INFOMALL

High Performance Computing and Parallel Processing in the Utility Industry

Geoffrey Fox, Paul Coddington, Mike Nilan, Salim Hariri

Northeast Parallel Architectures Center
Syracuse University

Executive Summary

We propose a major study and implementation project which addresses the role in the utility industry of High Performance Computing (HPC), and in particular, parallel processing. The project has three phases:

1) A study phase to evaluate the use of parallel computing for unit commitment, contingency evaluation and transient stability analysis. This will include a survey of available sequential and parallel algorithms for solving differential-algebraic and sparse matrix equations, and the development of new parallel algorithms for these problems which are optimized for power systems applications. A comparison of the performance of different algorithms for solving network and generator equations will be given. This phase will also include a more general survey on the role of HPC and parallel processing in the utility industry.

2) A proof-of-concept phase to demonstrate the feasibility of incorporating transient stability analysis into a fast contingency evaluation program by using parallel computing. At the end of this phase we will have incorporated a fast parallel sparse matrix solver into a prototype transient stability analysis program capable of being interfaced to contingency evaluation software. We will make recommendations on the appropriate hardware, software, and user interface for a production system. We will also investigate the creation of a consortium of interested corporations for developing this prototype into a product in a second phase of the project.

3) A production phase to produce a commercial package for unit commitment and on-line contingency evaluation and transient stability analysis with a high-level user interface incorporating a knowledge-based expert system. The package will include hardware, software, documentation, and training. We expect that the commercial development of this product will be done in collaboration with corporate partners.

Here we propose an initial project to carry out Phases 1 and 2 of this agenda.

- Phase 1 is proposed to run for 6 months from June 15, 1993 to December 16, 1993.
- Phase 2 is proposed to run for 3 months from December 17, 1993 to March 16, 1994.
- The implementation of Phase 3 will be dependent on the results of the previous phases, discussions between NPAC, NMPC and other possible commercial partners, and the acquisition of further funding. We will submit a further proposal concerning Phase 3 at the end of Phase 2. Phase 3 is expected to run for approximately 18 months.

The project will be part of our ACTION-NYS industrial outreach program and substantial cost sharing will be available due to state funding of this program.

Background

The rapid increase in computer performance has led to increased opportunities in the application of computing to the power systems industry. The report, "Comprehensive Study of Supercomputing and its Applicability to the Utility Industry" by TASC, concludes that by 1995 even small improvements in operations, peak reduction, transmission, and avoided customer interruption "can translate into total annual dollar savings that are nearly an order of magnitude greater than the annual cost of operating a supercomputer." Bill Merritt of the Operations group of Niagara Mohawk Power Corporation (NMPC) has suggested that improved contingency evaluation and transient stability analysis could save NMPC on the order of $1.7 million per year without sacrificing system reliability by intelligently reducing the present operating margins. The magnitude of this additional analysis will require supercomputing capabilities in order to produce near real-time modeling to enable operations staff to make better informed decisions.

It is widely recognized within the computer industry that increased supercomputer speeds will necessarily involve parallel computing. This is evidenced by the plans for parallel computing by such traditional supercomputing companies as Cray, Digital and IBM and by the success of parallel supercomputers like the Connection Machine. The Federal High Performance Computing and Communication Initiative (FHPCCI) both endorses this concept and will accelerate progress. The Syracuse University team is at the leading edge of this initiative and currently funded by DARPA to produce portable software systems for parallel machines. Furthermore, considerable experience has been gained in recent years on a variety of parallel computing architectures. For example, the Caltech Concurrent Computation Program, while it was under the direction of Geoffrey Fox, was able to successfully implement 50 large parallel application codes from all areas on the hypercube architecture.

Parallel computing has now been developed to the point where it is possible to examine computationally intense computer applications, for example those computer programs required to support daily power utility operations, to determine what parallel algorithms can implement those applications and what is the most effective parallel computer architecture to run those algorithms. We note that TASC acknowledges the advances of parallel processing but does not consider the implication of this in their report. Parallel processing will be needed to achieve the performance levels quoted by TASC (0.5×10^{12} operations per second) by the year 2000, and we expect at least an order of magnitude better performance than this. However, parallel processing requires the algorithms and software to be rethought and often reimplemented. This critical issue was not addressed in the TASC report and is a key point of our proposal.

We expect that the use of high performance parallel computers will permit faster and more accurate system modeling during daily operations, and this additional knowledge of the system status and security will have direct benefits for NMPC. Daily planning activities that examine stability analysis were previously impossible due to long run-times and the time required for engineers to interpret the results.

Significantly faster computational capabilities and an improved user interface for presenting the output data would enable near real-time system analysis, which should allow a reduction in the conservative operating margins currently required. Reducing excessive operating margins would reduce fuel consumption, thus saving money and reducing environmental impact while maintaining the present standards for reliable operations, or even improving system reliability.

The Syracuse team has excellent qualifications, indicated by the DARPA project mentioned above and by our membership in the prestigious NSF Science and Technology Center devoted to making parallel computation practical. In the latter project, we are involved in parallel software and direct all the parallel application work at the participating institutions (Caltech, Rice and Syracuse Universities and Los Alamos National Laboratory). Dr. Paul Coddington will be the project leader, and will work full time on co-ordination and technical aspects of the project, and liasing with NMPC staff. He has 8 years experience in computational physics and parallel computing. Our proposed project for NMPC also benefits from the State funded ACTION-NYS program, which provides infrastructure support and cost sharing for the parts of the project which have generic value to the power industry.

Project Description

We now detail the components of the project, which has three phases:

1) A study phase lasting 6 months, which will consist of research, study and evaluation in the areas listed below. This phase will also include a more general survey on the role of parallel processing in the utility industry.
This phase to run from June 15, 1993 to December 16, 1993.

2) A proof-of-concept phase lasting 3 months, which will consist of demonstrating the feasibility of using parallel processing and expert systems technology to improve speed, accuracy, and analysis of data for contingency evaluation incorporating transient stability analysis.
This phase to run from December 17, 1993 to March 16, 1994.

3) A production phase, in which we will develop hardware and software investigated in the study phase into a product for use at NMPC, in collaboration with commercial partners.
This phase is expected to take approximately 18 months.

The plan for the production phase will be dependent on the results obtained in the earlier phases. Work on a deliverable product will involve collaboration with corporate partners. During the proof-of-concept phase we will investigate the creation of a consortium of interested parties, including NPAC, NMPC and other companies from the utilities, software and computer industries, for developing the prototype system into a commercial product. At the end of this phase we will submit a proposal for the development of such a product in a third phase of the project. We are aware of at least two companies that would be interested in helping to develop this product.

1) Study Phase

The study phase will consist of study, research, and evaluation in the following interrelated research areas:

1. Unit Commitment
2. Contingency Analysis
3. Transient Stability Analysis
4. Parallel Sparse Matrix Algorithms
5. User Interface and Expert Systems
6. Parallelism and Computational Issues
7. Survey on HPC and Parallel Processing in the Utility Industry

Previous research into parallel unit commitment has shown promising results for shared memory multiprocessors. We will extend this work and investigate new optimization techniques such as simulated annealing, and implement and benchmark different algorithms on distributed memory parallel computers.

We will examine the present contingency analysis performed at the NMPC Operations center, and improve upon it by the application of parallel processing to provide more computing power, and by adding the capability to perform transient stability analysis.

A key research area required to implement efficient parallel transient stability analysis programs is parallel differential-algebraic equation (DAE) solvers. There has been much recent research into sequential DAE solvers, and we plan to capitalize upon that work as we identify the best of those algorithms to solve the transient stability equations concurrently.

Critical to any DAE solver are sparse matrix solvers. Extensive research will be required in this area because sparse matrix computations are some of the most challenging problems for parallel computers.

Another key research area for transient stability analysis is the user interface and the analytical software to classify generator stability. Introduction of transient stability into the security analysis functions in Operations will require computer assisted analysis of simulation output to allow timely interpretation of results.

Finally, computational issues related to the development of this work in the production phase will also be analyzed. We will also conduct a more general survey of High Performance Computing and parallel processing in the utility industry. This survey will extend the TASC report by discussing more carefully the computer architecture and software issues needed to implement HPC in utility applications.

I. Unit Commitment

Syracuse University has previously performed research under contract to NMPC to investigate the feasibility of utilizing parallel dynamic programming techniques on shared memory multiprocessors to perform electrical power system unit commitment. This research was performed by one of our graduate students for Professor Yehuda Wallach in 1990 [1]. This work illustrated that efficient parallel unit commitment algorithms exist for shared memory multiprocessors: efficiencies ranging from 56% to 91% were obtained. It is anticipated that these techniques can also be implemented efficiently on distributed memory multiprocessors.

The parallel dynamic programming unit commitment prototype that had the best efficiency was also the most coarse-grained. This implies that additional computations that improve the accuracy of the unit commitment calculations would also improve the efficiency of the parallel implementation. For example, it is possible that the additional calculations from adding new features such as environmental dispatch may be absorbed by better utilization of the available processors, so that little extra time will be required to perform a more valuable and realistic unit commitment study.

In addition to the dynamic programming implementation of unit commitment, there are other areas for research that could potentially improve the quality of unit commitment estimates. Dynamic programming implementations of unit commitment are affected by exponential growth in memory and computational time as the number of generators increases, which significantly limits the number of generator combinations that can be examined and may result in non-optimal solutions. One example of an optimization technique that could significantly reduce this problem is simulated annealing, which has efficient parallel implementations.

Parallel unit commitment techniques will significantly decrease the time to perform a unit commitment study, which in turn, may permit additional fidelity to be included into the models. Frequent unit commitment with high fidelity models should significantly reduce generation cost by accurately predicting the most efficient combinations of generators to supply adequate power for time-varying loads.

For this application we will:

- Implement the existing parallel unit commitment algorithms using dynamic programming on a distributed memory parallel computer.

[1] "Parallel Implementation of the Unit Commitment Problem on Shared-Memory Multiprocessors", by D.P. Koester, Fall 1990

- Carry out a research and implementation program for parallel unit commitment using simulated annealing.
- Investigate the use of higher fidelity models incorporating additional features such as environmental dispatch.
- Benchmark the performance of these different algorithms on parallel machines.

II. Contingency Evaluation

Contingency evaluation is currently done on a large number (of order 200) of independent cases. Contingency screening is first performed by running a fast contingency evaluation program which uses very approximate techniques. This orders the contingencies by increasing severity, and around 30 of the most severe cases are evaluated using a much more accurate program. Speeding up the contingency evaluation program would allow the use of a more accurate screening process, or even the elimination of screening if the program could be made fast enough to allow a full evaluation of all the contingencies.

Since the different contingencies are all independent, this problem is *embarassingly parallel*, meaning that different cases can be sent to different processors and processed in parallel with optimum efficiency, by just using the sequential algorithm on each processor. However before the contingency evaluation can be performed, the state estimator program must be run. To get near real-time performance for contingency evaluation, we may need to develop parallel algorithms to speed up state estimation.

For this application we will:

- Study the current contingency evaluation software used at the NMPC Operations center.
- Evaluate methods for incorporating parallelism into the state estimator and security screening process.

III. Transient Stability Analysis

Transient stability analysis involves the study of the electrical network and generator dynamics for several seconds following a disturbance. The dynamics of each generator in the system is described by multiple differential equations, while the coupling of the generators to the electrical network is described by algebraic equations. Together these lead to a system of differential-algebraic equations (DAEs) that must be solved simultaneously. Since a large system of DAEs needs to be solved at small time increments to minimize the errors in numerical integration of the differential equations, the transient stability analysis is a highly compute intensive problem.

Traditional approaches to solving transient stability DAEs, such as partitioning and linearization, have used simple trapezoidal integration techniques that require very small step sizes during the integration steps. Meanwhile, there has been significant research in the numerical analysis community on fast, accurate, and efficient DAE solvers. However, very little work has been done to date on using these more recent public domain DAE solvers for the transient stability problem. Numerous public domain DAE solvers exist, based on different computational techniques. These include:

1. DASSL – a multistep, backward differentiation formulas technique
2. DASRT – a variant of DASSL that handles discontinuities
3. CDASSL – a parallel variant of DASSL
4. RADAU5 – a multistep, implicit Runge-Kutta based technique
5. RODAS – a Rosenbrock technique
6. LIMEX – an extrapolation technique
7. SEULEX – an extrapolation technique

DASRT may be suitable for this problem, since it has the capability of locating discontinuities, which is a condition encountered in transient stability analysis. A parallel version of DASSL, CDASSL, has shown promising results in a petrochemical engineering application. Inherent in all DAE solvers are sparse matrix solvers. In general, most of the calculations – excluding the solution of the sparse matrices – are highly parallel. The research required in this task is to identify the most efficient DAE solver for the specific nature of the transient stability equations.

An additional source of parallelism in the power system transient stability problem is the examination of multiple, independent fault events simultaneously. Such a technique could significantly improve speed-up and efficiency on parallel computers by matching the simulation requirements at NMPC to available computer architectures.

For this application we will:

- Investigate the use of various DAE solvers and parallel sparse matrix algorithms, and the interfacing of these methods to create fast parallel DAE solvers.
- Benchmark the performance of these different algorithms on transient stability analysis using sample data from the NMPC power network and larger regional networks, since the performance of DAE and sparse matrix solvers can be highly problem dependent.
- Benchmark these programs on different parallel computers, in order to determine the most productive and cost-effective hardware for this problem.
- Develop the knowledge base required for an expert system and improved user interface to classify and display network output from the stability analysis program.

IV. Parallel Sparse Matrix Algorithms

Sparse matrix calculations are some of the most challenging problems for parallel machines, but they are key to many power system applications, including load flow and transient stability analysis, for which they represent the majority of the computation. Various organizations have efforts in progress to develop general concurrent sparse matrix solvers. There has already been significant progress in this area, especially by Michael Heath at Oak Ridge National Laboratory, Ian Duff in Europe and groups at Sandia National Laboratory.

However, specialized parallel sparse matrix solvers which can capitalize on any regular structure in the sparse matrices for a particular problem will generally give better performance than general solvers. In-house research at Syracuse University is currently focussing on the development of specialized sparse matrix solvers that exploit hierarchical structure in the underlying power systems problems.

There are distinct portions of the matrix that represent the generator equations and the power distribution network. Diakoptics theory and the actual matrix sparsity structure (the power distribution network structure) will be used to develop algorithms to reorder the sparse matrices resulting from power system applications into block-bordered diagonal form. The result of this reordering is that many of the operations that would be sequential in a general solver are now highly parallel. This will simplify the development of scalable, portable algorithms. Reordering can also be used to improve load balancing for both direct and iterative techniques for matrix solvers.

One of our graduate students is presently working on parallel sparse matrix techniques especially suited for power systems applications. Another graduate student in our team has worked with a group at Lawrence Livermore National Laboratory on a related petrochemical simulation using parallel DAE and sparse matrix solvers.

Parallel sparse matrix solvers are critically dependent on a proper match between the particular problem and the machine architecture and "general approaches" are often less than optimal. It is essential to have matrices of representative sizes to successfully research this real-life problem. We will work with NMPC to obtain the actual structure of the transmission network for three interconnected power grids:

1. the NMPC network
2. the New York State Power Pool network
3. the Eastern US regional power network

It would not be surprising if our research indicates that different sparse matrix techniques are applicable for the different size networks.

The problem size will also influence the applicable hardware. We expect that the sparse matrix solver for a system the size of the NMPC network, with around 300 generators and 2000 buses, would perform well on the 32 node MIMD parallel computers currently available at NPAC. However, a "small", say 10 generator and 50 bus network, system would not have enough parallelism to easily make use of such machines. Meanwhile, the large 10 teraflop (10^{13} floating point operations per second) parallel supercomputers available in the year 2000 will probably contain thousands of nodes. This will be too much parallelism for a simulation of the NMPC system, due to size of the sparse matrices, but could be well suited for the transmission system covering a large geographic area. The coordination of grids over a large region may be an exciting possibility opened up by the parallel supercomputers of the future.

For this application we will:

- Monitor work by Heath and others on general sparse matrix solvers.
- Continue in-house research at Syracuse University to develop specialized sparse matrix solvers that exploit the hierarchical structure in power systems applications.
- Compare the performance of direct and iterative methods, and investigate the use of composite direct/iterative methods.
- Benchmark different parallel sparse matrix solvers for different sized power networks and parallel computers.
- Interface the parallel sparse matrix solvers with the DAE solvers also being investigated.

V. User Interface and Expert Systems

The time required for performing transient stability analysis includes the time needed to interpret the large quantity of data produced by the program. Currently, stability classification requires lengthy scrutiny by a trained engineer. In order to obtain maximum benefit from an on-line transient stability analysis program, this data analysis procedure must also be sped-up so as to work in near real-time. This can only be achieved with some form of computerized tool to help with the complex task of interpreting the output of the stability program. We propose to implement a high-level user interface incorporating expert system and possibly neural network technology to greatly reduce the time required for this task.

Expert systems have been under intense study by research organizations for the last 10 years, and thousands of expert systems are now in use in a wide variety of applications in many large corporations. There has already been a considerable amount of research into the use of expert systems in the power industry, including transient stability analysis. Some experimental transient stability expert systems have been implemented, with promising results.

Professor Mike Nilan of the School of Information Studies at Syracuse University is an expert in cognitive science, user interfaces, data representation, and expert systems, and has worked successfully with Niagara Mohawk on previous projects. Under the direction of Professor Nilan, we will develop the basis of a high-level user interface for the transient stability analysis program. This will include an expert system to reduce the amount of program output which requires human interpretation, for example by filtering out "obviously" transient stable cases, and to represent the data which does require analysis by the operator in a form which can be more easily interpreted. This will reduce the training overhead for users of the program. Dr Salim Hariri of NPAC has experience at implementing expert systems and will be working on this aspect of the project.

The benefits of such a high-level user interface are not just in allowing near real-time analysis, but also simplification of operator training, minimization of human error, and time independent expertise, since the expert system is available on-line at any time. The system can also be improved by accumulating additional knowledge over time.

For this application we will:

- Work with NMPC engineers and operators in the Planning and Operations departments who are currently using transient stability analysis, or hope to use it in the developed product, in order to develop a user-based decision process model and undertake knowledge acquisition and analysis for the development of an expert system and its user interface.

- Study the recent research on the application of neural networks to the power industry, in particular to stability classification, and exploit the expertise in neural networks at Syracuse University to investigate their application to this problem.

VI. Parallelism and Computational Issues

STAG, the vendors of the software currently used by the Operations group at NMPC, are currently porting this software to a Local Area Network of IBM RS/6000 workstations. This environment appears to be more viable for use with parallel processing than the current STAG software run on IBM mainframes, since the new software will be more up-to-date and portable. Software running on a distributed workstation environment should allow easier interfacing to parallel machines, and both the software and hardware of such a system should also be easier to maintain and more cost-effective, robust, and reliable than the current system. This should thus allow a substantial reduction in system down-time, which is a major concern of the Operations group.

The new STAG contingency evaluation program is expected to be available in the third or fourth quarter of 1993, and most other features of their Energy Management System will be available before then. This release schedule will allow us to evaluate the new software during the proof-of-concept phase, and allow us the possibility of using their new software in the production system.

The parallel hardware and software used in a production system should be readily interfaced to existing and future systems, and needs to be scalable. This will allow the initial product to be used by NMPC, but should it become desirable, enable the system to be scalable in both hardware (the number of processors) and software to accommodate larger networks. Larger networks might incorporate additional detail by including lower voltage buses, the network for the New York Power Pool, or the network for the entire North East power grid.

For the power systems applications to be studied in this phase, we will:

- Conduct an extensive evaluation of the current software used at NMPC, as well as other available power systems software, to determine how easily these programs can be interfaced and modified to include parallel processing capability.

- Evaluate the relevant architecture and software issues for using parallel computers, including scalability and portability.

VII. Survey of High Performance Computing and Parallel Processing in the Utility Industry

During the study phase, we also propose to conduct a detailed survey of High Performance Computing and parallel processing in the utility industry. We will extend the TASC report mentioned above in two ways. Firstly, we will work with NMPC, especially in the load planning area and perhaps also the nuclear plant simulation problem, to extend and improve the analysis of the applications described in that report. Secondly, we will discuss more carefully the computer architecture and software issues needed to implement HPC in the TASC applications and further areas suggested by NMPC.

We have provisionally defined eight key HPC technologies of interest. These are:

1. Sparse Matrix Algebra

2. Signal Processing

3. Numerical Optimization

4. Component Modeling

5. Expert Systems

6. Neural Networks

7. Visualization and Virtual Reality

8. Computational Fluid Dynamics (CFD) and Continuum Simulations

We are familiar with the state-of-the-art in all the above technologies. Each technology has rather different computer architecture and software requirements. For instance, signal processing, neural networks and CFD are naturally suitable for massively parallel SIMD systems. Sparse matrix manipulation and component modeling requires the flexibility of the MIMD architecture. Expert systems, numerical optimization and virtual reality are naturally implemented in a heterogeneous network involving high performance workstations.

For this project, we will quantify the comments above and provide a survey relating utility applications to the most appropriate technologies, both hardware (SIMD, MIMD, number of processors, etc.) and software (existence or feasibility of parallel algorithms, etc.). We will indicate where it is realistic to use HPC and parallel processing in the electric utility industry over the next ten years.

2) Proof-of-Concept Phase

During this phase we will:

- Study the current PTI transient stability software used at the NMPC Planning department, and investigate the problems of interfacing this software with our parallel DAE solver and the STAG contingency evaluation software.
- Use the most promising parallel sparse matrix and DAE solvers to produce a prototype on-line transient stability analysis program.
- Develop the basis for an expert system and improved user interface to classify and display network output from the stability analysis program.
- Demonstrate the feasibility of incorporating a fast parallel on-line transient stability analysis capability into a contingency evaluation program.
- Evaluate the new workstation-based STAG EMS software to determine how easily it can be interfaced and modified to include parallel processing capability.
- Provide recommendations for hardware and software for the systems to be developed in the production phase. This will include the type of hardware to be used, the number of processors required, and an appropriate parallel software environment.
- Continue our research into more efficient parallel algorithms, and the application of neural networks to the project.
- Continue work with NMPC operators and engineers to develop and improve the knowledge base for the user interface and expert system.

3) Production Phase

During phase 2, we will talk with NMPC and other companies in the power utilities, software and computer industries about developing the prototype systems demonstrated in the proof-of-concept phase into a commercial product for use at NMPC. We envisage forming a consortium of interested corporations to fund and develop such a product.

Based on these negotiations with NMPC and other corporate partners, and the results obtained in phases 1 and 2, we will submit a further proposal for a third phase of this project, to develop a commercial product. We anticipate that this will include a system of parallel hardware and software for unit commitment and an on-line near real-time security model incorporating contingency evaluation and transient stability analysis. This system will include an improved user interface with an expert system for fast analysis of stability data. We expect that this production phase will take on the order of 18 months.

During this phase we will also continue to study more efficient parallel algorithms and new technologies which may benefit the final product. We will continue to work with NMPC operators and engineers to refine and improve the user interface and expert system.

Summary of NPAC and NMPC Deliverables

Meetings

- NPAC faculty, research staff and graduate students associated with this work will meet with NMPC staff and consultants for one to three meetings, as required, each month to review progress, discuss problem areas and discoveries, and for general information exchange.
- NMPC will arrange for appropriate application area personnel to meet with NPAC researchers in order to conduct the study and evaluation program, acquire knowledge for the expert system, and to provide software, data, and other requirements for the contingency evaluation and transient stability analysis package.

Reports

- Progress reports will be delivered by NPAC to NMPC at the end of each month. These reports will outline the status of each of the five main research and development areas:

    1. Unit Commitment
    2. Contingency Evaluation
    3. Transient Stability Analysis
    4. Parallel Sparse Matrix Algorithms
    5. User Interface and Expert System

  We will also report on the status of our research into computational issues for the above areas, and on the survey of applications of HPC in the utility industry.

  A work plan and goals for the following month will also be included.

- A final report for the study and evaluation phase will be delivered by NPAC to NMPC at the end of 6 months. The report will include detailed descriptions of the following items for each of the five research and development areas listed above:

    - Theoretical background.
    - Research issues and new results.
    - Methods and algorithms used.
    - Computational issues, including parallel algorithms, languages, software and hardware.

The report will also include:

  - An analysis of different DAE and sparse matrix solvers for power systems applications.
  - A general survey of the role of High Performance Computing and parallel processing in the utility industry.

- A final report for the proof-of-concept phase will be delivered by NPAC to NMPC at the end of 9 months. This will include software for the research and development prototype system for unit commitment, contingency evaluation and transient stability analysis. The report will also include:

  - Details of the design and implementation of the system.
  - Benchmarking of different software and hardware.
  - Evaluation of the performance of the system.
  - Feasibility study on the use of an expert system for fast analysis of transient stability data.
  - Documentation for the prototype software, including a user guide and reference manual detailing the major subroutines.
  - Recommendations for the design, hardware and software for the system to be developed in the production phase.
  - An analysis of a possible consortium of corporate partners for development of the prototype system into a product.