

COMPUTERWORLD

\$2.00 A COPY; \$44/YEAR

OCTOBER 21, 1985

VOL. XIX, NO. 42

BUSINESS LIBRARY
Univ. of Wisconsin - Madison

BUSINESS LIBRARY
NOV 21 1985

In Depth

E. F. Codd puts three relational DBMS to the test/49



In Depth

Are you spending the right amount on computers?/61

Change of pace

New Product section starts this week/65

CICS/VM

IBM fits CICS to VM/CMS/6

Intel sends 32-bit chips to vendors

By Clinton Wilder

Intel Corp. last week took the wraps off the initial entries in its much-heralded 32-bit product line, the 386 family.

Although mass production of the 80386 chip is not expected until late 1986 or early 1987, several vendors are expected to begin developing a new generation of computer hardware based on the Intel architecture. Representatives of a number of these commercial customers joined an Intel news conference in San Francisco last week.

Leading the parade of new Intel products was the 80386 microprocessor chip, the largest and most powerful ever made by Intel. Available now as a prototype to systems developers and expected to start shipping in volume in the second half of 1986, the 80386 contains processing power of 3 million to 4 million instructions per second.

Available in versions running at 12 or 16 MHz, the 80386 runs software compatible with all previous Intel microprocessors — the 8086, 8088, 80186 and 80286. The new chip supports simultaneous applications running on Microsoft Corp.'s MS-DOS or Xenix or on AT&T's Unix System V operating systems.

Representatives of several companies announced they have already made commitments to using the Intel chip in future product lines.

Conspicuously absent from the committed group was IBM. However, William Lowe, president of IBM's Entry Systems

Continued on page 4

IBM token-ring net bows; software prices boosted

Tariffs hiked on 1,100 packages; end-of-year profit relief expected

By John Gallant

ARMONK, N.Y. — Three days after confirming that its third-quarter profits were lower than those of a year earlier, IBM last week announced an average 10% price increase on more than 1,100 application and utility software packages for medium- and large-scale systems.

The price hike was the company's second major software pricing increase this year. In February, IBM increased by 7% the price of nearly all of its program products. With last Tuesday's increase, initial license charges and one-time charges for affected software rose by about 10%. Prices for program products supplied with a monthly license charge will increase by the same amount Feb. 1, 1986.

The price increase did not affect operating systems and many important systems software products such as DL/1, CICS or IMS. Analysts said that those products were spared because the cost of using them has already increased significantly over the past few years.

In its financial statement, IBM said that third-quarter profits were down 7% from those of the same period last year. For the quarter ended Sept. 30, IBM posted a \$1.47 billion profit, compared with \$1.58 billion in 1984. Revenue for the third quarter jumped by nearly 10%, from \$10.66 billion in 1984 to \$11.67 billion. Despite that rev-

Continued on page 7

Long-awaited network architecture connects micros, not mainframes

By John Dix

NEW YORK — The long-awaited debut of the IBM Token-Ring network played to a packed analysts' briefing last Tuesday morning and a standing-room-only press introduction that afternoon. But users and analysts were generally unimpressed with the network, which proved to be functionally similar to the previously available IBM PC Network.

The product followed along the lines of statements of direction issued by IBM more than a year ago, but as presented, it provides only for direct connection of Personal Computers.

IBM did not reveal any plans or timetables for connecting departmental computers to the network, but the company did provide gateways to the Series 1 processor and 370 mainframe hosts.

"This is a token-passing ring implementation of the PC Network," according to Dale Kutnick, an independent consultant in Wayland, Mass. "You can't do anything on this network that you could not have done on the PC Network."

Some users were expecting nothing more of the 4M bit/sec. network. "I'm not surprised that the Token-Ring is strictly for Personal Computers because IBM is losing a lot of network business to other technologies and firms," said William J. Johnson, director of telecommunications

Continued on page 10

Service firms in parts blockade

Leading vendors take hard line on maintenance

By Clinton Wilder

For players in the fast-growing game of third-party computer maintenance, there is a major occupational hazard. In practices ranging from withholding discounts on spare parts to outrightly refusing to do business with third parties, many leading hardware vendors follow a hard line when it comes to competition against independent service firms.

Such policies can constitute a significant roadblock as the demand for third-party service grows, but most third-party vendors concede there is little they can do about the policies — except find creative ways around them.

Although some third-party vendors believe the practices unfairly infringe on us-

that deal with this issue," according to Richard Donahue, an attorney with the U.S. Federal Trade Commission. "The unilateral refusal to deal in any commodity has always been allowed."

The only illegal situation would be collusion among vendors in jointly agreeing not to do business with third parties, which would violate antitrust laws, Donahue said. No third-party companies have alleged such collusion, and such a hypothetical charge would be extremely difficult to prove.

The notable exception to the rule is IBM, which is required by a 1956 consent decree negotiated with the U.S. Department of Justice to provide parts and documentation to any company that wants

TOP OF THE NEWS

Paul M. Villiere, executive vice-president at AT&T's Network Systems Group, said during Info '85 last week that 1984 industry modem sales matched the total of all previous years' sales. He predicted that 1987 will be the year in which corporations make volume purchases of Integrated Services Digital Network products. AT&T plans to incorporate ISDN capabilities into the company's 5ESS switch next year. For more on Info '85, see page 2.

As expected, Applied Data Research Chairman and CEO John Bennett had to field some hard questions about the current snafu at the New Jersey Department of Motor Vehicles with ADR's Ideal fourth-generation language product [CW, Sept. 30] at last week's Alex Brown & Sons computer services investment symposium in Baltimore. At least Benne **RPX Exhibit 1008** hu-

DEC86
1454
BUTTS PROG COORD
UNIVERSITY OF WI BUDGET/ANAL
BASCOM HALL
LINCOLN DRIVE
ISON WI 53706

IN DEPTH

This material may be protected by Copyright law (Title 17 U.S. Code)

Does your DBMS run by the rules?

To be "mid-80s" fully relational, a DBMS must support all 12 basic rules plus nine structural, 18 manipulative and all three integrity rules. There will be more requirements by the 1990s.

By E. F. Codd

Last week, the originator of the relational model described the 12 rules by which to measure any DBMS claiming to be relational. This week, Dr. E. F. Codd presents the practical consequences of his 12 rules as well as 30 additional features of a relational system. Then he asks vendors to measure up.

Part 2

No existing DBMS product that I know of can honestly claim to be fully relational at this time. The proposed ANSI standard does not fully comply with the relational model, so a DBMS' fidelity to the ANSI standard is no guarantee of relational capability. The standard could be modified, but already vendors are well advised to extend their products beyond the standard to support customers' DBMS needs fully.

In their ads and manuals, vendors have translated the term "minimally relational" to "fully relational," so more stringent criteria must be applied. Twelve rules (below) comprise a test to determine whether a product that is claimed to be fully relational is actually so. A grading scheme used to measure the degree of fidelity to the relational model follows.

A DBMS advertised as relational should comply with the following 12 rules:

1. The information rule.
2. The guaranteed access rule.
3. Systematic treatment of null values.
4. Active on-line catalog based on the relational model.
5. The comprehensive data sublanguage rule.
6. The view updating rule.
7. High-level insert, update and delete.
8. Physical data independence.
9. Logical data independence.
10. Integrity independence.
11. Distribution independence.
12. The nonsubversion rule.

E. F. Codd originated the relational model for data base management. He led the team that designed and implemented the first operating system with multiprogramming capability. This year he established two companies with Chris Date: The Relational Institute and the Codd & Date Consulting Group, both based in San Jose, Calif.

IN DEPTH/RELATIONAL DBMS

These rules are based on a single foundation rule. I call it Rule Zero:

For any system that is advertised as, or claimed to be, a relational data base management system, that system must be able to manage data bases entirely through its relational capabilities.

This rule must hold whether or not the system supports any nonrelational capabilities of managing data. Any DBMS that does not satisfy this Rule Zero is

not worth rating as a relational DBMS.

But compliance with Rule Zero is not enough. Failure to support the information rule, guaranteed access rule, systematic nulls rule and catalog rule can make integrity impossible to maintain. These four rules support significantly higher standards for data base administration and control (authorization and integrity control) than earlier DBMS supported. Users should remember that a

data base managed by a relational DBMS is likely to have both experienced and inexperienced users; it must be able to serve both.

Rule Zero not enough

Rules 1 and 4, the information and catalog rules, allow people with appropriate authorization (such as executives of the company) to find out easily via terminal what information is stored in a data base. I have encountered data base administra-

tors using nonrelational systems who were unable to determine if a specific kind of information was recorded in their data base.

Rule 3, which calls for the inclusion of systematic support for unknown and inapplicable information by means of null values that are independent of data type, should help users to avoid foolish and possibly costly mistakes. The treatment of nulls, when aggregate functions such as total and aver-

age are applied, holds considerable interest for users. The Oracle DBMS in particular has an outstanding approach to null values. The user may specify whether the aggregate function is to ignore null values or yield a null result if any null value is encountered.

In general, controversy still surrounds the problem of missing and inapplicable information in data bases. It seems to me that those who complain loudly about the complexities of manipulating nulls are overlooking the fact that handling missing and inapplicable information is inherently complicated. Going back to programmer-specified default values does not solve the problem.

Rule 5, the comprehensive


”

The ANSI standard as now proposed is quite weak. It fails to support numerous features users need to reap the advantages of the relational approach.

You may ask why Elgar's new UPS looks just like an IBM System/36? ...Why indeed!

- Our new UPS is indeed a perfect match for IBM. It provides the kind of power protection necessary for IBM systems. And only Elgar is plug-in compatible with both the CPU and peripherals for fast, easy installation.
- Our new UPS is packed with quality features for trouble-free performance. Only Elgar has $\pm 4\%$ dynamic power regulation. We deliver more precise power to your computer than any other UPS you can buy. Our new UPS costs less to install . . . to run . . . and maintain.
- Want more? It fits in your office. It's so quiet you'll hardly know it's running. And it's backed by a no-nonsense TWO YEAR WARRANTY from Elgar . . . the leader in quality power protection systems.
- Get the full story on our new T-Series UPS for mini computers.

Call toll free: 800-854-2213, Dept. 8.



ELGAR

An Onan Operating Unit

9250 Brown Deer Road
San Diego, California 92121
(619) 450-0085 Telex 211063



data sublanguage rule, is important for several reasons. First, it allows programmers to debug their data base statements interactively, treating them separately from whatever nondata base statements occur in their programs — a significant contributor to productivity. Second, it means that a single tool can be used for defining relations derived from the data base, whatever the purpose. The view updating rule, Rule 6, is vital for the system to support logical data independence.

Rule 7, which requires a multiple-record-at-a-time attack on insertion, update and deletion, can help save a good portion of the total cost of intersite communication in a distributed data base. If the system includes a good optimizer (an important component in relational DBMS performance), this rule can also result in substantial saving of CPU and I/O time, whether the data base is distributed or not.

Failure to support independence (Rules 8 through 11) can, and very likely will, result in skyrocketing costs in both money and time. Developing and maintaining applications programs and terminal activities will be more expensive. Managers may even be unwilling to consider changing certain business

IN DEPTH/RELATIONAL DBMS

relational data bases. All too frequently, I have seen situations in which data base administrators with nonrelational DBMS failed to control their data bases adequately; consequently, they could not maintain a state of integrity.

Domains

Many users confuse the domain concept with the concept of attribute of a relation or column of a table. Other people (often the vendors themselves) dismiss the domain concept as "academic." My reply to them is: The atom bomb was also academic!

In fact, the domain concept is very important, practical and simple. A domain consists of the whole set of legal values that can occur in a column. The column draws its values from the domain. Each column of a relational data base has precisely one domain, but any number of columns may share a domain. There are several reasons why domains should be supported.

For example, in a financial data base, there may be as many as 50 distinct columns (possibly, but not necessarily, in distinct tables) defined around the U.S. currency domain. Why repeat the definition of currency 50 times? In data bases supported by nonrelational systems, I frequently observe many inconsistent declarations of value type for fields that were intended to have the same type.

It is unreasonable to expect a DBMS to store all the legal values in a domain, unless there happen to be very few. However, it is entirely reasonable — and very worthwhile — to insist that a DBMS should store certain values:

- For each domain, a description of the type of values in that domain. This information is global since it applies to the entire data base, and it should of course be recorded in the catalog.

- For each column, the name of the domain from which that column draws its values. This domain name is a reference to the global definition.

Of course, the domain description can include range restrictions. For example, it could specify that quantities of parts in an inventory must not only be integers, they must also be non-negative.

Furthermore, individual columns may include additional range restrictions where these are semantically justifiable. In this example, the quantities of very expensive parts held in the inventory may be limited to some specified maximum

values is largely or even completely factored out. For example, when there are 50 distinct columns defined on U.S. currency, the data base is much easier to manage and manipulate if one avoids making 50 distinct declarations for U.S. currency.

Before the relational discipline arrived, users had to make separate declarations, and as a result, many of the 50 in the example would turn out to be incompatible with one another by acci-

dent. The factoring of declaration that prevents these errors is achieved in Digital Equipment Corp.'s RDB, which has a concept of "global field definition." But RDB fails to support domain constraints on certain operations, such as join.

Another benefit of supporting the domain concept is that relational operators, such as joins and divides, that involve comparison of values between different columns can be constrained by

the system. A DBMS can allow data base values to be compared only when they come from the same domain and are therefore comparable from the semantic viewpoint.

Such a constraint inhibits errors caused by interactive users of terminals who choose columns to be compared in such operations as joins. The wrong answers they obtain from these errors rarely uncover the errors themselves; meanwhile,

unwise business decisions may be made based on these wrong answers.

For various reasons, it is important to support as a qualifier in a command what I call "semantic override" — the ability to have the system ignore the usual comparison constraints. Users should be able to authorize this override qualifier separately from the operator involved and should authorize it rarely, reserving it chiefly for detective work.

TO INSURE THAT EVERYONE CAN ACCESS DATA, TRANSAMERICA USES INTELLECT™ AT EVERY LEVEL OF THE PYRAMID.

"We use INTELLECT because we want to give our users a better way to do business." —Mr. Carl Rahmqvist, Senior Systems Manager, Information Systems, Transamerica Insurance Group

Transamerica Insurance Group is recognized for the pyramid-shaped building of its parent, Transamerica Corporation, and for being one of the country's leading commercial and personal insurance companies. In some quarters, Transamerica is also recognized for its variety of innovative INTELLECT applications.

At Transamerica they see INTELLECT as more than a state-of-the-art natural language information retrieval system. They see it as a new way to do business: Giving all their end users—even those with no computer skills—instant access to more information than they'd ever had before.

"When they saw how easy, fast, and resource-effective it was to use, INTELLECT became a very popular tool!" —Ms. Sandra Dahlgren, Information Center System Supervisor

One Vice President uses INTELLECT to get the most current information on premiums and losses, and for longer term strategic planning with requests such as, "Give me the total June premiums and losses for each region." Regional offices use INTELLECT for a variety of tasks including asking INTELLECT to: "Tell me all about policy number 98579897." Personnel keeps track of employee records, EEO compliance, and more, by questioning INTELLECT in plain English. Payroll, Claims, Underwriting, and Services also use a variety of INTELLECT applications. So from the top of the pyramid right down to the mailroom, Transamerica is using INTELLECT to work faster and smarter.

"After just minimal training, people began using INTELLECT frequently and effectively."

—Mr. Carl Rahmqvist
After some initial implementation assistance from AIC,

Transamerica's Information Center took over, developing custom INTELLECT applications for each department. Mr. Rahmqvist and Ms. Dahlgren trained a group of Transamerica's "veteran" INTELLECT users to go into the field to train the company's 15 regional office personnel in using the company's many applications. Transamerica's INTELLECT Support Staff helps company personnel use INTELLECT more effectively. They've also devised several methods for enabling users to access their many INTELLECT applications. Easily understood menu screens help beginning users work faster, while experienced people can use an express mode for their application needs.

"We've found that we've saved both time and money by using INTELLECT to obtain pertinent information!" —Ms. Sandra Dahlgren

Anyone who can ask a question in everyday conversational English can get the information he or she needs. Immediately. You can imagine how much time a system like that can save an information-dependent organization like Transamerica. Combine that with the ability to get more people into the system and you can see how INTELLECT has increased the pyramid's power.

Find out how INTELLECT can get your organization into better shape. Write for a free demo diskette. Or for fast action, call AIC at (617) 890-8400.



I want to know more about INTELLECT!

Name _____

Title _____

Company _____

Address _____

State _____ Zip _____

Telephone (____) _____

O/S: VM/CMS _____ MVS _____ DBMS _____

_____ Please send me information about INTELLECT.

_____ Please have an AIC representative contact me.

CW1021

INTELLECT™ ARTIFICIAL

IN DEPTH/RELATIONAL DBMS

Even when the domain concept is restricted to assigning types to data, it should not be confused with the hardware-supported data type. Consider the example of a data base listing suppliers, parts and projects. Suppose the hardware-supported data types of supplier serial numbers and part serial numbers are identical: each type consists of fixed-length strings of 12 characters. The system still needs to keep these two data types distinct and remember which columns are defined on one and which columns are defined on the other.

If it can make these distinctions, then when a request comes in to delete or archive all records containing X3 as a supplier serial number, the system can handle such a transaction correctly. The system will not delete or archive any record that contains X3 as a part serial number and that also does not contain X3 as a supplier serial number.

Today, such a data type is often called an application data type. The concept is supported in Pascal but in very few other languages that enjoy current use. The Pascal support does not, of course, include constraints on sets, unions, joins and divides.

The domain concept is basically what makes all the meaningful selects, unions, joins and divides known to the DBMS. Thus, the domain makes the data base meaningfully integrated, and it does so without prejudicing distributability.

Contrast this with CODASYL links and IMS hierarchic links. They represent the CODASYL and IMS con-

Fidelity to the 12 rules (by DBMS)			
Rule	DB2	IDMS/R	Datacom/DB
1 Information rule	Yes	No	No
2 Guaranteed access rule	Partial	No	No
3 Systematic treatment of nulls	Partial	No	No
4 Active catalog based on resource management	Yes	No	No
5 Comprehensive data sublanguage	Yes	No	No
6 View-updating rule	No	No	No
7 High-level insert, update, delete	Yes	No	No
8 Physical data independence	Yes	Partial	Partial
9 Logical data independence	Partial	No	No
10 Integrity Independence	No	No	No
11 Distribution Independence	Yes	No	No
12 Nonsubversion rule	Yes	No	No
Score (1 for yes, 0 otherwise)	7	0	0

Figure 1

cept that a link "integrates an otherwise unintegrated data base," but they have several unfortunate restrictions. Most importantly, they obstruct data base distribution because of the constraints and complexity their data structures introduce into decisions regarding how the data should be deployed.

A second serious drawback of links is that they are only paths. Generation of a result such as a join requires traversal of these paths by the application program. It seems superfluous to cite other difficulties with this concept.

Many relational DBMS and languages including SQL do not support

the concepts of primary key and foreign key. I fail to see how these products can support the guaranteed access or the view updating rules without making the system aware of which column(s) constitute the primary key of each base table.

Furthermore, I fail to see how these products can support referential integrity or the view updating rule without offering clear support for both primary keys and foreign keys. For example, in SQL, the CREATE TABLE command should be extended to permit the user to declare which column or columns constitute the primary key and which constitute foreign keys. In addition, there

should be a new CREATE DOMAIN command in SQL.

Fidelity

Figure 1 shows fidelity to the 12 rules by IBM's DB2, Cullinet Software, Inc.'s IDMS/R and Applied Data Research, Inc.'s Datacom/DB — examples chosen for their wide differences. These scores represent counts of compliance with each rule (score one for "yes" and zero for either "partial" or "no").

Actually, the information rule is so fundamental to the relational approach that a system's compliance with this rule should receive a much higher score than one. Weighting it as high as 10 would not be excessive. However, I shall avoid assigning different points for different features, just as I avoided a fractional score for partial support of a feature: It is too easy to be subjective in these matters.

DB2 scores quite well on the fidelity evaluation. Very few other DBMS score higher on the 12 rules, although some others score equally well. Both IDMS/R and Datacom/DB allow information to be represented in the order of records in storage and in repeating groups — directly violating the information rule. In the case of IDMS/R, information may also be represented in links between record types (CODASYL calls them "owner-member sets") and also in "areas."

Some vendors of nonrelational DBMS have quickly added a few relational features — in some, cases, very few features — in order to be

TSO users DSM™ (Data Set Manager) Improves TSO Performance

```

-----Data Set Manager-----
PFK > 1-TOP, 2-BOTTOM, 3-PF, 4-PB, 5-FIND, 6-QUIT, 7-PRINT, 8-PFK, 9-HELP
CATALOG PREFIX= JONES
<PF> Global Command Input Area

<--DSNAME--> VOLSER ALLOC USED EX ORG RECFM BLKSIZE LRECL LASTREF
SPF  ANZ.CNTL  TSO010  3  2  1  PS  FB  3120  80  06/23/84
B   APR01.TEXT TSO012  8  8  1  PS  FB  3120  80  07/08/84
FSE  B13.COBOL  TSO008  10  8  1  PS  FB  3120  80  08/13/84
DEL  CP21.TEXT  TSO007  6  6  1  PS  VB  3120  255  08/23/84
REN  DEN.FDRIT TSO005  4  3  1  PS  FB  1000  80  09/20/84
COPY EZ.ASM    TSO002  3  3  1  PS  FB  3120  80  10/02/84
SUB  FM001.CNTL TSO006  1  1  1  PS  FB  3120  80  04/12/84
PRINT G20.LIST  TSO003  20  18  1  PS  VB  3120  255  06/20/84
CDMP  HCL.TEXT  TSO008  30  30  3  PD  VB  3120  255  10/15/84
RLSE  IM001.LOAD TSO010  50  18  1  PD  U   0  10088  08/03/84
LM    JCL1.CNTL  TSO002  20  16  1  PD  FB  3120  80  09/15/84
FREE  KA.COBOL  TSO004  6  5  1  PS  FB  3120  80  08/20/84
---- TSO.PL1    TSO003  10  8  1  PS  VB  3120  255  10/23/84
  
```

- DSM provides full-screen displays of cataloged files and supports TSO functions with simple, east-to-remember command names.
- DSM supports SPF, FSE, CLISTS and Command Processors. DSM's flexible design allows users to define their own commands.
- DSM file commands include: DELETE, RENAME, COPY, SUBMIT, PRINT, COMPRESS, Release unused space, List PDS members and deallocate.
- DSM supports PDS member displays, HSM and MSS.
- Security? - Our user exit lets you control all functions.

Call or Write for details:

Need to recruit people for your IBM systems? You'll find them reading the Computerworld Extra! on IBM.

Published December 4th and closing October 25th, this special edition of Computerworld Extra! will take a hard look at IBM's products and strategies. Anyone working in the IBM arena will certainly review this issue. So if you're looking for pros in IBM systems, get your ad in this special issue and be surrounded by in-depth editorial on IBM.

Computerworld Extra! will discuss IBM's strengths and weaknesses. We'll look at how SNA evolved, and how it will continue to evolve. And we'll discover whether IBM plans to provide a universal interconnect to SNA. Finally, we'll cover the alternatives, from PCs to mainframes, LANs, Communications. And, of course, the compatibility issue.

As you can see, this issue will have complete appeal for computer professionals working at IBM installed sites.

To reserve space call Al DeMille, National Recruitment Sales Manager, at (800) 343-6474, or (617) 879-0700 in Massachusetts.

There is no special classified section; all recruitment ads are considered display advertising for this issue.

COMPUTERWORLD

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.