UNIX IN THE OFFICE

PRODUCTS · TRENDS · ISSUES · ANALYSIS

The Informix RDBMS

Exploring Paths Less Traveled

By Laure Brown

WHEN YOU THINK about recent trends in relational database management system (RDBMS) technology, you tend to think of distributed transaction processing, referential integrity, or maybe multithreaded architectures. Informix, meanwhile, is making headway into alternative—though valuable—areas: multimedia and an expanded front-end toolset. This isn't to say that Informix is ignoring the prevailing RDBMS winds. Hardly. *(continued on page 3)*

• E D I T O R I A L •

Do the Right Thing! (Please)

Let's keep things in perspective and judge OSF and Unix International by their deliverables, not their posturing.

By Judith S. Hurwitz

FRUSTRATION AND confusion continue to dominate the computer landscape these days. The schism between the Open Software Foundation (OSF) and Unix International widens. We are forced back to the days when OSF was first formed, a time when emotions ruled and the future of Unix looked unstable. Ironically, things are not as bad as they appear. It was clear from the beginning that the cultures of OSF and Unix International were too different for the two to merge. Yes, we were intrigued by the possibility of a merger, but quite skeptical. The core of the difference between these two groups comprises the history of the companies involved, and their positions within the Unix industry specifically and the computer industry in general.

Companies that dominate Unix International tend to be old timers in the Unix market. Companies like AT&T, Unisys, Altos, and NCR have paid their dues, having been trailblazers at a time when no one paid much attention to Unix, except researchers working on some esoteric programs. These companies have taken a conservative approach to the operating system. They have no choice. They have large installed bases of customers who do not intend to migrate to new technology too quickly. Therefore, Unix International has taken on the mantle of protecting this base of customers, even at the expense of brave new and exciting technology. As a way to regain the upper hand, Unix International and AT&T wanted OSF to accept System V.4, including the Open Look user interface.

The Open Software Foundation, on the other hand, is driven by Unix upstarts like IBM, Digital, and Hewlett-Packard. For pragmatic reasons, these companies have taken on the open systems mantle. Most of their traditional customers use proprietary operating systems, not Unix. By pushing innovation within Unix, they can have the best of all possible worlds: They can protect their existing customers while offering new technology to an untapped base.

Given these different agendas and perspectives, it is not surprising that no progress has been made toward unification. However, there is also another subtle force at play. AT&T is not yet a significant power in the computer industry. Therefore, it was not able to contain and manage the players trying to gain control over a potentially critical operating system. Even though AT&T's partner, Sun Microsystems, is a significant player in the market, Sun is perceived as a renegade—a smart technology company, but not an industry leader. Contrast the Unix industry schism with the quieter and more contained picture of the DOS and OS/2 market. While there is plenty of debate over the longevity of DOS and the validity of OS/2, there is no move afoot to take control of either operating system away from Microsoft or IBM. These two companies are simply too powerful to mess with.

We would have preferred the two organizations to have resolved their differences in a more amicable way. The reality is that there is still common ground that both are committed to. They will continue to work with one another on their respective technology workgroups. They are committed to implementing the X/Open Portability Guide (XPG3) and the ever-expanding Posix definitions. Therefore, in the end, there should be few differences between what users will be able to get from either organization. The disappointment is that the rhetoric may continue to obscure the positive aspects of both. Let's try to keep things in perspective and judge OSF and Unix International by their deliverables, not their posturing. ◐

•INFORMIX•

(continued from page 1) The company has supported transaction processing since 1986 and has a distributed architecture in place. In fact, it has even renamed its database engine after OLTP to Informix-OnLine.

As a distributed server, OnLine leaves something to be desired, whereas its multimedia support and its desktop strategy are really quite good. In this article, we'll point out the specific weaknesses and strengths of the Informix information management system—its architecture and strategy as well as specific products.

Background

Originally called Relational Database Systems, the company's first RDBMS debuted in 1981. It was the first commercial Unix RDBMS—a claim that Informix is proud of. Unix has always been first and foremost in the company's strategy. Informix is almost smug about having latched onto Unix before it became trendy. Informix is second only to AT&T as a Unix software vendor.

That original RDBMS had a proprietary query language and was built on the Indexed Sequential Access Method (C-ISAM), which includes B-tree indexing, audit trails, and various file-locking capabilities. (C-ISAM was actually Informix's first marketable product and has become a standard in the Unix environment.) The current RDBMS evolved more from a 1985 version, an SQL-based engine that sported a menu-based interface and a database schema editor. The other major enhancement was the development in 1986 of Informix-4GL, a fourth-generation application development environment.

PICKING UP SMARTWARE AND WINGZ. Through its 1988 merger with Innovative Software, Informix inherited the SmartSeries (a.k.a. SmartWare) office system and Wingz, a graphical and powerful spreadsheet. Informix has been shrewd in integrating these two products with its database engines, making corporate data accessible from the desktop.

MARKETING STRATEGY. As a Unix software vendor, Informix has always marketed to smaller, multiuser environments. And it will continue to do so. However, since the Unix market is opening up, so is Informix—to larger MIS organizations. The company claims a fairly even distribution ratio: About 33 percent of its business comes from VARS, another 33 percent from OEMs, and the remainder from direct sales.

The company is riding its office automation toolset (i.e., the integration of SmartWareII and Wingz) as a marketing

> *Informix has always marketed to smaller, multiuser environments. And it will continue to do so. However, since the Unix market is opening up, so is Informix—to larger MIS organizations.*

vehicle. In addition, it hopes to lure new customers with its fault-tolerant, OLTP server. However, that won't be easy, because other Unix RDBMSs clearly offer more OLTP and distributed functionality than Informix does. While Informix has added features for fault tolerance, performance, and distributed querying, it has not touched server-side referential integrity (i.e., guarantees that data protection rules have been implemented correctly) or distributed transaction processing (i.e., updating multiple databases in a single transaction). The company reasons that neither feature has been standardized yet. Well, that may be the case, but both are requirements for mission-critical, OLTP environments. Both Sybase and Ingres already offer these capabilities.

Financial Standing. After teetering a little two years ago in the market, Informix looked good on last year's financial reports. Evidently, the company lost its equilibrium in 1988 after the merger with Innovative Software. Lots of marketing money was poured into SmartWareII and Wingz, which weren't released until 1989—severely late. However, in 1989, the company realigned its management under new CEO Phil White, and Informix made a profit, with revenues reaching $145 million. Currently, Informix and Oracle are the leading RDBMSs in the Unix marketplace. Informix shows a lead over Oracle in the United States, and, in Europe, its lead is large.

Product Line

In addition to the connectivity software that provides the foundation for its distributed architecture, Informix has three sets of products: database engines, application development tools, and tools for users. In its current release, Informix decided to break out the tools from the engine, so customers can select the products best suited for them.

THE ENGINES. Informix-OnLine, the SQL-based, back-end database management server, is the core of the Informix product line. It's actually a new release (Version 4.0) of Informix-Turbo, the company's original OLTP back-end database server. Aside from support for transaction processing, OnLine features support for distributed databases and multimedia.

The low-end engine, Informix-SE, is not as robust a system. It's SQL-based and built on C-ISAM, but geared toward smaller organizations, where performance and fault-tolerance are lesser issues than easy maintenance and administration.

DEVELOPER TOOLS. Informix offers developers two tools: Informix-4GL and ESQL (Embedded SQL), which embeds SQL into C, Cobol, Fortran, and Ada source programs.

Informix-4GL. In addition to basic 4GL programming functionality (e.g., store, retrieve, update, delete), Informix-4GL has a screen-building utility, a menu-building utility, a report writer, and a window manager.

You can use a compiled version of Informix-4GL, or you can opt for the 4GL Rapid Development System and Interactive Debugger (RDS/ID), which cuts down on application development time by eliminating the need for a C linker and C compiler. With it, you can interactively debug an application by viewing both the source code and the applications screens as the application is executed. If you've got a graphical workstation, you can view the application screen and the program screen simultaneously in two separate windows; otherwise, you can toggle between them. When you're ready to go into production with the application, you use the Informix-4GL Compiler to recompile RDS/ID code into standard C.

The most attractive of the Informix product line is Wingz. Essentially, it's a graphical spreadsheet with snazzy graphics—not exactly your traditional choice of data access.

USER TOOLS. The full suite of end-user products includes database access and manipulation tools—Informix-SQL and QuickStep, a menu-driven report generator—as well as office automation tools—Wingz and SmartWare.

Informix-SQL. Informix-SQL is comprises five modules:

- A forms package

- A report writer

- A command-driven, interactive SQL editor

- A menu builder for integrating forms, reports, SQL statements, and other menus

- A menu-driven database schema editor for creating and modifying databases and tables

QuickStep. QuickStep is a visual, intuitive way to go about reporting. You don't have to write or know any SQL code, and you create and edit your report via pull-down menus. Once you've designed the report on screen, QuickStep will generate the SQL code for you—or you have the option of generating Informix-4GL code so the report can be integrated into 4GL applications, as well.

SmartWareII. We'd be straying too far from our scope here if we were to go into detail on the entire SmartWareII office system. However, since SmartWare is part of Informix's front-end strategy, its highlights are worth mentioning.

From the database perspective, the value of Smart lies in its ability to transparently access the database from within the system's spreadsheet module (see "Front-End Strategy" below). As an office system, however, its best attribute lies in the integration and consistency among modules (word processor, spreadsheet, business graphics, personal database). Each module has direct access to global commands. The links among modules aren't live. You can, for instance, drop a spreadsheet into a word processing document, but it won't be a live spreadsheet—just an ASCII representation. If you want to include a graphic, Smart will include the graphic reference, along with the space that the graphic will fill. The Smart graphics are actually pretty good for an office system—with 3-D bar charts and pie charts that you can rotate and edit. The system also features macros and a programming language called SPL (Smart Programming Language) that enables you to create applications that combine the functionality of Smart's available modules.

Wingz. The most attractive of the Informix product line is Wingz. Essentially, it's a graphical spreadsheet with snazzy graphics—not exactly your traditional choice of data access. But Wingz offers additional functionality that makes it an appropriate tool—and a powerful one at that.

Aside from its advanced spreadsheet capabilities, Wingz includes HyperScript. And here Wingz gets interesting. HyperScript is a portable, fairly intuitive command- and event-driven programming language that has a DataLink option for accessing and manipulating database information (see Illustration 1). HyperScript in itself provides a strong application development environment. With it, you can design and customize graphical interfaces, create applications using regular programming constructs (e.g., global and local variables, CASE statements, and loops), or build a form (or screen or menu). Wingz also has a Learn mode, which records your actions as you navigate your way through point-and-click and dialogue box options. Thus, you don't need to know HyperScript to use it.

The DataLink option lets you reach a remote database from within the spreadsheet (see "DataLink" below). In other words, HyperScript can be used to create a customized front end for entering and updating data across a network. You can either type statements into an interactive SQL editor or use menus and dialogue boxes.

Until recently, Wingz has been a Mac product, but, earlier this year, it was ported to OS/2, NeXT, and a handful of graphical Unix systems—among them those of Apollo, Digital Equipment, and Sun. Since HyperScript code is portable, an application built on the Macintosh, for instance, can be recompiled and run on a NeXT cube.
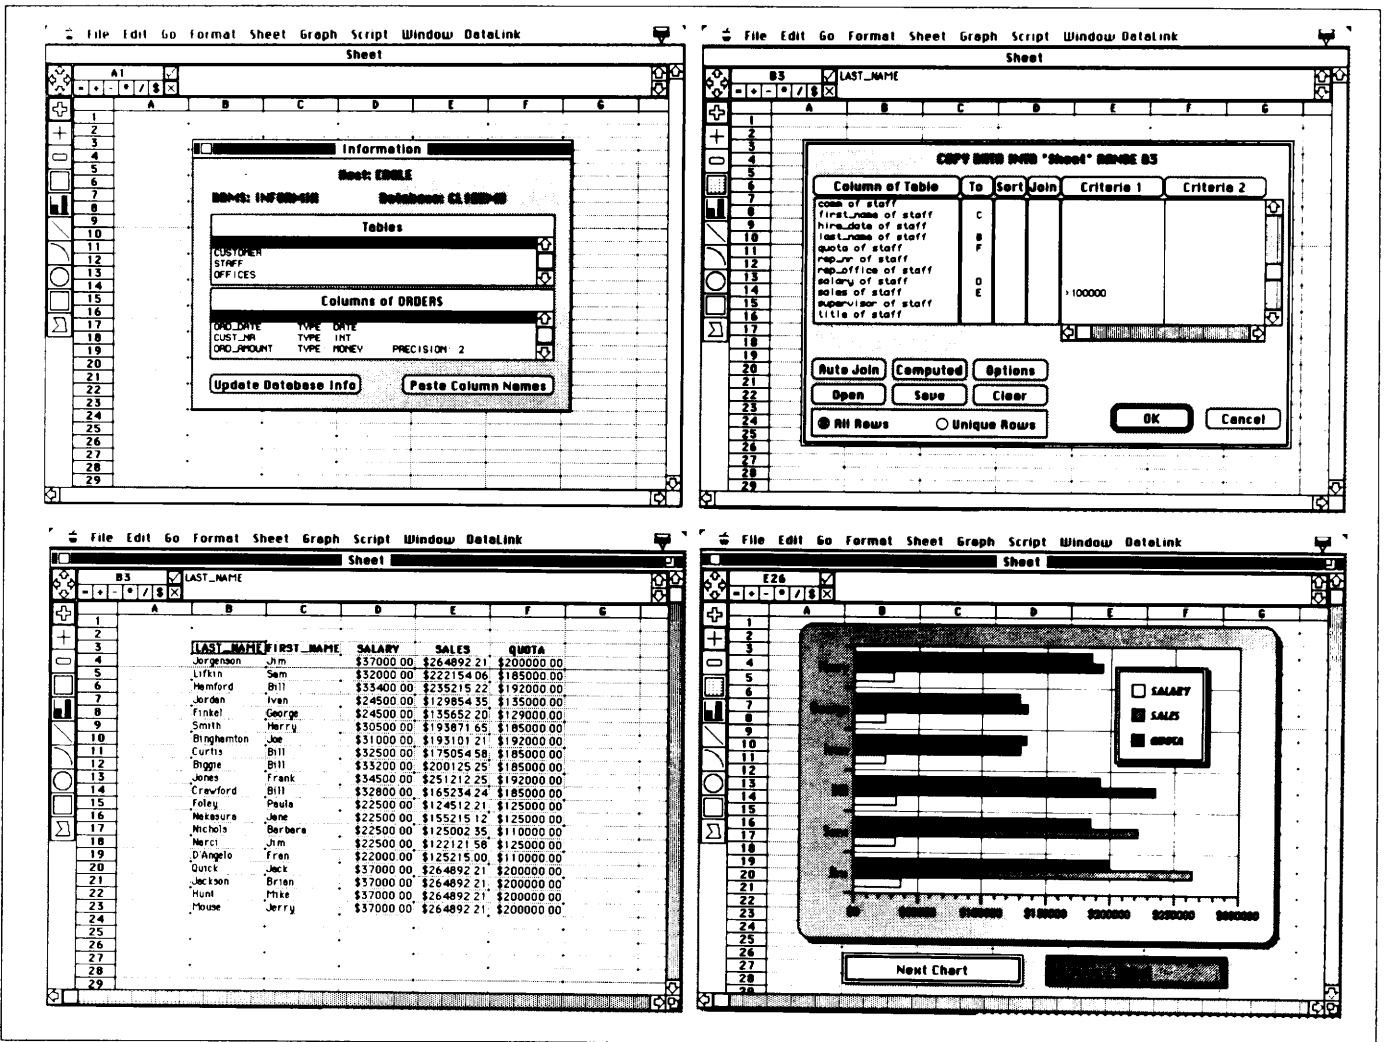
Illustration 1. These screen shots illustrate a typical set of activities that DataLink can provide—from browsing through available database tables (A), to creating SQL statements through point-and-click options (B), to viewing query results (C), to applying Wingz functionality to those results (D).

Architecture

FRONT-END STRATEGY. We stated earlier that Informix has taken a creative and practical approach to its tool strategy. Although it's not currently part of the integrated suite of RDBMS tools (which are character based), Wingz (and Hyper-Script) is being groomed as the Informix database graphical front end (form generator, application builder, menu generator, etc.). To let it participate in a distributed database environment, Informix has recently bundled DataLink with both Wingz and SmartWareII.

DataLink Dissapointement. Since Wingz was developed as a Macintosh product, DataLink has been based on the CL/1 network protocol, which connects Macs to remote VMS and MVS machines. In other words, if you have a CL/1 connection, you can use Wingz on a Macintosh as a front end to other databases supported by VMS—even, say, Oracle. How pleased we were when we heard this! How forward-thinking of Informix! Then we discovered that the Unix versions of Wingz use I-Net (Informix-Net), which only supports Informix databases. This decision stemmed more from the fact that CL/1 doesn't support Unix machines than the fact that Informix wants to limit Wingz's use of other database servers. Besides, CL/1 allows only rudimentary capabilities (e.g., it supports read-only data access). However, the I-Net protocol doesn't allow for any database interoperability whatsoever. Wingz is being marketed as an integral part of Informix's "end-to-end information management system," and the company doesn't want it supporting competitors' systems. We should mention that Informix is an active member of the SQL Access Group and hopes to attain interoperability through that channel. However, like most of its competitors, Informix is not pursuing RDBMS interoperability at the product level.

MULTIMEDIA. The other advantage of OnLine is its support for multimedia. The database can store BLOBs (Binary Large Objects) of up to 2 gigabytes in a field. BLOBs can be digitized information of any kind—documents, scanned images, faxes, video, audio, etc. Multimedia enhances a distributed database scheme by letting users share, store, and distribute more than just traditional data. (And it also gives Informix some slick demo material.)

OnLine's support for multimedia was enough to convince Plexus to use it as a foundation for its imaging system. The imaging system is impressive in its own right, with a distributed client-server architecture. Hewlett-Packard and AT&T, among others, recently licensed the product.

DISTRIBUTED ARCHITECTURE. With its solid implementation and positioning of Wingz and HyperScript and its support for multimedia, Informix seems to be addressing issues that some other database vendors haven't. At the same time, Informix is not up to snuff with the most pressing RDBMS issue—distributed databases. As we mentioned, OnLine is still missing centralized referential integrity—a must for distributed environments. Without some centralized deposit of integrity rules, you have no guarantee that applications have been implemented with the correct data protection. On-Line should also allow distributed transaction processing (i.e., multisite updates), which is usually handled with a two-phase commit proto-

> *With its solid implementation and positioning of Wingz and HyperScript and its support for multimedia, Informix seems to be addressing issues that some other database vendors haven't.*

col: After an issuing node makes a transaction request, all participating nodes must respond that they're ready before the transaction can actually occur. Although Informix plans to add these capabilities in the next release, OnLine currently supports multisite reads and single-site updates.

Informix-Star. Informix-Star, a client/server-based distributed environment, has two main characteristics: location transparency, and a distributed query optimizer, which decongests the network traffic by figuring out the best route for performing a query or a join (as opposed to just performing them at the first database server you connect to) and moves only the pertinent information to that site—not whole tables.

Informix-Star encompasses two pieces of communication software (see Illustration 2):

- Informix-Net, which lets applications access remote Informix databases on heterogeneous machines. The I-Net protocol supports various networks, and allows clients to perform only single-site reads and updates.

- LAN support, which lets networked DOS PCs act as a front end to an Informix database server. Tools installed on the

user workstation access a central RDBMS server. Informix supports Novell NetWare, 3Com 3+Open, IBM PC LAN, and AT&T StarGroup (see Illustration 3).

Gateways. In addition to Informix-Star, Informix offers gateways to IBM mainframe databases running under MVS, including DB2, IMS, IDMS, SQL/DS, Adabas, Total, and VSAM.

TRANSACTION PROCESSING. Distributed OLTP environments demand concurrency control and high availability as well as referential integrity and distributed transaction processing, and Informix has addressed these demands.

Concurrency Control. OnLine maintains a row-level locking mechanism to protect data from the inconsistencies that arise in a multiuser environment. If two processes try to get at the same row, the loser can either roll back or retry the transaction. For read consistency, OnLine allows *shared* locks, in which another process may read a locked row but may not update it.

Developers can choose from among several read-only locks, depending on the level of isolation they require:

- Dirty read, the lowest-level lock, reads without locking. Because it incurs no overhead, a dirty read won't guarantee that the data is accurate.

- Committed read guarantees that the data is accurate at the time the row is read. It also checks to see whether or not a shared lock can be obtained for that data.

- Cursor stability acquires a shared lock on a row only as it's being read, and releases it when the next row is read. It guarantees that the row won't be updated while it's being read, but not that it won't change as soon as it moves to the next row.

- Repeatable read acquires a shared lock on each row and holds it until the end of the transaction.

Availability. Online systems usually require that a database allow all backups, recovery, database changes, and other system maintenance to be performed while applications are running. Furthermore, if the server crashes, the database needs a reliable recovery system. Informix OnLine has a number of recovery features:

- Disk-mirroring, which lets you create an image of data in storage

Informix Connectivity Products

	Informix-Star	Informix-Net	Database Access	Informix on DOS LANs	Gateways
Type of Database Connectivity	• Distributed database and client/server • Application tool runs on client workstation • Database engines run on multiple servers	• Client/server • Application tools run on client workstation • Database engine runs on data-base server	• Client/server • Allows Wingz and SmartWareII applications to access data stored in SQL database	• Filesharing • Application tools and database engine run on the client work-station	• Database gateways • Informix's gate-way strategy will allow users to retrieve data from main-frames and store it in an Informix database on the client worksta-tion
Networks Supported	• TCP/IP • AT&T's StarGroup and StarLAN	• TCP/IP • AT&T's StarGroup and StarLAN • DECnet	• TCP/IP via Informix-Net • Ethernet, DECnet, 3270/SDLC, serial and asyn-chronous CL/1	• Novell NetWare • 3Com 3+Share • IBM PC LAN • AT&T's StarGroup and StarLAN	• LU6.2 protocol
Client Operating System Supported	• Unix • DOS (mid-1990)	• Unix • DOS • VMS	• Unix • DOS • OS/2 • Macintosh	• DOS	• Unix • DOS
Server Operating Systems	• Unix	• Unix • VMS	• Unix via Informix-Net • VMS and MVS under CL/1	• Not applicable	• Major main-frame operating environments
Features	• Multi-site read, single-site update access to distributed database • Cost-based (distributed) optimizer • Joins at any location • Location transparency • Site autonomy • Heterogeneous computers and operating systems	• Single-site read, single-site update • Location transparency • Heterogeneous computers and operating sys-tems	• Informix-Net provides interac-tive connectivity to any Informix DBMS support-ing Informix-Net • CL/1 provides access to a variety of host DBMS brands under VMS	• Low-cost, shared access to common databases • Location transparency • Multiple file server support • Optional passwords • Extended memory support	• Access to a variety of main-frame databases including IMS, DB2, IDMS, and SQL/DS • Access to VSAM files is also supported

Illustration 2. Informix connectivity software.

• Online and incremental archiving, which allows you to archive data while the system is online and users are accessing data

• Logical logs, which record all the changes to the database since the last backup

Should the system fail, the database can be recovered using either the mirrored data or the logical logs. Any interrupted transactions must be rolled back to the last checkpoint. (You can adjust the checkpoint rate to suit your needs; the default is 5 minutes.)

Using Informix

We evaluated the Informix system on a Sun 3/110. The OnLine database as well as its integrated tools—Informix-SQL Version 4.0, Informix-4GL Version 4.0, and Informix-ESQL Version 4.0—were installed on our workstation.

USER INTERFACE. Informix-SQL, the umbrella program that leads you into the Informix system, provides a menu-driven, integrated interface. Entering "isql" at the Unix prompt brings up the Informix-SQL main menu, which gives you access to its facilities: Form, Report, Query-Language, User-Menu, Database, and Table. Each component, incidentally, is also accessible at the operating system level.

Informix-SQL doesn't take advantage of graphical workstations (as we've mentioned, that's up to Wingz). The interface you get on a graphical machine is the same one you'd get on a character-based machine. You use Lotus-like ring menus and prompts to get around the system. In addition, Informix provides selection lists so you don't have to rack your brain trying to remember names of tables and forms. We liked the fact that the menus are intelligent. Once you've gone through the default form generation procedure, for instance, the system assumes you'll want to run it next, and thus highlights the "run" option. Then the selection list appropriately highlights the form you just generated.

Some Improvements. The major problem with the interface comes in using the form and report modules. You're forced into an editor to modify form and report specifications, which leaves you with no information about the task at hand.

Other than this obstacle, our complaints are minor. Unless you're in a text-entry screen, an "exit" option is available to back you out of a menu screen. If you're entering text, you rely on an interrupt key (Control-C). For consistency, we'd like to have an "exit" option on text-entry screens as well (and/or allow you to use Control-C to back out of menus).

We also found it somewhat counter-intuitive that, when filling in a form, the default is typeover mode, and most users would probably prefer insert mode. It would also be helpful if Informix-SQL made use of function keys to some extent—or at
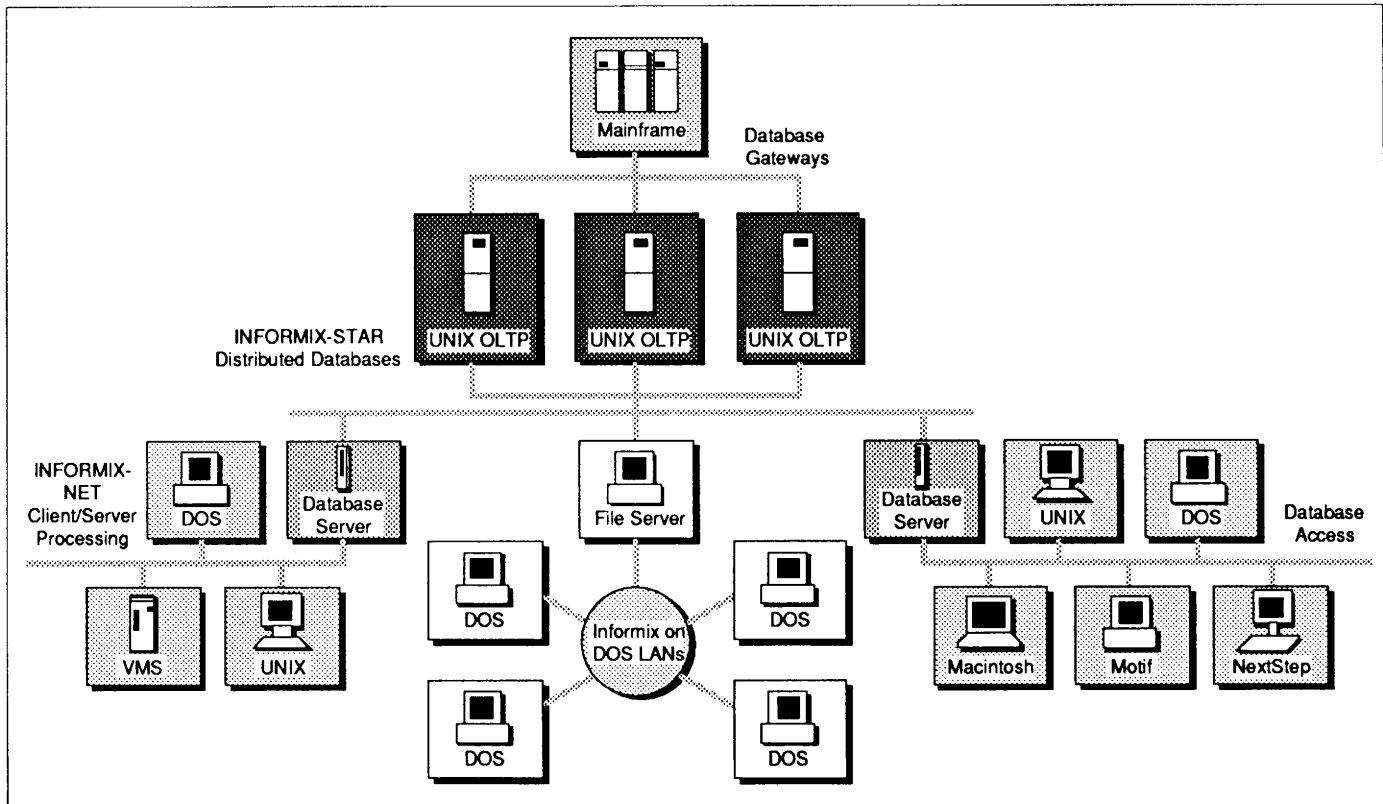

Illustration 3. The Informix distributed database architecture includes database gateways, a client/server architecture, SQL database access tools, and DOS LANs.

least had them available with a toggle key. While screens always tell you where to find help (Control-W), we had to hunt a little for the Interrupt key.

Another annoyance is that the system beeps at you a lot—often unnecessarily, such as when you hit Control-C to exit a text-entry screen.

Help. Informix provides full-screen contextual Help. In most cases, Help is adequate (although it's a little lean in some spots, like explaining transactions). Our reactions have more to do with manipulation than content. For example, you cannot back up to a previous Help screen. And, although we found the menus in the Help component of the query language module very useful, once you exit from one of these Help options, you should return to the Help menu rather than the query language menu. Otherwise, it's just a little keyboard-excessive.

Error-handling is well implemented. The messages we encountered were clear and specific. When you make a mistake, the error message appears at the bottom of the screen in reverse video. However, it doesn't beep at you. Why the system beeps for no apparent reason, and then remains silent when you make a mistake, we'll never know.

CREATING A DATABASE. The shortest route to database creation is through the menu-driven schema editor (although you can also use the query language). You select the database option from the main menu, choose "create", and name the database. For transaction logging, however, you'll need to write SQL statements specifying the full pathname of the transaction log. Then users can turn transaction-logging on or off.

CREATING A TABLE. Creating tables is also done through the schema editor. Prompts and menus guide you through the process. After naming the table, you enter the field name, its type, length, whether or not it's indexed, and whether or not it allows null (or unknown) values.

In addition to BLOB support, OnLine supports seven generic datatypes: character, number (including integer, small integer, float, and small float), serial, data, money, date/time, and time interval. OnLine also lets you define your own datatype—a very useful feature when you discover some oft-repeated field definition. User-defined datatypes not only save time when creating a new table, but they provide consistency among tables that use the same field definition.

Modifying Table Structure. You can change field definitions as you're creating a table. You can also do it later by using the SQL "alter table" command or by choosing "alter" from the Table menu. Either option allows you to add, rename, delete a column, or change the datatype—unless you're trying to

change serial type, that is. Informix won't let you do that from the menus (although you can do it in the query language). Serials don't allow nulls, and Informix automatically fills newly added fields with a null in existing records. The point is, if you want a serial type, make sure you do it when creating the table, and make sure you're not going to want to modify it.

FORMS. From within the Form option, you can either run a form for data entry or create one. The environment for creating a default form is similar to the database and table schema editor, where menus and prompts guide you through the process. To build one, you merely name it, identify the tables you want to include, and let the system automatically compile it.

Modifying Forms. The easiest way to design a form is to modify a default form, but you can always venture into your system editor and design one from scratch. To modify the form, you change the specifications and recompile it. You supply the database name, the layout of the way you want the screen to look (with prompts and brackets for field locations), table names, and field attributes.

> *It would also be helpful if
> Informix-SQL made use of function keys
> to some extent—or at least had them
> available with a toggle key.*

Added Function. You may include optional instructions in a forms specification to make the form more sophisticated. For example, you may include master-detail tables (which identify one-to-many relationships between two tables) to allow cross-table queries and subqueries. You may also establish composite joins and change field delimiters (the default is brackets), as well as control blocks, which is perhaps the most useful option. It allows you to incorporate pre- or post-processing instructions (including SQL commands and C functions) into your form for additional validation and database processing.

Not Exactly Intuitive. Designing forms is strictly a developer's task. Defining forms specifications is not fun. Informix leaves you in an editor with no direct access to the safe, comfortable confines of the forms generator. The default is the system editor—in our case, vi. Specifications mandate a specific syntax and format, so it's really a programming process (and a complicated one at that, if you're attaching instructions). While the system editor may be appropriate for the professional developer, we would also like the option to use a friendly editor with menus and prompts.

DATA ENTRY AND EDITING. Informix-SQL gives you two options for adding and editing data: screen forms or SQL commands. Those unfamiliar with SQL should obviously go the screen-form route. In screen forms, adding, updating, and removing rows from a form is simple enough, and it has the

Informix Features Chart

ARCHITECTURE	
Client/server	Yes
Multithreaded server	No
Support for symmetric multiprocessors	Yes
Open architecture (APIs available)	Yes

UNDERLYING FILE STRUCTURE	Raw I/O

DATABASE PARAMETERS	
Databases/server	No limit
Tables/database	No limit
Records/database	No limit
Fields/record	No limit
Record size	32K (except for BLOBs)

USER INTERFACE	
Menu bypass	Yes
Contextual help	Yes
Tutorial	No
Ability to customize standard menus	Yes

DATA TYPES	
Character (fixed/variable length)	Yes (fixed and variable length)
Integer	Yes
Float	Yes
Currency	Yes
Date/time	Yes (including time interval)
Binary (fixed/variable length)	Yes (support for 2GB BLOBs)
Long text	Yes
Image	Yes
Support for arrays	Yes
User-defined data types	Yes
User-defined functions and operators	Yes

B-TREE INDEXING	
Maximum number of indexes	16
Maximum number of fields/index	16
Maximum size of index key	256 characters
Order options	16 different order options
Unique index	Yes
Clustered index	Yes
Other file access methods	None

SCREEN FORMS	
Default form generator	Yes
Customized	Yes

SCREEN FORMS (continued)	
Multiple tables/form	Yes (up to 8 in default form)
Multiple screens/form	Yes
Embedded processing (if-then-else, display aggregates)	Yes

FIELD ATTRIBUTES ON FORMS	
Case conversion	Yes
Default value	Yes
Required field	Yes
Acceptable values	Yes
Verification (enter data twice)	Yes
Formatting of data	Yes
Calculated fields	Yes
Display only (no entry/update)	Yes
Hidden	Yes
Prompt (for data entry)	Yes
Error message	Yes
Customized help	Yes
Video display	Yes
Ability to change field attributes dynamically	Yes (in 4GL)

QUERY-BY-FORMS	
Exact match	Yes
Relational operators	Yes
Ranges	Yes
List of values	Yes
Wildcards	Yes
Maximum/minimum values	Yes
Print query results	Yes
Pass results to report writer	Yes
Text search	Yes

SQL	
Standard SQL statements	
Data definition language (DDL)	Yes
Data manipulation language (DML)	Yes
Query language	Yes
Extensions to SQL	
Commit/rollback transactions	Yes
Execute operating system commands	Yes
Load/unload data to/from ASCII file	Yes
Additional data definition statements	Yes
Control-of-flow logic	Yes (in 4GL)
Can be embedded in C/Cobol programs	Yes
Can create new table with query results	Yes
Stored queries	Yes
Case-insensitive (e.g., field names)	Yes
Can call C routines	Yes

SQL (continued)	
How create SQL queries/statements	Interactive SQL editor
Query optimizer	Yes
Syntax independent	Yes

REPORT WRITER	
Nonprocedural	Yes
Default report generator	Yes
Interactive report generator using screen forms	Yes (through QuickStep)
Interactive debugging	Yes
Input source	Yes
Multiple tables	Yes
Page formatting	Yes
Headers and footers	Yes
Data formatting	Yes
Sort data	Yes
Aggregate functions	Yes
Logical processing (if-then-else logic)	Yes
User variables	Yes
Prompt for input variables at run-time	Yes

APPLICATION GENERATOR	
Ability to design application menus	Yes
Default menu generator	Yes
Custom help	Yes

ABILITY TO CREATE VIEWS	Yes

TRANSACTIONS	
Logging	Yes
Commit/rollback transaction	Yes
Roll forward	Yes

INTEGRITY	
Referential integrity	No (at forms level only)
Field validation	No (at forms level only)
Support for business rules	No

STORED PROCEDURES	No

CONCURRENCY CONTROL— LOCKING LEVELS	
Database	Yes
Table	Yes
Record	Yes
Page	Yes
Other data isolation levels	Uncommitted read, dirty reads, cursor stability, and repeatable reads

DATABASE SECURITY	
Login password	Yes
Multilevel access control	
User	Yes
Group	No (programmable)
Application	No (programmable)
Database-level access	Yes
Table-level access	Yes
Record-level access	Yes
Field-level access	Yes
Access by time of day	No (programmable)
Access by location (workstation)	No (programmable)
Ability to define resource limits on user queries	No

AVAILABILITY	
Online backup	Yes
Online database changes	Yes

RAW INPUT/OUTPUT	Yes

I/O REDUCTION TECHNIQUES	fast commit, group commit, dual-log buffering

DATABASE CAN SPAN MULTIPLE PHYSICAL DEVICES (DISKS)	Yes

NETWORK SUPPORT	TCP/IP, async, X.25

DISTRIBUTED DATABASE CAPABILITY	
Location transparency	Yes
Distributed query processing	Yes
Distributed query optimizer	Yes
Distributed transaction processing (two-phase commit)	No
Support for data replication	No
Access to heterogeneous databases	Only through gateways (IMS, DB2, IDMS, VSAM, SQL/DS, Adabas, and Total)

INTERNATIONAL LANGUAGE SUPPORT	
Sorting sequences	No
Error messages	Yes
2-byte character set	Yes

additional benefits of optional instructions and validity checks. However, these checks won't affect data entered with an SQL command. Since OnLine doesn't support centralized data integrity rules, they must be included in the application.

INDEXING. Informix supports B-tree indexing, including a clustered index, where, for performance reasons, the physical order of the table matches that of its index. It also supports both ascending and descending ordering. The index can be based on multiple columns but cannot be longer than 120 characters. Informix indexes can be defined as unique and can be dynamically created or dropped.

QUERIES. You have two options for querying: query-by-forms or SQL commands. To query-by-form, you bring up the screen form, enter your search criteria, and select "query." Informix places all matching records in a temporary storage area called the "current list" and displays the first match it finds. From there, you can scroll back and forth through the list. If you want to select all the records in a table, you leave the search criteria blank. An "output" option writes the results to a file.

It'd be nice if the form also included the SQL rendition of your query. That way, users could learn SQL as they go. Some form-querying products we've seen do that; in fact, Wingz lets you view queries as SQL statements. Informix should consider adding this feature to its integrated tool as well.

SQL. You can create SQL queries within the menu system or through your system editor. Creating and modifying an SQL statement using the menus has been implemented in a very interactive and easy-to-use fashion. To begin with, you do not use the same commands as the Unix vi editor; instead, Informix has its own set of editing commands (done editing, delete character, typeover/insert toggle, delete rest of line, and redraw), which are displayed at the top of the screen.

If you make a mistake after you've run the command, Informix will show you the query again and give you an error message. As is the case with the other Informix-SQL modules, the menus lead you through the whole process. You can create very complex nested queries and then embed them in a C program. However, given the complex syntax and logical structure of nested queries, you won't learn how to do them in a day.

REPORTS. Informix's standard report writer is a thorough tool with advanced formatting capabilities. However, like the form generator, it is a little rough to use, because you're stuck with writing a program (of sorts). No doubt, QuickStep was developed as a result of the drudgery involved in creating even simple reports within Informix-SQL.

The first step is to create the report specifications: the database name, SQL statements that indicate the fields and rows to appear in the report, the tables where those fields and rows will live, instructions for sorting the tables, and formatting instructions (e.g., page header, control breaks, variable properties, etc.). Again, you have the option of starting with default report generation. It may not contain much information (data from just one table and a mere list of every field in the table for each record), but it will at least get you started. From there, you can add comments, data from multiple tables, page headers/footers, column headings, margins, page length, page numbers, group headings, arithmetic operations on numeric fields, fill characters, and aggregates.

Once you've finished the process, you can print it, write it to a file, or use it for input to another program.

CUSTOMIZING APPLICATIONS. Aside from Informix-4GL, the primary tools for application development are the Informix-SQL form generator, report writer, and user menu. We already mentioned that the form generator has optional instructions that add processing to a form. The report writer can also create different views of a database to meet specific requirements or restrictions. You wouldn't, for example, want just anybody peering into a database that contains salary information. So you would create a view that's actually a subset of the database that users can manipulate as if it were a separate table of information. (Views are created with the query language; they are not included on the menus.)

Other than that, each database can have one customized user menu. This menu can call submenus, other Informix-SQL modules, other programs, or operating-system utilities.

SECURITY. As distributed environments continue to proliferate, Informix should upscale its security options. The database provides various levels of data security using the standard "grant" and "revoke" commands, but it doesn't acceptably address access to groups or applications. You can reach these levels of access, but only if you program them.

Conclusion

FUTURE DEVELOPMENT. Informix has a number of plans on its agenda. The company acknowledges that it has work to do on its distributed architecture. It at least intends to implement multisite updates as well as referential integrity in the next release (most likely mid-year 1991). The next release will also include free-text search for variable-length text-based data.

Also on the drawing board are improved security and performance. Security in a distributed environment is an obvi-

> *While the Unix system editor may be appropriate for the professional developer, we would also like the option to use a friendly editor with menus and prompts.*

ous necessity. And, while the OLTP engine offers sufficient performance for typical transaction-based operations, speed will become more of an issue once people start trying to store and retrieve multimedia.

Informix also plans to release later this year an updated DOS engine. The company is positioning the engine to play in a distributed mixed-hardware architecture, although it will be worthwhile as a standalone machine as well, with transaction and BLOB support.

As far as front-end improvements go, Informix is working on pumping more functionality into HyperScript to make it a more appropriate data-access and application-building tool. It will add, for instance, a forms generator and schema editor. Likewise, Informix eventually plans to create a windows-based version of Informix-4GL.

OVERALL ASSESSMENT. Informix has some notable strengths, especially when you consider the completeness of its toolset. We've already drilled home the advantages of using HyperScript to front end the database. Also worth noting is Informix's rich implementation of SQL. When you couple the query languagae with Informix-4GL, C programming is rarely required. The Informix-SQL interface also impressed us. From the end-user's perspective, the interface rates favorably—particularly because of its menu implementation, interactive SQL processor, and integrated error-correction process when creat-

ing forms, reports, and SQL queries. We also found a comfortable level of consistency among the different functions. The processes of creating forms, reports, and queries have many similarities in content and structure.

But Informix could better integrate its modules. Each is a separate program connected through the user interface on the front end and the common database on the back end. For example, records extracted via query-by-forms cannot be sent to the report processor as input. This level of integration would be advantageous for extracting records and easily passing them through for formatting and output.

It's a little more difficult to assess the database core itself. While its support for multimedia is attractive, its distributed architecture limitations weigh the product down. When developing OnLine, Informix regarded BLOB support as a more immediate concern than distributed databases. We're not so sure about that decision. A distributed architecture—complete with multisite updates and referential integrity—is quickly becoming a requirement. In any case, Informix plans to complete its distributed architecture in the next release. And, when it does, that functionality will be available on just about every Unix platform—not just a select few (as is sometimes the case with other Unix software products). Once Informix has furnished these enhancements, the company will be well on its way to maintaining a competitive position in a challenging DBMS environment. **☾**

NEW SPECIAL RESEARCH STUDIES

from Patricia Seybold's Office Computing Group

Object Orientation: Concepts and Implementation
by John R. Rymer
$495

Compound Documents: Vendor Strategies, Architectures, and Applications
by Michael D. Millikin
$495

Open Software Foundation and Unix International: Implementing the Open Systems Agenda
by Judith S. Hurwitz
$495

Executive Support Systems: Concepts, Tools, Techniques
by Mickey Williamson
$495

The AS/400: A Midlife Assessment
by Norman Rasmussen
$495

Call 1-800-826-2424 to order or to receive more information on these reports (in Massachusetts, call 617-742-5200). Or fax your order or request to 617-742-1028.

NEWS

PRODUCTS • TRENDS • ISSUES • ANALYSIS

ANALYSIS

• DIGITAL •

Digging Deeper into Unix

On the face of it, Digital's DECstation and Ultrix announcement may appear to be yet another attack launched at the competitive and cutthroat RISC workstation market. More hardware, more MIPS and graphics power. Indeed, that is an important part of the story. But there is much more to this announcement. Unfortunately, Digital may have missed the opportunity to strut its technical prowess in such areas as symmetric multiprocessing and distributed computing by shoving too many pieces of its strategy and technology at its customers at once.

So what does the Digital announcement mean? Let's start with the obvious. Digital is attacking the workstation market from both the low and high ends. To compete against the low-end leader, Sun SPARCstation, Digital has reduced the prices on its existing products, bringing the DECstation 2100 workstation to below the magic $6,000 ($5,950). On the high end, Digital introduced a new workstation and server, the 24-MIPS DECstation/system 5000 and multiprocessor versions of the 5800 workstations and server family.

Unfortunately for Digital, the announcement comes right on the heels of IBM's RS/6000 announcement. Therefore, users automatically assume that the new products have been targeted as Digital's answer to the RS/6000. They are not. Rather, the new high-end products help fill out Digital's RISC family, while the price cuts are an attempt to extend the life of Digital's older workstations and servers. These products are enough to keep Digital treading water in the rough seas of the RISC workstation market. We expect that, by the end of the summer, Digital will be bringing out newer RISC-based machines to compete against IBM.

The most important parts of Digital's announcement are the enhancement's Digital has made to its Ultrix operating system. The most significant of these include the implementation of full symmetric multiprocessing and sophisticated distributed computing technology. Digital also inched closer to supporting AT&T System V by offering a version aimed at the telecommunications industry that insists on the System V brand. Digital may find that it will generate a lot of interest and speculation with this offering.

Ultrix 4.0

While Ultrix Version 4.0 is still based on BSD, Digital is emphasizing that it is adhering to all the standards it can get its hands on. The laundry list includes IEEE Posix 1003.1 (1988), ISO 9945-1:1990, and FIPS (Federal Information Processing Standard) 151-1. Digital also promises to demonstrate compliance through the NIST (National Institute of Standards and Technology) Posix conformance test suite, once it is available and stable. Ultrix already conforms to the base-level specification of the X/Open Portability Guide (XPG 3).

FULL SYMMETRIC MULTIPROCESSING FOR ULTRIX. Digital has provided sophisticated full symmetric multiprocessing (SMP) under VMS for several years. It has promised equivalent technology under Ultrix ever since. Now, Digital is delivering SMP as part of its new version of Ultrix (Version 4.0). Digital's implementation of SMP allows for N-way multiprocessing. Therefore, there is no theoretical limitation to the number of processors that can be tightly coupled. Initially, Digital is offering four-way multiprocessing, but we expect larger configurations to follow shortly. One particularly nice feature of SMP is the fact that the operating system does not have to run on all the processors and can be distributed among all the processors. Initially, SMP will be provided for the 5800 DECsystems, the VAX 6000 systems, and all multiprocessor VAX 8000.

DISTRIBUTED COMPUTING ENVIRONMENT. Digital is adopting HP/Apollo's Network Computing System's (NCS) RPC, including its Network Interface Definition Language compiler and its Location Broker. Digital continues to work with HP on a wide area network (WAN) implementation of NCS. This would mean allowing NCS to operate over X.25 and OSI networking protocols.

Digital has a rich set of services for distributed computing and is finally beginning to explain that fact to users. It is adding some additional technology, such as MIT's Project Athena, the Kerberos Authentication Service, and its Hesiod Name Service for managing a centralized password distributed database. Digital has also released a Network Time Protocol (NTP) that enables synchronization of all clocks in an Internet network. All three of these are important components in a true distributed computing environment (see *Network Monitor*, Vol. 5, No. 1).

While IBM also will provide Kerberos within AIX, Digital has gone further in providing these other services ahead of most competitors.

DECnet-Ultrix Version 4

Digital has enhanced the DECnet capabilities on Ultrix by improving installation procedures so that the kernel does not have to be rebuilt during installation. Also, Digital has added the License Management Facility (LMF) for Ultrix. Digital's LMF lets users increase the number of software licenses without actively installing that software. The LMF keeps track of the number of end-user licenses and bills the customer accordingly. Also added are enhancements to X.25 networks using the X.25 router 2000 and DECnet datalink mapping. Digital has added NCP to support X.25 protocol entities such as network, group, and destination.

One of the most significant aspects of the design of DECnet is that it provides the same Phase IV protocols across Ultrix, VMS, and non-Digital Unix systems that use TCP/IP.

VMS/Ultrix Connection V1.3

VMS/Ultrix Connection V1.3, a VAX/VMS layered software product, is intended to help users share resources, files, networks, and application development between VMS and Unix systems. It allows Unix systems to exploit the availability and reliability of VAXcluster systems using TCP/IP. Part of this software includes support within VMS for communication protocols traditionally found only on Unix systems, including HP/Apollo's NCS, Routing Information Protocol (RIP), and BSD Internet Name Domain Resolver (BIND).

SECURITY. Security under Ultrix 4.0 has been brought up to C2 level. While this is a good start, we expect that Digital will be under increasing pressure to surpass C2 and aim for B1 or B2.

OTHER ENHANCEMENTS. Digital has increased the number of disks supported from 32 to 96 per system (for up to 144 GB of storage). Inclusion of SQL run-time libraries is an operating system feature. This is expected to encourage developers to write to Ultrix because the developer does not have to purchase the run-time SQL library.

Database Access. Database access has been a weakness for Digital—especially in the Unix arena. Although Digital says it will develop a version of Rdb for Ultrix, it has been slow to provide that functionality. But, because of competitive pressures, Digital has hired Ingres to develop a relational SQL engine that can provide remote access to Rdb on VMS. Ingres also provides Digital with sorely needed front-end tools. As a marketing move, Digital is including Ingres's SQL run-time engine at no additional cost.

The downside of Digital's arrangement with Ingres is that it alienates other relational database vendors. We suspect that Informix, Sybase, and Oracle—to name a few—will not be pleased with this cozy partnership. Digital might have been wiser to pick an object database and allow any relational database to hook in.

Digital's next move in the database arena will be to provide its Builder product for Ultrix. Builder is a tool to allow users to define relationships among data elements from multiple data sources. It includes tools for creating cross-application macros. Because it relies on Rdb, this powerful tool is available only on VMS today.

System Management. Although Digital is not licensing Apollo's License Broker, it is providing software management tools to customers. Digital also has introduced Diskless Management Service (DMS) and the Remote Installation Service (RIS) to allow a single system to service both RISC and VAX systems at the same time. The RIS server, for example, performs the boot request on the local Ethernet and uses NFS to access RIS files on the full RIS server, as well as facilitating remote installations across gateways and wide area networks.

Unlike some of the operating systems available from other vendors, Digital OS for workstations (called Ultrix Workstation Software) bundles in a lot of good technology, including DECwindows, Adobe Display PostScript, Digital's Compound Document Architecture (CDA) libraries, DECnet, TCP/IP, NFS, a "quick screen refresh for overlapped windows," and support for monochrome, gray-scale, four-plane and eight-plane color systems. All of these are provided standard with the hardware. Digital is also preloading the operating system onto all the workstations it sells. This action is geared to helping Digital compete in the commercial Unix market.

OFFICE SOFTWARE. Digital is not ready with its home-grown Unix office products. In the meantime, it is offering two popular packages: Uniplex for the

New RISC System Hardware

The new members of Digital's RISC family are based on the 25 MHz MIPS R3000 processor and R3010 floating-point unit. The new systems include several workstations and several servers.

WORKSTATIONS

Model: DECstation 5000 Model 200

Focus: Compute-intensive 2-D and 3-D graphics applications

Performance: 24 MIPS; 6.4 Mflops (single precision), 3.7 (double precision)

Price: Ranges from $14,995 (for CX) to $51,100 (for PXG Turbo)

Configurations: 200CX (entry-level graphics), 200PX (2-D graphics), 200PXG (3-D modeling; includes an i860 graphics accelerator), 200PXG Turbo (3-D modeling and visualization)

SERVERS

Model: DECsystem 5000 Model 200

Focus: General purpose, multiuser RISC system, file server, or compute server

Performance: 24 MIPS; 6.4 Mflops (single precision), 3.7 (double precision)

Configurations: Base configuration includes 8MB of ECC Memory (expandable to 120MB), diskless

Price: $14,995 (including operating system and 8MB)

Model: DECsystem 5830 (three-way multiprocessor) and 5840 (four-way multiprocessor)

Focus: Multiprocessor implementations of the 5810 DECsystem

Performance: 5830, up to 48 MIPS; 5840, up to 62 MIPS

Configurations: Maximum memory is 192MB (5830); 128 (5840). Maximum Storage: In-cabinet 2.4GB (5830 and 5840), local 58GB (5830 and 5840)

Pricing: 5830 starts at $140,000; 5840 starts at $160,000

THE TURBOCHANNEL

Digital has added a new 32-bit synchronous I/O channel bus it calls TurboChannel. It operates at a speed of 100 MB per second at between 12.5 and 25 MHz, and connects to the existing VME or SCSI bus. Digital promises to be able to interface the TurboChannel to the IEEE Futurebus by 1992. The DECsystems include three TurboChannel slots. To encourage third parties to write to the interface, Digital is charging no royalty to board and system vendors.

multiuser systems and Applix's Alis for workstation users. Both products will fill a gap. Digital already has its DECwrite compound document editor available under Ultrix. In addition, DECwindows already provides a good set of facilities including calendar, mail, and a draw and paint program. This could be the basis for a respectable office package.

Conclusion

At a recent consultant's briefing when these systems were previewed, Digital's management stated that the company's goal is to become the number one vendor of open systems in the '90s by delivering "leadership RISC workstations." It is a lofty goal and one that will find some customers scratching their heads. Conventional wisdom says that Digital isn't very interested in the overall Unix market.

Speeches from company guiding light Ken Olsen help keep that perception alive. If one listens closely to what Olsen says, he makes some excellent points and philosophical observations about Unix and the confusion over the term "open." But it is inevitable that these talks will be misinterpreted. The business press is as interested in "sound bites" as the rest of the general press. Since Digital has no visible Unix spokesperson, any statement Olsen makes is taken the final word in corporate Unix strategy.

Digital has an incredible wealth of technology and know-how in both general computing and Unix. The company's biggest problem is that it assumes that if it engineers good systems and software, customers will find out. It's somewhat like the old saying, "goodness is its own reward." Unfortunately for Digital, in the cutthroat workstation and systems marketplace, goodness isn't always appreciated. Especially when the goodness is buried deeply so that customers don't even know what Digital has to offer them.

—*J. Hurwitz*

Enough Is Enough

The briefing invitation was from Sun Microsystems. In it, Sun promised to "shatter the myths" surrounding the Unix user interface squabbles.

Our reaction: "Oh, no! Another Open Look pitch." We had been relieved these past few months because peace had seemingly broken out in the graphical user interface (GUI) wars, which heretofore had been seriously interfering with more important developments. This may have been wishful thinking, as the GUI controversy is still slowing down application evolution, and also apparently helped nix the merger discussions between Open Software Foundation (OSF) and Unix International (UI).

Nevertheless, we went along. Sun *is* very vocal about the advantages of using Open Look (or OpenWindows) as an interface environment. And, to give credit where credit is due, OpenWindows does have its merits.

OPENWINDOWS MERITS. Open-Window's major selling point is that it's a more complete system than Motif. Other differences—like the appearance of the desktop metaphor, or the implementation of menus or buttons—are more trivial. Is it really essential whether or not a button has a 3-D representation? Or that a cute little hand drags icons and windows across a screen? The functionality of the interface environment, on the other hand, is essential. What OpenWindows has that Motif hasn't is desktop tools and font technology. Sun is also touting its developer's guide, an interface generator of sorts. (Motif has a developer's guide as well, but it's not an interface generator; it's more of a style guide.)

DeskSet. OpenWindows' DeskSet is a suite of Open Look-based desktop tools for Sun workstations only. Because it's designed for Sun's hardware, it sup-

ports a "drag-and-drop" metaphor, which is trickier to implement generically. Portable software would need to address a multitude of inconsistencies in various system resources and services to support drag-and-drop.

The DeskSet includes a graphical file manager, a calendar manager, print services, a tape tool (to read, write, or list data from a local or remote drive), a calculator, mail, a binder (which allows you to group applications or user actions into color-coded icons), a snapshot tool (which takes a snapshot of the screen or of part of the screen), a mouse-driven ASCII text editor, an icon editor, a system performance meters, and a clock.

These are the kinds of things offered in Motif-based Unix shells such as X.desktop from IXI or Looking Glass from Visix. Actually, many such tools are available from plain X Window (e.g., calendar, calculator, print services, mail, ASCII text editor, and clock). We should point out that OSF never intended Motif to be anything more than a user interface. Open Look doesn't have a toolset, either.

OpenFonts. Included in OpenWidows is OpenFonts, Sun's answer to Post-Script functionality for its XView toolkit. OpenFonts lets you chose from among 57 fonts, which are scalable and rotatable. OpenWindows uses Open-Fonts to generate bit-mapped images or grid-fitted outlines. To penetrate the market, Sun is licensing OpenFonts to font suppliers, hoping they will convert their outline libraries to OpenFonts format. And some of the leading type manufacturers are doing just that (among them, Linotype, Monotype, Bigelow & Holmes, and Berthold).

Developer's Guide. Sun is also offering a user interface prototyper—or screen generator—as a separate product to be used in conjunction with OpenWindows. With it, you can design your interface from a palette of icons that represent various objects (e.g., window control functions, scrollbars,

menus), and let the program bang out the Open Look code for you. This eliminates relying on a toolkit to construct a user interface.

LURING ENEMY DEVELOPERS. Within Sun's user interface environment is XView, Sun's portable Open Look-based X toolkit. Keeping track of Sun's window systems and toolkits is becoming daunting, isn't it? There's NeWS, X11/NeWS, Open Look, OpenWindows, and SunView—which has been rewritten for X and is now XView. XView will be bundled with every machine Sun sells. Sun essentially wrote XView to leverage the large number of SunView applications for an X environment. Since XView is compatible with SunView, SunView applications become X applications with a minimal amount of work. (Actually, it must have been a difficult task for Sun, since SunView was a kernel-based windowing system and XView is client/server-based.)

The company has started an XView standardization campaign of sorts by porting the toolkit to IBM, Digital, and Hewlett-Packard platforms—that's right, the big Motif supporters. These versions won't be available until later this year, and, to gain market share, Sun will make the licenses free of charge. In the meantime, XView is available on the X11 tape and is shipped with Unix System V.4.

CONSEQUENCES. A few days after our briefing with Sun, we attended a technical seminar on Motif, which made for an interesting counterpoint. Sun jumped all over the weaknesses of Motif, yet OSF representatives didn't mention Open Look or OpenWindows at all. Obviously, OSF has no doubts about Motif's status as a standard. Sun, of course, claimed that the perception that the industry is swarming around Motif as a standard is incorrect. And, in some ways, Sun has a point. Compared to Motif, Sun supports a signficant number of graphical applications. This is largely because SunView has been around a lot longer than Motif, as has

Sun's popular graphical workstation. However, you have to take into consideration that most workstations—and applications—out there are still character based. The flood of GUI applications has yet to materialize. We anticipate that, as graphical workstations become more mainstream, so too will Motif-based applications.

But producing graphical applications may take some time if application vendors need to write to two different APIs under Unix alone. We don't see why OSF and Unix International can't come to some settlement where user interface is concerned. As we said, this is a small issue. The differences in basic functionality between Motif and Open Look are not all that significant. The distinctions are application driven. Meanwhile, standardization is crucial. And, frankly, we're more interested in products that bridge rather than extend the GUI wars.   *—L. Brown*

▪ S O L B O U R N E ▪

Tackling GUI Portability

While some vendors are being stymied by the instability of the Unix graphical user interface environment, others, such as Solbourne Computer, are using it to their advantage. Solbourne, the vendor of high-end Sun SPARC clones, has recently developed a commercially available X-based toolkit that's independent of both Motif and Open Look (and, conceptually, any other X-based toolkit). Thus, an application doesn't need to be written to multiple toolkits; you can pick whatever interface you want at run-time.

MULTIPLE LOOK AND FEEL. Solbourne's toolkit, Object Interface, has a multiple look-and-fee. However, the product grew more out of Solbourne's need for a graphical application development environment than out of the industry's current graphical user inter-

face wrangling. The company started working on Object Interface in 1987—long before any interface squabbles erupted. Solbourne decided at the time that it didn't want to be dependent on another vendor's technology—which turned out to be a wiser decision than it may have seemed at first because it left Object Interface the opportunity to evolve into an open technology.

Applications are written using objects in an Object Interface class library. At run-time, the objects are bound to either Open Look or Motif instantiations. Therefore, the developer writes to a single API, and the user has the flexibility of using the interface of his or her choice.

Technically, Object Interface works with the concept of generic objects (see illustration). Applications have no notion of the particular model the object belongs to. All model-specific objects are derived from a generic object, and the application interface for all model-specific objects is the same. In other words, the appearance and behavior of the objects may be different, but the application interface to them is the same.

BUILT ON C++. Although Object Interface uses X as its underlying window system, its object library is based on C++ (V.2.0) rather than C, which is what most X toolkits use. Solbourne developers view C++ as a robust object environment—and it's somewhat more logical and pliable. Object Interface, for instance, features dynamic reparenting (i.e., objects are dynamically interoperable). It works like this: When you delete an object, its subobjects (or children) are moved to an "orphanage" so they can be used later (or reparented) rather than being destroyed. Dynamic reparenting allows for things like tear-off menus (which perhaps you've seen in the NeXT system), where you can "tear off" a submenu from a menu object and then use it as a main menu.

VIRTUAL DESKTOP. Included with Object Interface is the Solbourne Win-

dow Manager, which handles window manipulation and, unlike other window managers, has nothing to do with look and feel (you can configure it to Motif, Open Look, and TWM). In addition to window management, it provides basic session-management facilities.

However, the most interesting feature of the window manager is its Virtual Desktop facility. Essentially, Virtual Desktop makes the X root window larger than the physical limits of the display. Once the root window is expanded, you can put together a "rooms" type of environment, where groups of related windows can be moved into various spaces on the Virtual Desktop. Or you can make use of "sticky" windows—windows that don't move as the Virtual Desktop is scrolled. Sticky windows are a handy feature if you want to move to a different area or room without leaving a preexisting environment, or if you want to keep a standard environment (e.g., a clock, a mail notifier, etc.) visible at all times.

HYPERTEXT HELP. Object Interface also features a well-implemented hypertext Help object. Each Help item includes additional menu items for attaining extra information—subtopics or related topics. The Help system also has a couple of especially impressive aspects. One is its degree of context sensitivity. If you have Help running in a visible window, you can watch the Help window change according to what you're doing. Another is that the objects within a Help item are real objects—not just bit-mapped representations. Thus, you not only get a description of what an object is, you also get to view its behavior. For instance, the scrollbar Help item includes a sample scrollbar that's live—it actually scrolls.

AVAILABILITY. Solbourne has licensed the product to AT&T USO and also plans to bundle the Object Interface library on Solbourne machines. It currently runs on Solbourne and Sun (3 and 4) machines.

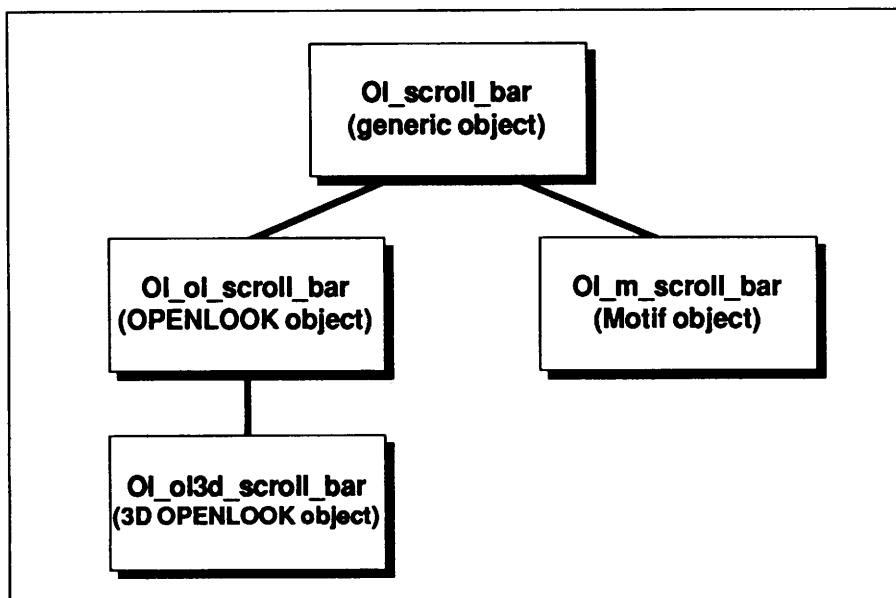Compliance to Motif and Open Look is not yet complete. Solbourne

Illustration. Sample class structure.

expects Open Look compatibility by summer and Motif by fall.

ISSUES. There are some inherent limitations in the implementation of a toolkit such as Solbourne's. Most obvious, perhaps, is the fact that Object Interface only accommodates X. Perhaps the GUI wars are most evident in the Unix marketplace, but developers have to address other systems as well—which means they're looking at a whole bunch of graphical development environments, not just Motif and Open Look. We're not holding Solbourne accountable for consolidating all graphical development systems out there; it's just something to be considered.

Furthermore, to accommodate both Open Look and Motif, Solbourne had to fudge a few object translations. Object Interface is actually a combination of the functionality of the two toolkits. Therefore, it has a superset of the objects that you'll find in either Open Look or Motif alone. If an Open look object, for instance, has no obvious Motif counterpart, Object Interface extends the Motif look and feel by using one of its superset objects.

But we're not about to complain. Since Solbourne is only dealing with Open Look and Motif, it won't suffer

much from the lowest common denominator syndrome. Toolkits that support additional window systems—such as the Extensible Virtual Toolkit from Advanced Programming Institute, which supports X, MS Windows, Macintosh, and Presentation Manager—have more of a problem here. But regardless, these kinds of products have a worthy goal—lowering the cost of porting graphical applications. Certainly 90 percent portability is better than no portability at all.   *—L. Brown*

• N O R T O N •

You Need Unix Utilities

As part of its multiplatform product strategy, Peter Norton Computing, Incorporated (Santa Monica, California), has announced its intention to develop products for Unix workstations.

In 1982, Peter Norton Computing introduced the Norton Utilities for the PC market. The software gave users the ability to recuperate from many kinds of computer errors, disk crashes, etc. with minimal loss of information. Today, Norton Utilities is probably the

most popular set of data protection software available. And soon those utilities will be available on Unix platforms.

MORE THAN JUST A PORT. Rather than simply porting its current products to Unix, Norton intends to redevelop its products to make use of Unix's capabilities. And in order to gain the know-how to accomplish this task, Norton has formed a partnership with Interactive Systems. Interactive will provide its knowledge of Unix, while Norton will provide knowledge of developing utility software.

Norton is smart to engage Interactive's help. Developing products for Unix will be far different from developing products for single-user systems like DOS or the Macintosh. For instance, one of the standard Norton Utilities is "UnErase," which recovers accidentally erased files. This is a (more or less) straightforward process in DOS or on a Mac, where the utility simply has to scan the hard disk(s) for information. But what happens in the multiuser Unix world? UnErase becomes a much more complicated process. The utility must start to take into account which version and which user's version of a file is meant to be unerased. And that starts to bring up security implications.

IS THERE A MARKET FOR UNIX UTILITIES? There are certainly enough DOS machines and Macs out there to warrant development of utilities for those platforms. But is there enough Unix usage to warrant this type of development on that platform? The answer is: Probably.

Unix may never become the operating system of choice on the world's desktops, but it will certainly be a strong contender. And the current developments in the Unix community of powerful and useful graphical interfaces will bring more and more "average" end users to the platform. Those users will need (even if they are not aware of it) the support of utilities such as Peter Norton's.

No matter what the platform, hard disks always crash, data on floppies always gets corrupted, glitches always occur. A user may do nothing until a major problem takes place, but, once a problem does, the user is likely to clamor for data protection support—like Norton Utilities.

CONCLUSIONS. Norton's intention to develop on Unix is both an indication of the platform's growing acceptance by the general business community and a strategic move by Norton to be the first major utilities developer on an up-and-coming platform. As more people use Unix, more support for those users is needed. And the better support there is for Unix users, the more people will think about selecting Unix as their platform of choice. So the better Norton supports Unix, the more Unix users Norton can look forward to having as potential customers. A beneficial Catch-22, if you will.   *—D. Freeman*

Patricia Seybold's Computer Industry Reports

Important: This report contains the results of proprietary research. Reproduction in whole or in part is prohibited. See back page for additional copy information.