Conf-950633--1
SAN094-2435C

AIAA-95-2180

**NUMERICAL SIMULATION OF
SUPERSONIC WAKE FLOW
WITH PARALLEL COMPUTERS**

**C. Channy Wong
Sandia National Laboratories
Albuquerque, NM**

**Moeljo Soetrisno
Amtec Engineering, Inc.
Bellevue, WA**

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

**26th AIAA Fluid Dynamics Conference
June 19-22, 1995/San Diego, CA**

For permission to copy or republish, contact the American Institute of Aeronautics and Astronautics
370 L'Enfant Promenade, S.W., Washington, D.C. 20024

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

MASTER

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

NUMERICAL SIMULATION OF SUPERSONIC WAKE FLOW WITH PARALLEL COMPUTERS

C. Channy Wong[†]

Sandia National Laboratories
Mail Stop 0825, P.O. Box 5800
Albuquerque, NM 87185-0825.

Moeljo Soetrisno[#]

Amtec Engineering, Inc.
3055 112th Ave. NE Suite 100
Bellevue, WA 98004.

Abstract

Simulating a supersonic wake flow field behind a conical body is a computing intensive task. It requires a large number of computational cells to capture the dominant flow physics and a robust numerical algorithm to obtain a reliable solution. High performance parallel computers with unique distributed processing and data storage capability can provide this need. They have larger computational memory and faster computing time than conventional vector computers.

We apply the PINCA Navier-Stokes code to simulate a wind-tunnel supersonic wake experiment on Intel Gamma, Intel Paragon, and IBM SP2 parallel computers. These simulations are performed to study the mean flow in the near wake region of a sharp, 7-degree half-angle, adiabatic cone at Mach number 4.3 and freestream Reynolds number of 40,600. Overall the numerical solutions capture the general features of the hypersonic laminar wake flow and compare favorably with the wind tunnel data. With a refined and clustering grid distribution in the recirculation zone, the calculated location of the rear stagnation point is consistent with the 2D axisymmetric and 3D experiments. In this study, we also demonstrate the importance of having a large local memory capacity within a computer node and the effective utilization of the number of computer nodes to achieve good parallel performance when simulating a complex, large-scale wake flow problem.

I. Introduction

Simulating a supersonic laminar/turbulent wake flow field behind a conical body is a difficult task because of the complexity of the flow physics involved¹⁻⁵. The supersonic wake flow field has a recirculation zone, a free shear layer, expansion waves and a recompression shock. Previous analyses have had only limited success simulating this supersonic wake flow field^{6,7}. Predicting

the location of the rear stagnation point or the size of the recirculation zone with confidence for flow and heating analysis is even harder. Most of the experimental or numerical results assume a two-dimensional axisymmetric behavior of the wake⁸⁻¹². Very few three-dimensional experimental data or numerical simulations exist¹³⁻¹⁵.

Figs. 1 and 2 illustrate general features of a supersonic wake behind a conical body. As the gas flows around the corner towards the base of the body, it expands and turns at the point of separation. In some cases, overexpansion occurs and the formation of a lip shock is seen. The separated boundary layer on the side of the body thus forms a free shear layer. This free shear layer will reattach downstream at about one and a half body diameters from the near wake base. This location, also known as the neck of the wake, marks the downstream end of the base flow region. This base flow region includes a low speed, recirculating flow bounded by the shear layers and the base wall. Between the base of a conical body and somewhere a few body diameters beyond the neck, the flow field is referred as the near wake. Far wake usually refers to the flow field with downstream distances greater than about ten body diameters.

The dividing streamline is the streamline which separates the recirculating flow from the external flow that continues downstream into the wake. Usually this dividing streamline stagnates on the axis at the rear

Figure 1 Representation of Flow Features of a Blunted Cone in Free Flight at Supersonic/Hypersonic Speeds.

[†]Senior Member of Technical Staff.

[#]Senior Engineer. Member AIAA.

Figure 2 Schematic of Near-Wake Flow Features.

stagnation point in the neck region. The other stagnation point on the axis will be at the center of the base wall.

Just near the dividing streamline, a thin viscous layer exists. These viscous layers originated from the base corner of the body and eventually coalesce to form the viscous core of the wake. A recompression shock also exists in this region. This recompression shock is generated from the compression waves whereby the supersonic portions of the shear layers are turned parallel to the axis. Beyond the recompression region, the fluid returns slowly to ambient conditions in a process governed mainly by the normal viscous mixing phenomenon.

Since the supersonic wake profile involves complex phenomena such as a recirculation zone, free shear layers, and shock waves, simulating this flow field with such a wide range of flow speed and pressure distribution requires a large number of computational cells in order to capture the dominant flow physics, and also a robust numerical algorithm to obtain a reliable solution. This makes the PINCA code¹⁶, a scalable, parallel compressible flow solver, very attractive to perform this flow simulation on a parallel computer.

Parallel computing is considered the future direction of high performance computing. Parallel computers can provide larger computational memory and faster computing time than conventional computer such as the Cray Y-MP. The common design of a massively parallel computer has many autonomous computer nodes with their own assigned tasks and local memory. This many nodes are interconnected through a high bandwidth communication network in a two-dimensional mesh or hypercube topology. Messages are passed along these communication interconnectors to share the needed information. In the following sections, we shall briefly describe the PINCA code and address a few technical issues associated with parallel computing. Then the

results of the supersonic wake simulation will be presented.

II. Description of the PINCA Code

PINCA¹⁶, a parallel compressible flow solver, is used to simulate the supersonic wake flow field. PINCA is a three-dimensional, finite-volume, multiple-zone, upwind, implicit Navier-Stokes solver with many capabilities. It solves the Reynolds-averaged Navier-Stokes equations with user-specified chemistry and an optional one-equation Baldwin-Barth or two-equation k- ϵ turbulence model. A wide range of inflow, outflow, and wall boundary conditions are available to model complex three-dimensional flow fields. The inviscid terms of the equations are approximated using either flux-vector or flux-difference splitting in combination with upwind biased or total variation diminishing (TVD) differencing.

PINCA solves the field equations using an LU-SGS implicit finite-volume method¹⁶. This type of algorithm has proven to be a robust and efficient relaxation procedure for obtaining steady state flow fields. Most implicit algorithms require solving a large sparse system of linear equations. The LU-SGS algorithm approximately solves this system using two sweeps of a point Gauss-Seidel relaxation. When there are no source Jacobian terms (i.e., for those cases with a perfect gas model), the LU-SGS scheme is very efficient because the main block of this linear system of equations is a diagonal matrix and hence is easy to invert. Details of the equations and numerical methods can be found in References 16.

This LU-SGS algorithm with the approximate source Jacobian terms is suitable for use in a parallel computing environment, since it requires less memory per computational cell and has a smaller communication-to-computation ratio compared to a direct inversion approach. Furthermore, the domain of dependence of the LU-SGS algorithm is smaller than the direct inversion approach, consequently, its convergence characteristics are not significantly altered by parallelization through domain decomposition.

III. Parallel Computing Approach

As part of an ongoing code assessment project, PINCA has been applied to simulate a hypersonic reentry problem and a shock-induced combustion problem¹⁶. Results of these analyses have demonstrated that when utilizing a parallel computer for a flow simulation, it is important to develop a parallel

computing strategy and to minimize communication overhead in order to achieve good parallel performance.

A. Domain Decomposition

A favorable approach to parallel computing in computational fluid dynamics is via domain decomposition^{17,18}. The basic idea of domain decomposition is to break up the overall spatial domain into subdomains (subzones) which are assigned to separate computer nodes (Figure 3). In this study, the number of

Figure 3 Grid Distribution and Layout of Subdomains for Parallel Computing. (PINCA Simulation on Intel Paragon: 18 Computer Nodes).

subzones are equal to the number of computer nodes used in the calculations. The size of the subzones must be small enough to fit within the available local memory of a computer node attached to that subzone. Calculations are driven using a global time step. Subzones are connected by means of interzone patches which provide communication between adjacent subzones. This communication between the subzones on different

nodes is done by copying the data from the sending subzone to an intermediate array on the same node. This array is then sent to an intermediate array on the receiving node and subsequently used as boundary conditions in the corresponding subzone on the receiving node.

There are many ways to decompose a spatial domain, for example, a strip or a patch decomposition scheme for a two-dimensional domain. Reference 18 has investigated different decomposition schemes and has concluded that the patch decomposition scheme works well for the numerical algorithm in the PINCA code.

B. Scalability and Load Balance

Scalability of the code addresses the scaled performance; whether the code can deliver an increase in performance proportional to an increase in problem size. In most simulations, the problem is set up as follows: locally, each computer node will process data within a subzone of interest. Each subzone will have a fixed and equal number of computational cells to maintain a balance loading. When the problem size increases, more computer nodes will be utilized, though the number of computational cells per node will remain constant. In theory, the ideal performance of a scalable code should be independent of the problem size (Figure 4). In practice, such an ideal performance is very difficult to achieve. However, using the patch decomposition scheme and the parallel computing strategy discussed in Reference 19, the PINCA code is able to achieve very good efficiency and scalable performance as shown in Figure 4.

Figure 4 Comparison of Scaled Performance Between a Serial Computer and a Parallel Computer.

C. Communication Overhead

One major development effort is directed towards improving the node-to-node and global communications. Minimizing the communication time and maximizing the computation time will reduce the communication overhead, thus improving parallel computing performance. In the present simulations, in order to ensure stability and consistency, a global communication is needed in PINCA to calculate and broadcast the minimum time step size. This procedure requires several steps. The earlier version of the PINCA code uses sequential message passing calls (do-loop format) to perform this task. Performing the task in this manner is very inefficient. It leads to too many wait periods between communications. A better way is to utilize a strategy based on a geometric progression technique for determining a global minimum. Incorporating this new global minimum utility helps save at least 20% of communication time.

D. Local Memory Capacity

Reference 19 reveals that it is important to fully utilize each computer node to achieve good scale performance. One way to fully utilize each computer node is to maximize the number of computational cells per node. This depends on the local memory capacity per node. Currently, two parallel computer architectures are emerging. One architecture, as in the Intel Paragon, has a large number of computer nodes, but each with a smaller local memory capacity (e.g. 32 Megabytes). To analyze a large problem such as a supersonic wake simulation with a large number of computational cells, users are required to utilize a large number of computer nodes. This results in higher communication overhead and it also leads to a few important parallel computing issues such as problem mapping and networking.

The other architecture, as in the IBM SP2, has a relatively smaller number of computer nodes, but each node has a much larger memory capacity. At Sandia National Laboratories and at NASA Ames Research Center, each node has minimum of 128 Megabytes memory. With a larger memory capacity per node, a smaller number of nodes are needed for the same large problem as in the supersonic wake simulation. Fewer computer nodes will lead to smaller communication overhead. However the debate is how to achieve an excellent parallel performance (in the range of TeraFlop) with this design. This is an ongoing active research area.

It is difficult to determine which parallel computer architecture works better for what problem. It is highly problem dependent. For most computational fluid

dynamics problems, the popular approach to parallelism is via domain decomposition. This creates fictitious boundary cells and may lead to high communication overhead. To minimize the communication overhead, the larger memory capacity per node works better. Sometimes incorporating a few extra nodes is required to enhance the parallel performance. One main thrust of performing this supersonic wake simulation is to investigate this issue of memory and nodes requirement.

IV. Simulation of a Supersonic Wake Flow Field

This section will discuss the PINCA simulation of a supersonic wake experiment. This study will demonstrate the importance of parallel computing and check out the parallel processing algorithm. In order to achieve our goal, it is important to have detailed flow field data in the supersonic wake region. One of the experiments available was performed by McLaughlin, et al. on a sharp cone model^{11,13}. Measurements were made to study the mean flow of the near wake of a sharp, 7-degree half-angle, adiabatic cone at Mach number 4.3 and a freestream Reynolds number of 40,600 based on the body diameter (D) of 0.01872 m. For this wind tunnel model, the length of the body is 0.0762 m and the wall temperature is 299K. These data are used to assess the code and to evaluate its parallel computing performance.

The flow simulations will be performed on three parallel computers: Intel iPSC/Gamma, Intel Paragon, and IBM SP2. The Sandia Intel iPSC/Gamma is a distributed memory, multiple-instruction, multiple-data (MIMD), hypercube communication machine. It has 64 i860 computer nodes, 16 nodes with 32 Megabytes of memory and 48 nodes with 8 Megabytes of memory. The Sandia Intel Paragon X/PS is also a distributed memory, MIMD machine, but its communication is set up via two-dimensional mesh topology. It has 1840 computational nodes, each with an i860 XP RISC processor for computation and an additional i860 XP processor devoted to inter-node communication. 512 of the computational nodes have 32 Megabytes of memory each while the remaining 1328 nodes have 16 Megabytes of memory.

IBM SP2 with RS6000/590 processors (66.7 MHz) is also a MIMD machine with distributed memory. Its unique feature is its enormous local memory capacity per node and the high speed communication switch of 40 Megabytes per sec bandwidth. The IBM SP2 at Sandia has 10 nodes. The memory capacity per node ranges from 512 Megabytes to 128 Megabytes. The IBM SP2 at NASA ARC has 160 nodes. Each has

minimum of 128 Megabytes main memory and two Gigabytes disk space.

With a larger memory capacity, each node can process more data. Hence more memory capacity per node will allow for more computational cells per subzone. For a fixed-size problem, utilizing more memory within a computer node will reduce the number of subzones and also the number of boundary cells for communication. This will reduce the communication overhead and make it more efficient and attractive to run on a parallel computer.

A. Case Study I - Axisymmetric Laminar Wake

The first simulation is performed on the Intel Gamma machine utilizing 8 computer nodes. Only those nodes with 32 Megabytes of memory are used. Thus each node processes data within a subzone that has 90 by 90 computational cells. Totally the global domain has 64,800 computational cells. Overall PINCA captures the main features of the hypersonic laminar wake flow and its comparison is qualitatively consistent with McLaughlin, et al.'s wind tunnel data¹¹. However the PINCA predicted location of the rear stagnation point is slightly further aft than the experimental location. This discrepancy may be due to the grid distribution in the recirculation zone not being fine enough to capture the flow physics in this region. The second simulation will apply a finer grid distribution in this base flow region.

The second simulation has the grids clustered in the recirculation zone and the grids re-aligned parallel to the free shear layer. To further improve the base flow solution, we also refined the global grid system. This global grid system has 160,000 computational cells; at least 20,000 in the recirculation zone (Figure 5). This simulation, performed on IBM SP2, produces the best result so far. Figure 6 shows the predicted pressure distribution. The calculated pressure field is consistent with the experimental data (Figure 7). Grid refinement helps capture the bow shock and free shear layer better (Figure 8), reducing the amount of numerical diffusion, thus improving the accuracy of the solution especially in the recirculation zone. Figure 9 plots the streamlines behind the cone. It shows the behavior of the free shear layer and the converging streamlines. PINCA predicts that the rear stagnation point occurs at the axial distance of 0.14 m. This location is closer to the base of the body than the coarse grid simulation. This gives the predicted ratio of $\Delta x/D$ to be 3.5, which is consistent with the measurement. Δx is the centerline distance between the rear stagnation point and the base of the body.

The backflow velocity in the recirculation zone is slightly overpredicted. The maximum calculated

backflow velocity is about 40% of the freestream velocity (Figure 8). In many backward facing step experiments, the backflow velocity is usually less than 30% of the freestream velocity²⁰. This discrepancy could be caused by underpredicting the level of mixing and entrainment. A follow-up question is whether the flow has transitioned into turbulent flow in the wake region. To address this question, a sensitivity study has been performed to simulate a turbulent wake flow to determine the fluid characteristics in the recirculation zone.

B. Case Study II - Axisymmetric Turbulent Wake

In many cases, the flow field behind a backward facing step can transition very quickly to turbulent flow. Since there is no experimental evidence that the flowfield behind McLaughlin's cone is laminar, therefore we investigate the effects of turbulence in the wake region. In particular, we are interested in the size of the recirculation zone and the range of backflow velocities.

The fine grid system is used in the turbulent flow calculations. The forebody of the cone is assumed to be laminar. This is achieved in PINCA by setting the transition point behind the base of the cone. The k-epsilon two-equation model with Zeman's compressibility correction is used to model the turbulent wake²¹.

Figure 10 shows the streamlines of the flow field. As expected, the rear stagnation point moves closer to the base than the laminar prediction resulting in $\Delta x/D$ of 1.8. However the maximum backflow velocity is about 28% of the freestream velocity. Hence this sensitivity study shows that in the McLaughlin's experiment, the wake flow could be in the transitional regime.

C. Case Study III - Three-Dimensional Wake

To further illustrate the importance of parallel computing for wake flow calculations, we computed the same cone at 5.5 degrees angle of attack. At this angle of attack, an interesting bifurcation in both the vertical and the horizontal planes is observed. The pressure profiles have four minima and the flow separates, leading to two vortices coming off the leeward surface of the cone.

A coarse and a medium grid were utilized for these calculations each consisting of 1.5 and 2.5 million cells, respectively. Calculations were performed using 32 nodes IBM SP2. A laminar flow is assumed.

Figures 11 through 13 show the streamlines near the cone surface, the Mach number contours at the cross-section at the base of the cone, and the streamlines at the

symmetry plane. As shown in Figure 11, we can see a small separation off the body which lead to a pair of vortices. The recirculation zone after the body is much smaller as compared to the zero degree angle of attack case. Unfortunately, there is not sufficient experimental data available for more comparisons.

As a comparison to the previous separated flow case, we also computed a $M=7.95$ flow around a 10-degree half-angle cone at 20 degrees angle of attack. Experiments on this cone were performed at the Graduate Aeronautical Laboratories at the California Institute of Technology by Tracy²².

Only the forebody is computed using approximately 1.2 million cells. Calculations were performed using 32 IBM SP2 nodes requiring approximately 1 hour (elapsed time) per 1000 steps. The resulting surface streamlines are shown in Figure 14. Unlike McLaughlin's case, this condition leads to a much bigger separation off the surface. Also since the Mach number is sufficiently high, the base flow influence on the forebody is limited in the boundary layer regions. The surface pressure profiles agree very well with the experimental data.

V. Parallel Performance Evaluation

Elapsed time and node-time are two popular parameters used to evaluate the performance of a parallel code. Elapsed time usually associates with the 'wall-clock' time minus the time required for input and output data processing. Node-time, a product of elapsed time and number of computer nodes, is an indication of the computational cost of a solution. A comparison of parallel performance between Intel Gamma and Intel Paragon will be discussed first. The first wake flow simulation which is performed on the Sandia Intel Gamma computer takes about 158 μs per computational cell per time step (Table 1). The same simulation when performed on the Sandia Intel Paragon computer requires 119 and 50 μs per computational cell per time step for 8 nodes and 18 nodes, respectively. This comparison reveals that simulating the wake problem on Intel Paragon computer with 18 nodes is about three times faster than simulating it on the Intel Gamma computer. The simulation using Intel Paragon is faster and has a better 'wall-clock' speed because (1) more computer nodes are utilized, i.e. 18 nodes compared to 8 nodes for this fixed-size problem, and (2) the Paragon has better, advanced, improved, and faster microprocessors, i860 XP.

Table 1 also compares the performance of Intel Paragon and IBM SP2 computers for the same fixed-sized problems. These results show that the IBM SP2

with 3 nodes performs better than the Intel Paragon with 18 nodes. This comparison demonstrates: (1) single processor performance in IBM SP2 is faster, and (2) a larger memory capacity per node will result in a smaller communication overhead, thus leading to a better parallel performance.

We have achieved good parallel performance when executing PINCA on the IBM SP2 for the three-dimensional wake problems. Table 1 shows the elapsed time and node-time requirement between two computational grid systems. With the addition of close to 60 thousand communication cells (globally 2.5 million computational cells compared to 1.5 million computational cells), the increase in the elapsed time is relatively small, at about 0.31 μs /cell/time step. That is only 4.5% increase. This demonstrates that PINCA has acceptable small communication overhead and scales reasonably well for the large problems (Figure 15). This finding also reveals the important aspect of large memory capacity per node.

VI. Conclusions and Recommendations

To accurately and efficiently simulate complex three-dimensional compressible flow problems such as a supersonic wake, a supercomputer with an enormous memory capacity and a super computing speed is needed. Parallel computers satisfy these requirements. The present study utilizes an efficient parallel computing strategy to analyze complex supersonic wake flow problems. This includes incorporating a domain decomposition technique, maintaining scalability and load balance, and minimizing communication overhead.

This paper presents very encouraging results of numerical simulations of supersonic wake flow fields using parallel computers. Several wake flow simulations have been performed. Overall the PINCA code can correctly capture the general features of the wake profile. By refining the grid distribution in the recirculation zone, the predicted location of the rear stagnation point compares very well with the measurement. These include axisymmetric as well as three-dimensional supersonic wake flow fields. These results demonstrate that parallel computing with its unique capabilities, can analyze a complex, large-scale supersonic wake flow problem in a cost-effective matter.

VII. References

1. Lykoudis, P. S., "A Review of Hypersonic Wake Studies," *AIAA Journal*, Vol. 4, No. 4, April 1966, p. 577-590.

2. Crane, R. I., "A Survey of Hypersonic Near Wake Studies," Aeronautical Journal of the Royal Aeronautical Society, Vol. 73, November 1969, p. 998-1006.
3. Chang, P. K., Separation of Flow, Pergamon Press, 1970.
4. Berger, S. A., Laminar Wakes, American Elsevier Publishing Company, Inc., New York, 1971.
5. Dutton, J. C., Herrin, J. L., Molezzi, M. J., Mathur, T., and Smith K. M., "Recent Progress on High-Speed Separated Base Flows," AIAA 95-0472, presented at the 33rd AIAA Aerospace Science Meeting, Reno, NV, January 9-12, 1995.
6. Patel, V. C., and Scheuerer, G., "Calculation of Two-Dimensional Near and Far Wakes," AIAA Journal, Vol. 20, No. 7, July 1982, p. 900-907.
7. Butler, G., et al., "Ballistic Range Flowfield Measurements of the Hypervelocity Near Wake of Generic Shapes and Correlation with CFD Simulations," AIAA 90-0621, presented at the 28th Aerospace Sciences Meeting, Reno, NV, Jan. 8-11, 1990.
8. Zakkay, V., and Cresci, R. J., "An Experimental Investigation of the Near Wake of a Slender Cone at $M_\infty = 8$ and 12," AIAA Journal, Vol. 4, No. 1, January 1966, p. 41-46.
9. Martellucci, A., Trucco, H., and Agnone, A., "Measurements of the Turbulent Near Wake of a Cone at Mach 6," AIAA Journal, Vol. 4, No. 3, March 1966, p. 385-391.
10. Murman, E. M., "Experimental Studies of a Laminar Hypersonic Cone Wake," AIAA Journal, Vol. 7, No. 9, September 1969, p. 1724-1730.
11. McLaughlin, D. K., Carter, J. E., Finston, M., and Forney, J. A., "Experimental Investigation of the Mean Flow of the Laminar Supersonic Cone Wake," AIAA Journal, Vol. 9, No. 3, March 1971, p. 479-484.
12. Herrin, J. L., and Dutton, J. C., "Supersonic Base Flow Experiments in the Near Wake of a Cylindrical Afterbody," AIAA Journal, Vol. 32, No. 1, January 1994, p. 77-83.
13. Browand, F. K., Finston, M., and McLaughlin, D. K., "Wake Measurements Behind a Cone Suspended Magnetically in a Mach Number 4.3 Stream," AGARD CP 19, May 1967.
14. Henderson, T. L., Dutton, J. C., and Loth, E., "Computational Study of the Inviscid, Nonlinear Physics of Supersonic Base Flows," AIAA 95-0783, presented at the 33rd AIAA Aerospace Science Meeting, Reno, NV, January 9-12, 1995.
15. Ober, C. C., Lamb, J. P., and Reinecke, W. G., "The Aerodynamics of Tandem Bodies at Mach 5. Part II: Computational Results," AIAA 95-0320, presented at the 33rd AIAA Aerospace Science Meeting, Reno, NV, January 9-12, 1995.
16. Wong, C. C., Soetrisno, M., Blottner, F.G., Imlay, S.T., and Payne, J.L., "PINCA: A Scalable Parallel Program for Compressible Gas Dynamics with Nonequilibrium Chemistry," SAND94-2436, Sandia National Laboratories, Albuquerque, New Mexico, April 1995.
17. Keyes, D., "Domain Decomposition: A Bridge Between Nature and Parallel Computers," in ASME AMD-Vol. 157: Adaptive, Multilevel, and Hierarchical Computational Strategies, presented at the Winter Annual Meeting of ASME, Anaheim, CA, Nov. 8-13, 1992, pp. 293-334.
18. Wong, C. C., Soetrisno, M., Blottner, F.G., and Payne, J.L., "A Domain Decomposition Study for Massively Parallel Computing in Compressible Gas Dynamics," AIAA 95-0572, presented at the 33rd AIAA Aerospace Science Meeting, Reno, NV, January 9-12, 1995.
19. Wong, C. C., Soetrisno, M., Blottner, F.G., and Payne, J.L., "Implementation of a Parallel Algorithm for Thermo-Chemical Nonequilibrium," AIAA 95-0152, presented at the 33rd AIAA Aerospace Science Meeting, Reno, NV, January 9-12, 1995.
20. Amatucci, V. A., "An Experimental Investigation of the Two-Stream, Supersonic, Near-Wake Flowfield Behind a Finite-Thickness Base," Ph. D. Thesis, University of Illinois at Urbana-Champaign, 1990 and also private communication, June 1995.
21. Zeman, O., "Dilatation dissipation: The Concept and Application in Modeling Compressible Mixing Layers," Physics of Fluids A, Vol. 2, No. 2, February 1990.
22. Tracy, R. R., "Hypersonic Flow over a Yawed Circular Cone," Galcit memo 69, California Institute of Technology, Pasadena, California, 1969.

Acknowledgment

This work was supported by the United States Department of Energy under contract DE-AC04-94AL85000. Simulations were performed at the DOE's Massively Parallel Computing Research Laboratory at Sandia, the Engineering Sciences Center Computing Facility at Sandia National Laboratories, and the Numerical Aerodynamic Simulation facility at NASA Ames Research Center. Their supports are really appreciated.

The authors would like to acknowledge William L. Oberkampf for providing the illustration of the supersonic wake structures shown in Figures 1 and 2. Many thanks to Vincent Amatucci, Fred Blottner, Scott Imlay, and Jeff Payne for their valuable comments and discussions.

Figure 5 Grid Distribution in the Recirculation Region and at the Corner of the Base. (See Figure 3 for the Setup of the Global Grid System.)

Figure 6 Pressure Distribution Predicted by PINCA. (Sharp Cone, Mach 4.3, Axisymmetric Geometry, 160,000 computational cells.)

Figure 7 Comparison of Pressure Distribution Between Experimental Data (symbols) and Prediction (solid line).

Figure 8 Velocity Ratio (Local Speed/Freestream Velocity) Predicted by PINCA.
(Sharp Cone, Mach 4.3, Axisymmetric Geometry, 160,000 cells.)

Figure 9 Streamline Profile Predicted by PINCA. (Sharp Cone, Mach 4.3,
Axisymmetric Geometry, 160,000 computational cells.)

Figure 10 Streamline Profile Predicted by PINCA with Turbulent Model. (Sharp
Cone, Mach 4.3, Axisymmetric Geometry, 160,000 computational
cells.)

Figure 11 Streamlines Near the Cone Surface Predicted by PINCA. (Sharp Cone, Mach 4.3, 5-degree angle-of-attack, 3D, 2,500,000 computational cells.)

Figure 12 Mach Number Contours at the Base of the Cone. (Sharp Cone, 5 degree angle-of-attack, 3D, 2,500,00 computational cells.)

Figure 13 Streamlines at the Symmetry Plane. (Sharp Cone, 5 degree angle-of-attack, 3D, 2,500,00 computational cells.)

Figure 14 Streamlines Near the Surface of Tracy Experiment. (Sharp 10-degree Cone, 20 degree of angle-of-attack, Mach 7.95)

Figure 15 Scaled Speedup Versus Number of Computer Nodes to Illustrate the Scaled Performance of PINCA.

Computer	No. of Nodes	Total Number of Computational Cells	Elapsed Time ($\mu\text{s}/\text{cell}/\text{time step}$)	Node-Time ($\mu\text{s}/\text{cell}/\text{time step}$)
Intel Gamma	8	64,800	158.0	1264.0
Intel Paragon	8	64,800	119.1	952.6
Intel Paragon	18	64,800	50.0	900.4
IBM SP2	3	64,800	38.6	115.8

Table 1 Parallel Performance of Different Computers for a Fixed-Sized Problem.

Computer	No. of Nodes	Total Number of Computational Cells	Elapsed Time ($\mu\text{s}/\text{cell}/\text{time step}$)	Node-Time ($\mu\text{s}/\text{cell}/\text{time step}$)
IBM SP2	32	1,534,464	6.52	208.7
IBM SP2	32	2,429,568	6.83	218.4

Table 2 Parallel Performance of IBM SP-2 for Three-Dimensional Problems with Different Sizes.