

## Cool'Em All

Platform for optimising the design and operation of modular configurable IT infrastructures and facilities with resource efficient cooling

Presented by Jean-Marc PIERSON  
IRIT / Toulouse University

Coordinated by Ariel Oleksiak,  
Poznan Supercomputing and Networking Center (PSNC)



## CoolEmAll Project

## CoolEmAll Project




- FP7 ICT Call 7
- Budget: 3614210€ (funded: 2645000€)
- Duration: 30 months
- Start date: 1<sup>st</sup> Oct 2011

# CoolEmAll Project



- FP7 ICT Call 7
- Budget: 3614210€ (funded: 2645000€)
- Duration: 30 months
- Start date: 1<sup>st</sup> Oct 2011

CoolEmAll consortium



- Consortium



## CoolEmAll Goal

## CoolEmAll Goal




- Goal: improve energy-efficiency of modular data centres by **optimisation** of their **design** and **operation** for a wide range of **workloads**, **IT equipment** and **cooling** options

## CoolEmAll Goal



- Goal: improve energy-efficiency of modular data centres by **optimisation** of their **design** and **operation** for a wide range of **workloads**, **IT equipment** and **cooling** options



Two main results:

- Open source simulation, visualization and decision support toolkit (**SVD Toolkit**)
- Open designs of computing building blocks (**ComputeBox Blueprints**)



## What CoolEmAll is NOT...





## What CoolEmAll is NOT...

- ...a new data center architecture and model
  - instead, various designs and models will be studied



## What CoolEmAll is NOT...

- ...a new data center architecture and model
  - instead, various designs and models will be studied
- ...just hardware realisation of a specific computing module



## What CoolEmAll is NOT...

- ...a new data center architecture and model
  - instead, various designs and models will be studied
- ...just hardware realisation of a specific computing module
  - instead, various ComputeBox designs will be used in simulations and made publicly available



## What CoolEmAll is NOT...


- ...a new data center architecture and model
  - instead, various designs and models will be studied
- ...just hardware realisation of a specific computing module
  - instead, various ComputeBox designs will be used in simulations and made publicly available
- ...data center management software
  - but will enable simulation of mngmt policies



## What CoolEmAll is NOT...

- ...a new data center architecture and model
  - instead, various designs and models will be studied
- ...just hardware realisation of a specific computing module
  - instead, various ComputeBox designs will be used in simulations and made publicly available
- ...data center management software
  - but will enable simulation of mngmt policies
- ...about new cooling methods
  - but existing cooling approaches will be modelled and simulated in diverse scenarios





- ComputeBox concept




- Modular, self-contained pool of machines
- Well defined by design parameters and energy efficiency metrics
- Rack- and container-level ComputeBox

- ComputeBox concept



- Modular, self-contained pool of machines
- Well defined by design parameters and energy efficiency metrics
- Rack- and container-level ComputeBox

- Open ComputeBox Blueprints



- Open designs of ComputeBoxes:
  - Architecture, cooling solution, node density and performance
- With energy-efficiency metrics with respect to
  - Workload, policies, conditions, and requirements





- Datacenter Efficiency Building Block (DEB)
  - 2 D schematic overview
  - 3 D model, reflecting a dataset for the SVD Toolkit
  - Datasheet including details of CPU and other components v performance (GHz, Cores) data and power consumption (TI



- Datacenter Efficiency Building Block (DEBB)
  - 2 D schematic overview
  - 3 D model, reflecting a dataset for the SVD Toolkit
  - Datasheet including details of CPU and other components v performance (GHz, Cores) data and power consumption (TI
- DEBB examples
  - Pizzabox-Design (1 Server in 1 U), Two Blades in 1 U, 4 Node in 1 U, 18 Nodes in 1 U, Bladecenters, Storage-modules



- Datacenter Efficiency Building Block (DEBB)
  - 2 D schematic overview
  - 3 D model, reflecting a dataset for the SVD Toolkit
  - Datasheet including details of CPU and other components v performance (GHz, Cores) data and power consumption (TDP)
- DEBB examples
  - Pizzabox-Design (1 Server in 1 U), Two Blades in 1 U, 4 Nodes in 1 U, 18 Nodes in 1 U, Bladecenters, Storage-modules
- Licensing and availability
  - The blueprints of the designed and evaluated ComputeBoxes for specific usage scenarios
  - DEBBs will be available as a free download on a website
    - Printed Circuit Board (PCB) layouts are not going to get published



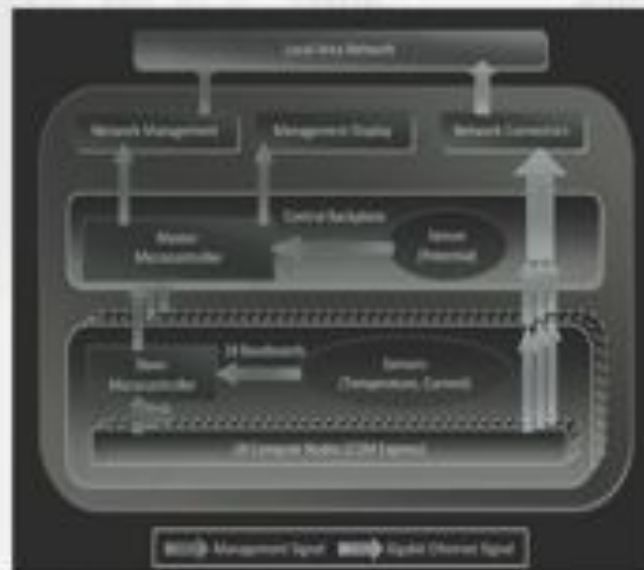


Skinless compute node for RECS (2 or 4 cores, max. 16 GB RAM)


- RECS - Resource Efficient Computing System
- Highly integrated monitoring and control infrastructure
- Airflow and liquid cooling, diverse density



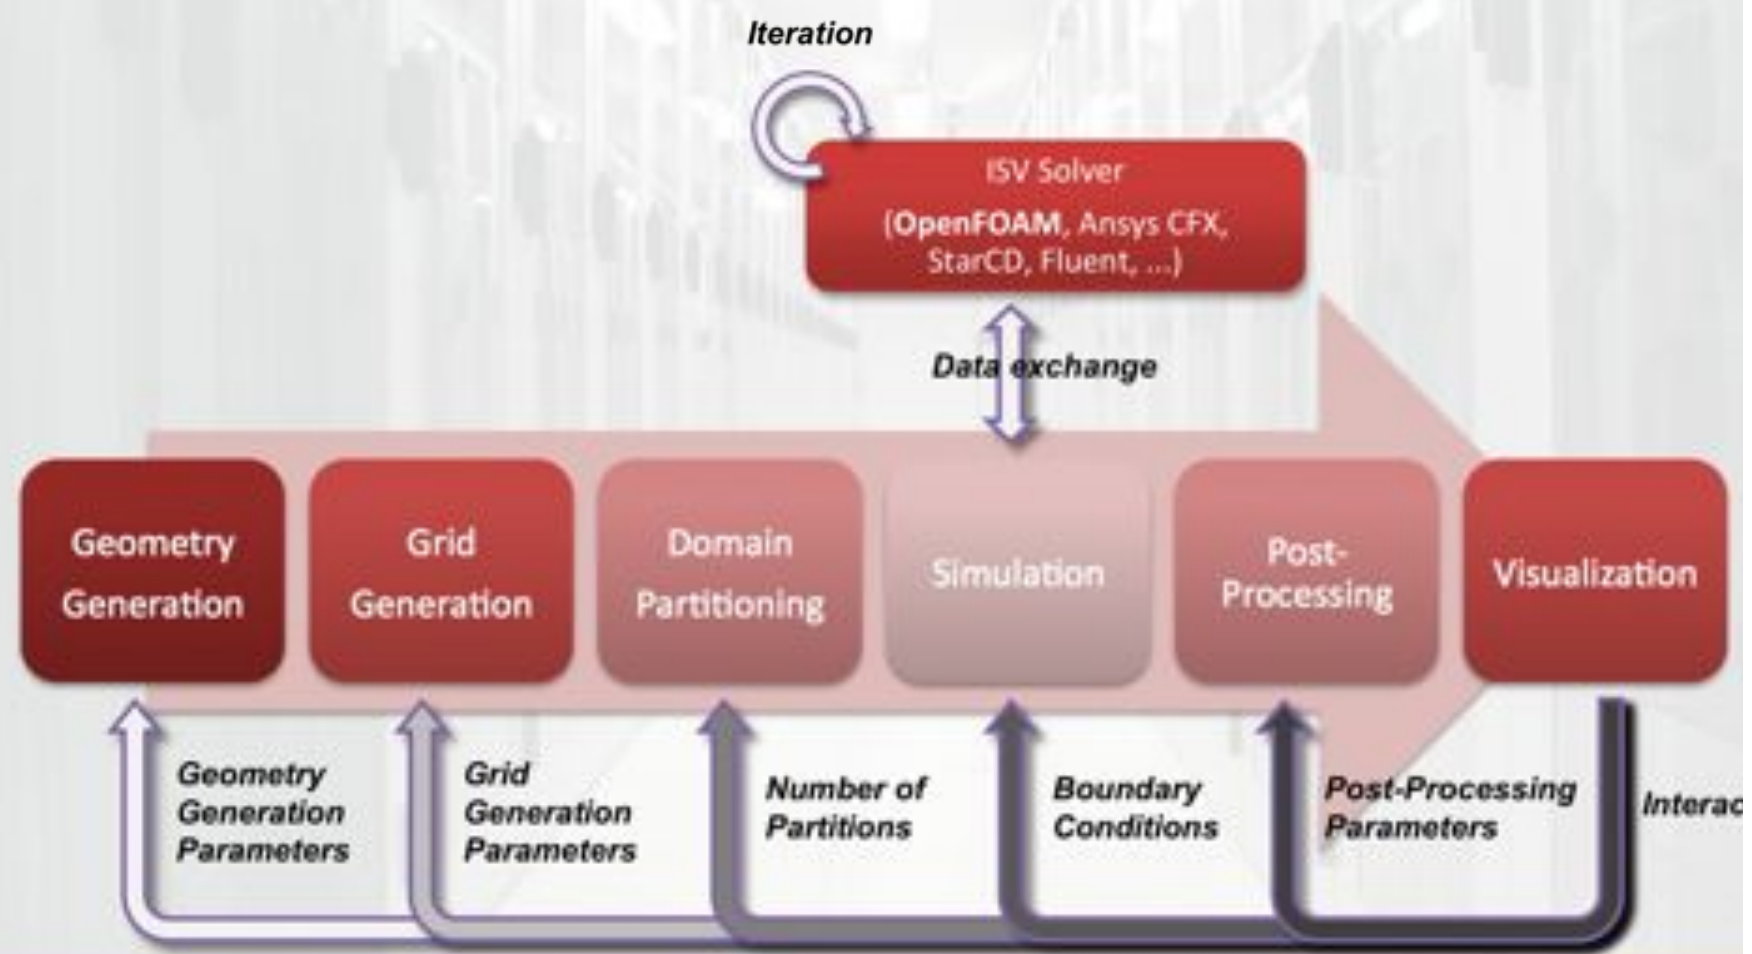
RECS backbone and baseboards  
January 2012

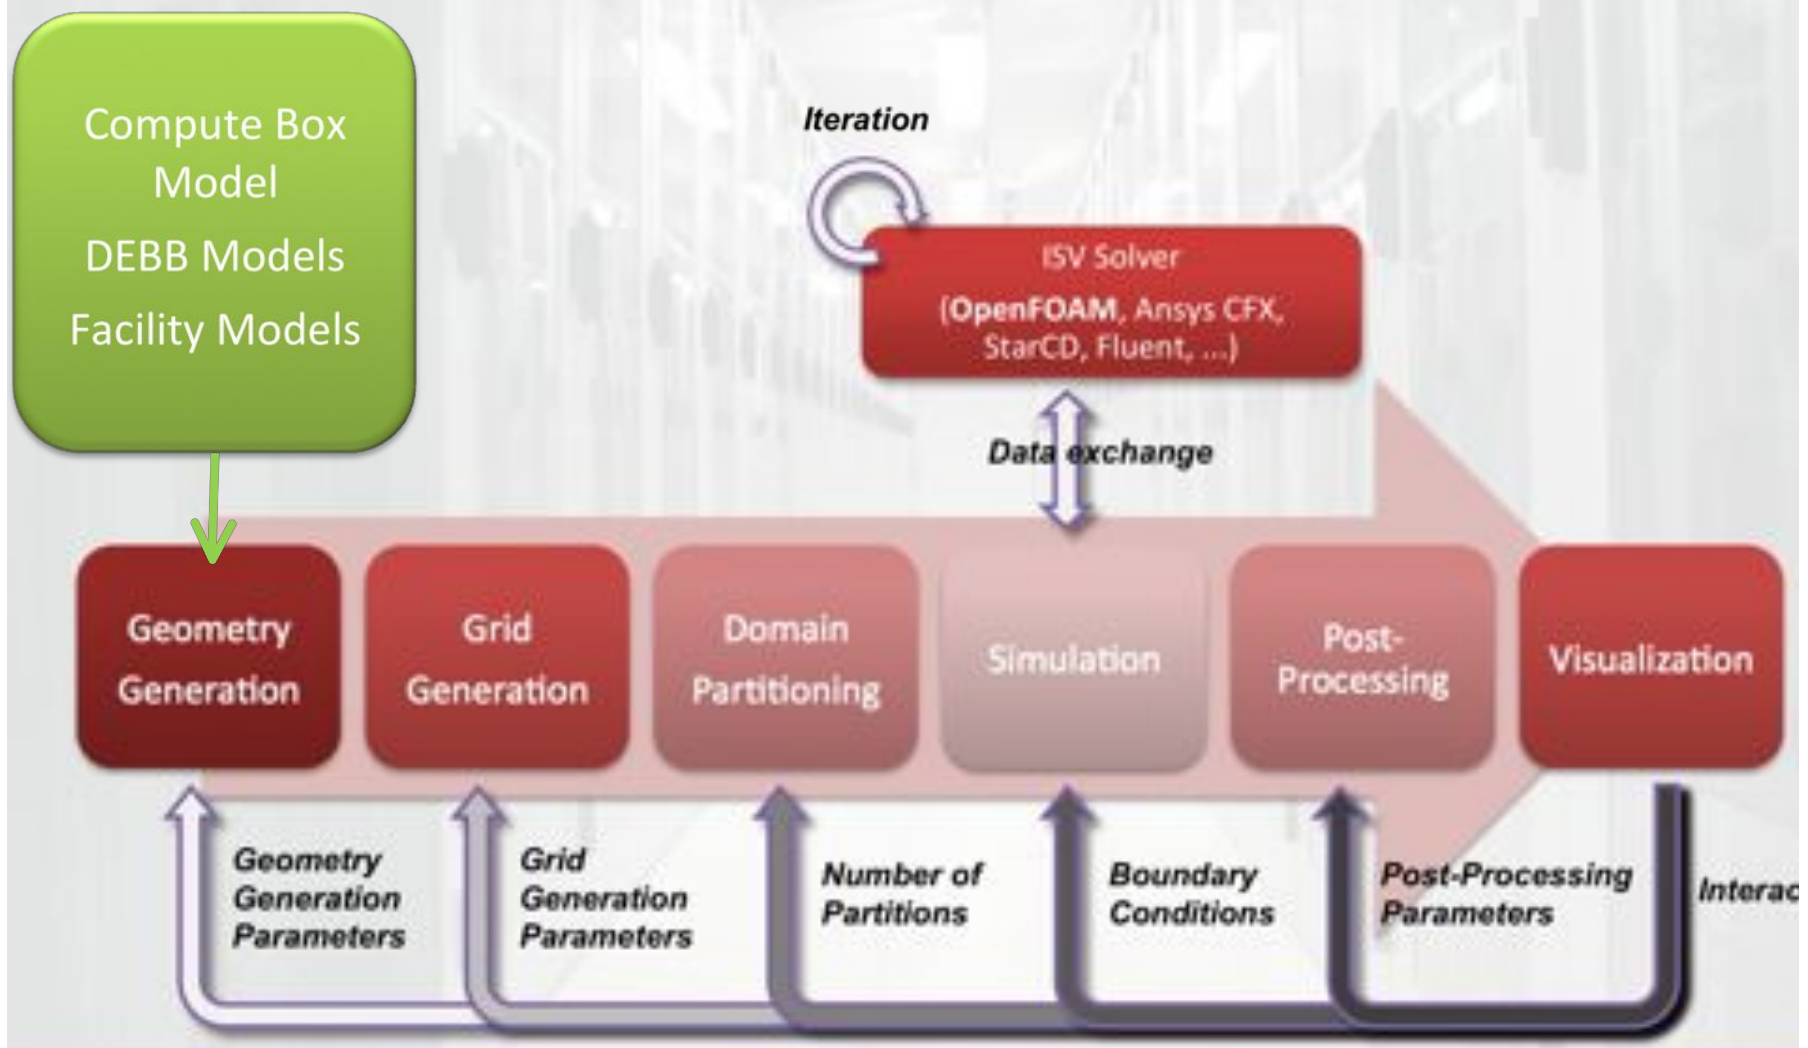


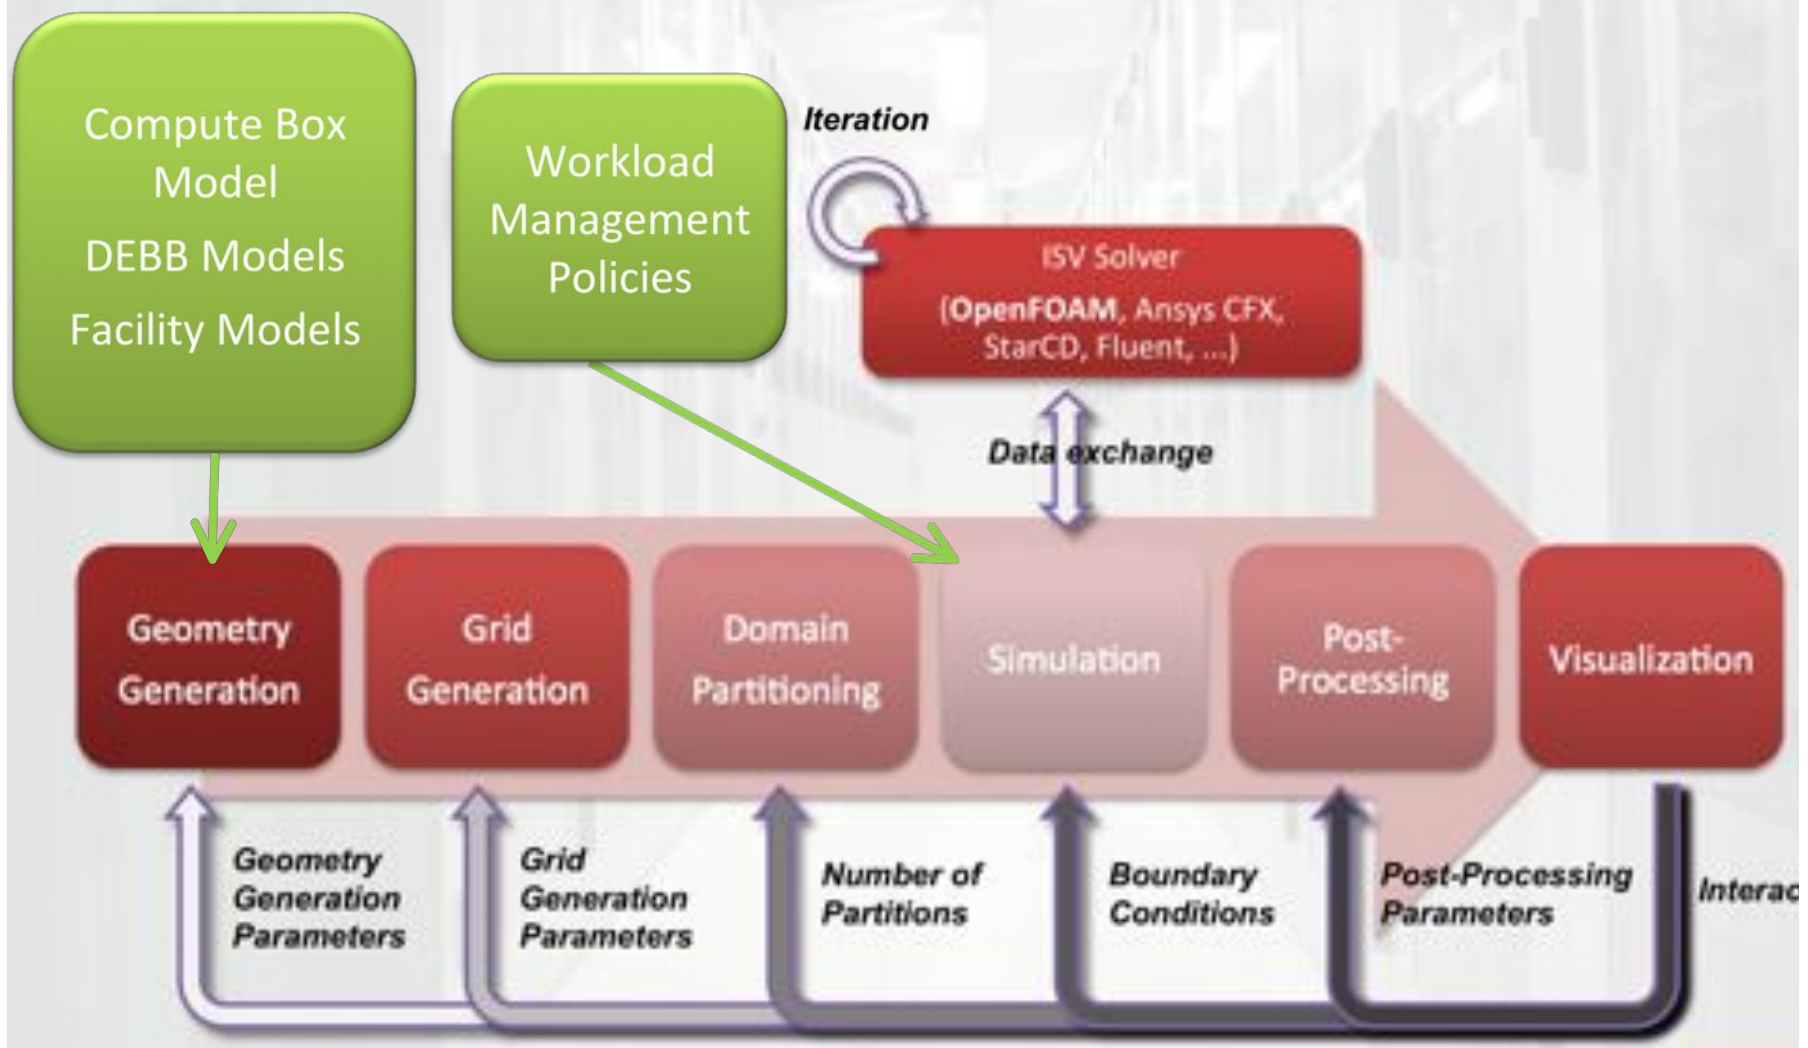
Architecture of the master-slave microcontroller system  
CoolEmAll - GreenDays@Lyon



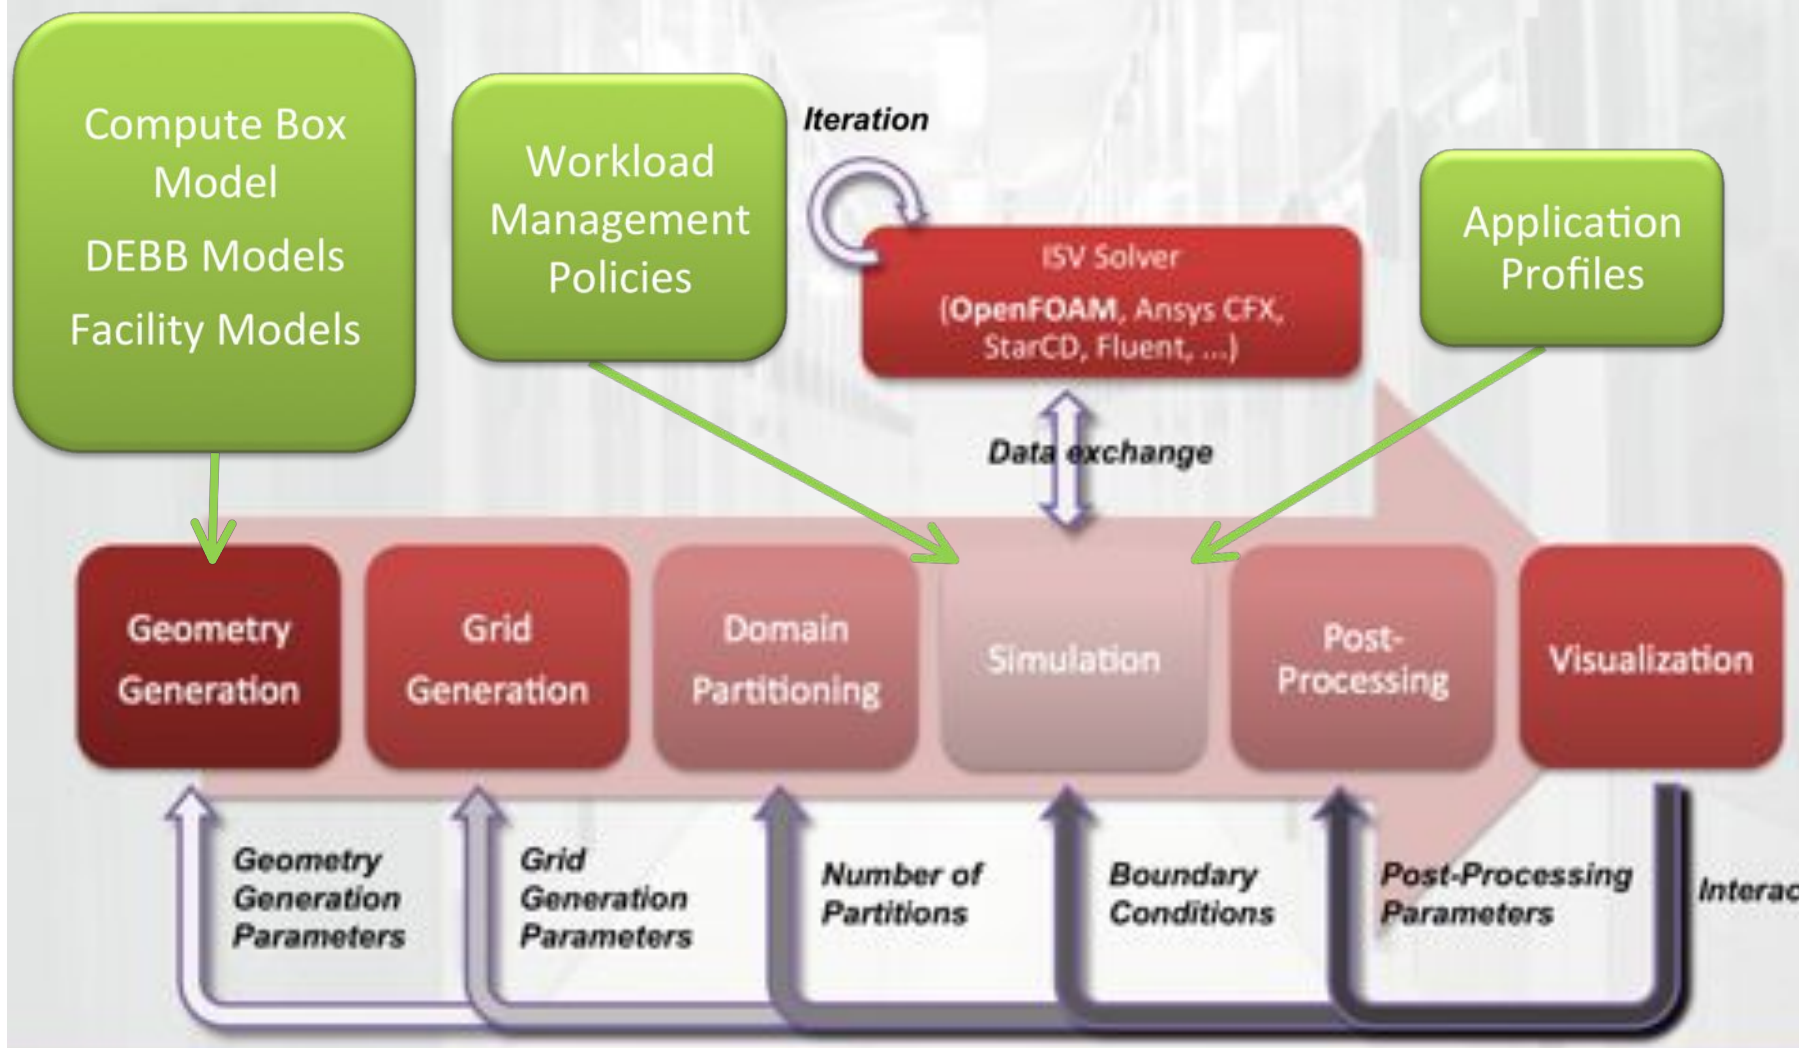
RECS microcontroller system

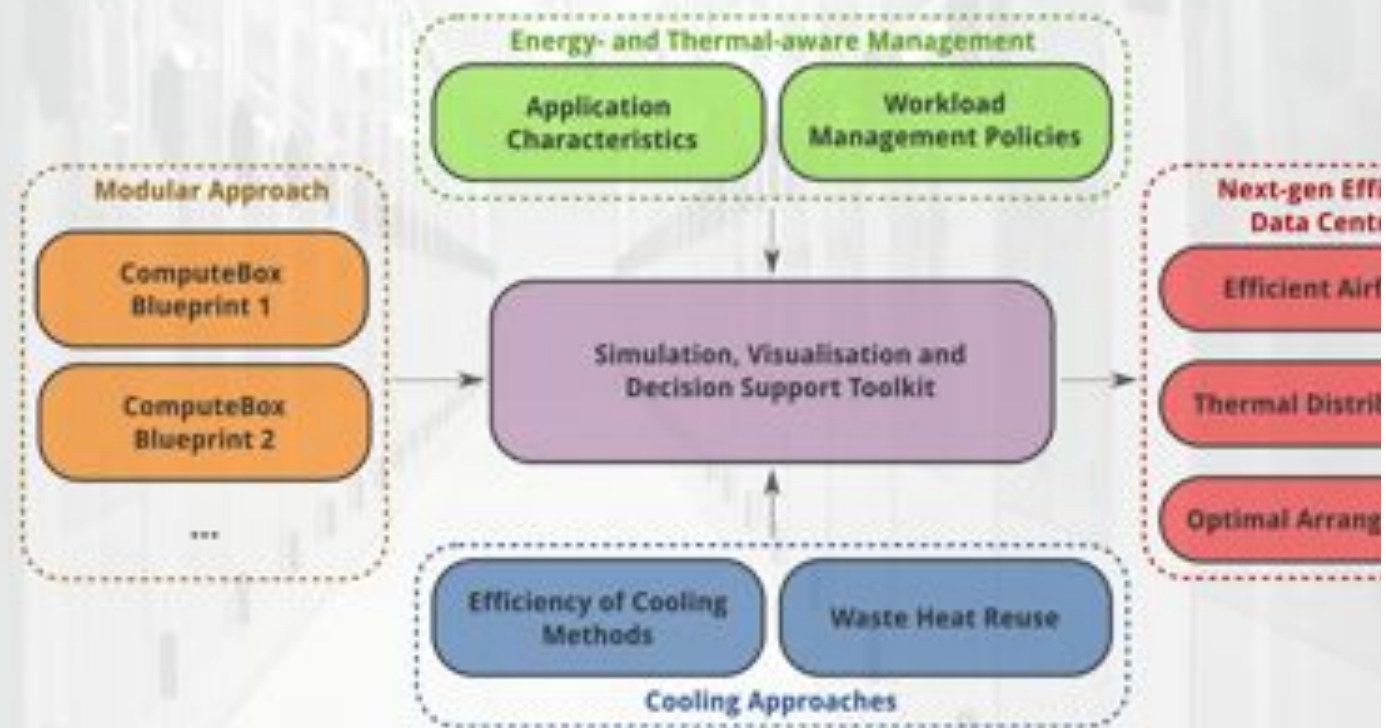












## Scale

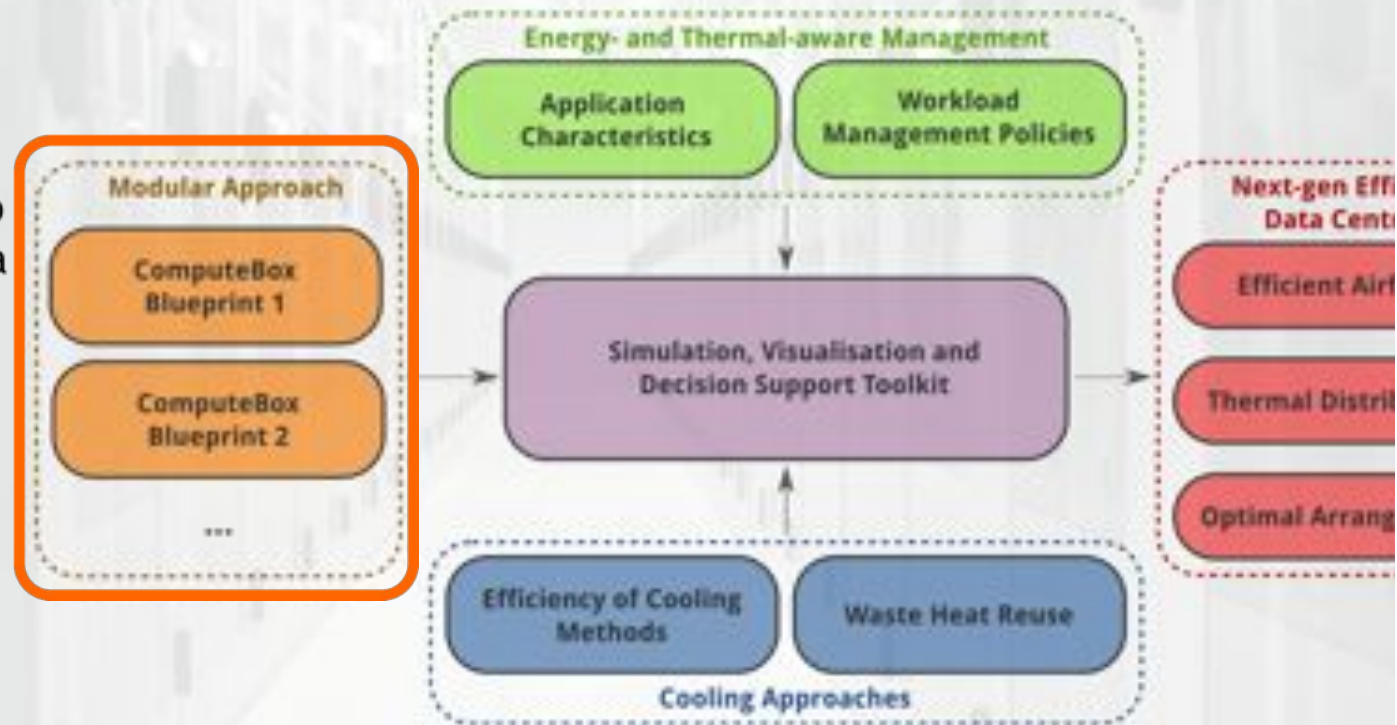
- Rack(s)
- Container(s)

## Density

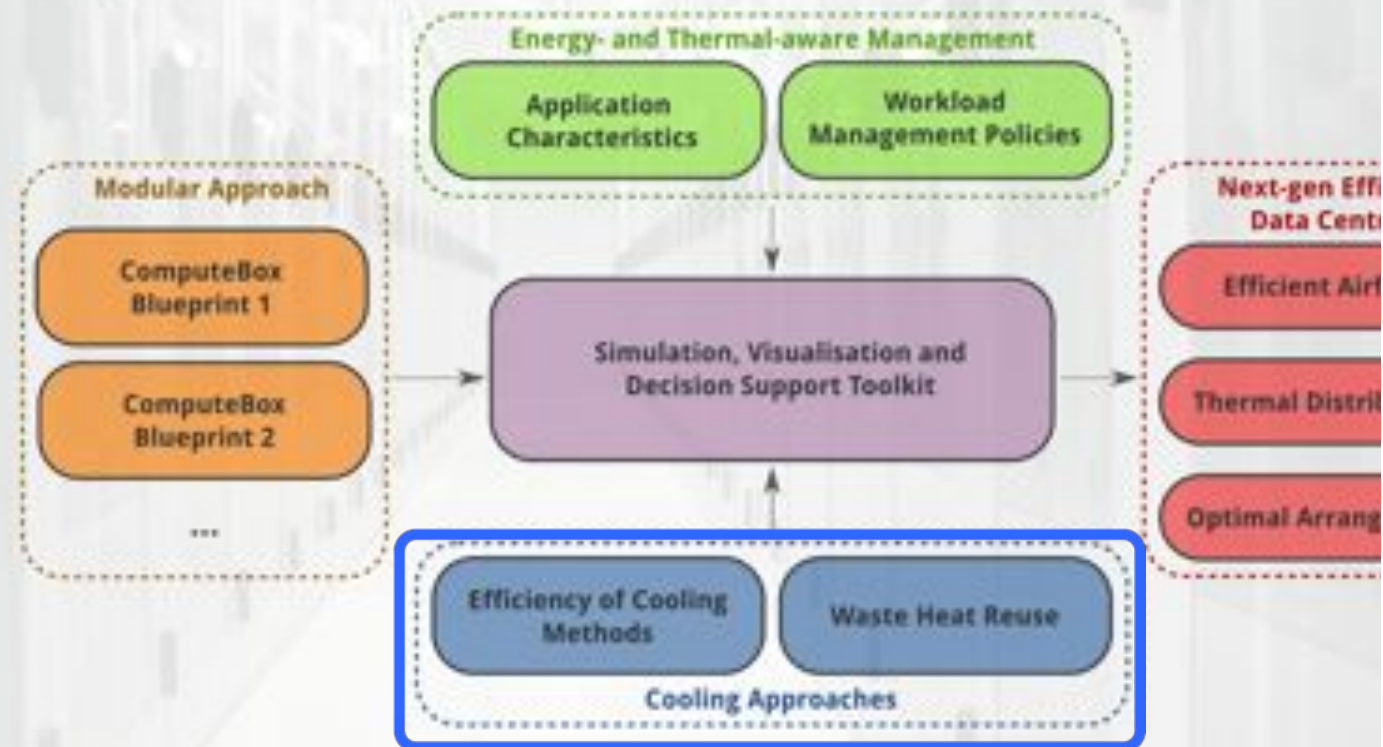
- High density (up to hundreds nodes in a rack)
- Low density

## Cooling

- Integrated
- No integrated cooling



- Direct expansion
- Air-cooled chiller
- Water-cooled chiller
- Integrated liquid cooling
- Free air cooling (air economizer)
- Heat re-use



## Application types

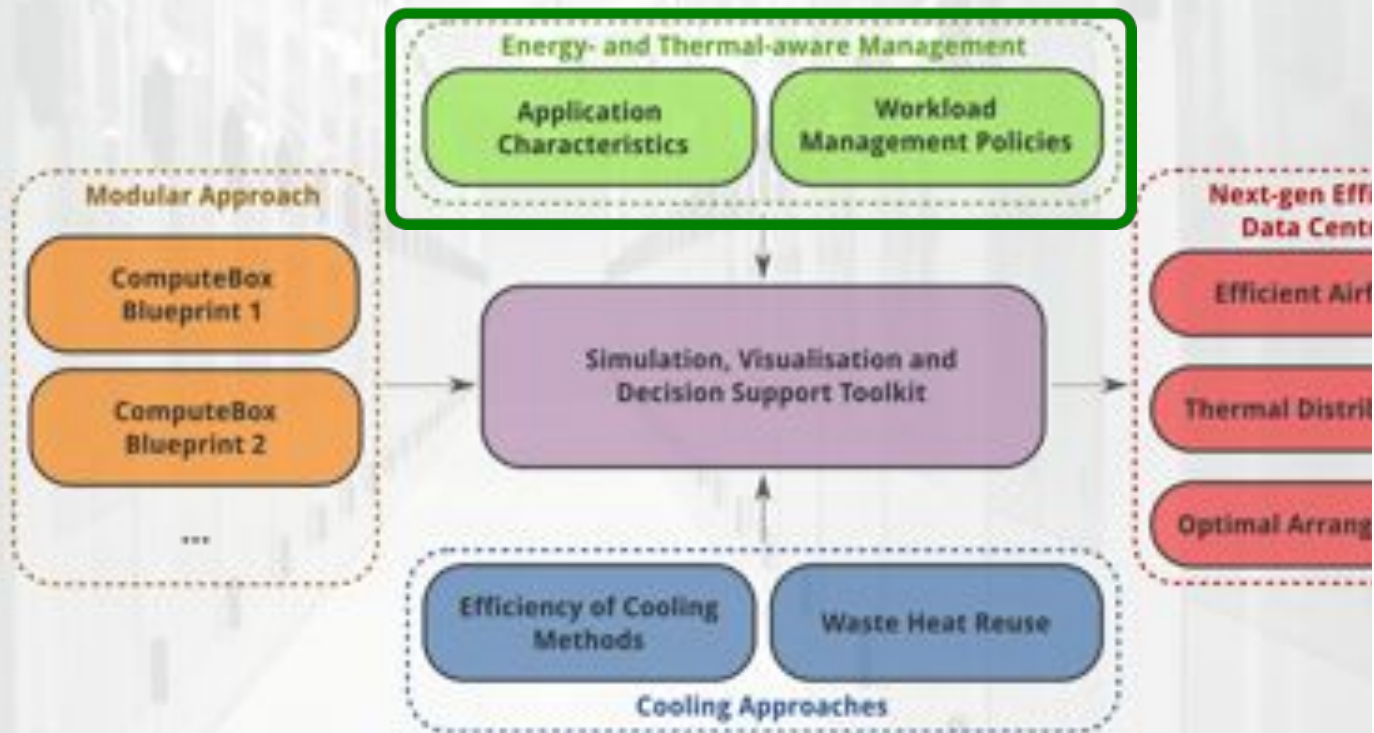
- HPC
- Virtual machines


## Application characteristics

- CPU-bound
- IO-bound
- Scale

## Workload mngmt policies


- Workload consolidation
- Energy-aware policies
- Thermal-aware policies









- Many data centers does not operates at full load all the time
  - Space for optimization and savings



- Many data centers does not operates at full load all the time
  - Space for optimization and savings
- Workload classes differ depending on center type
  - HPC applications, High throughput jobs, virtualization, services





- Many data centers does not operates at full load all the time
  - Space for optimization and savings
- Workload classes differ depending on center type
  - HPC applications, High throughput jobs, virtualization, services
- Energy consumption and heat generation for various applications differs
  - Up to 100W difference between applications (50W difference for similar CPU load)
  - 10-20C temperature differences






- Popular energy efficiency metrics proposed by Green Grid
  - PUE and DCiE







- Popular energy efficiency metrics proposed by Green Grid
  - PUE and DCiE
- Metrics to be studied within CoolEmAll
  - Rack-level
  - Workload/Application-specific
  - Total energy related
  - Heat re-use (ERE)



- Virtual thermal camera
- Heat transfer and air flow processes
- Remote, interactive visualization

- RECS configurations
  - RECS with a higher level of integrated management
  - Towards "data center as appliance"
- HLRS existing center
  - 5MW of power, ~2Petaflops
  - Dynamic, interactive visualization to optimize temperature distribution
- PSNC new center
  - Completely new center in design phase
  - 2MW-16MW power
  - 2x800m<sup>2</sup>, up to ~170 racks
- ATOS Origin data center
  - 45 data centres, total capacity ~40.000 m<sup>2</sup>
  - Greener Data Centres global program
  - Requirements from industrial DCs





- Use and extension of existing tools

COVISE

Open  FOAM

RECS™





- Use and extension of existing tools

COVISE

OpenFOAM

RECS™

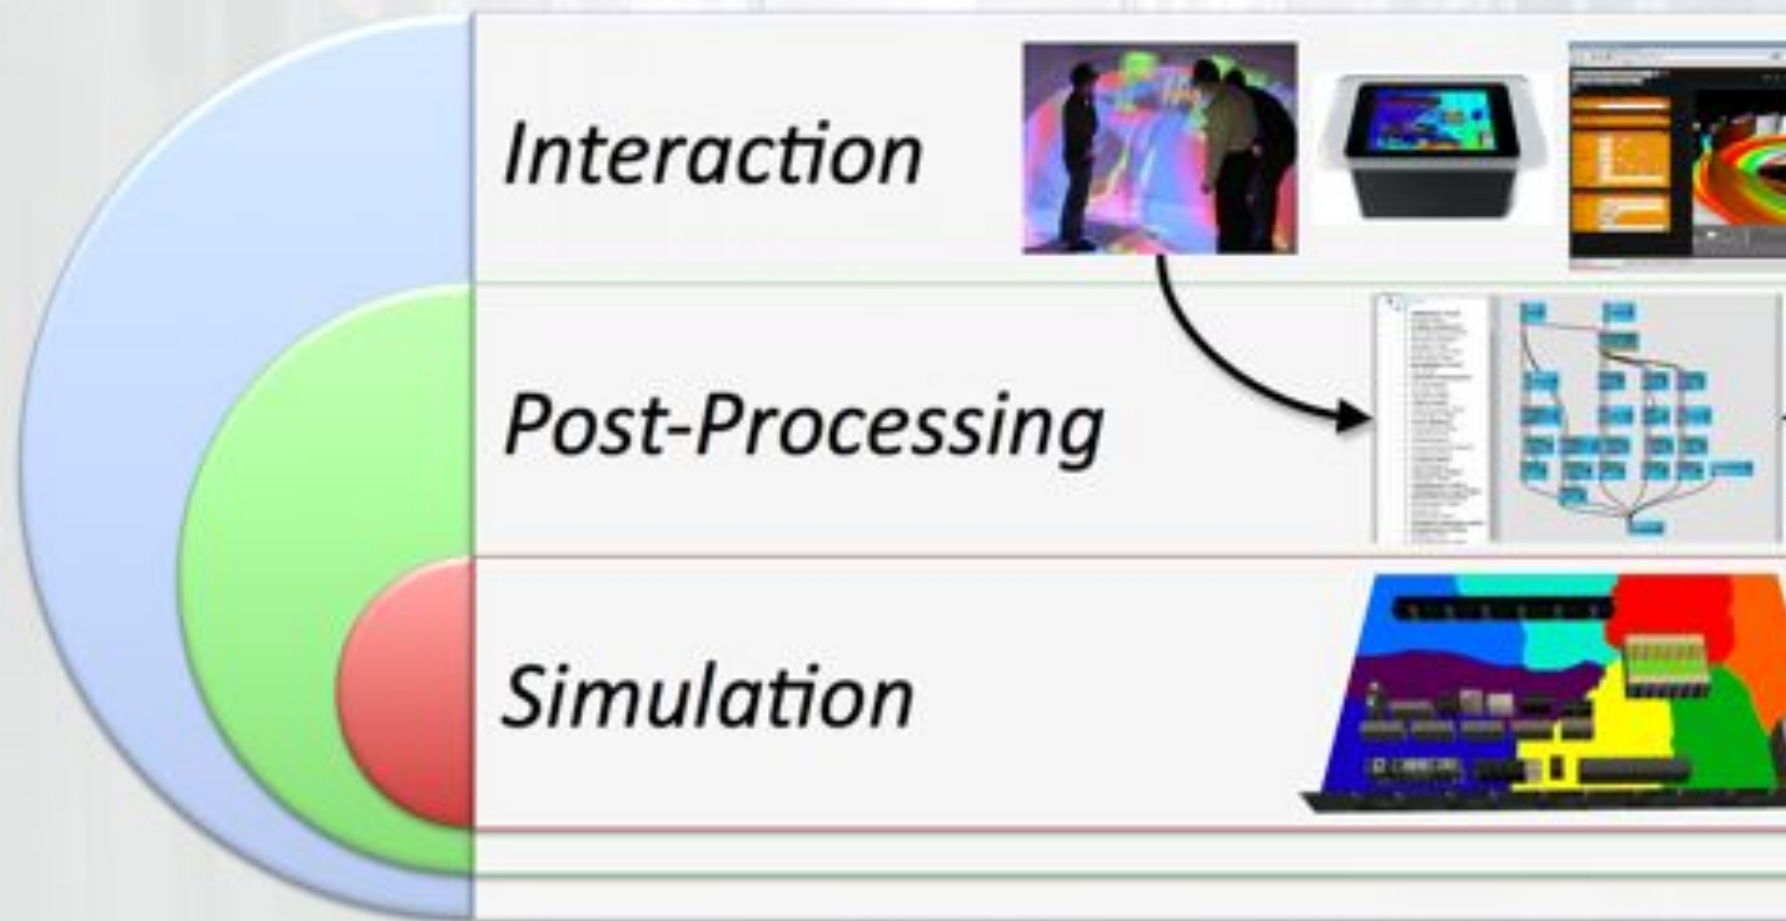


- Collaboration with projects and institutes



UptimeInstitute





*Interaction*



*Post-Processing*

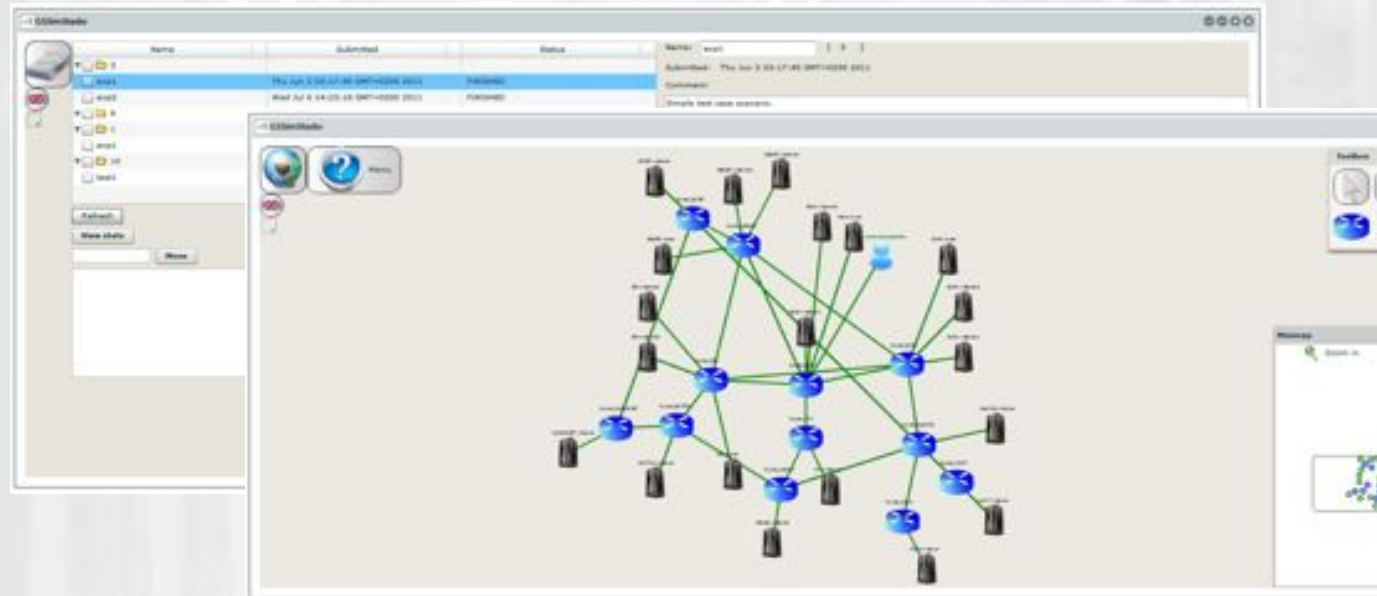


*Simulation*


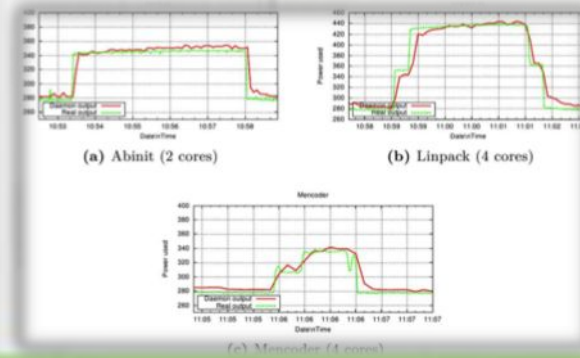
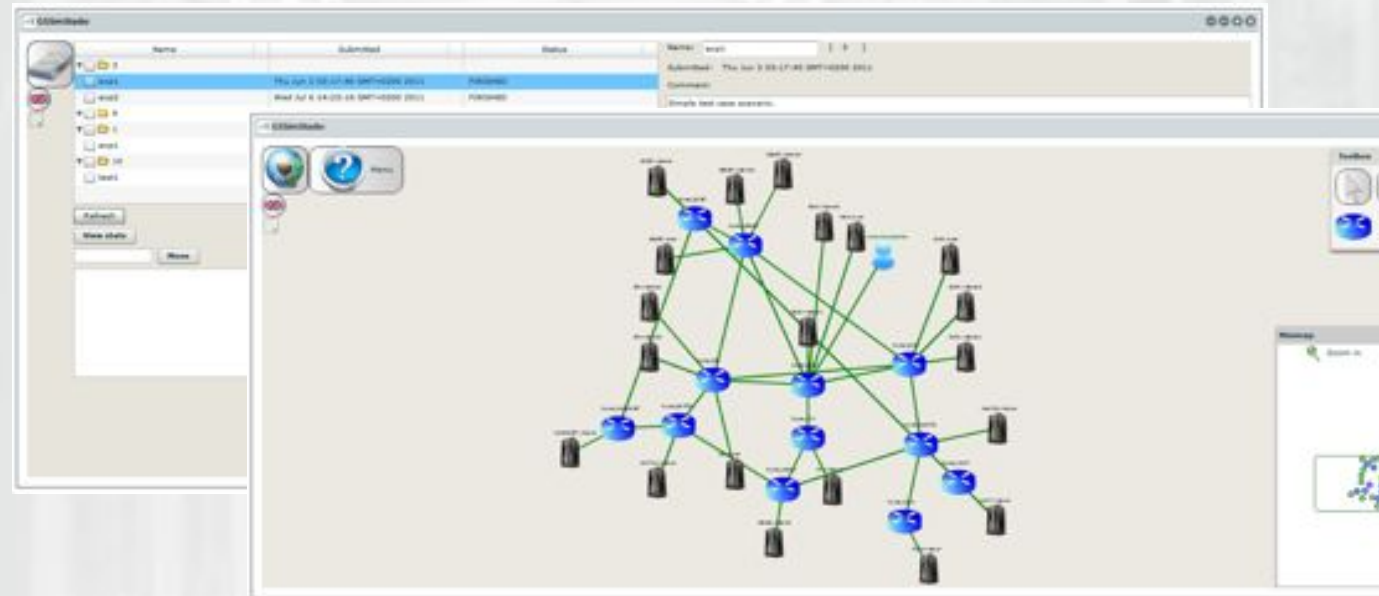




- GSSIM
  - workload modelling
  - Management policies
  - resource and network topology



- GSSIM
  - workload modelling
  - Management policies
  - resource and network topology
- Application models
  - estimated power usage
  - application characteristics





# What CoolEmAll may bring for you?



- For HPC center administrators, IT equipment providers, data center designers:
  - SVD toolkit to investigate data center energy-efficiency based on simulations
  - Best practises and open designs (CB blueprints, DEBB specifications)





- For HPC center administrators, IT equipment providers, data center designers:
  - SVD toolkit to investigate data center energy-efficiency based on simulations
  - Best practises and open designs (CB blueprints, DEBB specifications)
- For CoolEmAll:
  - Variety of validation scenarios
  - Specific tested ComputeBoxes



- WP 1 Project Management
- WP 2 Simulation, Visualisation and Decision Support Toolkit
- WP 3 ComputeBox Prototype
- WP 4 Workload and Resource Management Policies
- WP 5 Energy-efficiency Metrics (leader: IRIT)
  - Metrics
  - Monitoring of applications
- WP 6 Requirements, Verification and Validation Scenarios
- WP 7 Dissemination, Exploitation and RTD Standardization

- Takes specific characteristics of applications/workloads into consideration and specifies energy and thermal efficient metrics. This work package will be responsible for preparing benchmarks and classification of applications and workloads.
- Metrics, monitoring, benchmarking and Application characterization
  - Derive energy-efficiency metrics for computing modules extending existing power related metrics to energy related metrics (i.e. including time) taking also into account the runtime environment of the data centre (ambient temperature, heat re-use capacities [IREC]).
  - Design and develop a monitoring infrastructure adapted to energy- and heat-aware scheduling [IRIT]
  - Design a methodology for profiling applications in respect with their energy consumption [ATOS]
  - Develop benchmarks to evaluate derived metrics [IRIT]



At least :

- WP3: Design of compute box and blueprint
- WP4: Monitoring, scheduling
- WP5: Métriques, monitoring, profilage
- WP6: Validation processus
- WP7: Dissemination / standardization

2 PhD. recruited

1 Eng. to come.



- CoolEmAll will deliver two main products:
  - Open **ComputeBox Blueprints**
  - Open source **SVD Toolkit**





- CoolEmAll will deliver two main products:
  - Open **ComputeBox Blueprints**
  - Open source **SVD Toolkit**
- These outcomes will allow minimising energy consumption of data centers by
  - **Optimisation of their design and operation**
  - **Tuning to specific workloads and conditions**



Thank you for attention!  
Questions?