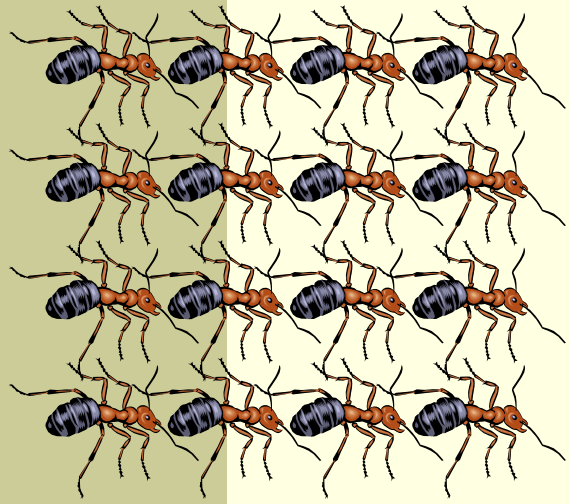

The Angstrom Project:
The Angstrom Project: Building 1000-Core Computer Systems

Anant Agarwal
CSAIL, MIT

CSAIL

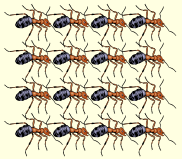

MIT's largest laboratory with ~1000 members

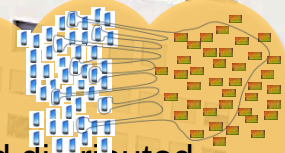

Systems:

- Parallel and distributed systems
- Wireless protocols and coding
- Mobile and mesh networks
- Relational databases
- Security & recovery
- Medical Telepresence

Stata Center

© 2006 http://philip.greenspun.com

MIT's largest laboratory with ~1000 members

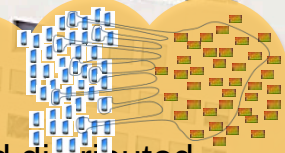

Systems:

- Parallel and distributed systems
- Wireless protocols and coding
- Mobile and mesh networks
- Relational databases
- Security & recovery
- Medical Telepresence

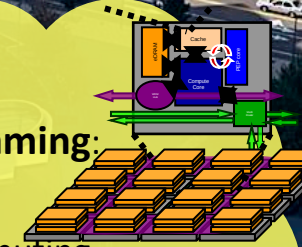

Architecture & Programming:

- Manycore architectures
- Organic or self-aware computing
- Languages for scalable computing
- Reconfigurable HW, Rapid Prototyping
- Provably Reliable Software
- Program analysis

© 2006 http://philip.greenspun.com

MIT's largest laboratory with ~1000 members

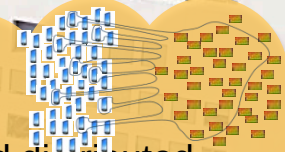

Systems:

- Parallel and distributed systems
- Wireless protocols and coding
- Mobile and mesh networks
- Relational databases
- Security & recovery
- Medical Telepresence

Theory:

- Theory of distributed systems
- Cryptography & Information Security
- Mechanism Design
- Quantum Information Science
- Computational Biology

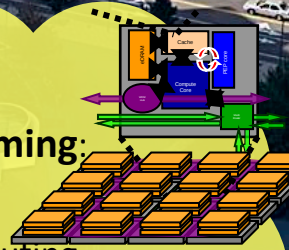

Architecture & Programming:

- Manycore architectures
- Organic or self-aware computing
- Languages for scalable computing
- Reconfigurable HW, Rapid Prototyping
- Provably Reliable Software
- Program analysis

© 2006 http://philip.greenspun.com

MIT's largest laboratory with ~1000 members

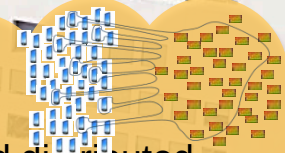

Systems:

- Parallel and distributed systems
- Wireless protocols and coding
- Mobile and mesh networks
- Relational databases
- Security & recovery
- Medical Telepresence

Theory:

- Theory of distributed systems
- Cryptography & Information Security
- Mechanism Design
- Quantum Information Science
- Computational Biology

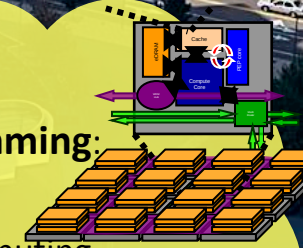

Architecture & Programming:

- Manycore architectures
- Organic or self-aware computing
- Languages for scalable computing
- Reconfigurable HW, Rapid Prototyping
- Provably Reliable Software
- Program analysis

Human/Computer Interactions:

- Spoken language systems
- Graphics, Vision, Image processing
- Natural language understanding
- Gesture-based interfaces
- Web automatio
- Crowd sourcing

© 2006 http://philip.greenspun.com

MIT's largest laboratory with ~1000 members

CSAIL

AI & Robotics:
- Intelligence Initiati
- Medical decision making
- Machine Learning
- Autonomous vehicles
- Robot locomotion & control

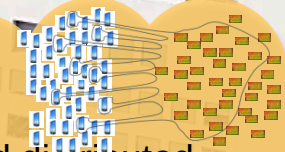

Systems:
- Parallel and distributed systems
- Wireless protocols and coding
- Mobile and mesh networks
- Relational databases
- Security & recovery
- Medical Telepresence

Theory:
- Theory of distributed systems
- Cryptography & Information Security
- Mechanism Design
- Quantum Information Science
- Computational Biology

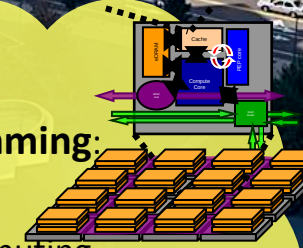

Architecture & Programming:
- Manycore architectures
- Organic or self-aware computing
- Languages for scalable computing
- Reconfigurable HW, Rapid Prototyping
- Provably Reliable Software
- Program analysis

Human/Computer Interactions:
- Spoken language systems
- Graphics, Vision, Image processing
- Natural language understanding
- Gesture-based interfaces
- Web automatio
- Crowd sourcing

© 2006 http://philip.greenspun.com

Project Angstrom:

Building 1000-Core Processor Systems

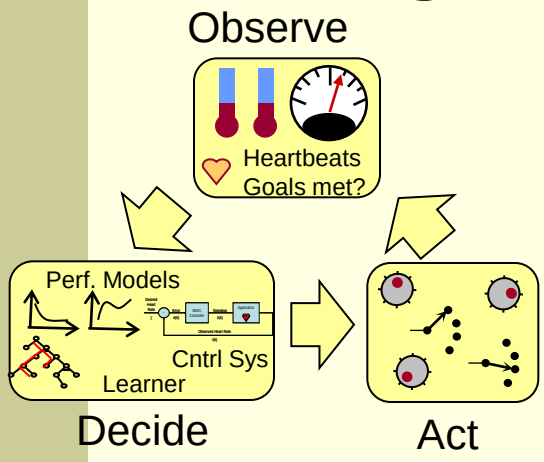

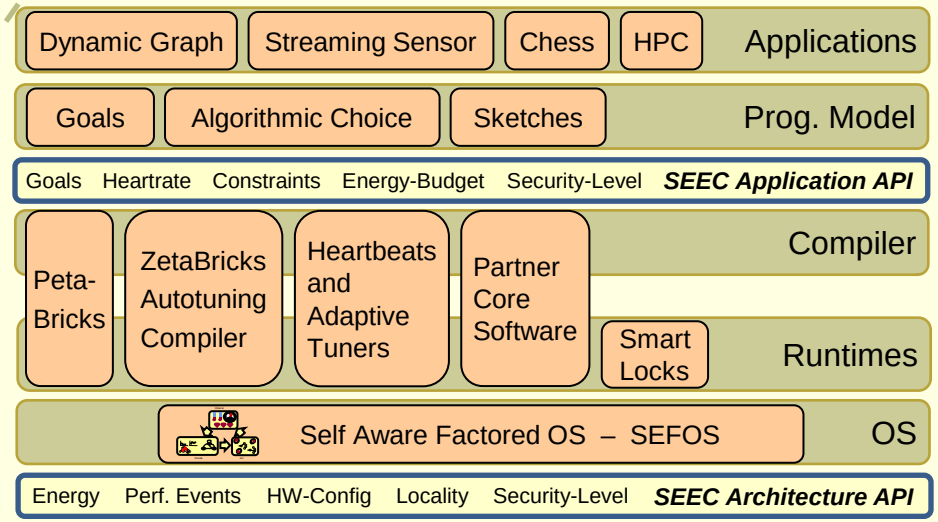

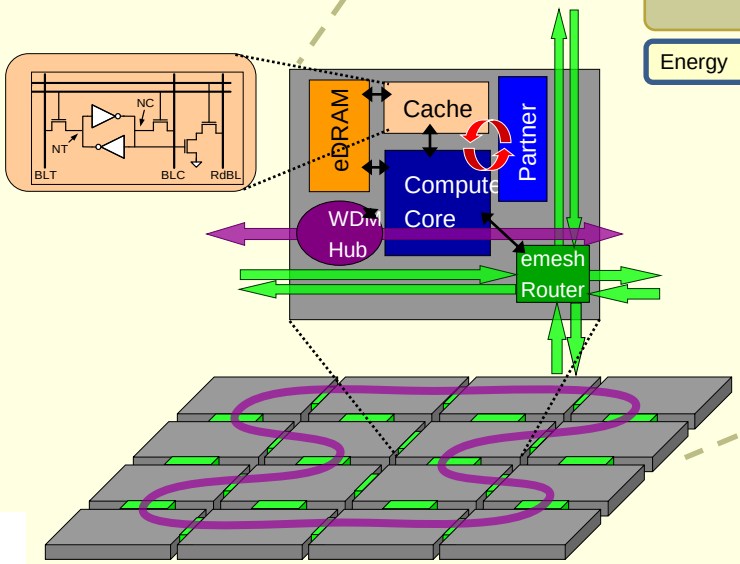

CSAIL

MTL microsystems technology laboratories massachusetts institute of technology

rle AT MIT

freescale semiconductor

MERCURY COMPUTER SYSTEMS

LOCKHEED MARTIN

UNIVERSITY OF MARYLAND 18 56

Challenges to Exascale Computing: The 3 P's

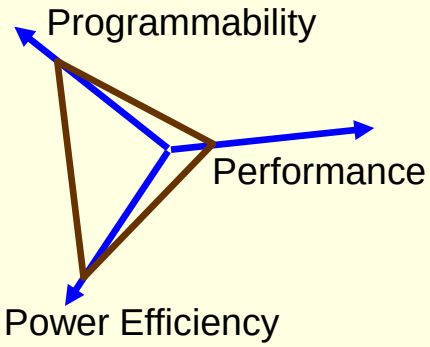

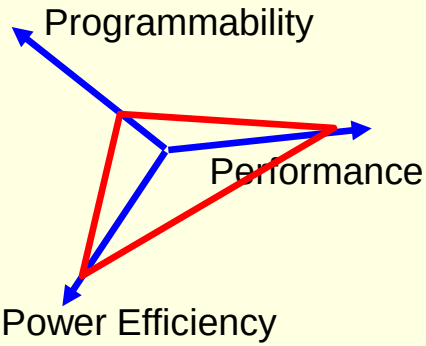

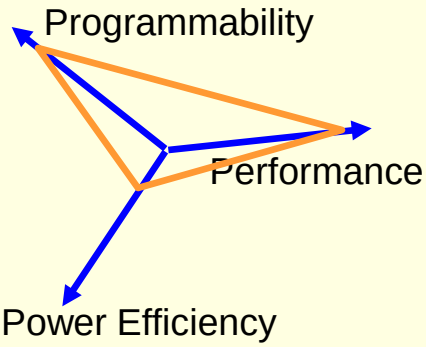

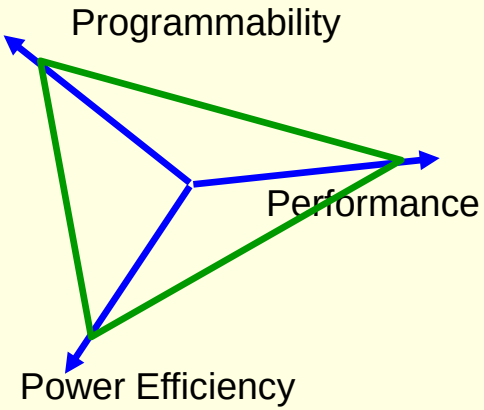

Easy to get two of three, hard to get all three

How to Get All Three

2. SEEC technology

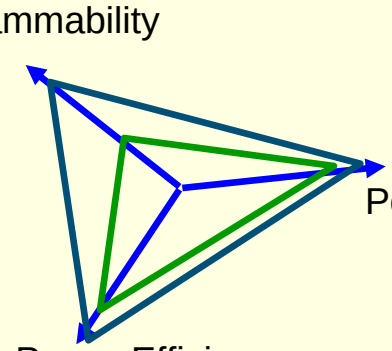

1. Fully distributed tiled architecture

1. Fully Factored Angstrom Chip Design – Yields Energy Efficiency and Scalability

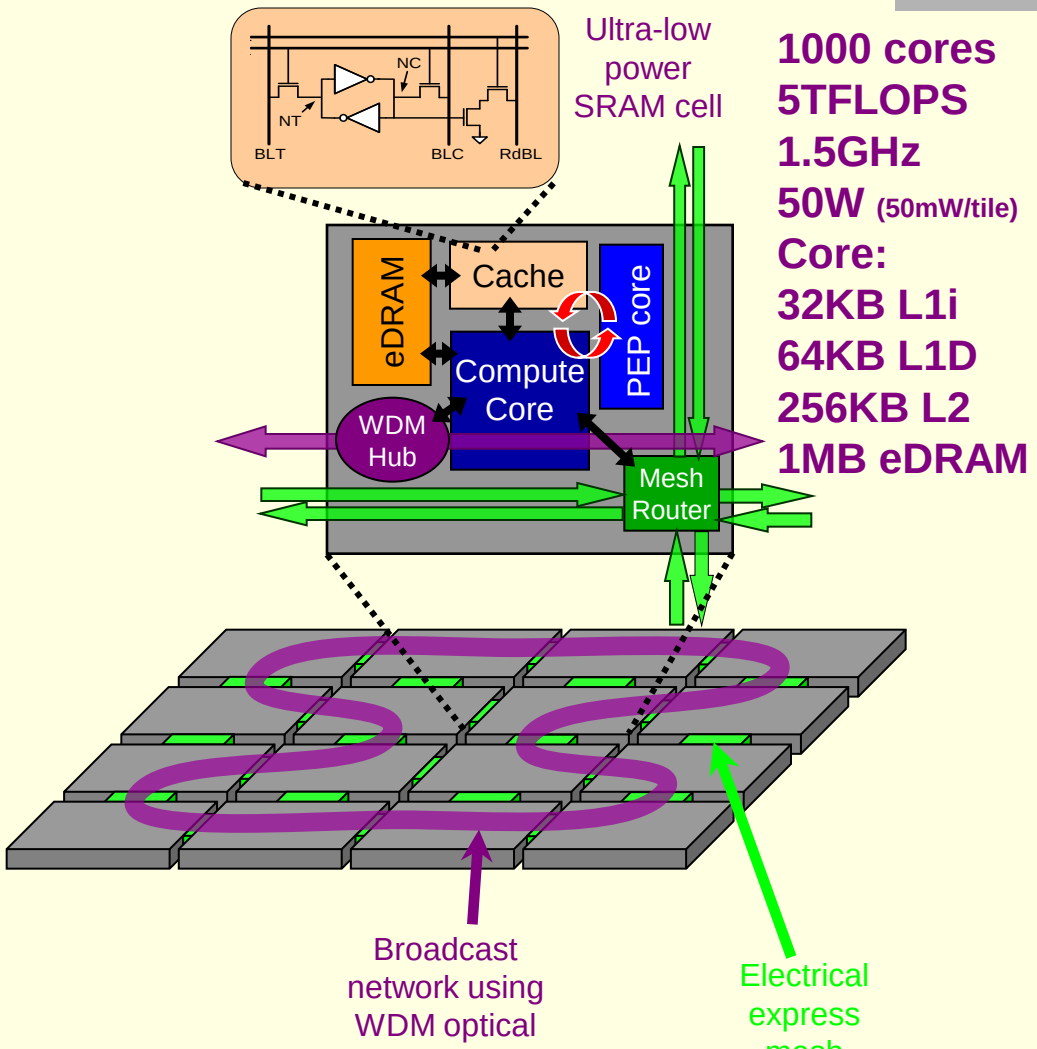

Distribute everything so things are close by
No large central structures

1. Fully Factored Angstrom Chip Design – Yields Energy Efficiency and Scalability

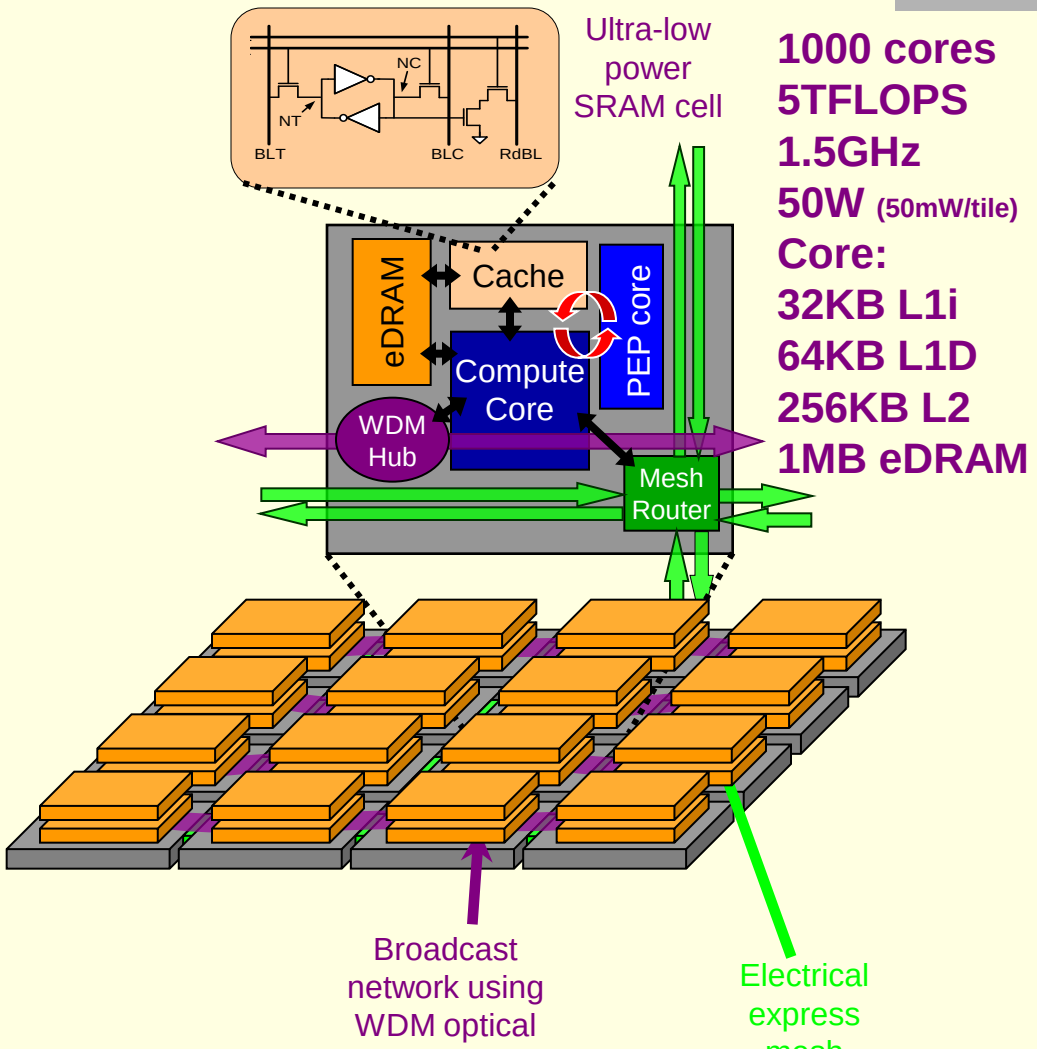

Distribute everything so things are close by
No large central structures

What Core to Use

10W/core to 50mW per core!

You don't

Start with embedded tile core
Go from 300mW to 50mW

Tiled Approach is Power Efficient and Scalable

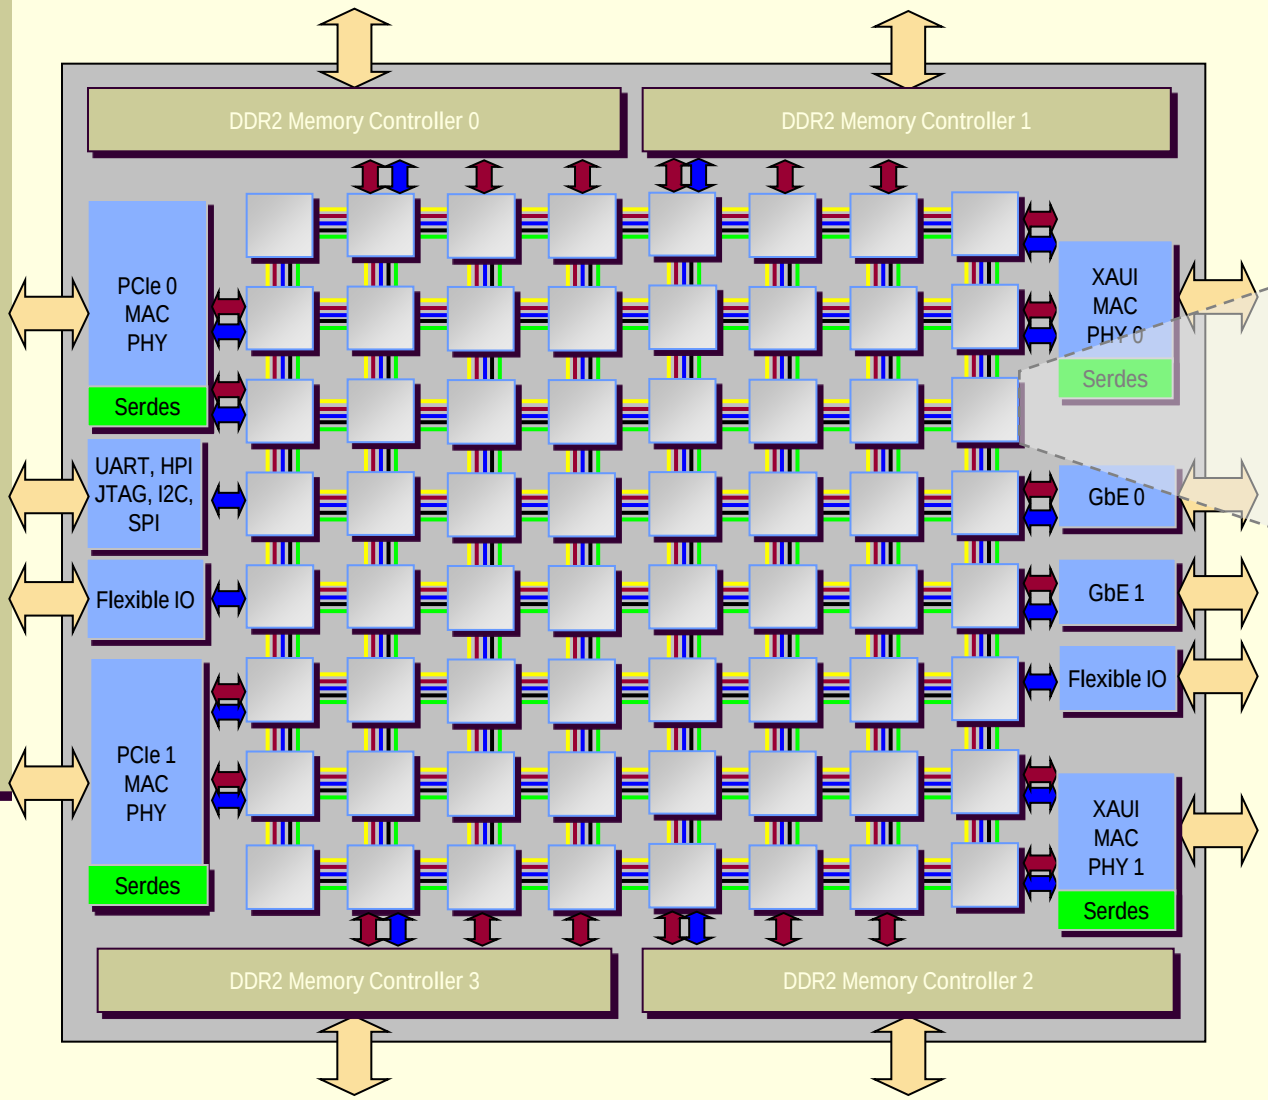

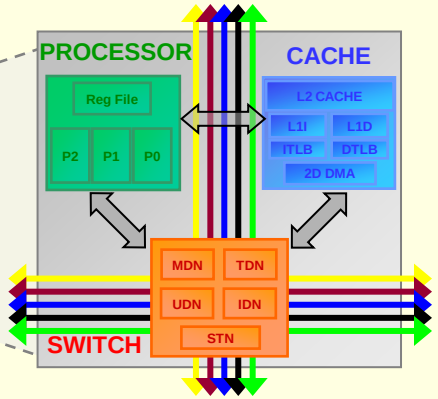

Example: TilePro64
200Gbps memory BW
40Gbps I/O
150GOPS
20W

2. SEEC: A New Computational Model

Self-Aware Execution (SEEC) – a computing paradigm in which systems observe their runtime behavior, learn, and take actions to meet desired goals

- User indicates performance or energy goals and provides alternatives of how to do things
- System hardware and software manage everything else (e.g., locality, resilience), meeting goals by adapting to changing conditions

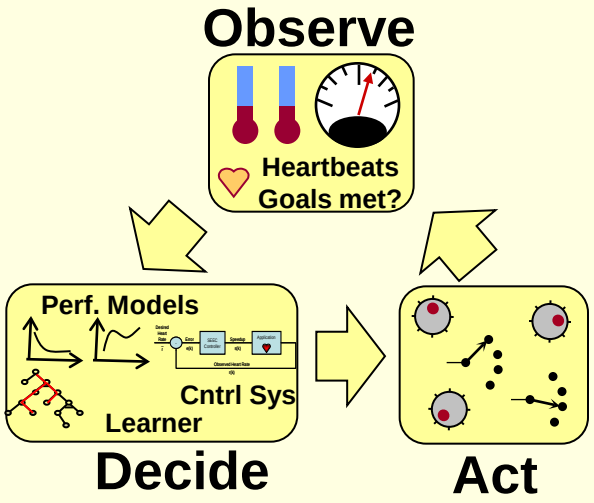

2. SEEC – Self Aware Computational Model

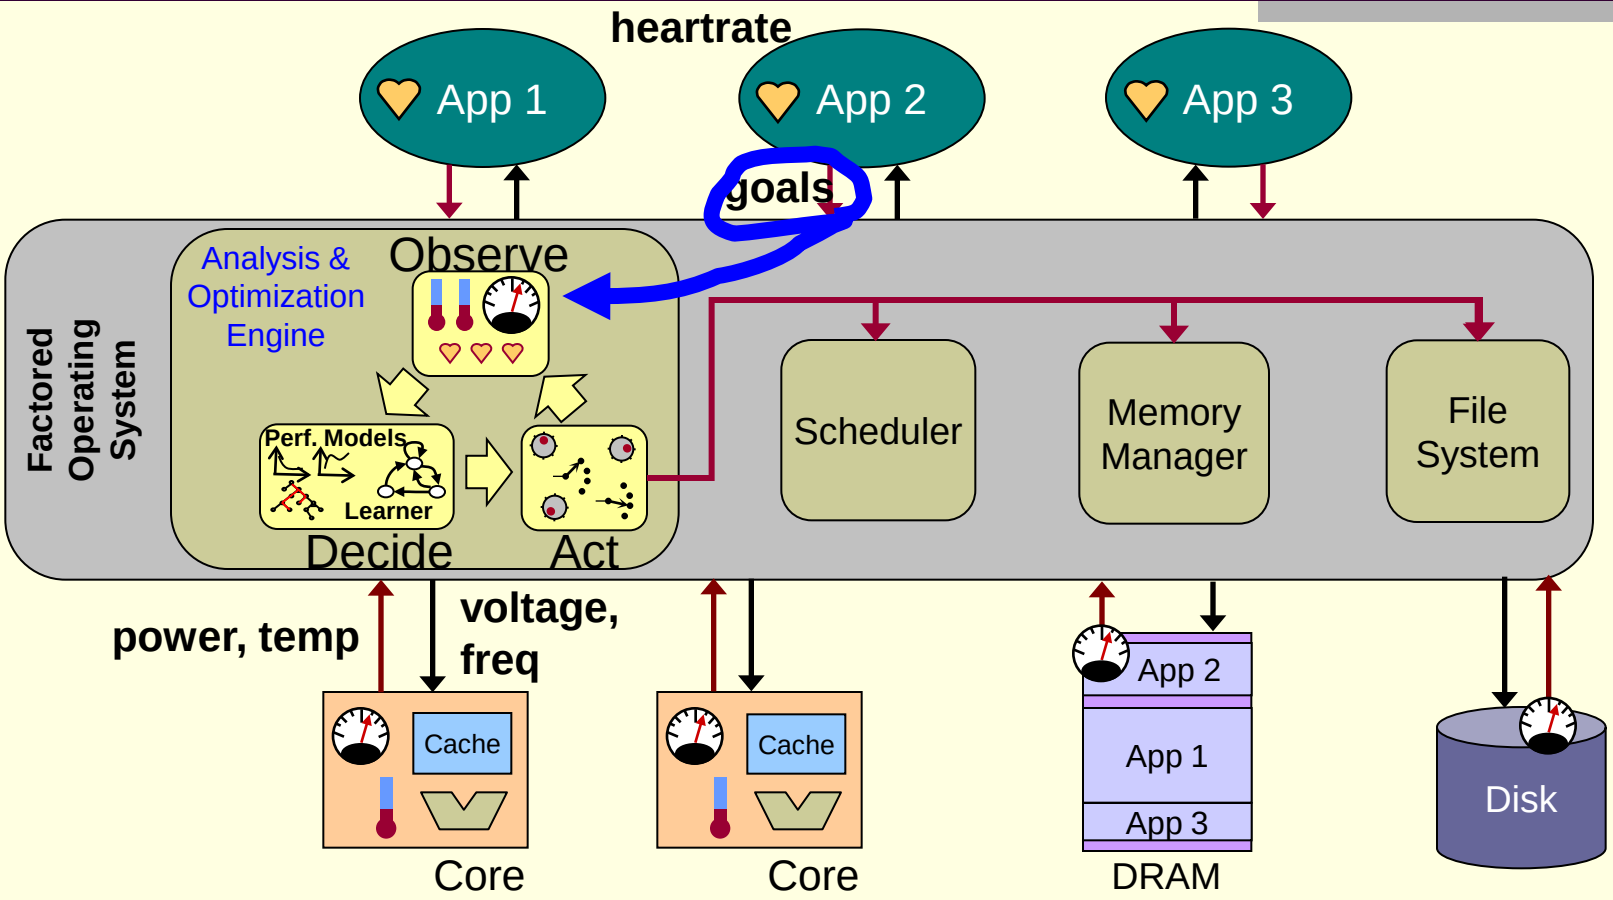

Computers should become more like humans

www.youtube.com/user/HeartbeatsAPI

2. SEEC – Self Aware Computational Model

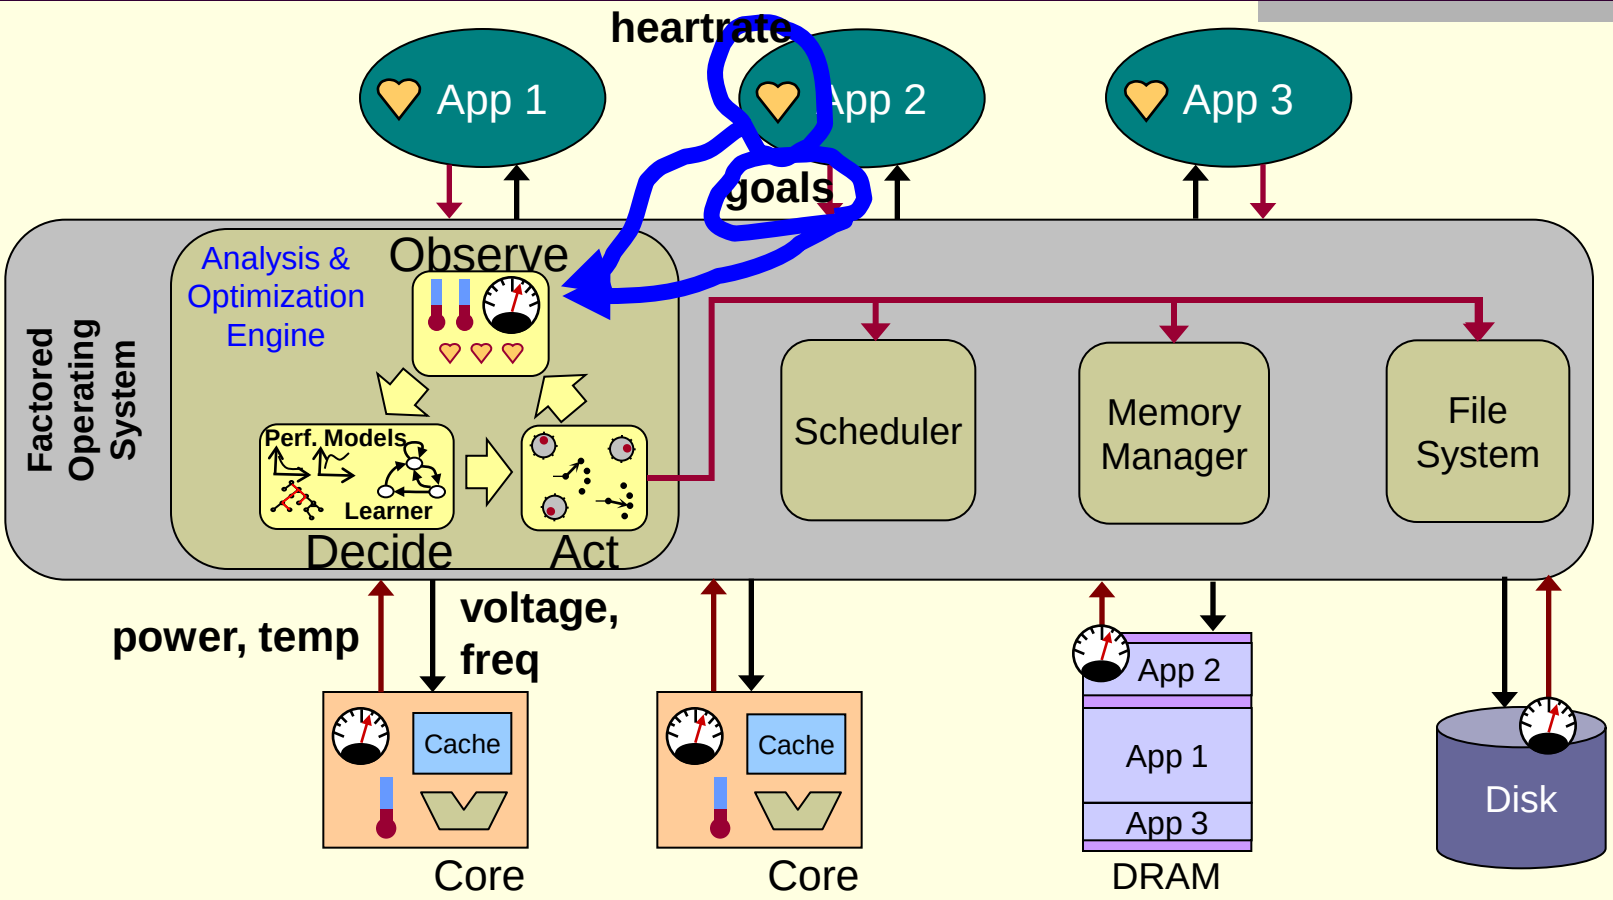

Computers should become more like humans

www.youtube.com/user/HeartbeatsAPI

2. SEEC – Self Aware Computational Model

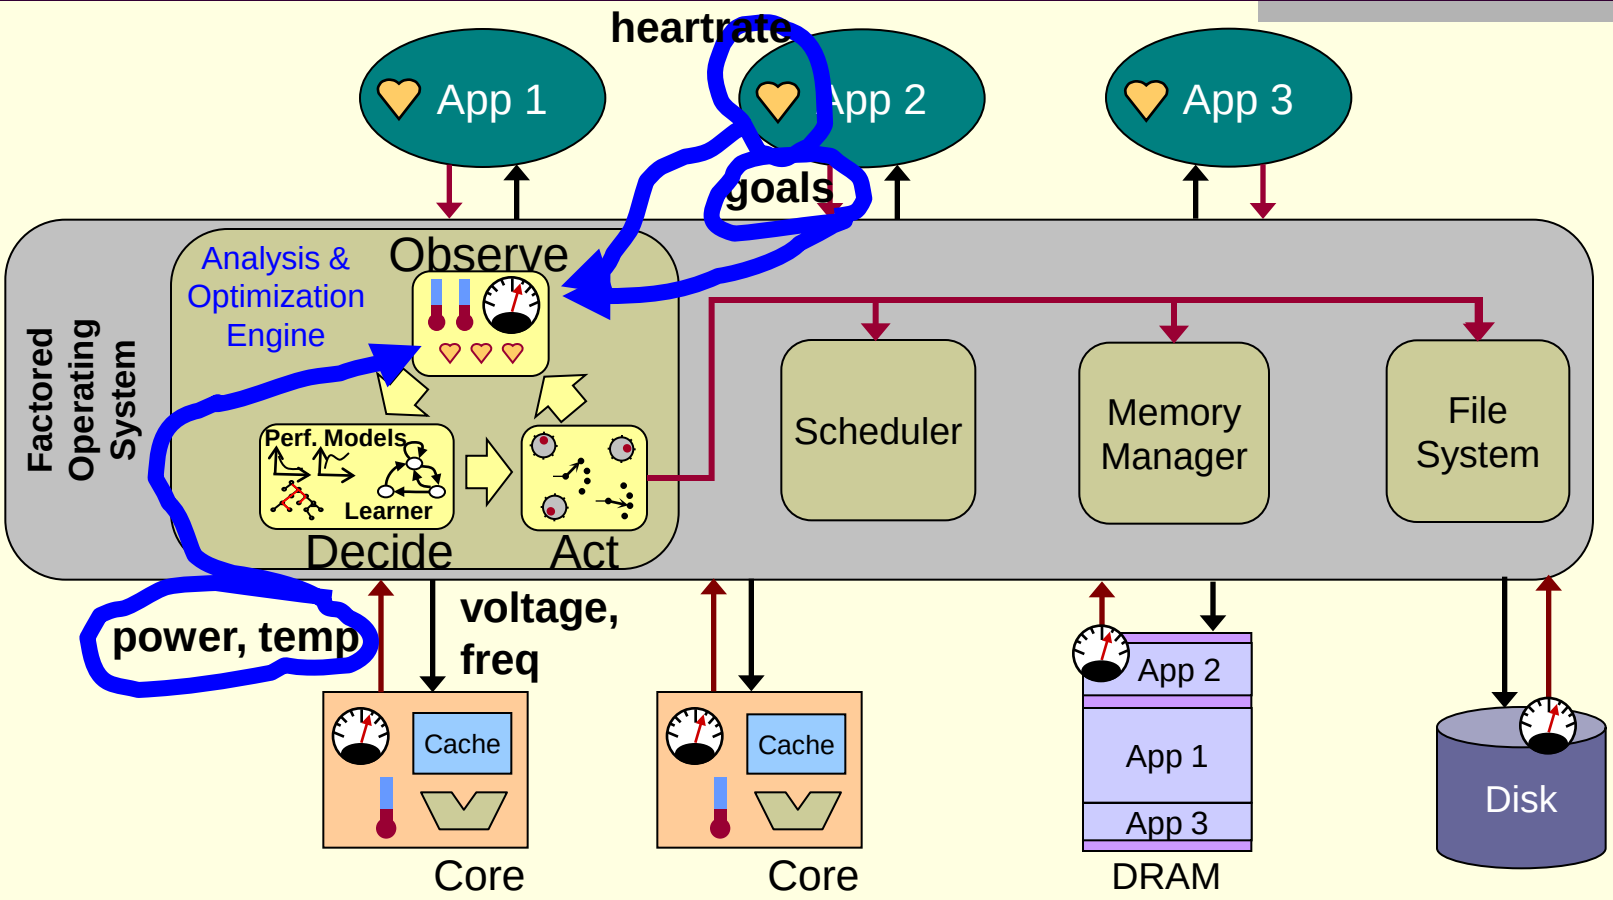

Computers should become more like humans

www.youtube.com/user/HeartbeatsAPI

2. SEEC – Self-Aware Computational Model

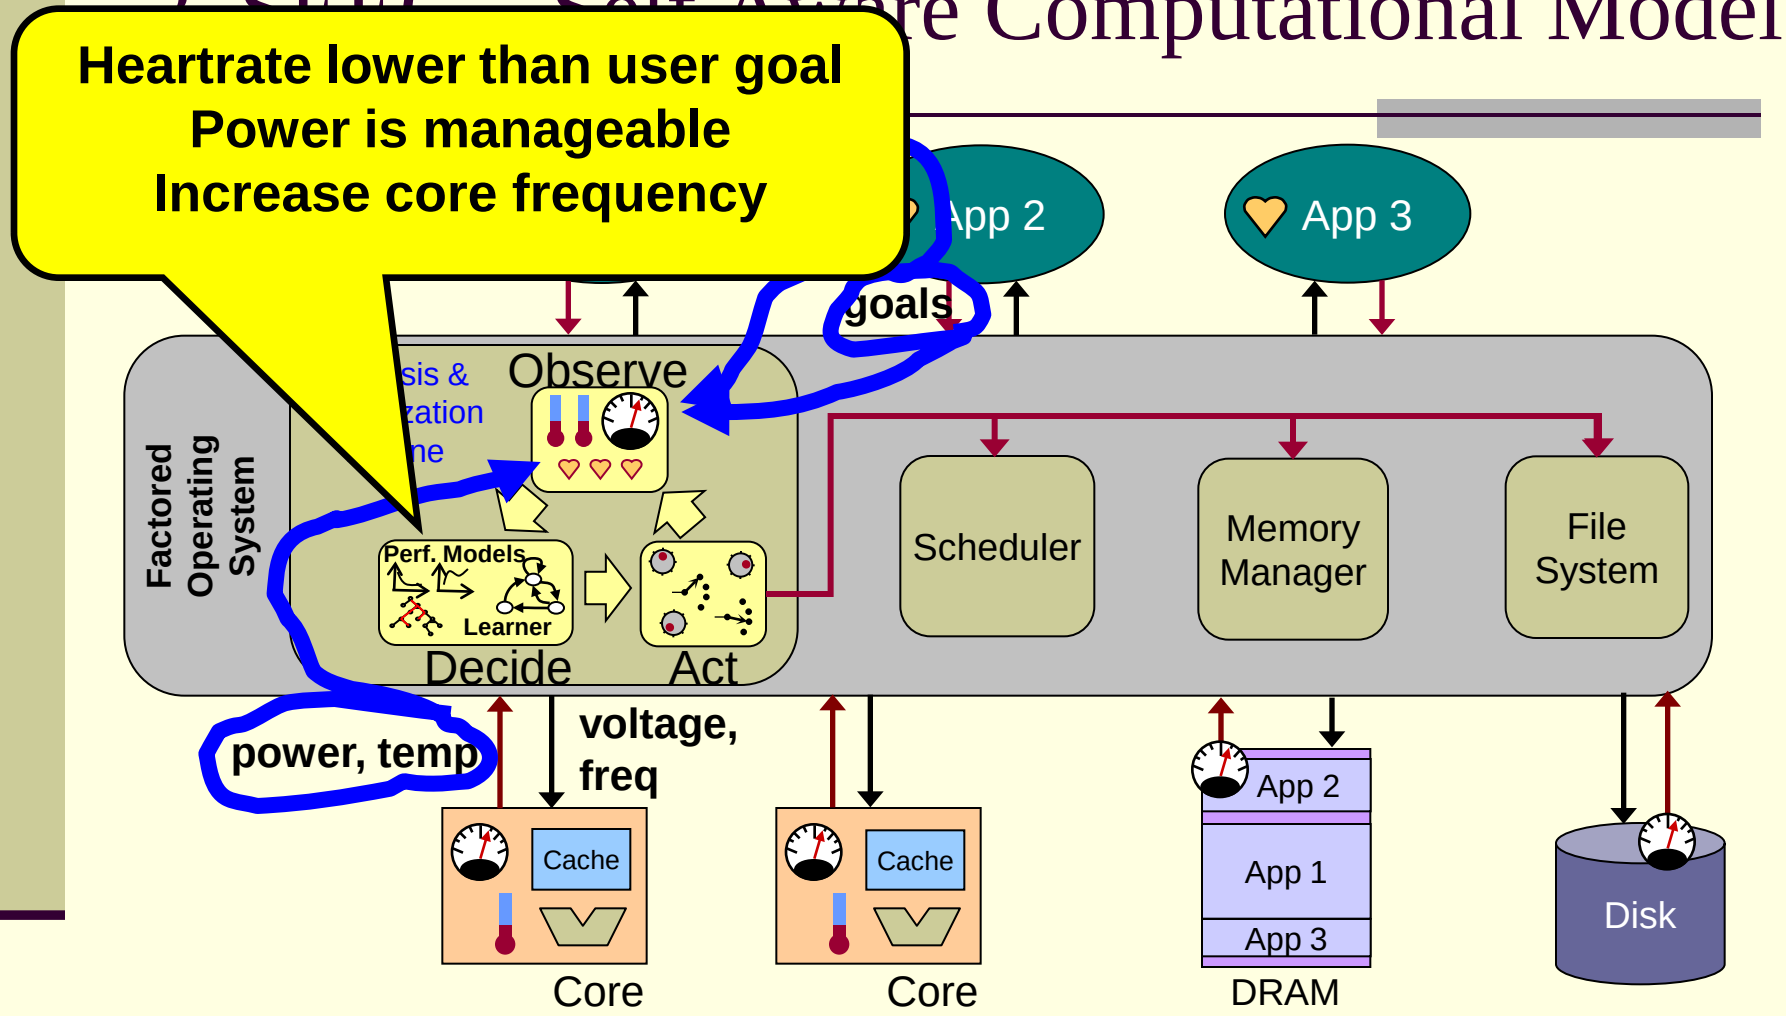

Computers should become more like humans

www.youtube.com/user/HeartbeatsAPI

2. SEEC – Self-Aware Computational Model

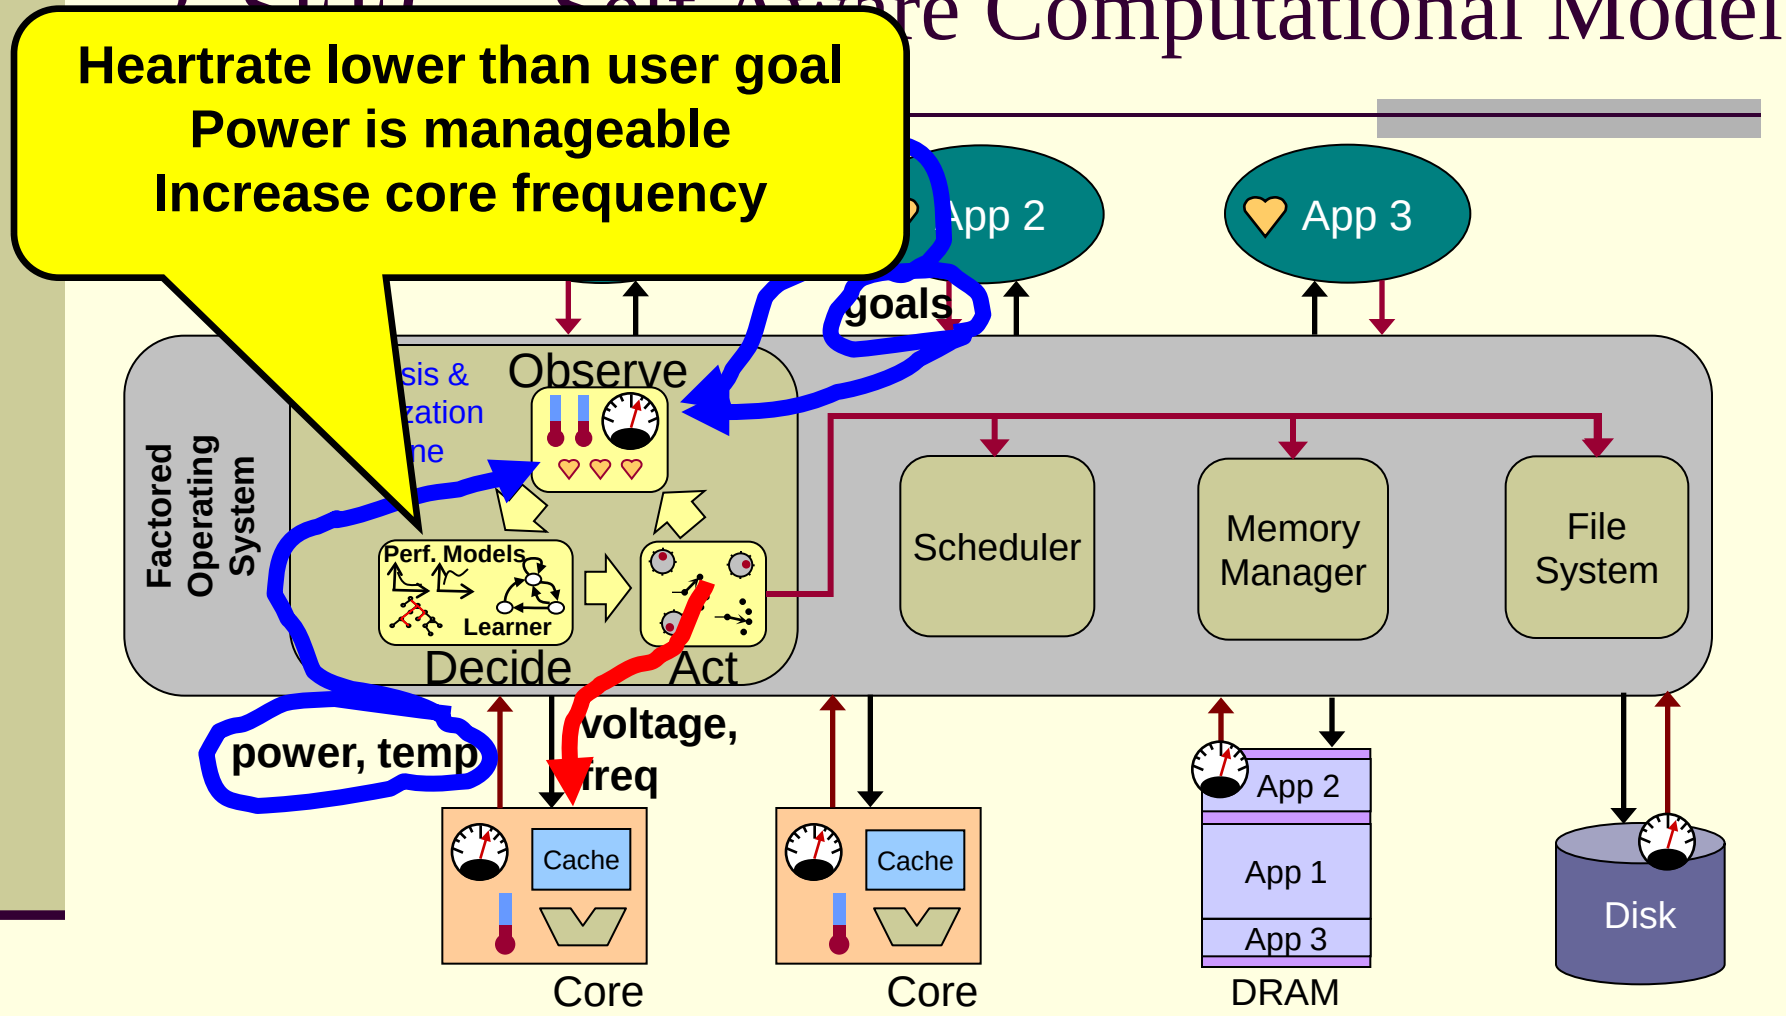

Computers should become more like humans

www.youtube.com/user/HeartbeatsAPI

2. SEEC – Self Aware Computational Model

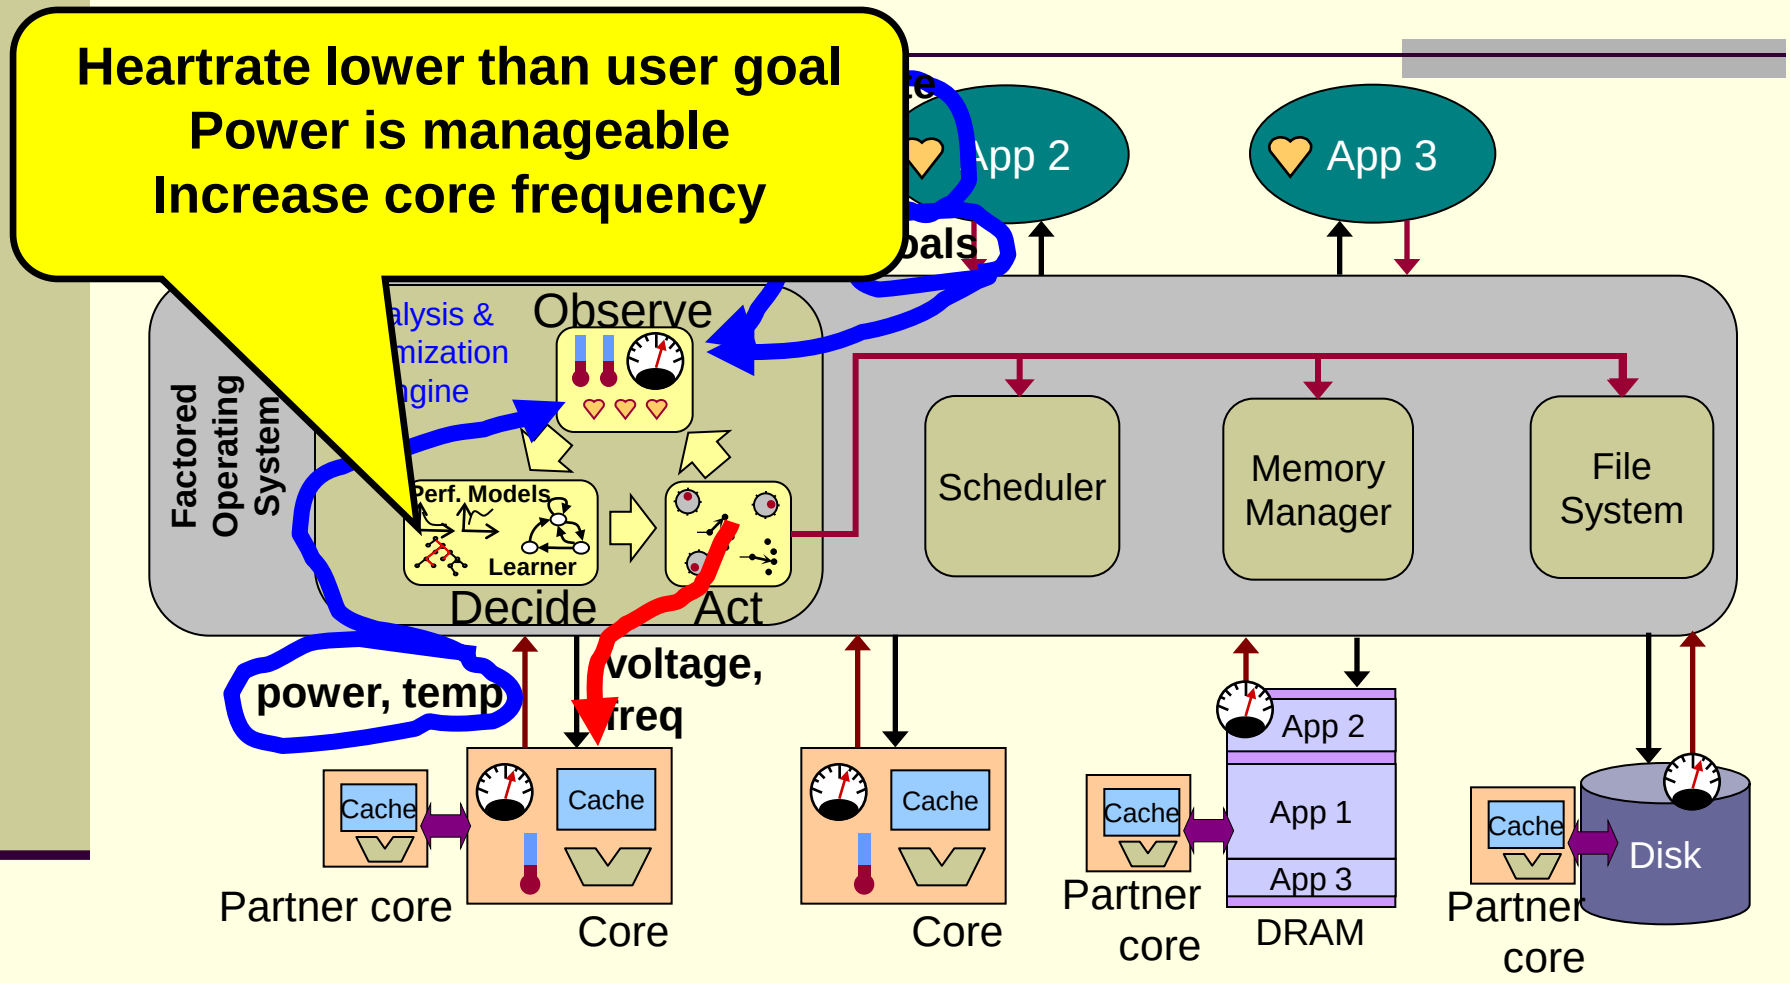

Computers should become more like humans

www.youtube.com/user/HeartbeatsAPI

2. SEEC – Self Aware Computational Model

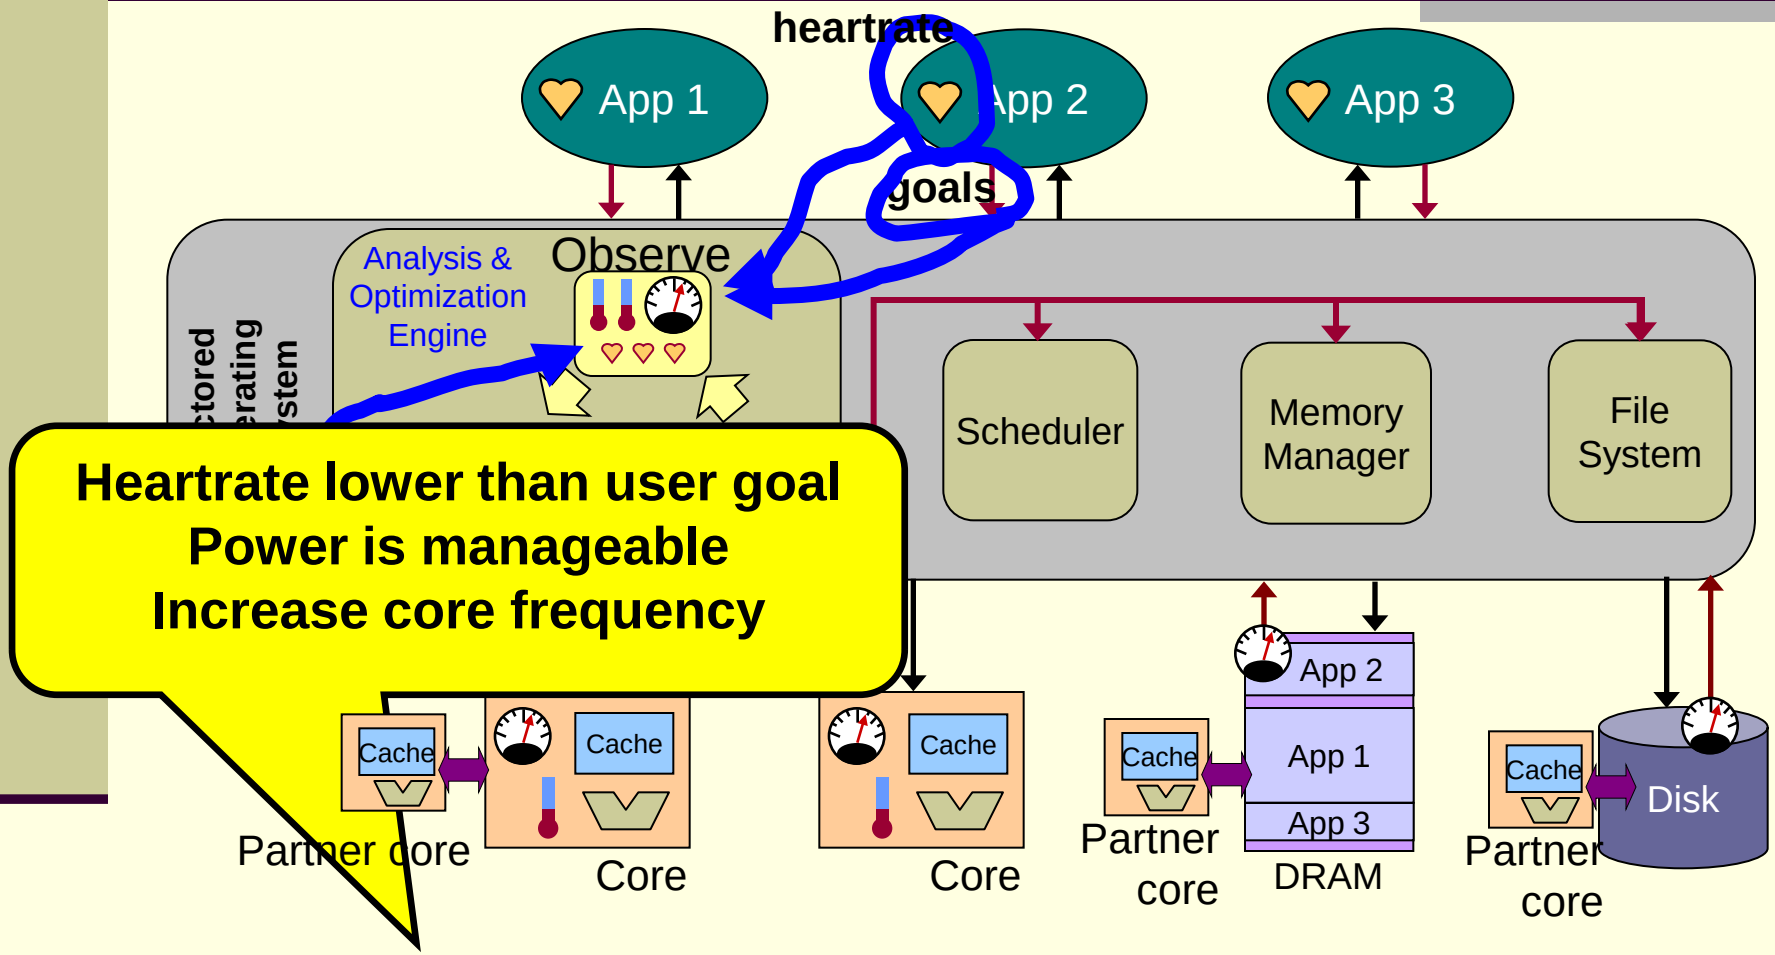

Computers should become more like humans

www.youtube.com/user/HeartbeatsAPI

Why SEEC?
Programming is Becoming very Hard

mini-SAR frontend

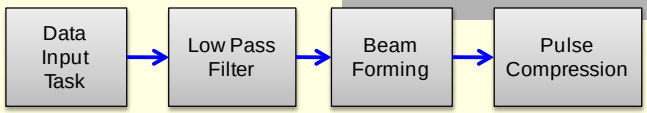

- Mini-SAR app has many configurations in existing or future machines
  - # threads/stage
  - Thread mapping → cores
  - Core frequencies and voltage
  - Memory controller mapping
  - Layout of threads and cores
  - Cache management
  - What if something breaks – I lose a core!
- Measured range of 0.07 - 0.7 pulses/sec/watt for various configurations (10X range) – 10X in energy efficiency on the table!
  - Programmer can easily make bad choices in configuration
  - Best configuration can also change with input
  - Communication and cache variability

SEEC technology achieved 90% of optimal with minimal programmer effort

Roles in the SEEC Model

	Application Developer	Systems Developer	SEEC System Infrastructure
	Express application goals and progress (e.g. frames/ second)		Read goals and performance
			Determine how to adapt (e.g. How to speed up the application)
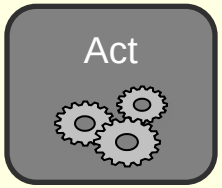		Provide a set of actions and a callback function (e.g. allocation of cores to process)	Initiate actions based on results of decision phase

The decision engine is key to enabling SEEC

23

ODA Control Loop

A control theory for Sefos

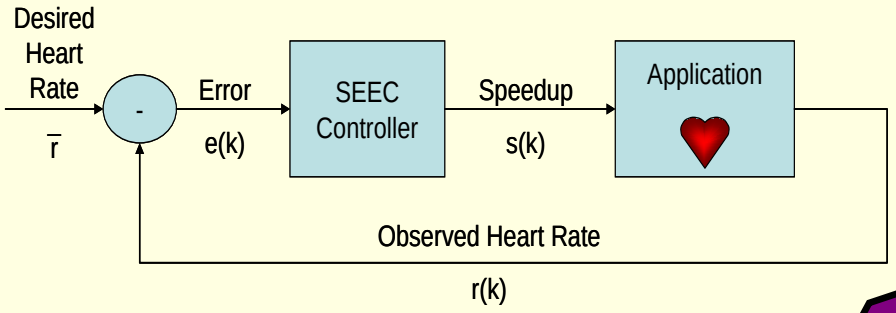

User can dial in desired behavior

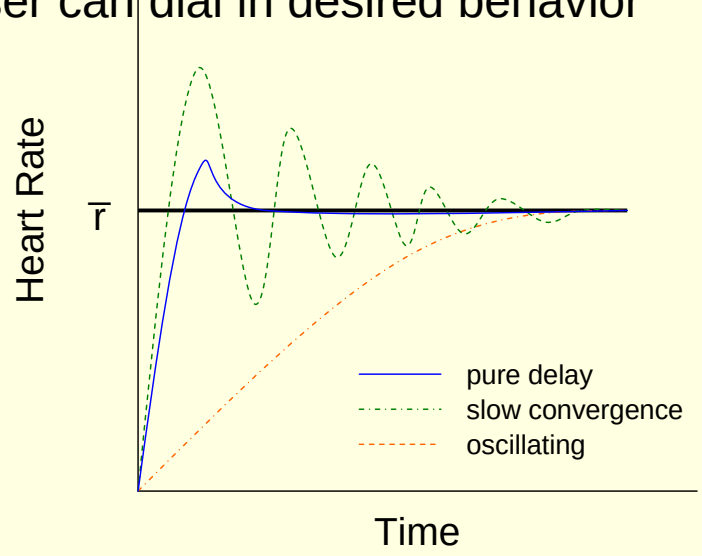

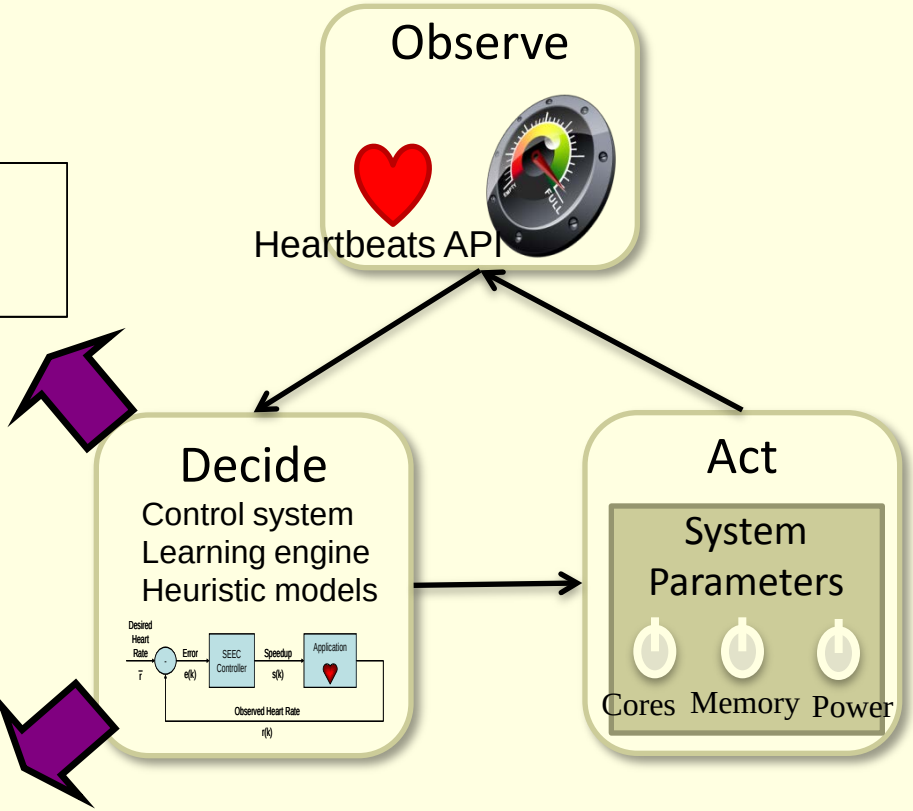

H.264 Video Encode: Procedural

25

H.264 Video Encode: Self-Aware using Heartbeats plus Heuristic Approach

26

Minimizing Power in a Self-Aware System

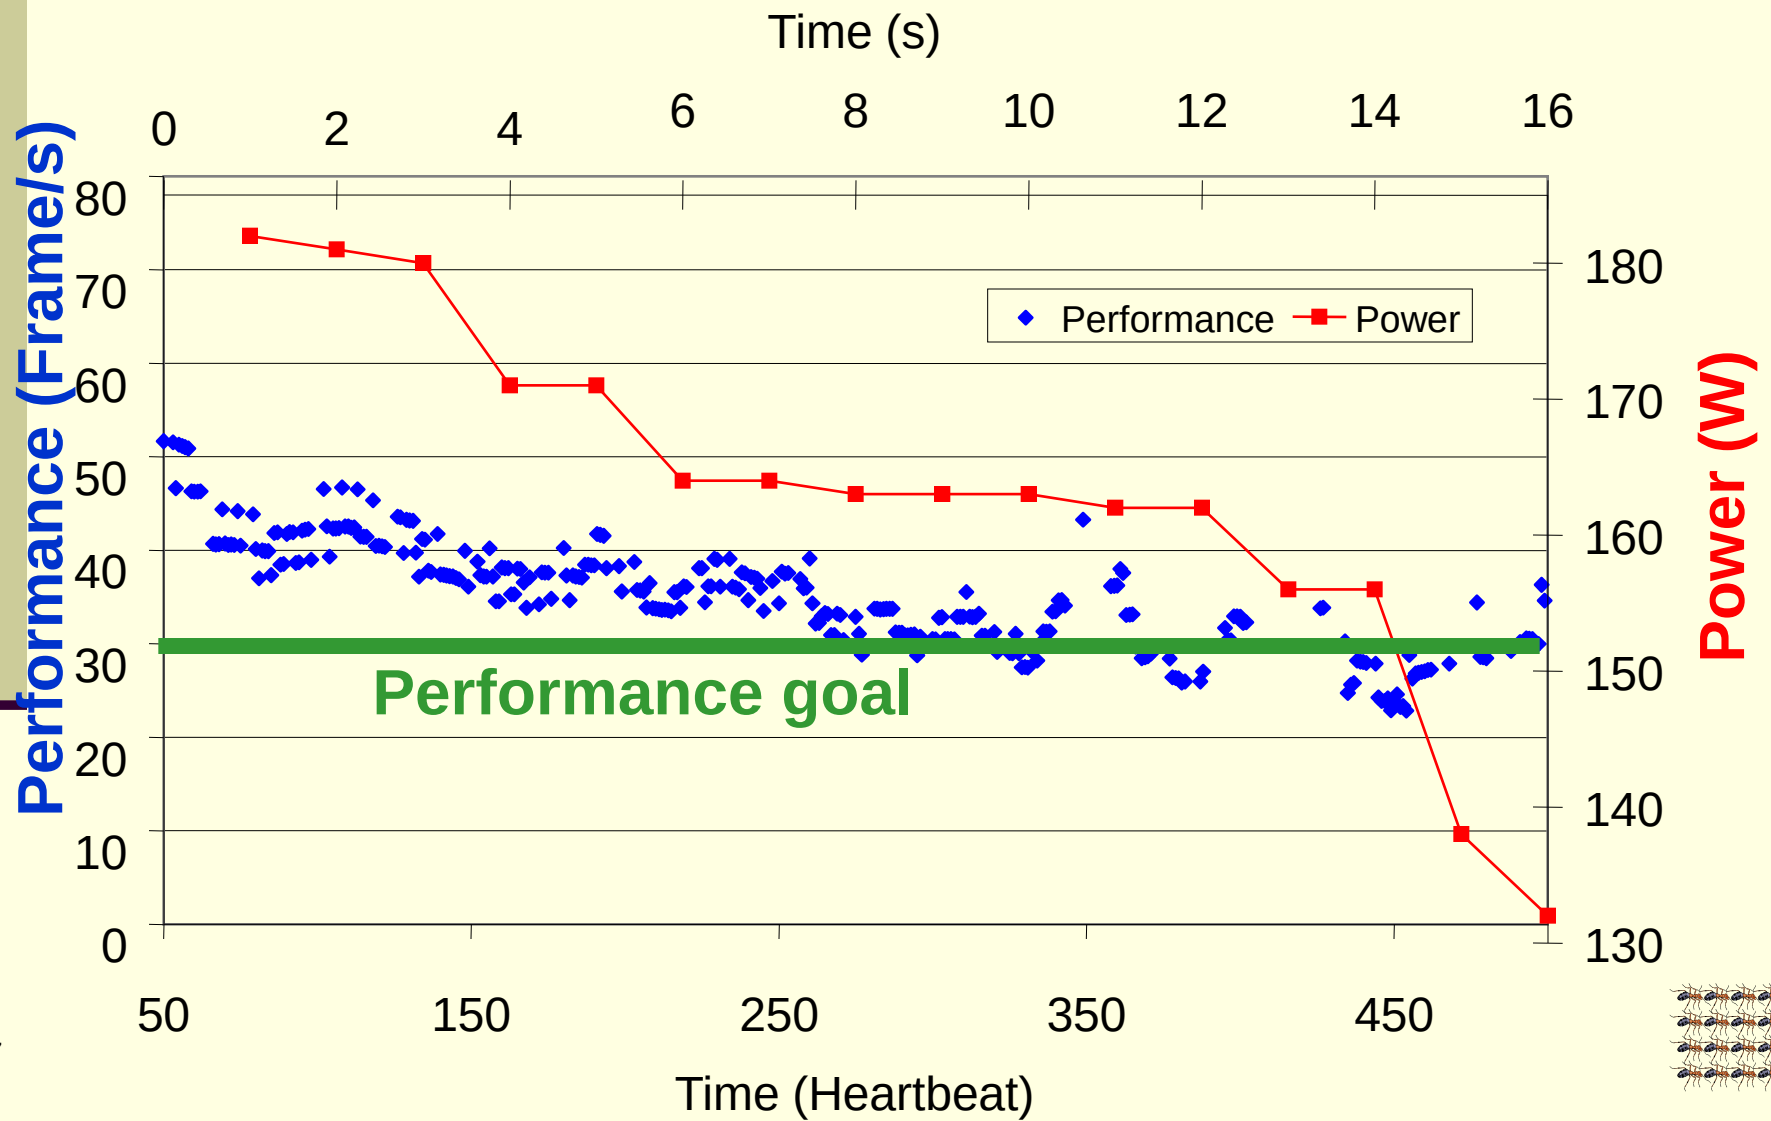

Minimizing Power in a Self-Aware System

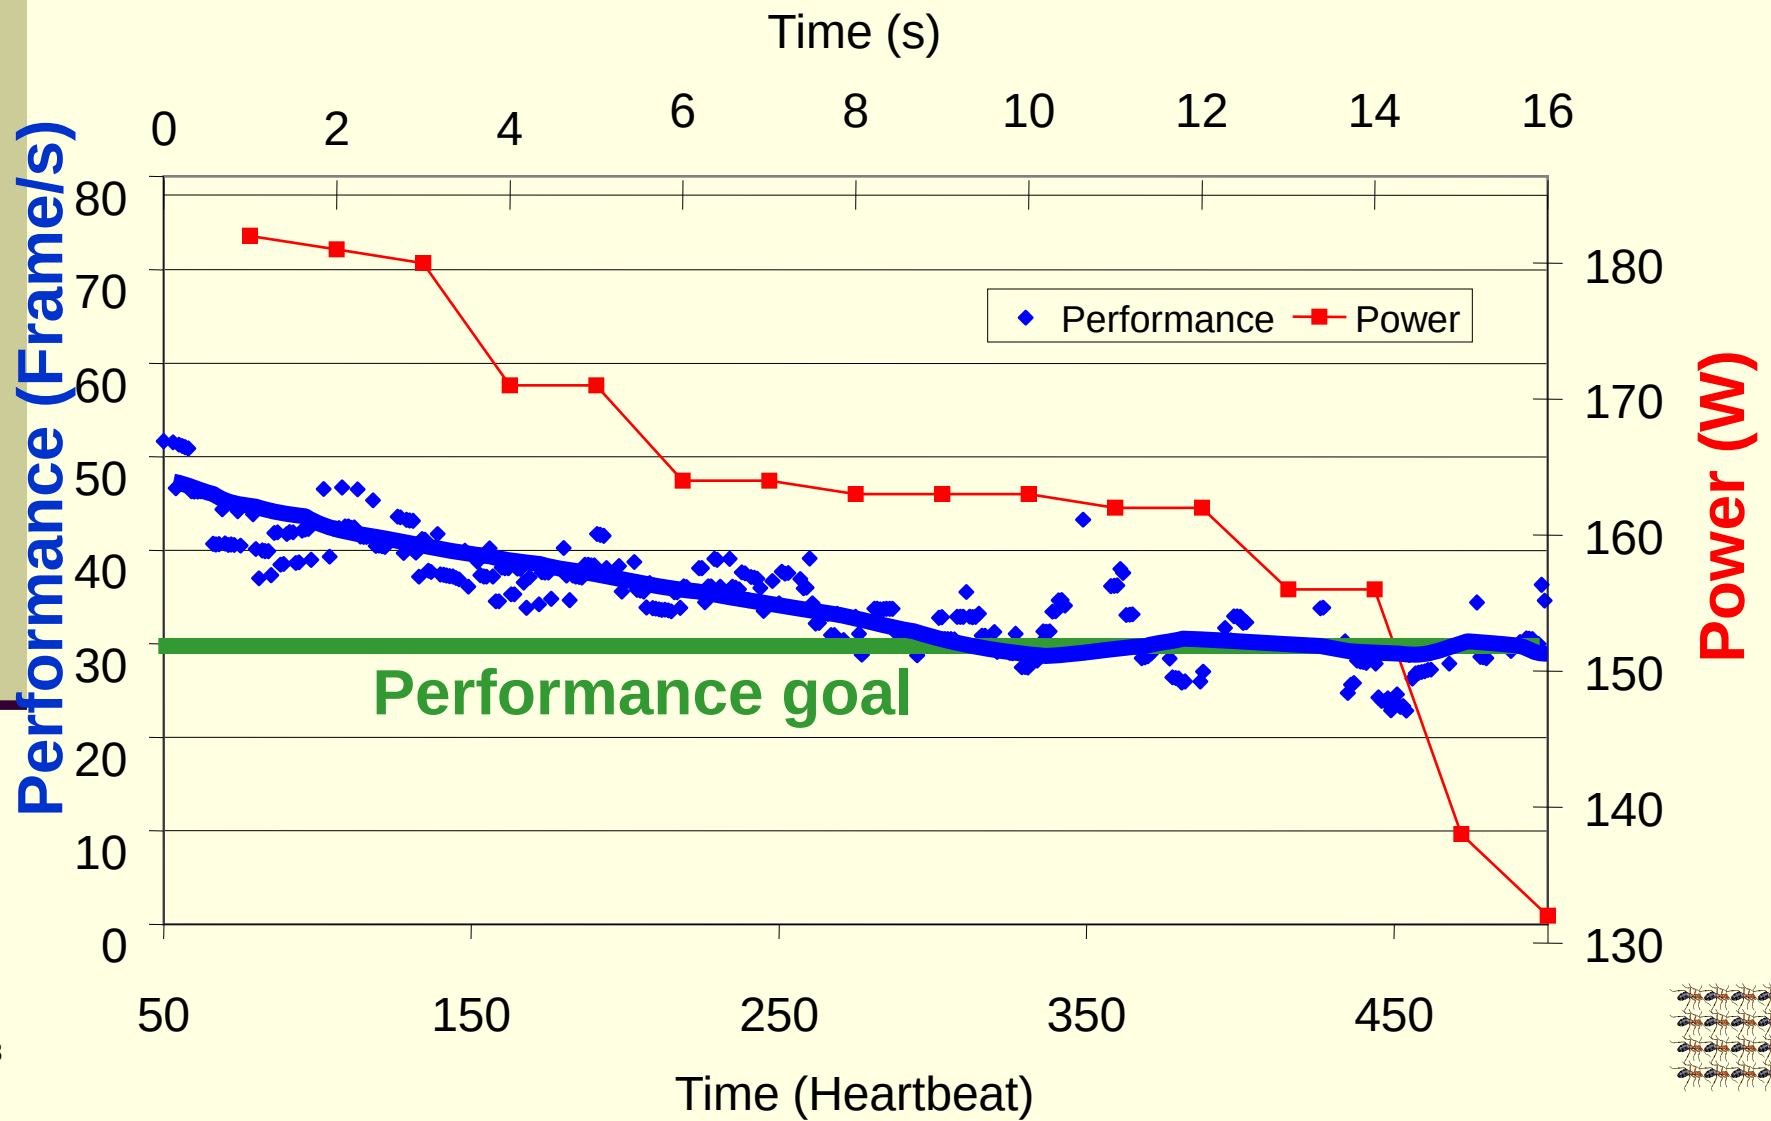

Minimizing Power in a Self-Aware System

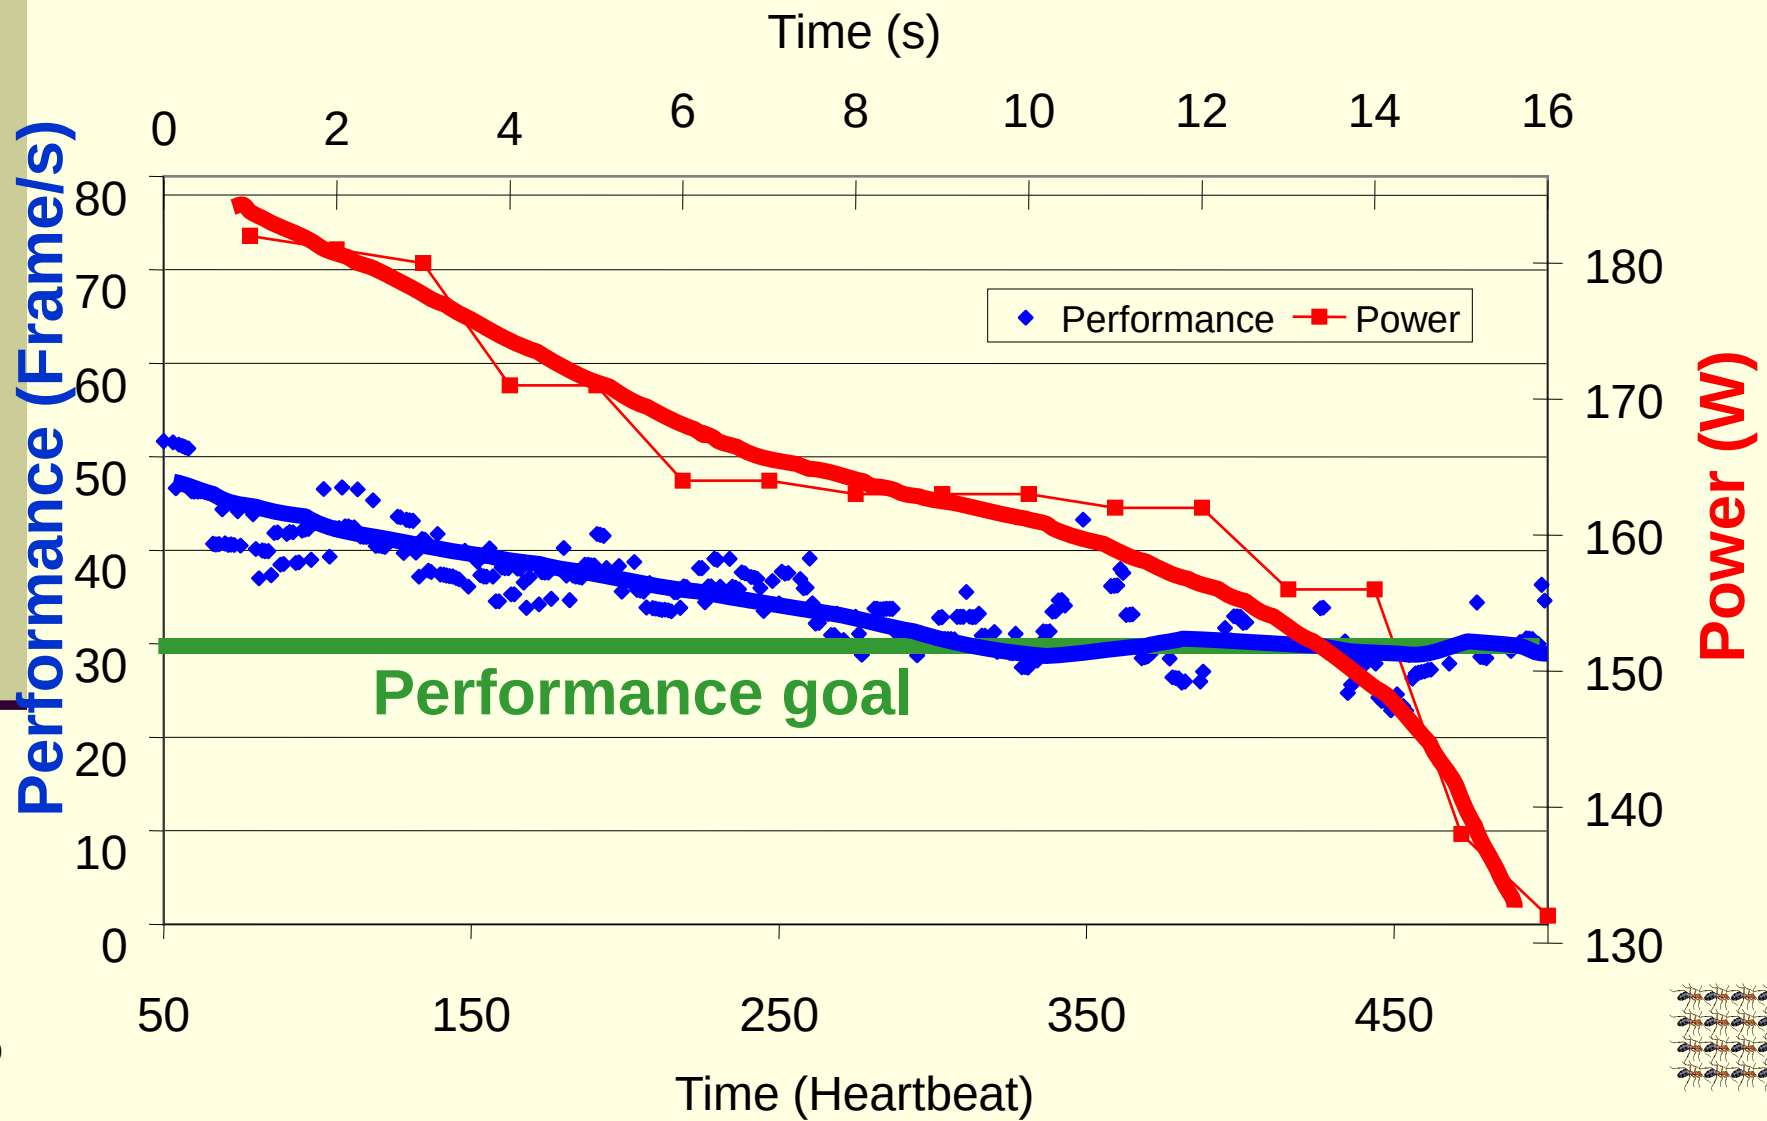

Multiple Applications

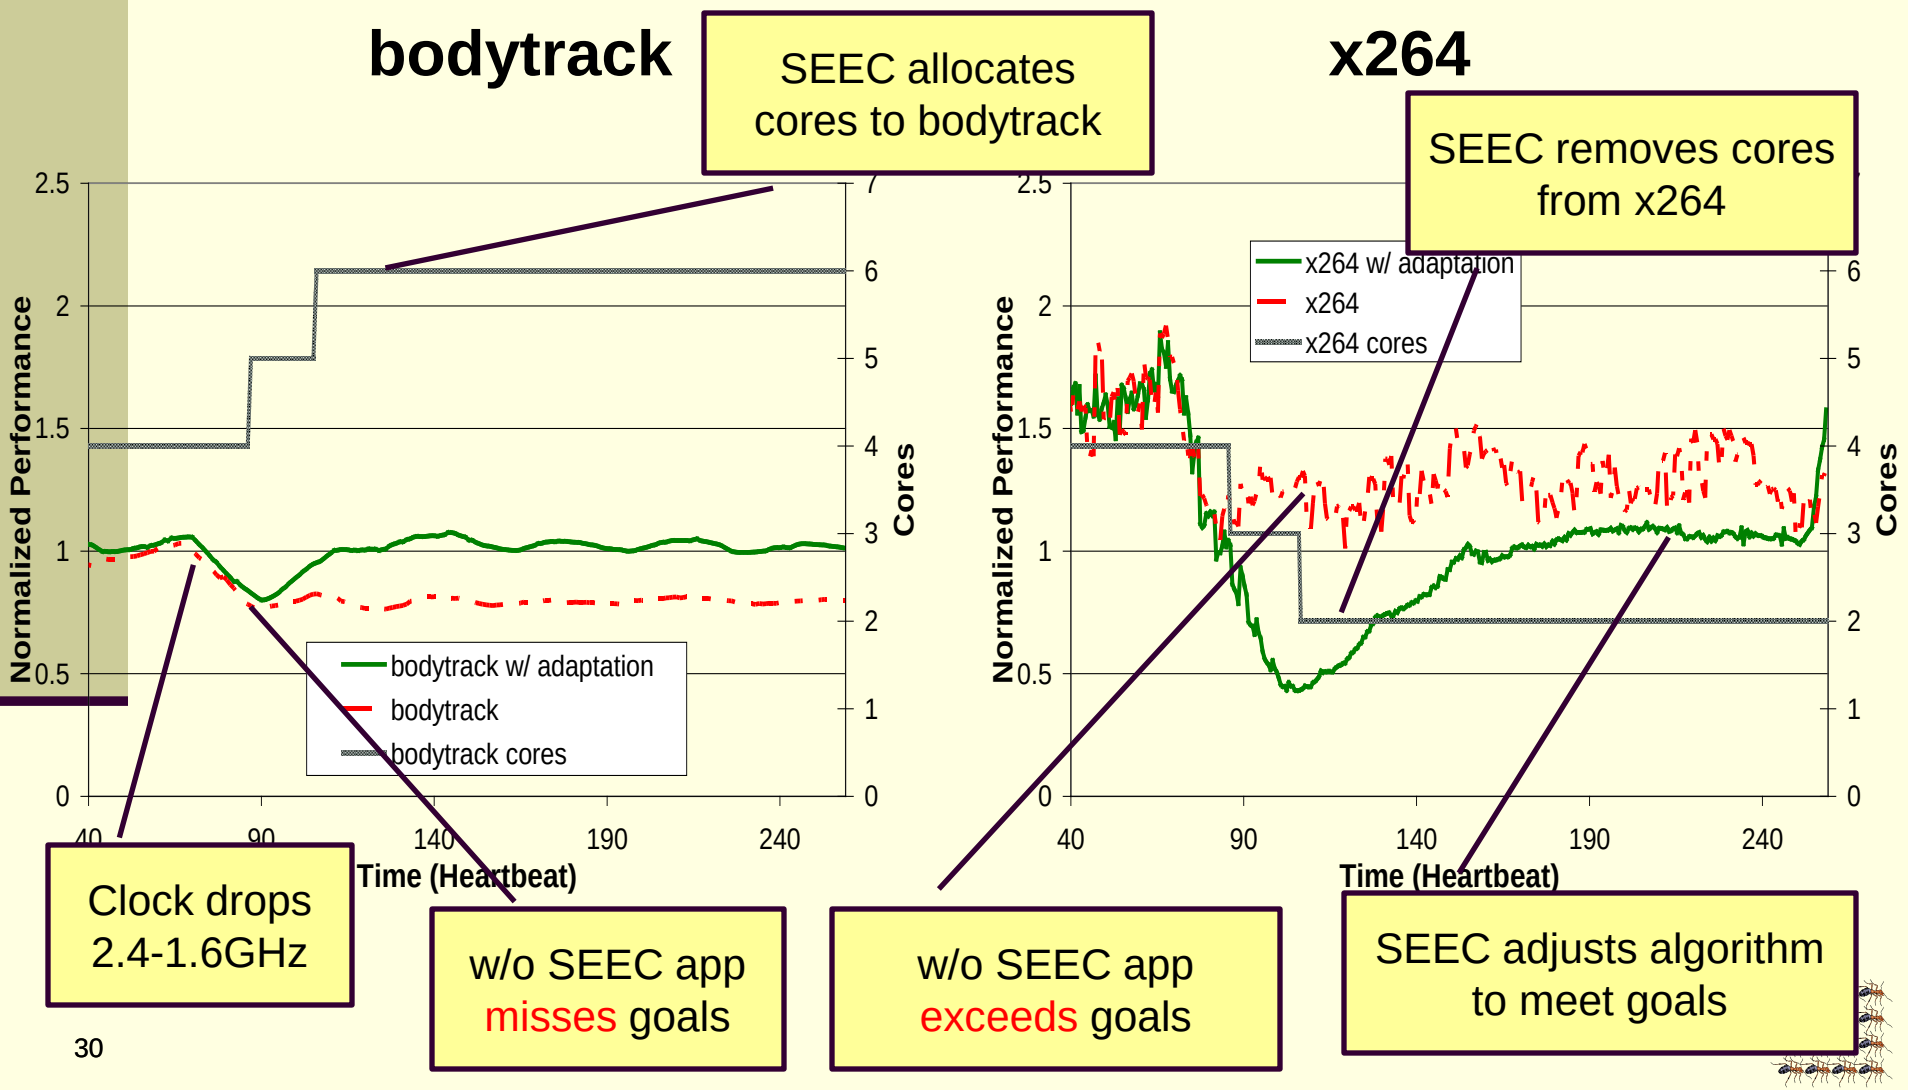

Decision Making Strategies

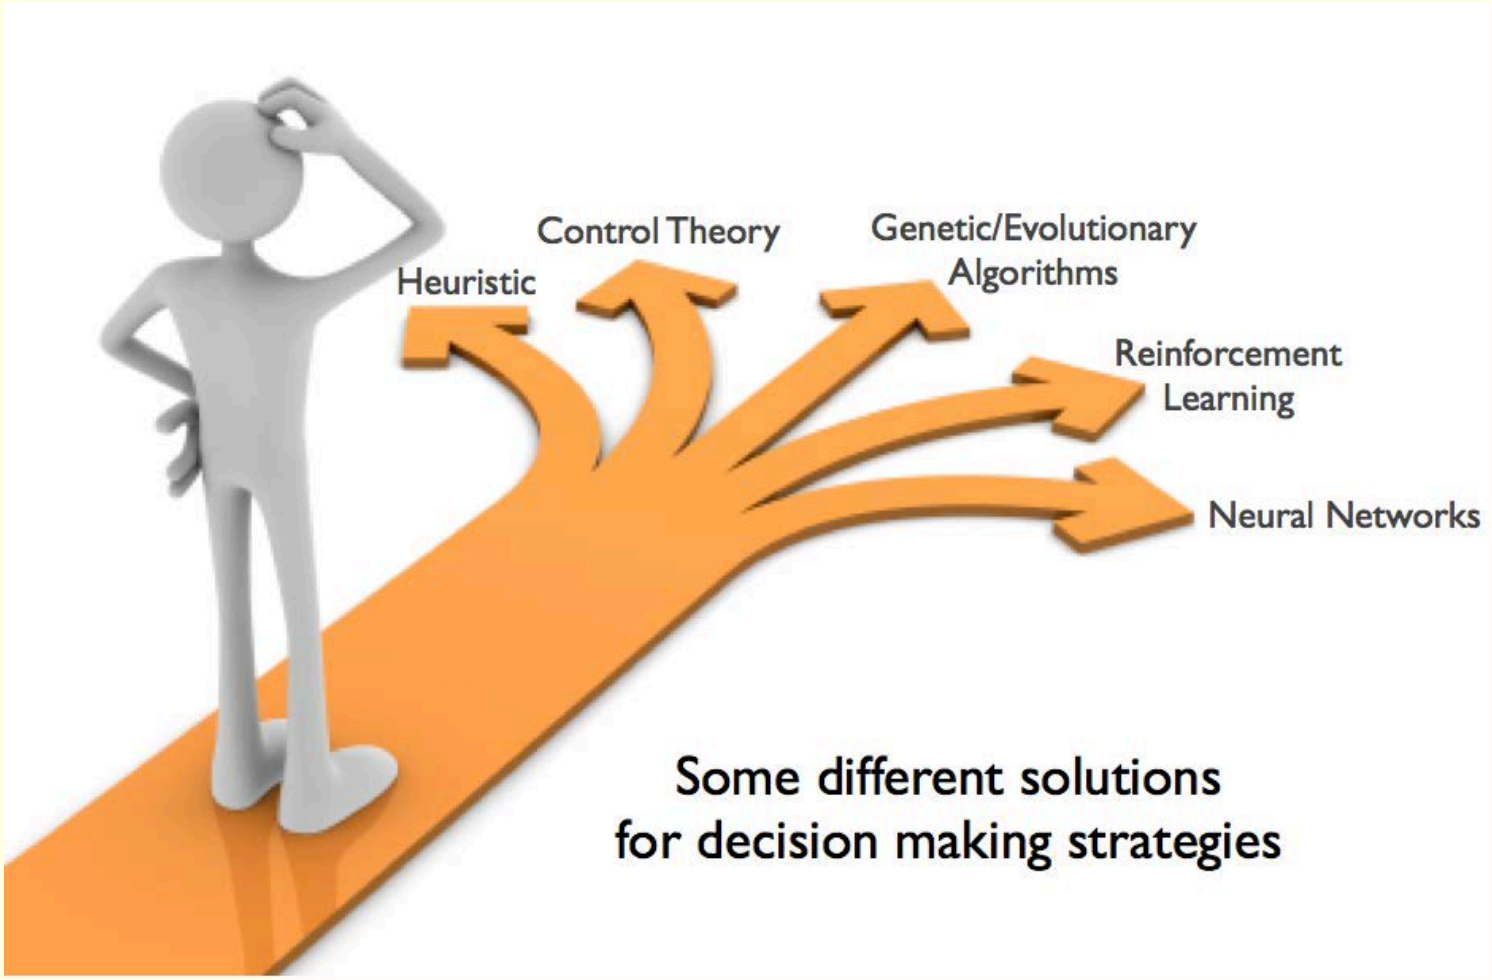

Decision Making Strategies

Comparison of Different Approaches

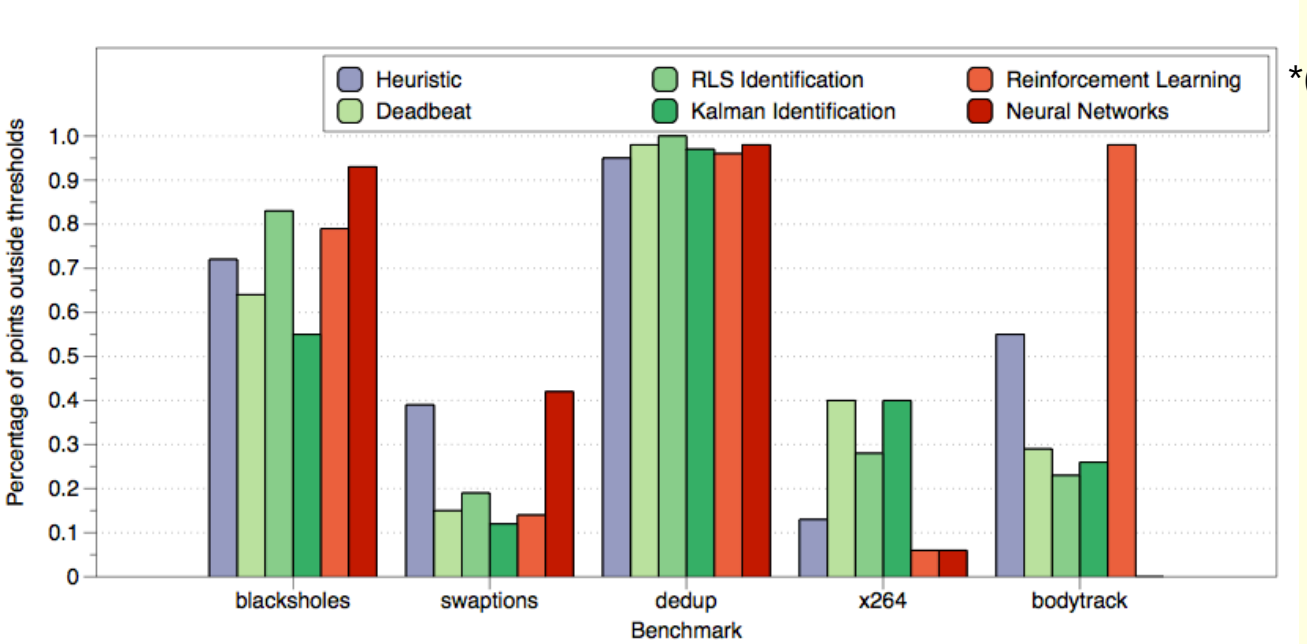

*(lower is better)

• WDP: measures the percentage of data points that are not in the desired performance interval

Summary

- Angstrom project is approaching the computing problem with two key ideas
  - Create a fully distributed architecture
  - Create a fundamentally new computational model – SEEC
- Angstrom approach has the potential to solve the power efficiency, performance and programmability challenges
- SEEC approach is showing promise as a new way of building computer ystems