InSIM

Proposal Abstract in Response to BAA 96-20: JFACC 24 June 1996

for Category: TECHNOLOGY BASE RESEARCH and DEVELOPMENT,
Subcategories (6) Planning in a Simulated (predictive) World
(8) Plan Evaluation and Visualization

Attn Mr. Alan Fredrick

Integrating Simulation Results into *JFACC after next* Systems (InSim)

Principal Investigator:
Gio Wiederhold, Stanford University
Implementation Subcontract:
Lynn Dean, Lockheed-Martin Idaho Technology Company (LMITCO)

Period of Performance: 1 January 1997 to 31 August 1999
Estimated Costs (POM): Base: $ 2,772,732; Options $ 639,500; Aggregate $ 3,411,232

Institutions:

Stanford University
Gio Wiederhold
Computer Science Dep.
Gates Computer Science Building 4A
Stanford CA 94305-9040
415 725-8363 fax 725-2588
<gio@cs.stanford.edu>

Lockheed-Martin Idaho Technology Company (LMITCO)
Lynn Dean
EROB building
2525 North Fremont Rd
Idaho Falls, ID 83415-3750
208 526-5316 fax 526-4313
<lad@inel.gov>

Myriad Software
Ted Linden
2245 Tasso street,
Palo Alto CA 94301
415 327-2973 fax 327-5509
<linden@myriadSfw.com>

Crystaliz Inc.,
Sankar Virdhagriswaran
686 Virginia Rd
Concord MA 01742
508 287-4511 fax 287-4512
<sv@@hunchuen.crystaliz.com>

Administrative contact: Marianne Siroker
Stanford University, Gates Computer Science Building 4A
Stanford CA 94305-9040
415 725-8363 fax 725-2588
<siroker@cs.stanford.edu>

Submitted to: Defense Advanced Research Projects Agency
c/o Col. Ed Patneaud
Information Systems Office, Planning and C3 Division
3701 North Fairfax Drive
Arlington, Virginia 22203-1714

This proposal abstract contains no restricted or proprietary information.

II Executive Summary

The goal of the **InSim** project is to make simulation results readily accessible to JFACC planners. The InSim project proposes research, development, and validation of a “Simulation Access and Query Language” or SimQL. SimQL will give planners access to information about simulated futures in the same paradigm that SQL provides to gain information about past events. Today simulations are initiated and their results processed outside the prime information flow in planning and plan assessment.

Military commanders must plan actions beyond the point-in-time where data is available. To deal with the complexity of alternative CoAs they increasingly use simulations. We quote from *"New World Vistas, Air and Space Power for the 21st Century"* [McCall:96, Air Force Scientific Advisory Board (AF SAB), April 1996], the two `Capabilities Requiring Military Investment' [Information technology volume, pp. 9] are:

1. Highly robust real-time SW modules for data fusion and integration
2. Integration of simulation software into military information systems

While this investigator has considerable experience in the first area, having initiated the I3 program at DARPA, and developed the concepts for integrating past information under earlier DARPA funding [W:86, W:89, W:92, W:95]; we perceive that today rapid progress is being made in data access and we are confident that JFACC will be able to select credible proposals in that area. Our InSim proposal focuses on the second capability cited by the AF SAB: integration of simulations. We have already started work (with small-scale support from DARPA DSO, Pradeep Khosla, PM) to define and implement a SimQL. Having an access language for simulations in *JFACC after next* will significantly augment the capability of intelligent agents in integrated information processing, by allowing concurrent and seamless access to factual data and to simulated projections; for example, integrated access to logistics status data with projection of future status.

SimQL is NOT intended for writing new simulations, but to provide access to results from simulations, via a wrapper infrastructure, as also used in the ACPT initiative for legacy data. The simulations supported by SimQL will range from spreadsheets, to weather forecasting, to fully distributed simulations as now performed in military training exercises (DIS). The concept of SimQL follows the experience gained with SQL for databases.

We will structure the SimQL interfaces to match closely the emerging conventions from the DARPA effort in Intelligent Integration of Information (I3) sharing the architectural concepts and standards. Current mediators being built at INEL, jointly by Stanford, ISX, and LMITCO use standard COTS (CORBA) technology for the interfaces, while using formal object models, suitable for flexible data integration, internally. SimQL will provide access to the metadate which define the exposed model of a simulation, allow setting of input parameters, and the acquisition and subsequent integration of simulation predictions.

To demonstrate the effectiveness of SimQL the InSim project will build mediators that demonstrate access to multiple heterogeneous data sources as well as to simulations, and access interfaces (APIs). We will integrate these closely with the architectural concept to be developed in JFACC, to assure seamless integration and effective delivery to the field.

Some useful options will provide hardened standard interfaces to DIS, KQML-based portability to popular machines as PC compatibles, and direct presentation capabilities.

III Problem Statement, Technical Goals, Operational Benefits

III.a Problem Statement

Decision-making in planning, both in the military and in business depends on knowing the past situation, the current situation, and the future. Computer technology has made great strides in accessing past information through improvements in databases, object-bases, and mediation technology. Communication for access to current events is also improving by leaps and bounds through the World-Wide Web, and is being integrated with the historical information provided by databases.

To project into the future, as every planner must, requires some form of simulation. In earlier generations, that simulation was done in the planner's mind, the planner would look over reasonable scenarios and develop mentally alternate courses-of-action, focusing on those that had been worked out in earlier exercises. When the plans become complex, tools are needed for pruning, presentation, and assessment. Sand tables are still used for the training of military planners, but are increasingly being replaced by computer-based simulations. In business, spreadsheets are used frequently. Simple heuristics help the planners in pruning, for instance, military doctrine demands air superiority, and that the forces and their armaments would be sufficient to overwhelm the enemy, while there would be enough uncommitted forces and logistical support to provide a reserve backup.

While these rules still hold, the world is getting more complex, the resources are fewer, and more specialized, so that traditional heuristics are often barely relevant. Rapid, ad hoc, access to information for planning, as exemplified by the use of DART during the preparations for Desert Storm, was extremely valuable, and superseded effectively many ponderous support planning systems that had been developed using older technology. However, DART in its planning could not execute arbitrary planning scenarios, and only one planning tool (originally developed for airport gate allocation by American Airlines) was used after some rapid adaptations. Today, we have crucial simulations in all aspects of military and commercial planning. Logistics plans are developed by simulating alternate transport modalities, capacities, and risks. Production planners execute simulations to see how they can best exploit their resources. Financial planners use spreadsheets to work out alternate budgets. Most importantly, tactical plans are fully or partially simulated in all military exercises. The expectation, as cited in [McCall:96], is that this technology will transfer into military operations in the future. Even limited situational assessment of current status requires projection. Since C2 data are often out-of-date, commanders must routinely make undocumented projections even to judge the current readiness situation and obtain a complete tactical picture.

Lacking in today's practice (and even explicitly in the *JFACC after next* planning process as outlined in Section 1.2 of the BAA96-20 PIP), is the ability to interoperate with simulations at a direct functional level. Although interchange standards have been developed for the exchange of data objects within the SimNET initiative operated by the Defense Modeling and Simulation Office (DMSO), there is today no direct access to simulation results, and no capability for interoperation with existing military data and information systems.

III.b Technical Goals

In the InSim project we plan to perform research, development, and implementation of a "Simulation Access Language" or SimQL, as a major contribution to interoperation in large-scale information systems. The InSIM project addresses the Technology Base category of BAA 96-20, JFACC, specifically subcategories (6) Planning in a Simulated (predictive) World and (8) Plan Evaluation and Visualization. It also intends to deliver tools, as requested in category Applications and Services. Since simulation is only implicitly called out in that second category, InSIM, in that context, would be relegated to subcategory (10) Other. Figure 1 shows the place of simulation access in the Objective-based Planning Process.

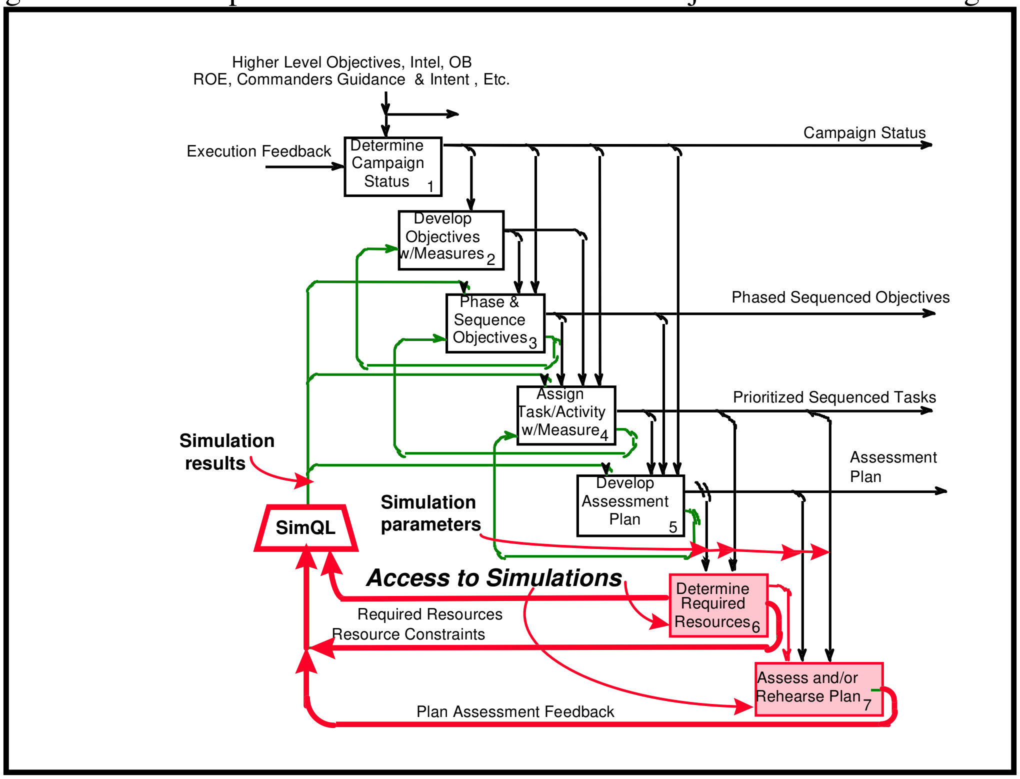

Figure 1 Planning path Infrastructure supported by InSim.

Access to simulations through SimQL will provide JFACC applications with powerful, generalized interface to simulations results. The Simulations themselves are likely to be pre-existing, and will be wrapped to provide robust access to the model parameters and execution results. Simulations may be written in any of the current simulation languages, provided that the systems written in those languages provide a compatible `machine-friendly' interface to exchange model information and respond to queries.

It needs to be re-emphasized that a SimQL is not a language to write simulations, just like SQL is not a language in which to write a database system. Just as SQL provides applications access to a wealth of database technology and to database services that are often maintained by others, we expect that SimQL will provide access to the growing portofolio of simulation technology and simulation services maintained by others, and break the bottleneck which is now experienced when simulations are to be a part of larger systems in planning for military, industry, or business.

Our vision is based on the model that has been successful in data access. Twenty years ago the major effort in writing software was in file programming, and most errors and ongoing maintenance concerned file operations. In the 1970's substantial effort and solid ideas made data sharable and accessible by more than one application. Eventually the SQL language enabled database systems to be built independently of the user applications, and in time database systems reached a level of performance so that today application programming of operations on data files is rare. Databases and the proposed wrapped simulations share with the applications a model, expressed in a schema, and provide algebraic capability in the interface. Having an algebra permits alternate execution sequences to be devised. The alternatives may produce equivalent results faster or better by some measure, and the major task of an SQL-based database is to translate the query into a sequence of operations that satisfies the query and executes as fast or as responsive as possible.

The broader context of SimQL is illustrated in Figure 2: the augmentation of the mediation layer to bring timely and relevant past and predictive information to the collaborative planning teams.

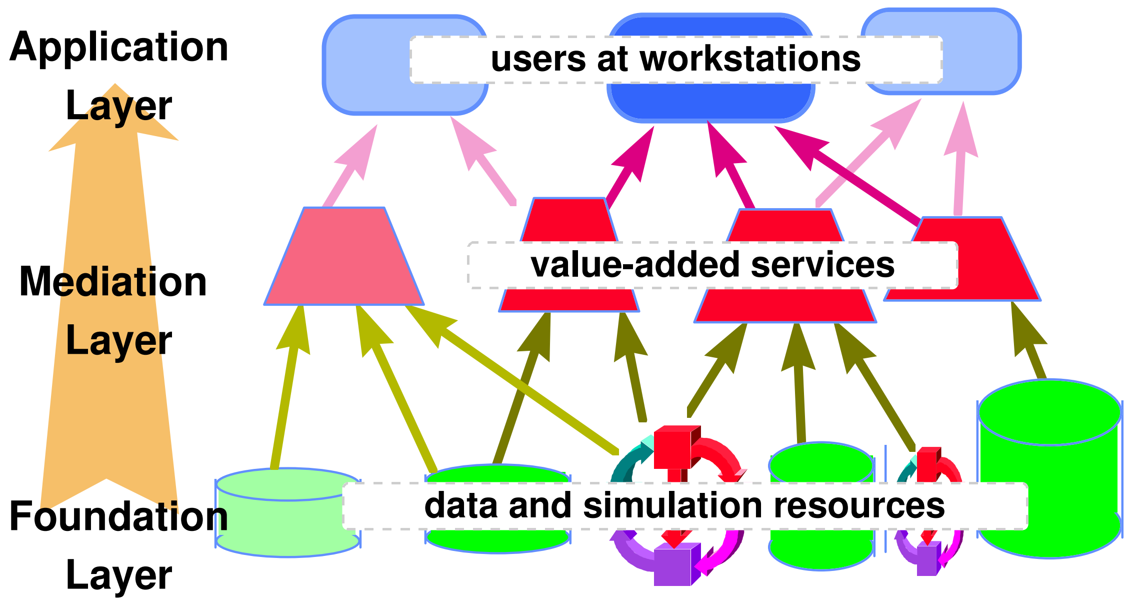

Figure 2: Transforming Data to Information

III.c Operational Benefits

The benefits of being able to integrate simulation results into *JFACC after next* are to provide a solid basis for truly objective planning, integrated with the many databases that other JFACC tasks will provide. Up-to-current-date situational assessments can be provided. Applications will be able to compare resource and effects specifics for multiple Courses-of-Actions, all based on comparable predictive models. Isolation of the source simulations through the formal SimQL interface will encourage the development and acquisition of better simulation systems as the technology develops. As the use of this technology broadens we also envisage that compliant COTS tools will begin to appear, drastically lowering the costs now incurred by DoD where nearly all the simulations are handcrafted to military-only specifications.

Simulation is an area of great ferment with the potential to impact information processing and utilization in a major way. Its progress, however, has been isolated because of the difficulties of access to simulation results from more general information systems. Today, the barriers to interoperating with multiple data sources are being overcome. It is a crucial time to make simulation results available to planners within the same paradigm.

Our task is not to write better simulations or devise better languages for writing simulations. If more people can use simulations more effectively, then market forces will take care of better simulations, and of the standards needed for them to interoperate.

Neither will we provide the evaluation and ranking analyses among alternate CoAs. That role is appropriate for the applications that SimQL will support with information. Actual integration, considering data heterogeneity and multiple ontologies, as proposed by other JFACC bidders will also not be the focus of InSim, although the models exposed through SimQL will provide the ontologies required for these tasks. The models will define information needed for planning that can be obtained from simulations and the required decision parameters to specify the results needed.

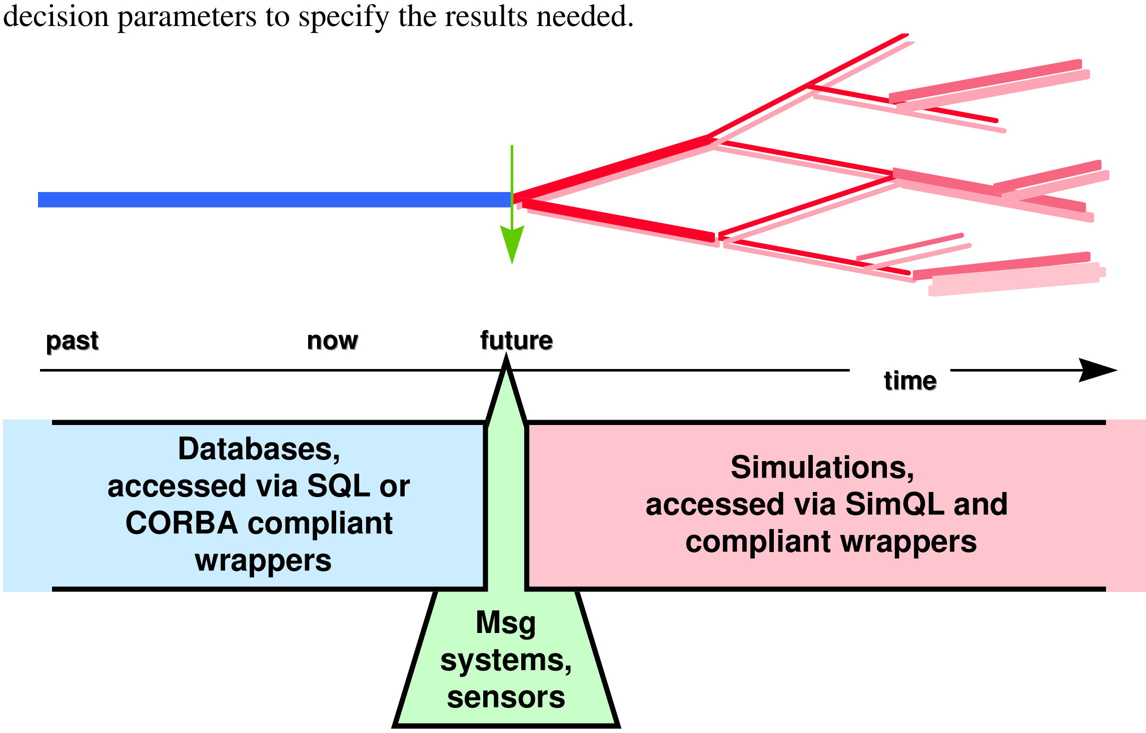

Figure 3: Seamless integration from the past to possible futures

IV Technical Approach

Our approach in InSIM is to develop the SimQL technology incrementally, and integrate it as soon as possible with other JFACC efforts. That will have a dual benefit: we will learn rapidly about weaknesses and irrelevancies in our formulation, and other JFACC tasks can familiarize themselves with the new capabilities which InSim can provide. Since we are treading on new ground here, it is important to augment our conceptual abstraction with real-world tasks, even if they will be initially limited and likely awkward.

We will identify some existing simulations that are candidates for interoperation with databases in the JFACC setting. We will seek to obtain systems written in different languages, so that SimQL will not be biased towards any particular simulation approach or methodology. Candidate simulations include military DIS training simulations from our MIFT project (Option 1), simulations of how to accomplish JFACC planning tasks, and some of the innumerable spreadsheet-based systems used for logistics and financial planning.

Our objective is then to:

1. Develop the SimQL language, which can be used to access, initiate, and query autonomous, wrapped simulation systems. The simulations themselves are typically executed on remote computers.

2. Devise the sharing of models between military customers and simulation developers. The models will not give the customer `programming level' control.

3. Validate the use of SimQL in a three-layer architecture, following the conventions being defined within the ARPA I3 architectural specification.

4. Provide a mediator that can take both simulation results and database results, following the JFACC architecture, and provide an integrated picture for the customer

By immediately focusing on the needs of end-user applications, we expect to avoid some problems which SQL encountered. A major problem in SQL is an awkward formal syntax and programming interface, because SQL was initially conceived as a `user-friendly' language. Today SQL is rarely used by end-users, but is typically embedded in an application program interface (API) to which data management is delegated. By using KQML, a `machine-friendly' language, designed for interoperation in a multi-customer, multi-site, and multi-resource environment as the transport mechanism for SimQL, we immediately gain interoperability with ongoing efforts in access to heterogeneous database and world-wide-web resources. Work leading to implementation of interoperation in the KQML context is requested in Option 2.

A simple illustration of our intent for the InSim project can be gleaned from showing two queries, the first one addressed to a database and the second one to a simulation:

1. SELECT Temperature, Cloudcover, Windspeed, Winddirection FROM WeatherDB WHERE Date = `yesterday' AND Location = `ORD'.

2. PREDICT Temperature, Cloudcover, Windspeed, Winddirection FROM WeatherSimulation WHERE Date = `tomorrow' AND Location = `ORD'.

The weather example is not intended to illustrate the actual proposed syntax of SimQL, although it should be easy to provide users familiar with SQL an interface that is syntactically consistent with SQL, or with the Remote Data Access (RDA) standard of SQL, as validated by the National Institute of Standards and Technology (NIST).

A SimQL will have to address several issues that do not arise when SQL deals with databases, because simulations have features that differ from databases. The primary one is that each returned value will have two parts: the centroid of the predicted value, and the uncertainty.

The models shared between the customer and the simulation will be of necessity more complex than database schemas, and a major objective of the research is to see how little need be exposed to the user's application programs. As a basis for communication we can build on languages as KQML, which support a multi-function (i.e., more than just SQL SELECT and UPDATE), multi-user and multi-resource system setting.

The algebra of SimQL must allow composition with results provided by databases, as by SQL or KQML. By default, such data will have a zero value for the associated uncertainty, although when multiple, overlapping databases are being accessed, those sources may be tagged with their reliability.

The research proposed in the InSIM provides two major strategic benefits:

1. Systems serving customer needs, can be rapidly built, and incorporate the best data system and simulation technology.

2. Simulation development and its technology can proceed independently, delivering incremental improvements.

Customers will no longer be bound to massive, obsolescing technology. Improvements that conform to standard interfaces can be incrementally introduced, avoiding the specter of JFACC becoming another legacy system ten years hence.

The resulting information systems paradigm will allow a seamless integration of the past history with the predictions for the future. There will of course be multiple futures if we take into account decisions that will lead to alternate courses-of-action (CoAs). It is important to assess costs, resource use, and risks at several points in the future for alternate CoAs. The lack of convenient access means that today such comparisons are handcrafted, and not directly supported by the information systems available to the military. The closest approximation to such facilities are multi-layered spreadsheets, which allow the results of alternate assumptions to be viewed by flipping through the layers.

Presentation of such results will require innovation. For the delivery to JFACC of innovative results in visual form we have formulated Option 3. Options 2 and 3 will present results to planners, independent of a particular application. Presentation will use common fieldable equipment, as PCs, using KQML or the standard agreed on in the JFACC program.. However, even without that option the base information can be fed to JFACC applications for integration with their needs.

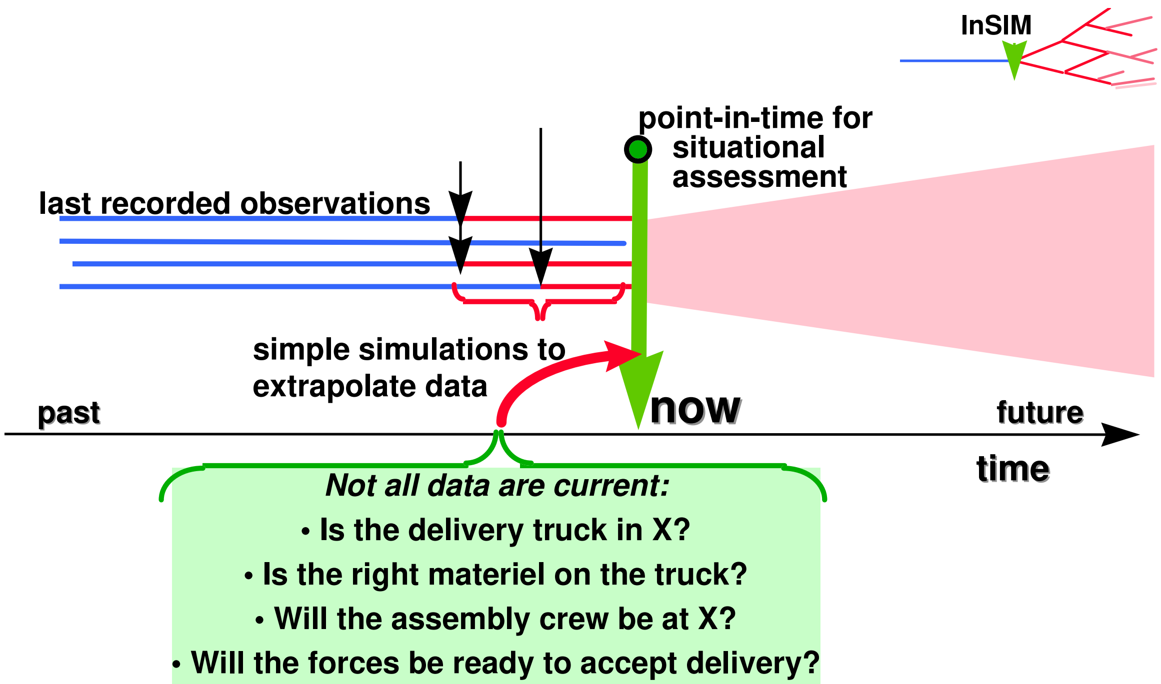

Figure 4. Use of simulation to bring data up-to a best current state.

Simulation may also be needed to bring somewhat-out-of-date situational information to the current point in time, for instance when logistics and shipping data do not reach up to the current instant in time. We propose that such a simulation be the initial milestone, since it is easy to verify and quite relevant to the JFACC commander. This proposal is quite novel and addresses problems that are just beginning to be recognized as we are overcoming the barriers to the integration of historical and current data. Planning was of course always forward looking. A major investment has been made in conceptual planning at ARPA for some time, and we expect to learn from them what the crucial model parameters are. These conceptual planning exercises, while valuable in providing insight about making tradeoffs in alternate courses-of-actions, have however not been well integrated with actual simulations or databases. Similarly, simulation is one of the major foundations for military training and development, and is making impressive progress. Work in the SimNet context, using the DIS standard, is now considering selective access to their data, in order to deal with problems of scale and feedback. On a DARPA CAETI-funded project (Mediation to Implement Feedback in Training, MIFT) Stanford and Myriad Software are accessing DIS-based training simulation data in order to develop model-based feedback. Integrating simulations is being performed on an ad hoc basis where needed, for instance in semi-conductor device simulations [Bob Dutton, Stanford]. We will learn from this effort, and provide the essential bridge that can make all of these systems synergistic and effective.

V Program Plan and Milestones

We will start from work currently being performed at Stanford, where a prototype SimQL is used to access some spreadsheets and weather information available on the world-wide-web. For this JFACC we will spend an initial time clarifying the joint objectives, so that all participants will follow the same path. At Stanford our resources will be mainly students and post-graduate researchers, who at the same time are performing academic research in related topics, typically within our active Database Group. Lockheed-Martin Idaho (LMITCO) will support this effort by providing knowledge and experience in the following areas: simulation software development, the development and implementation of mediators and distributed databases, and implementing and managing large software engineering projects for the Department of Defense (DoD) and other government agencies. LMITCO has collaborated with Stanford University and ISX during the past two fiscal years on the development and implementation of mediators to access data from existing heterogeneous databases. Furthermore, LMITCO specifically brings experience in simulation and software systems engineering to the InSIM project.

If possible, we would also like to work with the DMSO effort, or the principal scientists involved in SimNet developments at the Institute for Defense Analysis (IDA) and the Rand Corporation. However, we have not yet made contact with these organizations about this proposal, and have not included any funding requests for them.

The total period for which funding is estimated at this time is from January 1 1997 to August 31 1999, or 32 months. Our primary milestones will be:

1 Jan 97: Start of Contract. Kickoff meeting. Objective refinement.

1 Mar 97: Initial staffing and training of student research assistants; resource inventory.

1 July 97: A functional specification for SimQL in the JFACC context.

1 Oct 97: An initial language and interface specification, matching JFACC standards.
Definition of some candidate simulations relevant to JFACC to be accessed
Option 1: Wrapping a DIS simulation to help planners learn from simulations.
Option 3: Specification of a user interface and SW linkage.

1 Jan 98: Demonstration of a freestanding situational simulation accessed via SimQL.
Results of paper analysis of planning scenarios with other defined simulations.
Revised SimQL language and interface document.

1 Apr 98: An implementation of SimQL within the KQML framework.
Option 1: Access to DIS data via the SimQL prototype.
Option 2: Design of a PC- API interface for SimQL/KQML.

1 Jul 98: An evaluation of SimQL effectiveness in simulation access.
Measurements to evaluate actual use of SimQL.
Interoperation of SimQL and SQL within a mediator.

1 Oct 98: Demonstration of an operational SimQL accessing a simulation.
Option 2: Demonstration of a PC-API interface for SimQL/KQML.

1 Jan 99: Specification of SimQL in a realistic JFACC interoperation scenario.

1 Apr 99: Demonstration of SimQL interoperating with other KQML services.

Option 2: Operational prototype for SimQL/KQML interoperation.

1 Jul 99: Demonstration of interoperation among multiple simulations and data accesses.

Option 3: Demonstration of remote and email user interface for remote customers.
31Aug 99: Final analysis and report, including technology transition.
Delivery of an operational prototype to DARPA.

The options provide considerable added value to this proposal, since they will permit demonstrations of the results in a realistic setting, and hence speak more directly to the people who can make adoption and technology transition decisions. Without such support by professional organizations, who operate outside of the university setting, our means of transition will be predominantly through publication and presentations.

V.b Deliverables.

Unless otherwise negotiated, Stanford University and LMITCO will provide:

1. A Specification of SimQL, a Simulation Query Language, which is designed to interoperate with the emerging KQML standard.

2. An evaluation and assessment of a SimQL prototype and its effectiveness in accessing and retrieving simulation results from several legacy simulation programs.

3. A sample of a mediator program that utilizes a simulation as well as a database to deliver up-to-date situational information to a customer application program. All these programs will execute over the Internet, using standard protocols as far as feasible.

4. A demonstration of SimQL in the JFACC setting.

5. Documentation that will allow other JFACC and DoD parties to develop wrappers that allow SimQL accessibility, and specification that will enable the development of engineered versions of a SimQL processor. Documentation will be delivered on-line and be incrementally accessible to validated Web sites.
If the options listed below are funded, the SiMQL demonstration will be able to use SimNET DIS standard data, become portable, and involve multiple remote data and simulation sources. In respect to the options, Stanford and LMITCO's role will be in orchestrating the work of the contractors performing the optional tasks, and interacting with relevant parties within the JFACC initiative, DMSO, and DARPA.

6. Option 1: Myriad Software will implement SimQL as a means of accessing results of simulated exercises and rehearsals. The results of Distributed Interactive Simulation (DIS) training sessions will provide flexible feedback to customer applications. Simple simulations will guide users through the steps required for JFACC planning tasks.

7. Option 2: Crystalliz Corporation: Application Program Interfaces using KQML to transport SimQL results and database tuples to multiple sites, specifically including to the customers' personal computers. Workflow concepts developed at Crystalliz can help in synchronizing SimQL queries at the right stage in planning and replanning.

8. Option 3: LMITCO and Crystalliz Corporation: Service and presentation functions to provide visual 2D and 3D display within remote popular application workstations, covering both past data from a database as well as results from SimQL accesses.

VI Capabilities of Offerors

In this proposal we combine considerable intellectual and implementation resources, all with a strong background in large-scale interoperating software systems, experience in simulations, and insights into military requirements.

VI.a Capabilities at Stanford.

The Computer Science Department at Stanford University is rated as a prime computer Science Research center in the world. In addition to its faculty, staff, and students it provides excellent technical resources for experimental and theoretical research. The department supports a world-class Computer Science library, including various computerized search services. The department operates an infrastructure network connecting mall and large computer systems for research, education, and administration to each other and the world-wide Internet.

Within the Computer Science Department, the Stanford Database Group runs a growing workstation cluster with approximately 20 workstations. The Facilities staff maintains an up-to-date software environment, including a Sybase DBMS, an Oracle DBMS, CORBA software, and plans for near-term installation of one or more Object-Oriented DBMS's.

VI.b Capabilities at LMITCO.

The Idaho National Engineering Laboratory (INEL), managed by LMITCO, has a great deal of expertise in the development of computer simulations and the development and maintenance of large scale, distributed software. LMITCO is a major research and development facility with over 200 staff members in the software field.

Early simulations at INEL mostly dealt with computer models of nuclear reactor operations and safety analysis for the Nuclear Regulatory Commission (NRC). More recently, LMITCO has been developing computer simulations for the Department of Energy's (DoE) Office of Environmental Management (EM). These simulations support various planning and decision making activities at both the laboratory and national levels, and employ concepts and technology applicable to planning simulations in JFACC.

INEL computer simulations allow DOE decision makers to rapidly create and execute new scenarios and compare the results of these scenarios with each other. These simulations are models of highly complex systems which are difficult to evaluate using conventional methods. By simulating these models, users see the impacts of their decisions over time, and can determine the relative worth of those decisions. Typical output parameters of interest are cost, risk, schedule, public perception, etc. When the output parameters are given relative weighting, the simulations can optimize the objective function to determine the scenario which best fits the user's current goals.

The INEL Integrated Waste Stream Model integrated all types of hazardous and radioactive waste streams into one system, and interactively tracked them from source to sink. Decision makers can locate potential bottlenecks and reduce their impact.

The National Transuranic Waste Model simulates the complete national system of radioactive waste, including storage, treatment, transportation, and disposal. The model is a discrete event simulation, focusing on logistics, such as shipment scheduling, resource limitations, and bottleneck minimization.

The National Decommissioning Decision Analysis System helps national DOE decision makers to determine the best approach to scheduling the decontamination and decommissioning of its almost 500 facilities. These models were created using commercial off-the-shelf (COTS) simulation environments. Recent modeling efforts have utilized state-of-the-art object-oriented software development environments, which increase flexibility in terms of domain implementation, and scalability.

The Idaho National Engineering Laboratory and its management (LMITCO) has been involved with a number of software systems engineering efforts of command and control systems for the Department of Defense. Following is a listing of some of these efforts.

1. Contingency TACS Automated Planning System (CTAPS) is a modernization effort undertaken by the U.S. Air Force (USAF). The LMITCO was contracted to develop a modernized prototype for a Modular Tactical Air Control Center (MTACC), including hardware and software integration involving Government Furnished Equipment (GFE) software packages, off-the-shelf software, and custom developed software. The LMITCO software integrates all of these packages, establishing an open architecture that accommodates additional or replacement capabilities.
2. The Electronic Combat System Integration (ECSI) for the Air Force Information Warfare Center (AFIWC) at Kelly Air Force Base, in San Antonio, Texas has reengineered several AFIWC applications to an object-oriented, layered architecture, implemented using the Ada programming language. The project continues to be a key project at the INEL in the area of software engineering. Client/server architectures as well as software modeling and simulating functions of a given aircraft within a war games scenario are current areas of emphasis.
3. The AdaSAGE Development Toolkit facilitates rapid, professional construction of software applications. AdaSAGE is based on the concepts of object-oriented design and software reuse. At the 1994 ObjectWorld Conference in San Francisco, AdaSAGE received recognition with the Type Commander Readiness Management System (TRMS), developed by the U.S. Navy. TRMS won the award for Best Object-Based Application Developed Using Non-Object Tools.
4. Submarine Equipment Performance Analyzer - developed for use in maintenance monitoring by U.S. Navy shipyard engineers and managers.
5. Cutter Scheduling Assistant Program - a graphical scheduling assistant system for use by the U.S. Coast Guard for fleet management.
6. The Automated Container Offering System - used for optimizing the booking of military cargo shipments via commercial carriers.
7. The Reserve Component Automation System (RCAS), a multimedia, secure communications system. .
8. The Granite Sentry Video Distribution System (VDS) provides multi-level secure video information integration and distribution to upgrade C2 capabilities at the North American Air Defense Command (NORAD) Cheyenne Mountain Complex.

VI.c Biographies

Gio Wiederhold is a Professor of Computer Science at Stanford University. He also holds courtesy appointments in Medicine and Electrical Engineering. He is active in the application and development of knowledge-based techniques to database management, information systems, and software construction. This research direction, initiated with the KBMS project at Stanford in 1977, has now become an active research field in its own right. Results derived from this work are found in general system design, medicine, information systems, and planning applications.

During 1991 to 1994 Gio was the Program manager for Knowledge-Based Systems at ARPA, initiating programs in Intelligent Integration of Information (I3) and in Persistent-Object Bases (POB). He also participated in establishing directions for a Software Foundations program, specifically in composition technology for large, multi-site systems. Gio is a member of the ACM publications board, an advisory member of the ARPA-Rome Laboratory Planning Initiative (ARPI), a member of the Military Operations Research Society (MORS), a director of the National Geographic Information systems effort, and a other initiatives. He has participated in the data system designs for the joint services DMSO effort. Wiederhold has published 4 books and more than 250 publications in computing and medicine.

Prof. Wiederhold received a degree in aeronautical engineering in Holland in 1957 and shortly after came to the United States. After gaining 16 years of industrial experience, he earned in 1976 a PhD in Medical Information Science from the University of California in San Francisco. He is a member of the AAAS, AAMI, AIAA, AAAI, and is a fellow of the ACMI, the IEEE, and the ACM.

Dr. Ted Linden has broad experience as the architect, designer, and manager of extensible software applications that exploit advanced techniques such as intelligent agents, knowledge-based planning, and component-based software. He is the founder of Myriad Software where he has had contracts from the ARPA TRP program and from Stanford University to develop tools for adaptive training simulations. While with Advanced Decision Systems, he was the principal investigator leading five ARPA contracts to develop knowledge-based planning technologies. He led the design of the AI planning and scheduling component of the Air Force's Advanced Planning System (APS). APS is the Air Force's theater-level planning support tool now operational as a component of CTAPS. Designed in 1988 and operational in 1992, APS was one of the earliest successful applications of AI planning technology to military mission planning. In earlier work at Xerox in Palo Alto, Dr. Linden managed development work on the Star Workstation and Network Services which pioneered modern graphic user interfaces and client-server networks. His publications proposed object-oriented programming and software reuse techniques long before these concepts became widely recognized. He has contributed to the architecture and design of over 20 commercial products and military applications.

Dr. Sankar Virdhagriswaran is the founder of, Crystaliz, Inc. He has been building KQML service modules for interoperation, and interacting with OMG to achieve convergence of CORBA with advanced KQML concepts. The existing KQML server prototype has been distributed to about 50 universities worldwide.

VII Estimated Budget:

This is an unofficial budget. We believe that it represents fairly the cost of the proposed work and will not differ significantly from the final budget to be with the proposal, if selected. We are open to negotiating changes in scope of work, timing of milestones, and corresponding changes in the costs budgeted.

	Jan-Aug97	Sep97-Aug98	Sep98-Aug99.	Total
Stanford				
Salaries and benefits	$108,746	212,134	213,322	$537,202
Travel (domestic)	$ 10,000	20,000	20,000	50,000
Computing, software	25,000	40,000	21,000	86,000
Equipment	20,000	20,000	5,000	45,000
LMITCO subcontract				
Labor	$ 300,000	450,000	450,000	1,200,000
Travel	17,000	25,000	25,000	67,000
Subtotal direct	589,492	979.268	953,644	2,018.136
Stanford indirect cost	169,270	295,403	953,644	753,596
(61% on selected items)	----------	---------	----------	-----------
Total without options	$ 758,762	1,275,671	1,242,567	$2,771,732
	=======	=======	=======	=========
Option 1: Myriad Software	(Access to DI-based simulations)			
Salaries and overhead	$ 60,000	75,000	60,000	$195,000
travel	4,000	6,000	6,000	16,000
Stanford indirect	15,250	0	0	15,250
total Option 1	$ 79,250	81,000	66,000	226,250
	=======	=======	======	=========
Option 2: Crystalliz inc	(PC SiMQL/KQML API)			
Salaries and overhead	$ 50,000	100, 000	50,000	$200,000
travel	3,000	4,000	4,000	
Stanford indirect	15,250	0	0	15,250
total Option 2	$ 68,250	104,000	54,000	225,250
	=======	=======	======	=========
Option 3: Crystalliz and LMITCO	(Presentation module)			
Crystalliz labor	$ 20,000	30,000	40,000	$ 90,000
LMITCO labor	20,000	30,000	40,000	90,000
travel	2,000	2,000	4,000	8,000
Stanford indirect	0	0	0	0
total Option 2	$ 42,000	62,000	84,000	188,000
	=======	=======	======	=========
Aggregate total with all Options	$ 948,262	1,521,671	1,445,567	3,411,232
	=======	=======	=======	=========

VIII Bibliography (not counted towards 15page limit)

Corella, Francisco, S.J. Kaplan, G. Wiederhold, and L. Yesil: "Cooperative Responses to Boolean Queries"; *Proc. IEEE First Data Engineering Conference*, April 1984, pages 77-93.

INEL1: "Electronic Combat Model Reengineering", AJPO-MON-95-01, March 1995

INEL2: "Software Maintainability Metrics Models in Practice", Crosstalk, November/December 1995, Volume 8 Number 11.

INEL3: "Software Re-Engineering in Ada: A Practical Approach", TRI-Ada'94, TRI-Ada Proceedings, ACM ISBN: 0-89791-666-2

INEL4: "Ada Electronic Combat Modeling", OOPSLA'93, OOPSLA'93 Proceedings, ACM ISBN: 089791-661-1

Law, Kincho H., Gio Wiederhold, Thierry Barsalou, Niki Siambela, Walter Sujansky, David Zingmond, and Harvinder Singh: ``Architecture for Managing Design Objects in a Sharable Relational Framework";*International Journal of Systems Automation: Research and Applications (SARA)*, Vol. 1, No. 1, pages 47-65, 1991.

Lee, Byung Suk and Gio Wiederhold: "Efficiently Instantiating View-objects from Remote Relational Databases"; *The VLDB Journal*, Vol.3 No.3, July 1994, pages 289-323.

Linden, T.A. 1994a. "A New Generation of Adaptive Education and Training Software"; White Paper for NIST ATP Program, August, 1994.

Linden, T.A. 1994b. "Tools for a New Generation of Educational Software"; INQUIRY: Critical Thinking Across the Disciplines, Vol 13, 1-2, Feb./Mar. 1994.

Linden, T.A. & D. Gaw 1992, "JIGSAW: Preference-directed, Co-operative Scheduling," AAAI Spring Symposium: Practical Approaches to Scheduling and Planning, March 1992. (JIGSAW generalizes the AI scheduling techniques implemented in the Air Force's Advanced Planning System--which is operational as a component of CTAPS.)

Linden, T.A. 1991. "Representing Software Designs as Partially Developed Plans," Automating Software Design. AAAI/MIT Press, Menlo Park, CA. pp. 603-625.

Linden, T.A. 1990a. "A Meta-level Software Development Model that Supports V&V for AI Software," Expert Systems with Applications. Vol. 1, pp. 271-279.

Linden, T.A. 1990b "Transformational Synthesis: An Approach to Large-Scale Planning Applications," DARPA Workshop on Innovative Approaches to Planning, Scheduling, and Control., Morgan Kaufman Publ., San Mateo, CA, Nov. 1990.

Linden, T.A. 1989. "Planning by Transformational Synthesis," IEEE Expert, 4,2 Summer 1989.

McCall, Gene (chair): *New World Vistas, Air and Space Power for the 21st Century;* Air Force Scientific Advisory Board, April 1996.

Shuey, Richard and Gio Wiederhold: "Data Engineering and Information Systems"; *IEEE Computer*, Vol.19 No.1, January 1986, pages 18-30

Spooner, David, Arthur M. Keller, Gio Wiederhold, John Salasin, and Deborah Heystek: "Framework for the Security Component of an Ada DBMS"; *Proc. Twelfth VLDB Conference*, Kyoto, Japan, August 1986, Morgan Kaufmann pubs., pages 347 to 354;

also in *Proceedings of the Database Security Workshop*, National Computer Security Center, Fort Meade MD, June 1986.

Wiederhold, Gio and Gary Breitbard: "A Method for Increasing the Modularity of Large Systems"; *IEEE Computer*, March-April,1970, Vol.3 No.2.

Wiederhold, Gio: "A Spreadsheet Metaphor for Management of Distributed Data"; report to VisiCorp, San Jose, November 1983.

Wiederhold, Gio: "Views, Objects, and Databases"; IEEE Computer, Vol.19 No.12, December 1986, pages 37-44. Reprinted in Dittrich, Dayal, Buchmann (eds.): *Object-oriented Database Systems*, Springer Verlag, 1988.

Wiederhold, Gio, Peter Wegner, and Stefano Ceri: "Towards Megaprogramming"; Comm.ACM, Vol.35 No.11, November 1992, pages 89-99.

Wiederhold, Gio: "The Roles of Artificial Intelligence in Information Systems*"; Journal of Intelligent Information Systems*, Vol.11, No.1, 1992, pages 35-56.

Wiederhold, Gio: "The Role of Government in Standards"; *ACM Standards Bearer*, December 1993, pages 2-6.

Wiederhold, Gio, Sushil Jajodia, and Witold Litwin: "Integrating Temporal Data in a Heterogeneous Environment"; in Tansel, Clifford, Gadia, Jajodia, Segiv, Snodgrass: *Temporal Databases, Theory, Design and Implementation*; Benjamin Cummins Publishing, 1993, pages 563-579.

Wiederhold, Gio: "Intelligent Integration of Information*"; Proc. ACM-SIGMOD 93*, ACM, Washington DC, May 1993, pages 434-437.

Wiederhold, Gio: "Mediators in the Architecture of Future Information Systems"; *IEEE Computer*, March 1992, pages 38-49. Reprinted in A.R.Hurson, M.W.Bright, S.H.Pakzad (ed.s): Multi-database Systems: An Advanced Solution for Global Information Sharing; IEEE Press, 1993.

Wiederhold, Gio: "Intelligent Integration in Simulation"; MORS Mini-symposium, Fairfax VA, Military Operations Research Society, Alexandria VA, November 1993.

Wiederhold, Gio: "Digital Libraries, Value, and Productivity"; *Communications of the ACM*, Vol.38 No.4, April 1995, pages 85-96.

Winslett, Marianne, Keith Hall, David Knapp, and Gio Wiederhold: "Use of Change Coordination in an Information Processing System"; *Proceedings of the ACM Design Automation Conference*, July 1989.

----------------- o --- o ------------------------