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, financial markets, etc.