



**Argonne**  
NATIONAL  
LABORATORY

*... for a brighter future*

***ALCF***

*Argonne Leadership  
Computing Facility*




U.S. Department  
of Energy

UChicago ►  
Argonne<sub>LLC</sub>



A U.S. Department of Energy laboratory  
managed by UChicago Argonne, LLC



# Argonne Leadership Computing Facility

## *Overview*

*May 2008*



# Over 20 years of Advanced Systems for DOE and Others

## ■ ACRF period [1983-1992]

- DOE's founding ACRF
- Explored many parallel architectures, developed programming models and tools, trained >1000 people

## ■ HPCRC period [1992-1999]

- Production-oriented parallel computing for Grand Challenges in addition to Computer Science.
- Fielded 1<sup>st</sup> IBM SP in DOE



## ■ TeraGrid [2001-present]

- Overall Project Lead
- Defining, deploying and operating the integrated national cyberinfrastructure for NSF
- 9 sites, 22 systems, 200TF

## ■ LCRC [2003-present]

- Lab-wide production supercomputer service
- All research divisions, 56 projects, 380 users

## ■ BlueGene Evaluation [2005-present]

- Founded BlueGene Consortium with IBM
  - 67 institutions, >260 members
  - Applications Workshop Series
  - Systems Software Collaborations

# *ALCF Timeline*

## **2004**

- Formed of the Blue Gene Consortium with IBM
- DOE-SC selected the ORNL, ANL and PNNL teams for Leadership Computing Facility award

## **2005**

- Installed 5 teraflops Blue Gene/L for evaluation

## **2006**

- Began production support of 6 INCITE projects, with BGW
- Joined IBM and LLNL to design and develop next Blue Genes
- “Lehman” Peer Review of ALCF campaign plans

## **2007**

- Increased to 9 INCITE projects; continued development projects
- Install 100 teraflops Blue Gene/P system

## **2008**

- Begin support of 20 INCITE projects on Blue Gene/P
- Add 500T teraflops Blue Gene/P system

# *Mission and Vision for the ALCF*

## **Our Mission**

Provide the computational science community with a world leading computing capability dedicated to breakthrough science and engineering.

## **Our Vision**

A world class center for computation driven scientific discovery that has:

- outstandingly talented people,
- the best collaborations with computer science and applied mathematics,
- the most capable and interesting computers and,
- a true spirit of adventure.

See <http://www.alcf.anl.gov/> for additional information

# Argonne Leadership Computing Facility

*Dedicated to breakthrough science and engineering.*

## ■ Current Blue Gene/L Capabilities

- BGL: 1024 nodes, 2048 cores, 5.7 TF speed, 512GB memory
- Supports development + INCITE

## ■ Current Blue Gene/P Production System

- 111 TF Blue Gene/P system
- Fast PB file system
- Many PB tape archive
- Supports ~20 INCITE projects

## ■ Then for early 2009 production

- 556 TF Blue Gene/P
- 8PB next generation file system
- Supports even more challenging INCITE science projects



Blue Gene/L  
at Argonne

In 2004 DOE selected the ORNL, ANL and PNNL team based on a competitive peer review


- ORNL to deploy a series of Cray X-series systems
- ANL to deploy a series of IBM Blue Gene systems
- PNNL to contribute software technology



Blue Gene/P Engineering Rendition

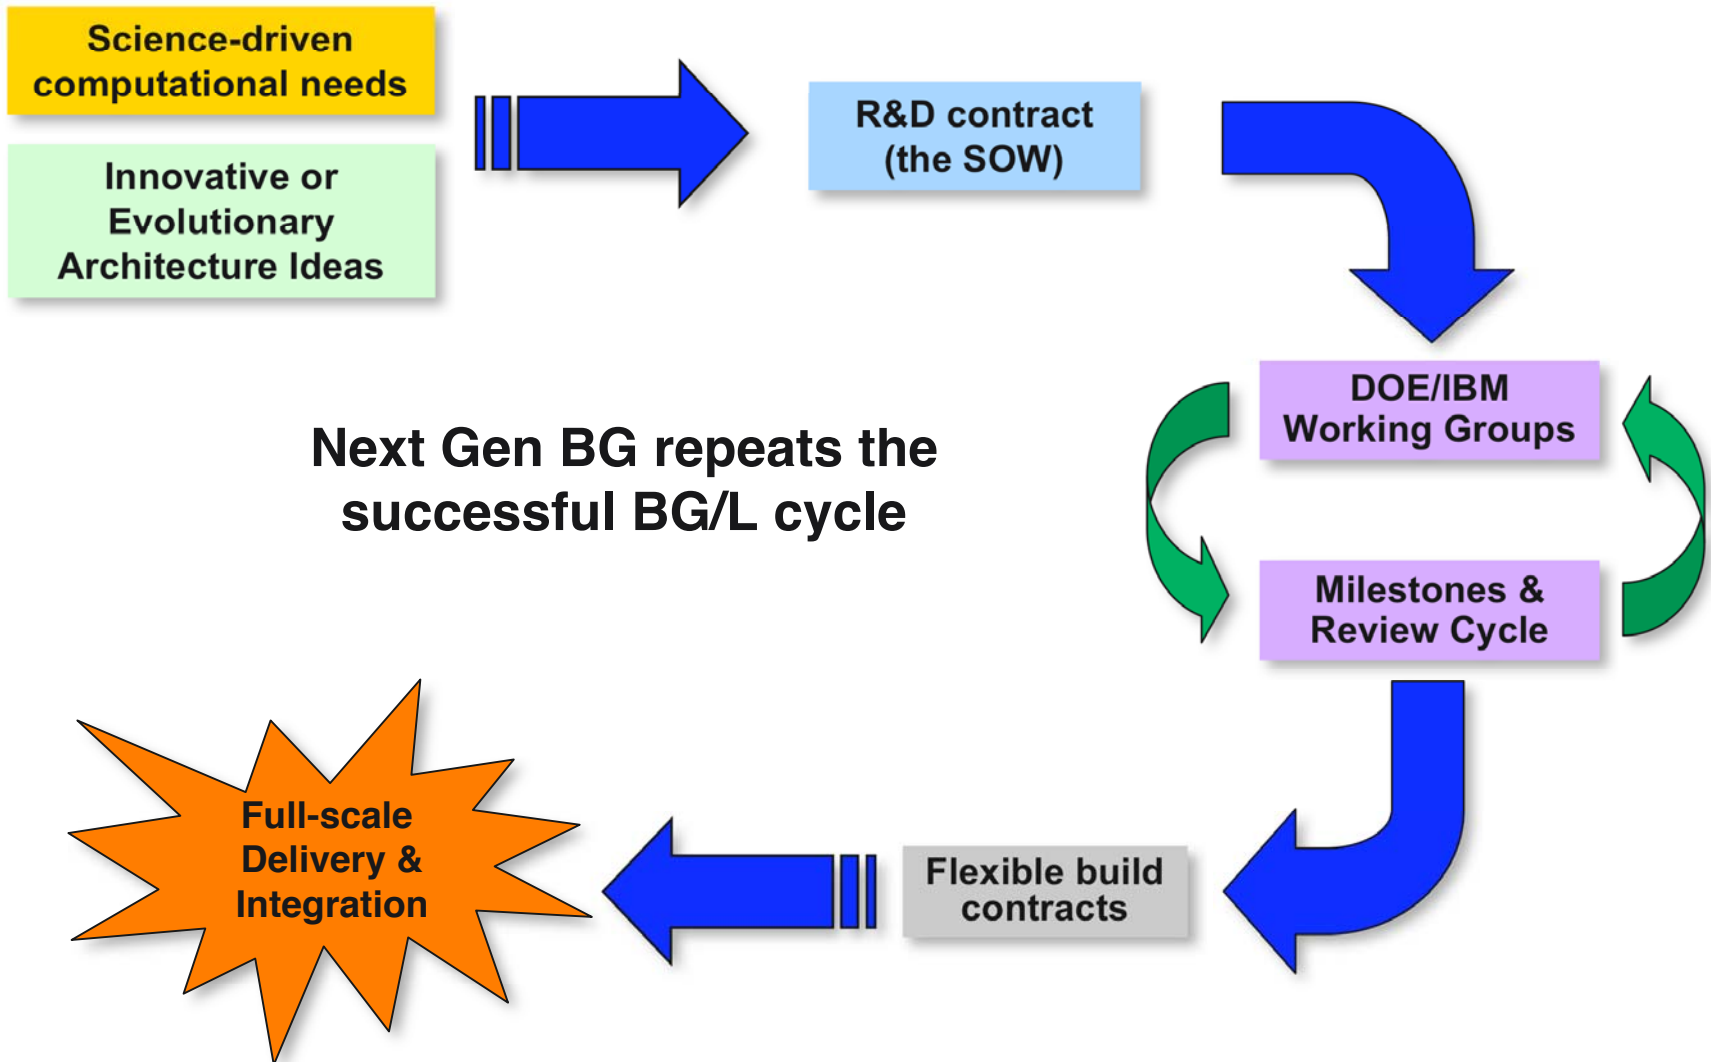
# Blue Gene DNA

- Low power design → massive parallelism
  - The leader in Green Computing
- System on a Chip (SoC)
  - Improves Price / Performance
  - Reduces system complexity & power
- Custom designed ASIC
  - Reducing overall part count, reducing errors
  - Permits tweaking CPU design to reduce soft errors
- Dense packaging
- Fast communication network(s)
- Sophisticated RAS (reliability, availability, and serviceability)
- Dynamic software provisioning and configuration



# DOE / IBM Partnership for Petascale Architectures

DOE NNSA/LLNL, DOE SC/ANL and IBM/Rochester/Watson



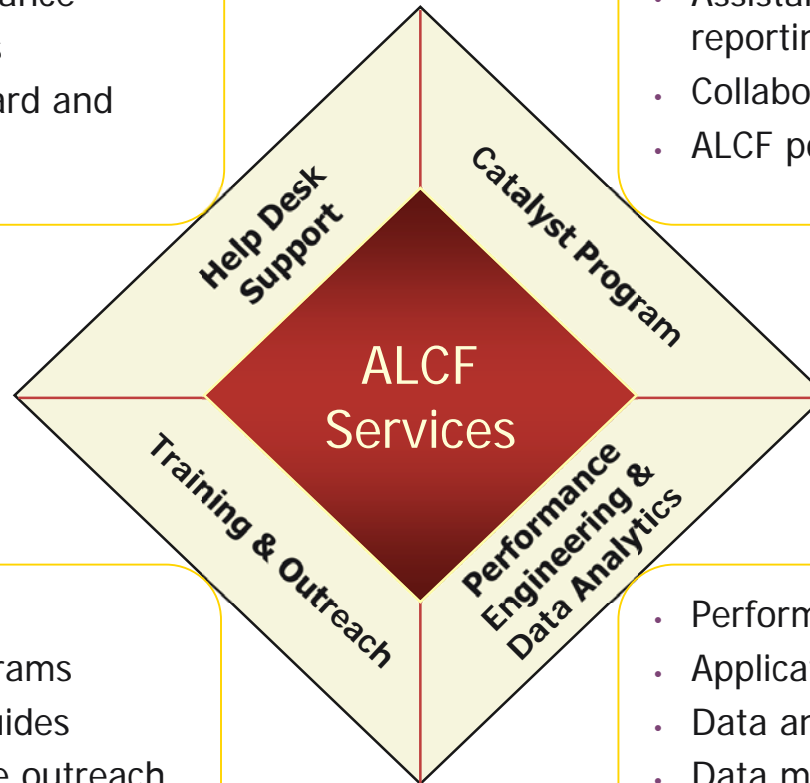
## *Decision to choose Blue Gene is Supported by*

- Applications community has reacted positively and some applications are doing extremely well
  - *For those applications that can make the transition, the BG platform provides outstanding scientific opportunity*
- Blue Gene has been remarkably reliable at scale
  - *The overall reliability/TF appears to be at least an order of magnitude better than other platforms for which we have data*
- Power consumption is 2x better than other platforms
  - *Lower cost of ownership and window to the future of lower power*
- System Cost
  - *The cost of deploying a balanced system is lower than other platforms*

# ALCF Service Offerings

- Startup assistance
- User administration assistance
- Job management services
- Technical support (Standard and Emergency)

- ALCF project management
- Assistance with proposals, planning, reporting
- Collaboration within science domains
- ALCF point of coordination



- Workshops & seminars
- Customized training programs
- On-line content & user guides
- Educational and corporate outreach programs

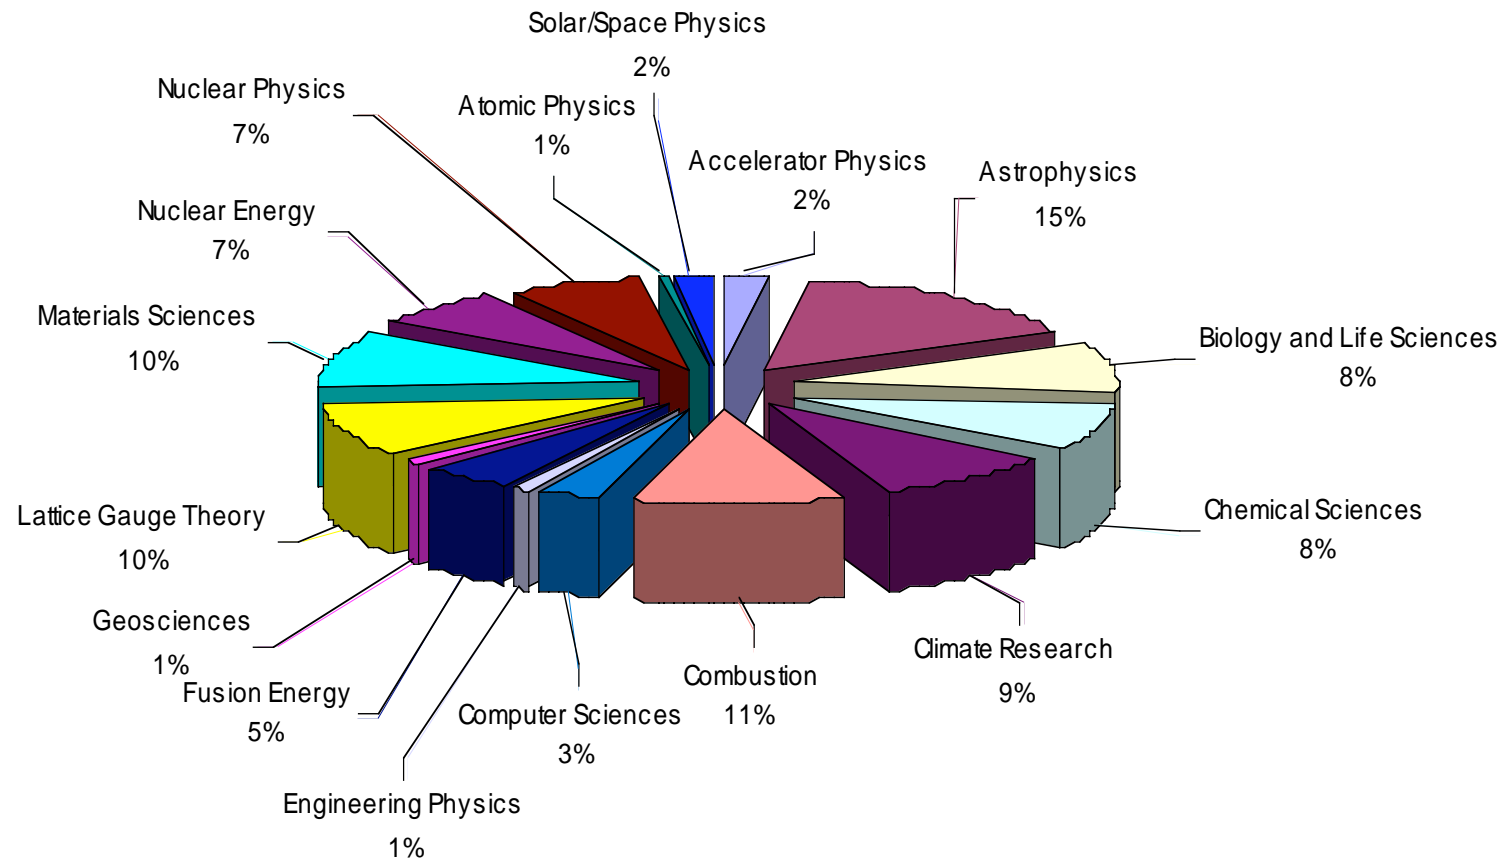
- Performance engineering
- Application tuning
- Data analytics
- Data management services

# *DOE INCITE Program*

## *Innovative and Novel Computational Impact on Theory and Experiment*

- **Solicits large computationally intensive research projects**
  - To enable high-impact scientific advances
- **Open to all scientific researchers and organizations**
  - Scientific Discipline Peer Review
  - Computational Readiness Review
- **Provides large computer time & data storage allocations**
  - To a small number of projects for 1-3 years
  - Academic, Federal Lab and Industry, with DOE or other support
- **Primary vehicle for selecting principal science projects for the Leadership Computing Facilities**

# 2008 INCITE Awards - Allocations by Discipline




# Modeling of Protofibril Structures Provides Insight into Molecular Basis of Parkinson's Disease


Igor Tsigelny, PI, U. California - San Diego



- Parkinson's Disease is the 2nd most common adult neurological disease
- Increased aggregation of *alpha-synuclein* protein is thought to lead to harmful pore-like structures in human membranes
- UCSD - SDSC team used molecular modeling and molecular dynamics simulations in combination with biochemical and ultrastructural analysis to show that *alpha-synuclein* can lead to the formation of pore-like structures on membranes
- Used NAMD and MAPAS on Blue Gene/L at ALCF and SDSC



*alpha-synuclein forming a dimer (above) and a completed pentamer (below) attached to a membrane*





# High-Resolution Protein Structure Prediction

David Baker, Srivatsan Raman, Univ. of Washington



- Proteins are the workhorse molecules of all biological systems
- ROSETTA uses a two-phase Monte Carlo algorithm to sample the extremely large space of possible structures to find the most favorable one
- Below: Electron density (black) was generated from phases determined from two starting models. The crystal structure (red) agrees very well with the phases determined from the Rosetta model.



# Many Routes to Blue Gene/P at ALCF

## ■ Annual INCITE Competition

- 1-3 year allocations
- 5-100M CPU hr/year typical
- 8-160K processors typical

## ■ SC Director's Discretionary

- Address special needs

## ■ Blue Gene Consortium


- Porting and testing allocations
- Develop on Surveyor
- Scaling days on big systems

## ■ Open Source Development

- Development allocations
- Develop and test on Surveyor

## ■ ALCF Staff / Discretionary

- Join the ALCF team !



# *ALCF - ISSF, TCS*





**Argonne**  
NATIONAL  
LABORATORY

*... for a brighter future*

***ALCF***

*Argonne Leadership  
Computing Facility*




U.S. Department  
of Energy

UChicago ►  
Argonne<sub>LLC</sub>



A U.S. Department of Energy laboratory  
managed by UChicago Argonne, LLC



# Argonne Leadership Computing Facility

## *Overview*

*May 2008*



# Over 20 years of Advanced Systems for DOE and Others

## ■ ACRF period [1983-1992]

- DOE's founding ACRF
- Explored many parallel architectures, developed programming models and tools, trained >1000 people

## ■ HPCRC period [1992-1999]

- Production-oriented parallel computing for Grand Challenges in addition to Computer Science.
- Fielded 1<sup>st</sup> IBM SP in DOE



## ■ TeraGrid [2001-present]

- Overall Project Lead
- Defining, deploying and operating the integrated national cyberinfrastructure for NSF
- 9 sites, 22 systems, 200TF

## ■ LCRC [2003-present]

- Lab-wide production supercomputer service
- All research divisions, 56 projects, 380 users

## ■ BlueGene Evaluation [2005-present]

- Founded BlueGene Consortium with IBM
  - 67 institutions, >260 members
  - Applications Workshop Series
  - Systems Software Collaborations

# *ALCF Timeline*

## **2004**

- Formed of the Blue Gene Consortium with IBM
- DOE-SC selected the ORNL, ANL and PNNL teams for Leadership Computing Facility award

## **2005**

- Installed 5 teraflops Blue Gene/L for evaluation

## **2006**

- Began production support of 6 INCITE projects, with BGW
- Joined IBM and LLNL to design and develop next Blue Genes
- “Lehman” Peer Review of ALCF campaign plans

## **2007**

- Increased to 9 INCITE projects; continued development projects
- Install 100 teraflops Blue Gene/P system

## **2008**

- Begin support of 20 INCITE projects on Blue Gene/P
- Add 500T teraflops Blue Gene/P system

# *Mission and Vision for the ALCF*

## **Our Mission**

Provide the computational science community with a world leading computing capability dedicated to breakthrough science and engineering.

## **Our Vision**

A world class center for computation driven scientific discovery that has:

- outstandingly talented people,
- the best collaborations with computer science and applied mathematics,
- the most capable and interesting computers and,
- a true spirit of adventure.

See <http://www.alcf.anl.gov/> for additional information

# Argonne Leadership Computing Facility

*Dedicated to breakthrough science and engineering.*

## ■ Current Blue Gene/L Capabilities

- BGL: 1024 nodes, 2048 cores, 5.7 TF speed, 512GB memory
- Supports development + INCITE

## ■ Current Blue Gene/P Production System

- 111 TF Blue Gene/P system
- Fast PB file system
- Many PB tape archive
- Supports ~20 INCITE projects

## ■ Then for early 2009 production

- 556 TF Blue Gene/P
- 8PB next generation file system
- Supports even more challenging INCITE science projects



Blue Gene/L  
at Argonne

In 2004 DOE selected the ORNL, ANL and PNNL team based on a competitive peer review

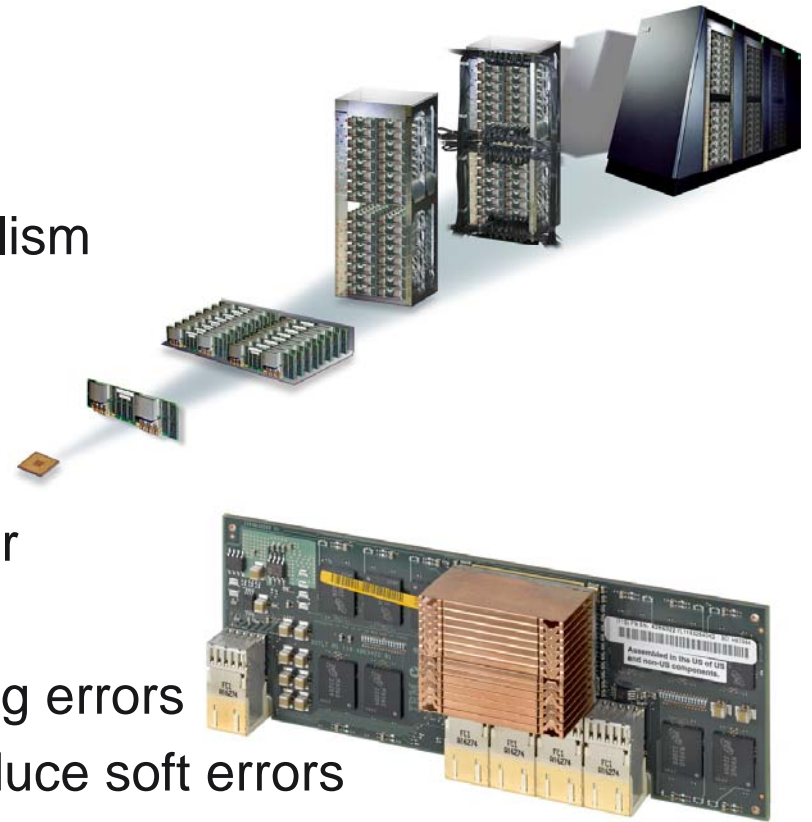
- ORNL to deploy a series of Cray X-series systems
- ANL to deploy a series of IBM Blue Gene systems
- PNNL to contribute software technology



Blue Gene/P Engineering Rendition

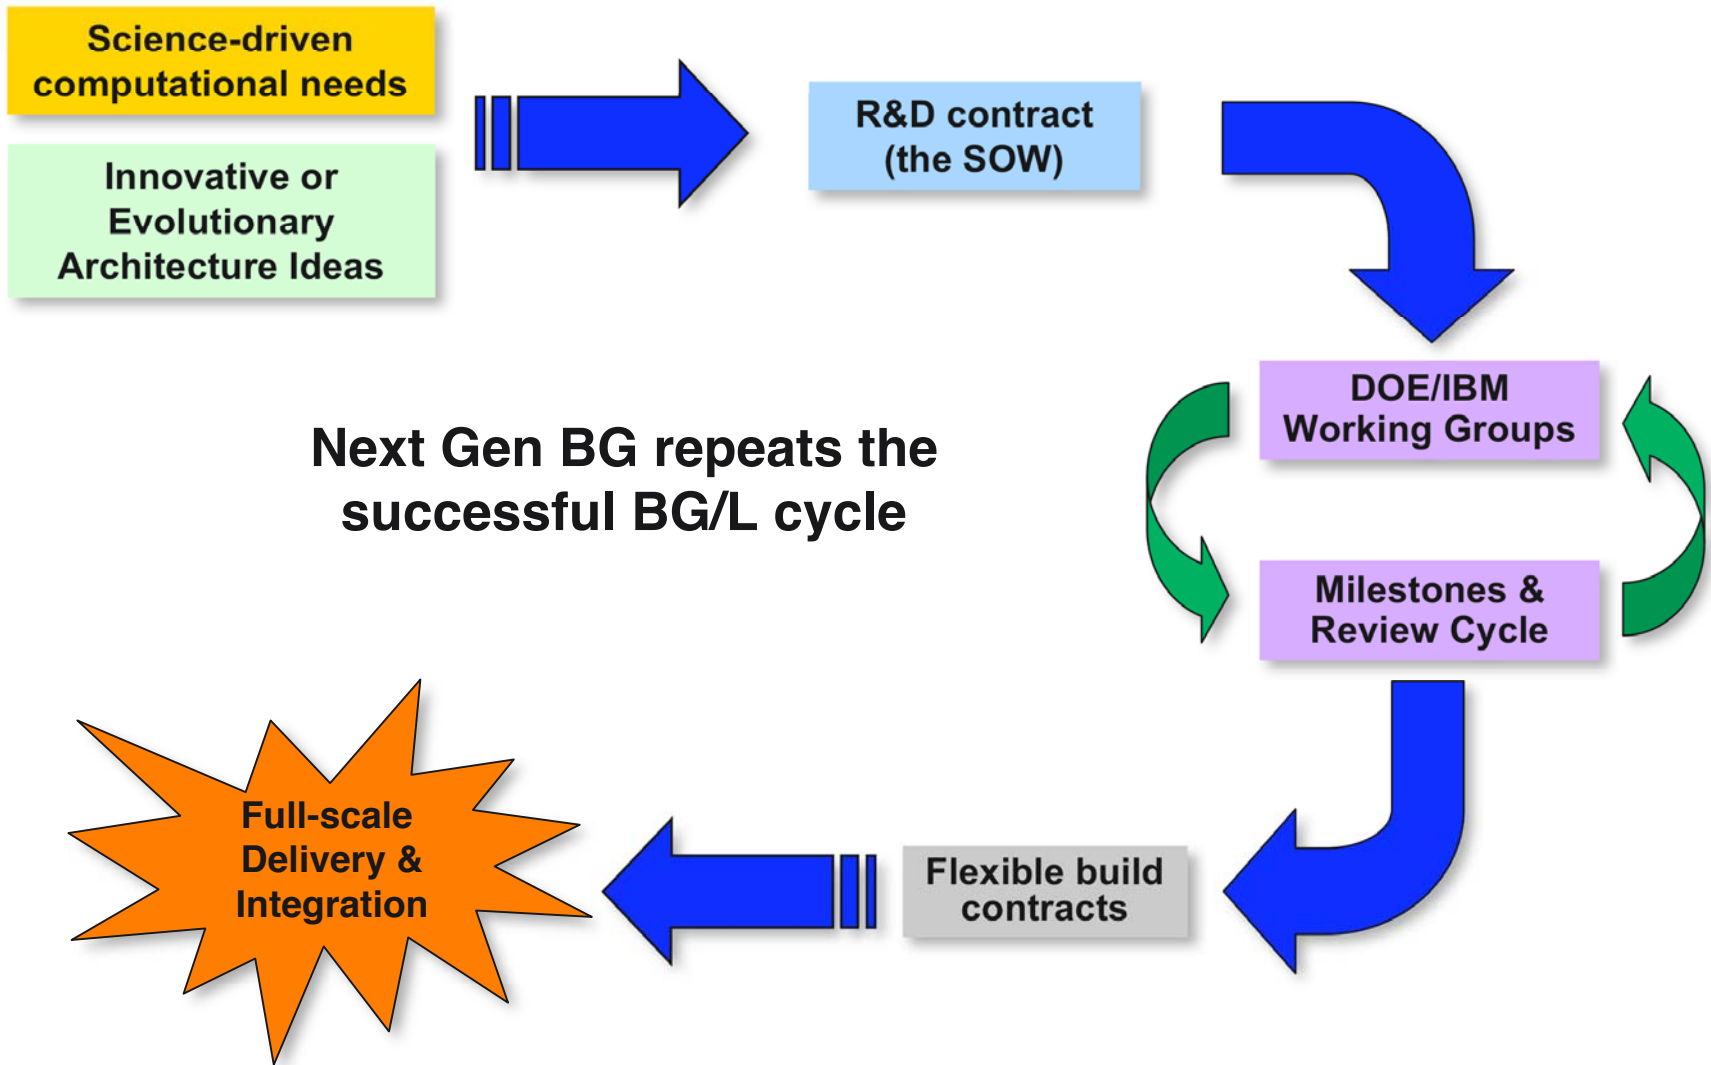
# Blue Gene DNA

- Low power design → massive parallelism
  - The leader in Green Computing
- System on a Chip (SoC)
  - Improves Price / Performance
  - Reduces system complexity & power
- Custom designed ASIC
  - Reducing overall part count, reducing errors
  - Permits tweaking CPU design to reduce soft errors
- Dense packaging
- Fast communication network(s)
- Sophisticated RAS (reliability, availability, and serviceability)
- Dynamic software provisioning and configuration



# DOE / IBM Partnership for Petascale Architectures

DOE NNSA/LLNL, DOE SC/ANL and IBM/Rochester/Watson



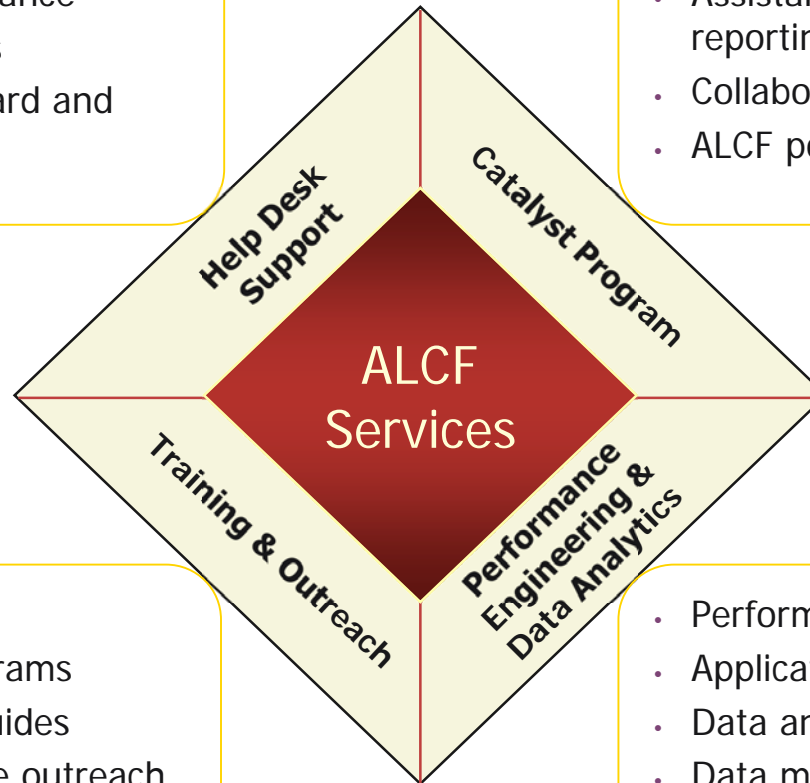
## *Decision to choose Blue Gene is Supported by*

- Applications community has reacted positively and some applications are doing extremely well
  - *For those applications that can make the transition, the BG platform provides outstanding scientific opportunity*
- Blue Gene has been remarkably reliable at scale
  - *The overall reliability/TF appears to be at least an order of magnitude better than other platforms for which we have data*
- Power consumption is 2x better than other platforms
  - *Lower cost of ownership and window to the future of lower power*
- System Cost
  - *The cost of deploying a balanced system is lower than other platforms*

# ALCF Service Offerings

- Startup assistance
- User administration assistance
- Job management services
- Technical support (Standard and Emergency)

- ALCF project management
- Assistance with proposals, planning, reporting
- Collaboration within science domains
- ALCF point of coordination



- Workshops & seminars
- Customized training programs
- On-line content & user guides
- Educational and corporate outreach programs

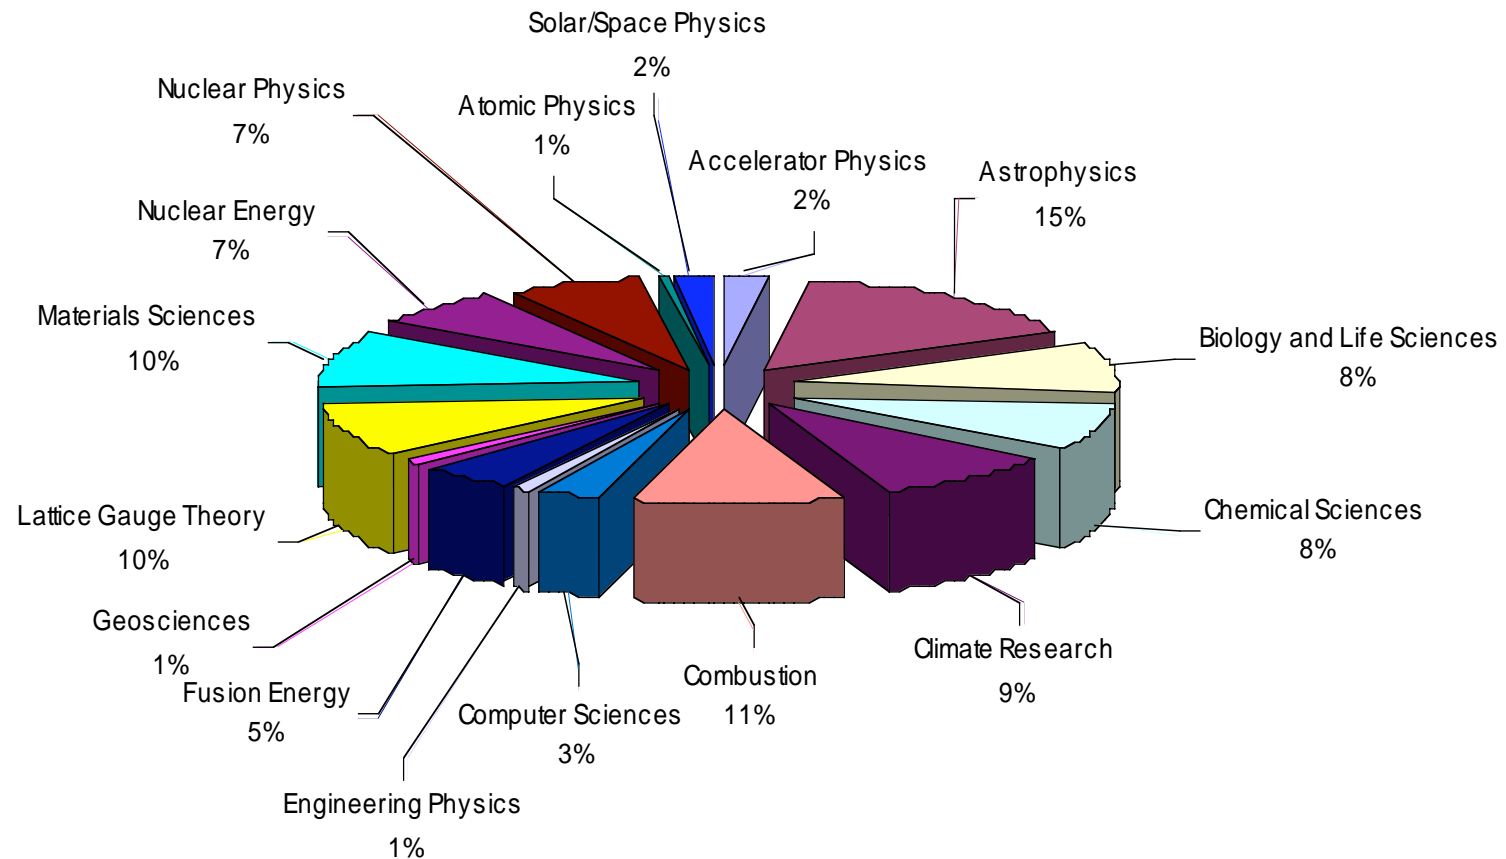
- Performance engineering
- Application tuning
- Data analytics
- Data management services

# *DOE INCITE Program*

## *Innovative and Novel Computational Impact on Theory and Experiment*

- **Solicits large computationally intensive research projects**
  - To enable high-impact scientific advances
- **Open to all scientific researchers and organizations**
  - Scientific Discipline Peer Review
  - Computational Readiness Review
- **Provides large computer time & data storage allocations**
  - To a small number of projects for 1-3 years
  - Academic, Federal Lab and Industry, with DOE or other support
- **Primary vehicle for selecting principal science projects for the Leadership Computing Facilities**

# 2008 INCITE Awards - Allocations by Discipline




# Modeling of Protofibril Structures Provides Insight into Molecular Basis of Parkinson's Disease


Igor Tsigelny, PI, U. California - San Diego



- Parkinson's Disease is the 2nd most common adult neurological disease
- Increased aggregation of *alpha-synuclein* protein is thought to lead to harmful pore-like structures in human membranes
- UCSD - SDSC team used molecular modeling and molecular dynamics simulations in combination with biochemical and ultrastructural analysis to show that *alpha-synuclein* can lead to the formation of pore-like structures on membranes
- Used NAMD and MAPAS on Blue Gene/L at ALCF and SDSC



*alpha-synuclein forming a dimer (above) and a completed pentamer (below) attached to a membrane*

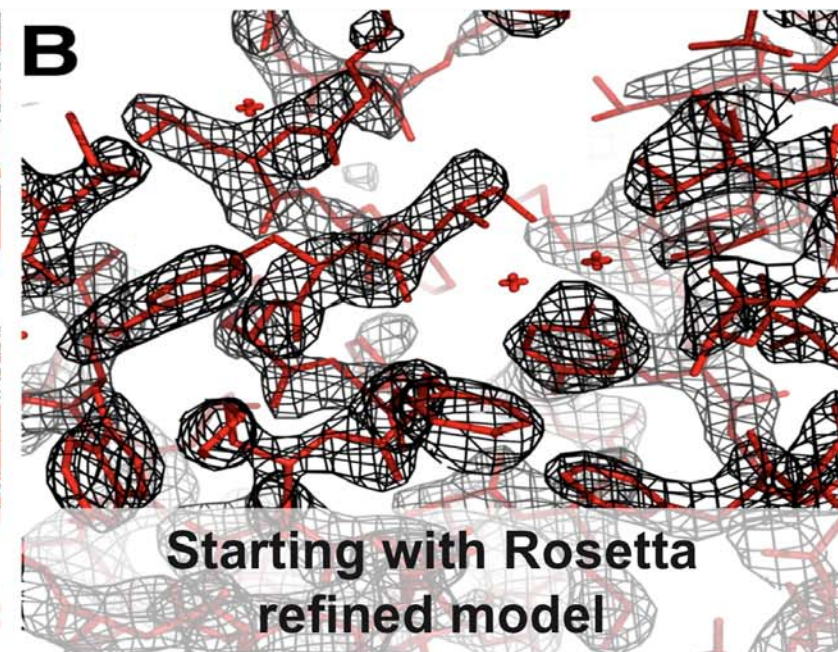



# High-Resolution Protein Structure Prediction

David Baker, Srivatsan Raman, Univ. of Washington



- Proteins are the workhorse molecules of all biological systems
- ROSETTA uses a two-phase Monte Carlo algorithm to sample the extremely large space of possible structures to find the most favorable one
- Below: Electron density (black) was generated from phases determined from two starting models. The crystal structure (red) agrees very well with the phases determined from the Rosetta model.



# Many Routes to Blue Gene/P at ALCF

## ■ Annual INCITE Competition

- 1-3 year allocations
- 5-100M CPU hr/year typical
- 8-160K processors typical

## ■ SC Director's Discretionary

- Address special needs

## ■ Blue Gene Consortium


- Porting and testing allocations
- Develop on Surveyor
- Scaling days on big systems

## ■ Open Source Development

- Development allocations
- Develop and test on Surveyor

## ■ ALCF Staff / Discretionary

- Join the ALCF team !



# *ALCF - ISSF, TCS*

