

# **Overview of HPC Computer Architecture:**

## ***A Long March Toward Exa-Scale Computing and Beyond***

\

**August 16, 2012**

**Guang R. Gao**

**ACM Fellow and IEEE Fellow**  
**Distinguished Professor, Dept. of ECE**  
**University of Delaware**

# Toward A *Codelet* Based Execution Model and Its Memory Semantics

*-- For Future Extreme-Scale Computing Systems*

\

August 16, 2012

Guang R. Gao


ACM Fellow and IEEE Fellow  
Distinguished Professor, Dept. of ECE  
University of Delaware

# Outline

- **Background and motivation**
- **Program execution models**
- **Evolution of codelet based execution models**
  - The EARTH project (1994 – 2004)
  - IBM Cyclops-64 project (2004 – 2010+ ): The TNT Experience
  - Intel-led UHPC/Runnemedede (2010 – 2012): The codelet concept and SWARM
- **Memory semantics the codelet model**
- **Conclusions and Future Directions**

# K (“KEI”) Computer

- "K" draws upon the Japanese word "Kei" for  $10^{16}$
- 3 times faster than Chinese Tianhe 1A
- 8.162 Pflops Rmax, 8.777 Pflops Rpeak
- 80,000 8-core 2GHz SPARC64 VIIIfx to deliver a total of more than 640,000 processing cores
- 1 PB memory
- 4<sup>th</sup> most energy-efficient system in the 500, with a performance-per-watt rating of 825 megaflops per Watt.
- Tofu : A 6D Mesh/Torus Interconnect





# Tianhe-1A 2.566 Petaflops Rmax



# Current Big Themes in Supercomputing

- Multi-core → Many-core
  - Exa-Scale is on horizon
- Heterogeneity and Accelerators
- Data-Intensive (big-data)
- Others ?

# Challenges

- **Challenges:**
  - **Big-compute** (performance demand on *massively parallelism*)
  - **Big-data** (massive, irregular, unstructured data need *big analytics*)
  - **Big chips** with architecture heterogeneity
  - **Energy efficiency and resiliency**

# A Fundamental Challenge - *Parallel Program Execution Models*



# Outline

- Background and motivation
- **Program execution models**
- Evolution of codelet based execution models
  - The EARTH project (1994 – 2004)
  - IBM Cyclops-64 project (2004 – 2010+ ): The TNT Experience
  - Intel-led UHPC/Runnemedede (2010 – 2012): The codelet concept and SWARM
- Semantics of the codelet model
- Conclusions and Future Directions

# A Quiz: Have you heard the following terms ?

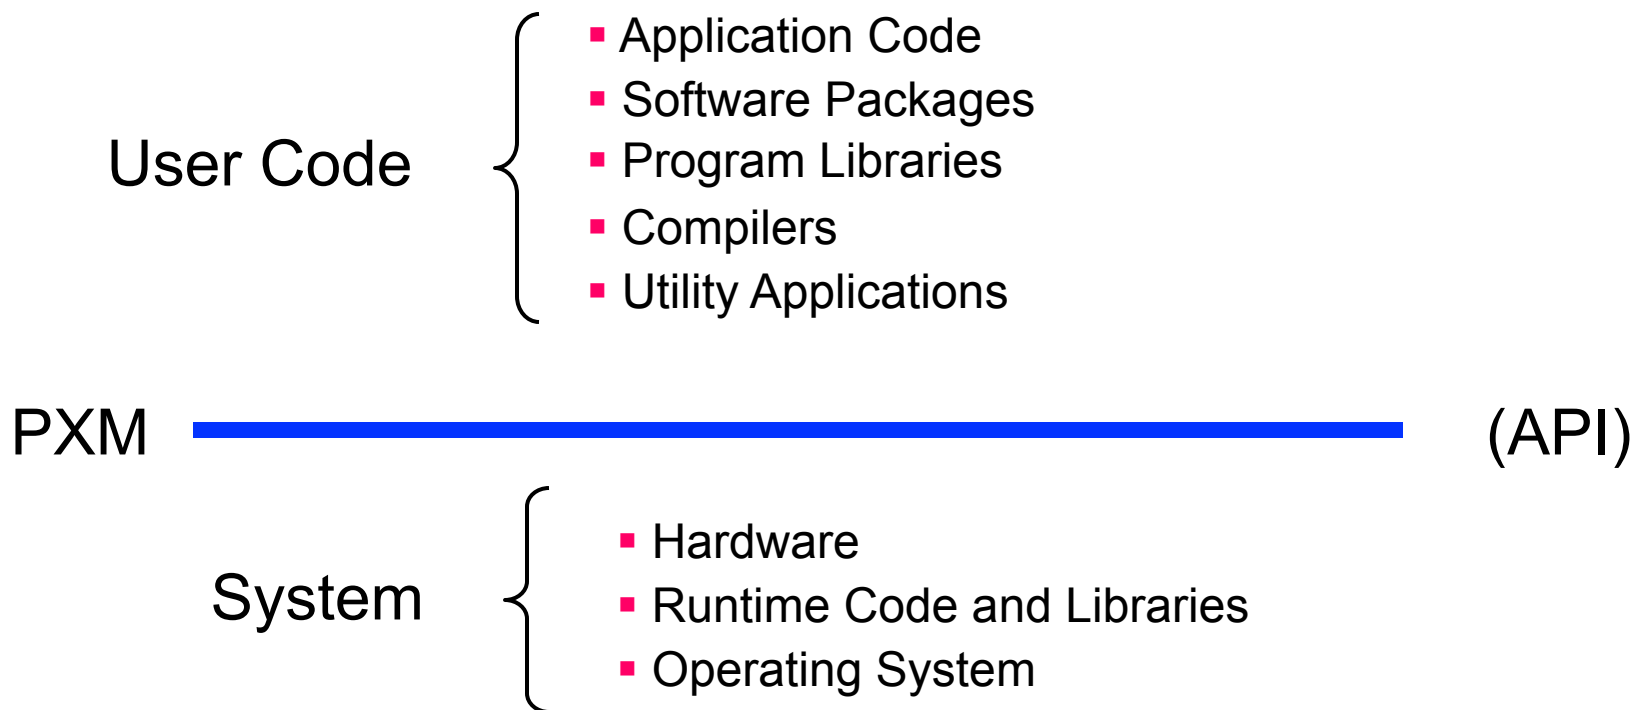
Actors (dataflow) ?


strand ?

fiber ?


codelet ?

# What is a Program Execution Model?





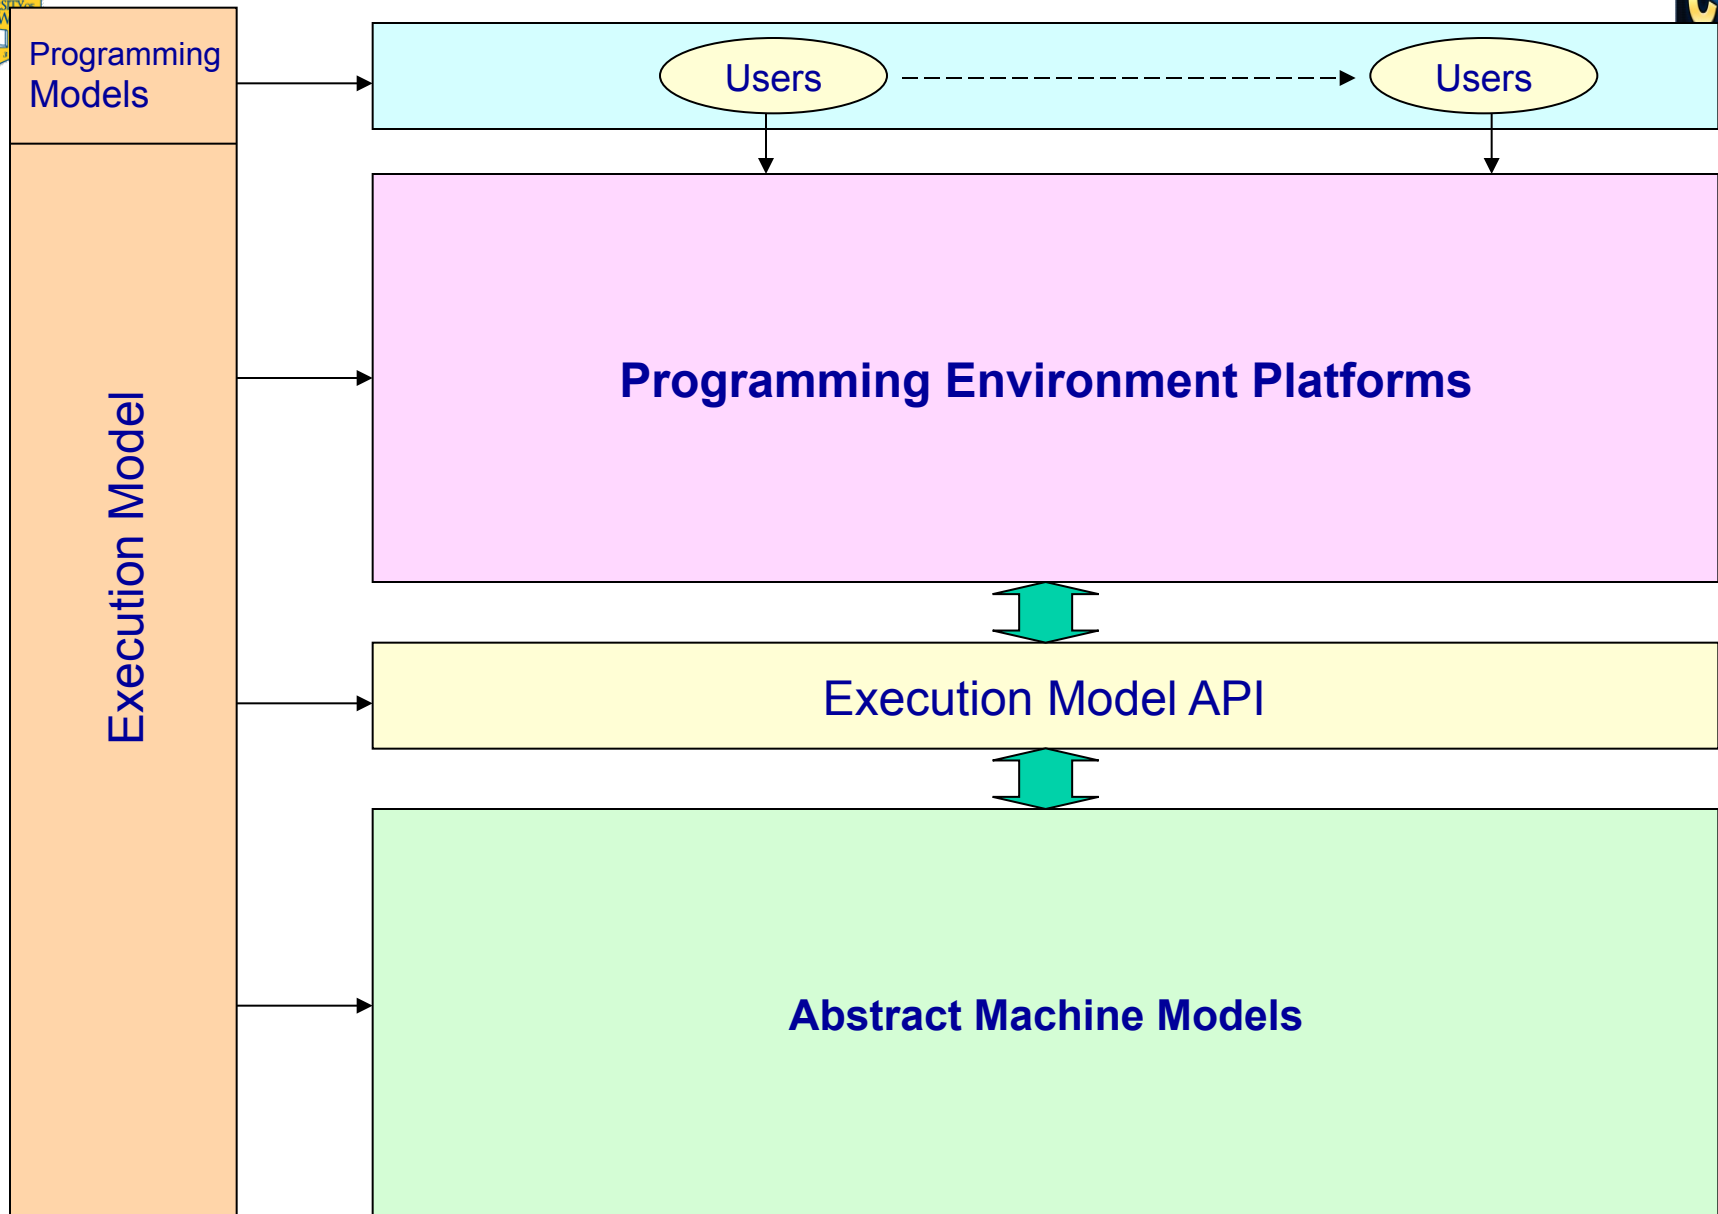
**Coarse-Grain thread-**  
The family home model



**Fine-Grain *non-preemptive* thread-**  
The “hotel” model

# Coarse-Grain vs. Fine-Grain Multithreading

[Gao: invited talk at Fran Allen’s Retirement Workshop, 07/2002]

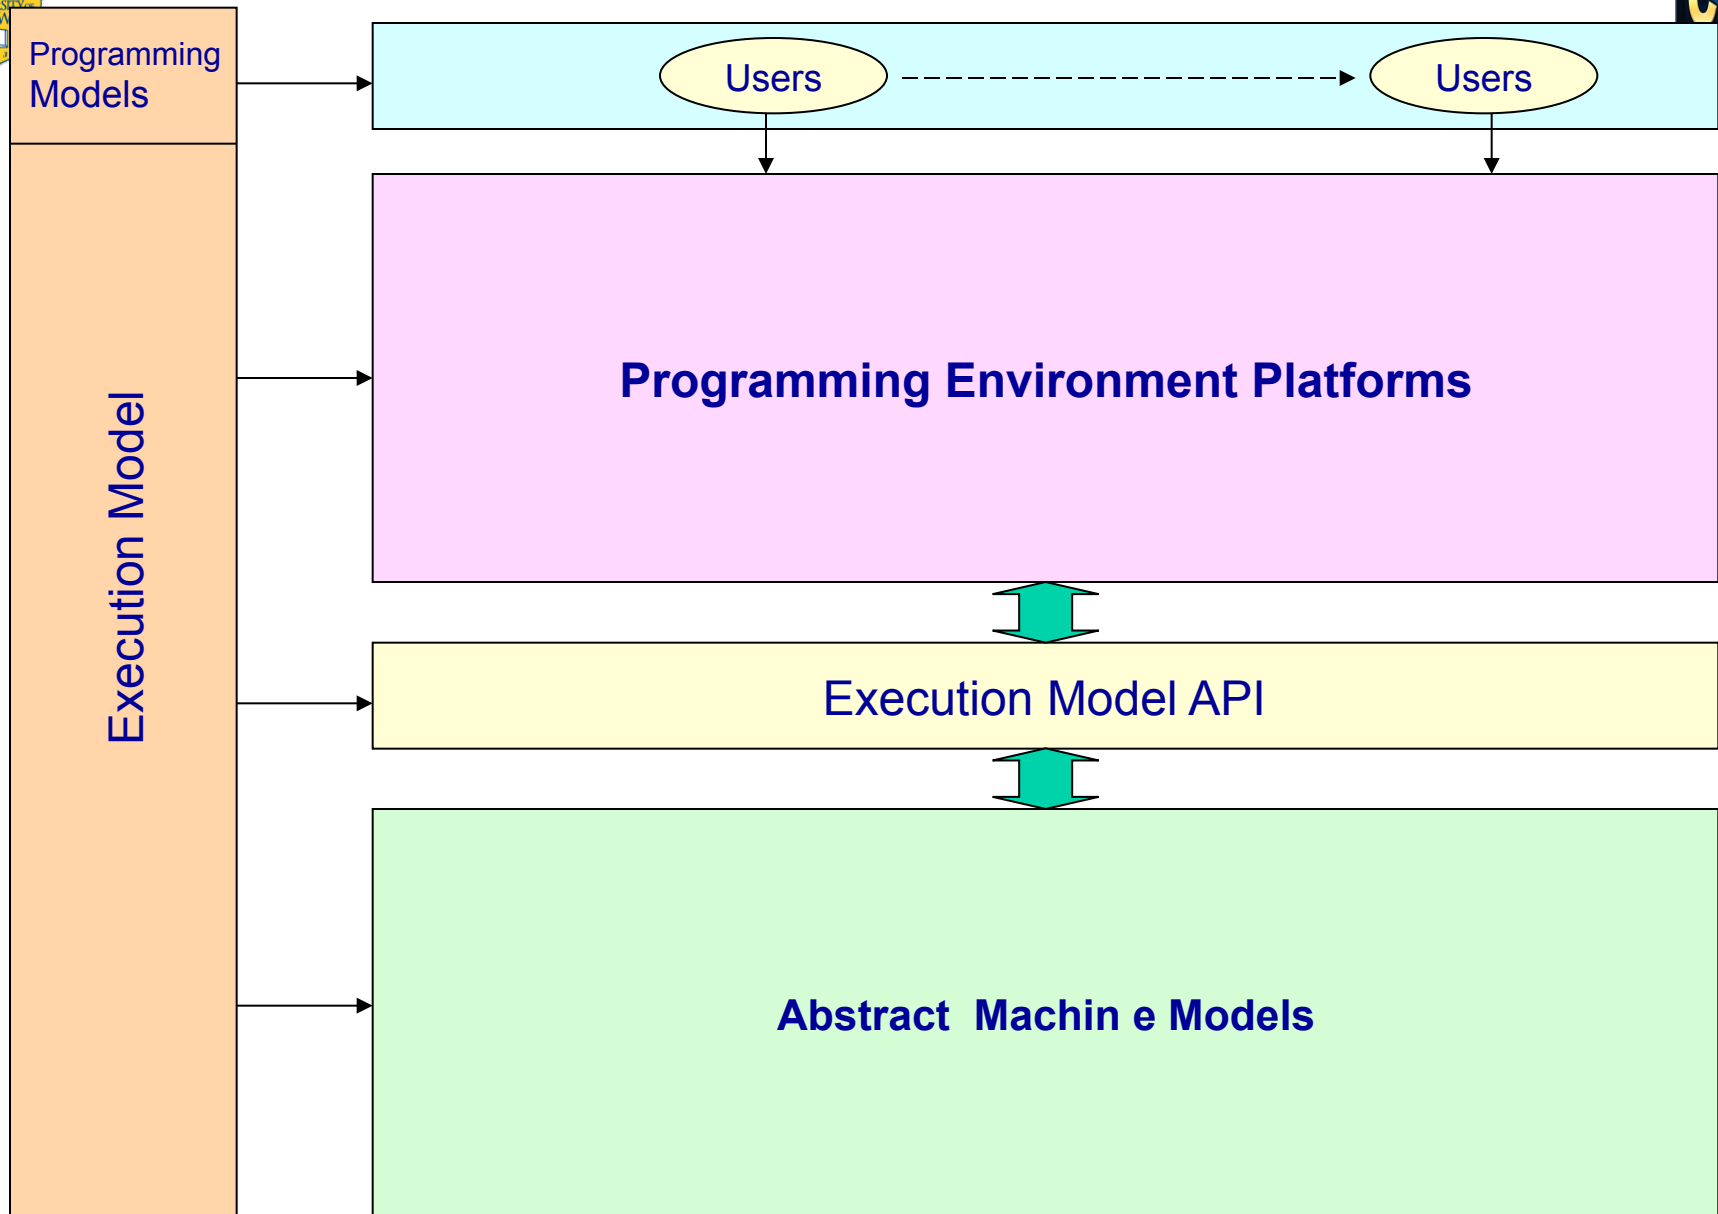


# Execution Model and Abstract Machines

# Outline

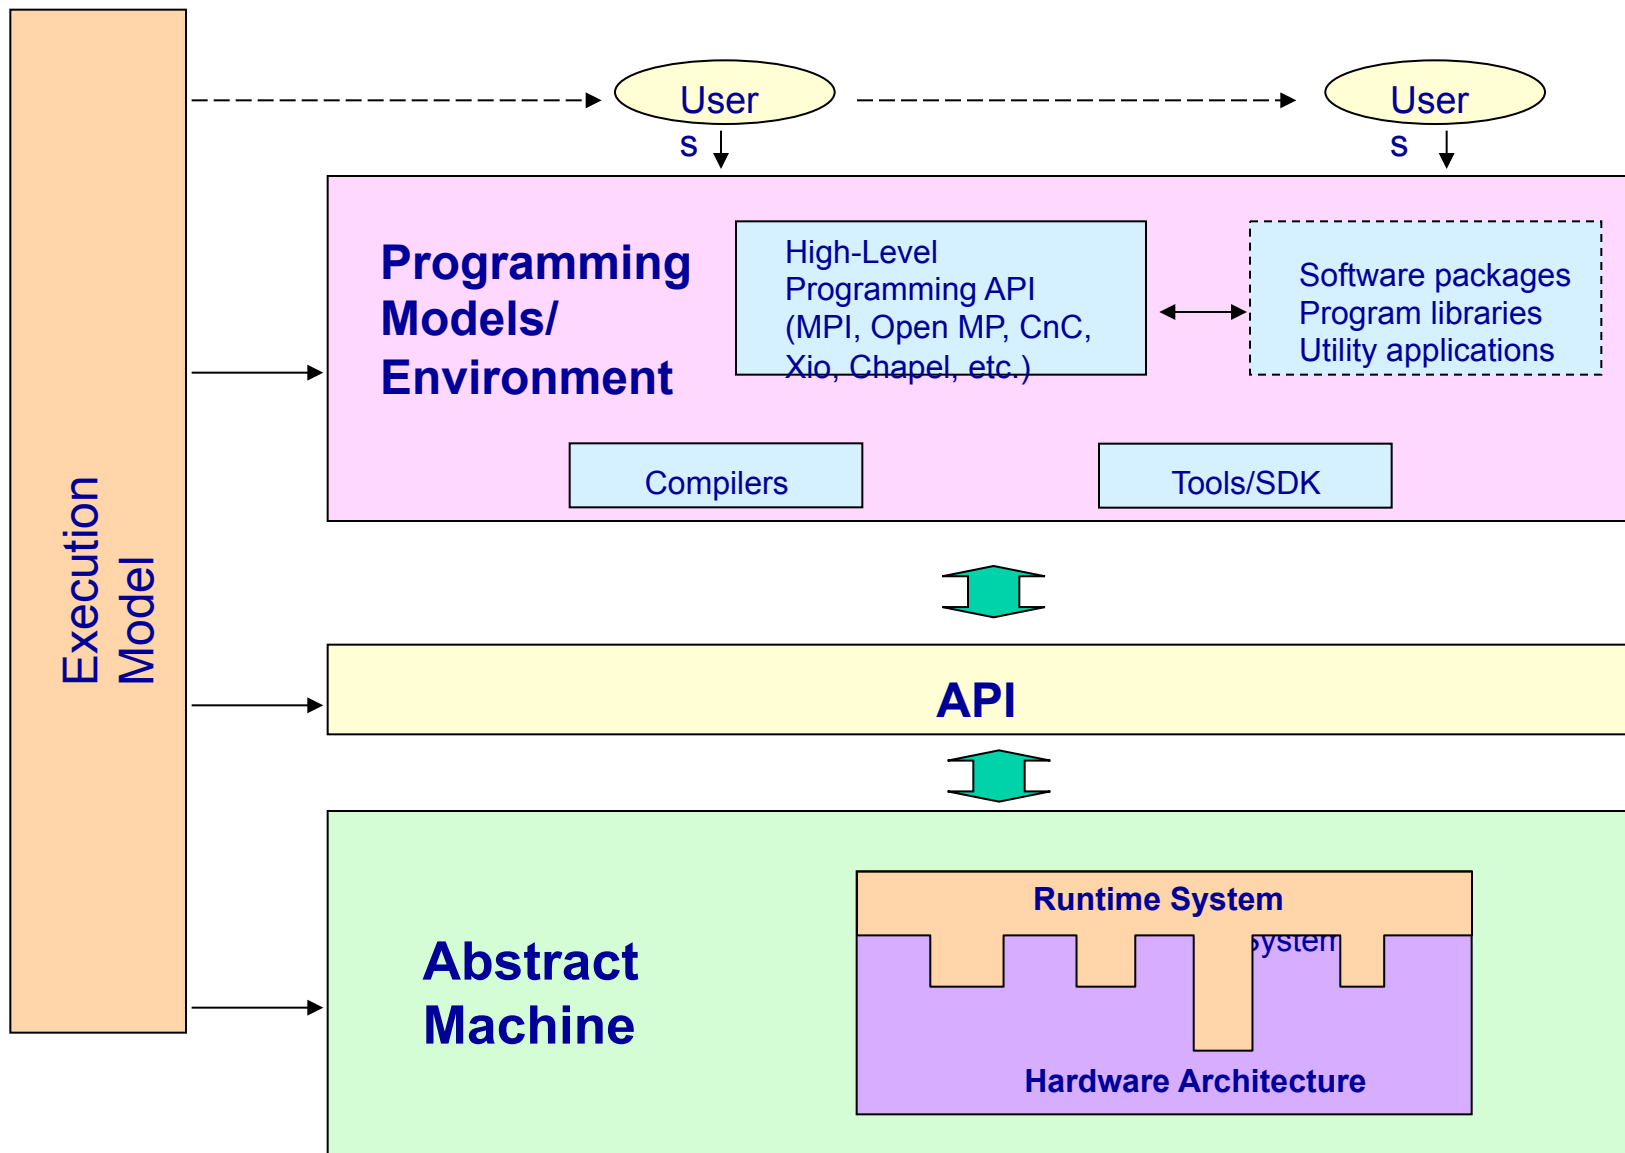
- Background and motivation
- Program execution and abstract machine models
- **Codelet based execution models**
  - **The EARTH project (1994 – 2004)**
  - IBM Cyclops-64 project (2004 – 2010+ ): The TNT Experience
  - Intel-led UHPC/Runnemedede (2010 – 2012): The codelet concept and SWARM
- Semantics of the codelet model
- Conclusions and Future Directions





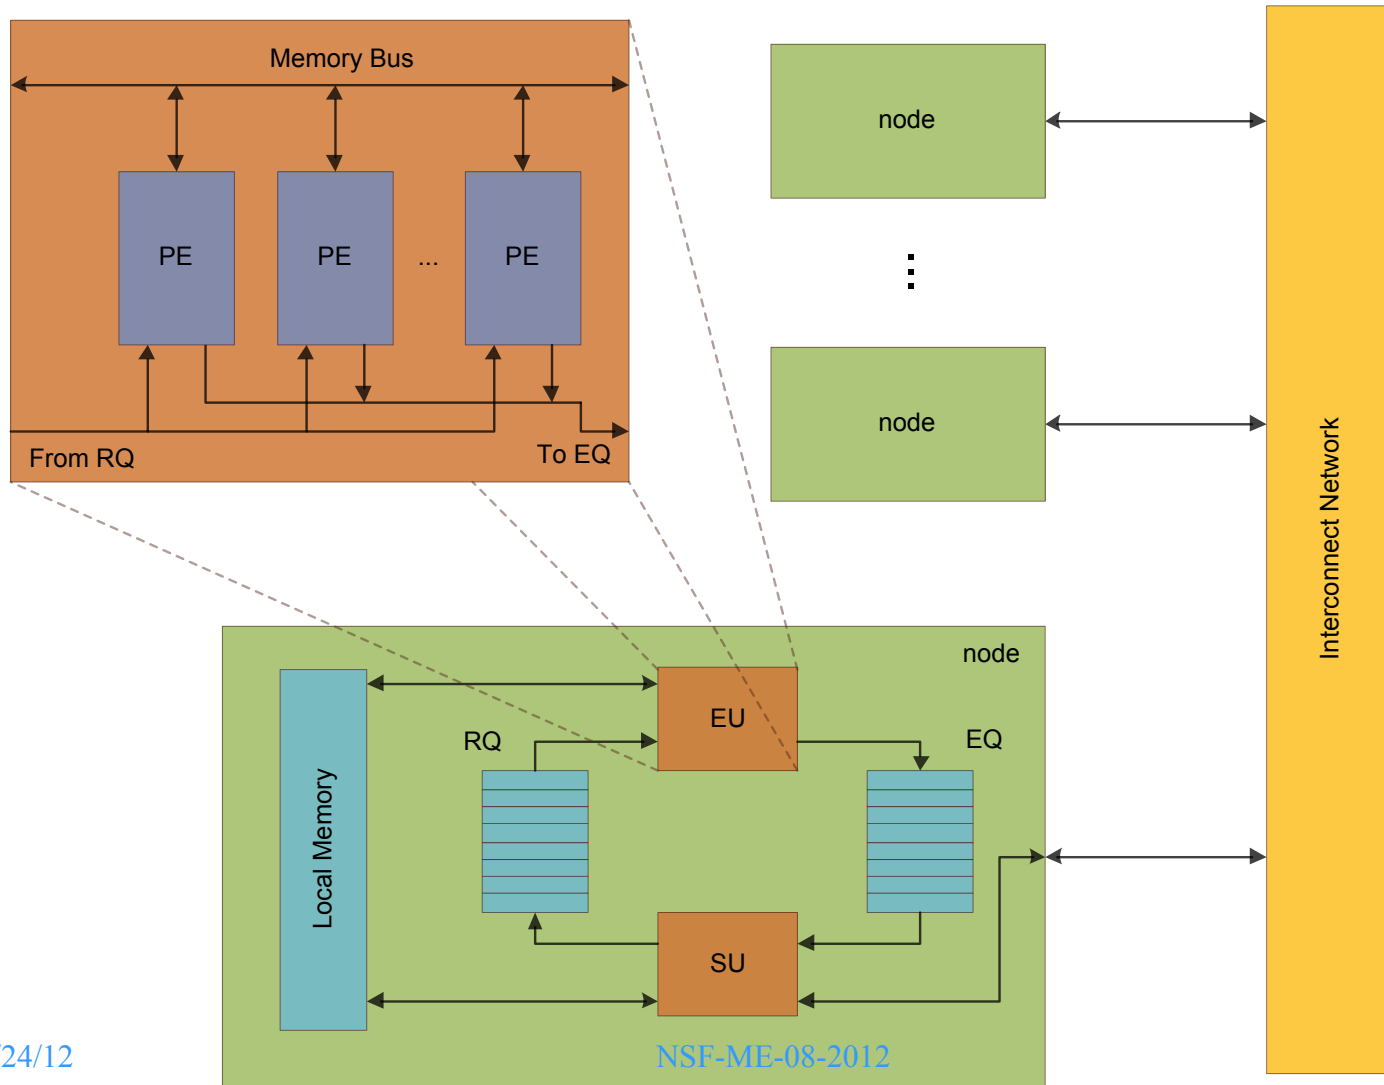
# Execution Model and Abstract Machines

# Abstract Machine Models May Be Heterogeneous!







## Execution Model and Abstract Machines


# EARTH Architecture




# The **EARTH** Multithreaded Execution Model (1993 – 200x)

Two Level of Fine-Grain Threads:  
 - threaded procedures  
 - **fibers**

-  **fiber** within a frame
-  **Aync.** function invocation
-  A sync operation
-  Invoke a threaded func



Fibers



2-level of threads

# Outline

- Background and motivation
- Program execution models
- Evolution of codelet based execution models
  - The EARTH project (1994 – 2004)
  - **IBM Cyclops-64 project (2004 – 2010+ ): The TNT Experience**
  - Intel-led UHPC/Runnemedede (2010 – 2012): The codelet concept and SWARM
- Semantics of the codelet model
- Conclusions and Future Directions



# Outline


- Background and motivation
- Program execution models
- Evolution of codelet based execution models
  - The EARTH project (1994 – 2004)
  - IBM Cyclops-64 project (2004 – 2010+ ): The TNT Experience
  - **Intel-led UHPC/Runnemedede (2010 – 2012): The codelet concept and SWARM**
- Semantics of the codelet model
- Conclusions and Future Directions

# Outline

- **Background and motivation**
- **Program execution models**
- **Evolution of codelet based execution models**
  - The EARTH project (1994 – 2004)
  - IBM Cyclops-64 project (2004 – 2010+ ): The TNT Experience
  - Intel-led UHPC/Runnemedes (2010 – 2012): The codelet concept and SWARM
  - DOE X-Stack (2012-2015): Continue the codelet path
- **Semantics of Codelet Models**
- **Conclusions and Future Directions**

Courtesy: Prof. Jack Dennis, 2001


# The Codelet: A Fine-Grain Piece of Computing



Supports Massively Parallel Computation!

Courtesy: Prof. Jack Dennis, 2001

# The Codelet: A Fine-Grain Piece of Computing



**This Looks Like Data Flow!!**

# Concept of Codelet

(Feb. 4<sup>th</sup>, 2011)

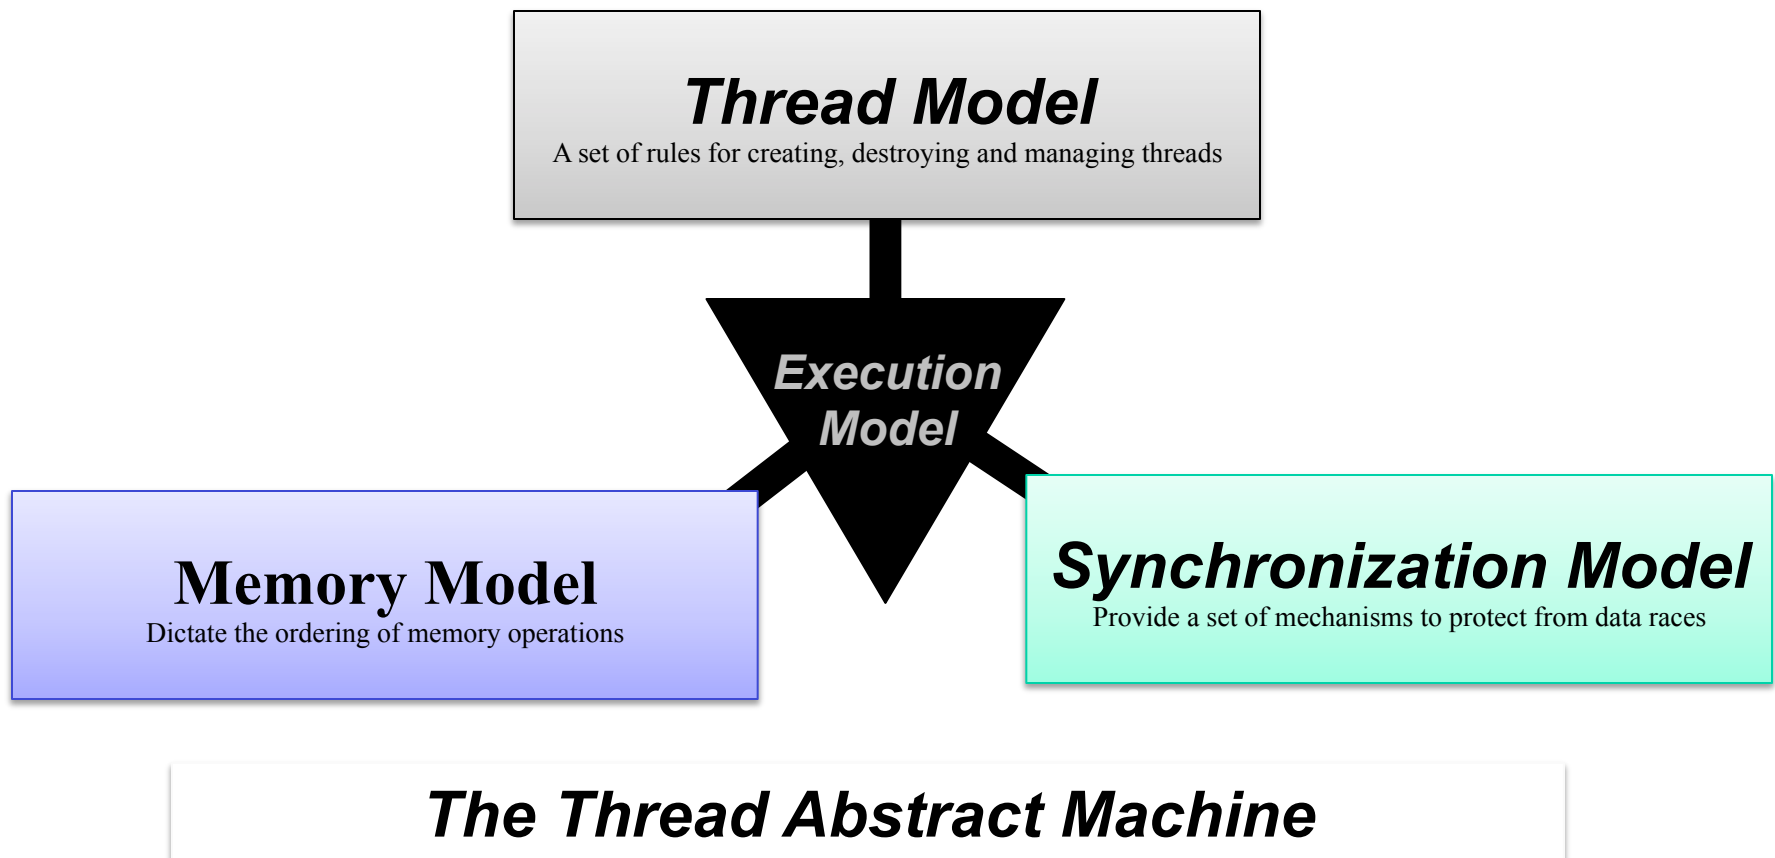
- Codelets are the principal *scheduling quantum* under our codelet based execution model. A codelet, once allocated and scheduled, will be kept usefully busy - since it is *non-preemptive*
- The underline hardware architecture and system software (e.g. compiler, etc.) are optimized to ensure such *non-preemption features* can be productively utilized.

# Outline

- Background and motivation
- Program execution models
- Codelet based execution models
  - The EARTH project (1994 – 2004)
  - IBM Cyclops-64 project (2004 – 2010+ ): The TNT Experience
  - Intel-led UHPC/Runnemedede (2010 – 2012): The codelet concept and SWARM
- Memory semantics of codelet models
- Conclusions and Future Directions



# What is A Shared Memory Execution Model?



# **“Memory Coherence”**

## **A Basic Assumption of SC-Derived Memory Models**

“...All writes to the same location are serialized in some order and are performed in that order with respect to any processor...”

[Gharacharloo Et Al 90]

# **Can We Break The Memory Coherence Barrier ?**

**No ?**

**Yes ?**

# Four Key Question on Memory Models

- What happens when two (or more) concurrent load/store operations happen (arrives) at the same memory location?
- Answers ?

# A Conjecture

The LC (Location Consistency) memory model belongs to the group of memory models that is *weakest* while still do not violate the *causality constraint*!

# Outline

- **Background and motivation**
- **Program execution models**
- **Evolution of codelet based execution models**
  - The EARTH project (1994 – 2004)
  - IBM Cyclops-64 project (2004 – 2010+ ): The TNT Experience
  - Intel-led UHPC/Runnemedede (2010 – 2012): The codelet concept and SWARM
- **The memory semantics of codelets**
- **Conclusions and Future Directions**



# DOE X-Stack Project

## July 2012 – June 2015

### ***Traleika Glacier***

*(Team Lead: Intel)*

*Universities: UIUC, **UD**, UCSD, Rice U)*

*Other Industries (**ETI**, Reservoir)*

*DOE Labs: (PNNL, Sandia, ORNL, ..)*

# Acknowledgements

- Our Sponsors
- Members of CAPSL
- Members of ETI
- Other Collaborators (T. Sterling, V. Sarkar, etc.)
- My Mentor - Prof. Jack B. Dennis
- My Host