/N -5 /

4 /86

P-69

ICASE

THE FIRST ICASE/LARC INDUSTRY ROUNDTABLE: SESSION PROCEEDINGS

Sharath S. Girimaji

(NASA-CR-195040) THE FIRST
ICASE/LARC INDUSTRY ROUNDTABLE:
SESSION PROCEEDINGS Final Report
(ICASE) 69 p

N95-25119

Unclas

G3/59 0047186

THE FIRST ICASE/LARC INDUSTRY ROUNDTABLE: SESSION PROCEEDINGS

Sharath Girimaji*

Institute for Computer Applications in Science and Engineering
NASA Langley Research Center
Hampton, VA 23681-0001

ABSTRACT

The first "ICASE/LaRC Industry Roundtable" was held on October 3 - 4, 1994, in Williamsburg, Virginia. The main purpose of the Roundtable was to draw attention of ICASE/LaRC scientists to industrial research agendas. The Roundtable was attended by about 200 scientists, of which 30% from NASA Langley; 20% from universities; 17% NASA Langley Contractors (including ICASE personnel); and, the remaining from federal agencies other than NASA Langley. The technical areas covered reflected the major research programs in ICASE and closely associated NASA branches. About 80% of the speakers were from industry. This report is a compilation of the session summaries prepared by the session chairmen.

FOREWORD

The first "ICASE/LaRC Industry Roundtable" was held on October 3 - 4, 1994, in Williamsburg, Virginia. The main purpose of the Roundtable was to draw attention of ICASE/LaRC scientists to industrial research agendas. The ICASE scientific staff consists primarily of recent Ph.D's whose general research focus is on the NASA agenda, which is evolving increasingly to include the U.S. industrial research agenda. However, there is presently no mechanism, besides certain professional meetings, whereby ICASE scientists are exposed to the industrial research needs.

In order to evaluate the potential need for the Roundtable, we conducted a survey of industrial scientists in May and June of 1994. The survey questions were aimed towards evaluating (i) industrial perception of the importance of basic and applied research, (ii) current capability of companies to do research 'in house', and (iii) industrial perception of the awareness of national laboratories and academia of the industrial agenda. It was clear from the responses that several basic research issues of industrial importance were put on the 'back burner' due to lack of resources, and that companies would like to see these problems addressed by institutions such as ICASE. The Roundtable was clearly an idea whose time had come. It is envisioned as just the first step towards establishing a direct communication between ICASE, LaRC and potential industrial collaborators and advisors.

The Roundtable was attended by about 200 scientists, of which 30% were from industry; 30% from NASA Langley; 20% from universities; 17% NASA Langley Contractors (including ICASE personnel); and, the remaining from federal agencies other than NASA Langley.

The attendees were welcomed by Dr. M. Y. Hussaini, Director of ICASE. This was followed by keynote speeches by Dr. Hans Mark, University of Texas at Austin and Dr. Paul Rubbert, Boeing Commercial Airplane Company. About 70 technical talks were presented in three parallel tracks. The areas covered reflected the major research programs in ICASE and closely associated NASA branches. About 80% of the speakers were from industry.

Since the Roundtable, ICASE has held workshops on computational acoustics and on adaptive grid methods, in which industrial scientists served as panel discussants. In future, ICASE intends to bring industrial perspective to bear on basic research through its regular workshop and short course program.

This report is a compilation of the session summaries prepared by the session chairmen. We would like to thank all the speakers, especially our industry 'guests', for educating us on industrial research needs.

1

In addition, we would like to express our gratitude to the various session chairmen for their contributions. Finally, we would like to thank Ms. Emily Todd for the smooth and efficient organization of the Roundtable.

Sharath S. Girimaji,
On behalf of the Organizing Committee

WELCOME ADDRESS

M. Y. Hussaini, Director, ICASE

Changing times, and changing national and international priorities, are bringing with them changes in the way those of us in research enterprises conduct our business. Shrinking resources for research drive the industries, national laboratories and academic communities together in hopes of getting the work done cooperatively.

The initial impetus for the ICASE/LaRC Industrial Roundtable arose from considerations independent of the financial restructuring of the research enterprise. We realized, in conversations amongst ourselves at ICASE, that several of us had individual industrial research collaborations which have provided intellectual stimulation that we would not have had apart from industrially-oriented motivation. Indeed, with high performance computing and communication, an important part of the ICASE research agenda, it would be unnatural not to be substantially motivated by grand challenge problems with industrial origins.

We also realize that we have, in many cases, to overcome substantial barriers to realize these productive collaborations: our communities do not mix well through professional technical meetings; existing personnel and travel policies discourage individual exchanges between workplaces; and, there are often restrictions of a proprietary nature that prevent detailed disclosures of technically necessary information and software. Because of these impediments and the random origins of the handful of successful collaborations that we currently enjoy, we thought that a systematic attempt to identify further areas of potential collaborations for ourselves and for our young postdoctoral staff would be in order. The Roundtable is our first step in this more systematic approach.

The technical areas represented at the Roundtable reflect the major research programs at ICASE and in closely associated NASA branches. However, we did not call on industrial participation, primarily, in order to have an audience for our own work. Rather, the breakdown of speakers on the program (80% from industry and the balance from universities, NASA, and ICASE) indicates our interest in the industrial research agenda. We recognize that the part of that agenda that can be discussed in public is necessarily limited to the medium- to long-term "precompetitive" category. We will be listening not only for fresh insights into important problems and technical directions, but also for practical mechanisms for the exchange of personnel and technical information.

3

KEYNOTE TALKS

The attendees heard keynote addresses from **Hans Mark** of the University of Texas at Austin and **Paul Rubbert** of the Boeing Commercial Airplane Company. Both praised the Roundtable as timely, but even as they welcomed the attention of those whose research agendas are typically open-ended to applied problems of industry, they affirmed the importance of federally sponsored *basic* research. According to the keynoters, as the nation's corporate research laboratories have restructured and assigned to even their most creative scientists and engineers a higher percentage of work related to short-term corporate missions, they must look more to the federal laboratories to maintain long-term thrusts.

Mark, who served more years than any other individual as the Director of the NASA Ames Research Center (from 1969 to 1977) and who oversaw one of the largest educational organizations in the world as Chancellor of the University of Texas from 1984 to 1992, devoted his talk to affirming NASA's historical dual mission in providing long-term leadership to space and aeronautics research. He amused attendees by reading excerpts from issues of *Collier's* dated 1952 to 1954 – before the acronym NASA had been coined – in which Wernher von Braun, Fred Whipple, and others had uncannily accurately predicted the means by which the moon would be explored fifteen years later. He also assigned NASA credit for the $20 B/year net exports of the nation's aviation industry.

Rubbert, a computational fluid dynamics (CFD) pioneer who led the revolution towards computationally-based design during 33 years at Boeing, has often testified in Congressional committees about the role NASA should play in the domestic aviation industry. He was recently freed by Boeing senior management to devote his full energies to enacting the ideas in his 1994 Wright Brothers Medal lecture on managing corporate R&D. Noting that in today's open global economy every aerospace company has fairly level access to computer power, good wind tunnels, and suppliers, the only remaining competitive advantage is *time*. Boeing has spent approximately five years to develop all of their passenger aircraft except for one. The aircraft that they rushed through in just three years – the 747 – has been and remains their greatest cash cow, largely because it won the race to be first in its market sector. Rubbert plotted contours of market share on axes of "goodness" versus "date available" in order to compare product development trajectories. He also set forth the concept of an R&D "food chain," a pyramid which is built on basic research. He argued that basic researchers must be even more informed about what the industrial customer needs than applied researchers closer to the top of the pyramid, since fundamental improvements must come from basic researchers, while the job of applied researchers is to improve current processes. Much of NASA's traditional CFD work has been motivated by the need to model flight performance

for wing design. However, with fuel costs down to 10% of the total operating costs of today's airlines, incremental advances in traditional airplane performance metrics have become less important than the ability to do "just in time" custom manufacturing.

SESSION A1 & A2: LAMINAR FLOW CONTROL

Ajay Kumar, NASA Langley

Mujeeb Malik, NASA Langley

SUMMARY

Two sessions of two hours each were held on laminar Flow Control (LFC). Each session had four speakers. Both the sessions were surprisingly well attended and there was extensive discussion on a number of issues related to laminar flow control technology. The speakers and the participants agreed upon the need of developing user-friendly, well-documented and calibrated LFC design and optimization tools that can be readily used by the industry. Areas of further research, such as receptivity, absolute amplitude method, supersonic attachment line stability, and roughness effects, for improved LFC configuration design were also emphasized. Apart from NASA and ICASE, attendees included representatives from Pratt & Whitney, Rockwell, Boeing, Northrop, Cessna, Beechcraft, etc.

INTRODUCTION

The first session had the following four talks:

- Overview of NASA's LFC Program for Subsonic and Supersonic Transport Aircraft – **Dr. Ajay Kumar**, NASA Langley

- Control of Transition for Drag Reduction – **Professor Eli Reshotko**, Case Western University

- Laminar Flow Research for General Aviation – **Mr. Randy Nelson**, Cessna Aircraft Company

- A PC-based, User-Friendly System for Airfoil Design and Boundary Layer Stability Analysis – **Mr. Jeff Viken**, Innovative Aerodynamic Technologies

The first talk basically described the current program that NASA is pursuing in Laminar Flow Control for subsonic and supersonic transport aircraft. The program's main objective is **to define and implement a national LFC technology development plan with US industry which will address and resolve key aerodynamic and subsystem issues to insure LFC technology readiness for application at reduced risk for both advanced subsonic and supersonic transport aircraft.** The program is not only

developing advanced LFC design tools but is also addressing some of the implementation issues, such as insect protection, anti-icing, and integration with high-lift systems.

The second talk addressed some of the issues related to the control of boundary layer transition to reduce drag of aircraft and underwater vehicles. It described techniques of transition control based on the linear stability theory and discussed vehicle and environmental factors, such as surface roughness, particulates, icing, insect impact, etc., that need to be overcome for practical applications.

The third talk addressed the issues related to natural laminar flow control for application in General Aviation. Cessna has an operational business jet with about 30% chord natural laminar flow (NLF) and is interested in developing a new moderate sweep business jet with extensive laminar flow. However, due to lack of their own resources, they would like to collaborate with NASA in this developmental effort. Their current and long-term needs are:

- Improved model and grid generation capability

- Workstation-based 2D Navier-Stokes codes

- 3D boundary-layer stability codes

- Improved user-interface and documentation of codes

- 3D aerodynamic optimization

- Stability and hinge moment prediction (Handling characteristics)

- 3D drag prediction

- 3D high-lift systems

- 3D icing physics

- Aeroelastic/structural interactions

- Multidisciplinary design and optimization

- Advanced/unconventional configurations

Cessna strongly supports basic research at NASA and considers itself as the applier of the basic research.

The fourth and final talk of the first session described a 2D incompressible, PC-based airfoil design and stability analysis package, LAMDA, which was developed under a NASA SBIR Phase I effort. It combines a flow analysis code, a boundary-layer mean flow code,

and a linear stability code for analysis and design of low-drag airfoils. The package provides expert input database to execute the stability code, proper data interfaces among the codes, on-line data analysis and plotting capability, and book keeping of multiple files produced by the codes. A demonstration of LAMDA package was also given on a PC after the session.

The second session had the following four talks:

1. Methods for Transition Prediction – **Dr. Mujeeb Malik**, High Technology Corporation

2. Boundary-Layer Receptivity in LFC Applications – **Dr. Meelan Choudhari**, High Technology Corporation

3. Laminar Flow Drag Reduction for High-Bypass Ratio Turbo Fan Nacelles – **Dr. Wesley Lord**, Pratt & Whitney

4. A General Aviation view of Laminar Flow – **Dr. Neil Pfeiffer**, Beechcraft

The first talk was an overview of the transition prediction methods where both the N-factor and absolute amplitude based methods were discussed. The N-factor method measures the relative growth of instability waves in boundary-layer flows. These instabilities take the form of Tollmien-Schlichting (TS) waves in two-dimensional boundary layers, crossflow disturbances in 3D boundary layers and Görtler vortices in flow along concavely curved bodies. Various examples were given where the N-factor method gives good correlation with experimentally observed transition front. The N-factor method is well-developed although there is a need to improve the computer codes for robust operation so that these codes can be used in the aerodynamic design with confidence. Work on this aspect is currently underway at NASA LaRC.

The newly developed parabolized stability equation (PSE) approach was also described. In the linear framework, it can be used to compute N factors with nonparallel and surface curvature effects. In the nonlinear framework, this technique is capable of yielding absolute amplitude of disturbances. Nonlinear disturbance modulation makes the boundary-layer flow amenable to secondary instabilities and the amplitude at which these instabilities appear varies from flow to flow (e.g., 1% for a TS-dominated 2D boundary-layer and 10% for crossflow dominated 3D boundary layers). The emergence of secondary instabilities may be used to develop a new criterion for transition prediction. Some examples of secondary instabilities and comparison with the experiments were also given. In the examples presented for nonlinear PSE, the initial amplitude of the disturbances was assigned arbitrarily and no attempt was made to link it to free-stream disturbances or surface inhomogeneities. However,

when such a connection is made, the nonlinear PSE will provide a tool for prediction of transition from first principles.

The connection between the free-stream disturbances and surface inhomogeneities with the generation of instability waves within the boundary layer is provided by the receptivity theory which was the subject of the talk given by Dr. Choudhari who presented an overview of the ongoing theoretical research in this area at NASA Langley Research Center. The current focus of this work is on the prediction of receptivity in laminar flow control applications, with the final objective being to couple the output of this stage with tools such as the PSE described above. Recent breakthroughs in the field have shown that receptivity can occur through multiple paths that include: (1) the mean-flow nonparallelism close to the leading edge as well as (2) any short-scale variations in the surface properties such as surface geometry, surface suction velocity, surface temperature, etc.

The surface nonuniformities are expected to provide the dominant source for the overall receptivity, except possibly at larger Mach numbers. A simple modular approach was described which uses localized receptivity as a building block for developing predicting capability for practical disturbance environments of varying degrees of complexity. This approach was applied to both localized and distributed nonuniformities and to different types of freestream disturbance fields. Both deterministic and stochastic aspects of the disturbance environment have been considered within this framework.

It has been found that the surface geometry variations are likely to be the most dangerous source of receptivity as compared to any other types of surface disturbances. Similarly, acoustic receptivity is found to be potentially more important than the receptivity to convected vortical disturbances. The disparity between the initial amplitudes of stationary and nonstationary instabilities that are excited through surface nonuniformities was also explained. Finally, the nonlinear enhancement of receptivity induced by surface projections (i.e., humps) with moderately large height perturbations was demonstrated for both two- and three-dimensional boundary layers.

This very illuminating talk on receptivity was followed by a talk on laminar flow drag reduction for high-bypass ratio turbo fan nacelles by Dr. Wesley Lord who pointed out that the application of LFC to about 60 percent of the fan cowl would result in about 2 percent reduction in aircraft drag which is quite significant. Dr. Lord presented a review of some of the aerodynamic design issues related to the development of laminar flow nacelle. As far as transition is concerned, the relevant instability mode is TS since the flow is mainly two dimensional at cruise conditions. It is well known that TS instability can be stabilized both by favorable pressure gradients and wall suction. A nacelle configuration can be designed where TS waves are stabilized by pressure gradient alone (NLF) during cruise. However,

9

the inlets must operate over a wide range of conditions (Mach number, angle of attack, etc.) during take-off and landing and the leading-edge design which is optimal for NLF during cruise is not desirable during off-design conditions, and for example, flow separation would occur during take-off for an NLF design. Therefore, a leading-edge design is needed which is suboptimal from LFC considerations but provides a functional inlet during take-off and landings. This leads to a hybrid approach which uses wall suction aft of the first pressure minimum to control TS waves in the adverse pressure gradient region. It was pointed out that there is need to fine tune the aero-design methods to better predict inlet separation and further enhance the capabilities of inverse design codes.

Two of the outstanding issues in the application of LFC to nacelle design were pointed out. First is the hole-size effects in perforated LFC surfaces including discrete hole-induced transition. Pratt & Whitney has recently performed some experiments in UTRC's Acoustics Research Tunnel where the effect of hole size was studied. The second problem is the effect of nacelle leading-edge joint on transition both from the point of view of instability enhancement and its interaction with the fan noise. In both these areas, NASA LaRC could provide theoretical support as well as enhancing some of the codes needed in the design.

The final talk was given by Dr. Neal Pfeiffer of Beechcraft who presented a General Aviation (GA) view of laminar flow. He appreciated the guidance and help that NASA had provided during last two decades to familiarize the GA community with both the performance improvements possible with laminar flow and the ability to obtain laminar flow with existing technology. In this respect, he specifically mentioned Dr. Bruce Holmes of NASA Langley.

General aviation airplanes range from small 2- and 4-seat trainers through turboprop business and commuter airplanes, up to business jets. These aircraft cover a wide range of speed and design conditions which creates a need for both subsonic and transonic computational tools. A range of 2D and 3D analysis codes are used regularly for GA airplanes. Dr. Pfeiffer pointed out that there is still ongoing need for improved and validated design tools to efficiently design new airplanes. A specific need is to have tools which include multipoint optimization so that designs are appropriate for a range of conditions such as climb, long-range cruise, and high-speed cruise for a range of weights.

In addition, improved and more comprehensive guidelines which reduce the reliance on detailed analysis at the advanced design stage would help reduce the schedule and cost. It was pointed out that there is a need for some studies of different manufacturing processes and whether they are compatible with laminar flow and how do they compare in cost. Then, how can they be optimized for laminar flow performance and minimum cost? Finally, the certification and operational issues need to be given some consideration.

The formal sessions were followed by roundtable discussions and notable among them

is the discussion between Rockwell, NASA Langley and NASA Ames on the subject of supersonic attachment-line.

LIST OF RESEARCH TOPICS

The areas of cooperative research include: (1) attachment-line stability, (2) interaction of fan noise and nacelle leading-edge joint on transition, (3) effect of suction hole size, including a revisit of Goldschmidt's tripping criterion, (4) effect of roughness/step/gap on transition, (5) improvement/ customization of the design codes and to provide help to the industry in the use of the design codes.

Pratt & Whitney is planning an experiment to study the problem of interaction of fan noise with leading-edge joint at a subsonic Mach number. An example of cooperative research would be to use Langley codes to correlate/validate this experiment. These codes can then be used for actual cruise conditions to provide a prediction which would give necessary input for the design of a future flight experiment.

LIST OF WORKSHOP TOPICS

An LFC workshop would be useful where industry participants come to ICASE and work with ICASE/LaRC researchers on problems of mutual interest and learn to use the design codes. In addition, 2-day workshops on special topics geared towards industrial applications will be quite useful. Some of the suggested workshop topics are:

- Attachment-line stability

- Effect of roughness/step/gaps and hole-induced transition

- Receptivity and amplitude method for transition prediction

SESSION A2: LAMINAR FLOW CONTROL II

Ajay Kumar, NASA Langley
Mujeeb Malik, NASA Langley

Please refer to SESSION A1.

SESSION A3: ACOUSTICS

Jay C. Hardin, NASA Langley

SUMMARY

The Acoustics in Industry session at the ICASE/LaRC Industry Roundtable consisted of four speakers from diverse fields–automotive, medicine, chemical and textiles–united by common problems and applications of acoustics. It was clear from their presentations that acoustics as a discipline spans the breadth of American industry.

PRESENTATIONS

The first speaker was Sam Ciray from Arvin North American Automotive. Although Arvin is involved in aerospace through their ownership of Calspan, they are also one of the country's leading manufacturers of new and replacement automobile exhaust systems. The talk dealt mainly with the design of mufflers, in which the two major problems are maximization of the noise reduction and minimization of the backpressure induced on the engine by the muffler. The backpressure caused by the muffler and catalytic converter can rob an engine of 10-15% of its horsepower.

A generic muffler consists of three distinct volumes separated by partitions. The exhaust flow is dumped into the volume farthest from the engine by the exhaust pipe. There it turns and enters a second pipe leading back through the middle volume to the volume nearest the engine where it is again dumped. Finally, it turns and enters the tailpipe which passes back through the middle and furthest volumes to exit to the atmosphere. Where the three pipes pass through the middle volume, they are perforated to allow further conversion of acoustic energy into vorticity. This complex, three dimensional flow and acoustic field is presently analyzed using one dimensional approximate flow and acoustic models in a finite element analysis, since efficient calculation is required.

In a typical design, the muffler designer is given the acoustic characteristics of the engine (pulse frequency and dB level), the size constraints where the muffler is to be located, and the desired noise levels. In order to meet the constraints while minimizing the backpressure of the system, several areas where refined models would be useful were detailed:

Flow

- Porous Media

- Flow Branching

- Expansions and Contractions

- Volumes

- "Bean Cans" (i.e. perforated pipes)

Acoustics

- Helmholtz Resonators

- Perforations

- Absorptive Materials

- Expansions and Contractions

- Volumes

In particular, for perforated pipes, the flow resistance and acoustic impedance under both grazing and normal flows as well as the orifice velocity were highlighted.

The second talk was by Dr. Robert Coleman of the Eastern Virginia Medical School who began with an impressive list of areas where acoustics is utilized in medicine, primarilly as a noninvasive diagnostic tool. These include Doppler blood flow monitoring, ultrasonic scanning, cardiac monitoring, and auscultation (stethoscope). He also mentioned applications in his specific area of expertise, otolaryngology, including head and neck surgery, voice analysis, audiology (particularly recent advances in hearing aid technology), and speech aids for non-communicative patients.

As for present research topics, the talk highlighted one success story and one area where further research is needed. The success story involved monitoring the opening in the airway of severely ill children. Previously this was accomplished by inserting a fiber-optic tube through the nose into the throat and visually observing the opening. This was not only frightening but could also further irritate the lining of the windpipe. Research carried out jointly with NASA Langley showed that the state of the airway could be continuously monitored by a non-invasive acoustic technique, which is presently being employed in the hospital setting. Data on hundreds of patients shows distinctive spectral changes as the airway swelling becomes critical.

The area where further research is needed has to do with the use of stethoscopes during emergency medical transport. In this critical period of an ambulance or med-evac helicopter ride, the stethoscope is useless in monitoring the patient's chest cavity as nothing can be heard over the background noise. Thus, some sort of signal enhancement or active noise cancelation scheme would be very helpful in saving lives.

14

Dr. Coleman also mentioned several generic problems in acoustic design for medicine. Successful designs must be non-invasive, simple to use, yield a real time read-out of information not available in other forms, operate in a noisy, sterile, and mobile environment, and allow high information rates for long time periods and long term storage with easy access to data. Such devices must also be inexpensive, due to the relatively low unit market potential, and overcome the difficulties inherent in performing tests on humans. Even obtaining consent to collect data can be a very time consuming process.

The third speaker was Dr. Terrence Dear from Dupont. As a large chemical manufacturer, Dupont has significant noise problems produced by electric motors, pressure control valves, rotary positive displacement blowers, axial fans, centrifugal fans and blowers, HVAC systems, gear boxes, size reduction machinery and combustion devices (e.g. flames and burners). In order to control these sources, Dupont has progressed from extensive retrofit to more stringent noise requirements in the procurement process. In this way, a high volume purchaser can apply considerable pressure toward noise reduction by the supplier, which then spills over throughout the supplier's product line. For instance, twenty years after Dupont started this policy, it is difficult to find a noisy electric motor!

Dr. Dear also mentioned several areas for further research. He noted that a textile machine operates on the basis of an instability over which there is no control, that short fiber fiberglass is going the way of asbestos and must be replaced as an insulating material, and that present acoustic programs that calculate noise on the basis of vibration data still have problems. There was also a discussion of the need for speed of sound measurements in two-phase flows.

The final speaker was Dr. Mansour Mohamed from the North Carolina State University College of Textiles. Dr. Mohamed related the tremendous progress which has been made in the textile industry in automating the looms as well as in reducing their noise radiation. However, he noted that even though every machine may meet the noise requirements, if enough of them are put together, high noise levels still result. Thus, the industry still relies on ear plugs for protection. Noise partitions around the machines are heavy and a great inconvenience in servicing the machines.

Dr. Mohamed also discussed the various types of looms in use. Some of these looms weave as many as 7000 threads, the breakage of any one of which stops the loom. Thus, a means for the detection of weakness in yarn before it enters the loom would be valuable. One type of loom actually blows the yarn of diameter 0.5mm through 3mm jets to replace the shuttle. Very little is known about the dynamics of flexible columns in jets.

LIST OF RESEARCH TOPICS

Out of this session came several suggestions for cooperative research efforts. Among these are:

- Behavior of perforated plates in the presence of flow, noise, and high temperature

- Optimization of flow and acoustic performance in complicated geometries

- Fluid dynamics in elastic pipes

- Signal enhancement/active noise control

- Acoustics of pressure relief valves

- Performance of axial fans

- Behavior of flexible structures in jets

- Continuous stress monitoring

LIST OF WORKSHOP TOPICS

- Multidisciplinary optimization of flow and acoustics

- Instabilities in industrial processes

- Fluid dynamics in inhomogeneous media

SESSION A4: HIGH-SPEED AND TURBULENT COMBUSTION

Sharath S. Girimaji, ICASE

J. Philip Drummond, NASA Langley

SUMMARY

The session on high-speed and turbulent combustion consisted of two speakers from industry and one from academia. Many research issues of interest to industry were discussed. Specific areas requiring basic research were identified and the manner in which research should be carried out in academia and national laboratories to best benefit industry were discussed. Industry was also made aware of the constraints involved in performing basic research in universities.

INTRODUCTION

This session addressed the research issues of importance to industry in the areas of high-speed and turbulent combustion. There were two presentations from researchers representing aerospace industry (Drs. R. J. Bakos and M. M. Sindir) and one speaker from academia (Dr. P. Givi) presenting an University perspective of the research in this area.

First of all, the industries' motivations for basic research in high-speed and turbulent combustion were discussed. Basic research was deemed crucial for improved design of the following important aerospace propulsion systems:

1. Propulsion systems for single stage to orbit vehicles.

2. Propulsion systems for hypersonic missiles.

3. Liquid fuel rocket engine propulsion.

4. Laser systems in aeronautics.

5. Vehicle propulsion system design and development.

Participants

- **Dr. R. J. Bakos**, *Principal Scientist, General Applied Sciences Laboratory (GASL), NY.*

- **Dr. M. M. Sindir**, *Manager, CFD Technology Center, Rockwell International, CA.*

- **Dr. P. Givi**, *Professor, State University of New York at Buffalo, NY.*

17

• **Audience** participation was very significant. It enhanced the discussions by raising several relevant issues which were further addressed by the speakers. The contents of each of the talks are summarized below.

PRESENTATIONS

Dr. R. J. Bakos. This presentation focused on hypervelocity propulsion systems, in particular, premixed shock-induced combustion and detonation. "Propulsion systems using shock-induced combustion and detonation wave combustion have the potential for enhanced performance over the entire air-breathing range (Mach number 0 - 25)." "Constructive use of the detonation phenomena appears to be leading to a new level of capability in hypervelocity aerothermal and propulsion simulation."

The talk started with an overview of shock-induced combustion and detonation phenomena as applied in aerospace propulsion. Then some recent experiments in premixed shock induced combustion performed at GASL were described. This was followed by a discussion of applications of detonation ' waves for enhanced pulse simulation facilities.

Research issues pertaining to the following four classes of wave engines were discussed: (i) Oblique detonation wave engine (ODWE); (ii) premixed shock induced combustion (PMSIC) engine, (iii) pulse detonation engine (PDE), and (iv) RAM accelerator. Physical phenomena of interest for propulsion system design were listed. For ODWE, PMSIC and RAM designs, the phenomena of interest are:

1. Shock-to-detonation transition.

2. Detonation wave stability and sensitivity to upstream disturbances.

3. Post detonation wave structure.

For PDE engine, the phenomena of interest are

1. Deflagration-to-detonation transition.

2. Ignition energies and transition distances.

3. Sensitivities to non-uniformities upstream.

"An understanding and prediction capability of the above phenomena must be in-hand before rational design of wave engines can proceed." Basic research to advance understanding in the above areas is hence crucial.

Dr. M. M. Sindir. This talk addressed the technology needs in the areas air-breathing and liquid-fuel propulsion systems and lasers. The various physical processes that need to be modeled were identified. The models currently used to simulate these physical processes in the design of propulsion systems, given the name state-of-the-practice models, were listed.

High-speed air-breathing propulsion system. This propulsion system is typically of relatively simple three-dimensional geometry and is intended to operate in a wide range of Mach numbers ($0 < M < 25$). Propulsion is produced by single phase gaseous combustion and the engine is made up of the following components: (i) forebody; (ii) inlet: (iii) combustor; (iv) internal nozzle; and, (v) external nozzle. Some of the important physical processes that need to be understood for effective and efficient design of the propulsion system and the relevant engine component (in parenthesis) are listed below:

1. Transition from laminar to turbulent flow (forebody and inlet).

2. Non-equilibrium air chemistry (forebody).

3. Turbulence with

   - Shock-on-shock interaction (inlet).

   - Shock-boundary layer interaction (inlet).

   - Separation and unstart (inlet).

   - Scalar mixing and combustion over a wide Mach number range (combustor).

   - Very high thermal loads (combustor).

   - Relaminarization (nozzle).

   - Anisotropy (nozzle).

4. Turbulence-chemistry interactions (combustor).

Liquid propellant rocket engine. This propulsion system typically features rotating machinery and complex three-dimensional flow geometry. This type of rocket engine is designed to operate in incompressible to low supersonic Mach number (< 2.0) regime. The main engine components in this case are: (i) Pump; (ii) Turbine; (iii) Injector/Combustion chamber (CC) ; and, (iv) Nozzle. The important physical processes are:

1. Transition (turbine).

2. Turbulence with

- Vortex shedding and losses in blade wakes (turbine).

- Transient effects on blade loading (turbine).

- Surface heat transfer (turbine).

- Large regions of separation (pump).

- Rotation induced secondary motions (pump).

- Large anisotropy (pump and unconventional 3D nozzles).

- Cavitation effects (pump).

- Sudden accelerations and decelerations (pump).

- Strong vortical motions (pump).

- Large density variations (CC).

- Effect of droplets/particle loading (CC).

- Dispersion of droplets/particles (CC).

- Non gradient-transport scalar phenomena (CC).

- Multi-phase mixing (film-cooled nozzles).

- Large separated flow regions (high aspect ratio nozzles).

3. Turbulence-chemistry interactions with:

- Combustion initiated/driven turbulence (CC).

- Intermittent reactions due to turbulent mixing (CC).

Laser Systems in aeronautics. These systems are of relatively simple geometries. The chemistry is characterized by fast kinetics and large heat releases. The important thermo-fluid processes are:

1. Turbulence with:

- Subsonic and transonic boundary layer/wake interactions.

- Atmospheric effects.

- Large density fluctuations.

- Relaminarization.

2. Turbulence-chemistry interactions.

In the design process, the models (state-of-the-practice models) used to simulate the above physical phenomena are typically:

1. Empirical correlations for transition modeling.

2. e^N models for non-equilibrium air chemistry.

3. Zero- and two- equation turbulence models (such as, Baldwin-Lomax, $k - \varepsilon$, $k - \omega$, and point-wise models) along with largely ad hoc compressibility corrections.

4. Laminar thermochemistry (i.e., no turbulence-chemistry interaction modeling) for chemical reactions.

The state-of-the-practice model used in integrated vehicle computations is the zero-equation turbulence model. Clearly, there is a large gap between the above state-of-the-practice models used for design and the state-of-the-art models developed in universities and research institutions such as ICASE.

As regards the role of universities and research institutions in the area of basic research and interaction with industry, Dr. Sindir had the following recommendations.

1. *Reynolds stress modeling.* Research to be continued at some level with the expectation of evolutionary improvements in the areas of (i) wall effects, (ii) compressibility and combustion, and (iii) multiphase flow interactions.

2. *Large eddy simulations.* This "represents the future of turbulence closures for engineering applications." Specific areas of possible research in this area include (i) unstructured grids, (ii) algorithms for massively parallel machines, and (iii) models of flow physics.

3. *Direct numerical simulations.* This approach is, at present, only a basic research tool and no direct engineering applications are anticipated in the foreseeable future. As a result, "no collaboration with industry can be expected until revolutionary changes in computer technology materializes."

Dr. P. Givi. He spoke on the research efforts of his group in SUNY Buffalo and pointed out some of the obstacles in the way of University-Industry interactions. Some of the state-of-the-art turbulence models developed in academia for the chemical manufacturing industry and propulsion industry were listed. Most notable among them are:

1. New algebraic Reynolds stress closure models that show distinct improvement over traditional gradient-diffusion models.

2. Scalar-velocity probability density function (pdf) models.

3. Reliable reduced kinetics schemes.

Some of the difficulties associated with doing academic research are:

1. Lack of enthusiasm among some industries to provide resources.

2. Impossibility of conducting unfunded research.

3. In some cases, significant emphasis among funding agencies on getting 'industry blessings'.

4. Conflict of basic research funding with quarterly profits.

LIST OF RESEARCH TOPICS

The two industry speakers provided comprehensive lists of research topics that they would like to see performed in ICASE/LaRC. The list topics can be found above and are not repeated here. One major message that came across loud and clear from the industry participants was that, as far as industry interest is concerned, how basic research is performed is as important as what basic research is performed. Basic research performed will be useful to industry, if, and only if, the ideas and models developed are thoroughly validated and their limitations well understood and documented. At the present time, industry does not have the time or resources to deploy its own personnel to examine the reliability of the ideas and models developed by academia and research institutions. *Any research that stops short of providing models that have been tried and tested in a variety of realistic benchmark problems is of little use to industry. Before a model can be employed as a design tool, its capabilities and limitations have to be well understood.*

LIST OF WORKSHOP TOPICS

Many of the research problems listed by the industry speakers would make good workshop topics. Two areas, in particular, are well suited for investigation in workshops: (i) evaluation of various turbulence-chemistry interaction models for a few benchmark cases; and, (ii) development of understanding of shock-to-detonation and deflagration-to-detonation phenomena.

SESSION A5: MODELING OF NEWTONIAN AND NON-NEWTONIAN TURBULENT (INDUSTRIAL) FLOWS

Thomas B. Gatski, NASA Langley

SUMMARY

The presentations in the Modeling of Newtonian and Non-Newtonian Turbulent (Industrial) Flows Session focused on aerodynamic flow over automobiles, hydro-dynamic flow over ship hulls, and artificial heart and assist devices for the human circulatory system. This wide diversity of turbulent flows sensitized both the speakers and audience to the wide application base for turbulent flow prediction and control. In addition, the complexity of the flow fields clearly indicated that higher-order closure models accounting for both turbulent stress anisotropy effects and turbulent history effects were necessary to properly predict the details of the flow.

INTRODUCTION

This diversity of topical areas was excellently represented and discussed by four speakers from a broad cross-section of industrial interests. The presentation on automotive flow fields was made by Dr. W. Brandstatter of AVL Powertrain Engineering, Inc. who also showed that automatic mesh generation methods were an important link in effective CFD research and development. Hydrodynamic applications were discussed by Dr. Nils Salvesen of Science Applications International Corporation. Increased use of CFD and advanced turbulence models in ship design has been necessitated by the need to know the near-field turbulent flow field and hydrodynamic loads experienced by the ship. Non-traditional applications were discussed by Dr. James Antaki of the University of Pittsburgh Medical Center and Dr. Gerson Rosenberg of the Hershey Medical Center. Dr. Antaki presented the work being conducted on a rotary pump assist device as an alternative approach to heart replacement, and Dr. Rosenburg presented the work on the development of a totally implantable artificial heart. In both cases, the complex flowfield is highly turbulent with the complicating feature of a non-Newtonian fluid medium.

PRESENTATIONS

Quantitative flow results were shown for the flow over an automobile with specific emphasis placed on the incorrect predictions obtained from isotropic eddy viscosity models relative to higher order closures. For example, the details of the vortical shedding from

the rear of the vehicle were completely in error with a two-equation turbulence model, but were computed correctly using a Reynolds stress model developed by Speziale, Sarkar and Gatski. Although drag calculations were less sensitive to the particular model used, lift on the vehicle was sensitive and found to be erroneously predicted by the simple two-equation model. Computational issues were also addressed with the overall conclusion that computational penalties using advanced closures such as Reynolds stress models were minimal and convergence characteristics similar to two-equation models could be obtained.

In the hydrodynamic application, a brief tutorial was presented of the problems facing the ship designer. The turbulent flow application focuses on the stern of the ship where Reynolds stress models are required to correctly predict the attached flow features. Utilizing isotropic eddy viscosity models predicted premature flow separation in contrast to experimental observations. While higher-order closures such as Reynolds stress models have not been traditionally used in the ship design area, new competitive forces require shipyards to become more involved with better design techniques. Whether this increased use of technology will lead to such sophistication as full Reynolds stress closures will be dictated by the economic impact of improved design techniques. At present, improved prediction of the near-field turbulent flow including the interaction between the ship's boundary layer, the nonlinear free-surface and the propeller is needed as well as improved prediction of the wave-induced ship motions, the hydrodynamic loads and ship responses.

Contrasting the traditional fluid dynamic applications were the presentations on non-Newtonian fluids. These talks focused on the critical problems of designing alternatives to a diseased or a failing human heart. Two approaches have been taken. As an alternative to heart replacement, a rotary, impeller-type pump component has been designed. The optimal design of such a device involves several fluid dynamic considerations including the accurate prediction of the turbulent fluid flow through the device. There appears to have been no computational studies of the turbulent blood flow through such pumps. Studies have begun on a CFD-shape optimization algorithm which automatically modifies the shape of an initial candidate blood path. Improved knowledge of the turbulent flow within the pump will aid in developing better criteria for design optimization.

Another approach to the heart problem is the development of a permanent, totally implantable artificial heart. At present, this device has been able to sustain animals for periods greater than one year. The goal is to have a device which will function for up to ten years. These devices contain turbulent flows which are extremely complex due to both the bounding geometries and the fluid itself. The crucial design problem related to fluid dynamics is the optimization of the thickness and shape of the blood sac to withstand about 5×10^7 cycles

per year that the pumping chamber experiences. The stresses experienced by the sac need to be examined through detailed fluid dynamic studies of blood flow within the valves and sac.

LIST OF RESEARCH TOPICS

The following are general flow features which may require additional research:

- Non-equilibrium Turbulent Flows

- Flows with Strong Swirl

- Flows with Separation

- Three Dimensional Turbulent Flows

- Turbulent Flow of Viscoelastic Fluids

LIST OF WORKSHOP TOPICS

Industry is constantly referring to the need for turbulence models to handle their complex flow problems. I now believe that it is time to quantitatively assess the performance of turbulence models on these complex industrial problems. It is not clear to me that anyone has performed an adequate evaluation of the closure models available to unambiguously conclude overall performance levels.

SESSION B1: MANUFACTURING PROCESSES IN INDUSTRY

Manuel Salas, NASA Langley

This report is unavailable.

SESSIONS B2 & B6: Cluster Computing

David E. Keyes, ICASE

SUMMARY

The ICASE/LaRC Industry Roundtable, October 3-4, 1994, in Williamsburg VA, had two sessions (B2 & B6) on cluster computing. They are summarized here for the purposes of documenting examples of ways in which cluster computing has already matured, providing some points of contact, creating dialog, and motivating future collaborative research between industry, academia, and ICASE.

PRESENTATIONS

There were seven speakers in the two two-hour sessions, one each on the first and second days of the Roundtable. Each speaker was invited to present material of their own choice for twenty minutes and then to guide discussion for another ten minutes. The seven presentations were followed by about an hour of free-form discussion after the final presentation.

A single cluster computing session was originally planned, to highlight instances in which cluster computing had already penetrated the industrial mainstream, particularly for field problems that are typical of NASA's computational fluid dynamics (CFD), computational structural mechanics (CSM), computational electricity and magnetism (CEM), and multi-disciplinary optimization (MDO). As we became aware of additional non-industry expertise and interest, we broadened the coverage of this topic by inviting some additional speakers.

Four of the speakers were from industry and the remaining three were from academia. Of the industrial speakers, three were from consumers of cluster technology, and one was from a vendor of cluster technology. All three of the academic speakers were also affiliated with government laboratories.

The speakers and their affiliations and presentation titles were, in order of presentation:

- **David Keyes**, *Computer Science Department, Old Dominion University and ICASE.*

- **Bill Coughran**, *Computational Mathematics Research Department, AT&T Bell Laboratories.*

- **Doug McCarthy**, *Boeing Propulsion Research, Boeing Commercial Airplane Group.*

- **Chae Rhie**, *Pratt & Whitney.*

- **Louis Turcotte**, *NSF Engineering Research Center for Computational Field Simulation, Mississippi State University and USAE Waterways Experiment Station.*

- **Andy Sherman**, *Scientific Computing Associates.*

- **Jay Jayasimha**, *Computer and Information Science Department, Ohio State University and NASA-LeRC.*

SYNOPSES OF PRESENTATIONS

Keyes introduced the cluster computing session by describing the three "eras" of parallel computing as, "Solve my problem faster, solve my problem bigger, and solve my problem cheaper." The first of these is a natural goal for a fixed-size problem, but is eventually defeated by Amdahl's Law. The second aims for good efficiency in the sense of Gustafson's Law – in a fixed amount of time larger and larger problems may be solved, in proportion to the number of processors contributing their memory. The third objective recognizes that the resources required to handle large problems are expensive, and can be better justified in many research environments if the same hardware can be used some of the time for parallel solution of large problems and at other times as a collection of independent workstations.

Keyes then proceeded to discuss a pair of recent ICASE technical reports on the feasibility of using workstations on an Ethernet for parallel computing. Many parallel algorithms have been demonstrated to perform in a scalable manner by theoretical complexity models and by experimental observation on tightly-coupled massively parallel processors (MPPs). On the other hand, some of the speaker's personal experiences in implementing parallel algorithms (for two-dimensional aerodynamic and geophysical flows) on workstation clusters have been disappointing. Furthermore, the problem of modeling workstation cluster performance is complicated relative to the tightly coupled, dedicated-processor/dedicated-link parallel machines by two factors: (1) owners must be allowed to preempt computation occurring on the workstations at which they sit, and (2) available networks are usually multilayer and subject to contention. The two ICASE reports #93-65 and #94-78, respectively, attempt to study each of these two factors in isolation.

S. Leutenegger and X.-H. Sun, in #93-65, assume that communication is instantaneous and report the results of discrete simulation of a computation of an abstract parallel task, consisting of independent load-balanced tasks on each processor and a terminal synchronization. (Such an abstract task arises repeatedly, as a loop body, in many parallel codes, such as applying a data-parallel preconditioner in a conjugate gradient method.) They assume that owner tasks are binomially distributed in length and preemptive, with the exception

28

that after each owner task completes the parallel task is allowed some minimum quantum of time before again being preempted. They introduce several dimensionless metrics and simulate from one to hundreds of workstations. The paper's quantitative performance data as a function of owner preemptiveness and the size of the parallel task quantum should be consulted for details. From Keyes' viewpoint, they show that even rather modest random owner preemptiveness prevents attainment of worthy parallel efficiencies, because of the accumulated effects of idling before the synchronization barrier. However, utilization, defined as the percentage of recovery of otherwise wasted cycles between owner tasks, is still impressive.

In #94-78, Keyes and two graduate students at ODU pursue the opposite extreme of dedicated nodes, but more realistic communication models. They develop a "communication calculus" for combining the latencies and transfer times of messages transiting down, across, and up the layers of a multilayer communication protocol, taking into account serial and parallel communication pathways and contention. For a variety of archetypal communication patterns (broadcast, nearest neighbor, global reduction, and synchronization), they fit parameters of a communication model to Sun sparcstations on an Ethernet over a range of message sizes from 16 Bytes to 16 KBytes, showing excellent regression coefficients with just a two-parameter model. They then apply the model to two PDE algorithms: domain decomposition and time-parallel multigrid. Domain decomposition, with small surface (communication) to volume (computation) should do well on a dedicated workstation network. Time-parallel MG, with equal order communication and computation volume, should be limited by contention, and is. By using the model to predict the communication times at maximum possible network contention in the "bursty" exchanges of these algorithms, and comparing the model with measurements, they are able to measure the degree to which the application drifts out of synchrony.

Keyes concluded with the observations that modeling cluster performance was considerably more complex than modeling dedicated parallel systems, and that the performance, itself, was vulnerable to non-dedicated processor effects and poor performance of a primitive network like Ethernet.

Next, the three industrial consumers of cluster computing were invited to describe their own experiences in porting practical computations to clusters, and to offer projections and recommendations for the future.

Bill Coughran described joint work with Petter Bjorstad and Eric Grosse at Bell Labs in which an idealized two-dimensional semiconductor device was modeled on SGI workstations connected by a Fiber Distributed Data Interface (FDDI). At close to 12 MB/s (100 Mb/s), FDDI is approximately ten times faster than Ethernet. (An upgrade of this testbed to

29

Asynchronous Transfer Mode (ATM) is planned in the near future.)

Coughran devoted the opening part of his presentation to displaying the the drift-diffusion equations of electron and hole transport, a multi-component system of highly nonlinear second-order PDEs, and described an iterative domain decomposition algorithm based on conjugate gradient iteration (BiCGSTAB) and Schwarz preconditioning. PVM was able to use only a small fraction of the available FDDI bandwidth, even for simple point-to-point communication, and so was abandoned. The communication needs of this algorithm (at least on a regular partition of the grid) are sufficiently simple and regular that a small set of communication programs built directly upon Unix sockets suffices. One particular performance enhancement worthy of dissemination is the batching together of multiple global reduction operations by a programmer instruction to the communication package, which indicates that there is no serialization intended by the conventional high-level programming style expressing each inner product with an individual subroutine call. In other words, the global summation communication is moved from inside to outside of the individual reductions. Fixed-size efficiencies of 75% were realized on cluster of four workstations.

Coughran testified that whereas specialized MIMD MPP architectures have had a mixed reception at Bell and have not succeeded in displacing the Cray for production runs, the workstation cluster is "changing things."

McCarthy offered observations from two years of experience on workstation clusters in Boeing propulsion, where the computing load consists primarily of multiple-parameter studies of workstation-scale problems, as opposed to "Grand Challenge" runs. He characterized the idle fraction of individual workstations as 80% of the week before Boeing began experimenting with parallel and distributed queuing systems, and showed sample workload statistics now routinely collected at Boeing indicating that the workstations in his cluster are now usefully busy about 80% of the week. Some of the cluster work is actual parallel production, using a not-expressly-load-balanced form of nonlinear overlapping Schwarz. Finding this sort of unbalanced parallelism is easy in block-structured Euler and Navier-Stokes codes, such as the wing-nacelle external compressible flow studies that constitute one of the staple problems of his research division. Since performance is less of an issue than convenience and cost-effectiveness, clusters work well. The main remaining part of the propulsion work is multiple-parameter studies, batched serially to free workstations in a globally coordinated way, but without synchronization requirements, except for terminal collection of statistics. This is also a natural setting for clusters.

McCarthy related various "social" problems caused by cluster computing, and also some effects of his group's adoption of cluster computing. One effect was that the Cray processing rates that his division was required to pay dropped to one-third of their original. Another

was that many of the most idle workstations in the company – those in marketing – were held off-limits to the cluster computing researchers for a time, because of the necessity of owners having exclusive use of the stations whenever the need was nontrivial. In conclusion, McCarthy proposed six areas of joint industry/academic investigation, which are included in Section 4 below.

Rhie followed with a third strongly enthusiastic endorsement of cluster computing, drawing upon a two- or more year history of porting compressible internal flow Navier-Stokes codes to networks of workstations at Pratt & Whitney. Today, up to approximately 16 workstations are ganged together into a distributed memory supercomputer each night, and a week of such nightly runs (stored into restart disk files during the day) may be required for converged solutions on the largest problems routinely solved in this manner, namely design studies of complete multistage compressors and turbines. Full engine simulation (including the combustor) is being contemplated, but the challenges are enormous and Rhie proposed strong teamwork between industry, NASA, workstation vendors, and academia.

Rhie noted that Pratt & Whitney's parent company, United Technologies, has done away entirely with in-house supercomputers of vector and MPP type, and now buys time from the National Center for Supercomputer Applications (NCSA) for the small fraction of large memory jobs that cannot conveniently be ported to workstations.

Turcotte, who has done an extensive literature review on workstation cluster demonstration projects and available coordination software, and who has been instrumental in the annual cluster computing workshop at Florida State University, analyzed the several independent forces behind the growing popularity of cluster computing. He identified four influences: technological, market, teraop environmental, and political. The principal technological influence is the doubling time of just two years for the number of transistors per die on memory and processor chips, permitting huge memories per processor, and allowing the clock rates for microprocessors to rise much faster than the processors at the top of the technology envelope. In fact, the specialized high-performance processors have nearly been overtaken. For instance, the TFP chip achieves nearly the same LINPACK-1000 benchmark rating as a single processor of a Cray Y-MP.

The fact that the worldwide market for MPPs was only \$300M in 1993, compared with a worldwide workstation market of \$12B (forty times larger), indicates that workstation technologies will "drive" development efforts. After a period of many years in which personal computers were the fastest growing segment of the market, workstation revenues are now growing just as fast. In confused markets, such as is caused by the current tightening of federal-funded procurements, conservative approaches win.

31

Turcotte also reminded listeners of the facts of life that have been discovered in the somewhat chaotic evolution of MPPs, in pursuit of the teraop. Few MPPs have been found balanced between memory, processor capability, communication speed, and I/O and storage capabilities. Clusters come "automatically balanced" except, possibly, for interprocessor I/O. Many of the MPPs can only be purchased and upgraded in inflexible increments, such as powers of two in processor number. Clusters, by not relying on a particular network topology, can be assembled more modularly. Whereas clusters use highly evolved commodity software, invested in MPP software is quite poor. The mean time between failures (MTBF) rating of most MPPs is dismal compared with workstations. Rapid obsolescence of MPPs allows little leverage of previous efforts, on the part of both vendors and users. It has been difficult to re-educate the user community. And little commercial software has been provided by third parties.

Turning to political considerations, he quoted liberally from four important recent federal studies or defining documents on high performance computing: a GAO study, a CBO study, an NSF study, and the recent "blue books" of the President's OS&T. In each can be found negative statements regarding the future of MPPs and/or encouraging remarks about the potential of cluster computing. Finally, Turcotte presented the "HPC pyramid" from the NSF study "From the Desktop to the Teraflop," and noted that clusters span the majority of the capability extent of the pyramid as well as the capacity extent. There are some problems at the top of the pyramid that are out of reach of cluster computing, but these are "Grand Challenges" being pursued by only a small percentage of computational science and engineering users.

Representing a vendor that licenses one of the earliest and best known coordination languages, Sherman presented the *Linda* and *Paradise* virtual shared memory (VSM) software environments. The current *Linda* is designed for parallel computations in which performance matters enough for the user to take control of data locality, and in which there is a definite beginning and end of each execution. In contrast, *Paradise* is a persistent associative data base that facilitates interprogram communication and synchronization. *Linda* has been ported to all of the major shared and distributed memory parallel systems, as well as network systems. It supports all of the major programming paradigms, including master/worker, data parallel, task parallel, and message-passing. Sherman briefly described the semantics whereby programs in high-level languages, such as FORTRAN and C, interact with the virtual shared memory, known as "tuplespace," through associative operations, including "wild card" matches.

Sherman concluded by showing speedup curves for six diverse applications (seismic inversion, vibration analysis, rotorcraft fluid dynamics, molecular simulation, ray tracing, and mortgage-backed securities pricing), the least successful of which achieved a speedup of five on eight processors.

Finally, Jayasimha presented a detailed case study of a single problem, an explicit compressible Navier-Stokes code based on MacCormack's method, on a single cluster, NASA's Lewis Advanced Cluster Environment (LACE), using four different networks on either 16 or 32 IBM RS6000s. The network environments tested included PVM on a public Ethernet, on a private Ethernet (10 Mb/s), on IBM's Allnode switch (32 Mb/s per channel), on FDDI (100 Mb/s), and on a Fiber Channel Switch. Sensitivity to the number of startups, overlapped communication and computation, and the staggering of otherwise bursty traffic permitted sublinear but monotonic speedup for this prototype CFD problem on the two fastest networks.

DISCUSSION

The seven presentations with questions took place in a little more than 3.5 hours over the two days. The Roundtable discussion, some of which spilled over after the meeting had officially ended, perhaps indicates that significantly more discussion time should have been built into the sessions. A (probably very incomplete) summary of the discussion topics is as follows:

1. Is a consensus emerging on the vocabulary of cluster computing? (Turcotte's talk had listed twelve near synonyms, and Keyes' introduction had been interrupted for a clarification of whether MPPs were meant to be included under "clusters.")

2. Is a consensus emerging on the most useful programming paradigms (e.g., user-controlled message passing, global virtual address space, global associative memory)?

3. Is a consensus emerging on the best packages for each of the paradigms, and are critical mass effects beginning to favor certain packages within certain user communities?

4. Which of the various networks and network protocols is the most cost-effective for a given granularity and communication load?

5. Will the new message-passing interface (MPI) standard provide a long-awaited platform independence for distributed memory programs and blur the distinction of developing for clusters versus developing for MPPs?

33

6. What is the largest practically achievable parallel granularity for a nearest neighbor code, such as an explicit CFD code?

7. What is the largest practically achievable parallel granularity for a code with global operations, such as an implicit CFD code?

8. What is the median degree of parallelism desired by engineering users, from the point of view of data structures and programming convenience?

9. What are the social and human factor impediments to workplace cluster computing?

10. Can a dedicated MPP-like machine be "built" by procuring a rack of fast commodity CPUs, without the associated monitors and user-interface peripherals?

11. Why are the "Big Three" automakers represented at the Roundtable adding conventional supercomputing capacity at the same time the aerospace companies are turning away from in-house support of a conventional supercomputer? Are there fundamental technical reasons, rooted in different problem classes, or are the reasons rooted in corporate culture and central computing price-structuring?

12. What provisions are there in cluster computing for working with large-scale distributed data sets, in terms of archiving distributed restarts, distributed data base inquiry, and real-time and post-processing visualization of distributed sets?

PROBLEMS FOR COLLABORATION

The industrial speakers were asked to identified some "pre-competitive" problems for collaborative research with ICASE/LaRC and with academia. This request was undertaken with considerable thoroughness by some, and a variety of interesting problems were proposed:

1. Making cluster computing systems safe for owners: permitting rapid preemptiveness, insuring security, minimizing the possibility of crashes or deadlocks brought about by parallel applications gone awry, and prioritizing network traffic so that cluster jobs do not degrade remote file or remote host access during certain work periods.

2. Making cluster computing systems safe for users: facilitating migration of processes that are shut out on a preempted processor, monitoring distributed (perhaps migrating) jobs, user-directed checkpointing and automatic routine checkpointing for restart recovery, restart techniques, and standards

3. Cluster computing tools: flexible process scheduling and mapping, parallel emulation, debugging and performance visualization tools, tools for working with distributed data sets.

4. Cluster computing on networks of multiple-CPU computers, including heterogeneous MPPs.

5. General algorithmic research motivated by the high-latency non-dedicated environment: useful paradigms for PDEs besides domain decomposition, chaotic relaxations, dependence of convergence rates on interprocessor update frequency, partitioning strategies (including adaptive solution-dependent schemes) that minimize communication needs, dynamic repartitioning, and implicit versus explicit methods.

6. General systems research motivated by the high-latency non-dedicated environment: multiple threads, prefetching, message bundling.

7. The use of parallel networks to cost-effectively and flexibly "grow" network capacity.

Not all of these projects would be appropriate for ICASE or NASA. Some would seem most appropriate for vendors; others for systems software research in the abstract.

LIST OF WORKSHOP TOPICS

Two types of workshop topics arise naturally from the Roundtable discussions: supplier-oriented and consumer-oriented.

Multiple Threads. A supplier-oriented workshop seems both desirable and feasible on the topic of multiple threads. The communication penalty for cluster computing with the networks readily available today is too great for parallel scalability (even defined very liberally), and the penalty is widening as processing rates grow more rapidly and more cost-effectively than reliable communication rates. Multiple threads, perhaps along with other latency-hiding or latency-reducing programming and operating system techniques, such as software or hardware prefetching, and compile-time and run-time message bundling, might help solve this problem. ICASE has had a multiple threads research program for over a year, which is currently attracting its first (academic) customers. Workstation and software suppliers might like to experiment to determine the potential payoff of and to collaborate in the development of practical multiple-threads environments.

35

The consumer-oriented workshop topics include:

A "B.Y.O.C." Workshop on Implicit Cluster Computing Methods. User training in the extra software of cluster computing is required before important scientific and engineering codes are ported into the cluster environment. ICASE could sponsor a "bring-your-own-code" workshop, particularly focused on implicit methods, which are harder than explicit methods to effectively parallelize. Implicit methods are increasingly important in aerodynamics and other areas because envelope problems for which increased computer power is needed tend to have larger ranges of spatial and temporal scales than their quotidian counterparts. Recommended environments for ICASE to use and assist others in using are MPI and the Argonne National Lab Portable, Extensible Toolkit for Scientific Computing (PETSc).

Cluster Computing Support for MDO. Multidisciplinary optimization has communication characteristics that may be very different from those of other large-scale scientific computing tasks. MDO codes often call individual large-scale disciplinary codes as inner procedures. These disciplinary codes are likely to store very different data sets in different formats, and they may run most naturally on different machines. Cluster computing tools that can interface such data sets and control the different disciplinary programs would be useful, even apart from considerations of performance.

SESSION B3: COMPUTATIONAL FLUID DYNAMICS

Jim Thomas, NASA Langley

SUMMARY

Three papers were presented in this session on computational fluid dynamic (CFD) activities in industry, following a short description of ongoing work at Langley Research Center by the session chairperson. The first was a presentation from an aerospace perspective by Dr. R. Agarwal of McDonnell Douglas Aerospace (MDA); the second and third were from automotive perspectives by Dr. T. Gielda of Ford Motor Company and Dr. R. Sun of Chrysler Corporation, respectively.

PRESENTATIONS

R. Agarwal of MDA presented single discipline issues, as well as multidisciplinary issues, in aerospace industrial use of CFD codes. The CFD elements addressed were surface modeling, grid generation, flow solution, physical modeling, validation, and multidisciplinary coupling. Application to MDA configurations, including F-18, C-17, and High Speed Civil Transport, demonstrated the inroads that CFD has made in recent years.

MDA has a well-developed grid capability, including both structured- and unstructured-grid capability; the structured-grid capability is mature and embodied in a single code MACGS, which has a user interface that is widely used in the company. The largest unresolved issue is adaptation for structured-grid codes since the relative efficiency of such codes, as compared to unstructured-grid codes, is better for high Reynolds number viscous flow applications. The flow solver requirements for improvements at MDA are in the areas of efficiency, turbulence modeling, parallel processing, and multidisciplinary coupling. From their perspective, industry needs better numerical algorithms for Reynolds-Averaged Navier-Stokes equations, as opposed to more codes, echoing a similar message in the keynote address of Rubbert. Multigrid methods have been demonstrated to achieve $O(N)$ convergence in some cases, but better smoothing operators are needed. There are recent advances in matrix solvers which are finding use as direct solvers or as preconditioners. In the area of parallel processing, a 5 million point CFD solution with 57 zones was demonstrated on a cluster of 15 HP715 workstations to achieve 40 microseconds/grid point-iteration, as opposed to 28 microseconds/grid point-iteration on a Cray C90. Scalability remains an issue, although a 66 percent efficiency was achieved in this case.

The physical modeling needs are in the areas of turbulence, chemistry, electromagnetics, and acoustics. In regards to turbulence modeling, the physical model desired is the simplest

one that captures the physics. The industrial use of direct numerical simulations for large-scale computations is very far off in the future. Assessments of given models suffer currently from lack of grid convergence; more benchmark validation case studies are needed, such as the NASA Langley high-lift workshop. Aeroelastic effects are also becoming more important for assessment of turbulence models, as tolerances for validation are decreasing.

Multidisciplinary coupling of aerodynamics and controls was demonstrated for the six-degree-of-freedom release of a store. Dr. Agarwal envisioned the multidisciplinary activities as in their infancy in the company, especially in regards to usage of large-scale CFD codes for the aerodynamics methodology.

T. Gielda from Ford Motor Company discussed CFD and its role in interior analysis for cars. CFD is beginning to play a more important role; grid generation and CFD analysis is one of the nine areas addressed in the cooperative agreements recently signed between Ford and NASA Langley Research Center. Since tooling costs are huge, CFD is used as a screening tool for the design of interiors from the standpoint of thermal comfort, acoustics, and ergonomics. Styling, rather than function, rules in the current design environment and trend predictions from CFD with a 75 percent accuracy are acceptable. Often zeroth order effects are needed and calibration is secondary. The savings from mass production can be enormous; a 10 pound per square inch reduction in compressor head pressure can translate to a 1 dollar savings in warranty, which translates to a 15 million dollar savings over the total engine sales. Ford's needs differ from much of the CFD community, since their need for robust shock-capturing algorithms is very small. They have 32 Cray C90 processors at Ford; algorithm improvements are required in the areas of high order algorithms, time accuracy, efficiency, and variable fluid properties.

Two case studies were presented. The first was for an underhood cooling device. The initial CFD studies did not predict the flow correctly as determined from wind-tunnel tests. However, a simple ejector was found which improved air flow and reduced the cooling inlet temperature. The external drag of the vehicle at cruise was reduced,, which was ultimately the deciding factor in the installation onto a production vehicle. The device was installed on the Thunderbird vehicle and also resulted in a transmission fluid temperature decrease.

Interior analyses of airflow predict the rate of change in temperature with time as a function of vent design and location. This is an area where time accuracy and turbulence modeling improvements are needed. The actual interior flow is transitional in the vehicle. There is too large a range in predicted time from a design viewpoint to attain a given setting of interior temperature between assuming either fully laminar or fully turbulent flow calculations; thus, transition models are necessary. The high-Reynolds number versions of two-equation models are not sufficient. The large-eddy simulation technique may find

applications at Ford, especially in the acoustics area. Two-phase flows and combustion model improvements are also needed.

Dr. Sun presented an extensive list of CFD requirements and needs at Chrysler, including engine intake and exhaust systems airflow, water jacket coolant flow, in-cylinder airflow and combustion, underhood airflow and thermal management, climate airflow, and wind noise. Currently, 80 percent of the CFD time is devoted to grid generation. The surface models are generated with CDRS and CATIA and ICEM-CFD is used for grid generation. A number of analysis codes are used, including AIRFLO3D, FIDAP, CFD-ACE, STAR-CD, FIRE, NASTRAN, and WAVE. Third-party developers are used extensively because of the available software support. Chrysler has 8 Cray C90 processors and a large number of workstations for visualization. Crash simulations account for 35 percent of the computing, with CFD and linear structures the next major two. A video was shown illustrating CFD solutions for the external flow over a Dodge Intrepid, a windshield de-icing simulation, an exhaust manifold airflow, and a catalytic converter simulation.

A short-term project which Chrysler viewed as an important technology transfer area with ICASE/Langley is for unstructured-grid external automotive applications in ground effect. To accomplish this transfer, site visits by principal developers from each organization are desirable. The long-term areas were wind noise predictions, database accumulation for assessment of codes, and continuing interactions of lead developers between industry and government labs.

LIST OF RESEARCH TOPICS

- Evaluation of unstructured-grid methods for automotive external simulations with ground effect

- Simulation of wind noise for automobiles

LIST OF WORKSHOP TOPICS

- Turbulence modeling

- Massively parallel processing/network computing

- Computational aeroacoustics

- Receptivity in transition

- Neural nets for active control multidisciplinary optimization

SESSIONS B4 & B5: MULTIDISCIPLINARY DESIGN OPTIMIZATION

Natalia Alexandrov, NASA Langley

Jaroslaw Sobieski, NASA Langley

SUMMARY

Multidisciplinary Design Optimization (MDO) was represented at the ICASE/ LaRC Industry Roundtable, October 3-4, 1994, in Williamsburg, VA, by two lecture sessions and an informal roundtable discussion. This report provides a summary of the presented topics and the roundtable exchanges which served to identify the problems faced by the practitioners of MDO and to motivate collaborative pursuit by the industry, ICASE, academia, and NASA of the research directions that will address these problems.

INTRODUCTION

The two MDO sessions took place on the second day of the Roundtable and contained seven speakers. The organizers planned to give brief introductory presentations, while the speakers from industry were invited to give 25-30 minute presentations. Each presentation was followed by a short question-and-answer period. The first MDO session was followed by a well-attended and productive informal roundtable discussion that lasted for over an hour and served as a transition to the second MDO session.

The organizers requested that the speakers acquaint the audience with the status of MDO in their industries, to provide an overview of important unsolved problems, and to propose problems for collaborative research with the ICASE and NASA colleagues that would facilitate further implementation of MDO in industry.

The two organizers were affiliated with ICASE and NASA LaRC. The five industrial speakers represented a range of companies, such as automotive, aerospace, and chemical engineering.

The speakers and their coordinates follow in the order of presentation:

- **Natalia Alexandrov**, *MDO Branch, Mail Stop 159, NASA Langley Research Center, Hampton, VA 23681-0001.*

- **Jim Bennett**, *Principal Research Engineer, Engineering Mechanics Department, General Motors Research Laboratories, 30500 Mound Road, Warren, MI 48090-9055.*

- **Bhieng Tjoa**, *Research Scientist, Dynamic Matrix Control Corporation, 9896 Bissonnet, Houston, TX 77036.*

- **Jaroslaw Sobieski**, *MDO Coordinator, Mail Stop 242, NASA Langley Research Center, Hampton, VA 23681-0001.*

- **Peter Finnigan**, *Research Manager, The General Electric Research Center, Hampton, VA 23681-0001.*

- **Joseph Giesing**, *Chief of Structural Engineering, McDonnell Douglas Aerospace Corp., 3855 Lakewood Blvd., MC 71-34, Long Beach, CA 90846.*

- **John Volk**, *Leader of CAD Development Group, Northrop Grumman Corporation, B-2 Division, 8900 East Washington Blvd., Pico Rivera, CA 90660-3737.*

PRESENTATIONS

The following summary is based on the notes taken during the presentations and the presentation abstracts.

N. Alexandrov opened the MDO session by introducing the emerging MDO program at ICASE. The objective at ICASE is to identify unsolved problems and promising research areas and to define a systematic program in applied mathematics for the development of techniques that would facilitate wider implementation of the MDO methodology.

The discovery process consists of collecting information during interaction with engineering and applied mathematics researchers, as well as during formal gatherings of MDO specialists, such as the Roundtable and the ICASE/LaRC MDO workshop organized to take place on March 13-16, 1995 in Hampton, VA.

Since, from an applied mathematician's point of view, MDO problems are very large, complex, constrained nonlinear programming problems, by the date of the Roundtable, the research efforts at ICASE were concentrated in the area of multilevel methods for large-scale constrained optimization (ICASE Report # 94-53, Alexandrov and Dennis). These methods provide a robust, practical optimization algorithms that allow for natural decomposition of large-scale constrained problems into smaller ones. The smaller problems are solved by progressively smaller dimensional, independent subproblems. The algorithms, which, in principle, can handle problems of unlimited size, are convergent for arbitrarily partitioned problems. In practice, the partitioning is expected to be controlled by the applications.

Another direction actively pursued at ICASE was the development of algorithms and their theory for general multilevel optimization with application to multiobjective optimization (ICASE Report # 94-77, Alexandrov and Dennis). Any design problem is multiobjective in nature. The proposed algorithms provide a way of regularizing the ill-posed problem of multiobjective optimization without any need to quantify the relative importance of the

41

objectives in any way other than specifying the order of importance. In addition, general multilevel procedures have been applied to MDO problems for a long time, and there is a need to provide theoretical foundations for these procedures.

The third direction of research at ICASE has been the study of foundations and applicability of various MDO formulations with the purpose of evaluating the formulations and providing guidelines for their use in specific applications.

Among the research areas that will have a large impact on MDO, two stand out as very important and amenable to the applied mathematics research methods: approximations concepts and optimization with both discrete and continuous components. Any optimization procedure applied to practical problems will likely be faced with extreme expense of function evaluations. Research into the question of approximating the objective and constraint functions and the effect of the approximations on the solution processes is vital. The expense of evaluating functions enters also into the area of nonlinear mixed integer programming (MIP). MIP addresses optimization problems that include both continuous and discrete variables. It is expected that the current approaches to MIP, such as genetic algorithms and simulated annealing, while successful for a number of problems, will not be satisfactory for many MDO problems because they require a large number of function evaluations. The very difficult area of optimization with discrete/continuous variables must be addressed.

The Roundtable has played an important role in bringing the unsolved MDO problems to the attention of the ICASE researchers. The MDO program at ICASE is expected to grow further following the ICASE/LaRC MDO Workshop in March 1995.

J. A. Bennett's lecture dealt with the need to reestablish the engineer as an integrator as opposed to a specialist in a single discipline. He stated that integrated design is the crucial R & D issue in engineering today.

One of the main difficulties in engineering R & D is the weak connection between industry and research. This disparity needs to be recognized and managed.

The work needs to be carried out in several areas, such as modeling and data representation strategies, automated discretization methods, directional information generation and decision methods.

The following engineering formalisms need to be developed to form a systematic infrastructure:

- Design representations:

  Design is to be represented by feature-based design context, in which all the disciplines can extract models.

- Discipline specific model generation:

For all disciplines, models must be generated rapidly, in less than 5% of the analysis time. Automatic discretization and the assembly of parts must be provided.

- Discipline specific analyses:

  All discipline analyses must be represented at various degrees of sophistication.

- Discipline directional information:

  The following information must be available: efficient and implemented mathematical sensitivity information; sensitivity information obtained by heuristic methods; physics and data based approximations. There must be measures of variability and accuracy.

- Design strategies:

  The strategies must allow for solving very large-scale problems with continuous and discrete variables. They must account for slow and expensive function evaluations and provide inexpensive intermediate models. The practitioner needs information about variability and access to models of various levels of accuracy and confidence.

An infrastructure that addresses these requirements must be able to handle large volumes of data, mathematical, heuristic, and test; provide a heterogeneous and flexible computing environment, be suited for efficient change and maintenance.

In summary, the area of multidisciplinary design is extremely complex and does not lend itself to complex isolated efforts. In integrated design, one cannot take advantage of great strides in one area, while another remains at the initial stages of development. Thus, work must be carried out uniformly.

I. B. Tjoa discussed modeling and optimization strategies of on-line Closed Loop Real Time Optimization (CLRTO), which is a procedure successfully applied to the design of large-scale petrochemical processes.

The design of such processes is a multidisciplinary task subject to rigorous demands on the accuracy of the model. Complexity of reaction mechanisms and a large number of frequently unknown chemical components make the modeling of a chemical reactor an extremely difficult task and produce a highly nonlinear model. However, an accurate model is crucial to the success of the optimization procedure, and much effort goes into producing as rigorous a model as possible.

CLRTO is an example of a highly successful optimization procedure applied to a problem with expensive function evaluations.

Dynamic Matrix Control Corporation has used the sequential quadratic programming algorithms (SQP) to develop a proprietary version of the SQP method for solving sparse optimization problems.

43

Tjoa indicated that effective sparse linear algebra tools would provide a most useful contribution toward using optimization techniques in larger-scale refinery-wide petrochemical design problems.

Following the informal Roundtable discussion, J. Sobieski opened the second MDO session by providing a perspective of the development of MDO in the past several years.

J. P. Giesing described current MDO activities at McDonnell Douglas Transport Division, including the Aeroelastic Design Optimization Program (ADOP), Advanced Integrated Loads System (AILS), and a local panel optimization code (ADVISOR). Applications of these systems were demonstrated for the HSCT.

The description of MDO activities was followed by the statement of current MDO challenges to provide ICASE and NASA LaRC with research and development topics of use to industry. The areas of interest included three major directions:

- Practical procedures for large problems.

  Developing practical procedures for large problems should include the development of procedures robust with respect to size, i.e., robustness should not deteriorate as the size of the problem increases. (Flutter optimization served as an example of the effect of size on robustness.) Large cycle times that drain time and resources must be reduced. The procedures should handle realistically large number of design variables, and they must be verified on large practical problems, such as HSCT.

- Efficient MDO processes.

  The emphasis here was on multidisciplinary structural design optimization. The following processes must be developed:

  - Interface between local panel optimization and vehicle-level structural optimization.

  - Practical global MDO schemes for systems with large suboptimization elements.

  - Efficient and robust processes for optimization with discrete and continuous variables and expensive or noisy sensitivities and function evaluations.

  - Include global value, cost, and profit objectives into the MDO technology.

44

- Disciplinary and multidisciplinary tools.

    The following specific tools for structural MDO must be addressed:

    - Structural loads and flutter to include CFD and wind tunnel data.

    - Develop rapid FEM generators and sensitivities.

    - Develop cost and constraint models for maintenance and manufacturing.

    - Develop revenue models.

    - Durability and damage tolerance to be included at the panel optimization level.

    - Use built-up weights instead of conceptual weights.

    - Extend cruise point optimization to the entire mission.

    - Develop accurate mappings from CFD to FEM and from FEM to CFD.

The following specific suggestions were made for ICASE/LaRC researchers:

- Install a large structural MDO program at ICASE/LaRC, such as ADOP or ASTROS.

- Interface a local panel suboptimizer, such as VICONOPT, with the large program.

- Implement several optimization methods, for example "Controlled Growth Optimization", and determine the one most efficient for aircraft applications. (AIAA paper 81-0549, 1981, Hajela, Sobieski)

- Extend VICONOPT to D&DT optimization.

- Perform an MDO profit or ROI (return-of-investment) optimization on a large HSCT model, taking into account FEM, CFD, and other disciplines.

- Develop and assess a parametric or rapid FEM generator.

J. A. Volk of Grumman Northrop concluded the MDO sessions with the description of MDO activities at Northrop Grumman, the company has a long history of development and implementation of multidisciplinary analysis and optimization, with an emphasis on structural optimization. Northrop Grumman was the prime contractor on the development of such MDO related tools as FASTOP and ASTROS.

The MDO experience of Northrop Grumman has led to the formulation of requirements for future air vehicle design environments. The next generation of concurrent engineering systems would benefit from the following characteristics:

- The system is flexible and heterogeneous and is capable of handling easily interchangeable analysis modules.

- The system supports distributed computing via a extensive range of hardware platforms.

- The system has advanced visualization and virtual reality capabilities.

Based on the experience of implementing structural optimization tools, Northrop Grumman holds the view that the transition to use of MDO technology on a regular basis faces bigger challenges than technical development, complex as the development is in itself. The following steps were proposed to facilitate the transition.

- The industry user base should be maintained and expanded, thus providing strong support of the existing MDO systems.

- University MDO curricula should be developed with the government and industrial support. University education should expose future engineers to MDO methodologies before they enter industrial jobs.

- Multidisciplinary analysis (MDA) should be developed and implemented before MDO to "provide earlier deliverables and minimize culture shock". (From the abstract.)

SUMMARY OF THE ROUNDTABLE DISCUSSION

The informal Roundtable discussion lasted for slightly over an hour. The discussion was very productive, but short. The experience indicates that at future Roundtables, it would probably be very useful to allocate longer times for the informal discussions, scheduled after the formal presentations take place, and at the times not conflicting with other presentations.

The following is likely an incomplete list of questions raised at the informal discussion and some of the comments and answers provided by the participants. The comments are not verbatim.

- There is a need for a study of the potential payoffs in MDO. Now there is no information on high-payoff areas. It is necessary to determine this, because the investments into implementing MDO will be large. Some Pareto-optimization work will be needed to look at the management issues.

  We must to identify the critical areas in need of accuracy. That is, we must discover what accuracies are needed for each discipline to obtain an MDO payoff.

- There must be some safety buffer if a discipline is insufficiently accurate.

- What is the role of ICASE?

  A brief introduction to the ICASE activities given by several participants. MDO at ICASE is its formative stages, but ICASE is committed to developing an active MDO program.

- Simultaneous Analysis and Design is a very promising area in MDO. There, design and optimization are inseparable, and the analysis PDE's serve as constraints.

- An immediate need is to produce models for manufacturing and operations. These models would make the MDO product applicable to the entire life-cycle.

- In general, the use of optimization is difficult. We need to discover whether optimization is an appropriate tool for engineering.

  What is the difficulty?

  The weakness of optimization is, in a way, its strength: if optimization is used, the user needs to state his requirements rigorously. This is difficult. The modeling is weak. It may not be possible for engineers to state their requirements as precisely as optimization techniques demand.

- There are techniques, called "exploratory optimization" that allow for preliminary optimization studies in the absence of rigorous information. Once the user knows more about the design space via exploratory optimization, methods requiring more complete information can be applied.

- We also need optimization methods that allow us to modify the optimization strategy as we proceed. Flexibility is most important.

- Why is there a resistance to optimization in industry? Sometimes there is a perception that optimization is mysterious once the human factor is taken out of it. To make the optimization processes better understood, we need visualization tools.

  The perception that optimization takes away the human factor does not hold in all companies. However, visualization is very necessary. If it could be done in more than 2-D, it would be great, but even 2-D snapshots would be very helpful. Visualization would take a lot of problems out of optimization.

  There is a special need to provide visualization tools for MDO. MDO will not lend itself easily to visualization, but it is a necessity, especially because it is difficult to visualize the effects of many disciplines.

LIST OF RESEARCH TOPICS

The participants gave much thought to the request to describe problems of interest to industry for research and development at ICASE and NASA.

Several areas were mentioned repeatedly. The major research topics follow.

- Optimization:

  1. MDO formulations and strategies;

  2. Robust, efficient algorithms for optimization with discrete and continuous variables of the problems with a large number of variables or expensive or noisy function and derivative evaluations;

  3. Large-scale optimization algorithms, robust with respect to problem size.

- Modeling and approximation concepts: development of strategies for handling expensive function and derivative evaluations by introducing cheaper models and studying their effect on the solution process.

- MDO environments and tools: development of flexible, modular, heterogeneous computational systems, easy to maintain and modify.

- Development of discipline specific rapid model generation capabilities and of analyses of various fidelity.

The topics are listed in detail in the summary of the session. Not surprisingly, given the complexity of the MDO problem, the research topics are very extensive.

LIST OF WORKSHOP TOPICS

In addition to the scheduled ICASE/LaRC MDO Workshop scheduled for March 1995, the participants expressed interest in further Industry Roundtables, as most attendees found them very productive.

The organizers solicited feedback on month-long summer workshops with industrial, academic, ICASE, and NASA participants with a focus on specific industrial problems. ICASE has been conducting such workshop with academic participants. While the members of industry expressed interests in such longer, highly technical workshops because of the expectations of high productiveness, they also expressed reservations about their ability to leave work for such a long period of time. Another difficulty is that specific problems would probably involve working with proprietary information. Despite these serious reservations,

the participants who were asked this question stated that, given a sufficiently long-term and careful preparation, such longer and more technical workshops would be potentially very useful and productive.

SESSION B5: MULTIDISCIPLINARY OPTIMIZATION IN INDUSTRY II

Natalia Alexandrov, NASA Langley
Jaroslaw Sobieski, NASA Langley

Please refer to SESSION B4.

SESSION B6: CLUSTER COMPUTING II

David Keyes, ICASE

Please refer to SESSION B2.

SESSION C1: ALGORITHMIC AND ARCHITECTURAL ASPECTS OF LARGE-SCALE SCIENTIFIC COMPUTING

Chester Grosch, ICASE

This report is unavailable.

SESSION C2: SYSTEMS ENGINEERING

David Nicol, ICASE

The Systems Engineering session of the Roundtable involved participants who, for the most part, are already involved in systems engineering research jointly with ICASE. In this session we identified by example different models of joint ICASE/Industry interaction.

Andres Haynes and Fred Wieland of MITRE presented their work supporting Federal Aviation Authority studies of the U.S. aviation system. A key part of this support is the development of analysis tools, particularly simulation-based tools. ICASE Staff scientist Linda Wilson is working with MITRE to port the parallel simulation system MITRE uses (SPEEDES) to the Intel Paragon, and to evaluate MITRE's code (DPAT) on that platform.

Philip Heidelberger of IBM Research presented the results of the sabbatical he spent (Sept. 93 - June 94) at ICASE. Heidelberger, ICASE Staff Scientist Phillip Dickens, and David Nicol (also on sabbatical at ICASE during that period) developed a tool called LAPSE that provides parallelized direct execution simulation of computer codes developed for the Intel Paragon. LAPSE demonstrates the potential of using simulation to study scalability properties of existing codes, as well as providing a basis for intrusion free instrumentation of parallel codes, or for using application codes as on-line traffic generators in fast studies of operating system policies or network designs. Heidelberger is now back at IBM investigating the transfer of the ideas generated during his sabbatical to IBM products. ICASE hopes to hire a staff scientist to be involved with this effort.

Andy Ogielski of Bellcore talked of Bellcore's work in monitoring network traffic, and in analyzing the patterns found within. He proposed the development of a publically accessible database of network measurements that could form the basis of traffic models for simulation studies of network systems.

Albert Greenberg of AT&T presented techniques for the massively parallel simulation of wireless communication networks. This work was conducted jointly with David Nicol, while Nicol was either a consultant at ICASE or on sabbatical there. ICASE's involvement here is in supporting Nicol's time on this project, and publication of technical reports resulting from the collaboration.

Cooperative research topics shall continue to be focused on the use of discrete-event simulation as a means of studying computer and communication systems. Future workshop topics would include the results of these on-going collaborations.

SESSION C3: PARALLEL PROGRAMMING ENVIRONMENTS

Piyush Mehrotra, ICASE

INTRODUCTION

Participants

- **Ian Angus**, *Research Scientist, Boeing Computer Services.*

- **Mark Furtney**, *Group Leader, Cray Research.*

- **John Levesque**, *President, Applied Parallel Research.*

- **Vijay Naik**, *Research Scientist, IBM T. J. Watson Research Center*

Porting scientific codes to parallel architectures still faces some acute challenges. The focus of this session (and the accompanying roundtable discussion) was the research required to solve some of the significant problems that arise when scientists and engineers attempt to parallelize their codes. In the session, each speaker provided an overview of work going on in their organization and then discussed future research requirements from the perspective of their organization. High Performance Fortran (HPF), a set of extensions for Fortran 90, has recently been designed by a group of researchers from academia, industry and research labs. HPF has received a lot of attention since, it provides a first portable method for implementing data parallel codes on parallel architectures. Thus a lot of the discussion was centered around issues arising out of HPF.

PRESENTATIONS

Vijay Naik, a research scientist at IBM, gave a synopsis of IBM's efforts in the area parallel programming software. IBM has supported the HPF effort from its inception and currently provides Applied Parallel Research's HPF compiler with SP series machines. They have a prototype compiler effort in their Research division but do not have plans to release it immediately. They would like to see research focussed on issues such as communication optimization and HPF related tools and libraries. Another area that they are interested in pursuing, is that of threads in a distributed memory environment.

Ian Angus, research scientist in the Computer Services division of Boeing, has observed the HPF development from outside and presented his own opinions, rather than the official Boeing position. Boeing has concentrated on distributed computing, since parallel machines

have not been able to satisfy the needs of scientists and engineers in a production environment. According to Angus, the current definition of HPF is inadequate since it supports only about 10% of their applications. In his view, the focus should be on management of data and programming in the large issues and in particular extensive testing of the ideas needs to be done before any standardization is attempted.

John Levesque discussed the need for tools required to transport legacy F77 codes to parallel machines. Since APR is a small software house, it views universities and research labs as its research divisions. Levesque would like the research community *not* to duplicate what commercial houses are doing, but rather forge ahead so that companies like his can leverage off of research done elsewhere to build commercial products. To that end, he is willing to collaborate with any research organization by providing the source code of his compiler, with appropriate safeguards.

Mark Furtney is the Group leader of the Tools, Libraries, Commands group and the MPP Software group at Cray Research. He provided an overview of Cray's hardware and software trends. At this point Cray is not providing an HPF compiler for the T3D (they do have APR's HPF compiler running on the T3D), but are pursuing their own version called Craft. Craft is very similar to HPF, differing in details rather than philosophy, and also extends HPF to take advantage of the T3D's architectural model. Furtney said that his remarks on HPF were based on their experience with Craft and apply equally both HPF and Craft. From his perspective, a parallel programming language should provide a balance between performance, portability, and ease of use. HPF conceptually provides a degree of portability not found in current environments. However, its usability is still not clear. Also, current HPF compilers are woefully inadequate from the point of performance of the generated code. He would like to see research done on an integrated tool set, which allows information to flow back and forth between the different pieces to exploit maximally the underlying architecture.

LIST OF RESEARCH TOPICS

Since a couple of the speakers were leaving right after the session, the roundtable discussion was held before the session and hence could not be heavily advertised. The discussion at the roundtable concentrated on areas which need to be further studied and suggestions for possible workshops.

One of the top issues raised by a couple of the speakers was support for porting legacy Fortran 77 codes to parallel environments. There is a need for tools to help the "naive" (from the point of view of parallelism) user understand large legacy codes and aid the user in taking decisions about the distribution of data and computation.

A related issue is debuggers, performance monitoring, and prediction tools. Debuggers need to maintain a link to the original source code even after translators have completely transformed the code. There is a critical need for "what if" tools which allow the user to quickly test out choices without going through a complete compile/execute cycle. The issue with performance monitoring tools is the presentation of information, since an information overload, tends only to confuse the user. The attendees saw a need for a standard runtime interface, so that all the tools could work with each other and pass information to each other. For example, the compiler may produce the instrumented code which dumps the execution trace. As the trace is analyzed, information about the particular execution could be fed back to the compiler so that it can further optimize the code.

Another issue that was discussed was the support for object oriented methodology for parallel scientific codes. The use of C++ for data structure manipulation is on the rise, though the code to do the actual computation code is still usually in Fortran. There needs to be support systems for clean interfaces between languages, such as HPF and C++, in a parallel environment.

Threads are the focus of current ICASE research in runtime support systems and attendees saw a need for studying the use of threads not only for coarse grained task parallelism but also for fine grained functional parallelism and communication-computation overlap. Threads integrated with objects were seen as facilitating the balancing of load in both massively parallel and heterogeneous environments.

Other issues discussed at the roundtable in passing included parallel I/O, distributed databases and management of large sets of data, heterogeneity and distributed computing, and real time and embedded environments.

LIST OF WORKSHOP TOPICS

Several topics were suggested for workshops and symposiums: implementation of threads and their use both as compiler targets and as targets for direct user encoding; performance and debugging tools and their integration into a cohesive integrated environment; and object-oriented methodology in the service of parallelism.

Due to common interests, there arose several possibilities for collaboration. In particular, a staff scientist at IBM has been looking at the issues of integrating MPI and pthreads, an issue in which ICASE has already gained some experience. The implementation of such a system on the IBM SP2 could be done jointly. As noted above, APR is willing to provide their HPF compiler including source code to research groups. They are particularly interested in jointly investigating possible HPF compiler optimizations when targetting a thread-based runtime environment.

SESSION C4: SAFE SYSTEMS

Wayne Bryant, NASA Langley

This report is unavailable.

SESSION C5: RELIABLE TOOLS & ANALYSIS

Kishor S. Trivedi, ICASE

SUMMARY

The main theme of the session 'Reliable Systems: Tools and Analysis,' was that a number of tools for reliability, performance and performability analysis are available and yet their use in practice is limited at best. Speakers provided ideas as to how to increase the use of tools in practice. The need for dissemination of the capabilities of tools was pointed out. Transferring technology from research tools to commercial tools was deemed to be important. Integration of different tools, techniques and automated translation from design database was recommended.

INTRODUCTION

In the session on 'Reliable Systems: Tools and Analysis,' three speakers were scheduled: Kevin Prodromides of Allied Signal Aerospace, Andrew Rindos of IBM Networking Systems and Tilak Sharma of Boeing Commercial Airplane Group. The first speaker could not come and so the session chair (Kishor Trivedi of Duke University) said a few words instead.

PRESENTATIONS

Dr. Tilak Sharma narrated his experience in selecting and acquiring tools for reliability analysis suitable for Boeing engineers. He also pointed out the importance of training the engineers in the use of the tools. If multiple tools/techniques are to be adopted, the need for a clear criteria for selecting the right techniques and tools for a given application is paramount. It is then clear that an automated selection of the tool/technique will be ideal. The investment made in training and usage of earlier tools needs to be protected when new tools are acquired. The interface of the tools need to be close to the designer. The Boeing approach to solve some of these problems is to create the Boeing Integrated Reliability Environment. In this environment, the user is provided a high-level modeling language and the selection of the technique/tools is automated. A large number of diverse tools (SETS, CAFTA, HARP, EHARP, SHARPE, and SPNP) are integrated in the environment. A prototype version of the environment is being tested.

Dr. Andy Rindos related his experience with performance analysis of computer communication systems. He felt that real problems are extremely complex and the standard approach of using discrete-event simulation is unacceptable. At the same time, the use of

product-form networks is inappropriate as real networks are far from product-form. He suggested the uses of fast yet accurate approximations as an effective method to solve realistic problems. Ward Whitt's technique is one of the chosen ones. The automated generation (via stochastic Petri nets) to compute both steady-state and transient performance measures is another approach. A tool is being developed within IBM Networking Systems based on these ideas.

Dr. Trivedi observed that the use of state-of-the-art tools for performance and reliability prediction is minimal in practice. It is necessary to provide information about the capabilities and limitations of the tools via frequent seminar and short courses. He pointed out that much duplication of effort should be avoided by increased sharing of tools and interfaces. At the same time usability of the tools should be increased by means of several types of integration:

1. Integration of design database with analysis tools.

2. Integration of various system design factors such as reliability, performance and performability in the same tool.

3. Integration of experimental evaluation with model-based evaluation.

4. Integration of qualitative evaluation with quantitative evaluation.

5. Integration of discrete-event simulation with analytic-numeric methods.

LIST OF RESEARCH TOPICS

- Automated mapping from the design database into a reliability, performance, performability prediction tool.

- Integration of hardware and software reliability and performance.

- Collection of high-speed network traffic measurements and the analysis of the traffic pattern so as to determine the nature of traffic processes.

LIST OF WORKSHOP TOPICS

- Tools for Reliability and Performability

- Solving Large Reliability and Performability Problems

- Dealing with Models with Non-Exponential Distributions

59

SESSION C6: FORMAL METHODS AND SOFTWARE ENGINEERING

Ricky W. Butler, NASA Langley

SUMMARY

The goal of this session was to identify major problems industry is facing in the development of software and the potential for formal methods to eliminate or ameliorate these problems. The first speaker was John Rushby (SRI International) who convincingly argued that formal methods offers a solution to the growing problem of design errors in software and digital hardware. He explained that formal methods can play a similar role in computer design as Computational Fluid Dynamics (CFD) plays in aeronautical design, by providing a means of calculating, and hence predicting what the behavior of a digital system will be prior to its implementation. Steve Miller (Rockwell Collins) described the successful application of formal methods in pilot studies done at Rockwell Collins under NASA LaRC sponsorship. He argued that formal methods has a significant potential in several important areas of Collins' commercial work. In particular, he believes that formal methods capabilities should be injected into their requirements specification and analysis tools, their testing methodology, their logical partitioning (i.e. integrated modular avionics) work and their hardware design. Jeff Voas (Reliable Software Technologies Corporation) presented a novel approach to software testing that has been funding by NASA LaRC via the NASA SBIR program. He argued that there was a significant opportunity to advance software engineering by integrating the NASA efforts in formal methods and software testing.

INTRODUCTION

Participants

- **Ricky W. Butler**, *Mail Stop 130, NASA Langley Research Center, Hampton, Virginia 23681-0001.*

- **John Rushby**, *Computer Science Laboratory, SRI International, Menlo Park, CA 94025.*

- **Steven P. Miller**, *Collins Commercial Avionics, Rockwell International, Cedar Rapids, IA 52498.*

- **J. M. Voas**, *Reliable Software Technologies Corporation, Reston, VA.*

PRESENTATIONS

Computer systems that interpret pilot commands, sample sensors, evaluate control laws, and command the actuators are used on all modern warplanes, and increasingly on commercial airplanes (Airbus A320, Boeing 777). Aircraft certification requires assurance that catastrophic failures will be "extremely improbable" (FAA AC 25.1309A) that is, not expected to occur in lifetime of the fleet.

The Problem

Software for aircraft is developed to extremely rigorous standards, and subjected to massive testing. There is evidence that coding bugs in sequential components are not as frequent or serious as the design problems that are introduced earlier in the life-cycle. Concern centers on the following:

- Requirements (e.g., JPL data for Voyager and Galileo spacecraft: only 3 of 197 mission-critical defects were programming problems; IBM data for Space Shuttle: 400 "user notes" documenting requirements anomalies: 6 potentially life-threatening implementation defects known to have flown)

- The intrinsically hard problems (e.g., coordination of distributed computations, timing, synchronization, fault tolerance)

Some examples where software design errors have had serious effect:

- AFTI F-16 Flight Test (redundancy management)

- X29A Flight Test (redundancy management)

- HiMAT Flight Test (timing anomaly)

- X31 and C17 channel dropoffs in flight tests

- Saturation of AT&T network on 15 January 1990

- Failure to launch STS-1 (synchronization of backup computer)

- Loss of data from Voyager at Jupiter

- Patriot failure at Dharan (clock drift)

61

Eliminating faults from redundancy management software requires achieving complete understanding of the behavior of distributed, concurrent, real time systems, operating in the presence of faults.

Formal Methods is Mathematical Modeling

One way to predict behavior of a system is to construct a mathematical model and calculate it. For this to be effective, the model must be reasonably accurate and the calculations must be performed without error. For many continuous systems of traditional engineering, well-developed mathematical theories (e.g., Navier-Stokes equations for aerodynamics) are available. For computer systems, we must use discrete mathematics (logic, set theory) and build our own theories. Proofs of theorems take the place of numerical calculation. This is the essence of formal methods. Thus formal methods can be thought of as the applied mathematics of software engineering. Formal methods use the techniques from logic and discrete mathematics to model the requirements, specification, design, and implementation of computer systems, in a way that supports analysis of certain properties (e.g., consistency, completeness), and enables prediction of (modeled) behavior through systematic processes that resemble calculation.

SRI International Perspective

John Rushby argued that formal methods offers the following advantages:

- Enable faults (of assumptions, requirements, design) to be detected earlier than otherwise due to greater precision and explicitness early in the lifecycle.

- Enable faults to be detected with greater certainty than otherwise because they replace reviews (consensus) by analyses (calculations).

- Can provide total coverage (of selected, modeled properties).

- Guarantee absence of specified faults (subject to accuracy of modeling employed) because the calculations (proofs) can be checked mechanically (by a theorem prover).

Tools for formal methods should provide a collection of *graduated* techniques to help in the detection and elimination of faults

- Syntax: parser

- Formal consistency: typechecker

- Sanity check: direct execution ("animation")

- Reviews: prettyprinter, cross-reference generator

- Deeper sanity check: formal challenges

- Dark corners: state exploration/model checking

- Hazard Analysis/Safety: backwards execution

- Correctness: formal verification

Systems designed for the more difficult analyses can be extended to support the simpler ones; vice-versa is almost impossible.

Dr. Rushby believes that the benefits of Formal Verification are much more than "proof of correctness":

- Debugging (i.e., discovery of incorrectness)

- Complete enumeration of assumptions

- Sharpened statements of assumptions and lemmas

- Streamlined arguments

- Enhanced understanding that can lead to further improvements

- Support for reliable and safe exploration of alternative assumptions and designs, and adaptation to changed requirements.

- Like CFD, formal methods can allow the design space to be explored more completely and more cheaply than prototyping.

Rockwell Collins Perspective

To improve performance and reliability, the commercial aerospace industry has moved many critical functions previously performed manually or by mechanical systems into digital avionics. Examples include autopilot and autoland systems, the Traffic Alert and Collision Avoidance System (TCAS), fly-by-wire primary flight control systems, Integrated Modular Avionics (IMA), satellite navigation and communication, and increasingly sophisticated displays. Dr. Steve Miller believes that this trend will only accelerate as the need for improved efficiency, performance, and safety is met by making aircraft more intelligent. An order of magnitude increase in the size and complexity of airborne digital systems over the next decade is not infeasible.

Dr. Miller argued that this demand can only be met in a manner that maintains existing levels of safety and reliability through the development of improved methods for the design

and verification of software and hardware. He believes that formal methods can play an important role at all levels, from requirements capture and analysis to test case generation to maintenance. However, just as existing methods will need to change to incorporate greater formality, formal methods will need to change to meet the specific needs and realities of the avionics industry.

Dr Miller stated that a better method for requirements analysis is the single greatest need at Rockwell Collins. In particular, Rockwell Collins needs a methodology for requirements that:

- is accessible by domain experts and customers,

- can accommodate frequent change,

- is precise enough to support formal analysis and development of tools,

- supports generation of test cases,

- identifies criticality of each function,

- avoids overspecification,

- and provides mechanisms to determine when requirements are inconsistent or incomplete.

Software testing is the single most costly activity of the Advanced Technology & Engineering Department at Rockwell Collins. Structural unit testing can consume over 50% of the total project budget.

Dr. Miller described the excellent progress that is being made towards completing the formal verification of the microcode of the Rockwell Collins AAMP5 Microprocessor. The formal specification of the instruction set (2,550 lines of PVS) and the micro-architecture (2,679 lines of PVS) has been completed. Eleven instructions ranging over three instruction classes have been formally verified. Two design errors were discovered during the specification and one design inconsistency was discovered during the specification. The formal verification systematically uncovered two "seeded" errors in the microcode delivered to SRI. Rockwell Collins covertly planted the bug to test the efficacy of the SRI formal analysis method. They had been anxiously waiting to see if the verification team would discover the errors they had carefully planted and were very pleased when they were discovered.

The AAMP5 microprocessor is the newest member of the CAPS/AAMP family of microprocessors that have been widely used by the commercial and military aerospace industries.

This family has been used for (1) Boeing 747-400 Integrated Display System (IDS) (2) Boeing 737-300 Electronic Flight Instrumentation System (EFIS) (3) Boeing 777 Flight Control Backdrive, (4) Boeing 757, 767 Autopilot Flight Director System (AFDS) (5) military and commercial Global Positioning (GPS) Systems.

Reliable Systems Technology Perspective

Software testability is a software metric with many definitions and perspectives. For the past 6 years, Jeff Voas has been working with one such definition that he developed during his graduate years and post-doctoral fellowship at NASA. This definition has allowed for a more precise understanding of why software faults hide during testing, and what if anything can be done during design and coding to prevent programs from hiding errors during testing.

Testing and proofs are defined with respect to some "authority", usually a specification written in English, that decides what is correct. Testability says nothing about correctness nor does it attempt to quantify the unreliability of a system. Instead testability measures the propensity of a program to hide its flaws from a given testing strategy. Thus, high testability is desirable during testing and low testability during operation. The Voas testability methodology can also identify the particular place in a program where the fault masking is occuring. Thus, Dr. Voas argued that there is a natural synergism between formal methods and testability. The techniques of formal methods can be concentrated on the parts of the program where testability indicates faults could be lurking. In particular, Dr. Voas suggested that assertions appear to be a good starting point from which we can begin to create new achievement/assessment methodologies that are both formal and empirical.

LIST OF RESEARCH TOPICS

1. Formalization of useful requirements specification languages using a mechanical theorem prover such as PVS.

2. Application of formal methods to train signaling systems.

3. Research project to add formal methods techniques such as verification condition generation and assertions to the Voas testability method.

LIST OF WORKSHOP TOPICS

1. Workshop dedicated to Software Engineering and Formal Methods.

ICASE

The Institute for Computer Applications in Science and Engineering (ICASE) is operated at the NASA Langley Research Center (LaRC) by the Universities Space Research Association (USRA) under a contract with the Center. USRA is a non-profit consortium of major colleges and universities.

The Institute conducts an unclassified basic research program in three major areas: i) applied and numerical mathematics, including numerical analysis and algorithm development; ii) theoretical, computational, and experimental research in fluid mechanics in areas of interest to LaRC, including acoustics and combustion; iii) applied computer science and parallel computing.

The Institute helps promote cooperative research activities among scientists and their institutions, and in addition, enhances communication among researchers in related disciplines by providing an academic setting within which academic and industrial scientists can collaborate with NASA scientists and engineerings on problems of mutual interest.

The Institute serves as NASA's gateway to a large pool of prominent scientists from major universities. During the last year about 160 researchers from over 100 universities visited ICASE. Individual stays range from a few days to a full year, typical activities include formal and informal seminars, workshops, and interaction with NASA scientists. In addition, ICASE has approximately 40 longer-term members on its scientific staff. The synergy between scientists from ICASE, NASA, and academia advantageously positions ICASE for dealing with multidisciplinary problems. Opportunity now exists for interested industries to utilize the synergy of this group to conduct fundamental research of economic importance.

ICASE's vision of industrial cooperation is that ICASE and Langley Research Center be used as a focal point where ICASE affiliates make industrial representatives aware of on-going research interests and talents, while industrial representatives share fundamental pre-competitive problems. The goal is for industrial input to spark interest and help guide ICASE affiliates in the selection and formulation of problems they choose to attack. ICASE welcomes extended visits (e.g., sabbaticals) by industrial representatives.

ICASE benefits from the experimental and computational facilities and resources of Langley Research Center. Particularly important are extensive scientific computing facilities, which include machines of advanced and novel architectures, e.g., a Cray-2, an Intel Paragon, and (shortly) an IBM SP-2. To complement these facilities, ICASE operates its own network of SUN and Silicon Graphics workstations.

REPORT DOCUMENTATION PAGE

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE February 1995	3. REPORT TYPE AND DATES COVERED Contractor Report

4. TITLE AND SUBTITLE
THE FIRST ICASE/LARC INDUSTRY ROUNDTABLE: SESSION PROCEEDINGS

5. FUNDING NUMBERS
C NAS1-19480
WU 505-90-52-01

6. AUTHOR(S)
Sharath Girimaji

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Institute for Computer Applications in Science
and Engineering
Mail Stop 132C, NASA Langley Research Center
Hampton, VA 23681-0001

8. PERFORMING ORGANIZATION REPORT NUMBER
ICASE Interim Report No. 26

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
National Aeronautics and Space Administration
Langley Research Center
Hampton, VA 23681-0001

10. SPONSORING/MONITORING AGENCY REPORT NUMBER
NASA CR-195040
ICASE Interim Report No. 26

11. SUPPLEMENTARY NOTES
Langley Technical Monitor: Dennis M. Bushnell
Final Report

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Unclassified–Unlimited

Subject Category 59

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

The first "ICASE/LaRC Industry Roundtable" was held on October 3 - 4, 1994, in Williamsburg, Virginia. The main purpose of the Roundtable was to draw attention of ICASE/LaRC scientists to industrial research agendas. The Roundtable was attended by about 200 scientists, of which 30% from NASA Langley; 20% from universities; 17% NASA Langley Contractors (including ICASE personnel); and, the remaining from federal agencies other than NASA Langley. The technical areas covered reflected the major research programs in ICASE and closely associated NASA branches. About 80% of the speakers were from industry. This report is a compilation of the session summaries prepared by the session chairmen.

14. SUBJECT TERMS
Industry Roundtable

15. NUMBER OF PAGES
68

16. PRICE CODE
A04

17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT

NSN 7540-01-280-5500