Introduction to the Argonne Training Program on Extreme-Scale Computing (ATPESC)

Paul Messina

Director of Science

Argonne Leadership Computing Facility

Argonne National Laboratory

Outline

- Welcome
- A few words about Argonne National Laboratory
- Motivation of the ATPESC
- The curriculum
- Logistics

Argonne – a part of DOE National Laboratory System

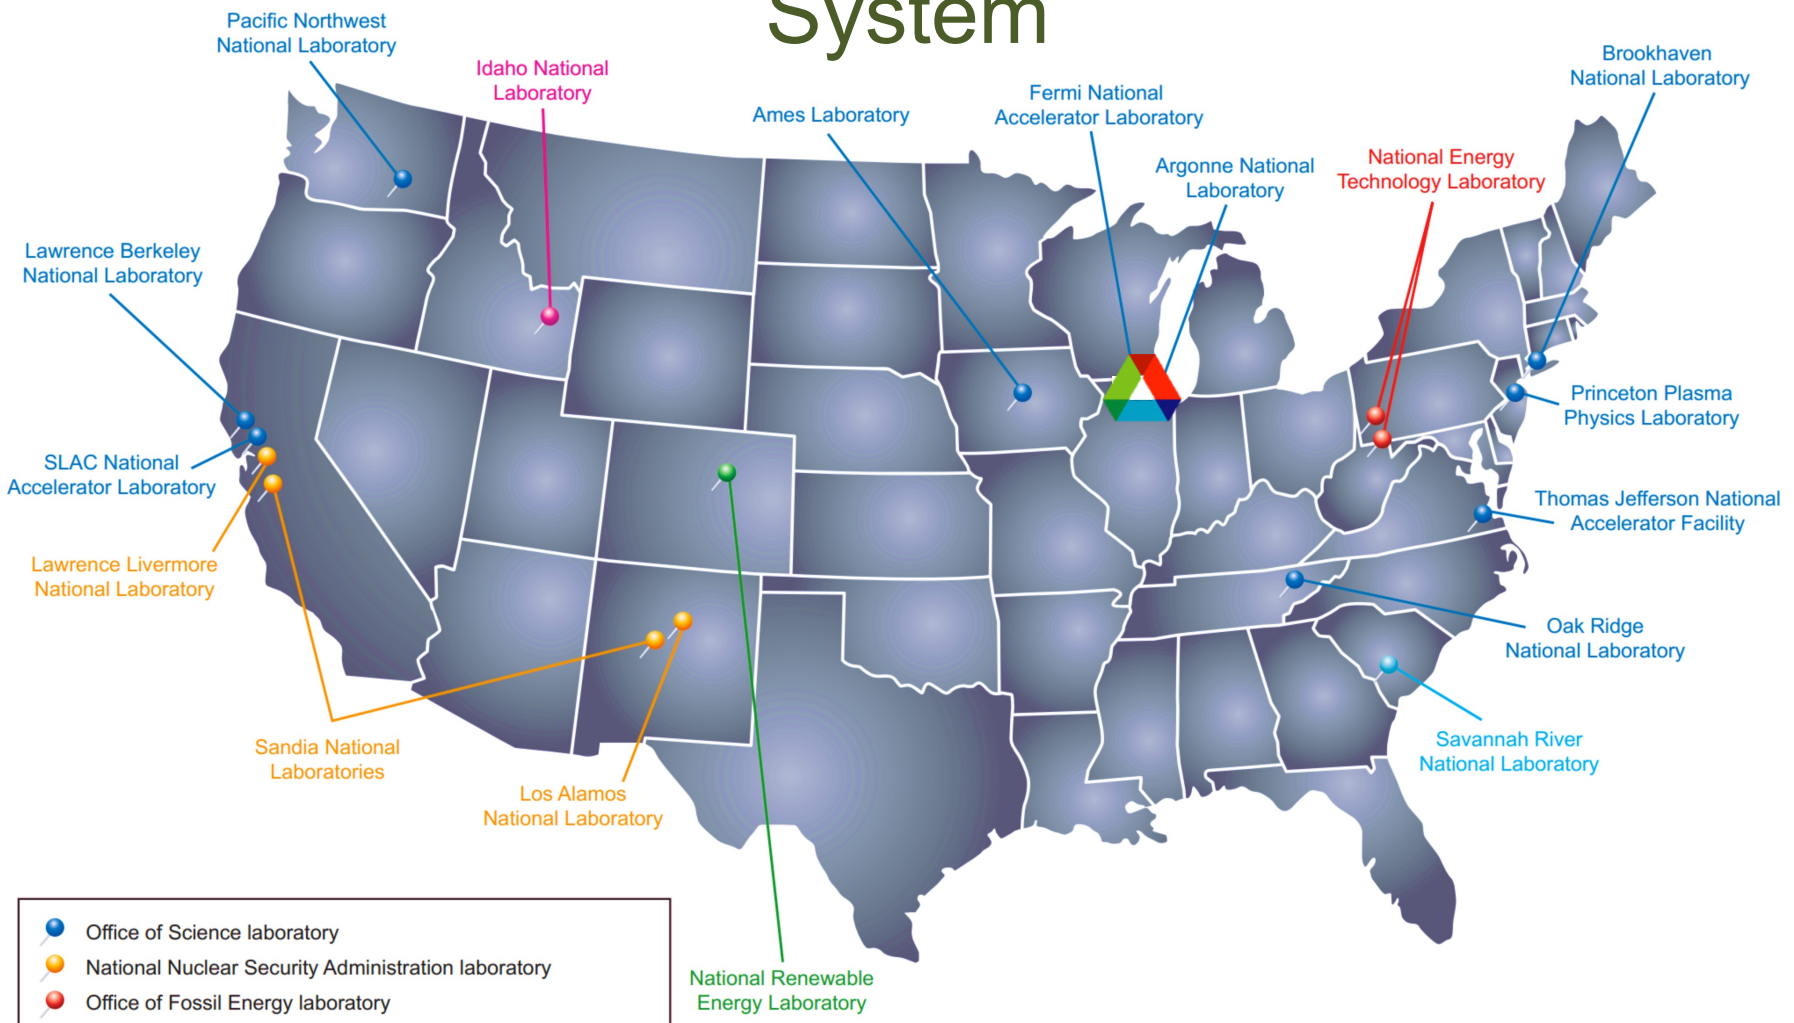

U.S. DEPARTMENT OF ENERGY

Argonne's mission: To provide science-based solutions to pressing global challenges

Energy Science

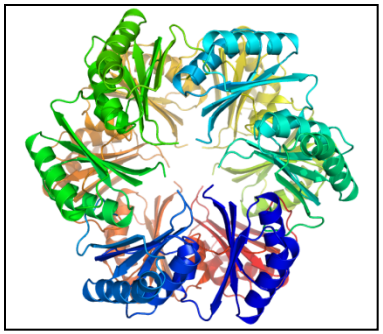

Environmental Sustainability

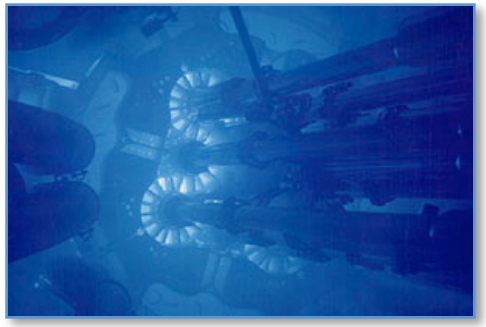

Nuclear and National Security

Use-Inspired Science and Engineering…

…Discovery and Transformational Science and Engineering

Major User Facilities

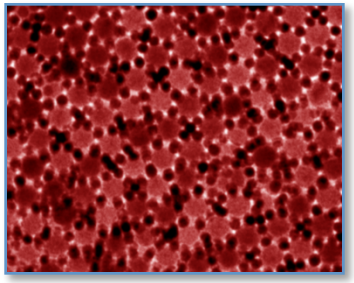

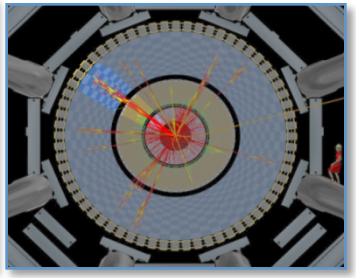

Science and Technology Programs

AVIDAC: Argonne's Version of the Institute's Digital Arithmetic Computer: 1949-1953

"Moll" Flanders, Director
Jeffrey Chu, Chief Engineer

- **AVIDAC: based on prototype at the Institute for Advanced Study in Princeton**
- **Margaret Butler wrote AVIDAC's interpretive floating-point arithmetic system**
  - Memory access time: 15 microsec
  - Addition: 10 microsec
  - Multiplication: 1 millisec
- **AVIDAC press release: 100,000 times as fast as a trained "Computer" using a desk calculator**

Early work on computer architecture

Margaret Butler helped assemble the ORACLE computer with ORNL Engineer Rudolph Klein. In 1953, ORACLE was the world's fastest computer, multiplying 12-digit numbers in .0005 seconds (2Kop/s). Designed at Argonne, it was constructed at Oak Ridge.

Motivation for the ATPESC

- **Today's most powerful supercomputers have complex hardware architectures and software environments**
  - and even greater complexity is on the horizon from next-generation and exascale systems –
- **The scientific and engineering applications that are tackled with these systems are themselves complex**
- **There is a critical need for specialized, in-depth training for the computational scientists poised to facilitate breakthrough science and engineering using these systems**

The DOE Leadership Computing Facility

- Collaborative, multi-lab, DOE/SC initiative ranked top national priority in ***Facilities for the Future of Science: A Twenty-Year Outlook.***
- Mission: Provide the computational and data science resources required to solve the most important scientific & engineering problems in the world.
- Highly competitive user allocation program (INCITE, ALCC).
- Projects receive 100x more hours than at other generally available centers.
- LCF centers partner with users to enable science & engineering breakthroughs (Liaisons, Catalysts).

OAK RIDGE National Laboratory

Argonne NATIONAL LABORATORY

Leadership Computing Facility systems

	Argonne LCF	Oak Ridge LCF
System	IBM Blue Gene/Q	Cray XK7
Name	Mira	Titan
Compute nodes	49,152	18,688
Node architecture	PowerPC, 16 cores	AMD Opteron, 16 cores NVIDIA K20x (Kepler) GPU
Processing Units	786,432 Cores	299,008 x86 Cores + 18,688 GPUs
Memory per node, (gigabytes)	16	32 + 6
Peak performance, (petaflops)	10	27

10 Petaflops Blue Bene/Q - Mira

- ***Mira*** **– BG/Q system**
  - 49,152 nodes / 786,432 cores
  - 786 TB of memory
  - Peak flop rate: 10 PetaFLOPs
  - 3,145,728 hardware threads
- ***Vesta (T&D)*** **- BG/Q system**
  - 2,048 nodes / 32,768 cores
  - 32 TB of memory
  - Peak flop rate: 420 TF

- ***Tukey*** **– Nvidia system**
  - 100 nodes / 1600 x86 cores/ 200 M2070 GPUs
  - 6.4 TB x86 memory / 1.2 TB GPU memory
  - Peak flop rate: 220 TF
- **Storage**
  - Scratch: 28.8 PB raw capacity, 240 GB/s bw (GPFS)
  - Home: 1.8 PB raw capacity, 45 GB/s bw (GPFS)
  - Storage upgrade planned in 2015

Mira's Ecosystem

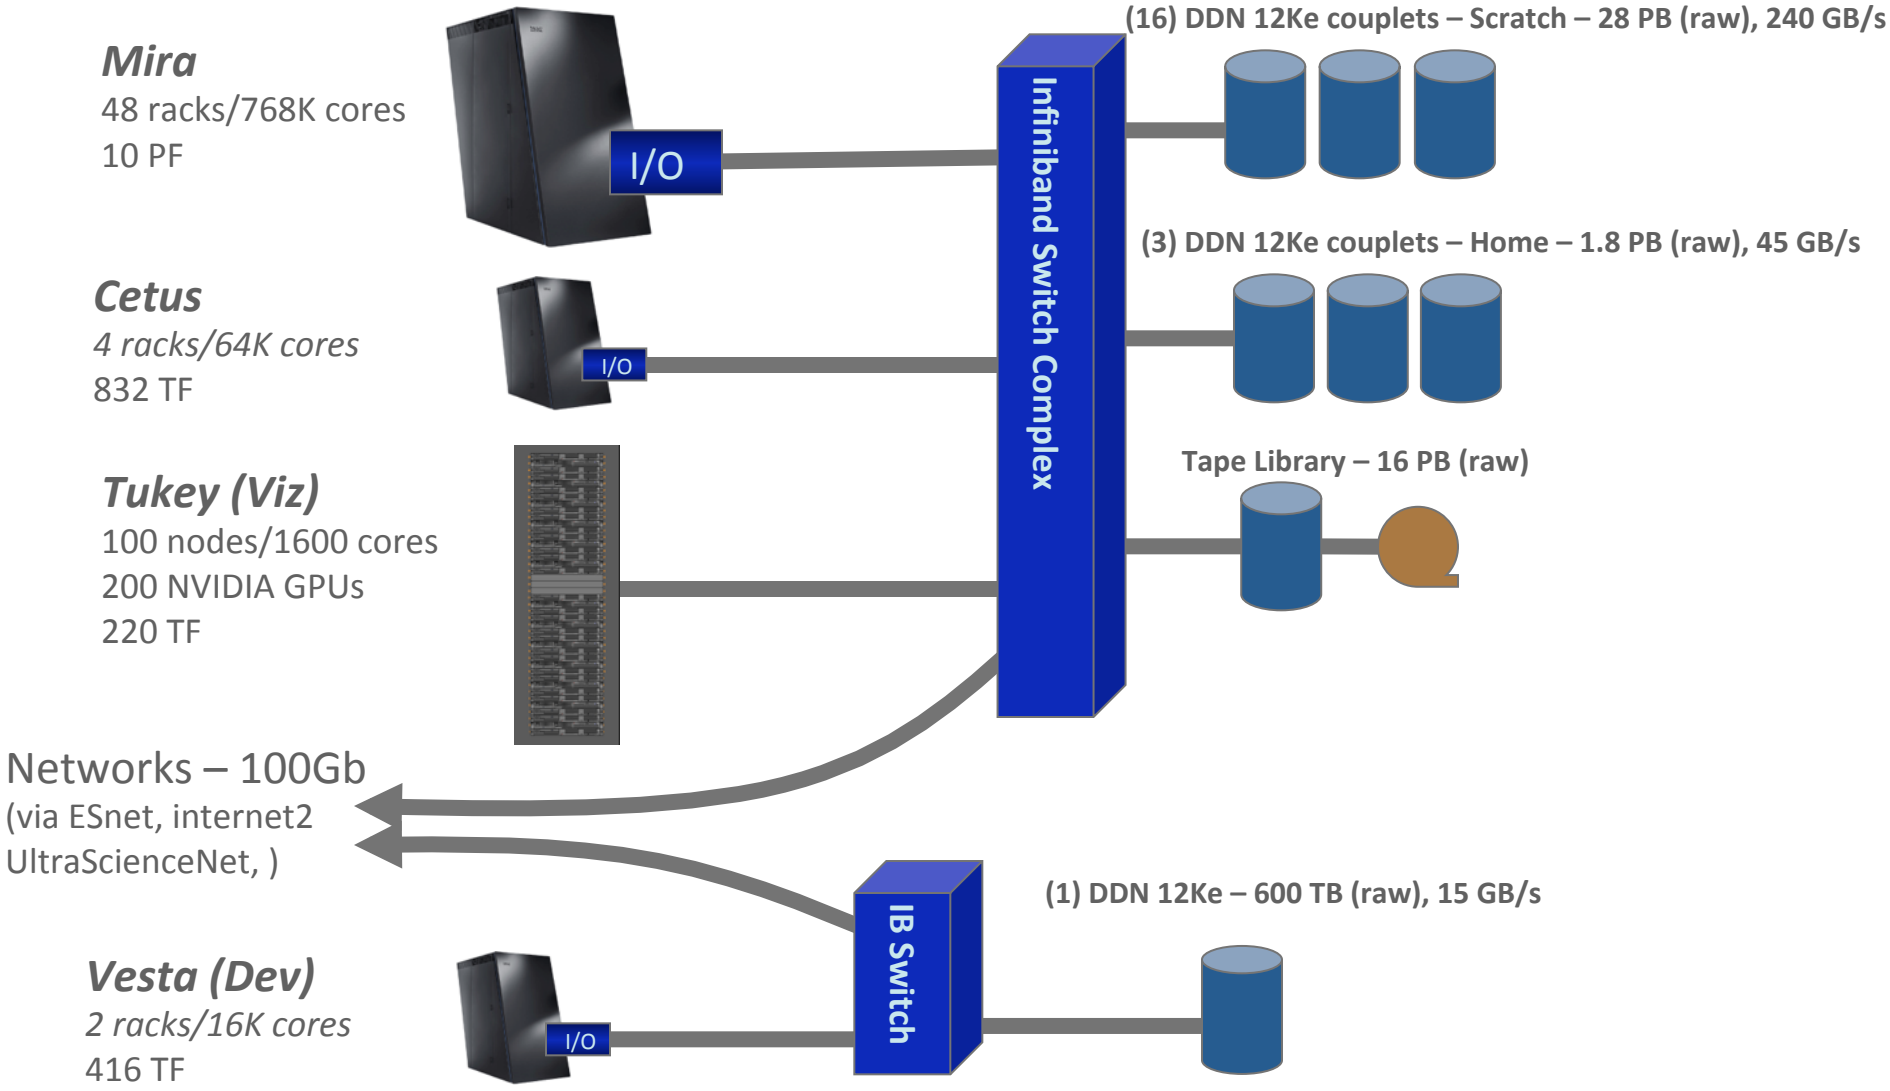

BlueGene/Q Compute Chip

System-on-a-Chip design : integrates processors, memory and networking logic into a single chip

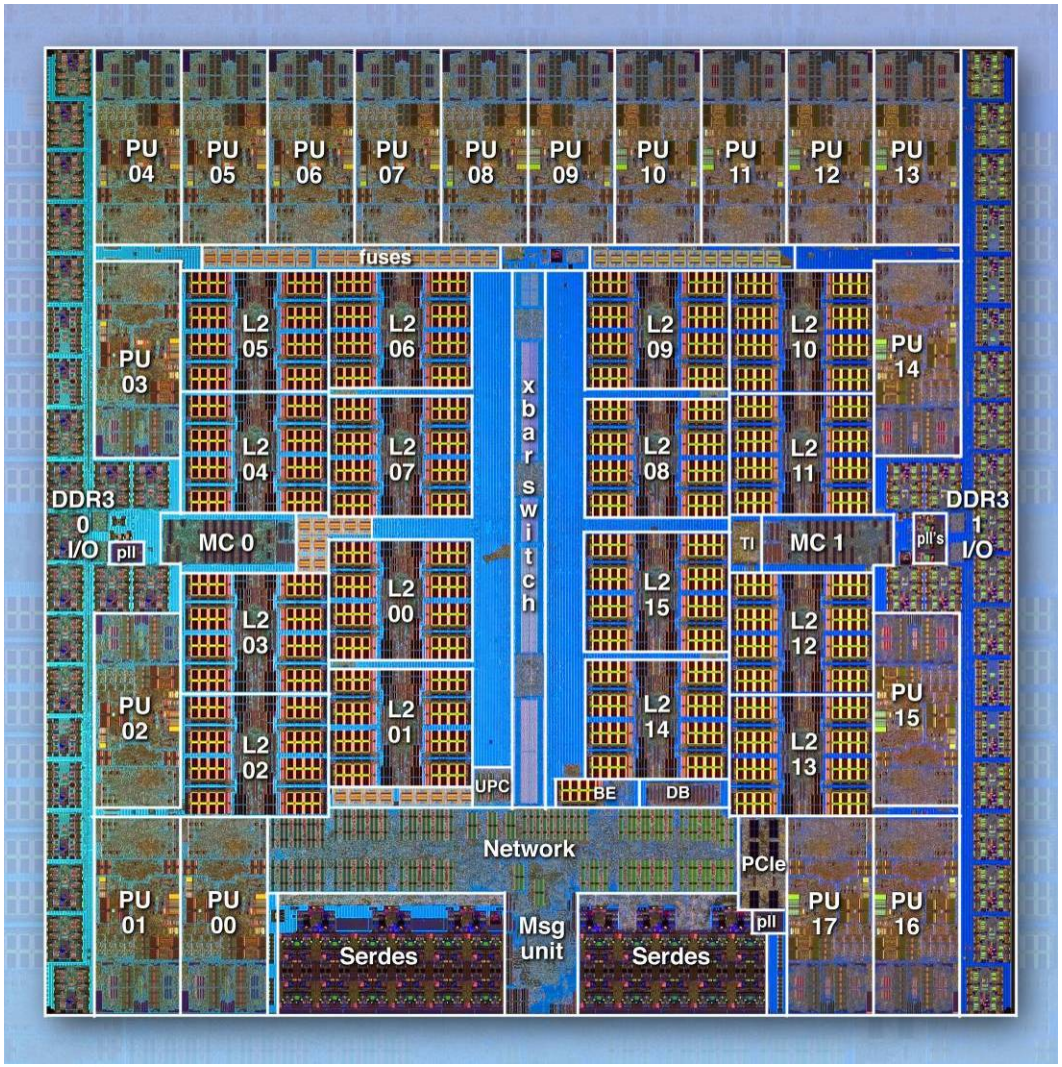

- **360 mm² Cu-45 technology (SOI)**
  - ~ 1.47 B transistors
- **16 user + 1 service processors**
  - plus 1 redundant processor
  - all processors are symmetric
  - L1 I/D cache = 16kB/16kB
  - L1 prefetch engines
- **Crossbar switch**
  - Connects cores via L1P to L2 slices
- **Central shared L2 cache**
  - 32 MB eDRAM
  - 16 slices
- **Dual memory controller**
  - 16 GB external DDR3 memory
  - 42.6 GB/s
- **Chip-to-chip networking**
  - Router logic integrated into BQC chip
  - DMA, remote put/get, collective operations
  - 11 network ports
- **External IO**
  - PCIe Gen2 interface

CORAL Joint NNSA & SC Leadership Computing Acquisition Project

Current DOE Leadership Computers

Titan (ORNL)

Sequoia (LLNL)

Mira (ANL)

Objective - Procure 3 leadership computers to be sited at ANL, ORNL, and LLNL in CY17-18

Leadership Computers run the most demanding DOE mission applications and advance HPC technologies to assure continued US/DOE leadership

Approach

Competitive process - one RFP (issued by LLNL) leading to 3 computer procurement contracts

For risk reduction and to meet a broad set of requirements,
2 architectural paths will be selected

Each lab manages and negotiates its own computer procurement contract, and may exercise options to meet their specific needs

Understanding that **long procurement lead time may impact architectural characteristics and designs** of procured computers

ALCF-3 Project Description

- **Acquire and deploy a 75 – 200PF Leadership Computing system for the Department of Energy's Office of Science Leadership Computing Facility**
  - Procure the compute and storage system
  - Prepare site for system power and cooling needs
  - Select and prepare Early Science Program code teams for the system
  - Prepare for general users
  - Deploy, test, and accept the system integrated in the ALCF environment

NERSC's Cori* System: on the path to exascale

**Gerty Cori was the first American woman to receive the Nobel Prize in Science*

- **Gerty Cori was the first American woman to receive the Nobel Prize in Science (bio-chemistry)**
- **Cori will be a Cray XC system whose nodes are based on Intel's next-generation Intel® Xeon Phi™ processor –- code-named "Knights Landing"**
  - on-package high bandwidth memory
  - 3 teraFLOPS of double-precision peak performance per single socket node
- **The Knights Landing processor will have 72cores, each with multiple hardware threads**
- **Cori is cheduled for delivery in mid-2016**

Curriculum tracks/sessions and their leaders

- Architectures – Pete Beckman
- Programming models and languages – Rusty Lusk and Rajeev Thakur
- Numerical algorithms and software -- Lois McInnes and Lori Diachin
- Toolkits and frameworks – Kalyan Kumaran and Scott Parker
- Visualization and data analysis – Mike Papka and Joe Insley
- Data-intensive computing and I/O – Rob Ross and Rob Latham
- Community codes and software engineering – Katherine Riley and Anshu Dubey

Dinner talks

- **Purpose: present additional topics that will probably be relevant to your research at some point in your career – but in any case interesting**
- **Nine dinner talks**

Yes, the ATPESC is an intense program

- **Many lectures every day, followed by evening hands-on sessions**
- **Ideally we would cover all topics in more depth but the result would be a six-week program**
  - But few people's schedules would allow them to participate
- **Note the 8:30 a.m. starting time, dinner at 5:30 p.m. right after the end of the afternoon lectures, evening sessions**
- **Slides will be posted online as soon as available**

Thank you, DOE Office of Advanced Scientific Computing Research (ASCR)

- **This training program was made possible by funding from the Research Division of the Advanced Scientific Computing Research program of the Department of Energy's (DOE) Office of Science**
- **The funding is for three years**
  - The training program will be offered again in 2015
  - And we may ask for funding for 2016ff
- **Help us improve the training program**
  - Track evaluations
  - Overall program evaluation
  - Conversations or emails to any of us

Surveys

- **Help us improve the training program**
  - Track evaluations
  - Overall program evaluation
  - Conversations or emails to any of us
- **Please fill out the online evaluation surveys on each track and the overall program**
  - at the end of each track, you will receive an email from Chel@alcf.anl.gov with a link to that track's evaluation
  - Respond by the morning of the next day to be eligible for the prize raffle
  - Chel Lancaster is coordinating the evaluations and will be available to answer questions or help

Suggestions from last year's surveys adopted this year

- **Suggested reading before the ATPESC**
- **Participant introductions**
  - Details during dinner

Paul Messina

- **Position: Ph.D. student, applied mathematics, University of CIncinnati**
- **Research background:**
  - Solution of elliptic PDEs with singularities
  - Mathematical software
- **Research interests:**
  - Parallel computer architectures
- **Personal interests**
  - Sailing
- **Personal background:**
  - Had lived in four countries by the time I was 14 years old

Logistics

- **All lectures and hands-on sessions in Rembrandt room**
- **All meals in Utrillo room on ground floor**
  - Lunch and dinner presentations will be in this room
- **Other rooms in ground floor might be used as needed**
- **Wi-fi SSID in this building is Argonne**
- **Password is argonne14**

Diagram of Meeting Rooms: Second Floor

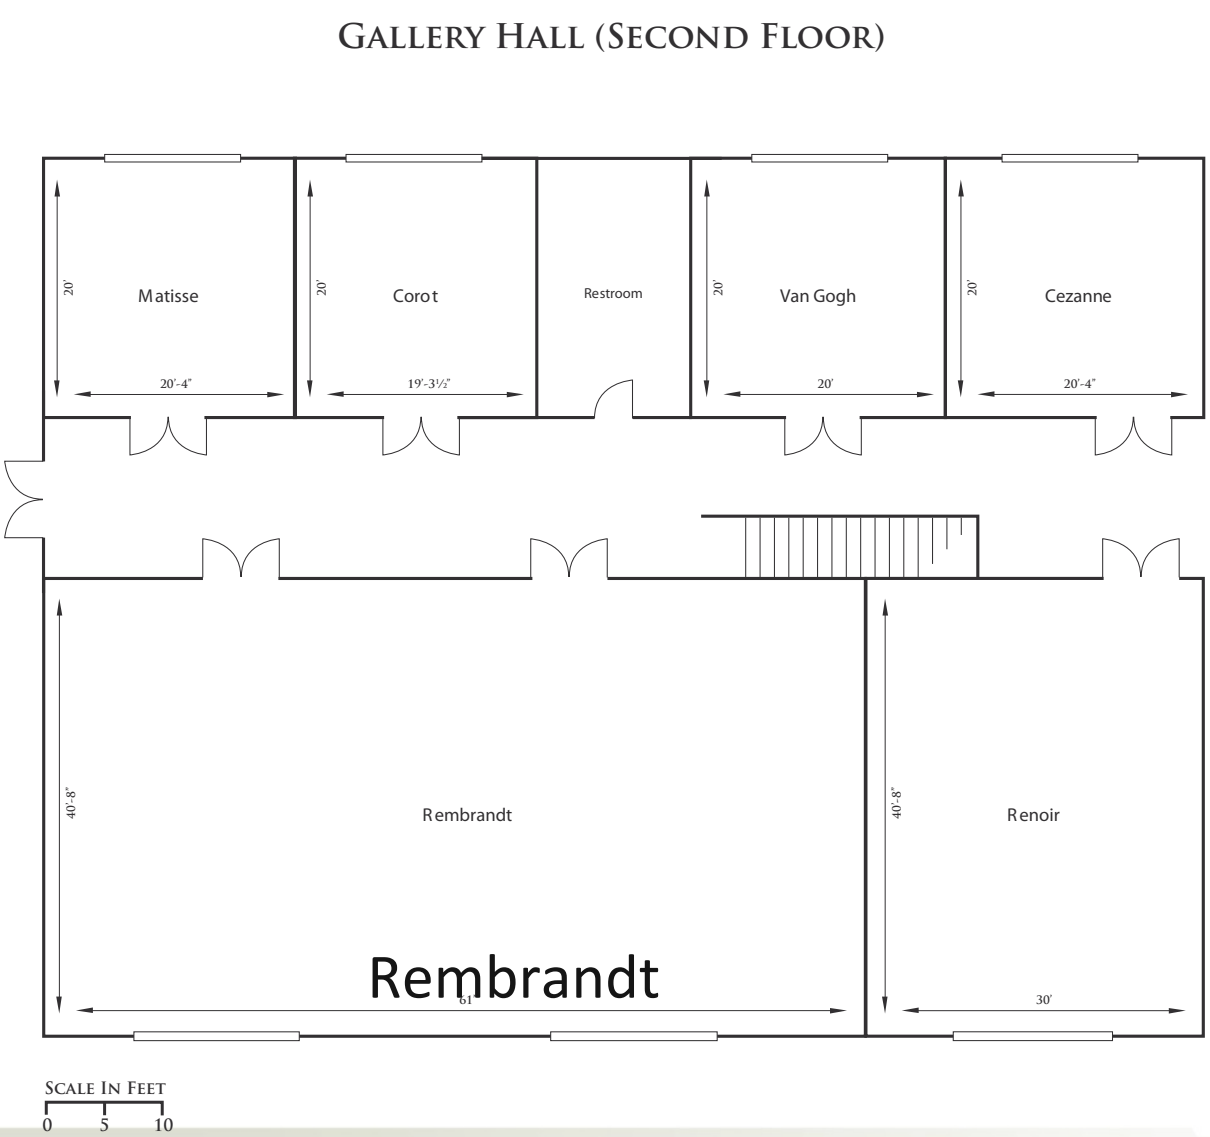

Diagram of meeting rooms: Ground floor

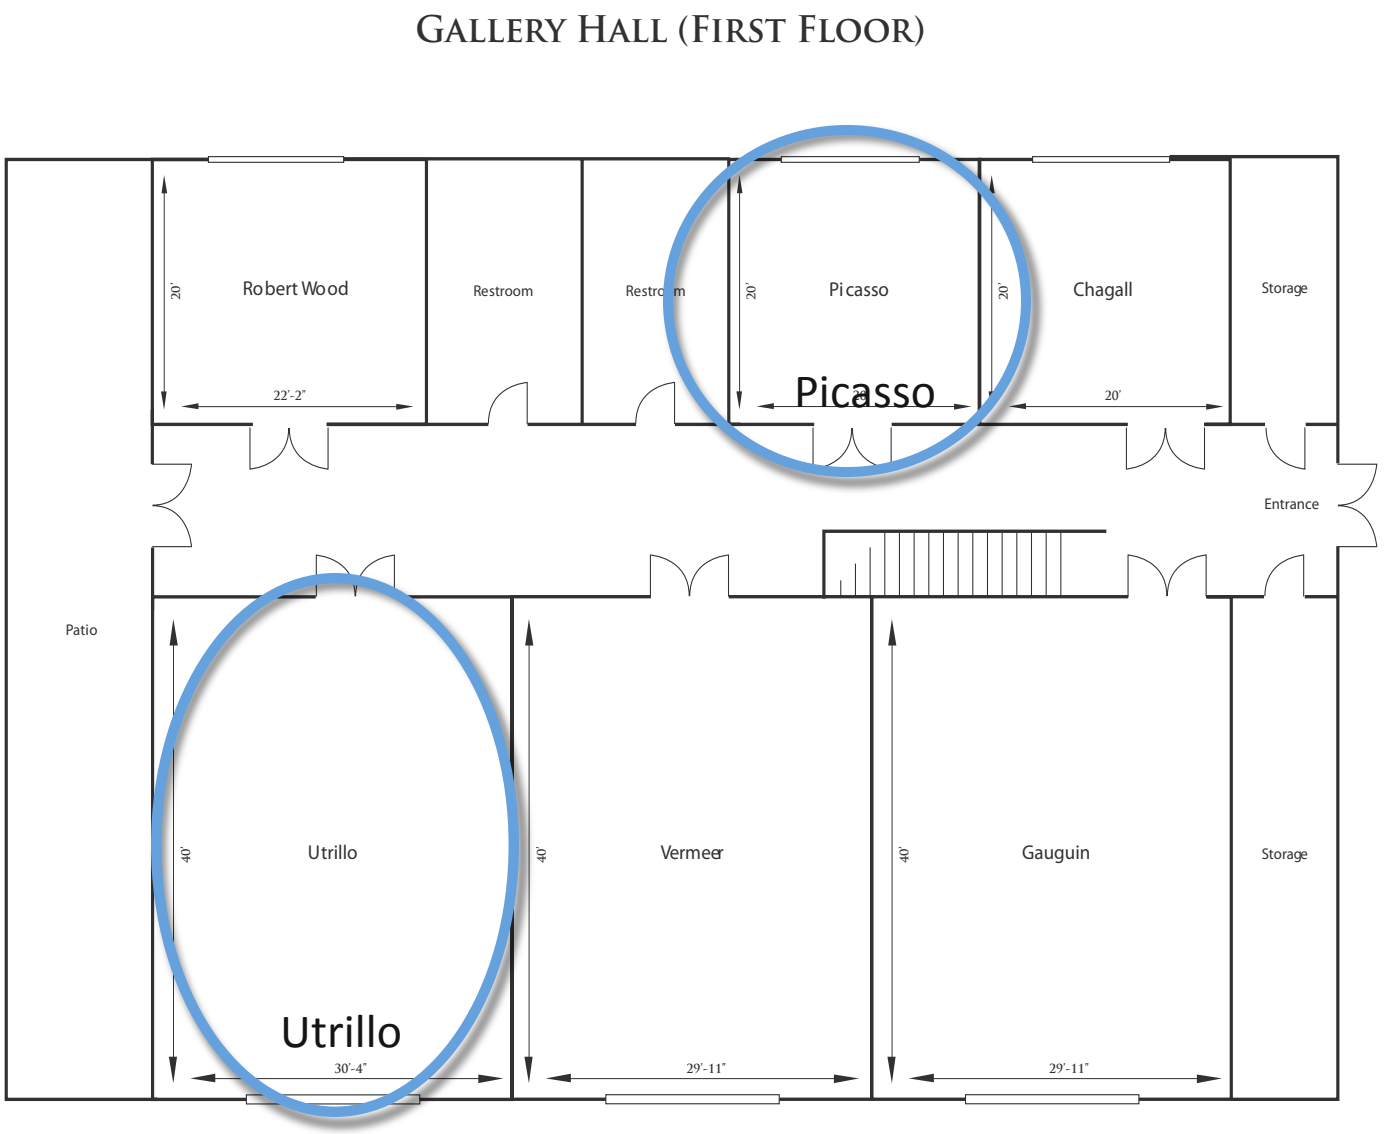

Whom to ask for help

- Local arrangements
  - Cheryl Zidel
  - Ashley Boyle
  - Ginny Doyle
  - (sometimes they will be in Picasso room on ground floor)
- Surveys
  - Chel Lancaster
- Computing issues
  - Ray Loy
  - Robert Scott
  - Adam Scovel
  - Others TBD

Summary

- Thanks in advance to all of you for taking two weeks of your summer to participate in this program
- Questions?