

# When Are You Ready to Submit an ERC (Advanced) Grant Proposal?

Lieven Eeckhout  
Ghent University

FWO/NCP ERC Info Session  
Brussels – June 15, 2017



**GHENT  
UNIVERSITY**

# Who Am I?

PhD in 2002 from Ghent University

2003-2009: FWO postdoc

2006-2009: assistant professor (10%)

2009-2012: associate professor (100% tenured)

2012-present: professor

**Field of expertise: computer architecture**

Faculty of Engineering and Architecture at Ghent University

My work fits ERC PE6: Computer Science and Informatics

# My ERC Trajectory

## **2010: ERC Starting Grant**

- *Dependable Performance on Many-Thread Processors*

## **2012: ERC Proof-of-Concept**

- *Data Center Monitoring for Improving Insight and Efficiency*

## **2015: ERC Proof-of-Concept**

- *High-Speed Architectural Simulation of ARM-based Systems*

## **2016: ERC Advanced Grant**

- *Load Slice Core: A Power and Cost-Efficient Microarchitecture for the Future*

# When Are You Ready to Submit an ERC (Advanced) Grant Proposal?

Your CV and Track Record

The Project

*[Disclaimer: This is just based on my perspective and experience on the whole process]*

# Your CV and Track Record

## Key mission

- Convince the panel that you are the forefront of your research field
- And this may be (very) different for everybody
  - Highlight your key strengths and accomplishments

# Some Suggestions

## **Publish in top-tier venues**

- *Quality is way more important than quantity!*

## **Explain your key contributions and how they have impacted the field**

- *Changed current practice, moved the state-of-the-art, industry usage of your technology, citations, downloads, awards, patents, spin-offs*

## **Demonstrate that you are recognized as a world expert by your peers**

- *Serve on or chair technical program committees, associate editor, editor-in-chief, expert service, etc.*

## **Demonstrate you are internationally active**

- *Research mobility, international collaborations*

## **Demonstrate that you can manage research**

- *List prior research endeavors and funding, explain your role and contribution*

# Is Now the Right Time to Apply?

**I knew my strengths (and weaknesses)**

**But I had some concerns**


- Am I senior/old enough?
- Is my research group big enough?
- Is my h-index high enough?

**... should I apply now or wait a little longer?**

# Am I Old/Senior Enough?

I was 40 years old on Jan 1, 2016

## 2016 STG-COG-ADG Calls Age of grantees





# Is My Research Group Big Enough?

People in my faculty with an ERC AdG

	#professors	#postdocs	#PhD students
Roel Baets	8	17	60
Piet Demeester	22	27	82
Geert De Schutter	6	40	
Guy Marin	10	15	55

My research group: 1 postdoc + 8 PhD students

# Is My H-index High Enough?

My survey of 2014 and 2015 ERC AdG PE6 grant holders

- H-index (Google Scholar) ranging from 35 to 60+

My h-index in summer 2016: 37

BUT I'm young...

# Bottom line

**I knew my strengths, weaknesses and concerns  
but in the end I decided to go for it**

- Take-away message: Don't self-sensor

**Convince your panel that you are at the forefront  
of your research field (in Europe)**

- With your particular strengths and accomplishments

**Try to impress them**

- but don't overdo it

# The project

**The most important thing**

**Key idea of proposal must**

- create a ‘wow’-feeling
- be relevant
- be high impact
- be high-risk/high-gain

**You must be the ideal person for the job**

# What is a high-risk/high-gain proposal?

What follows is my own experience

- Not just with ERC but also with other funding agencies

Pose high-impact hypothesis and objective

- And provide preliminary data to support this

*Examples that follow are taken from my 2010 StG and 2016 AdG applications*

# Fundamental problem in many-thread processors

System software assumes threads make equal progress

Threads do not make equal progress due to resource sharing

Major concern for future applications

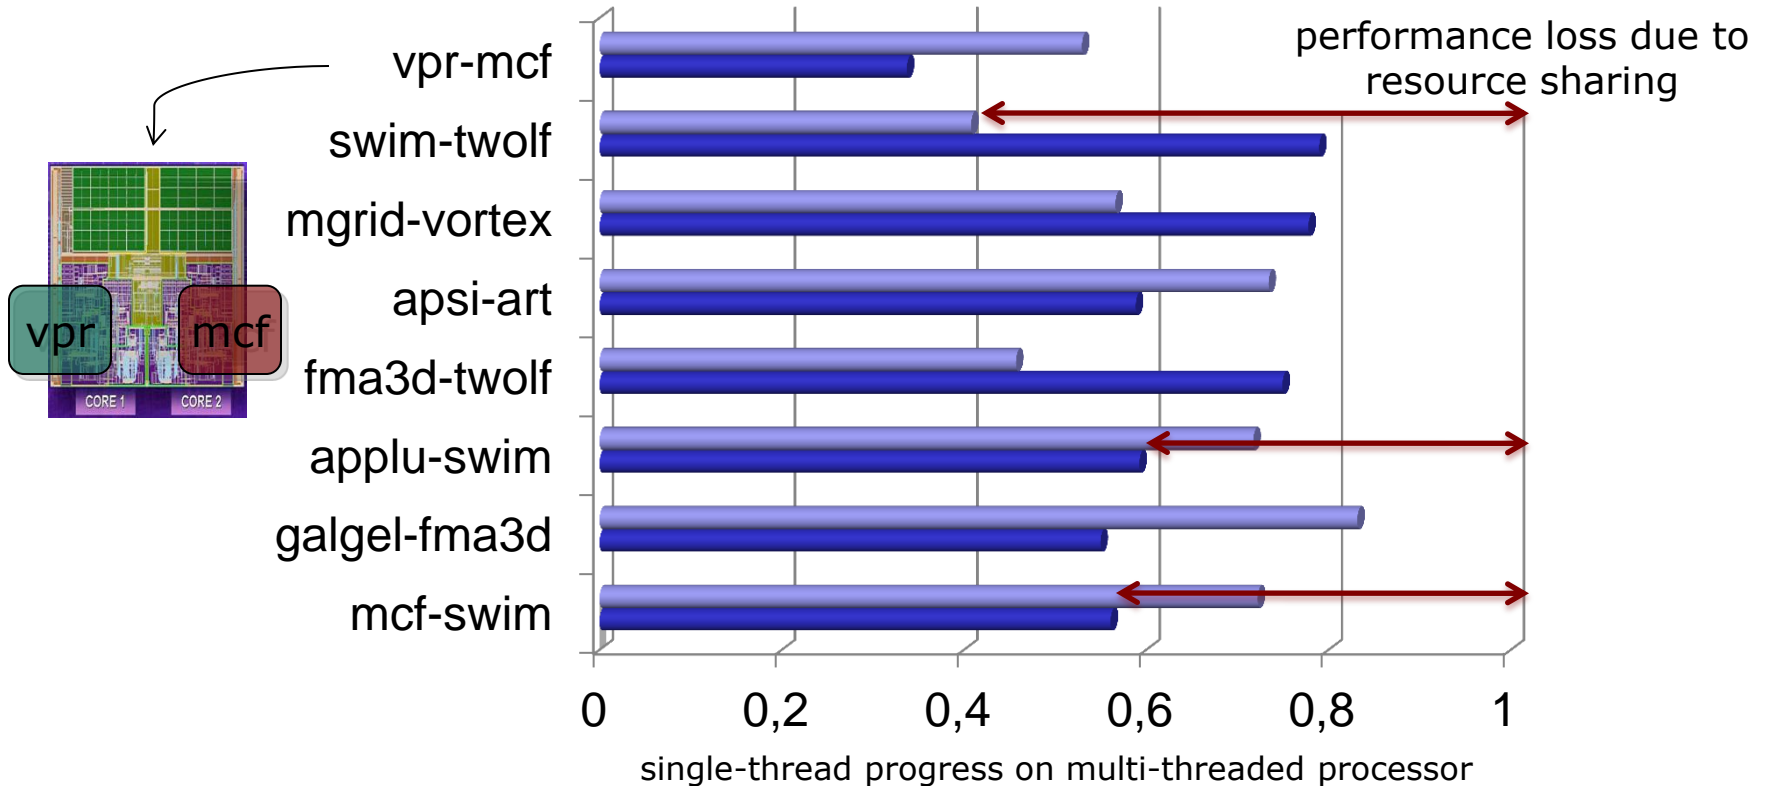
**Real-time embedded:** missed deadlines, uneconomical safety margins

**High performance computing:** load imbalance in parallel workloads

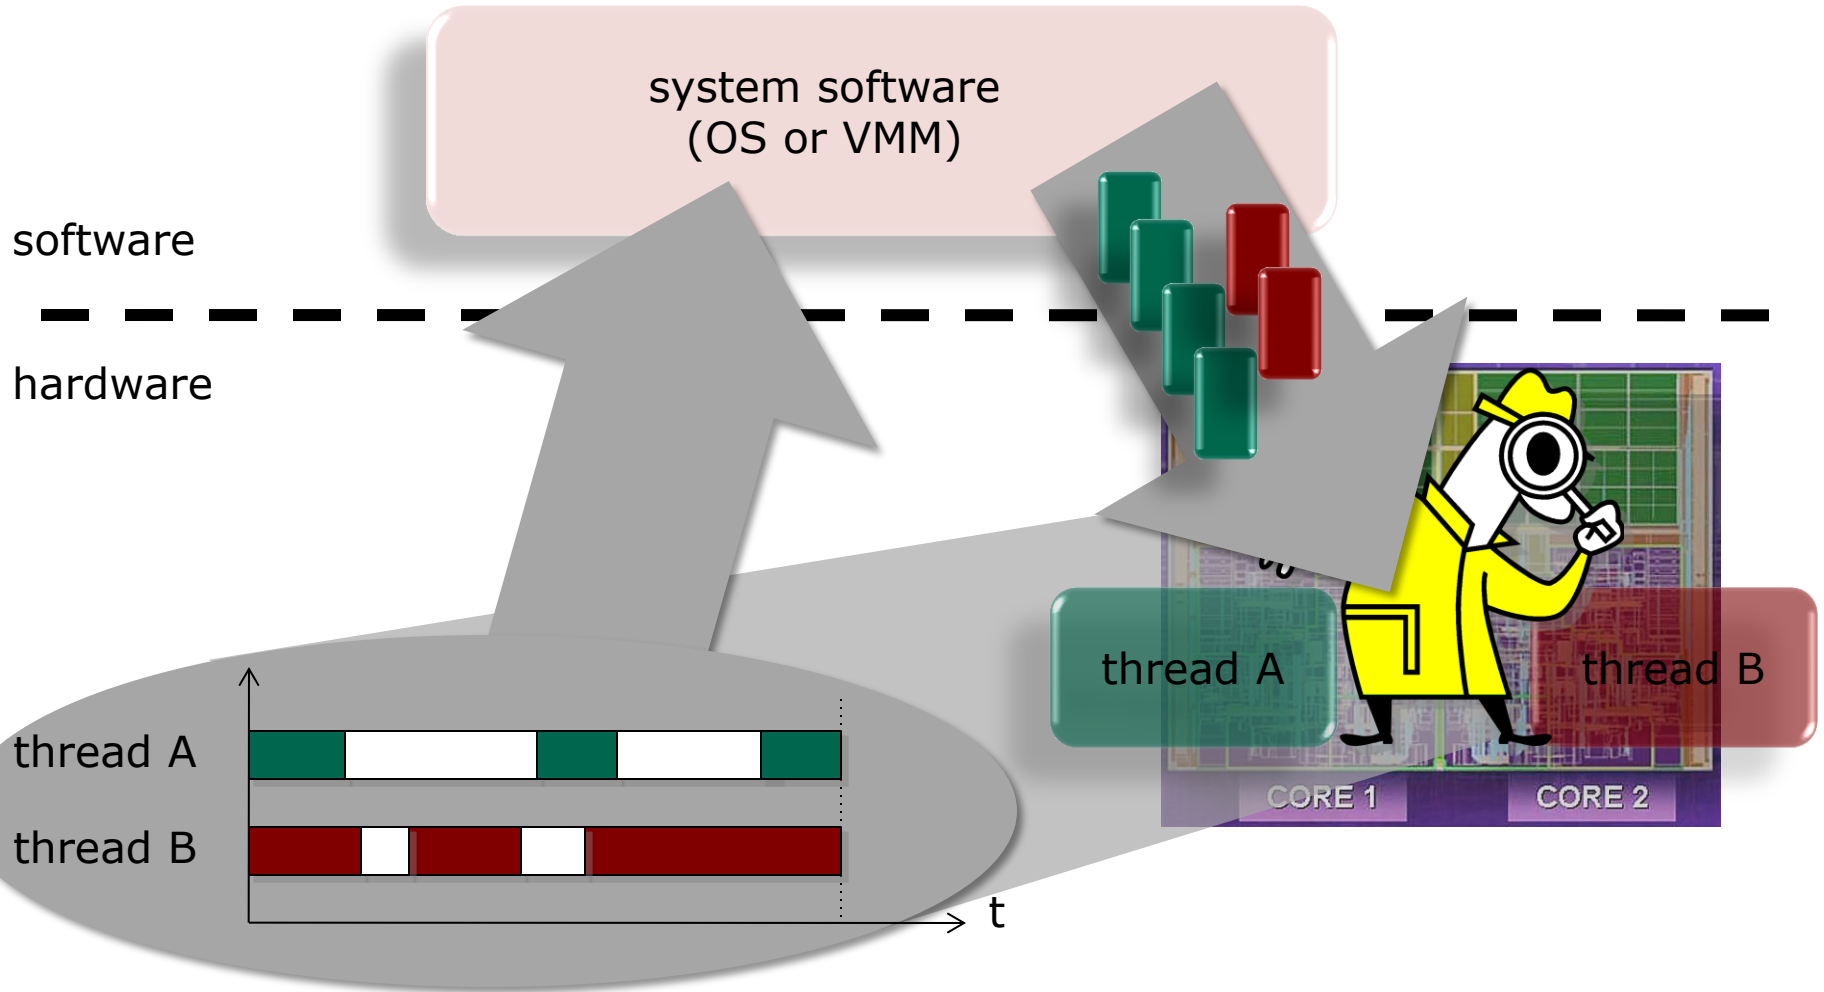
**Datacenters, the cloud:** large and variable response times

# Non-dependable performance

*Some threads make considerably faster progress than others depending on the execution context*



# The DPMP proposal





# Novel paradigm for HW/SW performance interaction on many-thread processors

## Fundamental problem

- Non-dependable performance on many-thread processors
- System software is unaware of thread progress

## Solution

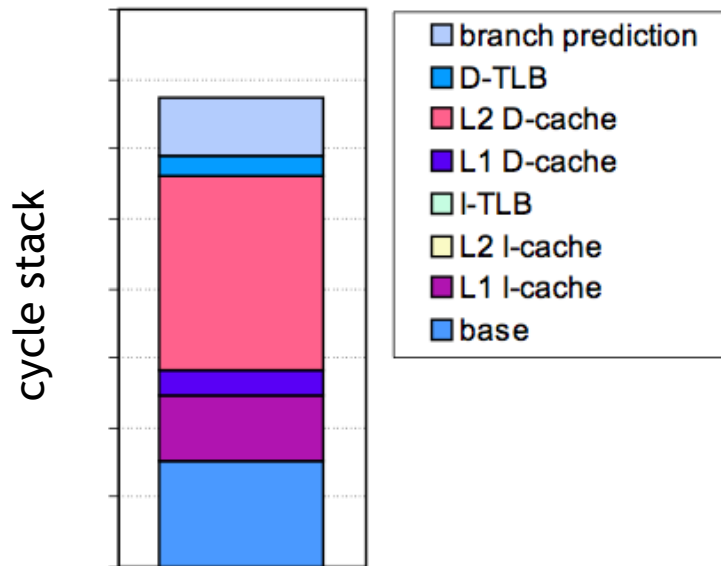
- Key ideas: performance introspection and thread progress aware scheduling and resource management
- Key novelty: based on well-founded analytical modeling

## Impact

- Novel paradigm for HW/SW performance interaction on many-thread processors
- Better system throughput, bounded response times, meet deadlines, balanced parallel performance, better QoS and SLA on future many-thread processors

# Key idea: Performance introspection

Per-thread cycle accounting: estimate per-thread progress during multi-threaded execution



*Analytical modeling based on first principles:*

$$T = \frac{N/D + \frac{m_{L1I\$} \times l_{L1I\$} + m_{br} \times l_{br} + m_{L2D\$} \times l_{L2D\$}}{MLP}}{\quad} \quad // \text{ useful work}$$

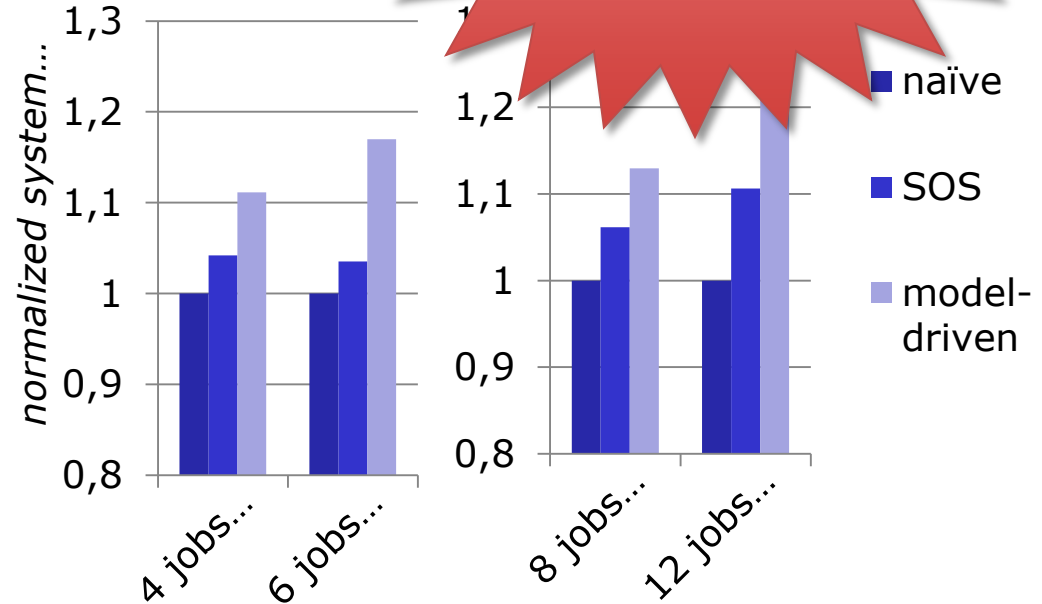
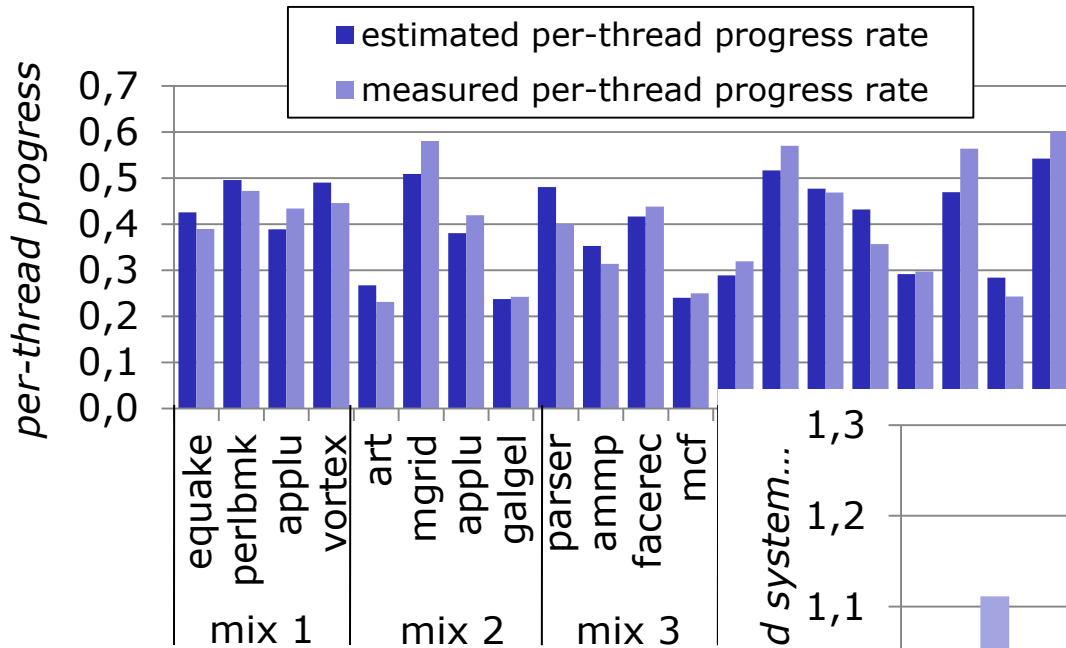
*m<sub>L1I\$</sub> × l<sub>L1I\$</sub> // L1 I-cache misses*

*m<sub>br</sub> × l<sub>br</sub> // branch mispredicts*

*m<sub>L2D\$</sub> × l<sub>L2D\$</sub> // L2 D-cache misses*

[ACM Transactions on Computer Systems, 2009  
IEEE Micro Top Picks, 2007]

# Preliminary results on SMT processor cores are very promising



[ASPLOS, 2009 & 2010]


# ERC AdG: Load Slice Core

## In-order processor

- *High power-efficiency*
- *High cost-efficiency*
- *4 decades old*

## Out-of-order processor

- *High performance*
- *2 decades old*




**Given current design constraints: What we really need is high performance in a cost and power-efficient way**

# ERC AdG: Load Slice Core

“We propose the *Load Slice Core (LSC) microarchitecture* [...]

Experimental results published at the 2015 International Symposium on Computer Architecture (ISCA), the flagship conference in the field of computer architecture, report that the Load Slice Core delivers 4.7 times higher performance per Watt than an out-of-order core [6]. Taking cost into account as well, we find that the *Load Slice Core delivers nearly 8 times higher performance per Watt per euro compared to an out-of-order core.* [...]

These preliminary results suggest that *the Load Slice Core could potentially be a game-changing core microarchitecture, which is the key motivation for submitting this project proposal.*”



# Take Your Time

## Developing key idea in proposal takes time

- A year is normal
- Needs to be a ‘big’ idea, high-risk/high-gain
  - E.g., 10× improvement, paradigm shift, novel solution to a long-standing problem, fundamentally new contribution, etc.
- Needs to be timely, relevant, high-impact
- Convince the panel you are the ideal person for the job
  - Expertise
  - Preliminary data to support the hypothesis and to demonstrate the objectives are achievable

**The actual proposal writing takes much less time: 3 to 4 weeks**

# Some additional thoughts

**Make sure your proposal is written with both the expert and not-so-expert in mind**

- B1 reviewed by panel; B2 goes to external (expert) reviewers

**First write B2 – then write B1**

**Try to make your proposal visually attractive and different from other proposals**

- It needs to stand out!
- Use typographic elements and figures on every page

# Thank you

And good luck!



**GHENT  
UNIVERSITY**

Lieven Eeckhout