How to give good talks

François Trahay

Very much inspired by F. Suchanek soft skills seminar
https://suchanek.name/work/teaching/topics/good-talks/index.html

What kind of talk ?

- This talk focuses on academic talk, including :
  - Presenting your research work during a seminar/conference/project defense
  - Presenting someone else’s work during a reading group

Overview

- Know why you talk
- Identifying the audience
- Identifying the key elements of the talk
- Talk structure
- Presenting an idea/concept
- Experiments

Not in this talk:

- Overcoming shyness/stress
- Creating beautiful slideshows
- Improving your body language

More info on these topics in Fabian Suchanek's softskills seminar
https://suchanek.name/work/teaching/topics/good-talks/index.html

Know why you talk

- Do you want to
  - Teach or explain ?
    → The audience understands the key ideas
  - Make people interested ?
    → The audience wants to read the paper
  - Sell an approach ?
    → The audience wants to collaborate with you
  - Show your capacities ?
    → The audience wants to give you a grant/promotion

Identifying the audience

- What does the audience probably know ?
  - PDS seminar : knows about distributed systems, machine learning, operating systems
    - May need to explain what MPI is
  - At a HPC conference : knows HPC
    - Everybody knows what MPI is
  - At a conference/workshop on MPI :
    - Everybody knows how MPI is implemented

Identifying the key elements of the talk
The scienfic method

- A method for acquiring knowledge
  - Can be applied to CS too
- Presenting the key elements of the method is **crucial**
  - What is the problem ?
  - What is the proposed solution ?
  - Does the evaluation show that the solution fixes the problem ?

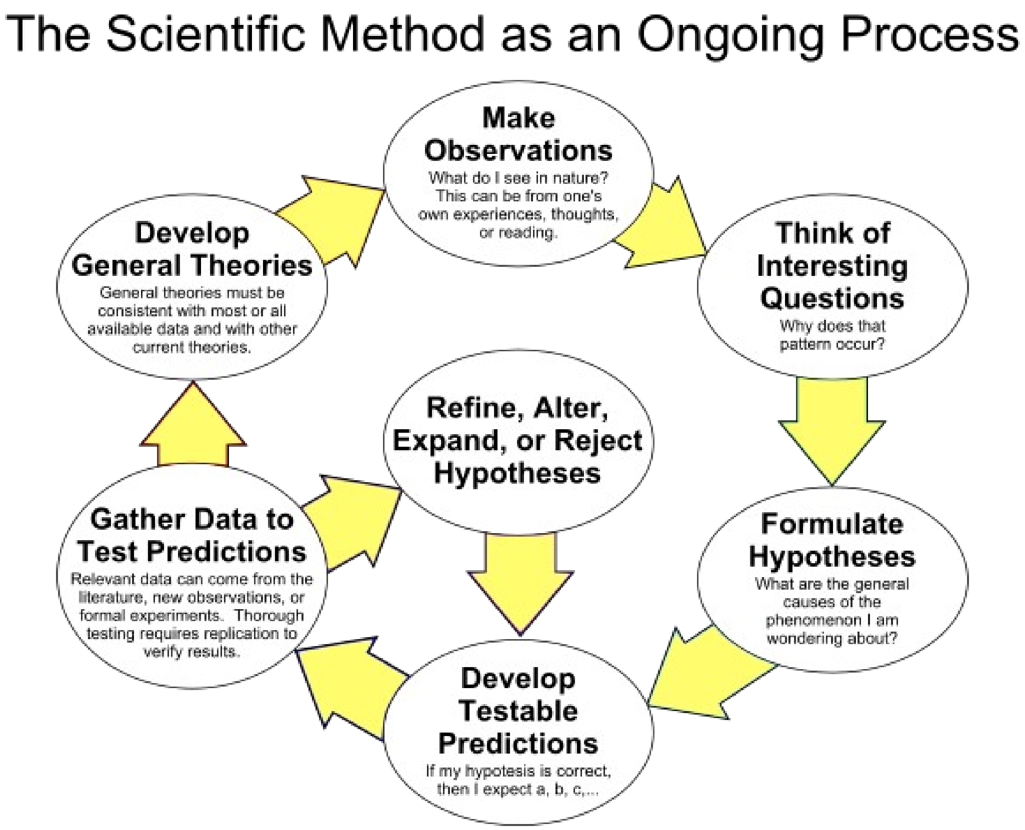

Source: wikipedia

Structure of a research presentation

The *Double funnel* scheme

- Goal of this scheme
  - Everybody gets the key ideas
  - Experts get technical details
  - Non-experts don’t get too many details

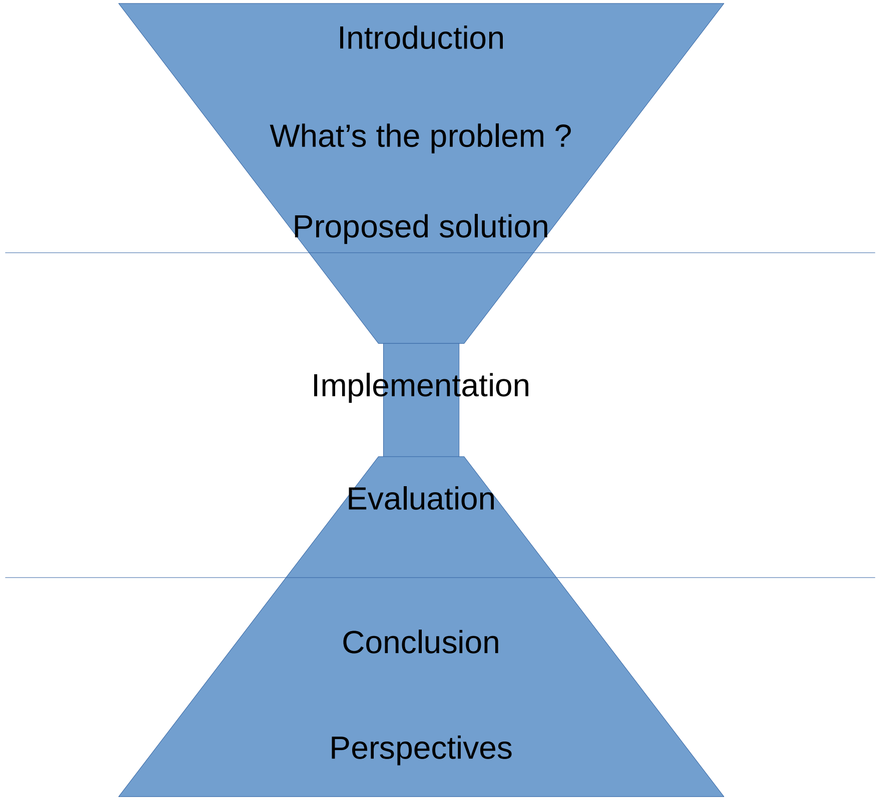

The *Double funnel* scheme : example

Fibers are not (P)Threads: The Case for Loose Coupling of Asynchronous Programming Models and MPI Through Continuations (EuroMPI'20)

- Simulation is increaslingly costly
  - Need to parallelize application and distribute them
  - Computer topology requires to mix programming models
- **Asynchronous programing model mess with MPI**
- **Proposal: introduction asynchronous tasks into MPI**
  - **Add a continuation feature to MPI**
  - **Description of the continuation API**
  - **Use a common task scheduler for both MPI and the application**
- **Evaluation :**
  - **Tasks don't regrade MPI performance in the worst case**
  - **Tasks improve MPI performance in case of high concurrence**
- **Conclusion:**
  - **Coupling MPI and tasks improves performance of asynchronous programing models**
  - **Rethink how programming models interact**

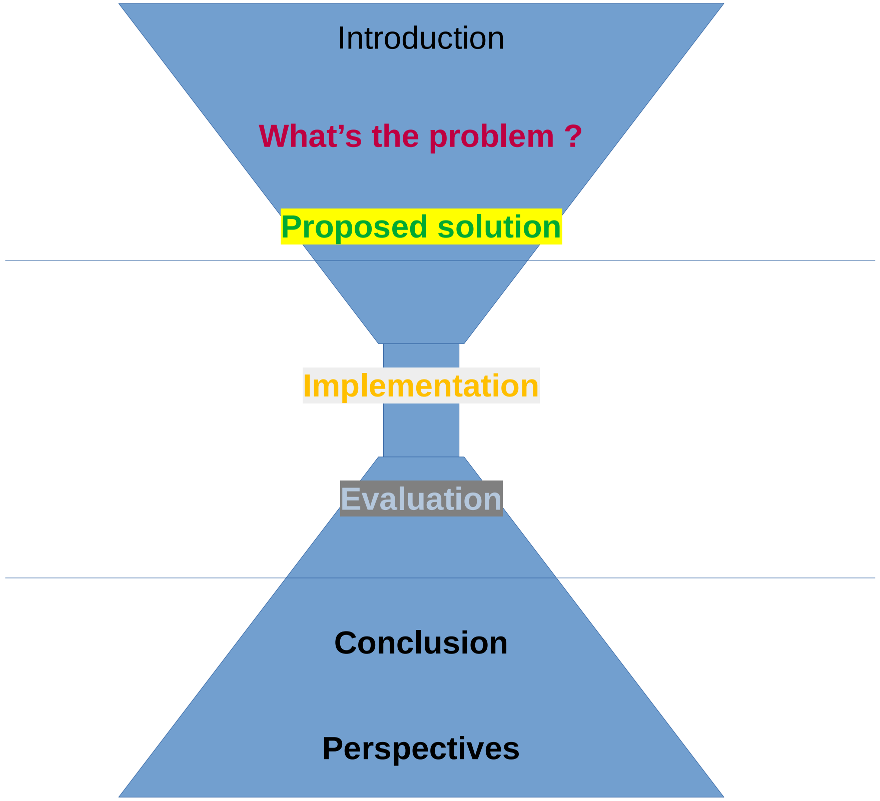

Presenting an idea/concept

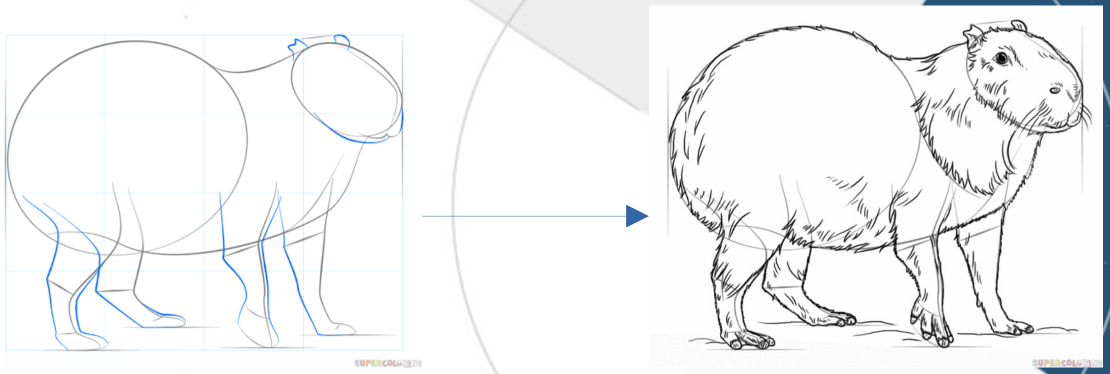

- The *Big picture* approach
  - First sketch the idea
    - *A car is a big box with wheels*
  - Then refine and add details
    - *A steering wheel allows to choose the direction. The wheels are propelled by a motor*
- The *flow* approach
  - First describe
    - *Wheels are efficient for moving things and they can be propelled by motor*
  - Then summarize
    - *Combining wheels with a large box result in a vehicle*

Running experiments

SIGPLAN Empirical Evaluation Checklist

- Evaluation is here to validate the proposed solution
  - Show that it solved the problem
- Each experiment validates one of the claims
- Adequate data analysis
  - Fair comparison with state of the art (eg. same level of optimization)
  - Repeat experiments and report variability

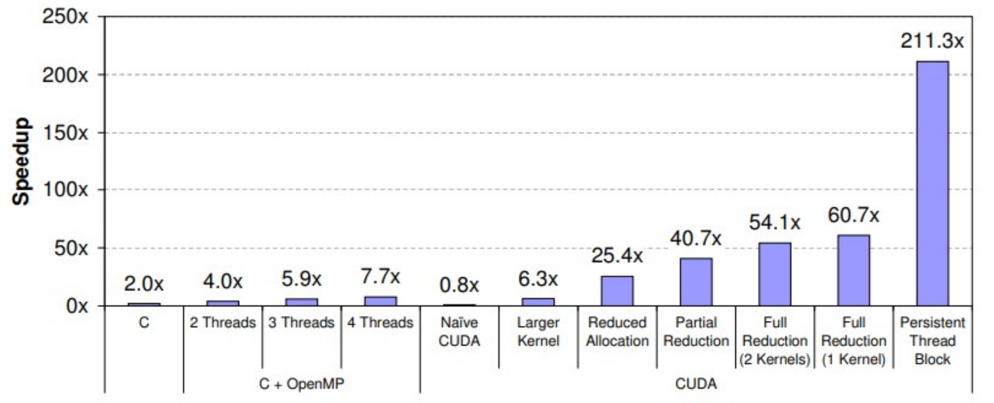

Figure 3. Speedup of the different implementations of the tracking stage over the original MATLAB implementation

Example of (possible) unfair comparison: Comparing a matlab implementation with a highly tuned CUDA implementation

Accelerating leukocyte tracking using CUDA: A case study in leveraging manycore coprocessors. In IPDPS 2009

Describing experiments

SIGPLAN Empirical Evaluation Checklist

- Experimental setup
  - Hardware/software configuration
    - *Dual-socket 12C Intel Haswell, ConnectX-3, OpenMPI 3.1.4*
  - What data is reported ?
    - *mean duration over 10 runs + standard deviation (or min/max)*
- Visualize data correctly
- Describe, then analyze
  - *For small messages, X latency is 20 % lower than Y*
  - *This is due to X that batches messages which reduce the number of system calls*
- Experiments should be reproducible
  - Some conferences require *Artifacts*

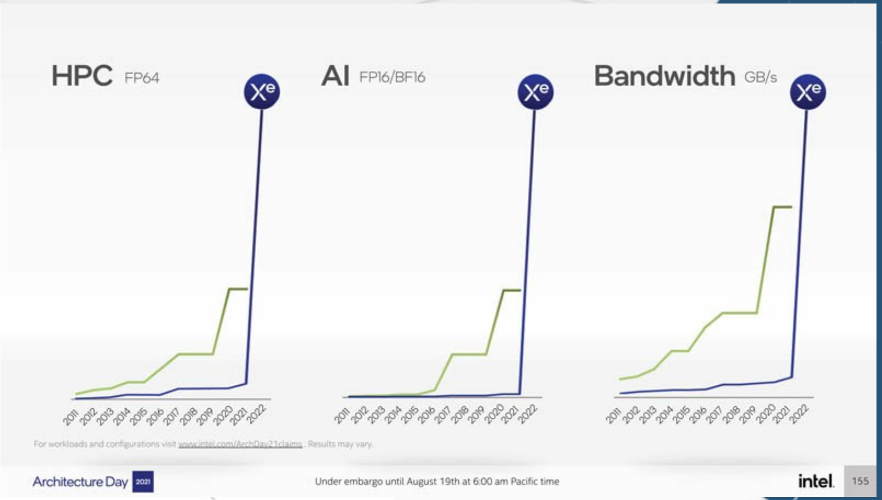

Source : Intel Architecture Day 2021
https://www.servethehome.com/intel-ponte-vecchio-is-a-spaceship-of-a-gpu/

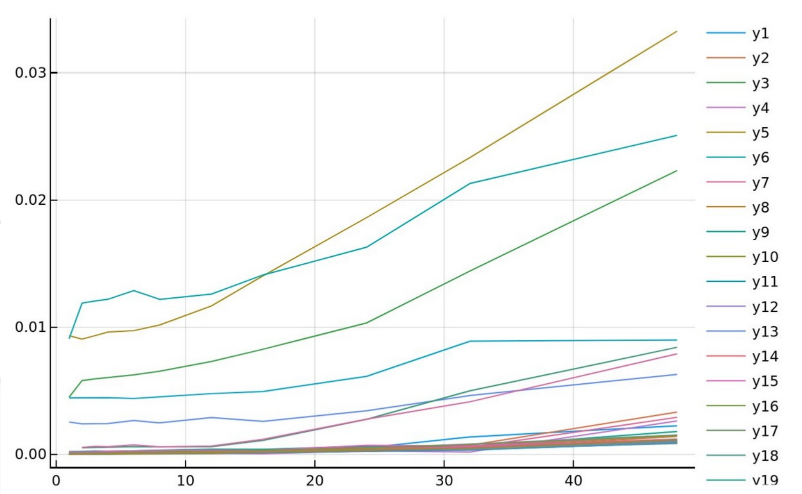

Reproducible Research in Computer Science *https://hal.inria.fr/hal-01110206/*

Conclusion

- Make sure the audience identifies
  - The problem you try to solve
  - The proposed solution
- As part of the reading group, you are evaluated on
  - Your understanding of the paper
  - Your ability to show the key points of the paper
  - Your communication skills
- This applies to other research presentations