2000 Annual Technical Report

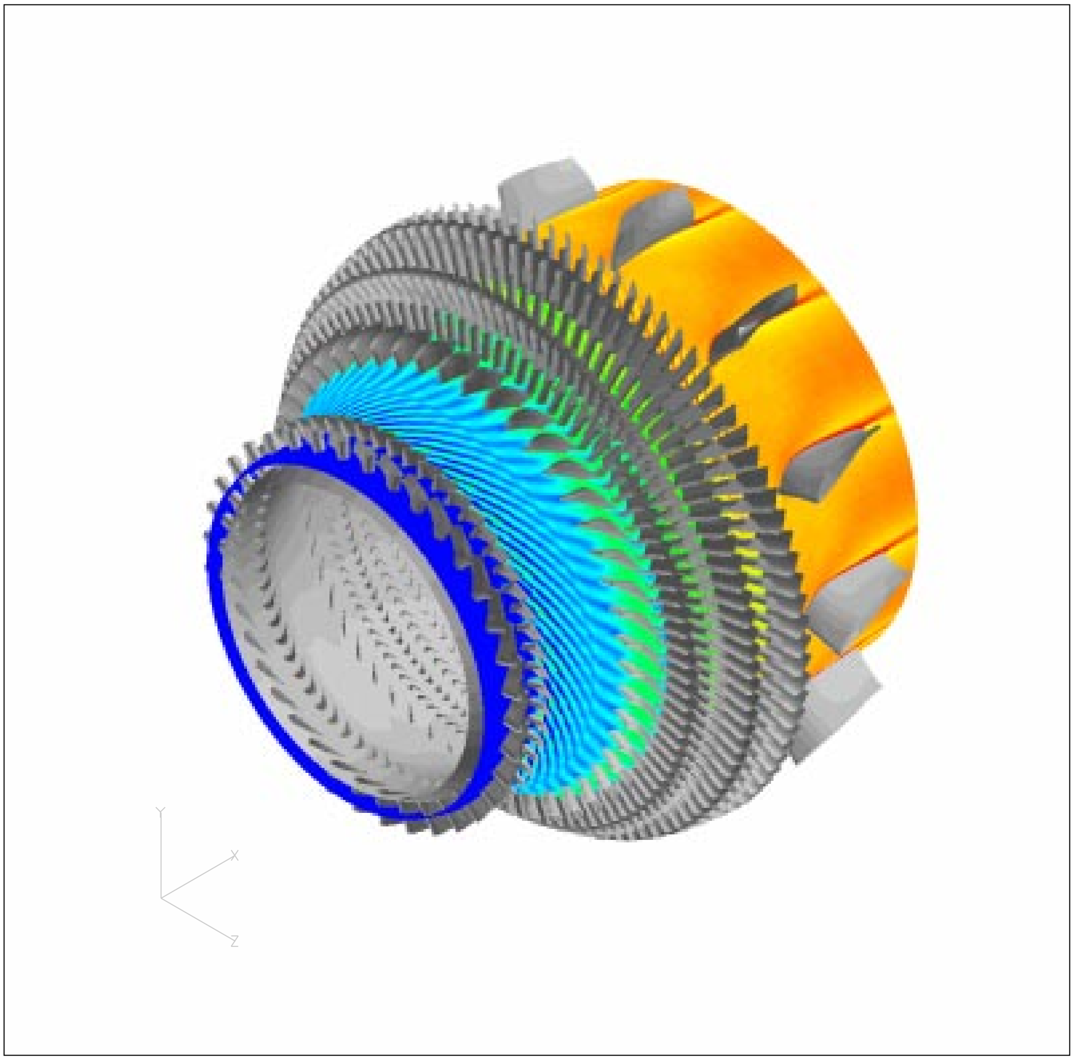

Center for Integrated Turbulence Simulations

Stanford University

Abstract

The Center for Integrated Turbulence Simulations was established at Stanford University in September, 1997, as one of five university centers in the Academic Strategic Alliances Program of the Department of Energy's Accelerated Strategic Computing Initiative. This report outlines the Center's technical work during the third year.

Contents

Chapter 1

Overview

1.1 ASCI and CITS

The Center for Integrated Turbulence Simulations (CITS) is one of five university research efforts established by the Department of Energy (DOE) as part of their Accelerated Strategic Computing Initiative (ASCI) under the Academic Strategic Alliances Program (ASAP). The purpose of the ASCI program is to develop advanced general-purpose technology for very large-scale multi-physics scientific computations. The DOE will use this simulation technology to maintain the safety and reliability of the U.S. nuclear weapons stockpile without resort to further physical testing. Each of the Alliances is working on an important non-weapons application that provides a framework for technical advances in this sort of computing. The powerful supercomputing capabilities in the DOE laboratories are being used by the Alliances for this work.

The CITS is building on Stanford's strengths in flow physics and computation to develop an integrated high-fidelity simulation technology for aircraft gas turbine engines. The major U.S. aircraft engine manufacturers are partners in this effort. They providing realistic geometries, data, experience, evaluations against their own flow simulation codes, and people. They are interested in seeing what this sort of envelope-pushing simulation could do for them. If the potential is significant, their interest would add further market push for the computing systems necessary for such simulations. The success of the CITS effort in engine simulation ultimately will be measured by its impact on the aircraft engine industry, and the involvement of industry with the CITS program is a good indication that there is potential for significant impact.

We believe that the future of large-scale scientific computing lies in the creative use of low-cost commodity computing components in large-scale, highly parallel computing systems. The commercial computer industry today builds server computers of up to 128 processors. While very capable at the data management and engineering simulation for which they are designed, these commercial machines are an order of magnitude slower and smaller in scale than the supercomputers required for ASCI simulations. Research in Stanford's Computer Systems Laboratory (CSL), conducted as part of the Stanford ASCI program, aims to augment this commercial technology with scalable computing technology that will enable the design of very large high-performance computers and the development of efficient programs for these machines.

The research in CSL spans the spectrum from hardware to operating systems and software tools. On the hardware side, a group in CSL is developing interconnect technology that will dramatically increase the global bandwidth of large-scale parallel computers. Operating systems research uses virtual machine technology to apply commercial operating systems to reliable, scalable operation of large-scale machines. Research on compilers, simulation tools, static program checking tools, and visualization enables more rapid development of efficient codes for large-scale machines.

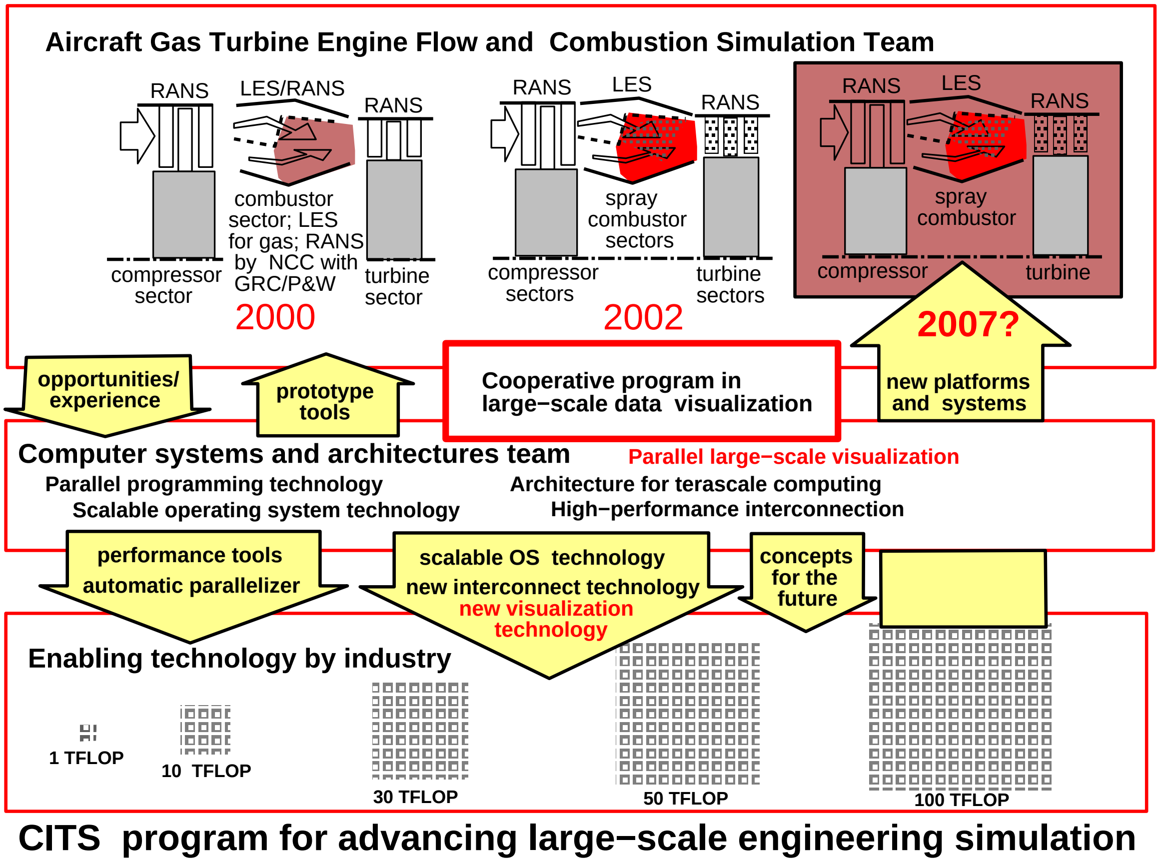

Figure 1.1: Overall CITS roadmap

The coupling between the CSL team and the engine simulation team has increased as the work progressed. Some of the programming tools developed in CSL are now being used on an exploratory basis by the engine simulation team. Needs identified in the engine simulation are beginning to influence research in the CSL. And a promising new cooperative program in visualization of large data bases in taking shape.

However, the primary impact of the CSL research for ASCI will be on the industry that will produce the next generation of scalable computing components, which in due course will be assembled to build the 100 teraflop and faster supercomputers to be used by the DOE and the Alliances. Hence the success of the CSL effort should be measured in part by the speed of transition of the new technologies to industry. The entrepreneurial environment of the Silicon Valley lends itself to this rapid transition, which often involves the faculty who made the research advance.

The roadmap for the overall CITS program is shown in Fig. 1.1. This roadmap suggests that, with the advances enabled by CSL and other computer systems research, it should be possible to do a full integrated engine simulation on hardware expected to become available in the 2007 timeframe. With such capability, the engine designers could use simulation to look at complex off-design situations that are crucial for engine performance and reliability, reducing the time and money that must be spent on experimental development, thereby increasing competitiveness of the U.S engine industry. The importance of this to the national economy cannot be overstated.

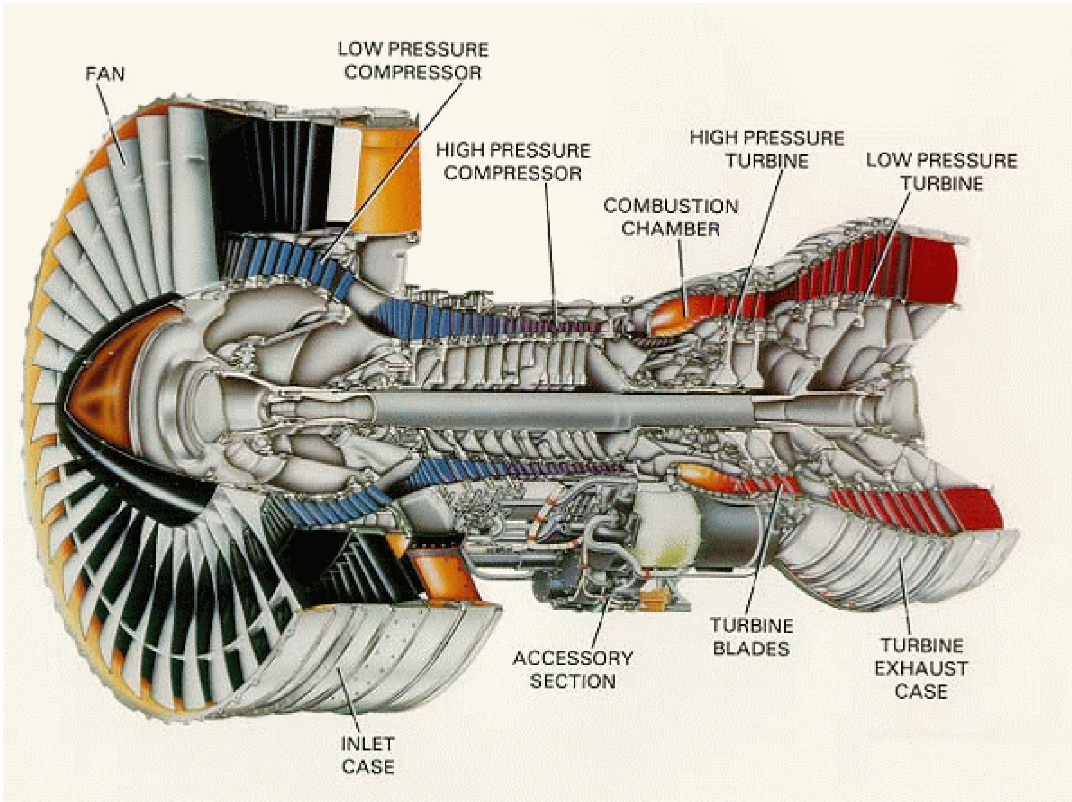

Figure 1.2: Gas turbine engine.

1.2 Engine simulation plan

An aircraft gas turbine engine is a very complex system, an example of which is shown in Fig. 1.2. Flow enters at the left through the fan, is compressed by a series of rotating blades and stationary vanes in the low- and high-pressure compressors, and then is heated by burning fuel in the combustor. The flow exiting the combustor is discharged through the high-pressure turbine, which drives the high-pressure compressor, then passes through the low pressure turbine, which drives the fan and low-pressure compressor, and finally exits through the jet engine nozzle. The goal of the CITS program is to produce high-fidelity integrated simulations of the complete flow path for a modern engine, and thereby to enable study important component interactions.

The turbomachines (compressor and turbines) may have several stages (blade rows) and many more blade passages per blade row, for a total of several hundred blade passages. Current industry simulations typically involve only one blade passage per blade row, and are made using steady-flow codes that assume complete mixing of the discharge of each blade row before it is ingested

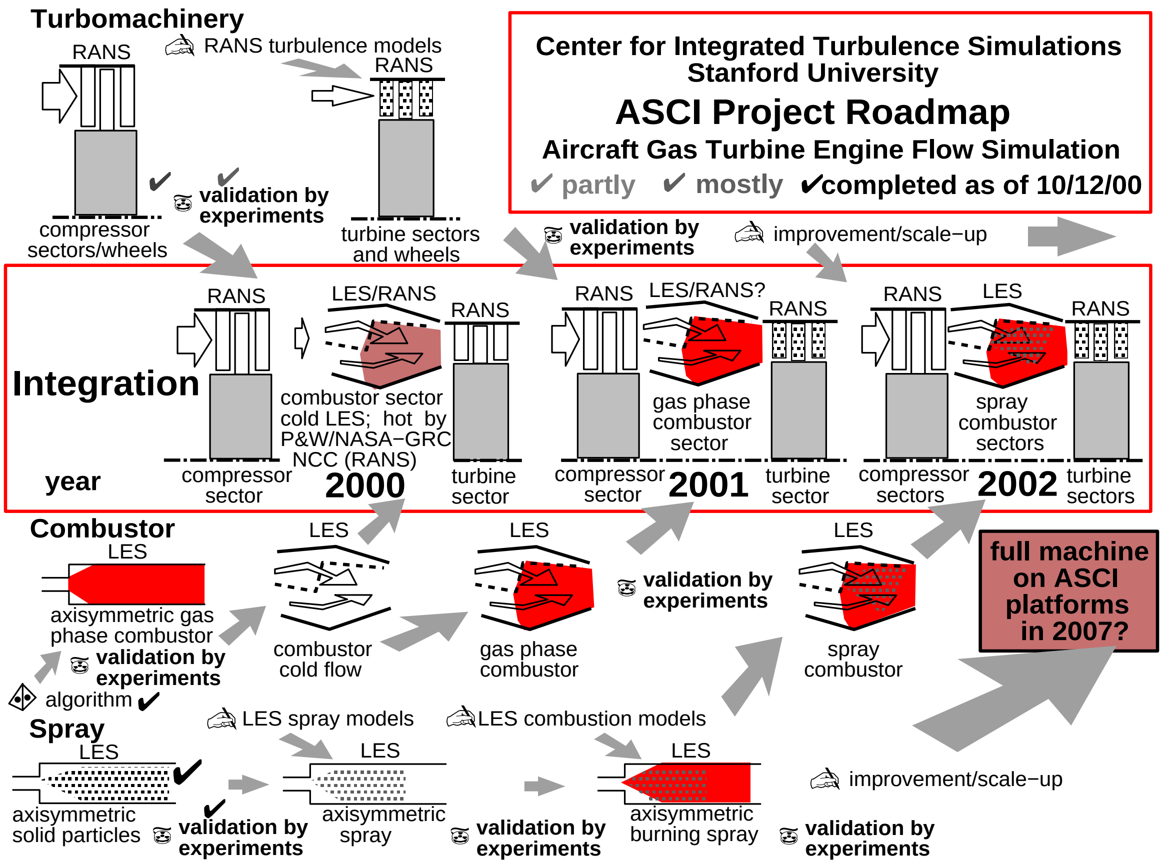

Figure 1.3: Simulation roadmap

into the next. Detailed unsteady calculations show that the shock waves and moving wakes of upstream blades have a pronounced effect on the flow in the subsequent blade passages. Hence unsteady calculations that can handle shocks and wakes are needed to address many important design problems. Moreover, some problems, such as rotating stall in the compressor, can only be addressed by considering every blade passage in the machine. Problems caused by localized hot streaks from the combustor entering the turbine can be addressed only with integrated combustor-turbine simulations. The CITS program provides the first opportunity to carry out calculations of this scope and scale, which is why the industry is so interested.

The roadmap for the engine simulation is shown in Fig. 1.3. The effort is organized in two teams, one focusing on the turbomachinery and the other on the combustor. The approach is to handle the turbomachines using the unsteady compressible Reynolds-Averaged Navier-Stokes (RANS) equations with advanced models for the wall and wake turbulence. Because the combustor flow is very complex, and RANS models of combustion have not been very adequate, the CITS approach is to use Large Eddy Simulation (LES) for the combustor. The combustor is a critical element in the system. The designs are geometrically very complex, with passages for dilution air, recirculating flow, fuel spray insertion and evaporation, and highly turbulent combustion. The Mach numbers in the combustor are typically low, but the density changes due to heat release are very substantial. Hence the combustor LES code is currently based on low Mach number compressible flow equations. The RANS calculations can use various methods (such as upwinding, fourth-order dissipation, *etc.*) to achieve adequate flow predictions. However, LES must use non-dissipative conservative methods in order to capture the turbulence dynamics correctly. Hence research on algorithms and sub-grid modeling have been required to develop LES for the combustor.

The research plan includes work on turbulence modeling for rotating systems, spray modeling and simulation, drop breakup and evaporation, and other flow physics that needs to be improved for a high-fidelity engine simulation. It also involves the development of new numerical methods for flow computation. At Stanford there is a long tradition and record of contributions to numerical analysis from flow physics faculty, who teach the principal courses in numerical analysis, have written books and many papers on the subject, and are journal editors in the field. Hence in the ASCI program at Stanford this fundamental numerical analysis research is carried out by the turbomachinery and combustor simulation teams as an integrated part of their programs. The relatively small group in the Computer Science Department that also does numerical analysis assists in this effort.

1.3 RANS turbomachinery simulation

The CITS turbomachinery code TFLO was developed starting with well-established aircraft flow codes built previously by Prof. Antony Jameson, the leader of the TFLO team. This big head-start enabled TFLO to be developed quickly. It has since undergone an extensive set of verification (against other codes used by industry and NASA) and validation (against data) tests by the extended TFLO team, which includes several industrial participants.

TFLO is a scalable code that is used for large-scale steady and unsteady compressible RANS simulations. Coupled to a combustor code, it can be used to study complex component interaction problems. A goal of the program is to develop numerical simulation techniques that make this type of calculation computationally affordable, and to demonstrate this capability by simulating typical interactions of engineering interest.

TFLO uses a multiblock structured multi-grid mesh with double halos at interfaces. The mesh is body-fitted to each stage, with a mixing-plane interface (steady flow) or sliding-mesh interface (unsteady flow) between stator and rotors. Unsteady simulations use a dual time step algorithm. Jameson-Schmidt-Turkel (JST) or Convective Upwind Split Pressure (CUSP) algorithms are used to control numerical stability. A variety of turbulence models have been implemented, including Durbin's V2F method developed at Stanford. A special preprocessor provides domain decomposition and load balancing in parallel processing systems. The parallel implementation uses a standard Message Passing Interface (MPI), and has been shown to be quite scalable and efficient with over 1000 processors.

Key TFLO developments made during the past year include new conservative sliding mesh interface treatments, the implementation of parallel file input/output (Sect. 2.1.2), and a generic algorithmic framework for dealing with various turbulence models (Sect. 2.2) including V2F. Section 2.3 provides an overview of some of the important validation and verification work involving tens of millions of meshpoints. Additional information on large-scale TFLO performance is provided in the section on large-scale integrated simulations (Sect. 4.3.2 and 4.3.3). Prof. Juan Alonso and Dr. Roger Davis, who is on assignment to the Stanford TFLO team from the United Technologies Research Center (UTRC), and Dr. Jixan Yao, are key contributors to these efforts.

Transition from laminar to turbulent flow is an important phenomena in the low-pressure turbine. This process is very complicated because of the unsteady wakes from upstream blades that are injected into each blade passage (rotor or stator). In order to develop understanding of the physics, and in the hope that this would enable modifications of V2F to account for transition, landmark Direct Numerical Simulations (DNS) for a low-pressure turbine blade passage has been carried out under the direction of Prof. Paul Durbin as part of this program. This work (Sect. 2.4) has revealed new phenomena of lasting scientific value and immediate engineering importance.

Other fundamental studies are ongoing as part of the turbomachinery program. Professors Joel Ferziger and Sanjiva Lele are working with a student on LES of a single turbine blade passage. The goal here is to produce a simulation that can be used to assess and improve turbulence models. Prof. W.C. Reynolds, NASA's Drs. Alan Wray and Karim Shariff, and Dr. Stavros Kassinos are conducting large-scale DNS of homogeneous turbulent shear and straining flows in rotating frames on the ASCI Red machine in order to develop a database for assessing and improving turbulence models. This work will be reported next year.

1.4 LES combustor simulation

Large Eddy Simulation was pioneered at Stanford starting in the mid 1970s. It has been a focal activity of the NASA/Stanford Center for Turbulence Research (CTR) since its inception in the late 1980s. The dynamic sub-grid scale modeling approach, which does not require empirical constants and has become the workhorse of modern LES, was developed largely through work of the CTR. CTR work has also included considerable numerical simulations of turbulent reacting flows. Thus, the Stanford team is well qualified to take on the development of an LES framework for complex combustors such as found in modern aircraft gas turbine engines.

LES requires conservative, non-dissipative algorithms. This presents a challenge for the complex geometries of typical combustors, where unstructured meshes are preferred. Once it had been decided that LES would be necessary for high-fidelity combustor simulation, the combustor group under Prof. Parviz Moin set out to develop a conservative, non-dissipative algorithm for complex unstructured meshes. This algorithm, by Dr. Krishnan Mahesh, was a major contribution from the first year of the CITS program. It is patterned after the staggered approach of Harlow and Welch, which had become the workhorse for LES on Cartesian structured meshes.

The algorithm, which was first developed for incompressible flows, has now been extended to low-Mach number variable density flows and validated against other simulations for a variety of steady and unsteady flows. The dynamic Smagorinsky model has been extended to unstructured grids and implemented in the code. Section 3.1.1 describes recent algorithm improvements. Section 3.1.2 reports some of the validations against laminar and turbulent flows, and outlines our current work in setting the code up to run cold flow LES in a modern engine combustor. The code is written in Fortran 90 and designed for effective parallel implementation. The work has progressed to the point where the algorithm is beginning to be tested on large numbers of processors.

In parallel with the combustor flow code development, we are developing new methods for handling the chemistry in LES. The first is an unsteady Lagrangian flamelet approach of Dr. Heinz Pitsch. A somewhat different flamlet-progress variable approach by Mr. Charles Pierce and Prof. Moin also shows promise for use in LES. These are described in Sect. 3.2.

Another parallel effort is developing technology for high-fidelity spray simulation. This must describe the spray injection, sheet and droplet breakup, and evaporation that precede the gas-phase combustion process. Dr. Joseph Oefelein, a graduate of Penn State University with considerable practical aircraft engine experience, came to us after his PhD on LES for spray simulation. He has used both the Pierce/Moin gas phase combustor LES code and his own code developed at Penn State to study various aspects of spray modeling in simple axisymmetric and planar geometries. Both codes use structured grids and staggered data. The Pierce/Moin code is based on the low Mach number equations and uses a Cartesian axisymmetric mesh, whereas Oefelein's code can handle fully compressible flow in generalized curvilinear coordinates. The first work, reported last year, showed that standard drag models for solid particles produce excellent replications of the particle distributions (compared to data) when tracked individually in a LES of a gas that is fairly lightly loaded with particles. Current work using his own code is extending these evaluations to more highly loaded channel flow, for which there is experimental data from Stanford, and to a more complex axisymmetric swirling combustor geometry in which spray and combustion experiments are being conducted now at Stanford. This work is described in Sect. 3.3.

As part of the development of improved spray simulation technology, we have two efforts in DNS related to droplets and droplet interactions. The first (Sect. 3.4), being conducted by Dr. Brian Helenbrook, studies flow in and around deformable axisymmetric drops using a multigrid spectral element method. The second (Sect. 3.5), conducted by PhD student Tristan Burton, investigates the interaction of droplets in clusters using Chimera grids. The expectation is that these studies will lead to improvements in the modeling of spray droplet behavior, and that this will improve the fidelity of the LES spray combustor simulations.

In addition to these developments in numerical analysis being carried out by the flow physics group, the Scientific Computing and Computational Mathematics (SCCM) group, led by Prof. Eugene Golub, has been working with Prof. Joel Ferziger to assist the spray modeling effort. SCCM Prof. Andrew Stuart initiated work with a student on particle collision detection before leaving Stanford to return to the U.K. This work has continued in collaboration with Prof. Ferziger and SCCM's Dr. Wing-Loc Wan. Dr. Wan has also worked on fast multigrid methods for flows with discrete interfaces, and on other problems of interest to our ASCI program. This SCCM work is summarized in Sect. 3.6.

All of the work outlined above is intended to give us the capability for high-fidelity simulation of the very complex physics of turbulent combustion in the very complex geometries of modern aircraft engine combustors. The contributions from this work should have lasting impact on both the science and simulation of turbulent spray combustion.

Several Stanford experimentalists are involved to be sure that we use the most current and best experimental data to guide and validate the combustor model. Prof. John Eaton continues his experiments on particle-laden turbulent flow under other support in parallel with his guidance of Mr. Burton's ASCI droplet/turbulence interaction research. Prof. Chris Edwards, a spray expert, is conducting spray injection and spray combustion experiments designed to complement the ASCI simulation under NASA suppor, and is guiding the ASCI work of Dr. Helenbrook. Profs. Tom Bowman, a combustion and kinetics expert, and Prof. Godfrey Mungal, a turbulent mixing expert, consult with the combustion team. Other combustion experts (*e.g.* Norbert Peters, George Kosaly) visiting the CTR also consult with the team.

1.5 Large scale multicomponent integrated simulations

The ultimate objective of our program is the integrated coupled simulation of the compressor, combustor, turbine, and secondary air flows in a realistic aircraft gas turbine engine. The computing power to do this simulation does not presently exist, but it should by the end of this decade. Because TFLO had a head start and is now operational, but the LES combustor code is still under development, we entered into a partnership with the NASA/Glenn Research Center to use their National Combustor Code (NCC). The goal is to couple the NCC to TFLO to study various integration issues. Prof. Juan Alonso is leading this effort in collaboration with Dr. Roger Davis, our partner from UTRC in residence at Stanford.

The NCC is a RANS code, so its fidelity in combustion simulation is not expected to be great, but it has enough in common with our LES combustor code to be useful in gaining some early integration experience. It uses an unstructured mesh that can deal with complex geometries, and can run in an unsteady-RANS mode similar to LES. The NCC is export-controlled, hence because we have foreign nationals involved (faculty, RAs, and students) we can not obtain the source code. We have not provided TFLO source code to NASA/Glenn, so that the two parties must find a way to couple their codes without one knowing the details of the other. Each party has written its side of the interface, which exchanges information between the NCC and TFLO. Stanford has tested the interface by coupling two TFLO calculations through the interface. NASA/Glenn is making the same sort of tests of its side of the interface. TFLO and the NCC have been coupled, but some technical problems need to be solved and so it is too soon to report any details.

At Stanford, Dr. Roger Davis is learning to run the NCC without access to its source. As a key member of the TFLO team, he is also running very large scale turbine calculations involving some of the important secondary air flows using TFLO. These very large scale simulations have been done in dedicated time on the DOE Blue Pacific computer at LLNL, using over 1000 processors. These have the potential to become landmark simulations for the engine industry, and if successful they should generate considerable enthusiasm for this sort of large-scale simulation. This work is reported in Sect. 4.

The attempts to begin integrating the NCC with TFLO have been seriously hampered by the lack of time that knowledgeable NASA/Glenn personnel have been able to devote to this collaborative effort. We are hopeful that this will improve in the coming fall. NCC-TFLO coupling would provide a valuable RANS-RANS baseline with which future LES-RANS coupling could be compared. However, if the NCC-TFLO effort proves to be unsuccessful or slow in maturation, the ASCI effort will be redirected to accelerate development and integration of the LES combustor code.

Coupled combustor-turbine simulations will provide new insight into the origins of "hot streaks" observed in the first turbine stages. These probably arise as a result of non-uniformities in the combustor, and are a critical factor in determining turbine blade lifetime. The results could make an important near-term impact on engine design, and ultimately such calculations could become an important step in the engine design process.

1.6 Full system integration at small scale

The DOE platforms will not be sufficient for full-scale simulation of large aircraft gas turbine engines until the 2007 timeframe. The best we are able to do now is a sector of an engine. Therefore, we are considering starting a parallel project on a much smaller gas turbine engine manufactured for educational purposes by Turbine Technologies, Ltd. Their little SR-30 engine (Fig. 1.4) is only about 7 inches in diameter. It has a single-stage radial compressor, an annular combustor with six injectors, and a single-stage axial flow turbine. We have just obtained two of these engines at Stanford. One is sectioned for instructional viewing, and the other is fully instrumented and running in an undergraduate teaching laboratory. The objective would be to make fully-integrated simulations at this scale to predict the overall engine performance, which could then be compared with experiments as part of the laboratory course.

There would be several advantages to using this engine for a first integration step. It is of a scale that it could be handled completely with high fidelity using TFLO and the LES combustor code on current ASCI platforms, and possibly even on our in-house clusters. We would encounter most if not all of the RANS/LES/RANS coupling problems needed to be solved to do large-scale full-engine simulations. We would not have the problems of data display that we have encountered when using proprietary commercial geometries from industry. And if this simulation can be done on our in-house cluster, the use of this simulation tool could become an integral part of our teaching

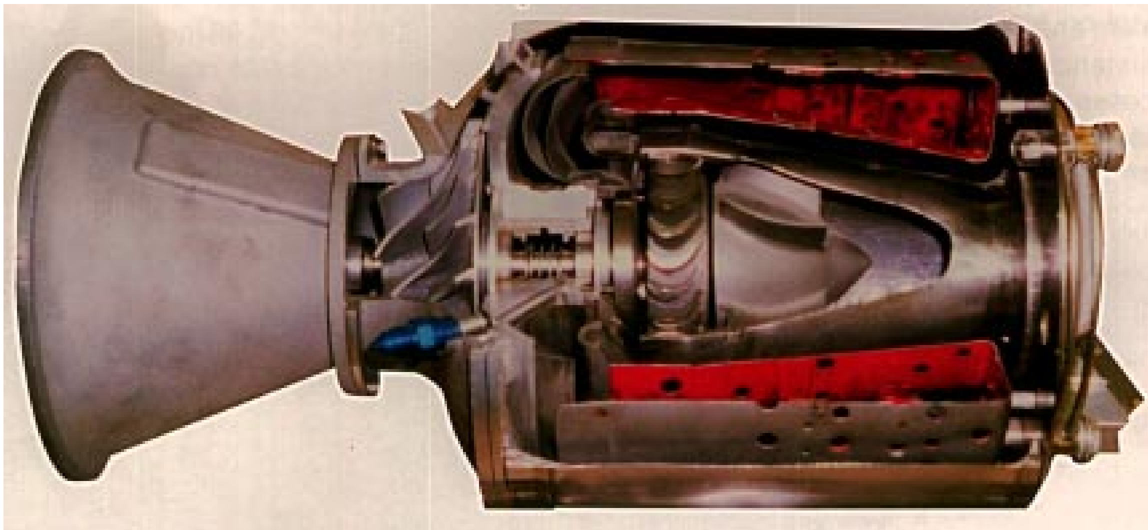

Figure 1.4: Turbine Technologies SR-30 Mini-Lab Engine. Reproduced by permission of Turbine Technologies, Ltd.

program in Mechanical Engineering, thereby stimulating student interest in simulation, which is another objective of the ASCI Alliance program.

We have discusses this new idea with Turbine Technologies, which expressed interest in collaborating with us on this project by providing the geometrical information that we need to construct the flow grid. Further discussions among our simulation group and with the DOE are needed before a final decision on this plan is made.

1.7 Computer systems and architectures research

Faculty and students in the Computer Systems Laboratory (CSL) have been an integral part of the ASCI program since its inception. The CSL vision that made its participation in ASCI attractive to the DOE has not changed, although the rapid elevation of the team leader, then Dean John Hennessy, to Provost and then to President of the University, has had some impact. His vision was recently reaffirmed by the current CSL faculty leaders, who see the future of large-scale scientific computation as involving many thousands of interconnected processors assembled from commodity components, sharing their memory at high bandwidth, and operating in a fully-integrated manner. The CSL ASCI work has been, and continues to be, in support of this vision.

Prof. Hennessy's work in ASCI has included Flash, a special machine with a programmable node controller (MAGIC) that handles cache coherence and all other node traffic. This machine provides a very flexible platform for study of different cache control strategies, for simulation of alternative machine architectures, and for instrumentation of application codes. FlashPoint is a software tool that uses Flash to assess performance of Distributed Shared Memory (DSM) codes. The engine simulation group has begun to use it to study performance of its OpenMP codes running on SGI machines. Flash was described previously and an additional report is not included this year.

The current faculty lead for the CSL ASCI program is Prof. William Dally. His work on interconnection networks and high-speed signaling is described in Sect. 5. High bandwidth interconnections will simplify the programming of large parallel computers, and Dally has come up with new router architecture and signaling technology that offers a large potential increase in bandwidth. He is already taking this technology to industry, a first step in getting this into commodity components for supercomputing.

Simulation of very large supercomputers is an important aspect of scalable operating system design. This requires big computers to simulate much bigger computers. Prof. Mendel Rosenblum's group completed the development of this technology during the first year of the ASCI program. His work on SimOS, a large-scale operating system simulation tool, was reported previously. He has been on leave for the past two years bringing this technology to market. The work has been continued by his students and faculty colleagues, and is not reported on this year.

A concept for virtual clusters has been explored by Prof. Mendel Rosenblum and his team. The nodes of a large shared-memory machine are partitioned into groups that use virtual machine technology with each running a separate operating system image. The groups are then clustered using the memory system of the underlying system for communication. This approach offers scalable performance, dynamic resource partitioning, fault containment, and low implementation cost. Hence it addresses several needs of large-scale shared memory multiprocessors. This work is ongoing and hence will be reported next year.

Prof. Dawson Engler and his students are developing tools for automatic error checking. They have developed a *meta compiler* that has detected hundreds of errors in contemporary operating systems including Linux. The system works by statically checking a number of simple rules, *e.g.* "every lock that is acquired must be released", against software. While these tools have been primarily applied to system software, they hold the potential to check for synchronization and consistency errors in parallel applications. This is new work to be reported next year.

Prof. Monica Lam's work on automatic parallel compilers is building a firm mathematical foundation for automatic parallel compilation. She has demonstrated her Stanford University Intermediate Format (SUIF) compiler on a number of ASCI codes, including some provided by the engine simulation team. Section 6 gives a detailed report on her work.

As a result of a suggestion provided to Prof. Reynolds at the annual DOE Conference on High Speed Computing this February, we initiated a collaboration between the engine simulation team and visualization research of Profs. Pat Hanrahan and Ron Fedkiw. Prof. Hanrahan has been developing technology for parallel visualization on clusters under support from the ASCI Views program. His team has extended the Visualization Tool Kit to parallel clusters, and he will make his cluster system available for use by the simulation team. Dr. Massimiliano Fatica of the engine simulation group is now writing parallel VTK scripts to be run on Hanrahan's cluster, and the CTR/CITS is upgrading its cluster for compatibility with Hanrahan's so that these scripts can be tested at CITS. The visualization researchers and engine simulators have defined some specific objectives for the collaborative effort and identified research needed to advance visualization capabilities. We hope to grow this cooperative visualization effort over the coming year, and look forward to including a report on this in next year's Annual Technical Report.