

Charles W. Bachman: Database Software Pioneer

Thomas Haigh

University of Wisconsin, Milwaukee

Editor: Thomas Haigh



Charles W. Bachman was born in Manhattan, Kansas, in 1924 but did not stay there long. His father, a Hall of Fame football coach, moved their family several times before settling in East Lansing, Michigan, for a long spell working at what was then Michigan State College. The younger Bachman retained a lifelong interest in sports, playing on several

teams during high school. He earned strong grades in most subjects but decided on an engineering career at an early age. At 11 years old, he had triumphed in a local Soap Box Derby, in which children custom build and race small, unpowered cars.¹ When I interviewed him in 2004, he told me, "I was never interested in anything else. I was always going to be an engineer."

Early Education and Career

Before Bachman could realize this dream, there was the little matter of World War II, which was raging throughout his high school years. Resolving to enter the military with a little additional preparation, he graduated a semester early from high school to spend the spring and summer of 1943 satisfying freshman college requirements at Michigan State University. He spent two years in the Army Antiaircraft Artillery Corp, serving in New Guinea, Australia, and the Philippines. Antiaircraft guns of this era used a simple mechanical computer to target a plane's predicted future

position based on a straight-line extrapolation of its previous trajectory. This convinced him both of the importance of predictive control loops and the difficulty of accurately hitting future targets. After the war ended, he returned to Michigan State and completed his mechanical engineering degree in 1948.

It was time to leave East Lansing again. An interest in thermodynamics led Bachman to graduate school at the University of Pennsylvania, although "it wasn't until I arrived there that I had realized that the engineering graduate school at Penn was a night school" and that his classmates all had day jobs. He filled the daylight hours with management courses from Penn's Wharton School, satisfying most of the coursework requirements for an MBA even though there was no formal joint degree in engineering and business. In the summer of 1949, he married his college sweetheart, Connie Hadley, two weeks after her graduation from Michigan State.

The next year he graduated amid the uncertainty that accompanied the outbreak of the Korean War. A search for work in the Great Lakes region culminated in a job with Dow Chemical in Midland, Michigan.

Bachman's career at Dow began with a nine-month spell rotating through the plant's various departments. After that, he settled down to his role as "an engineer, at a drafting table, in a room with more than one hundred engineers." The chemical industry is oriented around networks of processes, as the products of one reaction become the inputs for another. He worked in a

Background of Charles W. Bachman

Born: 1924, Manhattan, Kansas

Education: BSc (mechanical engineering), Michigan State University, 1948; MS (mechanical engineering), University of Pennsylvania, 1950.

Professional Experience: Dow Chemical, assignments in engineering, finance, manufacturing, and systems work, 1950–1955, manager of corporate data processing, 1956–1959; and manager of data processing research, 1960; General Electric, Integrated Systems Project in Production Control Services, 1960–1964; Computer Department (titles include manager of software product planning, manager of applied

technology subsection, and manager of data management software for the advanced system division), 1964–1970; Honeywell Information Systems, chief staff engineer, 1970–1980; Cullinane Database Systems, 1980–1983; Bachman Information Systems, president and CEO, 1983–1988, and chairman, 1988–1996; Cayenne Software, president, 1996–1997; Constellar, consultant, 2002–2006; InfiniteIQ, consultant, 2009–2010.

Honors and Awards: ACM Turing Award, 1973; Distinguished Fellow of the British Computer Society, 1977.

group responsible for the water, steam, and electrical power systems that kept these processes running at its Midland plants. Bachman's work centered more on economic trade-offs than purely technical designs, analyzing design choices within the context of larger set of assumptions to optimize expected returns.

He recalled having independently reinvented concepts such as discounted cash flow as part of an effort dubbed the "Equivalent Capital Project." By 1952, this led to work analyzing the capital cost of building all Dow manufacturing plants, for which Bachman commandeered time with the mechanical punched-card machines in the Accounting Department to iteratively solve networks of simultaneous equations representing data collected from several Dow plants. He wired plug boards and operated the machines himself. In 1954, he worked in the Pricing Department and was responsible for the pricing of agricultural chemicals.

In 1955, he shifted from analysis to production for two years, working first as a process engineer at the Saran plant and then as the number two manager of the Styron 475 plant. These experiences exposed him to many different aspects of business, reinforcing his tendency to see activities as part of a unified system.

Share and 9PAC

In 1957 Dow Chemical decided to join the computer age. On the recommendation of its treasurer, with whom Bachman had worked closely on his capital accounting project, Bachman was chosen to head its new corporate Data Processing Department. His first job was to recommend a machine to purchase. As he recalled,

I don't think we ever did a feasibility study. The idea was that "Everyone else is doing it. We should be doing it." Which one should we buy was the question. . . . It wasn't clear what was going to happen with these computers, what they were good for.²

He was courted by IBM, with a tour of its Poughkeepsie factory, and by Univac which provided a two-week programming course. Impressed by its ability to handle both business data processing and technical computing, he recommended renting an IBM 709.

Computer customers generally waited a year or two for delivery, during which time they hired and trained data processing staff

for the new operation. By May 1958, Bachman had hired 30 people and the department was beginning to plan its applications. But Dow ran into economic trouble and began to lay off its newest hires, slashing 10 positions from his group. Others on his team quit in search of more promising positions, and Dow eventually cancelled the IBM order, although not before constructing an empty computer room in its new headquarters building.

Ordering a 709 had qualified Dow Chemical for membership in Share, the group for users of large IBM scientific computers. At a time when IBM itself provided only a limited selection of system programs for its machines and there was no independent software industry, Share played a vital role in establishing a common library of programs and subroutines between sites. Its committees sponsored the joint development of new projects, including the ambitious SOS project to produce a standard and powerful operating system for the 709.

Bachman quickly became a leader among the smaller community of computing centers intending to use the 709 primarily for business applications. At his second Share meeting, he met Harry Tellier of General Electric, manager of its computer team working in the Hanford Atomic Reservation. Tellier's group had created an ingenious system for the earlier IBM 702 computer that generalized the laborious programming work involved in maintaining master data files on tape and generating reports based on the data contained in them. This inspired many subsequent systems, including IBM's own widely used Report Program Generator (RPG) family.³

Bachman quickly recognized the value of such a system to Dow, and Tellier and he became founding members of the Share Data Processing Committee. The committee sponsored a collaborative project to create what became known as 9PAC, an updated version of the Hanford report generator package for the 709. (I have previously discussed the development of 9PAC and its importance to the development of data management software, so I will not do so in detail here.⁴ It entered use in 1959 and was widely deployed. Among its most important innovations, which Bachman credits to discussion between himself and Russ McGee of GE's Hanford group, was support for hierarchical parent-child relationships within a magnetic-tape file—for example, following each

customer record with records for the orders placed by that customer.

It is a tribute both to the practical flexibility of Bachman's intelligence and Share's effectiveness in allowing participants to learn from the experiences of others that he could play a significant role in the 9PAC design without having written a data processing application or ever actually taking delivery of a computer. Share also established a Committee on Theory Information Handling, with Bachman as an active member. The committee tried to abstract a concept of data away from the constraints of its physical representation in tape or card files. It formulated the idea of data as a set of "triples"—so called because each was defined by the unique combination of a single data value, an identifying name, and a designated owner.⁵ He also worked on the design of a still more powerful successor, Ingen, that would have provided more flexibility to restructuring data hierarchies as needed for particular operations.⁶ Bachman's participation in Share was cut short by Dow's decision to cancel its order with IBM, but these efforts clearly influenced his subsequent work on database management.

GE and the Integrated Data Store

Bachman left Dow Chemical in 1960, moving to New York City for a corporate-level staff position with General Electric. GE was admired throughout corporate America both for its leading position in high-technology industries and its equally innovative managerial culture. It was the leading exemplar of what business historian Alfred Chandler was about to dub the "M-form" corporation, with a small headquarters staff governing a sprawling range of different business divisions manufacturing products from light bulbs to nuclear-powered steam turbines.⁷ Responsibility for operational decisions was decentralized to these divisions, or as they were known within the company, "departments." In 1956, the firm had opened an internal business school in Crotonville, New York, to educate a new cohort of managers with the skills needed to run a decentralized firm. Successful managers were moved frequently between different parts of the company, and they were held accountable for their performance on a set of standard financially based metrics. This reflected a conviction that management was a generalized skill that could be separated from expertise in any particular industry.⁸

General Electric was therefore investing heavily in the latest management methods and technologies. One of the responsibilities of its corporate-level teams was to research areas such as operations research, simulation, forecasting, and automation so that effective new practices could be spread across its diverse empire. Bachman was recruited to its prosaically named Production Control Services Group, headed by H. Ford Dickie, who was an early and enthusiastic advocate for the "total systems approach" of using computers and cybernetics to revolutionize the control of industrial production.⁹ Bachman's experience with data processing, engineering, management, and cost accounting made him a perfect fit.

In a 1962 presentation prepared for GE's Information Systems Task Force, Bachman wrote that "the introduction of computers into our business as a new resource" had "created an atmosphere in which the entire outlook of information systems is changing." He went on:

Many of the old equilibriums between marketing, manufacturing, finance and engineering are being upset by brash young men who think they know all about everybody else's business and who propose to change it. The orderly absorption of these men with their tools is today's order of business.

By analogy with his experience with the separation of design and production within engineering, Bachman argued that a new "information system designer" role should be established to oversee the system as a whole. He identified the "information inventory" as a crucial subsystem in this new vision of computerized management, writing that "the maintenance of information inventories is one of the largest work assignments given to computers today. It is also a thing that currently they are doing very poorly."¹⁰

The group was using a factory in Philadelphia as a testbed for its Integrated Systems Project #2, which eventually produced an integrated system known as MIACS to automate its planning, scheduling, operational control, and inventory control. Running on a GE 225 computer, it integrated capabilities such as production scheduling, economic order calculations, order dispatching, receiving, inventory management, and parts explosion. Bachman claims it was the first integrated manufacturing control system, and the direct ancestor of the programs

later marketed as material requirement planning (MRP) and bill of materials processing systems. MIACS was used internally by a number of General Electric departments, and was offered as a software package to purchasers of the firm's more powerful GE 600 models that succeeded the 225 later in the 1960s.

During the 1960s, many firms attempted to create what were sometimes called "totally integrated management information systems," but MIACS was one of the few ambitious integrated systems projects to succeed.¹¹ A small team produced a useful and powerful integrated system in about four years. In the process, they invented some of the most fundamental technologies behind today's large-scale administrative computer systems. This reflected the exceptional caliber of the people involved, a testament to the scale of GE's investment in systems work and operations research during the 1950s and early 1960s. (One of Bachman's predecessors in its production control services group, Harry Markowitz, later won the Nobel Prize for economics).¹²

Successful as it was, MIACS is remembered primarily because of the foundation Bachman created for it with another piece of software, the Integrated Data Store (IDS).¹³ This reflected Bachman's concept of an "information inventory." IDS was the first package to adopt what was later called the "network data model," meaning that the system would support and enforce relationships between records in a much more flexible manner than earlier packages such as 9PAC. This exploited the new power of disk storage and was an essential capability for the storage of data in a form that could support applications for different functional areas. IDS provided application programmers with a set of powerful commands to manipulate data, an early expression of what would soon be called a Data Manipulation Language. It maintained a separate data dictionary, tracking information on the different kinds of records in the system and their relationships.

Only the boldness of Bachman's IDS design, and the remarkable efficiency with which he squeezed IDS and the MIACS applications into a computer with the equivalent of 40 Kbytes of memory, made this possible. This tight coupling meant that IDS and the Problem Controller, a transaction-oriented operating system produced by the team, almost entirely replaced General Electric's rudimentary operating system and other systems software tools. Bachman was uniquely well

qualified to tackle the job; he had the skills and mindset of an engineer, combined with years of systems work, deep immersion in Share's efforts to produce reusable generalized data management software, access to prototype hardware and software, and a supportive corporate culture.

Today's database designers rely on graphical tools to depict the relationships between different tables as a web of connections. These data structure diagrams are the direct descendents of a technique Bachman devised to illustrate the complex data structures required to picture the adaptive manufacturing control logic of MIACS.

Turning IDS from a prototype to something robust and flexible enough for GE to offer to its internal customers required some changes. Bachman began taking the subway from his office in GE's headquarters on Lexington Avenue in midtown Manhattan to International General Electric, located about 20 blocks south, where this work was being carried out. He delivered two boxes of punched cards holding the system code and returned to assist the team as needed. Changes included the addition of a backup and recovery system, disentanglement from GE's proprietary GECOM business application programming language, and improved handling of memory paging.¹⁴ GE publicly announced IDS as a supported package for its computer customers in October 1964.¹⁵

The GE Computer Department

In 1964, Bachman transferred to GE's Computer Department in Arizona. He was originally assigned to its marketing team; a letter sent to him summarizing his intended responsibilities promised that he would assist with the launch of IDS as a (bundled) product by pitching its benefits to potential customers, "communicate system design techniques" to key technical managers in user firms, and assist in the preparation of technical proposals involving IDS.¹⁶ But GE's computer department was in constant flux, and over the next six years, Bachman worked on a number of projects within its Project Planning Group, and eventually he found himself in the new Advanced Software Group.

A few years after MIACS was completed, Bachman helped GE build a still more ambitious management information system project on the same foundation. Known as WEYCOS 2 (for Weyerhaeuser Comprehensive Operating Supervisor), it was produced for Weyerhaeuser's Wood Products division

in Tacoma, Washington, by a team composed of employees from both companies. This was a sequel to WEYCOS, fully operational by 1965, which integrated the Problem Controller transaction-processing system developed for IDS with a smaller computer connected to teletype units across the country. Using WEYCOS, employees in remote offices could send messages directly to the computer to update records or request information.

Weyerhaeuser credited the WEYCOS system with dropping order-processing time from weeks to days, greatly improving its efficiency and competitiveness. However, the GE 235 system was slow, relying on the ability to queue requests and clear the backlog of lower priority transactions over nights and weekends. WEYCOS 2 was intended to serve the whole company and handle a growing volume of requests. One initial goal, ultimately abandoned, was for it to run on a cluster of four powerful GE 600 series computers with a shared IDS database and a common memory bank to coordinate transactions. The team did successfully adapt IDS so that a single database could be shared by up to eight application programs running simultaneously—another key step in the evolution of the database management system (DBMS).

Other major additions to the standard GE operating system of the era added support for automatic job scheduling, execution of programs from remote locations, and message transmission between locations. Development of WEYCOS 2 ran from 1966 to 1968. In 1972, WEYCOS was still in use, costing around \$140,000 a year to support and maintain. Weyerhaeuser stated that its capabilities at that time were still not matched by standard operating systems, though the gap was closing.¹⁷

As a product, IDS was considerably more flexible than IBM's rival Information Management System (IMS) package and was released several years earlier. However, software packages in those days, particularly those from computer manufacturers, were tied to particular platforms. GE had less than one tenth of the market share held by IBM, and its computer department suffered from a relative lack of resources and managerial stability. Although IDS was not the most widely used package of the 1960s, it was one of the most influential.

This influence came in large part through Bachman's participation in the Data Base Task Group organized by Codasyl, the data processing standards group best known for

its promulgation of Cobol as the industry standard applications programming language. As I have discussed elsewhere,⁴ this group's reports popularized the concept of a "data base management system" and established its key features such as separate data-definition and data-manipulation languages. Bachman was an active member in the group's early history, from 1966 to 1968, and his influence on its final reports and draft language specification was profound. In particular, the group endorsed Bachman's network data model approach over the less flexible hierarchical data model backed by IBM. Bachman himself believes that the data-manipulation language proposed by the task group had few substantive differences from the one he had designed for IDS. He told me that

I wrote, rewrote, and rewrote the specification for the DBTG specifications for a database management system. The original text was the text for the GE IDS database management system. The "master" and "detail" records of IDS became the "owner" and "member" records of the DBTG specifications. The word, "set," was in the [phrase owner/member set]. These were good improvements. Some of the changes were really to satisfy some member's desire to have something that they could point to as their contribution. My challenge was to preserve the IDS functionality; not to preserve the original IDS data manipulation language in all of its details.¹⁸

Bachman's other main project of this era was DataBASIC, implemented by a small team working under his direction in 1967–1968. Basic had originally been designed at Dartmouth College as an interpreted programming language for interactive use on GE computers by people without significant computer expertise. DataBASIC extended the language to include the data-manipulation capabilities needed to use it effectively to create and query databases. This reflected general faith during the late 1960s that time-sharing technology would offer low-cost computing services to professionals via telephone lines and cheap terminals. Bachman wrote that it was "designed to do for the IDS/COBOL world that which BASIC had done for the FORTRAN world" and "permitted business professionals who are not programmers to construct, maintain and operate their own private data bases."¹⁹ DataBASIC was based on the data triple idea Bachman had earlier explored within Share. It was




Figure 1. Charles Bachman teaching a group of Honeywell staff, probably in Vienna, Austria in 1978. (Image courtesy of Charles Babbage Institute.)

a completely ad hoc system, with no record types and no predefined field names. A decade later, systems such as dBase II were being widely used on personal computers for such purposes. DataBASIC was ahead of its time, but even though GE supplied the software for its 400 and 600 product lines, it was not widely used.

Within the GE Computer Division, Bachman's role as "Manager – Application Technology Subsection" gave him an opportunity to develop his position as one of the industry's leading experts on data management and information systems. However, he was less comfortable with the managerial side of the job. His 1969 performance appraisal, while unstinting in its praise of his technical leadership, noted that "Charlie needs to improve his communication channels with other competent people, who do not understand or agree with his ideas without compromising what he believes is right. Charlie's primary interests are technical rather than managerial."²⁰

Honeywell and the 1970s

In 1970 Honeywell acquired GE's computer business. Bachman worked for the merged operation in Honeywell's advanced research

group in Boston (see Figure 1). The era's surviving mainframe companies were then engaged in massive integrated projects to deliver next-generation computer architectures and advanced software features. Bachman was concerned with the design of data management software for these systems, including communication facilities and the creation of databases distributed across multiple computers. For Honeywell, as for IBM, these massive projects consumed the attention of many top technical people for years before being cancelled in favor of a more evolutionary approach. Bachman moved to Boston, Massachusetts, in August 1970 to assume his NPL responsibilities. Bachman recalled

a thing called Shangri-La, which was a big project that GE had to design their new computer line. That new computer project was taken over by Honeywell in the fall of 1970, so my main responsibilities for a couple years after the GE/Honeywell merger had to do with that new product line (NPL) development.²¹

In 1973, Bachman received the ACM Turing Award, the most prestigious in the computing field. Its previous winners all

had backgrounds in mathematics or physics, rather than engineering or business. This was the first time the award had been presented to someone whose work focused on business data processing or to someone whose entire career was spent in industry rather than academia. He remains one of a handful of winners not to have a PhD.²² Bachman does not believe the award had much impact on his subsequent career progress, although it caused him to develop an interest in Alan Turing's life and even to visit his mother, then in an English nursing home.

His award acceptance lecture, "The Programmer as Navigator," was an influential declaration of a new world in which complex data structures provided the structure of corporate computing systems, around whose topographies individual application programs would navigate.²³ The award cemented Bachman's position within the industry as a leading expert on DBMSs and the most respected advocate for the network data model and the systems influenced by the Codayl approach that were rapidly gaining market share in the mainframe world. As such, he stood in opposition to the ideas of Edgar F. ("Ted") Codd, a mathematically inclined IBM research scientist whose relational model for database manipulation had attracted a growing band of supporters and was beginning to legitimate database systems as a theoretically respectable research field within computer science.

Bachman had been arguing for several years with Codd and his leading supporters, although he insists that their rivalry was friendly, noting that despite their "strong differences in ideas" they would "go out and have dinner together" afterwards.²⁴ Matters came to a head when Bachman delivered a preview of his Turing Award lecture at another meeting. As C.J. Date, Codd's primary supporter, told me

I could feel Ted gripping his chair, and the instant that Charlie sat down, Ted leapt up... he said of course, "I'd like to congratulate Charlie on the Turing Award. That's really great, and he deserves it; he's a very nice guy... None of that affects the fact, however, that what he is saying in his paper is dead wrong." That was typical Ted! Then he said, "There are so many things that need to be said in order to deconstruct Charlie's argument... This is not the right forum to do it... So instead, I would like to issue a challenge that Charlie... have a formal debate on these matters."²⁵

The debate was the centerpiece of the ACM SIGFIDET Workshop on Data Description, Access, and Control in Ann Arbor, Michigan, held 1–3 May 1974. By January 1974, plans for the debate were underway and letters were exchanged between Codd, Date, other team members, and the workshop organizers to negotiate the format. A full afternoon was reserved for "Debate: The Differences and Similarities Between the Data Base Set and Relational Views of Data." Discussion spilled over into the evening.

Bachman's strategy was to argue that the relational model was mathematically equivalent to his network approach, making the choice a matter of implementation efficiency and stylistic preference. The debate was ultimately inconclusive, as neither side managed to convince the other. At this point, the first practical relational system was still several years away, and relational products would not begin to sell in significant numbers for another decade. The choice of approaches was therefore also one between scientific rigor and engineering pragmatism, and between elegant, but unproven theory and functioning high-performance technology.²⁶

Data Models and Schemas

While unconvinced by the relational approach, Bachman did not view his own network model as the endpoint of database evolution. In the late 1970s, he worked on what he called the "role data model," extending the network model used by IDS to enhance its data-modeling capabilities. When several database records described aspects of the same real-world entity (for example, records for the same individual in roles as employee and customer), these records could be tied together as members of a set.²⁷

The role data model had similar motivations to the widely adopted technique of entity relationship modeling, introduced by Peter Chen in the 1970s. (Bachman himself called the resemblance "superficial" in my oral history, although he knew Chen, having worked with him on a joint paper on data models in this era.²⁸) However, the additional information represented in these models was not usually integrated into the data-definition languages of DBMSs themselves, as Bachman had hoped to achieve by replacing the network model.

Bachman's work on the role data model was one expression of his broader interest

in the concept of data independence. This is the principle that application programs should be shielded from the data's actual organization and location so that changes can be made to the way data is stored without requiring reprogramming. IDS had been an important but incomplete step in that direction. Database textbooks typically introduce the concept of data independence via the cryptically named ANSI-SPARC Three-Level Architecture, distinguishing between internal, conceptual, and external representations of data.

In fact, the group's promulgation of the three-tier model was something of a consolation prize. The Systems Planning and Resources Committee (SPARC) had been established as a steering committee within the American National Standards Institute (ANSI) to deal with computer standards. Its Study Group – Database Systems was established in 1972 to “review the current state of development in the database systems field with the objective to determining whether standardization activities were appropriate.”²⁹ Bachman was vice chairman of the committee and took the lead in attempting to draft a suitable standard “in such a way that the Basic Concepts [behind] IDS and the DBTG are preserved.”³⁰ The committee met frequently for years, but no consensus could be reached, something Bachman blames particularly on the influence of IBM, which at that point backed both the hierarchical and relational approaches but not the network model. As Bachman recalled this during our interview, “The sense of the group was that there was nothing worth doing in terms of standardizations, so what other interesting thing could we do.”³¹

A 1977 paper summarized the committee's attempt to at least identify the different levels and interfaces involved in a DBMS.³² The most abstract level, the conceptual schema, was supposed to represent real-world entities on a level abstracted not only from the details of physical storage but also from the design constraints that the choice of either the network or relational data model imposed at the “internal level.” The committee had agreed to differ with respect to choice of data model, but offered hope that standardization could eventually be accomplished at this new and more abstract level. Bachman sees computerized data modeling tools, able to generate database schemas for actual systems from a higher-level design, as a realization of the conceptual schema concept.

Computing Communications Protocols

Although Bachman is known primarily for his database work, his contribution to the field of computer communications was equally significant. This came when he served as chair of the Open Systems Interconnection (OSI) subcommittee of the International Organization for Standardization (ISO) from 1978 to 1982. Data communication was becoming increasingly important for corporate computer users in the 1970s, and Bachman recognized that this would take place not just between users and computers, or even from one computer to another, but from one program to another. This would require a new approach, more flexible than the proprietary, terminal-oriented communications protocols then used in the mainframe industry. Bachman had been working within his company on something he called the Honeywell Distributed System Architecture. This ultimately became Honeywell's submission to the American national committee, chaired by Bachman, which in turn formed the basis for the American submission to the international subcommittee, also chaired by Bachman. Regarding the final product, he said, “Well, I would say that 75% was based on my work at Honeywell. It was all well documented. We worked, as you sometimes have to do, to cover the traces.”³³

By 1978, Bachman's subcommittee had already produced a conceptual diagram of the communications process.³⁴ Today the entire OSI effort is remembered primarily for this seven-layer model, endorsed by the ISO's Technical Committee on Data Processing the following year.³⁵ This was an abstracted model of communication, isolating protocols at each level from the details of the levels beneath them. For example, a transport protocol (level 4) would not have to deal with network routing (level 3), physical addressing (level 2), or bit-by-bit transmission over a physical medium (level 1). As long as interfaces between levels were standardized, it would be possible to switch protocols at any level without disturbing the functioning of the other levels.

By the early 1990s, this was the organizing framework of computer communications textbooks, and posters of the model were fixtures in the cubicles of IT specialists everywhere. OSI eventually proposed standard protocols for particular layers, which gained considerable institutional backing but were ultimately swept away by the unexpected

mass adoption of Internet standards. Internet protocols do not map exactly onto the seven-layer model, but they follow its general approach of layering and abstraction.

Cullinane Database Systems

In 1980, Bachman went to work for Cullinane Database Systems, a mainframe software company. He felt underappreciated by his managers and had grown frustrated with Honeywell, which was unable to invest significant resources to develop new database software. In 1979, he wrote to a newly installed manager that "We have wasted more engineering resources and frustrated and lost more engineers by start/stop and switch operation than we have applied to new products during the last few years."³⁶

In contrast, Cullinane was rapidly growing and located close to Bachman's house. Its main product, IDMS, followed the Codsyl model that Bachman had helped to define. Indeed, it was a direct descendent of his IDS, having originally been created by a team at B.F. Goodrich attempting to produce a compatible replacement for IDS that could run on the firm's new IBM hardware. Acquired by John Cullinane, a skilled salesman, IDMS had become one of the leading software products of the early 1980s. Bachman called Cullinane and was quickly hired as the vice president of applied research.

Bachman came to view this as something of a figurehead position, telling me that "I sometimes think the role that John really had in mind for me was for me to sit by the front door and meet the visitors coming in" to show that "Charlie Bachman is here now." "I was supposed to go out and give John's speeches," he continued. "The problem was that I would go out and give my speeches, and that wasn't quite what John wanted. I think I disappointed John in some ways."³⁷

Bachman pursued several projects at Cullinane, including an attempt to revive the idea of a personal DBMS, the development of computer tools to support data-modeling work, and an extension of his ongoing work on new data models to something he called the "Partnership Data Model."³⁸ But Cullinane's growth was beginning to slow as enthusiasm switched toward relational DBMSs. Facing new pressures, the firm launched a futile effort to rebrand IDMS as a relational system and, in 1983, fired Bachman as irrelevant to its new direction.

Bachman Information Systems

Bachman left Cullinane Database Systems with a "consulting agreement" that provided a year of salary and access to computer time. On 1 April 1983, he founded Bachman Information Systems with his wife Connie. The company built on his long-standing interest in data models and targeted widespread enthusiasm for what was then known as computer-aided software engineering (CASE).

An initial project with the Norwegian firm Nordata to build a graphical design tool for IDMS did not go well, but by 1984, Bachman had partnered with Phil Cooper, a software industry entrepreneur and received US\$2 million in venture capital from two pillars of that industry: Venrock and Kleiner Perkins, Caufield and Byers. Its first product, Data Analyst, was released in 1987. This was a PC-based graphical tool to diagram data models as conceptual views, in a notation based on Bachman's partnership data model. It could translate these models into physical designs for mainframe database systems such as IDMS and could import existing metadata from databases. As a tool to support mainframes, it was priced at \$15,000 for each user, much higher than the typical cost of personal computer software.

Data Analyst was a success. Later versions ran on OS/2 rather than MS-DOS and added support for IBM DB/2 systems. This led Bachman Information Systems into a strategic partnership with IBM, which was pursuing a grand vision of metadata repositories and enterprise data models. Bachman himself had become chairman and was focusing on this work and its technical challenges while leaving management mostly to others. Data Analyst was extended to model processes and data structures, and Bachman Information Systems was soon listed on the NASDAQ exchange. The Data Analyst product name eventually encompassed four products: the Data Analyst, itself, the Database Administrator, a Forward Engineering tool, and a Reverse Engineering tool.

The IBM relationship deteriorated, however, after a rival was selected as its key partner in the area. Data Analyst was eventually ported to Windows, and the company itself merged with another firm to create Cayenne Software. The late move to Windows and its high cost meant that Data Architect was ultimately eclipsed by lower cost rivals, most notably Rational Rose.

A downward spiral began, and in 1998, the company was sold to Sterling Software

for a fraction of its earlier value. Bachman views this as the biggest disappointment of his career, saying

failure came about because I wasn't the manager the company needed. The company didn't learn to live within its revenue. Secondly... our business did not understand what it was doing and how its world was changing, fast enough to survive and grow and be a success. We blew the opportunity that Rational Software seized.³⁹

Although by then in his mid-70s, Bachman was not ready for full-time retirement. He began work as a consultant, helping companies model their data. He worked for two years with a startup called Constellar, which was trying to produce a package to help companies integrate information systems acquired during mergers. At the time of our interview in 2004, he was engaged in an elaborate data-modeling project for Cbr System, operator of the Cord Blood Registry (see Figure 2). He retained a keen interest in new ideas and technologies, although it was somewhat incongruous to see the man who invented the DBMS working with gusto to prototype a design in Microsoft Access 2003. He is currently working on a book about data modeling and his enhanced data-structure diagrams, documenting their evolution since 1964 and application to today's database systems. He considers data structure diagrams to be a new algebra.

Summary

Bachman's career was long, varied, and distinguished. He ran a chemical plant, created a capital accounting systems, headed an early data processing group, pioneered the application of computers to manufacturing control, led efforts to standardize database and computer communication concepts, won the highest honor in computer science, and founded a publicly traded company. Innovation is a collaborative process, but by creating IDS and advocating forcefully for the concepts behind it Bachman did more than any other person to invent the DBMS as we know it today.

Through all of this he retained an engineer's zest for the elegant solution of difficult problems and faith in the power of careful analysis and the systems approach to make the world a better place. As he wrote in a note at the end of our oral history transcript, "My work has been my play."




Figure 2. Charles Bachman in 2004, at home with his wife, Connie.

References and Notes

1. These details are taken from C.W. Bachman, "Oral History Interview by Thomas Haigh," 25–26 Sept. 2004, ACM Oral History Interviews collection. Material in this biography not explicitly referenced should be assumed to come from this interview, referred to hereafter as CWB TH OH.
2. CWB TH OH, p. 16.
3. SHARE, 9PAC and IDS are discussed in more detail in T. Haigh, "How Data Got its Base: Generalized Information Storage Software in the 1950s and 60s," *IEEE Annals of the History of Computing*, vol. 31, no. 4, 2009, pp. 6–25.
4. T. Haigh, "How Data Got its Base: Generalized Information Storage Software in the 1950s and 1960s," *IEEE Annals of the History of Computing*, vol. 31, no. 4, 2009, pp. 6–25.
5. CWB OH TH, pp. 23–24. B. Orchard-Hays, "SHARE Committee on Theory of Information Handling: Report TIH:1," Charles W. Bachman, Papers, CBI 125, Charles Babbage Inst., 1958.
6. C.W. Bachman, "INGEN Proposal," Charles W. Bachman, Papers, CBI 125, Charles Babbage Inst., 1 May 1959 or 1960.
7. A.D. Chandler, Jr. *Strategy and Structure: Chapters in the History of the American Industrial Enterprise*, MIT Press, 1962.
8. For a snapshot of General Electric's managerial culture in the period see R.J. Cordiner, *New Frontiers for Professional Managers*, McGraw-Hill, 1956.
9. H.F. Dickie, "Integrated Systems Planning at G.E.," *Management Control Systems*, D.G. Malcolm and A.J. Rowe, eds., John Wiley & Sons, Inc., 1960, pp. 137–155.
10. C.W. Bachman, "Presentation Prepared for the Information Systems Task Force, September 4,"

- Charles W. Bachman, Papers, CBI 125, Charles Babbage Inst., 1 May 1962.
11. T. Haigh, "Inventing Information Systems: The Systems Men and the Computer, 1950–1968," *Business History Rev.*, vol. 75, no. 1, 2001, pp. 15–61.
 12. According to Bachman SIMSCRIPT, one of the accomplishments for which Markowitz won the prize, was a continuation of work begun on General Electric's first Integrated Systems Project.
 13. C.W. Bachman, "The Origin of the Integrated Data Store (IDS): The First Direct-Access DBMS," *IEEE Annals of the History of Computing*, vol. 31, no. 4, 2009, pp. 42–54.
 14. CWB OH TH, pp. 71–75.
 15. J.R. Elliot, Jr. "More Information Please," *Baron's*, 2 Nov. 1964.
 16. V.S. Cooper, "Letter to Charles Bachman, July 14," Charles W. Bachman, Papers, CBI 125, Charles Babbage Inst., 1 May 1964.
 17. WEYCOS is discussed in J.P. Fichten, "The Weyerhaeuser Information Systems – A Progress Report," Proc. Fall Joint Computer Conference, part II, Am. Federation of Information Processing Societies, 1972; and CWB TH OH, pp. 82–91.
 18. CWB TH OH, p. 77.
 19. C.W. Bachman, "Biographical Sketch (attached to May 6, 1977 letter to Chris Bunyan)," Charles W. Bachman, Papers, CBI 125, Charles Babbage Inst., 1 May 1977.
 20. S.B. Williams, "Performance Appraisal: C.W. Bachman, Manager - Application Technology Subsection, June 1968–May 1969," Charles W. Bachman, Papers, CBI 125, Charles Babbage Inst., 1 May 1969.
 21. CWB TH OH, p. 101.
 22. Of the 55 other Turing Award winners through 2010, only four resemble Bachman in having spent all or most of their careers in industry without earning a PhD: John Backus, Ken Thompson, Frances E. Allen, and Charles Thacker. Several other winners were career academics without PhDs (such as Robert Floyd) or industrial researchers with PhDs (such as Denis Ritchie).
 23. C.W. Bachman, "The Programmer as Navigator," *Comm. ACM*, vol. 16, no. 11, 1973, pp. 653–658.
 24. CWB TH OH, p. 108.
 25. C.J. Date, "Oral History Interview with Thomas Haigh," Computer History Museum, 13 June 2007; <http://www.computerhistory.org/collections/accession/102658166>.
 26. Papers from debate participants were published in R. Rustin, ed., *Proc. 1974 ACM SIGFIDET Workshop on Data Description, Access and Controls: Data Models: Data-Structure-Set versus Relational*, ACM Press, 1974. The Bachman collection at CBI includes the program for the debate and a collection of related correspondence. The event features prominently in discussion of the period by database researchers.
 27. C.W. Bachman and M. Daya, "The Role Concept in Data Models," *Proc. 3rd Int'l Conf. Very Large Data Bases*, VLDB Endowment, 1977, pp. 464–476.
 28. An incomplete draft of "On the Basic Similarity of Three Major Data Models" is in the Bachman collection at CBI. It does not appear to have been published.
 29. C.W. Bachman, "Summary of Current Work: ANSI/X3/SPARC/Study Group - Database Systems, revised January 29," Charles W. Bachman, Papers, CBI 125, Charles Babbage Inst., 1 May 1974.
 30. C.W. Bachman, "Memo Subject: ANSI/X3/SPARC/Sudy Gorup – Database Systems Status Report," Charles W. Bachman Papers, CBI 125, Charles Babbage Inst., 24 Sept. 1974.
 31. CWB TH OH, p. 129.
 32. D. Tschritzis and A.C. Klug, "The ANSI/X3/SPARC DBMS Framework Report of the Study Group on Database Management Systems," *Information Systems*, vol. 3, no. 3, 1978, pp. 173–191.
 33. CWB TH OH, p. 114.
 34. ISO/TC97/SC16, "Provisional Model of Open-Systems Architecture," Charles W. Bachman, Papers, CBI 125, Charles Babbage Inst., c. 1 May 1978. Similar five- and six-layer models appear in the group's internal materials for 1977.
 35. H. Zimmerman, "OSI Reference Model—The ISO Model of Architecture for Open Systems Interconnection," *IEEE Trans. Comm.*, vol. 28, no. 4, 1980, pp. 425–432.
 36. C.W. Bachman, to R.A. Douglas, 12 Feb. 1979, Charles W. Bachman Papers, CBI 125, Charles Babbage Inst.
 37. CWB TH OH, p. 135.
 38. The partnership data model is described in US patent 4631664, in several documents included in the Bachman collection at CBI, and along with other aspects of Bachman's career in C.W. Bachman "A Personal Chronicle: Creating Better Information Systems, with Some Guiding Principles," *IEEE Trans. Knowledge and Data Eng.*, vol. 1, no. 1, 1989, pp. 17–32.
 39. CWB TH OH, p. 162

Thomas Haigh is a historian of computing and an associate professor in the School of Information Studies at the University of Wisconsin, Milwaukee. Contact him at thaigh@computer.org.

 Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.