



Estimating the population density of unmarked animals using camera traps

A simulation-based evaluation


Louis von Leitner, Thomas Hay

Problem





How many deer are in the forest?




Simulations!



Movement Model



Movement Model




Movement Model




Movement Model

$$\text{Probability} = 1/a * \exp (d*b)$$




Weight computation takes $\frac{2}{3}$ of the runtime!

Camera trapping



Camera trapping





Camera Trapping


- Happens rarely
- Takes on average double the time as doing one animal step

Sequential Solution Ideas


Small steps and always check everything




Small steps and always check everything



Animal inside?



Small steps and always check everything



Remarks

Idea:

- No path camera tracking necessary → decrease runtime?

Answer:

- This increases runtime by a lot

Quick Math

ca. 6 month period simulation

One step every minute:

⇒ 2000 steps x 120 = 240.000

x 100 animals = 24.000.000
steps in total

One step every 2 hours:

⇒ 2000 steps

x 100 animals = 200.000
steps in total

Average time of camera
trapping = 2x time of moving
animal

Quick Math

ca. 6 month period simulation

$$\Rightarrow 24.000.000 - 200.000 = 23.800.000$$

$$23.800.000 / 2 = 11.900.000$$



Amount of Photos to be taken so that simulations take the same amount of time

Average of about 800 photos per simulation (max 200.000)


Quick Math

ca. 6 month period simulation


Sparse modelling is a lot more efficient!

Back to our method!


10 animals
100 steps



Movement times for number of steps for different numbers of animals



Movement times for numer of total steps



Performance in Sequential

Linear Scaling with amount of total steps


2000 Steps for 100 Animals takes 45 minutes.
→ Too long!

First parallelization Idea: Parallel Agents


- Initialize multiple threads
- Distribute animals evenly to threads
- Let every thread compute steps for their animals
- Gather results (photos, location data, ...)

(like the computation of pi)


Parallel Agents: Master-Slave



Parallel Agents Performance over different amounts of Threads




Parallel Agents: God-Master-Slaves



God-Master-Slave Performance

90 animals / 1000 steps each



Performance Analysis

- looks exponential
- expectation: linear

45 mins = 2700 seconds

2700 seconds / 100 threads = 27 seconds

- our model run took 52 seconds → not linear

What's going on?

Animal partitioning

- animals are discrete units and cannot be split in half
- with 100 animals and 99 threads:
 - thread 1 computes 2 animals
 - thread 2-100 compute 1 animal
 - Thread 2-100 are idle half of the time

Parallel Agents

- Number of nodes and thus performance capped by number of animals (100)
- embarrassingly parallel!
- Not linear scaling with amount of threads used

Gain: Good base comparison!


Other Idea: Optimization of thread use

Given n threads for computation, how can you leverage them best for most efficient computation?


Optimization of thread use

From last time:

add_relocation processes' times



Subprocess Times for get_step()



Optimization of thread use: Idea

Workmaster and worker nodes

Workmaster:

- Generate potential steps
- Send steps to Workers
- Do Bookkeeping (Camera detections and more data collection)
- Get Weights from Workers and determine next step

Worker:

- Get potential steps from Workmaster
- Calculate Weights
- Send Weights for steps to Workmaster


Optimization of thread use: Idea

Workmaster and worker nodes


Weight calculation $\frac{2}{3}$ of time $\Rightarrow$ 1 Master, 2 Workers

Master Worker Approach

Master Thread Worker Thread Worker Thread



Vampir Performance Analysis




 = waiting

 = doing something

computing next position and generating steps

bookkeeping



still a lot of waiting

A lot of waiting?


0.002 seconds of waiting x 2000 steps = 4 seconds of waiting per animal

⇒ 4 seconds per animal x 100 animals = 400 seconds of waiting

Master Worker v2


depends on weights

depends on location




Idea: Can Master do preparatory substeps while he is waiting?


Movement Model





Movement Model





Movement Model






generate possible steps during weight calculation



Our Chart looks like this by now ...



What we achieved:

Workers are barely idle

but Master is Idle a lot

Solution: Find optimal amount of Workers per
Master

10 Workers per Master is optimal


Lots of green! (little idling)







Performance

Master Worker Performance over different amounts of Threads






- Master Worker more efficient for lower thread numbers
- Master Worker possibility for more threads employed
- Parallel Agents more efficient around 50-100 threads



Performance

- Not linear either
- Capped by number of Animals x Threads per Workgroup
($100 \times 11 = 1100$)
- Number of Threads must be divisible by Threads per Workgroup
- 300 threads → 24 seconds runtime

What about photos?



Problem

10 Workers waiting for 1 Master for a “long” time

Idea:

For multiple Workgroups (Master + 10 Workers) assign one
“Photographer”

Master Worker v3


taking photo while animals keep running

catching up with other bookkeeping


Photographer

Master

Workers



Is this quicker?



Is this quicker?

No.

Challenges:

- a lot of communication overhead
- volatile and instable
- hard to implement (MPI buffer overwriting problem)

⇒ Previous solution was better

Conclusion

Conclusion

- didn't get linear growth in speed by resources
- reached goal: 24 seconds < 45 seconds
- Parallel Agents approach is already very quick
- Master Worker v2 is the most efficient and can scale longer

Key Lessons

- parallel programming is highly variable and unpredictable (race conditions, etc.)
 - ⇒ Make the program Waterproof
- Vampir is great to see weaknesses and if the program even works properly
- environment setup on the cluster is complex (python)
- Often the easiest solutions are the best (80:20 rule)
- dividing into subfunctions is useful for parallelizing already existing sequential code
- parallel programming is fun

Thank you!

Questions?