

ASCAC MEETING



40TH ANNIVERSARY OF DOE

PARTNERSHIP BETWEEN NNSA AND SC IN HPC

PAUL MESSINA
Argonne National Laboratory

September 26, 2017
Arlington, VA

OUTLINE

- Algorithms and software libraries
- Programming languages and models
- Design and use of advanced architecture computers
- Data and related topics
- Networking and distributed computing
- Other collaborations

ALGORITHMS AND SOFTWARE LIBRARIES

- Mathematical algorithms and libraries (Barry Smith will cover)
 - Historically a major contribution of ASCR predecessors, work carried out at SC and NNSA labs
- TOOLPACK
- POOMA
- Visit -- initially developed at LLNL as part of ASCI, now developed by ORNL and LBNL, as well as many other institutions and has been supported by SciDAC
- FastBit (John Wu will cover)
- Categorization of methods: seven dwarves
- etc., etc.

PROGRAMMING LANGUAGES AND MODELS

- The Language Working Group
 - The LWG was formed November 1976 by the DOE-established Advanced Computing Committee to recommend a common programming language for the national laboratories.
 - Motivated by the Cray-1 at Los Alamos and at NCAR
 - The LWG work involved close collaboration with individuals from ANL, LBNL, LANL, NCAR, NMFEC, ORNL, SNL
 - Had substantial impact on Fortran 88 standard
- 1984: “We shall synthesize the information gathered from users over the past two years and develop specific recommendations for vendors, standards committees, and other interested implementors of programming languages for parallel computers.”

PROGRAMMING LANGUAGES AND MODELS

- Staff from the national labs have been active members of standards committees:
 - E.g., Fortran, PL/1, IEEE 754 arithmetic
- Design and implementation of languages and models for early experimentation at universities and national laboratories
 - CrOS message passing routines => Express
 - p4
 - Sisal
 - ...
- PVM (developed at ORNL, widely used)
- MPI (Rusty Lusk will cover), OpenMP, OpenACC, ...

DESIGN AND USE OF ADVANCED ARCHITECTURE COMPUTERS

- In the early 1980s AMS funded Caltech, UIUC, and NYU Courant to develop parallel architectures and software and applications to use them
 - NYU Ultracomputer
 - UIUC CEDAR
 - Caltech Cosmic Cube
- *“AMS acquired the first experimental parallel machine in the DOE (the Denelcor HEP at LANL) and it supported the design and construction of the first hypercube (the Cosmic Cube at Caltech), as well as other novel computer architectures such as the Ultracomputer at NYU and the CEDAR cluster at UIUC.”* MICS report c. 1995
- NNSA labs had similar activities, e.g., PuPS at LANL
- Many interactions and collaborations among SC and NNSA labs (and universities), including providing access to each others' computers, developing software, short courses

ARGONNE RESEARCH COMPUTING FACILITY

Later renamed Advanced Computing Research Facility

One of several grass-roots efforts to experiment with parallel architectures

Excerpt from the 1984 FTP for the Argonne MCS Division:

SPECIAL PROJECTS

The Research Computing Facility project, which began in FY 1982, has as its goal the configuration and operation of state-of-the-art computing systems tailored to support our work and to facilitate collaborative research among AMS-funded groups. Hardware and software are evaluated for suitability to serve our research needs and, where desirable, acquired and installed. Emphasis is given to exploring the use of innovative systems that have the potential of being useful to other DOE programs as well as our own.

CO-DESIGN IN 1984

One example: Argonne MCS division FTP

B. ADVANCED COMPUTING RESEARCH

Summary: We propose a major thrust in advanced computing research that emphasizes the interaction between algorithms, the software environment, and advanced computer architectures, for both numerical and reasoning tasks. This initiative will feature two activities: one an Experimental Computing Facility (ECF) supporting our research activities and operating as a national user facility, the other an expanded research program building on existing expertise in software engineering, computational mathematics, and applied analysis.

The ECF will select and operate computers with innovative designs that are likely eventually to be effective for a wide range of computational and reasoning tasks. Their use as experimental research tools in a national user facility will provide the understanding needed to realize their performance potential and to design software environments for eventual production use. The systems chosen will initially have hardware and software usable for

DESIGN AND USE OF ADVANCED ARCHITECTURE COMPUTERS

- Argonne's Advanced Computing Research Facility
 - Funded by ASCR (then called AMS) but made available to NNSA labs
- Quotes from Argonne MCS Division April 1984 FTP
 - “A new program has also been initiated in advanced scientific computing; an Experimental Computing Facility has been established and a broad research program begun, with a focus on parallel architectures.”
 - “The Research Computing Facility continues as a special project. This facility provides a computing environment tailored to mathematical and computer science research and collaborative studies among AMS-funded groups.”

HIGH PERFORMANCE COMPUTING RESEARCH CENTERS HPCRC 1991

- Transition to production HPC on advanced systems, research on global climate science, materials, groundwater transport
- Advanced Computing Laboratory (LANL)
- Center for Computational Sciences (ORNL)
- NERSC (LLNL) formerly known as MFECC, formerly known as CTRCC

DESIGN AND USE OF ADVANCED ARCHITECTURE COMPUTERS

- Intel Touchstone Delta/Concurrent SuperComputing Consortium (CSCC); involved other institutions but important funding and participation by Argonne, LANL, LLNL, PNNL, SNL
- By the early 1990s, most labs had collections of systems with advanced architectures, mostly for research in algorithms and computer science but also experimentation with full applications, e.g.,
 - Ames Laboratory Scalable Computing Laboratory
 - Argonne High-Performance Research Facility
 - LANL Advanced Computing Laboratory
 - ORNL Center for Computational Sciences
 - Sandia Massively Parallel Computing Research Laboratory

DESIGN AND USE OF ADVANCED ARCHITECTURE COMPUTERS

- A recurring activity at labs is partnership with vendors to advance technologies, both hardware and software
- A few examples:
 - Sandia-Cray partnership for Red Storm
 - LLNL-ANL-IBM partnership for Blue Gene
 - ASCI
 - ECP

DATA AND RELATED TOPICS

- HiPPI
- HPSS (LANL, NERSC, ORNL, LLNL, SNL, IBM) widely used, Buddy Bland will cover
- Scalable I/O project (Caltech, ANL, LANL, LLNL, SNL, Princeton, Syracuse), led to ROMIO and MPI-IO, among other developments
- LANL High Performance Data Storage project
- LLNL National Storage Laboratory: a collaborative project created to remove storage system bottlenecks, provide much-needed storage functionality, and facilitate the establishment of national storage system standards.
- ADIOS (Matt Wolf will cover)

NETWORKING AND DISTRIBUTED COMPUTING

- CASA Gigabit Network Testbed (included in MICS report, as a highlight; LANL, JPL, SDSC, Caltech)
 - World record 500 Mbits/s (TCP/IP), 792 Mbits/s raw HiPPI transfers on a 2,000 kilometer link
 - Much more important were the applications and the software environment that were developed in the CASA project
- MFEnet, Esnet
 - 1994 contract with Sprint for Asynchronous Transfer Mode (ATM) over SONET services at 622 Mbits/s
- I-way
- Distributed Collaboratory Experiment Environments (DCEE) at LBNL
 - Collaboratory testbeds included ANL, LLNL, ORNL, PNNL

OVERVIEW OF THE I-WAY: WIDE-AREA VISUAL SUPERCOMPUTING (1995)

Thomas A. DeFanti, Ian Foster, Michael E. Papka, Rick Stevens, Tim Kuhfuss

- This paper discusses the I-WAY project and provides an overview of the papers in this issue of IJSA. The I-WAY is an experimental environment for building distributed virtual reality applications and for **exploring issues of distributed wide-area resource management and scheduling**. The goal of the I-WAY project is to enable researchers to use multiple internetworked supercomputers and advanced visualization systems to conduct very large scale computations. By connecting 12 ATM testbeds, 17 super computer centers, 5 virtual reality research sites, and over 60 applications groups, the I-WAY project has created an extremely diverse wide-area environment for exploring advanced applications. This environment has provided a glimpse of the future for advanced scientific and engineering computing.

OTHER COLLABORATIONS

- ASCI
 - Started supporting CSGF in 1999
 - Support for some MPI activities (via labs)
 - Visualization R&D
 - Etc.
- CORAL and LBNL-LANL-SNL procurements
- Workshops on exascale applications, software, technologies, etc. , etc.
- ECP

SUMMARY

- The activities mentioned in this talk are but a sample of the collaborative HPC activities supported by SC and NNSA
 - For example, the myriad training and human resources activities at the labs, in addition to the CSGF program
- Most of them were self-organizing, based on common vision and interests, perception of mutual benefits
- Apologies for my oversights or factual errors
 - I had limited time to prepare this talk

THANK YOU!