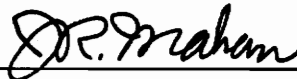# ACCELERATION TECHNIQUES FOR THE RADIATIVE ANALYSIS

OF GENERAL COMPUTATIONAL FLUID DYNAMICS SOLUTIONS

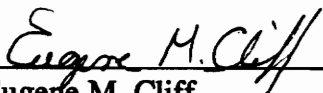# USING REVERSE MONTE-CARLO RAY TRACING

by

Jeffrey A. Turk

Dissertation submitted to the Faculty of the

Virginia Polytechnic Institute and State Univeristy

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in

Aerospace Engineering
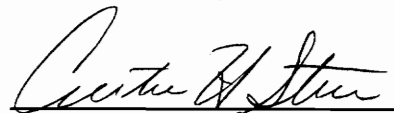
APPROVED

J. Robert Mahan, Chairman

Eugene M. Cliff

Wayne C. Durham

William H. Mason

Curtis H. Stern

December, 1994

Blacksburg, Virginia

c.2

LD
5655
V856
1991
T875
c.2

ACCELERATION TECHNIQUES FOR THE RADIATIVE ANALYSIS

OF GENERAL COMPUTATIONAL FLUID DYNAMICS SOLUTIONS

USING REVERSE MONTE-CARLO RAY TRACING

by

Jeffrey A. Turk

J. Robert Mahan, Chairman

Aerospace Engineering

(ABSTRACT)

A reverse Monte-Carlo ray trace capable of performing a radiative analysis on arbitrary multiple overlapping structured computational fluid dynamics solution sets is developed. In order to make effective use of time, a method based on a set of simplifying assumptions but using the same calculation procedures is developed for comparison and study purposes.

Three acceleration techniques are tried. One acceleration technique reduces the grid dimensions to reduce the number of volumes intersected. The second acceleration technique develops a version of the code for execution in a parallel processing environment. The third acceleration technique mixes an orthogonal, evenly spaced grid with the computational fluid dynamics grids to obtain fast ray traversal of low variance areas while retaining the higher resolution of the computational fluid dynamics grids in the high variance areas.

Two experimental data sets are used for comparison and as test cases during these studies: an exhaust plume from an auxiliary power unit, and a Boeing 747 in flight. Timing for the baseline and accelerated analyses is provided as well as numerical comparisons for a selected subset.

ACKNOWLEDGEMENTS

fruitful one. The practical experiences in this field of research that I obtained at Ames will doubtlessly be useful to me for the rest of my life.

I would also like to thank my fellow graduate students, Ed Nelson and Pierre Villeneuve, for always being there with answers to my interminable questions. I know they will always be out there somewhere with answers to more of my questions.

Finally, I would also like to express my deepest thanks to my parents and brothers for their heart-felt support and encouragement over this long time. The support of one's family is often so complete and all-pervasive that one takes it for granted. On special occasions it is good to reflect and give thanks for that support.

TABLE OF CONTENTS

Table of Contents   x

LIST OF FIGURES

LIST OF TABLES

NOMENCLATURE

Symbols:

A_i Area of element i (m^2)

a_λ Absorption coefficient (cm^{-1})

B Scanner calibration constant   (K)

b Offset value   (-)

C Species concentrations   (-)

C_1 Constant in equation 1 $(0.54544 \times 10^8 \; W \cdot \mu m^4 \cdot m^{-2})$

C_2 Constant in equation 1 $(14\,388 \; \mu m \cdot K)$

D_{ij} Distribution factor from element i to element j   (-)

dE_i Energy emitted or absorbed by element i within a differential wavelength band $d\lambda$ centered about λ (W)

dE_{ij} Energy absorbed by element i and originating from element j within a differential wavelength band $d\lambda$ centered about λ (W)

ERR   Difference between floating point representation of ray location and rounded integer representation of ray location   (-)

F Scanner calibration constant   (-)

I First component of mesh labeling triplet   (-)

i First component of grid labeling triplet   (-); counter   (-)

$i'_{\lambda b}$ Spectral directional blackbody intensity $(W \cdot \mu m^{-1} \cdot m^{-2} \cdot sr^{-1})$

J   Second component of mesh labeling triplet (-)

j   Second component of grid labeling triplet (-); counter (-)

K   Third component of mesh labeling triplet (-)

k   Third component of grid labeling triplet (-); counter (-)

m   Nondimensionalized direction (-); nondimensionalized mesh spacing (-)

N   Count quantity (-)

n   Similarity profile parameter (-)

P   Pressure (atm)

R   Scanner calibration constant (-)

s   Path length (cm)

T   Temperature (K)

t   Ray step size (-); time (s)

tr   Temperature ratio trigger (-)

VS   Scanner isothermal value (-)

x   x-coordinate in nondimensionalized global coordinates (-)

y   y-coordinate in nondimensionalized global coordinates (-)

z   z-coordinate in nondimensionalized global coordinates (-)

α_{λ} Spectral absorptance (-)

α_{λ}' Spectral directional absorptivity (-)

ϵ_λ Spectral emittance (-); spectral hemispherical emissivity (-)

$\epsilon_\lambda{}'$ Spectral directional emissivity (-)

θ Polar angle of incident radiation (-)

θ_r Polar angle of reflected radiation (-)

λ Wavelength (micrometers)

$\rho_\lambda{}'$ Spectral directional-hemispherical reflectivity (-)

$\rho_\lambda{}''$ Spectral bidirectional reflectivity (-)

τ_λ Spectral transmittance (-)

$\tau_\lambda{}'$ Spectral directional transmissivity (-)

ϕ Circumferential angle of incident radiation (-)

ϕ_r Circumferential angle of reflected radiation (-)

ω_{ij} Solid angle subtended by element j at element i (sr)

[✳]   Representative symbol for multiplicative set (-)

Subscripts:

align   Refers to a value used for alignment

all   Refers to summation over both solid and volume elements

b   Refers to a blackbody quantity

INTER Refers to intersection quantities

i   Refers to: element i; counter i; mesh grid i component

j   Refers to: element j; counter j; mesh grid j component

Nomenclature   xvii

k	Refers to: element k; counter k; mesh grid k component
MESH	Refers to mesh grid quantities
MOD	Refers to radiative model quantities
m	Refers to a volume element
n	Refers to the first n elements, refers to second volume element
o	Refers to an observer quantity
p	Refers to a projected quantity
pt	Refers to the current ray value
r	Refers to reflected quantities
ray	Refers to quantities for a given ray
S	Refers to a surface element
SYS	Refers to computer system quantities
TOT	Refers to a total value
V	Refers to a volume element
x	Refers to the ray x-component
y	Refers to the ray y-component
z	Refers to the ray z-component
λ	Refers to a spectral quantity

CHAPTER 1

INTRODUCTION

Infrared radiation is defined as any electromagnetic radiation with wavelengths between 0.78 and 1000.0 μm [1]. For comparison, visible light is composed of wavelengths between 0.4 and 0.78 μm. This infrared wavelength range is fairly broad but in the current application the most interest is in the 3-to-5- and the 8-to-12-μm wavebands. These subbands are called the *mid wavelength* and *long wavelength* bands, respectively. The reason for the interest in infrared radiation stems from the fact that all objects whose temperatures are above absolute zero emit infrared radiation as a consequence of their molecular agitation. Humans perceive this radiation as heat. Passive technology can be used to detect and measure this radiation from a distance, thereby posing no risk of perturbing the source. Additionally, the velocity of propagation of the radiation is sufficiently high that it allows virtually instantaneous measurement.

"Heat-seeking" missiles use passive detection from a distance to enable them to track aircraft. As opposed to radar tracking, the aircraft itself is the emitter of the thermal radiation and so cannot detect when it is being tracked. This vulnerability is displayed by

a study [2] that attributes 308 of the 408 aircraft destroyed by missiles between 1958 and 1990 to infrared-seekers and by Klass [3], who estimates that 80 percent of the allied aircraft lost during operation Desert Storm fell victim to infrared missiles. Aircraft makers have an incentive therefore in being able to perform a radiative analysis on existing and new aircraft designs in order to establish and possibly mitigate any infrared vulnerability. Several radiative analysis programs exist which can accomplish this goal to varying degrees of accuracy. One of the more successful programs is SPectral InfraRed Imaging of Targets and Scenes (SPIRITS) which performs a radiative analysis on a plume and then incorporates that analysis into one involving a wireframe model of the vehicle [4]. This approach has the weakness that the plume and the solid vehicle exist as totally separate and non-interacting entities. Neither the influence of surface radiation passing through the plume, nor plume reflections in the solids is modeled. Many radiative analysis methods make similar or stronger assumptions concerning the relevant radiative phenomena and the data set to be analyzed.

These problems would be eliminated if the radiative analysis started with the results from a Computational Fluid Dynamics (CFD) analysis. The CFD results contain the vehicle geometry, the appropriately modeled plume, the aerodynamic surface heating, and an accurate representation of the air flow about the vehicle. By using the CFD results, the infrared analysis becomes another step in the design process using a common set of

information and retaining a comparable level of fidelity instead of being a completely separate process operating under different assumptions and modeling accuracies.

An additional use for a radiative analysis which uses CFD results would be to provide an alternative means to validate those CFD results. Birckelbaw and Nelson at the National Aeronautical and Space Administration (NASA) Ames Research Center have successfully demonstrated that infrared imaging of aircraft in flight can be an effective tool for flow visualization [5]. The infrared images can be compared on a qualitative level with particle traces from CFD results. The natural extension of infrared imaging of flowfields is to compare the measured infrared images to images obtained by a radiative analysis of a comparable CFD analysis to obtain a quantitative comparison. This quantitative comparison could then be used to establish the validity of the CFD analysis for the associated flow conditions. While infrared imaging of the flowfield is not always the best temperature measurement technique, its strengths allow it to obtain measurements under actual operating conditions where conventional measurement techniques are awkward or impossible to apply.

Both of these applications - the design and analysis of aircraft and the CFD validation effort - require a radiative analysis tool capable of handling arbitrary CFD grid geometries and correctly modeling any radiative phenomena that might be encountered under actual flight conditions. These phenomena can include scattering, refraction, multiple

bidirectional reflections, and rapid spectral variation of gaseous properties. The approach best able to fulfill these requirements is a Monte-Carlo ray trace. A variation on the technique, the reverse Monte-Carlo ray trace, is developed here since it uses a reduced workload to provide a computed infrared image from a fixed observer viewpoint. A serious drawback to ray-tracing methods in general is the large execution time requirement. The reverse Monte-Carlo ray trace method proves not to be an exception, requiring relatively long execution times. For the reverse Monte-Carlo ray trace to be a viable analysis method, a means must be found for accelerating the analysis.

The goal of the research reported here is two-fold. The first goal is to develop a reverse Monte-Carlo ray trace capable of handling any arbitrary CFD grid geometry and its corresponding solution files as inputs. The analysis must be capable of supporting a wide range of radiative phenomena across a specified wavelength band for an observer with a given spatial resolution. The second goal is to develop techniques capable of accelerating the reverse Monte-Carlo ray trace without the imposition of unnecessary assumptions or the unacceptable loss of modeling accuracy. A set of experimental infrared images along with their CFD simulations is used as test cases to establish the accuracy and time requirements of the acceleration techniques, and to serve as representative CFD data sets.

CHAPTER 2

REVIEW OF RADIATION PHYSICS

This chapter provides a quick review of some key concepts and terminology in radiation physics. The review is not intended to be exhaustive but to provide a common background and basic terminology. More detail is available from any thermal radiation textbook or the contemporary literature. A notable example is Siegel and Howell's **Thermal Radiation Heat Transfer** [1]. Throughout this dissertation all radiative properties used are spectral rather than band averaged because the use of broad-band radiative properties requires the adoption of restrictive assumptions. Finally, standard texts on thermal radiation use the word *energy*, when the word *power*, energy per unit time, is more appropriate physically. The work presented in this dissertation follows that tradition and used the word energy rather than power.

2.1 Fundamentals

Two key concepts to introduce are *local thermodynamic equilibrium* and the *blackbody*. Local thermodynamic equilibrium is a state of an element where any absorbed energy from the incident radiation is quickly redistributed throughout the element before re-

emission can take place. If this condition does not hold, highly localized parts of the element would radiate energy that does not follow the spectral pattern for an object at its indicated temperature. The radiative analysis developed here is based on the assumption that local thermodynamic equilibrium applies within the modeled space.

A *blackbody* is a conceptual idealization defined as an object that absorbs all the radiation incident upon it from every direction and at every wavelength. A consequence of this definition is that no object can emit more radiation than a blackbody at the same temperature. This characteristic is true for all wavelengths and for all emissive directions. Blackbodies radiate into a vacuum following Planck's blackbody distribution function,

$$i'_{\lambda b}(\lambda) = \frac{2 C_1}{\lambda^5 \left(e^{C_2/\lambda T} - 1 \right)} .$$

(1)

As a consequence of this characteristic, the intensity of a blackbody is used as a yardstick to compare the radiant intensity of real objects.

2.2 Surface Radiation

Radiation incident upon a solid surface has three possible fates: it can be absorbed, transmitted, or reflected. The *spectral, directional absorptivity*, $\alpha'_{\lambda}(\lambda,\theta,\phi,T)$, is defined as the ratio of energy absorbed at the surface to the energy incident on the surface from

direction (θ,ϕ) within a wavelength interval $d\lambda$ centered about wavelength λ. It can range from zero to unity in value. For a given surface material and surface finish, the absorptivity varies with the wavelength involved, the direction of the incident radiation, and the temperature of the solid. Absorptivity is usually not measured directly but is obtained from reflectivity measurements.

The *spectral, directional transmissivity*, $\tau'_\lambda(\lambda,\theta,\phi,T,s)$, is defined as the ratio of energy transmitted through the material to the original energy incident upon the surface from direction (θ,ϕ) within a wavelength interval $d\lambda$ centered about wavelength λ. It also ranges from zero to unity and is a function of the material itself, the wavelength, the incident direction, the material temperature, and the thickness traversed. Generally transmission through a material will involve a change in the index of refraction resulting in a change in the direction and wavelength of the radiation due to the change in velocity of propagation. This change in direction follows Snell's law. The frequency of the radiation remains constant across the interface between the two materials.

Reflection is a bit more complicated. The most general reflectivity is *spectral, bidirectional reflectivity*, $\rho''_\lambda(\lambda,\theta,\phi,\theta_r,\phi_r,T)$. It is defined as the ratio of radiation intensity reflected into direction (θ_r,ϕ_r) to the energy incident upon the surface in the incident direction (θ,ϕ) per unit area per wavelength interval $d\lambda$ centered about wavelength λ. This

reflectivity is a surface property dependent upon the wavelength, the incident direction, the reflected direction, and the temperature of the material. Two special limiting cases of spectral bidirectional reflectivity are *diffuse* and *specular*. Diffuse reflection occurs when the incident energy is reflected with the same intensity in all directions. Specular reflection is when the incident energy is only reflected like a mirror. In other words the reflected beam is $180°$ circumferencially from the incident direction and the incident and reflected beams make the same angle relative to the local surface normal. Experimental measurements show that many materials exhibit a similar type of bidirectional reflection pattern [6-9]. This pattern is characterized by having a large fraction of the reflected energy being found at or near the specular angle, with a variable amount being reflected in the other directions. The amount and angular extent of the reflected energy that approximates the specular response is strongly dependent on the incident beam's direction.

One important property of bidirectional reflectivity is the relationship:

$$\rho''_\lambda(\lambda,\theta,\phi,\theta_r,\phi_r,T) = \rho''_\lambda(\lambda,\theta_r,\phi_r,\theta,\phi,T) .\qquad(2)$$

This relationship is called *reciprocity* and means that the reflectivity from one incident direction into a specific reflected direction is equal to the reflectivity from that reflected direction into the original incident direction.

One last definition needed for the discussion of reflectivity is the *spectral, directional-hemispherical reflectivity*, $\rho'_\lambda(\lambda,\theta,\phi,T)$. This reflectivity is defined as the ratio of the sum of the energies in all the reflected directions to the energy in the one incident direction for a wavelength interval $d\lambda$ centered about wavelength λ. The spectral, directional-hemispherical reflectivity ranges from zero to unity in value.

These three radiative quantities are related by

$$\alpha'_\lambda(\lambda,\theta,\phi,T) + \tau'_\lambda(\lambda,\theta,\phi,T,s) + \rho'_\lambda(\lambda,\theta,\phi,T) = 1.0 \ . \tag{3}$$

Equation 3 is simply a restatement of the first sentence in this section; namely that all the incident energy must be either absorbed, transmitted, or reflected. Most solid objects encountered during the radiative analysis of aircraft are opaque and have a transmissivity of zero. Taking this into account, equation 3 reduces to

$$\alpha'_\lambda(\lambda,\theta,\phi,T) + \rho'_\lambda(\lambda,\theta,\phi,T) = 1.0 \ . \tag{4}$$

This relationship also explains how experimental data from spectral, directional-hemispherical reflectivity measurements can be used to indirectly obtain the spectral, directional absorptivity of a material.

Besides interacting with incident radiation, a solid also emits radiation. The *spectral, directional emissivity*, $\epsilon'_\lambda(\lambda,\theta,\phi,T)$, is defined as the ratio of energy emitted by a solid in

(θ,ϕ) direction within a wavelength interval $d\lambda$ centered about wavelength λ to the energy emitted by a blackbody at the same temperature in the same direction and within the same wavelength interval. This quantity is dependent on the material involved, the wavelength, the direction, and the temperature of the solid. Kirchhoff's law, based on energy equilibrium, establishes that the spectral, directional emissivity of an object is equal to its spectral, directional absorptivity. Using this relation equation 4 becomes

$$\epsilon'_\lambda(\lambda,\theta,\phi,T) + \rho'_\lambda(\lambda,\theta,\phi,T) = 1.0 \ . \tag{5}$$

2.3 Volume Radiation

Both gaseous media and suspended particles require physical depth to interact with radiation. The gaseous media can range from atmospheric gases to the highly heated and species-laden jet engine exhaust plume. No differentiation is made here between the nomenclature of a volume of plume gases and a volume of atmospheric gases since both are subject to the same physical phenomena. Radiation incident on a gaseous volume is either absorbed or transmitted. Scattering of radiation at the wavelengths studied in this work by molecular scattering processes, i.e. Rayleigh scattering, is negligible and so scattering by the gases themselves is ignored.

Bouguer's law defines the basic radiative property of a *spectral absorption coefficient*, $a_\lambda(\lambda, C, T, P)$, for a gas volume as

$$i'_\lambda(s) = i'_\lambda(0)\text{EXP}\left\{-\int_0^s a_\lambda(\lambda, C(s^*), T(s^*), P(s^*))ds^*\right\}, \tag{6}$$

where $i'_\lambda(s)$ is the radiant intensity remaining after traversing through a gaseous volume with a path length s, and $i'_\lambda(0)$ is the radiant intensity originally incident on the volume. The absorption coefficient is a function of the wavelength of the radiation, the gaseous species present, the temperature, and the pressure. Typically these last three quantities will vary with location in a gaseous volume, so the absorption coefficient is indirectly a function of position also. While the other radiative quantities defined thus far have been nondimensional, the absorption coefficient has dimensions of cm^{-1}.

The *spectral absorptance*, $\alpha_\lambda(\lambda, C, T, P, s)$, of a column of gas of length s is defined as the ratio of energy absorbed by the volume to the energy incident within a wavelength interval $d\lambda$ centered about wavelength λ upon the volume. Absorptance is related to the absorption coefficient by

$$\alpha_\lambda(\lambda, C, T, P, s) = 1.0 - \text{EXP}\left\{-\int_0^s a_\lambda(\lambda, C(s^*), T(s^*), P(s^*))ds^*\right\}. \tag{7}$$

This dissertation uses the absorptance in its derivations, even though gaseous thermal radiation is more traditionally expressed in terms of the absorption coefficient. Absorptance depends on the wavelength, the concentrations of gaseous species present, the temperature, the pressure, and the path length s. Even if the absorption coefficient is not a function of position, as for a homogenous gas volume, the absorptance still depends on the path length traversed by the radiation. Note that the absorptance is a nondimensional quantity which ranges in values from zero to unity. Research [10-12] is continually underway to obtain accurate measurements and to predict the absorptance for a wide variety of conditions.

Since all energy incident on a gaseous volume must be accounted for, and since gases do not reflect and have negligible scattering at these wavelengths, equation 3 may be written for a gas as

$$\alpha_\lambda(\lambda,C,T,P,s) + \tau_\lambda(\lambda,C,T,P,s) = 1.0 \ . \tag{8}$$

Equation 8 allows the *spectral transmittance*, $\tau_\lambda(\lambda,C,T,P,s)$, to be defined in terms of the absorptance.

A gaseous volume also emits radiation, and Kirchhoff's law also applies to volumetric emission and absorption:

$$\epsilon_\lambda(\lambda,C,T,P,s) = \alpha_\lambda(\lambda,C,T,P,s) \ . \tag{9}$$

The *spectral emittance*, $\epsilon_\lambda(\lambda,C,T,P,s)$, defined by equation 9, is a nondimensional quantity ranging from zero to unity which depends upon the wavelength of the radiation involved, the species concentrations, the temperature and pressure, and the path length traversed through the gaseous volume. It can also be related to the fundamental quantity of the absorption coefficient in the same manner as is the absorptance in equation 7.

A gaseous volume may also have variations in density that lead to variations in the local index of refraction. Usually the variation in index of refraction is sufficiently small that the change in wavelength and direction of the radiation involved is negligible. However, as the path length increases, the small direction change due to refraction can amount to a large change in relative position of the ray. Research into refractive index variation of atmospheric gases [13-15] usually involves long distances such as that encountered by radiation traversing the entire thickness of the atmosphere.

Similar radiative quantities can be defined for both the solid and gaseous elements. The solid has an absorp*tivity* while the gas has an absorp*tance* for traversing a given path length. More than just terminology is involved here. The absorptivity of a solid is treated as a quantity that is dependent on the surface condition, not interacting with any material below the surface. The absorptance of a volume depends on the depth of the radiating substance. In reality surface radiation is actually a volumetric phenomena also, depending

Review of Radiation Physics   13

on substrate emissions and absorptions for the actual radiation. The general depth of the substrates involved in this radiation exchange is sufficiently small that modeling the radiation as occurring at the surface does not engender serious discrepancies. The discussion of this aspect of radiation is beyond the scope of this work. However, having made the distinction here, the remainder of this dissertation refers to the radiative quantity as absorptivity, with the understanding that when dealing with a volume element the proper term is absorptance, and when dealing with a solid the proper term is absorptivity.

Suspended particles within the gaseous volume can also be a source of interaction with thermal radiation. These suspended particles can be soot [16,17], metals [18], water droplets or ice [19,20], or any number of chemical aerosols [21]. The particles can absorb, emit, or scatter thermal radiation. Property values are a function of the material itself; the wavelength, volumetric density, distribution of particle sizes and shapes, particle temperature, and thickness traversed by the radiation. This dissertation does not focus on scattering by particles but allows for the possibility of including scattering in the radiative analysis method developed.

CHAPTER 3

REVIEW OF RADIATIVE ANALYSIS

This chapter presents the basics of the radiative analysis methods used during the elaboration of this work. Most of the information comes from established radiative analysis procedure but the reverse Monte-Carlo ray trace requires the extension of the reciprocity relations to include volume-to-surface calculations. Additionally, an alternate method, the line-of-sight method, is presented. The details of the actual implementation of the two methods are reserved for a later chapter. The current chapter begins with a description of a formal analysis of the problem.

3.1 Problem Statement

The problem to be solved is the determination of the amount of radiative energy arriving at a specific plane observer location in space integrated over some finite wavelength band, observer area, and viewing solid angle. The temperature, pressure, species concentration, and geometry are known for a large but finite number of points in the region of interest. These points, called *nodes*, can be associated with gaseous volumes, or solid surfaces, or volumes with suspended particles, each of which are referred to as an *element*

and identified by an individual number. The element type determines its composition and arrangement of nodal points. For example, if an element is flagged as a solid surface, it consists of four nodes whose orientation in space defines the boundaries and orientation of the solid surface.

For the moment we assume that all the radiative properties are uniform across an element. Take one element as an emitter. If the element is directly adjacent to the observer location, the direct contribution of spectral, directional intensity is simply the product of the emissivity of that element with the blackbody intensity at the element's temperature. To obtain the integrated energy, this quantity must be integrated over the wavelength interval of interest, over the solid angle subtended by the element from the observer, and finally multiplied by the area of the observer projected along the direction of the radiation.

If the element is not directly adjacent to the observer, the calculations become more difficult. Every intervening volume attenuates and scatters the radiation passing through it. Solid elements obstruct the view between the emitting element and the observer. All of these factors must be incorporated into the derivation of the spectral, directional intensity before integrating to obtain the energy arriving along the direct emission path from the emitting element to the observer.

The emitting element contributes energy to the observer along several indirect paths also. These paths involve radiation which intersects a second element before being redirected to the observer. The redirection takes the form of reflection for solid elements, and scattering or refraction for volume elements. The spectral, directional intensity along this secondary path is a function of the emissive intensity of the emitter in the direction to the intersected redirecting element, the attenuation, scattering or obstruction encountered on the way to this redirecting element, the reflectivity, scattering, or refraction of the redirecting element itself from the incident direction of the emitter to the direction towards the observer, and finally the attenuation, scattering or obstruction encountered on the way to the observer. Again the spectral, directional intensity is integrated across wavelengths to obtain the radiant energy arriving along this secondary route. All possible paths involving the redirection of radiation from the emitter via other elements must be considered.

It is possible for radiation from the emitting element to encounter one element which redirects the radiation towards a second element which redirects the radiation towards the observer. The specular directional intensity arriving at the observer along this tertiary path is similar to the specular directional intensity for a secondary path except now it includes the attenuation, scattering, or obstruction between the first and second redirecting elements, as well as the reflection, scattering, or refraction from the direction of the first redirecting element into the direction of the observer. Again all possible combinations of two redirecting elements must be checked to determine the contribution along these new paths.

This process continues until all the emitted radiation from the emitting element that arrives at the observer, regardless of path, is obtained. The sum of all the radiant energy from all these paths is then the cumulative energy from the emitting element to the observer. At that point a new emitting element is chosen and the whole process is repeated until all the possible emitting elements have been considered. The final answer is the sum of the radiant energies from the individual elements which arrives at the observer. Notice that since the current research is based on the possibility of solid and gaseous radiation, the list of emitting and redirecting elements includes all of the elements.

This calculation outline above is possible if one is sufficiently patient. Practical analysis techniques make several simplifying assumptions to reduce the number of calculations required. For example, assuming that the intervening medium does not participate radiatively, and that the surfaces are specular reflectors leads to calculations involving only direct emission to the observer from the solid elements and any multiple specular reflections in the observer direction. Since only one reflected direction must be considered in the case of specular reflections, the workload is tremendously reduced. The calculations can still become quite extensive even for this highly restrictive set of assumptions. For example, Villeneuve, Chapman, and Mahan [22] present a case where reflections form the dominant means of radiation transfer in an engine cavity with many hundreds of surfaces. Few situations occur in actual practice where such broad assumptions are justified. The only restriction within the Monte-Carlo method on the type of radiative

phenomena modeled is the user's ability to formulate a model of the radiation interacting with matter.

3.2 The Monte-Carlo Method

Halton [23] provides a good historical account of the general Monte-Carlo method and a description of its fundamental theory. Although its use has been traced far back in history, he attributes its current name and rise in popularity to the mathematicians and physicists working at Los Alamos during the late 40's and early 50's, in particular von Neumann, Ulam, Metropolis, Kahn, Fermi and others. Halton also defines the Monte-Carlo method as:

> ...representing the solution of a problem as a parameter of a hypothetical population, and using a random sequence of numbers to construct a sample of the population, from which statistical estimates of the parameter can be obtained.

The problem we wish to solve is defined in the previous section. The parameter which forms our solution is the *distribution factor* from element i to element j, D_{ij}, which represents the fraction of radiation energy emitted from element i that is finally absorbed at element j due to direct emission and to all possible reflections, scattering or refractions. In our case we actually want D_{io} where i ranges over all elements and o stands for the observer element. The hypothetical population is then the final absorption sites of the radiant energy being emitted by each element. Note that all the energy must be assigned a final absorption site.

Radiation reaching the boundaries of the modeled space can either be absorbed by an infinite extension of the atmosphere, by a simulation of the actual background, or by a blackbody at 0 K.

3.2.1 The Monte-Carlo Ray Trace

Now that the problem, the hypothetical population, and the parameter of that population have all been defined, all that remains in our application of the Monte-Carlo method is to construct a sample of that population via a random sequence. The ray-trace portion of the method makes available the tools to construct that sample population. Since the hypothetical population is comprised of the final absorption sites, the idea is to construct the possible paths of the radiation by tracing rays representing the radiation being emitted and recording where each ray ends up.

The basic algorithm starts by dividing the modeled space into a finite number of elements. This step is already accomplished by the CFD grids in the present application. Allow the first element to emit a ray in a direction chosen at random and determine the first element encountered by the ray. The radiative properties $(\alpha'_\lambda, \tau'_\lambda, \rho''_\lambda)$ of the encountered element are then obtained. As explained in Chapter 2 these properties range between zero and unity. A uniformly distributed random number between zero and unity is drawn. If the

absorptivity is greater than or equal to this random number then the ray is absorbed by that element. In this eventuality a counter representing emission by the first element and absorption by this element is incremented. If the random number is greater than the absorptivity then the element type determines what happens next. The ray can be reflected, scattered, refracted, or simply transmitted. For the reflected and scattered cases, a set of random numbers is used in an appropriately weighted directional distribution function to generate the new direction. For the refracted or transmitted case, the physics involved determine a unique direction for the ray. The ray is traced to the next element encountered and the process is repeated. This ray trace continues until the ray is absorbed, or until it crosses a boundary of the modeled space which signifies absorption by the outside environment.

At that point a new ray is emitted from the emitter element and a new ray trace is started. The number of rays traced determines the sample population size, which in turn determines the reliability of the value of the average parameter. The total number of rays represents all the radiant energy emitted by the first element. Therefore the ratio of the number of rays absorbed by any given element, N_{ij}, to the total number of rays, N_i, represents the fraction of energy emitted by the emitting element that is finally absorbed by that given element after participation in a chain of statistically determined radiative events. In other words,

$$D_{ij} \equiv \frac{E_{ij}}{E_i} = \frac{N_{ij}}{N_i} \; . \tag{10}$$

The random numbers used in the Monte-Carlo ray trace must have certain characteristics. As already mentioned, one such characteristic is a uniform distribution between zero and unity. Obviously if the sequence was biased either towards the high or low end it would similarly bias the absorption sites. Another characteristic is that successive random numbers should not be correlated. In other words the sample population can be biased by being constructed using a sequence where a particular random number is always followed by the same second random number. This point applies not only to immediately successive numbers but also to sequences of numbers. If a particular sequence of numbers leads to the creation of a particular ray path, another sequence consisting of the identical numbers generates the identical path, again biasing the population.

For the problem here the radiative transfer in a specified wavelength interval from the elements in modeled space to the observer is desired. Only those rays absorbed by the observer contribute to the calculation of the desired energy transfer. If one assumes for the moment that each absorption site has an equal probability of absorbing a ray then the probability of a ray being absorbed at any element is $1/n$, where n is the total number of absorption sites, including the observer element. Since n is typically on the order of several

million, the probability of a ray being absorbed by a specific element, in our case the observer, is very small. While the assumption of equal probability of absorption does not generally apply, it does illustrate the fact that only a small percentage of the rays traced are absorbed at the observer. Consequently an enormous number of rays must be traced to ensure that a statistically significant number of rays is absorbed by the observer and that the corresponding radiant energy transfer can be reliably estimated. A statistically significant number is defined by having a standard deviation of the cumulative rays' energies fall below some user-defined threshold. A large majority of the rays traced does not contribute directly to the determination of D_{io}.

3.2.2 The Reverse Monte-Carlo Ray Trace

A solution to the problem pointed out in the previous section is to begin the ray trace at the observer element and trace it back into the modeled space. Since the ray trace starts with the observer, this method guarantees that each ray is a member of the population of rays which involves the observer. The distribution factors calculated in this manner represent the fraction of energy that the observer emits that is absorbed by some element in the modeled space, which is the inverse of our desired solution. However, reciprocity relations relate these distribution factors to the distribution factors representing the fraction of energy

emitted by the elements in modeled space that is finally absorbed by the observer. In a sense the reciprocity relations reverse the ray trace.

It is important to establish that the process of reversing the ray trace in this manner is conceptually sound. In a reverse ray trace the observer is the emitting element. Rays leaving the observer are reflected, scattered, refracted, or transmitted, eventually reaching an absorption site. Now, Kirchhoff's law from Chapter 2 establishes that the spectral, directional emissivity is equal to the spectral, directional absorptivity. The absorptivity and emissivity have little intuitive meaning within the context of the Monte-Carlo method which deals in probabilities. Consider then the spectral, directional absorptivity as the probability that a ray corresponding to a given wavelength and incident upon an element in a given direction is absorbed. Similarly the spectral, directional emissivity can be viewed as the probability that a ray is emitted from the same site along the opposite direction for the same wavelength. Using this interpretation, Kirchhoff' law states that these two probabilities are equal. Reversing the path then leads through the same elements traversed to the absorption site. The reciprocity relation for spectral, bidirectional reflectivity means that the probability of the ray being reflected back along the same reverse path is equal to the probability of the ray being reflected along its original path. Similar relations hold for the scattering phase function, refraction, and transmission. Once back at the observer, Kirchhoff's law again implies that the absorption of the reversed ray at the observer is equal in probability to the ray's emission in the first place.

So following the reverse path of the original ray from the absorption site back to the observer is a valid and legitimate representation of the exchange of radiant energy which actually occurs. Establishing the equivalence of the ray's forward and reversed paths does not mean that one is free to incorporate the answer directly from the reverse ray trace, especially since the total number of rays now represents the total energy emitted by the observer. The reciprocity relations which are developed in the next section relate the obtained distribution factors, D_{oi}, to the desired distribution factors, D_{io}.

3.2.3 Reciprocity Relations

The reciprocity relations are a set of equations that relates the distribution factor from one element to a second element to the distribution factor from the second element back to the first element. The appropriate form of the reciprocity relations depends on whether the elements involved are both surfaces, both volumes, or a mixture of one each. The reciprocity relations are derived in the Appendix, while this section presents the final results necessary for the solution of the problem under consideration. These reciprocity relations are based on the assumptions that the index of refraction does not vary significantly from unity for the elements being modeled, and that properties are uniform across an element.

The reciprocity relation between two finite surface elements, i and j is

$$D_{ij}\epsilon_\lambda(\lambda,T_i)A_i = D_{ji}\epsilon_\lambda(\lambda,T_j)A_j \ . \tag{11}$$

The reciprocity relation between a finite surface element i and a finite volume element m is

$$\pi D_{im}\epsilon_\lambda(\lambda,T_i)A_i = D_{mi}\int_{4\pi}\int_{V_m \rightarrow A_{pm}} \epsilon_\lambda(\lambda,C_m,T_m,P_m,s_m)dA_m d\omega_m \ . \tag{12}$$

where A_{pm} is the projected area of the volume V_m along the direction of emission.

The reverse Monte-Carlo ray trace obtains the distribution factors from the observer to the various elements in the modeled space. Since the observer is a surface element, equations 11 and 12 can be used to obtain the distribution factors from the various elements in modeled space to the observer. The energy is given by

$$dE_o(\lambda) = \sum_{i-1}^{N_S} D_{oi} \frac{\epsilon_\lambda(\lambda,T_o)A_o}{\epsilon_\lambda(\lambda,T_i)A_i} \pi\, \epsilon_\lambda(\lambda,T_i) i'_{\lambda b}(\lambda,T_i)A_i d\lambda \ +$$

$$\sum_{m-1}^{N_V} D_{om} \frac{\pi\, \epsilon_\lambda(\lambda,T_o)A_o}{\int_{4\pi}\int_{V_m \rightarrow A_{pm}} \epsilon_\lambda(\lambda,C_m,T_m,P_m,s_m)dA_m d\omega_m}$$

$$\left\{ i'_{\lambda b}(\lambda,T_m)d\lambda \int_{4\pi}\int_{V_m \rightarrow A_{pm}} \epsilon_\lambda(\lambda,C_m,T_m,P_m,s_m)dA_m d\omega_m \right\} , \tag{13}$$

where the first sum is over all the surface elements and the second sum is over the volume elements. Equation 13 can be simplified to

$$dE_o(\lambda) = \sum_{k-1}^{N} \pi D_{ok} \epsilon_\lambda(\lambda,T_o) i'_\lambda(\lambda,T_k) A_o d\lambda \; , \qquad (14)$$

where the sum is now over all elements.

Equation 14 appears deceptively simple since it implies the answer is some weighted combination of blackbody intensities using the temperatures from the modeled space, and is independent of the actual dimensions of the modeled space dimensions. The key is the distribution factors, which correctly attribute the relative probability of radiation from a particular element actually arriving at the observer through the many possible paths. Those distribution factors have been obtained after long and arduous simulation of the radiation process.

The distribution factors obtained by the reverse Monte-Carlo ray trace as outlined above are spectral distribution factors; that is, they are created based on a simulation using band-limited property values. A multispectral analysis must include a recalculation of the distribution factors using an independent ray trace for each spectral interval in the wavelength band of interest. Generally the narrower the spectral sub-band, the more

accurate are the calculated distribution factors due to the rapid spectral variation of the radiation properties in the gaseous components.

3.3 The Line-Of-Sight Method

It soon became apparent that the reverse Monte-Carlo ray trace had such a slow rate of convergence that it would be difficult to test new acceleration techniques in a meaningful manner. The decision was made to develop an alternative code which shared the same calculational procedures but offered a better convergence rate by making use of a set of simplifying assumptions. This alternative code could be then used to test accelerational techniques. The acceleration techniques that work for the new code also work for the reverse Monte-Carlo ray trace, provided that the acceleration techniques are not based on a simplification resulting from the assumptions. Furthermore, the two methods should converge to the same answer when modeling a case where the necessary assumptions for the second code are applicable. This alternate code is based on the *line-of-sight* method.

The line-of-sight method is primarily a deterministic method that calculates the radiative transfer along lines of sight that are generated by a ray trace. All the radiative events encountered by the ray during its traversal of the modeled space contribute to the radiative model. For the general line-of-sight method, a radiative event such as scattering

which scatters the incoming ray into many directions, results in a set of sub-rays being traced in all directions to account for the cumulative energy from the scattering center. In order to avoid the time penalty associated with tracing a large number of sub-rays, an assumption is made that no spectral radiative models are permitted that can result in multiple outcomes. This assumption precludes scattering, diffuse reflection, and bidirectional reflection. Wavelength-dependent refraction is also disallowed by assuming that the radiative phenomena do not lead to different outcome directions as a function of wavelength. This second assumption allows a single ray to represent all wavelengths since the path is no longer wavelength dependent. The remaining models, gaseous absorption, and absorption and specular reflection by solid surfaces, are still functions of wavelength and incident direction. As a result the direction of the path is wavelength independent but its length varies as a function of wavelength.

The method starts by emitting a ray from the observer in a random direction. The ray starts with no radiative energy and a spectral transmissivity of 1.0. Note that in this method the transmissivity is a property attributed to the ray as well as a property of the path traversed by the ray. If the first element encountered is a gaseous volume, its emitted energy and spectral transmissivity are calculated. The emitted energy of the volume element is multiplied by the transmissivity of the ray and then attributed to the ray. The transmissivity of the path traversed through the volume element is multiplied by the transmissivity of the ray to obtain the transmissivity of the ray once it has traversed the volume. If the element

is a solid then its emitted energy and spectral reflectivity are obtained. Recall that since only specular reflections are allowed, the one incident direction uniquely identifies both the reflected direction and the reflectivity. The emitted energy of the solid is multiplied by the transmissivity of the ray and the resulting energy increment is added to the energy of the ray. The reflectivity of the surface is multiplied by the transmissivity of the ray to obtain the transmissivity of the ray after reflection. The ray trace is then continued in the reflected direction. At any point in time, the spectral energy of the ray within wavelength interval $d\lambda$ centered about wavelength λ can be written

$$dE_{o_{ray}}(\lambda) = \sum_{i=1}^{n} \tau_{\lambda,i-1} \epsilon_{\lambda,i} i'_{\lambda b}(\lambda, T_i) A_o \cos\theta \, d\lambda \, d\omega_{ray} , \qquad (15)$$

where the sum is over the number of elements encountered so far, the emissivity, $\epsilon_{\lambda,i}$, is the emissivity of the i^{th} element encountered by the ray, and the ray transmissivity, $\tau_{\lambda,i-1}$, is defined in equation 16. The energy is that amount arriving at the projected area A_o of the observer within the differential solid angle $d\omega_{ray}$ subtended by the ray within some differential wavelength interval $d\lambda$ centered about the wavelength λ of the analysis. The spectral transmissivity of the ray after it has encountered i-1 elements is

$$\tau_{\lambda,i-1} = \prod_{j=0}^{n} \tau_{\lambda_j} , \qquad (16)$$

where $\tau_{\lambda,0}$ is 1.0 and $\tau_{\lambda,j}$ is the transmissivity or reflectivity of the j[th] element encountered by the ray. This ray trace is continued either until the ray exits the modeled space or until the transmissivity back along the ray for each wavelength falls below some prescribed value.

Recall that in the reverse Monte-Carlo ray trace the total number of rays represents the summed energy emitted from a surface. In the line-of-sight method, however, the total number of rays represents the total solid angle viewable by the observer. Consequently, $d\omega_{ray}$ is the total solid angle divided by the number of rays. The more rays that are traced, the smaller the solid angle represented by an individual ray. The smaller the solid angle, the more realistically the corresponding ray represents a true ray since a true ray ideally has no thickness. The summed energy of all the rays is then the energy that arrives at the observer. A separate ray transmissivity is tracked for each spectral interval being analyzed. As mentioned previously, the path of all thermal radiation is wavelength independent under the current assumptions. However, the transmissivities and absorptivity are wavelength dependent and these values can make each wavelength have a different length along that path before the ray's transmissivity at that wavelength becomes essentially zero. The line-of-sight method continues tracing a ray until the spectral transmissivity along the ray is essentially zero for all wavelengths, even if the transmissivity for some wavelengths is already zero.

CHAPTER 4

REVIEW OF CURRENT LITERATURE

This chapter provides an overview of the current literature that applies to the problem being addressed. It is divided into a section dealing with radiative issues and a section dealing with graphical ray traces. The reader interested in a more classical approach to radiation heat transfer analysis is urged to consult Mahan and Eskin [24]. References 25 through 29 provide good examples of applications of Monte-Carlo methods and are not considered further here.

4.1 Radiative Analysis

The literature presented in this section deals with models of radiative phenomena, radiative analysis, and Monte-Carlo ray traces. A section of code capable of calculating the gaseous radiative properties was already available. Nelson [30] modified an existing NASA band model along with the necessary data from the North Atlantic Treaty Organization Infrared Air Target Model to provide a compact model capable of predicting the gaseous spectral absorptivity and transmissivity of a volume element given the volume's pressure,

temperature, species concentrations, path length, and the analysis wavelength. The data available to the model is restricted from 2 to 25 µm in constant increments of five wavenumbers. Wavenumbers can be obtained by dividing 10 000 by the wavelength in micrometers. The constant wavenumber increment translates to an increment of around 0.002 µm at 2 µm and an increment of 0.3 µm at 25 µm. This level of resolution is sufficient for the intended purposes. One additional restriction made by using this band model is that the pressure, temperature and species concentrations are assumed constant across the path length. Because the path length used here is across an individual volume element, that assumption is not too restrictive. Goody [31] addresses the Curtis-Godson approximation for inhomogeneous gases while Lindquist and Simmons [32], and Young [33,34] each have developed their own versions of band models dealing with inhomogeneous path lengths. Hartmann, Levi Di Leon, and Taine [35] describe both line-by-line and narrow-band models in their paper, but unfortunately it was not possible to implement any of these alternative gaseous calculation methods for comparisons with Nelson's due to time limitations.

Various authors have developed bidirectional reflectivity models. Some of the more recent include Neumann and Neumann [36], Sillion et al. [37], Westin, Arvo and Torrance [38], and He et al. [39,40]. A bidirectional reflectivity model should be capable of reproducing the limiting cases of specular and diffuse reflectivities, correctly incorporate reciprocity, and allow adjustment to duplicate experimental measurements of the reflectivity. Another important feature is that the model should be amenable to use in the Monte-Carlo

environment. In the Monte-Carlo ray trace, the incident direction along with a set of random numbers is typically used to generate the reflected directions. So the reflectivity must be converted to a weighting function that correctly maps the uniform random number distribution into the appropriate reflection pattern for a given incident angle. Ward [9] does this with his reflectivity function; however, it is a theoretically derived function and has difficulty capturing all the features present in an actual reflectance pattern.

Numerous papers describe the modeling of radiative transfer problems involving scattering. These include Thynell and Özişik [41], Subramanian and Mengüç [42], and Kamiuto [43]. These scattering models are all derived from theoretical considerations and, though they contain coefficients which can be adjusted, no way is prescribed for measuring the coefficients experimentally.

Models for refractive index of gases in the infrared are extremely scarce. Edlen's [44] paper is about the only one that models the refractive index as a function of the local conditions and includes the infrared band.

References to the Monte-Carlo method and the Monte-Carlo ray trace abound in the current literature. Voytishek and Prigarin [45] present the theoretical proof behind the convergence of various classes of Monte-Carlo problems. Shamsundar, Sparrow and Heinisch [46] develop a modification to the general Monte-Carlo ray trace called energy

partitioning and display the effect that the number of rays has on accuracy. Their application is the analysis of the emission from the solid interior of a cavity that exits through an aperture. Though the partitioning method is not applicable to the current problem, their observation that "the output corresponding to a specific pre-selected number of ray bundles may not always produce a proper representation of the results," is noteworthy. Maltby and Burns [47] describe the application of the Monte-Carlo ray-trace method to an enclosure with non-participating media and make similar observations concerning the convergence rate.

Niederreiter [48] presents the theory behind the use of quasi-random number sequences in a Monte-Carlo solution. Quasi-random number sequences emphasis the uniform distribution of the sequence across its range above the absence of correlation between the values of successive subsets of the sequence. The uniform distribution can lead to improved convergence in problems where the correlation factor is not a concern. O'Brien [49] presents an accelerated version of a Monte-Carlo analysis using a quasi-random sequence to analyze the scattering of an incident beam of light. His acceleration consists of using the quasi-random sequence instead of a random sequence, and of making a simplifying assumption concerning the nature of scattering beyond a certain number of scattering events. The simplifying assumption is that after several scattering events the radiation no longer has a significant directional component and so is essentially diffuse ambient radiation.

4.2 Graphical Ray Tracing

This section deals with articles concerning *graphical ray tracing*. Graphical ray tracing refers to a ray trace that is performed for visual display purposes. One set of papers deals with multiprocessor ray tracing while the other covers methods used to improve and accelerate the ray trace based on single processor systems. For a multiprocessor environment effectively two ways exist to divide up the problem. Each processor can handle individual groups of rays since the path of any ray is independent of the path followed by others rays. Woodwark [50] describes such a setup where the number of processors available is less than the number of groups of rays. In order to ensure that each processor has the same overall workload each bundle of rays is assigned on a random basis to the available processors. Cleary et al. [51] describe the second method of problem division where each processor models a fixed volume of the modeled space. In this approach each processor contains a description of objects within its modeled volume against which it performs the ray trace. Once the ray leaves the processor's allotted space, a message is sent to the processor containing the adjoining volume to continue the ray trace. The first approach to multiprocessor analysis depends on each processor being able to maintain the geometry of the entire problem in local memory and does not require any communication between the processors other than at the start and finish. The second approach requires only sufficient local memory for the geometry of a small section of the complete problem, but

requires extensive communication between the different processors. Gaudet et al. [52] demonstrate the application of the multiprocessor scheme to a set of processors specifically designed for ray tracing calculations.

Amdahl [53] derives a theoretical limit on the speedup possible from parallel processing based on the amount of serial and parallel tasks in a code; however, Lukkien [54] states that Amdahl's limit does not take into account the benefits gained from distributed memory systems. Based on these two accounts a limiting factor exists to the acceleration gained by adding more processors but it is not clear that the limiting factor is consistent across all systems or for all problem types.

Graphical ray traces differ from the method developed in this dissertation due to the fact that graphical ray traces are based on visible light and so do not model participating media and actually have only a limited number of sources of radiation. In addition, the principal objective of the graphical ray trace is to create an image that looks realistic to the human eye when viewed from a computer screen. That is why Lee et al. [55] report using ray traces consisting typically of 2 to 256 rays while Shamsundar, Sparrow and Heinisch [46] report using 270 000 rays for a relatively simple radiative problem. Thus it is appropriate to keep in mind that the acceleration techniques applicable in the graphical ray trace arena might not be appropriate in a radiative analysis.

A good citation to start with is İşler and Özgüç's [56] survey paper on current acceleration techniques. One of the most popular techniques involves spatial subdivision [57-63]. The main idea is to subdivide the domain of interest, keeping track of what scene objects are present in which subdivision. Any subdivision that contains more than one object is recursively subdivided until each subdivision only contains one scene object. If the volume is subdivided into two equal subdivisions it is called a *binary space partitioning* subdivision. If the volume is subdivided into eight equal subvolumes it is called *oct-tree* subdivision. Tracing rays through the subdivisions is relatively simple and fast, and the more complex intersection calculation between a ray and a scene object is carried out only when a ray enters a subdivision containing a scene object. A further improvement is gained at the object intersection calculation level because the ray is tested only against one object at a time. The subdivision intersection is so efficient because the orientation and spacing of the subdivisions is known ahead of time. Due to this feature an algorithm called the graphical line generator can be extended to three dimensions to determine which subdivisions are intersected by the ray [57]. The line generator is used to determine the integer pixels which represent a given geometric line. The idea of a line generator is elaborated in Chapter 6.

It is possible to combine both spatial division and parallel processing to profit from the benefits of both. Pitot [64] and McNeill et al. [65] both incorporate a ray trace using

spatial division in a parallel processing environment. McNeill et al. also study the effects of using a uniform spatial subdivision, a binary space partitioning subdivision, or an oct-tree subdivision. They arrive at the conclusion that all three have similar timing requirements but that the binary subdivision and the oct-tree subdivision require less memory to describe the same geometry data. In addition they find that the oct-tree requires fewer traverses since the same information is available through fewer nodes.

Akimoto, Mase, and Suenaga [66] establish another acceleration technique based on subsampling the image area. Instead of starting with the ray trace using the same pixel resolution as the final image, they do a ray trace with a coarser resolution. By analyzing which of these coarse pixels show variations from their neighbors, they can infer which areas might contain an object. Those areas are traced using a slightly less coarse pixel set, and again variation between neighboring pixels suggests a need for more pixels. Finally the limit of resolution of the final image is reached, but that level of resolution in the areas has only been applied where it is needed. All the other areas are represented using fewer pixels and so consume less time. The potential drawback is that the coarse set of pixels may miss elements in the modeled space and so fail to resolve that something is there. The method essentially trades time for image accuracy. The acceleration method would result in the same time usage for scenes that have a high level of complexity across the entire scene.

The choice of initial ray direction is an important issue in the literature of graphical ray tracing. Regular spacing tends to introduce aliasing artifacts into the final image because the modeled space is sampled only from a fixed number of observer positions. This problem can be avoided by randomly distributing the location of the starting point for the ray within the pixel area. This process is termed jittering and is described by Cook [67], Mitchell [68], and Kirk and Avro [69]. Shirley [70], on the other hand, advocates using a quasi-random distribution since it ensures that the rays are not clumped together. Since the modeled path length is independent of the distribution that produces the initial directions of the rays, it is possible to incorporate a quasi-random method to calculate the initial ray directions. This issue is addressed in Chapter 6.

The ray tracing papers have some interesting ideas. Shinya, Takahasi, and Naito [71] propose the idea of pencil tracing. Instead of tracing rays with infinitesimal thicknesses, a bundle or rays is traced with a central ray surrounded by other rays. The potential acceleration comes from the fact that more than one ray is being traced but the intersection cost is the same as one ray. The paper describes transmittance, refraction, and specular reflection of a pencil but does not discuss the potential for scattering. Once again graphical ray tracing does not involve participating media so the only contribution comes from solid objects. In this sense dispersive scattering of the pencil is reduced, and the concept does not apply to environments where scattering is important.

Another interesting idea is presented by Glassner [72] when he describes using ray tracing for animation purposes. The idea of adding motion to the scene complicates the ray trace since now the time-of-flight of a ray must be recorded in order to determine the appropriate location of objects in the modeled space. This idea could be incorporated in a radiative analysis. If a CFD analysis was performed across some finite time, objects in the domain of interest would remain fixed relative to the grids, however flow properties would change with time. A radiative ray trace could start in the first CFD solution and keep track of the actual time difference as it progressed. Obviously the time would be derived from the space crossed and the velocity of the radiation, not the time elapsed during the ray trace analysis. Once sufficient time had elapsed, the analysis could update the CFD solution to a later solution, and the ray trace could continue. Since the length scale involved in most analysis problems is small compared to the speed of light, little will change as the radiation traverses the scene. This ability to model time is important however, since most infrared imaging devices have some processing-time delays that affect the measured image. It is the measuring instrument time-delay that is being modeled here.

Dias [73] and Heckbert [74] describe simple algorithms for handling ray tracing involving semi-transparent material with refractive indices. Dias goes as far as describing the conditions necessary to produce interference patterns and models them. Her method, however, only traces the rays to the interface. A numerical model is applied after the interface that models the ultimate outcome of multiple thin-surface reflections. The

mechanics of combining a numerical solution to a statistical simulation should be considered further.

Finally, Sweeney and Bartels [75] discuss the fundamentals of ray tracing B-spline curves. One issue that is a potential problem is that the CFD grids might be sufficiently smooth for the aerodynamic calculations while lacking the level of resolution necessary to accurately represent the surface during the radiative ray trace. Large surface elements lead to regions where the surface normal remains uniform. This lack of variation leads to large reflection values in a given direction where in reality very little of the actual surface is appropriately oriented to reflect energy in that direction. One option is to use the CFD grids for the gaseous part of the ray trace and a B-spline surface representation for the solids. The B-spline representation offers a smoother variation of surface normal with position instead of trying to model curved surfaces by a set of flat planes, as is done in CFD grids. One key point is that Sweeney and Bartels' method first intersects the ray with a set of bounding points. These bounding points serve as a rough representation of the B-spline surface and allow an approximate intersection location of the ray with the B-spline surface to be obtained. If B-spline representations of surfaces are incorporated into the radiative model then this first intersection step will already be completed since the ray will have intersected the CFD grid representation of the surface.

CHAPTER 5

PRESENTATION OF TEST CASES

In this chapter the test cases used as a baseline for the code are presented along with a description of the infrared imaging system, CFD solution format, and the computer systems used for the research.

5.1 The Auxiliary Power Unit

Exhaust plume data obtained from a specially modified auxiliary power unit are used most frequently throughout this research. They represent a simple case, containing only gaseous emissions, and as a result offer an analysis that can be completed in relatively little time.

5.1.1 General Description

A Federal Aviation Authority category II, class B gas turbine auxiliary power unit was used to create a vertical, heated exhaust plume for study during the current research.

The auxiliary power unit is an Airesearch Model GTCG30-142C manufactured by the Garrett Auxiliary Power Division of the Allied Signal Aerospace Division [76]. The auxiliary power unit has been modified by Hardman [77] by attaching a flow-turning and straightening assembly to the nozzle exit. This assembly turns the exit flow from horizontal to vertical and has screens to remove most large scale eddies from the flow before reaching the 11.43-cm diameter circular nozzle. The nominal exit conditions are 681 K and 1.044 atm.

Nelson [78] conducted a temperature, pressure, and infrared imaging survey of the exhaust plume generated by the auxiliary power unit. Figure 1 shows the experimental setup. In the foreground and to the left is the mechanical traverse used to position the combination temperature/pressure probe in the flow. To the right is part of the flow straightener and nozzle of the auxiliary power unit. In the background is a sheet of aluminized mylar set at a 45° angle to the horizontal. This sheet is used to reflect the cold night sky directly overhead to the detector to provide an apparent cold and uniform thermal background for the infrared imaging of the plume. The infrared imaging system, which is not shown in the figure, is positioned immediately to the right of the camera viewpoint.

The experiment was conducted at night to reduce wind disturbances and to eliminate thermal radiation contamination due to sunlight. The auxiliary power unit was run continuously over the course of the measurements to avoid any fluctuations caused by
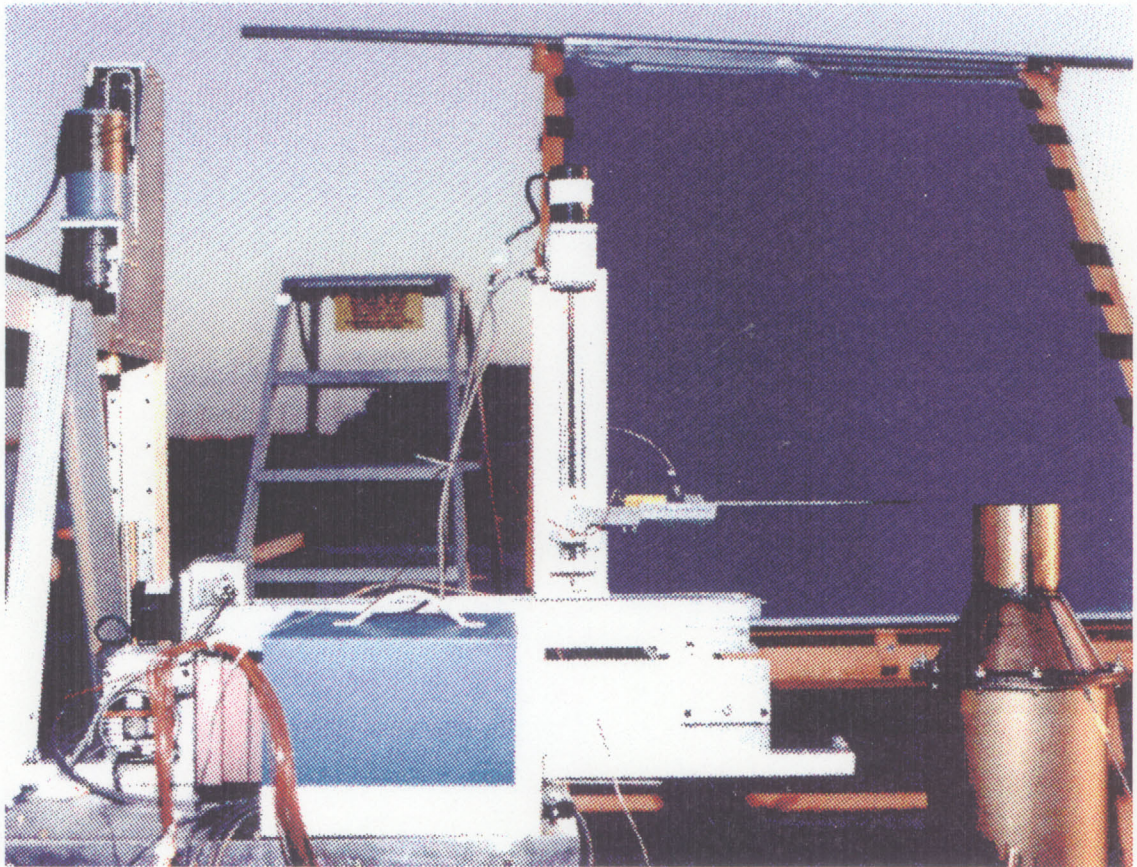
Figure 1   Experimental Setup for the Auxiliary Power Unit Exhaust Plume Surveys.

starting and stopping. The temperature and pressure measurements are obtained at a series of axial stations in planes normal to the nozzle axis located at various distances downstream. The spacing of the data points was arranged to provide a set of measurements in the core, across the mixing layer, and exterior to the plume. In addition to these measurements the local atmospheric temperature, pressure, wind velocity and relative humidity were recorded every half-hour. The error bound on the temperature meaurements is $\pm(0.3 + 0.001\ T)$ K where T is the measured temperature in kelvins, while the error bound on pressure measurements is ± 0.001 atm.

5.1.2 Infrared Data

Two sets of infrared images exist: a set covering wavelengths 2.0 to 5.0 μm, and a filtered set covering the wavelength interval from 4.0 to 4.8 μm. Both sets were taken with a 20° field-of-view lens at a horizontal distance of 2.23 m and a height of 0.33 m above the exhaust exit plane. Figure 2 shows three consecutive images of the auxiliary power unit plume in the 4.0 to 4.8 μm range. The exit plane of the auxiliary power unit is located in the bottom center of the image, and the plume flows upwards towards the top of the image. Each image consists of 140 pixels horizontally by 70 pixels vertically and the time required to scan from the upper left corner to the lower right corner is 0.04 s. The color scale used is based on the wavelength-integrated intensity, with the minimum value corresponding to

40.7

20.5

0.2

W/m² · sr

Time
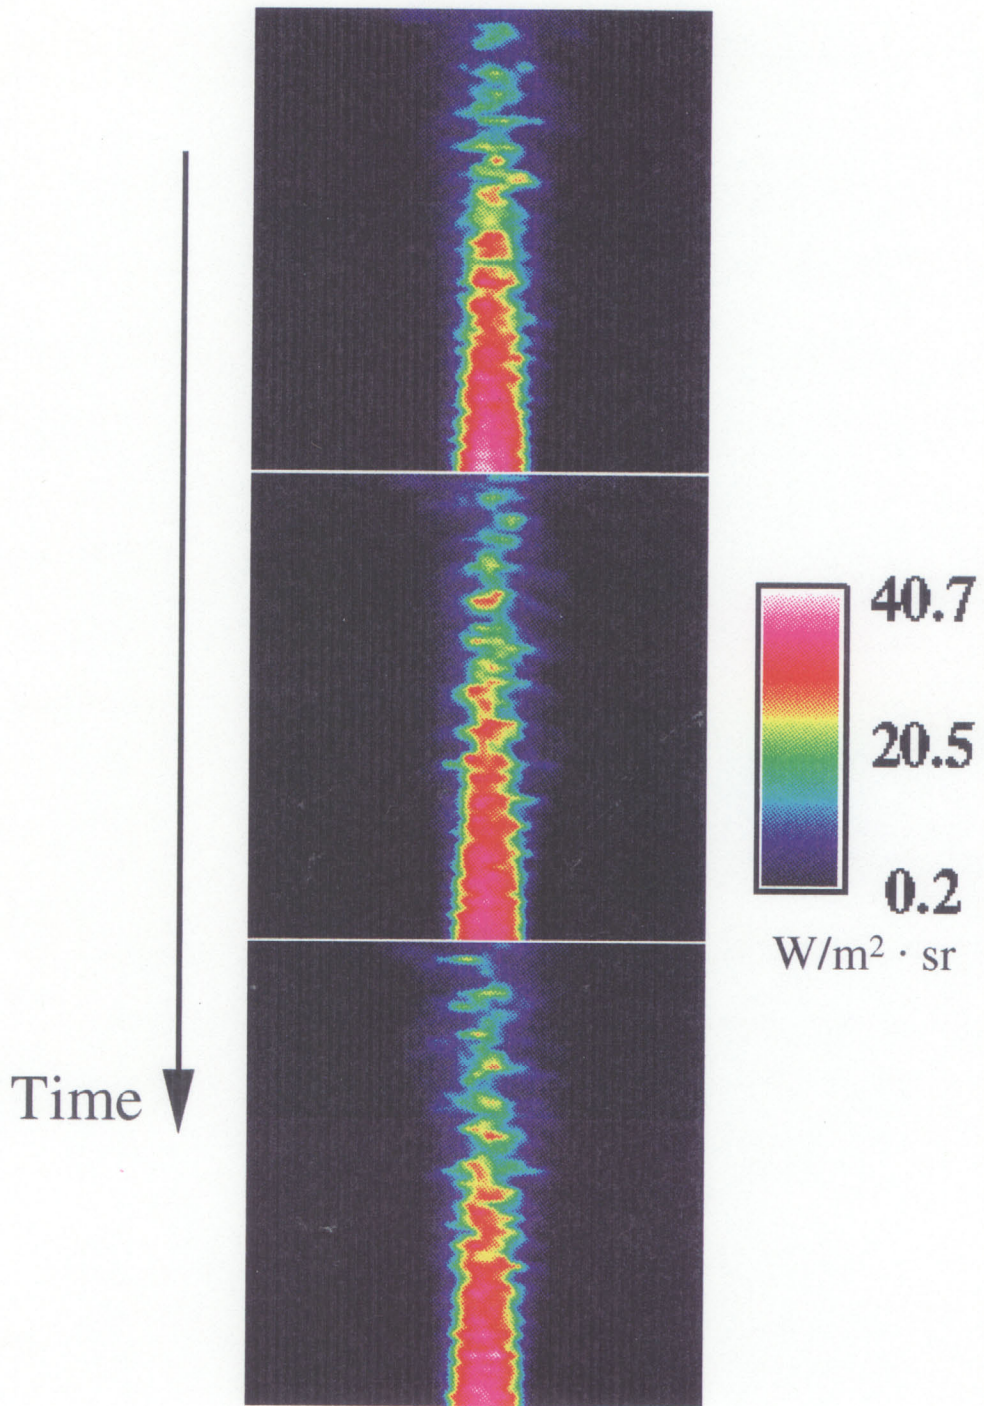
Figure 2   Instantaneous Filtered Mid Wavelength Infrared Images of the Auxiliary Power Unit Exhaust Plume.

Presentation of Test Cases   47

black and the maximum value corresponding to white. From these images it is apparent that the plume experiences large-scale turbulent fluctuations and, given the time scale of these images, these fluctuations are very rapid. A series of such images can be averaged together to obtain the average radiant energy from the plume. This is done in figure 3, which shows an image of the auxiliary power unit's exhaust plume between 4.0 and 4.8 µm. This image is averaged over 44 images representing a total time of about two seconds. In figure 3 the image now has 140 pixels horizontally by 280 pixels vertically. Figure 3 is closer to the classical image of a plume but the combination of scanning method and plume turbulence contributes to produce jagged edges along the boundary of the plume. Figure 4 is an image in the 2.0 to 5.0 µm band with 140 pixels horizontally and 280 pixels vertically which has been averaged over 2000 individual images. It represents a time sequence of about 1.3 min. Clearly the more images that are averaged together the smoother the average plume appears. The 4.0 to 4.8 µm band images that are used the most in the current research contain a set of 44 consecutive images.

Figure 5 is an unaveraged infrared image of the reflective backdrop covering the 2.0-to-5.0-µm wavelength range with 140 pixels horizontally by 280 pixels vertically. The color scale is based on the intensity integrated over the appropriate wavelength interval. The range of the scale is not as broad as the previous images. The image displays the effectiveness of the reflective backdrop for providing a cool, uniform background for the

0.2 20.4 40.7

W/m$^2 \cdot$ sr
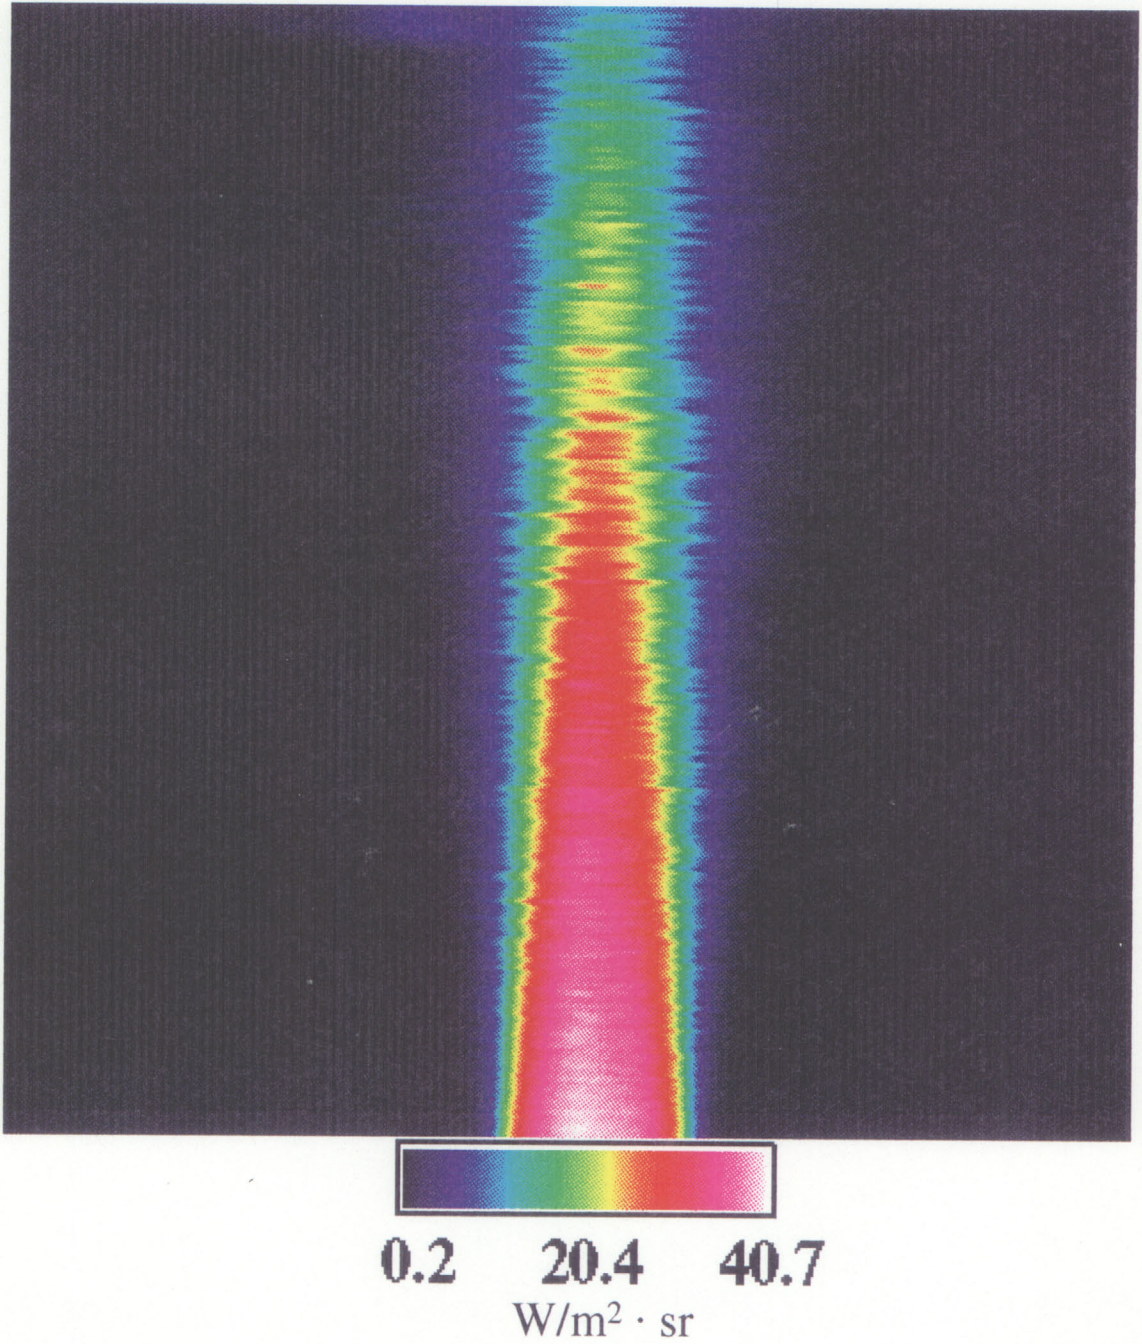
Figure 3 Averaged Filtered Mid Wavelength Infrared Image of the Auxiliary Power Unit Exhaust Plume.
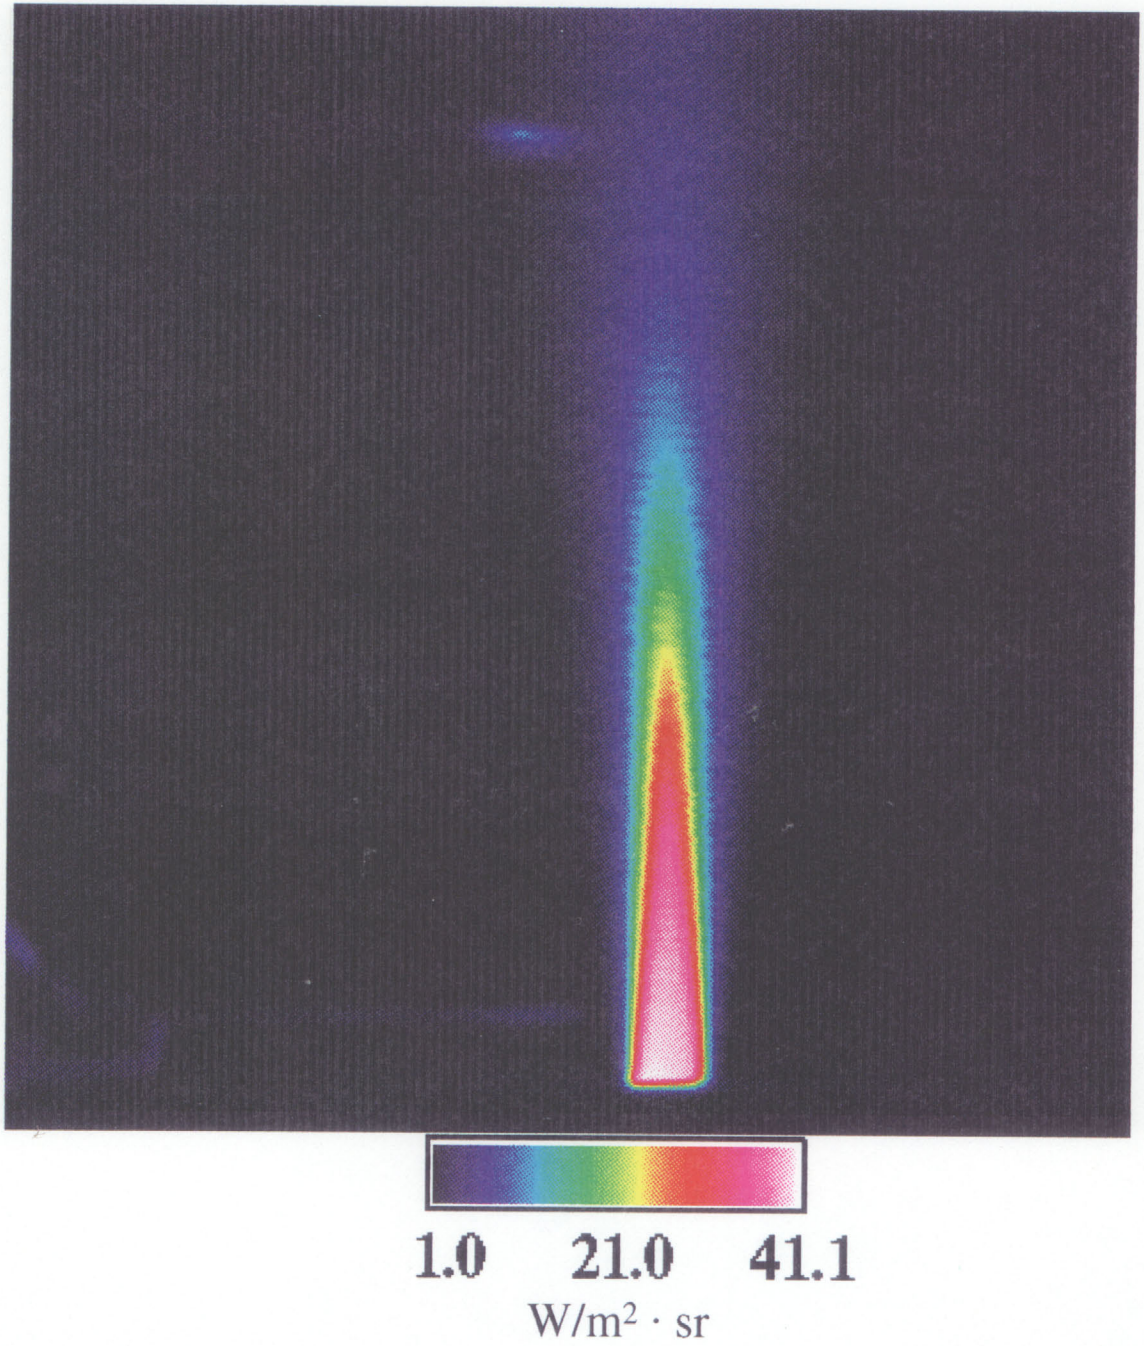
Figure 4   Averaged Mid Wavelength Infrared Image of the Auxiliary Power Unit Exhaust Plume.

9.1   10.6   12.1
$W/m^2 \cdot sr$
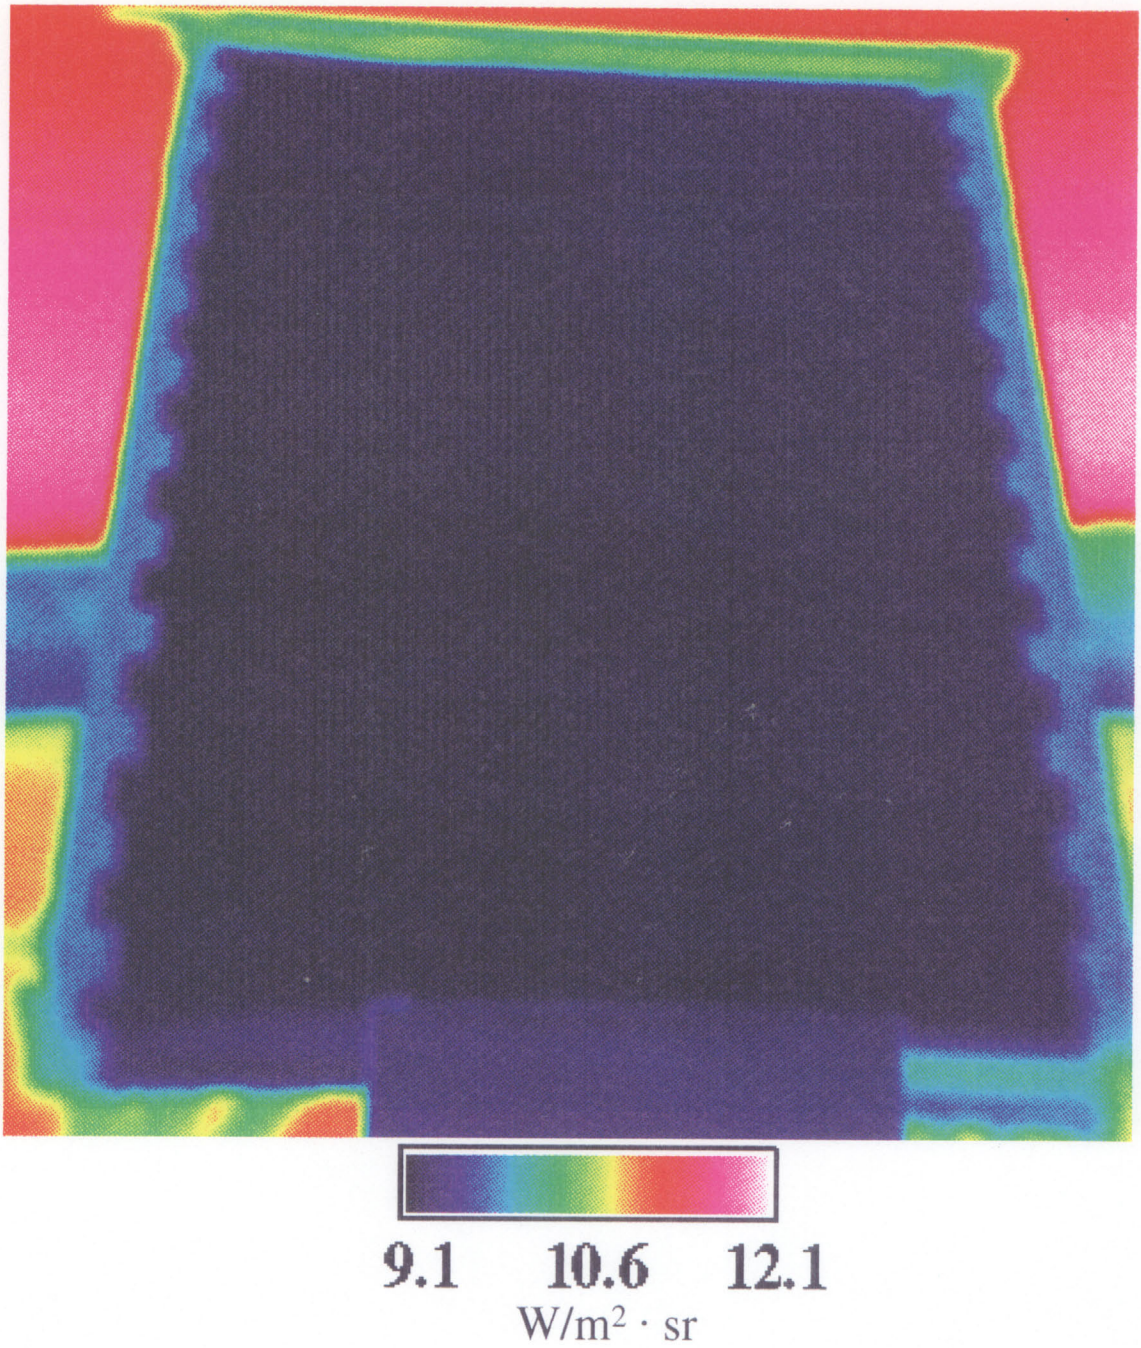
Figure 5   Mid Wavelength Infrared Image of the Reflective Backdrop.

Presentation of Test Cases   51

plume. This uniform background contrasts sharply with the thermal gradients present in the nighttime sky along the horizon, visible along the left and right-hand borders of the image.

5.1.3 CFD Representation

No CFD models of the auxiliary power unit's exhaust plume were available at the time of this research. However, the temperature and pressure measurement data described in Subsection 5.1.1 could be used to simulate the results from a CFD analysis. By using the experimental data directly, any question concerning discrepancies between the measured infrared image and the computed infrared image would then be attributed to the radiative analysis and not to the CFD methodology. The plane of measurement positions, along with some additional points located outside the plume to represent the ambient conditions, was rotated about the plume centerline to create a grid representation of the plume. The grid consists of 89 radial positions by 25 circumferencial positions by 91 downstream positions for a total of 202 475 nodes. Further measurements by Nelson show that the plume is fairly axisymmetric, thereby allowing the temperature and pressure measurements from the measurement plane to be rotated about the plume axis. The one remaining quantity needed to carry out a radiative analysis is the species concentrations throughout the plume. The resources necessary to measure the species concentration were not available during the time of the experiment, so Crocco's [79] analogy is used to relate species distributions within the

plume to the measured temperatures by assuming that the species concentration distributions are similar to the total temperature distribution. Crocco's analogy is valid as long as the turbulent Lewis number, a dimensionless number representing the ratio of the turbulent mass transport length to the turbulent thermal energy transport length, is near unity. Although it is difficult to reliably estimate this parameter for the experiment at hand, it is reasonable to assume that Crocco's analogy holds. The extent to which the subsequent computational results match the experimental measurements will provide an indication of the appropriateness of this assumption. Crocco's analogy provides a way of relating the species distribution between two sets of known concentrations at two known temperatures to the corresponding total temperature distribution. The ambient concentration of species is available and provides the concentration at ambient temperatures. The only relevant data available concerning species concentration in a combustion plume come from Dinger et al. [80], who report on the species concentration for the exhaust flow from a JT9D-7 Pratt and Whitney engine. The concentration of the major species, along with the atmospheric concentrations are listed in Table 1. At present these concentrations are taken as the best estimate available.

5.2 The Stratospheric Observatory For Infrared Astronomy

The Stratospheric Observatory For Infrared Astronomy (SOFIA) is a NASA program that flies a large infrared telescope aboard a modified Boeing 747 to make infrared astronomy measurements above a large portion of the earth's attenuating atmosphere. Recently a study done by Dinger et al. [80] was completed to determine the best location for the telescope aboard the aircraft. A telescope positioned towards the aft section of the aircraft has several benefits but risks potential contamination of the infrared signal by the engines and exhaust plumes. Inflight infrared measurements of an aircraft similar to the SOFIA aircraft, the NASA Shuttle Carrier Aircraft 905, were obtained in order to establish the extent of the radiation from the engines and exhaust plumes.

This data set is not as complete as the auxiliary power unit data set but it does have an associated CFD solution, and it represents the typical complexity of the CFD grids, along with a situation where both solid and gaseous radiation features are present.

5.2.1 General Description

In order to evaluate the potential contamination from the engines, the infrared imaging system was flown aboard a Lear Jet flying in formation with the Shuttle Carrier,

Table 1   Species Concentrations.

Constituent	CO_2	H_2O	N_2	O_2	Other	Total
Ambient Mole Fraction	0.000	0.009	0.781	0.209	0.001	1.000
Core Mole Fraction	0.031	0.030	0.769	0.161	0.009	1.000

a modified Boeing 747-100. Images were taken as the Shuttle Carrier flew at an altitude of 35 000 feet with an approximate 5° angle-of-attack at a Mach number of 0.6. The Shuttle Carrier was used since its wing and engines match those of the SOFIA Boeing 747SP. Figure 6 shows the Shuttle Carrier during a routine maintenance flight.

5.2.2 Infrared Data

Figure 7 is an infrared image in the 2.0 to 5.0 μm band obtained with the 20° field-of-view lens of the Shuttle Carrier during flight. The image is 140 pixels horizontally by 70 pixels vertically, and uses a color scale based on arbitrary units. Black on the scale corresponds to small values of radiant energy while white corresponds to large values of radiant energy. The scale has been modified to enhance the visibility of the fuselage and solid surfaces against the background. It is difficult to assign reliable intensities to the image since the infrared scanning system had not been calibrated when the data were obtained. Since the normal aircraft window is opaque in the infrared, it was replaced with a calcium fluoride window. Distances between the Lear Jet and the Shuttle Carrier were not recorded. The lack of actual infrared values means than these data can be compared qualitatively but not quantitatively.

Figure 6   The NASA Shuttle Carrier Aircraft 905, a Modified Boeing 747-100.

0.0   0.5   1.0
Arbitrary
Units
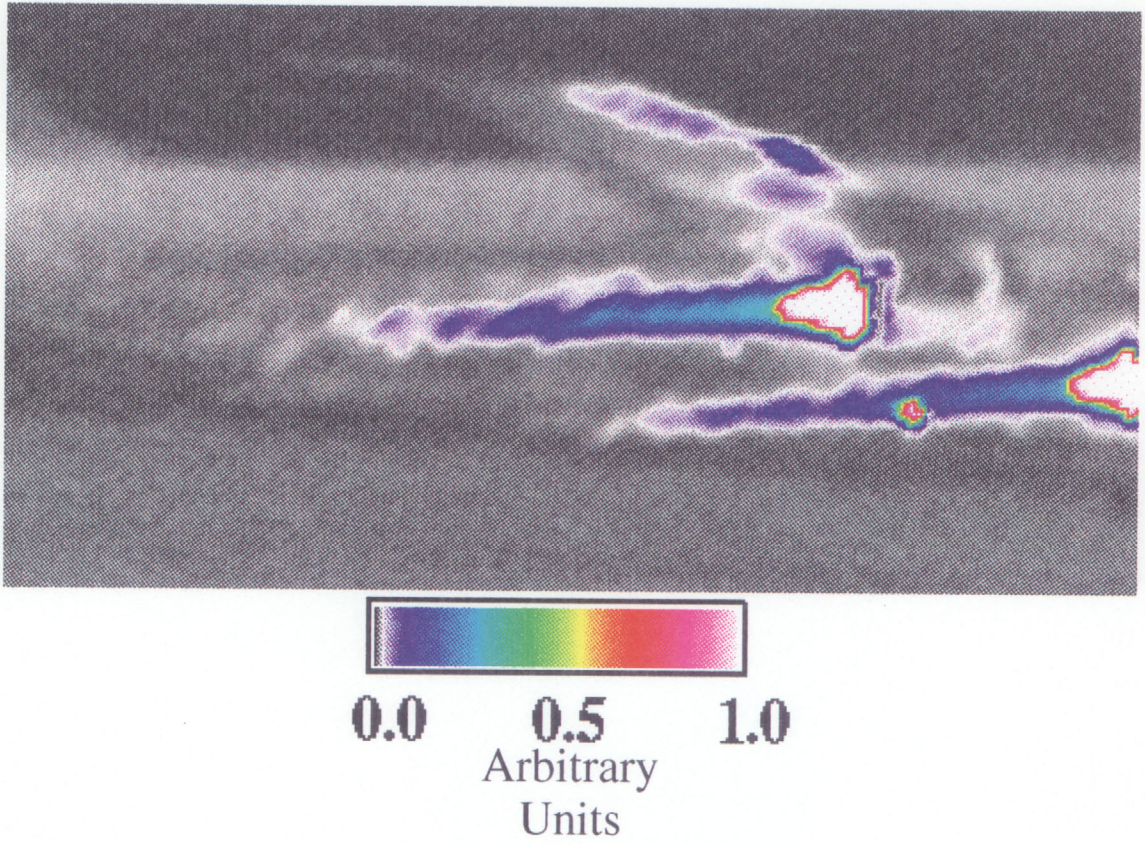
Figure 7   Mid Wavelength Infrared Image of the Shuttle Carrier.

5.2.3 CFD Representation

Atwood [81], and Atwood and Van Dalsem [82] developed a CFD model of the SOFIA 747 for comparison with flowfield measurements from wind tunnels. This model is comprised of 13 separate grids, with a total of 869 914 nodal points. It is a half-model, representing only the right-hand side of the aircraft, and lacks a tail since it is used for comparison with wind tunnel models. Figure 8 shows some grids used in the computational fluid dynamics calculations. The solid elements are shaded in a solid color with different colors corresponding to different grids. The CFD solution is based on the same flight conditions as those listed for the SOFIA test. The solution methodology uses an algebraic turbulence model that uses a modified Baldwin-Lomax model at the fuselage wall, and a Prandtl mixing layer model for shear layer regions such as the plume boundaries and the wing wake.

5.3 Infrared Imaging System

The infrared imaging system used during the current study is the Agema 880 dual waveband infrared scanning system [83]. The infrared scanning system consists of a data acquisition computer and two cryogenically-cooled scanners; one sensitive in the mid wavelength range (from 2.0 to 5.0 µm) and the second sensitive in the long wavelength
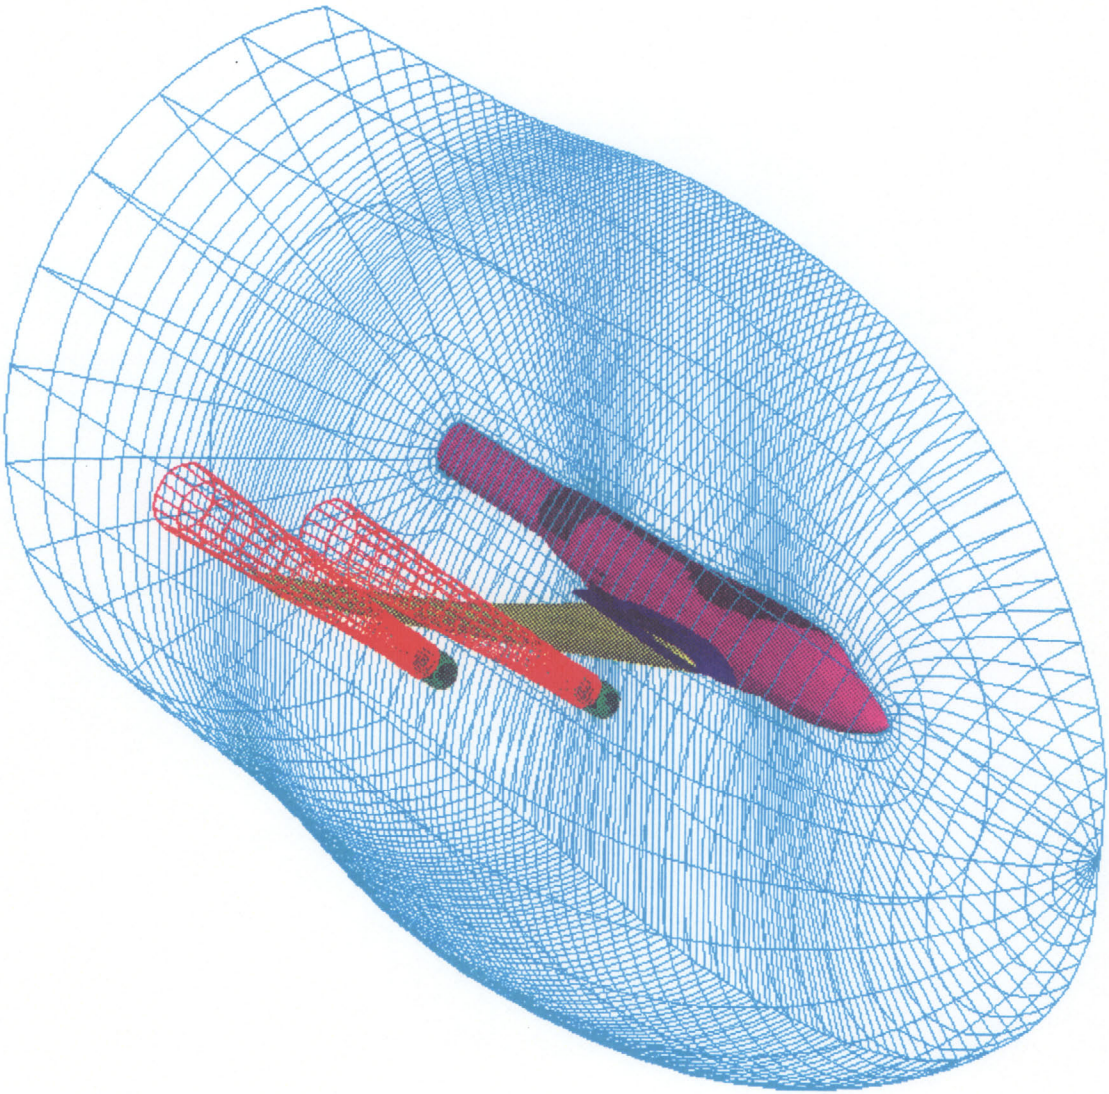
Figure 8   CFD Grid Representation of the Boeing 747.

Presentation of Test Cases   60

range (from 8.0 to 12.0 µm). The mid wavelength scanner uses an Indium Antimonide detector and has a sensitivity of 0.1°C at 30°C with a system response given in figure 9. The long wavelength scanner uses a Mercury Cadmium Telluride detector with a sensitivity of 0.05°C at 30°C with a system response shown in figure 10. A filter is available for the mid wavelength scanner which attenuates radiation outside the 4.0-to-4.8-µm band. With the filter in place the mid wavelength scanner has the system response given in figure 11. A set of 20° field-of-view and 7° field-of-view lenses are available for both scanners. Inside each scanner is a set of heated reference blackbodies and internal temperature sensors that allow the drift and changes in gain to be accounted for as the system is being used and help achieve the high level of accuracy claimed by the manufacturer.

Since a single detector is used in each scanner, each has a set of computer-controlled horizontal and vertical scanning mirrors. These mirrors scan the image presented by the lens and reflect portions to the detector. Each portion of the image that is reflected onto the detector in this manner is considered an individual pixel of the image. The image is divided up into a set of 280 horizontal lines each containing 140 pixels. The scanning sequence starts by scanning the first line, dropping to scan the fifth line, and then the ninth line, and so on until the end of the image space is reached. The scanner then returns to scan the second line, the sixth line, the tenth line, and so on. On the fifth scan though the image space, the scanner begins again with the first line. The reason behind this scan pattern is that it allows the user to obtain four images of 140 pixels horizontally by 70 pixels vertically
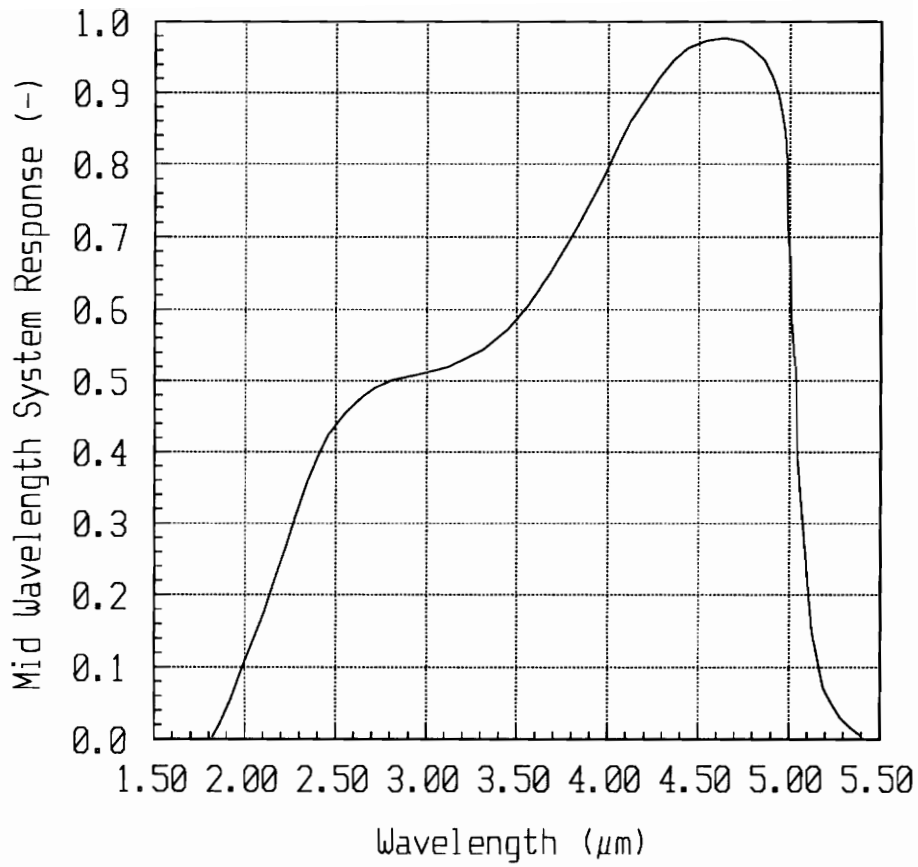
Figure 9   Mid Wavelength Scanner Response Curve.
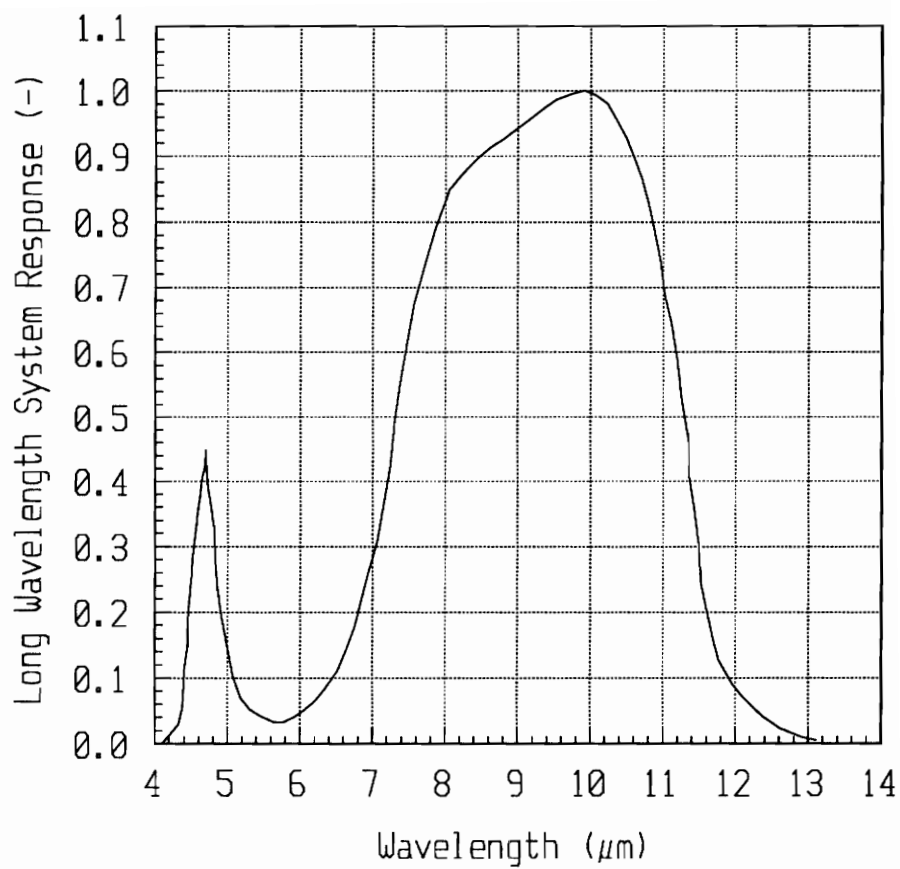
Figure 10   Long Wavelength Scanner Response Curve.

Presentation of Test Cases   63
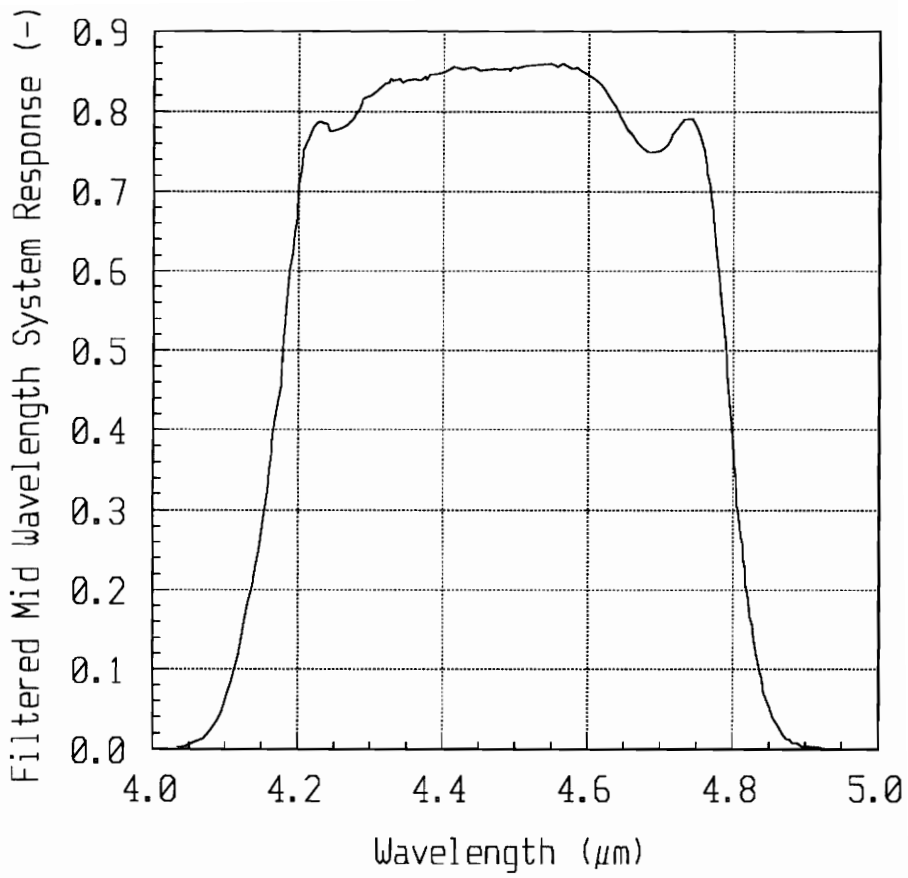
Figure 11   Filtered Mid Wavelength Scanner Response Curve.

which approximately cover the entire field of view in a short time. Each of these images takes about 0.04 s to complete, offering images with excellent time resolution but poor spatial resolution. The four images can be interlaced to provide an overall image of 140 pixels horizontally by 280 pixels vertically to improve the spatial resolution but at the cost of the temporal resolution since each successive horizontal line is delayed by 0.04 s with respect to its neighboring lines. The interlaced image then represents a time difference of 0.16 s between the start of the measurement and its end and requires 80 kilobytes of memory.

The thermal radiation incident on the detector produces a proportional voltage. The scanner converts this voltage to a 12-bit integer that is recorded by the data acquisition unit. This integer is subsequently convert a detector signal during processing. This analog detector signal can be related to the equivalent blackbody temperature according to

$$T = \frac{B}{\ln\left(\dfrac{R}{DS} + F\right)} , \qquad (17)$$

where R, B, and F are calibration constants, DS is the detector signal, and T is the equivalent blackbody temperature. The calibration constants are obtained by imaging a blackbody at a variety of known temperatures and recording the accompanying detector signal. A curve-fit program then calculates the appropriate calibration values. The system was calibrated by

Agema just before the auxiliary power unit test and again afterwards. An example of the calibration curve for the filtered mid wavelength scanner is presented in figure 12. The horizontal axis is proportional to the nondimensional detector signal from the system, while the vertical axis represents the temperature of the blackbody from the calibration test. The detector produces a signal proportional to 800 units when the radiation emitted by a blackbody source at 460 K and within the wavelength sensitivity band of the scanner strikes the detector. The same amount of radiant energy across the same wavelength band from any object will elicit the same detector signal.

Note that equation 17 provides an equivalent blackbody temperature and not the energy incident on the detector for a given detector signal. To obtain the value of incidentradiant energy instead of the object's apparent temperature, Planck's blackbody radiation distribution function, equation 1, can be integrated over the appropriate wavelength range. A word of caution: the instantaneous detector signal is due to all the energy incident at that moment. Most of this energy comes from the object being imaged, but some can also come from background or surrounding objects through scattering or reflection.
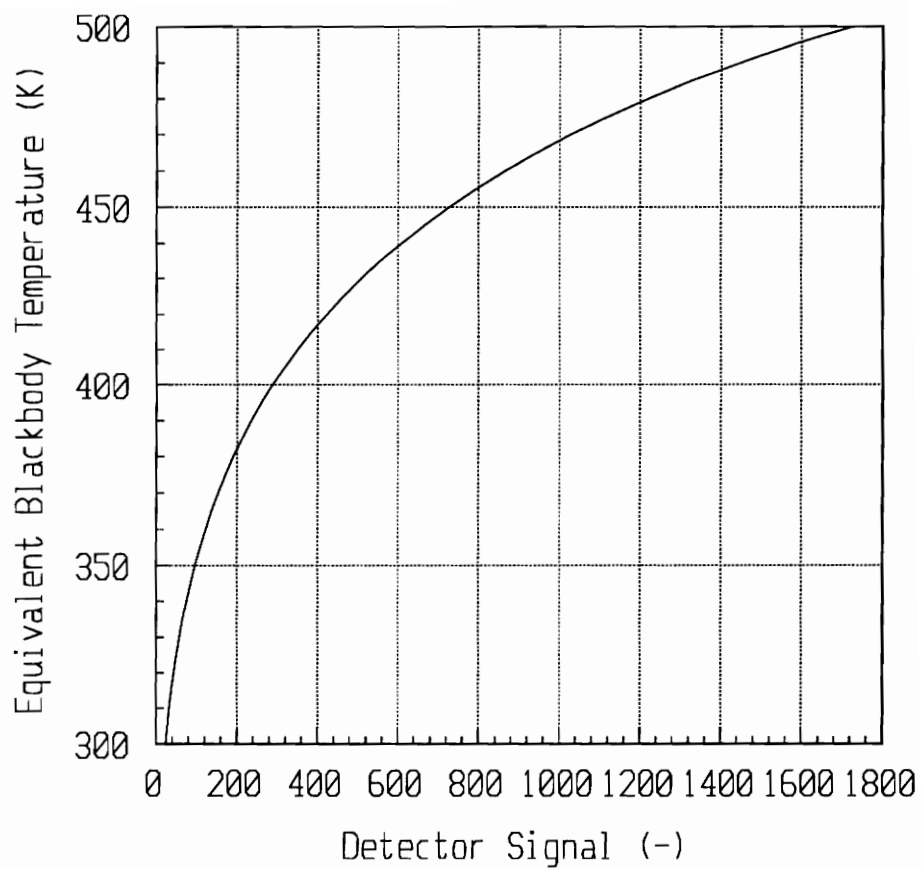
Figure 12 Filtered Mid Wavelength Scanner Calibration Curve.

5.4 Computational Fluid Dynamics Format

The basic input into the radiative analysis code is the Computational Fluid Dynamics (CFD) solutions and grids. These define both the geometry and values used in deriving the radiative properties of the elements. Although many CFD grid types exist, the format chosen for the current research centers on the multiple overlapping structured grids that are popular at NASA Ames Research Center. A structured grid contains nodal points that are numbered using a set of ordered *triplets*. A triplet is a set of three integer values used to label nodal points. The numbering of the triplets indicates how the points are related. For example the point identified by the triplet (i,j,k) has as its surrounding neighbors the points identified as (i-1,j,k), (i+1,j,k), (i,j-1,k), (i,j+1,k), (i,j,k-1), and (i,j,k+1). The only two restrictions placed upon the triplets are that each component of the triplet must have a consistent range throughout the grid, and the ordering of the triplets must follow a right-hand rule.

The first restriction specifies that if there are 20 i-nodes identified at one point in the grid, then there must be 20 i-nodes everywhere. This restriction does not mean that each of the nodes must be physically distinct. An example is a grid where several nodes having the same physical coordinates is a grid for an axisymmetric object. The grid is generated by rotating a planar distribution of nodes around the axis of rotation. Those nodes located on

the axis are overlapped when the plane rotates, resulting in an individual node for each plane of rotation modelled but sharing the same coordinate on the axis of rotation.

The second restriction implies that the direction defined by the difference in position between point (i,j,k) and point $(i+1,j,k)$ crossed with the direction defined by the difference in point (i,j,k) and point $(i,j+1,k)$ must have a non-negative dot product with the direction defined by the difference in point (i,j,k) and $(i,j,k+1)$. This restriction does not imply that the directions are constant or orthogonal throughout the grid, but only that they do not violate the right-handed system principle.

While the user is free to orient the grids in any manner, certain common practices have evolved. One is that the direction of increasing values for the first component of the triplet is along the downstream direction. The direction of increasing values for the third component of the triplet is along the normal away from any solid surface, with the solid located at the first k value. The direction of the second element of the triplet is then fixed by the restriction for a right-handed coordinate system. Again these are common practices and not requirements.

Usually the distribution of nodal points is arranged to easily model some feature or phenomenon; however, when dealing with a complex model the existence of several features and modeling issues makes the choice of a given distribution pattern difficult. It is for

situations like this that multiple grids prove useful. A good example of the use of multiple grids is evident in the model for the Boeing 747 already shown in figure 8. Here one grid is used to capture the fuselage definition, another is used to define the wing and the airflow about it, and still others are used to model the engines and exhaust plumes. The use of multiple grids allows the user to break up the problem of modeling a complex situation into simpler and more manageable sub-problems. In addition, the user can use grids of different refinement in areas where it might be more critical to resolve more accurately the flow characteristics.

With the use of multiple grids, however, comes the problem of communicating values between overlapping grids during the CFD analysis. Associated with each node is a flag that indicates if the node represents a solid, a gas, or lies in a transition region between overlapping grids. In a transition region the nodes are identified with the grid number of the grid that overlaps them. In addition, a communications file contains the neighboring nodes in the overlapping grid that are associated with each transition node in the first grid. During the CFD analysis phase it is the communication of values between these neighboring transition nodes that results in interaction of the various modeled components.

The CFD results consist of a set of three files. The first file contains the triplet dimensions of all the grids and the geometric coordinates of all the nodal points. The second file contains a listing of the mutual neighboring nodes for all transition points between

overlapping grids. The final file contains the actual solution of the CFD analysis. The file usually contains the temperature, pressure, density, and velocity of the nodal point. Withington [84] explains one possible method in his dissertation that can be used to include the species calculations in a typical CFD calculation.

5.5 Computer Resources

During the course of the current research radiative analysis and bench marking were carried out on two computer systems. The first is a Cray C-90 supercomputer with 128 megawords of main memory and eight processors, each with a peak speed of one gigaflop. The second computer system used is a Silicon Graphics Indigo 2 Extreme workstation operating with a single 150 megahertz IP2 processor, an R4010 Floating Point chip, an R4400 central processing chip, and 64 megabytes of main memory.

CHAPTER 6

ACCELERATION STUDIES

In this chapter the details of the implementation of the reverse Monte-Carlo ray trace and the line-of-sight methods are presented along with three different acceleration techniques. The general results of the methods and the acceleration strategies are documented. Considering the two test cases available it was decided to limit the radiative modeling to gaseous emission and absorptions and opaque solids with specular or diffuse reflectivities. For the auxiliary power unit case, this restriction does not greatly affect the ability to perform a radiative analysis, because the main object is a gaseous plume. The SOFIA data require more complex radiative modeling capabilities to capture accurately all the features present; however, since this data set is qualitative instead of quantitative in nature it is sufficient to model the dominant features, which are plume emissions and specular solid-surface reflections.

6.1 General Algorithms

Both the reverse Monte-Carlo ray trace and the line-of-sight method use the same radiative models and same ray tracing algorithms. The main difference between the reverse Monte-Carlo ray trace method and the line-of-sight method are that the former tracks only the final absorption site of the ray while performing a different ray trace for each wavelength subinterval. Beyond this difference, however, the two methods work the same and so are described together through this and the next chapter.

6.1.1 Observer Model

The first element of the ray trace is the observer. Since the objective is to obtain the same radiative results as the infrared measuring system, it is necessary to perform the analysis in a way that the pixelized image obtained from the infrared measurement is recreated. To summarize from Chapter 5, the Agema infrared imaging system uses a scanner which collects radiation arriving at the lens and passes it through a set of scanning mirrors. The scanning mirrors subdivide the imaged scene into distinct sections, or pixels, that are presented to a single infrared detector. One simplification is to replace the scanning mirrors and single detector by a plane containing the equivalent number of detectors. Such a device, called a *focal plane array*, already exists. However, in a focal plane array all the

detectors obtain the radiation from the scene at the same time, while in the Agema system the single detector receives radiation from different parts of the scene at different times. For the moment it is assumed that the effects of the time difference are negligible since the CFD solutions available as test cases are steady state solutions with no time variance. For a time-varying CFD solution set the time difference between detectors would be accounted for by using the appropriate data set during the analysis, as discussed in Chapter 4. Thus, it is possible to replace the single observer element by an equivalent observer with the required number of detector elements to produce a comparable pixelized image. Each of these detector elements will then correspond to a pixel in the final infrared image. From this viewpoint, the words "detector" and "pixel" are used interchangeably in this chapter.

A further simplification is to model only the scanner lens. This simplification permits the radiation loss due to system optics and detector sensitivity to be applied to the radiative analysis results as a post-processing step. The same radiative analysis can be used for different scanner systems without requiring the radiative analysis to be repeated since the system response is incorporated afterwards. This simplification, however, eliminates the correspondence between the radiation incident at a discrete and separate detector, and the radiant energy attributed to an individual pixel. All radiation contributing to the final image is modeled as being incident on the common area of the lens. The distinguishing feature that links the radiation passing through the lens to a given pixel is its direction relative to the lens' surface normal at the point of intersection. Consider for a moment figure 13. Light
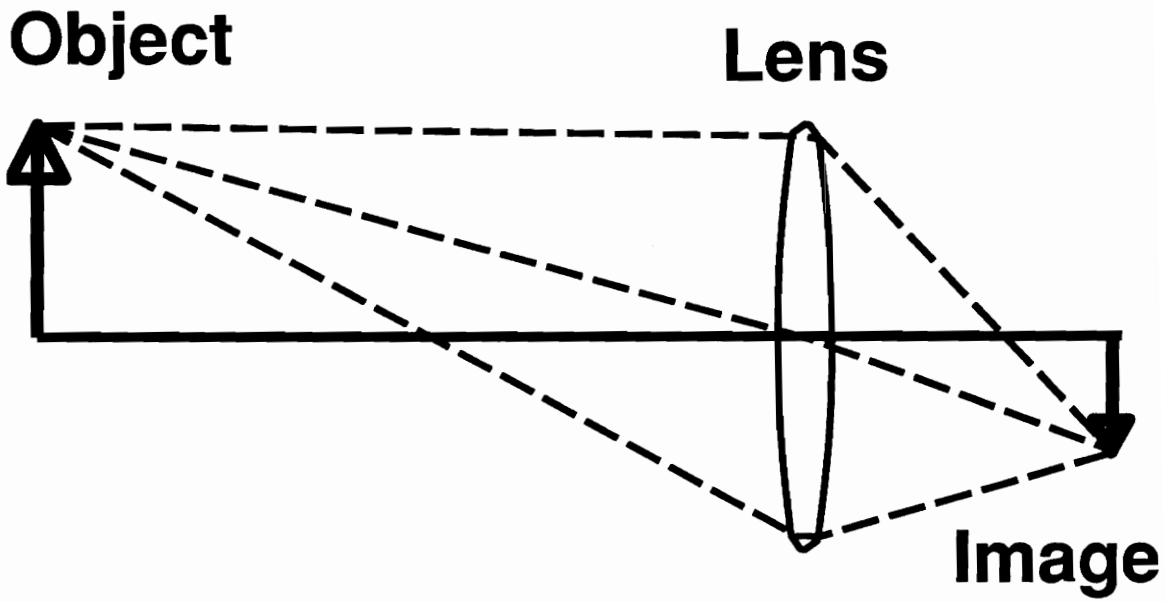
Object

Lens

Image

Figure 13 Example of a Simple Condensing Lens.

from an object is emitted in all directions but only that light which passes through the lens contributes to the final image. Furthermore that light must have a certain direction relative to the local normal of the lens surface in order to be refracted, according to the laws of optics, in the appropriate direction to contribute to the same point on the image. Light arriving at the same surface point on a lens but from a different direction will be refracted in a different direction. The same laws that apply to visible radiation also apply to thermal radiation.

Now consider for a moment an array of detectors. Each detector occupies a certain area, and is a certain distance away from the lens. If the detectors are considered as an object then the lens forms an image of the detectors along the plane in modeled space on which the lens is focused. Any radiation at the surface of the lens that appears to come from the direction of the image of a detector will be focused onto the corresponding detector. The radiation does not have to originate from the area occupied by the image of the detector; rather it only has to arrive at the lens surface along a direction defined by linking a point within the image of the detector with the arrival point on the surface of the lens, as shown in figure 14.

• **Monte–Carlo Event**

▪ **Starting Position**

Image of Detectors

Detectors

Lens

Radiation Path
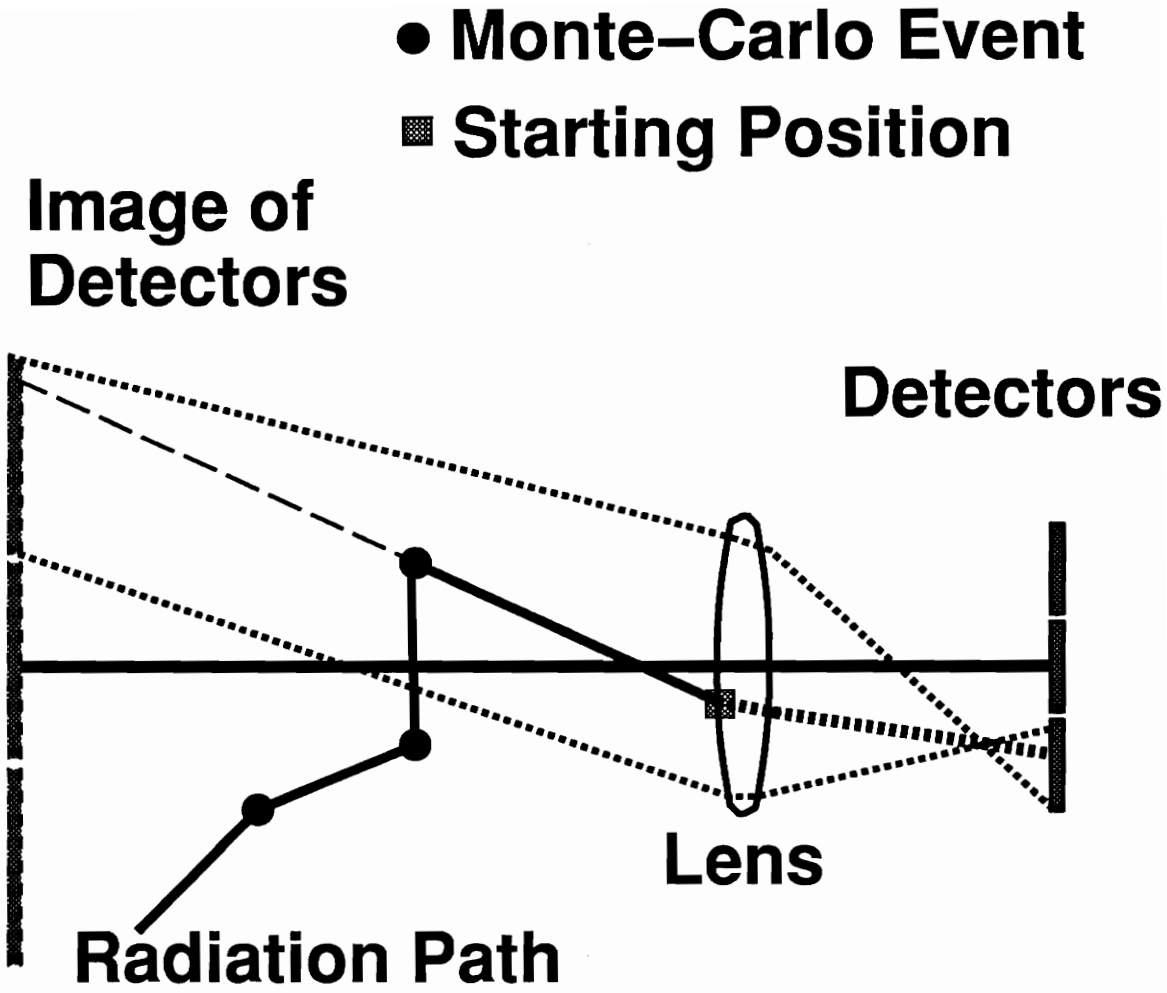
Figure 14   Example of a Detector Image Formed by a Simple Condensing Lens.

The effective lens' area forms the basis for the probable starting location of the rays. Once a starting point for the ray is chosen at random from within this area, another point is chosen at random from the area defined by the image of the detector under analysis. Together these two points define an initial direction for a ray associated with the detector under analysis. Since a detector is not capable of spatial resolution across its detecting surface, the entire area defined by the image of the detector must be sampled. As mentioned in Chapter 4, the best approaches are to distribute the location of the points across the two areas by using either a random number sequence or a quasi-random number sequence [67-70]. Both were tried, using a pseudo-random number generator based on a linear recursive congruent method, and a quasi-random number generator described by Halton [85] and Halton and Smith [86]. Neither offered an advantage in terms of time or convergence so the pseudo-random number generator was retained. This pseudo-random number generator is called *drand48()* and generates a sequence of random numbers uniformly distributed between zero and unity in double precision. No information is available concerning the generator's period, but a simpler version called *rand()* has a minimum period of 2^{48}. Therefore, it is assumed the more powerful *drand48()* has at least an equivalent period. The pseudo-random number generator is first used to distribute the starting point randomly across the lens' surface. Then it is used to select a point randomly from the area of the image of the detector currently being analyzed. The initial direction of

the ray is defined by the direction from the starting point located on the lens surface to the point on the image of the detector.

An additional benefit is gained by using the image of the detectors to determine initial ray direction. The term *depth-of-focus* in optics describes an effect caused by tolerances of imaging systems to a certain level of defocusing. An object at a fixed distance from the lens will generate a crisp image at a distance determined by the focusing power of the lens. An object located closer to the lens will generate a crisp image further back, but a blurred image at the previous image point. However photographic film, the human eye, and infrared detectors can all tolerate a certain level of blur in an object while still producing a recognizable image. This level of blur translates into a distance that an object can be shifted away from the original object and still generate an acceptable image. This distance is known as depth-of-focus. This phenomenon is recreated by the observer model as a side-effect of determining the rays' direction by the same physical logic that governs image creation in an actual imaging system. An object too far from the plane defined by the image of the detectors will appear indistinct due to lack of resolution as the cone defined by the rays traced through the image of the pixel are spread out. Objects within a certain distance of the actual focused plane will be distinguished as individual objects. Additionally, as the area of the lens decreases, the depth-of-focus will increase.

6.1.2 Volume and Surface Definition

Each radiative element in the CFD grids is identified by the element's smallest triplet. For example, a solid identified by (i,j,k) and oriented with its surface normal along the local k+ direction includes the nodal points (i,j,k), (i+1,j,k), (i,j+1,k) and (i+1,j+1,k). The volume element identified by (i,j,k) includes (i,j,k), (i+1,j,k), (i,j+1,k), (i+1,j+1,k), (i,j,k+1), (i+1,j,k+1), (i,j+1,k+1), and (i+1,j+1,k+1). The volume includes 12 half-planes, as shown in figure 15. The two i- half-planes are (i,j,k), (i,j+1,k), and (i,j,k+1); and (i,j+1,k+1), (i,j,k+1), and (i,j+1,k). The two j- half planes are (i,j,k), (i,j,k+1), and (i+1,j,k); and (i+1,j,k+1), (i+1,j,k), and (i,j,k+1). The two k- half-planes are (i,j,k), (i+1,j,k), and (i,j+1,k); and (i+1,j+1,k), (i,j+1,k), and (i+1,j,k). The two i+ half-planes are (i+1,j,k), (i+1,j,k+1), and (i+1,j+1,k); and (i+1,j+1,k+1), (i+1,j+1,k), and (i+1,j,k+1). The j+ and k+ half-planes are identified in the same manner. The planes, which are defined by having triplet index in common among themselves, are divided into half-planes because there is no requirement that a plane defined by triplets in this manner have four nodes who are geometrically coplanar.

These definitions of volume and surface elements hinge on the nodes being used to define bounding corners. Another possible approach is to define the volume elements as being centered on existing nodes. This latter approach has the benefit that the volume
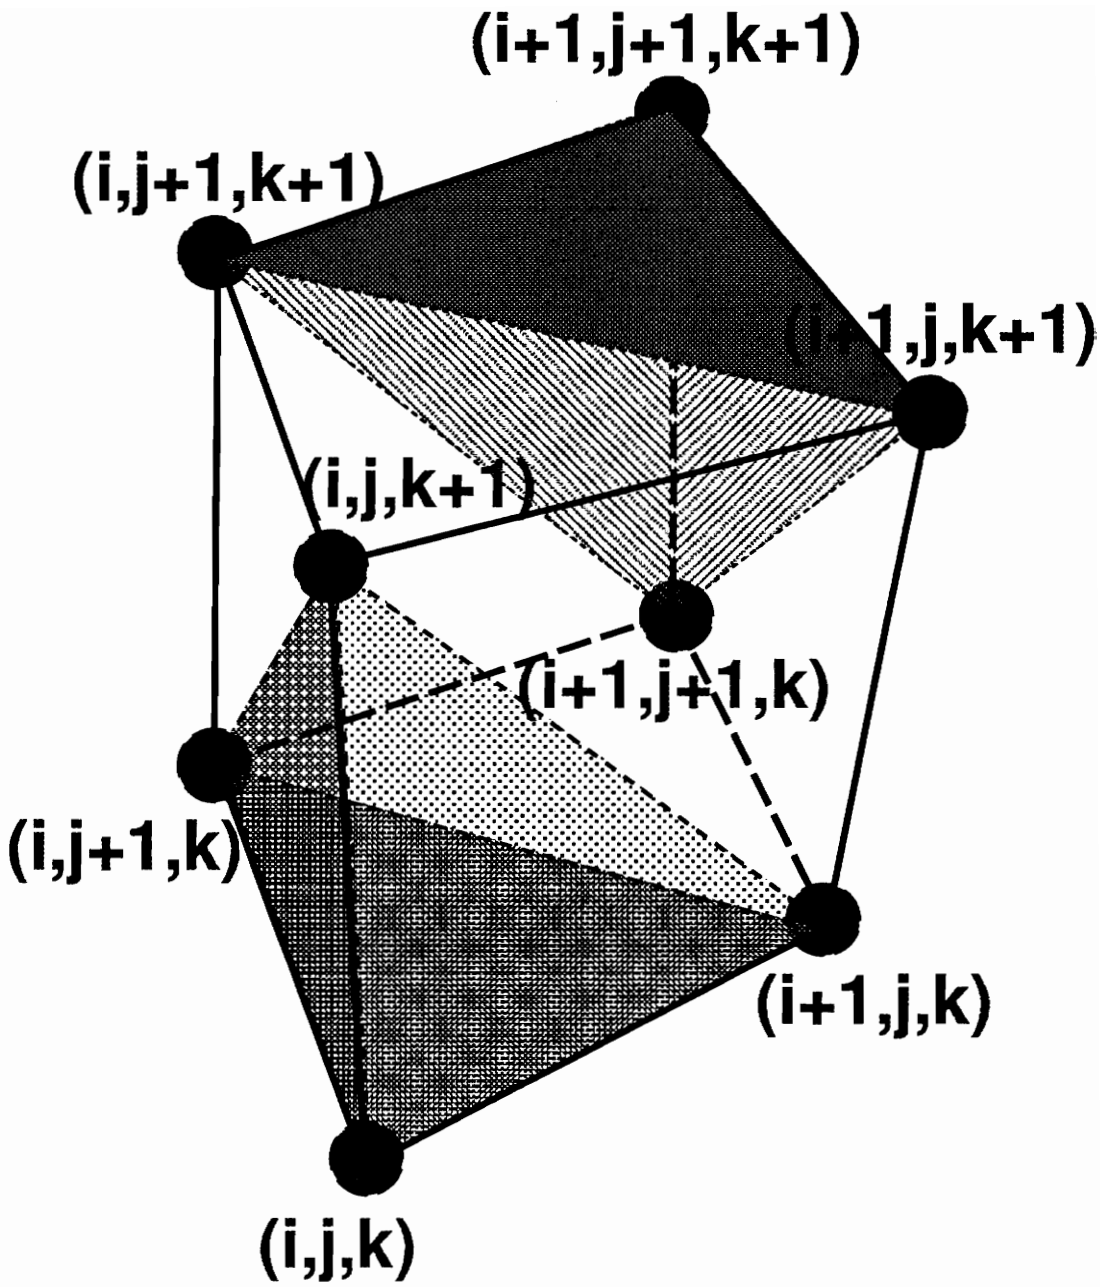
Figure 15   A Typical Volume Element and Its Included Half-Planes.

temperature and pressure are known values of that central node. The former approach must ascribe a value averaged from the eight nodes representing corners to the volume as a whole. The latter method though lacks a means to uniquely identify the boundaries of a volume element. As a result the former method, which defines the boundaries of the volume element by using the nodes to represent corners on the boundaries, is used here. Note that a surface element is defined as the four nodes forming the boundary for a volume element.

6.1.3 Implementation

Once the initial position and direction of the ray have been determined, it is necessary to determine the first element encountered. The user provides the number of the first grid encountered. Assuming that the grids are nonconcave relative to the observer location, a gradient descent method starting with an initial estimate is used to determine the closest grid point to the ray's starting point. A gradient descent method uses the gradients in geometric locations of neighboring points to determine the appropriate change in values of the triplets in order to reduce the distance. If the grid is concave it is possible for this method to obtain a node that is the closest local node to the starting point but not the grid's closest global point. The method does a search over neighboring points whenever it locates a minimum point to detect local minima but the method can still fail. The only method that

would guarantee the closest node point would be a search through all nodal points, however this method would require too much time to be practical.

The eight nodal points belonging to the volume identified by the closest nodal point are checked to see if they bound the starting point. If they do not bound the starting point, then the volume which does bound the starting point is determined. This bounding volume will be one of the four volume elements containing the closest node as a corner.

It is possible for the ray's initial location to be outside the CFD modeled space. The ray is outside the modeled space when the bounding volume is located in the outer boundary layer of the volume elements forming the grid. When the ray is outside the grid, the dot product of the ray's direction and the direction of the outward normal of the grid at that point is calculated. If the dot product is positive then the ray is heading away from modeled space and the ray trace is terminated. If the dot product is negative then the ray is heading toward the modeled space. When outside the CFD modeled space the ray is assumed to be traveling through a gas at ambient temperature, pressure and species concentration. A set of ten half-planes is defined for the exterior grids, two being the exterior surface grid, and the other eight being defined by extension of the planes from the volume below the exterior boundary volume. The ray's intersection point with the planes is calculated. A half-plane is intersected if the ray intersects the infinite plane defined by the three nodal points comprising a half-plane, and the potential intersection point is bounded in the plane by the

triangle formed by the three defining nodal points. This intersection point and the ray's starting point provide a line segment with a given length. Using this segment's length, along with the ambient temperature, pressure, and species concentration, the band model calculates the absorptivity and transmissivity of the volume. For the reverse Monte-Carlo ray trace, a pseudo-random number is generated and compared to the absorptivity. If the random number is smaller than or equal to the absorptivity then the ray is absorbed and a new ray is emitted. Otherwise the ray trace continues. For the line-of-sight method the radiative energy is attributed to the ray while the ray's transmissivity is multiplied by the volume's transmissivity. Recall that in this method the transmissivity is considered a property of the ray, as described in Chapter 3. At this point the ray's current location is advanced to the intersection point. If the ray is still outside the modeled space, the whole process is repeated with the new bounding volume until either the ray enters modeled space or heads away from modeled space.

When the ray finally enters modeled space the ray tracing process is pretty much the same as described above. The ray's intersection with one of the 12 half-planes is calculated. Again the ray's current point and the intersected point provide a length for input into any radiative model. The temperature and pressure are determined for the volume's centroid based on the modified Shepard's method for interpolating volumetric data [87,88]. For the test cases reported here, the species concentration is obtained by similarity profile from the total temperature as explained in Chapter 5. The band model calculates the absorptivity and

transmissivity based on the temperature, pressure, species concentration, path length, and wavelength of the analysis. Again the reverse Monte-Carlo ray trace uses a pseudo-random number to determine if the ray is absorbed while the line-of-sight method incorporates the volume's energy and transmissivity into the ray's values. The ray's current point is advanced to the intersection point. The four nodal points which comprise the intersected triplet plane are checked to see if the nodes represents a solid or a transition boundary between overlapping grids.

If the four nodal points are flagged as being solid nodes, then the ray has encountered a solid. The solid flag also indicates the solid type. Based on user definitions for the indicated solid type, the emissivity of the solid is calculated for the given incident direction relative to the intersected half-plane's surface normal, the temperature of the plane defined by the average of the four solid nodal points, and the wavelength of the analysis. A pseudo-random number is generated which is compared to the calculated emissivity. If the random number is less than or equal to the emissivity then the ray is absorbed by the solid; otherwise the ray is reflected. Using a weighting function based on the solid's bidirectional reflectivity distribution function, the incident direction, and a set of random numbers, the reflected direction is determined and the ray trace continues. For the line-of-sight method, the energy emitted by the solid is transmitted back along the ray while the reflectivity is multiplied by the ray's transmissivity. Note that the reflected ray usually will be traveling through the same volume it just traversed.

If the four nodal points of the intersected plane are transition points for overlapping grids, then the neighboring grid numbers and the neighboring nodal points corresponding to the four nodal points are obtained. The ray trace then executes another gradient descent method starting at the neighboring nodal points to determine the closest point in the overlapping grid. Once the closest nodal point is determined then the bounding volume is identified. If this bounding volume is valid, then the ray has successfully transferred to the new grid and the ray trace continues there.

If the bounding volume is not valid, meaning the ray is exterior to the overlapping grid or bounded by a volume with null values, then the ray trace tries the neighboring points defined by the other nodal points in the intersected plane. If these nodes fail to yield a valid volume in the new grid, then the ray is displaced by incremental values of its previous path-length forward. Again the overlapping grids are checked for valid bounding volumes. If this attempt fails then the ray is displaced backwards from its initial location and the overlapping grids checked again. If this attempt fails to provide valid bounding volumes then the ray is assumed "orphaned" between the grids. A warning message is printed, the values from the ray discounted, and a replacement ray is started from the observer. Normally a CFD analysis does not require a one-to-one correspondence between overlapping grids. A ray trace, if allowed enough rays, does traverse every single geometric point between the two overlapping grids. As a result subtle discontinuities, such as interfaces between a gaseous grid and a hollow solid, which do not affect the CFD analysis are

uncovered during the ray trace. Due to the complexity of the overlapping grids in such situations, the ray can become "trapped" in the hollow, with no practical means to determine its correct location. Generally the transition regions between overlapping grids are defined to ensure that interface zones between a particular solid and its surrounding gaseous elements are all contained within the same grid.

When a ray intersects a single half-plane then the new bounding volume is uniquely defined. For example if a ray located in volume (i,j,k) intersects one or both of the half-planes forming the $k+$ plane, then the new ray position will be located in volume $(i,j,k+1)$. Similarly if the ray intersects the $i-$ plane the new volume will be $(i-1,j,k)$. If the ray intersects two or more half-planes not forming a common boundary plane, then the next volume is indeterminate. When the volume is indeterminate the candidate volumes must be checked to determine which is the correct bounding volume. For example if a ray intersects the $i-$ and the $j+$ planes then volumes $(i-1,j,k)$, $(i,j+1,k)$ and $(i-1,j+1,k)$ must be checked to determine which is the proper bounding volume.

The current ray is traced thought the grids until it exits the modeled space or is absorbed. For the line-of-sight method the process is repeated until the current ray exits modeled space or the individual transmissivities for all the wavelengths being analyzed falls below some user-defined limit. At this point if more rays need to be traced for the current pixel, then a new observer location and a new direction point in the image of the current

detector are obtained and the ray is traced. If the current pixel is completed then the next pixel is started, starting the ray at the lens but using the image of the next detector. If the last pixel is completed then the line-of-sight method is finished, and the needed information is written out. The reverse Monte-Carlo ray trace must perform a complete ray trace for each wavelength of interest within the wavelength interval being studied before it is finished.

The information written out is a set of energies for each pixel number and for each wavelength of the analysis. This information can be displayed using the Flow Analysis Software Toolkit, FAST, a visualization program developed at NASA Ames [89]. This program displays an image with a user-defined color scale. The images displayed in this work were created using FAST.

6.1.4 Convergence Results

In order to verify that both the reverse Monte-Carlo ray trace and the line-of-sight methods converged to the same final solution when operating under the same set of assumptions, a set of five pixels was selected for simulation using the auxiliary power unit's exhaust plume. One pixel is located outside the flow in ambient conditions. A set of three pixels brackets the exhaust plume's mixing layer. The final pixel is located on the centerline

of the exhaust plume. The five pixels were used in the radiative analysis of the auxiliary power unit's plume at a wavelength of 4.17 µm. The number of rays and the instantaneous energy reported by both methods were recorded. The line-of-sight method was run for 10 million rays per pixel, with the final answers not varying beyond the sixth decimal place. The reverse Monte-Carlo ray trace was run for 10 million rays, and for each pixel the same answer was obtained as the line-of-sight method to three decimal places. This study confirms that the two methods converge to the same final answer after a sufficient number of rays.

The study also provides a measure of the two methods' convergence rates, which are shown in figures 16 and 17. Figure 16 shows the convergence rate for the line-of-sight method for the single pixel aligned with the center of the auxiliary power unit's plume. The horizontal axis is the number of rays traced while the vertical axis is the accuracy defined by the current value of energy assigned to the pixel divided by the final convergence value. Figure 17 shows the convergence for the reverse Monte-Carlo ray trace method for the same conditions. Because the other four pixels demonstrated comparable rates of convergence, these results are not shown. Both of these convergence rates are also similar to those given in the literature [46,47].
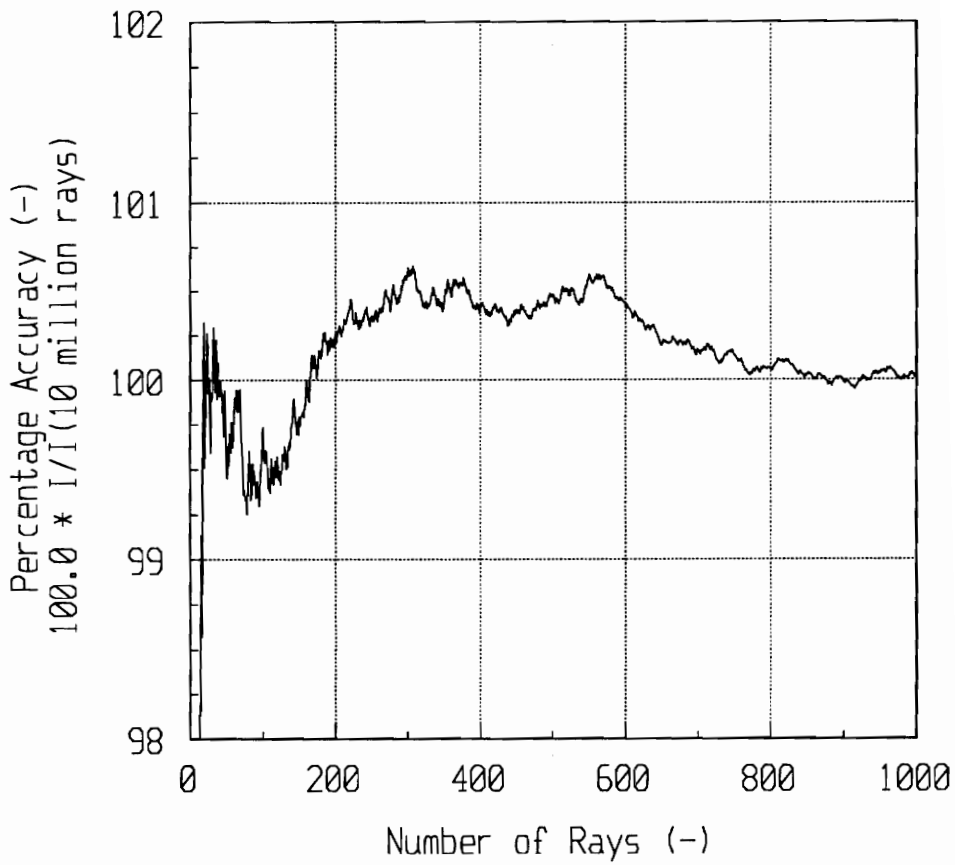
Figure 16   Convergence Rate for the Line-Of-Sight Method.
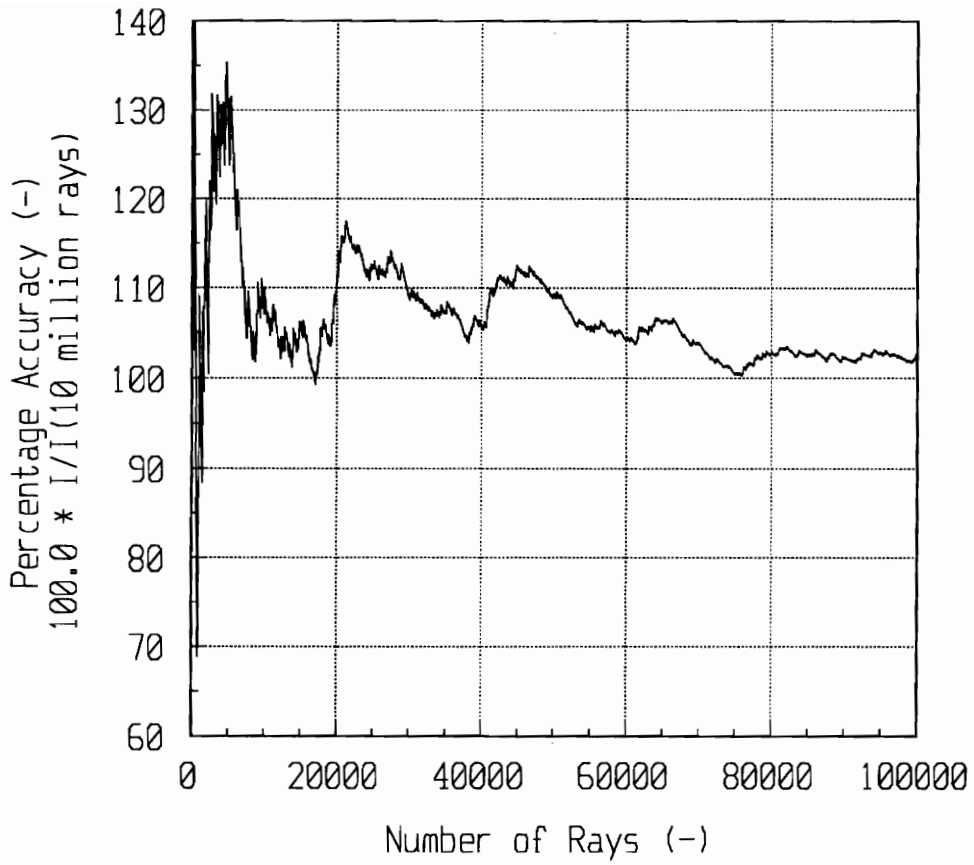
Figure 17   Convergence Rate for the Reverse Monte-Carlo Ray Trace Method.

The difference in the convergence rate for the two methods is the result of the difference in relative population sizes from which the methods obtain their statistical averages. The line-of-sight method is sampling the populations composed of all the possible viewing directions from a given pixel. The reverse Monte-Carlo ray trace is sampling from the population composed of all the possible viewing directions along with all the possible absorption sites along those different directions. Consequently, the line-of-sight method can converge to within one percent of the final answer after 20 rays while the reverse Monte-Carlo ray trace method must trace five million rays before getting within one percent!

Another point is that both methods are of a statistical nature and so the accuracy associated with a given number of rays cannot be predicted a priori. The fluctuations in the convergence rates are the result of the momentary biasing of the answer by the random sequence. The average of a sample population centered about the true average will experience a shift in that average as new members are added. This shift in the average will continue until the subpopulation of new members achieves the correct distribution among the possible states. As the sample population grows, the effect of adding new members to the population decreases because the large number of old members moderates the effect of the new members. This is the reason why the fluctuations decrease as the number of rays increases. This fluctuation is the point that Shamsundar, Sparrow and Heinisch [46] make when stating that a fixed number of rays does not always provide the best estimate. One could trace three rays and have the answer exactly and then trace 1000 rays and be off by

more than 20 percent. While the number of rays determines the error bound on the calculated answer, the problem itself and the sequence of random numbers used determine how close an individual answer is to the actual answer. No method currently exists which predicts the best sequence to use for a given number of rays so the only alternative is to keep increasing the number of rays. This increase in the number of rays will decrease the error bound on the calculated answer, so that the calculated answer approaches the actual value in the limit as the number of rays approaches infinity.

A final point is that the convergence rate is case dependent. Though the general convergence pattern will be the same, the convergence rate will depend on the number and type of radiative models, the specific radiative property values being modeled, the extent of the modeled space, the instantaneous field-of-view of the pixel, the wavelength, and even on the nature of the random sequence used in the analysis.

6.1.5 Timing Model

It is convenient to think of the time requirement for these programs as being determined by

$$t_{TOT} = t_{INTER}N_{INTER} + t_{MOD}N_{MOD} + t_{STS} . \tag{18}$$

Equation 18 is not intended for calculation purposes but as a means to comprehend the factors behind the methods' time consumption. Equation 18 expresses the total execution time as a sum of three components: (1) the time required per intersection calculation multiplied by the number of intersections, (2) the time required per model multiplied by the number of modeling calculations, and (3) the time required by the system for initiation and file input and output. The number of intersections cannot be predicted, but for the reverse Monte-Carlo ray trace this quantity is determined by the number of intersections per volume, the number of volumes encountered per ray, the number of rays per pixel, the number of pixels modeled, and the number of wavelength intervals. For the line-of-sight method, this quantity is a product of only the first four of these factors. Similarly the number of modeling calculations depends on the number of radiative models consulted per volume, the number of volumes encountered per ray, the number of rays traced, the number of pixels modeled, and the number of wavelength intervals for both methods.

Fundamentally equation 18 states that an acceleration technique can focus on many different aspects of the problem to reduce execution times. The total time can be reduced by reducing the times required for intersections and radiative model calculation or by reducing the number multiplying those times. The graphical ray trace literature states that the intersection part constitutes 70 to 95 percent of the total time [57,59] and a computer flow analysis of the programs described in this dissertation reveals that 90 percent of the time is spent in the subroutines involved in intersection calculations. Since it is such a major

contributor to the overall work load, the obvious target for reduction is the intersection component. The subroutines involved were tuned as much as possible, but the N_{INTER} factor is still enormous, making the time for the intersection calculations very high. While t_{INTER} can be reduced, efforts should be made to reduce N_{INTER} as much as possible also.

It is possible to reduce the number of rays; however, the number of rays is related to the accuracy of the solution. While the reduction in N_{INTER} will result in a reduction in analysis time, it will also lead to a reduction in accuracy. Generally the user will not be satisfied with very low accuracies and so other means of reducing the computer time must be found. It is possible to track the ray's histories and after a suitable number of rays have been traced that pass only through ambient or near-ambient conditions, the ray tracing for that pixel can be halted. The assumption is that the pixel "views" ambient space and is consequently further refinement of the answer for that pixel is not important. Even better would be the elimination of all pixels deemed unimportant by the user. For example, instead of analyzing a total of 39 200 pixels, one could analyze a distribution of 1000 pixels that are known to be important. Though this approach has merit, if one is analyzing a conceptually new aircraft there is no way to know which pixels will be important before the analysis is done. Consequently the number of rays traced must remain flexible.

Another option is to reduce the number of intersection calculations per volume. One can reduce the number of intersections calculated per volume in exchange for a higher level of uncertainty of which volume actually bounds the ray. This uncertainty could lead to the ray's actual position being different from the position the program conceives the position to be. As a result radiation would be incorrectly attributed to the ray that would lead to an incorrect image. A solution would be to analyze only those CFD grids sufficiently smooth to avoid the uncertainty; however, that only accomplishes a restriction on the type and number of solutions that can be analyzed.

The number of wavelength intervals could be reduced either by studying narrower wavelength intervals, or by subdividing the interval into fewer subintervals. Studying narrower intervals does reduce the analysis time but real detectors remain broad band despite analysis restrictions. Spectral resolution is sacrificed if the number of wavelength subintervals is decreased. Spectral resolution is very important for gaseous radiation since the radiative properties of a gas vary significantly with wavelength.

One possibility not discussed so far is the restriction of the reverse Monte-Carlo ray trace to simpler problems. For example, if the same assumptions are made as for the line-of-sight method, the paths of the rays are no longer wavelength dependent. This means that the same ray trace can be utilized for multiple wavelengths. Additionally the histories of the rays will be similar, allowing the construction of a set of possible ray histories as a pre-

processing step. Then one could simply process the rays' histories to perform a radiative analysis. In fact, the assumption could be made that solid surfaces are not present, reducing the analysis to simple straight paths through a gas volume. A curve-fit to the data could be obtained and the resulting equations numerically integrated directly. The problem with making restrictive assumptions is that part of the motivation behind this work is to develop a reverse Monte-Carlo ray trace capable of modeling any radiative phenomenon to any level of accuracy and using as its base any CFD solution available. Restrictive assumptions reduce that capability and restrict the applicability of the method, so they should remain a last resort.

6.2 Reduced Grid Sets

One possibility not yet mentioned is to reduce the number of volume elements the ray encounters. CFD grids use a given local spacing in order to increase the accuracy of the flow model but also to assure the numerical stability of the basic differencing algorithms that form the heart of the CFD methodology. Many volume elements are at or near ambient conditions and do not contribute significantly to the radiative transfer. If these volume elements could be identified and eliminated, the time required should decrease because fewer volume elements would be encountered and so fewer intersections would have to be calculated. This idea forms the basis of the first acceleration technique.

The SOFIA data set was used to test the technique because it contains both surface and volume elements that must be considered for the work reduction. It was decided to retain the structured grid in the new grids instead of going to an unstructured grid format. This requirement leads to certain restrictions concerning which nodes can be eliminated. The first restriction is that if one node is to be eliminated then all nodes along a plane in the structured grid must also be eliminated to maintain the structure of the grid. For example if the node (1,2,3) is to be eliminated a plane must be selected. Here a plane is identified as all nodes having a common triplet the value. If the k plane is to be eliminated for the example node, then all nodes fitting the triplet (i,j,3) will be eliminated, where i and j cover the full range of possible values. The second restriction is that i and j planes of the grid cannot be eliminated. This restriction is due to the fact that the i and j components of the grid are used in the geometric definition of the solid surfaces. If an i or j plane is eliminated then the surface will lose some of the surface definition. This loss will have a strong impact on the solid radiation contribution. This restriction means only k planes can be eliminated. Finally, care must be taken so that the elimination of a plane of nodal points does not disrupt the communication scheme between the overlapping grids.

6.2.1 Implementation

The first step in reducing the grids is to determine which nodal planes are radiatively insignificant. One way to determine if a node is important is to calculate the radiative absorptivity for a set of consecutive nodes, then temporarily eliminate the intervening nodal plane and repeat the calculations. If the absorptivity does not change by more than a certain percentage then one can eliminate the intervening node without changing the outcome; otherwise the intervening node is too important to remove.

The general process starts by examining nodes (i,j,k), (i,j,k+1) and (i,j,k+2). If the (i,j,k+1) node is a solid node then it cannot be eliminated and the process moves up one k value, considering (i,j,k+1), (i,j,k+2) and (i,j,k+3). If the (i,j,k+1) node is not a solid or transition node then the average temperature, pressure and species concentrations between (i,j,k) and (i,j,k+1) and between (i,j,k+1) and (i,j,k+2) are determined. Also, the distance between the two sets of nodes is calculated. These values are used in the band model to determine the two transmissivities. The transmissivity for a ray passing through both sets of nodes will be the product of the individual transmissivities. Transmissivities must be used since they represent the ratio of energy that passes through a volume to the energy incident on the volume, whereas absorptivity represents the ratio of energy that is absorbed within the volume to the energy incident on the volume. Once the overall transmissivity is

calculated, equation 8 can be used to obtain the overall absorptivity. If node (i,j,k+1) is temporarily eliminated, the path length and average temperature, pressure, and species concentration will be determined between nodes (i,j,k) and (i,j,k+2). Using these new values an absorptivity can be obtained and compared to the previous overall absorptivity. If the values change by more than 10 percent, say, then node (i,j,k+1) is considered radiatively significant and the next set of three nodes, (i,j,k+1), (i,j,k+2) and (i,j,k+3) are considered. If the node is insignificant then the process is repeated for all values of i and j. If all the nodes along the k plane are insignificant and do not contain solid nodes, then the entire plane can be eliminated and the process then considers nodes (i,j,k), (i,j,k+2) and (i,j,k+3).

One detail neglected in the above discussion is the wavelength for which the absorptivity is calculated. Since absorptivity is a function of wavelength, a node which is insignificant at one wavelength might become significant at another wavelength. This consideration requires that a wavelength interval be identified over which the radiative analysis is going to be performed. Then each wavelength subinterval in this interval must be used, and only those nodes deemed insignificant at all wavelengths in the interval can be eliminated. Once a list of all the insignificant nodes is finalized, the grids and data can be renumbered. During this renumbering process care must be taken so that the transition regions are properly maintained and that the communication files are updated.

6.2.2 Results

This grid reduction process was applied to the SOFIA data over the 4.0 to 5.0 µm range using one variant of the process that retained the basic shape of the transition regions, and another variant that aggressively reduced the grids whatever the change in the transition regions. The first process resulted in a 22 percent reduction in the number of nodes, going from 869 514 to 682 458 nodes. The run time was 9434 s for the line-of-sight program on the Cray C-90 for an analysis of a 100 by 100 pixel image simulating a 20° field-of-view lens with 20 rays per pixel for a single spectral calculation at 5.0 µm. For comparison, the unaccelerated method required 10 152 s for the same problem. This acceleration produces a seven percent decrease in execution time. Figure 18 (a) shows the resulting image produced. The observer is located just in front of the right outboard engine and is looking back and up across the tailcone and underside of the wing. The color scale is based on arbitrary units where zero represents the minimum calculated intensity and unity represents the maximum calculated intensity. The color scale has been biased to enhance the difference between the fuselage and solid surfaces and the background. The more aggressively reduced data set produced a 54 percent reduction, going from 869 514 to 393 457 nodes. The run time for the same analysis conditions was 7728 s for a 24 percent reduction in execution time, and the resulting computed infrared image is shown in figure 18 (b). The irregularity in the plume apparent in figure 18 (b) is the result of the breakdown of the communication

(a) Normal Reduction
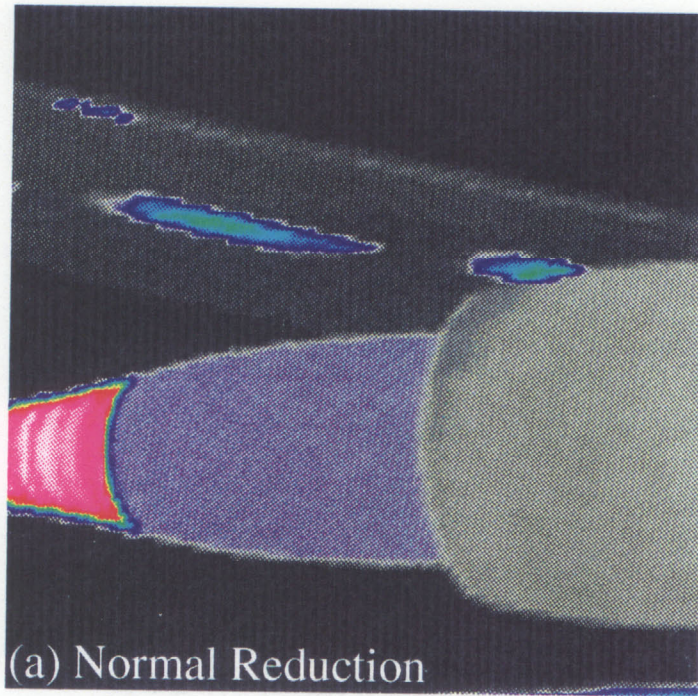
(b) Aggressive Reduction
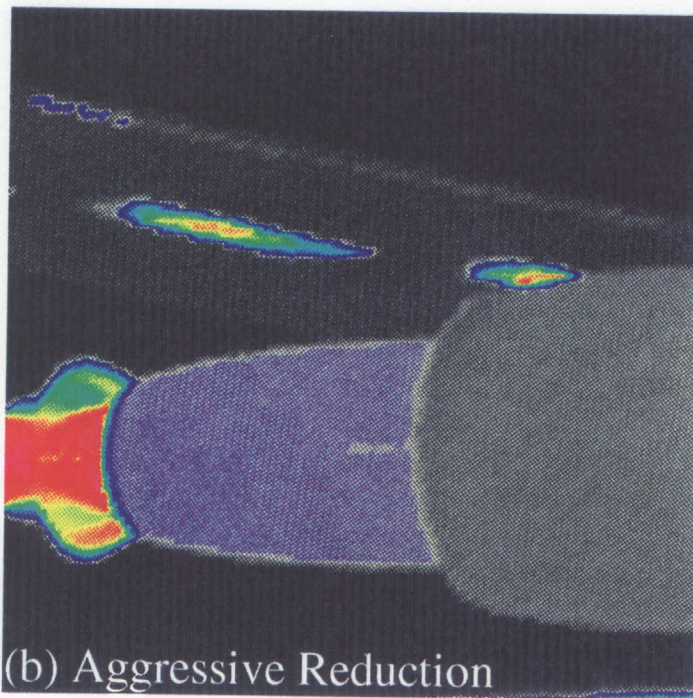
Figure 18   Computed Infrared Images of the Boeing 747 Right Outboard Engine Using Reduced Grid Sets.

Acceleration Studies   102

scheme between the overlapping grids due to the loss of transition markers caused by the aggressive reduction algorithm. Clearly this demonstrates that there is a limit to how far the grids can be reduced before enough information is lost that the grids become useless for a ray trace.

The reverse Monte-Carlo ray trace had similar results, going from a baseline execution time of 10 256 s for the unreduced data set to run times of 9302 s for the first reduced data set and 7519 s for the aggressively reduced data set. These times equate to a nine percent and a 27 percent reduction in execution time for the two data sets, respectively. Since the reductions in execution times are modest, the grid reduction approach is not applied to the data for the auxiliary power unit.

6.3 Parallel Processing

Since each ray is independent of the path followed by the other rays, both the reverse Monte-Carlo ray trace and the line-of-sight method can be parallelized to a high degree. The only machine available for conducting the research in parallel processing was the Cray C-90. The code was broken up so that each processor was responsible for a group of pixels. The

lack of any significant memory limitation allows each processor to work using a copy of the entire data set.

6.3.1 Implementation

The Cray handles parallel code in a specific manner. The section of code that is to be processed in parallel is identified using processor commands. The compiler recognizes the embedded commands and develops two sets of machine code: a master code and a slave code. The master code contains the sections of nonparallel code and additional code for the parallel section. This additional code allows the executing processor, or master processor, to request and control additional processors. When these processors become available, the master processor provides them with a copy of the slave code and the slave processor goes through that code with the data provided from the master processor. The only communication between the master and slave processors is at the start and end of the slave processor's task. At that time the slave processor can be released and can return to general use, or it can be assigned a new set of data to process. It is possible for the master processor to assign the data singly, or in chucks, depending upon the user's directions. Since the C-90 has only eight processors, it operates more efficiently when each slave processor is given several pixels to analyze at once, reducing the overhead associated with communication between the master and slave processors. One complication in using the C-90 is that several

jobs compete for the available processors and any job that requires a large number of processors must wait a long time for the processors to become available.

6.3.2 Results

The primary data set analyzed is the SOFIA data set. The radiative analysis modeled an observer comprised of 100 by 100 pixels located at a distance of 165 m from the center of the fuselage and aimed at a point just aft of the wing. The analysis was spectral at 5.0 μm and used 20 rays per pixel. The line-of-sight method was used predominately. The results are given in figure 19, where the horizontal axis represents the number of processors used while the vertical axis is the execution time in seconds. The solid line represents elapsed execution time, while the dotted line represents the total Central Processing Unit (CPU) execution time. An example of the difference between elapsed time and CPU time is that a code running three processors for one second has an elapsed time of one second while the total processor time is three seconds. The values from the parallel study are repeated in Table 2. The reverse Monte-Carlo ray trace was executed for two and four processor configurations, leading to similar execution times of 10 526 s total and 7491 s elapsed for the two-processor run and 10 471 s total and 4120 s elapsed for the four-processor run. These results compare with an execution time of 10 256 s for the single-processor version of the reverse Monte-Carlo ray-trace code. The two-processor analysis resulted in a 27
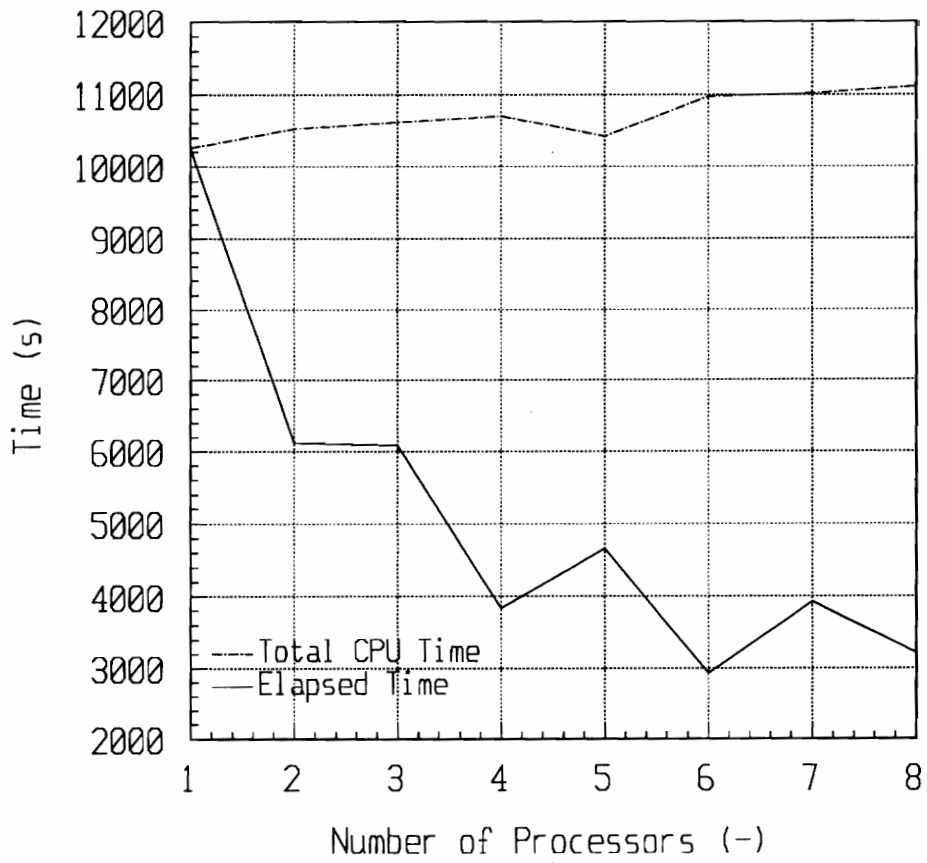
Figure 19 Parallel Processing Results.

Table 2   Parallel Processing Elapsed and Total CPU Times.

Number of Processors	Elapsed Time (s)	Total CPU Time (s)	Reduction in Elapsed Run Time (%)	Increase in Total CPU Run Time (%)
1	10 256	10 256	0.0	0.0
2	6119	10 520	40.3	2.6
3	6085	10 612	40.7	3.5
4	3838	10 700	62.7	4.3
5	4664	10 417	54.5	1.6
6	2929	10 969	71.4	7.0
7	3929	11 016	61.7	7.4
8	3209	11 108	68.7	8.3

percent reduction in elapsed time but a three percent increase in total CPU time while the four-processor analysis resulted in a 60 percent reduction in elapsed time and a two percent increase in total CPU time.

One good point is that ray traces are extremely good candidates for parallel processing. This point was apparent when the system monitor, Hampe [90], stated that the parallel version of my code was running at 96.7 percent efficiency, which was one of the best efficiencies ever attained on that system.

Figure 19 shows mixed results. If the user is concerned only with speeding up the code to obtain the answer in a shorter amount of time, then the solid line based on elapsed time should be followed. In this case equation 18 is modified by dividing by some value representing the division of the problem into several sub-problems. Since the division must include the processing time for requesting and assigning processors, the value is not simply the number of processors involved, but is related to that number. In the situation encountered with the C-90, the master processor requests additional processors and starts using them when they become available. So even though 3209 s of elapsed time where used to analyze the problem for eight processors, all eight processors were concurrently working on the analysis during only five percent of that time, or 160 s. If no other jobs were running then the master processor would have immediate access to the seven remaining processors and all the processors would work concurrently for a longer period of time.

It appears from figure 19 that the maximum efficiency is attained with six processors for this particular computer system architecture and level of usage. Since communication is not an instantaneous process, there should be an upper limit to the actual number of processors that can be profitably handled in this fashion. This upper limit comes about because the master processor spends a finite amount of time initiating and starting the other processors, and at some point the master processor needs to return and check all the processors currently running to determine if any have completed their tasks and need to be restarted. Unfortunately a dedicated system is needed to test this hypothesis.

The negative aspect of figure 19 is that while the answer is obtained in a shorter span of time, the actual computer work done to obtain the answer actually increases. This aspect of the problem comes from consideration of the percentage increase in execution time based on the total processor time. The computer work is increasing because, besides performing the radiative analysis, at least one processor is spending its time running the other processors and distributing and collecting tasks from them. Obviously as more processors are added, the one processor must spend an increasing amount of time just monitoring and communicating with other processors. This time is wasted from the standpoint that it is computer time that does not contribute directly to the solution of the problem. In today's economy it is not sufficient to get the answer in a reasonable amount of time; the answer must also be cost-effective. Since the actual monetary cost is linked to the total processing time and not elapsed time, the parallel processing analysis becomes less desirable in

intensive computer environments such as those involving supercomputers. It is possible though to run the problem through a network of inexpensive personal computers or even desktop workstations during idle time. Generally the time use associated with this level of machine does not fall under the strict scrutiny of time use on a supercomputer. If a method can be developed for initiating and communicating across the diverse machines, the radiative analysis can be done in an efficient and cost-effective manner. The Parallel Virtual Machine (PVM) [91] is an example of one such possibility for using networked machines.

6.4 Meshed Grid Sets

Returning to equation 18, the large number of rays necessary for a reasonable accuracy results in a large value of N_{INTER} for the analysis. Furthermore, t_{INTER} cannot be reduced because it is needed to navigate the rays through the complex and arbitrary orientations encountered in the CFD grids. One obvious solution then is to replace the complex grids by a simpler grid through which rays can be intersected with a simpler algorithm, such as that used by Fujimoto et al. [57]. The strength of the simpler grid is that whatever the problem involved, the grid is always oriented in the same direction, thereby leading to assumptions that greatly simplify the intersection process. Many problems are introduced by replacing the CFD grid with a simpler grid. The interpolation to the new grid can introduce errors into the analysis. Secondly, surface definition is lost since the grid can

only have lines at predetermined locations. The only available surfaces in this case would be those lying within an x, y or z plane. Thirdly, the replacement grid must either be exceedingly fine, leading to a large memory requirement, or else some of the flowfield detail is lost due to the replacement grid being coarser than the CFD grid at some locations.

It is possible to combine the two grids and retain the good features of both. The idea here is that the ray trace can start in an orthogonal, evenly spaced, grid that is oriented along the global coordinate system. This grid is called the *mesh grid* to distinguish it from the original CFD grid. The ray trace proceeds along the mesh grid using the values it contains until it reaches a region where a higher level of resolution is necessary. At that point the ray trace switches over to the corresponding point in the CFD grid and continues the ray trace in that grid. Once the ray exits the higher resolution region, it returns to the mesh grid and continues the ray trace there. The mechanics of the radiative phenomena are not tied to one set of grids or another, so no problems are introduced by the two analysis methods. Also because the ray is either in the mesh grid or the CFD grid at any given point during the analysis, then the question of which grid values to use is avoided.

The only remaining issues are the assignment of values to the mesh grid, how to handle the transition of a ray from the mesh grid to the CFD grid and back, and how to define when a ray needs to make the transition from one grid to another. First we define the

mesh grid. The mesh grid will be an orthogonal, evenly spaced, structured grid oriented along the global axis defined by the CFD grid. The direction identified by mesh nodal points (i_m, j_m, k_m) and (i_m+1, j_m, k_m) will be along the x direction and at right angles to the direction identified by mesh nodal points (i_m, j_m, k_m) and (i_m, j_m+1, k_m) for the y direction, and similarly for the z direction. The nodes in the mesh grid must be evenly spaced; however, that spacing does not need to be the same in all three directions. For example, the mesh grid could use a step of m_x for the x direction, a step of m_y for the y direction, and step of m_z for the z direction. All that is required to define completely the mesh grid are a single point, the spacing in the three directions, and the number of mesh nodes in each direction. Any desired nodal point can then be generated rapidly by

$$\begin{Bmatrix} x_{m_i} \\ y_{m_j} \\ z_{m_k} \end{Bmatrix} = \begin{Bmatrix} x_{m_0} \\ y_{m_0} \\ z_{m_0} \end{Bmatrix} + \begin{Bmatrix} i\ m_x \\ j\ m_y \\ k\ m_z \end{Bmatrix}. \tag{19}$$

Equation 19 can also be inverted to convert any floating-point coordinate into its designating triplet. Generally a triplet obtained in this manner will not be a set of three integers unless the coordinate being converted happens to lie on an existing mesh node. If the floating-point triplet is rounded to the nearest integer then the integer triplet will be the closest mesh node to that point.

The mesh node temperature and pressure values are assigned from the values of the CFD nodes located in the space around the mesh node. First, all CFD nodes located within the distance established by

$$\frac{(x - x_{m_i})^2}{\left(\frac{m_x\sqrt{3}}{2}\right)^2} + \frac{(y - y_{m_j})^2}{\left(\frac{m_y\sqrt{3}}{2}\right)^2} + \frac{(z - z_{m_k})^2}{\left(\frac{m_z\sqrt{3}}{2}\right)^2} \leq 1 \qquad (20)$$

are identified. Equation 20 is an ellipsoid centered at the mesh node and extending half the distance along the volumetric diagonal from node (i_m, j_m, k_m) to (i_m+1, j_m+1, k_m+1). The reason for using this ellipsoid is that it ensures that all of the space between nodes is covered without gaps, and allows for overlap of the CFD nodes which are used in calculating the mesh node values. Notice that the volume about the mesh node can contain nodes from many different CFD grids. Once all the CFD nodes within the ellipsoid are identified, the mesh node is assigned a temperature, pressure, and species concentrations by averaging all the CFD values without any weighting factors. If no computational fluid dynamic nodes are located within the ellipsoid then the mesh node is assigned the ambient temperature, pressure, and species concentration provided the closest CFD node is also sufficiently close to ambient values.

While calculating the values associated with the mesh node, two other pieces of information are also obtained. The first piece of information is the nearest valid CFD node

and the number of the grid to which it belongs. The second is the existence of any nodes representing a solid within a region of space covering twice the volume of the original ellipsoid. The first piece of information is used in transferring the ray from the mesh grid into the CFD grids. When the ray needs to make the transition to the CFD grid, the numbered grid and closest node associated with the current mesh node are used as the starting point for the gradient descent method to determine the nearest CFD node. Once the closest valid node is determined then the ray trace proceeds as described in the Section 6.1 of this chapter. The transition from the CFD grid back to the mesh grid is effected by using equation 19 to convert the current ray coordinate into a corresponding mesh triplet. Once the nearest mesh node is obtained, which is a simple process of rounding the components of the triplet, the ray trace can proceed in the mesh grid as outlined below.

The existence of solid elements within the vicinity of the mesh node is one indication that a ray should transition from the mesh grid to the CFD grid. Whenever a ray comes close to a solid surface, it should transfer to the CFD grid at that point in order to have access to the better-defined surfaces in the CFD grids. Another time when the ray should transfer from the mesh grid to the CFD grid is when the ray is crossing a region that involves changes in the flow properties away from ambient conditions. These changes indicate that something is occurring in the flow that the more accurate CFD grid will better resolve. For the current research, since the thermal radiation is affected primarily by the temperature, changes in the temperature away from ambient are used as the trigger to initiate

the transition of the ray from the mesh grid to the CFD grid. The ratio of the local node absolute temperature to the ambient absolute temperature is calculated, and if the ratio is larger than some prescribed value, then transition to the CFD grid is started. If the ratio drops below the prescribed value then the ray is passed back to the mesh grid. The value of this ratio is called the *temperature ratio trigger*.

6.4.1 Implementation

The ray trace proceeds in the mesh grid in much the same way as in the CFD grid. First of all the ambiguity present when starting the ray trace with multiple overlapping grids is gone. Only one mesh grid exists and it is easy to determine the location of the observer relative to it. If the observer is exterior to the mesh grid, then any steps the ray makes heading towards the grid are at ambient conditions. If the ray heads away from the mesh grid, then it is terminated at that point. Once the ray enters the mesh grid then it is handled in much the same way as for the CFD grids. One distinction is that instead of using the nodes to describe the boundaries of the volume, the volumes are rectangular boxes centered on the mesh nodes. Because of this approach, no re-interpolation is necessary to obtain the temperature, pressure, and species concentrations. Note that the mesh grid is used for the volumetric radiative calculations, which in these test cases means the gaseous absorption and transmission. In a general analysis, however, scattering and refraction could also be

involved. Any solid surface radiation is performed in the CFD grid, which has a better surface resolution.

The temperature, pressure, and species concentration of the center mesh node is used along with the path length of the ray across the volume to calculate the absorptivity and transmissivity. The ray's progress is not between bounding planes on the volume but between planes passing through the centroid of the mesh volume element. Note that since all the volumes have the same dimensions, the path length across a volume is a constant as long as the direction remains the same. Consequently the path length only needs to be calculated once. If the mesh node is in a transitional region, as described above, then the absorptivity and transmissivity for half the path length are calculated and used to determine if the ray is absorbed as it approaches the current volume's center. If the ray is not absorbed then it is transferred into the CFD grid at the current coordinate as described above and the ray trace continues from that point. If the ray is absorbed then a new ray is started from the observer location. If the mesh node is not in a transitional region then the entire path length across the volume is used in calculating the absorptivity and transmissivity. If the ray is absorbed then the appropriate counter is incremented; otherwise the ray proceeds to the next volume.

6.4.2 Mesh Intersection Algorithm

One aspect of this new approach not discussed so far is the intersection algorithm. Following the approach discussed by Fujimoto et al. [57], a simplified line-generation method can be used which involves integer arithmetic and a constant floating arithmetic to track the ray as it traverses the grid. A parametric representation of the ray is given by

$$
\begin{Bmatrix} x_{pt\cdot1} \\ y_{pt\cdot1} \\ z_{pt\cdot1} \end{Bmatrix} = \begin{Bmatrix} m_x \\ m_y \\ m_z \end{Bmatrix} t + \begin{Bmatrix} x_{pt} \\ y_{pt} \\ z_{pt} \end{Bmatrix}, \tag{21}
$$

which states simply that the next ray coordinate is the ray direction multiplied by some non-negative step size added to the current ray coordinate. We assume that all the mesh spacings are nonzero and positive. A *driving axis* is determined by examining m_x/m_i, m_y/m_j, and m_z/m_k and choosing the one with the largest absolute magnitude. A driving axis is simply that component along which increments will occur most frequently. If all three are the same absolute magnitude then the choice of driving axis between them is arbitrary. Once a driving axis is chosen, the current ray point must be moved to lie on the associated nodal plane. For the sake of illustration, we will assume that the driving axis is x in the following derivation, but a similar derivation can be obtained for a driving axis along the other two directions. The nodal plane that the ray must step towards in order to align with the mesh grid is

$$i_{align} = \text{TRUNC}\left(\frac{x_{pt}-x_{m_0}}{m_i}\right) + b , \tag{22}$$

where b is unity if m_x is positive, and zero if m_x is negative. Once the corresponding x mesh coordinate is obtained through equation 19 the new ray coordinate becomes

$$\begin{Bmatrix} x_{pt+1} \\ y_{pt+1} \\ z_{pt+1} \end{Bmatrix} = \begin{Bmatrix} m_x \\ m_y \\ m_z \end{Bmatrix}\left(\frac{x_{i_{align}}-x_{pt}}{m_x}\right) + \begin{Bmatrix} x_{pt} \\ y_{pt} \\ z_{pt} \end{Bmatrix}. \tag{23}$$

These coordinates become the new ray position after calculating the absorptivity and transmissivity encountered by the ray moving to its new position. The ray is now located on an x-plane that contains the mesh node, but the ray coordinate is not necessarily coincident with that mesh nodal point.

Now that the ray is aligned on a mesh plane along the driving direction, all subsequent steps will be from one driving mesh plane to the next. In other words,

$$t = \frac{m_i}{|m_x|} . \tag{24}$$

Rewriting equation 21, the new ray coordinates become

$$
\begin{Bmatrix} x_{pt+1} \\ y_{pt+1} \\ z_{pt+1} \end{Bmatrix} = \begin{Bmatrix} m_x \\ m_y \\ m_z \end{Bmatrix} \frac{m_i}{\mid m_x \mid} + \begin{Bmatrix} x_{pt} \\ y_{pt} \\ z_{pt} \end{Bmatrix}.
\tag{25}
$$

The ray coordinates are used primarily to aid in locating the closest current node. By using equation 19, equation 25 can be rewritten to provide the nodal triplet directly, yielding

$$
\begin{Bmatrix} \dfrac{x_{pt+1} - x_{m_0}}{m_i} \\[2mm] \dfrac{y_{pt+1} - y_{m_0}}{m_j} \\[2mm] \dfrac{z_{pt+1} - z_{m_0}}{m_k} \end{Bmatrix} = \begin{Bmatrix} \dfrac{m_x}{\mid m_x \mid} \\[2mm] \dfrac{m_y m_i}{m_j \mid m_x \mid} \\[2mm] \dfrac{m_z m_i}{m_k \mid m_x \mid} \end{Bmatrix} + \begin{Bmatrix} \dfrac{x_{pt} - x_{m_0}}{m_i} \\[2mm] \dfrac{y_{pt} - y_{m_0}}{m_j} \\[2mm] \dfrac{z_{pt} - z_{m_0}}{m_k} \end{Bmatrix}.
\tag{26}
$$

The x component of equation 26 indicates that the i component of the ray will be changing by a constant integer value of unity. The change will be negative if the direction of the ray is negative, and positive if the direction of the ray is positive. At this point a set of definitions concerning the y component is made. We define

$$
J = \text{ROUND} \left(\frac{y_{pt} - y_{m_0}}{m_j} \right),
\tag{27}
$$

$$ERR_j = \frac{y_{pt} - y_{m_0}}{m_j} - J \, , \tag{28}$$

and

$$ERR_{j+1} = \frac{y_{pt+1} - y_{m_0}}{m_j} - J \, . \tag{29}$$

Analogous definitions can be made for the z components, and introducing all of these into the y and z components of equation 26 leads to

$$\begin{Bmatrix} J + ERR_{j+1} \\ K + ERR_{k+1} \end{Bmatrix} = \begin{Bmatrix} \dfrac{m_y m_i}{m_j |m_x|} \\ \dfrac{m_z m_i}{m_k |m_x|} \end{Bmatrix} + \begin{Bmatrix} J + ERR_j \\ K + ERR_k \end{Bmatrix} . \tag{30}$$

Equation 30 further reduces to

$$\begin{Bmatrix} ERR_{j+1} \\ ERR_{k+1} \end{Bmatrix} = \begin{Bmatrix} \dfrac{m_y m_i}{m_j |m_x|} \\ \dfrac{m_z m_i}{m_k |m_x|} \end{Bmatrix} + \begin{Bmatrix} ERR_j \\ ERR_k \end{Bmatrix} . \tag{31}$$

Acceleration Studies

The current mesh triplet associated with the ray is (I,J,K) where the values are obtained from equation 27 and the analogous version of equation 27 for the x and z components. ERR_j and ERR_k are both within the range from -0.5 to 0.5, but ERR_{j+1} and ERR_{k+1} can have values within the range -1.5 to 1.5. Whenever an ERR term reaches or exceeds an absolute magnitude of 0.5, the corresponding component of the triplet must be updated, by unity if the ERR term is positive, and by negative unity if the ERR term is negative. Once the triplet has been updated the ERR terms are returned to a range between -0.5 and 0.5 by adding or subtracting unity as appropriate.

To summarize, the mesh intersection algorithm starts by finding the driving axis for the current ray direction. It then moves the ray by some small amount to align it with the driving axis mesh planes and obtains the associated mesh node (I,J,K). Up to this point the calculation has featured a series of floating point multiplications, divisions, and additions. Once this point in the calculation is reached, all subsequent calculations involve an integer addition, two floating point additions, two comparisons, and possibly two integer additions. This simplification results from equation 26, which states that the driving axis component of the triplet will always change by a constant integer value on each step. The sign of the integer value will depend on the direction of the ray but the magnitude will always be unity. Furthermore, equation 31 says that the deviation of the ray coordinate from the current mesh node is the sum of the previously calculated deviations with a constant floating-point term. If this deviation exceeds certain bounds then the mesh node triplet is updated by integer

addition or subtraction, and the deviation term is readjusted. Consequently, the intersection algorithm provides a highly efficient, quick, and accurate computation of the mesh nodes encountered as the ray traverses the mesh grid. Note that the ray's current coordinate is no longer calculated but it can be obtained by using the $(I+ERR_i, J+ERR_j, K+ERR_k)$ values in equation 19. Whenever the ray changes direction, the new driving axis must be obtained, a step aligning the ray with the driving axis mesh plane must be made, and then the ray can proceed as described above.

6.4.3 Results

A set of six computer runs was performed on the Silicon Graphics machine to determine the effectiveness of the mesh grid in accelerating the methods. The line-of-sight code was set to radiatively analyze the auxiliary power unit data with an observer composed of 100 pixels horizontally by 200 pixels vertically with 20 rays per pixel. The analysis was performed at a single wavelength of 4.17 µm. Table 3 provides the percentage reductions in time.

The three mesh grids used have the same spacing in all three directions: 0.5 grid units, 1.0 grid unit, and 5.0 grid units. Since there are no solid surfaces in the flow the only time that a ray will transfer from the mesh grid to the CFD grid is when the conditions vary

Table 3   Percentage Run Time Reduction Based on Mesh Spacing and Transitional Trigger Values.

Percentage Reduction in Run Times $100.0*(t_{baseline} - t_{current})/t_{baseline}$		Mesh Spacing		
		0.5	1.0	5.0
Transition Trigger $T/T_{ambient}$	3.00	89.1	94.0	96.1
	1.02	67.7	70.1	62.1

away from ambient conditions. When the ratio of nodal temperature to ambient temperature rises above 3.0 in the first case, and above 1.02 in the second case, the ray passes into the CFD grid. The ray will transfer back into the mesh grid when the ratio of local CFD temperature to ambient temperature drops below these factors. The value of 3.0 was chosen for the first case because it ensures that the ray remains in the mesh grid, allowing the execution time for a pure mesh grid to be determined, along with the resulting loss in resolution. The value of 1.02 allows the ray trace to cross the ambient volumes using the mesh grid, but then forces a transition to the CFD grid as the ray approaches the exhaust plume mixing layer.

The results in Table 3 are encouraging. For comparison purposes the baseline line-of-sight method without the mesh grid but using the same radiative problem has an execution time of 24 844 s on the Silicon Graphics. In the worst case this acceleration technique results in a 62 percent reduction in run time. In the best case the reduction is as high as 96 percent. To better understand what is occurring, equation 18 is now rewritten by taking into account the mesh grid contribution:

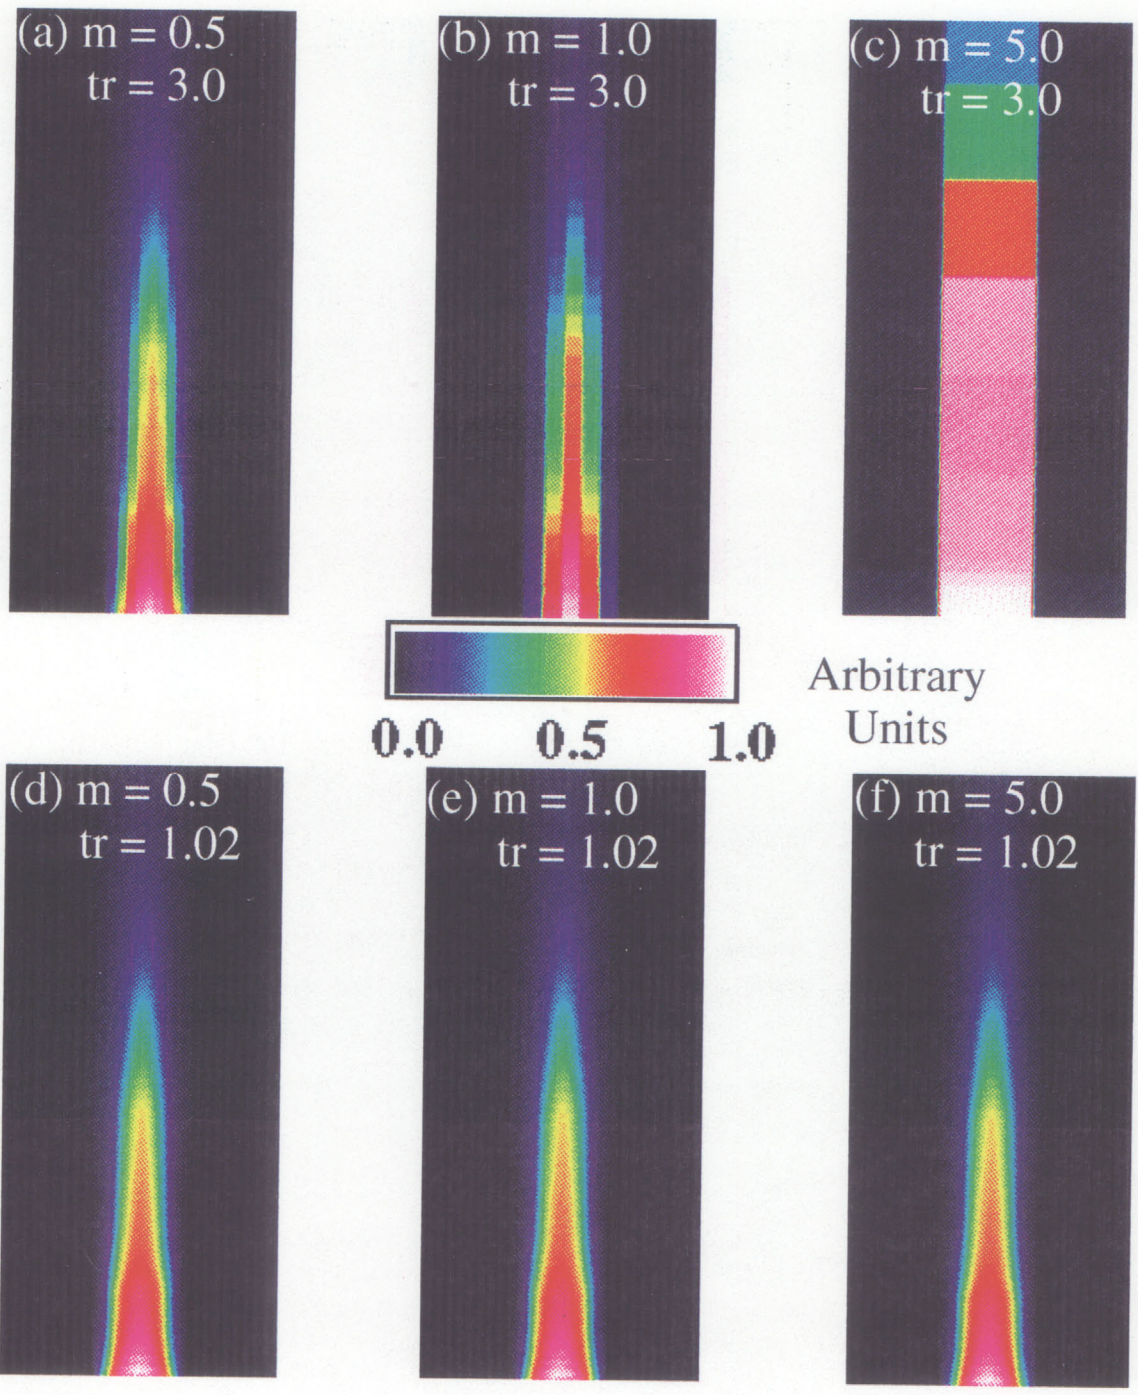$$t_{TOT} = t_{MESH}N_{MESH} + t_{CFD}N_{CFD} + t_{MOD}N_{MOD} + t_{STS} . \qquad (32)$$

In equation 32 the total time is the sum of the amount of time spent moving through the mesh grid, the amount of time spent moving through the CFD grids, the amount of time spent modeling the radiative phenomena, and the amount of time needed for initiation and

input/output. The less time spent in the CFD grid then the smaller is N_{CFD} and the less time is lost by the complex intersection algorithm. In the test case where the temperature ratio trigger is 3.0, no rays are traced through the CFD grid so the great efficiency of the mesh traversal becomes apparent through the high reduction in run times. Decreasing the spacing increases the number of mesh volumes traversed, so the run times are not reduced as much. In the cases where the temperature ratio trigger is 1.02, the reduction in run time drops because now a portion of the ray trace occurs in the CFD grid, with its slower intersection routine, rather than exclusively in the mesh grid.

One unusual feature is that for this data set and transition trigger of 1.02, the best reduction in run time is obtained for a mesh grid with a spacing of 1.0 grid units. One possible explanation for this result is that there is a trade-off involved between the number of mesh nodes the ray must traverse and how closely the mesh node corresponds to the local variations of the CFD grid. In a coarse grid the ray will reach a mesh node that has a temperature at the transition value sooner, but because that node represents a larger volume of space, the ray might have to trace through a larger segment of the CFD grid before it arrives at the part of the CFD grid that contributes significant radiation. The time gained by arriving at a transition area in the mesh grid sooner due to larger spacing can easily be lost by then forcing the ray to traverse several volumes in the CFD grid. The more refined mesh grid requires the ray to spend more time before reaching a transition area, but the transition area is likely to be closer to the plume, providing a time savings since the ray now has to

traverse fewer CFD grid volumes. As the mesh grid size decreases further, the number of

mesh volumes traversed before arriving at the transition site increases but the corresponding

location in the CFD grid does not change. The result is that the ray spends a longer time in

the mesh grid to arrive at the same transition point in the CFD grid. It appears though that

the optimal mesh grid sizing and placement will depend on the choice and value of the

transitional trigger and the problem being modeled. These results suggest that the

optimization of the mesh grid is a promising candidate for future research.

Figure 20 provides the computed infrared images of the test cases. The images are

all 100 pixels horizontally by 200 pixels vertically. The same color scale based on arbitrary

units is used, where unity represents the maximum calculated intensity, and zero represents

the minimum calculated intensity. The top images, (a), (b) and (c), are produced by the

algorithm with a temperature ratio trigger of 3.0, while the bottom images, (d), (e) and (f),

are produced with a temperature ratio trigger of 1.02. Images (a) and (d) correspond to a

mesh spacing of 0.5 grid units, while (b) and (e) correspond to a mesh spacing of 1.0 grid

units, and (c) and (f) correspond to a mesh spacing of 5.0 grid units. The vertical and

horizontal striations present in the top images are due to the resolution of the mesh grids

being inferior to the resolution of the pixels of the observer. Note that even the bottom

images have horizontal striations where the CFD grid is coarser than the level of resolution

of the observer. The top images show that large run time reductions obtained by using only

the mesh grid are offset by the poor resolution of the mesh grids. In order to match the

(a) m = 0.5
tr = 3.0

(b) m = 1.0
tr = 3.0

(c) m = 5.0
tr = 3.0

Arbitrary
Units

0.0   0.5   1.0

(d) m = 0.5
tr = 1.02

(e) m = 1.0
tr = 1.02

(f) m = 5.0
tr = 1.02

Figure 20   Computed Infrared Images Showing Effects of Mesh Spacing and Transition
Trigger Values.

Acceleration Studies   127

resolution possible by the CFD grid, the mesh grid must have the same spacing as the smallest spacing used in the CFD grid. Usually this spacing is exceedingly small and since the mesh grid must use this spacing everywhere, the increase in resolution is accompanied by a tremendous increase in number of mesh volumes traversed. The bottom images are essentially all identical to each other, as well as to the baseline line-of-sight result which is not shown. This result demonstrates the viability of mixing the two sets of grids to achieve an improvement in the method's overall speed while retaining the same accuracy and resolution.

With the acceleration obtained by using mesh grids, it now becomes possible to do multispectral images using the line-of-sight method. The radiative analysis of the auxiliary power unit data for 100 pixels horizontally by 200 pixels vertically with 40 rays per pixel, at 4.17 µm using the mesh grid with uniform spacing of 1.0 grid units and a temperature ratio trigger of 1.2 on the Cray C-90 requires a total of 7726.7 s, or about 2 hr. This provides a run time reduction of 82.4 percent. In comparison a line-of-sight analysis of the auxiliary power unit data for an image of 50 pixels horizontally by 100 pixels vertically with 5 rays per pixel for a wavelength interval from 4.0 to 5.0 µm and using the same mesh grid requires a total of 11 598.1 s, or slightly over 3 hr. A line-of-sight analysis performed without the benefit of the mesh grid for a 50-pixel horizontally by 100-pixels vertically image with five rays per pixel across 4.0 to 5.0 µm had a total execution time of 21 809.5 s, about 6 hrs, on the Cray C-90. Using this time as the multispectral baseline time, the

accelerated multispectral analysis experienced a run time reduction of 46.8 percent. The reason the run time reduction is not as large as for the spectral case is due to the fact that a multispectral analysis increases the time spent consulting radiative models. This increase in modeling time is not affected by the current acceleration techniques. Note the maximum spectral resolution possible with the band model data results in 100 spectral divisions across the 4.0 to 5.0 µm wavelength band.

In terms of the reverse Monte-Carlo ray trace, a multispectral image is generated by repeating the spectral ray trace at each wavelength division. The total number of rays is multiplied by the number of wavelength divisions, while the total time is also a multiplication of the time for a spectral ray trace multiplied by the number of wavelength divisions. Consequently if a reverse Monte-Carlo ray trace analysis is performed under the same conditions as the multispectral line-of-sight analysis outlined above, and it has a run time comparable to the spectral line-of-sight analysis, the total time required on the Cray C-90 would be 1 159 810 s, or a little over 13 days. Also note that due to the difference in convergence rates between the two methods, the reverse Monte-Carlo ray trace image generated after 13 days would have a much poorer accuracy bound than the comparable line-of-sight image.

CHAPTER 7

RESULTS

This chapter presents a comparison of the computed infrared images with the experimental infrared data described in Chapter 5. Information regarding the sensitivity of the analysis to changes in temperature, pressure, species concentration, and the similarity profile is also presented.

7.1 Comparison With Experimental Data

Again, the experimental data available consists of infrared images of the auxiliary power unit's exhaust plume in the 2.0-to-5.0- and the 4.0-to-4.9-µm wavelength bands and infrared images of the 747 Shuttle Carrier in the 2.0-to-5.0-µm wavelength band. The auxiliary power unit's data are properly documented and calibrated, allowing a derivation of the intensity associated with the measured image. The 747 data are not calibrated and are included for a qualitative comparison of the method operating on a typical CFD grid.

7.1.1 Auxiliary Power Unit

An infrared analysis between 2.0 and 5.0 µm involves 600 separate wavelength calculations at the highest spectral resolution possible with the band model data available. Time constraints did not permit an analysis across this wavelength band at this level of resolution. It was decided to concentrate on the 4.0-to-5.0-µm wavelength band since the band model data has only 100 spectral divisions there. The calculation across this narrower wavelength band involves a feasible commitment of time. These data can then be compared with the 4.0-to-4.9-µm measured images after the application of the filter response curve, shown in figure 11, to the calculated data.

In preparation for this comparison, the averaged filter data, shown in figure 3, were manipulated by selecting those pixels that contain an image of the auxiliary power unit exhaust plume. This selection resulted in an image with 35 pixels horizontally and 280 pixels vertically. Once this section of the image was obtained, a vertical section representing the background contributions without any plume contributions was also selected. This vertical section was used to simulate the background contributions behind the plume. By subtracting the simulated image of the measured background contributions from the measured image of the plume, an image is obtained which contains only the measured radiative contributions from the plume itself. This image is presented in figure 21 (a).
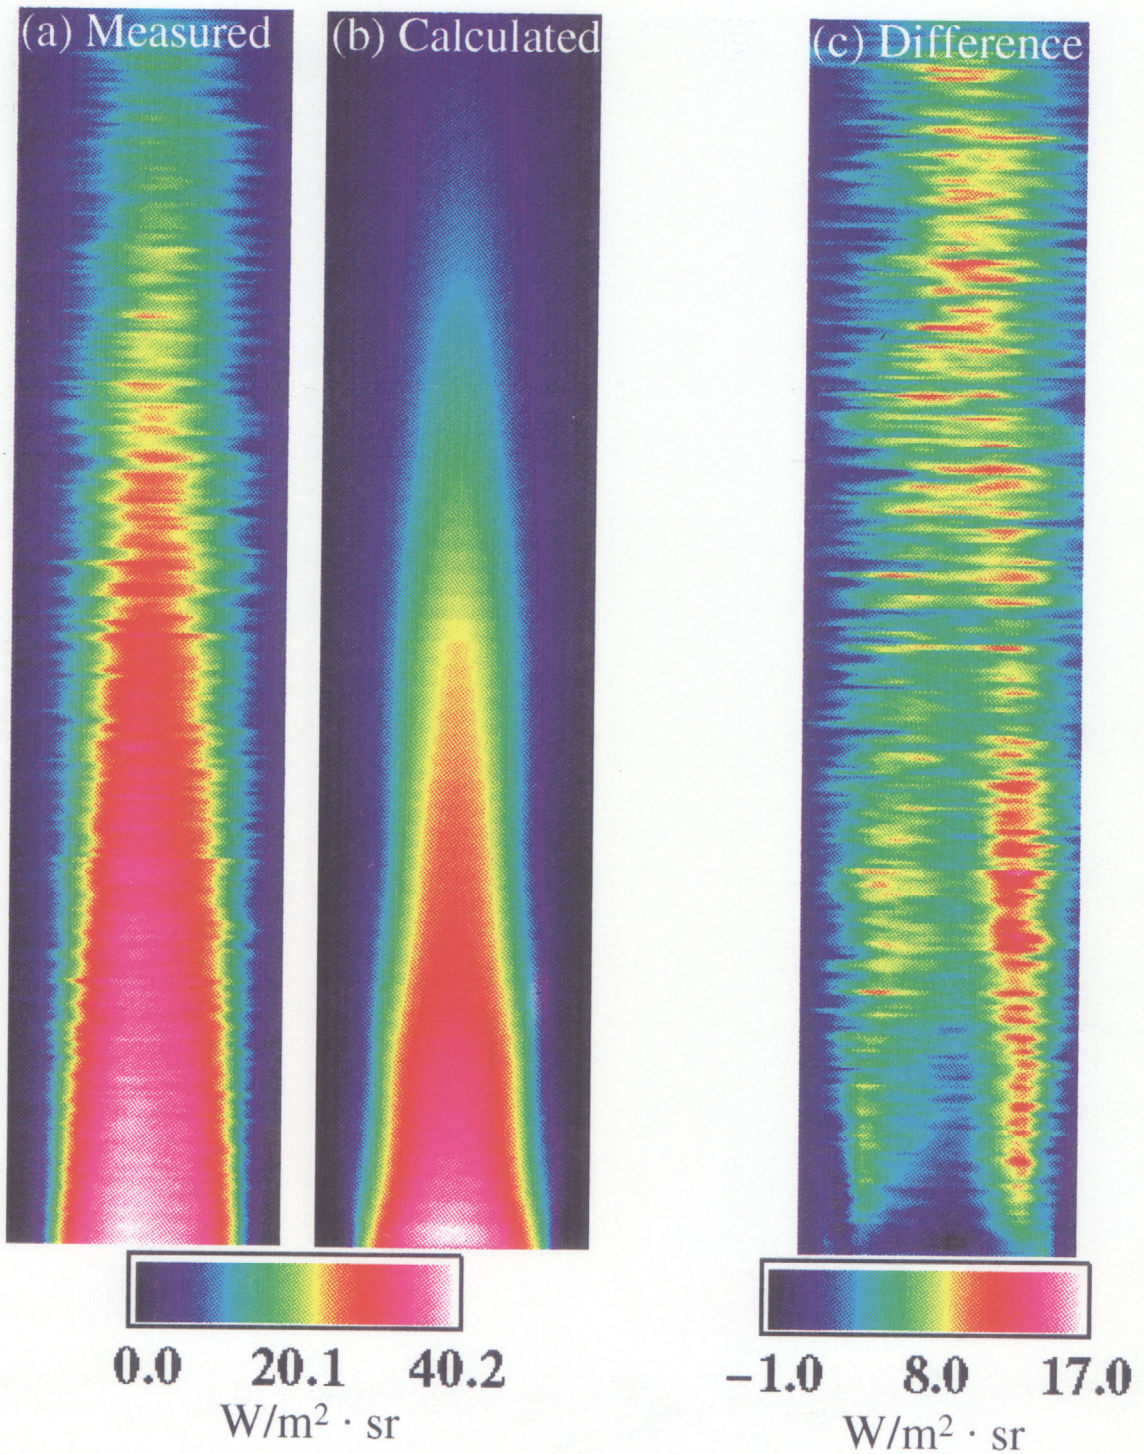
Figure 21   Comparison of Measured and Computed Infrared Images of the Auxiliary Power Unit Exhaust Plume.

Figure 21 (b) is the computed infrared image. A line-of-sight analysis was executed on the auxiliary power unit data for an observer with 35 pixels horizontally, with 280 pixels vertically, with five rays per pixel across 4.0 to 5.0 µm, and with 100 wavelength divisions. A mesh grid with 1.0-grid-unit spacing in all three directions was used along with a temperature ratio trigger of 1.2. The analysis required 22 298.0 s on the Cray C-90, a little over 6 hr. Once completed, the background contributions due to ambient gases is subtracted off and the filtered system response applied to the spectral data. The background contributions in this calculated image consist of the radiation from the atmospheric gas surrounding the plume, because no solid surfaces were included in the background. The atmospheric gas radiation is expected not to vary with location across the image scene. The spectral data were then numerically integrated across the wavelength band and the resulting answer assigned to the appropriate image pixel.

Figure 21 (c) is generated by subtracting the computed infrared image from the measured infrared image on a pixel-by-pixel basis. This subtracted image highlights the difference between the two images, providing positive numbers where the radiance of the measured image is higher than that of the computed image, and negative where the radiance of the measured image is smaller. The color scale for the three images is printed below them. The two images exhibit a similar range in intensities and image structure.

Two differences apparent in the images are the horizontal scan lines in the measured infrared image and the downstream extent of the plume. The first difference is an artifact of the scanner and related limitations in averaging and can be dismissed. The second difference is more serious for it indicates that the actual plume is longer than the computed plume. Figure 22 is a plot of the number of pixels as a function of their reported value of the difference between the two images. The average of the difference between the measured and computed images is 5.18 $W/m^2 \cdot sr$ while the standard deviation is 2.86 $W/m^2 \cdot sr$. Out of a total of 9800 pixels, 2326 have values within 1.0 $W/m^2 \cdot sr$ of the average value. Another 3869 pixels have intensity differences below 4.18 $W/m^2 \cdot sr$ and 3605 pixels have intensity differences above 6.18 $W/m^2 \cdot sr$.

Overall the results are encouraging because the average difference is 12.9 percent of the maximum measured intensity. The main detail missing that is needed to make the measured data set complete is an independent description of the species concentration distribution across the plume. As is described later in this chapter, the analysis is very sensitive to the presence of carbon dioxide across this wavelength interval, and the difference between the assumed species distribution and the actual distribution can be at least partly responsible for the observed differences in the images.

Another factor that could be responsible for the difference between the measured and calculated images is the presence of scattering in the actual exhaust plume. After each run
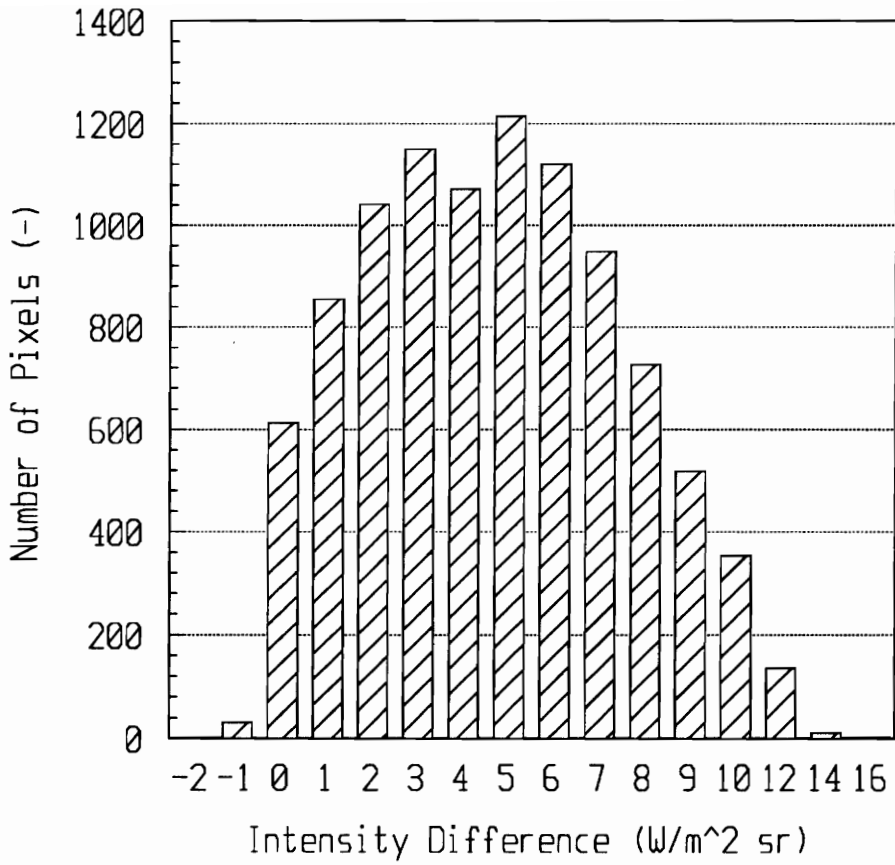
Figure 22   Pixel Radiance Probability Distribution of Differenced Image of Auxiliary Power Unit Exhaust Plume.

Results   135

the interior of the nozzle was covered in soot. If enough soot particles were present in the plume during the measurements then they could have acted as emitters operating at or near the local temperature. The soot particles would therefore extend the range of intensities as they continue emitting at or near the local temperature beyond the point where the gaseous species are sufficiently diluted that they no longer contribute a measurable signal above ambient. The relative validity of these speculations can only be established by further experimental testing.

Figure 23 demonstrates the result of a reverse Monte-Carlo ray trace analysis of the auxiliary power unit data for an observer with 25 pixels horizontally by 50 pixels vertically with 5000 rays per pixel at a wavelength of 4.17 µm. The reverse Monte-Carlo ray trace used a mesh grid with a spacing of 1.0 grid units, and a temperature ratio trigger of 1.2. Since the image is spectral and cannot be quantitatively compared, a color scale based on arbitrary units is used. The turbulent appearance of the image is due to the different levels of convergence of the individual pixels. Recall from the convergence plot of figure 17 that for 5000 rays the possible error bound is around 30 percent. However, the actual accuracy will vary from one pixel to the next, resulting in a patchwork of values where an underestimation will cause a cool spot to appear next to a hot spot representing an overestimated pixel. It is interesting to note that though this incompletely converged image bears little resemblance to the measured image that has been averaged, figure 3, it bears a striking resemblance to figure 2, the instantaneous plume images. This resemblance is
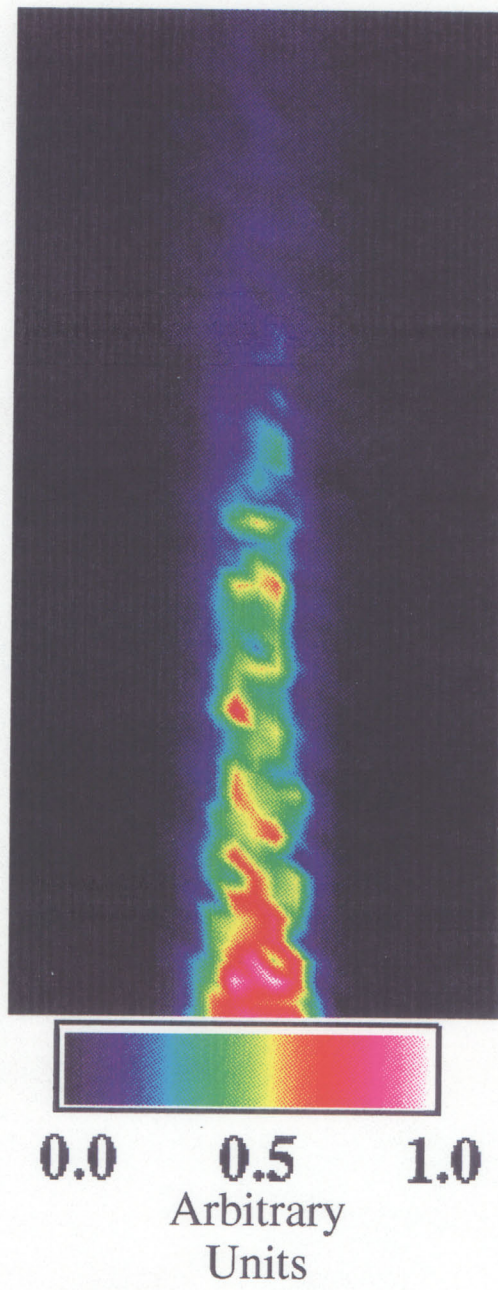
Figure 23   Reverse Monte-Carlo Ray Trace Infrared Image of the Auxiliary Power Unit Exhaust Plume.

unintentional and is the result of the random nature of turbulence fluctuations in a fluid flow being mimicked by the random nature of the convergence rate of the ray traces for the individual pixels.

7.1.2 SOFIA

The lack of calibrated data for the Boeing 747 Shuttle Carrier means that only the radiative features visible in both images can be compared. Figure 7, the measured infrared image of the 747, exhibits low background radiation, low overall solid radiation, and a broad but rapidly attenuated plume. Also visible is the large signal from the heated tailcone, and the reflection of the plumes in the underside of the wing. Figure 24 duplicates most of these features. Figure 24 is the image produced by the line-of-sight analysis of the Boeing 747 data using an observer with 70 pixels horizontally, with 70 pixels vertically, and with 90 rays per pixel at 4.17 µm. The analysis used a mesh grid with a spacing of 1.0 grid units along the three axes with a temperature ratio trigger of 1.2. The analysis was performed on the Cray C-90 and required 11 739.1 s, or a little over 3 hr. The run time reduction over the baseline execution is 48.1 percent. The color scale used with the image is based on arbitrary units, where unity corresponds to the maximum calculated intensity and zero corresponds to the minimum calculated intensity. The color scale has been biased to enhance the
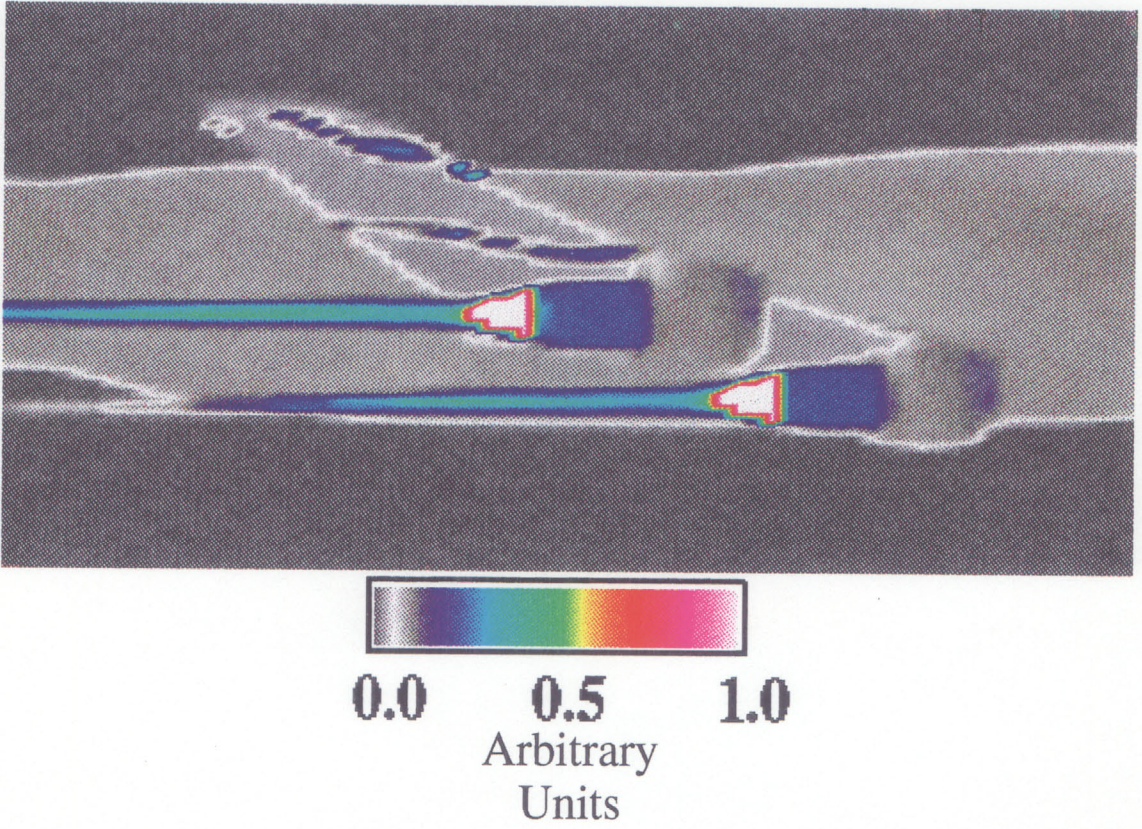
Figure 24   Computed Infrared Image of the Boeing 747 Shuttle Carrier.

visibility of the fuselage and solid surfaces from the background, as is the case for the measured image.

Even though this image results from a spectral analysis, it shows the same low background and low fuselage intensities. In addition, it shows the large signal from the heated tailcone. The computed infrared image of the exhaust plume extends further downstream and does not show the turbulent fluctuations present in the instantaneous measured image. The actual exhaust plumes encounter a lot of mixing with the surrounding airflow which quickly disperses the emitting species. The CFD solution fails to model these eddies in the plume. The computed image shows a reflection of the plume in the underside of the wing that is similar in shape and extent to that of the measured reflection. The correlation between the two images illustrates the fact that even the relatively simple radiation models use in the line-of-sight method can give results approximating the actual physics encountered. Two noteworthy features in the computed image are the tapered appearance of the inboard plume downstream, and the aliasing of the leading and trailing edges of the wing. Both of these features are a result of a limit of resolution, but the source of that resolution is different. The inboard plume appears tapered downstream because the CFD grid is stretched in the downstream direction, resulting in large volume elements at that location which produce a "block" edge to the plume image. The image of the wing is aliased not because of a lack of resolution in the CFD grid but due to the coarse resolution of the observer. The edge of the CFD wing grid is too finely resolved for the observer pixels to

distinguish. This image serves to illustrate the point that the infrared analysis has its own resolution limit, the observer resolution, and that the data on which it operates has an independent resolution. The final image will be a convolution of these two resolutions.

7.2 Analysis Sensitivity

A study was carried out to examine the sensitivity of the radiative analysis to various parameters. The parameters varied are temperature, pressure, species concentrations, and the similarity profile between temperature and species. The study involved doing a line-of-sight analysis on the auxiliary power unit data with a single pixel aimed at the core region slightly above the exit plane. The analysis was carried out across the 4.0-to-5.0-μm wavelength band, with a 1000 rays per pixel and using a mesh grid with a spacing of 1.0 grid unit in every direction, and a temperature ratio trigger of 1.2. The large number of rays traced ensures that the calculated answer is converged and that changes in intensity are the result of a change in a studied parameter.

The temperature sensitivity was obtained by varying the temperature of the volume elements along the ray by adding or subtracting a fixed amount, ΔT, at each point in the exhaust plume analysis while holding all other values constant. Note that the species concentrations were calculated before modifying the temperatures so that the species

distribution does not vary from the baseline distribution. Similarly the pressure of the volume elements encountered by the ray were also varied by adding or subtracting a fixed amount, ΔP, for each analysis during the pressure sensitivity study. The molar concentrations of carbon dioxide and water vapor were studied since they are the principal infrared emitters modeled in this work and they are present in the jet exhaust gases. The species distribution is established by using the assumed concentrations at ambient and core locations along with a linear similarity profile based on total temperature. The molar fractions of the emitting species at the core location were varied independently while holding the ambient concentrations constant. The baseline concentrations for both the centerline location and the ambient location are given in Table 1.

The effect of assuming a linear similarity profile for the species distribution was also studied. A general similarity profile can be written as

$$C_i = [C_i(core) - C_i(ambient)]\left[\frac{T - T(ambient)}{T(core) - T(ambient)}\right]^n + C_i(ambient) . \qquad (33)$$

The quantity C_i is the molar concentration of species i, T is the total temperature either at the current point, at the core, or at ambient conditions, and n is a parameter. When n is unity then the linear similarity profile is obtained. If n is varied from unity, then similarity profiles as shown in figure 25 result. These profiles lead to the same initial and final concentrations but skew the distribution through the rest of the plume either towards the core

Results   142
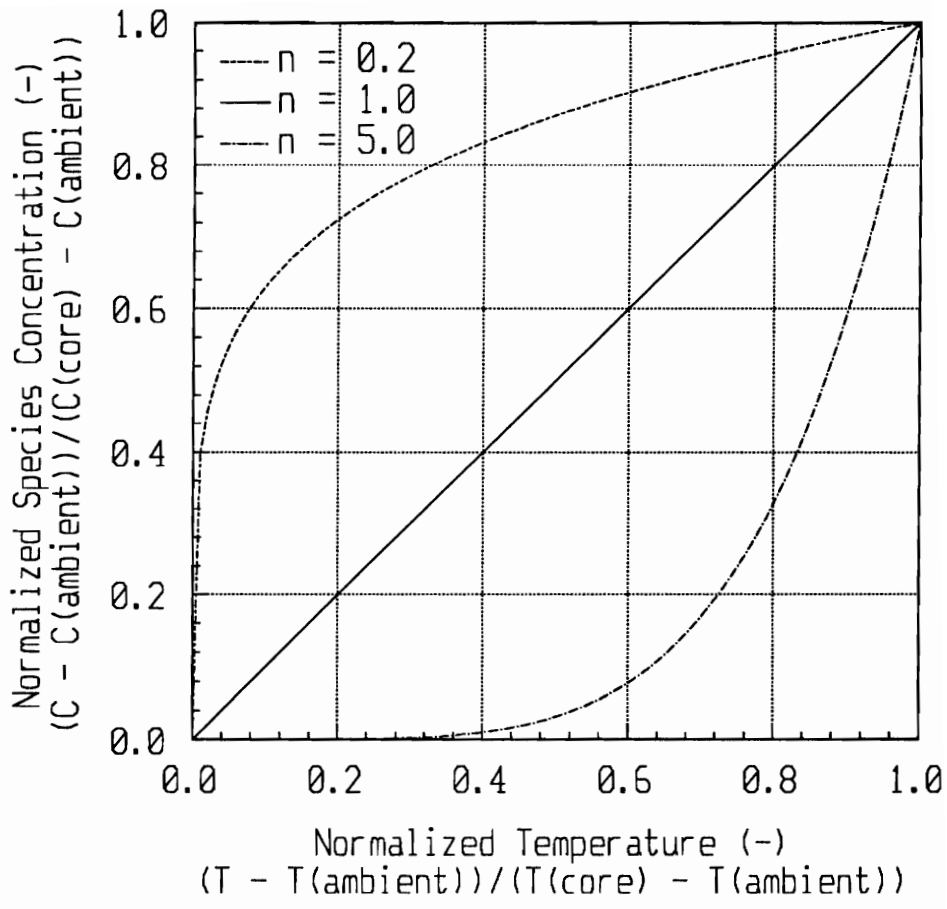
Figure 25   Representative Similarity Profiles Based On Equation 33.

concentrations, as for a value of n of 0.2, or away from the core concentrations, as for a value of n of 5.0. The baseline conditions are a core temperature of 681 K, a core pressure of 1.044 atmospheres, and core molar concentration of 0.03 carbon dioxide, 0.031 water vapor, 0.161 oxygen, and 0.778 nitrogen. The baseline radiative analysis leads to a radiance of 41.92 W/m^2·sr in this case.

The figures and tables showing the results of the sensitivity studies are gathered at the end of this chapter for ease of reference. Table 4 presents the values of the temperature sensitivity study while figure 26 shows the same results in graphical form. The horizontal axis is the difference in temperature from the baseline state normalized by the baseline state, while the vertical axis is the difference in intensity from the baseline intensity normalized by the baseline intensity, both presented as percentages. From figure 26 it is apparent that the temperature plays an important role in the infrared analysis. This result is to be expected since the temperature will not only determine the absorptivity but also the energy contribution through the Planck blackbody distribution function. The temperature sensitivity can be placed in perspective by comparing the error bound in the experimental measurements with the possible error introduced into the radiative analysis by using these measurements. The temperature error bound is ± (0.3+0.001 T) K which translates to a 0.1 percent normalized difference in temperature at the maximum temperature. This uncertainty

in the temperature measurement can result in a 1.2 percent change in intensity, which is a 0.5 W/m²·sr change out of 41.9 W/m²·sr.

Table 5 and figure 27 present the results of the pressure sensitivity study. Again the difference in pressures is normalized by the baseline state while the difference in intensities is normalized by the baseline intensity. The pressure error bound is ±0.001 atm which translates to a 0.1 percent normalized difference in pressure at the maximum pressure. The analysis essentially registers no change in intensity for that small of a change in pressure. This insensitivity is probably due to the fact that the pressure is used to calculate the gaseous absorptivity but not the direct energy contribution. Note that the figure 27 does not depict a perfect straight line due to statistical variations. Due to the relative insensitivity of the analysis method to pressure variations, the variability of the method due to its random sampling causes a slight difference in the values from their final value.

Table 6 and figure 28 present the results of the sensitivity study for the carbon dioxide concentration. The analysis is strongly dependent upon the carbon dioxide concentrations because 4.0 to 5.0 μm covers a fundamental emission band of carbon dioxide. The change in intensity is greatest for changes in carbon dioxide concentrations at low concentrations. The change in intensity starts to level off after approaching the baseline concentration. As the concentrations of carbon dioxide increase, more rays are absorbed

over a smaller depth. At some point the inner core region of the exhaust plume becomes irrelevant because all the rays are absorbed in the opaque outer areas of the exhaust plume.

Table 7 and figure 29 present the sensitivity of the analysis to water vapor concentration. Water vapor has no emission bands located close to the 4.0-to-5.0-µm waveband of this analysis. Consequently, a large change in water vapor concentration produces relatively little change in the intensity. From the physical point of view, water vapor does not even contribute to the radiative calculations at these wavelengths beyond acting as a neutral broadening agent for the infrared emissions from the carbon dioxide molecules. Again the insensitivity of the method to water vapor within this waveband means that any actual changes are at the same level as the deviation due to the random sampling of the rays.

Figures 30 shows a set of three images calculated by the line-of-sight method on the auxiliary power unit data for an observer with 35 pixels horizontally, with 280 pixels vertically from 4.0 to 5.0 µm, with five rays per pixel, a mesh grid with a 1.0 grid unit spacing, and a temperature ratio trigger of 1.2. The image on the left, figure 30 (a), is for a similarity profile with a value of n of 5.0; the image in the middle, figure 30 (b), is for the linear similarity profile; and the image on the right, figure 30 (c), is for a similarity profile

with a value of n of 0.2. The color scale associated with the three images is based on the maximum and minimum calculated intensities.

One would intuitively expect that as the relative concentration of the emitting species decreases in the area surrounding the core the overall intensity would also decrease since fewer rays get absorbed by the plume, while increasing the relative concentrations would have the opposite effect. However, the actual images lead to a different conclusion. The image for a similarity profile with a value of n of 5.0 actually has a more intense core region that drops off faster than the baseline case although the similarity profile leads to a sharp decrease of emitting species in the region surrounding the core. The image for a similarity profile with a value of n of 0.2 has less intensity in the core region but it extends slightly further out than the baseline case. These differences can be explained if one considers that a decrease in emitting species concentration decreases the absorptivity while simultaneously increasing the transmissivity. The image of figure 30 (a) therefore has a higher intensity in the core region because more of the radiation from the core region is getting through to the observer. Once the core has decayed away further downstream, the similarity profile ensures that the species concentration is so low that there is essentially nothing present to emit any radiation above the ambient value.

In figure 30 (c) the case is reversed because there the similarity profile now forces the outer layers of the plume to become species rich. The emission occurring in the cooler areas increases so the areas downstream show an increase in intensity. That increase in emission in the cooler areas is accompanied by an increase in absorption of radiation entering those areas. As a result the high intensity radiation from the core region is absorbed by the cooler region before it has a chance to reach the observer.

These results suggest the possibility that the intensity persists further downstream in the measured image than in the computed image because the centerline concentrations of species have not decayed away in that location at the same rate as the temperature has decayed. It might be that the species concentration drops off at a different rate for a line traced radially from the center out through the mixing layer than along the centerline. If the species concentration decays rapidly in the radial direction but slowly in the downstream direction, then the resulting change in the computed image will be a better estimate of the measured infrared image.
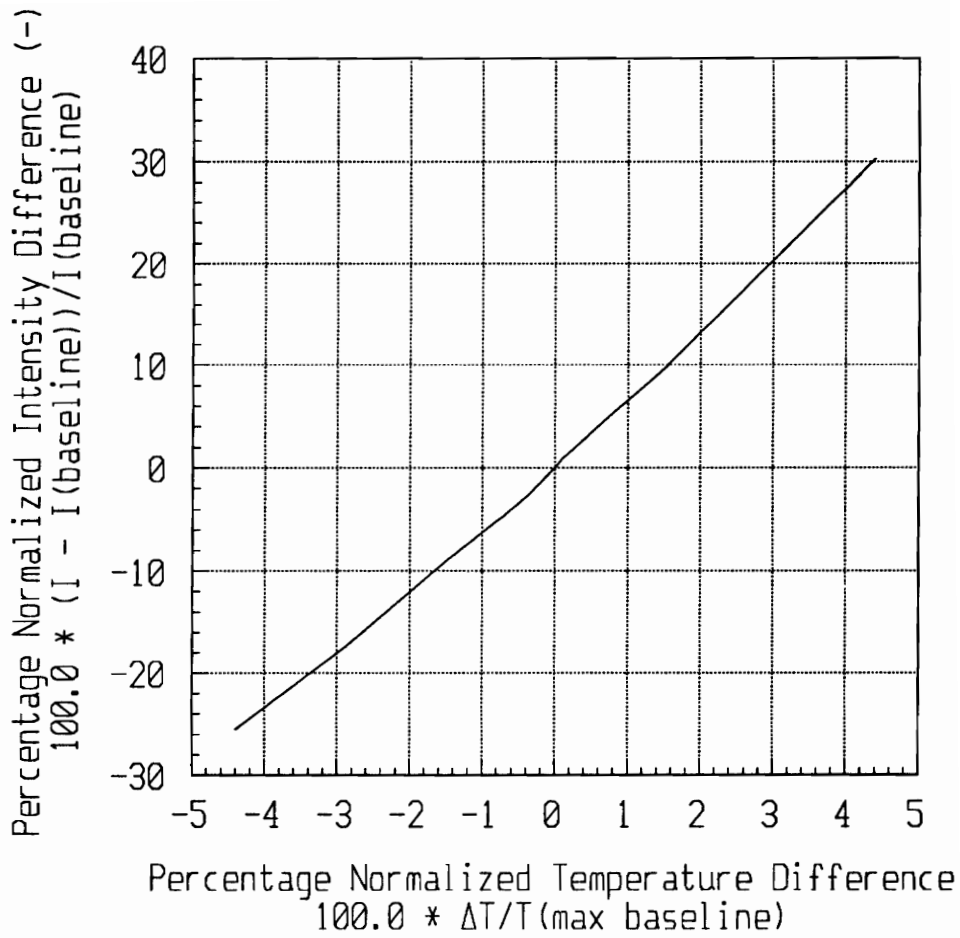
Figure 26   Results of the Temperature Sensitivity Study.
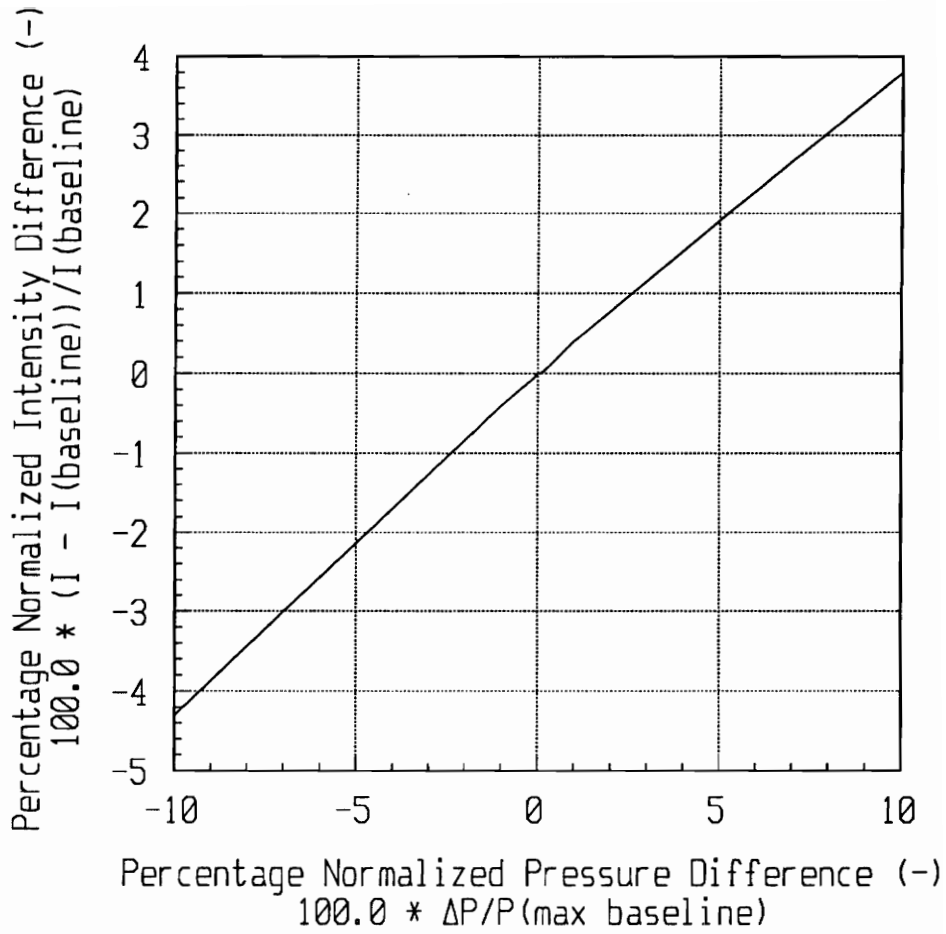
Figure 27   Results of the Pressure Sensitivity Study.
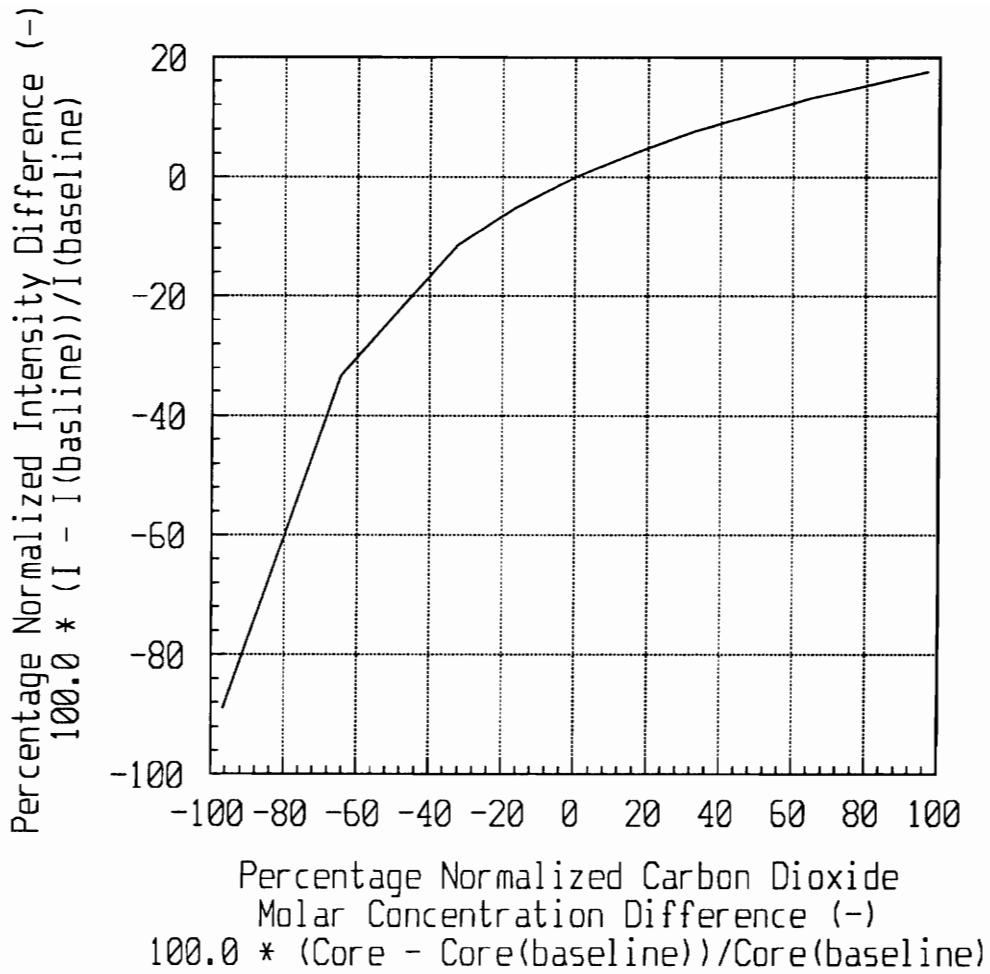
Figure 28 Results of the Carbon Dioxide Concentration Sensitivity Study.
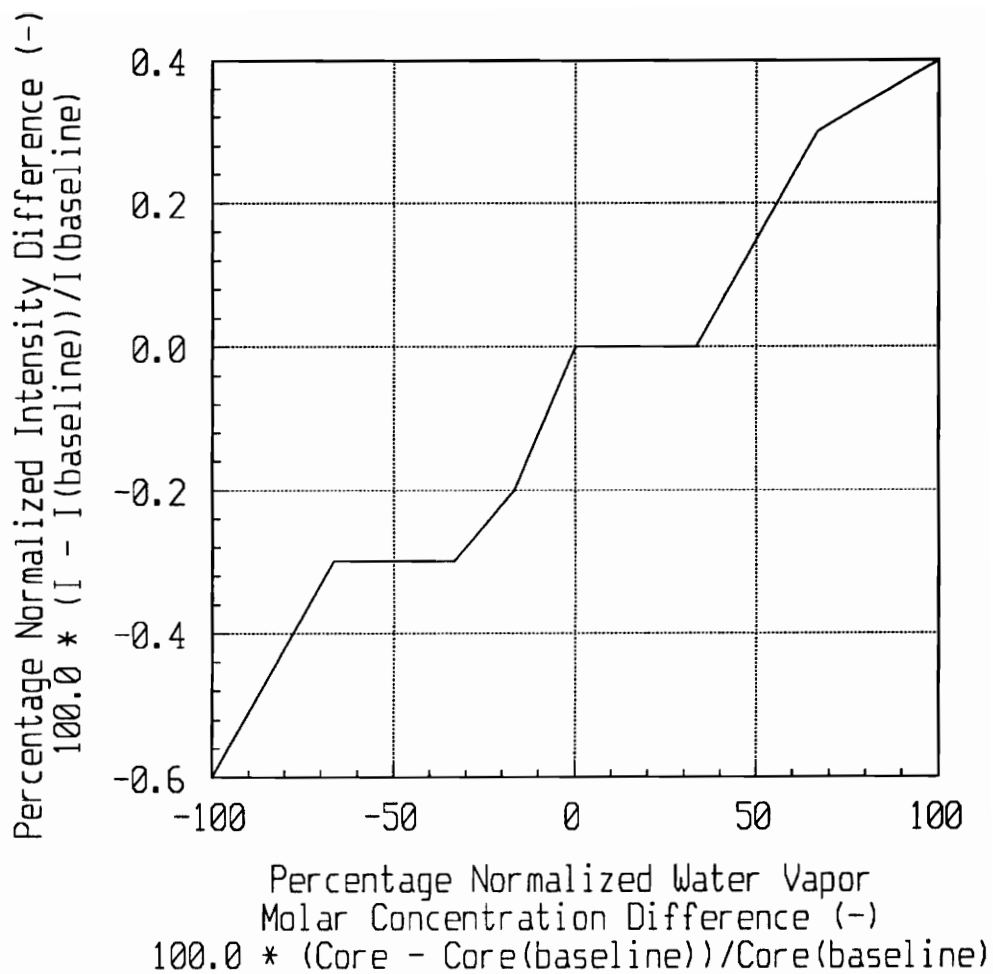
Figure 29   Results of the Water Vapor Concentration Sensitivity Study.

(a) n = 5.0   (b) n = 1.0   (c) n = 0.2

W/m² · sr

0.00   20.13   40.26
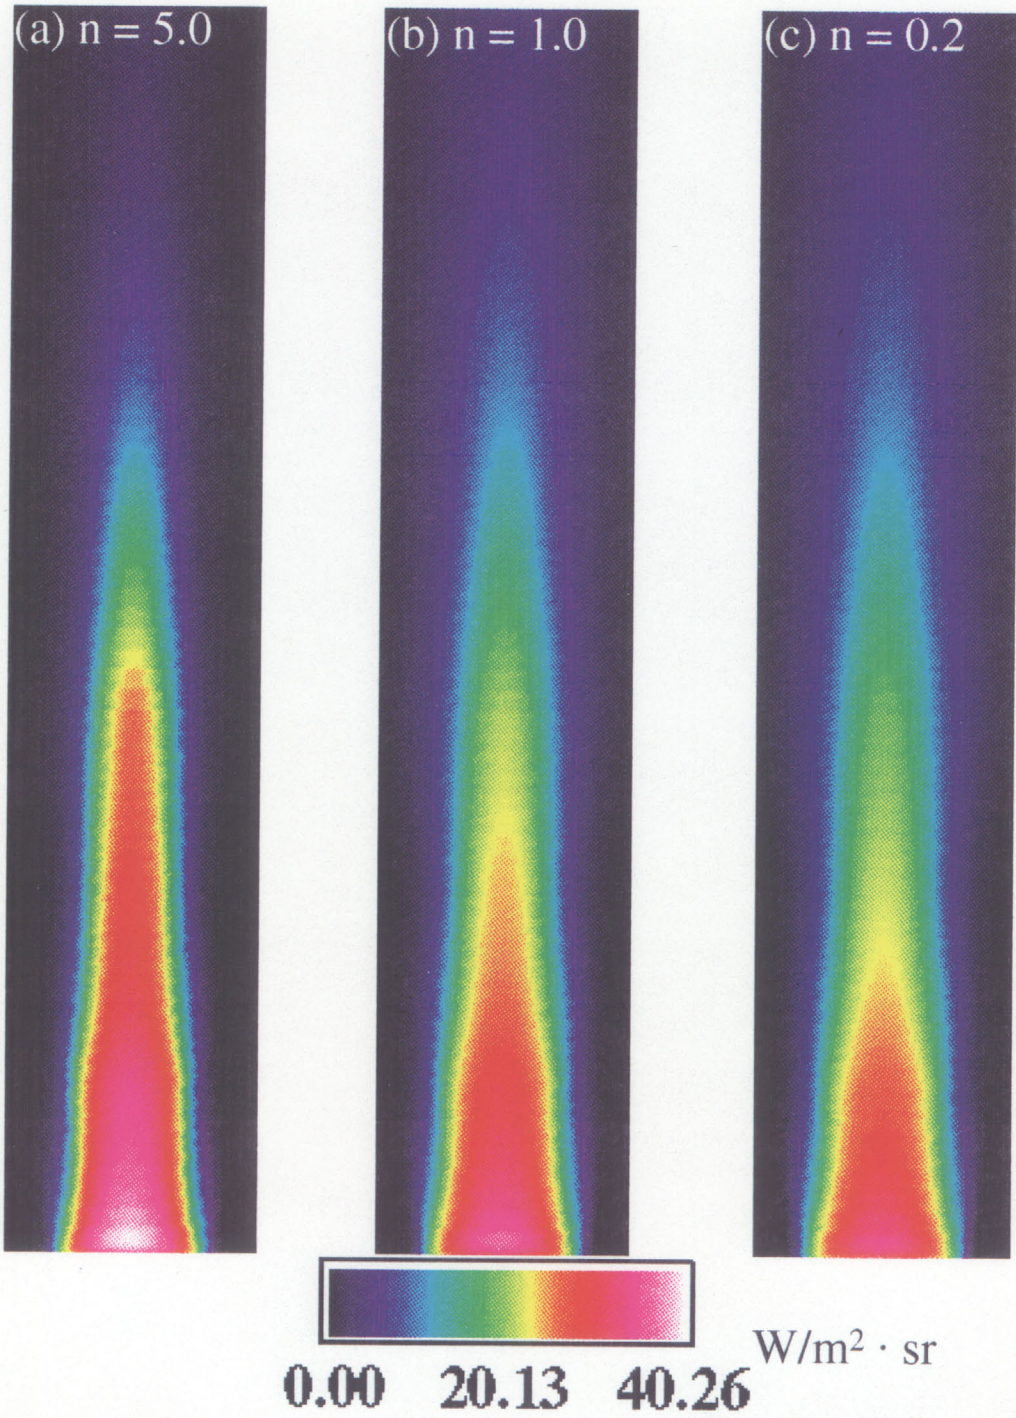
Figure 30   Infrared Images Computed for Similarity Profile Study.

Table 4   Results of the Temperature Sensitivity Study.

Temperature Difference (K)	Percentage Temperature Difference	Intensity (W/m^2·sr)	Percentage Intensity Difference
-30.0	-4.4	31.24	-25.5
-20.0	-2.9	34.60	-17.5
-10.0	-1.5	38.13	-9.0
-5.0	-0.7	39.98	-4.6
-3.0	-0.4	40.76	-2.8
0.0	0.0	41.92	0.0
+1.0	0.1	42.28	0.9
+3.0	0.4	43.11	2.8
+5.0	0.7	43.90	4.7
+10.0	1.5	45.93	9.6
+20.0	2.9	50.17	19.7
+30.0	4.4	54.62	30.3

Table 5   Results of the Pressure Sensitivity Study.

Pressure Difference (atm)	Percentage Pressure Difference	Intensity (W/m²·sr)	Percentage Intensity Difference
-0.1	-10.0	40.12	-4.3
-0.01	-1.0	41.75	-0.4
-0.001	-0.1	41.90	-0.05
-0.0001	-0.01	41.91	0.0
-0.00005	-0.005	41.92	0.0
0.0	0.0	41.92	0.0
+0.00005	0.005	41.91	0.0
+0.0001	0.01	41.89	-0.07
+0.001	0.1	41.90	-0.05
+0.01	1.0	42.09	0.4
+0.1	10.0	43.50	3.8

Results   155

Note: The table header units are in W/m²·sr.

Intensity column unit: $(W/m^2 \cdot sr)$

Table 6   Results of the Carbon Dioxide Concentration Sensitivity Study

Mole Fraction CO_2	Percentage Change from 0.031	Intensity (W/m^2·sr)	Percentage Intensity Difference
0.001	-96.8	4.62	-89.0
0.011	-64.5	27.95	-33.3
0.021	-32.3	37.06	-11.5
0.026	-16.1	39.81	-5.1
0.031	0.0	41.92	0.0
0.036	16.1	43.61	4.0
0.041	32.3	45.09	7.6
0.051	64.5	47.40	13.1
0.061	96.8	49.31	17.6

Table 7   Results of the Water Vapor Concentration Sensitivity Study.

Molar Fraction H_2O	Percentage Change from 0.03	Intensity (W/m²·sr)	Percentage Intensity Difference
0.00	-100.0	41.67	-0.6
0.01	-66.7	41.78	-0.3
0.02	-33.3	41.79	-0.3
0.025	-16.7	41.85	-0.2
0.03	0.0	41.92	0.0
0.035	16.7	41.90	0.0
0.04	33.3	41.95	0.0
0.05	66.7	42.03	0.3
0.06	100.0	42.08	0.4

CHAPTER 8

CONCLUSION

8.1 Discussion

The first goal of this research is to create a radiative analysis tool capable of treating arbitrary multiple overlapping structured grids and the accompanying computational fluid dynamics solutions as input and generating a pixelized image of the infrared radiation arriving at an observer point. This tool has applications in aircraft design analysis, flow visualization and potential computational fluid dynamics validation. Currently most analysis tools available bypass the computational fluid dynamics results entirely and make various other assumptions that can lead to significant differences between computed results and measured results. The analysis tool developed here makes use of the computational fluid dynamics results and has the capability of modeling the various radiative phenomena encountered under actual operating conditions.

This first goal was achieved, resulting in a reverse Monte-Carlo ray trace which operates directly on the computational fluid dynamics grids to compute the infrared energy arriving at an observer. The reverse Monte-Carlo ray trace is capable of modeling gaseous

and solid surface emissions, scattering, refraction, and many other phenomena. Because the reverse Monte-Carlo ray trace models the physics of thermal radiation itself, all that it requires is a model of what a particular radiative phenomenon does to generate or alter the radiative energy. As a by-product of this research a line-of-sight code was also developed. The reverse Monte-Carlo ray trace is a statistical method, requiring the sampling of a large simulated population to achieve numerical convergence, while the line-of-sight method is a deterministic method operating under strict assumptions which requires a moderate number of rays to achieve convergence due to the spatial resolution of the observer. Besides this difference, both methods use the same geometric ray tracing techniques and radiative phenomena models, and as a result both methods converge to the same final solution when operating under the same restrictive assumptions. This similarity between the two methods enabled us to test techniques on the line-of-sight method to a higher level of convergence and remain confident of achieving the same results with the reverse Monte-Carlo ray trace.

The second goal of this research is to develop acceleration techniques. Since the Monte-Carlo method is a statistical modeling method, it requires a large amount of simulation to achieve a reasonable accuracy in its answer. For the reverse Monte-Carlo ray trace, this simulation is performed by using a ray trace that steps along an arbitrary grid. The combination of large simulation populations, and time-intensive intersection sorting of the ray trace, result in prohibitively long execution times. Acceleration techniques must be

developed which retain the important radiative features from the computational fluid dynamics results and do not impose unnecessary restrictions on the radiative analysis.

The first acceleration technique involves replacing radiatively insignificant small gaseous volume elements with a radiatively equivalent larger volume. The idea is that the larger volume will require fewer ray intersections than the smaller volume elements and lead to a comparably faster algorithm. The reduction of the grid is performed in such a manner as to maintain the structured nature of the grid and not modify the resolution of the solid surfaces or the radiatively significant volume elements. Due to these restrictions, the reduction in volume elements for the case studied is small, on the order of 22 percent, and this results in a seven percent decrease in execution times for the same problem. The case studied is fairly representative of the typical application of this method and so this acceleration is not deemed sufficient to make it a viable acceleration technique.

The second acceleration technique involved developing a parallel version of the code. Due to the nature of the ray trace, and the Monte-Carlo method, the disposition of each ray is independent of that of the other rays and so rays can be analyzed separately. This makes the solution of this type of problem exceedingly efficient when it is executed in parallel. The best parallel code resulted in a 71 percent decrease in elapsed execution time but a seven percent increase in total central processing unit (CPU) time. The reason a change in both elapsed times and CPU times are given is because the actual cost involved depends on

whether clock time or computer resources are more important in a given application. Parallel code creates the dilemma that elapsed execution times decrease while the total computer work increases due to overhead associated with the parallel code.

The final acceleration technique involves the introduction of an orthogonal, evenly-spaced, axis-oriented grid called a mesh grid. The mesh grid allows easy ray traversal but has poor spatial and surface modeling resolution, while the computational fluid dynamics grids lead to difficult ray traversal but have good spatial and surface modeling resolution. The idea behind this acceleration technique is to mix the two grid types so that the analysis can use the mesh grid in the areas where it is appropriate and the computational fluid dynamics grids in the areas where they are appropriate. The execution time is reduced by 70 percent by using the mesh grid for the case studied. This acceleration technique appears to be the most effective of the three studied since it involves an actual reduction in both time and computational effort.

Unfortunately none of the acceleration techniques tested are sufficiently powerful to permit the use of the reverse Monte-Carlo ray trace on a typical practical radiative analysis. A typical problem might involve a 140 horizontal by 280 vertical pixel image across a 2.0-to-5.0-µm wavelength band, which has 600 wavelength divisions at the highest resolution currently available. Assuming the answer must be within ±10 percent of the final answer, the number of rays required per pixel is one million. This problem would require

38 000 years for the unaccelerated analysis on a Cray C-90! Assuming that both the parallel and mesh grids accelerations can be used together, the problem would require 3424 years instead. Clearly the acceleration techniques are not adequate to make this type of analysis viable. Reducing the image size, the level of acceptable accuracy, or the number of wavelength divisions would result in further time reductions, but each has implications for the answer that must be evaluated by the user to determine if those compromises are acceptable.

8.2 Recommendations For Future Work

The recommendations fall into three categories: future measurement work, radiative models, and possible acceleration techniques. The key limitation encountered during this research was the lack of data concerning species concentrations. Though it is relatively simple to add the mass transport equation to the computational fluid dynamics analysis, the results obtained will still require experimental verification. It is important to obtain species distributions not only for the simple conditions encountered with the auxiliary power unit, but also under supersonic and shocked flow conditions. It is also important to determine experimentally how important other radiative phenomena are and how to measure and quantify those phenomena. If certain radiative features can be shown experimentally to be insignificant, then they can be omitted from the analysis. Inclusion of scattering, for

example, might have a significant impact of the difference between a computed image and a measure image. Another experimental measurement that should be obtained is the spectral energy distribution. Most of the experimental data available is relatively broad band and, although the narrow band experimental and theoretical images agree in magnitude, the spectral distribution of energy within the band may be different between the experimental and theoretical images.

One of the recommendations for the radiative models has already been addressed. That recommendation is the determination of the importance of various phenomena and their quantification and measurement. The other recommendations are to address the lack of experimental data and models for the various radiative phenomena. For example, no studies exist which involve the measurement of bidirectional reflectivity for a wide range of materials at a wide range of temperatures and for a wide range of wavelengths. This type of information is needed, however, in order to develop realistic radiative models involving solid surfaces. Another area is the modeling of inhomogeneous paths for gaseous radiation. Though some work has been done on nonhomogeneous band models [32,33], these models are not based on experimental data for the short path lengths and high temperature and pressure gradients encountered in aircraft exhaust plume analysis. The analysis techniques developed here are based on the assumption that local thermodynamic equilibrium holds throughout the modeled space. This assumption is applicable in a wide variety of analysis situations involving aircraft; however, for flows involving shock structures the assumption

of local thermodynamic equilibrium breaks down. Further research is needed in this area to measure the radiative contributions of shocks and to develop the necessary computational models.

The final set of recommendations involves possible acceleration techniques. Much work still remains to be done on the mesh grid method to determine the optimal grid spacing and transition trigger, but in the limit it represents the fastest ray trace currently available during an analysis. The next development would be to divorce the ray tracing from the radiative analysis. In this method, a ray trace would be performed where the nodes encountered by a ray trace are recorded. The ray trace would continue until it encounters a solid surface or exists the grid. This ray trace would be done just once for all the possible directions associated with the detectors at the observer position, but also for all the directions for rays being emitted from each node. Histories would be accumulated that relate the conditions encountered along a set of directions from every node in space. The radiative analysis then would consist of accessing the histories for the appropriate direction from the analysis' current location in space. For example, the analysis could start by following a certain direction from the observer using the accumulated histories of rays along that direction. When the analysis encounters a scattering event, it would obtain a new set of ray histories based on the new direction and its current location that it would then use in its calculations. Since the ray trace is independent of wavelength, it is only calculated once, so

the expense of performing multiple intersections would only be accrued once regardless of the number of radiative analyses made or their complexities.

Another acceleration method would be to look into using a set of hybrid analysis techniques. For example, a neural net could be set up and trained using the reverse Monte-Carlo ray trace for simple problems. Once trained the neural net could handle the broad analysis of aircraft in a fast manner. Any problem areas or features requiring additional accuracy could be solved using the reverse Monte-Carlo ray trace. In this manner the final image is a combination of a large number of pixels obtained from a fast analysis with a smaller set of pixels highlighting special areas obtained with the full reverse Monte-Carlo ray trace.

REFERENCES

1   Siegel, R., and Howell, J, R., **Thermal Radiation Heat Transfer**, 2nd ed., Hemisphere Publishing Company (1981).

2   Ethell, J. L., *"Radar Combat and the Illusion of Invincibility,"* **Aerospace America**, vol. 28, no. 1, pp. 14-18 (1990).

3   Klass, P. J., *"Pentagon Weighs MAWS Merger,"* **Aviation Week & Space Technology**, vol. 141, no. 15, pp. 38-40 (1994).

4   Scott, H., *"Simulations 'Picture' IR Scenes,"* **Aerospace America**, vol. 28, no. 1, pp. 20-23 (1990).

5   Birckelbaw, L. D., and Nelson, E. L., *"Infrared Flow Visualization of V/STOL Aircraft,"* AIAA Paper No. 92-4253, American Institute of Aeronautics and Astronautics, New York (1992).

6   Torrance, K. E., and Sparrow, E. M., *"Off-Specular Peaks in the Directional Distribution of Reflected Thermal Radiation,"* **Journal of Heat Transfer**, vol. 88, pp. 223-230 (1966).

7   Elson, J. M., Bennett, J. M., and Stover, J. C., *"Wavelength and Angular Dependence of Light Scattering from Beryllium: Comparison of Theory and Experiment,"* **Applied Optics**, vol. 32, no. 19, pp. 3362-3376 (1993).

8   Bawolek, E. J., Mohr, J. B., Hirleman, E. D., and Majumdar, A., *"Light Scatter from Polysilicon and Aluminum Surfaces and Comparison with Surface-Roughness Statistics by Atomic Force Microscopy,"* **Applied Optics**, vol. 32, no. 19, pp. 3377-3400 (1993).

9   Ward, G. J., *"Measuring and Modeling Anisotropic Reflection,"* **Computer Graphics**, vol. 26, no. 2, pp. 265-272 (1992).

10   Scutaru, D., Rosenmann, L., Taine, J., Wattson, R. B., and Rothman, L. S., *"Measurement and Calculations of CO₂ Absorption at High Temperatures in the 4.3 and 2.7 μm Regions,"* **Journal of Quantitative Spectroscopy and Radiative Transfer**, vol. 50, no. 2, pp. 179-191 (1993).

11   Parker, R. A., Esplin, M. P., Wattson, R. B., Hoke, M. L., Rothman, L. S., and Blumberg, W. A. M., *"High Temperature Absorption Measurements and Modeling of CO₂ for the 12 Micron Window Region,"* **Journal of Quantitative Spectroscopy and Radiative Transfer**, vol. 48, no. 5/6, pp. 591-597 (1992).

12   Hartmann, J. M., Perrin, M. Y., Ma, Q., and Tipping, R. H., *"The Infrared Continuum of Pure Water Vapor: Calculations and High-Temperature Measurements,"* **Journal of Quantitative Spectroscopy and Radiative Transfer**, vol. 49, no. 6, pp. 675-691 (1993).

13   Hill, R. J., Clifford, S. F., and Lawrence, R. S., *"Refractive-Index and Absorption Fluctuations in the Infrared Caused by Temperature, Humidity, and Pressure Fluctuations,"* **Journal of the Optical Society of America**, vol. 70, no. 10, pp. 1192-1205 (1980).

14   Wesely, M. L., *"The Combined Effect of Temperature and Humidity Fluctuations on Refractive Index,"* **Journal of Applied Meteorology**, vol. 15, no. 1, pp. 43-49 (1976).

15   Spuckler, C. M., and Siegel, R., *"Refractive Index Effects on Radiative Behavior of a Heated Absorbing-Emitting Layer,"* **Journal of Thermophysics and Heat Transfer**, vol. 6, no. 4, pp. 596-604 (1992).

16   Charalampopoulus, T. T., and Chang, H., *"Effects of Soot Agglomeration on Radiative Transfer,"* **Journal of Quantitative Spectroscopy and Radiative Transfer**, vol. 46, no. 3, pp. 125-134 (1991).

17   Stull, V. R., and Plass, G. N., *"Emissivity of Dispersed Carbon Particles,"* **Journal of Optical Society of America**, vol. 50, no. 2, pp. 121-129 (1960).

18   Iskander, M. F., Olson, S. C., Benner, R. E., and Yoshida, D., *"Optical Scattering of Metallic and Carbon Aerosols of High Aspect Ratio,"* **Applied Optics**, vol. 25, no. 15, pp. 2514-2529 (1986).

19   Bissonnette, L. R., Smith, R. B., Ulitsky, A., Houston, J. D., and Carswell, A. E., *"Transmitted Beam Profiles, Integrated Backscatter, and Range-Resolved Backscatter in Inhomogeneous Laboratory Water Droplet Clouds,"* **Applied Optics**, vol. 27, no. 12, pp. 2485-2494 (1988).

20   Rockwitz, K. D., *"Scattering Properties of Horizontally Oriented Ice Crystal Columns in Cirrus Clouds, Part 1,"* **Applied Optics**, vol. 28, no. 19, pp. 4103-4110 (1989).

21   Toon, O. B., Pollack, J. B., Khare, B. N., *"The Optical Constants of Several Atmospheric Aerosol Species: Ammonium Sulphate, Aluminum Oxide, and Sodium Chloride,"* **Journal of Geophysical Research**, vol. 81, no. 33, pp. 5733-5748 (1976).

22   Villeneuve, P. V., Chapman, E. L., Mahan, J. R., *"Use f the Monte-Carlo Ray-Trace Method as a Design Tool for Jet Engine Visibility Suppression,"* **6th AIAA/ASME Thermophysics and Heat Transfer Conference**, Colorado Springs, Colorado, June 20-23 (1994).

23   Halton, J. H., *"A Retrospective and Prospective Survey of the Monte Carlo Method,"* **SIAM Review**, vol. 12, no. 1, pp. 1-63 (1970).

24   Mahan, J. R., and Eskin, L. D. "The Radiation Distribution Factor - Its Calculation Using Monte-Carlo and An Example of Its Application," **First U.K. National Heat Transfer Conference**, Yorkshire, England, July 4-6, (1984).

25   Bongiovi, R. P., **A Parametric Study of the Radiative and Optical Characteristics of a Scanning Radiometer for Earth Radiation Budget Applications Using the Monte-Carlo Method**, Master of Science Thesis, Virginia Polytechnic Institute and State University (1993).

26   Chapman, D. D., **A Monte-Carlo-Based Simulation of Jet Exhaust Nozzle Thermal Radiative Signatures**, Master of Science Thesis, Virginia Polytechnic Institute and State University (1992).

27   Eskin, L. D., **Application of the Monte Carlo Method to the Transient Thermal Modeling of a Diffuse-Specular Radiometer Cavity**, Master of Science Thesis, Virginia Polytechnic Institute and State University (1981).

28   Meekins, J. L., **Optical Analysis of the ERBE Scanning Thermister Bolometer Radiometer Using the Monte Carlo Method**, Master of Science Thesis, Virginia Polytechnic Institute and State University (1990).

29   Villeneuve, P. V., **A Paramtric Study of the Validity of the Weak-Line and Strong-Line Limits of Infrared Band Absorption**, Master of Science Thesis, Virginia Polytechnic Institute and State University (1992).

30   Nelson, E. L., **Development of an Infrared Gaseous Radiation Band Model Based on NASA SP-3080 for Computational Fluid Dynamic Code Validation Applications**, Master of Science Thesis, Virginia Polytechnic Institute and State University (1992).

31   Goody, R. M., **Atmospheric Radiation, Theoretical Basis**, vol. 1 (1964).

32   Lindquist, G. H., and Simmons, F. S., "*A Band Model Formulation for Very Nonuniform Paths*," **Journal of Quantitative Spectroscopy and Radiative Transfer**, vol. 12, pp. 807-820 (1972).

33   Young, S. J., "*Nonisothermal Band Model Theory*," **Journal of Quantitative Spectroscopy and Radiative Transfer**, vol. 18, pp. 1-28 (1977).

34   Young, S. J., "*Evaluation of Nonisothermal Band Models for H_2O*," **Journal of Quantitative Spectroscopy and Radiative Transfer**, vol. 18, pp. 29-45 (1977).

35   Hartmann, J. M., Levi Di Leon, R., and Taine, J., "*Line-by-line and Narrow-Band Statistcial Model Calculations for H_2O*," **Journal of Quantitative Spectroscopy and Radiative Transfer**, vol. 32, no. 2, pp. 119-127 (1984).

36   Neumann, L. and Neumann, A., "*Photosimulation: Interreflection with Arbitrary Reflectance Models and Illumination*," **Computer Graphics Forum 8**, pp. 21-34 (1989).

37   Sillion, F. X., Arvo, J. R., Westin, S. H., and Greenberg, D. P., "*A Global Illumination Solution for General Reflectance Distributions*," **Computer Graphics**, vol. 25, no. 4, pp. 187-196 (1991).

38   Westin, S. H., Arvo, J. R., and Torrance, K. E., "*Predicting Reflectance Functions from Complex Surfaces*," **Computer Graphics**, vol. 26, no. 2, pp. 255-264 (1992).

39 He, X. D., Heynen, P. O., Philips, R. L., Torrance, K. E., Salesin, D. H., and Greenberg, D. P., "*A Fast and Accurate Light Reflection Model,*" **Computer Graphics**, vol. 26, no. 2, pp. 253-244 (1992).

40 He, X. D., Torrance, K. E., Sillion, F. X., and Greenberg, D. P., "*A Comprehensive Physical Model for Light Reflection,*" **Computer Graphics**, vol. 25, no. 4, pp. 175-186 (1991).

41 Thynell, S. T., and Özişik, M. N., "*Radiation Transfer in Absorbing, Emitting, Isotropically Scattering, Homogeneous Cylindrical Media,*" **Journal of Quantitative Spectroscopy and Radiative Transfer**, vol. 38, no. 6, pp. 413-426 (1987).

42 Subramaniam, S., and Mengüç, M. P., "*Solution of the Inverse Radiation Problem for Inhomogeneous and Anisotropically Scattering Media Using a Monte Carlo Technique,*" **International Journal of Heat and Mass Transfer**, vol. 34, no. 1, pp. 253-266 (1991).

43 Kamiuto, K., "*An Iterative Method for Inverse Scattering Problems,*" **Journal of Quantitative Spectroscopy and Radiative Transfer**, vol. 49, no. 1, pp. 1-13 (1993).

44 Edlen, B., "*The Refractive Index of Air,*" **Metrologia**, vol. 2, no. 2, pp. 71-80 (1966).

45 Voytishek, A. V., and Prigarin, S. M., "*The Functional Convergence of the Limits and Models in the Monte Carlo Method,*" **Computational Maths Math. Phys.**, vol. 32, no. 10, pp. 1471-1479 (1992).

46 Shamsundar, N., Sparrow, E. M., and Heinisch, R. P., "*Monte Carlo Radiation Solutions - Effect of Energy Partitioning and Number of Rays,*" **International Journal of Heat and Mass Transfer**, vol. 16, pp. 690-694 (1973).

47 Maltby, J. D., and Burns, P. J., "*Performance, Accuracy, and Convergence in a Three-Dimensional Monte Carlo Radiative Heat Transfer Simulation,*" **Numerical Heat Transfer**, Part B, vol.19, pp. 191-209 (1991).

48 Niederreitter, H., "*Quasi-Monte Carlo Methods and Pseudo-Random Numbers,*" **Bulletin of the American Mathematical Society**, vol. 64, no. 6, pp. 957-986 (1978).

References   170

49   O'Brien, D. M., "*Accelerated Quasi Monte Carlo Integration of the Radiative Transfer Equation*," **Journal of Quantitative Spectroscopy and Radiative Transfer**, vol. 48, no. 1, pp. 41-59 (1992).

50   Woodwark, J. R., "*A Multiprocessor Architecture for Viewing Solid Models*," **Displays**, vol. 6, pp. 97-103 (1984).

51   Cleary, J. G., Wyvill, B. M., Birtwistle, G. M., and Vatti, R., "*Multiprocessor Ray Tracing*," **Computer Graphics Forum 5**, vol. 5, pp. 3-12 (1986).

52   Gaudet, B., Hobson, R., Chilka, P., and Calvert, T., "*Multiprocessor Experiments for High-Speed Ray Tracing*," **ACM Transactions on Graphics**, vol. 7, no. 3, pp. 151-179 (1988).

53   Amdahl, G., "*The Validity of the Single Processor Approach to Achieving Large Scale Computing Capabilities*," **AFIPS Conference Proceedings 30**, pp. 483-485 (1967).

54   Lukkien, J. J., **On the Performance of Parallel Algorithms**, Technical Report, Department of Computing Science, Groningen University, the Netherlands (1989).

55   Lee, M. E., Redner, R. A., and Uselton, S. P., "*Statistically Optimized Sampling for Distributed Ray Tracing*," **SIGGRAPH 85**, vol. 19, no. 3, pp. 61-65 (1985).

56   İşler, V., and Özgüç, B. "*Fast Ray Tracing 3D Models*," **Computer & Graphics**, vol. 15, no. 2, pp. 205-216 (1991).

57   Fujimoto, A., Tanaka, T., and Iwata, K., "*ARTS: Accelerated Ray-Tracing System*," **IEEE Computer Graphics & Applications**, vol. 6, no. 4, pp. 16-26 (1986).

58   Yagel, R., Cohen, D., and Kaufman, A., "*Discrete Ray Tracing*," **IEEE Computer Graphics & Applications**, vol. 12, no. 9, pp. 19-28 (1992).

59   Glassner, A. S., "*Space Subdivision for Fast Ray Tracing*," **IEEE Computer Graphics and Applications**, vol. 10, no. 4, pp. 15-22 (1984).

60   Arvo, J., and Kirk, D., "*Fast Ray Tracing by Ray Classification*," **Computer Graphics**, vol. 21, no. 4, pp. 55-63 (1987).

61   Woo, A., *"Efficient Shadow Computations In Ray Tracing,"* **IEEE Computer Graphics & Applications,** vol. 13, no. 6, pp. 78-83 (1993).

62   Spackman, J., and Willis, P., *"The Smart Navigation of a Ray Through an Oct-Tree,"* **Computer & Graphics,** vol. 15, no. 2, pp. 185-194 (1991).

63   Levoy, M., *"Efficient Ray Tracing of Volume Data,"* **ACM Transactions on Graphics,** vol. 9, no. 3, pp. 245-261 (1990).

64   Pitot, P., *"The Voxar Project,"* **IEEE Computer Graphics & Applications,** vol. 13, no. 1, pp. 27-33 (1993).

65   McNeill, M. D. J., Shah, B. C., Hebert, M. P. Lister, P. F., and Grimsdale, R. L., *"Performance of Space Subdivision Techniques in Ray Tracing,"* **Computer Graphics Forum,** vol. 11, no 4., pp. 213-220 (1992).

66   Akimoto, T., Mase, K., and Suenaga, Y., *"Pixel-Selected Ray Tracing,"* **IEEE Computer Graphics & Applications,** vol. 11, no. 9, pp. 14-22 (1991).

67   Cook, R. L., *"Stochastic Sampling in Computer Graphics,"* **ACM Transactions on Graphics,** vol. 5, no. 1, pp. 51-72 (1986).

68   Mitchell, D. P., *"Spectrally Optimal Sampling for Distributed Ray Tracing,"* **Computer Graphics,** vol. 25, no. 4, pp. 157-164 (1991).

69   Kirk, D., and Arvo, J., *"Unbiased Sampling Techniques for Image Synthesis,"* **Computer Graphics,** vol. 25, no. 4, pp. 153-156 (1991).

70   Shirley, P. S., **Physically Based Lighting Calculations for Computer Graphics,** PhD Dissertation, University of Illinois at Urbana-Champaign (1991).

71   Shinya, M., Takahashi, T., and Naito, S., *"Principles and Applications of Pencil Tracing,"* **Computer Graphics,** vol. 21, no. 4, pp. 45-54 (1987).

72   Glassner, A. S., *"Spacetime Ray Tracing for Animation,"* **IEEE Computer Graphics & Applications,** vol. 8, no. 3, pp. 60-70 (1988).

73   Dias, M. L., *"Ray Tracing Interference Color,"* **IEEE Computer Graphics & Applications,** vol. 11, no. 3, pp. 54-59 (1991).

74   Heckbert, P. S., *"Ray Tracing Jell-O® Brand Gelatin,"* **Computer Graphics**, vol. 21, no. 4, pp. 73-74 (1987).

75   Sweeney, M. A. J., and Bartels, R. H., *"Ray Tracing Free-Form B-Spline Surfaces,"* **IEEE Trans Computer Graphics and Applications**, vol. 5, no. 2, pp. 41-49 (1986).

76   Anonymous, **Model Specification, FAA Category II, Class B Gas Turbine Auxiliary Power Unit Airesearch Model GTCG30-142C Part 380946-1**, Report Number SC-5899, Airesearch Manufacturing Company of Arizona (1969).

77   Hardman, R. R., **Infrared Imaging: A Proposed Validation Technique for Computational Fluid Dynamics Codes Used in STOVL Applications**, Master of Science Thesis, Virginia Polytechnic Institute and State University (1990).

78   Nelson, E. L., **Temperature, Pressure, and Infrared Image Survey of an Axisymmetric Heated Exhaust Flow**, PhD Dissertation, Virginia Polytechnic Institute and State University (1994).

79   Stephens, S. E., *"Evaluation of Air Injection in a Sudden Expansion Diffuser for Real-Time Afterburning Application,"* Arnold Engineering Development Center Technical Report 90-29 (1991).

80   Dinger, A. S., Dunham, E. W., Erickson, E. F., Klotz, D., and Yee, R., *"Evaluation of Jet-Exhaust Scattering for the Aft Telescope Configuration on SOFIA,"* Internal Report, NASA Ames Research Center, (1992).

81   Atwood, C. A., *"Flowfield Simulation Above SOFIA Airborne Observatory,"* AIAA Paper 92-0656, Presented at the AIAA 30th Aerospace Sciences Meeting, January (1992)

82   Atwood, C. A., and Van Dalsem, W. R., *"Flowfield Simultation about the Stratospheric Observatory For Infrared Astronomy,"* **AIAA Journal of Aircraft**, pp. 719-727 (1993).

83   Anonymous, **User's Manual to Version 4.2.1**, Agema Infrared Systems, Danderyd Sweden, (1991).

84   Withington, J. P., **A Time Accurate Numerical Method for Chemically Reacting Flows at All Mach Numbers**, PhD Dissertation, Penn State University (1994).

85   Halton, J. H., *"On the Efficiency of Certain Quasi-Random Sequences of Points in Evaluating Multi-Dimensional Integrals,"* **Numerische Mathematik**, vol. 2, pp. 84-90 (1960).

86   Halton, J. H., and Smith, G. B., *"Radical Inverse Quasi-Random Point Sequence,"* **Communications of the ACM**, vol. 7, no. 12, pp. 701-702 (1964).

87   Gordon, W. J., and Wixom, J. A., *"Shepard's Method for "Metric Interpolation" to Bivariate and Multivariate Interpolation,"* **Mathematics of Computation**, vol. 32, no. 141, pp. 253-264 (1978).

88   Nielson, G. M., *"Scattered Data Modeling,"* **IEEE Computer Graphics & Applications**, vol. 13, no. 1, pp. 60-70 (1993).

89   Anonymous, **FAST User's Manual Version Beta 2.0**, Sterling Federal Systems, Inc., Sterling RND Report 01-011 (1991).

90   Hampe, A. F., NASA Ames Research Center, Moffett Field, CA, Personal Communication (1994).

91   Beguelin, A., Dongarra, J. J., Geist, G. A., Manchek, R., and Sunderam, V. S., **A User's Guide to PVM Parallel Virtual Machine**, Oak Ridge National Laboratory, Technical Report ORNL/TM-11826 (1991).

APPENDIX

DERIVATION OF RECIPROCITY RELATIONS

This appendix presents the derivation of the reciprocity relations. All the reciprocity relations are derived here for the sake of completeness, even though the current work only requires the reciprocity relations involving exchanges between finite surfaces, and between finite surfaces and finite volumes. This derivation is based on the assumptions that the index of refraction of the transmitting medium is sufficiently close to unity, that the properties are uniform across any element, and that local thermodynamic equilibrium holds.

A.1 Cumulative Energies

A differential surface element, i , with area dA_i at temperature T_i emits in direction (θ, ϕ) within a wavelength band $d\lambda$ centered about wavelength λ

$$d^3E_i = \epsilon'_{\lambda i}(\lambda, \theta, \phi, T_i) i'_{\lambda b}(\lambda, T_i) \cos\theta_i \, dA_i \, d\omega_i \, d\lambda \ . \tag{34}$$

The cumulative energy emitted by the differential element is obtained by integrating equation 34 over all possible solid angles, $d\omega_i$. For a solid surface this leads to an integration over 2π sr producing

$$d^2E_i = \int_{2\pi} \epsilon'_{\lambda i}(\lambda,\theta,\phi,T_i)i'_{\lambda b}(\lambda,T_i)\cos\theta_i d\omega_i dA_i d\lambda \quad . \tag{35}$$

At this point, the *spectral, hemispherical emissivity* is introduced as

$$\pi \epsilon_{\lambda i}(\lambda,T_i) = \int_{2\pi} \epsilon'_{\lambda i}(\lambda,\theta,\phi,T_i)\cos\theta_i d\omega_i \quad . \tag{36}$$

Substituting the spectral, hemispherical emissivity into equation 35, the cumulative energy emitted by a differential surface element in all directions in wavelength band $d\lambda$ centered about wavelength λ with area dA_i is

$$d^2E_i = \pi \epsilon_{\lambda i}(\lambda,T_i)i'_{\lambda b}(\lambda,T_i)dA_i d\lambda \quad . \tag{37}$$

The cumulative energy emitted from a finite surface element is obtained by integrating equation 37 over the area, A_i, to obtain

$$dE_i = \int_{A_i} \pi \epsilon_{\lambda i}(\lambda,T_i)i'_{\lambda b}(\lambda,T_i)dA_i d\lambda = \pi \epsilon_{\lambda i}(\lambda,T_i)i'_{\lambda b}(\lambda,T_i)A_i d\lambda \quad . \tag{38}$$

A differential volume element, m, emits in direction (θ,ϕ) within a wavelength band $d\lambda$ centered about wavelength λ

$$d^3E_m = \int_{dV_m \rightarrow dA_{pm}} \epsilon_{\lambda m}(\lambda,C_m,T_m,P_m,s_m)i'_{\lambda b}(\lambda,T_m)dA_m d\omega_m d\lambda \quad . \tag{39}$$

The integral in equation 39 is across the area formed by the projection of the volume in direction (θ,ϕ). Note that the path length, s, depends on the direction of the solid angle as well as the location of the integrating differential area dA_m. The cumulative energy from a volume is found by integrating over the possible 4π sr. Thus the cumulative energy for a differential volume is

$$d^2E_m = \int\int_{4\pi'dV_m \to dA_{pm}} \epsilon_{\lambda m}(\lambda, C_m, T_m, P_m, s_m)i'_{\lambda b}(\lambda, T_m)dA_m d\omega_m d\lambda . \qquad (40)$$

The cumulative energy for a finite volume with projected area along direction (θ,ϕ), A_{pm}, is

$$dE_m = \int\int_{4\pi'V_m \to A_{pm}} \epsilon_{\lambda m}(\lambda, C_m, T_m, P_m, s_m)i'_{\lambda b}(\lambda, T_m)dA_m d\omega_m d\lambda . \qquad (41)$$

A.2 Surface-to-Surface Reciprocity Relations

Figure A.1, located at the end of this appendix, shows two surfaces and the associated variables used in this section. Note that $d\omega_{ij}$ is the solid angle subtended by the projection of the differential area, dA_j, at the surface element i. In other words

$$d\omega_{ij} = \frac{\cos\theta_j dA_j}{R^2} . \qquad (42)$$

The spectral energy emitted by the first differential area, dA_i, that arrives at the second differential area, dA_j, is

$$d^3E_{ij} = \alpha'_{\lambda j}(\lambda,\theta_j,\phi_j,T_j)[\ast]\epsilon'_{\lambda i}(\lambda,\theta_i,\phi_i,T_i)i'_{\lambda b}(\lambda,T_i)\cos\theta_i dA_i d\omega_{ij} d\lambda \quad . \tag{43}$$

The factor [$\ast$] reminds us that in the general case the energy emitted by surface element i is subject to reflection, transmission, and scattering before being absorbed by surface element j. For example if the energy is simply transmitted through an attenuating medium then

$$[\ast] = \tau_\lambda(\lambda,C_m,T_m,P_m,s_m) \quad , \tag{44}$$

whereas if the energy is reflected by an intervening surface element while traveling though an attenuating medium,

$$[\ast] = \tau_\lambda(\lambda,C_m,T_m,P_m,s_m)\rho''_\lambda(\lambda,\theta_i,\phi_i,\theta_j,\phi_j,T_n)\tau_\lambda(\lambda,C_q,T_q,P_q,s_q) \quad . \tag{45}$$

No matter how complex [$\ast$] becomes, it retains the following properties: it is a valid path for the reverse direction, and it has a reciprocity relation between the initial emissive direction (θ_i,ϕ_i) and the final absorptive direction (θ_j,ϕ_j). The concept behind these properties for [$\ast$] is presented in Chapters 2 and 3.

Equation 43 may be rewritten by using equation 41 and Kirchhoff' law from Chapter 2 to obtain

$$d^3E_{ij} = i'_{\lambda b}(\lambda, T_i) d\lambda [\ast] \alpha'_{\lambda i}(\lambda, \theta_i, \phi_i, T_i) \alpha'_{\lambda j}(\lambda, \theta_j, \phi_j, T_j) \frac{\cos \theta_i \cos \theta_j}{R^2} dA_i dA_j . \qquad (46)$$

Similarly the spectral energy emitted from the second differential area element, dA_j, that is absorbed by the first differential element, dA_i, is

$$d^3E_{ji} = \alpha'_{\lambda i}(\lambda, \theta_i, \phi_j, T_i) [\ast] \epsilon'_{\lambda j}(\lambda, \theta_j, \phi_j, T_j) i'_{\lambda b}(\lambda, T_j) \cos \theta_j dA_j d\omega_{ji} d\lambda , \qquad (47)$$

which may be rewritten as

$$d^3E_{ji} = i'_{\lambda b}(\lambda, T_j) d\lambda [\ast] \alpha'_{\lambda i}(\lambda, \theta_i, \phi_i, T_i) \alpha'_{\lambda j}(\lambda, \theta_j, \phi_j, T_j) \frac{\cos \theta_i \cos \theta_j}{R^2} dA_i dA_j . \qquad (48)$$

At this point we introduce the differential distribution factors, dD_{ij} and dD_{ji}, by

$$dD_{ij} \equiv \frac{d^3E_{ij}}{d^2E_i} = \frac{i'_{\lambda b}(\lambda, T_i) d\lambda [\ast] \alpha'_{\lambda i}(\lambda, \theta_i, \phi_i, T_i) \alpha'_{\lambda j}(\lambda, \theta_j, \phi_j, T_j) \frac{\cos \theta_i \cos \theta_j}{R^2} dA_i dA_j}{i'_{\lambda b}(\lambda, T_i) d\lambda \pi \epsilon_{\lambda i}(\lambda, T_i) dA_i}, \qquad (49)$$

and

$$dD_{ji} \equiv \frac{d^3E_{ji}}{d^2E_j} = \frac{i'_{\lambda b}(\lambda, T_j) d\lambda [\ast] \alpha'_{\lambda i}(\lambda, \theta_i, \phi_i, T_i) \alpha'_{\lambda j}(\lambda, \theta_j, \phi_j, T_j) \frac{\cos \theta_i \cos \theta_j}{R^2} dA_i dA_j}{i'_{\lambda b}(\lambda, T_j) d\lambda \pi \epsilon_{\lambda j}(\lambda, T_j) dA_j}. \qquad (50)$$

Appendix 179

Eliminating the common factors between equations 49 and 50 leads to the reciprocity relation between two differential surface elements:

$$dD_{ij}\epsilon_{\lambda i}(\lambda,T_i)dA_i = dD_{ji}\epsilon_{\lambda j}(\lambda,T_j)dA_j \; . \tag{51}$$

If one of the surface elements is finite instead of differential, the reciprocity relation may be obtained by integrating the above equations. For this section we assume that surface element j is the finite element. Equation 46 becomes

$$d^2E_{ij}=\int_{A_j} d^3E_{ij}=i'_{\lambda b}(\lambda,T_i)d\lambda\int_{A_j}[*]\alpha'_{\lambda i}(\lambda,\theta_i,\phi_i,T_i)\alpha'_{\lambda j}(\lambda,\theta_j,\phi_j,T_j)\frac{\cos\theta_i\cos\theta_j}{R^2}dA_i\,dA_j, \tag{52}$$

and equation 48 becomes

$$d^2E_{ji}=\int_{A_j} d^3E_{ji}=i'_{\lambda b}(\lambda,T_j)d\lambda\int_{A_j}[*]\alpha'_{\lambda i}(\lambda,\theta_i,\phi_i,T_i)\alpha'_{\lambda j}(\lambda,\theta_j,\phi_j,T_j)\frac{\cos\theta_i\cos\theta_j}{R^2}dA_i\,dA_j. \tag{53}$$

Incorporating these equations into the distribution factor definitions leads to

$$D_{ij}\equiv\frac{d^2E_{ij}}{d^2E_i}=\frac{i'_{\lambda b}(\lambda,T_i)d\lambda\int_{A_j}[*]\alpha'_{\lambda i}(\lambda,\theta_i,\phi_i,T_i)\alpha'_{\lambda j}(\lambda,\theta_j,\phi_j,T_j)\dfrac{\cos\theta_i\cos\theta_j}{R^2}dA_i\,dA_j}{i'_{\lambda b}(\lambda,T_i)d\lambda\,\pi\,\epsilon_{\lambda i}(\lambda,T_i)dA_i}, \tag{54}$$

and

$$dD_{ji} \equiv \frac{d^2E_{ji}}{dE_j} = \frac{i'_{\lambda b}(\lambda,T_j)d\lambda \int_{A_j} [*] \alpha'_{\lambda i}(\lambda,\theta_i,\phi_i,T_i)\alpha'_{\lambda j}(\lambda,\theta_j,\phi_j,T_j)\dfrac{\cos\theta_i\cos\theta_j}{R^2}dA_i \, dA_j}{i'_{\lambda b}(\lambda,T_j)d\lambda \, \pi \, \epsilon_{\lambda j}(\lambda,T_j)A_j} . \tag{55}$$

Eliminating the common factors in equations 54 and 55 leads to the reciprocity relation between a differential surface element i and a finite surface element j:

$$D_{ij}\epsilon_{\lambda i}(\lambda,T_i)dA_i = dD_{ji}\epsilon_{\lambda j}(\lambda,T_j)A_j . \tag{56}$$

Finally integrating over the remaining area, A_i, leads to the reciprocity relation between finite surface areas. Equation 52 becomes

$$dE_{ij} = \int_{A_i} d^2E_{ij} = i'_{\lambda b}(\lambda,T_i)d\lambda \int_{A_j}\int_{A_i} [*] \alpha'_{\lambda i}(\lambda,\theta_i,\phi_i,T_i)\alpha'_{\lambda j}(\lambda,\theta_j,\phi_j,T_j)\frac{\cos\theta_i\cos\theta_j}{R^2}dA_i \, dA_j, \tag{57}$$

and equation 53 becomes

$$dE_{ji} = \int_{A_i} d^2E_{ji} = i'_{\lambda b}(\lambda,T_j)d\lambda \int_{A_j}\int_{A_i} [*] \alpha'_{\lambda i}(\lambda,\theta_i,\phi_i,T_i)\alpha'_{\lambda j}(\lambda,\theta_j,\phi_j,T_j)\frac{\cos\theta_i\cos\theta_j}{R^2}dA_i \, dA_j. \tag{58}$$

Propagating these changes through the distribution factor definitions leads to

$$D_{ij} \equiv \frac{dE_{ij}}{dE_i} = \frac{i'_{\lambda b}(\lambda,T_i)d\lambda \iint_{A_j,A_i}[\ast]\alpha'_{\lambda i}(\lambda,\theta_i,\phi_i,T_i)\alpha'_{\lambda j}(\lambda,\theta_j,\phi_j,T_j)\dfrac{\cos\theta_i\cos\theta_j}{R^2}dA_i\,dA_j}{i'_{\lambda b}(\lambda,T_i)d\lambda\,\pi\,\epsilon_{\lambda i}(\lambda,T_i)A_i} \qquad (59)$$

and

$$D_{ji} \equiv \frac{dE_{ji}}{dE_j} = \frac{i'_{\lambda b}(\lambda,T_j)d\lambda \iint_{A_j,A_i}[\ast]\alpha'_{\lambda i}(\lambda,\theta_i,\phi_i,T_i)\alpha'_{\lambda j}(\lambda,\theta_j,\phi_j,T_j)\dfrac{\cos\theta_i\cos\theta_j}{R^2}dA_i\,dA_j}{i'_{\lambda b}(\lambda,T_j)d\lambda\,\pi\,\epsilon_{\lambda j}(\lambda,T_j)A_j} \qquad (60)$$

Finally, eliminating the common factors between equations 59 and 60 produces the reciprocity relation between finite surface areas:

$$D_{ij}\epsilon_{\lambda i}(\lambda,T_i)A_i = D_{ji}\epsilon_{\lambda j}(\lambda,T_j)A_j \,. \qquad (61)$$

A.3 Surface-to-Volume Reciprocity Relations

Figure A.2, located at the end of the appendix, shows a surface and volume element and the definition of the variables used in this section. The spectral energy emitted from differential surface element i and absorbed by differential volume element m is

$$d^3E_{im} = \int_{dV_m \to dA_{pm}} \alpha_{\lambda m}(\lambda, C_m, T_m, P_m, s_m)[*]\epsilon'_{\lambda i}(\lambda, \theta_i, \phi_i, T_i)i'_{\lambda b}(\lambda, T_i)\cos\theta_i dA_i d\omega_{im} d\lambda, \quad (62)$$

which may be rewritten as

$$d^3E_{im} = i'_{\lambda b}(\lambda, T_i)d\lambda \int_{dV_M \to dA_{pm}} [*]\alpha'_{\lambda i}(\lambda, \theta_i, \phi_i, T_i)\alpha_{\lambda m}(\lambda, C_m, T_m, P_m, s_m)\frac{\cos\theta_i}{R^2}dA_i dA_m. \quad (63)$$

Similarly the spectral energy emitted by the volume element that is absorbed by the surface element can be written as

$$d^3E_{mi} = i'_{\lambda b}(\lambda, T_m)d\lambda \int_{dV_m \to dA_{pm}} [*]\alpha'_{\lambda i}(\lambda, \theta_i, \phi_i, T_i)\alpha_{\lambda m}(\lambda, C_m, T_m, P_m, s_m)\frac{\cos\theta_i}{R^2}dA_i dA_m. \quad (64)$$

Introducing the distribution factors leads to

$$dD_{im} \equiv \frac{d^3E_{im}}{d^2E_i} =$$

$$\frac{i'_{\lambda b}(\lambda, T_i)d\lambda \int_{dV_m \to dA_{pm}} [*]\alpha'_{\lambda i}(\lambda, \theta_i, \phi_i, T_i)\alpha_{\lambda m}(\lambda, C_m, T_m, P_m, s_m)\frac{\cos\theta_i}{R^2}dA_i dA_m}{i'_{\lambda b}(\lambda, T_i)d\lambda \, \pi \, \epsilon_{\lambda i}(\lambda, T_i)dA_i}, \quad (65)$$

and

$$dD_{mi} \equiv \frac{d^3E_{mi}}{d^2E_m} =$$

$$\frac{i'_{\lambda b}(\lambda,T_m)d\lambda \int_{dV_m \to dA_{pm}} [*]\alpha'_{\lambda i}(\lambda,\theta_i,\phi_i,T_i)\alpha_{\lambda m}(\lambda,C_m,T_m,P_m,s_m)\frac{\cos\theta_i}{R^2}dA_i dA_m}{i'_{\lambda b}(\lambda,T_m)d\lambda \int\int_{4\pi\, dV_m \to dA_{pm}} \epsilon_{\lambda m}(\lambda,C_m,T_m,P_m,s[d,p])dA_m d\omega_m} \quad . \tag{65}$$

Eliminating the common factors leads to the reciprocity relation between a differential surface element and a differential volume element:

$$\pi\, dD_{im}\epsilon_{\lambda i}(\lambda,T_i)dA_i = dD_{mi}\int\int_{4\pi\, dV_m \to dA_{pm}} \epsilon_{\lambda m}(\lambda,C_m,T_m,P_m,s_m)dA_m d\omega_m \quad . \tag{66}$$

The reciprocity relation between a differential area element and a finite volume element may be obtained by using the projected area of the finite volume along the direction of the integration. The spectral energies then become

$$d^2E_{im} = i'_{\lambda b}(\lambda,T_i)d\lambda \int_{V_m \to A_{pm}} [*]\alpha'_{\lambda i}(\lambda,\theta_i,\phi_i,T_i)\alpha_{\lambda m}(\lambda,C_m,T_m,P_m,s_m)\frac{\cos\theta_i}{R^2}dA_i dA_m, \tag{67}$$

and

$$d^2E_{mi} = i'_{\lambda b}(\lambda,T_m)d\lambda \int_{V_m \to A_{pm}} [*] \alpha'_{\lambda i}(\lambda,\theta_i,\phi_i,T_i) \alpha_{\lambda m}(\lambda,C_m,T_m,P_m,s_m) \frac{\cos\theta_i}{R^2} dA_i dA_m \ . \ (68)$$

The reciprocity relation between a differential surface element and a finite volume is

$$\pi D_{im}\epsilon_{\lambda i}(\lambda,T_i)dA_i = dD_{mi} \int\int_{4\pi} \int_{V_m \to A_{pm}} \epsilon_{\lambda m}(\lambda,C_m,T_m,P_m,s_m)dA_m d\omega_m \ , \qquad (69)$$

while similar derivations produce the reciprocity relation between a finite surface element and a differential volume as

$$\pi dD_{im}\epsilon_{\lambda i}(\lambda,T_i)A_i = D_{mi} \int\int_{4\pi} \int_{dV_m \to dA_{pm}} \epsilon_{\lambda m}(\lambda,C_m,T_m,P_m,s_m)dA_m d\omega_m \ . \qquad (70)$$

Finally the reciprocity relation between a finite surface element and a finite volume element is

$$\pi D_{im}\epsilon_{\lambda i}(\lambda,T_i)A_i = D_{mi} \int\int_{4\pi} \int_{V_m \to A_{pm}} \epsilon_{\lambda m}(\lambda,C_m,T_m,P_m,s_m)dA_m d\omega_m \ . \qquad (71)$$

A.4 Volume-to-Volume Reciprocity Relations

The spectral energy emitted by one differential volume element and absorbed by a second differential volume element is

$$d^3E_{mn} = i'_{\lambda b}(\lambda, T_m) d\lambda \int_{dV_n \to dA_{pn}} \int_{dV_m \to dA_{pm}} [\ast] \alpha_{\lambda m}(\lambda, C_m, T_m, P_m, s_m)$$

$$\alpha_{\lambda n}(\lambda, C_n, T_n, P_n, s_n) \frac{1}{R^2} dA_m dA_n \ . \tag{72}$$

Using a similar equation for the spectral energy emitted by the second differential volume element and absorbed by the first differential volume element, and then using the distribution factor definitions produces the reciprocity relation between two differential volume elements as

$$dD_{mn} \int\int_{4\pi\, dV_m \to dA_{pm}} \epsilon_{\lambda m}(\lambda, C_m, T_m, P_m, s_m) dA_m d\omega_m =$$

$$dD_{nm} \int\int_{4\pi\, dV_n \to dA_{pn}} \epsilon_{\lambda n}(\lambda, C_n, T_n, P_n, s_n) dA_n d\omega_n \ . \tag{73}$$

Following derivations similar to those described above leads to the reciprocity relation between a differential volume and a finite volume being

$$D_{mn} \int\int_{4\pi\, dV_m \to dA_{pm}} \epsilon_{\lambda m}(\lambda, C_m, T_m, P_m, s_m) dA_m d\omega_m =$$

$$dD_{nm} \int\int_{4\pi\, V_n \to A_{pn}} \epsilon_{\lambda n}(\lambda, C_n, T_n, P_n, s_n) dA_n d\omega_n \ , \tag{74}$$

Appendix 186

and the reciprocity relation between two finite volumes being

$$D_{mn} \int\int_{4\pi \, V_m \to A_{pm}} \epsilon_{\lambda m}(\lambda, C_m, T_m, P_m, s_m) dA_m d\omega_m =$$

$$D_{nm} \int\int_{4\pi \, V_n \to A_{pn}} \epsilon_{\lambda n}(\lambda, C_n, T_n, P_n, s_n) dA_n d\omega_n \; . \tag{76}$$

A.5 Simplified Relations

If the modeled volume elements are all known to be optically thin, then the above relations can be greatly simplified. A volume is considered optically thin when $a_\lambda ds \ll 1.0$. When this relation holds then the emittance from equation 7 in Chapter 2 can be written as

$$\epsilon_{\lambda m}(\lambda, C_m, T_m, P_m, ds) = 1.0 - \text{EXP} \, (-a_{\lambda m}[\lambda, C_m, T_m, P_m) ds] \approx a_{\lambda m}(\lambda, C_m, T_m, P_m) ds \; . \tag{77}$$

Substituting equation 76 into equations 38 produces

$$d^3 E_m = a_{\lambda m}(\lambda, C_m, T_m, P_m) i'_{\lambda b}(\lambda, T_m) \int_{dV_m \to dA_{pm}} ds \, dA_m d\omega_m d\lambda \; , \tag{78}$$

which reduces to

$$d^3 E_m = a_{\lambda m}(\lambda, C_m, T_m, P_m) i'_{\lambda b}(\lambda, T_m) dV_m d\omega_m d\lambda \; . \tag{79}$$

The cumulative spectral energy emitted by the differential volume element is then

$$d^2E_m = 4\pi\, a_{\lambda m}(\lambda, C_m, T_m, P_m) i'_{\lambda b}(\lambda, T_m) dV_m d\lambda \ , \tag{80}$$

and by a finite volume is

$$dE_m = 4\pi\, a_{\lambda m}(\lambda, C_m, T_m, P_m) i'_{\lambda b}(\lambda, T_m) V_m d\lambda \ . \tag{81}$$

Pursuing this change through the reciprocity relations leads to the reciprocity relation between a finite surface element and a finite volume element as

$$D_{im}\epsilon_{\lambda i}(\lambda, T_i) A_i = 4 D_{mi} a_{\lambda m}(\lambda, C_m, T_m, P_m) V_m \ , \tag{82}$$

and between two finite volume elements as

$$D_{mn} a_{\lambda m}(\lambda, C_m, T_m, P_m) V_m = D_{nm} a_{\lambda n}(\lambda, C_n, T_n, P_n) V_n \ . \tag{83}$$

All the other reciprocity relations involving volumes elements are similarly simplified, but it must be kept in mind that this simplification of the reciprocity relations only applies for optically thin volumes.
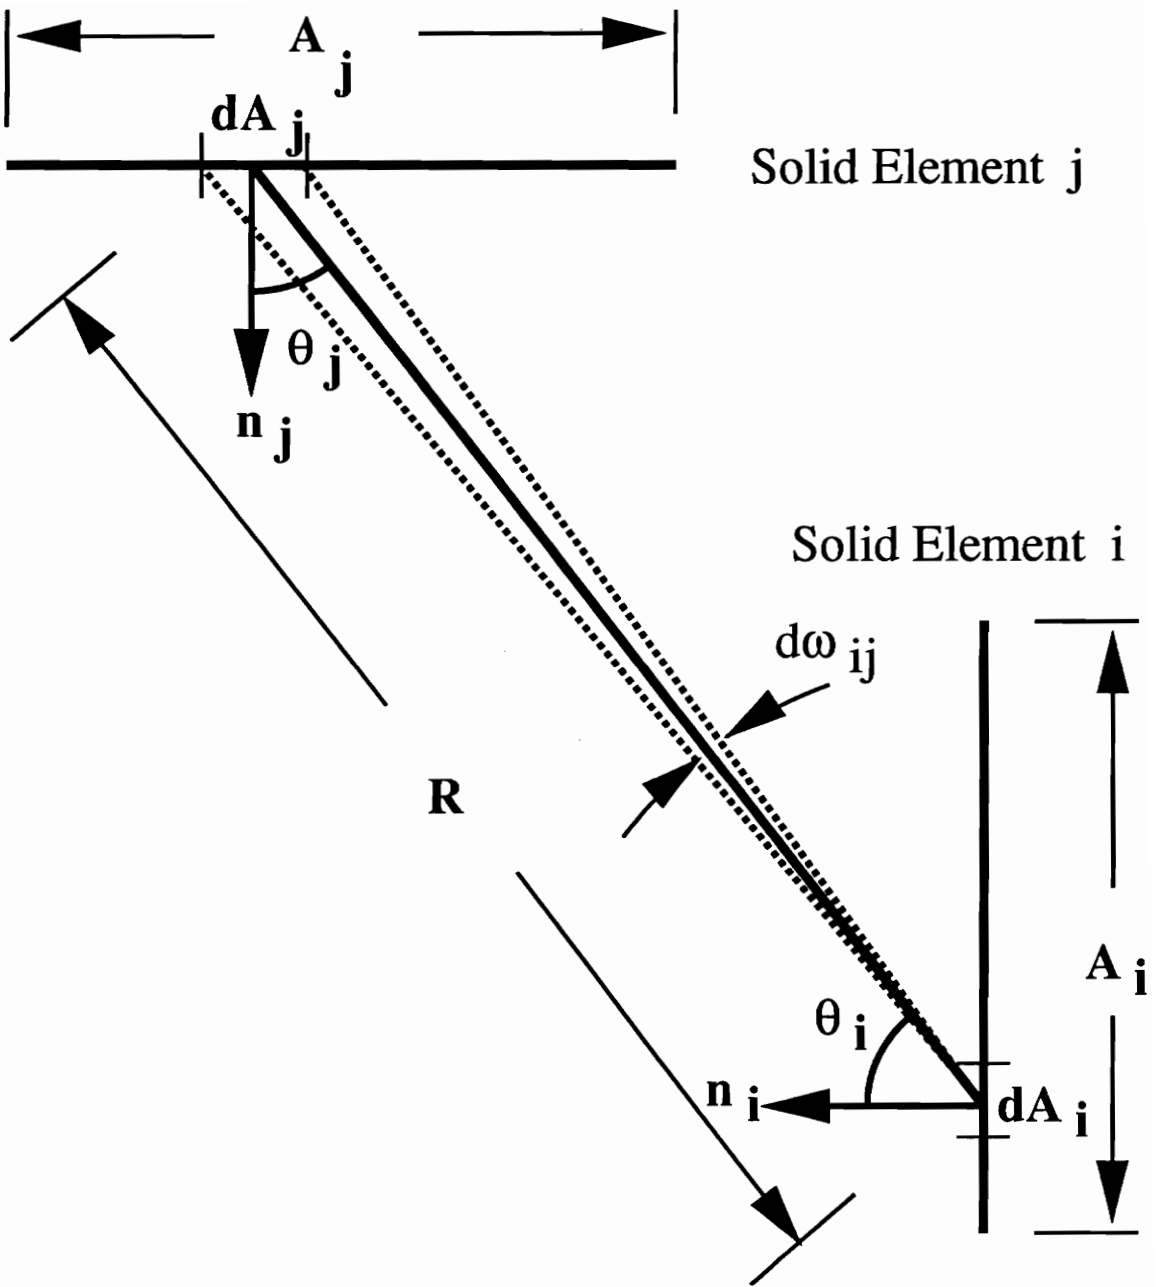
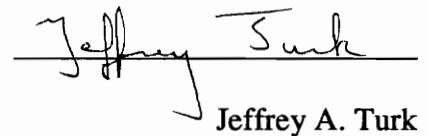**Figure A.1**   Representation of Surface-to-Surface Radiant Energy Exchange.
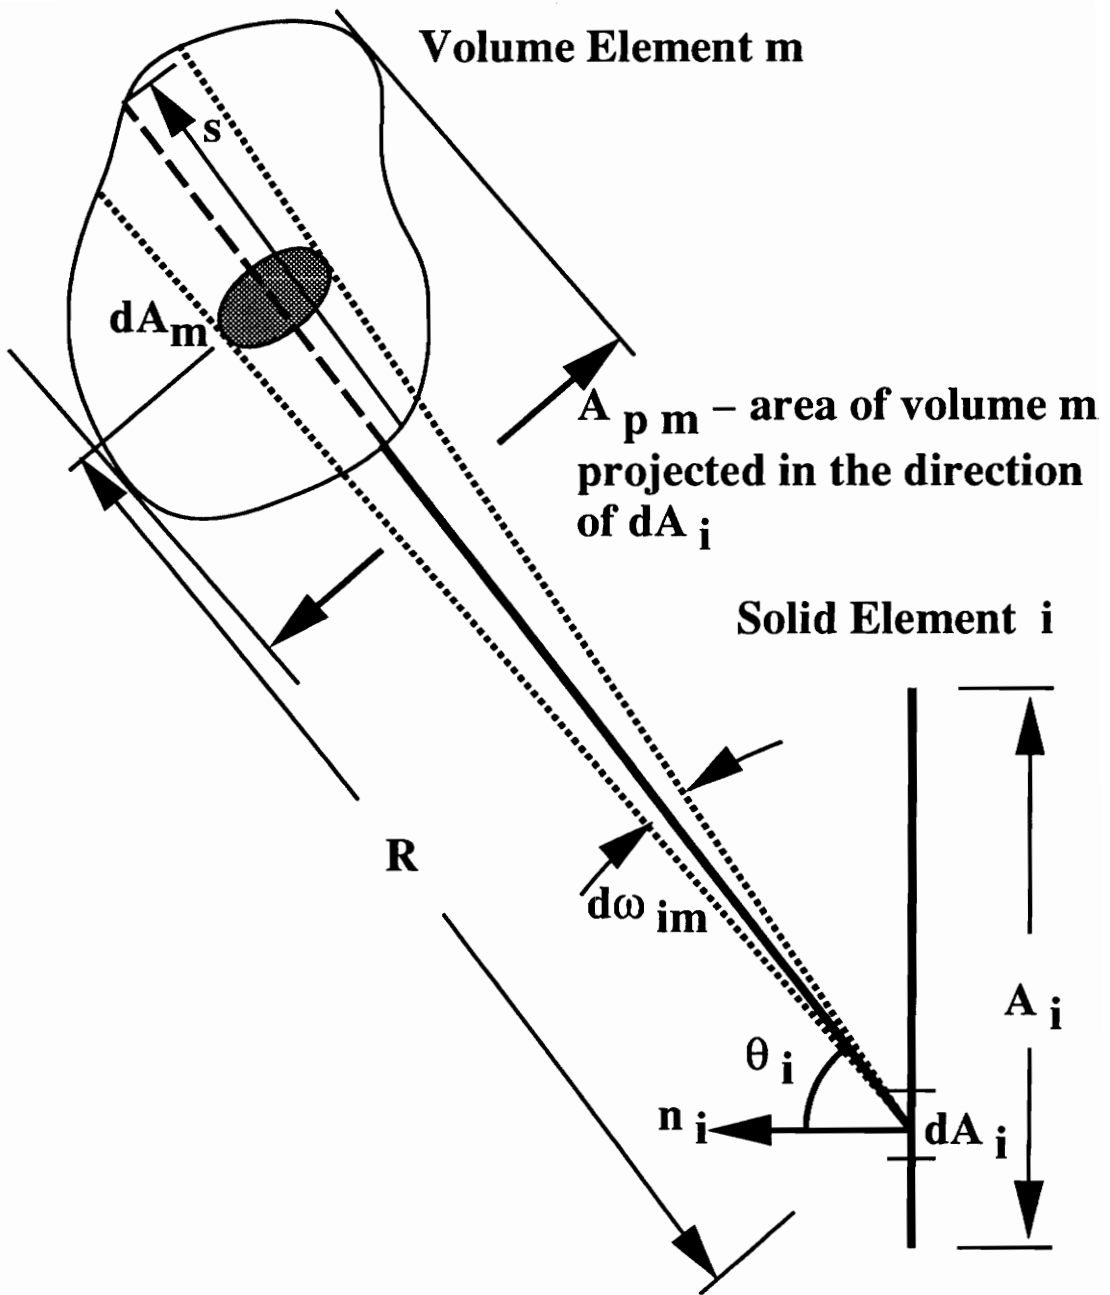
Figure A.2   Representation of Surface-to-Volume Radiant Energy Exchange.

VITA

Jeffrey Turk was born on May 10, 1968, in Greenville, South Carolina. He graduated as Valedictorian from Strongsville Senior High School in Strongsville, Ohio, in 1986 and attended Virginia Tech that fall. While at Tech, he majored in aerospace engineering. He studied abroad during his senior year at *l' Université de Compiègne* and returned to graduate Summa Cum Laude. He is a member of Tau Beta Pi, Phi Beta Kappa, and Sigma Gamma Tau honor societies.

He returned to Virginia Tech in June of 1990 in pursuit of the Ph.D. in aerospace engineering. During the course of the dissertation work he spent a couple of years at NASA Ames Research Center as a graduate coop.

Jeffrey A. Turk