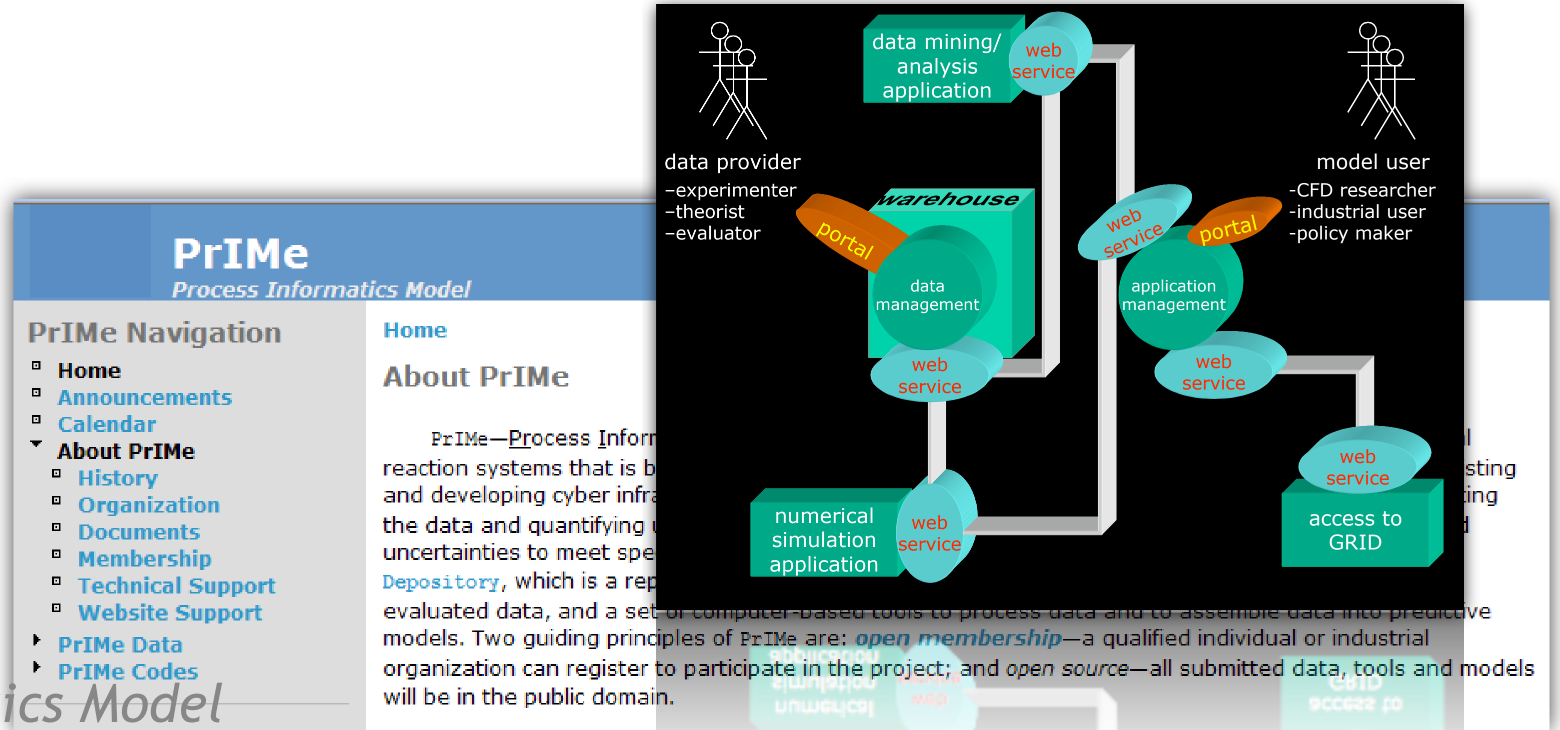

PrIMe—Process Infor
reaction systems that is b
and developing cyber infra
the data and quantifying
uncertainties to meet spe
Depository, which is a rep
evaluated data, and a set of computer-based tools to process data and to assemble data into predictive models. Two guiding principles of PrIMe are: *open membership*—a qualified individual or industrial organization can register to participate in the project; and *open source*—all submitted data, tools and models will be in the public domain.

Process Informatics Model

PrIMe

cast of 100+ (including you)

Recommendation from:
President’s Information Technology Advisory Committee

National Data and Software Repositories

- long-term support for computational science community data repositories
  - defined frameworks, metadata structures
  - algorithms, data sets, applications
  - review and validation infrastructure

- Government must require funded researchers to deposit their data and research software in these repositories or with access providers that respect any necessary or appropriate security and/or privacy requirements

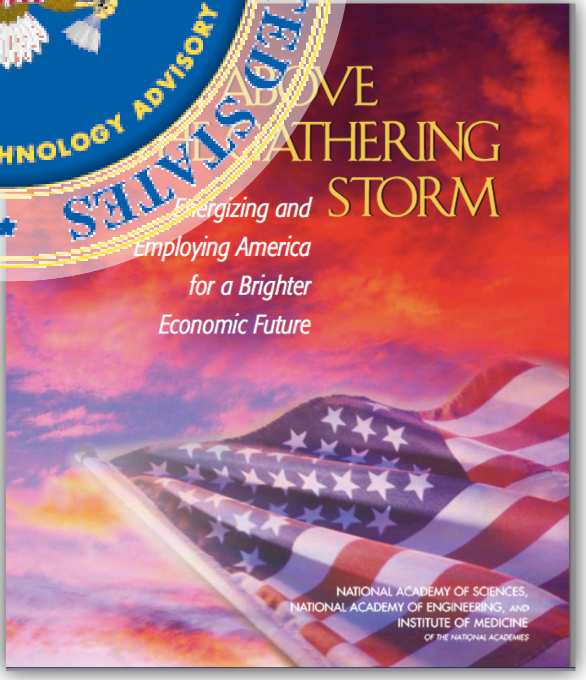

Recommendation from:
President’s Information Technology Advisory Committee

National Data and Software Repositories

- long-term support for computational science community data repositories
  - defined frameworks, metadata structures
  - algorithms, data sets, applications
  - review and validation infrastructure

- Government must require funded
res
res
wit
nec
pri

Executive Office of the President of the United States

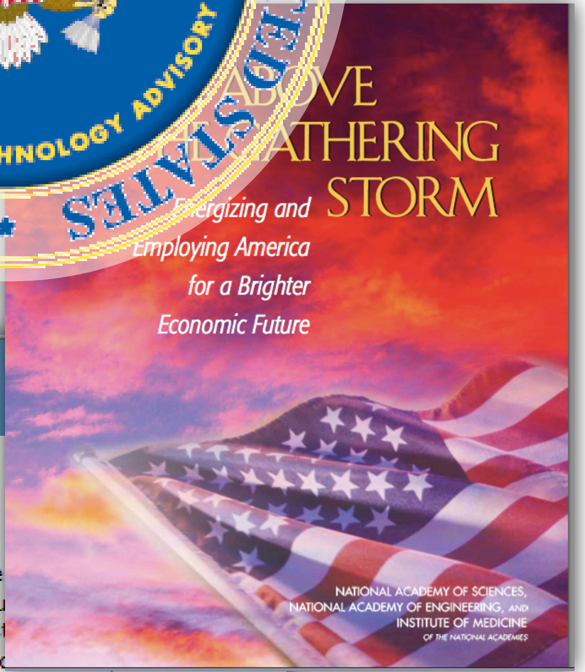

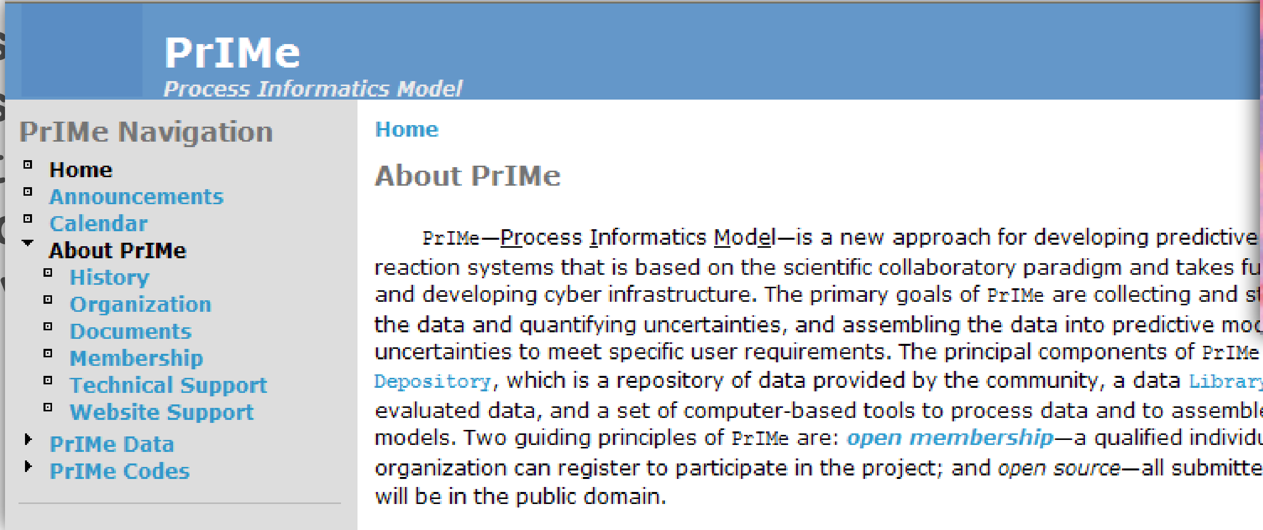

About PrIMe

PrIMe—Process Informatics Model—is a new approach for developing predictive
reaction systems that is based on the scientific collaboratory paradigm and takes fu
and developing cyber infrastructure. The primary goals of PrIMe are collecting and s
the data and quantifying uncertainties, and assembling the data into predictive mo
uncertainties to meet specific user requirements. The principal components of PrIMe include: a data Depository, which is a repository of data provided by the community, a data Library for storage of evaluated data, and a set of computer-based tools to process data and to assemble data into predictive models. Two guiding principles of PrIMe are: *open membership*—a qualified individual or industrial organization can register to participate in the project; and *open source*—all submitted data, tools and models will be in the public domain.

Combustion VO

“international data & software repository”

- multi-physics
- data intensive
- simulation tools
- error quantification

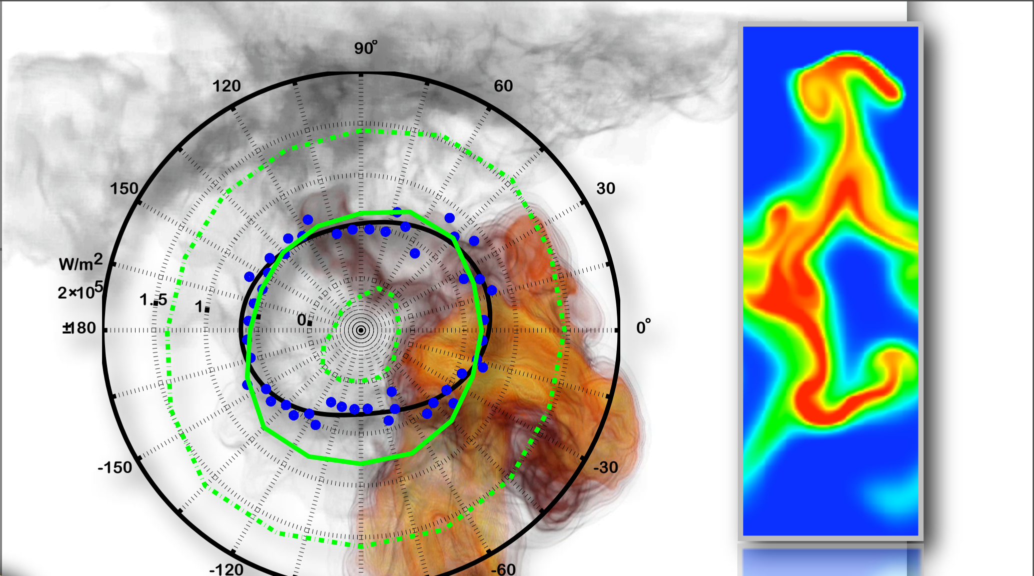

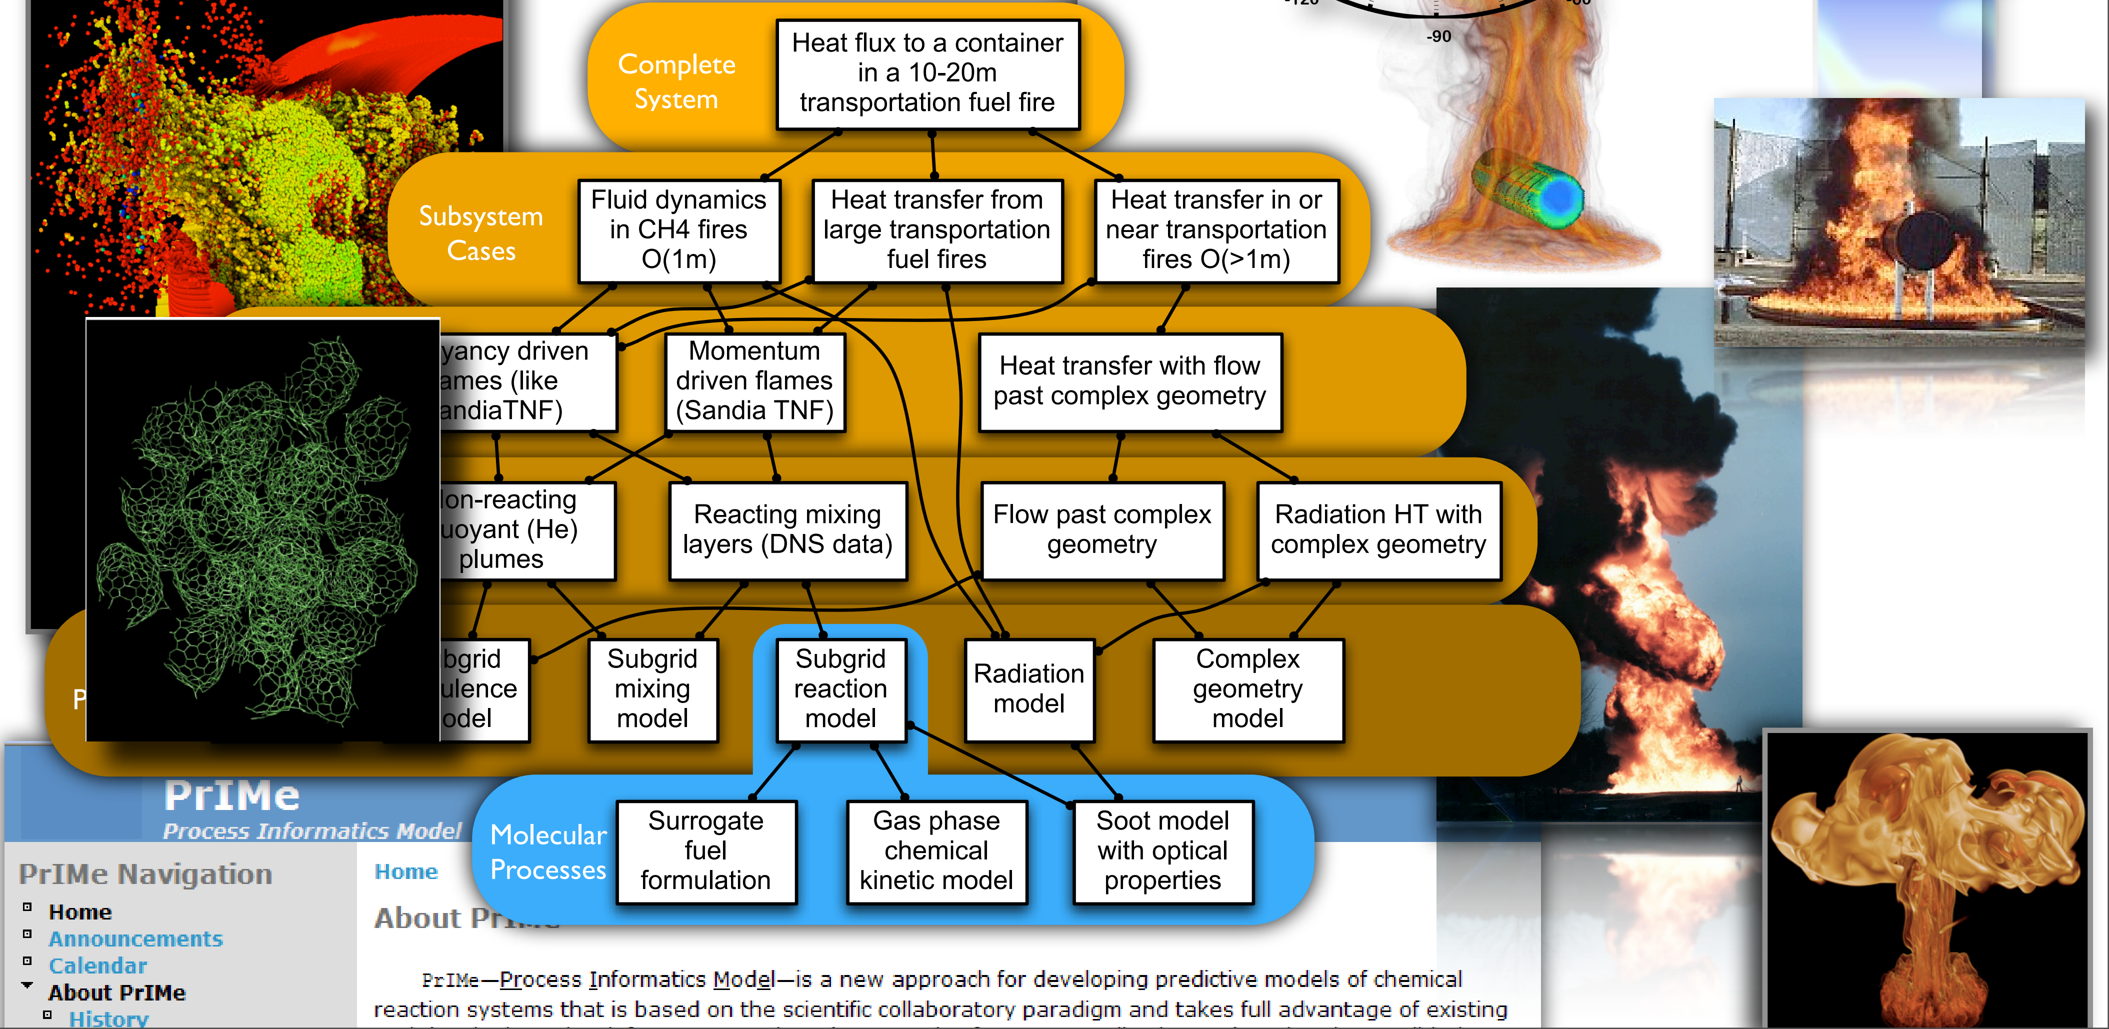

Complete System

Heat flux to a container in a 10-20m transportation fuel fire

Subsystem Cases

Fluid dynamics in CH4 fires O(1m)

Heat transfer from large transportation fuel fires

Heat transfer in or near transportation fires O(>1m)

yancy driven ames (like andiaTNF)

Momentum driven flames (Sandia TNF)

Heat transfer with flow past complex geometry

on-reacting uoyant (He) plumes

Reacting mixing layers (DNS data)

Flow past complex geometry

Radiation HT with complex geometry

P

bgrid ulence odel

Subgrid mixing model

Subgrid reaction model

Radiation model

Complex geometry model

Molecular Processes

Surrogate fuel formulation

Gas phase chemical kinetic model

Soot model with optical properties

PrIMe

Process Informatics Model

PrIMe Navigation

- **Home**
- Announcements
- Calendar
- **About PrIMe**
  - History

Home

About Pr

PrIMe—Process Informatics Model—is a new approach for developing predictive models of chemical reaction systems that is based on the scientific collaboratory paradigm and takes full advantage of existing

Combustion VO

“international data & software repository”

- multi-physics
- data intensive
- simulation tools
- error quantification

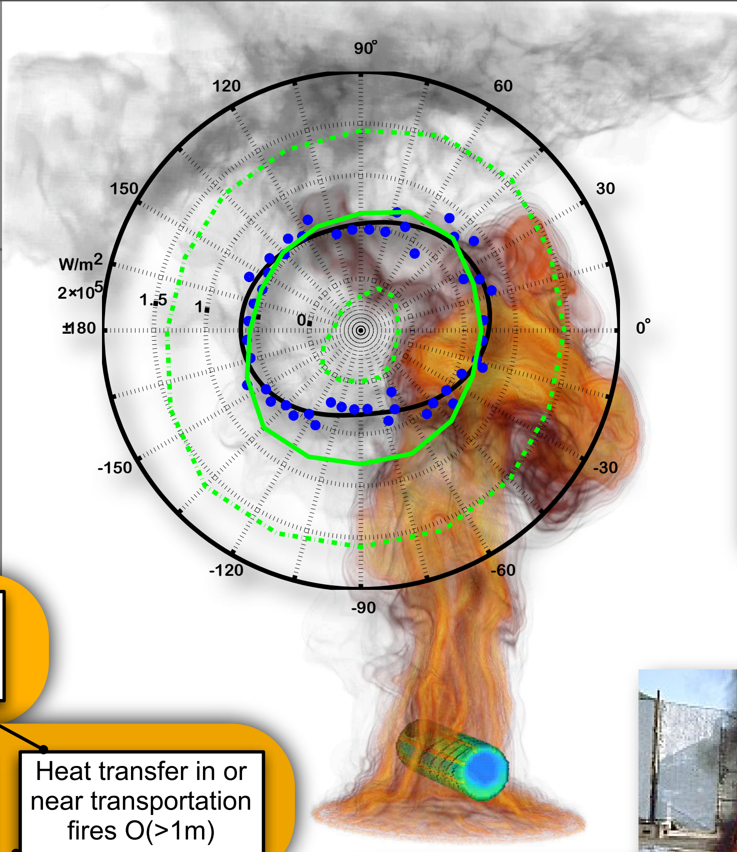

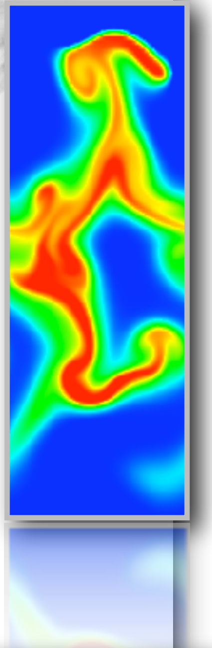

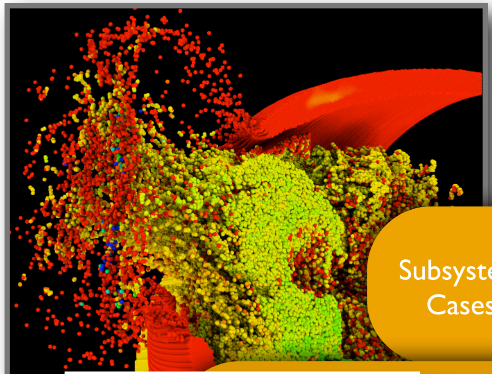

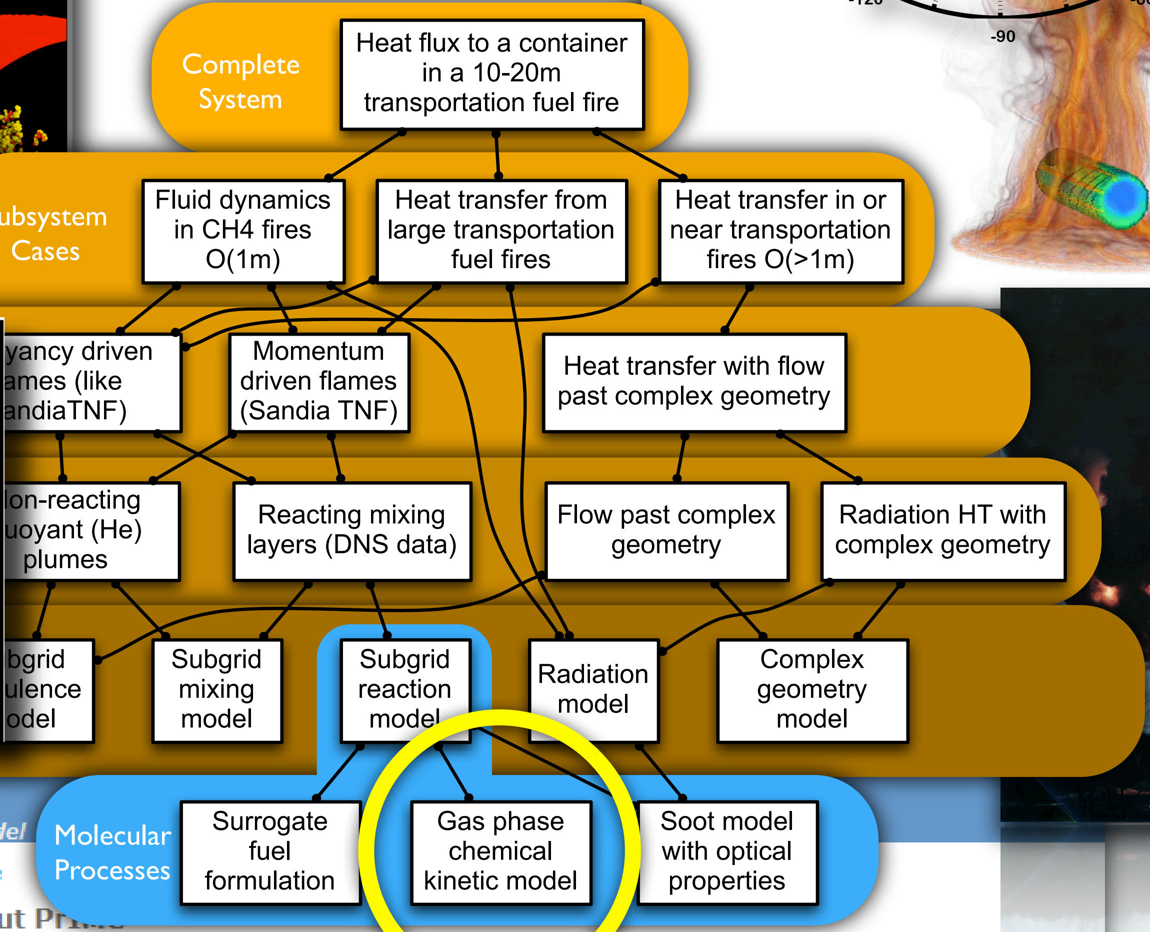

Complete System

Heat flux to a container in a 10-20m transportation fuel fire

Subsystem Cases

Fluid dynamics in CH4 fires O(1m)

Heat transfer from large transportation fuel fires

Heat transfer in or near transportation fires O(>1m)

yancy driven ames (like andiaTNF)

Momentum driven flames (Sandia TNF)

Heat transfer with flow past complex geometry

on-reacting uoyant (He) plumes

Reacting mixing layers (DNS data)

Flow past complex geometry

Radiation HT with complex geometry

P

bgrid ulence odel

Subgrid mixing model

Subgrid reaction model

Radiation model

Complex geometry model

Molecular Processes

Surrogate fuel formulation

Gas phase chemical kinetic model

Soot model with optical properties

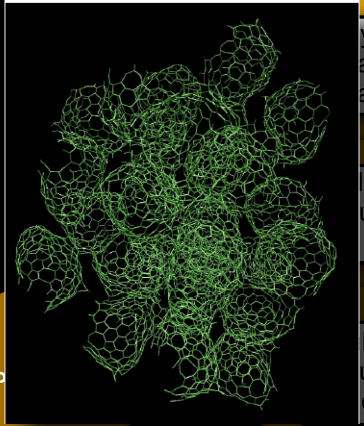

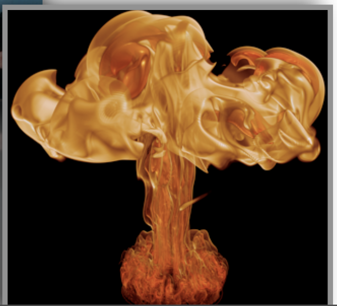

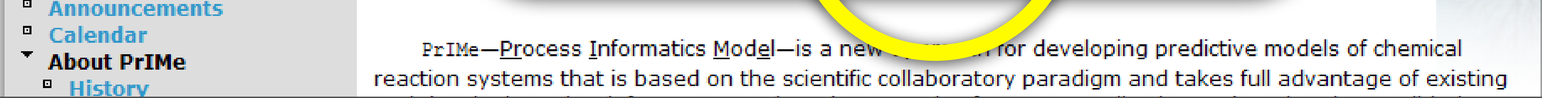

PrIMe

Process Informatics Model

PrIMe Navigation

- Home
- Announcements
- Calendar
- About PrIMe
  - History

Home

About Pr

PrIMe—Process Informatics Model—is a new for developing predictive models of chemical reaction systems that is based on the scientific collaboratory paradigm and takes full advantage of existing

Combustion VO

“international data & software repository”

- multi-physics
- data intensive
- simulation tools
- error quantification

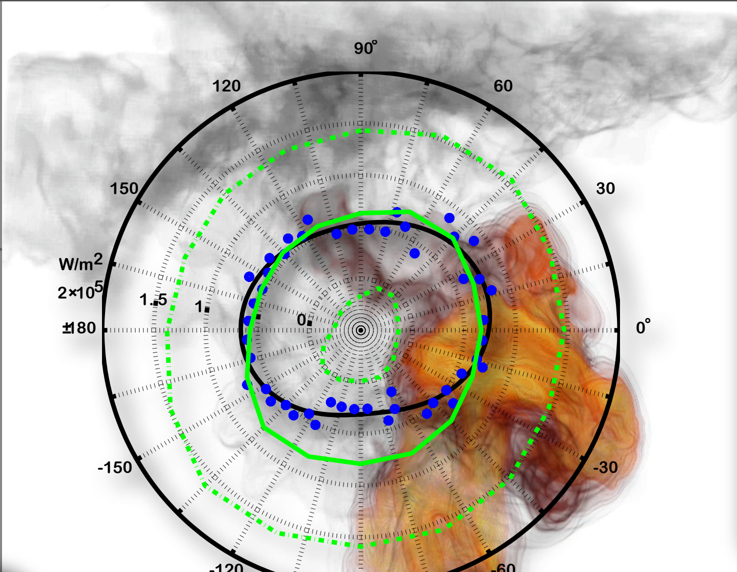

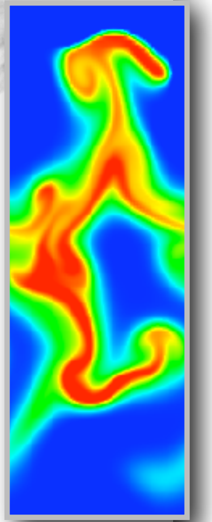

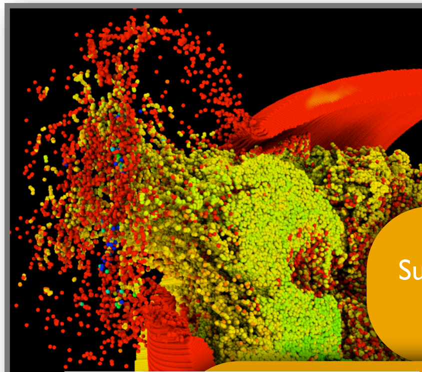

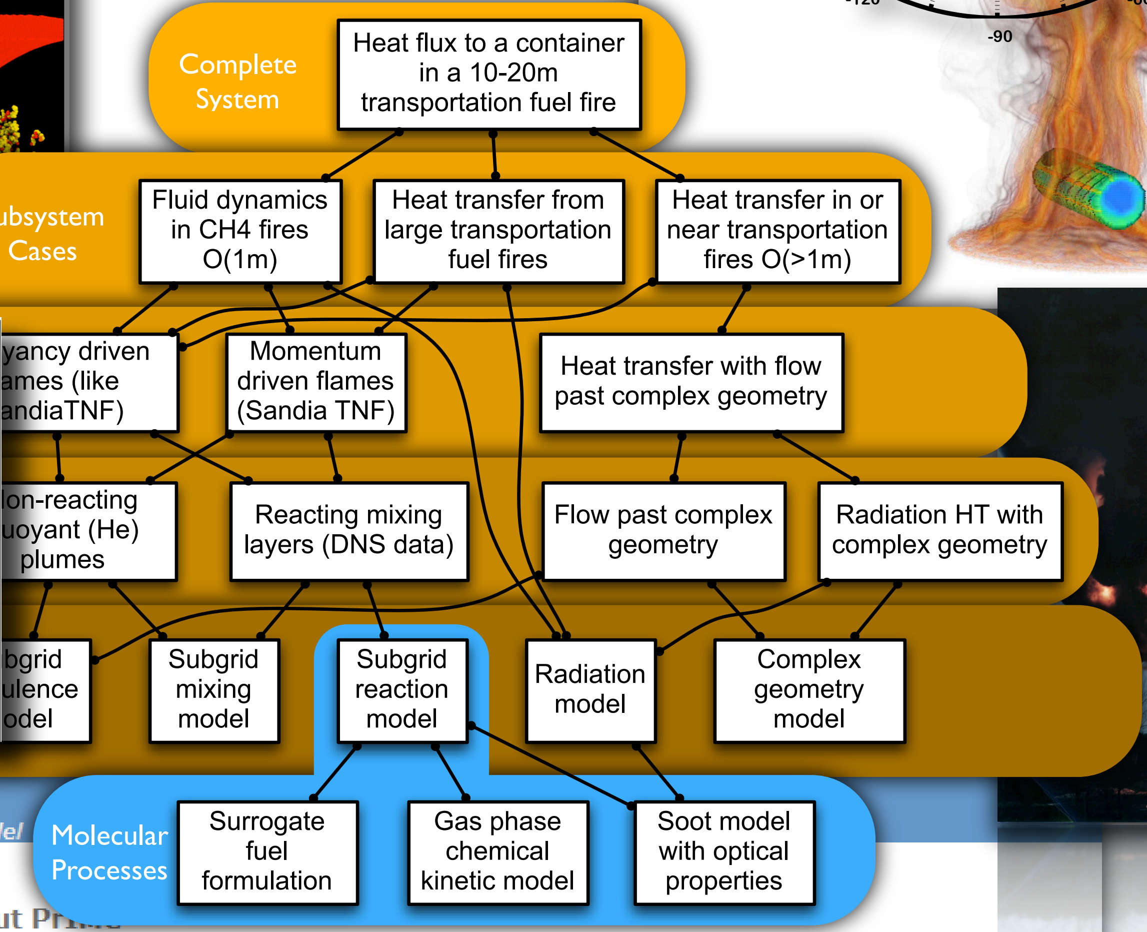

Complete System

Heat flux to a container in a 10-20m transportation fuel fire

Subsystem Cases

Fluid dynamics in CH4 fires O(1m)

Heat transfer from large transportation fuel fires

Heat transfer in or near transportation fires O(>1m)

yancy driven
ames (like
andiaTNF)

Momentum driven flames (Sandia TNF)

Heat transfer with flow past complex geometry

on-reacting
uoyant (He)
plumes

Reacting mixing layers (DNS data)

Flow past complex geometry

Radiation HT with complex geometry

P

bgrid
ulence
odel

Subgrid mixing model

Subgrid reaction model

Radiation model

Complex geometry model

Molecular Processes

Surrogate fuel formulation

Gas phase chemical kinetic model

Soot model with optical properties

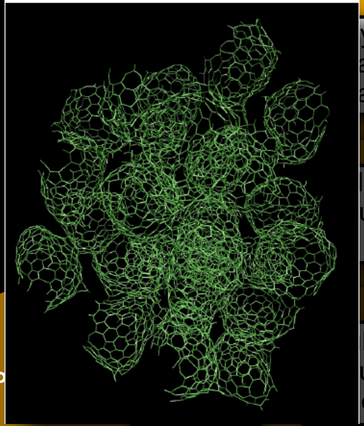

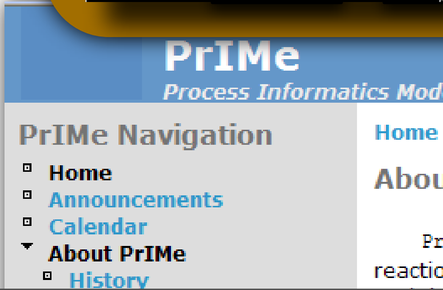

PrIMe

Process Informatics Model

PrIMe Navigation

- **Home**
- Announcements
- Calendar
- **About PrIMe**
  - History

Home

About PrIMe

PrIMe—Process Informatics Model—is a new approach for developing predictive models of chemical reaction systems that is based on the scientific collaboratory paradigm and takes full advantage of existing

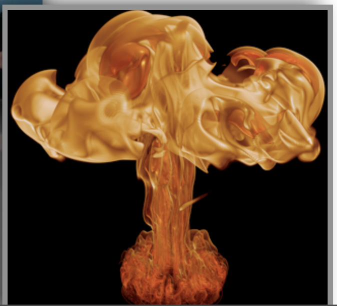

combustion community “services”

quantification & predictivity → discovery

PrIMe
Process Informatics Model

PrIMe Navigation

- Home
- Announcements
- Calendar
- About PrIMe
  - History
  - Organization
  - Documents
  - Membership
  - Technical Support
  - Website Support
- PrIMe Data
- PrIMe [illegible]

Home

About PrIMe

PrIMe—Process Informatics Model—is a new approach for developing predictive models of chemical reaction systems that is based on the scientific collaboratory paradigm and takes full advantage of existing and developing cyber infrastructure. The primary goals of PrIMe are collecting and storing data, validating the data and quantifying uncertainties, and assembling the data into predictive models with quantified uncertainties to meet specific user requirements. The principal components of PrIMe include: a data Repository, which is a repository of data provided by the community, a data Library for storage of evaluated data, and a set of computer-based tools to process data and to assemble data into predictive models. Two guiding principles of PrIMe are: *open membership*—a qualified individual or industrial organization can register to participate in the project and *open source*—all submitted data, tools and models will be in the public domain.

- historic opportunity for combustion simulations & experimental data to be used for rigorous validation and modeling error quantification.

- enabling technologies
  - tera- to peta-scale computing produces multi-scale simulations of combustion process for the discovery of new technologies
  - high-resolution diagnostic/ measurement techniques produces unprecedented volumes of experimental combustion data

- simultaneous analysis of simulation data & experimental data

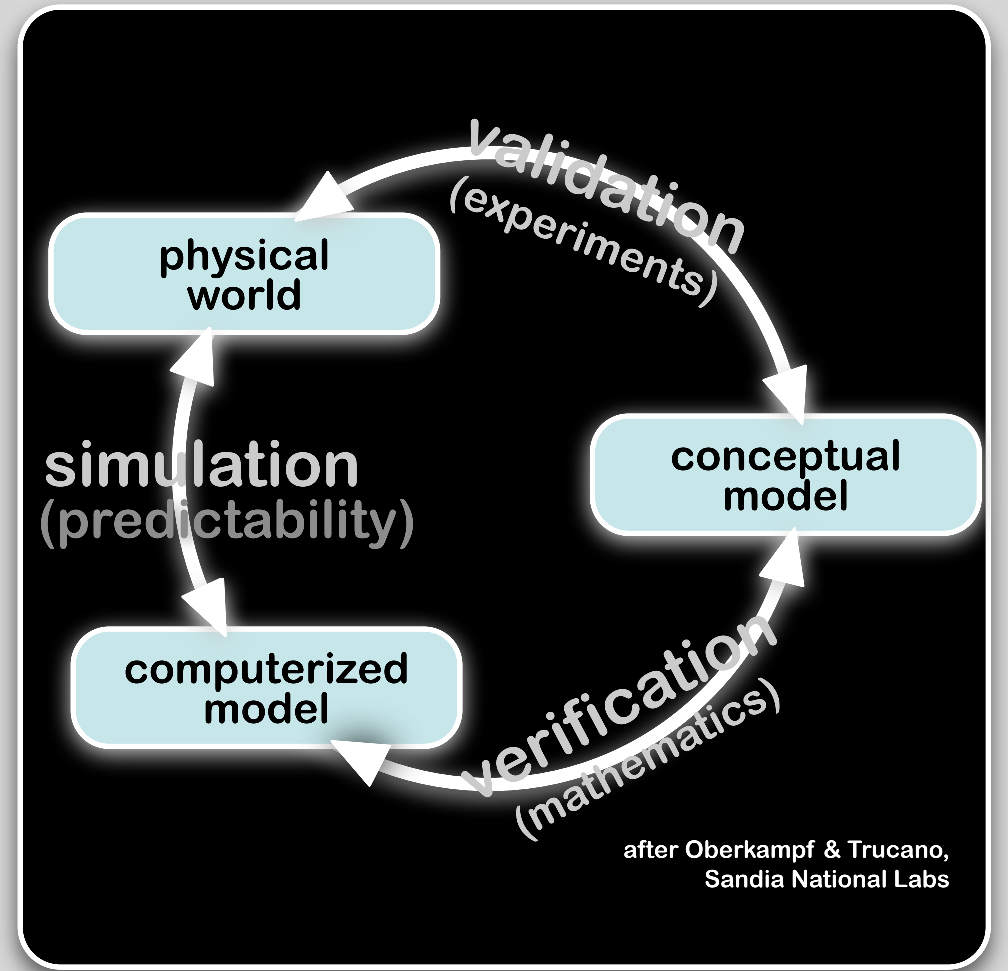

after Oberkampf & Trucano, Sandia National Labs

PrIMe—Process Informatics Model—is a new approach for developing [illegible] scientific collaboratory paradigm an[illegible] and developing cyber infrastructure. The primary goals of PrIMe are collect[illegible] the data and quantifying uncertainties, and assembling the data into pred[illegible] uncertainties to meet specific user requirements. The principal components [illegible] of data provided by the community, a da[illegible] [illegible]-based tools to process data and to [illegible] models. Two guiding principles of PrIMe are: *open membership*—a qualifie[illegible] and *open source*—al[illegible] will be in the public domain.

combustion community "services"

quantification & predictivity → discovery

- historic opportunity for combustion simulations & experimental data to be used for rigorous validation and modeling error quantification.

- enabling technologies
  - tera- to peta-scale computing produces multi-scale simulations of combustion process for the discovery of new technologies
  - high-resolution diagnostic/ measurement techniques produces unprecedented volumes of experimental combustion data

- simultaneous analysis of simulation data & experimental data

Computational Science Demands a New Paradigm

The field has reached a threshold at which better organization becomes crucial. New methods of verifying and validating complex codes are mandatory if computational science is to fulfill its promise for science and society.

Douglass E. Post and Lawrence G. Votta

Computers have become indispensable to scientific research. They are essential for collecting and analyzing experimental data, and they have largely replaced pencil and paper as the theorist's main tool. Computers let theorists extend their studies of physical, chemical, and biological systems by solving difficult nonlinear problems in magnetohydrodynamics; atomic, molecular, and nuclear structure; fluid turbulence; shock hydrodynamics; and cosmological structure formation.

Beyond such well-established aids to theorists and experimenters, the exponential growth of computer power is now launching the new field of computational science. Multidisciplinary computational teams are beginning to develop large-scale predictive simulations of highly complex technical problems. Large-scale codes have been created to simulate, with unprecedented fidelity, phenomena such as supernova explosions (see figures 1 and 2), inertial-confinement fusion, nuclear explosions (see the box on page 38), asteroid impacts (figure 3), and the effect of space weather on Earth's magnetosphere (figure 4).

Computational simulation has the potential to join theory and experiment as a third powerful research methodology. Although, as figures 1–4 show, the new discipline is already yielding important and exciting results, it is also becoming all too clear that much of computational science is still troublingly immature. We point out three distinct challenges that computational science must meet if it is to fulfill its potential and take its place as a fully mature partner of theory and experiment:

▶ *the performance challenge*—producing high-performance computers,
▶ *the programming challenge*—programming for complex computers, and
▶ *the prediction challenge*—developing truly predictive complex application codes.

The performance challenge requires that the exponential growth of computer performance continue, yielding ever larger memories and faster processing. The programming challenge involves the writing of codes that can efficiently exploit the capacities of the increasingly complex computers. The prediction challenge is to use all that computing power to provide answers reliable enough to form the basis for important decisions.

The performance challenge is being met, at least for the next 10 years. Processor speed continues to increase, and massive parallelization is augmenting that speed, albeit at the cost of increasingly complex computer architectures. Massively parallel computers with thousands of processors are becoming widely available at relatively low cost, and larger ones are being developed.

Much remains to be done to meet the programming challenge. But computer scientists are beginning to develop languages and software tools to facilitate programming for massively parallel computers.

The most urgent challenge

The prediction challenge is now the most serious limiting factor for computational science. The field is in transition from modest codes developed by small teams to much more complex programs, developed over many years by large teams, that incorporate many strongly coupled effects spanning wide ranges of spatial and temporal scales. The prediction challenge is due to the complexity of the newer codes, and the problem of integrating the efforts of large teams. This often results in codes that are not sufficiently reliable and credible to be the basis of important decisions facing society. The growth of code size and complexity, and its attendant problems, bears some resemblance to the transition from small to large scale by experimental physics in the decades after World War II.

A comparative case study of six large-scale scientific code projects, by Richard Kendall and one of us (Post),[1] has yielded three important lessons. Verification, validation, and quality management, we found, are all crucial to the success of a large-scale code-writing project. Although some computational science projects—those illustrated by figures 1–4, for example—stress all three requirements, many other current and planned projects give them insufficient attention. In the absence of any one of those requirements, one doesn't have the assurance of independent assessment, confirmation, and repeatability of results. Because it's impossible to judge the validity of such results, they often have little credibility and no impact.

Part of the problem is simply that it's hard to decide whether a code result is right or wrong. Our experience as referees and editors tells us that the peer review process in computational science generally doesn't provide as effective a filter as it does for experiment or theory. Many things that a referee cannot detect could be wrong with a computational-science paper. The code could have hidden defects, it might be applying algorithms improperly, or its spatial or temporal resolution might be inappropriately coarse.

Douglass Post is a computational physicist at Los Alamos National Laboratory and an associate editor-in chief of *Computing in Science and Engineering*. **Lawrence Votta** is a Distinguished Engineer at Sun Microsystems Inc in Menlo Park, California. He has been an associate editor of *IEEE Transactions on Software Engineering*.

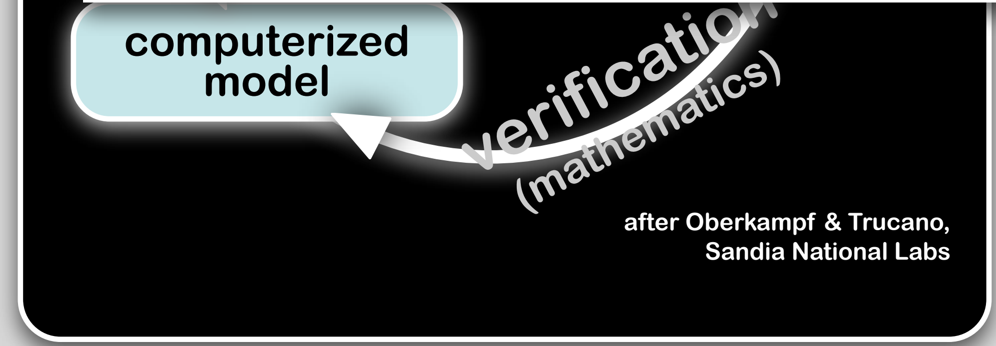

after Oberkampf & Trucano, Sandia National Labs

PrIMe—Process Informatics Model—is a new approach

scientific collabora

and developing cyber infrastructure. The primary goals o

the data and quantifying uncertainties, and assembling

uncertainties to meet specific user requirements. The pri

of data provided by the

r-based tools to pr

will be in the public domain.

combustion community "services"

quantification & predictivity → discovery

- historic opportunity for combustion simulations & experimental data to be used for rigorous validation and modeling error quantification.

- enabling technologies
    - tera- to peta-scale computing produces multi-scale simulations of combustion process for the discovery of new technologies
    - high-resolution diagnostic/ measurement techniques produces unprecedented volumes of experimental combustion data

- simultaneous analysis of simulation data & experimental data

Computational Science Demands a New Paradigm

The field has reached a threshold at which better organization becomes crucial. New methods of verifying and validating

efficiently exploit the capacities of the increasingly complex computers. The prediction challenge is to use all that

VIEWPOINT

Service-Oriented Science

Ian Foster

New information architectures enable new approaches to publishing and accessing valuable data and programs. So-called service-oriented architectures define standard interfaces and protocols that allow developers to encapsulate information tools as services that clients can access without knowledge of, or control over, their internal workings. Thus, tools formerly accessible only to the specialist can be made available to all; previously manual data-processing and analysis tasks can be automated by having services access services. Such service-oriented approaches to science are already being applied successfully, in some cases at substantial scales, but much more effort is required before these approaches are applied routinely across many disciplines. Grid technologies can accelerate the development and adoption of service-oriented science by enabling a separation of concerns between discipline-specific content and domain-independent software and hardware infrastructure.

Paul Erdös claimed that a mathematician is a machine for turning coffee into theorems. The scientist is arguably a machine for turning data into insight. However, advances in information technology are changing the way in which this role is fulfilled—by automating time-consuming activities and thus freeing the scientist to perform other tasks. In this Viewpoint, I discuss how service-oriented computing—technology that allows powerful information tools to be made available over the network, always on tap, and easy for scientists to use—may contribute to that evolution.

The practice of science has, of course, already been affected dramatically by information technology and, in particular, by the Internet. For example, the hundreds of gigabytes of genome sequence available online means that for a growing number of biologists, "data" is something that they find on the Web, not in the lab. Similarly, emerging "digital observatories" [already several hundred terabytes in dozens of archives (*1*)] allow astronomers to pose and answer in seconds questions that might previously have required years of observation. In fields such as cosmology and climate, supercomputer simulations have emerged as essential tools, themselves producing large data sets that, when published online, are of interest to many (*2*). An exploding number of sensors (*3*), the rapidly expanding computing and storage capabilities of federated Grids (*4*), and advances in optical networks (*5*) are accelerating these trends by making increasingly powerful capabilities available online.

Sometimes, however, the thrill of the Web seems to blind us to the true implications of these developments. Human access to online resources is certainly highly useful, putting a global library at our fingertips. But ultimately, it is automated access by software programs that will be truly revolutionary, simply because of the higher speeds at which programs can operate. In the time that a human user takes to locate one useful piece of information within a Web site, a program may access and integrate data from many sources and identify relationships that a human might never discover unaided. Two dramatic examples are systems that automatically integrate information from genome and protein sequence databases to infer metabolic pathways (*6*) and systems that search digital sky surveys to locate brown dwarfs (*7*).

The key to such success is uniformity of interface, so that programs can discover and access services without the need to write custom code for each specific data source, program, or sensor. Electric power–transmission standards and infrastructure enabled development of the electric power grid and spurred the development of a plethora of electric tools. In a similar manner, service technologies enable the development of a wide range of programs that integrate across multiple existing services for purposes such as metabolic pathway reconstruction, categorization of astronomical objects, and analysis of environmental data. If such programs are themselves made accessible as services, the result can be the creation of distributed networks of services, each constructed by a different individual or group, and each providing some original content and/or value-added product (*8*).

We see this evolution occurring in the commercial Internet. As the Web has expanded in scale, so the preferred means of finding things has evolved from Yahoo's manually assembled lists to Google's automatically computed indices. Now Google is making its indices accessible, spurring development of yet other services. What makes Google's indices feasible is the existence of large quantities of data in a uniform format (HTML, HyperText Markup Language) and—two important factors that must be considered when we turn to science—smart computer scientists to develop the algorithms and software required to manage the 100,000 computers used (at last count) to analyze Web link structure, and smart businesspeople to raise the money that pays for those computers!

The term "service-oriented architecture" refers to systems structured as networks of loosely coupled, communicating services (*9*). Thus, "service-oriented science" refers to scientific research enabled by distributed networks of interoperating services. [The term "e-Science," coined by John Taylor, has a similar but broader connotation (*10*).]

Creating and Sharing Services

Creating a service involves describing, in some conventional manner, the operations that the service supports; defining the protocol used to invoke those operations over the Internet; and operating a server to process incoming requests. A set of technologies called Web services (*9*) are gaining wide acceptance for these purposes. A variety of commercial and open-source Web services tools exist for developing services, deploying and operating services, and developing client applications. A fair amount of experience has been gained with the creation of services and applications in different science domains. Although problems remain (e.g., efficiency, interoperability of different vendor offerings), the technology is well beyond the experimental stage. Nevertheless, it can still be a big step to realize the full potential of service-oriented science, for reasons that I now discuss.

Interoperability. Services have little value if others cannot discover, access, and make sense of them. Yet, as Stein has observed (*11*), today's scientific communities too often resemble medieval Italy's collection of warring city states, each with its own legal system and dialect. Web services mechanisms for describing, discovering, accessing, and securing services provide a common alphabet, but a true lingua franca requires agreement on protocols, data formats, and ultimately semantics (*12*). For example, the definition of VOTable, a standard XML (eXtensible Markup Language)–based representation for tabular data (*13*), has been a powerful force for progress in astronomy.

Scale. Services must often deal with data volumes, computational demands, and numbers of users beyond the capacity of a typical PC. Responding to a user request—or to the arrival of new data—can involve large amounts of computation. For example, the Argonne GNARE system searches periodically through DNA and protein databases for new and updated genomes and then computes and pub-

---

Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439, USA, and Department of Computer Science, University of Chicago, Chicago, IL 60637, USA. E-mail: foster@mcs.anl.gov

combustion community “services”

quantification & predictivity → discovery

- historic opportunity for combustion simulations & experimental data to be used for rigorous validation and modeling error quantification
- enabling technologies
  - tera- to peta-scale computing produces multi-scale simulations of combustion process for the discovery of new technologies
  - high resolution diagnostic measurement techniques produces unprecedented volumes of experimental combustion data
- simultaneous analysis of simulation data & experimental data

PrIMe—Process Informatics Model—is a new approach for developing reaction systems … the scientific collabora… and developing cyber infrastructure. The primary goals o… the data and quantifying uncertainties, and assembling … uncertainties to meet specific user requirements. The pri… of data provided by the … evaluated data, and a set of computer-based tools to pr… models. Two guiding principles of PrIMe are: open mem… will be in the public domain.

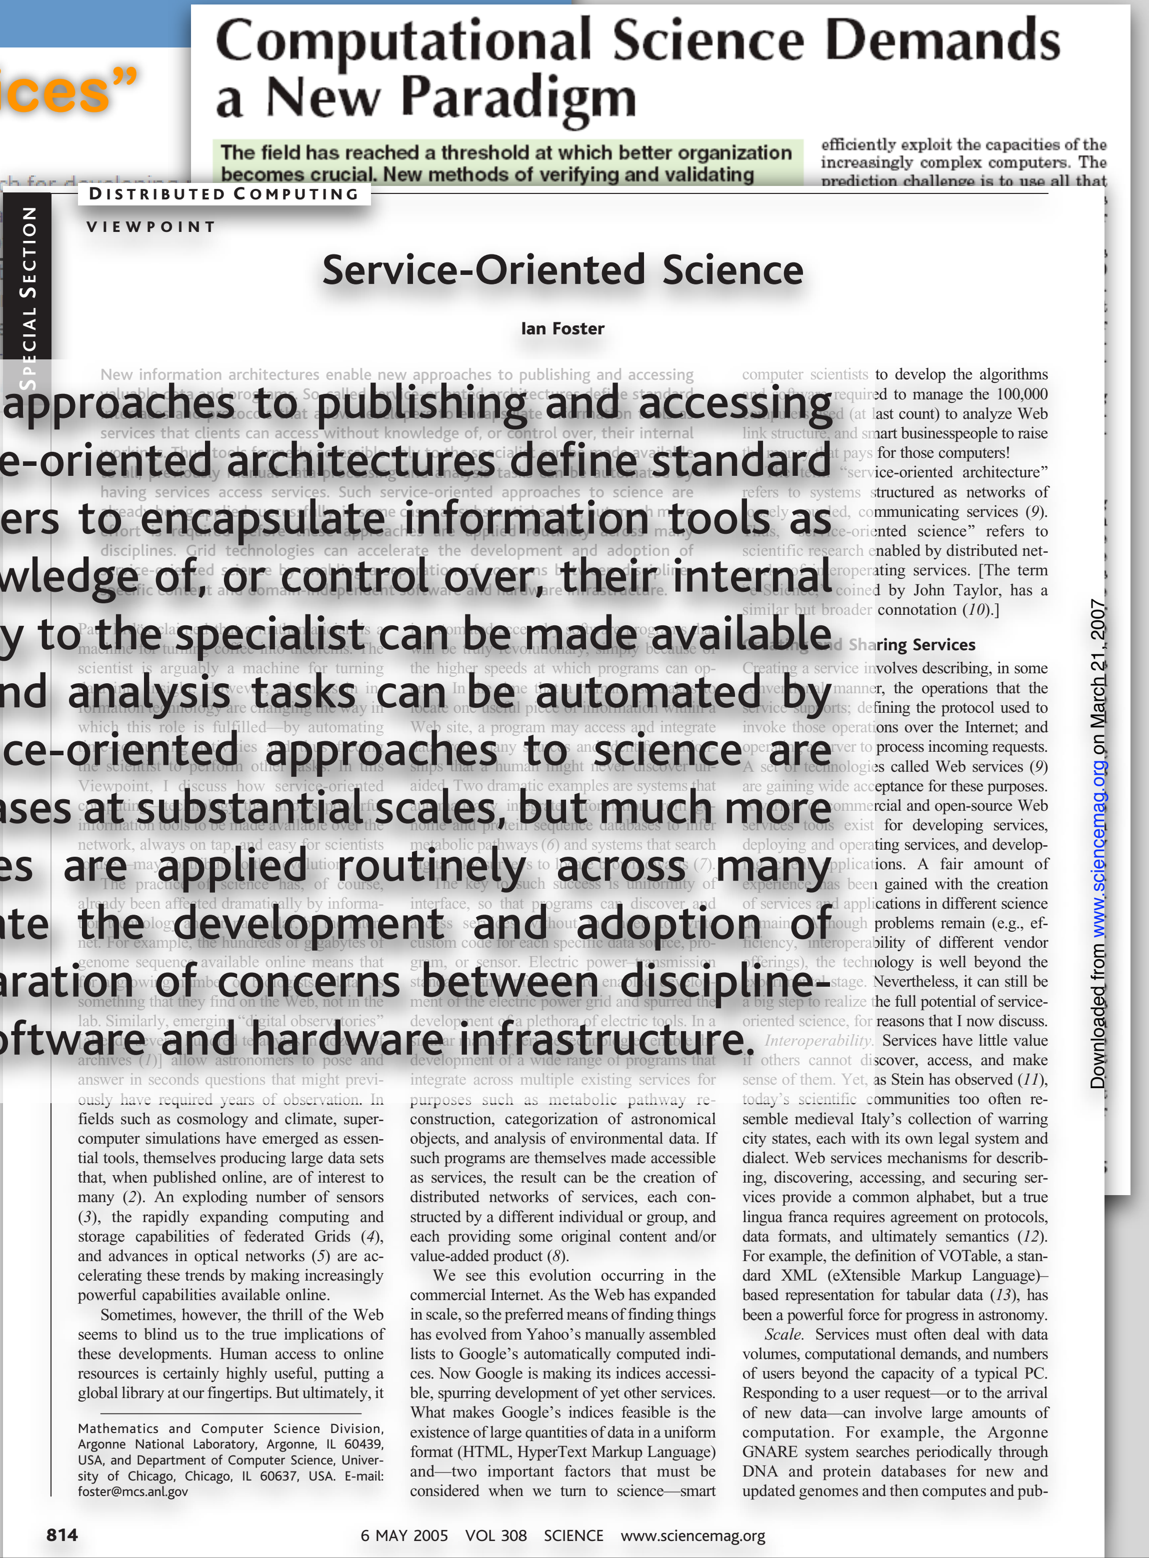

Computational Science Demands a New Paradigm

The field has reached a threshold at which better organization becomes crucial. New methods of verifying and validating

efficiently exploit the capacities of the increasingly complex computers. The prediction challenge is to use all that

---

DISTRIBUTED COMPUTING

SPECIAL SECTION

VIEWPOINT

Service-Oriented Science

Ian Foster

New information architectures enable new approaches to publishing and accessing valuable data and programs. So-called service-oriented architectures define standard interfaces and protocols that allow developers to encapsulate information tools as services that clients can access without knowledge of, or control over, their internal workings. Thus, tools formerly accessible only to the specialist can be made available to all; previously manual data-processing and analysis tasks can be automated by having services access services. Such service-oriented approaches to science are already being applied successfully, in some cases at substantial scales, but much more effort is required before these approaches are applied routinely across many disciplines. Grid technologies can accelerate the development and adoption of service-oriented science by enabling a separation of concerns between discipline-specific content and domain-independent software and hardware infrastructure.

[illegible] machine for turning coffee into theorems. The scientist is arguably a machine for turning [illegible] However, advances in information technology are changing the way in which this role is fulfilled—by automating [illegible] the scientist to perform other tasks. In this Viewpoint, I discuss how service-oriented [illegible] information tools to be made available over the network, always on tap, and easy for scientists [illegible] evolution.

The practice of science has, of course, already been affected dramatically by informa[illegible] net. For example, the hundreds of gigabytes of genome sequence available online means that [illegible] something that they find on the Web, not in the lab. Similarly, emerging “digital observatories” [illegible] archives (1)] allow astronomers to pose and answer in seconds questions that might previously have required years of observation. In fields such as cosmology and climate, supercomputer simulations have emerged as essential tools, themselves producing large data sets that, when published online, are of interest to many (2). An exploding number of sensors (3), the rapidly expanding computing and storage capabilities of federated Grids (4), and advances in optical networks (5) are accelerating these trends by making increasingly powerful capabilities available online.

Sometimes, however, the thrill of the Web seems to blind us to the true implications of these developments. Human access to online resources is certainly highly useful, putting a global library at our fingertips. But ultimately, it [illegible] will be truly revolutionary, simply because of the higher speeds at which programs can op[illegible] locate one useful piece of information within a Web site, a program may access and integrate [illegible] ships that a human might never discover unaided. Two dramatic examples are systems that [illegible] genome and protein sequence databases to infer metabolic pathways (6) and systems that search [illegible] (7).

The key to such success is uniformity of interface, so that programs can discover and [illegible] custom code for each specific data source, program, or sensor. Electric power–transmission [illegible] ment of the electric power grid and spurred the development of a plethora of electric tools. In a [illegible] development of a wide range of programs that integrate across multiple existing services for purposes such as metabolic pathway reconstruction, categorization of astronomical objects, and analysis of environmental data. If such programs are themselves made accessible as services, the result can be the creation of distributed networks of services, each constructed by a different individual or group, and each providing some original content and/or value-added product (8).

We see this evolution occurring in the commercial Internet. As the Web has expanded in scale, so the preferred means of finding things has evolved from Yahoo’s manually assembled lists to Google’s automatically computed indices. Now Google is making its indices accessible, spurring development of yet other services. What makes Google’s indices feasible is the existence of large quantities of data in a uniform format (HTML, HyperText Markup Language) and—two important factors that must be considered when we turn to science—smart computer scientists to develop the algorithms and software required to manage the 100,000 computers used (at last count) to analyze Web link structure, and smart businesspeople to raise the money that pays for those computers!

The term “service-oriented architecture” refers to systems structured as networks of loosely coupled, communicating services (9). Thus, “service-oriented science” refers to scientific research enabled by distributed networks of interoperating services. [The term “e-Science,” coined by John Taylor, has a similar but broader connotation (10).]

Creating and Sharing Services

Creating a service involves describing, in some conventional manner, the operations that the service supports; defining the protocol used to invoke those operations over the Internet; and operating a server to process incoming requests. A set of technologies called Web services (9) are gaining wide acceptance for these purposes. A variety of commercial and open-source Web services tools exist for developing services, deploying and operating services, and developing client applications. A fair amount of experience has been gained with the creation of services and applications in different science domains. Although problems remain (e.g., efficiency, interoperability of different vendor offerings), the technology is well beyond the experimental stage. Nevertheless, it can still be a big step to realize the full potential of service-oriented science, for reasons that I now discuss.

Interoperability. Services have little value if others cannot discover, access, and make sense of them. Yet, as Stein has observed (11), today’s scientific communities too often resemble medieval Italy’s collection of warring city states, each with its own legal system and dialect. Web services mechanisms for describing, discovering, accessing, and securing services provide a common alphabet, but a true lingua franca requires agreement on protocols, data formats, and ultimately semantics (12). For example, the definition of VOTable, a standard XML (eXtensible Markup Language)–based representation for tabular data (13), has been a powerful force for progress in astronomy.

Scale. Services must often deal with data volumes, computational demands, and numbers of users beyond the capacity of a typical PC. Responding to a user request—or to the arrival of new data—can involve large amounts of computation. For example, the Argonne GNARE system searches periodically through DNA and protein databases for new and updated genomes and then computes and pub-

Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439, USA, and Department of Computer Science, University of Chicago, Chicago, IL 60637, USA. E-mail: foster@mcs.anl.gov

combustion community “services”

quantification & predictivity → discovery

- historic opportunity for combustion simulations & experimental data to be used for rigorous validation & error quantification
- enabling technologies
  - tera- to peta-scale computing produces multi-scale simulations of combustion process for the discovery of new technologies
  - high-resolution diagnostic measurement techniques produces unprecedented volumes of experimental combustion data
- simultaneous analysis of simulation data & experimental data

New information architectures enable new approaches to publishing and accessing valuable data and programs. So-called service-oriented architectures define standard interfaces and protocols that allow developers to encapsulate information tools as services that clients can access without knowledge of, or control over, their internal workings. Thus, tools formerly accessible only to the specialist can be made available to all; previously manual data-processing and analysis tasks can be automated by having services access services. Such service-oriented approaches to science are already being applied successfully, in some cases at substantial scales, but much more effort is required before these approaches are applied routinely across many disciplines. Grid technologies can accelerate the development and adoption of service-oriented science by enabling a separation of concerns between discipline-specific content and domain-independent software and hardware infrastructure.

service-oriented architecture

Computational Science Demands a New Paradigm

The field has reached a threshold at which better organization becomes crucial. New methods of verifying and validating

efficiently exploit the capacities of the increasingly complex computers. The prediction challenge is to use all that

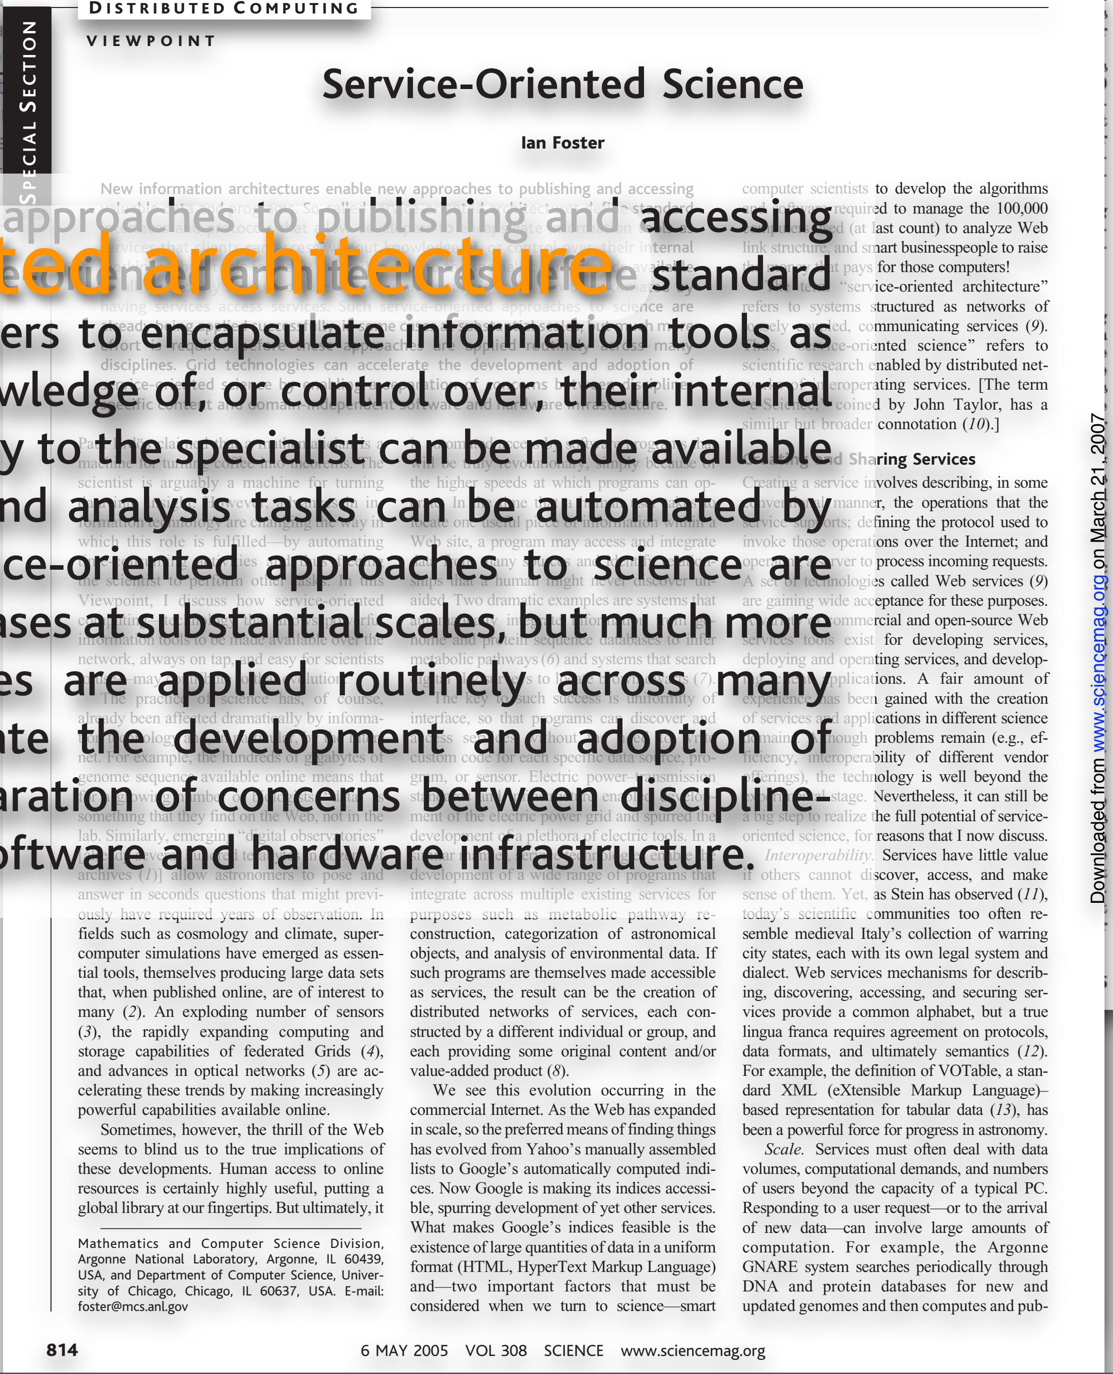

SPECIAL SECTION

DISTRIBUTED COMPUTING

VIEWPOINT

Service-Oriented Science

Ian Foster

New information architectures enable new approaches to publishing and accessing valuable data and programs. So-called service-oriented architectures define standard interfaces and protocols that allow developers to encapsulate information tools as services that clients can access without knowledge of, or control over, their internal workings. Thus, tools formerly accessible only to the specialist can be made available to all; previously manual data-processing and analysis tasks can be automated by having services access services. Such service-oriented approaches to science are already being applied successfully, in some cases at substantial scales, but much more effort is required before these approaches are applied routinely across many disciplines. Grid technologies can accelerate the development and adoption of service-oriented science by enabling a separation of concerns between discipline-specific content and domain-independent software and hardware infrastructure.

Paul Erdős claimed that a mathematician is a machine for turning coffee into theorems. The scientist is arguably a machine for turning data into insight. However, advances in information technology are changing the way in which this role is fulfilled—by automating time-consuming activities and thus freeing the scientist to perform other tasks. In this Viewpoint, I discuss how service-oriented computing—technology that allows powerful information tools to be made available over the network, always on tap, and easy for scientists to use—may contribute to that evolution.

The practice of science has, of course, already been affected dramatically by information technology and, in particular, by the Internet. For example, the hundreds of gigabytes of genome sequence available online means that for a growing number of biologists, “data” is something that they find on the Web, not in the lab. Similarly, emerging “digital observatories” [already several hundred terabytes in dozens of archives (*1*)] allow astronomers to pose and answer in seconds questions that might previously have required years of observation. In fields such as cosmology and climate, supercomputer simulations have emerged as essential tools, themselves producing large data sets that, when published online, are of interest to many (*2*). An exploding number of sensors (*3*), the rapidly expanding computing and storage capabilities of federated Grids (*4*), and advances in optical networks (*5*) are accelerating these trends by making increasingly powerful capabilities available online.

Sometimes, however, the thrill of the Web seems to blind us to the true implications of these developments. Human access to online resources is certainly highly useful, putting a global library at our fingertips. But ultimately, it is automated access by software programs that will be truly revolutionary, simply because of the higher speeds at which programs can operate. In the time that a human user takes to locate one useful piece of information within a Web site, a program may access and integrate data from many sources and identify relationships that a human might never discover unaided. Two dramatic examples are systems that automatically integrate information from genome and protein sequence databases to infer metabolic pathways (*6*) and systems that search digital sky surveys to locate brown dwarfs (*7*).

The key to such success is uniformity of interface, so that programs can discover and access services without the need to write custom code for each specific data source, program, or sensor. Electric power–transmission standards and infrastructure enabled development of the electric power grid and spurred the development of a plethora of electric tools. In a similar manner, service technologies enable the development of a wide range of programs that integrate across multiple existing services for purposes such as metabolic pathway reconstruction, categorization of astronomical objects, and analysis of environmental data. If such programs are themselves made accessible as services, the result can be the creation of distributed networks of services, each constructed by a different individual or group, and each providing some original content and/or value-added product (*8*).

We see this evolution occurring in the commercial Internet. As the Web has expanded in scale, so the preferred means of finding things has evolved from Yahoo’s manually assembled lists to Google’s automatically computed indices. Now Google is making its indices accessible, spurring development of yet other services. What makes Google’s indices feasible is the existence of large quantities of data in a uniform format (HTML, HyperText Markup Language) and—two important factors that must be considered when we turn to science—smart computer scientists to develop the algorithms and software required to manage the 100,000 computers used (at last count) to analyze Web link structure, and smart businesspeople to raise the money that pays for those computers!

The term “service-oriented architecture” refers to systems structured as networks of loosely coupled, communicating services (*9*). Thus, “service-oriented science” refers to scientific research enabled by distributed networks of interoperating services. [The term “e-Science,” coined by John Taylor, has a similar but broader connotation (*10*).]

Creating and Sharing Services

Creating a service involves describing, in some conventional manner, the operations that the service supports; defining the protocol used to invoke those operations over the Internet; and operating a server to process incoming requests. A set of technologies called Web services (*9*) are gaining wide acceptance for these purposes. A variety of commercial and open-source Web services tools exist for developing services, deploying and operating services, and developing client applications. A fair amount of experience has been gained with the creation of services and applications in different science domains. Although problems remain (e.g., efficiency, interoperability of different vendor offerings), the technology is well beyond the experimental stage. Nevertheless, it can still be a big step to realize the full potential of service-oriented science, for reasons that I now discuss.

Interoperability. Services have little value if others cannot discover, access, and make sense of them. Yet, as Stein has observed (*11*), today’s scientific communities too often resemble medieval Italy’s collection of warring city states, each with its own legal system and dialect. Web services mechanisms for describing, discovering, accessing, and securing services provide a common alphabet, but a true lingua franca requires agreement on protocols, data formats, and ultimately semantics (*12*). For example, the definition of VOTable, a standard XML (eXtensible Markup Language)–based representation for tabular data (*13*), has been a powerful force for progress in astronomy.

Scale. Services must often deal with data volumes, computational demands, and numbers of users beyond the capacity of a typical PC. Responding to a user request—or to the arrival of new data—can involve large amounts of computation. For example, the Argonne GNARE system searches periodically through DNA and protein databases for new and updated genomes and then computes and pub-

---

Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439, USA, and Department of Computer Science, University of Chicago, Chicago, IL 60637, USA. E-mail: foster@mcs.anl.gov

combustion community “services”

quantification & predictivity → discovery

- historic opportunity for combustion simulations & experimental data to be used for rigorous [illegible]
- enabling technologies
  - tera- to peta-scale computing produces multi-scale simulations of combustion process for the discovery of new technologies
  - high-resolution diagnostic measurement techniques produces unprecedented volumes of experimental combustion data
- simultaneous analysis of simulation data & experimental data

New information architectures enable new approaches to publishing and accessing valuable data and programs. So-called service-oriented architectures define standard interfaces and protocols that allow developers to encapsulate information tools as services that clients can access without knowledge of, or control over, their internal workings. Thus, tools formerly accessible only to the specialist can be made available to all; previously manual data-processing and analysis tasks can be automated by having services access services. Such service-oriented approaches to science are already being applied successfully, in some cases at substantial scales, but much more effort is required before these approaches are applied routinely across many disciplines. Grid technologies can accelerate the development and adoption of service-oriented science by enabling a separation of concerns between discipline-specific content and domain-independent software and hardware infrastructure.

service-oriented architecture

define standard interfaces and protocols

PrIMe
Process Informatics Model

PrIMe Navigation
Home
Announcements
Calendar
About PrIMe
History
Organization
Documents
Membership
Technical Support
Website [illegible]

Home
About PrIMe

PrIMe—Process Informatics Model—is a new approach for developing [illegible]
reaction systems [illegible] based on the scientific collabora[illegible]
and developing cyber infrastructure. The primary goals o[illegible]
the data and quantifying uncertainties, and assembling [illegible]
uncertainties to meet specific user requirements. The pri[illegible]
Repository, which is a repository of data provided by the [illegible]
evaluated data, and a set of computer-based tools to pr[illegible]
models. Two guiding principles of PrIMe are: open mem[illegible]

Computational Science Demands a New Paradigm

The field has reached a threshold at which better organization becomes crucial. New methods of verifying and validating

efficiently exploit the capacities of the increasingly complex computers. The prediction challenge is to use all that

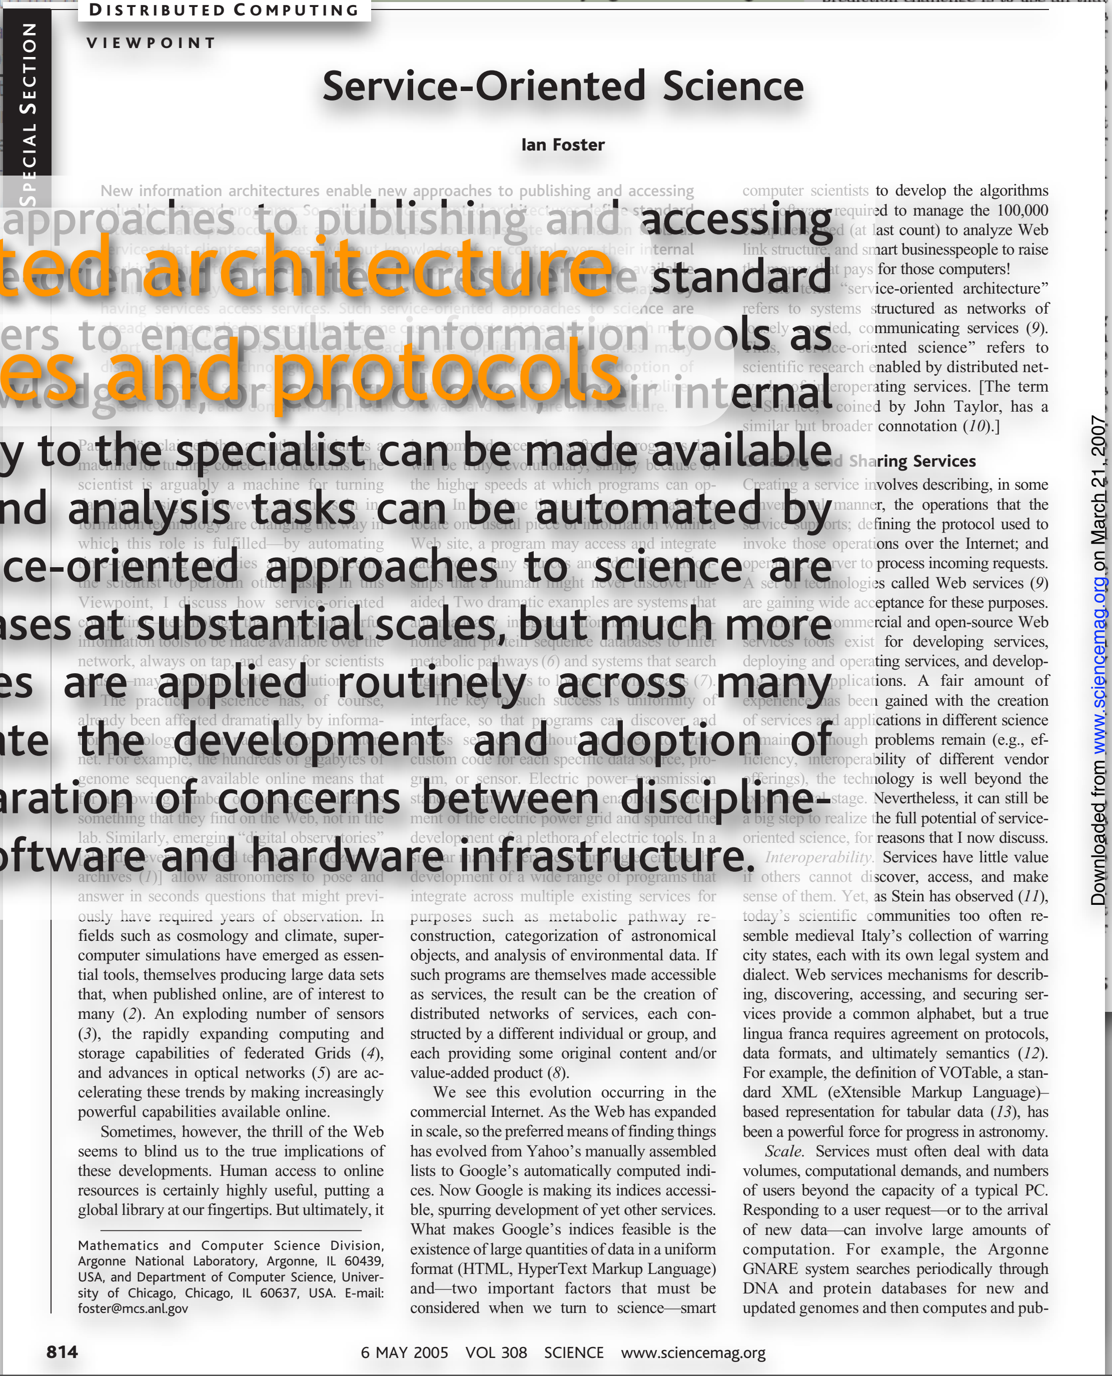

VIEWPOINT

Service-Oriented Science

Ian Foster

New information architectures enable new approaches to publishing and accessing valuable data and programs. So-called service-oriented architectures define standard interfaces and protocols that allow developers to encapsulate information tools as services that clients can access without knowledge of, or control over, their internal workings. Thus, tools formerly accessible only to the specialist can be made available to all; previously manual data-processing and analysis tasks can be automated by having services access services. Such service-oriented approaches to science are already being applied successfully, in some cases at substantial scales, but much more effort is required before these approaches are applied routinely across many disciplines. Grid technologies can accelerate the development and adoption of service-oriented science by enabling a separation of concerns between discipline-specific content and domain-independent software and hardware infrastructure.

Paul Erdös claimed that a mathematician is a machine for turning coffee into theorems. The scientist is arguably a machine for turning data into insight. However, advances in information technology are changing the way in which this role is fulfilled—by automating time-consuming activities and thus freeing the scientist to perform other tasks. In this Viewpoint, I discuss how service-oriented computing—technology that allows powerful information tools to be made available over the network, always on tap, and easy for scientists to use—may contribute to that evolution.

The practice of science has, of course, already been affected dramatically by information technology and, in particular, by the Internet. For example, the hundreds of gigabytes of genome sequence available online means that for a growing number of biologists, “data” is something that they find on the Web, not in the lab. Similarly, emerging “digital observatories” [already several hundred terabytes in dozens of archives (*1*)] allow astronomers to pose and answer in seconds questions that might previously have required years of observation. In fields such as cosmology and climate, supercomputer simulations have emerged as essential tools, themselves producing large data sets that, when published online, are of interest to many (*2*). An exploding number of sensors (*3*), the rapidly expanding computing and storage capabilities of federated Grids (*4*), and advances in optical networks (*5*) are accelerating these trends by making increasingly powerful capabilities available online.

Sometimes, however, the thrill of the Web seems to blind us to the true implications of these developments. Human access to online resources is certainly highly useful, putting a global library at our fingertips. But ultimately, it is automated access by software programs that will be truly revolutionary, simply because of the higher speeds at which programs can operate. In the time that a human user takes to locate one useful piece of information within a Web site, a program may access and integrate data from many sources and identify relationships that a human might never discover unaided. Two dramatic examples are systems that automatically integrate information from genome and protein sequence databases to infer metabolic pathways (*6*) and systems that search digital sky surveys to locate brown dwarfs (*7*).

The key to such success is uniformity of interface, so that programs can discover and access services without the need to write custom code for each specific data source, program, or sensor. Electric power–transmission standards and infrastructure enabled development of the electric power grid and spurred the development of a plethora of electric tools. In a similar manner, service technologies enable the development of a wide range of programs that integrate across multiple existing services for purposes such as metabolic pathway reconstruction, categorization of astronomical objects, and analysis of environmental data. If such programs are themselves made accessible as services, the result can be the creation of distributed networks of services, each constructed by a different individual or group, and each providing some original content and/or value-added product (*8*).

We see this evolution occurring in the commercial Internet. As the Web has expanded in scale, so the preferred means of finding things has evolved from Yahoo’s manually assembled lists to Google’s automatically computed indices. Now Google is making its indices accessible, spurring development of yet other services. What makes Google’s indices feasible is the existence of large quantities of data in a uniform format (HTML, HyperText Markup Language) and—two important factors that must be considered when we turn to science—smart computer scientists to develop the algorithms and software required to manage the 100,000 computers used (at last count) to analyze Web link structure, and smart businesspeople to raise the money that pays for those computers!

The term “service-oriented architecture” refers to systems structured as networks of loosely coupled, communicating services (*9*). Thus, “service-oriented science” refers to scientific research enabled by distributed networks of interoperating services. [The term “e-Science,” coined by John Taylor, has a similar but broader connotation (*10*).]

Creating and Sharing Services

Creating a service involves describing, in some conventional manner, the operations that the service supports; defining the protocol used to invoke those operations over the Internet; and operating a server to process incoming requests. A set of technologies called Web services (*9*) are gaining wide acceptance for these purposes. A variety of commercial and open-source Web services tools exist for developing services, deploying and operating services, and developing client applications. A fair amount of experience has been gained with the creation of services and applications in different science domains. Although problems remain (e.g., efficiency, interoperability of different vendor offerings), the technology is well beyond the experimental stage. Nevertheless, it can still be a big step to realize the full potential of service-oriented science, for reasons that I now discuss.

Interoperability. Services have little value if others cannot discover, access, and make sense of them. Yet, as Stein has observed (*11*), today’s scientific communities too often resemble medieval Italy’s collection of warring city states, each with its own legal system and dialect. Web services mechanisms for describing, discovering, accessing, and securing services provide a common alphabet, but a true lingua franca requires agreement on protocols, data formats, and ultimately semantics (*12*). For example, the definition of VOTable, a standard XML (eXtensible Markup Language)–based representation for tabular data (*13*), has been a powerful force for progress in astronomy.

Scale. Services must often deal with data volumes, computational demands, and numbers of users beyond the capacity of a typical PC. Responding to a user request—or to the arrival of new data—can involve large amounts of computation. For example, the Argonne GNARE system searches periodically through DNA and protein databases for new and updated genomes and then computes and pub-

Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439, USA, and Department of Computer Science, University of Chicago, Chicago, IL 60637, USA. E-mail: foster@mcs.anl.gov

Computational Science Demands a New Paradigm

The field has reached a threshold at which better organization becomes crucial. New methods of verifying and validating

efficiently exploit the capacities of the increasingly complex computers. The prediction challenge is to use all that

DISTRIBUTED COMPUTING

SPECIAL SECTION

VIEWPOINT

Service-Oriented Science

Ian Foster

Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439, USA, and Department of Computer Science, University of Chicago, Chicago, IL 60637, USA. E-mail: foster@mcs.anl.gov

combustion community "services"

quantification & predictivity → discovery

- historic opportunity for combustion simulations & experimental data to be used for rigorous validation & error quantification
  - tera- to peta-scale computing produces multi-scale simulations of combustion process for the discovery of new technologies
  - high-resolution diagnostic measurement techniques produces unprecedented volumes of experimental combustion data
- simultaneous analysis of simulation data & experimental data

New information architectures enable new approaches to publishing and accessing valuable data and programs. So-called service-oriented architectures define standard interfaces and protocols that allow developers to encapsulate information tools as services that clients can access without knowledge of, or control over, their internal workings. Thus, tools formerly accessible only to the specialist can be made available to all; previously manual data-processing and analysis tasks can be automated by having services access services. Such service-oriented approaches to science are already being applied successfully, in some cases at substantial scales, but much more effort is required before these approaches are applied routinely across many disciplines. Grid technologies can accelerate the development and adoption of service-oriented science by enabling a separation of concerns between discipline-specific content and domain-independent software and hardware infrastructure.

service-oriented architecture

define standard interfaces and protocols

encapsulate information tools as services

combustion community "services"

quantification & predictivity → discovery

- historic opportunity for combustion simulations & experimental data to be used for rigorous [illegible]
  - tera- to peta-scale simulations of combustion process for the discovery of new technologies
  - high-resolution diagnostic measurement techniques produces unprecedented volumes of experimental combustion data
- simultaneous analysis of simulation data & experimental data

New information architectures enable new approaches to publishing and accessing valuable data and programs. So-called service-oriented architectures define standard interfaces and protocols that allow developers to encapsulate information tools as services that clients can access without knowledge of, or control over, their internal workings. Thus, tools formerly accessible only to the specialist can be made available to all; previously manual data-processing and analysis tasks can be automated by having services access services. Such service-oriented approaches to science are already being applied successfully, in some cases at substantial scales, but much more effort is required before these approaches are applied routinely across many disciplines. Grid technologies can accelerate the development and adoption of service-oriented science by enabling a separation of concerns between discipline-specific content and domain-independent software and hardware infrastructure.

service-oriented architecture

define standard interfaces and protocols

encapsulate information tools as services

specialist tools made available to all

PrIMe—Process Informatics Model—is a new approach for developing

Computational Science Demands a New Paradigm

The field has reached a threshold at which better organization becomes crucial. New methods of verifying and validating

efficiently exploit the capacities of the increasingly complex computers. The prediction challenge is to use all that

---

DISTRIBUTED COMPUTING

SPECIAL SECTION

VIEWPOINT

Service-Oriented Science

Ian Foster

The practice of science has, of course, already been affected dramatically by information technology and, in particular, by the Internet. For example, the hundreds of gigabytes of genome sequence available online means that for a growing number of biologists, "data" is something that they find on the Web, not in the lab. Similarly, emerging "digital observatories" [already several hundred terabytes in dozens of archives (*1*)] allow astronomers to pose and answer in seconds questions that might previously have required years of observation. In fields such as cosmology and climate, supercomputer simulations have emerged as essential tools, themselves producing large data sets that, when published online, are of interest to many (*2*). An exploding number of sensors (*3*), the rapidly expanding computing and storage capabilities of federated Grids (*4*), and advances in optical networks (*5*) are accelerating these trends by making increasingly powerful capabilities available online.

Sometimes, however, the thrill of the Web seems to blind us to the true implications of these developments. Human access to online resources is certainly highly useful, putting a global library at our fingertips. But ultimately, it

Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439, USA, and Department of Computer Science, University of Chicago, Chicago, IL 60637, USA. E-mail: foster@mcs.anl.gov

The key to such success is uniformity of interface, so that programs can discover and access services without the need to write custom code for each specific data source, program, or sensor. Electric power–transmission standards and infrastructure enabled development of the electric power grid and spurred the development of a plethora of electric tools. In a similar manner, service technologies enable the development of a wide range of programs that integrate across multiple existing services for purposes such as metabolic pathway reconstruction, categorization of astronomical objects, and analysis of environmental data. If such programs are themselves made accessible as services, the result can be the creation of distributed networks of services, each constructed by a different individual or group, and each providing some original content and/or value-added product (*8*).

We see this evolution occurring in the commercial Internet. As the Web has expanded in scale, so the preferred means of finding things has evolved from Yahoo's manually assembled lists to Google's automatically computed indices. Now Google is making its indices accessible, spurring development of yet other services. What makes Google's indices feasible is the existence of large quantities of data in a uniform format (HTML, HyperText Markup Language) and—two important factors that must be considered when we turn to science—smart computer scientists to develop the algorithms and software required to manage the 100,000 computers used (at last count) to analyze Web link structure, and smart businesspeople to raise the money that pays for those computers!

The term "service-oriented architecture" refers to systems structured as networks of loosely coupled, communicating services (*9*). Thus, "service-oriented science" refers to scientific research enabled by distributed networks of interoperating services. [The term "e-Science," coined by John Taylor, has a similar but broader connotation (*10*).]

Creating and Sharing Services

Creating a service involves describing, in some conventional manner, the operations that the service supports; defining the protocol used to invoke those operations over the Internet; and operating a server to process incoming requests. A set of technologies called Web services (*9*) are gaining wide acceptance for these purposes. A variety of commercial and open-source Web services tools exist for developing services, deploying and operating services, and developing client applications. A fair amount of experience has been gained with the creation of services and applications in different science domains. Although problems remain (e.g., efficiency, interoperability of different vendor offerings), the technology is well beyond the experimental stage. Nevertheless, it can still be a big step to realize the full potential of service-oriented science, for reasons that I now discuss.

Interoperability. Services have little value if others cannot discover, access, and make sense of them. Yet, as Stein has observed (*11*), today's scientific communities too often resemble medieval Italy's collection of warring city states, each with its own legal system and dialect. Web services mechanisms for describing, discovering, accessing, and securing services provide a common alphabet, but a true lingua franca requires agreement on protocols, data formats, and ultimately semantics (*12*). For example, the definition of VOTable, a standard XML (eXtensible Markup Language)–based representation for tabular data (*13*), has been a powerful force for progress in astronomy.

Scale. Services must often deal with data volumes, computational demands, and numbers of users beyond the capacity of a typical PC. Responding to a user request—or to the arrival of new data—can involve large amounts of computation. For example, the Argonne GNARE system searches periodically through DNA and protein databases for new and updated genomes and then computes and pub-

PrIMe—Process Informatics Model—is a new approach for developing … and developing cyber infrastructure. The primary goals … the data and quantifying uncertainties, and assembling … uncertainties to meet specific user requirements. The pri… of data provided by the … tools to pr…

combustion community "services"

quantification & predictivity → discovery

- historic opportunity for combustion simulations & experimental data to be used for rigorous validation & error quantification
  - tera- to peta-scale simulations of [illegible] technologies
  - high-resolution diagnostic/ measurement techniques produces unprecedented volumes of experimental combustion data
- simultaneous analysis of simulation data & experimental data

New information architectures enable new approaches to publishing and accessing valuable data and programs. So-called service-oriented architectures define standard interfaces and protocols that allow developers to encapsulate information tools as services that clients can access without knowledge of, or control over, their internal workings. Thus, tools formerly accessible only to the specialist can be made available to all; previously manual data-processing and analysis tasks can be automated by having services access services. Such service-oriented approaches to science are already being applied successfully, in some cases at substantial scale, but much more effort is required before these approaches are applied routinely across many disciplines. Grid technologies can accelerate the development and adoption of service-oriented science by enabling a separation of concerns between discipline-specific content and domain-independent software and hardware infrastructure.

service-oriented architecture

define standard interfaces and protocols

encapsulate information tools as services

specialist tools made available to all

grid technologies

Computational Science Demands a New Paradigm

The field has reached a threshold at which better organization becomes crucial. New methods of verifying and validating

efficiently exploit the capacities of the increasingly complex computers. The prediction challenge is to use all that

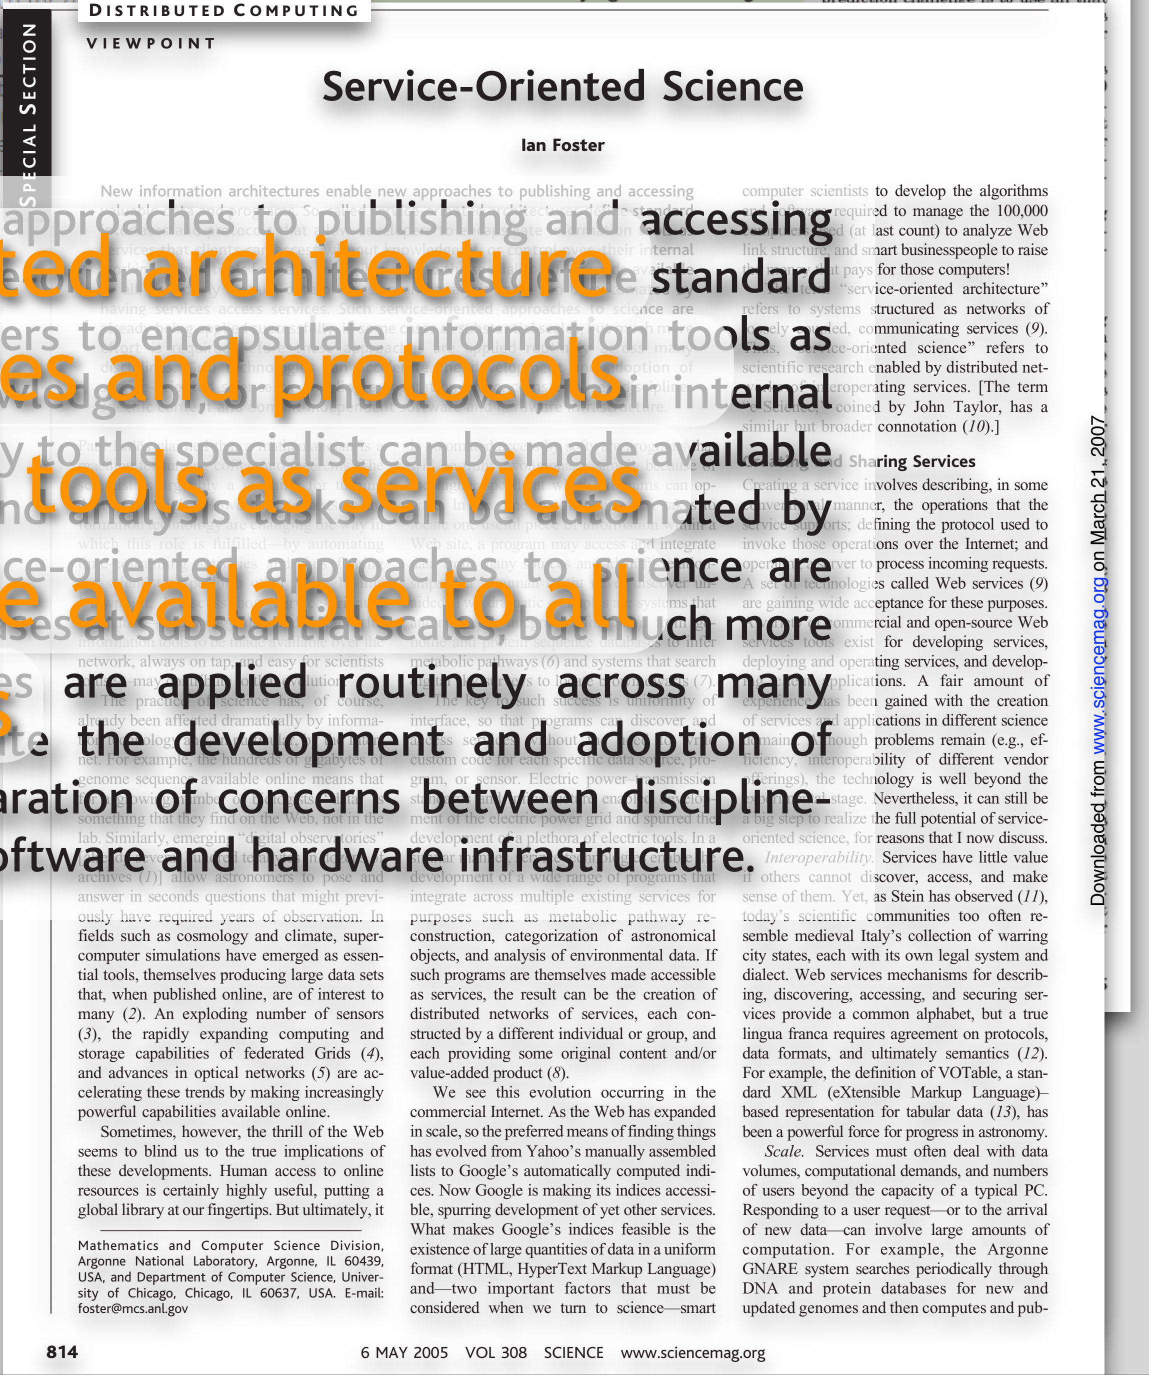

DISTRIBUTED COMPUTING

SPECIAL SECTION

VIEWPOINT

Service-Oriented Science

Ian Foster

… computer scientists to develop the algorithms and software required to manage the 100,000 computers used (at last count) to analyze Web link structure, and smart businesspeople to raise the money that pays for those computers!

The term "service-oriented architecture" refers to systems structured as networks of loosely coupled, communicating services (9). Thus, "service-oriented science" refers to scientific research enabled by distributed networks of interoperating services. [The term "e-Science," coined by John Taylor, has a similar but broader connotation (10).]

Creating and Sharing Services

Creating a service involves describing, in some conventional manner, the operations that the service supports; defining the protocol used to invoke those operations over the Internet; and operating a server to process incoming requests. A set of technologies called Web services (9) are gaining wide acceptance for these purposes. A variety of commercial and open-source Web services tools exist for developing services, deploying and operating services, and developing client applications. A fair amount of experience has been gained with the creation of services and applications in different science domains. Although problems remain (e.g., efficiency, interoperability of different vendor offerings), the technology is well beyond the experimental stage. Nevertheless, it can still be a big step to realize the full potential of service-oriented science, for reasons that I now discuss.

Interoperability. Services have little value if others cannot discover, access, and make sense of them. Yet, as Stein has observed (11), today's scientific communities too often resemble medieval Italy's collection of warring city states, each with its own legal system and dialect. Web services mechanisms for describing, discovering, accessing, and securing services provide a common alphabet, but a true lingua franca requires agreement on protocols, data formats, and ultimately semantics (12). For example, the definition of VOTable, a standard XML (eXtensible Markup Language)–based representation for tabular data (13), has been a powerful force for progress in astronomy.

Scale. Services must often deal with data volumes, computational demands, and numbers of users beyond the capacity of a typical PC. Responding to a user request—or to the arrival of new data—can involve large amounts of computation. For example, the Argonne GNARE system searches periodically through DNA and protein databases for new and updated genomes and then computes and pub-

… archives (1)] allow astronomers to pose and answer in seconds questions that might previously have required years of observation. In fields such as cosmology and climate, supercomputer simulations have emerged as essential tools, themselves producing large data sets that, when published online, are of interest to many (2). An exploding number of sensors (3), the rapidly expanding computing and storage capabilities of federated Grids (4), and advances in optical networks (5) are accelerating these trends by making increasingly powerful capabilities available online.

Sometimes, however, the thrill of the Web seems to blind us to the true implications of these developments. Human access to online resources is certainly highly useful, putting a global library at our fingertips. But ultimately, it

… If such programs are themselves made accessible as services, the result can be the creation of distributed networks of services, each constructed by a different individual or group, and each providing some original content and/or value-added product (8).

We see this evolution occurring in the commercial Internet. As the Web has expanded in scale, so the preferred means of finding things has evolved from Yahoo's manually assembled lists to Google's automatically computed indices. Now Google is making its indices accessible, spurring development of yet other services. What makes Google's indices feasible is the existence of large quantities of data in a uniform format (HTML, HyperText Markup Language) and—two important factors that must be considered when we turn to science—smart

---

Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439, USA, and Department of Computer Science, University of Chicago, Chicago, IL 60637, USA. E-mail: foster@mcs.anl.gov

814 6 MAY 2005 VOL 308 SCIENCE www.sciencemag.org

PrIMe

Process Informatics Model

PrIMe Navigation

- Home
- Announcements
- Calendar
- About PrIMe
  - History
  - Organization
  - Documents
  - Membership
  - Technical Support
  - Website Support
- PrIMe Da
- PrIMe Codes

Home

About PrIMe

PrIMe—Process Informatics Model—is a new approach for developing
and developing cyber infrastructure. The primary goals o
the data and quantifying uncertainties, and assembling t
uncertainties to meet specific user requirements. The pri
Depository, which is a repository of data provided by the
evaluated data, and a set of computer-based tools to pr
organization can register to participate in the project; ar
will be in the public domain.

combustion community "services"

quantification & predictivity → discovery

- "Combustion Data & Software Repository"
  - shared infrastructure
  - development tools
  - data and compute services
  - quality control (provenance)
  - interoperability
  - scale
  - standards (data archives, I/O, ...)
  - authentication procedures (certificates)
- an environment such that every combustion scientist would always want to contribute all his/her data and and all her/his software tools

Computational Science Demands a New Paradigm

The field has reached a threshold at which better organization becomes crucial. New methods of verifying and validating

efficiently exploit the capacities of the increasingly complex computers. The prediction challenge is to use all that

DISTRIBUTED COMPUTING

SPECIAL SECTION

VIEWPOINT

Service-Oriented Science

Ian Foster

New information architectures enable new approaches to publishing and accessing valuable data and programs. So-called service-oriented architectures define standard interfaces and protocols that allow developers to encapsulate information tools as services that clients can access without knowledge of, or control over, their internal workings. Thus, tools formerly accessible only to the specialist can be made available to all; previously manual data-processing and analysis tasks can be automated by having services access services. Such service-oriented approaches to science are already being applied successfully, in some cases at substantial scales, but much more effort is required before these approaches are applied routinely across many disciplines. Grid technologies can accelerate the development and adoption of service-oriented science by enabling a separation of concerns between discipline-specific content and domain-independent software and hardware infrastructure.

Paul Erdös claimed that a mathematician is a machine for turning coffee into theorems. The scientist is arguably a machine for turning data into insight. However, advances in information technology are changing the way in which this role is fulfilled—by automating time-consuming activities and thus freeing the scientist to perform other tasks. In this Viewpoint, I discuss how service-oriented computing—technology that allows powerful information tools to be made available over the network, always on tap, and easy for scientists to use—may contribute to that evolution.

The practice of science has, of course, already been affected dramatically by information technology and, in particular, by the Internet. For example, the hundreds of gigabytes of genome sequence available online means that for a growing number of biologists, "data" is something that they find on the Web, not in the lab. Similarly, emerging "digital observatories" [already several hundred terabytes in dozens of archives (1)] allow astronomers to pose and answer in seconds questions that might previously have required years of observation. In fields such as cosmology and climate, supercomputer simulations have emerged as essential tools, themselves producing large data sets that, when published online, are of interest to many (2). An exploding number of sensors (3), the rapidly expanding computing and storage capabilities of federated Grids (4), and advances in optical networks (5) are accelerating these trends by making increasingly powerful capabilities available online.

Sometimes, however, the thrill of the Web seems to blind us to the true implications of these developments. Human access to online resources is certainly highly useful, putting a global library at our fingertips. But ultimately, it is automated access by software programs that will be truly revolutionary, simply because of the higher speeds at which programs can operate. In the time that a human user takes to locate one useful piece of information within a Web site, a program may access and integrate data from many sources and identify relationships that a human might never discover unaided. Two dramatic examples are systems that automatically integrate information from genome and protein sequence databases to infer metabolic pathways (6) and systems that search digital sky surveys to locate brown dwarfs (7).

The key to such success is uniformity of interface, so that programs can discover and access services without the need to write custom code for each specific data source, program, or sensor. Electric power–transmission standards and infrastructure enabled development of the electric power grid and spurred the development of a plethora of electric tools. In a similar manner, service technologies enable the development of a wide range of programs that integrate across multiple existing services for purposes such as metabolic pathway reconstruction, categorization of astronomical objects, and analysis of environmental data. If such programs are themselves made accessible as services, the result can be the creation of distributed networks of services, each constructed by a different individual or group, and each providing some original content and/or value-added product (8).

We see this evolution occurring in the commercial Internet. As the Web has expanded in scale, so the preferred means of finding things has evolved from Yahoo's manually assembled lists to Google's automatically computed indices. Now Google is making its indices accessible, spurring development of yet other services. What makes Google's indices feasible is the existence of large quantities of data in a uniform format (HTML, HyperText Markup Language) and—two important factors that must be considered when we turn to science—smart computer scientists to develop the algorithms and software required to manage the 100,000 computers used (at last count) to analyze Web link structure, and smart businesspeople to raise the money that pays for those computers!

The term "service-oriented architecture" refers to systems structured as networks of loosely coupled, communicating services (9). Thus, "service-oriented science" refers to scientific research enabled by distributed networks of interoperating services. [The term "e-Science," coined by John Taylor, has a similar but broader connotation (10).]

Creating and Sharing Services

Creating a service involves describing, in some conventional manner, the operations that the service supports; defining the protocol used to invoke those operations over the Internet; and operating a server to process incoming requests. A set of technologies called Web services (9) are gaining wide acceptance for these purposes. A variety of commercial and open-source Web services tools exist for developing services, deploying and operating services, and developing client applications. A fair amount of experience has been gained with the creation of services and applications in different science domains. Although problems remain (e.g., efficiency, interoperability of different vendor offerings), the technology is well beyond the experimental stage. Nevertheless, it can still be a big step to realize the full potential of service-oriented science, for reasons that I now discuss.

Interoperability. Services have little value if others cannot discover, access, and make sense of them. Yet, as Stein has observed (11), today's scientific communities too often resemble medieval Italy's collection of warring city states, each with its own legal system and dialect. Web services mechanisms for describing, discovering, accessing, and securing services provide a common alphabet, but a true lingua franca requires agreement on protocols, data formats, and ultimately semantics (12). For example, the definition of VOTable, a standard XML (eXtensible Markup Language)–based representation for tabular data (13), has been a powerful force for progress in astronomy.

Scale. Services must often deal with data volumes, computational demands, and numbers of users beyond the capacity of a typical PC. Responding to a user request—or to the arrival of new data—can involve large amounts of computation. For example, the Argonne GNARE system searches periodically through DNA and protein databases for new and updated genomes and then computes and pub-

Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439, USA, and Department of Computer Science, University of Chicago, Chicago, IL 60637, USA. E-mail: foster@mcs.anl.gov

814 6 MAY 2005 VOL 308 SCIENCE www.sciencemag.org

Downloaded from www.sciencemag.org on March 21, 2007

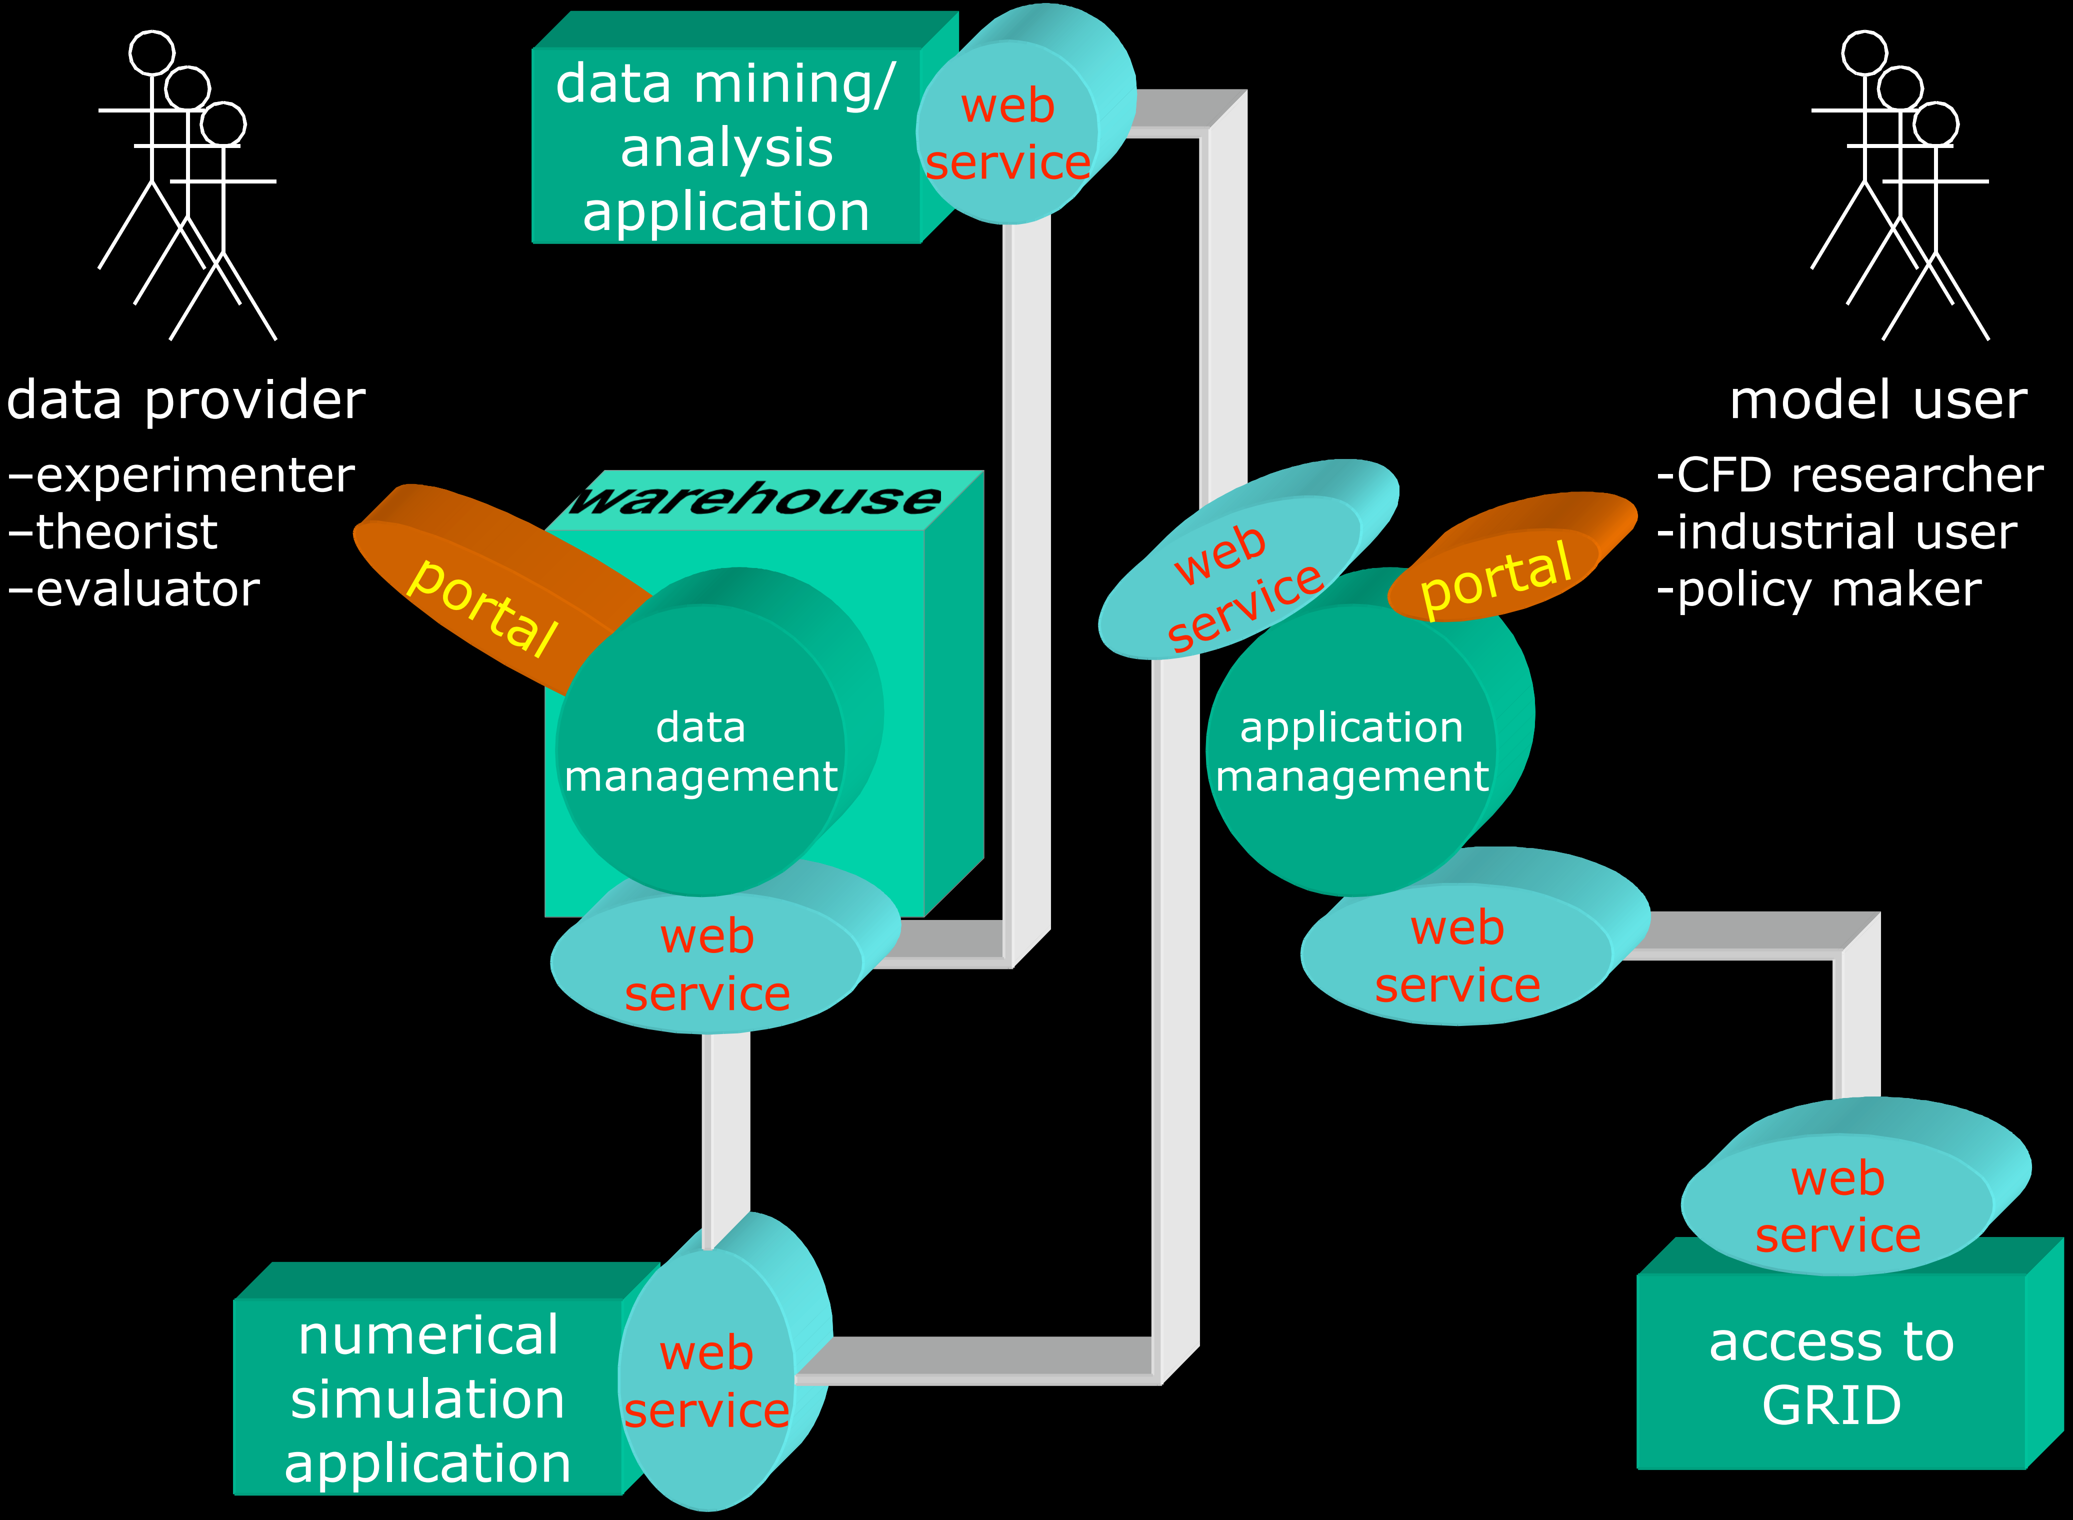

data provider

- –experimenter
- –theorist
- –evaluator

model user

- -CFD researcher
- -industrial user
- -policy maker

data mining/ analysis application

web service

warehouse

portal

data management

web service

numerical simulation application

web service

web service

portal

application management

web service

web service

access to GRID