

**INFORMATION ABOUT PRINCIPAL INVESTIGATORS/PROJECT DIRECTORS(PI/PD) and
co-PRINCIPAL INVESTIGATORS/co-PROJECT DIRECTORS**

Submit only ONE copy of this form for each PI/PD and co-PI/PD identified on the proposal. The form(s) should be attached to the original proposal as specified in GPG Section II.B. Submission of this information is voluntary and is not a precondition of award. This information will not be disclosed to external peer reviewers. **DO NOT INCLUDE THIS FORM WITH ANY OF THE OTHER COPIES OF YOUR PROPOSAL AS THIS MAY COMPROMISE THE CONFIDENTIALITY OF THE INFORMATION.**

PI/PD Name: Michael P Frank

Gender: Male Female
Ethnicity: (Choose one response) Hispanic or Latino Not Hispanic or Latino

Race:
(Select one or more)
 American Indian or Alaska Native
 Asian
 Black or African American
 Native Hawaiian or Other Pacific Islander
 White

Disability Status:
(Select one or more)
 Hearing Impairment
 Visual Impairment
 Mobility/Orthopedic Impairment
 Other
 None

Citizenship: (Choose one) U.S. Citizen Permanent Resident Other non-U.S. Citizen

Check here if you do not wish to provide any or all of the above information (excluding PI/PD name):

REQUIRED: Check here if you are currently serving (or have previously served) as a PI, co-PI or PD on any federally funded project

Ethnicity Definition:

Hispanic or Latino. A person of Mexican, Puerto Rican, Cuban, South or Central American, or other Spanish culture or origin, regardless of race.

Race Definitions:

American Indian or Alaska Native. A person having origins in any of the original peoples of North and South America (including Central America), and who maintains tribal affiliation or community attachment.

Asian. A person having origins in any of the original peoples of the Far East, Southeast Asia, or the Indian subcontinent including, for example, Cambodia, China, India, Japan, Korea, Malaysia, Pakistan, the Philippine Islands, Thailand, and Vietnam.

Black or African American. A person having origins in any of the black racial groups of Africa.

Native Hawaiian or Other Pacific Islander. A person having origins in any of the original peoples of Hawaii, Guam, Samoa, or other Pacific Islands.

White. A person having origins in any of the original peoples of Europe, the Middle East, or North Africa.

WHY THIS INFORMATION IS BEING REQUESTED:

The Federal Government has a continuing commitment to monitor the operation of its review and award processes to identify and address any inequities based on gender, race, ethnicity, or disability of its proposed PIs/PDs. To gather information needed for this important task, the proposer should submit a single copy of this form for each identified PI/PD with each proposal. Submission of the requested information is voluntary and will not affect the organization's eligibility for an award. However, information not submitted will seriously undermine the statistical validity, and therefore the usefulness, of information received from others. Any individual not wishing to submit some or all the information should check the box provided for this purpose. (The exceptions are the PI/PD name and the information about prior Federal support, the last question above.)

Collection of this information is authorized by the NSF Act of 1950, as amended, 42 U.S.C. 1861, et seq. Demographic data allows NSF to gauge whether our programs and other opportunities in science and technology are fairly reaching and benefiting everyone regardless of demographic category; to ensure that those in under-represented groups have the same knowledge of and access to programs and other research and educational opportunities; and to assess involvement of international investigators in work supported by NSF. The information may be disclosed to government contractors, experts, volunteers and researchers to complete assigned work; and to other government agencies in order to coordinate and assess programs. The information may be added to the Reviewer file and used to select potential candidates to serve as peer reviewers or advisory committee members. See Systems of Records, NSF-50, "Principal Investigator/Proposal File and Associated Records", 63 Federal Register 267 (January 5, 1998), and NSF-51, "Reviewer/Proposal File and Associated Records", 63 Federal Register 268 (January 5, 1998).

**INFORMATION ABOUT PRINCIPAL INVESTIGATORS/PROJECT DIRECTORS(PI/PD) and
co-PRINCIPAL INVESTIGATORS/co-PROJECT DIRECTORS**

Submit only ONE copy of this form for each PI/PD and co-PI/PD identified on the proposal. The form(s) should be attached to the original proposal as specified in GPG Section II.B. Submission of this information is voluntary and is not a precondition of award. This information will not be disclosed to external peer reviewers. **DO NOT INCLUDE THIS FORM WITH ANY OF THE OTHER COPIES OF YOUR PROPOSAL AS THIS MAY COMPROMISE THE CONFIDENTIALITY OF THE INFORMATION.**

PI/PD Name: Abdelsalam A Helal

Gender: Male Female
Ethnicity: (Choose one response) Hispanic or Latino Not Hispanic or Latino

Race:
(Select one or more)
 American Indian or Alaska Native
 Asian
 Black or African American
 Native Hawaiian or Other Pacific Islander
 White

Disability Status:
(Select one or more)
 Hearing Impairment
 Visual Impairment
 Mobility/Orthopedic Impairment
 Other
 None

Citizenship: (Choose one) U.S. Citizen Permanent Resident Other non-U.S. Citizen

Check here if you do not wish to provide any or all of the above information (excluding PI/PD name):

REQUIRED: Check here if you are currently serving (or have previously served) as a PI, co-PI or PD on any federally funded project

Ethnicity Definition:

Hispanic or Latino. A person of Mexican, Puerto Rican, Cuban, South or Central American, or other Spanish culture or origin, regardless of race.

Race Definitions:

American Indian or Alaska Native. A person having origins in any of the original peoples of North and South America (including Central America), and who maintains tribal affiliation or community attachment.

Asian. A person having origins in any of the original peoples of the Far East, Southeast Asia, or the Indian subcontinent including, for example, Cambodia, China, India, Japan, Korea, Malaysia, Pakistan, the Philippine Islands, Thailand, and Vietnam.

Black or African American. A person having origins in any of the black racial groups of Africa.

Native Hawaiian or Other Pacific Islander. A person having origins in any of the original peoples of Hawaii, Guam, Samoa, or other Pacific Islands.

White. A person having origins in any of the original peoples of Europe, the Middle East, or North Africa.

WHY THIS INFORMATION IS BEING REQUESTED:

The Federal Government has a continuing commitment to monitor the operation of its review and award processes to identify and address any inequities based on gender, race, ethnicity, or disability of its proposed PIs/PDs. To gather information needed for this important task, the proposer should submit a single copy of this form for each identified PI/PD with each proposal. Submission of the requested information is voluntary and will not affect the organization's eligibility for an award. However, information not submitted will seriously undermine the statistical validity, and therefore the usefulness, of information received from others. Any individual not wishing to submit some or all the information should check the box provided for this purpose. (The exceptions are the PI/PD name and the information about prior Federal support, the last question above.)

Collection of this information is authorized by the NSF Act of 1950, as amended, 42 U.S.C. 1861, et seq. Demographic data allows NSF to gauge whether our programs and other opportunities in science and technology are fairly reaching and benefiting everyone regardless of demographic category; to ensure that those in under-represented groups have the same knowledge of and access to programs and other research and educational opportunities; and to assess involvement of international investigators in work supported by NSF. The information may be disclosed to government contractors, experts, volunteers and researchers to complete assigned work; and to other government agencies in order to coordinate and assess programs. The information may be added to the Reviewer file and used to select potential candidates to serve as peer reviewers or advisory committee members. See Systems of Records, NSF-50, "Principal Investigator/Proposal File and Associated Records", 63 Federal Register 267 (January 5, 1998), and NSF-51, "Reviewer/Proposal File and Associated Records", 63 Federal Register 268 (January 5, 1998).

**INFORMATION ABOUT PRINCIPAL INVESTIGATORS/PROJECT DIRECTORS(PI/PD) and
co-PRINCIPAL INVESTIGATORS/co-PROJECT DIRECTORS**

Submit only ONE copy of this form for each PI/PD and co-PI/PD identified on the proposal. The form(s) should be attached to the original proposal as specified in GPG Section II.B. Submission of this information is voluntary and is not a precondition of award. This information will not be disclosed to external peer reviewers. **DO NOT INCLUDE THIS FORM WITH ANY OF THE OTHER COPIES OF YOUR PROPOSAL AS THIS MAY COMPROMISE THE CONFIDENTIALITY OF THE INFORMATION.**

PI/PD Name: Jorg Peters

Gender: Male Female
Ethnicity: (Choose one response) Hispanic or Latino Not Hispanic or Latino

Race:
(Select one or more)
 American Indian or Alaska Native
 Asian
 Black or African American
 Native Hawaiian or Other Pacific Islander
 White

Disability Status:
(Select one or more)
 Hearing Impairment
 Visual Impairment
 Mobility/Orthopedic Impairment
 Other
 None

Citizenship: (Choose one) U.S. Citizen Permanent Resident Other non-U.S. Citizen

Check here if you do not wish to provide any or all of the above information (excluding PI/PD name):

REQUIRED: Check here if you are currently serving (or have previously served) as a PI, co-PI or PD on any federally funded project

Ethnicity Definition:

Hispanic or Latino. A person of Mexican, Puerto Rican, Cuban, South or Central American, or other Spanish culture or origin, regardless of race.

Race Definitions:

American Indian or Alaska Native. A person having origins in any of the original peoples of North and South America (including Central America), and who maintains tribal affiliation or community attachment.

Asian. A person having origins in any of the original peoples of the Far East, Southeast Asia, or the Indian subcontinent including, for example, Cambodia, China, India, Japan, Korea, Malaysia, Pakistan, the Philippine Islands, Thailand, and Vietnam.

Black or African American. A person having origins in any of the black racial groups of Africa.

Native Hawaiian or Other Pacific Islander. A person having origins in any of the original peoples of Hawaii, Guam, Samoa, or other Pacific Islands.

White. A person having origins in any of the original peoples of Europe, the Middle East, or North Africa.

WHY THIS INFORMATION IS BEING REQUESTED:

The Federal Government has a continuing commitment to monitor the operation of its review and award processes to identify and address any inequities based on gender, race, ethnicity, or disability of its proposed PIs/PDs. To gather information needed for this important task, the proposer should submit a single copy of this form for each identified PI/PD with each proposal. Submission of the requested information is voluntary and will not affect the organization's eligibility for an award. However, information not submitted will seriously undermine the statistical validity, and therefore the usefulness, of information received from others. Any individual not wishing to submit some or all the information should check the box provided for this purpose. (The exceptions are the PI/PD name and the information about prior Federal support, the last question above.)

Collection of this information is authorized by the NSF Act of 1950, as amended, 42 U.S.C. 1861, et seq. Demographic data allows NSF to gauge whether our programs and other opportunities in science and technology are fairly reaching and benefiting everyone regardless of demographic category; to ensure that those in under-represented groups have the same knowledge of and access to programs and other research and educational opportunities; and to assess involvement of international investigators in work supported by NSF. The information may be disclosed to government contractors, experts, volunteers and researchers to complete assigned work; and to other government agencies in order to coordinate and assess programs. The information may be added to the Reviewer file and used to select potential candidates to serve as peer reviewers or advisory committee members. See Systems of Records, NSF-50, "Principal Investigator/Proposal File and Associated Records", 63 Federal Register 267 (January 5, 1998), and NSF-51, "Reviewer/Proposal File and Associated Records", 63 Federal Register 268 (January 5, 1998).

**INFORMATION ABOUT PRINCIPAL INVESTIGATORS/PROJECT DIRECTORS(PI/PD) and
co-PRINCIPAL INVESTIGATORS/co-PROJECT DIRECTORS**

Submit only ONE copy of this form for each PI/PD and co-PI/PD identified on the proposal. The form(s) should be attached to the original proposal as specified in GPG Section II.B. Submission of this information is voluntary and is not a precondition of award. This information will not be disclosed to external peer reviewers. **DO NOT INCLUDE THIS FORM WITH ANY OF THE OTHER COPIES OF YOUR PROPOSAL AS THIS MAY COMPROMISE THE CONFIDENTIALITY OF THE INFORMATION.**

PI/PD Name: Sanjay Ranka

Gender: Male Female

Ethnicity: (Choose one response) Hispanic or Latino Not Hispanic or Latino

Race: American Indian or Alaska Native

(Select one or more)

Asian

Black or African American

Native Hawaiian or Other Pacific Islander

White

Disability Status: Hearing Impairment

(Select one or more)

Visual Impairment

Mobility/Orthopedic Impairment

Other

None

Citizenship: (Choose one) U.S. Citizen Permanent Resident Other non-U.S. Citizen

Check here if you do not wish to provide any or all of the above information (excluding PI/PD name):

REQUIRED: Check here if you are currently serving (or have previously served) as a PI, co-PI or PD on any federally funded project

Ethnicity Definition:

Hispanic or Latino. A person of Mexican, Puerto Rican, Cuban, South or Central American, or other Spanish culture or origin, regardless of race.

Race Definitions:

American Indian or Alaska Native. A person having origins in any of the original peoples of North and South America (including Central America), and who maintains tribal affiliation or community attachment.

Asian. A person having origins in any of the original peoples of the Far East, Southeast Asia, or the Indian subcontinent including, for example, Cambodia, China, India, Japan, Korea, Malaysia, Pakistan, the Philippine Islands, Thailand, and Vietnam.

Black or African American. A person having origins in any of the black racial groups of Africa.

Native Hawaiian or Other Pacific Islander. A person having origins in any of the original peoples of Hawaii, Guam, Samoa, or other Pacific Islands.

White. A person having origins in any of the original peoples of Europe, the Middle East, or North Africa.

WHY THIS INFORMATION IS BEING REQUESTED:

The Federal Government has a continuing commitment to monitor the operation of its review and award processes to identify and address any inequities based on gender, race, ethnicity, or disability of its proposed PIs/PDs. To gather information needed for this important task, the proposer should submit a single copy of this form for each identified PI/PD with each proposal. Submission of the requested information is voluntary and will not affect the organization's eligibility for an award. However, information not submitted will seriously undermine the statistical validity, and therefore the usefulness, of information received from others. Any individual not wishing to submit some or all the information should check the box provided for this purpose. (The exceptions are the PI/PD name and the information about prior Federal support, the last question above.)

Collection of this information is authorized by the NSF Act of 1950, as amended, 42 U.S.C. 1861, et seq. Demographic data allows NSF to gauge whether our programs and other opportunities in science and technology are fairly reaching and benefiting everyone regardless of demographic category; to ensure that those in under-represented groups have the same knowledge of and access to programs and other research and educational opportunities; and to assess involvement of international investigators in work supported by NSF. The information may be disclosed to government contractors, experts, volunteers and researchers to complete assigned work; and to other government agencies in order to coordinate and assess programs. The information may be added to the Reviewer file and used to select potential candidates to serve as peer reviewers or advisory committee members. See Systems of Records, NSF-50, "Principal Investigator/Proposal File and Associated Records", 63 Federal Register 267 (January 5, 1998), and NSF-51, "Reviewer/Proposal File and Associated Records", 63 Federal Register 268 (January 5, 1998).

**INFORMATION ABOUT PRINCIPAL INVESTIGATORS/PROJECT DIRECTORS(PI/PD) and
co-PRINCIPAL INVESTIGATORS/co-PROJECT DIRECTORS**

Submit only ONE copy of this form for each PI/PD and co-PI/PD identified on the proposal. The form(s) should be attached to the original proposal as specified in GPG Section II.B. Submission of this information is voluntary and is not a precondition of award. This information will not be disclosed to external peer reviewers. **DO NOT INCLUDE THIS FORM WITH ANY OF THE OTHER COPIES OF YOUR PROPOSAL AS THIS MAY COMPROMISE THE CONFIDENTIALITY OF THE INFORMATION.**

PI/PD Name: Mark S Schmalz

Gender: Male Female
Ethnicity: (Choose one response) Hispanic or Latino Not Hispanic or Latino

Race:
(Select one or more)
 American Indian or Alaska Native
 Asian
 Black or African American
 Native Hawaiian or Other Pacific Islander
 White

Disability Status:
(Select one or more)
 Hearing Impairment
 Visual Impairment
 Mobility/Orthopedic Impairment
 Other
 None

Citizenship: (Choose one) U.S. Citizen Permanent Resident Other non-U.S. Citizen

Check here if you do not wish to provide any or all of the above information (excluding PI/PD name):

REQUIRED: Check here if you are currently serving (or have previously served) as a PI, co-PI or PD on any federally funded project

Ethnicity Definition:

Hispanic or Latino. A person of Mexican, Puerto Rican, Cuban, South or Central American, or other Spanish culture or origin, regardless of race.

Race Definitions:

American Indian or Alaska Native. A person having origins in any of the original peoples of North and South America (including Central America), and who maintains tribal affiliation or community attachment.

Asian. A person having origins in any of the original peoples of the Far East, Southeast Asia, or the Indian subcontinent including, for example, Cambodia, China, India, Japan, Korea, Malaysia, Pakistan, the Philippine Islands, Thailand, and Vietnam.

Black or African American. A person having origins in any of the black racial groups of Africa.

Native Hawaiian or Other Pacific Islander. A person having origins in any of the original peoples of Hawaii, Guam, Samoa, or other Pacific Islands.

White. A person having origins in any of the original peoples of Europe, the Middle East, or North Africa.

WHY THIS INFORMATION IS BEING REQUESTED:

The Federal Government has a continuing commitment to monitor the operation of its review and award processes to identify and address any inequities based on gender, race, ethnicity, or disability of its proposed PIs/PDs. To gather information needed for this important task, the proposer should submit a single copy of this form for each identified PI/PD with each proposal. Submission of the requested information is voluntary and will not affect the organization's eligibility for an award. However, information not submitted will seriously undermine the statistical validity, and therefore the usefulness, of information received from others. Any individual not wishing to submit some or all the information should check the box provided for this purpose. (The exceptions are the PI/PD name and the information about prior Federal support, the last question above.)

Collection of this information is authorized by the NSF Act of 1950, as amended, 42 U.S.C. 1861, et seq. Demographic data allows NSF to gauge whether our programs and other opportunities in science and technology are fairly reaching and benefiting everyone regardless of demographic category; to ensure that those in under-represented groups have the same knowledge of and access to programs and other research and educational opportunities; and to assess involvement of international investigators in work supported by NSF. The information may be disclosed to government contractors, experts, volunteers and researchers to complete assigned work; and to other government agencies in order to coordinate and assess programs. The information may be added to the Reviewer file and used to select potential candidates to serve as peer reviewers or advisory committee members. See Systems of Records, NSF-50, "Principal Investigator/Proposal File and Associated Records", 63 Federal Register 267 (January 5, 1998), and NSF-51, "Reviewer/Proposal File and Associated Records", 63 Federal Register 268 (January 5, 1998).

List of Suggested Reviewers or Reviewers Not To Include (optional)

SUGGESTED REVIEWERS:

Not Listed

REVIEWERS NOT TO INCLUDE:

Not Listed

COVER SHEET FOR PROPOSAL TO THE NATIONAL SCIENCE FOUNDATION

PROGRAM ANNOUNCEMENT/SOLICITATION NO./CLOSING DATE/if not in response to a program announcement/solicitation enter NSF 03-2					FOR NSF USE ONLY	
NSF 03-533			04/24/03		NSF PROPOSAL NUMBER	
FOR CONSIDERATION BY NSF ORGANIZATION UNIT(S) (Indicate the most specific unit known, i.e. program, division, etc.)						
ANI - Strategic Tech for Internet						
DATE RECEIVED	NUMBER OF COPIES	DIVISION ASSIGNED	FUND CODE	DUNS# (Data Universal Numbering System)	FILE LOCATION	
				969663814		
EMPLOYER IDENTIFICATION NUMBER (EIN) OR TAXPAYER IDENTIFICATION NUMBER (TIN)		SHOW PREVIOUS AWARD NO. IF THIS IS <input type="checkbox"/> A RENEWAL <input type="checkbox"/> AN ACCOMPLISHMENT-BASED RENEWAL		IS THIS PROPOSAL BEING SUBMITTED TO ANOTHER FEDERAL AGENCY? YES <input type="checkbox"/> NO <input checked="" type="checkbox"/> IF YES, LIST ACRONYM(S)		
596002052						
NAME OF ORGANIZATION TO WHICH AWARD SHOULD BE MADE			ADDRESS OF AWARDEE ORGANIZATION, INCLUDING 9 DIGIT ZIP CODE			
University of Florida			University of Florida			
AWARDEE ORGANIZATION CODE (IF KNOWN)			219 Grinter Hall			
0015354000			Gainesville, FL. 32611			
NAME OF PERFORMING ORGANIZATION, IF DIFFERENT FROM ABOVE			ADDRESS OF PERFORMING ORGANIZATION, IF DIFFERENT, INCLUDING 9 DIGIT ZIP CODE			
PERFORMING ORGANIZATION CODE (IF KNOWN)						
IS AWARDEE ORGANIZATION (Check All That Apply) (See GPG II.C For Definitions)						
<input type="checkbox"/> SMALL BUSINESS <input type="checkbox"/> MINORITY BUSINESS <input type="checkbox"/> IF THIS IS A PRELIMINARY PROPOSAL THEN CHECK HERE <input type="checkbox"/> FOR-PROFIT ORGANIZATION <input type="checkbox"/> WOMAN-OWNED BUSINESS						
TITLE OF PROPOSED PROJECT OCEAN: A Peer-to-Peer Market for Allocating Grid Computing Resources						
REQUESTED AMOUNT		PROPOSED DURATION (1-60 MONTHS)		REQUESTED STARTING DATE		SHOW RELATED PRELIMINARY PROPOSAL NO. IF APPLICABLE
\$ 982,484		36 months		01/01/04		
CHECK APPROPRIATE BOX(ES) IF THIS PROPOSAL INCLUDES ANY OF THE ITEMS LISTED BELOW						
<input checked="" type="checkbox"/> BEGINNING INVESTIGATOR (GPG I.A) <input type="checkbox"/> HUMAN SUBJECTS (GPG II.C.11) <input type="checkbox"/> DISCLOSURE OF LOBBYING ACTIVITIES (GPG II.C) Exemption Subsection _____ or IRB App. Date _____ <input type="checkbox"/> PROPRIETARY & PRIVILEGED INFORMATION (GPG I.B, II.C.6) <input type="checkbox"/> INTERNATIONAL COOPERATIVE ACTIVITIES: COUNTRY/COUNTRIES INVOLVED (GPG II.C.9) <input type="checkbox"/> HISTORIC PLACES (GPG II.C.9) <input type="checkbox"/> SMALL GRANT FOR EXPLOR. RESEARCH (SGER) (GPG II.C.11) <input type="checkbox"/> VERTEBRATE ANIMALS (GPG II.C.11) IACUC App. Date _____ <input type="checkbox"/> HIGH RESOLUTION GRAPHICS/OTHER GRAPHICS WHERE EXACT COLOR REPRESENTATION IS REQUIRED FOR PROPER INTERPRETATION (GPG I.E.1)						
PI/PD DEPARTMENT			PI/PD POSTAL ADDRESS			
Dept. of Comp. & Information Sci. & Eng.			CSE Bldg, PO Box 116120			
PI/PD FAX NUMBER			Gainesville, FL 32611			
352-392-1220			United States			
NAMES (TYPED)		High Degree	Yr of Degree	Telephone Number	Electronic Mail Address	
PI/PD NAME						
Michael P Frank		Ph.D	1999	352-392-6888	mpf@cise.ufl.edu	
CO-PI/PD						
Abdelsalam A Helal		Ph.D.	1991	352-392-6845	helal@cise.ufl.edu	
CO-PI/PD						
Jorg Peters		PhD	1990	352-392-1226	jorg@cise.ufl.edu	
CO-PI/PD						
Sanjay Ranka		PhD	1988	352-392-1526	ranka@cise.ufl.edu	
CO-PI/PD						
Mark S Schmalz		PhD	1996	352-392-6831	mssz@cise.ufl.edu	

CERTIFICATION PAGE

Certification for Authorized Organizational Representative or Individual Applicant:

By signing and submitting this proposal, the individual applicant or the authorized official of the applicant institution is: (1) certifying that statements made herein are true and complete to the best of his/her knowledge; and (2) agreeing to accept the obligation to comply with NSF award terms and conditions if an award is made as a result of this application. Further, the applicant is hereby providing certifications regarding debarment and suspension, drug-free workplace, and lobbying activities (see below), as set forth in Grant Proposal Guide (GPG), NSF 03-2. Willful provision of false information in this application and its supporting documents or in reports required under an ensuing award is a criminal offense (U. S. Code, Title 18, Section 1001).

In addition, if the applicant institution employs more than fifty persons, the authorized official of the applicant institution is certifying that the institution has implemented a written and enforced conflict of interest policy that is consistent with the provisions of Grant Policy Manual Section 510; that to the best of his/her knowledge, all financial disclosures required by that conflict of interest policy have been made; and that all identified conflicts of interest will have been satisfactorily managed, reduced or eliminated prior to the institution's expenditure of any funds under the award, in accordance with the institution's conflict of interest policy. Conflicts which cannot be satisfactorily managed, reduced or eliminated must be disclosed to NSF.

Drug Free Work Place Certification

By electronically signing the NSF Proposal Cover Sheet, the Authorized Organizational Representative or Individual Applicant is providing the Drug Free Work Place Certification contained in Appendix A of the Grant Proposal Guide.

Debarment and Suspension Certification

(If answer "yes", please provide explanation.)

Is the organization or its principals presently debarred, suspended, proposed for debarment, declared ineligible, or voluntarily excluded from covered transactions by any Federal department or agency?

Yes

No

By electronically signing the NSF Proposal Cover Sheet, the Authorized Organizational Representative or Individual Applicant is providing the Debarment and Suspension Certification contained in Appendix B of the Grant Proposal Guide.

Certification Regarding Lobbying

This certification is required for an award of a Federal contract, grant, or cooperative agreement exceeding \$100,000 and for an award of a Federal loan or a commitment providing for the United States to insure or guarantee a loan exceeding \$150,000.

Certification for Contracts, Grants, Loans and Cooperative Agreements

The undersigned certifies, to the best of his or her knowledge and belief, that:

(1) No federal appropriated funds have been paid or will be paid, by or on behalf of the undersigned, to any person for influencing or attempting to influence an officer or employee of any agency, a Member of Congress, an officer or employee of Congress, or an employee of a Member of Congress in connection with the awarding of any federal contract, the making of any Federal grant, the making of any Federal loan, the entering into of any cooperative agreement, and the extension, continuation, renewal, amendment, or modification of any Federal contract, grant, loan, or cooperative agreement.

(2) If any funds other than Federal appropriated funds have been paid or will be paid to any person for influencing or attempting to influence an officer or employee of any agency, a Member of Congress, an officer or employee of Congress, or an employee of a Member of Congress in connection with this Federal contract, grant, loan, or cooperative agreement, the undersigned shall complete and submit Standard Form-LLL, "Disclosure of Lobbying Activities," in accordance with its instructions.

(3) The undersigned shall require that the language of this certification be included in the award documents for all subawards at all tiers including subcontracts, subgrants, and contracts under grants, loans, and cooperative agreements and that all subrecipients shall certify and disclose accordingly.

This certification is a material representation of fact upon which reliance was placed when this transaction was made or entered into. Submission of this certification is a prerequisite for making or entering into this transaction imposed by section 1352, Title 31, U.S. Code. Any person who fails to file the required certification shall be subject to a civil penalty of not less than \$10,000 and not more than \$100,000 for each such failure.

AUTHORIZED ORGANIZATIONAL REPRESENTATIVE		SIGNATURE	DATE
NAME			
TELEPHONE NUMBER	ELECTRONIC MAIL ADDRESS	FAX NUMBER	

*SUBMISSION OF SOCIAL SECURITY NUMBERS IS VOLUNTARY AND WILL NOT AFFECT THE ORGANIZATION'S ELIGIBILITY FOR AN AWARD. HOWEVER, THEY ARE AN INTEGRAL PART OF THE INFORMATION SYSTEM AND ASSIST IN PROCESSING THE PROPOSAL. SSN SOLICITED UNDER NSF ACT OF 1950, AS AMENDED.

OCEAN (the Open Computation Exchange & Auctioning Network) – A Peer-to-Peer Market for Allocating Grid Computing Resources

B. Project Summary

The OCEAN (Open Computation Exchange & Auctioning Network) project [1,2,3,4,5,6,7,8,26,21] was conceived at MIT in early 1997 by PI Michael Frank, with the goal of providing a system to automate the buying and selling of remote access to distributed computing resources, so as to both defray everyone's costs of ownership of underutilized computing resources and to provide instant on-demand access to enormous computational power for everyone from academic researchers to commercial end-users.

Intellectual Merit. Although the OCEAN project's high-level goals are primarily focused on practical aims of infrastructure development and standardization, a number of very interesting academic research issues of fundamental importance arise naturally, and are being addressed in the course of this effort:

1. **Adaptive P2P search algorithms.** One important research question is how to efficiently perform an effective search (for resources matching stated requirements) among a dynamically-changing, peer-to-peer network of resource providers. We have invented and implemented several new peer-to-peer matching algorithms (Tobias [6], PLUM [7] and MarcoPolo [8]) using machine-learning capabilities for dynamic adaptation to continually optimize the network's efficiency. A paper on one of the algorithms was accepted for HiPC-02 [8], and many additional conference & journal articles will be forthcoming.
2. **Market design.** OCEAN will provide a realistic testbed for experimenting with the real-world effectiveness of different market/auction mechanisms. Fixed-price markets as well as double-bid auctions have already been implemented for OCEAN, and alternative mechanisms such as English and Dutch auctions and novel proposed market mechanisms all can be implemented within the OCEAN infrastructure, and their effectiveness thereby compared in a real-world setting.
3. **Collaborative language extension.** OCEAN's language for describing resource requirements and capabilities must be standardized for interoperability, while remaining arbitrarily extensible to express any desired kinds of features and constraints. This leads to a novel research effort, to design and develop an online system for the continual, collaborative extension of resource description languages & ontologies, represented using XML schemas [9].
4. **Security validation.** The real-world security needs of OCEAN provide an environment in which a number of new security paradigms can be effectively developed and tested. We have already implemented the XML digital signature standard, to prevent forgery and repudiation of sales contracts. Java and .NET CLR [10] sandbox security mechanisms will be significantly exercised and stress-tested by real-world use in OCEAN. Finally, sophisticated *Computation Certificate* algorithms, proposed in the 1990's [11,12], might find their first practical application in automated verification of computations in OCEAN. In these algorithms, a computation yields as a side-effect a short probabilistically-checkable proof that the requested computation was in fact performed correctly.
5. **Distributed scientific applications.** Finally, many issues arise in the design of distributed algorithms for solving particular scientific problems of interest that can perform robustly in a distributed environment such as OCEAN with no centralized control, *e.g.*, purchasing new resources and reallocating subtasks to them when a resource becomes unavailable. The co-PIs Schmalz, Peters, and Hammer will explore these and other issues in the specific contexts of image processing, 3D modeling, and database processing respectively.

Broader Impacts. The OCEAN project has a potentially enormous and broad impact, both on users in the academic and industrial computational science & engineering communities, but also on users in other areas, such as computer animation, financial analysis, web testing, cryptanalysis, computer gaming, *etc.* It has the potential to significantly lower the average cost of computing for everyone, by reducing waste of unused cycles. Furthermore, the same infrastructure, with further extensions to the resource description language, can also be used to exchange not just computational resources, but all manner of other commodities, goods, services, financial instruments, *etc.*, and as a result, it can potentially eventually revolutionize all of e-commerce and the financial markets, by providing a universal, open standard for automated, peer-to-peer trading with no centralized control, in contrast to the existing centralized business-to-business exchanges, stock markets, *etc.*

TABLE OF CONTENTS

For font size and page formatting specifications, see GPG section II.C.

	Total No. of Pages	Page No.* (Optional)*
Cover Sheet for Proposal to the National Science Foundation		
Project Summary (not to exceed 1 page)	1	_____
Table of Contents	1	_____
Project Description (Including Results from Prior NSF Support) (not to exceed 15 pages) (Exceed only if allowed by a specific program announcement/solicitation or if approved in advance by the appropriate NSF Assistant Director or designee)	15	_____
References Cited	2	_____
Biographical Sketches (Not to exceed 2 pages each)	12	_____
Budget (Plus up to 3 pages of budget justification)	8	_____
Current and Pending Support	10	_____
Facilities, Equipment and Other Resources	2	_____
Special Information/Supplementary Documentation	2	_____
Appendix (List below.) (Include only if allowed by a specific program announcement/ solicitation or if approved in advance by the appropriate NSF Assistant Director or designee)	_____	_____
Appendix Items:		

*Proposers may select any numbering mechanism for the proposal. The entire proposal however, must be paginated. Complete both columns only if the proposal is numbered consecutively.

C. Project Description - C.1. Introduction

The name "OCEAN" is an acronym for Open Computation Exchange and Auctioning Network. The purpose of OCEAN is to provide (as middleware, layered on the internet) a world-wide automated commodities market for the on-demand purchase of remote access rights (compatibly with Grid standards) to distributed computing resources (such as CPU cycles, memory, disk space, network bandwidth, access to specific remote software/databases, *etc.*) for use by high-performance distributed computing applications, mobile agents, or thin (*e.g.* wireless) clients.

The R&D strategy of the OCEAN project is heavily geared towards impelling rapid adoption and installed-base growth via creating extreme usefulness. The design is intended to emphasize making it as easy as possible for users to deploy OCEAN server nodes, and to develop and run OCEAN applications at will, with the application users paying the server providers incrementally for resources utilized. OCEAN is intended to exhibit the following key design characteristics, not necessarily in order of importance.

1. **Open.** Every aspect of OCEAN operation will be based on public, open, freely reimplementable standards, APIs, data formats, and protocols. All OCEAN software will come with free source code, as a reference implementation of the standard. Anyone will be free to modify or reimplement the OCEAN standards as long as they conform to its specification. Also, anyone can extend the standards, with approval from the OCEAN governing body. Non-approved modifications of the standard would constitute a violation of copyright.
2. **Profitable.** Companies and individuals who are involved in helping build the OCEAN infrastructure and providing servers, services, and applications for it should be able to profit financially from doing so. This kind of true economic incentive system seems to be a prerequisite for rapid growth of the technology, its democratization, and the emergence of new industries based upon it.
3. **Portable.** OCEAN server software should run on the widest possible variety of host platforms so that it can be quickly deployed on the largest possible numbers of existing machines.
4. **Interoperable.** OCEAN should be compatible with other systems, tools, and protocols based on major emerging web-services (*e.g.*, XML [9], SOAP [13], WSDL [14]) and Grid (*e.g.*, Globus [15], OGSA [16]) standards.
5. **Secure.** The OCEAN infrastructure should offer as much security for providers and users, and accountability for transactions, as is feasibly achievable. We address security issues in more detail in section C.7 below.
6. **Efficient.** The OCEAN infrastructure should put as little overhead as possible in the way of the basic task - getting distributed computations done - and meanwhile it should allocate resources in a way that maximizes the overall economic efficiency of the resource allocation as much as possible.
7. **Scalable.** OCEAN servers should be able to be deployed on arbitrarily many hosts, and support arbitrarily high levels of usage, while degrading in performance as little as possible as the scale of the network increases. This requirement needs to the use of a self-optimizing peer-to-peer system for resource-request matching.
8. **Easily deployable.** It should be trivial for anyone with a computer to deploy an OCEAN node; *e.g.*, any high school or college student armed with a PC and a credit card should be able to deploy an OCEAN node, as easily as they today deploy KaZaa file-sharing clients. With only a modest amount of network administration work, it should also be relatively easy to deploy OCEAN nodes even behind firewalls and within private IP address spaces.
9. **Configurable.** With only a little extra effort on the resource provider's part, OCEAN server nodes should be able to be flexibly configured in a variety of ways, for example to be available only at certain times of day, or to set a certain pricing policy, or to provide access to special types of computational resources.
10. **Unobtrusive.** OCEAN server software should coexist unobtrusively with other programs running on the user's machines. So, for example, it should run at lower priority (with higher "niceness" on Unix) than interactive user processes. It should (if possible) never hog all of the memory or other resources of a machine.
11. **Easily programmable.** It should be extremely easy to write applications for OCEAN, so that, for example, any programmer who is fluent in C++ & Globus, Java, or C# can download the API libraries, read a reasonably short and simple API document (or a few "hello world" examples), and be able right away to start writing distributed OCEAN applications that can purchase and harness a variety of nodes to perform a computation.
12. **Monitorable.** It should (ideally) be easy for human users and providers to monitor the state and interesting statistics of both their own nodes, and the OCEAN market as a whole.
13. **Automatic.** The fundamental operation of an application deployer launching an OCEAN application, the application purchasing the resources it needs on behalf of its deployer, spawning child tasks on remote server nodes, performing its computation, and the resulting transfer of funds from the buyer to the seller (along with collection of any transaction fees by the market operator), should not require any human intervention.
14. **Dynamic.** Of course the set of available resources, the set of outstanding resource requests, and the set of jobs running on the OCEAN will all be constantly changing, and the system should adapt to these changes as rapidly as can be feasibly accommodated given the other requirements.
15. **Robust.** Insofar as possible, the system should be able to tolerate large fractions of its nodes going down or becoming disconnected from each other due to network failures.

C.2. Specific Relevance to the STI Funding Program

In this section, we quote some key phrases from the Strategic Technologies for the Internet Program Announcement and address them specifically, to emphasize the high relevance of the proposed effort.

“The primary purpose of this program is to support work that will have impact on the network infrastructure within 3-5 years, by supporting projects that advance the network infrastructure technology frontier.”

→ OCEAN is exactly a network-based infrastructure project. It has potential near-term impact, since it is developing rapidly. We plan to release a prototype locally this year, and to the public a year afterwards. (See our proposal for local deployment [19], and our 5-year marketing plan [17].) It clearly advances the frontier, since there is currently no easy-to-use internet-based open market framework available for use by distributed computing applications.

“This program is designed to stimulate the development of critical infrastructure for the commodity Internet and/or advanced research networks...”

→ OCEAN is designed as an infrastructure layer that can be installed on top of the commodity Internet, as well as on corporate intranets. However, since it is a higher-layer protocol, there is nothing preventing it from also working with newer high-speed networks, such as Internet2.

“have outcomes not currently under development, meet an infrastructure need not met elsewhere”

→ No known existing or proposed competing project will provide *exactly* the key capabilities of OCEAN, specifically, an open, low-barrier-to-entry, dynamic, scalable, peer-to-peer market suited for exchange of access to computing resources, or any other commodity. (See §C.4.) Yet, the need for such an infrastructure seems clear.

“...are high risk...”

→ OCEAN is indeed very high risk from an *industry* perspective, because significant applications will not be ported to use its capabilities until the infrastructure supporting those capabilities already exists, but the infrastructure development won't be funded by short-sighted industry firms until some applications that harness it exist and the demand for the infrastructure is realized. That is exactly why government support for this project is crucially needed to create this new piece of critical infrastructure. Meanwhile, the open, scalable, extensible architecture of OCEAN, and its ease of use for developers, will together ensure that network effects are maximally leveraged, so that the grass-roots growth and adoption of OCEAN standards and technology occurs as rapidly as possible. Previous similar ventures we know (see §C.4) were much more cumbersome for developers to use; such design flaws effectively prevented their success.

“...advance the national research agenda and conduct of science...”

→ The OCEAN project was originally motivated by the PI's desire for easy access to massive computing resources for his own particular computational science applications (quantum mechanical simulations and 3-D rendering for scientific visualization). Providing a framework that is well-suited to meet the needs of the national academic (and industry) research community remains the primary focus of the OCEAN project. We have elicited support from key faculty in several local science & engineering departments (Physics, Aerospace & Mechanical, Nuclear & Radiological, Electrical & Computer) to have students port specific research applications to our platform. For example, Ali Haghghat, chair of the Nuclear & Radiological Engineering department here, and entrepreneur of radiological computing services, has assigned a Masters student to work with the project. The UF College of Engineering's committee on High Performance Computing [18] is beginning a major initiative to create an integrated cluster and grid computing infrastructure here on campus. UF is currently soliciting a donation of several thousand machines for this purpose from Dell and IBM. OCEAN is being promoted locally for campus-wide deployment [19] to support scientific computing here at UF, with approval of UF's vice-provost for IT.

“...have outcomes of interest to a broad range of the user community.”

→ OCEAN has interest to the entire internet-using community. Anyone who owns an underutilized desktop or server PC, or who manages an underutilized, possibly obsolete cluster would likely be willing to eke some additional benefit from those machines by selling their excess capacity on the OCEAN network. Meanwhile, any academic researcher or developer of commercial software who wants the capability to harness any desired amount of computing power on demand can do so (once OCEAN is widely deployed) with a minimum of effort. We are proposing to Microsoft to incorporate OCEAN capabilities into the standard developers' toolkit for future releases of the Microsoft .NET platform. As a result, even naïve end-users of internet technology may find themselves clicking a “Use OCEAN” checkbox to enhance the computing power of their shrink-wrapped, Windows-based consumer applications.

“The focus areas for FY 2003 are: Advanced applications for the Internet, particularly collaboration tools, remote ... access and grid-based applications, etc....”

→ OCEAN might itself be considered the epitome of an advanced Internet application, as well as being a platform that supports the operation of more domain-specific applications. It can be viewed as a collaboration tool, in the sense that it provides resource sharing, although the exchange of resources is automated and not human-mediated. One collaboration-oriented aspect of the project is our planned system for collaborative extension of resource-description and transaction-constraint ontologies, coded in XML. OCEAN is explicitly a tool for automatically obtaining rights for remote access to resources via payment, and it is explicitly designed to support grid-based applications.

“...techniques for identifying and reducing vulnerabilities in the Internet and advanced networks”

→ If OCEAN becomes widely deployed, by its very nature it will expose such vulnerabilities, as a result of opening up and even advertising access to resources across organizational boundaries. The risks of its deployment will nevertheless be outweighed by its benefits, for a large class of potential users. Due to its transaction-automation and remote-access-authorization capabilities, OCEAN will significantly stress-test certain aspects of network security, such as the unforgeability of XML digital signatures used for our contracts, the integrity of the Java SecurityManager sandbox framework used in our Java grid system, and Microsoft’s similar sandbox environment in the Common Language Runtime virtual machine in the .NET Framework. With widespread use of OCEAN, the most serious holes in these security mechanisms should be quickly identified and repaired. Furthermore, for protection of users of remote resources, OCEAN plans to implement new schemes for certifying the correctness of remotely-performed computations, using some theoretical techniques that have been proposed in the literature [11,12,20].

“...technologies for enhancing the edges and network ubiquity: wireless (mobile, broadband, pervasive computing)”

→ OCEAN has a significant application in mobile and pervasive computing, namely to allow lightweight mobile clients to have on-demand access to any needed amount of computing power, to support more sophisticated applications, via broadband wireless links to more powerful landline-wired systems. Even the less power-limited desktop systems located at the edges of the network will still benefit from having easy on-demand access to even more powerful remote computing, like a utility. OCEAN will therefore extend the range of viable applications for all peripheral network users.

In summary, OCEAN is a perfect fit both to the overall goals of the STI program, as well as to several of the focus areas for the current fiscal year. The OCEAN project was already proposed for STI support two years ago (in FY 2001) [21], and received very positive reviews, but did not end up getting funded. However, even without any funding, the project has progressed significantly in the two years since our last STI proposal, producing a number of conference papers and Masters’ theses, implementing many key components, and (via trial implementations and empirical simulations) clarifying many core research issues, such as those relating to the design of efficient peer-to-peer resource/request matching algorithm.

In this new proposal, rather than repeating our previous material, we have focused on carefully addressing the issues raised by the FY2001 reviewers. It is our fervent hope that this time, the project will receive funding, so that we can continue pursuing OCEAN as an open, public system, rather than pursuing an alternative privately-funded approach, which would most likely require declaring intellectual property, distributing OCEAN code as proprietary binaries rather than open-source, *etc.* Not only would this private approach make OCEAN more difficult for the academic community to use freely for research purposes, but it would also, in our opinion, hamper the overall level of grass-roots support and thus ultimately growth potential for OCEAN, in comparison with the alternative of it being a publicly-funded open standard, free for anyone to extend, reimplement, and deploy.

C.3. Overview of OCEAN Architecture

Figure 1 below gives an overview of the software architecture of OCEAN, currently under development. OCEAN is designed to be easily interoperable with popular Grid technologies currently under development. OCEAN provides a value-added service of resource discovery and allocation in an open-market fashion that is not currently supported by any of the existing freely-available Grid systems.

Figure 2 below shows an abstract, high-level diagram of network communications within the OCEAN system. The cloud in the middle represents a network of peering arrangements between OCEAN nodes (§C.6.1). This P2P network comprises a complex graph of connections which are self-optimizing, in the sense that each node continually adjusts its peer list based on past experience, using a simple and efficient Bayesian statistical machine learning algorithm, so that eventually each node becomes most likely to be peered most directly with those nodes that are most likely to lead to a successful match. (See §C.6.2.) Eventually, in a later version, intermediate nodes whose work leads to successful matches will even be paid a small “finder’s fee” out of the overhead fees for the transaction.

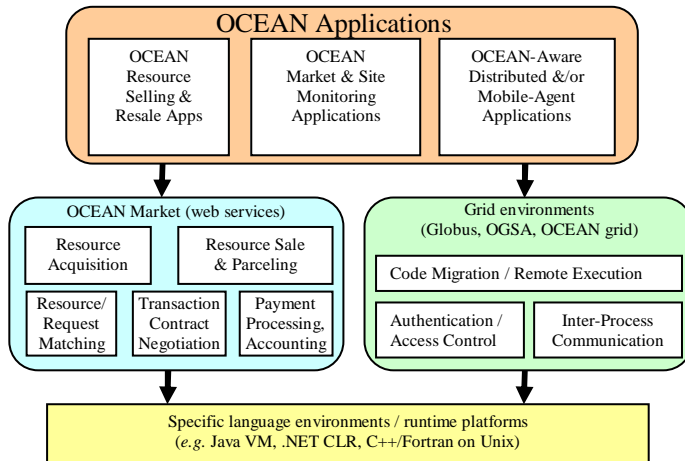


Figure 1. OCEAN Software Architecture. OCEAN will interoperate with other Grid systems, such as the Globus open-source Grid toolkit, the GriPhyN (Grid Physics Network) project, and the Global Grid Forum’s emerging Open Grid Services Architecture. OCEAN will be a value-added service within the context of the growing installed base of such grid technologies. We are also developing our own, very easy suite of Grid tools for Java and C# users who might rather not install the mainstream Grid tools, most of which are for legacy C++/Fortran/Unix environments. Java and C# providers benefit from the enhanced security of a “sandboxed” execution environment. OCEAN’s facility for market-based resource trading is implemented in a Web-Services model, using the XML and SOAP open standards being supported by the major industry players. Applications that use, monitor, or sell resources using OCEAN can be built on top of the Grid and Market facilities.

C.4. Comparison Matrix of Related Projects

There are a large number of past and present academic projects and commercial ventures that leverage one or more of the distributed (cluster and/or grid), market-based, and peer-to-peer computing concepts. The below listing does not even include most of the application-specific projects, such as many special-purpose distributed computing applications, and various P2P-oriented file-sharing, file searching, and collaborative communication ventures. Some additional narrative comparing OCEAN to some of the most important of the past & present Grid and market-based projects was given in the FY2001 edition of this proposal [21]. However, absolutely *none* of the current projects, so far as we have been able to determine to date, actually combines all of these important capabilities in a single, complete, application-independent, market-based, wide-area-capable grid system, including especially the crucial qualities of extremely low barriers-to-entry and high ease of use for resource providers, application developers and end users, as well as support for multiple important language/runtime platforms, both the traditional legacy environments (such as Unix/C++/Fortran) and key state-of-the-art, emerging environments (Java 2 Enterprise

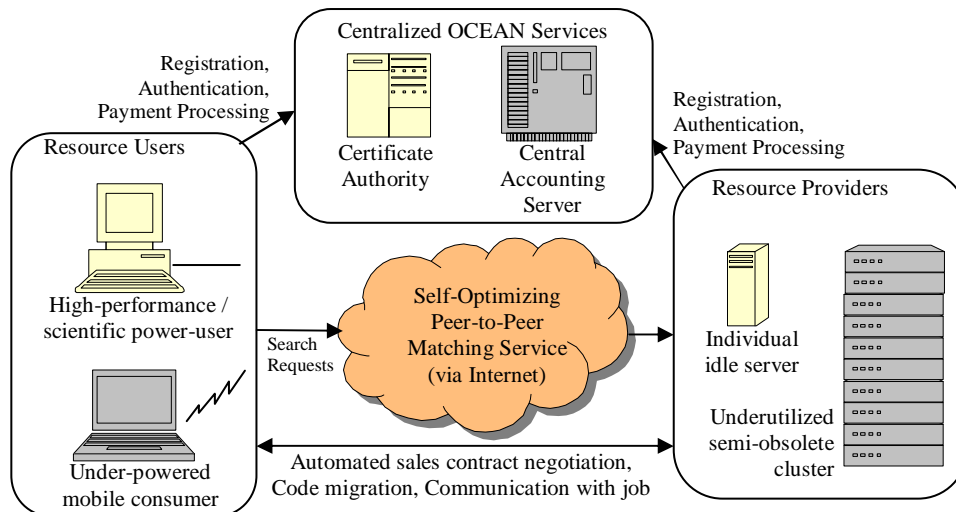


Figure 2. OCEAN Network Diagram. The left and right boxes illustrate some different types of resource users and providers, respectively, but nothing prevents a given machine from serving in both roles at different times, or even simultaneously. OCEAN works as follows. All nodes initially register with the central OCEAN services, to (1) obtain a digital certificate for later use in creating verifiable digital signatures, *e.g.*, on sales agreement contracts, (2) set up any accounts on the OCEAN Central Accounting Server (CAS) needed to facilitate later payments, and (3) initialize the node’s participation in the peer-to-peer network (cloud). When resources are needed, the buyer sends a search request to the self-optimizing, peer-to-peer matching service, consisting of other nearby OCEAN nodes. The peers use certain patentable machine learning techniques we have developed to quickly and intelligently route search queries to the provider nodes most likely to be able to fulfill them. A provider with matching available resources gets the query, and starts a sales contract negotiation directly with the buyer’s machine. A successful negotiation yields a sales agreement contract that is digitally signed by both parties. Payments are then processed according to terms specified in the contract, possibly utilizing services of the Central Accounting Server, depending on the specified payment mechanism. The client job code for remote execution is transferred and installed using a negotiated method (*e.g.*, GridFTP or Java classloaders) and executed on the provider machine, as per the contract terms. The buyer agent that initiated the transfer gets a handle to the remote process, which can be used for process control & communication. The application code responds to messages from its parent and sibling processes and/or pushes intermediate or final results of the computation back to the user’s machine as needed, using pre-specified protocols (*e.g.*, MPI [1]).

Edition / .NET Common Language Runtime). In other words, a system that is *truly* OCEAN-equivalent (or OCEAN-subsuming) simply does *not* yet exist, and furthermore has not even been *proposed* by our competitors.

Key: <input checked="" type="checkbox"/> = Has the given feature <input checked="" type="checkbox"/> = Does not have the given feature [?] = Status not easily determinable. <input type="checkbox"/> = Either not yet researched, or just not yet recorded.	Project/Venture Still Active	Enterprise / Local-area Clusters / Grids	Wide-area / Internet Distributed Grids	Supports Mobile Agents	Market-Based System	Free, Open Standards, Source, and Market	Scalable Peer-to-Peer Technology	Application-Independent, Flexible	Low Barriers-to-Entry for Providers	Low Barriers-to-Entry for Developers	Platform / Language / Runtime Independent
OCEAN	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Agorics, Inc.	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	[?]	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	[?]	[?]	[?]	<input checked="" type="checkbox"/>	[?]
Avaki (Applied Meta)		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>			<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Beowulf	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>					
Compute Power Market	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>			<input checked="" type="checkbox"/>
Condor	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Data Synapse		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Distributed.Net [22]	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Distributed Science	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>			<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
EconomyGrid (now CPM)	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>				
Entropia.com		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
FreeNet			<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>			<input checked="" type="checkbox"/>			
Globus Project	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>					<input checked="" type="checkbox"/>
Gnutella			<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>			<input checked="" type="checkbox"/>			
Google Compute	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>				<input checked="" type="checkbox"/>
Grid Resource Broker	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>					
GriPhyN	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>			
IBM Aglets		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>			<input checked="" type="checkbox"/>			
Juxtapose (JXTA)	<input checked="" type="checkbox"/>				<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>				
KaZaa gridware	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Legion		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>						
Napster	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Nimrod/G & GRACE	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>						
OceanStore			<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>			<input checked="" type="checkbox"/>			
OpenCola, Inc.		<input checked="" type="checkbox"/>			<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>			<input checked="" type="checkbox"/>
Open Grid Services Arch.	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Parabon Computing, Inc.		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Platform Computing, Inc.		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>					<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
PopCorn	<input checked="" type="checkbox"/>				<input checked="" type="checkbox"/>						<input checked="" type="checkbox"/>
Popular Power	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>			<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	
Porivo			<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	
ProcessTree	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	
SETI@Home	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Tryllian Agent Dev. Kit	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
United Devices, Inc.		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Figure 3. Comparison matrix between OCEAN and selected distributed/P2P computing projects and ventures. Others that we have also looked at but removed from the chart for conciseness include: 2AM Inc., Alien@GRID, Apples, BaseOne, Bond, CAN, CapCal, Centrata, CERN DAS, Data Grid, Chord, cJVM, Compaq Grid Solutions, CORBA, Covise, DesignDrug@Home, DISCWorld, DOE Science Grid, Ejasent Inc., eMikolo Networks, Enron exchanges, EROPPA, ExactOne Inc., FirstPeer, Folding@Home, FreeMarkets, GCTI, GLUE, Gossip, Globe, Gridbus Project, GridInfrastructure, Grid Interoperability Project, HARNESS, HEPGrid, HTC, InfoSpheres, Javelin, JaCo3, JaWs, Juno, JXTA, Lightshare, MetaMPI, METHODIS, Mithral (COSM), MOL, NSF Middleware Initiative, NASA Information Power Grid, NetSolve, NINF, Noemix, Oculus Technology Corp., OpenDesign, Inc. PARDIS, Pastry, Poznan Metacomputing, Punch, RSA Factoring-by-Web, Semantic Grid, Simulant Fearnor, Spawn, Sun Grid Engine, Tapestry, Ubero, UNICORE, WAMM, WebFlow, WebSubmit, and XtremWeb. Despite the large number of vaguely related projects, we have not yet found any that actually provides all the key features of OCEAN.

C.5. Application Areas and Application Development Plan

Like other Grid initiatives, OCEAN has an enormous range of potential applications, from all areas of computational science and engineering, including medicine & biosciences, to applications in business, finance, computer animation, *etc.* However, we believe that the success potential is much greater for OCEAN than other ventures, due

to its peer-to-peer, market-based aspect, as well as its emphasis on ease of use and its support for both modern and legacy application environments.

The total size of the commercial and academic market for OCEAN-like facilities is difficult to estimate, but clearly the potential areas for application are significant in size, and comprise a fast growing market. Eventually, OCEAN might even serve a significant fraction of the world's computing resources. But first, let us take a look at some specific vertical application areas of interest. Perhaps the best way to estimate the potential market size is to examine the individual application areas and industries. The current market opportunity for distributed processing and high-performance technical systems is about \$2.4 billion, according to a market overview by Stearns and Co. [23]. Lehman Brothers has also researched the market for direct applications of peer-to-peer technologies [24]. Some of the important potential application areas that we anticipate for OCEAN include:

- **Science:** All of computational science, including computational genomics/proteomics, protein folding, quantum computer simulations (Dr. Frank), high-energy physics (collaborator Dr. Avery) and astrophysics simulations (*e.g.*, QCD lattice models), quantum chemistry, medical radiology (collaborator Dr. Haghghat) & diagnosis, astronomical/geophysical data/image analysis, weather forecasting.
- **Engineering:** Computer graphics rendering/animation (Dr. Peters), engineering visualization, finite-element modeling for stress analysis and computational fluid dynamics (collaborator Dr. Shyy), VLSI circuit simulations, quantum mechanics simulations for nanodevice engineering (Dr. Frank), image processing, (Dr. Schmalz), pattern recognition and machine learning (Dr. Frank), combinatorial search and optimization algorithms, genetic algorithms for automated design evolution, etc.
- **Government:** Nuclear stockpile simulation, missile defense, cryptography and cryptanalysis, political and military simulation/analysis, and signal processing and intelligence analysis.
- **Consumer:** Home multimedia production, multiuser gaming, speech and handwriting recognition, computer vision for home security systems and household robotics, high-powered remote applications for mobile users.
- **Business:** Software distribution; financial trading/analysis/settlement, competitive analysis, distributed database query systems & web service testing (Dr. Hammer), data mining (Dr. Ranka).

Note that for highly sensitive government and commercial applications, OCEAN could be deployed internally within organizational intranets for added security.

Some projected 2004 market sizes, as sampled below in Figure 4, further underscore the point that application markets will require even greater resources in the future as they continue to grow.

Industry	2004 Projected Market Size	Data Source
Biosciences IT	\$11.6 billion	IDC Research
Animation	\$8 billion	Roncarelli Report (20% of \$40B market)
Bioinformatics	\$2 billion	Oscar Gruss and Sons
Web testing	\$450 million	WitSoundview 6/4/01 report
All Windows software	\$60 billion	IDC Research
Automated SW Quality (ASQ) tools	\$2.2 billion	IDC Research

Figure 4. Summary of selected results from preliminary application market research performed by OCEAN members in partnership with collaborators in the Virtual Asset Exchange [25] project.

The market clearly exists for an effective scalable solution. Distributed networks can provide supercomputing power for appropriate applications at personal computer per-CPU prices, or less due to the harnessing of underutilized resources. The evidence points to distributed and peer-to-peer computing as a growth market. What has not yet been achieved is a fully automated and scalable open auctioning system that both provides benefit to the user through increased computing power and to the resource suppliers through fair compensation. The use of standardized Grid architectures and easy entry pathways into the market will allow OCEAN to benefit a wide range of customers. The actual double-bid auctioning system run by the software could potentially trade any commodity or facilitate any definable market.

C.5.1. Application Development Plan

Obviously, a key requirement for the success of OCEAN will be the rapid development of applications for it. This in turn leads to our requirement that OCEAN be extremely flexible, yet simple to understand and use for developers, since otherwise, they will not bother to learn and use it. An adequate API design that is highly optimized for usability, as well as associated documentation, training materials and support must be produced to elicit grass-roots enthusiasm for the technology from among developers.

To evaluate the quality of our APIs and documentation, we will begin by using them for computational science & engineering projects within our team of senior personnel and students in our department. In addition, one of our

former graduate students, Chaitanya Chokkareddy, who is now a professor in India, will recruit his students to develop various OCEAN-based applications for their student projects. Once the early kinks are worked out, we will distribute our materials more widely to the entire large high-performance computing community here at UF, starting with our informal collaborators in other departments at UF who have agreed to work with OCEAN, including Ali Haghghat (Nuclear & Radiological Eng.), Wei Shyy (Aerospace & Mechanical Eng.), Alan George (Elec. & Computer Eng.), so that their students can try using the OCEAN development environment for their custom computational science & engineering research codes, and give additional feedback on its design. Finally, we will make OCEAN materials available to all of academia, as well as the general public, publicize its capabilities at conferences, and approach major potential commercial customers in the various target industries, to give them free training seminars on the use and benefits of OCEAN.

C.6. Samples of OCEAN Technology in Progress

Although the detailed design of the OCEAN protocols remains in flux, significant progress has been made in the last two years towards refining and implementing various specifications and reference implementations for components of the OCEAN system. In this section, we provide samples of technical details of several components, to give the flavor of the work that is being accomplished.

C.6.1. Resource/Request Description Language

In OCEAN, we use an XML-based constraint language for matching the needs of resource users to the offerings of resource providers. The language is very general and can be easily extended to describe any type of resource requirements, including constraints on the CPU, memory, disk, network bandwidth, auxiliary hardware, software runtime environment, auxiliary libraries/databases, as well as various constraints on the acceptable sales agreement details, such as price details (currency, units), schedule for usage of resources, limits on use, method of payment (play money, OCEAN account transfer, intraorganizational journal transfer, OFX, PayPal, e-Gold, digital cash, credit card), means of payment (direct transfer, OCEAN-mediated, escrow through OCEAN CAS) and schedule of payments (up-front, on-delivery, post-hoc, periodic). Each software application using or providing resources can decide how many constraints to specify up-front in a “trade proposal” for use in the distributed P2P matching/auction system.

To reduce the computational demand on the matching system, the semantics of the constraint language is kept very simple (specifying only a Boolean combination of minimum constraints that *must* be satisfied). More involved

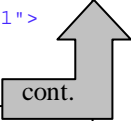
<pre><?xml version="1.0" encoding="utf-8" ?> <trade-proposal type="buy"> <access-agreement> <resource-package name="system1"> <machine type="linux" name="cpul"> <os type="linux" vendor="RedHat"> <version number="7.3 constraint=">="/> </os> <cpu type="x86" vendor="Intel"> <model name="Pentium4" constraint=">="/> <clockspeed number="2.0" unit="GHz" constraint=">="/> </cpu> <gridware type="Globus"> <version number="3.0 constraint=">="/> </gridware> <software type="linpack"/> <access mode="exclusive"/> </machine> <cluster type="beowulf" name="cluster1"> <access mode="shared"/> <machine type="linux" count="10" constraint=">="/> </cluster> <storage type="network" name="disk1"> <capacity number="10 units="GB" constraint=">="/> </storage> </resource-package> <network type="Ethernet"> <speed number="100 units="Gbps" constraint=">="/> <usage type="unlimited"/> <linkage item="cpul" item="cluster1" item="disk1"/> </network> </resource-package> <lease name="lease1" resource="#system1"> <start-time type="asap"/> <usage pattern="as-needed"/> <timespan type="flexible"/> <conjunction> <timespan number="2" unit="hours" constraint=">="/> <timespan number="1" unit="day" constraint="<="/> </conjunction> <pricing type="flat-rental"> <price num="1 currency="USD" constraint="<="/> <timeunit num="1 unit="hour"/> </pricing> </lease> <payment item="#lease1" type="via-OCEAN" currency="USD" method="OCEAN-acct-tfr"> <timing type="post-hoc"> <delay number="1" unit="day" constraint=">="/> </timing> </payment> </access-agreement> </trade-proposal></pre>	<p>cont.</p> 
---	--

Figure 5. Example of a trade proposal requesting to buy resources in OCEAN’s planned constraint-based resource description language.

(Turing-universal) custom matching criteria can be applied later by the buyer and seller in the negotiation stage, if desired. The matching process eventually leads to a one-on-one negotiation between buyer and seller, at which point complex custom strategies for deal optimization can be used. A very simple “yes-or-no” negotiation strategy is provided in the system, as a starting point for developers.

Although the precise design of our resource-description / constraint language is still evolving, Figure 5 shows a mocked-up example to give the idea of the type of language features that we are considering. This is a hypothetical request from a buyer for an access agreement that allows him to lease a certain package of resources, consisting of (1) exclusive access to a ≥ 2 GHz Intel Pentium4 or better machine running RedHat (v.7.3 or greater) Linux and with Globus ≥ 3.0 and the linpack library pre-installed, (2) shared access to a Beowulf cluster of at least 10 machines, also running linux, (3) network storage of ≥ 10 GB, and (4) unlimited usage of a ≥ 100 GBps Ethernet connecting the machines and the network storage. The buyer wants to access this system immediately in any desired usage pattern, for a period not less than 2 hours, but no more than 1 day, paying by the hour at a flat rate of not more than US\$1 per hour. Payment will be made via a direct transfer from his OCEAN account to the seller’s account, but the buyer wants to be able to wait at least 1 day between completion of the resource usage and payment, presumably to give a human user a chance to manually inspect and dispute whether the resources were provided as agreed. Namespace qualifiers and schema URNs are omitted for readability.

Note the use of inequality constraints on the numeric values of various parameters. Such constraints, together with Boolean combinations (conjunction/disjunction) thereof, will be automatically processed by the matching system at each node, to determine whether any of the locally-available trade proposals are compatible with the given proposal. For “buy” trade proposals, any unspecified details are assumed to be unconstrained, whereas for “sell” trade proposals, only the explicitly specified details are assumed to be available. The matching node will propagate the buyer’s trade proposal to a subset of its peers, in a fashion intended to maximize the number of matches found, while limiting the resources consumed in the matching system. (More details below.) A successful match leads to communication of the matching trade proposal back to the buyer, and possible initiation of detailed negotiations between the buyer and seller.

C.6.1. MARCOPOLLO P2P Matching/Searching Protocol

MARCOPOLLO [26] (named after the swimming-pool game) is our newest experimental peer-to-peer resource/request matching algorithm. It is based on undirected peering relationships between nodes, which constitute mutual agreements to directly process and/or forward network search requests from each other. A node joins the network by querying an initial peer, which, if its peer quota is full, pushes the new node further out towards the network edges. A node can conveniently retrieve a list of all nodes that are located within n peering hops (see Figure 6), via a *Marco* (“Who’s there?”) broadcast, to which all nodes within the given range reply with a *Polo* (“Hello I’m here.”) response. As in our earlier P2P algorithms, a node can gradually optimize its “location” in the P2P network by modifying its set of active peering arrangements, so as to be located nearby (within few hops) of nodes that have a record of past matching success. See §C.6.2.

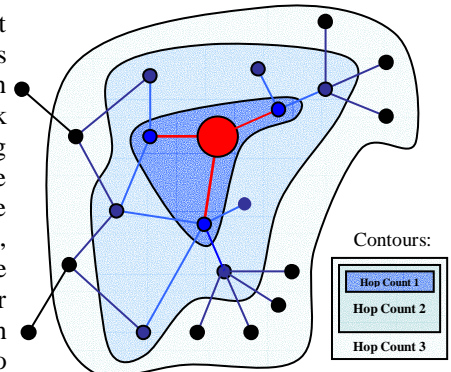


Figure 6. Hop Contours in MarcoPolo P2P Net

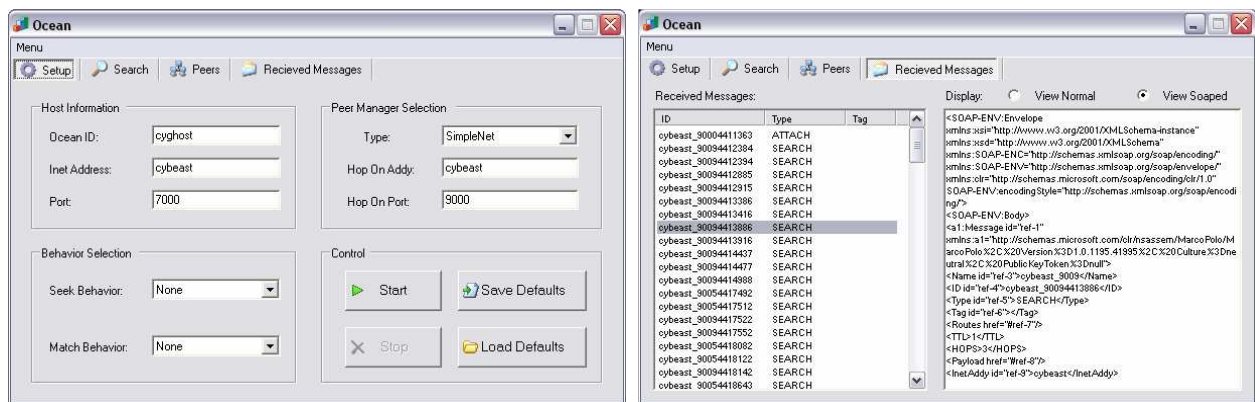


Figure 7. Screenshots of MarcoPolo experimental P2P simulator testbed (coded in Microsoft .NET).

C.6.2. UNDERTOW P2P Network Self-Optimization Algorithm

In peer-to-peer message-passing systems for resource/request matching, we wish the network to be adaptively optimized in response to dynamically-changing patterns of availability of resources that meet the needs of currently popular requests. Given fast-changing availability and pricing information, maintaining a static index (whether centralized or replicated) of all nodes on the network would be prohibitive. Instead, we use a P2P architecture in which each node only maintains local knowledge about its immediate logical neighbors (direct peers). Nevertheless, by collecting statistics about the past matching performance of its neighbors, a node can adaptively modify and customize its set of peers over time, in effect “moving” logically through the network. If there have been $nTrials$ queries to a neighbor, of which $nSuccesses$ have resulted in successes, then the correct Bayesian expectation of future successes (obtained by integrating the Beta-multinomial distribution) [27] is given by the expression $(nTrials + trialBias)/(nSuccesses + successBias)$, where $trialBias/successBias$ corresponds to the prior probability of success before collecting any data. This machine learning approach can be tuned to be maximally sensitive to the most recent data by exponentially decaying the $nTrials$ and $nSuccesses$ statistics over time. In UNDERTOW, which is a variation on this general approach, peers are rated not just on the total number of successes resulting from queries passed through them, but more delicately, based on the number of successes *per unit distance* (in hops) that the message was forced to travel before success. This is one way to bias the results towards nodes that not only match successfully, but that do so most efficiently. An alternative tried earlier in another experimental algorithm, WAVE, was to count as successes only those matches that resulted in an eventually successful transaction, which naturally tended to be the ones that were found most quickly. However, UNDERTOW empirically tends to evolve much faster than WAVE, because it utilizes more of the available data produced during the matching process.




Figure 8. UNDERTOW algorithm. When any specific immediate peer (colored circle) is selected as a candidate for removal from our peer list, we consider replacing it with one of the best-performing peers from among the tree of nodes reachable within a few hops of the removed peer (highlighted in correspondingly-colored boxes).

C.6.3. Lightweight Code-Migration Facility for Java

In addition to supporting compatibility with existing Grid frameworks such as Globus, in OCEAN we are also implementing our own, lightweight facility for code migration in Java and C# which can be utilized by users who prefer to avoid the overheads associated with a complete Grid infrastructure. The following stripped-down version of a working example in Java is intended to illustrate just how ridiculously easy we intend OCEAN to be for developers to learn and begin using. Although the detailed capabilities of OCEAN are quite sophisticated and powerful, a developer who is happy to accept all the built-in system defaults (for negotiation strategies, payment methods, etc.) can begin writing applications in Java and C# that buy and utilize remote resources with very little effort. Even somewhat more customized uses of OCEAN are intended to be comparably easy as well.

```
import ocean.tsm.*, ocean.trading.*;
/** The Hello world agent - Continually just jumps from one machine to another,
    Paying for access as it goes, at market prices, and printing "Hello" on the
    console of each machine used. */
public class Hello extends DefaultOceanTask implements Serializable {
    public Hello() { }
    public void run() {
        System.out.println("Hello World!");
        /* Buy generic OCEAN code-migration access to some remote machine, using
           default trade-proposal parameters, then use the Task Spawning &
           Migration system (TSM) to migrate (transfer) the current task to
           that machine. */
        TSM.migrateTask(this, Trader.ProcureRemoteTSMHost());
    }
    public static void main(String args[]) {
        TSM.spawnTask(new Hello());} // Start up the Hello mobile-agent.
```

C.7. Security Concerns in OCEAN

Security is of course a major concern for any internet-connected system. Users want security, and OCEAN attempts to make all reasonable accommodations to provide it. But we believe that the art of good systems design lies in creating an intelligent compromise between security and other needs for the majority of the most important applications. The Internet itself is wildly successful despite the fact that its own core low-level protocols

(TCP,IP,DNS) provided essentially no support for security in their original incarnations, and to a large degree even today. The Internet's benefits of interoperability, simplicity, and reliability outweighed its security issues for millions of early adopters. We view OCEAN in the same way; that is, we feel that the basic service it provides is so beneficial that many people will use it even if its security provisions are not perfect. Our goal is to implement enough security so that the majority of plausible applications of OCEAN are facilitated with reasonably sufficient protections, but not so much security that the system's design and day-to-day operation becomes bogged down with constant security-related headaches. In view of this overall stance towards security, we have implemented or are planning to implement the following security provisions in our primary public release of OCEAN:

1. **Registration validation.** Verification of a user's real-world identity before automatic generation of digital certificates by the OCEAN Certificate Authority will initially be cheaply approximated using confirmation-soliciting email messages sent to the user's declared email address, rather than by some more labor-intensive approach (*e.g.*, having a human employee telephone the registrant). This approach is vulnerable to man-in-the-middle attacks through email interception and spoofing, as well as to the use of email anonymizers, but it is nevertheless considered secure enough to be the basis of registration procedures for many existing commercial websites. For many applications, the low cost of this approach will outweigh its risks, practically speaking.

2. **Digital signatures.** Sales agreement contracts will be protected against forgery and repudiation by the use of standard XML digital signatures, which we have already implemented. The signatures are verified using public-key certificates that are created during the registration process. This is easy and it provides a degree of protection that is as strong as the method used for identity verification during the registration process. Signatures are also required for performing manipulation of OCEAN accounts that involve real money (*e.g.*, transfers and payouts).

3. **Encryption.** Standard TLS (Transport Layer Security) and/or SSL (Secure Socket Layer) stream encryption techniques will be used to ensure privacy of sensitive data, *e.g.*, during registration and account creation protocols.

4. **Access control.** Access to computational resources made available on the OCEAN will be carefully guarded by requiring each access request to be associated with (signed with the same public key as) a digitally-signed purchase agreement for use of those resources. If a user abuses the purchased resources, it will be possible to trace them through the email address verified in the registration process, and hold them accountable for the abuses.

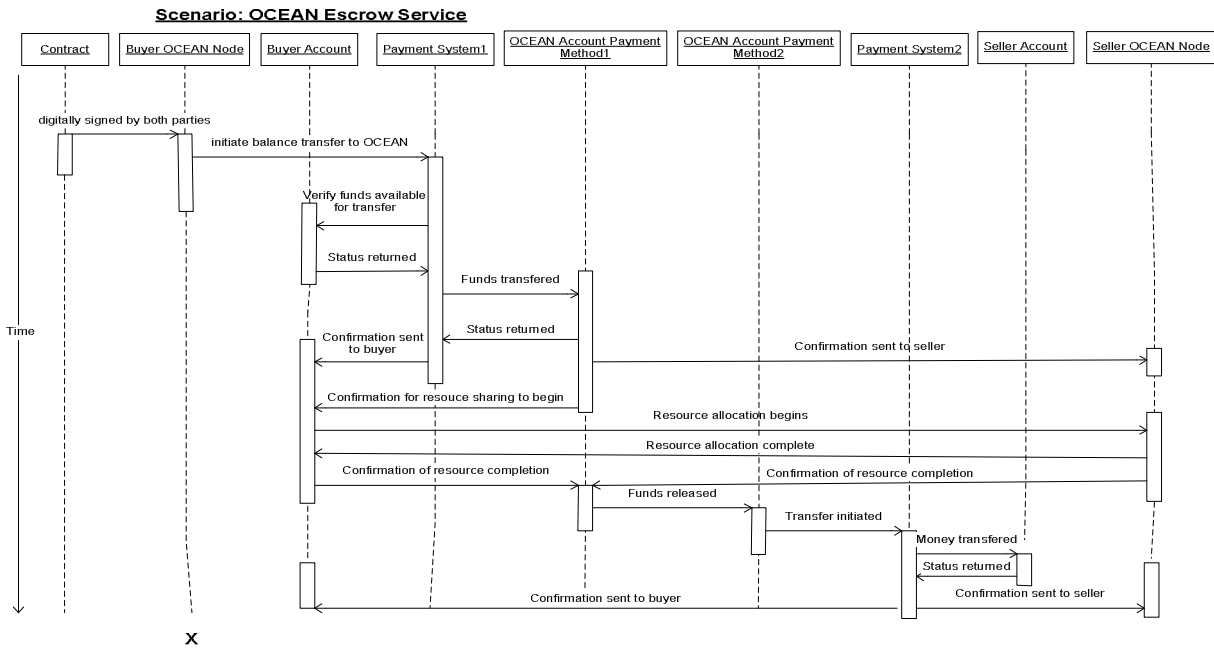
5. **Sandbox environments.** Providers of Java and CLR (Microsoft .NET Common Language Runtime) runtime environment platforms will benefit from virtual-machine-level security guarantees. The activities of client jobs running in these environments are carefully circumscribed, and in theory cannot step outside the specific access rights granted within the `SecurityManager` context of the particular sandbox. Of course, vulnerabilities in the sandboxes are constantly being found and patched. But OCEAN will be at least as secure as other Java or .NET mobile-code systems, and we feel that the use of these environments constitutes a reasonable accommodation that is in practice fairly unlikely to be compromised in any specific instance. Unfortunately, users of legacy environments (C++/Fortran) cannot depend on language-level security. There are other ways to provide protection, *e.g.*, by intercepting system calls (such as with `ptrace` on Unix) and using OS-level resource control mechanisms (*e.g.*, non-root Unix accounts). However, these are more cumbersome to implement than the Java/C# approaches. Realistically, a tradeoff between security and functionality is needed, and ultimately, it is up to the individual resource provider how much to trust remote users not to abuse the resources, and to negotiate an access agreement that provides the seller sufficient assurances against and compensation for the risk. If a system is compromised, the access control provisions of (4) will often allow the offending user to be found and punished, to deter abuses.

6. **Payment escrow.** One easy provision to discourage transaction fraud is an escrow service where payments for resources are held by OCEAN until the buyer declares satisfaction, or a timeout expires. This prevents either the buyer or seller from being able to quickly "take the money and run" if a dispute arises over whether resources were actually provided as per the sales contract. In case of a dispute over a large-value transaction, a human arbiter can step in and discuss the situation with both parties. Several commercial internet-based payment-escrow services exist already. We would utilize one of these services or integrate a similar functionality into OCEAN directly. We have already done some design work in this direction. (See fig. 9 below.)

7. **Computation certificates.** Certain theoretical techniques (see *e.g.* [11,12,20]) offer the prospect of automatically resolving disputes over computational contracts, by requiring the resource provider to provide a certificate (in the form of a certain type of probabilistically-checkable proof) that actually *proves* (with high confidence) that the specific job submitted by the buyer was in fact run correctly and to completion. However, currently known techniques impose a significant constant-factor overhead on job processing. Finding ways to reduce these overheads is an interesting fundamental research problem that we will investigate.

8. **Intranet deployment.** For organizations most concerned about security breaches, OCEAN can be deployed in local fashion, behind the corporate firewall, on an enterprise intranet. If the organization is small, this will reduce the advantages of OCEAN somewhat, since the organization's internal computing needs may not be very level, and

Figure 9. Message sequence diagram for the OCEAN Escrow Service, an optional payment mechanism that reduces the opportunity for a buyer or seller to gain by committing funds to fraudulent transactions. Funds fronted for the transaction are held “in escrow.” If the buyer and seller disagree whether the conditions of the access agreement were satisfied, the funds are held by OCEAN until the dispute is resolved or one party defaults (fails to reply to inquiries within a certain time).



so resources may still end up being underutilized at times. Nevertheless, it is expected still to be beneficial relative to the alternative of no resource sharing.

C.I. Statement of Work to be Undertaken - C.I.a. Objectives and Significance

The below table outlines our key objectives for each quarter-year before and during the award period. Below each year’s summary, the significance of that year’s work is noted.

Year	Quarter	Research/Theory Goals	System Development Goals
Year 0 (Jan.-Dec. 2003) (pre-award)	1st (Jan.-Mar.)	Business/accounting research ✓	Test P2P system ✓
	2nd (Apr.-Jun.)	Design API for science users	Implement OCEAN accounting system
	3rd (Jul.-Sep.)	Prototype small science apps.	System integration / documentation
	4th (Oct.-Dec.)	Test prototype apps., publish results	Internal OCEAN deployment in our lab
<i>Significance of year-0 work: Even before the funded work period starts, the initial OCEAN prototype and simple scientific test applications will be working and preliminary results published.</i>			
Year 1 (Jan.-Dec. 2004)	1st (Jan.-Mar.)	Design major computational science apps. (Peters/Schmalz/Hammer)	Deployment of OCEAN throughout UF’s College of Engineering
	2nd (Apr.-Jun.)	Science apps. – Development	Local alpha-testing using play money
	3rd (Jul.-Sep.)	Science apps. – Testing	Features/bugfixes based on test results
	4th (Oct.-Dec.)	Publish results of science apps.	Public beta release of OCEAN using real money, offered for free by UF
<i>Significance of year-1 work: After year 1, a working beta release of OCEAN will be publicly available. Subsequent rapid growth of usage due to design advantages, publicity from conferences / discussion boards / word-of-mouth.</i>			
Year 2 (Jan.-Dec. 2005)	1st (Jan.-Mar.)	Investigate market/auction theory	Incorporate beta-test feedback.
	2nd (Apr.-Jun.)	Experiment with alternative auction/negotiation mechanisms	Complete reference implementation.
	3rd (Jul.-Sep.)	Develop tools for collaborative extension of XML resource schemas	Submit OCEAN protocol specs to international standards bodies.
	4th (Oct.-Dec.)	Publish market/collaboration results	Publish final results; plan privatization.
<i>Significance of year-2 work: By the end of year 2, the reference implementation of OCEAN should be a stable, reliable high-performance platform based on sound market theory and flexible, collaborative extension of its standards. At this point we begin to see a large impact (e.g., significant new projects based on OCEAN springing up).</i>			

Year 3 (Jan.-Dec. 2006)	1st (Jan.-Mar.)	Research new advanced security techniques for OCEAN (computation certificates, etc.).	Design commercially-viable version of OCEAN, for operation e.g. by a self-sufficient non-profit subsidiary of UF.
	2nd (Apr.-Jun.)	Implement new security measures	Implement & release commercial vers.
	3rd (Jul.-Sep.)	Test new security measures	Testing & enhancement of release.
	4th (Oct.-Dec.)	Publish new results on advanced security measures within the computer security community.	Try to get papers & press articles about OCEAN into widely-read magazines (e.g., Scientific American, Wired).
<i>Significance of year-3 work: Due to the important new features from theory research in years 2+3, the range of applicability and popularity of OCEAN is further increased. It is now viable as a self-sufficient non-profit entity in which minimal transaction fees cover the costs of operation of the central services. It starts to be a household word.</i>			

C.I.b. Relation to Longer-Term Goals

The proposed work plan directly targets the long-term goals of the OCEAN project, and is therefore rather ambitious, but we believe this speed is necessary due to the current explosion in attention to Peer-to-Peer and Grid technologies. We would like to see that our approach (of easy market-based computation) becomes widely available sooner rather than later, in a flexible and compatible form, so that users will not be faced with a possibly burdensome migration from a less flexible alternative resource-allocation approach to our better one as the technology evolves. In the OCEAN project, we wish to short-cut the technology evolution, and provide the best solution that we can currently envision, with as little delay as possible.

C.I.c. Relation to State of Field and to Work in Progress Elsewhere

This was already addressed extensively in section C.4 above, as well as in the FY2001 edition of this proposal [21]. Many distributed computing projects are currently ongoing, but none that we have seen are nearly as comprehensive as OCEAN in aiming of accomplish all of the key requirements that we believe will be essential in order for distributed computing and Grids to progress from being the exception to being the rule for (perhaps, eventually) the majority of commercial computing.

C.II. General Plan of Work – C.II.a. Activities to be Undertaken

C.II.a.i. Infrastructure Work - Michael Frank (PI), Sanjay Ranka (co-PI), and Abdelsalam Helal (co-PI) will have primary responsibility for supervising design and implementation of the infrastructure. **Dr. Frank**, who started the project and has already supervised a large amount of student work on the project over the last 3 years, will be the lead designer on OCEAN itself. **Dr. Ranka** will focus on the important issue of compatibility with other emerging Grid projects and standards, such as the GriPhyN (Grid Physics Network) project [28] with which he is already involved. **Dr. Helal** will apply his extensive expertise in Internet Computing to refine the OCEAN infrastructure design in a variety of ways, for example to be compatible with the Mobile-Agent, Workflow Management and Pervasive Computing infrastructures that his group has developed. Dr. Meera Sitharam, an informal collaborator in our department, has done some work [29,30,31] on non-cooperative game theory as a market model for network resource allocation, as well as on the effects of pricing, and is willing to serve as a local source of expertise in that area, as we strive to improve our market & pricing mechanisms. In addition, the infrastructure design process will continually incorporate input from the applications-oriented senior personnel and other informal collaborators at UF, as to their needs and requirements for infrastructure capabilities, as they proceed to implement significant applications that also further their own scientific research goals (see below) on top of the OCEAN infrastructure. Also, we are involved in collaborations with other researchers who are working on related infrastructure projects; continuing interactions between the projects will likely influence the design of our OCEAN, as it incorporates ideas from other projects or is designed to be compatible with them. For example, we are collaborating with Rajkumar Buyya of Monash University, Australia regarding interoperability between OCEAN and his Compute Power Market project [32], for example to use OCEAN's peer-to-peer matching system as a pricing mechanism for his system.

C.II.a.ii. Applications Work - This work ensures that OCEAN is truly a strategically beneficial infrastructure that facilitates useful research and development of all manner of applications. Here, we break down the applications work by the senior personnel and other collaborators who will be primarily involved in it.

Sanjay Ranka (co-PI) - **A. Data Mining:** In many practical environments, the data/execution resources required for mining and exploration tools are distributed among sites which may belong to different organizations and may have policies limiting access to data and other resources. We are developing middleware that will enable distributed data mining and exploration in such an environment. Users will be able to explore and mine information using state-of-the-art data exploration and data mining primitives. Data owners will be able to ensure that all users are given access to data based on their privileges, and guarantee that a user does not compromise privacy constraints. System administrators will be able to effectively manage the available resources and provide users

constrained access to them based on the organization's policy. The use of OCEAN will allow bartering of storage, compute resources and access rights for this middleware, facilitating execution in a collaborative environment.

B. Resource scheduling is a critical issue in executing large-scale data intensive applications in a grid. Due to the characteristics of grid resources, traditional scheduling algorithms are not suitable for grid computing. We are developing a fully distributed, fault-tolerant system called SPHINX that provides new distinct functionalities, such as execution time estimation, dynamic workflow planning, enforcement of policy and QoS requirements, and VO-wide optimization of throughput. Sphinx complements OCEAN, which will provide an important set of services that will allow multiple users and organizations to barter resources.

Jorg Peters (co-PI) – A. Computational Geometry Application: Database of vertex neighborhoods in a 3D space partition. A vertex in a 3D space partition can have four neighbors or five neighbors in two different configurations or six neighbors in twelve combinatorially unisomorphic configurations. Enumerating all such connectivity-distinct configurations and pruning out all isomorphic cases is compute-intensive. The result, even though exponentially growing with the number of neighbors, can be stored in compressed form and made accessible as a data base entry. At present, determining the entries for n=15 neighbors requires 15 hours on the best computational facilities in the department. We will leverage the OCEAN system, graph isomorphism algorithms and eigenanalysis packages to analyze associated 3D subdivision fields. The 3D fields will be used for particle distribution simulations in closed environments. Both simulation and visualization require extensive computational resources. **B. Computer Graphics Application: Web-based rendering farm for realistic animation.** Currently high quality animation is achieved by dedicated rendering farms that take the graphics scene through various filters (include ray-tracing, shadows, etc.) Public tools are available but computational resources are still expensive. The OCEAN network would replace this dedicated computational environment to allow realistic rendering of a small movie on a low budget.

Mark Schmalz (co-PI) – Image Processing & Image Algebra Applications - Dr. Mark Schmalz (Co-PI) specializes in image and data compression, image and signal processing (ISP), and computer vision (CV). In the proposed OCEAN project, Dr. Schmalz would apply his expertise in adaptive systems design, network computing, ISP/CV, and performance analysis to (a) develop ISP/CV applications for the OCEAN system, and (b) measure and analyze the performance of a wide variety of image-related tools and applications on the OCEAN platform. Such applications are of key importance to rapidly expanding areas of technology including Internet video, telemedicine, and distance learning. Efficient porting of these applications to OCEAN would be facilitated by an image algebra application programming interface. The proposed effort leverages over 15 years of experience at UF in porting image algebra to a wide variety of sequential and parallel workstations, massively parallel processors, digital signal processors, and adaptive or reconfigurable computing systems [33,34]. After porting image algebra system to the proposed OCEAN platform, Dr. Schmalz further proposes to implement, test, analyze, and evaluate a wide variety of proven ISP/CV applications in image algebra, such as object recognition, video segmentation, image and signal compression, automated error analysis and profiling routines, and medical image processing algorithms [35,36]. This diverse suite of applications is both I/O- and compute-intensive, and provides a flexible yet highly specifiable suite of performance measurement and analysis benchmarks for the OCEAN platform. The proposed benchmarks would extend more common benchmarking program suites such as SPEC or iCOMP benchmarks to include multi-level data abstraction with I/O- and compute-challenged behavior designed specifically for the OCEAN paradigm. Performance would be analyzed in terms of computational error, as well as space and time complexity [37,34].

Joachim Hammer (senior personnel) – Mobile Agents & Data Grids. Dr. Hammer will develop mobile software agents that use the OCEAN infrastructure for migration and communication, as well as a distributed computing platform. In the context of B2B commerce, such agents can search for and gather corporate resources by visiting the relevant trading partners, and if necessary, carry out negotiations on location. This model of bringing the computation to the data is often more efficient than moving large data sets of which only a fraction is really needed. We will also develop a *data grid* to perform massively parallel alignment searches of expressed sequence tags (EST) against proprietary and public genome repositories. Such a data grid, which would make all of the distributed databases and analysis tools appear as if they were available locally, has many challenges including moving and tracking large amounts of data and dealing with heterogeneities at the data, storage, and application level. However, such a grid would enable the solution of problems no existing distributed technology is capable of addressing.

Several informal collaborators at UF and elsewhere have also expressed interest in porting their distributed and high-performance applications to OCEAN, but detailed descriptions of their specific projects are omitted for brevity.

C.II.b. Experimental Methods and Procedures

The design of the OCEAN standards and the development of its prototype will involve normal structured software engineering procedures, with modular design and code reuse, regression testing of components, and acceptance

testing of the complete system versus requirements. OCEAN nodes will first be deployed widely on computing resources throughout the department and university, and the market mechanisms tested in real practice to ensure their stable operation. However, the infrastructure cannot be thoroughly tested and benchmarked without a wide variety of user applications; that is the reason for the application work described above. Standard methods of performance characterization will be applied to determine how efficiently the OCEAN network runs these applications, and how well its performance scales with increasing numbers of processors. The results will be compared to the performance of running comparable applications on pre-existing platforms. The market mechanisms will be tested initially using accounts of “play money” assigned to the various application development groups; multiple applications will be run simultaneously on the OCEAN, and the market mechanisms will be examined for features of price stability, fairness, *etc.* Results of all OCEAN experiments will be recorded methodically for use in research reports.

C.II.c. Preservation/Documentation/Sharing of Research Products

All standards, design documents, source code, application code, experimental data, and results of the OCEAN project will be made permanently publicly available via the web by the UF CISE department. In addition, the major novel elements of the design and the experimental results (*e.g.*, on new P2P algorithms) will be archived via journal publications. A main goal of OCEAN is for it to become a universal standard; if this goal is achieved, some standards organization such as W3C, ANSI, ISO or other consortium eventually will take over the primary responsibility for maintaining and distributing the standard.

C.III. Broader Impacts - C.III.a. Integration of Research and Education

OCEAN has already been the focus of dozens of student-semester’s worth of unfunded graduate and undergraduate research and design projects, and so has played a significant educational role already. We anticipate that the high enthusiasm already displayed by students for this project will increase further once funding becomes available. The student work plays an important role in helping students to learn about critical emerging internet technologies. The other collaborators working in various aspects of OCEAN also have students who will also become trained in distributed programming and the use of OCEAN-like infrastructures. We believe that programming for such platforms will inevitably become an important software engineering skill sooner or later, whether or not OCEAN itself becomes the eventual standard. **New course development.** Dr. Frank is designing a new course for Spring 2004 called *Distributed Computing Technologies*; it will survey key distributed computing technologies, including web-service, cluster, Grid, Peer-to-Peer, and Market-based technologies, emphasizing OCEAN-related technologies, as well as OCEAN itself. The course will bring lessons learned from OCEAN to the classroom.

C.III.b. Participation by Underrepresented Groups

Dr. Frank is active in encouraging participation by women, minorities, and other underrepresented groups in his projects. He often serves in UF’s Undergraduate Minority Mentoring Program and he has worked to help ensure diversity in the student body through his role on the CISE department’s graduate committee. This committee has often striven to ensure that women and students from underrepresented foreign countries are especially encouraged to attend the UF CISE graduate program. The other senior personnel and collaborators on OCEAN themselves hearken from a wide variety of nations and backgrounds, and will help to ensure diversity in the project as a whole.

C.III.c. Enhancements to Infrastructure for Research and Education

The explicit goal of OCEAN is to provide an extensive, scalable, easily usable infrastructure for distributed computing in all its forms. As such, it will directly facilitate an enormous variety of new, large-scale research projects that would not otherwise have been possible, projects that harness large numbers of distributed servers, or widely-wandering mobile agents. It also facilitates education in the design and use of computationally-intensive tools and technologies, since it provides an environment in which students can cheaply harness vast amounts of computational power for their computational projects for short periods of time. Many schools might not be able to afford to purchase large clusters of computing servers for students to run their big projects on – but schools or even individual students might easily be able to afford to rent the needed power for short periods. The educational applications are especially appropriate since their needs are so “bursty” - student projects often are run in large waves late in the semester, rather than continually over time – so school-owned, fixed resources for student projects are currently chronically underutilized or overloaded during prime seasons. OCEAN would solve this problem by leveling resources between all types of users (students, researchers, industry) who have different schedules.

C.III.d. Dissemination of Results

The OCEAN project has received preliminary press exposure in PC World magazine, the ABC News local TV affiliate in Orlando, and various newspapers around Florida, which led to inquiries from people eager to begin using the system. As the project develops, it will be further publicized through additional conference publications,

announcements on relevant Usenet newsgroups and mailing lists, link sharing with related projects (which has started already), and so forth. Third-party application development houses will be contacted and provided with free training sessions to teach them how to write applications for OCEAN, by showing them examples that our grant-funded applications work will develop. Once the OCEAN services are available, and anyone can begin to earn money by running applications for end-users, we expect that word-of-mouth will contribute significantly to further dissemination and growth of the OCEAN concepts and standards. As grass-roots attention is drawn to the project, our large and growing collection of original research on market-based distributed computing will be more widely seen and understood.

C.III.e. Potential Benefits to Society at Large

In making computational resources a more liquid commodity, OCEAN will have some extraordinary benefits for society at large: (1) It will make computation cheaper for everybody, by reducing the waste of unused cycles caused by varying individual usage. (2) At the same time, it will aid the computer industry by making computers themselves more intrinsically valuable on average – if a user can potentially recoup the cost of a computer by selling its unused cycles on the OCEAN, rather than wasting them, he may be willing to pay a higher price for the machine to begin with. (3) It will enable new kinds of computing applications that would not otherwise have been feasible, most especially, applications that require large amounts of parallel power for relatively short periods of time, and applications that require mobile agents. Due to the OCEAN's open architecture, applications that harness the power of OCEAN can be developed by third-party vendors and then used by anybody – so that in the future, the financial analysis package on one's desktop, for example, may silently harness the power of thousands of distributed servers so as to return an answer more quickly. Or, an airline-reservation web site might harness a vast, distributed database running on top of OCEAN to rapidly process an itinerary-search transaction. Finally, by supporting important research using computationally intensive basic science & engineering applications, OCEAN will facilitate the more rapid advancement of basic science and technology, with the resulting indirect benefits to society.

C.IV. Results from Prior NSF Support

Dr. Frank and Dr. Ranka have not had related NSF support within the last 5 years. However, Dr. Frank's graduate work was supported by an NSF Graduate Fellowship from 1992-1995, and it was during this period that Frank first studied market-based computing under Dr. Jon Doyle, then of the MIT Lab for Computer Science. A fellow student of Doyle during this period (R.J. Bodkin) suggested the idea of bringing market-based computing to the internet to Frank in 1996. Frank's NSF graduate fellowship thus led indirectly to the very conception of OCEAN, as well as supporting Frank's Master's thesis work [27] on decision-theoretic techniques in AI. Some of the ideas in this thesis led to the machine-learning techniques that are being used in OCEAN's experimental Peer-to-Peer network today.

Dr. Peters' prior NSF support: **National Young Investigator Award CCR-9457806, *Surface Splines Over Irregular Meshes, 1994-2000*** advanced fundamentals and techniques for modeling free-form surfaces, in particular developed the theory of Geometric Continuity. The grant supported 5 graduate students and 7 undergraduate research assistants. The grant resulted in 20 journal publications and 10 refereed conference articles. **NSF grant CCR-9901894, 1999-2003** produced the first explicit derivation of Gauss and mean curvature of subdivision surfaces at their extraordinary points [38] and reconciled subdivision and polynomial surface definitions [39,40]. The fundamental problem of constructing piecewise polynomial C^k surfaces of arbitrary topology with least degree was solved in [41]: C^k surfaces with any n patches joining with geometric flexibility d to fill an n -sided hole in a C^k spline surface can now be constructed with patches of degree $d+k$. [42] defined N-patches in the DX8 graphics standard used widely by computer game artists. Several other publications have also resulted from this grant.

Dr. Helal & Dr. Hammer: Abdelsalam (Sumi) Helal and Joachim Hammer have a CISE CCR award, for Proposal number 0100770, titled "Adaptive Synchronization Framework Supporting Device-independent Mobile Computing." (\$150,000, 9/01-8/03). The objective of this research is to develop a three tier architecture that enables device-independent disconnected operation and transparent synchronization of data. The Coda system has been used as a basis for this research and new on-the fly algorithms have been developed to compute the user working data set more efficiently. Publications from this research so far include [43,44,45,46].

Dr. Hammer and Dr. Schmalz have been awarded an exploratory NSF grant titled "Scalable Enterprise Systems: Theory and Methodologies to Support the Operation of Flexible Production Networks." The goal of this research is to develop the next-generation wrapper, data engineering, and analysis technologies to support extended enterprise collaboration. So far, we have developed the initial concept and demonstrated an early prototype that can extract resource and cost information from legacy project resource planning software. Initial results from our prototype implementation demonstrate feasibility. This has allowed us to begin work with industrial partners including Pratt&Whitney Space Propulsion Operations. Preliminary results have been described in [47].

D. References Cited

- 1 The OCEAN group, OCEAN project homepage, <http://www.cise.ufl.edu/research/ocean>.
- 2 Chaitanya Chokkareddy, "Automated Negotiations in O.C.E.A.N.," Masters thesis, CISE Dept., Univ. of Florida, 2002, <http://www.cise.ufl.edu/research/ocean/master.pdf>.
- 3 Nitin Chawla, "Registration and Authentication protocol for OCEAN," Masters thesis, CISE Dept., Univ. of Florida, 2002, http://www.cise.ufl.edu/research/ocean/Nitin_Chawla/Chawla_N.pdf.
- 4 Sriram Kumar, "Trading System Design and Implementation in OCEAN," Masters thesis, CISE Dept., Univ. of Florida, 2002, <http://www.cise.ufl.edu/research/ocean/Thesis.pdf>.
- 5 Sahib Singh Wadhwa, "Signing and Validating Contracts in OCEAN," Masters thesis, CISE Dept., Univ. of Florida, 2002, http://www.cise.ufl.edu/research/ocean/sahib/Wadhwa_S.pdf.
- 6 Mark Joseph Tobias, "Resource and Requirement Schemas Applied to Auctioning in a Computational Market," Master's thesis, CISE Dept., Univ. of Florida, 2001, http://www.cise.ufl.edu/research/ocean/MarkTobias/-mjtobias_th.pdf.
- 7 Heeyong Park, "Peer List Update Manager (PLUM) Implementation in OCEAN," Master's thesis, CISE Dept., Univ. of Florida, 2002, http://www.cise.ufl.edu/research/ocean/Heeyong_P/Park_H.pdf.
- 8 Sama Govindaramanujam, Cyrus Harrison, Erwin Jansen, Sriram Kumar Nallan, Sahib Singh, and Michael P. Frank, "Locating Suitable Resources in OCEAN," paper accepted for poster presentation at HiPC-02, the International Conference on High-Performance Computing, Bangalore, India, Dec. 18-21, 2002, http://www.cise.ufl.edu/research/ocean/paper_v2.pdf.
- 9 World-Wide Web Consortium, "Extensible Markup Language (XML)," <http://www.w3.org/XML/>.
- 10 Microsoft Corporation, "About the Common Language Runtime (CLR)," http://www.gotdotnet.com/team/clr/-about_clr.aspx.
- 11 László Babai, Lance Fortnow, Leonid A. Levin, and Mario Szegedy, "Checking Computations in Polylogarithmic Time," preprint, 1991, <http://citeseer.nj.nec.com/babai91checking.html>.
- 12 Silvio Micali, "Computationally Sound Proofs," *SIAM J. on Computing* **30**(4):1253-1298, <http://epubs.siam.org/sam-bin/dbq/article/28495>.
- 13 World-Wide Web Consortium, "Simple Object Access Protocol (SOAP) 1.1", <http://www.w3.org/TR/SOAP>.
- 14 World-Wide Web Consortium, "Web Services Description Language (WSDL) 1.1", <http://www.w3.org/TR/-wsdl>.
- 15 "The Globus Project," <http://www.globus.org/>.
- 16 Open Grid Services Architecture Working Group, home page, <http://www.ggf.org/ogsa-wg/>.
- 17 Karen J. Campbell, "Marketing Plan for the Online Computational Exchange & Auctioning Network," term project for Marketing Management, MBA565, St. Leo University, April 2003, <http://www.cise.ufl.edu/-research/ocean/business.htm>.
- 18 Alan George, "HPC Forum in the College of Engineering," website, <http://www.hcs.ufl.edu/coe-hpc/>.
- 19 Michael P. Frank, "Brief Proposal for Experimental Campus-Wide Deployment of OCEAN: a P2P market for dynamic allocation of computational resources," proposal to UF HPC community, <http://www.cise.ufl.edu/-research/ocean/UF-proposal.doc>.
- 20 Jonathan D. Bright, "Checking and Certifying Computational Results," Ph.D. dissertation, Johns Hopkins, 1994, <http://citeseer.nj.nec.com/470539.html>.
- 21 Michael P. Frank, Joachom Hammer, Mark Schmalz, Abdelsalam Helal, and Jorg Peters, "OCEAN: A Highly Liquid Market for Distributed Computing," proposal to *NSF/CISE/ANIR/STI – Strategic Technologies for the Internet*, June 2001, <http://www.cise.ufl.edu/research/ocean/proposals.html>.
- 22 The Distributed.net homepage, <http://www.distributed.net>.
- 23 Bear Stearns, "Internet Infrastructure & Services: Introducing Internet 3.0: Distribution and Decentralization Change Everything," market analysis report, May 2001.
- 24 Lehman Brothers, "Peer to Peer: Is There Hope for the Hype?," equity research report, May 2001.
- 25 Philip Pidot and Edmund Wilson, "Vax: The Virtual Asset Exchange," Business Plan, Mar. 2002.
- 26 Cyrus Dashiell Harrison, "Self Evolving Peer to Peer Networks for Service Discovery: Design of MarcoPolo Matching Network for the Open Computation Exchange and Arbitration Network (OCEAN)", senior honors thesis, CISE Dept., Univ. of Florida, Apr. 2003, http://www.cise.ufl.edu/research/ocean/cdh/cyrus_thesis.doc.
- 27 Michael P. Frank, *Advances in Decision-Theoretic AI: Limited Rationality and Abstract Search*, Master's thesis, MIT Lab for CS, May 1994, <http://www.cise.ufl.edu/~mpf/papers/Frank/Frank-94/Frank-94.html>.
- 28 GriPhyN—Grid Physics Network, <http://www.griphyn.org>.

-
- 29 K. Park, M. Sitharam, S. Chen, "Quality of service provision in noncooperative network environments," Purdue University Technical Report, Aug. 1997.
 - 30 Kihong Park, Meera Sitharam, Shaogang Chen, "Quality of Service Provision in Noncooperative Networks: Heterogeneous Preferences, Multi-Dimensional QoS Vectors, and Burstiness," Proceedings of the International Conference on Information and Computation Economics (ICE), Oct. 1998.
 - 31 Kihong Park, Meera Sitharam, Shaogang Chen, "Quality of service provision in noncooperative networks with diverse user requirements," *Decision Support Systems, Special Issue on Information and Computation Economics*, **28**:101-122, 2000.
 - 32 "The Compute Power Market (CPM) Project," <http://www.compute-power.com>.
 - 33 Ritter, G.X. and J.N. Wilson. Handbook of Computer Vision Algorithms in Image Algebra, Second Edition, CRC Press (2001).
 - 34 Schmalz, M.S., G.X. Ritter, and F.M. Caimi. "Space, time, error, and power optimization of image compression transforms", Proceedings SPIE 4122:176-189 (2000).
 - 35 Schmalz, M.S., Y.S. Natour, C.M. Sapienza, and G.X. Ritter. "Accuracy of Laryngeal Measurements from Digital Endoscopic Imagery. 2. Effects of Compression", Accepted for publication in Proceedings of the SCI 2001 Conference, Orlando, FL (2001).
 - 36 Schmalz, M.S. and G.X. Ritter. "Detection and characterization of motion in video compression", *Proceedings SPIE 4793*:254-266 (2002).
 - 37 Schmalz, M.S. "Automated analysis and prediction of accuracy and performance in ATR algorithms", Proceedings SPIE 3079:249-272 (1997).
 - 38 J. Peters and G. Umlauf. Gaussian and mean curvature of subdivision surfaces. In R. Cipolla and R. Martin, editors, *The Mathematics of Surfaces IX*, pages 59-69. Springer, 2000.
 - 39 J. Peters. Patching Catmull-Clark meshes. In Kurt Akeley, editor, *Siggraph 2000, Computer Graphics Proceedings*, Annual Conference Series, pages 255-258. ACM Press / ACM SIGGRAPH / Addison Wesley Longman, 2000.
 - 40 J. Peters. Smooth patching of refined triangulations. *ACM Transactions on Graphics*, 2001.
 - 41 J. Peters, "C² Free-Form Surfaces of degree (3,5)," *Computer Aided Geometric Design*, **19**(2):113-126, 2002.
 - 42 Alex Vlachos, J. Peters, C. Boyd, J. L. Mitchell, "Curved PN Triangles," Bi-Annual Conference Series 2001 Symposium on Interactive 3D Graphics, pp. 159-166, ACM Press, 2001.
 - 43 Jinsuo Zhang, Abdelsalam (Sumi) Helal and Joachim Hammer, "UbiData: Ubiquitous Mobile File Service," Proceedings of the ACM Symposium on Applied Computing (SAC), Melbourne, Florida, March 2003
 - 44 M. Lanham, A. Kang, J. Hammer and A. Helal , "Format-Independent Change Detection and Propagation in Support of Mobile Computing," Proceedings of the XVII Brazilian Symposium on Databases (SBBD), Gramado, Rio Grande do Sul, Brazil, October, 2002
 - 45 Khushraj, A. Helal and J. Zhang, "Incremental Hoarding and Reintegration in Mobile Environments," Proceedings of the IEEE/IPSJ International Symposium on Applications and the Internet (SAINT), February 2002 in Nara, Japan
 - 46 Helal, J. Hammer, A. Khushraj and J. Zhang, "A Three-tier Architecture for Ubiquitous Data Access" Proceedings of the First ACS/IEEE International Conference on Computer Systems and Applications, June 2001, Beirut, Lebanon.
 - 47 W. O'Brien and J. Hammer, "Robust mediation of construction supply chain information," to appear in Proceedings of the ASCE Specialty Conference on Fully Integrated and Automated Project Processes (FIAPP) in Civil Engineering, Blacksburg, VA, September 2001.

Michael P. Frank

Biographical Sketch

(i) Professional Preparation

- | | | |
|--------------------------------|------------------------------|--------------------------|
| • Stanford University | Symbolic Systems | B.S. w/distinction, 1991 |
| • Massachusetts Inst. of Tech. | Elec. Eng. & Comp. Sci. | M.S., 1994 |
| • Massachusetts Inst. of Tech. | Elec. Eng. & Comp. Sci. | Ph.D., 1999 |
| • Massachusetts Inst. of Tech. | Artificial Intelligence Lab. | Postdoc, summer 1999 |

(ii) Appointments

- Univ. of Florida, Electrical & Computer Engineering Dept., affiliate assistant prof., July 2001-present.
- Univ. of Florida, Computer & Information Science & Eng. Dept., assistant prof., Aug. 1999-present
- MIT, Artificial Intelligence Lab., postdoctoral researcher, Summer 1999.
- MIT, AI Lab, graduate research assistant on the Pendulum reversible computing project, 1995-1999.
- NASA Ames Research Center, Deep Space One project, summer research intern on simulator, 1996.
- IBM T. J. Watson Res. Ctr., summer research internships developing handwriting recognition, 1994-95.
- NEC Research Institute, summer research internship on decision-theoretic search techniques, 1993.
- MIT, Lab for Computer Science, graduate research assistant on two projects, 1992-95.
- SRI International, summer research internship developing Agent & HCI testbed environments, 1990-91.
- Stanford University, Center for the Study of Language & Information, summer research intern, 1989.

(iii) Publications

Up to 5 related publications:

- Sama Govindaramanujam, Cyrus Harrison, Erwin Jansen, Sriram Kumar Nallan, Sahib Singh, and Michael P. Frank, "Locating Suitable Resources in OCEAN," paper accepted for poster presentation at HiPC-02, the International Conference on High-Performance Computing, Bangalore, India, Dec. 18-21, 2002. http://www.cise.ufl.edu/research/ocean/paper_v2.pdf.
- Michael P. Frank, "OCEAN: The Open Computation Exchange & Auctioning Network," talk presented to the Harris Lab research group, CISE dept., summer 2001.
- Michael P. Frank, "OCEAN: The Open Computation Exchange & Arbitration Network: An Open Platform and Commodities Market for Distributed Computation," business proposal to Cenetec technology incubator firm, Nov. 2000.
- Michael P. Frank, "The O.C.E.A.N. Project: An Open Computation Exchange & Arbitration Network," MIT AI Lab student seminar, Feb. 1997.

Up to 5 other significant unrelated publications:

- Michael P. Frank, *Nanocomputer Systems Engineering*, forthcoming book, CRC Press, 2004.
- Michael P. Frank, "Nanocomputers—Theoretical Models," invited article (review chapter) for the *Encyclopedia of Nanoscience and Nanotechnology*, Hari Singh Nalwa, ed., American Scientific Publishers, 2003. Manuscript at <http://www.cise.ufl.edu/research/revcomp/Nanocomputers.doc>.
- Michael P. Frank, "Nanocomputer Systems Engineering," proceedings of the 2003 Nanotechnology Conference and Trade Show, held Feb. 23-27, San Francisco, CA. <http://www.cise.ufl.edu/research/revcomp/theory/NanoTech2003/Frank-NanoTech-2003.doc>.
- Michael P. Frank, "Physical Limits of Computing," *Computing in Science and Engineering*, 4(3):16-25, May/June 2002. <http://www.cise.ufl.edu/research/revcomp/physlim/plpaper.html>.
- Michael P. Frank and Thomas F. Knight, Jr., "Ultimate Theoretical Models of Nanocomputers," *Nanotechnology*, 9(3):162-176, 1998, also presented at the 5th Foresight Conference on Molecular Nanotechnology, Palo Alto, CA, Nov. 1997. <http://www.cise.ufl.edu/~mpf/Nano97/abstract.html>.

(iv) Synergistic Activities

- **Course Development:** I am developing a course based on the OCEAN project's research called *Distributed Computing Technologies*, to be taught Spring 2004, which will cover distributed computing techniques with a focus on emerging technologies such as web services technologies, cluster technologies, grid technologies, peer-to-peer technologies, and market-based computing. The course will give students hands-on experience with the latest state-of-the-art technologies. I also developed a new course on *Physical Limits of Computing* spun out of my other research project, and have authored and contributed to the academic community complete sets of web-accessible course materials on 4 different courses I have taught.
- **Graduate Student Training:** I currently have 1 Ph.D. student on the project (Erwin Jansen) who is being trained in theoretical work and is also getting experience with project management by supervising the Masters' students. I am training many Masters' students in system design and development; I have graduated 5 Masters' students who have produced theses reporting on a component of the OCEAN system that each student designed and developed. Several more Masters' theses are forthcoming. About 10 other Masters' students have also contributed to OCEAN in individual study projects. I also have several graduate students on my other (non-OCEAN-related) research projects.
- **Involvement of Undergraduates** in OCEAN project research: I have had about 8 senior-year engineering students work on OCEAN for their capstone senior design project. These students have developed software components for OCEAN and written senior design theses on the topic. I have had one student (Cyrus Harrison) also do a senior highest-honors thesis describing his work on OCEAN.
- **Encouragement of Underrepresented Groups.** I have twice served in UF's *UMMP* (Undergraduate Minority Mentoring Program) to personally assist members of minority groups in adjusting to UF during their Freshman year. I have twice served on the department's Graduate Admission Committee and helped push for the admission of students from underrepresented foreign countries and for disadvantaged students from within the US. I have had many minority students work on my projects. My best student so far was a minority student who I recommended for a Hertz graduate fellowship.
- **Service to Scientific Community.** I have been helping the scientific community by reviewing papers for a number of conferences. I have also served on an NSF review panel, for the SBIR program.

(v) Collaborators & Other Affiliations

(a) Collaborators & Co-Editors, last 48-months:

- Hari Singh Nalwa (American Scientific Publishers) – Editor of Encyclopedia I published in.
- M. Josephine Ammer (UC Berkeley) – Co-author on a paper.
- Sama Govindaramanujam, Cyrus Harrison, Erwin Jansen, Sriram Kumar Nallan, Sahib Singh (graduate students, U. of Fla.) – Co-authors on a paper.
- Matthew Laundon (NSTI) – Nanotechnology executive director, published my conf. paper.
- Meera Sitharam (UF CISE), Paul Avery (UF Physics), Alan George (UF ECE), Ali Haghghat (UF NRE), Wei Shyy (UF AeME), Chris Carothers (RPI) – Informal collaborators on OCEAN.

(b) Graduate and Postdoctoral Advisors:

- Jon Doyle (NC State) – Master's thesis advisor, then at MIT.
- Peter Szolovits (MIT LCS) – Group leader during Masters' work.
- David Gifford (MIT LCS) – Temporary Ph.D. research advisor.
- Thomas F. Knight, Jr. (MIT AI Lab, Permabit) – Ph.D. & postdoc advisor
- Norman H. Margolus (MIT, Boston Univ., Permabit) – Ph.D. co-advisor
- Gerald J. Sussman, Gill Pratt (MIT LCS) – Ph.D. committee members

ABDELSALAM (SUMI) HELAL, Ph. D.

Computer & Information Science & Engineering Department
University of Florida, Gainesville, FL 32611
<http://www.cise.ufl.edu/~helal>

A. PROFESSIONAL PREPARATION

- B.S. in Computer Science and Automatic Control, Alexandria University, 1982
- M.S. in Computer Sciences, Purdue University, 1989
- Ph.D. in Computer Sciences, Purdue University, 1991

B. APPOINTMENTS

- Aug 1998 – present: Associate Professor of Computer & Information Science & Engineering, University of Florida (tenured).
- June 1996-May 1998, Senior Member of the Technical Staff and a Project Leader, Microelectronics and Computer technology Corporation (MCC), Austin, Texas.
- Aug 1995-May 1996, Visiting Professor, Department of Computer Sciences, Purdue University, West Lafayette, IN.
- Aug 1991-Aug 1995, Assistant Professor, Dept of Computer Science Engineering, U. of Texas at Arlington.

C. PUBLICATIONS

Related Publications

- Sumi Helal and Jingting Lu, "E-service Based Information Fusion: A User-Level Information Integration Framework," Proceedings of the Fourth Workshop on Technologies for E-Services, In conjunction with VLDB 2002, Hong Kong, August 2002
- Sumi Helal, Stanley Su, Jie Meng, Raja Krithivasan, and Arun Jagatheesan, "The Internet Enterprise," Proceedings of the Second IEEE/IPSJ Symposium on Applications and the Internet (SAINT'02), Jan/Feb 2002, Nara, Japan
- Stanley Y. W. Su, Jie Meng, Raja Krithivasan, Seema Degwekar and Sumi Helal, "Dynamic Inter-enterprise Workflow Management in a Constraint-based E-Service Infrastructure," Accepted for publication in Kluwer Electronic Commerce Research Journal. Special Issue on Business Process Integration and E-Commerce Infrastructure. Accepted July 2002.
- J. Meng, R. Krithivasan, S. Su and A. Helal, "Flexibility and Dynamism in E-Service based Workflow Management Systems," Kluwer Electronic Commerce Research Journal (special issue on "Business Process Integration and E-Commerce Infrastructure).
- Amar Nalla, Abdelsalam (Sumi) Helal and Vidya Renganarayanan, "aZIMAs: Almost Zero Infrastructure Mobile Agents System," Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC), Orlando, Florida, March 2002

Other Significant Publications

- Vidya Renganarayanan, Amar Nalla and Abdelsalam (Sumi) Helal, "Internet Agents for Effective Collaboration," Proceedings of the IEEE Fourth International Workshop on Mobile Agents for Telecommunication Applications (MATA'2002), Barcelona, Spain, October 22-24, 2002,
- J. Jing, A. Helal, and A. Elmagarmid, "Client-Server Computing in Mobile Environments," ACM Computing Surveys, volume 31, Number 2, pp117-157. June 1999.
- X. Liu, A. Helal, and W. Du, "Multi-view Access Protocols for Large-Scale Replication," ACM Transactions on Database Systems (TODS), Volume 23, Number 2, pp158-198 June 1998.

- D. Buszko, W. Lee, and A. Helal, “Decentralized Ad-Hoc Groupware API and Framework for Mobile Collaboration,” Proc. of the ACM Inter’l Conference on Supporting Group Work, Sept. 2001, Boulder, CO.
- C. Lee and A. Helal, "Context Attributes: An Approach to Enable Context-awareness for Service Discovery," Submitted to the Third IEEE/IPSJ Symposium on Applications and the Internet, to be held in Orlando, Florida, January 2003

D. SYNERGISTIC ACTIVITIES

- *Main Activities:* Dr. Helal is directing the Harris Mobile Computing and Networking laboratory at the CISE department (www.harris.cise.ufl.edu). He is also Director of Technology Developments of the newly established DOE/NIDRR RERC Center on Pervasive Technology for Successful Aging (www.rerc.ufl.edu). He is currently managing several projects on wireless and mobile networks, Internet Computing, as well as pervasive and proximity-based computing. A complete list of ongoing research conducted by Dr. Helal can be found from the Harris lab web site (<http://www.harris.cise.ufl.edu>).
- *PI of Related Research:* Dr. Helal is currently funded (or has received funding) by the National Science Foundation, Motorola, Harris Corporation, Microsoft Corporation, DOE National Institute on Disability and Rehabilitation Research, and the US Department of Veteran Administration.
- *Professional Leadership:*
 - Chair of the IEEE Computer Society Technical Committee on the Internet (Sept 2000 - Present)
 - Member of the Executive Committee of the IEEE TC on Operating systems (Jan 1995 - Present)
 - Steering Committee Member of: the Symposium on Applications and the Internet (SAINT), The IEEE Workshop on Mobile Computing Systems and Applications (WMCSA), the IEEE Workshop on Internet Applications (WIAPP), and The IEEE Inter’l Symp. on Performance Analysis of Systems and Software.
- *Technical Conference/Workshop Organization:*
 - General Chair of the 2nd IEEE Workshop on Mobile Computing Systems and Applications (WMCSA'99)
 - Program Chair of the 19th IEEE International Performance, Computing and Comm. Conf. (IPCCC-2000).
 - Program Co-Chair of the 10th IEEE Workshop on Research Issues in Data Engineering (RIDE-2000)
 - General Chair of the 21st IEEE International Performance, Computing and Comm. Conf. (IPCCC-2002).
 - Program Co-Chair of the 3rd IEEE/IPSJ Symposium on Applications and the Internet (SAINT'03)
- *Editorials:*
 - Associate Editor and Department Editor of the “Standards, Tools and Best Practices” department of the newly established IEEE Magazine on Pervasive Computing. Also, an associate editor of the magazine.
 - Department Editor of the “Web Technologies” department of the IEEE Computer Magazine.
 - Associate Editor of the Wireless Systems and Mobile Computing journal, Wiley & Sons
 - Associate Editor of the Parallel and Distributed Database Systems Journal, Kluwer Publishers
 - Guest Editor of several special issues (ACM-Baltzer Journal on Mobile Networks and Applications, Kluwer Journal of Intelligent Information Systems, ACM SIGMOD Record, and the International Journal on Cluster Computing, Baltzer Science Publishers.

E. COLLABORATORS

- *UMBC (Anupam Joshi, Yelena Yesha) GE TRearch (Paul Cadehy) Motorola (Jaime Borrás, Dan Lee and Dominik Buszko), Honeywell (Rose May Richardson), IBM Watson Research Center (Giovanni Pacifici, Alla Yousef, Magda Mourad), Fujitsu Labs (Jason Yao), DoCoMo USA labs (Nayeem Islam).*

F. CURRENT AND PAST PH.D. STUDENTS

- *PAST:* Choonhwa Lee (U of Florida), Jinsuo Zhang (Microsoft Corporation), Edwin Hernandez (Microsoft Corporation), Jie Meng (RightOrder Corporation), Yousry Taha (Alexandria University), and Magdy Ahmed (Alexandria University).
- *CURRENT:* Youssef Kaddoura, Wenzheng Zhang, Junhyo Bae, Bekir Arslan, and Ahmed Abukmail.

G. ADVISOR

- *GRADUATE ADVISOR:* Professor Bharat Bhargava, Department of Computer Science, Purdue Univ.

E Biographical Sketch — Jörg Peters

E.a Professional Preparation

Technische Hochschule Darmstadt	Computer Sciences	BS equivalent 1984
	Business,Economics	BS equivalent 1984
University of Wisconsin, Madison	Computer Sciences	PhD 1990
IBM, T. J. Watson Research Center	Manufacturing Research	1990-1991

E.b Appointments

1998—	Assoc. Professor, Computer and Information Sciences and Engineering, tenured	University of Florida, Gainesville, FL
1997–98	Assoc. Professor, CS, tenured	Purdue University, West Lafayette, IN
1992–97	Asst. Professor, CS	Purdue University, West Lafayette, IN
1991–92	Asst. Professor, Math. Sciences	Rensselaer Polytechnic Inst, Troy, NY

E.c Publications

see also <http://www.cise.ufl.edu/research/SurfLab/papers/>

(i) Publications most closely related to the proposed project:

- J. Peters. C^2 free-form surfaces of degree (3,5). *Computer Aided Geometric Design*, 19(2):113–126, 2002. Lowest-degree, piecewise polynomial curvature continuous surface construction to date.
- J. Peters and M. Wittman. Blending basic implicit shapes using trivariate box splines. In Tim N. T. Goodman and Ralph R. Martin, editors, *The Mathematics of surfaces VII*, pages 409–426, Winchester, UK, 1997. Information Geometers. Modeling with the zero set of trivariate, 7-direction boxesplines.
- J. Peters and U. Reif. Analysis of generalized B-spline subdivision algorithms. *SIAM Journal on Numerical Analysis*, 35(2):728–748, April 1998. First complete analysis of Catmull-Clark and Doo-Sabin subdivision; defined the general toolkit for analyzing spline subdivision surfaces at extraordinary points.
- J. Peters and G. Umlauf. Gaussian and mean curvature of subdivision surfaces. In R. Cipolla and R. Martin, editors, *The Mathematics of Surfaces IX*, pages 59–69. Springer, 2000. First explicit derivation of Gauss and mean curvature of subdivision surfaces.
- J. Peters. Geometric continuity. Chapter 8 in *Handbook of Computer Aided Geometric Design*. Elsevier, 2002. Definitive account of concepts and use of geometric continuity in geometric design.

(ii) other significant publications

- D. Lutterkort and J. Peters. Smooth paths in a polygonal channel. In *Proceedings of the 15th annual symposium on Computational Geometry*, pages 316–321, 1999. First solution to the Channel problem for functions.
- D. Nairn, J. Peters, and D. Lutterkort. Sharp, quantitative bounds on the distance between a polynomial piece and its Bézier control polygon. *Computer Aided Geometric Design*, 16(7):613–633, Aug 1999.

Pioneered the study of the exact (non-asymptotic) distance between nonlinear functions and their control structures.

- Alex Vlachos, Jorg Peters, Chas Boyd, and Jason L. Mitchell. Curved PN triangles. In *2001, Symposium on Interactive 3D Graphics*, Bi-Annual Conference Series, pages 159–166. ACM Press, 2001.

Defines the N-patches of the DX8 graphics standard; widely used by computer game artists.

- J. Peters. Smooth patching of refined triangulations. *ACM Transactions on Graphics*, 20(1):1–9, 2001.

Together with Patching Catmull-Clark Surfaces, Siggraph 2000 gives a way tightly approximate quad and triangle meshes by a smooth surface in standard polynomial representations.

- C. Gonzalez and J. Peters. Localized hierarchy surface splines. In S.N. Spencer J. Rossignac, editor, *ACM Symposium on Interactive 3D Graphics*, 1999. First hierarchical spline modeling system that allows for change of topology at and between each level of refinement.

E.d Synergistic Activities

- TOMS Algorithm 783 at Netlib
- Associate Editor of Computer Aided Geometric Design
- Associate Editor of ACM Transactions on Graphics
- Secretary of SIAM Activity Group on Geometric Design

E.e Collaborators and Affiliations

Collaborators: D. Nairn (U Delaware), A. Nasri (American U Beirut), U. Reif (TH Darmstadt), G. Umlauf (U of Kaiserslautern).

Graduate and Postdoctoral Advisors: Carl de Boor (U Wisconsin)

Thesis Advisor and Postgraduate-Scholar Sponsor:

4 PhD students: C. Gonzalez (Disney), D. Lutterkort (Red Hat), X. Wu (current), A. Shiue (current),

4 visitors: G. Umlauf (U of Mannheim, Germany), U. Reif (U of Stuttgart, Germany), Xuli Han (Academia Sinica, China), Kestutis Karciauskas (U of Vilnius, Lithuania).

Sanjay Ranka

Professor of Computer Information Science and Engineering

University of Florida, Gainesville

Email: ranka@cise.ufl.edu Telephone: 352-392-1467

Sanjay Ranka research interests are the areas of large-scale software systems, grid computing, high performance computing, data mining, knowledge based systems and CRM. Most recently, he was the Chief Technology Officer at Paramark, where he developed real-time optimization software for optimizing marketing campaigns. He has also held positions as a tenured faculty positions at Syracuse University and as a researcher/visitor at IBM T.J. Watson Research Labs and Hitachi America Limited.

He has coauthored 2 books, 6 book chapters, 40+ journal articles and 100+ conference articles. He was one of the main architects of the Syracuse Fortran 90D/HPF compiler. He served on the MPI Standards committee. He is a fellow of the IEEE, advisory board member of IEEE Technical Committee on Parallel Processing and a member of IFIP Committee on System Modeling and Optimization.

Education

1988	Ph.D.	Computer and Information Science	University of Minnesota, Minneapolis
1987	M.S.	Computer and Information Science	University of Minnesota, Minneapolis
1985	B.Tech.	Computer Science and Engineering	Indian Institute of Technology, Kanpur

Professional Experience

2002--	Professor, Computer Science, University of Florida, Gainesville, FL
1999-2002	Chief Scientist and CTO, Paramark Inc., Sunnyvale, CA
1995-99	Associate Professor, Computer Science, University of Florida, Gainesville, FL
1993-95	Associate Professor, Computer Science, Syracuse University, Syracuse, NY
1988-93	Assistant Professor, Computer Science, Syracuse University, Syracuse, NY

Professional Activities

1. Editorial Board, Journal of Parallel and Distributed Computing, 1994-present.
2. Member, IFIP Committee on System Modeling and Optimization, 2000-present.
3. Guest Editor, Journal of Parallel & Distributed Computing, Special Issue on High Performance Data Mining, 1999.
4. Member, MPI Standards Committee
5. Served as co-chair and vice-chair for many workshops on parallel computing and High Performance Data Mining.

Books and Book Chapters

1. K. G. Mehrotra, S. Ranka, and C. K. Mohan. Elements of Artificial Neural Networks, MIT Press, October 1996
2. S. Ranka and S. Sahni. Hypercube Algorithms with Applications to Image Processing and Pattern Recognition, Springer Verlag, June 1990.
3. M. Sreenivas, K. Alsabti and S. Ranka, Parallel Out of Core Decision Classifiers, Advances in Distributed And Parallel Knowledge Discovery (Editors: Kargupta and Chan), 2000, pp 319-338.
4. Z. Bozkus, A. Choudhary, G. Fox, T. Haupt, and S. Ranka. Compiling HPF for Distributed memory MIMD computers, Impact of Compilation Technology on Computer Architecture. David Lilja and Peter Bird, eds. Kluwer Academic, Norwell, MA, October 1993, pp. 191-217.
5. G. Fox and S. Ranka. Impact of Grand Challenge Applications on HPCC Software and Tools, System Software and Tools for High-Performance Computing Environments. Paul Messina and Thomas Sterling, eds. SIAM Press, April 1993, pp. 25-34.

Five Publications Related to the Proposed Research

1. K. Alsabti, S. Ranka and V. Singh. A One-Pass Algorithm for Accurately Estimating Quantiles for Disk-Resident Data In Proceedings of VLDB'97 Conference, 1997.
2. M. Kaddoura and S. Ranka. Runtime Support for Parallelization of Data-Parallel Applications on Adaptive and Nonuniform Environments, Journal of Parallel and Distributed Computing (Special Issue on Workstation Clusters and Network-based Computing), June 1997, pp. 163-168.
3. M. Kaddoura, S. Ranka and A. Wang. Array Decompositions for Non-uniform Computational Environment, Journal of Parallel and Distributed Computing, August 1996, pp. 91-105..
4. C. W. Ou and S. Ranka, Parallel Incremental Graph Partitioning, IEEE Transactions on Parallel and Distributed Systems, August 1997, pp. 884-896.
5. Z. Bozkus, A. Choudhary, G. Fox, T. Haupt, and S. Ranka. Compiling Fortran 90D/HPF for Distributed Memory MIMD Computers, Journal of Parallel and Distributed computing, April 1994, pp. 15-26.

Five other Publications

1. K. Alsabti, S. Ranka and V. Singh. An Efficient Space-Partitioning Based Algorithm for the K-Means Clustering, PAKDD 1999, pp. 355-359.
2. K. Alsabti, S. Ranka and V. Singh. An Efficient Parallel Algorithm for High Dimensional Similarity Join, Proc. of 1998 International Parallel Processing Symposium, pp. 556-560.
3. K. Alsabti, S. Ranka and V. Singh, CLOUDS: A Decision Tree Classifier for Large Datasets, Conference on Knowledge Discovery and Data Mining (KDD-98), August 1998..
4. S. Thomas, S. Bodagala, K. Alsabti and S. Ranka. An Efficient Algorithm for the Incremental Updation on Association Rules in Large Databases, Proceedings of the 3rd International Conference on Knowledge Discovery and Data Mining, 1997.
5. S. Su, S. Ranka and X. He, Performance Analysis of Parallel Query Processing Algorithms for Object-Oriented Databases. IEEE Transactions on Knowledge and Data Engineering, November 2000, pp. 979-997.

Synergistic Activities

1. Fortran 90D/HPF Compiler (1991-1993): Co-architected a compiler for Fortran90D, while at Syracuse University. Fortran 90D provided significant input to the High Performance Fortran Specification. Technology developed was transferred to PGI, a compiler developer for large scale parallel machines.
2. PILOT, an interactive marketing optimization software (1999-2002). Architected and developed real-time production software for optimizing marketing campaigns for Paramark Inc. The novel techniques allowed for on the fly learning and optimization of marketing campaigns without significant historical data. This made it applicable to a large number of direct marketers because their products and offers, media buys and customer responses change dramatically at a daily to weekly basis. PILOT had a 99.999% uptime and was architected to deploy seamlessly with existing email servers, commerce systems and Ad Servers. Most deployments took less than a few weeks. PILOT delivered 20-100% improvements in marketing spend for marketers such as Capitol One, Chase Manhattan Bank, Next Card, Ameritrade, Intel, Coke and Omnisky.

Collaborators. Stanley Su (Florida), Srinivas Aluru (Iowa State), Geoffrey Fox (Indiana), Ananth Grama (Purdue)

Graduate Students. Todd Heywood (IBM), Anand Rangachari (IBM), Yeh-Chin Chung (National Taiwan University), Jhy-Chun Wang (IBM), Harpal Maini (Startup), David Koester,(MITRE), Chao-Wei Ou (Startup), Ravi Shankar (Lucent), Maher Kaddoura (Architecture Technology Corporation), Seungjo Bae (ETRI (Korea)), Jang Sun Lee(ETRI (Korea)), Khaled Alsabti (King Fahd University (Saudi Arabia)), Ibraheem Al-furaih (National Labs, Saudi Arabia), Hankil Yoon (Oracle), Scott Winterstein (Private Consultant)

Advisor. Sartaj Sahni

E. __. MARK S. SCHMALZ

Research Interests

Image Processing, Pattern Recognition, Computer Vision, Error Analysis, Optical Systems Design, Artificial Intelligence, Medical Imaging

Selected Research Accomplishments

Researched and developed theory and algorithms for reconstruction of 3-D object from multiple views, with error analysis. Established theory and developed software for nonparametric error analysis of discrete algorithms. Developed theory and algorithms for dynamic scheduling of image and signal processing and pattern recognition algorithms, on networks of faulting embedded or reconfigurable processors. Developed image compression transforms with small, efficient kernels and local operations for visually acceptable 200:1 compression of underwater and video imagery. Developed image quality measures for image compression transforms. Developed techniques for significantly increasing computational efficiency by processing compressed imagery, without decompression.

Education

<i>Year</i>	<i>Degree</i>	<i>Institution</i>
1996	Ph.D., Computer Science	University of Florida
1991	M.S., Computer Science	University of Florida
1979	O.D., Optometry	Pacific University
1974	B.S., Applied Physics	Michigan Technological University

Previous and Current Positions

- Assistant Scientist, Department of Computer and Information Science and Engineering (CISE), University of Florida (1998-present)
- Visiting Assistant Professor, Department of CISE, University of Florida (1996-1998)
- Graduate Research/ Teaching Assistant, Department of CISE, University of Florida (1989-1996)
- Electrical Engineer, Harbor Branch Oceanographic Institution (1985-1988)
- Research Assistant, Harbor Branch Foundation, Department of Physical Oceanography (1982-1985)
- Private Practice of Optometry (1979-1984)

Selected Research Funding

Over \$2.6 million in research funding as PI or Co-PI in the past six years. Relevant grants include:

- Nonlinear Signal Processing for Mine Detection (Co-PI, P. Gader as PI), US Army, 2001, \$158,166
- BIOCAM – Biomimetic Camouflage for Adaptive Target Concealment (PI), Subcontract to Frontier Technology, Inc. on US Army SBIR Phase II Grant, 2001, \$183,520
- Flexible Production Networks Phase II: SEEK – Scalable Extraction of Enterprise Knowledge (Co-PI with J. Hammer as PI), NSF, 2001, \$638,925
- Instrumentation for Image/Signal Processing Research and Education (Co-PI with G.X. Ritter), DURIP-2000, \$99,263
- Mine Detection with Morphological Data Fusion (PI), Frontier Technology, Inc., 2000, \$22,010
- MLI – Errors Inherent in 3D Target Recognition from Multiple Airborne Images, Subcontract to RDL on AFOSR Contract #F49620-93-C-0063, 1998, \$24,619
- AIM – Adaptive Image Manager for Reconfigurable High-Performance Computing Systems (Co-PI with G.X. Ritter), DARPA, 1997, \$775,666
- Compressive Underwater Video Camera with Acoustic Uplink (Co-PI with G.X. Ritter), ONR, 1996, \$114,645
- Modelling, Image Simulation, and Analytical Support for ML(A) System (Co-PI with G.X. Ritter), US Navy Contract #N61331-94-M8684, NSWC, 1996, \$21,200
- Surface Mine Detection and Localization with Differential Multispectral Data Fusion, US Navy Contract #N61331-93-C0048 (Co-PI with G.X. Ritter), USMC/NSWC, 1993, \$377,153

Awards

- Best Paper Award (Medical Imaging), SCI-2002 World Multiconference (Orlando, FL), 2002
- AFOSR Summer Faculty Fellowship, AFRL/IFEC (Rome, NY), 1997
- Selected for NASA Summer Faculty Fellowship, Goddard Space Flight Center, 1997

Consulting

- Triton Alliance, Montreal PQ Canada – Underwater Imaging and Pattern Recognition

Professional Activities

- Reviewer for *Journal of Electronic Imaging*, *Journal of Mathematical Imaging and Vision*, *Optical Engineering*, *Pattern Analysis and Applications*, *VLSI Design Journal*
- Co-Chair, *Program on Mathematical Imaging* (SPIE Annual Meeting, San Diego CA, July 2003)
- Conference Chair, *Mathematics and Applications of Data Coding, Compression, and Encryption V* (SPIE, Seattle WA, July 2002).
- Conference Chair, *Mathematics and Applications of Data Coding, Compression, and Encryption III & IV* (SPIE, San Diego CA, August 2000 and August 2001).
- Conference Chair, *Mathematics and Applications of Data Coding, Compression, and Encryption I & II* (SPIE, Denver CO, August 1999), (SPIE, San Diego CA, August 1998).
- Session Chair at numerous conferences (*Florida Artificial Intelligence Research Society*, *SPIE Aerosense 1997*, *IEEE Southeastcon*, etc.)

List of Recent Collaborators

- Dr. Gerhard Ritter, Dr. Paul Gader, Dr. Joachim Hammer, University of Florida Department of CISE
- Dr. Frank Caimi, Harbor Branch Oceanographic Institution
- Mr. Gary Key, Frontier Technology, Inc.
- Dr. Christine Sapienza, Dr. Joseph Geunes, Dr. Raymond Issa, Dr. Sherman Bai, University of Florida
- Dr. Wen-Chen Hu, Auburn University

Five Publications Pertinent to Proposed Research

- Schmalz, M.S. and G.X. Ritter. “Detection and characterization of motion in video compression”, *Proceedings SPIE 4793*:254-266 (2002).
- Schmalz, M.S. “Space, time, error, and power optimization of image compression transform design”, in *Proceedings SPIE 4122* (2000).
- Schmalz, M.S. “Errors inherent in reconstruction of obscured targets for multilook imagery - Error analysis and simulation”, *Selected SPIE Papers on CD-ROM, Volume 6, Automated Target Recognition*, 12 pages (1999).
- Schmalz, M.S. “Reconstruction of submerged targets from multiple airborne images”, in *Proceedings SPIE 3761*:84-107 (1999).
- Schmalz, M.S. “Effects of ocean optical processes on reconstruction of submerged targets from multiple underwater images”, *Proceedings of the Ocean Optics XIV Conference, Kona Hawaii*, vol. 3, sec. 4, pp.1-6 (1998).

Five Other Recent Publications.

- Schmalz, M.S. “Recent advances in error analysis techniques and systems for astronomical image processing”, in *Proceedings of the 2002 AMOS Conference, Maui Hawaii*, 11 pages (in press).
- Schmalz, M.S. “Analysis of error in digital image processing algorithms using coincidence measures”, *Proceedings of the SCI-2002 World Multiconference on Systemics, Cybernetics, and Informatics 13*:457-462 (2002).
- Schmalz, M.S. “Errors inherent in reconstruction of obscured targets from multi-look imagery. 1. Background and theory”, *Proceedings SPIE 3371*:488-499 (1998).
- Schmalz, M.S. “Errors inherent in reconstruction of obscured targets from multi-look imagery. 2. Error analysis and simulation results”, *Proceedings SPIE 3371*:92-104 (1998).
- Schmalz, M.S. “Stereoscopic imaging through the sea surface – Theory and error analysis”, *Proceedings SPIE 3392*:124-135 (1998).

Summary of Publications

More than 115 book chapters, or papers in conference proceedings and refereed journals

a. Professional Preparation

B.S. in Computer Science and Applied Mathematics; University of Rochester, May 1988
 M.S. in Computer Science; University of Southern California, December 1991
 Ph.D. in Computer Science; University of Southern California, May 1994
 Thesis: Resolving Semantic Heterogeneity in a Federation of Autonomous, Heterogeneous Database Systems

b. Academic Appointments

Current

Assistant Professor, Department of Computer and Information Science and Engineering, University of Florida, since August 1997
Visiting Professor, Department of Electrical Engineering, The University of Hong Kong, since Sept. 2001
Assistant director, E-Commerce Research Center, University of Florida, since January 2000
Areas of research: heterogeneous and semistructured information systems; federated database systems; data transformations; data warehousing; database system modeling; knowledge engineering; bioinformatics
Areas of teaching: data structures and algorithms; introduction to database management systems; data warehousing and decision support; information integration and XML

Previous

Department of Computer Science, Stanford University, 1994–1997
 Research scientist, database group
Computer Science Department, USC, 1989–1994
 Lecturer, Summer 1993 and Spring 1994
 Head teaching assistant for advanced database course, 1991–1993; teaching assistant, 1989–1990
 Research assistant, 1990–1991

c. List of Publications (available online at <http://www.cise.ufl.edu/~jhammer/publications.html>)

Related to Project

1. J. Hammer and C. Pluempitiwiriyawej. "Element Matching across XML Sources using a Multi-strategy Clustering Technique." To appear in *Data and Knowledge Engineering (DKE)*. Elsevier Science. 2003.
2. W. O'Brien, R. R. Issa, J. Hammer, M. S. Schmalz, J. Geunes, and S. X. Bai, "SEEK: Accomplishing Enterprise Information Integration Across Heterogeneous Sources." *Journal of Information Technology in Construction*, 7, pp. 101-124, 2002.
3. A. Helal, J. Hammer, A. Khushraj, and J. Zhang, "A Three-tier Architecture for Ubiquitous Data Access." In *Proceedings of the First ACS/IEEE International Conference on Computer Systems and Applications*, Beirut, Lebanon, pages 177-180, July 2001.
4. J. Hammer and D. McLeod, "On the Resolution of Representational Diversity in Multidatabase Systems," in *Management of Heterogeneous and Autonomous Database Systems*, A. K. Elmagamid, M. Rusinkiewicz, and A. P. Sheth, Eds., Morgan Kaufmann Publishers, San Francisco, CA, 1998.
5. D. Fang, J. Hammer, and D. McLeod. "A Mechanism and Experimental System for Function-Based Sharing in Federated Databases." In D.K. Hsiao, E.J. Neuhold, and R. Sacks-Davis, editors, *Interoperable Database Systems (DS-5) (A-25)*, pages 239-253. Elsevier Science Publisher B.V. (North-Holland), 1993.

Significant Other Publications

1. J. Hammer and L. Fu*. "CubiST⁺⁺: Evaluating Ad-Hoc CUBE Queries Using Statistics Tree." *Distributed and Parallel Databases*. To appear in Kluwer Academic Publishers. 2003.
2. M. Lanham, A. Kang, J. Hammer, A. Helal and J. Wilson. "Format-Independent Change Detection and Propagation in Support of Mobile Computing." In *Proceedings of the SBBD2002 - XVII Brazilian Symposium on Databases*, Gramado, Brazil, pages 27-41, 14-17 October 2002.
3. S. Su, C. Huang, J. Hammer, Y. Huang, H. Li, L. Wang, Y. Liu, C. Pluempitiwiriyawej, and H. Lam, "An Internet-Based Negotiation Server for E-Commerce," *VLDB Journal*, 10:2, 2001.
4. M. Lee and J. Hammer, "Speeding Up Warehouse Physical Design Using A Randomized Algorithm," *International Journal of Cooperative Information Systems (IJCIS) - Special Issue on Design and Management of Data Warehouses*, 10:3, pp. 327-354, 2001.
5. Y. Zhuge, H. Garcia-Molina, J. Hammer, and J. Widom, "View Maintenance in a Warehousing Environment," in *Materialized Views*, A. Gupta and I. S. Mumick, Eds. MIT Press, 1999, pp. 616 pp.

d. Synergistic Activities***Relevant Research Experience***

PI for NSF-sponsored project to develop theories and algorithms for (semi)automatic extraction of enterprise knowledge from legacy systems (SEEK); goal is to enable information sharing among heterogeneous enterprises in complex supply chains in order to enhance coordination; research is carried out jointly with domain experts in building construction (duration: 2001–2004)

PI for state-funded project to conduct feasibility study for data warehouse containing information on transportation infrastructure including traffic conditions on all major highways in Florida (duration: 2002)

Co-PI for NSF-sponsored project to develop algorithms to support ubiquitous computing irrespective of the device that is used to access and manipulate the data (duration: 2001–2003)

PI for government contract to develop algorithms for managing geo-spatial data from multiple terrain sources in relational data warehouse; focus is on middleware to produce a variety of query output formats (e.g., XML, GML, VRML, X3D) at run-time in support of real-time terrain visualization (duration: 2000–2002)

Co-PI for Univ. of Florida component of a NIST/ATP sponsored project for design and development of an integrated collaborative manufacturing systems; focus is on support for knowledge sharing and negotiation among the collaborating components of a supply chain system (duration: 1997–2000)

Project manager for Stanford's DARPA-sponsored *TSIMMIS Project*, overseeing and participating in the research, design, and implementation of methodologies for providing integrated access to autonomous, heterogeneous information sources (duration: 1994–1997)

Project manager for the DARPA-sponsored *WareHousIng Project* at Stanford (WHIPS); involved in the design of the initial architecture as well as the development of two new algorithms for incrementally maintaining views in data warehouses (duration: 1994–1995)

Memberships and Committees

Member executive committee (2001-2004):

Elected Secretary and Treasurer for the Special Interest Group Management Of Data (SIGMOD) of ACM

Member editorial staff:

Data and Knowledge Engineering (DKE), International Journal of Computer and Information Systems (IJCIS), Gallery of Courses Taught with Technology

Program committee member (list of most recent):

SIGMOD, VLDB, Data Engineering, Cooperative Information Systems, Data Warehousing and OLAP, Design and Management of Data Warehouses, Cooperative Information Agents

Conference and workshop organizer:

2004 ACM Conference on Inform. and Knowledge Management (CIKM) – program committee co-chair

2001 ACM International Workshop on Data Warehousing and OLAP (DOLAP) – program committee chair

1999 ACM SIGMOD International Conference on Management of Data – registration chair

e. Collaborators and Other Affiliations***Recent Collaborators on Sponsored Projects***

Ron Wakefield – *Virginia Tech University*

Sherman Bai, Joseph Guenes, Ray Issa, William O'Brien, Herman Lam, Stanley Su – *University of Florida*

Graduate and Postdoctoral Advisors

Ph.D. Advisor: Prof. Dennis McLeod – Computer Science Department, University of Southern California

Postdoctoral Advisor: Prof. Hector Garcia-Molina – Computer Science Department, Stanford University

Thesis Advisor

Current Students at the Univ. of Florida:

Ph.D.: Lucian Badea, Mingxi Wu

M.S. (with thesis): Prakash Bulusu, Nikhil Haldavnekar, Sangeetha Shekar, Gargi Sur

Graduated Ph.D Students:

Lixin Fu (University of North Carolina, Greensboro); Charnyote Pluempitwiriyawej (University of

Thailand, Bangkok); Hankil Yoon (co-advised with Sanjay Ranka); Chunbo Huang (co-advised with Stan

Su – BEA Systems, San Jose, CA)

SUMMARY PROPOSAL BUDGET YEAR 1

ORGANIZATION University of Florida				FOR NSF USE ONLY			
				PROPOSAL NO.	DURATION (months)		
PRINCIPAL INVESTIGATOR / PROJECT DIRECTOR Michael P Frank				AWARD NO.	Proposed	Granted	
				A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title, A.7. show number in brackets)			
				CAL	ACAD	SUMR	
1.	Michael P Frank - Assistant Professor			0.00	0.00	1.00	\$ 8,044
2.	Joachim Hammer - Assistant Professor			0.00	0.00	1.00	8,813
3.	Abdelsalam A Helal - Associate Professor			0.00	0.00	1.00	10,210
4.	Jorg Peters - Associate Professor			0.00	0.00	1.00	10,135
5.	Sanjay Ranka - Professor			0.00	0.00	1.00	13,733
6.	(1) OTHERS (LIST INDIVIDUALLY ON BUDGET JUSTIFICATION PAGE)			0.00	0.00	1.00	6,132
7.	(6) TOTAL SENIOR PERSONNEL (1 - 6)			0.00	0.00	6.00	57,067
B. OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)							
1.	(0) POST DOCTORAL ASSOCIATES			0.00	0.00	0.00	0
2.	(0) OTHER PROFESSIONALS (TECHNICIAN, PROGRAMMER, ETC.)			0.00	0.00	0.00	0
3.	(6) GRADUATE STUDENTS						93,960
4.	(0) UNDERGRADUATE STUDENTS						0
5.	(0) SECRETARIAL - CLERICAL (IF CHARGED DIRECTLY)						0
6.	(0) OTHER						0
TOTAL SALARIES AND WAGES (A + B)							151,027
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)							13,244
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A + B + C)							164,271
D. EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM EXCEEDING \$5,000.) See Budget Justification						\$ 20,000	
TOTAL EQUIPMENT							20,000
E. TRAVEL 1. DOMESTIC (INCL. CANADA, MEXICO AND U.S. POSSESSIONS)							7,000
2. FOREIGN							0
F. PARTICIPANT SUPPORT COSTS							
1.	STIPENDS	\$	0				
2.	TRAVEL		0				
3.	SUBSISTENCE		0				
4.	OTHER		0				
TOTAL NUMBER OF PARTICIPANTS (0)				TOTAL PARTICIPANT COSTS			0
G. OTHER DIRECT COSTS							
1. MATERIALS AND SUPPLIES							0
2. PUBLICATION COSTS/DOCUMENTATION/DISSEMINATION							0
3. CONSULTANT SERVICES							0
4. COMPUTER SERVICES							5,250
5. SUBAWARDS							0
6. OTHER							40,500
TOTAL OTHER DIRECT COSTS							45,750
H. TOTAL DIRECT COSTS (A THROUGH G)							237,021
I. INDIRECT COSTS (F&A)(SPECIFY RATE AND BASE) MTDC (Rate: 45.5000, Base: 176521)							
TOTAL INDIRECT COSTS (F&A)							80,317
J. TOTAL DIRECT AND INDIRECT COSTS (H + I)							317,338
K. RESIDUAL FUNDS (IF FOR FURTHER SUPPORT OF CURRENT PROJECTS SEE GPG II.C.6.j.)							0
L. AMOUNT OF THIS REQUEST (J) OR (J MINUS K)							\$ 317,338
M. COST SHARING PROPOSED LEVEL \$ 0				AGREED LEVEL IF DIFFERENT \$			
PI/PD NAME Michael P Frank				FOR NSF USE ONLY			
ORG. REP. NAME*				INDIRECT COST RATE VERIFICATION			
		Date Checked	Date Of Rate Sheet	Initials - ORG			

SUMMARY PROPOSAL BUDGET COMMENTS - Year 1

Other Senior Personnel

Name - Title	Cal	Acad	Sumr	Funds Requested
-----	---	-----	-----	-----
Schmalz, Mark S - Assistant Scientist	0.00	0.00	1.00	6132

SUMMARY PROPOSAL BUDGET YEAR 2

ORGANIZATION University of Florida				FOR NSF USE ONLY		
				PROPOSAL NO.	DURATION (months)	
PRINCIPAL INVESTIGATOR / PROJECT DIRECTOR Michael P Frank				AWARD NO.	Proposed	Granted
					NSF Funded Person-months	
A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title, A.7. show number in brackets)				CAL	ACAD	SUMR
1. Michael P Frank - Assistant Professor				0.00	0.00	1.00
2. Joachim Hammer - Assistant Professor				0.00	0.00	1.00
3. Abdelsalam A Helal - Associate Professor				0.00	0.00	1.00
4. Jorg Peters - Associate Professor				0.00	0.00	1.00
5. Sanjay Ranka - Professor				0.00	0.00	1.00
6. (1) OTHERS (LIST INDIVIDUALLY ON BUDGET JUSTIFICATION PAGE)				0.00	0.00	1.00
7. (6) TOTAL SENIOR PERSONNEL (1 - 6)				0.00	0.00	6.00
B. OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)						
1. (0) POST DOCTORAL ASSOCIATES				0.00	0.00	0.00
2. (0) OTHER PROFESSIONALS (TECHNICIAN, PROGRAMMER, ETC.)				0.00	0.00	0.00
3. (6) GRADUATE STUDENTS						
4. (0) UNDERGRADUATE STUDENTS						
5. (0) SECRETARIAL - CLERICAL (IF CHARGED DIRECTLY)						
6. (0) OTHER						
TOTAL SALARIES AND WAGES (A + B)						96,780
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)						0
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A + B + C)						0
D. EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM EXCEEDING \$5,000.) See Budget Justification					\$ 20,000	
TOTAL EQUIPMENT						20,000
E. TRAVEL 1. DOMESTIC (INCL. CANADA, MEXICO AND U.S. POSSESSIONS)						7,000
2. FOREIGN						0
F. PARTICIPANT SUPPORT COSTS						
1. STIPENDS \$ _____				0		
2. TRAVEL _____				0		
3. SUBSISTENCE _____				0		
4. OTHER _____				0		
TOTAL NUMBER OF PARTICIPANTS (0) TOTAL PARTICIPANT COSTS						0
G. OTHER DIRECT COSTS						
1. MATERIALS AND SUPPLIES						0
2. PUBLICATION COSTS/DOCUMENTATION/DISSEMINATION						0
3. CONSULTANT SERVICES						0
4. COMPUTER SERVICES						5,250
5. SUBAWARDS						0
6. OTHER						43,488
TOTAL OTHER DIRECT COSTS						48,738
H. TOTAL DIRECT COSTS (A THROUGH G)						244,879
I. INDIRECT COSTS (F&A)(SPECIFY RATE AND BASE) MTDC (Rate: 45.5000, Base: 181391)						
TOTAL INDIRECT COSTS (F&A)						82,533
J. TOTAL DIRECT AND INDIRECT COSTS (H + I)						327,412
K. RESIDUAL FUNDS (IF FOR FURTHER SUPPORT OF CURRENT PROJECTS SEE GPG II.C.6.j.)						0
L. AMOUNT OF THIS REQUEST (J) OR (J MINUS K)						\$ 327,412 \$
M. COST SHARING PROPOSED LEVEL \$ 0				AGREED LEVEL IF DIFFERENT \$		
PI/PD NAME Michael P Frank				FOR NSF USE ONLY		
ORG. REP. NAME*				INDIRECT COST RATE VERIFICATION		
				Date Checked	Date Of Rate Sheet	Initials - ORG

SUMMARY PROPOSAL BUDGET COMMENTS - Year 2

Other Senior Personnel

Name - Title	Cal	Acad	Sumr	Funds Requested
Schmalz, Mark S - Assistant Scientist	0.00	0.00	1.00	6316

SUMMARY PROPOSAL BUDGET YEAR 3

ORGANIZATION University of Florida				FOR NSF USE ONLY		
				PROPOSAL NO.	DURATION (months)	
PRINCIPAL INVESTIGATOR / PROJECT DIRECTOR Michael P Frank				AWARD NO.	Proposed	Granted
					NSF Funded Person-months	
A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title, A.7. show number in brackets)				CAL	ACAD	SUMR
1. Michael P Frank - Assistant Professor				0.00	0.00	1.00
2. Joachim Hammer - Assistant Professor				0.00	0.00	1.00
3. Abdelsalam A Helal - Associate Professor				0.00	0.00	1.00
4. Jorg Peters - Associate Professor				0.00	0.00	1.00
5. Sanjay Ranka - Professor				0.00	0.00	1.00
6. (1) OTHERS (LIST INDIVIDUALLY ON BUDGET JUSTIFICATION PAGE)				0.00	0.00	1.00
7. (6) TOTAL SENIOR PERSONNEL (1 - 6)				0.00	0.00	6.00
B. OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)						
1. (0) POST DOCTORAL ASSOCIATES				0.00	0.00	0.00
2. (0) OTHER PROFESSIONALS (TECHNICIAN, PROGRAMMER, ETC.)				0.00	0.00	0.00
3. (6) GRADUATE STUDENTS						99,684
4. (0) UNDERGRADUATE STUDENTS						0
5. (0) SECRETARIAL - CLERICAL (IF CHARGED DIRECTLY)						0
6. (0) OTHER						0
TOTAL SALARIES AND WAGES (A + B)						160,224
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)						13,933
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A + B + C)						174,157
D. EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM EXCEEDING \$5,000.) See Budget Justification				\$	20,000	
TOTAL EQUIPMENT						20,000
E. TRAVEL 1. DOMESTIC (INCL. CANADA, MEXICO AND U.S. POSSESSIONS)						7,000
2. FOREIGN						0
F. PARTICIPANT SUPPORT COSTS						
1. STIPENDS \$ _____				0		
2. TRAVEL _____				0		
3. SUBSISTENCE _____				0		
4. OTHER _____				0		
TOTAL NUMBER OF PARTICIPANTS (0) TOTAL PARTICIPANT COSTS						0
G. OTHER DIRECT COSTS						
1. MATERIALS AND SUPPLIES						0
2. PUBLICATION COSTS/DOCUMENTATION/DISSEMINATION						0
3. CONSULTANT SERVICES						0
4. COMPUTER SERVICES						5,250
5. SUBAWARDS						0
6. OTHER						46,512
TOTAL OTHER DIRECT COSTS						51,762
H. TOTAL DIRECT COSTS (A THROUGH G)						252,919
I. INDIRECT COSTS (F&A)(SPECIFY RATE AND BASE) MTDC (Rate: 45.5000, Base: 186407)						
TOTAL INDIRECT COSTS (F&A)						84,815
J. TOTAL DIRECT AND INDIRECT COSTS (H + I)						337,734
K. RESIDUAL FUNDS (IF FOR FURTHER SUPPORT OF CURRENT PROJECTS SEE GPG II.C.6.j.)						0
L. AMOUNT OF THIS REQUEST (J) OR (J MINUS K)				\$	337,734	\$
M. COST SHARING PROPOSED LEVEL \$ 0				AGREED LEVEL IF DIFFERENT \$		
PI/PD NAME Michael P Frank				FOR NSF USE ONLY		
ORG. REP. NAME*				INDIRECT COST RATE VERIFICATION		
				Date Checked	Date Of Rate Sheet	Initials - ORG

SUMMARY PROPOSAL BUDGET COMMENTS - Year 3

Other Senior Personnel

Name - Title	Cal	Acad	Sumr	Funds Requested
Schmalz, Mark S - Assistant Scientist	0.00	0.00	1.00	6505

SUMMARY PROPOSAL BUDGET Cumulative

ORGANIZATION University of Florida				FOR NSF USE ONLY		
				PROPOSAL NO.	DURATION (months)	
PRINCIPAL INVESTIGATOR / PROJECT DIRECTOR Michael P Frank				AWARD NO.	Proposed	Granted
					NSF Funded Person-months	
A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title, A.7. show number in brackets)				CAL	ACAD	SUMR
1. Michael P Frank - Assistant Professor				0.00	0.00	3.00
2. Joachim Hammer - Assistant Professor				0.00	0.00	3.00
3. Abdelsalam A Helal - Associate Professor				0.00	0.00	3.00
4. Jorg Peters - Associate Professor				0.00	0.00	3.00
5. Sanjay Ranka - Professor				0.00	0.00	3.00
6. (1) OTHERS (LIST INDIVIDUALLY ON BUDGET JUSTIFICATION PAGE)				0.00	0.00	3.00
7. (6) TOTAL SENIOR PERSONNEL (1 - 6)				0.00	0.00	18.00
B. OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)						
1. (0) POST DOCTORAL ASSOCIATES				0.00	0.00	0.00
2. (0) OTHER PROFESSIONALS (TECHNICIAN, PROGRAMMER, ETC.)				0.00	0.00	0.00
3. (18) GRADUATE STUDENTS						290,424
4. (0) UNDERGRADUATE STUDENTS						0
5. (0) SECRETARIAL - CLERICAL (IF CHARGED DIRECTLY)						0
6. (0) OTHER						0
TOTAL SALARIES AND WAGES (A + B)						466,809
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)						40,760
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A + B + C)						507,569
D. EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM EXCEEDING \$5,000.)						
\$ 60,000						
TOTAL EQUIPMENT						60,000
E. TRAVEL 1. DOMESTIC (INCL. CANADA, MEXICO AND U.S. POSSESSIONS)						21,000
2. FOREIGN						0
F. PARTICIPANT SUPPORT COSTS						
1. STIPENDS \$ _____				0		
2. TRAVEL _____				0		
3. SUBSISTENCE _____				0		
4. OTHER _____				0		
TOTAL NUMBER OF PARTICIPANTS (0) TOTAL PARTICIPANT COSTS						0
G. OTHER DIRECT COSTS						
1. MATERIALS AND SUPPLIES						0
2. PUBLICATION COSTS/DOCUMENTATION/DISSEMINATION						0
3. CONSULTANT SERVICES						0
4. COMPUTER SERVICES						15,750
5. SUBAWARDS						0
6. OTHER						130,500
TOTAL OTHER DIRECT COSTS						146,250
H. TOTAL DIRECT COSTS (A THROUGH G)						734,819
I. INDIRECT COSTS (F&A)(SPECIFY RATE AND BASE)						
TOTAL INDIRECT COSTS (F&A)						247,665
J. TOTAL DIRECT AND INDIRECT COSTS (H + I)						982,484
K. RESIDUAL FUNDS (IF FOR FURTHER SUPPORT OF CURRENT PROJECTS SEE GPG II.C.6.j.)						0
L. AMOUNT OF THIS REQUEST (J) OR (J MINUS K)						\$ 982,484 \$
M. COST SHARING PROPOSED LEVEL \$ 0				AGREED LEVEL IF DIFFERENT \$		
PI/PD NAME Michael P Frank				FOR NSF USE ONLY		
ORG. REP. NAME*				INDIRECT COST RATE VERIFICATION		
				Date Checked	Date Of Rate Sheet	Initials - ORG

C *ELECTRONIC SIGNATURES REQUIRED FOR REVISED BUDGET

Budget Justification Page

For each year of the award period, the PI and each co-PI are allocated the following:

- (1) One month's worth of summer support, in order to compensate for their share of the work on the project, although it is unlikely that all the faculty work on the project will actually be concentrated within that month. (Note that one month of support corresponds to an annual average FTE of ~8.3%.)
- (2) One year-round, 50% FTE graduate research assistant to help the faculty member carry out their share of the work.
- (3) Enough travel money (\$1,167) that the faculty member can travel to an estimated 1 conference in his field to present results of this research and disseminate news about the OCEAN project to the broader research community.

In addition, the budget also includes enough equipment money (\$20,000 per year) for the purchase each year of at least one new mid-range server machine on which to deploy high-performance OCEAN nodes and services for testing purposes. (Note that the equipment purchases are spread out over the three years, so that in each year's work we can take advantage of the most up-to-date technology.)

The "other" direct cost category is the tuition for the graduate students.

Current and Pending Support

(See GPG Section II.D.8 for guidance on information to include on this form.)

The following information should be provided for each investigator and other senior personnel. Failure to provide this information may delay consideration of this proposal.

Investigator: Michael Frank	Other agencies (including NSF) to which this proposal has been/will be submitted.
Support: <input type="checkbox"/> Current <input checked="" type="checkbox"/> Pending <input type="checkbox"/> Submission Planned in Near Future <input type="checkbox"/> *Transfer of Support Project/Proposal Title: OCEAN - A Peer-to-Peer Market for Allocating Grid Computing Resources	
Source of Support: NSF Total Award Amount: \$ 982,484 Total Award Period Covered: 01/01/04 - 12/31/06 Location of Project: University of Florida Person-Months Per Year Committed to the Project. Cal: 0.00 Acad: 0.00 Sumr: 1.00	
Support: <input type="checkbox"/> Current <input type="checkbox"/> Pending <input type="checkbox"/> Submission Planned in Near Future <input type="checkbox"/> *Transfer of Support Project/Proposal Title:	
Source of Support: Total Award Amount: \$ Total Award Period Covered: Location of Project: Person-Months Per Year Committed to the Project. Cal: Acad: Sumr:	
Support: <input type="checkbox"/> Current <input type="checkbox"/> Pending <input type="checkbox"/> Submission Planned in Near Future <input type="checkbox"/> *Transfer of Support Project/Proposal Title:	
Source of Support: Total Award Amount: \$ Total Award Period Covered: Location of Project: Person-Months Per Year Committed to the Project. Cal: Acad: Sumr:	
Support: <input type="checkbox"/> Current <input type="checkbox"/> Pending <input type="checkbox"/> Submission Planned in Near Future <input type="checkbox"/> *Transfer of Support Project/Proposal Title:	
Source of Support: Total Award Amount: \$ Total Award Period Covered: Location of Project: Person-Months Per Year Committed to the Project. Cal: Acad: Sumr:	
Support: <input type="checkbox"/> Current <input type="checkbox"/> Pending <input type="checkbox"/> Submission Planned in Near Future <input type="checkbox"/> *Transfer of Support Project/Proposal Title:	
Source of Support: Total Award Amount: \$ Total Award Period Covered: Location of Project: Person-Months Per Year Committed to the Project. Cal: Acad: Summ:	

*If this project has previously been funded by another agency, please list and furnish information for immediately preceding funding period.

Current and Pending Support

(See GPG Section II.D.8 for guidance on information to include on this form.)

The following information should be provided for each investigator and other senior personnel. Failure to provide this information may delay consideration of this proposal.

Investigator: Abdelsalam (Sumi) Helal	Other agencies (including NSF) to which this proposal has
---------------------------------------	---

Support: Current Pending Submission Planned in Near Future *Transfer of Support

Project/Proposal Title:
Rehabilitation Engineering Research Center on Tech. for Successful Aging

Source of Support: US Department of Education
 Total Award Amount: \$4,500,000 Total Award Period Covered: 10/01/01-09/30/06
 Location of Project: Univ. of Florida
 Person-Months Per Year Committed to the Cal: 5.2 Acad: Sumr:

Support: Current Pending Submission Planned in Near Future *Transfer of Support

Project/Proposal Title:
Adaptive Synchronization Framework Supporting Device-independent Mobile Computing

Source of Support: National Science Foundation
 Total Award Amount: \$150,000 Total Award Period Covered 07/09/01-06/09/03
 Location of Project: University of Florida
 Person-Months Per Year Committed to the Cal:0 Acad: 0 Sumr: 1.0

Support: Current Pending Submission Planned in Near Future *Transfer of Support

Project/Proposal Title:
Engineering Smart Phones for the Elderly

Source of Support: Motorola Corp.
 Total Award Amount: \$85,000 Total Award Period Covered: 07/01/03-06/01/04
 Location of Project: University of Florida
 Person-Months Per Year Committed to the Cal: Acad: Sumr: 1.0

Support: Current Pending Submission Planned in Near Future *Transfer of Support

Project/Proposal Title:
Mobile Computing Smart Phone Projects

Source of Support: Motorola Corp.
 Total Award Amount: \$20,046 Total Award Period Covered: 10/07/02-08/31/03
 Location of Project: University of Florida
 Person-Months Per Year Committed to the Cal: 1 Acad: Sumr:

Support: Current Pending Submission Planned in Near Future *Transfer of Support

Project/Proposal Title:
Low LDL Monitoring Program

Source of Support: The Department of Veteran Affairs
 Total Award Amount: \$700,000 Total Award Period Covered: 07/20/02-07/19/05
 Location of Project: University of Florida
 Person-Months Per Year Committed to the Cal: 1 Acad: Sumr:

*If this project has previously been funded by another agency, please list and furnish information for immediately preceding funding period.

Current and Pending Support

(See GPG Section II.D.8 for guidance on information to include on this form.)

The following information should be provided for each investigator and other senior personnel. Failure to provide this information may delay consideration of this proposal.				
Investigator: Abdelsalam (Sumi) Helal		Other agencies (including NSF) to which this proposal has		
Support:	<input checked="" type="checkbox"/> Current	<input type="checkbox"/> Pending	<input type="checkbox"/> Submission Planned in Near Future	<input type="checkbox"/> *Transfer of Support
Project/Proposal Title: UPnP Service Discovery in Ad-hoc Networks				
Source of Support: Microsoft Corpration				
Total Award Amount: \$70,000		Total Award Period Covered: 04/06/01-03/05/03		
Location of Project: University of Florida				
Person-Months Per Year Committed to the		Cal:	Acad:	Sumr:
Support:	<input checked="" type="checkbox"/> Current	<input type="checkbox"/> Pending	<input type="checkbox"/> Submission Planned in Near Future	<input type="checkbox"/> *Transfer of Support
Project/Proposal Title: Interdisciplinary Rehabilitation Research Post-Doctoral Program				
Source of Support: National Institute on Disability and Rehabilitation Research				
Total Award Amount: \$750,000		Total Award Period Covered: July 1, 2002 to June 30, 2007		
Location of Project: Univ. of Florida				
Person-Months Per Year Committed to the		Cal:.05	Acad:	Sumr:

Current and Pending Support

(See GPG Section II.D.8 for guidance on information to include on this form.)

The following information should be provided for each investigator and other senior personnel. Failure to provide this information may delay consideration of this proposal.

Investigator: Abdelsalam Helal	Other agencies (including NSF) to which this proposal has been/will be submitted.
Support: <input type="checkbox"/> Current <input checked="" type="checkbox"/> Pending <input type="checkbox"/> Submission Planned in Near Future <input type="checkbox"/> *Transfer of Support Project/Proposal Title: OCEAN - A Peer-to-Peer Market for Allocating Grid Computing Resources	
Source of Support: NSF Total Award Amount: \$ 982,484 Total Award Period Covered: 01/01/04 - 12/31/06 Location of Project: University of Florida Person-Months Per Year Committed to the Project. Cal: 0.00 Acad: 0.00 Sumr: 1.00	
Support: <input type="checkbox"/> Current <input type="checkbox"/> Pending <input type="checkbox"/> Submission Planned in Near Future <input type="checkbox"/> *Transfer of Support Project/Proposal Title:	
Source of Support: Total Award Amount: \$ Total Award Period Covered: Location of Project: Person-Months Per Year Committed to the Project. Cal: Acad: Sumr:	
Support: <input type="checkbox"/> Current <input type="checkbox"/> Pending <input type="checkbox"/> Submission Planned in Near Future <input type="checkbox"/> *Transfer of Support Project/Proposal Title:	
Source of Support: Total Award Amount: \$ Total Award Period Covered: Location of Project: Person-Months Per Year Committed to the Project. Cal: Acad: Sumr:	
Support: <input type="checkbox"/> Current <input type="checkbox"/> Pending <input type="checkbox"/> Submission Planned in Near Future <input type="checkbox"/> *Transfer of Support Project/Proposal Title:	
Source of Support: Total Award Amount: \$ Total Award Period Covered: Location of Project: Person-Months Per Year Committed to the Project. Cal: Acad: Sumr:	
Support: <input type="checkbox"/> Current <input type="checkbox"/> Pending <input type="checkbox"/> Submission Planned in Near Future <input type="checkbox"/> *Transfer of Support Project/Proposal Title:	
Source of Support: Total Award Amount: \$ Total Award Period Covered: Location of Project: Person-Months Per Year Committed to the Project. Cal: Acad: Summ:	

*If this project has previously been funded by another agency, please list and furnish information for immediately preceding funding period.

Current and Pending Support

(See GPG Section II.D.8 for guidance on information to include on this form.)

The following information should be provided for each investigator and other senior personnel. Failure to provide this information may delay consideration of this proposal.

Investigator: Mark Schmalz	Other agencies (including NSF) to which this proposal has been/will be submitted.
----------------------------	---

Support: <input checked="" type="checkbox"/> Current <input type="checkbox"/> Pending <input type="checkbox"/> Submission Planned in Near Future <input type="checkbox"/> *Transfer of Support Project/Proposal Title: BIOCAM - Biomimetic Camouflage for Adaptive Target Concealment	Source of Support: USubcontract to Frontier Technology, Inc., Army SBIR Phase II Total Award Amount: \$183,520 Total Award Period Covered: 10/01/01-09/30/03 Location of Project: Univ. of Florida Person-Months Per Year Committed to the Project. Cal: Acad: 3.0 Sumr:
---	---

Support: <input checked="" type="checkbox"/> Current <input type="checkbox"/> Pending <input type="checkbox"/> Submission Planned in Near Future <input type="checkbox"/> *Transfer of Support Project/Proposal Title: Flexible Production Networks Phase II: SEEK - Scalable Extraction of Enterprise Knowledge Source of Support: NSF Total Award Amount: \$638,925 Total Award Period Covered: 10/01/01 through 09/30/04 Location of Project: Univ. of Florida Person-Months Per Year Committed to the Project. Cal: Acad: Sumr: 1.0	
--	--

Support: <input checked="" type="checkbox"/> Current <input type="checkbox"/> Pending <input type="checkbox"/> Submission Planned in Near Future <input type="checkbox"/> *Transfer of Support Project/Proposal Title: Spinning Projectile Imaging and Target Tracking System Source of Support: Subcontract to Frontier Tech, Inc. on Army SBIR 2002.2 Total Award Amount: \$20,900 Total Award Period Covered: 12/20/02 through 05/08/03 Location of Project: Univ. of Florida Person-Months Per Year Committed to the Project. Cal: Acad: 1.0 Sumr:	
--	--

Support: <input type="checkbox"/> Current <input checked="" type="checkbox"/> Pending <input type="checkbox"/> Submission Planned in Near Future <input type="checkbox"/> *Transfer of Support Project/Proposal Title: DYVID - Dynamic Video Imaging and Detection System Source of Support: Subcontract to Frontier Tech, Inc. on DARPA SBIR 2003.1 Total Award Amount: \$24,946 Total Award Period Covered: Six Months Location of Project: University of Florida Person-Months Per Year Committed to the Project. Cal: Acad: 1.2 Sumr:	
---	--

Support: <input type="checkbox"/> Current <input checked="" type="checkbox"/> Pending <input type="checkbox"/> Submission Planned in Near Future <input type="checkbox"/> *Transfer of Support Project/Proposal Title: DFPAT - Data Fusion and Pattern Recognition Source of Support: Subcontract to Frontier Tech, Inc. on Army/MDA SBIR 2003.1 Total Award Amount: \$20,900 Total Award Period Covered: Six Months Location of Project: University of Florida Person-Months Per Year Committed to the Project. Cal: Acad: 1.0 Sumr:	
---	--

*If this project has previously been funded by another agency, please list and furnish information for immediately preceding funding period.

Current and Pending Support

(See GPG Section II.D.8 for guidance on information to include on this form.)

The following information should be provided for each investigator and other senior personnel. Failure to provide this information may delay consideration of this proposal.

Investigator: Mark Schmalz	Other agencies (including NSF) to which this proposal has been/will be submitted.
----------------------------	---

Support: Current Pending Submission Planned in Near Future *Transfer of Support
 Project/Proposal Title:
 FEATS - Feature Extraction, Analysis and Tracking via Similarity
 Source of Support: Subcontract to Frontier Technology, Inc., Air Force STTR 2003
 Total Award Amount: \$39,845 Total Award Period Covered: 08-28-03 through 01-01-04
 Location of Project: Univ. of Florida
 Person-Months Per Year Committed to the Project. Cal: Acad: 1.2 Sumr:

Support: Current Pending Submission Planned in Near Future *Transfer of Support
 Project/Proposal Title:
 RTAD - Riverine Traffic Analysis and Detection
 Source of Support: Subcontract to Justak R&D Corp., Army STTR 2003
 Total Award Amount: \$30,427 Total Award Period Covered: 08-28-03 through 01-01-04
 Location of Project: Univ. of Florida
 Person-Months Per Year Committed to the Project. Cal: Acad: 1.3 Sumr: 1.0

Support: Current Pending Submission Planned in Near Future *Transfer of Support
 Project/Proposal Title:
 GASP - Generation of Adaptive Statistical Pattern Recognition Systems
 Source of Support: Subcontract to Frontier Tech, Inc. on Phase-II Army STTR 2002
 Total Award Amount: \$199,875 Total Award Period Covered: 07-01-03 through 06-30-05
 Location of Project: Univ. of Florida
 Person-Months Per Year Committed to the Project. Cal: 3.0 Acad: Sumr:

Support: Current Pending Submission Planned in Near Future *Transfer of Support
 Project/Proposal Title:
 SPITT - Spinning Projectile Imaging and Target Tracking
 Phase-II
 Source of Support: Subcontract to Frontier Tech, Inc. on Army SBIR 2002
 Total Award Amount: \$220,000 est Total Award Period Covered: 24 Months
 Location of Project: University of Florida
 Person-Months Per Year Committed to the Project. Cal: 3.0 Acad: Sumr:

Support: Current Pending Submission Planned in Near Future *Transfer of Support
 Project/Proposal Title:
 OCEAN - A Peer-to-Peer Market for Allocating Grid Computing Resources
 Source of Support: NSF
 Total Award Amount: \$982,484 Total Award Period Covered: 01-01-04 through 12-31-06
 Location of Project: University of Florida
 Person-Months Per Year Committed to the Project. Cal: Acad: Sumr: 1.0

*If this project has previously been funded by another agency, please list and furnish information for immediately preceding funding period.

Current and Pending Support

(See GPG Section II.D.8 for guidance on information to include on this form.)

The following information should be provided for each investigator and other senior personnel. Failure to provide this information may delay consideration of this proposal.			
Investigator: Joachim Hammer	Other agencies (including NSF) to which this proposal has been/will be submitted.		
Support: <input type="checkbox"/> Current <input checked="" type="checkbox"/> Pending <input type="checkbox"/> Submission Planned in Near Future <input type="checkbox"/> *Transfer of Support Project/Proposal Title: OCEAN: The Open Computation Exchange & Auctioning Network A Peer-to-Peer Market for Allocating Grid Computing Resources Source of Support: NSF Total Award Amount: \$982,484 Total Award Period Covered: Jan 1, 2004 – Dec. 31, 2006 Location of Project: University of Florida Person-Months Per Year Committed to the Project. Cal: Acad: Sumr: 1.0			
Support: <input type="checkbox"/> Current <input type="checkbox"/> Pending <input type="checkbox"/> Submission Planned in Near Future <input type="checkbox"/> *Transfer of Support Project/Proposal Title: Source of Support: Total Award Amount: Total Award Period Covered: Location of Project: Person-Months Per Year Committed to the Project. Cal: Acad: Sumr:			
Support: <input type="checkbox"/> Current <input type="checkbox"/> Pending <input type="checkbox"/> Submission Planned in Near Future <input type="checkbox"/> *Transfer of Support Project/Proposal Title: Source of Support: Total Award Amount: Total Award Period Covered: Location of Project: Person-Months Per Year Committed to the Project. Cal: Acad: Sumr:			
Support: <input type="checkbox"/> Current <input type="checkbox"/> Pending <input type="checkbox"/> Submission Planned in Near Future <input type="checkbox"/> *Transfer of Support Project/Proposal Title: Source of Support: Total Award Amount: Total Award Period Covered: Location of Project: Person-Months Per Year Committed to the Project. Cal: Acad: Sumr:			
Support: <input type="checkbox"/> Current <input type="checkbox"/> Pending <input type="checkbox"/> Submission Planned in Near Future <input type="checkbox"/> *Transfer of Support Project/Proposal Title: Source of Support: Total Award Amount: Total Award Period Covered: Location of Project: Person-Months Per Year Committed to the Project. Cal: Acad: Sumr:			
*If this project has previously been funded by another agency, please list and furnish information for immediately preceding funding period.			



FACILITIES, EQUIPMENT & OTHER RESOURCES

FACILITIES: Identify the facilities to be used at each performance site listed and, as appropriate, indicate their capacities, pertinent capabilities, relative proximity, and extent of availability to the project. Use "Other" to describe the facilities at any other performance sites listed and at sites for field studies. USE additional pages as necessary.

Laboratory: Several of the co-PIs have access to lab space in the CISE department which can be used by students working on this project. For example, we have The Harris Mobile and Pervasive Computing and Networking Laboratory. This lab, which occupies over a 1,000 sq. ft., is directed by Dr. Helal,

Clinical:

Animal:

Computer: The CISE department has a significant amount of existing computational and networking resources which can be harnessed as testbeds for deploying the OCEAN server software (in addition to equipment we buy ourselves). In later stages, we have cooperation from UF's Vice-Provost for IT (Dr. Charles Frazier) as well as from the PIs of two major

Office: All co-PIs have sufficient office space and supplies. Students in our department with research or teaching assistantships have desks and computers in shared offices.

Other: _____

MAJOR EQUIPMENT: List the most important items available for this project and, as appropriate identifying the location and pertinent capabilities of each.

No large, monolithic resources are needed - the advantage of a distributed computing project. However, there is an old IBM SP2 mini-supercomputer in our department (as well as other large machines around campus) which we may utilize as another test platform.

OTHER RESOURCES: Provide any information describing the other resources available for the project. Identify support services such as consultant, secretarial, machine shop, and electronics shop, and the extent to which they will be available for the project. Include an explanation of any consortium/contractual arrangements with other organizations.

FACILITIES, EQUIPMENT & OTHER RESOURCES

Continuation Page:

LABORATORY FACILITIES (continued):

is part of the CISE department and supports research in wireless and mobile networking, pervasive and context-aware computing, ubiquitous data and information access, power-aware computing and mobile e-commerce applications. The Harris Lab hosts a variety of wireless networks including several 802.11b wireless LANs, Digianswer Bluetooth network devices and dongols, iDEN wireless packet data and GDM wireless data networks. The facilities in the lab also include 5 Jini boxes (book-size PC's with Jini lookup services and wireless connections), several wireless PDAs including iPAQ, HandSpring Visors, Palms, a number of Java-enabled, IP-based smart phones (Motorola i85s phones), several hand-held computers including 5 Vadem Clio and HP Journada. In addition, a wireless Xybernaut wearable computer and 15 laptops (IBM Thinkpads) are wirelessly connected to the lab network. In addition, the lab hosts a collection of sensors and micro-controllers. Sensors include a variety of X10 based wireless video cameras, motion detectors, thermo-sensors, among others. Micro-controllers include several Motorola M68HCIs, several Dallas-Semiconductor's I-button and four Tini web-based micro-controllers. A Sun Solaris cluster (E450 server, 12 Ultra 60 workstations), a Linux cluster, and an NT cluster are connected to a Cisco router in the lab, which is equipped with state-of-the-art firewall and security software. The Harris lab hosts the current version of the RAMON prototype. More detailed information about the projects and resources in the Harris laboratory can be obtained from <http://www.harris.cise.ufl.edu>.

COMPUTER FACILITIES (continued):

high-performance computing ventures on campus (Dr. Alan George, ECE, and Dr. Paul Avery, Physics) to support OCEAN's deployment on their clusters in concert with the campus's present High-Performance-Computing Grid-coordination initiative.

Office of the Vice Provost and
Senior Associate Vice President

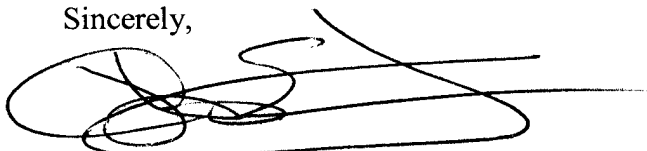
235 Tigert Hall
PO Box 113175
Gainesville, FL 32611-3175
(352) 392-1301
Fax: (352) 392-8735
frazier@ufl.edu

April 23, 2003

Dear NSF Review Panel,

The University of Florida is pleased to support the development and experimental deployment of the OCEAN project, the "Open Computation Exchange & Auctioning Network", led by Dr. Michael Franks of the UF Computer Science and Engineering Department. As one of many High Performance Computing projects under development at UF, OCEAN provides an important opportunity to advance our research in grid computing. With support from the National Science Foundation, the project will contribute to standards for an automated, open, peer-to-peer market for the sale and use of distributed computing resources over the Internet. The combined support of the NSF and UF for project OCEAN can determine whether or not this 'liquid market' for distributed computing will yield the efficiencies we seek.

Sincerely,



Chuck Frazier

Vice Provost for Information Technology

Institute for High Energy Physics and Astrophysics (IHEPA)
Department of Physics

P.O. Box 118440
Gainesville, Florida 32611-8440
(352) 392-9264
Fax (352) 392-8863

April 24, 2003

To whom it may concern,

As Director of the GriPhyN Project and International Virtual Data Grid Laboratory (iVDGL), I wish to extend my strong support for the OCEAN proposal of Michael Frank, Sanjay Ranka, and others in the CISE department at the University of Florida.

GriPhyN and iVDGL are deeply involved in the development and deployment of Grid infrastructure for four frontier data-intensive experiments. These experiments together will utilize Petaflops of CPU power and generate hundreds of Petabytes of data that must be archived, processed, and analyzed by thousands of researchers at laboratories, universities and small colleges and institutes spread across the world. The scale and complexity of this "Petascale" science are driving GriPhyN's research program to develop a Grid-based architecture based on "virtual data" as a unifying concept and to package Grid tools for the research community through our Virtual Data Toolkit (VDT). The iVDGL is deploying a global Grid laboratory where these and other grid tools can be tested at sufficient scale to be interesting and where computational experiments can be conducted by other scientific and engineering disciplines

The OCEAN project aims to provide a market-based infrastructure for resource discovery and allocation that is compatible with existing Grid technologies, such as those we have developed in GriPhyN and deploying in iVDGL. This work is complementary to and synergistic with our work in virtual data.

I agree to provide advisement to the OCEAN team on Grid compatibility issues, as time permits, including having myself and/or one of my students examine the forthcoming OCEAN specifications and Application Programming Interfaces closely, and provide feedback as to how they could be improved so as to better meet the needs of users in the wider Grid communities that I participate in. Once the OCEAN design and implementation reach a stage deemed satisfactory to us, I also agree to have a student or one of our staff participate in experiments in which OCEAN node software is installed and tested on clusters on our persistent Grid Laboratory, and perhaps ported to applications of interest to our community.

I am pleased to offer our participation with the OCEAN team for the benefit of both our endeavors and look forward to working with them.

Sincerely,



Paul Avery
Professor of Physics
Director, GriPhyN and iVDGL