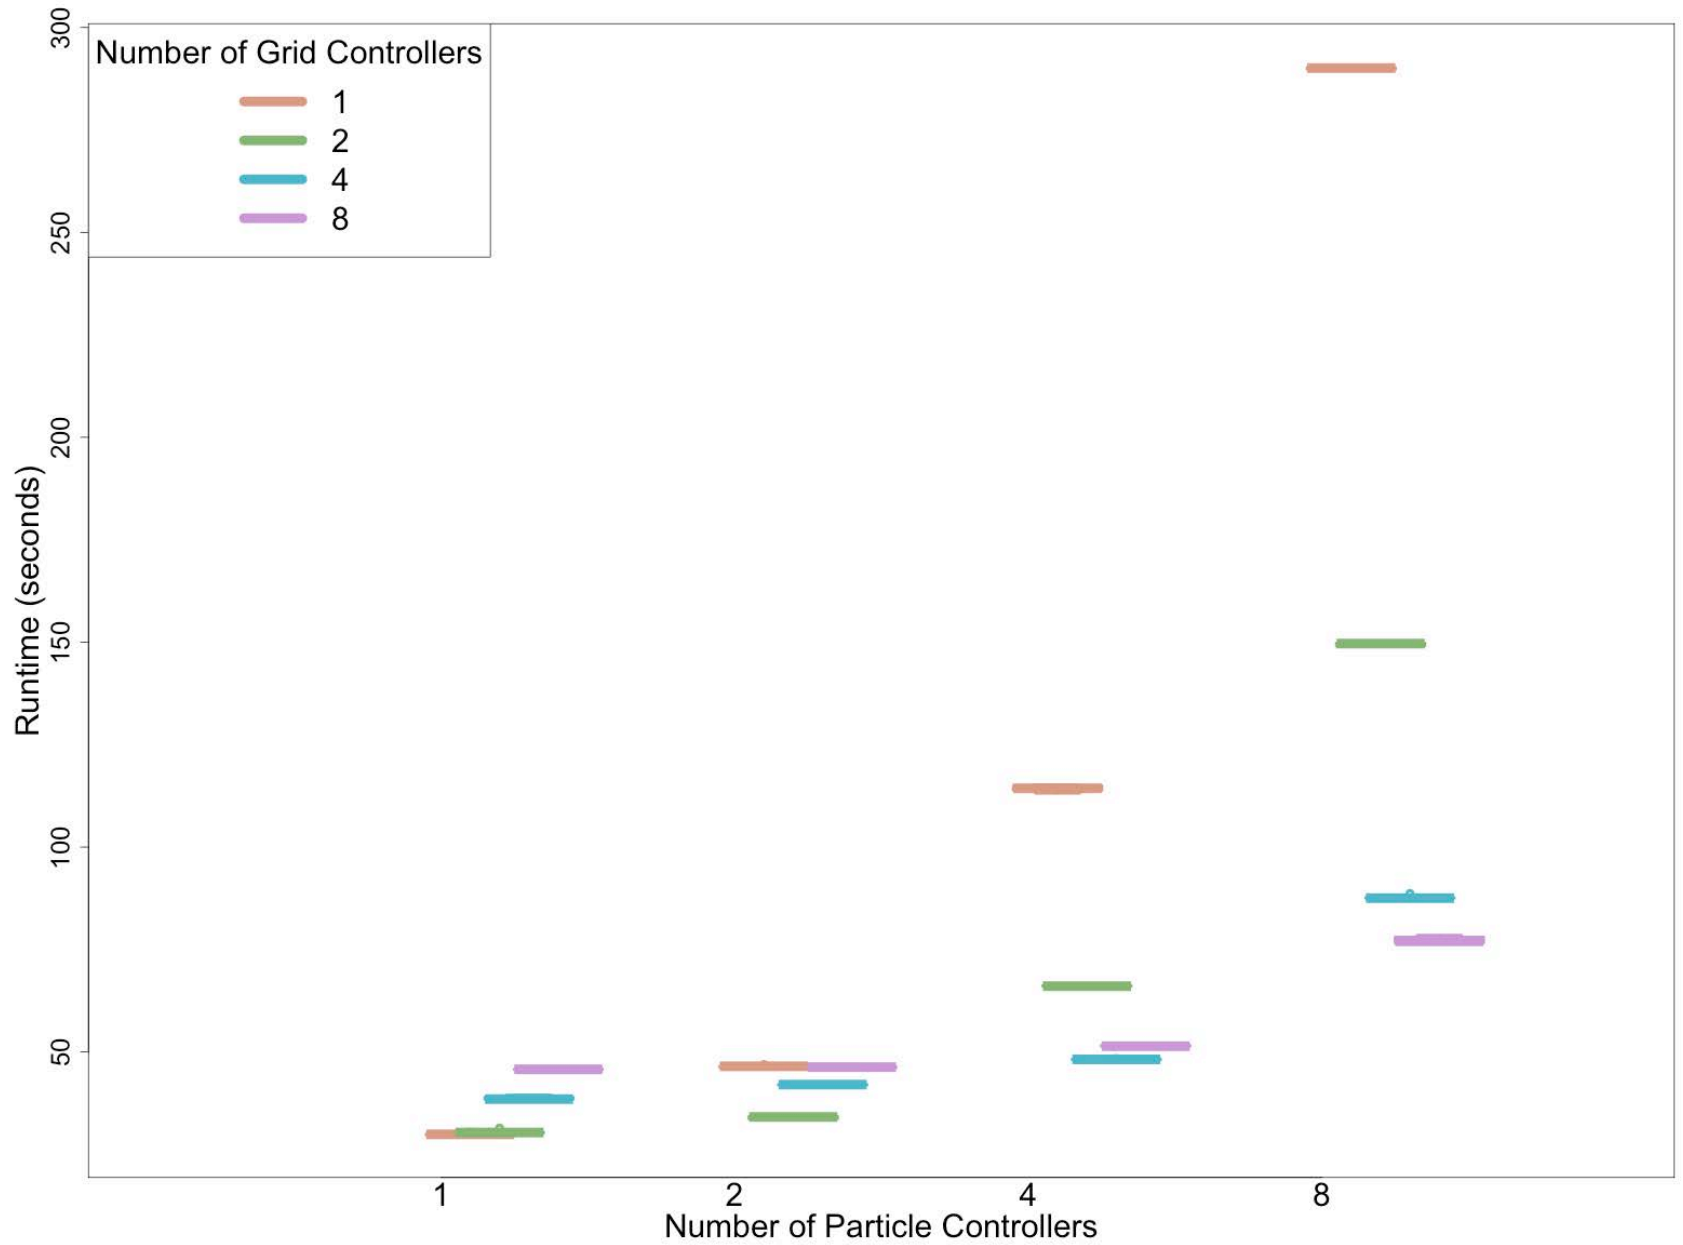
IO

Forcing

Integration

Communication





# Possible Extensions: Features

- ⦿ Support for real ocean models
  - FVCOM \*
  - HYCOM
  - ROMS
- ⦿ Advection-diffusion module \*
- ⦿ Biological module
  - Density-dependence \*
  - Growth
  - Movement

# Possible Extensions: MPI

- Try MPICH2
- Minimize communication
- Optimize arrangement of nodes
- OpenMP within each node
- Dynamically balance number of grid controllers and particle controllers



# Possible Extensions: CUDA

- ⦿ Tune threads per block vs. number of blocks
- ⦿ Reduce data movement
- ⦿ Make use of shared or thread memory
- ⦿ Move temporal interpolation to GPU
- ⦿ Integrate analyses into the model

# Conclusions

- ⦿ Optimization level matters
- ⦿ GPUs can offer cheap parallelism
- ⦿ Beware of context switching
- ⦿ Communication is expensive
- ⦿ Lots of opportunity for improvement