Heterogeneity in Computing: Now and in the Future

Anne Benoit

LIP, Ecole Normale Supérieure de Lyon, France

Anne.Benoit@ens-lyon.fr

http://graal.ens-lyon.fr/~abenoit/

HCW workshop, in conjunction with IPDPS
Rio de Janeiro, Brazil, May 20, 2019

HCW Panel - *Heterogeneity in Computing: Now and in the Future*
A few words about me

Grenoble, France
1995-1997: Math studies
1997-2000: Engineer school
2000-2003: PhD thesis
Performance evaluation, Markov chains

Edinburgh, UK
2003-2005: Post-doc
Algorithmic skeletons

ENS Lyon, France
2005-Present: Associate Prof.
Multi-criteria scheduling, resilience, energy, memory, …
Georgia Tech, Atlanta, USA
2017-2018: Visiting Ass. Prof.

Anne

Julie, 2012 Sophie, 2014

Program (Papers) Chair for HiPC'16, ICPP'17, SC'17, IPDPS'18

Head of Fundamental CS Master @ ENS Lyon (2015-2017)
Head of Third-year students (2018-Present)

AE (in Chief) of Parco, AE of TPDS

Question 1: Past examples of HC

- *What are examples of HC* **(Heterogeneity in Computing)** *that began as research ideas and are now mainstream?*
- Where did we start?

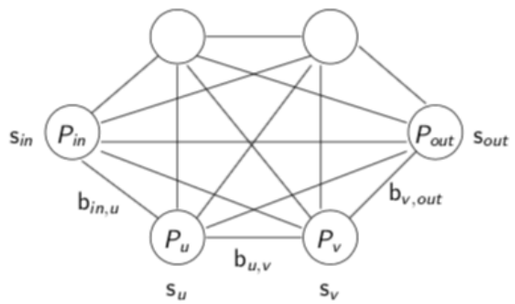

- p processors P_u, $1 \leq u \leq$ p, fully interconnected
- s_u: speed of processor P_u
- bidirectional link $link_{u,v} : P_u \rightarrow P_v$, bandwidth $b_{u,v}$
- one-port model: each processor can either send, receive or compute at any time-step

- General heterogeneous platform model I used (2005-2012)

Question 1: Past examples of HC

- Different levels of heterogeneity

Fully Homogeneous – Identical processors ($s_u = s$) and links ($b_{u,v} = b$): typical parallel machines

Communication Homogeneous – Different-speed processors ($s_u \neq s_v$), identical links ($b_{u,v} = b$): networks of workstations, clusters

Fully Heterogeneous – Fully heterogeneous architectures, $s_u \neq s_v$ and $b_{u,v} \neq b_{u',v'}$: hierarchical platforms, grids

- Heterogeneous computing system: diverse computing resources, either local or geographically distributed
- Using these resources $\rightarrow$ cluster computing, grid computing, cloud computing

Question 1: Past examples of HC

- Grids and Clouds are now mainstream
  $\rightarrow$ Theoretical and practical research on heterogeneous computing environments has been leading the way towards efficient use of these platforms
- Look up *heterogeneous systems* on Google scholar since 2018/2015: 64k / 772k references
- What about clusters, grids, clouds, fogs? (in k references, since 2018/2015)

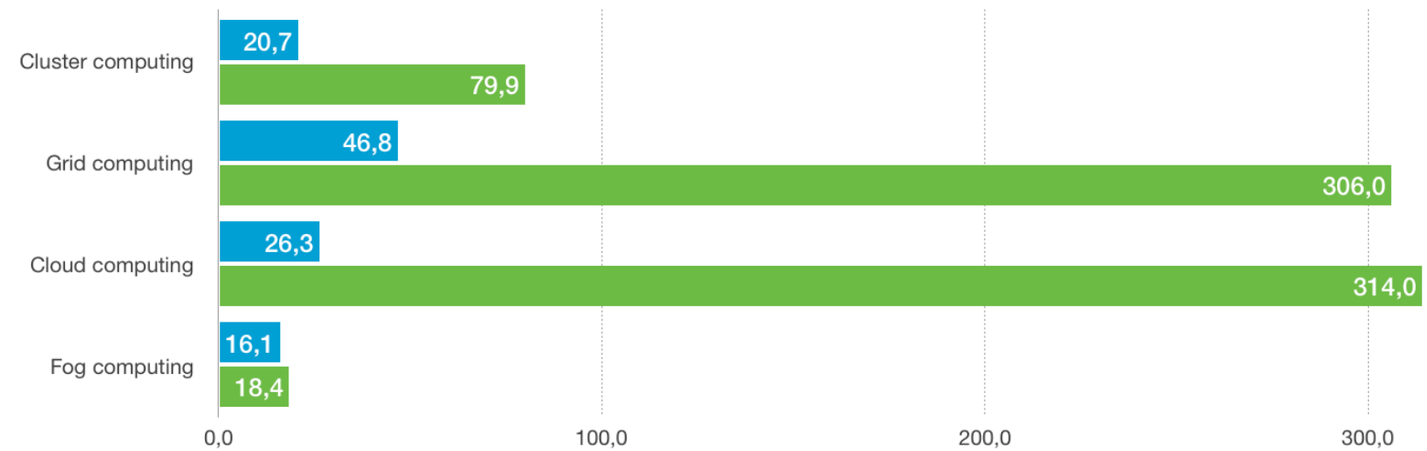

Question 1: Past examples of HC

- *From the past to the present...*

- Besides these distributed heterogeneous platforms, clusters and supercomputers have more and more homogeneous nodes/cores
- Heterogeneity through GPUs: the first two top-500 supercomputers (Summit and Sierra) are IBM-built supercomputers, powered by Power9 CPUs and NVIDIA V100 GPUs
- GPU computing Google scholar count since 2018: 22k
- CPU and GPU approach: combine the best features of both PUs

Question 1: Past examples of HC

- *From the past to the present...*

- Besides these distributed heterogeneous platforms, clusters and supercomputers have more and more homogeneous nodes/cores
- Heterogeneity through GPUs: the first two top-500 supercomputers (Summit and Sierra) are IBM-built supercomputers, powered by Power9 CPUs and NVIDIA V100 GPUs
- GPU computing Google scholar count since 2018: 22k
- CPU and GPU approach: combine the best features of both PUs

Question 2: Future of HC

- *What are the future aspects of HC that will be critically important for next generation computing systems?*
- I have two answers: energy and resilience!
- Back in 2014, Advanced Scientific Computing Advisory Committee (ASCAC) published top ten Exascale research challenges to achieve the development of an Exascale system. Energy and resilience appear as major challenges!

Question 2: Future of HC

- *What are the future aspects of HC that will be critically important for next generation computing systems?*
- I have two answers: energy and resilience!
- Back in 2014, Advanced Scientific Computing Advisory Committee (ASCAC) published top ten Exascale research challenges to achieve the development of an Exascale system. Energy and resilience appear as major challenges!

Question 2: Future of HC - Energy

"The internet begins with coal"


- Nowadays: more than 90 billion kilowatt-hours of electricity a year; requires 34 giant (500 megawatt) coal-powered plants, and produces huge CO_2 emissions
- Explosion of artificial intelligence; AI is hungry for processing power!
  Need to double data centers in next four years $\rightarrow$ how to get enough power?

Energy and power awareness $\rightsquigarrow$ crucial for both environmental and economical reasons
Heterogeneous computing: may help by clever mix of CPUs and GPUs

Question 2: Future of HC - Resilience

- Consider one processor (e.g. in your laptop)
  - Mean Time Between Failures (MTBF) = 100 years
  - (Almost) no failures in practice ☺

  Why bother about failures?

- **Theorem:** The MTBF decreases linearly with the number of processors!
  With 36500 processors, a failure per day on average!

A large simulation can run for weeks, hence it will face failures ☹
And then, consume even more **energy** ☹

- Heterogeneous computing: Account for different kinds of processors (with different failure rates/speeds) and be even more reliable

Question 2: Future of HC - Resilience

- Consider one processor (e.g. in your laptop)
  - Mean Time Between Failures (MTBF) = 100 years
  - (Almost) no failures in practice ☺

  Why bother about failures?

- **Theorem:** The MTBF decreases linearly with the number of processors!
  With 36500 processors, a failure per day on average!

A large simulation can run for weeks, hence it will face failures ☹
And then, consume even more **energy** ☹

- Heterogeneous computing: Account for different kinds of processors (with different failure rates/speeds) and be even more reliable

Question 2: Future of HC - Resilience

- Consider one processor (e.g. in your laptop)
  - Mean Time Between Failures (MTBF) = 100 years
  - (Almost) no failures in practice ☺

  Why bother about failures?

- **Theorem:** The MTBF decreases linearly with the number of processors!
  With 36500 processors, a failure per day on average!

A large simulation can run for weeks, hence it will face failures ☹
And then, consume even more **energy** ☹

- Heterogeneous computing: Account for different kinds of processors (with different failure rates/speeds) and be even more reliable

Question 2: Future of HC - Resilience

- Replicate work on two platforms running at different speed:

Assume $\sigma_1 > \sigma_2$ (P_1 is faster than P_2).

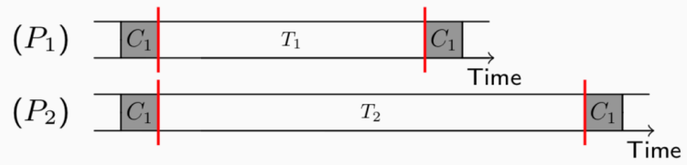

	P_1	P_2
Speed	σ_1	σ_2
MTBF	μ_1	μ_2
Period length	$T_1 = \frac{W}{\sigma_1}$	$T_2 = \frac{W}{\sigma_2}$

Objective: Find the optimal period W^*.

- Optimal period length? See [Benoit et al., Optimal checkpointing period with replicated execution on heterogeneous platform, FTXS'2017]
- Aim at minimizing energy consumption
- Still a lot of open problems, and a lot to do for our planet...

Question 3: Other HC

- *Please feel free to briefly discuss an additional important topic related to HC that is not incorporated by your answers to questions 1 and 2.*

- Dynamic environments: unpredictable execution times, failures...
- Leads to even more heterogeneity

- For instance, you do not know for how long a task will take to execute on a given processor, and whether it will be hit by a failure

- ... And if not mentioned before, of course, dealing with data distribution in heterogeneous environments!
- Beaumont et al.: *Partitioning a square into rectangles* (2002), *Matrix partitioning for parallel computing on heterogeneous platforms* (2018), and Ravi's HCW'19 talk ☺

Question 3: Other HC

- *Please feel free to briefly discuss an additional important topic related to HC that is not incorporated by your answers to questions 1 and 2.*

- Dynamic environments: unpredictable execution times, failures...
- Leads to even more heterogeneity

- For instance, you do not know for how long a task will take to execute on a given processor, and whether it will be hit by a failure

- ... And if not mentioned before, of course, dealing with data distribution in heterogeneous environments!
- Beaumont et al.: *Partitioning a square into rectangles* (2002), *Matrix partitioning for parallel computing on heterogeneous platforms* (2018), and Ravi's HCW'19 talk ☺