Notes on the Technical Sessions of the Second HPCD Workshop July 20-21, 1995 Rutgers University New Brunswick, NJ

John Keane
Mark Schwabacher

October 24, 1995

1 Welcoming Remarks – Joseph Seneca, University VP for Academic Affairs

1.1 Summary of Presentation

Dr. Seneca addressed two important dimensions of the HPCD work: substance *and* process. *He noted that the project is carrying out work that is both scientifically significant, and that has a large potential for economic impact. He observed that the goal of the project – to increase productivity – is consistent with national economic goals of a higher standard of living and reduced scarcity.*

In addition, Dr. Seneca remarked on the multi-disciplinary, multi-institutional and collaborative nature of the project. In this regard we are challenged to give up more "traditional" views of project success as accruing to a single department or institution. If the nature of research is, in fact, becoming more multi-disciplinary and collaborative, institutions must adjust their thinking to reflect this.

2 Introduction to the HPCD Project – Saul Amarel

2.1 Summary of Presentation

Dr. Amarel welcomed everyone to the second HPCD workshop, being held at about the midpoint of the second of four years of the project. He laid out the goals of the workshop: to review the project goals; to bring together the people related to the project; and to get a good "big-picture" overview of project progress.

The project intends to build on work in high-performance computing, artificial intelligence, computer-aided design, and modeling technology to develop new generations of engineering tools that can significantly boost industrial productivity. To do this the focus has been on complex engineering design problems that are not decomposable, have conflicting goals, and are computationally intensive, requiring multidisciplinary modeling. These problems include: the design of surface ships; the design of jet engines and nozzles; the conceptual design of aerospace vehicles; photolithography; and multi-level microprocessor design. However, the main thrust of the work is the development of methodologies and systems that will be applicable beyond the scope of these domains.

We wish to explore the *intelligent* use of high-performance computing, with a focus on a speedup of the design process from concept to realizable design. We would also like to improve the level of design automation, and the ability of designers to react to changes in technology or mission. In doing this, we are developing prototype systems and testbeds for the exploration of ideas and to provide a basis for generalization. Technology transfer will be ensured by building realistic design systems in areas of industrial significance in broad collaboration with industry.

Problems of design involve a mapping from a point in a *specification space* to a point (or region) in *design space*. The mapping from specifications to design space may be constrained by *a priori* knowledge, but search in the space of designs is normally required. Among the project goals is the development of a generic framework for the solution of design problems, which has been named **DA/MSA**. The components of this system concretely embody elements of the design process. A key technical idea of the DA/MSA is that though high-performance computing is required for these design problems, raw computing power is insufficient – modeling approximations and reductions are essential to make the design process tractable.

Dr. Amarel concluded by summarizing the personnel involved in the project, and the university and industry collaborators.

3 Ship Design I – Nils Salvesen, SAIC

3.1 Summary of Presentation

The objective of this design problem is the development of a methodology for using advanced physics in a design process that is currently dependent on physical scale-modeling.

The design problem that has been chosen is that of *bow flare* of container ships, for which only a single set of physics codes is needed. Both linear and non-linear codes are available, with advanced 3-D codes (the **LAMP** codes) under ongoing development. The design of a bow flare for maximum profitability requires analysis between the tradeoffs of maximizing the cargo capacity of the ship and the cost of building and operating the ship.

4 Ship Design II – Tom Ellman

4.1 Summary of Presentation

The DA/MSA has been instantiated for the problem of ship design by the incorporation of domain-specific models, including a simple *profit model*, and sophisticated linear and non-linear codes to model the bending moments encountered by ships under varying seas. These codes are computationally expensive, requiring an estimated 5 CPU-days (275 MHz DEC Alpha) per design optimization for the linear version (LAMP-2), and 90 CPU-days for the non-linear version (LAMP-4).

Two approaches are being taken in managing the complexity of the design process: parallelization of the LAMP code (to be presented later), and the use of multi-level modeling. Since only the most extreme 10% of bending moments encountered need be analyzed in detail, it would be a substantial improvement in complexity if less-expensive evaluation could be used for the remaining 90%. To

test this, a *hybrid controller* has been constructed that will run a less-expensive code first, analyze the results, run a more expensive code for a portion of the solution, and "splice" the results together.

The system has been tested for a combination of MIT (linear) codes and LAMP2, and for pre-computed LAMP2 and LAMP2. The MIT codes were found to be good predictors of peak location, but not of peak height. The pre-computed LAMP2 (run against the "base case" ship design) proved to be a very good predictor of critical bending moments, and resulted in a five-fold speedup in analysis with a 7.6% error rate. However, though the absolute error is low, it is sufficient to shift the optimum design significantly, an undesirable result. Further work needs to be done to explore tradeoffs between the amount of analysis done with the more expensive model and impact on the optimal design. Future work will also include extending the design problem to include the entire bow of the ship.

4.2 Questions and Discussion

The value of multi-level modeling was questioned. If the cost of producing an optimal design is as much as "90 CPU-days", isn't this still a relatively small cost?

Dr. Ellman responded that producing an optimal design is virtually never the result of a single optimization run. Multiple designs, reflecting multiple tradeoffs, goals, and modeling assumptions, will be produced. A faster system allows for much greater exploration of possible designs.

5 Parallel Computations in Ship Design; Tools and Algorithms for Parallel Processing – Apostolos Gerasoulis

5.1 Summary of Presentation

The current state of the art in automating parallelization is still far from being able to take unadorned input of a complex system such as the LAMP codes and produce an efficient parallelization of those codes. Recent work with the PYRROS system for parallel scheduling, and PLUSPYR, a tool for assisting in the analysis of potential parallelization in sequential programs, have enabled the useful parallelization and scheduling of a portion of LAMP for different architectures.

The PLUSPYR system takes as its input code segments plus user-provided annotations, and reports on the potential parallelism. The user can change the granularity of the parallelism in order to find the most useful level. The output of PLUSPYR can be fed to the PYRROS system, which will automatically

schedule it for any of a number of parallel architectures. LAMP-4 was chosen as a good candidate to test this system, because of the potential for parallelism in the code, and because in sequential form it runs about 20 times slower than LAMP-1.

Analysis indicated that parallelization of the time-dependent portion of LAMP-4 held the greatest potential benefit. Initial tests using the PVM software and a workstation farm gave a much smaller speedup than originally predicted – about 5 times. This problem is still under investigation, but single-processor slowup caused by the PVM software seems a likely culprit.

A parallelization for the Cray T3D, however, gave much better results. The efficiency of the parallelization was found to increase with the number of iterations of the simulation, approaching linear speedup in the number of processors for large iteration runs. Theoretically, it should be possible to reduce the current 6 days of processing per evaluation to about 1 hour with the next generation of Cray machine and the parallelized software.

6 Inlet Design – Doyle Knight, Andrew Gelsey et al

6.1 Summary of Presentation

This research has two objectives: to improve both the *efficiency* and *quality* of high speed inlet design through the integration of computational fluid dynamics and artificial intelligence, and to improve the *quality* of high speed inlet design through the development of more accurate numerical algorithms for the compressible Navier-Stokes equations.

In the work on final design of high-speed inlets, the NPARC Reynolds-averaged Navier-Stokes code was used in a systematic redesign of the NASA P2 hypersonic inlet. The first phase of the work involved computational experiments to determine appropriate grid densities, etc. for using NPARC to achieve grid-converged simulations of the P2 inlet which adequately matched published experimental data. The second phase of the work involved formulating the redesign of the P2 inlet as a numerical optimization problem which was attacked using state-of-the-art numerical optimization software. The resulting P2 inlet design is significantly superior to the original design. In particular, the static pressure distortion at the throat was reduced by more than a factor of five.

The redesign of the P8 inlet is currently in progress, with promising preliminary results. Collaboration has begun with United Technologies Research Center on the design of a supersonic inlet for a missile for the Navy, known as Project Fast Hawk (formerly Project Cheap Shot), which is intended to replace the Tomahawk. Effort will begin in September 1995 on the preliminary design of novel 3-D high speed inlets.

In inlet analysis, the development and validation of both 2-D and 3-D Navier-Stokes codes has been completed. The 3-D code is currently being applied to the P8 inlet. In the 2-D code, work is currently being done on preconditioning and combustion.

7 Design of Nozzles and Aerospace Vehicles – Andrew Gelsey

7.1 Summary of Presentation

The goal of this portion of the project is to develop innovative AI-based technologies and test them on realistic problems in aerospace design. Because the cost of evaluating designs is relatively low, this is a good testbed for ideas concerning AI and design. Current efforts are focusing on the development of the Modeling-Simulation Associate, techniques for AI-enhanced optimization, a search-space toolkit (SST), problem reformulation, and multi-level design.

Briefly, the goal of the aerospace vehicle design problem is to choose values for a set of key design parameters such that the cost of the vehicle is minimized for a particular mission. The take-off mass of the vehicle for provides a good estimator for both the acquisition and operating costs.

Work on this problem has motivated some of the work on the MSA. Earlier versions (used for the nozzle design problem) implemented a failure-driven multi-start strategy to handle modeling failures, but this strategy was not sufficient for the full vehicle problem. The MSA has been extended to greatly increase the communication between the DA and the model, by providing information about *modeling constraint violations*. When used with a constrained optimizer (such as CFSQP) in a 7-dimensional parameter space, this allows about one-third of designs that are initially model-infeasible to be refined to feasibility.

Another approach tested in this domain is that of multi-level design for partially-decomposable problems. This addresses the problem of finding the global optimum in design spaces with local optima where "brute-force" methods (such as genetic algorithms or simulated annealing) would be prohibitively expensive. Preliminary results in the aerospace vehicle domain indicate that design-quality improvement can be achieved at far lower cost than using genetic algorithms.

The **Search Space Toolkit** (SST) is a set of tools to allow systematic exploration of the design space and extraction of useful knowledge that can drive problem reformulation. Components provide the ability to do a number of different experiments, including a Monte-Carlo-like multi-start, line search between local optima, a variety of analyses of local regions of the search space, fixed-grid sampling, and various visualization tools.

The issue of problem reformulation is also being addressed in this domain. The objective is the transformation of the search space to one in which there is

a much higher proportion of physically meaningful designs, and which can be searched more efficiently. These transformations can be carried out through the addition of "wrappers", which allow for quick construction and modification. Results of this kind of transformation in the nozzle design domain have been shown to be effective in significantly reducing the cost of producing an optimal design, when compared to the original problem formulation.

8 Design Associate and Modeling-Simulation Associate – Tom Ellman

8.1 Summary of Presentation

There are many choices faced by the system builder in the process of developing a system for automated design. It is not the case that a naive composition of analysis models and numerical tools is likely to produce a usable design system. To address this problem, the **Design and Modeling Process Language** (DMPL) has been developed. This allows the entire design system to be systematically transformed. Tests on this in the 12-meter yacht domain have demonstrated the value of one type of transformation, constraint incorporation.

Another objective of the DA/MSA system is the development of a design record that captures the evolution of the design process as a series of computational experiments. This is a fundamentally new approach to the design record, and allows the possibility of using a Machine Learning approach to the automated exploration of design strategies. Work in this area has shown that it is possible to learn useful rules for constraint incorporation transforms from data collected in the design space.

Yet another difficulty encountered in the development of automated design systems is the brittleness of analysis models, particularly those comprised of "legacy" codes. When coupled with "off-the-shelf" generic numerical tools, failures in the analysis models will abort the design process. Direct modification of the models, however, is time-consuming and may necessitate their re-validation. The DA/MSA provides a means for problem resolution to be added to existing systems by means of code wrappers, leaving the original system code intact. Problem resolution strategies, called *intercessor schema* are generically defined in a declarative language, and can be instantiated for specific problems within design systems.

9 Design of "Voice Mimic" I, Overview – Jim Flanagan

9.1 Summary of Presentation

The Voice Mimic system is a synthesis-driven system for speech understanding. By modeling the human vocal tract, speech input is transformed into a set of parameters defining the shape of the tract corresponding to the instantaneous speech element (a non-unique inversion). These parameters are then used to drive the recognition process.

10 Design of "Voice Mimic" II, Closed-loop Optimization – Qiguang Lin

10.1 Summary of Presentation

The present model of the vocal tract is 2-dimensional and axisymmetric. For the purposes of vowel recognition, three parameters have been deemed most important. This model has been tested on both synthetic vowel input and natural vowel input. Close matches were obtained for synthetic vowel sounds, less close for natural vowel sounds, a likely consequence of the current linear nature of the model.

A demonstration tape of vowel synthesis was played.

10.2 Questions and Discussion

The question of whether the synthesis process was currently a "real-time" process was asked. Dr. Flanagan responded that it is not, it currently runs at about 100:1.

11 Design of "Voice Mimic" III, Open-loop Steering – Fredrick Zussa

11.1 Summary of Presentation

The purpose of the open-loop steering component of the system is to provide a good "initial guess" as input to the optimizer. A "nearest-neighbor" matching algorithm is used against a database of approximately 42,000 precomputed spectra, each encoded by only 4 complex numbers. The physical constraints of the problem make it possible to use a hierarchical search strategy to speed matching.

A demonstration tape of vowel matches was played.

12 Design of "Voice Mimic" IV, Articulatory-based Vowel Recognition – Gael Richard

12.1 Summary of Presentation

Each vowel sound is encoded as 8 centroids in articulatory space using the synthesis portion of the system. Incoming speech is broken up into 100msec samples, and these are further divided into 10msec samples. Each 10msec sample is matched using the 2-norm against the database. The best subsample candidate match is then chosen. Initial results exceeded 96% success against a test dataset, a number that may be improved with a more complex vocal tract model.

13 Design of "Voice Mimic" V, Speech Synthesis From Fluid Dynamic Formulation – Daniel Sinder

13.1 Summary of Presentation

A Navier-Stokes solver is being applied to a gridded model of the vocal tract (axisymmetric and with some simplifications) to predict sound resulting from pulse-trains originating at the vocal cords. While the results match well with LP synthesis, more work is needed to extend the model. Computational cost is a significant barrier: currently about 24 CPU-hours of workstation time are required for each millisecond of synthesis.

A demonstration tape of Navier-Stokes vowel synthesis was played.

14 Design of Microprocessors – Alvin Despain, Louis Steinberg et al

14.1 Summary of Presentation

The goal of this research is to study design across multiple levels of artifact representation, in a domain that already has a well-developed set of levels, and a CAD technology for each level. It is hoped that this research will be able to connect with other domains in the HPCD project.

The current state of the art in microprocessor design has many levels, many tools, and much use of CAD. The problems are that there are "islands of automation" with little support for the overall design process, and that feedback from lower up to higher levels is not supported at all.

The main effort this year was on automating the upward flow of information, using two systems. LADDER focuses on overall system structure. The initial emphasis here was on the use of feedback to calibrate estimators. PDAS focuses on specific estimators.

Multi-level design can mean at least three things: levels of modeling, levels of search space, or levels of artifact description. The research on microprocessor design has focused on levels of artifact description. Each level describes a complete microprocessor in term of different kinds of parts.

LADDER uses feedback among levels, which is necessary to approximate lower levels when making decisions at higher levels. Forms of feedback include success/failure, cost estimates, and critiques such as "to reduce time, don't do add5 in stage 4." Feedback can be used both within the same problem, and from previous problems. LADDER uses feedback to compare alternatives. PDAS uses feedback for tradeoff curves between time and area constraints.

A problem with feedback is that when comparing alternatives, taking the design deeper for each one produces better quality estimates, but is more expensive. The solution used in LADDER is to use feedback from only a few alternatives to *calibrate* a model to the current problem.

PDAS (Processor Design Automation System) reasons about the geometry at more levels, and uses feedback to get trade-offs between goals.

Future work includes combining LADDER and PDAS, extending multi-level design to lower levels and to higher levels, reasoning about the geometry at more levels, adding critique feedback, and parallel circuit simulation.

The expected impact of this work is more flexible automation of microprocessor design, better integration of levels, and automated control of multi-level design in other domains. Expected deliverables include methodologies for controlling mutli-level design, and a prototype microprocessor design system to implements those methodologies.

This research is linked to all other HPCD domains in its emphasis on multi-level analysis, and to the aircraft domain in its emphasis on multi-level design.

It is linked to lithography in its emphasis on better information flow between the electrical design level of microprocessor design and the layout design level of lithography.

15 Design of Microlithographic Processes; related Simulations and their Parallelization – Eytan Barouch, Steve Orszag, Don Smith, Jerry Richter

15.1 Summary of Presentation

15.1.1 Microlithography

Efficient, accurate analysis is critical to chip design. At the 0.25 micron level, electrical design and chip manufacture are tightly coupled. Detailed knowledge of fabrication processes is necessary to design a manufacturable microchip. Analysis tools capable of simulating the optical/chemical fabrication processes must be integrated into the chip design environment. Improved process models, efficient parallel algorithms, and enhanced representations are critical to the development and integration of these tools.

Microlithography involves shining light through a mask onto a photosensitive surface. This research has focused on the analysis of this process and the design of the mask.

Currently electrical design and lithographic manufacture are, at best, weakly coupled. At feature sizes of 4 microns, lithography is easy and lack of coupling is not an issue. Below 1 micron, lithography becomes difficult, but advanced techniques for mask design have been effective at allowing electrical design and lithographic manufacture to remain decoupled. At feature sizes of roughly 0.25 micron, it is necessary to couple the two phases.

The focus is on problems faced by industry when designing microchips. These problems and their solutions form a testbed for investigation of real-world design issues. The plan is to develop and validate fast and accurate models of the fabrication process, develop a design environment for automatic mask design, and couple electrical design and lithography.

There are several ways to perform the analysis. Direct evaluation requires $O({N_x}^3 {N_y}^3)$ operations. Fast Fourier Transform (FFT) methods are faster, requiring only $O[N_x N_y log(N_x N_y)]$ operations, but are inaccurate due to sharp edges of features and random locations of features. FAIM (Fast Aerial Image Model) does an *exact* evaluation of Fourier *integrals* over the mask, and uses a non-uniform FFT grid to account for random location of features on the mask.

OPTIMASK (Mask Optimizer *via* Optical Proximity Correction) is applicable to bright field masks and dark field masks, uses a hard correction "wall" to

prevent optical coupling, uses constrained Newton optimization for rapid convergence, uses constraints to allow optimization over the full defocus budget, is user configurable, and is able to use parallel analysis programs such as FAIM.

FAIM has been tested on the SP2. For one problem, using 16 processors allowed a solution to be computed in 45 minutes, compared with 12 hours when run on one processor.

Technology Transfer: Industrial users include IBM, AMD, HP, DEC, LSI Logic, Micron, SVGL, and SEMATECH. Mask bias designs are actively being benchmarked on parallel computers from a variety of vendors. Special purpose software is being developed that addresses specific design problems. Collaboration also exists with the following DoD Laboratories: Lincoln Laboratory, Wright-Patterson AFB, Brooks AFB, Phillips Lab, the Naval Research Laboratory, and NSA/Fort Meade.

FAIM has been used to simulate large chips on the IBM SP2 parallel supercomputer. It takes less than an hour to simulate a 4Mb DRAM. It has been estimated that previous algorithms would have taken more than 1000 CPU years on the SP2. FAIM is now used as a guide for technology decision making by leading IC manufacturers.

In conclusion: Full chip optimization/analysis capability has been achieved through the use of domain independent decomposition techniques. Efficient, fast computer codes have been parallelized and tested on three parallel supercomputers: IBM SP2, Cray T3D, and SGI Power Challenge. A general package of flexible, easy-to-use codes has been developed: FAIM, CDLOSS (finds locations where designs fail specifications), SEED (constructs phase-shifted masks to improve designs), E-OPTIMASK (iterative optimizer to eliminate failures), and GDSUTIL (general interface between flattened format for codes and hierarchical CAD description of masks).

15.1.2 Microlithography: Design for/of Manufacture

Optimask uses feature-based decomposition of the flattened electrical layout to automatically design masks. Cell-based decomposition of layout hierarchy produces additional gains, both in improved performance of design tools, and by providing a framework for the exchange of information between electrical and mask design.

A layout hierarchy describes objects to be manufactured as a hierarchy of cells, each of which is composed of subcells or arrays of subcells. The designed chip is "produced" by flattening the hierarchy. A lithographic hierarchy is different from a layout hierarchy in that the printed image depends on optically-close masks, and identical printed images may require different masks. A lithographic hierarchy defines masks used in manufacturing, rather than objects to be manufactured.

To obtain a lithographic hierarchy from a layout hierarchy, the arrays of subcells are expanded based on context. Only the subcells with identical contexts

will have identical masks.

A reconfigured hierarchy for manufacture combines design and manufacture information. It enables reuse of designed components, provides compact representations, enables efficient scaling of designs, and reduces the computational cost of lithographic analysis. The layout hierarchy can be automatically augmented to include lithographic information.

OPTIMASK's "best-possible" mask cannot always meet the specification. Redesign of the desired image can be accomplished using knowledge of the relationship between the desired image and the "best-possible" mask. The one-layer redesign approach uses CDLOSS to find the "new failures" introduced by OPTIMASK, identifies the deformations that cause these failures, and inserts "spacers" in the desired image to decouple the interacting features. The multiple-layer redesign approach uses spacers to construct a failure-specific constraint satisfaction problem in the desired image space.

The cell re-design phase must share information with the electrical layout phase, since the lithographic design phase only has geometric information, and speed and power must stay within specifications.

The current status of cell redesign is that OPTIMASK is being integrated into redesign, a credit/blame algorithm for spacer placement is being implemented, and a failure-based constraint satisfaction framework is being developed.

15.2 Questions and Discussion

Dr. Despain asked, What about interactions among masks?

Dr. Baruch answered that normally most of the work focused on the poly layer, since that is where the smallest features are. If the job is done well on the poly layer, then 99% of the job is done. If you can't do the job well on the poly layer, there may be interactions.

16 Visiometrics and Reduced Modeling – Norman Zabusky, Deborah Silver, et al

16.1 Summary of Presentation

16.1.1 Overview of recent research

In the 1994-1995 period, research in core visiometrics (visualization and quantification) included research on interactive visiometric environments using AVS/DAVID, and extraction, tracking, and quantifying coherent structures, using the CM5 and T3D supercomputers with AVS/DAVID.

Research on application visiometrics included ships and vortex and coherent structure dynamics. In the ship domain, research focused on interactive visio-

metrics for ship analysis and design, and running LAMP 4 using a wide range of Froude number and a single harmonic sea state.

Research on vortex and coherent structure dynamics included 3D vortex collapse using multifilament Biot-Savart with intermittency and vorticity intensification. This research includes parallel diagnosis, topology change (or reconnection), and space-time adaptiveness. This research has utilized the CM5 and T3D parallel supercomputers. Other research on vortex and coherent structure dynamics focused on reduced vortex models of the Richtmyer-Meshkov for laser fusion pellet instability.

16.1.2 Stages of reduced modeling

A reduced model is one in an asymptotically linked hierarchy, which is used to obtain the solution to a cogent problem and minimizes an error-cost function. Cogent means within a narrow-but-useful physical parameter range. The error-cost could be the product of solution error and computational cost, where error is a combination of truncation, filtering, round-off and physical (omission of processes), and computational cost relates to algorithm and parallelization, etc.

An early stage reduced model embodies an initial concept and design goals. It typically is linear and has a low number of degrees of freedom. Examples from the ship domain include LAMP 1 and LAMP 2.

An intermediate stage reduced model is typically nonlinear, has inhomogeneities, and has an intermediate number of degrees of freedom. LAMP 4 is an example from the ship domain.

An advanced stage reduced model has a large number of degrees of freedom, strong nonlinearity and "breaking", and merges different disciplines, by, for example, considering fluid-structure interactions. LAMP X is an example from the ship domain.

Dr. Zabusky showed various slides and videos of visualizations, including a video of a ship bouncing in waves.

17 Wrap-up; Discussion – Saul Amarel

17.1 Summary of Presentation

There have been lots of interactions at this workshop among groups, among domains, and among institutions. The key technical issues can be considered from three major *perspectives*: design processes, domain knowledge, and high performance computing. Design processes include formulation of design requirements, evaluation of solutions, formulation of search spaces, control of search and optimization, and construction of system frameworks. Domain knowledge is used for model formulation, development of simulation algorithms based on models, analysis of models, and model reduction. Work in high performance computing includes parallelization of simulation and optimization algorithms,

and development of software tools for programming in parallel and distributed environments.

Handling a design problem is an evolutionary process. All the components of a design process are progressively refined as they become better integrated into a *design loop*. The design loop is the essential core of design processes. HPC is needed to support problem reformulations by running simulations, optimizations, and computational experiments.

There was a major re-planning effort early this year, in response to substantial reductions in the expected level of funding for the year. The guidelines for re-planning included retaining the key technical orientation in the clusters, retaining all key investigators in the project, and giving priority to efforts that involve synergy, that advance the development of automatic design functionalities in significant applications, and that require and challenge HPC capabilities. The project is concentrating on a smaller number of tasks than originally planned for this year, and is trying to capitalize on the momentum attained in Year 1 to attain short term results and to advance longer term goals.

Different areas of the project are at various stages in the development of advanced design capabilities. The first stage focuses on domain knowledge acquisition. The second stage incorporates the domain into a design loop. The third stage focuses on advanced strategies, such as multi-level modeling, search, and optimization. The fourth stage involves reformulation of the design problem and solution strategy. Masks and Speech are currently in stages 1 and 2. Inlets are currently in stage 2. Ships, Nozzles, Aerospace Vehicles, and Microprocessors are currently in stage 3. Nozzles, Aerospace Vehicles, and Ships are currently in stage 4.

In conclusion, there has been substantial technical progress in the second year of the project, we have clear ideas about the future directions of the project, and we expect considerable progress in the coming years.