Cyber-Based Combustion Science

Report on the NSF Workshop on Cyber-Based Combustion Science

April 19-20, 2006
National Science Foundation Headquarters, Arlington VA

Sponsored by the National Science Foundation (NSF)
Directorate for Engineering - Division of Chemical and Transport Systems
Thermal Systems Program (Combustion and Plasma Systems)

Authors: A. Trouvé, D. C. Haworth, J. H. Miller, L. K. Su & A. Violi
http://www.nsf-combustion.umd.edu/

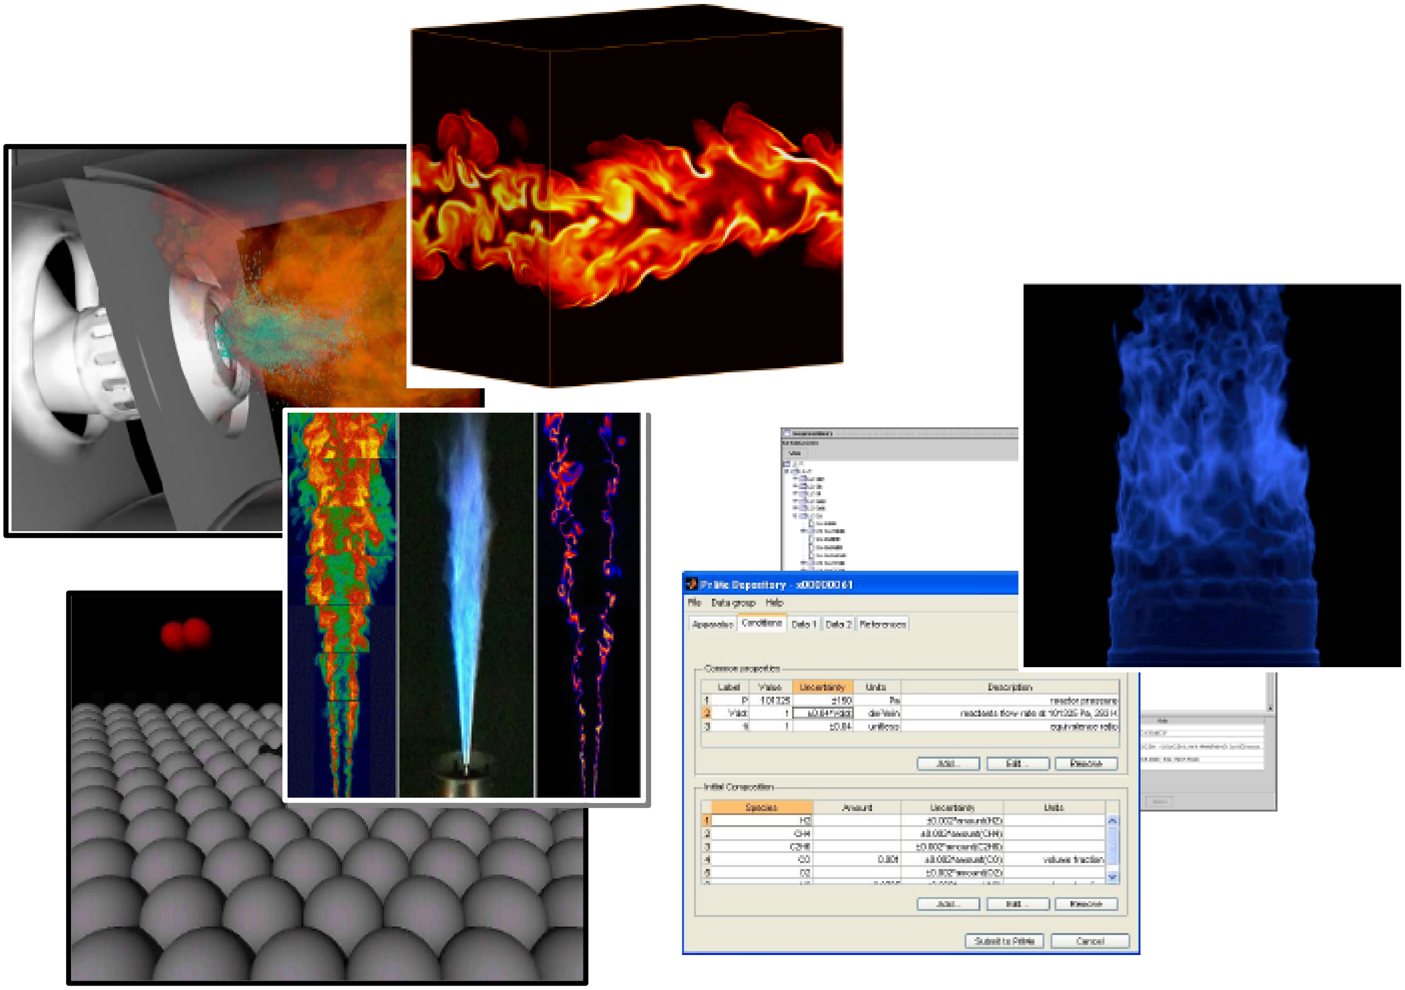

Cyber-Based Combustion Science

Report on the NSF Workshop on Cyber-Based Combustion Science

List of authors:

Prof. Arnaud Trouvé, *Chair*
Tel: (301) 405-8209 / Email: atrouve@eng.umd.edu
Department of Fire Protection Engineering
University of Maryland
1131L Glenn L. Martin Hall
College Park, MD 20742-3031 (USA)

Prof. Daniel C. Haworth, *Panel Chair*
Tel: (814) 863-6269 / Email: dch12@psu.edu
Department of Mechanical and Nuclear Engineering
The Pennsylvania State University
232 Research Building East
University Park, PA 16802 (USA)

Prof. J. Houston Miller, *Panel Chair*
Tel: (202) 994-7474 / Email: houston@gwu.edu
Department of Chemistry
George Washington University
Corcoran Hall, room 107
725 21st Street NW
Washington, DC 20052 (USA)

Prof. Lester K. Su, *Panel Chair*
Tel: (410) 516-8637 / Email: lsu@jhu.edu
Mechanical Engineering Department
The Johns Hopkins University
229 Latrobe Hall
34000 North Charles Street
Baltimore, MD 21218 (USA)

Prof. Angela Violi, *Panel Chair*
Tel: (734) 615-6448 / Email: avioli@umich.edu
Mechanical Engineering Department
University of Michigan
2150 G.G. Brown
2350 Hayward
Ann Arbor, MI 48109-2125 (USA)

Contents:

Notations:

AFOSR	Air Force Office of Scientific Research
ASC	Advanced Simulation and Computing
CBCS	Cyber-based combustion science
CFD	Computational fluid dynamics
CI	Cyber-infrastructure
DNS	Direct numerical simulation
DDDAS	Dynamic data-driven application systems
DOE	U.S. Department of Energy
HPC	High-performance computing
IT	Information Technology
ITR	Information Technology Research
LES	Large eddy simulation
NASA	National Aeronautics and Space Administration
NSF	National Science Foundation
ONR	Office of Naval Research
PC	Personal computer
RAS	Reynolds-averaged simulation
SciDAC	Scientific Discovery through Advanced Computing
TEAM	Training, Education, Advancement and Mentoring

1. Executive summary

The focus of the Cyber-Based Combustion Science (CBCS) Workshop is the cyber-infrastructure, defined herein as a set of base technologies for computation, storage, communication, and data processing services. The cyber-infrastructure (CI) has become a reality in recent years; its development has already impacted many, if not all, scientific fields; and there is a growing consensus that future progress in computer and network technologies, combined with the dissemination of these technologies on a global scale, will lead to further revolutionary changes in scientific methods and organizations. Mindful of this new world of possibilities, the National Science Foundation (NSF) is actively promoting the development of a nation-wide strategy to build a cyber-infrastructure in support of 21st century scientific discovery and engineering innovations [1-2].

The CBCS Workshop took place on April 19-20 2006 at NSF headquarters in Arlington, Virginia [3]. It brought together representatives of the U.S. combustion science community at large, primarily from academia, but also from national laboratories, government agencies, and industry. The participants were charged with reviewing current accomplishments, as well as identifying emerging opportunities and challenges associated with the cyber-infrastructure. More specifically, the objectives of the CBCS Workshop were to: (1) increase the visibility of the combustion science community by reviewing its strong history as a CI user/contributor, and by promoting its potential role as a leading application field for an enhanced cyber-infrastructure; (2) identify specific CI capabilities and needs in support of combustion science and combustion engineering education; (3) develop and articulate a community-wide vision for combustion science, that not only responds to the revolutionary changes in methods and organizations that will accompany the new cyber-infrastructure, but also actively contributes to shape these changes.

The Workshop therefore represents a collective contribution of the U.S. combustion science community to the wider ongoing CI effort at NSF. It is worth emphasizing that combustion science has already a long history of using the cyber-infrastructure. For instance, combustion science has become in recent years one of the largest users/beneficiaries of high-performance computing centers (through a number of national programs, including the ASC and SciDAC programs sponsored by the U.S. Department of Energy [4-5], and NSF’s ITR program [6]). Combustion science also features a high level of community awareness/readiness, that results from previous community-wide efforts based on early cyber-technologies (including the CHEMKIN library [7], the GRI-Mech project [8], the TNF workshop [9], and more recently Cantera [10] and the PrIMe initiative [11]).

The format of the 1.5 day Workshop consisted of a mix of invited talks and breakout group discussion sessions. Three themes were emphasized: (1) sub-community-specific, cyber-based combustion science applications (*i.e*. high-performance computing, and sensor-driven modeling); (2) community-wide, cyber-based combustion science applications (*i.e*. chemical libraries, and Web-enabled collaboratories); and (3) cyber-based combustion education. Lists of speakers and participants, and electronic copies of the invited talks can be found in Reference [3].

The list of recommendations that came out from the Workshop includes the following main points:

- The goal of achieving quantitative (predictive) capabilities for engineering-level simulations of combustion systems will require a cyber-infrastructure-enabled framework built around high-performance computing (HPC) and collaborative science infra-structures:
    - Tremendous progress has been made over the past fifteen years in the area of numerical modeling of combustion systems (including research-level fine-grained approaches and engineering-level coarse-grained approaches). Ongoing developments in numerical combustion [4-6], driven in part by continued access of combustion scientists to high-end HPC centers, should continue to be encouraged.
    - Much progress has also been made in the emerging area of chemical digital libraries. These libraries play a dual role as a data/software store and a collaboratory, and thereby function as a much needed coordination framework in the area of combustion chemistry, as well as an interface between combustion chemists and reacting flow researchers. Ongoing developments in chemical digital libraries [11] should also continue to be encouraged.
    - Progress is also needed to coordinate efforts in and across other sub-communities. The goal of achieving predictive combustion models will require a hierarchical approach to integrate multi-scale, multi-physics, multi-solver capabilities, and the adoption of systematic data and software verification, validation, and certification processes. This also implies an unprecedented integration of a wide range of scientists and engineers in a common framework, *i.e*. integration of different areas of combustion expertise – from nano-scales to engineering device scales –, integration of different skill sets – whether theoretical, experimental or computational –, and integration of different research interests – from fundamental sciences to practical applications. While some components of this framework have already been established [9], much remains to be done. Progress would best be pursued as community-wide projects coordinated across different funding agencies (with NSF, the DOE Office of Science, the Air Force Office of Scientific Research – AFOSR –, the Office of Naval Research – ONR –, and NASA as leading sponsors).
- The CI offers unprecedented opportunities to improve combustion education, including: a better integration of combustion science and combustion engineering; the promotion of combustion as a multi-scale/multi-physics discipline; and a better integration of combustion science with both computer and computational sciences. We recommend that combustion educators take advantage of the new CI-enabled possibilities and work to transform the combustion curriculum. Important ideas that should guide this transformation include: a renewed emphasis on establishing stronger pedagogical ties between fundamentals (*i.e*. chemistry, fluid mechanics, heat transfer) and applications (for instance, engine design); the promotion of combustion as a multi-scale discipline (from nano-scales to engineering device scales); the integration of data science (a computer science topic) and scientific computing (a computational science topic) into the combustion curriculum.

2. Background

2.1 Opportunities

Tremendous progress has been made during recent years in the development and large-scale dissemination of powerful computer and networking technologies. The status of the current and/or near-future cyber-infrastructure may be characterized by the following figures: state-of-the-art computational rates between 1 Tera- (10^{12}) and 1 Peta- (10^{15}) Flops; storage capacity up to 1 Peta-bytes; network bandwidth up to 1 Tera-bits per second. This technological progress has triggered in turn profound changes in the way science operates. For instance, the cyber-infrastructure can already be credited for the following series of transformations in the scientific landscape:

- The emergence of computational research as a new scientific approach, along with the traditional theoretical and experimental approaches. Computational research has been fueled in the past two decades by a remarkable growth (*i.e*. speed-up at an affordable cost) in high-performance computing (HPC) technology. There is a general consensus that this growth in HPC technology will be sustained in the foreseeable future, and will therefore allow for a continuing expansion of the domain of application of computational research. The expansion of computational research is taking place along different directions, including the development of new capabilities that provide a more complete (*i.e*. more realistic) description of previously established scientific problems, and the development of capabilities that provide a description of previously unexplored scientific problems (usually characterized by a high-level of complexity and/or a cross-disciplinary nature). Note that progress in computational research often relies on the ability to assemble multi-institutional, cross-disciplinary teams of applied mathematicians, computer scientists and application scientists (see for instance the SciDAC program [5]).

- The development of open-source cyber-based scientific data and software libraries. Until recently, lack of access to generic data/information was a major barrier often responsible for slowing down progress in scientific endeavors; technologically-driven open-source distribution of data/information has since then emerged as a powerful way to remove these barriers and empower entire scientific communities. It is noteworthy that the combustion science community has a long history of working with open-source libraries, a history that started with the pioneering years of CHEMKIN [7]. The CHEMKIN library is a software toolbox developed nearly thirty years ago by Sandia National Laboratories; the toolbox contains a suite of numerical solvers developed for detailed chemical kinetics calculations of academic problems corresponding to chemical reactors and laminar flame configurations. The solvers use thermodynamic and chemical kinetics data as inputs and provide information relevant to combustion modeling as outputs. Beyond the specific values of individual CHEMKIN calculations, the CHEMKIN library represents a valuable interface between two separate sub-communities in combustion science: the combustion chemists and the laminar/turbulent flame scientists; and one important legacy of CHEMKIN has been to bring together these two sub-communities, and establish standards and tools to allow a cross-exchange of information. More recent examples of similar community-wide efforts to construct open-source software or data libraries for combustion applications are the GRI-Mech project [8], Cantera [10] and the PrIMe initiative [11].

- The development of distance collaborations and the formation of new cyber-based communities. This is somewhat similar to the topic described in the previous paragraph, except that in addition to providing data/information, the objective here is to provide a common forum to bring together experts from geographically distributed locations and/or separate technical disciplines, and thereby to form new non-traditional cyber-based communities. The cyber-infrastructure provides powerful tools to reduce the effects of geographical or institutional barriers, and promote interactions between scientists on a global scale. An excellent example of a cyber-based combustion community is the turbulent nonpremixed flame (TNF) workshop [9]. The TNF workshop is an international collaboration between experimental and computational researchers, focused on fundamental issues associated with turbulence-chemistry interactions, and organized as a coupling interface between state-of-the-art experimental and numerical simulation databases. The GRI-Mech project [8] and the PrIMe initiative [11] are other examples of cyber-based community-building efforts in the combustion science community.

Further cyber-based transformations in the scientific landscape are also anticipated, for instance:

- The development of the grid, defined as a set of distributed and heterogeneous computational resources (super-computers, clusters, workstations) and data stores assembled into a transparent set of services via high-speed networks. The grid is expected to lead to a dramatic increase in cyber-performance, in terms of computational power, storage capacity and communication speeds. This increase will in turn benefit cyber-based science and engineering, and open the door to new applications, such as real-time simulations of complex systems, and/or the coupling of sensor technologies with high-performance computing capabilities [1].

2.2 Challenges

While the ongoing development of a cyber-infrastructure offers an exciting range of new possibilities, and is likely to accelerate the pace and scope of scientific discovery, there are also a number of challenges that result from new technical difficulties and/or cyber-induced transformations in scientific methods and organizations. The list of upcoming challenges includes:

- The need to provide the growing cyber-users community with a stable environment and manage the rapidly increasing volume of cyber-based data. This implies a certain level of standardization and the development of suitable (and community-specific) tools that can automatically process raw data (defined as disorganized collections of facts) and turn them into information (defined as organized collections of data), and then into knowledge (defined as insightful uses of information). The development of suitable data management tools is a key to control the cyber-traffic that may otherwise create data flooding and overwhelm the whole infrastructure.
- The need to reduce the growing gap between the current knowledge base of domain scientists and engineers, and the future needs of a cyber-based work environment. This implies developing stronger ties between computer science, computational science and domain science. For instance, combustion science and engineering students will need to be educated in the concepts and tools of the cyber-infrastructure. These include a variety of Information

Technology (IT) methods such as software design, data structures, data visualization, network architectures, in a distributed and heterogeneous (grid-like) environment, as well as a variety of computational science methods such as numerical methods, parallel software design, and parallel computing optimization. In addition, combustion science and engineering students will need to be exposed to an increasingly wider, cross-disciplinary, technical framework, for instance to the concepts and tools of a multi-scale approach to combustion. The framework of multi-scale combustion includes molecular dynamics occurring at nano-scales, laminar flame chemistry occurring at millimeter scales, turbulent flame dynamics occurring at centimeter scales, and the overall engineering systems performance characterized by scales on the order of tens of centimeters or more. In combustion as in other engineering fields, the development of a cyber-based infrastructure will lead to profound transformations in the work environment and professional practice; these transformations must be accompanied by corresponding changes in educational programs.

- The need for individual science and engineering communities to step up and adapt to the new cyber-induced transformations. Not all scientific fields and communities are equally prepared to meet the challenges of cyber-induced transformations. In this regard, combustion has many attractive features. For instance, one feature is a high level of technical readiness that results from the fact that combustion is a scientific field in which the basic governing equations are known (and are based on first principles rooted in thermodynamics, chemical kinetics, fluid mechanics and heat transfer). The difficulty in combustion science often comes from the complexity and nonlinear coupling between all the basic equations, a difficulty that remains a major handicap for theoretical analysis, but that may be overcome by a scientific computing approach. Consistent with this argument, combustion science has become in recent years one of the largest users/beneficiaries of high-performance computing centers [4-6]. In addition, another attractive feature of the combustion landscape is a high level of community awareness/readiness, that results from previous community-wide efforts based on early cyber-technologies (*i.e.* the CHEMKIN library [7], the GRI-Mech project [8], the TNF workshop [9], Cantera [10] and the PrIMe initiative [11]). Strengthened by these experiences, it may be argued that the combustion science community is a prime candidate to act as a leading application field for an enhanced cyber-infrastructure.

3. Workshop program

The CBCS Workshop took place on April 19-20 2006 at NSF headquarters in Arlington, Virginia [3]. It brought together representatives of the U.S. combustion science community at large, primarily from academia, but also from national laboratories, government agencies, and industry. The format of the 1.5 day Workshop consisted of a mix of invited talks and breakout group discussion sessions.

NSF's CI vision was described in three keynote presentations given by Drs. R. O. Buckius, Acting Assistant Director for Engineering [14], J. Munoz, Deputy Office Director in the Office of Cyberinfrastructure [15], and M. Heller, Program Director in the Office of Cyberinfrastructure [16].

The presentations and group discussions were otherwise organized around the following three general themes:

- *Sub-community-specific, cyber-based combustion science applications*. This first theme was focused on applications of CI-enabled computational science to combustion, and included the topics of high-performance computing and sensor-driven modeling:
    - **High-performance computing**. This topic was illustrated with presentations in the following areas: direct numerical simulation (DNS) [17], large eddy simulation (LES) [19], visualization [18], verification and validation of scientific software [20], and data and software repositories [20];
    - **Sensor-driven modeling**. This topic was illustrated with presentations in the following areas: dynamic data-driven application systems (DDDAS) [21], and sensor-driven grid computing [12,22]. The FireGRID project presented in Reference [22] is part of the e-Science program developed in the U.K.
- *Community-wide, cyber-based combustion science applications*. This second theme was focused on the CI-driven changes in methods and organizations that transform the combustion science community, and included the topics of digital libraries and collaborative science:
    - **Chemical data and software libraries**. This topic was illustrated with presentations in the following areas: CHEMKIN [7,23], the PrIMe initiative [11,24], detailed and reduced chemical kinetic mechanisms [25,27], data uncertainty [26], and multi-scale modeling [27,28]. While presented in the Session on education, the Cantera project [10,32] does also belong to the topic of chemical software libraries.
    - **Collaboratories**. This topic was illustrated with presentations in the following areas: general-purpose collaboratories (*i.e*. the CMCS project [13,30]) and topical collaboratories (*i.e*. the TNF workshop [9,29]). There is a significant cross-over between the topics of digital libraries and collaborative science, as illustrated by the PrIMe initiative [11,24].
- *Cyber-based combustion education*. This third theme was focused on the opportunities and needs for combustion education in a cyber-based environment:
    - **Education**. This topic was illustrated with presentations in the following areas: Web-based educational tools [31], curriculum changes [34], grid computing [31], chemical software libraries [32], and education of the general public and policymakers [33].

4. Theme 1: High-performance computing and sensor-driven modeling

This first theme (labeled "*sub-community-specific, cyber-based combustion science applications*") was focused on applications of CI-enabled computational science to combustion, and included the topics of high-performance computing (HPC) and sensor-driven modeling.

4.1 Oral presentations

Well-known combustion-related HPC/CI beneficiaries include DNS [17], LES [19], and computational chemistry. Much progress has been made over the past fifteen years in the area of numerical modeling of combustion systems, both at the research and engineering levels. However, despite the progress made to date, engineering-level simulation tools for combustion systems remain essentially non-predictive. This does not mean that these simulation tools cannot be quantitative or predictive, but rather means that they are only made quantitative after a costly calibration phase based on detailed comparisons with experimental data. This calibration phase is itself problem-specific and limited in scope, and therefore only intended to certify a predictive capability over a limited domain of application.

The lack of predictive capability of current engineering combustion models was emphasized during the presentations. While DNS has been applied successfully to generate new physical insight and develop/calibrate engineering models [17], further progress will require better organization within the combustion community to take advantage of the emerging CI. A disciplined hierarchical approach with data and software verification, validation, and certification at each stage will be required for the data intensive multi-scale/multi-physics/multi-solver integrated simulations that are on the horizon [19,20]; experimental data will also play an important role in this certification process. The establishment of a national combustion data and software repository was recommended [20]. The importance of community data sets, data management tools, and data processing tools (for instance, tools developed for feature recognition and extraction from large three-dimensional, time-dependent HPC data [18]) was emphasized by all presenters.

In addition to requiring a better integration of data and software, the goal of achieving predictive capabilities for engineering combustion models also requires a better integration of people (this point was echoed in the discussion on collaborative science, see Section 5). The wide gap that often separates combustion scientists from combustion engineers is responsible for a wasteful disconnect between high-end research programs developed at the basic science level and research needs expressed at the level of engineering applications. By promoting unprecedented levels of integration, the cyber-infrastructure can provide new ways to bridge this gap between different sub-communities (from communities of scientists who generate scientific ideas, to communities of engineers who generate design ideas and thereby express new scientific needs) and enlist these sub-communities into a common framework.

The last two presentations ([21-22]) in the Theme 1 Session were focused on the topic of sensor-driven modeling. Sensor-driven modeling is an emerging area of great promise that can benefit from the Peta-scale CI (and arguably depends on the successful development of a Peta-scale infrastructure). The two presentations in References [21] and [22] both dealt with systems to mitigate adverse effects of fires. The intended users are firefighters or other emergency responders in the field, not combustion scientists or engineers. Therefore, there is a premium on robustness, simple and unambiguous output, and user training. Super-real-time models are required, steered by real-time, multiple-source field data and sensor streams. Fast movement of large data sets, integration of disparate data and models, and hierarchical software design are particular issues. The importance of accounting explicitly for data uncertainties, and propagation of uncertainties through the models, was emphasized (the topic of data uncertainty was also the focus of the presentation in Reference [26]).

4.2 Break-out session

Four questions were posed to initiate the discussion: What has been the CI impact to date on combustion science and engineering? What are the combustion community's specific CI requirements? What new opportunities and high-growth areas does CI open for combustion science and engineering? What can the combustion community do for CI?

Much of the initial discussion centered on DNS, one of the "obvious" CI beneficiaries. DNS has been used successfully to formulate and calibrate engineering models for Computational Fluid Dynamics (CFD), *i.e*. Reynolds-averaged simulation (RAS) and/or LES. It provides high-fidelity data (albeit for idealized configurations and limited parameter ranges) that can be used in much the same way as experimental data. Although combustion DNS has been practiced for 20 years, the ability to perform three-dimensional simulations with somewhat realistic thermo-chemistry in parameter ranges (Reynolds number, Damköhler number) that approach those of practical interest (and/or approach those used in laboratory-scale experiments) is just becoming possible now, enabled by the emerging CI. Progress in this area will continue and will be driven by continued progress in HPC technologies and continued access of combustion scientists to high-end HPC centers [4-6].

It was noted that CFD-based combustion modeling (RAS and LES) is being used heavily and increasingly worldwide in all relevant industrial sectors including piston engines, gas turbines, burners, and boilers. This is in spite of known and significant shortcomings in currently available combustion models for engineering-level CFD; today's models generally cannot be considered to be predictive. Robust, truly predictive models for multi-phase turbulent combustion are urgently needed. Progress in this area will require a CI-enabled framework with HPC (*i.e*. DNS and LES, merged into a hierarchical approach with data and software verification, validation, and certification) and collaborative science. The latter refers to the integration of combustion scientists (theoretical, experimental and computational researchers) with combustion engineers (CFD practitioners and design engineers), as well as to the integration of combustion computational scientists, applied mathematicians, and computer scientists, as promoted for instance by the SciDAC program [5].

The management of huge data sets is recognized as a bottleneck in DNS, LES, and sensor-driven simulations. This problem includes archiving, compressing, documenting, moving, and mining/analyzing Tera-byte and Peta-byte data sets. In the case of sensor-driven simulations, it also includes processing data in real time. This problem is shared by many CI applications and is better addressed at a generic level; it may then be overcome by the application of generic IT solutions to combustion science needs.

Solver/model/data integrity is another area of increasing concern. Strong and divergent views were expressed concerning the appropriateness of making available open-source software for high-end simulations (*i.e*. DNS and LES solvers). The open-source chemical software libraries cited earlier are usually limited to zero- or one-dimensional solvers; with few exceptions, multi-dimensional solvers are not available as open-source software. Solver/model/data integrity, pedigree and structure for inter-operability and sharing (*i.e.* encapsulation, interfaces, standards) are increasingly important issues, particularly as the community moves towards coupled multi-scale/multi-physics/multi-solver integrated simulations. Progress in this area would be strongly facilitated by the establishment of a national combustion data and software repository [20].

The break-out session also included some discussion on the relative importance of access to larger and faster super-computers (a "bigger hammer" approach) versus algorithm improvements or other CI infrastructure issues (*i.e*. visualization, open source, documentation, integration of experimental data and theory, solver/model/data integrity). The general consensus was that greater benefits would be realized by addressing the algorithmic and infrastructure issues. Specific examples were cited for which algorithm improvements resulted in greater benefits than larger/faster computers. These include improved CFD algorithms for DNS [17] and LES [19,20], and various "chemistry acceleration" schemes (*i.e.* faster algorithms enhanced by advanced chemical reduction and/or tabulation techniques) that allow larger (*i.e.* more complete) chemical mechanisms to be accommodated in multi-dimensional simulations.

4.3 Summary and recommendations

Three key points were distilled from the discussion; these points are made below along with specific recommendations. These recommendations are primarily directed at funding agencies.

- CFD-based simulation tools for combustion systems have become ubiquitous in combustion engineering practice. The transformation of these tools from qualitative to quantitative (predictive) capabilities is a challenging endeavor that will require a CI-enabled framework built around HPC and collaborative science infra-structures. Such an effort would best be pursued as a community-wide project coordinated across different funding agencies (with NSF, the DOE Office of Science, the Air Force Office of Scientific Research – AFOSR –, the Office of Naval Research – ONR –, and NASA as leading sponsors).
- The framework required to achieve predictive capabilities in research- and engineering-level combustion models requires a hierarchical approach with standardized management of combustion science data (whether numerical or experimental) and software (including the possible development of open-source multi-dimensional laminar and turbulent combustion solvers). This requirement may be met by either expanding the role of emerging chemical data and software libraries (see Section 5 below) and/or establishing a new combustion data and software repository (as discussed in Reference [20]).
- The development of increasingly powerful HPC centers must be balanced by continued investments into algorithm developments (*i.e.* developments of numerical algorithms and chemical/physical models) and infrastructure developments (*i.e.* tools to collect community data/software, data management tools, data processing tools, *etc*).

It is worth emphasizing again that historically, combustion has had a strong connection with high-performance computing. The multi-scale, multi-physics, multi-disciplinary nature of combustion makes it particularly well suited to the emerging CI, since so much in combustion science depends on integration of "imported" knowledge. Combustion has benefited from, and has contributed to CI through numerous programs, including the Advanced Simulation and Computing (ASC) [4] and Scientific Discovery through Advanced Computing (SciDAC) [5] programs sponsored by DOE, and through the Information Technology Research (ITR) [6] program sponsored by NSF, among others.

5. Theme 2: Chemical data/software libraries and collaboratories

This second theme (labeled "*community-wide, cyber-based combustion science applications*") was focused on the CI-driven changes in methods and organizations, and included the topics of digital libraries and collaborative science.

5.1 Oral presentations

Eight invited presentations pertaining to Theme 2 were presented at the Workshop: four on chemical libraries [23-26], two on multi-scale modeling [27-28] and two on collaboratories [29-30]. The topics of digital libraries and collaboratories are strongly coupled, as demonstrated by the PrIME initiative [11,24].

Much of the progress made in combustion science has been brought by the unique tradition within the combustion community to explain combustion phenomena, from homogeneous chemical reactions to turbulent combustion, in a quantitative and predictive manner. Because even the simplest flame process involves coupled chemical reactions, heat release, and mass and heat transport, combustion simulations are inherently multi-scale and multi-physics. The obvious goal of multi-scale modeling is to predict macroscopic behavior of a process from first principles. To date, combustion simulation is certainly one of the most successful story of multi-scale modeling applied to complex nonlinear problems [27], and approaches developed in combustion science have been increasingly used by other scientific communities (for instance, in the areas of catalysis [28], chemical vapor deposition, material processing, atmospheric chemistry, astrochemistry, *etc*).

By its nature, multi-scale modeling is highly inter-disciplinary, with developments occurring independently across fields. There is therefore a need not only to extend our ability to communicate across several sub-disciplines, bringing people with different backgrounds and expertise together, but also to find a common language that can be used to exchange data and ideas. At the moment, several examples exist in the community of Collaboratories – virtual centers in which scientists can perform research, interact with colleagues, access remote instrumentation in virtual laboratories, access information in digital libraries, and share data and computational resources without regard to geographic location –. The concept was born in the mid-1990s out of the Department of Energy's desire to augment its capabilities through advanced computing and communications technologies, thereby promoting access to, and sharing of nationwide scientific resources.

The Collaboratory for Multiscale Chemical Science (CMCS) is developing an informatics-based approach to synthesizing multi-scale information and supporting interconnected collaborations across disciplines [13,30]. An open source multi-scale informatics toolkit has been developed by CMCS that addresses a number of issues, core to the emerging concept of knowledge grids including tracking of data provenance. The CMCS portal is in use by a number of pilot groups.

Another example of open collaboration among research groups is represented by the International Workshop on Measurement and Computation of Turbulent Nonpremixed Flames (TNF) [9,29]. The objectives of this initiative are to establish a Web-based library of well documented turbulent flames that can be used for numerical model validation, and to provide a framework for joint computational and experimental studies. The TNF workshop has proven to be an extremely successful framework for the development and promotion of the field of turbulent combustion.

A more recent example of community-wide effort to construct open-source chemical data libraries is the Process Informatics Model (PrIMe) initiative [11,24]. PrIMe introduces a new paradigm: a community-maintained “ever-green database”, and on-the-fly-generated models to answer specific questions. This approach relies on three major ingredients: (1) a systematic collection and proper organization of the scientific data; (2) the availability of scientific tools for analysis and processing of these data; (3) and the engagement of the entire scientific community in the data collection and analysis. The PrIMe infrastructure has thus two principal components: a data warehouse and a collection of tools. The tools currently in place or under development are of two general kinds: those enabling the collection, transfer, organization, display and mining of the data; and those enabling processing and analysis of the data along with assembly of the data into models.

An early example of chemical digital libraries and community-building effort is represented by CHEMKIN [7,23]. CHEMKIN was developed in the 1980s at the Combustion Research Facility of Sandia National Laboratories, to assist the incorporation of large chemical reaction mechanisms into numerical combustion solvers. The toolbox contains a suite of numerical solvers developed for detailed chemical kinetics calculations of problems ranging from chemical reactors to laminar flame configurations. Because of the modular architecture of the software, new capabilities can be easily implemented, enabling applications in new technical areas (including new technologies, such as catalytic combustion, fuel cells, *etc*). In the process of developing tools for combustion calculations with detailed chemistry models, CHEMKIN also established standards in the formatting of chemical kinetics data. Finally, as mentioned in Section 2.1, through its history, CHEMKIN has also played the role of a valuable collaborative interface between combustion chemists and reacting flow researchers.

The two presentations in References [25] and [26] provided additional examples of software libraries developed in support of data analysis and/or computational chemistry applications. Reference [25] discussed software packages for generation and/or reduction of chemical kinetic mechanisms (for instance, RIOT - Range Identification & Optimization Tool - is a mechanism reduction package made available on the Web through the CMCS portal), while Reference [26] focused on the topic of data and model uncertainty. Quantification of uncertainty is a major ingredient in the emerging verification, validation, and certification process proposed to build predictive combustion simulation capabilities (Section 4).

5.2 Break-out session

Four questions were posed to initiate the discussion: What has been the CI impact to date on combustion science and engineering? What are the combustion community's specific CI requirements? What new opportunities and high-growth areas does CI open for combustion science and engineering? What can the combustion community do for CI?

The initial discussion centered on the problem of managing large and often heterogeneous data sets. The participants discussed the need for chemical libraries to provide a stable environment and manage the rapidly increasing volume of cyber-based data. The difficulties found when managing large volumes of data in the context of chemical libraries are similar to those found when manipulating computational data in HPC applications (see Section 4.2). The need for a stable environment for users of chemical libraries also implies a certain level of standardization in data formats and software. These problems of capacity and standardization are shared by many CI applications and are therefore better addressed at a generic level; they may then be overcome by the application of generic IT solutions (developed outside of the combustion domain) to combustion science needs.

For instance, while different languages and data structures are used within the current CI infrastructure, there is a growing consensus that future developments are likely to focus on XML for data formats and Web-based tools for data exchange. As a result, the PrIMe initiative [11,24] has adopted XML as a way to represent descriptions of parameter uncertainties. One objective of PrIMe is to develop more systematic ways to collect data about parameter uncertainties and represent this information using modern relational database tools.

The discussion also turned to the topic of collaborative science. It is worth noting that collaboratories created as part of chemical library projects are becoming an increasingly important component of combustion science. These collaboratories function as a much needed coordination framework among combustion chemists, as well as an interface between combustion chemists and reacting flow researchers. For instance, one traditional role of combustion chemistry includes the development of detailed chemical kinetic models that can then be used by CFD modelers for multi-dimensional numerical simulations of laminar or turbulent flame problems. Presently, it is not unusual to find different research groups from around the world studying independently the same chemical system; it is also not unusual to find that these groups end up producing different chemistry models. These chemistry models end up different because they are based on different sets of experimental/theoretical data, because they correspond to different choices made in the evaluation of data uncertainty, or different choices made in the adoption of a verification/validation/certification process. The development of multiple chemistry models for a given problem generates in turn some confusion and uncertainty among potential users (*i.e.* CFD modelers). The role of chemical collaboratories (for instance, the GRI-Mech project [8], the PrIMe initiative [11], *etc*) is therefore to: (1) bring coordination and avoid a wasteful duplication of efforts; (2) establish a framework for integration of distributed knowledge/expertise and promote consensus-building and standardization; (3) allow the development of chemistry models approved by an entire scientific community; (4) act as a representative body for interactions with other technical sub-disciplines.

As discussed in Section 4.2, the role of collaboratories is expected to dramatically increase in the future, as the combustion science community moves towards the goal of achieving predictive capabilities for engineering combustion models. This goal requires an unprecedented level of integration of people and expertise, and in particular the integration of combustion scientists (theoretical, experimental and computational researchers) with combustion engineers (CFD practitioners and design engineers).

Another point that was made during the Workshop is the realization that as the scientists adapt their professional practice in response to the pressures and opportunities created by the cyber-infrastructure, there will be a need to adapt the performance metrics applied to researchers and educators. NSF and the cyber-based community will need to work together to identify new performance measures that are consistent with the team-based work environment promoted by the CI. The new vision of changing science based on the cyber-infrastructure will also have implications for the peer review and tenure/promotion processes within academia.

5.3 Summary and recommendations

Two key points were distilled from the discussion; these points are made below along with specific recommendations. These recommendations are primarily directed at funding agencies.

- The development of chemical data/software libraries should continue to be encouraged. These libraries play a dual role as a data/software store and a collaboratory, and function as a much needed coordination framework in the area of combustion chemistry, as well as an interface between combustion chemists and reacting flow researchers. They correspond to a single point of entry for otherwise distributed chemical data/information/knowledge, and produce state-of-the-art chemical kinetic models that are approved by an entire scientific community.
- The goal of achieving predictive capabilities for engineering combustion models will also require additional collaborative science infra-structures. There is a need to emulate the example and/or expand the scope of emerging chemical collaboratories, and to thereby enlist the entire combustion science and engineering community into a common framework. By thus promoting unprecedented levels of integration across areas of expertise (from nano-scales to engineering device scales), skill sets (theoretical, experimental or computational), and research interests (from fundamental sciences to practical applications), the cyber-infrastructure can provide the common forum and critical mass needed by the combustion community to achieve its goal. Such efforts would best be pursued as community-wide projects coordinated across different funding agencies (with NSF, the DOE Office of Science, the Air Force Office of Scientific Research − AFOSR −, the Office of Naval Research − ONR −, and NASA as leading sponsors).

6. Theme 3: Education

This third theme (labeled "*cyber-based combustion education*") was focused on the opportunities and needs for combustion education in a cyber-based environment.

6.1 Oral presentations

The education Session included presentations that encouraged a range of thinking about how the emerging CI might affect combustion education, and also explored different interpretations of the term "education." The issues under consideration can broadly be grouped into three categories: (1) how the cyber-infrastructure can augment existing approaches to combustion education; (2) how the cyber-infrastructure can suggest and facilitate new directions and structure for combustion education; (3) and how the combustion community might apply the unique capabilities of the cyber-infrastructure to expand the reach of its educational mission.

In the first presentation on the education theme [16], Dr. M. Heller, Program Director in the NSF Office of Cyberinfrastructure, described NSF's CI efforts both generally, and specifically in the education area. These efforts are well described in the 2003 NSF Report "Revolutionizing Science and Engineering through Cyberinfrastructure" [1], which partially motivated the creation of the Office of Cyberinfrastructure. Dr. Heller also mentioned the "2020 Science" Report issued by Microsoft Research, which touches on many relevant themes. Among these themes is the notion that computation should be embedded in all aspects of the scientific process, rather than remaining a separate sub-field in the way that theoretical and experimental work have. The Office of Cyberinfrastructure supports TEAM (Training, Education, Advancement and Mentoring) grants that specifically aim to prepare a "science and engineering workforce with the knowledge and skills needed to create, advance and take advantage of cyberinfrastructure over the long term" (from www.nsf.gov).

As may be expected, the topic of digital libraries already omnipresent in the Sessions on combustion research (see Sections 4 and 5) was also represented in the Session on combustion education. *thermofluids.net* [31] and Cantera [32] provided two different examples of building software libraries for combustion education. *thermofluids.net* is a Web-accessible and Web-executable tool for thermodynamic calculations, while Cantera is an open-source software library for chemically-reacting flow problems that can be used in conjunction with a variety of programming environments. An issue common to both *thermofluids.net* and Cantera (and which is certainly familiar in all areas of computational science) is that of ensuring broad compatibility for users featuring different hardware, operating systems and software packages. In the case of *thermofluids.net*, these issues of compatibility are sidestepped by taking advantage of universal Web standards, meaning that any user with an up-to-date Web browser can successfully use the service. Hence, *thermofluids.net* is a very effective tool for reference or classroom purposes. As a research tool, it is constrained somewhat in that its capabilities need to be fairly concretely defined on the provider side, which reduces its flexibility for research users.

Cantera's strengths are complementary to those of *thermofluids.net*. The barrier to entry for users is somewhat higher for Cantera, since the user needs to install Cantera for his/her specific implementation, then call on packages through a programming environment. While this makes it less optimal than *thermofluids.net* for users who are seeking a reference tool or an aid to classroom instruction, the limitations of Cantera for those users are precisely its strengths for research users. The flexibility of Cantera, in its ability to interface with different programming environments, is ideal for researchers whose computational routines may have been written with little concern for broad compatibility.

The usage patterns that have emerged for both *thermofluids.net* and Cantera highlight both the challenges and opportunities for the combustion cyber-infrastructure in the education area. *thermofluids.net* is intended for students without significant programming experience (for instance, undergraduate students) and educators looking for providing their students with an enhanced learning experience on a topic – thermodynamics – generally considered to be difficult. In contrast, Cantera is intended for more advanced users (graduate students, researchers or professional engineers) with programming savvy, and educators looking to challenge their students with projects that combine computational and combustion sciences.

The last two presentations on the education theme were made by two recent NSF CAREER grant awardees [33-34]. Prof. Sung emphasized the close, cyber-based integration of research and teaching, whereby research results are used to inform instruction [34]. He discussed his own extensive efforts in effecting this integration, which has the advantages of maintaining currency in the instructional material, emphasizing the applications of the classroom material, and broadening the pedagogical base.

Prof. Su, a former science and engineering fellow on Capitol Hill, steered the theme in a different direction by discussing the potential of the cyber-infrastructure to support combustion education efforts directed at the general public and policymakers [33]. Specifically, Prof. Su identified three challenges faced by the combustion community. First, there is a strongly negative and ill-informed perception of combustion among the general public. Combustion is often associated with environmental degradation and "old technology," which affects recruitment of new practitioners to the field, for example. The second, related challenge is the negative and ill-informed opinion of combustion among policymakers, which critically affects funding opportunities. Third, the academic and research study of combustion may be in the initial stages of marginalization that has been faced by fields such as nuclear and petroleum engineering, or traditional chemical engineering, where a disconnect emerges between fundamental research and the needs of industry (this topic is also discussed in Section 4.1). The cyber-infrastructure could help the combustion community address these challenges. The first two challenges could be addressed by using the CI to disseminate research results and also improve understanding of fundamental energy issues among the public, which would require, for example, that the combustion community become increasingly involved in such activities as secondary school curriculum reform, or advocacy before policymakers. The marginalization of combustion research in the eyes of policymakers and industry could be addressed by exploiting the CI to increase connections between different research areas (*i.e.* from fundamental sciences to practical applications), which would help the community articulate its goals and capabilities in a more cohesive fashion. Improved communication will also reduce redundancy of research efforts and help focus the allocation of intellectual and financial resources.

6.2 Break-out session

The discussion regarding the education theme focused on three broad questions: What are the objectives of building a cyber-infrastructure for combustion education? What skills do we expect future combustion engineers to need, and what will student training look like in the presence of a combustion cyber-infrastructure? How can the community integrate different efforts and assess pedagogical effectiveness, and who will do this work?

Question 1: What are the objectives of building a cyber-infrastructure for combustion education?

The participants envisioned three major purposes of a CI for combustion education. The first is to reinforce current pedagogical approaches to providing the relevant skills for combustion research and practice, for example by making the current state of knowledge more widely accessible. The second major purpose is to redefine pedagogical approaches. The participants envisioned that the CI could facilitate collaborative class environments, and reintroduce discovery to pedagogy through, for example, a “virtual combustion laboratory” that would allow online demonstrations and experiments. The final major purpose for a CI for combustion education is to broaden the participant base in combustion science. This purpose plays upon the previous two; the increased accessibility of combustion knowledge, and improved methods of pedagogy, would help to reduce the barriers of entry into the field, perhaps by allowing combustion to be introduced to students earlier in their curricula, *i.e.* in sophomore- or junior-level introductory thermodynamics or fluid mechanics.

Question 2: What skills do we expect future combustion engineers to need, and what will student training look like in the presence of a combustion cyber-infrastructure?

The Workshop participants agreed that one significant change effected by the advent of the combustion CI is the ready accessibility of enormous amounts of data, both experimental and computational. This will require working engineers to be comfortable with what might be called “data science”, which encompasses an understanding of IT tools for data management and data processing.

The CI will allow or require major changes in the combustion curriculum in addition to the pedagogical changes discussed in the context of Question 1. Four major changes are envisioned: (1) a better integration of combustion science and combustion engineering; (2) increased topical and pedagogical connections with related fields; (3) the promotion of combustion as a multi-scale/multi-physics discipline; (4) and a better integration of combustion science with both computer and computational sciences.

First, the CI offers an opportunity to reduce the gap between fundamental research and engineering practice, and will allow a better integration in the classroom of combustion fundamentals with application-oriented material. One possible change is to introduce the tools of combustion science and engineering together with the instruction in design methodology that is offered to engineering undergraduates. While many of the Workshop participants are interested in the scientific aspects of combustion, it is well-understood that combustion applications are firmly in the engineering realm, and that combustion must also be understood as one component in an engineering system, and a solution to a technological problem (for instance, producing heat or power).

Second, a CI could help to emphasize the links between combustion and related applications/disciplines in engineering and the sciences, by facilitating easy access to relevant materials in those areas. A CI could also allow for collaborative programs that incorporate a diversity of experiences across fields related to combustion, including in industry, academic research and the classroom. In that case, the CI would allow instructors and mentors working in the different fields to coordinate efforts and share teaching materials and information.

Third, the CI offers an opportunity to redefine combustion as a multi-scale/multi-physics discipline. Students will need to be exposed to an increasingly wider, cross-disciplinary, technical framework, for instance to the concepts and tools of a multi-scale approach to combustion, *i.e.* molecular dynamics occurring at nano-scales, laminar flame chemistry occurring at millimeter scales, turbulent flame dynamics occurring at centimeter scales, and the overall engineering systems performance characterized by scales on the order of tens of centimeters or more. The CI would enable the integration of these disparate technical areas by simplifying the handling of large amounts of digital information from different disciplines, and by simplifying the coordination of instruction in the different topics.

Finally, important changes will be driven by the need to prepare future engineers for an evolving cyber-based environment. This means developing stronger ties between domain (combustion) science and both computer and computational sciences. For instance, engineering students will need to become CI-literate, *i.e.* will need to be educated in the concepts and tools of a variety of IT methods such as software design, data structures, data visualization, and network architectures in a distributed and heterogeneous (grid-like) environment, and will need to be versed in a variety of computational science topics such as numerical methods, parallel software design, parallel computing optimization, and grid computing.

Question 3: How can we integrate different educational efforts and assess pedagogical effectiveness?

One suggestion put forth for assessing effectiveness was to track the usage patterns of on-line materials, then to tailor the development of these materials to reflect demand. The assumption in this case is that the community would eventually, and unprompted, settle on optimal use and configuration of these on-line materials. Partly, this "invisible hand" approach is motivated by necessity, given that it is impractical to envision the creation of a centralized body for organizing combustion education from the top down. This point immediately suggested two others, which the Workshop did not have time to address: first, who, or what group, will, even informally, take on the responsibility of the integration and assessment efforts, and second, how can we ensure the continued currency of pedagogical tools in the future?

6.3 Summary and recommendations

Two key points were distilled from the discussion; these points are made below along with specific recommendations. These recommendations are primarily directed at the combustion science and education community:

- The CI offers unprecedented opportunities to improve combustion education, including: a better integration of combustion science and combustion engineering; the promotion of combustion as a multi-scale/multi-physics discipline; and a better integration of combustion science with both computer and computational sciences. We recommend that combustion educators take advantage of the new CI-enabled possibilities and work to transform the combustion curriculum. Important ideas that should guide this transformation include: a renewed emphasis on establishing stronger pedagogical ties between fundamentals (*i.e*. chemistry, fluid mechanics, heat transfer) and applications (for instance, engine design); the promotion of combustion as a multi-scale discipline (from nano-scales to engineering device scales); the integration of data science (a computer science topic) and scientific computing (a computational science topic) into the combustion curriculum.
- The discussion among the Workshop participants of the potential of the CI to affect public and governmental perceptions of combustion resulted in a clear consensus that the Combustion Institute, as the *de facto* representing body for combustion researchers, was ideally placed to take the lead in public policy matters. Various other professional technical societies (including the American Physical Society, American Chemical Society, American Society of Mechanical Engineers, *etc*) were cited for their extensive involvement in government affairs at both the federal and state levels. The participants recommended that the Combustion Institute look into the possibility of establishing a public policy arm, whether internal to the Institute, or contracted out to an external body.

7. Conclusion

As suggested above and noted throughout the Workshop, it was the opinion of the majority of the Workshop participants that the cyber-infrastructure provides an exceptional opportunity for the combustion science community. In the course of discussing this opportunity, the goal of building and promoting the CI could not be separated from the goal of defining the critical position of combustion science in our nation’s economic vitality. By its nature our discipline is highly interdisciplinary and a broad range of scientific and engineering skill sets are required to maintain progress. The combustion science community is now shifting its focus from fossil fuel utilization technologies to new hybrid, energy extraction systems that may involve chemical–tailoring of fuels, multi-phase catalytic chemistry and electrochemical reactors. Thus new expertise will be called upon to bring these technologies to market. At the same time, stakeholders in the general population, ranging from K-12 students to political leaders, will need to develop a realistic vision of energy needs for the next several decades, and to appreciate that combustion will continue to play a dominant role in our country's economic vitality for several more generations. Perhaps our greatest challenge lies in attracting the best young minds to professional roles in combustion science and engineering. That may best be accomplished by reminding the broader scientific community of the tremendous technological advancements made in the name of combustion technology in fields as diverse as supercomputing and advanced laser diagnostics. The cyber-infrastructure will play a role in this educational process.

The cyber-infrastructure also provides a unique environment to foster wide-ranging collaborations among scientists. In fact, as the Workshop highlighted, combustion researchers have already demonstrated the value of distributed high-performance computing, well-defined laboratory benchmarks, and the use of multi-platform data/programming objects in international scientific partnerships. Perhaps more than many scientific communities, we understand the promise, but also the current limitations in implementation of the CI. Our goal is to develop the ability to model length/mass scales ranging from nano to kilo/mega, and to seamlessly capture complexity from quantum through continuum to mechanical. In all three "realms" of technology (hardware, software, and experimentation), we are not there yet. Even within the more traditional combustion paradigm, in which we wish to describe both chemical and fluid mechanical complexity, our research tools fail to fully capture the physical world. Figure 1 illustrates this disparity between scientific approaches, whether computational or experimental, and real-world applications. A great deal of our discussion, as reflected in the material above, was focused on moving technology towards solving real-world, complex problems, and towards connecting scientists and engineers who traditionally have focused more narrowly on separate regions in the maps of Figure 1.

The dialog that we began pointed out some roadblocks ahead. For example, as noted above, language barriers exist at the most basic level between scientists and engineers. Communication will grow more, not less, complex as we reach out to biotechnologists, electrochemists, *etc*, for assistance in the development of next generation energy systems. An efficient cyber-infrastructure will demand formats and standards for data exchange. Even the simplest questions (for instance, how do we define a molecule?), have not been fully fleshed out. There will always be tension between simplicity and functionality. Storing, accessing, and mining Tera-bytes of data are problems that we will share with all scientific disciplines in CI. There are also concerns about disenfranchising valued members of the community. Although there will be problems uniquely suited to solution by the cyber-infrastructure, there will clearly always be value in single PI "sandbox" projects that do not neatly fit into the concerted "Big Science" models that the cyber-infrastructure will likely promote.

Despite these concerns, the combustion science community will not only embrace the emerging cyber-infrastructure, but will likely lead the way in demonstrating its value. Several goals have been articulated during the Workshop; these goals correspond to an emerging community-wide vision for combustion science in a fast-evolving cyber-based environment. Two of these goals come out as the dominant themes of the Workshop: the goal of achieving quantitative (predictive) capabilities for engineering-level simulations of combustion systems; and the goal of redefining combustion as a multi-scale/multi-physics discipline. We recommend that NSF and other funding agencies support the combustion science community in its efforts to achieving those goals.

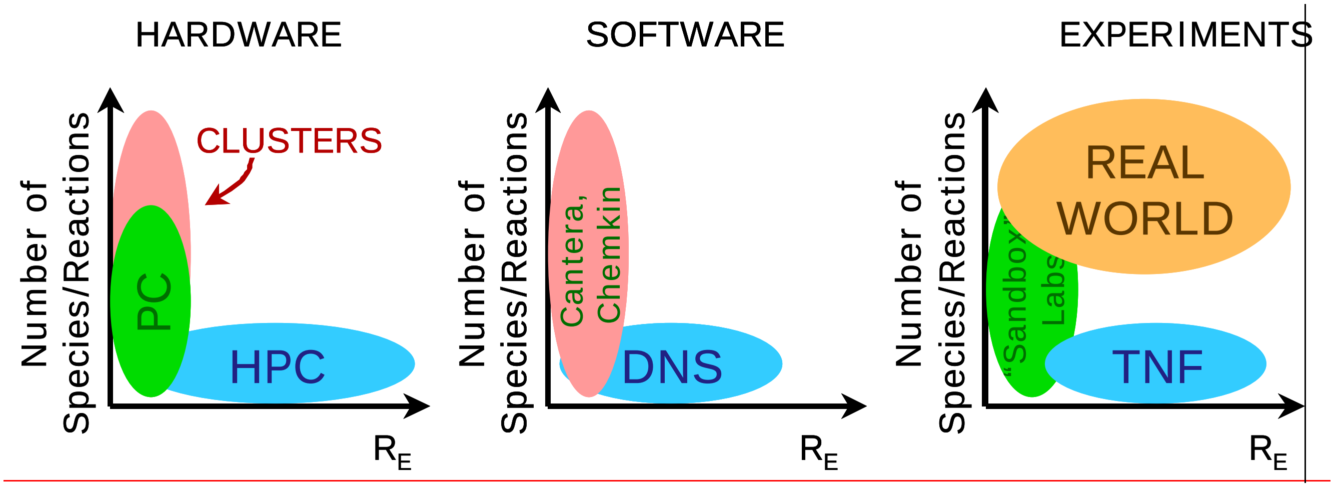

Figure 1: Illustration of chemical and fluid mechanical complexity in combustion science problems. R_E may be interpreted as a flow Reynolds number and represents flow complexity.

Acknowledgements

This Workshop was sponsored by the National Science Foundation (Directorate for Engineering, Division of Chemical and Transport Systems, Thermal Systems Program, Combustion and Plasma Systems Program). Support and encouragements from Dr. Linda G. Blevins, the NSF Program Director for Combustion Science (Thermal Systems Program, Combustion and Plasma Systems Program), are gratefully acknowledged. Additional encouragements and logistic support from Dr. Michael W. Plesniak, NSF Program Director (Fluid and Particle Processes Program, Fluid Dynamics and Hydraulics Program), are also gratefully acknowledged.

8. References

General References

[1] "Revolutionizing Science and Engineering through Cyberinfrastructure", *Report of the National Science Foundation Blue-Ribbon Advisory Panel on Cyberinfrastructure*, 2003, http://www.communitytechnology.org/nsf_ci_report, also available at http://www.cise.nsf.gov/sci/reports/atkins.pdf.

[2] "Identifying Major Scientific Challenges in the Mathematical and Physical Sciences and their Cyberinfrastructure Needs", *National Science Foundation Workshop*, 2004, http://www.nsf.gov/attachments/100811/public/CyberscienceFinal4.pdf.

[3] "A NSF Workshop on Cyber-Based Combustion Science", *National Science Foundation Workshop*, 2006, http://www.nsf-combustion.umd.edu.

[4] "Advanced Simulation and Computing (ASC)", *Department of Energy, National Nuclear Security Administration*, http://www.sandia.gov/NNSA/ASC/.

[5] "SciDAC: Scientific Discovery through Advanced Computing", *Department of Energy, Office of Science*, http://www.scidac.org.

[6] “Information Technology Research program (ITR)”, *National Science Foundation*, http://www.itr.nsf.gov.

[7] CHEMKIN, http://www.ca.sandia.gov/chemkin/.

[8] GRI-Mech, http://www.me.berkeley.edu/gri_mech/.

[9] “International Workshop on Measurement and Computation of Turbulent Nonpremixed Flames (TNF Workshop)”, http://www.ca.sandia.gov/TNF.

[10] Cantera. http://www.cantera.org/.

[11] “Process Informatics Model (PrIMe)”, http://primekinetics.org.

[12] FireGrid, http://www.firegrid.org/.

[13] “Collaboratory for Multi-scale Chemical Science (CMCS)”, http://cmcs.org/.

Workshop References

Presentations made at the *NSF Workshop on Cyber-Based Combustion Science*, April 19-20, 2006, Arlington, VA. Presentations are available at http://www.nsf-combustion.umd.edu/.

[14] Buckius, R. O.”NSF directorate for engineering”.

[15] Munoz, J., “Cyberinfrastructure and office of cyberinfrastructure (OCI)”.

[16] Heller, M., “Cyber-enabled combustion science and engineering education”.

[17] Chen, J. H., “Cyberinfrastructure requirements for terascale direct numerical simulation and modeling of turbulent combustion”.

[18] Ma, K.-L., “Visualization technologies for combustion science”.

[19] Pitsch, H., “What would I do if I had a really big computer…”.

[20] Smith, P., “High performance computing: Simulation – verification – validation – innovation”.

[21] Mandel, J., “Coupled weather-wildfire modeling driven by sensor and image data”.

[22] Torero, J. L., “Sensor driven computations for emergency response: FireGRID”.

[23] Kee, R. J., “Chemically reacting flow modeling and the CHEMKIN experience”.

[24] Frenklach, M., “Process informatics for combustion science”.

[25] Green, W. H., “RIOT and a vision for predictive modeling of combustion”.

[26] McRae, G. J., “Uncertainty propagation in combustion kinetics – Implication for database design and computation”.

[27] Wang, H., “Chemical libraries: multi-scale modeling”.

[28] Vlachos, D., “The emerging field of multi-scale simulation. Relation to cyber-infrastructure and educational needs”.

[29] Barlow, R. S., “Overview of the TNF workshop (International Workshop on Measurement and Computation of Turbulent Nonpremixed Flames)”.

[30] Myers, J. D., “Cyberenvironments for research and education: Collaboratory for Multi-scale Chemical Science (CMCS) and beyond”.

[31] Bhattacharjee, S., “A cyber-based collaborative framework for thermodynamic education and research”.

[32] Goodwin, D., “Cantera”.

[33] Su, L. K., “Issues in cyber-based combustion education”.

[34] Sung, C.-J., “Cyber-infrastructure enabled integration of research and education”.

9. Appendix 1: List of Workshop participants

Robert Barlow Email: barlow@ca.sandia.gov	Sandia National Laboratories
Subrata (Sooby) Bhattacharjee Email: subrata@thermo.sdsu.edu	San Diego State University
Linda Blevins Email: linda.blevins@science.doe.gov	Department of Energy, Basic Energy Sciences
Harsha Chelliah Email: harsha@virginia.edu	University of Virginia
Jacqueline Chen Email: jhchen@sandia.gov	Sandia National Laboratories
Jean-Pierre Delplanque Email: delplanque@ucdavis.edu	University of California, Davis
Ofodike Ezekoye Email: dezekoye@mail.utexas.edu	University of Texas
Michael Frenklach Email: myf@me.berkeley.edu	University of California at Berkeley
Graham Goldin Email: gmg@fluent.com	Fluent Inc.
David Goodwin Email: dgoodwin@caltech.edu	California Institue of Technology
Jayavant P. Gore Email: gore@purdue.edu	Purdue University
William Green Email: whgreen@mit.edu	Massachusetts Institute of Technology
Daniel Haworth Email: dch12@psu.edu	Pennsylvania State University
Richard Hilderbrandt Email: Richard.Hilderbrandt@science.doe.gov	Department of Energy, Basic Energy Sciences
E. David Huckaby Email: huckaby@netl.doe.gov	DOE, National Energy Technology Laboratory
Hong Im Email: hgim@umich.edu	University of Michigan
Robert Kee Email: rjkee@MINES.EDU	Colorado School of Mines
Merrill King Email: merrill.king@nasa.gov	NASA

Michael Klassen Email: kcalise@csefire.com	Combustion Science & Engineering
Chung Law Email: cklaw@princeton.edu	Princeton University
Kwan-Liu Ma Email: ma@cs.ucdavis.edu	University of California at Davis
Jan Mandel Email: jmandel@math.cudenver.edu	University of Colorado
Jeffrey Manion Email: Jeffrey.manion@nist.gov	National Institute of Standards and Technology
Andre Marshall Email: awmarsh@eng.umd.edu	University of Maryland
Henry McDonald Email: henry-mcdonald@utc.edu	University of Tennessee
J. Thomas McKinnon Email: jmckinno@mines.edu	Colorado School of Mines
Gregory John McRae Email: mcrae@MIT.edu	Massachusetts Institute of Technology
J. Houston Miller Email: Houston@gwu.edu	George Washington University
Jim Myers Email: jimmyers@ncsa.uiuc.edu	National Center for Supercomputing Applications
Heinz Pitsch Email: H.Pitsch@stanford.edu	Stanford University
Michael Plesniak Email: mplesnia@nsf.gov	National Science Foundation
Ishwar Puri Email: ikpuri@vt.edu	Virginia Polytechnic Institute and State University
Larry Rahn Email: rahn@sandia.gov	Sandia National Laboratories
Sutanu Sarkar ssarkar@ucsd.edu	University of California at San Diego
David Schmidt schmidt@ecs.umass.edu	University of Massachusetts
Balu Sekar Email: balu.sekar@wpafb.af.mil	Air Force Research Laboratory
Philip Smith Email: philip.smith@utah.edu	University of Utah

Lester Su Email: lsu@jhu.edu	Johns Hopkins University
Peter Sunderland Email: pbs@umd.edu	University of Maryland
Chih-Jen Sung Email: cjs15@case.edu	Case Western Reserve University
Doug Talley Email: Douglas.Talley@edwards.af.mil	Air Force Research Laboratory
Jose Torero Email: j.torero@ed.ac.uk	The University of Edinburgh (U.K.)
Arnaud Trouve Email: atrouve@eng.umd.edu	University of Maryland
Stephen Turns Email: srt@psu.edu	Pennsylvania State University
Angela Violi Email: avioli@umich.edu	University of Michigan
Dionisios Vlachos Email: vlachos@udel.edu	University of Delaware
Anthony Vodacek Email: vodacek@cis.rit.edu	Rochester Institute of Technology
Hai Wang Email: haiw@usc.edu	University of Southern California
Phillip Westmoreland Email: westm@ecs.umass.edu	University of Massachusetts
Carole Womeldorf Email: womeldor@ohio.edu	Ohio University
P. K. Yeung Email: yeung@peach.ae.gatech.edu	Georgia Institute of Technology

10. Appendix 2: Workshop program (April 19-20, 2006)

Wednesday, April 19

7:30-8:15 am		*Breakfast* (provided)
8:15-8:30	(15’)	Introduction (R. Buckius, National Science Foundation)
8:30-8:45	(15’)	Welcome (A. Trouvé, University of Maryland)
8:45-9:00	(15’)	Introduction (Linda Blevins and Michael Plesniak, National Science Foundation)
9:00-9:30	(30’)	NSF vision for the Cyber-Infrastructure (Jose Munoz, National Science Foundation)

Theme 1: sub-community-specific cyber-based combustion science applications

Session Chair: A. Marshall (University of Maryland)

9:30-9:50	(20’)	High-performance computing (J. H. Chen, Sandia National Laboratories)
9:50-10:10	(20’)	Visualization (K.-L. Ma, University of California Davis)
10:10-10:30	(20’)	*Break*
10:30-10:50	(20’)	High-performance computing (H. Pitsch, Stanford University)
10:50-11:10	(20’)	High-performance computing (P. Smith, University of Utah)
11:10-11:30	(20’)	Sensor-driven modeling (J. Mandel, University of Colorado Denver)
11:30-11:50	(20’)	Sensor-driven modeling, FireGrid (J. Torero, University of Edinburgh, U.K)

Theme 2: community-wide cyber-based combustion science applications

Session Chair: H. Im (University of Michigan)

11:50-12:10	(20’)	Chemical Libraries, CHEMKIN (R. J. Kee, Colorado School of Mines)
12:10-13:10	(60’)	*Lunch* (provided)
13:10-13:30	(20’)	Chemical Libraries and collaborative science, PrIME (M. Frenklach, University of California Berkeley)
13:30-13:50	(20’)	Chemical Libraries, RIOT (W. H. Green, Massachusetts Institute of Technology)
13:50-14:10	(20’)	Chemical Libraries, uncertainty (G. J. McRae, Massachusetts Institute of Technology)
14:10-14:30	(20’)	Chemical Libraries, multi-scale modeling (H. Wang, University of Southern California)

14:30-14:50	(20')	Chemical Libraries, multi-scale modeling (D. Vlachos, University of Delaware)
14:50-15:10	(20')	Collaboratories, TNF Workshop (R. Barlow, Sandia National Laboratories)
15:10-15:30	(20')	Collaboratories, CMCS (J. D. Myers, National Center for Supercomputing Applications)
15:30-15:45	(15')	*Break*
15:45-16:45	(60')	Break-out sessions – Themes 1 and 2 Moderator for Theme 1: D. Haworth (Pennsylvania State University) Moderator for Theme 2: A. Violi (University of Michigan)
16:45-17:45	(60')	Joint session – Report from break-out sessions
17:45-18:45	(60')	*Break*
18:45-22:00		*Dinner (Matsutake Restaurant)*

Thursday, April 20

7:30-8:30 am	*Breakfast* (provided)

Theme 3: cyber-based combustion education

Session Chair: P. Sunderland (University of Maryland)

8:30-9:00	(30')	NSF vision for the Cyber-Infrastructure, engineering education (Miriam Heller, National Science Foundation)
9:00-9:20	(20')	Chemical Libraries, distributed grid-enabled applications (S. Bhattacharjee, San Diego State University)
9:20-9:40	(20')	Chemical Libraries, Cantera (D. G. Goodwin, California Institute of Technology)
9:40-10:00	(20')	Short presentations by recent NSF CAREER awardees: a cyberinfrastructure to support combustion Education including towards policymakers and the broader public (L. Su, the Johns Hopkins University), and cyber-based integration of Research and Education (J. Sung, Case Western Reserve University)
10:00-10:15	(15')	*Break*
10:15-11:00	(45')	Joint session – Theme 3: cyber-based combustion education Moderator for Theme 3: L. Su (The Johns Hopkins University)
11:00-12:00	(60')	Joint session – Articulation of a community-wide vision for cyber-based combustion science and recommendations to NSF Moderator for Joint Session: H. Miller (George Washington University)

12:00-12:15	(15')	Closing remarks (Linda Blevins and Michael Plesniak, National Science Foundation, and A. Trouvé, University of Maryland)
12:15-13:15	(60')	*Lunch* (provided)