

# Information Engineering

**David N. Hunt,  
Thrust Area Leader**

This is an informal report intended primarily for internal or limited external distribution. The opinions and conclusions stated are those of the author and may or may not be those of the Laboratory.

Work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

**February 1997**

### Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This report has been reproduced  
directly from the best available copy.

Available to DOE and DOE contractors from the  
Office of Scientific and Technical Information  
P.O. Box 62, Oak Ridge, TN 37831  
Prices available from (615) 576-8401, FTS 626-8401

Available to the public from the  
National Technical Information Service  
U.S. Department of Commerce  
5285 Port Royal Rd.,  
Springfield, VA 22161

THRUST AREA REPORT • UCRL-ID-125477

# Information Engineering

**David N. Hunt,  
Thrust Area Leader**

Reprinted from Engineering Research, Development and Technology  
FY 96 UCRL 53868-96

**February 1997**



David N. Hunt, Thrust Area Leader

The Information Engineering thrust area develops information technology to support the programmatic needs of Lawrence Livermore National Laboratory's (LLNL) Engineering Directorate. The mission of the thrust area is to maximize the information bandwidth between the system and its users.

A major requirement to fulfill this mission is "intelligent" access to the information served by the system, calling for advances in compute platforms, networks, user interfaces and applications.

In particular, the development of systems that can act as intelligent agents for their users is critical to maximizing the benefits of this technology.

Our approach has two major components. First, we are matching the needs of the programs at LLNL with developments in leading edge information technology. Our efforts include work in high-speed computing and networking, data acquisition and control, and the collection, storage, collation, and analysis of massive amounts of scientific data. Progress in this area requires collaboration with programs and other disciplines at LLNL, as well as with the industrial and academic sectors.

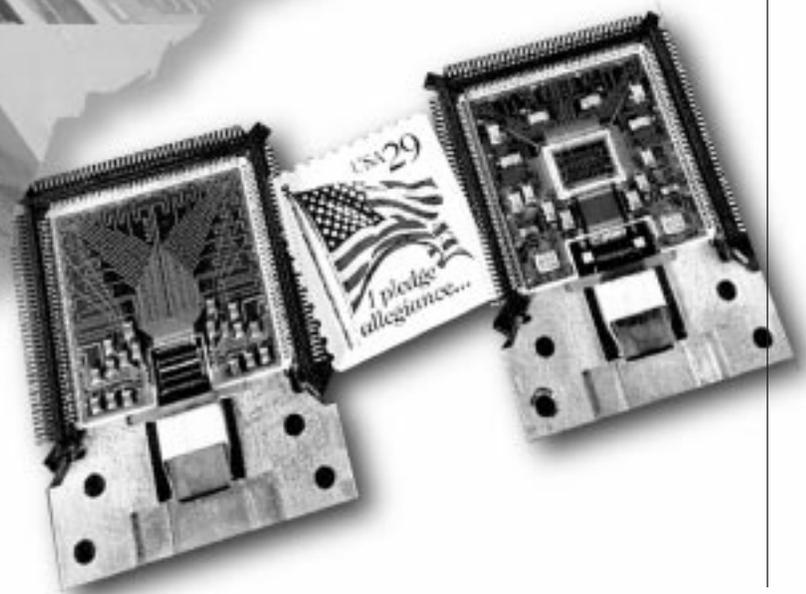
The second major component of the thrust area effort is focused on leveraging information technology by applying it to the LLNL infrastructure. Much of the technology used to meet programmatic needs can be used to increase the efficiency of day-to-day operations. In these days of fiscal consciousness, proper application of information technology can conserve both time and money.

As we close FY-96, we see benefits of the work we have pursued over the past two years: we have produced patentable work in subcarrier multiplexed networks, and LLNL's Procurement Department is embracing the Zephyr concept in rapid prototype procurement and is investigating its broader applications.

The thrust area activities for FY-96 are described in the following five papers: (1) Three-Dimensional Object Creation, Manipulation, and Transport; (2) Zephyr: A Secure Internet-Based Process to Streamline Engineering Procurements; (3) Subcarrier Multiplexing: Optical Network Demonstrations; (4) Parallel Optical Interconnect Technology Demonstrations; and (5) Intelligent Automation Architecture.

# Information Engineering

8



## 8. Information Engineering

### **Overview**

*David N. Hunt, Thrust Area Leader*

### **Three-Dimensional Object Creation, Manipulation, and Transport**

*Scott D. Nelson.....8-1*

### **Zephyr: A Secure Internet-Based Process to Streamline Engineering Procurements**

*Cecil W. Jordan, Rudy E. Cavitt, William A. Niven, Nikola Mitschkowetz, Timothy M. Sharick,  
Sandra S. Taylor, Donald L. Vickers, Fred E. Warren, and Richard L. Weaver.....8-7*

### **Subcarrier Multiplexing: Optical Network Demonstrations**

*Paul D. Sargis and Bruce D. Henderer.....8-13*

### **Parallel Optical Interconnect Technology Demonstrations**

*Robert F. Hills, Ronald E. Haigh, and James M. Zumstein.....8-17*

### **Intelligent Automation Architecture**

*Scott A. Couture .....8-21*

# **T**hree-Dimensional Object Creation, Manipulation, and Transport

Scott D. Nelson  
*Defense Sciences Engineering Division*  
*Electronics Engineering*

The effective creation, manipulation, communication, and transport of 3-D objects has historically been a problem because of competing or specialized data formats, the requirements for sophisticated tools on the end user's part, or a steep hardware requirement for the display and transport of objects. This effort combines several parallel activities into a unified medium, taking advantage of recent economies of scale in the commercial sector that allow for the efficient distribution of tools to a broad class of users.

## Introduction

There are several kinds of 3-D data types used to convey information about solids. The newer versions of these are specific to various programs, and at Lawrence Livermore National Laboratory (LLNL) they are used in a series of computational modeling codes. The key to the approach presented here is the creation of a standard 3-D infrastructure to aid in the creation, manipulation, and transport of 3-D data objects.

The exchange and definition mechanism was dealt with by creating the new high-level MIME format for 3-D data types. The software portion of the transport issue is handled as part of this new primary MIME type. The methodology used for the creation aspect is different from what others are working on.<sup>1</sup> Our effort has the advantage of using general images taken from positions surrounding the body of interest, with little information or control used during the acquisition phase. The difference is that more information is known about the object and less about the camera positions.

By contrast, this approach is different from simplistic commercial systems<sup>2</sup> that require the user to specify object geometries in coordinate planes. For complex objects such as missiles and helicopters, this is beyond the scope of such software.

## Progress

### **New Primary MIME Type Registered**

The Glossary in Reference 3 contains definitions and abbreviations.

As part of this activity, the MIME type for the general class of problems was created and is an Internet standard as specified in Reference 4. This allows the easy conversion of various 3-D types by grouping the data types in a common framework of the model/\*MIME type. This type presently includes the VRML 3-D data type, the IGES data type, and a generic mesh data type.

The principal focus of this effort was to create a standard "umbrella" mechanism that would allow other 3-D/4-D data types to be encapsulated underneath. As a consequence, the base set of parameters for these data types was also created and includes requirements on scale, interaction, abstraction, and the flow of time.

A tree diagram and outline for the model type is as follows:

```

model/
  vrml
  iges
  mesh

```

A model primary MIME type is an electronically exchangeable behavioral or physical representation within a given domain. Each sub-type in the model structure has unique features, just as does each sub-type in the other primary types. The important fact is that these various sub-types can be converted between each other with less loss of information than by converting to other primary types. This fact groups these sub-types together into the model primary type.

All of the expected sub-types have several features in common that are collectively unique to this primary type. To loosely summarize: models are multi-dimensional structures composed of one or more objects. If there are multiple objects, then one object defines the arrangement, setting, and relationship of the others. These objects all have calibrated coordinate systems, but these systems need not be in the same units, nor need they have the same dimensionality. In particular:

1. They have three or more dimensions that are bases of the system and form an orthogonal system (any orthogonal system is sufficient). This system is specifically defined in terms of an orthogonal set of basis functions (for a subspace of the L2 function space) over a coordinate system of dimension three or more. Note that this does not preclude regular skewed systems, elliptical coordinates, or different vector spaces.
2. They contain a structural relationship between model elements.
3. They have scaling or calibration factors that are related to physical units such as force, momentum, time, velocity, acceleration, and size. Thus, an IGES file will specify a building of non-arbitrary size; computational meshes and VRML models will have real spatial/temporal units. This allows for differing elements to be combined non-arbitrarily.
4. The models can be single objects or composed of a collection of objects. These normally independent objects are arranged in a master/slave scenario so that one object acts as the reference, or primary object, and defines how the other objects interrelate and behave. This allows for the creation of mathematical, physical, economic, or behavioral models that typically are composed of different elements.

The key is in the description: these types describe how something “behaves,” in contrast to typical data types which describe how something “is.” The inclusion of this “collective” system works in a way similar to the e-mail system’s

multi-part/related type which defines the actions of the individual parts. Further specification of the model/\*subtypes using these properties is left to the sub-type authors.

With these assumptions:

1. The default dimensionality will be spatial and temporal (but any are allowed).
2. Models will contain underlying structure which may or may not be immediately available to the user (for example, fluid dynamics vector fields, electromagnetic propagation, interrelated IGES dimensional specifiers, VRML materials and operators).
3. Basis set conversion between model domains is lossless. The interpretation of the data may change but the specification will not, that is, convert the model of the U.S. Gross Domestic Product into a VRML model and navigate it to explore the variances and interrelationships. The model has many dimensions but also “passages” and “corridors” linking different parts of it. A similar situation is true for meshes and CAD files. The key is identifying the basis set conversion that makes sense.
4. Models are grouped to assure less loss of information between the model sub-types than to sub-types of other primary types.

### Data Conversion

The conversion part of this activity grew out of standard activities and included the SBIR<sup>5</sup> made by NIST on display and conversion programs. Additional efforts were made by the community at-large to develop the appropriate suite of conversion tools for VRML, and these are mostly in place and operational. At the beginning of the project, quite a bit of effort was put into a set of functions that operated with the LLNL mesh format, but since then the broader class of STEP<sup>6</sup> mesh descriptions is going through the International Standards Organization (ISO) process and is expected to be the more generic alternative. Conversion functions for this format will be created as part of future program-specific activities as needed.

Examples of the existing formats include a mesh model of a shipping container,<sup>7</sup> a VRML model of an aircraft carrier, and an IGES model of a bone (Fig.1).

### Object Creation Definition and Approach

This problem involves locating the coordinates of N points in 3-space contained in a collection of M 2-D images taken from different perspectives. Not

all points must be visible in each of the images; however, each point must be visible in at least two images. It is also assumed that the field of view (FOV) of each image and the point correspondence between images are known.

To make acquiring the images in the field as simple as possible and to reduce the effects of measurement error, it is desirable to solve for the  $N$  point locations with the minimum required information about the camera positions and orientations, while using the measured data to gain as much information as possible. The orientations and locations of all cameras are solved for whose images contain these points.

Next, all uncomputed points in each of the cameras' images are computed. This process is then repeated until all camera information is known and the locations of all points are known. Note that

depending on how many images a point appears in, the location of a point may be computed many times. These locations are averaged to provide the ultimate estimate of the point location.

### Problem Solution

We define the geometry for an image as shown in **Fig. 2**. In local coordinates,  $x', y', z'$ , the normal vector of the image plane  $\hat{n}' = \hat{z}'$  originates at the center of the image plane. Each image is  $w \times h$  pixels in size, with known angular field of view in both the  $x$  and  $y$  directions. The coordinate location  $(x_{ij}, y_{ij})$ , of the  $j^{\text{th}}$  point in the  $i^{\text{th}}$  image, along with point correspondence between points in the images, are given.

From these definitions, the direction vector from the  $i^{\text{th}}$  camera location to the  $j^{\text{th}}$  point,  $\hat{u}'_{ij}$  follows from

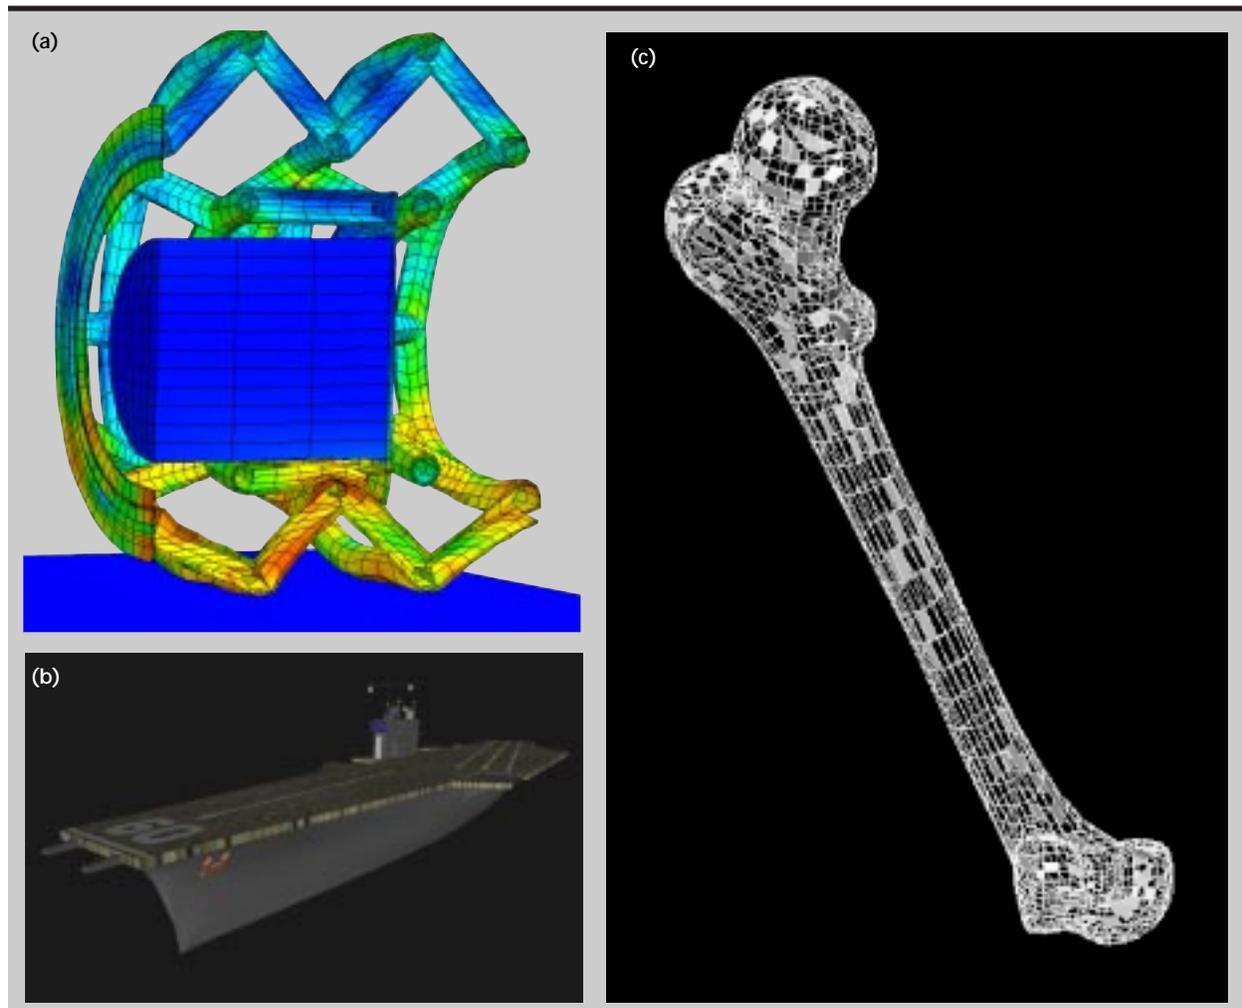


Figure 1. Mechanical model of (a) a shipping container; (b) a VRML model of an aircraft carrier; and (c) an IGES model of a bone.

$$\theta_{xz} = \text{sign}(x_{ij}) \sqrt{\left(\frac{FOV_x}{\omega} x_{ij}\right)^2 + \left(\frac{FOV_y}{h} y_{ij}\right)^2},$$

$$\theta_{xy} = \begin{cases} \tan^{-1}\left(\frac{y_{ij}}{x_{ij}}\right) & x_{ij} \neq 0 \\ 90^\circ & x_{ij} = 0 \end{cases}$$

giving

$$\hat{u}'_{ij} = \begin{bmatrix} \cos(\theta_{xy}) & \sin(\theta_{xz}) \\ \sin(\theta_{xy}) & \sin(\theta_{xz}) \\ \cos(\theta_{xz}) \end{bmatrix}.$$

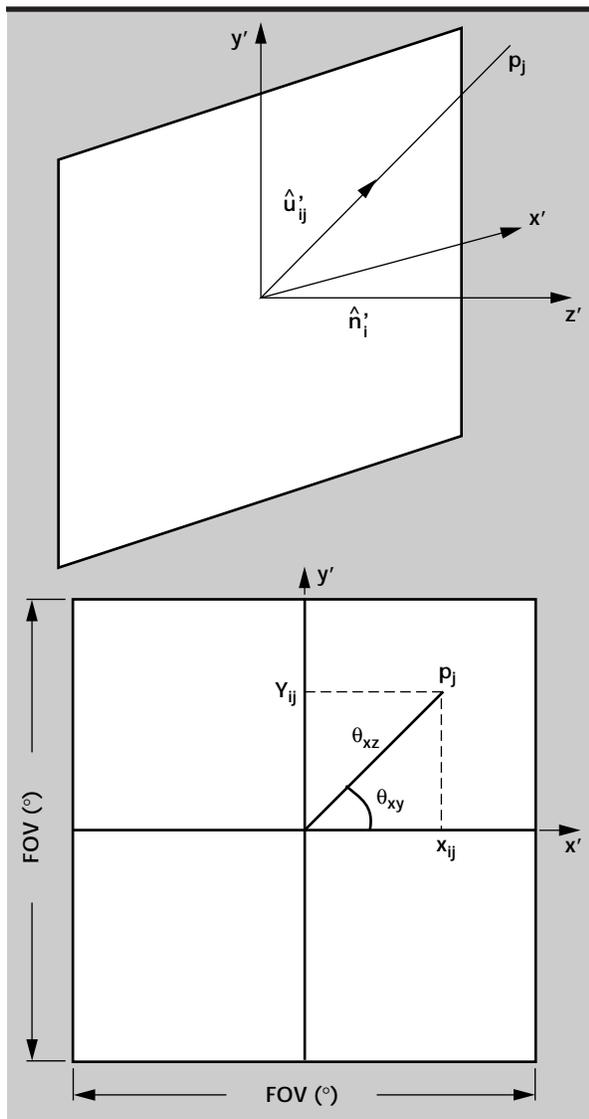


Figure 2. Point location defined in image geometry and in 3-D space relative to camera location and direction. The direction vector of the camera is normal to the image plane and originates at the center of the image. The image provides angular information to the point relative to the direction vector of the camera.

In global coordinates, the location of the direction vector to the  $j^{th}$  point from the  $i^{th}$  camera is

$$\hat{u}_{ij} = \begin{bmatrix} \cos(\theta_{xy}) [n_{xi} \cos(\theta_{xz}) + n_{zi} \sin(\theta_{xz})] - n_{yi} \sin(\theta_{xy}) \\ \sin(\theta_{xy}) [n_{xi} \cos(\theta_{xz}) + n_{zi} \sin(\theta_{xz})] + n_{yi} \cos(\theta_{xy}) \\ -n_{xi} \sin(\theta_{xz}) + n_{zi} \cos(\theta_{xz}) \end{bmatrix}$$

where

$$n_i = [n_{xi}, n_{yi}, n_{zi}]^T.$$

Defining the translation vector,  $\bar{t}_p$ , denoting the position of the  $i^{th}$  camera we have the relationship between corresponding points contained in two images as depicted in Fig. 3. The two angles  $\theta_1$  and  $\theta_2$  follow straightforwardly from the fact that the direction vectors from the cameras to the point and the location vectors are known. The distances to the point from the respective cameras then follow from the Law of Sines. Each point location is computed in this manner for each corresponding stereo image pair.

We now use the information about the location of the points to compute the unknown position and direction information for all images which contain these points. For a given point in an image, the relationship between the point location and the camera position and direction is

$$\bar{p} = \bar{t} + R\alpha\hat{u}',$$

where R is defined in terms of the three Euler rotation matrices

$$R = R_{xy}R_{yz}R_{zx}$$

$$R_{xy} = \begin{bmatrix} \cos(\theta) & -\sin(\theta) & 0 \\ \sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 1 \end{bmatrix},$$

$$R_{yz} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\rho) & -\sin(\rho) \\ 0 & \sin(\rho) & \cos(\rho) \end{bmatrix},$$

$$R_{zx} = \begin{bmatrix} \cos(\phi) & 0 & -\sin(\phi) \\ 0 & 1 & 0 \\ \sin(\phi) & 0 & \cos(\phi) \end{bmatrix}.$$

With R having three degrees of freedom, the translation vectors being unknown, and the distance from the image plane to the points being unknown, this vector equation contains seven unknowns. If we take two additional points within the image such

that the three points are not collinear, then we obtain three linearly independent vector equations in 3-space with nine unknowns

$$\begin{aligned}\bar{p}_1 &= \bar{t} + R\alpha_1\hat{u}'_1, \\ \bar{p}_2 &= \bar{t} + R\alpha_2\hat{u}'_2, \\ \bar{p}_3 &= \bar{t} + R\alpha_3\hat{u}'_3.\end{aligned}$$

To solve the above equations, we first subtract pairs of equations to remove dependency upon the translation vector  $\bar{t}$ . Using the fact that magnitude is invariant to rotation, we can remove the dependency of R, giving:

$$\begin{aligned}|\bar{p}_2 - \bar{p}_1|^2 &= (\alpha_2\hat{u}'_2 - \alpha_1\hat{u}'_1)^T (\alpha_2\hat{u}'_2 - \alpha_1\hat{u}'_1), \\ |\bar{p}_2 - \bar{p}_3|^2 &= (\alpha_2\hat{u}'_2 - \alpha_3\hat{u}'_3)^T (\alpha_2\hat{u}'_2 - \alpha_3\hat{u}'_3), \\ |\bar{p}_3 - \bar{p}_1|^2 &= (\alpha_3\hat{u}'_3 - \alpha_1\hat{u}'_1)^T (\alpha_3\hat{u}'_3 - \alpha_1\hat{u}'_1).\end{aligned}$$

These three equations are then solved numerically for  $\alpha_1$ ,  $\alpha_2$ , and  $\alpha_3$ .

At this point, we need to solve for the rotation and translation unknowns. We assume the two vector sets  $\{\bar{p}_1, \bar{p}_2, \bar{p}_3\}$  and  $\{\alpha_1\hat{u}'_1, \alpha_2\hat{u}'_2, \alpha_3\hat{u}'_3\}$  are linearly independent and not orthogonal. This assumption is reasonable, given the geometry and the way the points were selected. Noting that both multiplication by an orthogonal matrix and translation are angle-preserving transformations and that the two sets of vectors form congruent triangles, we proceed, letting  $\bar{q}_1 = \bar{p}_2 - \bar{p}_1$ ,  $\bar{q}_2 = \bar{p}_3 - \bar{p}_1$ ,  $\bar{q}_3 = \bar{q}_1 \times \bar{q}_2$ , and  $\hat{q}_4 = \bar{q}_3 \times \bar{q}_1$ . Similarly, we create  $r_1, r_2, r_3, r_4$  from the vectors  $\alpha_1\hat{u}'_1, \alpha_2\hat{u}'_2, \alpha_3\hat{u}'_3$  respectively. Normalizing, we now complete the solution for the rotation and translation of the camera by solving

$$R = [\hat{q}_1 \quad \hat{q}_4 \quad \hat{q}_3] \begin{bmatrix} \hat{r}_1^T \\ \hat{r}_4^T \\ \hat{r}_3^T \end{bmatrix},$$

and

$$\begin{aligned}\bar{t} &= \\ \frac{1}{3} &(\bar{p}_1 + \bar{p}_2 + \bar{p}_3 - R(\alpha_1\hat{u}'_1 + \alpha_2\hat{u}'_2 + \alpha_3\hat{u}'_3)).\end{aligned}$$

The algorithm for determining the collection of M points is as follows:

**Step 1.** For each image with known position and direction taken two at a time, compute the location of each point in the images. Iterate through all possible combinations of pairs of images.

**Step 2.** Compute the average value for each computed point.

**Step 3.** For each image without known location or rotation and at least three common known points, compute R and t for these images. Repeat for each group of three points in each image.

**Step 4.** Compute average value for each image of the direction vector and translation vectors computed in Step 3.

**Step 5.** Repeat Steps 1 through 4 until all points in all images are computed.

This solution is being applied to ideal and realistic 3-D objects with compound curved surfaces. The key is that information about the camera locations is limited or unknown thus allowing the surveyor more latitude in acquiring the data. Also note that this is applicable post-fact from a series of images, assuming that they meet certain criteria.

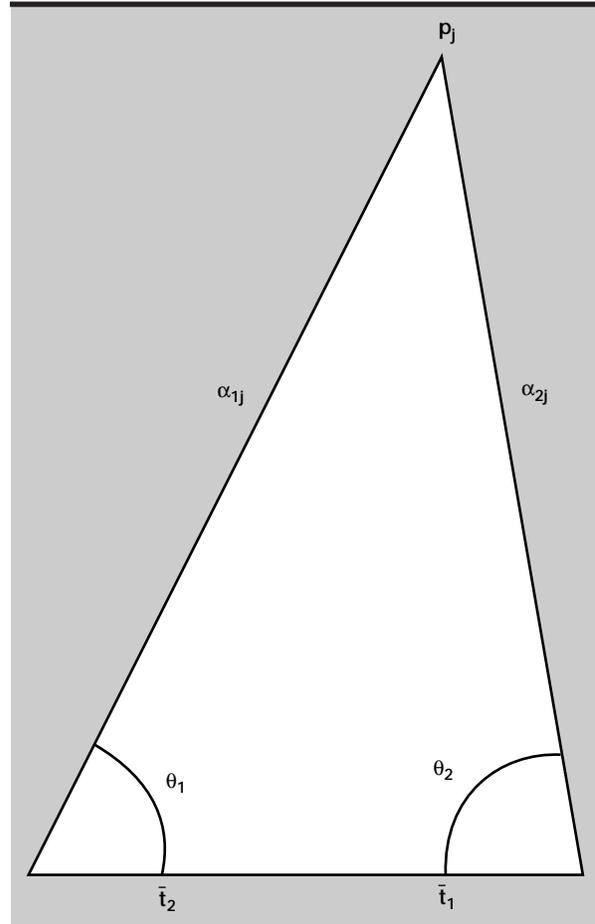


Figure 3. Stereo image geometry for solving for distance to point camera locations.

## Acknowledgments

Thanks go to J. E. Mast and S. Lehman for their work in the solution of the 3-D vector encoding method. Additional thanks go to M. Henderson, of the Naval Air Warfare Center at China Lake, Calif., for his efforts in assisting in the initial verification of this system.

## References

1. Lu, S-Y., J. E. Hernandez, and R. K. Johnson (1994), *Real-time Computer Vision*, Lawrence Livermore National Laboratory, Livermore, Calif., (UCRL-53868-94).
2. 3D Construction Company, (<http://www.3dconstruction.com/>).
3. Nelson, S. D., J. Kallman, and E. Brugger (1995), *Design and Implementation of 3D MIME Data Access and Exchange*, Lawrence Livermore National Laboratory, Livermore, Calif., (<http://www-dsed.llnl.gov/documents/vis/sdnieta95.html>), (UCRL-53868-95).
4. Nelson, S. D., M. Mitra, and C. Parks (1996), *The Model Primary Content Type for Multipurpose Internet Mail Extensions*, Lawrence Livermore National Laboratory, Livermore, Calif., (UCRL-ID-123343).
5. National Institute of Standards and Technology (1995, 1996), *Small Business Innovation Research (SBIR) for IGES Figure Viewer*, (<http://www.eeel.nist.gov/iges/viewerSpec.html>).
6. International Standards Organization (1996), *Integrated Application Resources: Finite Element Analysis*, (ISO 10303-104).
7. Hoover, C. (1994), *Shipping Containers for Hazardous Materials—DYNA3D Crash Simulation*, Lawrence Livermore National Laboratory, Livermore, Calif., ([http://www.llnl.gov/liv\\_comp/meiko/apps/dyna3d/dyna3d.html](http://www.llnl.gov/liv_comp/meiko/apps/dyna3d/dyna3d.html)), (UCRL-MI 118882). 

# Zephyr: A Secure Internet-Based Process to Streamline Engineering Procurements

Cecil W. Jordan, Rudy E. Cavitt, and William A. Niven  
*Laser Engineering Division*  
*Electronics Engineering*

Nikola Mitschkowetz  
*New Technologies Engineering Division*  
*Mechanical Engineering*

Timothy M. Sharick and Sandra S. Taylor  
*Administrative Information Systems*

Donald L. Vickers and Fred E. Warren  
*Computations Organization*

Richard L. Weaver  
*Distributed Procurement Program*  
*Procurement and Materiel Department*

Lawrence Livermore National Laboratory (LLNL) is piloting an Internet-based paperless process called "Zephyr" to streamline engineering procurements. Programs at LLNL are benefiting from the efficiencies introduced since implementing Zephyr's engineering and commerce on the Internet. Major benefits include reducing procurement time, speeding the engineering development cycle, facilitating industrial collaboration, and reducing overall costs.

## Introduction

Zephyr is being developed by LLNL's Engineering Directorate in partnership with its departments of Procurement, Administrative Information Systems, and Computations. Procurements that use Zephyr can move 60 to 90% faster from engineering design release, through contract award, to product delivery. Zephyr links people in LLNL's Engineering, Procurement, and Finance Organizations to pre-qualified small and medium enterprises (SME's) throughout the nation in a practical, simple, secure way. Procurements that used to take 30 to 60 days are now being done in 2 to 3 days, from initial request to delivery.

All the process steps are conducted by a World-Wide Web browser (Netscape or Internet Explorer) and e-mail. These steps are secure, that is, encrypted and password-protected. Zephyr's Web-based workflow process moves projects away from the sluggish paperwork path. LLNL's programs benefit

from a compressed time-frame for engineering deliverables by ensuring valid designs, minimizing schedule impact, and reaching project goals sooner.

## Concept

Zephyr is part of a system concept called CERPS, for Concurrent Engineering Rapid Prototyping System<sup>1,2</sup> (**Fig. 1**). CERPS integrates LLNL's engineering and business systems and links them to multiple development partners to speed all aspects of the engineering development cycle (design, procurement, fabrication, testing, and evaluation) independent of geographic location. Zephyr specifically focuses on speeding engineering procurements.

The Zephyr concept builds upon the ready availability of computer networks and new browser software technology, the Internet and the World-Wide Web. Engineering's computer-aided design systems are integrated with institutional business systems that support electronic commerce, linking procurement and

finance, commercial vendors, and end-users. Zephyr's systems integration and workflow facilitate a very rapid procurement and delivery cycle that meets all DOE purchasing requirements. With centralized, sharable data sources, and Zephyr's e-mail workflow process, people in LLNL's programmatic, engineering, procurement and finance organizations and our development partners are assured of accurate and up-to-date drawings, manufacturing models, and financial information.

Zephyr's workflow process (Fig. 2) is initiated upon a user request to procure engineering goods or services once a design package is approved for engineering release. Soon after engineering design release, a buyer prepares the electronic request for quotation (RFQ) package. Then Zephyr's workflow process e-mails a vendor announcement indicating the presence of the RFQ in a secure section of the Zephyr World-Wide Web "home page." After the announcement, solicitations for bid-quote, award,

technical data exchange, payment, delivery tracking, and record-keeping are provided by Zephyr. At all transactional levels, the benefits of compressed cycle time, ease of use, and secure business information are realized.

RSA encryption through the use of Netscape's Commerce Server allows secure exchange of engineering drawings, models, specifications, project planning information, bid-quote procurement, and corporate credit card information. Computer databases provide the record-keeping and audit trails required by LLNL programs, industrial partners, and DOE.

The Zephyr public home page is at <http://zephyr.llnl.gov/>.

### Progress

One of Zephyr's early project milestones was the first meeting of LLNL's Engineering and Commerce on the Internet (ECI) Working Group. This working

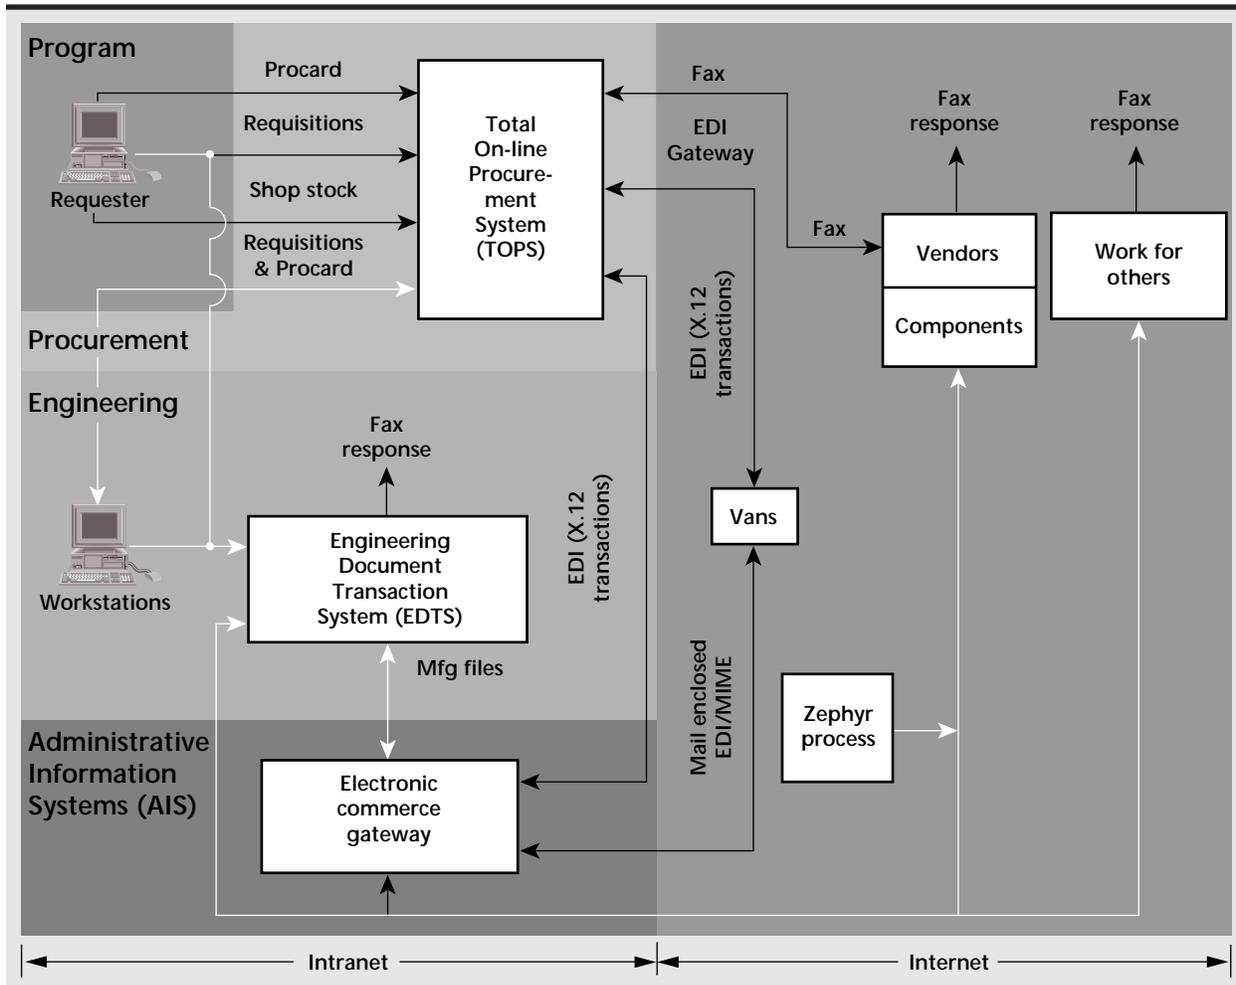


Figure 1. Concurrent Engineering Rapid Prototyping System (CERPS). CERPS integrates LLNL's Engineering Document Transaction System (EDTS) and business systems Total On-line Procurement Systems (TOPS) and Electronic commerce gateway, linking them to multiple development partners to speed all aspects of the engineering development cycle.

group consisted of representatives from Procurement, Engineering, Computations, Administrative Information Systems, and the Technical Information Department.<sup>3</sup> We reached an agreement to work together to promote Internet-based approaches to lower LLNL's cost of doing business. Corporate-wide "working group" agreements are critical to implementing institutional cultural change. Without such agreements, many excellent technical solutions lose out to "time honored" techniques.<sup>4</sup>

The Zephyr project has received three years of funding through DOE. It is built on a decade of work funded by the Department of Defense (DoD). Projects that have contributed to its success include Lawrence Berkeley National Laboratory's Web-based computer-aided design system, the U.S. Air Force's Computer-Aided Acquisition and Logistics Support (CALs) Program, the Air Force's Government Acquisition Through Electronic Commerce (GATEC) System, and the Department of Commerce (DOC) Technology Reinvestment Program's CommerceNet.

Zephyr has been in constant operation since its inception as a prototype in March, 1995, providing immediate access to a wide variety of vendors in many locations across the country.

**Metrics**

In fiscal years 1995 and 1996, LLNL's engineering programs saved time dramatically in a series of prototype fabrication tests. For example, Internet-based procurement of an assembly fixture dropped to a remarkably low 5.5 days, from a traditional (paper-based) cycle of 56 days, representing a 90% reduction. In other tests, we purchased printed circuit boards in less than two days as opposed to our customary 20 to 25 days. We conducted these tests for procurements of less than \$5,000. LLNL's corporate purchasing card, ProCard, can complete the transaction with electronic payment to the vendor in only one or two days, compared with the traditional one or two months.

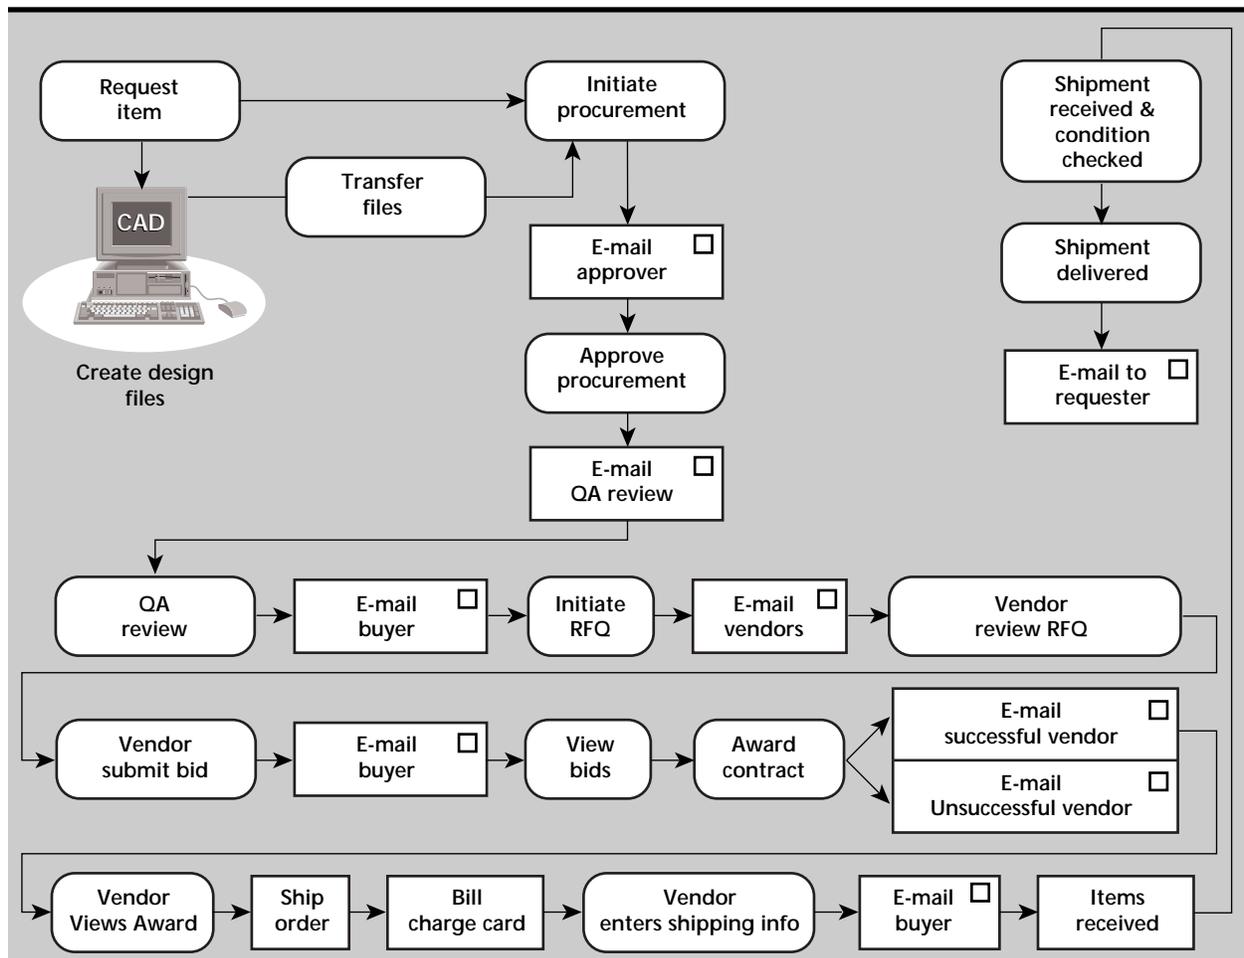


Figure 2. Zephyr's e-mail driven Web-workflow process to speed engineering procurements by communicating with users.

In a recent series of pilot procurements, Zephyr's procurement officer purchased non-standard optical components using Zephyr's paperless bid-quote process and ProCard. The bid package, RFQ, and award-process cycle took less than five days, as opposed to the traditional 30 to 60 days in a paper-based process. Actual procurement processing time was less than 15 minutes, reduced from the typical five to seven days.

Because of ProCard's prompt payment, our procurement officer negotiated "same as cash" discounts amounting to \$120,000 in just two months. Zephyr's workflow process allows rapid accrual of such savings.

Zephyr's workflow process has been successfully used to procure engineering goods and services for several projects and programs at LLNL. Our early testing involved the production of a part designed by EG&G in Nevada, giving us a chance to verify that our rapid prototyping process works in collaboration with another DOE site. Zephyr eliminated paper from every step of the design and procurement process while using engineering models to directly manufacture parts.

### Future Work

LLNL will continue to improve its internal business and engineering practices. Procurement and Administrative Information Systems will simplify both administrative and technical support systems to gain increased productivity while reducing overall

costs. Zephyr's pilot engineering procurement phase will be expanded to explore synergistic ways to complement ongoing programs like the National Ignition Facility (NIF) and Nuclear Weapons Stockpile Stewardship.

LLNL's Zephyr experience is being applied to two far-reaching external initiatives.

In the first initiative, the DOE Nuclear Weapons Complex (NWC), with LLNL participation, is developing new strategies for operations that automate traditional engineering and business processes, while significantly reducing cycle time and costs within. Several DOE projects are tasked with implementing these strategies, including Advanced Manufacturing using National Information Infrastructure (AM-NII).

The NWC's electronic concurrent engineering and business processes will use the Internet to connect the AM-NII national laboratories (LLNL, Sandia National Laboratory, Oak Ridge National Laboratory, Los Alamos National Laboratory, and Allied Signal, Kansas City, Mo.) and their vendors. Given the sensitive nature of the work, the infrastructure must provide strong authentication, non-repudiation, message integrity, and privacy for information being exchanged over the Internet. The technology and strategies for implementing these capabilities over a large community of users has only recently become available.

AM-NII's Public Key Infrastructure (PKI) project will deploy a solution that spans five DOE sites. This solution will enable secure, authenticated exchange

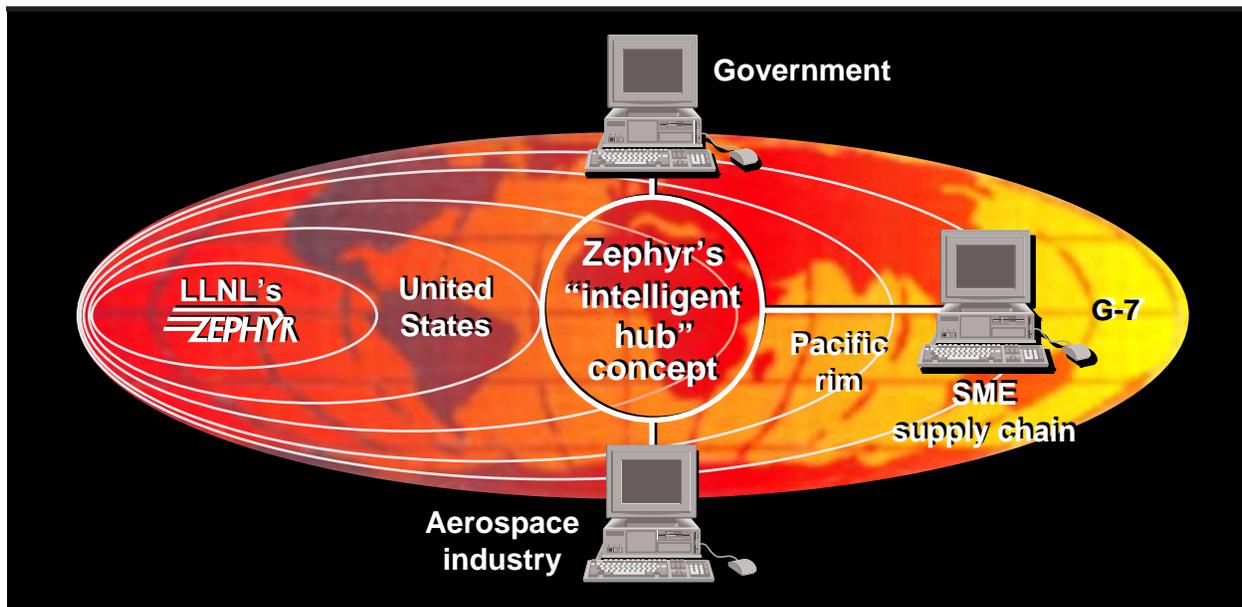


Figure 3. Zephyr's "intelligent hub" concept, linking the government, the aerospace industry, and the aerospace small and medium-sized enterprise supply chain. The goal is the support of ad hoc trading partnerships for an aerospace industry virtual supply network using collaborative engineering practices via the Web.

of e-mail, technical data, electronic forms, and other information required to conduct complex engineering and business processes over the Internet.

In the second initiative, LLNL has joined the CommerceNet-CALS Working Group pilot to develop and test the architecture for an intelligent hub. This server solution will eventually support information and operation of *ad hoc* trading partnerships for an aerospace industry virtual enterprise supply chain (with small and medium-sized enterprises) using collaborative engineering practices via the Web.

The CommerceNet-CALS pilot has received "global economic powers" (G7) endorsement as a potential candidate for world-wide trade and commerce. Zephyr's participation will provide LLNL with experience in collaborative commerce relationships with other members of the CommerceNet-CALS Working Group. It will also provide a way to access an additional base of SME vendors that will be important to growing programs such as NIF and NWC.

### Acknowledgments

The authors wish to thank L. Gottlieb, L. Reed, and LLNL's Technical Information Department for communications support (for example, home page development and graphic design); S. Hawkins for demonstrating solid modeling collaborative

exchanges between LLNL, EG&G/North Las Vegas, Allied Signal/Kansas City, and commercial vendors; W. Edmonson (recently deceased) for beta-testing ProCard and Unix Systems Management; B. Weber for code supporting file transfers to LLNL's Engineering Records Center; J. Miller for our printer/plotter capabilities; and R. Moore for benchmarking CERPS network performance with ISDN using IBM/PC clones from a remote location.

### References

1. Jordan, C. W., W. A. Niven, R. E. Cavitt, G. Pavel, G. Obegi, F. E. Warren, C. J. Buchanan, and J. M. Taylor (1996), *Using the World-Wide Web as a Medium for Concurrent Engineering and Rapid Prototyping*, Lawrence Livermore National Laboratory, Livermore, Calif., (UCRL-ID-122242).
2. Carter, D. E., and B. S. Baker (1992), *CE Concurrent Engineering. The Product Development Environment for the 1990s*, Addison-Wesley Publishing Company, Reading, Mass.
3. Palmer, C. (1996), *Bringing Design and Procurement Together*, Mechanical Engineering, Lawrence Livermore National Laboratory, Livermore, Calif, April.
4. Walton, R. E. (1989), *Up and Running: Integrating Information Technology and the Organization*, Harvard Business School Press, Boston, Mass. 



# Subcarrier Multiplexing: Optical Network Demonstrations

Paul D. Sargis and Bruce D. Henderer  
*Defense Sciences Engineering Division*  
*Electronics Engineering*

We have increased the channel capacity of our fiber-optic transport technology. We have also conducted field demonstrations on commercially-installed fiber to prove its viability and to discover its limitations. Our computer model now includes the effects of noise and cross-talk to better simulate network performance.

## Introduction

Subcarrier multiplexing (SCM) is a multi-channel, multi-gigabit per channel transport technology that breaks the traditional bit rate  $\times$  distance barrier of conventional multi-gigabit, long-haul fiber-optic transmission systems. Hybrid transmission systems using SCM and wavelength-division multiplexing (WDM) technology may allow for 500 Gbit/s-plus data transmission over long distances (1000 km).

In FY-95, we introduced our phase one prototype of SCM, which was capable of carrying two 2.5 Gbit/s data streams. Most of our efforts were then focused on studying SCM's inherent dispersion-reduction feature using laboratory demonstrations and computer models. We showed that our dispersion model was in close agreement with laboratory measurements.<sup>1</sup>

When we realized that SCM could have a revolutionary impact on long-distance telecommunications, we considered the issue of technology transfer. We quickly discovered that the telecommunications industry is not impressed with demonstrations which are confined to a sanitized laboratory environment. SCM would have to prove its viability on existing telephone company fiber routes, complete with lossy fiber and poor connections.

## Progress

In the past year, our efforts have focused on hardware and software improvements and on network demonstrations.

In our second hardware phase, we designed a four-channel SCM prototype, but were limited to a three-channel implementation because neither of

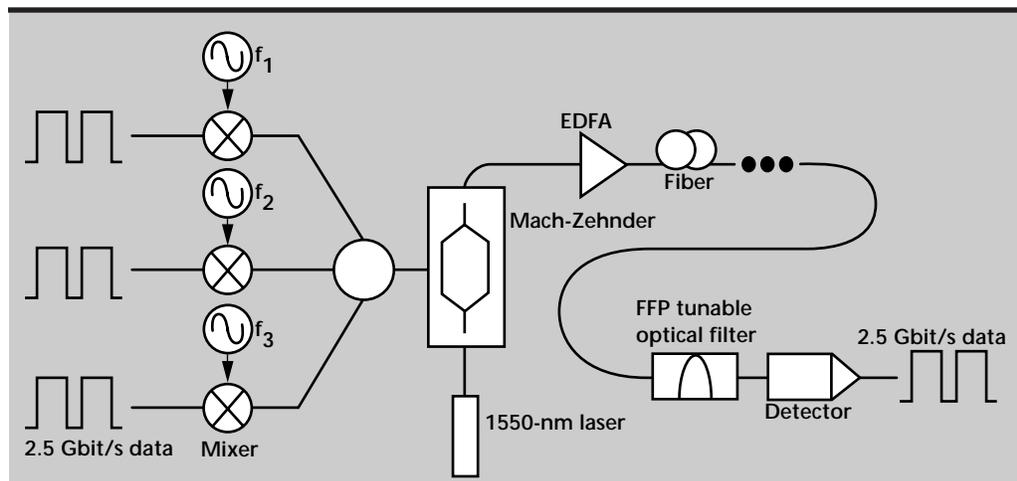


Figure 1. Block diagram of SCM system.

our suppliers of high-bandwidth Mach-Zehnder external modulators was able to deliver a device. Our 7.5 Gbit/s SCM system is diagrammed in Fig. 1.

In our second modeling phase, we added the capability to simulate transmission through multiple spans of fiber and multiple erbium-doped fiber amplifiers (EDFAs). Also, we have expanded the model to include the effects of adjacent-channel cross-talk, four-wave mixing, and stimulated Raman scattering. Inclusion of these noise sources and non-linear effects allows us to more accurately model the performance limitations of SCM under a variety of network conditions. The SCM model was created using LabVIEW.

We conducted demonstrations over the National Transparent Optical Network (NTON) around the San Francisco Bay. A map of the network is shown in Fig. 2. SCM occupied one of the four wavelengths present on the NTON. The network was configured so that we could loop back the SCM wavelength at San Ramon (100 km round trip) or at Burlingame (500 km round trip).

For our first network demonstration, we wanted to challenge the distance capability of SCM by adding 200 km of spooled fiber to the 500 km round trip over the network. Theoretically, fiber dispersion should begin to cause signal degradation at a fiber length of 640 km.<sup>2</sup> The results of this test are shown in Fig. 3. Figure 3(a) indicates slight degradation of the waveform due to dispersion. In

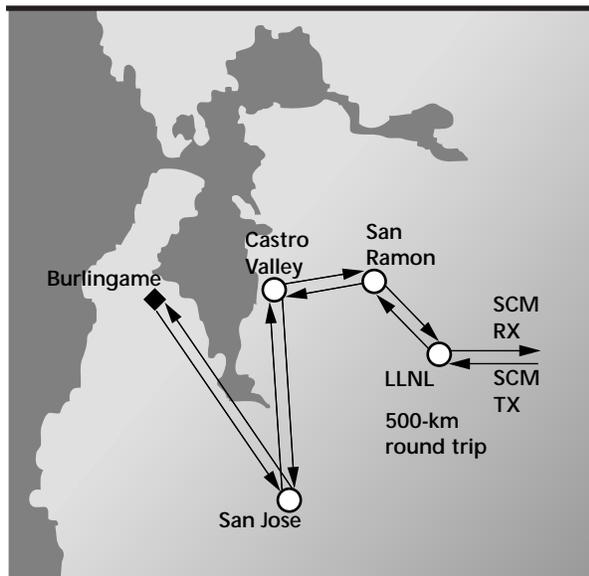


Figure 2. NTON fiber route for SCM demonstrations.

contrast, Fig. 3(b) shows what happens to the data when the dispersion-reduction enabling component is left out of the SCM receiver.

Next, we used the network to validate our enhanced computer model. The network was configured to loop back the SCM wavelength at San Ramon. We varied the optical attenuation at the SCM receiver to generate bit-error rate (BER) curves. Data was collected on a 2.5 Gbit/s subcarrier by itself, with cross-talk from one other subcarrier, and with cross-talk from two other

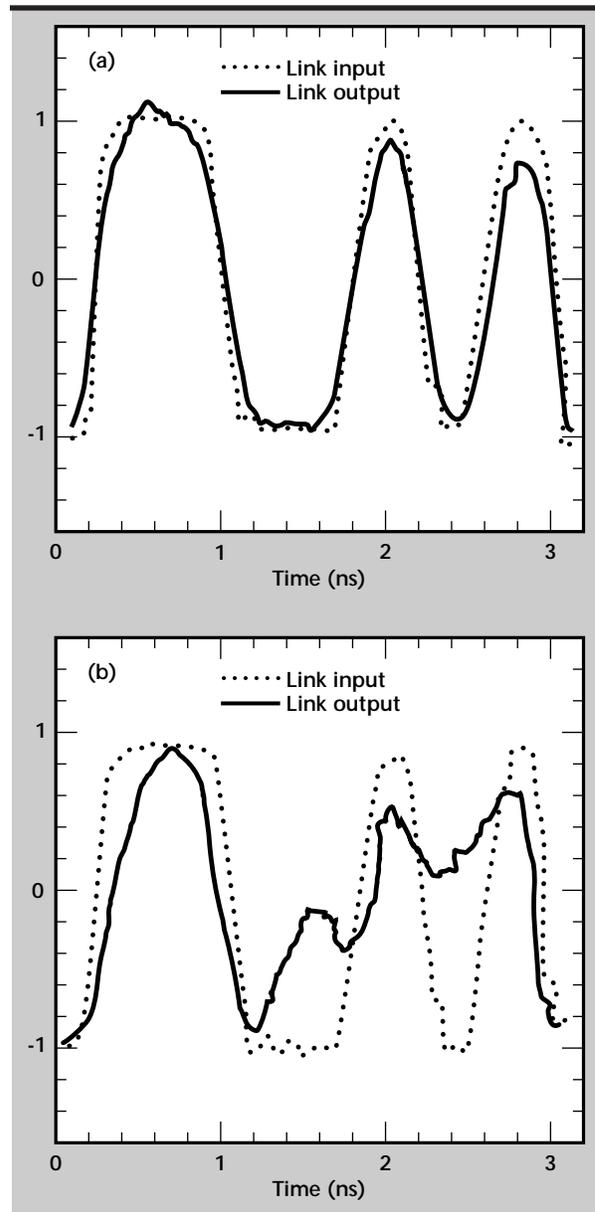


Figure 3. 700-km-SCM link results with (a) dispersion reduction and (b) no dispersion reduction.

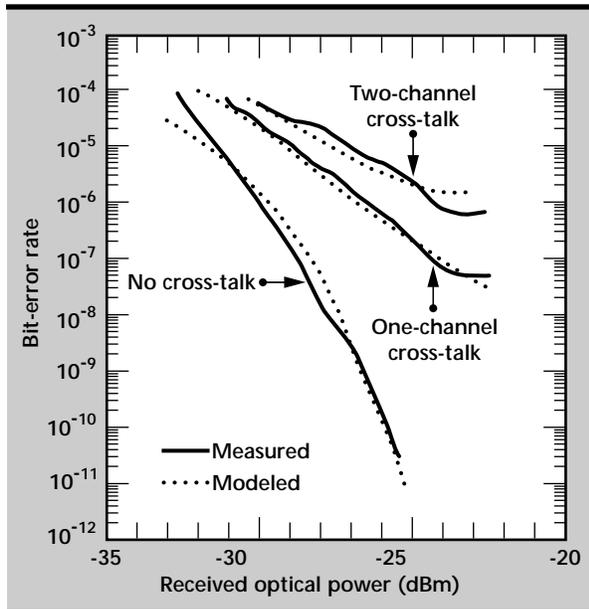


Figure 4. Results from SCM field trial.

subcarriers. Results from the computer model are compared with measured data in **Fig. 4**. High fiber loss between Livermore and San Ramon caused a low signal-to-noise ratio on the link, resulting in BER noise floors in the data. Minor differences between modeled and measured data can be attributed to measurement error and repeatability of fiber connections.

Representatives from Northern Telecom, GTE Sprint, Pacific Bell, and Uniphase Telecommunications Products witnessed a demonstration of SCM on the NTON. We are in the process of following up on those who expressed an interest in the technology.

### Future Work

When we are able to obtain a higher-bandwidth Mach-Zehnder external modulator, we will expand SCM to four 2.5 Gbit/s channels. In addition, we plan to move our SCM receiver to a remote location on the optical network and demonstrate remote channel selection. Also, we will continue to focus on establishing industrial partnerships with telecommunications companies and soliciting interest from government agencies.

### References

1. Sargis, P. D., R. E. Haigh, and K. G. McCammon (1995), *Dispersion-reduction technique using subcarrier multiplexing*, SPIE Vol. **2614**, pp. 244-251.
2. Sargis, P. D., R. E. Haigh, and K. G. McCammon (1995), *Subcarrier Multiplexing for Fiber-Optic Networks*, Lawrence Livermore National Laboratory, Livermore, Calif. (UCRL 53868-95). 



# **P**arallel Optical Interconnect Technology Demonstrations

Robert F. Hills, Ronald E. Haigh, and James M. Zumstein  
*Defense Sciences Engineering Division*  
*Electronics Engineering*

We report on the development of functional prototype communication and computing systems that use parallel optical interconnect (POI) technology developed under DARPA sponsorship. These prototype systems were developed and tested in field demonstrations with industrial partners to accelerate the adoption of POI technology.

## Introduction

Parallel optical interconnect technology uses an array of multimode fiber optic waveguides to transmit light from semiconductor laser diodes, usually vertical surface emitting lasers (VSELs), to an array of photodiodes that can be monolithically integrated with pre-amplifiers and decision circuitry.<sup>1</sup> The multimode ribbon fiber and an array connector simultaneously interconnect all of the fiber waveguides in a single ribbon connector.

The size of the multimode fiber, usually 62.5  $\mu\text{m}$ , allows for low cost, passive alignment packaging techniques. High levels of device integration allow compact, efficient optoelectronic transceiver modules to be realized.<sup>2</sup> Typical channel bandwidths in a POI module are 1 to 2 Gbits/s. The information capacity of a single POI module can approach several Gbytes/s.

The adoption of high-performance POI technology is essential to enabling high-performance communication and computing systems envisioned by several programs in the national interest; for example, Joint Strike Fighter (JSF) and the Accelerated Strategic Computing Initiative (ASCI).

For ASCI, microprocessors are seen as the building blocks of high-performance, parallel processor computer systems. As the number of processors on a system bus, and the instruction rate of each processor increases, electronic buses interconnecting these systems become a transaction-rate-limiting bottleneck for increasing computational performance. Optical interconnects will allow new bus interconnect fabrics to eliminate this bottleneck and will supersede the traditional electronic buses found in today's computers. An optical

interconnect eases the distance constraints between parallel processing nodes, and removes many of the limitations of copper interconnects at high data rates.

Functional prototypes of optically interconnected parallel processing and communication systems were co-developed with industry during this project. The POI technology was demonstrated with Lockheed-Martin under the DoD JSF technology development and maturation program, FORE Systems, with a prototype medical communications network at the Mayo Clinic, and a workstation cluster with Dolphin Interconnect Solutions. We have worked jointly with the DARPA-sponsored Optoelectronic Technology Consortium (OETC), as a neutral partner, to demonstrate optical interconnects in industry, provide testing and evaluation, and enable industry device developments.

## Progress

The POIs used in the technology demonstrations with industry (**Fig. 1**) are 32 channels wide with a total aggregate communication bandwidth of



Figure 1. 32-channel OETC Parallel Optical Interconnect transceiver modules.

2 Gbytes/s. The devices were developed by the OETC, a DARPA-sponsored consortium of IBM, Honeywell, Lockheed-Martin, and AT&T. The links were produced with two electrical interfaces, ECL and low-voltage differential signals (LVDS). LVDS is an IEEE standard (1596.3-1994), and specifies a low-voltage differential signal (250 mV swing) that is compatible with CMOS, BiCMOS, and GaAs circuitry. Bit-error tests on the LVDS POI's yielded a bit error rate of better than  $10^{-15}$ , which is acceptable for most applications. Since no errors were detected, the error rate may be better than this.

Three cooperative demonstrations with industry were the primary deliverables of this project. The necessary hardware to interface the POIs to the various applications was developed at LLNL. Each demonstration had unique system requirements, electrical interfaces, and transmission protocol. A common interface to POI was developed to allow the interface to be re-used in multiple demonstrations.



Figure 2. Parallel Optical Interconnect with 50 m of fiber optic ribbon cable configured for insertion into Joint Strike Fighter demonstration.

Application-specific interfaces were developed and served as the glue logic between the interconnect and the custom applications environments.

### DoD JSF Technology Insertion

The POIs were integrated into an SCI-based parallel processing system developed by Lockheed-Martin for JSF avionic applications.

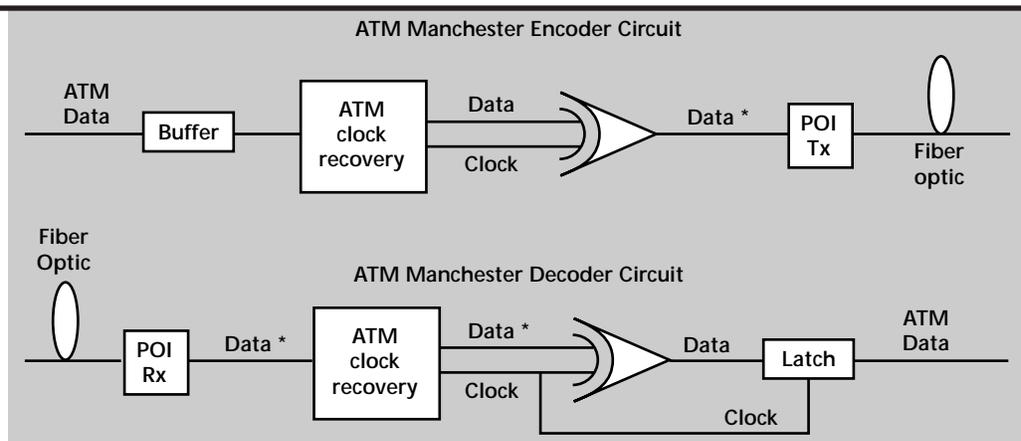
In this demonstration the POIs were integrated into the SCI backplane to allow parallel optical data transmission of 18 data and control bits at data rates of 250 Mbits/s per channel. POIs have application in JSF avionics by allowing processing and sensor resources to be distributed strategically across the air frame. During this demonstration, critical processing resources were physically isolated by 50 m of fiber optic ribbon cable with no degradation in computing performance (see Fig. 2).

### ATM Switch Transceiver Integration

The POIs were integrated into a commercial ATM communications switch developed by FORE Systems. The ATM backbone at the Mayo Clinic served as the site for the field trial. Essential telemedicine services are being developed by researchers at the Mayo clinic to allow physicians real-time access to digitized medical records, for image enhancements, and to enable distributed collaborations and consultation between physicians.

Due to the nature of the ATM protocol and the transmission characteristics of the POIs, an encoder/decoder circuit was developed to allow multiple ATM streams from asynchronous systems, to be transmitted across the interconnects. The

Figure 3. ATM Manchester encode/decode circuitry.



circuit Manchester encodes the data so that a data transition occurs at least once every bit period, and a complimentary decode circuit recovers the original data (Fig. 3).

### Optically Interconnected Computing Cluster

An optically interconnected computing cluster has been assembled at LLNL. The cluster consists of two SPARC workstations with SBus to SCI (Scalable Coherent Interface) adapter cards developed by Dolphin Interconnect Solutions. The SCI protocol allows applications to allocate and share memory across workstations, thus creating a powerful multiprocessing environment with commercially available workstations, a computing paradigm commonly referred to as cluster computing. The integration of POI technology into the physical layer allows these workstations to be distributed among offices and buildings, and across small business complexes. Logic translation from ECL to PECL was successfully implemented to interface the POIs to the SPARC adapter cards.

Future software engineering will reduce the communications overhead in the operating system, device drivers, and protocol stacks. Next generation silicon will push the interconnect bandwidth to beyond 500 Mbytes/s. Applications are being ported to this environment and we expect the optically interconnected computing cluster to be showcased at Supercomputing '96.

### Future Work

We are continuing this work by developing first generation wave division multiplexing (WDM) POI technology. WDM transceivers enable many channels

of data to be transmitted on a single optical fiber. Unique WDM technologies developed at LLNL will enable many interconnect fabrics that are not possible with copper interconnect. It is expected that these novel fabrics will be useful in advanced communication and computing systems.<sup>3</sup>

The objectives of these technology demonstrations were to accelerate the adoption and insertion of optical interconnect technology into the future communication and computing systems. Market acceptance and development, and insertion of this technology into COTS components are critical to the missions of DOE ASCI and DOD JSF programs and others.

### References

1. Choi, J., D. L. Rogers, D. M. Kuchta, F. H. Kwarck, H. Ainspan, K. G. Stawiasz, and J. D. Crow (1996), "High-performance high-yield, uniform 32 channel optical receiver array," *Optical Fiber Conference (OFC) '96, Maui, Hi., FC4*, October.
2. Wong, Y.-M., D. J. Muehlner, C. C. Faudskar, D. B. Buchholz, M. Fishteyn, J. L. Brandner, W. J. Parzygnat, *et al.* (1995), "Technology development of a high-density 32-channel 16-Gb/s optical data link for optical interconnect applications for the optoelectronic technology consortium (OETC)," *IEEE Journal of Lightwave Technology*, **13**, pp. 995-1016.
3. DeGroot, A. J., R. J. Deri, R. E. Haigh, F. G. Patterson, and S. P. DiJaili (1996), "High-Performance Parallel Processors Based on Star-Coupled WDM Optical Interconnects," *1996 Conference on Massively Parallel Processing with Optical Interconnects, Maui, Hi.* 





# Intelligent Automation Architecture

Scott A. Couture  
*Applied Research Engineering Division*  
*Mechanical Engineering*

This short-term exploratory project focused on identifying an expedient and efficient mechanism to encourage collaboration among a wide variety of Lawrence Livermore National Laboratory (LLNL) researchers working on various aspects of intelligent computing machines. The primary project goals were to identify the set of applicable competencies at LLNL, review the literature in the field, and determine the most practical way to merge the results to encourage collaboration. A careful review of the subsumption architecture, coupled with a review of critical theoretical work in the fundamentals of system complexity lead to the conclusion that the subsumption architecture differs in focus and distribution of system complexity, but not in the overall level of complexity required to achieve intelligent behavior. This facilitates less rigorous approaches to collaboration and will allow the integration of a wider set of competencies than would be practical otherwise.

## Introduction

Constructing machines that act intelligently has been a pursuit of researchers around the world for many years. Expert systems have been used to capture rules and knowledge that human experts use when solving problems in specific areas. Research in artificial intelligence and reasoning has been used to develop "intelligent" planning systems. However, little of this work has made it into practical use for real-time control of autonomous systems. Currently available controller technology is insufficient to build "intelligent agents" for even simple tasks in the real, unstructured world.

A number of researchers at LLNL have been advancing the state of the art in many key technologies needed to develop more intelligent machines. These include machine controls; high-speed collision avoidance and real-time path planning; high-speed 3-D image processing; neural nets; fuzzy logic; task planning; parallel processing and supercomputing; intelligent sensor development; data representation; and data mining. This research has been supported by many different sponsors and has been applied to many different programs. One of the key missing pieces has been a unifying architecture or "backbone" that facilitates integration of this wide variety of competencies. Determining an expedient and effective method to develop that backbone is the objective of this research.

In the short term, we plan to identify LLNL researchers in all the key areas needed to pursue development of an interactive intelligent computer system, to identify the state of the art in intelligent machine architectures, and to select an architecture that will facilitate the researchers' collaboration toward implementing intelligent behavior in a particular application. In the longer term, our goal is to develop a uniform computing paradigm for autonomous systems that allows tasks to be described in a high-level goal definition language suitable for use in an industrial environment, while incorporating advanced control features being developed at LLNL and other research institutes.

The definition of the architecture and interfaces will help foster collaboration across diverse research groups and facilitate integration of different research results. Hopefully, the process will also serve to kindle increased interest in intelligent systems and spawn follow-on work to implement a core intelligent "backbone."

## Progress

### **Problem Definition**

The lack of widespread use of traditional artificial intelligence approaches in real-time control is primarily due to their inefficiency and their failure to meet the fast reaction times needed for real-time

systems. Reaction times suffer from the vertical (or hierarchical) control approach that is the basis for current machine controllers. Passing data up and down through control layers requires an incredible amount of vertical bandwidth to communicate input data to the high-level model and output data back to the actuator. Similarly, nearly all of the data interpretation and behavior rules reside at the highest level of the control, leading to very complex context-dependent knowledge representation schemes and rules.

A symptom of this focused complexity is difficulty maintaining real-time performance when reacting to these types of situations. A quick comparison of our computer systems' sensor/processor/actuator bandwidth with some of nature's biological counterparts shows something is awry. The biological systems have much less vertical communication bandwidth and less processing, but perform much more intelligently. Hence, a possible change in approach seems warranted.

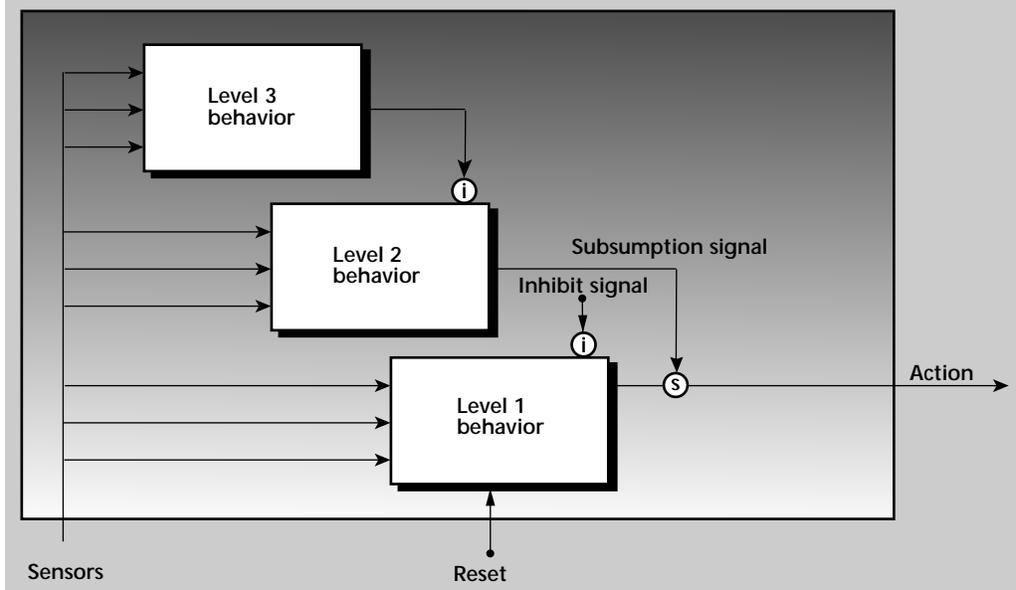
In fact, a number of different approaches have been proposed to address the deficiencies in current "intelligent" machines. These include the biologically-inspired subsumption architecture (SA) by Brooks,<sup>1</sup> the Perception-Action model by Kaelbling,<sup>2</sup> a number of other biologically inspired approaches<sup>3,4,5</sup> and a variety of hybrid architectures that blend goal- and behavior-oriented approaches.<sup>6,7</sup> Of these, the biggest stimulus to move away from goal-oriented, plan-based, "good-old-fashioned AI" systems towards behavior-oriented "new AI" was provided by Brook's work in SA at MIT beginning in 1986. Since that time, the SA approach has been widely applied to (primarily) low-level control of autonomous agents in real-world situations.

Behavior-based approaches still dominate the literature and are the focus of many related approaches. Therefore, our work began with developing an understanding of the SA approach, its strengths and weaknesses, and its applicability to LLNL's needs. As it turns out, careful examination of the SA approach as it has evolved over the past ten years reveals telling similarities to other biologically-inspired research. Further, a theoretical understanding of the conservation of complexity principle points to reasons why the architecture yields robust systems in rich environments and also bounds the possible complexity of the "emergent" intelligent behavior.

### Subsumption Architecture

SA is a modular, behavior-based approach to building machine controllers to work in the real-world. The inspiration for SA comes, in part, from a study of insect behavior and a realization that even very neurologically simple insects behave "intelligently," or at least appropriately, in the real world. In general, each module in a behavior-based system exploits the physics of the environment as much as possible, describes an interaction between the agent and the environment, and attempts to use emergent behavior whenever possible.<sup>8</sup> As shown in Fig. 1, modules have very simple interfaces with other modules and the world. Modules are reactive—they read the state of the world as input, determine an appropriate behavior, and output that behavior. What is appropriate behavior is dependent upon the internal state and sensors measuring the state of the environment. There is no global symbolic representation of the world or plan.

Figure 1. Basic behavior module arrangement in Brook's subsumption model.



Behavior modules are implemented as augmented finite-state-machines (FSMs) with timers. Additional control signals allow the module's behavior to be integrated with other behaviors. The primary control signals are a suppression signal, which causes the output of the module to be suppressed while it is active, and a subsumption signal, which causes the output of the module to be replaced, or subsumed, by the behavior supplied by another module. More complex behaviors build upon simple behaviors. For example, a *walk* behavior may subsume the *stand* behavior's "leg down" action, to allow the insect to walk. Each module implements a specific behavior. Collections of behaviors interacting with the environment determine the agent's collective competence level.

Since Brooks introduced SA in 1986, numerous systems have been built using its basic principles, with applications varying from mobile robots to vision systems and theorem proving.<sup>9-12</sup> Results have shown that when SA is applied to low-level reactive control problems, like collision avoidance during mobile vehicle navigation or wall following (where the physics of the environment can be exploited and sensor information is rich), the approach is more compact (requires less computer code) and generally more robust than monolithic control approaches.

However, a number of shortcomings must be addressed before SA approaches are likely to address higher level, reasoning behaviors.<sup>13</sup> In particular, strict SA approaches have no representation of the world, which will make higher level reasoning based on symbolic knowledge representations difficult. In SA, each layer works independently without consideration of the total strategic plan, so systems are often deadlocked trying the same reaction many times.<sup>14</sup> A number of these drawbacks have been addressed through refinement and extension of the SA approach.

### Subsumption Architecture Variants

Gottschalk addressed the deadlock problem that occurs when a reactive SA system is caught between two opposing stimuli using a hybrid architecture that resulted in a system which mimics the brain's instinctive (brain stem), behavioral (limbic system) and goal (cerebral cortex) levels.<sup>3</sup> The behavior level uses situational if/then rules to resolve the deadlock and the goal level to address longer term goals.

One of the factors that inhibits extending the SA approach to higher level reasoning is that the state information about the world is distributed throughout the behaviors and is not visible to other behaviors.

To communicate the information between modules, more elaborate mechanisms will need to be developed than the simple control signals now in use.

One of the current favorite mechanisms to communicate global information in SA systems is the use of hormones—another concept borrowed from the biological community.<sup>15,16</sup> Hormones are global state variables that are influenced by the environment and are responsible for establishing which behaviors are dominant in a particular situation. For example, when sensors determine that a walking surface is uneven, a hormone related to *fear* increases. In response, the *walk* behavior causes the walking rate to slow so the system is more "careful." Researchers have noted similarities between biological systems in panic and fear situations to their mechanical equivalents after hormones have been added.

Another severe criticism of SA is that there is no way to arrive at a compromise solution. When one module inhibits the behavior of another, the concerns of the inhibited module are ignored, so there is no way to arbitrate between them.<sup>17</sup> The inhibit function often undermines the architecture by shutting off low-level survival type behaviors. Instead of forcing one behavior to be dominant and subsume another behavior altogether, an alternative approach researchers have taken is to modify the SA architecture connectors to allow cooperation between behaviors.

The Action Selection Dynamics architecture allows all behaviors to be active at the same time, but applies variable priority (or weighting) to competencies depending on the situation.<sup>18,19</sup> The weightings can also be changed based on reinforcement to accomplish learning, similar to neural nets. With outputs weighted between  $-1$  and  $1$ , one behavior can override another completely, but the system responds in a more continuous way, and the approach allows a sharing of information between modules that is not possible with the SA inhibit signal.

Kaelbling's perception-action model, shown in **Fig. 2**, is also behavior-based inasmuch as it maps input information and agent state to an output action.<sup>2</sup> However, Kaelbling explicitly addresses information in the context of reaching a goal. As more state information becomes available, perception of the state of the environment improves and the action is more specific to the situation at hand. The perception/action model can be implemented in an SA approach, where each perception/action pair is a behavior module, or in a more traditional vertical approach where information acquisition (perception) is separate from action.

Kaelbling's theoretical treatment of information, goals and behaviors applies whether a vertical or horizontal implementation of the perception-action model is pursued. If that is the case, how does "emergent" intelligent behavior arise from a SA approach when it presumably does not in traditional systems?

### Conservation of Complexity

Proponents of the SA approach claim that intelligent behavior "emerges" from collections of simple behaviors. How complex the resulting behavior can be is a key question that is rarely answered in application papers. This is an issue that goes to the core of behavior-based approaches—if there is no "emergent" intelligence in a behavior-based approach, then the difference between goal and behavioral approaches is one of structure and focus and not substance.

This issue is addressed in detail by Werner in his theoretical analysis of the complexity of systems.<sup>20</sup> As Kaelbling also points out, agents and their environments must be treated together as a system. The complexity of a task depends in large part on the complexity of the environment that the task is taking place in. For example, if all parts in an assembly are placed in the environment in their correct relationships, an agent does not have to have the internal complexity necessary to figure out their relationships prior to assembly.

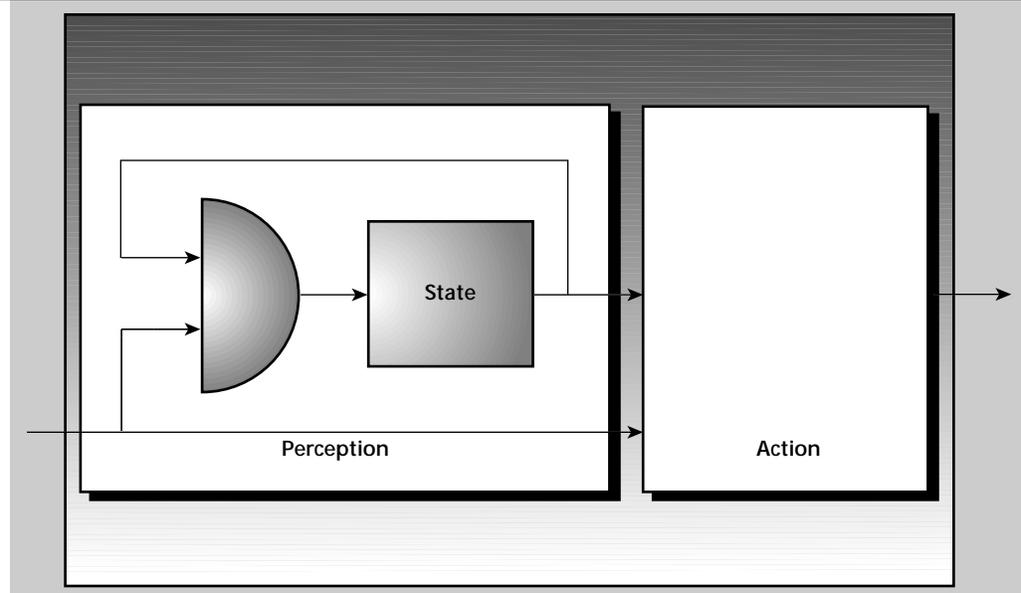
The first SA design principle—take advantage of the physics of the environment—is a conscious effort to reduce agent complexity by transferring agent complexity to the environment. If the designer

knows there will always be a line on the floor to follow, the navigation agent can be very simple even if the pattern on the floor is very complex. Given that the assumptions about the physics of the environment always hold true, the agent is competent to behave intelligently whenever it is in such an environment.

The fundamental conservation of complexity principle states that the complexity of an event generated by a set of agents cannot be greater than the complexity of the agents plus the complexity of the influencing environment. Complexity is additive, so the greater the number of agents with distinct roles (competencies), the greater the possible complexity of the emergent behavior. Again, behavior-based approaches that use a broad collection of specific agents can be expected to have fairly complex "emergent" behavior. However, collections of identical agents, as some have proposed, cannot exceed the complexity level of the individual agent unless the complexity arises from information "encoded" in the environment or a plan (program). For some agents, the environment is the program.

Information reduces possibilities and increases abilities of agents. The representation the agent has of the world increases in complexity with increasing information, allowing the agent to generate more complex output. Reactive, behavior-based systems are sensor- and information-rich compared to many of their traditional AI counterparts. The information is distributed, perhaps appropriately, much closer to the behavior that is influenced by the information, leading to faster reaction times than traditional global model based systems.

Figure 2. Division between the perception and action components in Kaelbling's model.



After careful consideration of the conservation of complexity principles, it appears that "emergent" behaviors from SA are intelligent only insofar as the behaviors are designed to properly account for the physics of the environment and the expected interaction between the collection of agents. It is very likely that behavior-based systems can, and often do, exhibit inappropriate, and not-so intelligent behavior, given either poorly defined behaviors, or improper assumptions regarding the environment.

One inherent benefit from the SA approach is the established design discipline that focuses the designer on the environment, the agent's behavior, and the agent's interaction with the uncertainties of the environment. Given the same analysis and the same set of information-gathering tools (sensors), alternate architectures will likely behave as intelligently as SA approaches. In addition, alternative approaches may benefit from a more extensive set of communication and symbolic data representation capabilities than the restrictive set provided by the classic SA approaches.

There is no fundamental reason why higher level symbolic reasoning systems and low-level reactive systems cannot be used together to achieve more complex goals. In fact, according to the conservation of complexity principle, improved performance is certain.

## Conclusions

By understanding how evolution has successfully encoded environmental complexity and data interpretation into our biological systems, computer analogs can potentially use similar tricks to avoid a lot of agent complexity. The jury is still out regarding how far up the reasoning chain the SA approach can be taken without encountering serious stumbling blocks. Perhaps current work in biological equivalents will assist in the transition from reactive to symbolic reasoning paradigms when that bridge must be crossed.

## Future Work

This exploratory effort was funded during the last part of FY-96 and is complete. A task analysis identified a core set of sensors and behaviors needed to implement an intelligent system in the selected task and environment scenario. Given available funding, there is sufficient interest and competency at LLNL to do further work this area.

## Acknowledgments

I would like to thank a number of researchers who were unable to take time from their job commitments so late in the year, but chose to participate on their own time. In particular, thanks to B. Maurer, K. Wilhelmsen, R. Hurd, and especially C. Mason of the University of California, Berkeley, whose skill mix we did not have at LLNL.

## References

1. Brooks, R. A. (1986), "A Robust Layered Control System for a Mobile Robot," *IEEE Journal of Robotics and Automation*, Vol. RA-2, (1).
2. Kaelbling, L. P. (1988), "Goals as Parallel Program Specifications," *Proceedings of the Seventh National Conference on Artificial Intelligence*, St. Paul, Minn.
3. Gottschalk, M. A. (1996), "Intelligent motion controller mimic's brain's structure," *Design News*, March, pp. 119-121.
4. Mura, F., and N. Franceschini (1994), "Visual control of altitude and speed in a flying agent," *Proceedings of the Third International Conference on Simulation of Adaptive Behavior, (From Animals to Animats)*, Cambridge, Mass., 3, pp. 91-99.
5. Murray, D. (1995), "Developing reactive software agents," *AI Expert*, Vol. 10, (3), March, pp. 26-29.
6. Innocenti, C., G. Mondino, P. Regis, and G. Sandini (1994), "Trajectory planning and real-time control of an autonomous mobile robot equipped with vision and ultrasonic sensors," *Proceedings of the IEEE/RSJ/GI International Conference on Intelligent Robots and Systems (IROS '94)*, New York, N.Y., Vol. 3, pp. 1861-1866.
7. Lindley, C. A. (1993), "Autonomous satellite architecture integrating deliberative reasoning and behavioural intelligence," *Telematics and Informatics, United Kingdom*, Vol. 10, (3), pp. 231-250.
8. Steels, L. (1994), "A case study in the behavior-oriented design of autonomous agents," *Proceedings of the Third International Conference on Simulation of Adaptive Behavior, (From Animals to Animats)*, Cambridge, Mass., 3, pp. 445-452.
9. Stein, M. R., and R. P. Paul (1994), "Operator interaction, for time-delayed teleoperation, with a behavior-based controller," *Proceedings of the 1994 IEEE International Conference on Robotics and Automation, Los Alamitos, Calif.*

10. Gomi, T., and K. Ide (1994), "Vision based navigation for an office messenger robot," *Proceedings of the IEEE/RSJ/GI International Conference on Intelligent Robots and Systems, (Advanced Robotic Systems and the Real World), (IROS '94), New York, N.Y., Vol. 3*, pp. 2015–2022 .
11. Pinhanez, C. S. (1994), "Behavior-based active vision," *International Journal of Pattern Recognition and Artificial Intelligence*, Vol. **8**, (6), pp. 1493–1526.
12. Brady, A. (1994), "A subsumption architecture for theorem proving?" *Philosophical Transactions of the Royal Society, Series A (Physical Sciences and Engineering), United Kingdom*, Vol. **349**, (1689), pp. 71–84.
13. Arnold, J. E. (1989), "Experiences with the subsumption architecture," *Proceedings of The Fifth Conference on Artificial Intelligence Applications, Washington, D.C.*, pp. 93–100.
14. Yamamoto, M. (1994), "SOZZY: a hormone-driven autonomous vacuum cleaner," *Proceedings of the SPIE—The International Society for Optical Engineering*, Vol. **2058**, pp. 211–223.
15. Brooks, R. A., and P. A. Viola, (1990), "Network based autonomous robot motor control: from hormones to learning," *Advanced Neural Computers, Amsterdam, Netherlands*, pp. 341–348.
16. Gomi, T., and J. Ulvr (1993), "Artificial emotions as emergent phenomena," *Proceedings of 1993 2nd IEEE International Workshop on Robot and Human Communication, New York, N.Y.*, pp. 420–425.
17. Rosenblatt, J. K., and D. W. Payton (1989), "A fine-grained alternative to the subsumption architecture for mobile robot control," *International Joint Conference on Neural Networks, New York, N.Y.*, Vol. **2**, pp. 317–323.
18. Webber, A. D., and D. L. Bisset (1994), "Competition and co-operation—a model for behaviour-based robot controllers," *Proceedings of PerAc '94, (From Perception to Action), Los Alamitos, Calif.*, pp. 384–387.
19. Gomi, T., and K. Ide (1994), "Emulation of emotion using vision with learning," *Proceedings of 1994 3rd IEEE International Workshop on Robot and Human Communication, New York, N.Y.*, pp. 210–215.
20. Werner, E. (1994), "What ants cannot do," *Proceedings of 6th European Workshop on Modelling Autonomous Agents in Multi-Agent World, Distributed Software Agents and Applications*, pp. 19–39. 





Technical Information Department  
Lawrence Livermore National Laboratory  
University of California  
Livermore, California 94551