A FULL-YEAR COURSE SEQUENCE IN REAL-TIME COMPUTING

D. A. MELLICHAMP
University of California
Santa Barbara, CA 93106

THE CHEMICAL ENGINEERING program at U.C. Santa Barbara began in the mid-1960's with an intended emphasis on the process control part of the curriculum. The use of digital computers for data acquisition and control, what often is referred to as real-time computing, at that time was an object of relatively intense interest in industry. However, except for one or two schools with the financial resources to acquire an industrial-scale computer (e.g. IBM 1800), there was little opportunity for universities to give students hands-on experience with data acquisition and control computers.

The appearance of the minicomputer in the late 1960's made it possible for virtually any department to acquire or build a real-time computing facility which could be used in teaching and in research. The more recent introduction of the microprocessor (and microcomputer) has merely accelerated the trend of chemical engineering departments to install real-time facilities and to introduce elements of real-time computing into the curriculum.

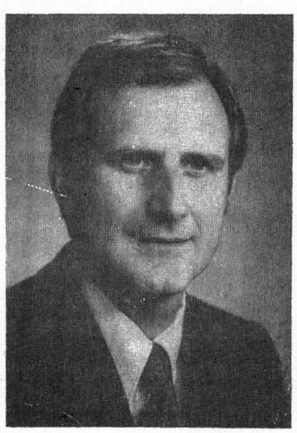

Duncan A. Mellichamp is Professor of ChE at the University of California, Santa Barbara, where he was initially responsible for development of the process dynamics and control program. An early interest in computers led subsequently to activities with the CACHE Real-Time Task Force (Chairman 1974-76), to the CACHE Corp. Board of Trustees (President 1977-78), to the Editorship of the CACHE Monograph series in Real-Time Computing, and to a long succession of aggravations associated with building and maintaining a real-time laboratory facility. He would like to forswear all future activities involving computers, but cannot.

In the past few years, the early fascination with hardware and software among real-time users has decreased to some extent, and the more important educational questions of what, where, and how to teach this new subject area are receiving more attention. Since real-time computing was and is tied to the subject of process control, it seems reasonable to explore some ideas concerning the teaching of real-time computing within that context. The chemical engineering department at U.C. Santa Barbara has developed what must be one of the most extensive teaching and research programs in real-time computing in this country (at least among chemical engineering departments); and I would like to describe it, to discuss some of the background behind its development, and to note how it is changing. To begin with, however, it might be useful to discuss the process dynamics and control program which now coordinates with the more recently developed real-time computing program.

TABLE 1
Topics Covered in Undergraduate Control Courses
(2 hr. lecture and 3 hr. lab per week for two quarters)

- Derivation of process dynamic models
- Transfer function models
- Open- and closed-loop systems
- Frequency-response methods for controller design
- Process applications
- Overview of advanced control methods

THE PROCESS CONTROL PROGRAM

IN DESCRIBING THE PROCESS control program at UCSB, it probably is fair to say that it is a traditional one. Two one-quarter undergraduate courses are required of all seniors and the usual

If there is any non-traditional aspect of the undergraduate program it would be the emphasis on "practical" experiments. The associated laboratory contains several experimental units for the study of liquid level and stirred tank heating dynamics and control, pneumatic and electronic control systems, and simulation facilities . . . .

range of topics is covered. Depending on who is teaching the course, there may be more or less emphasis on process dynamics and on "advanced control" topics; nevertheless the core areas listed in Table 1 are fairly rigorously and extensively covered.

If there is any non-traditional aspect of the undergraduate program it would be the emphasis on "practical" experiments. The associated laboratory contains several experimental units for the study of liquid level and stirred tank heating dynamics and control, pneumatic and electronic control systems, and simulation facilities for other, more complex processes. The experimental units are all bench scale and designed to have time constants on the order of one to two minutes. Since individual dynamics and control experiments can be carried out in 15 to 30 minutes, a whole range of experiments (summarized in Table 2) can be run during the two quarters. Students write up their results in brief (i.e. memo) form; although theory underlies all experimentation, emphasis is on the use of theory to evaluate practical consequences such as "Will the surge tank system overflow?", "Is the catalyst mixing system adequately stirred?", etc. By the end of two quarters students will have designed and tested

TABLE 2
Undergraduate Process Dynamics and Control Laboratory Experiments

FIRST COURSE

Liquid level system:
1. **Step response**
2. **Pulse response**

Stirred tank heating system:
3. **Step response**
4. **Transportation lags**

Stirred tank reactor (simulated):
5. **Model parameter fitting**
6. **Steady-state optimization**

SECOND COURSE

Controllers:
1. **Dynamic characteristics of three-mode controllers**
2. **Control of 1st- and 2nd-order systems**

Stirred tank heating system:
3. **Frequency response**
4. **Closed-loop control system design**

Liquid level system:
5. **Closed-loop control system design**

control systems for each of the laboratory bench scale processes using: (1) theory only, (2) empirically-determined process models, and (3) online (loop-tuning) methods.

At the graduate level we presently offer two courses in process dynamics and control. In recent years the first course has covered both time and frequency-domain methods with emphasis on state space techniques used in conjunction with the computer for analysis and design. The course deals substantially with multivariable systems and, at least in part, parallels the undergraduate courses at an advanced level (Table 3).

The second graduate course covers advanced control topics exclusively (also summarized in Table 3). Although the graduate offerings are traditional in nature (there are no laboratory experiments at the graduate level), the rigorous yet extensive coverage of advanced materials reflects the same approach as at the undergraduate level.

TABLE 3
Topics Covered in Graduate Control Courses
(3 hr. lecture each week, per course)

FIRST COURSE
- **Derivation of models for multivariable systems**
- **Formulation of state space models**
- **Solution of multivariable system models (matrix methods)**
- **Modal analysis**
- **Design of controllers using modal theory**
- **Simulation and computer-aided controller design**
- **Discrete systems analysis**

SECOND COURSE
- **Sampled data systems**
- **Sampled data controller design (or digital control algorithms)**
- **Decoupling control systems**
- **Controllability, observability, etc.**
- **Optimal control: quadratic and time-optimal**
- **Analysis of stochastic systems**
- **Observers, filters and state estimators**

THE REAL-TIME COMPUTING PROGRAM

ONE OF THE MAIN PROBLEMS in bringing new material into the curriculum is that (usually) an equivalent amount of old material will have to come out. In developing the real-time teaching program we began with four basic tenets:

1. **Real-time computing instruction will be offered to both**

undergraduate and graduate students on an elective basis.

2. Real-time computing course work will supplement, not replace, existing process control course work.
3. Real-time computing will be taught from a fundamental point of view. Students will be expected to understand basic hardware and software structures and how they are used.
4. Lectures in real-time computing will be paralleled by a "hands-on" laboratory with appropriately-designed experiments.

We have followed these principles substantially down to the present day; hence a few words of discussion might be appropriate: Tenet 1 arose out of an early realization that many chemical engineers, not just a few, will be involved with on-line process computing as part of their professional careers. Tenet 2 was based on a natural reluctance to tamper with established control courses, in particular to remove some significant amount of material so as to introduce lectures on real-time computing. Tenet 3 might well be open to argument but has its parallel in the controls area: twenty years ago many people felt that control theory was not chemical engineering; probably no one today would argue in favor of a control course based totally on an empirical approach. Our experience with chemical engineering students is that they do not, in general, like to spend time on computer fundamentals; nevertheless those who do find that practical applications are much easier to understand and are able to transfer their knowledge to other real-time systems much easier. Tenet 4 will need no explanation.

.... the ratio of "outsiders" to chemical engineers was about 30/70 in this second year; through the last academic year it has been more like 70/30, with the number of chemical engineering students relatively constant at 12-15.

Historically, we began our real-time course offerings with a one-quarter graduate seminar in Spring 1972 and followed it immediately in the Fall of that year with a senior-level elective course open to graduate students. Development of the course and the associated laboratory were underwritten by the National Science Foundation through two grants totalling almost $100,000 over a three and one-half year period. In the Fall of 1973 the course was elected by about twice as many chemical engineering students (ten). Several electrical engineering students also took the course; in fact the ratio of "outsiders" to chemical engineers was about 30/70 in this second year; through the last academic year it has been more like 70/30 with the number of chemical engineering students relatively constant at 12-15. In the past three years we have had to restrict enrollment because of the limitation in our real-time laboratory facilities.

One of the interesting developments that came out of the real-time computing course resulted from the large number of student requests for sequel courses in the same area, but covering more advanced topics. Many of the requests came from electrical engineering or computer science students who claimed that there were no equivalent applications courses within their own departments. Additional requests came from some of our own students, both undergraduate and graduate, who planned to work in the process computer control applications areas. The real-time field naturally divides into three applications areas: (1) single process/single computer, e.g. the topics covered in our first course, (2) multiple processes/single computer (multitasking or multiprogramming applications), (3) multiple processes/multiple computers (multiprocessing or networking). So far as the author can determine, it generally is the case that the real-time instruction offered by most computer science departments is (a) non-existent, (b) concerned only with on-line systems, e.g. airline reservation systems, (c) theoretically- rather than practically-oriented, e.g. concerned with hypothetical job scheduling problems in a multiprocessor environment. Condition (c) holds on our campus; hence in order to accommodate student requests we decided to add two additional courses to our offerings to cover substantially multitask programming and operating systems for real-time applications, and networking and digital computer control systems.

Several points are worth noting here concerning the decision to expand the real-time computing course to a full-year sequence:

- Some of the specialized computer-oriented material we now teach is outside the area of expertise of most chemical engineering faculty even though the applications-oriented material is not. We have avoided potential problems somewhat by using a Teaching Associate, a Ph.D. candidate in computer engineering, to share teaching responsibilities and to supervise the laboratory. In the four years we have offered the full sequence, several students working on joint research projects involving the real-time laboratory facilities have been supported financially in this way and have contributed significantly to the development of our teaching and research program.

- Attenuation of students enrolled in the sequence historically was relatively high, running 40-50% per quarter. Hence by the third quarter the enrollment might have dropped from approximately 40 to about 10; most students continuing through the entire sequence have been our own graduate students, chemical engineering undergraduates who have accepted jobs involving a process control starting assignment, computer science undergraduates, and undergraduate or graduate electrical engineering students with an interest in computer applications.
- Chemical engineering students who have taken the real-time sequence along with the required courses in dynamics and control have relatively little difficulty finding employment in process-control-related areas. Several process-oriented companies now recruit process control engineers actively at Santa Barbara and, if statements from recruiters can be believed, would have hired about twice as many students for control work last year if they had been available.
- The mixing of chemical engineering students who have relatively little computer background (in general only experience in programming a higher-level language, i.e. FORTRAN) together with computer science students who have little or no experience with physical equipment never was totally satisfactory. The distribution of abilities in any particular prerequisite subject area is invariably bimodal: e.g. ChE students will have relatively little background in binary arithmetic and logic (computer science students will feel they have mastered the subject); the reverse situation is true in the area of physical measurements and measurement errors. This situation led to major problems in the introductory course where so much of the lecture material must cover topics which will appear to be elementary to a computer science major.

This year, for the first time, we have not permitted computer science students to take the introductory course. The entire sequence has been rearranged somewhat to reflect these new developments. These actions represented an attempt to return the first course to what it originally was—an introduction for chemical engineers. At the same time we hoped to retain a reasonable enrollment of "outside" students in the two following courses. This hope did not grow out of any purely altruistic motivations; rather the presence of outside students furnished the department with a claim on the additional teaching staff resources necessary for us to offer such an extensive program. Also, in a rapidly changing field such as real-time computing, the presence of relatively advanced computer science students in our courses has kept the discussions lively and the lectures more nearly "state-of-the-art." The success of these changes is now apparent as will be noted in the sequel. In any case, this rather lengthy description of the development of the real-time sequence is intended to motivate the description of the present courses which follows immediately.

TABLE 4
A First Course in Real-Time Computing
(3 hr. lecture and 2 hr. lab per week)

- Introduction to BASIC and to real-time BASIC
- Structure of real-time systems
- Measurements, transducers, and signal handling
- Number systems and computer arithmetic
- Introduction to computer architecture and hardware
- Input/output systems: ADCs and DACs
- ISA FORTRAN
- Device controllers and device drivers

Real-Time Computing Courses. There is no "traditional" first course in real-time computing; we have, after much experimentation, settled on coverage of the topics listed in Table 4. From the table it can be seen that we spend considerable time on computer fundamentals; number systems and digital arithmetic, digital logic and hardware, computer architecture, interfacing, assembly language programming, interrupt handling, etc. We also spend time on some topics which have long since been dropped from most process control courses; measurements and measurement errors, transduction, signal transmission, etc. In a first course of this sort the emphasis is on single process/single computer systems and the coverage must, unfortunately, be light. Our purpose is to develop a basic understanding of all the elements in a real-time system, how these interact, and how they comprise the whole.

Our purpose in teaching this course has not been to treat real-time computing as an isolated subject area but to teach it so that the material can be integrated into the control courses, at least into the undergraduate process dynamics and control laboratory. Since the introductory real-time course is taught in Fall quarter and precedes the two-quarter sequence in dynamics and control, students normally are in a position to make immediate application. Those students who have elected to take the real-time course are "permitted" to run all of their dynamics (data logging) and control experiments using one of the real-time computers. Although this normally requires more work of the student—outside reading, programming, debugging programs, etc.—our experience shows that they often take this oppor-

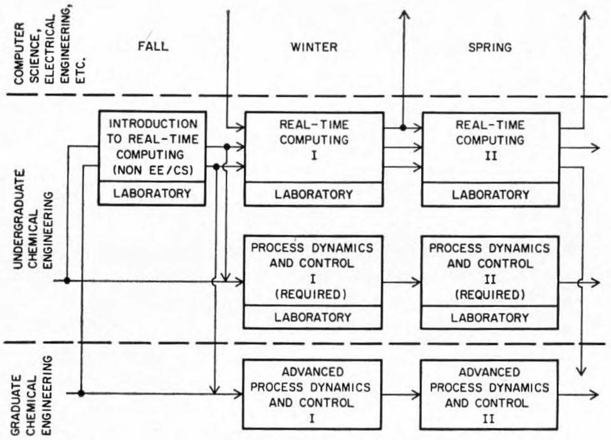

FIGURE 1. Process Dynamics and Control and Real-Time Computing Offerings.

tunity. Figure 1 indicates how the real-time sequence fits together with the undergraduate and graduate process control program, also how the three different groups of students—undergraduate and graduate chemical engineers and "out-of-department" (C.S., E.E., etc.)—can be accommodated.

The remaining two courses in the sequence (labeled "Real-Time Computing I and II" in Figure 1) cover the major areas of multitask and multiprocessor applications, respectively, with an emphasis on applications of computers, either singly or in networks, for control purposes. Tables 5 and 6 furnish a brief description of the course content for each of these courses; there necessarily must be a small degree of repetition to bring entering students up to operating speed.

Real-Time Computing Laboratory. The present real-time laboratory (shown in Figure 2 with one of the undergraduate process dynamics and control experiments visible in the background) contains three minicomputers, two of which are configured for real-time operations. These facilities, built up over the past ten years, will be substantially replaced in early 1980 by the single-computer system shown schematically in Figure 3. This multiprogrammed system will accommodate up to six real-time user programs in main memory simultaneously and, potentially, can be expanded to handle many more if program swapping using the fast disk can be tolerated. Features of the new system will include a link with the main campus computer, full graphics capabilities, a dial-up facility for one remote user, and one or more terminals in remote study rooms and laboratories.

TABLE 5
Second Course in Real-Time Computing
(2 hr. lecture and 2 hr. lab per week)

- **Overview of real-time computing**
- **Introduction to real-time FORTRAN**
- **Analog and digital input/output**
- **Operating systems and schedulers**
- **Introduction to multitask programming**
- **Multitask program design**
- **File handling and bulk storage**
- **Assembly language device driver routines**
- **Multitasking applications**

As part of the instructional laboratory we have constructed several interesting auxiliary units:

- **a set of input and display panels for experiments involving input, output and conversion of analog and digital quantities;**
- **an air pressure experiment with binary inputs (solenoid-operated valves) and outlays (pressure-operated) relays for instruction in digital I/0 and simple on-off control;**
- **a metal bar heated at one end, with eight temperature**

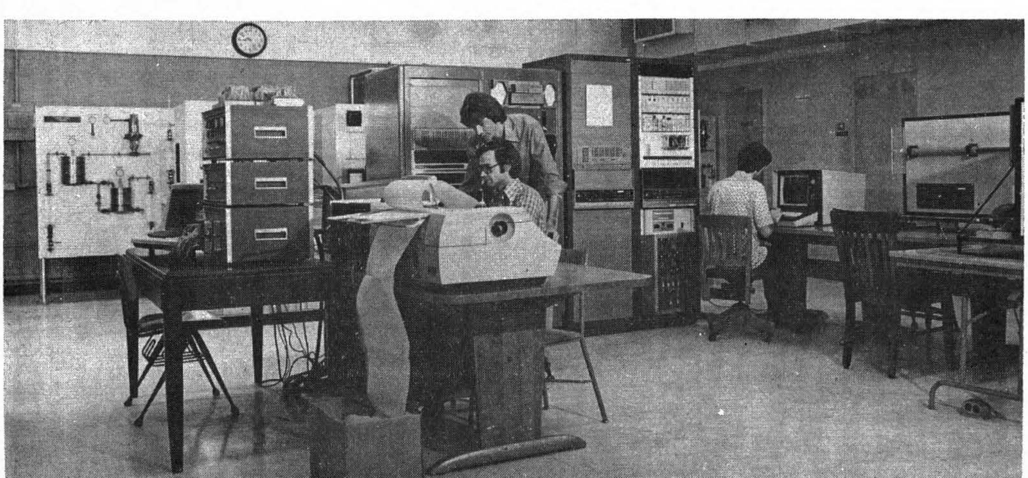

FIGURE 2. The Real-Time Computing Laboratory (Foreground).

TABLE 6
Third Course in Real-Time Computing
(2 hr. lecture and 2 hr. lab per week)

- Real-time computers in process control
- Controller and filtering algorithms
- Controller design and applications
- Overview of computer networks
- Network architectures
- Interprocessor communications
- Distributed processing
- Networks in process control

sensing elements located along the unheated section for multipoint data logging studies;

- a fully interfaced model railroad designed to demonstrate the control of multiple, largely-random processes.

A list of experiments which typically would be performed as part of the introductory real-time course is given in Table 7. The model railroad is used as the basis of a sequence of five experiments in the second course. In the third course, several of our present computers as well as the new system will form the basis for the networking portions of the course. The stirred-tank heating systems are used for the process control portions. Students completing the laboratory sequence will, as a final project, put together a two-computer real-time system (one computer for data acquisition and control, the other for process operator communications and report generation) with interprocessor communications carried out over an existing multiprocessor bus.

SUMMARY AND CONCLUSIONS

CHEMICAL ENGINEERS WHO plan to work closely with digital computer control systems need a much more fundamental exposure to real-time systems principles than can be obtained through a brief exposure as part of a senior-level control course. Even a single-quarter course in real-time computing cannot cover important advanced topics in the field such as real-time operating systems, multitasking, multiprogramming, and

TABLE 7
First Course in Real-Time Computing: Laboratory Experiments

- Calibration of a resistance thermometer for a stirred tank heating system
- Estimation of dynamic measurement error in the stirred tank temperature transducer
- Automated number conversions
- Digital input/output: "super pong"
- Analog input/output: simulation of a staircase ADC
- Data logging of the heated bar temperature profile
- Data logging and control of the pressure tank

. . . in a rapidly changing field such as real-time computing, the presence of relatively advanced computer science students in our courses has kept the discussions lively and the lectures more nearly "state of the art."

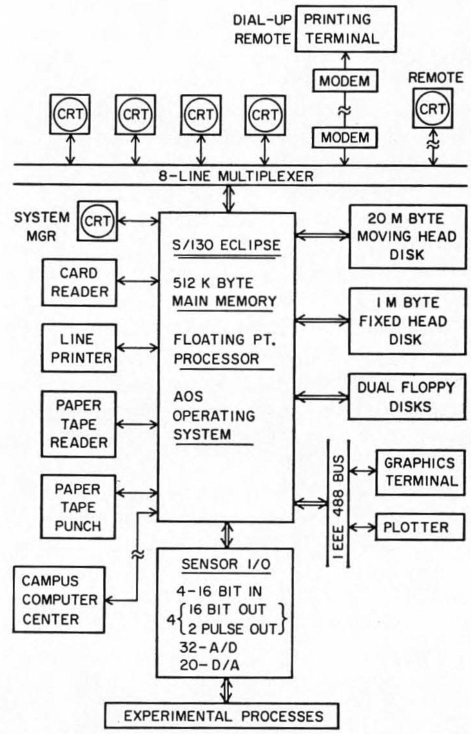

FIGURE 3. The UCSB Real-Time Computing System.

networking. At Santa Barbara we have expanded our offerings to a three-quarter sequence with heavy emphasis on laboratory exercises and experiments. As a service course, the sequence attracts enough outside students to warrant augmented teaching support staff from the college. Also, chemical engineering students appear to benefit from the experience of working with computer specialists; still it is clear that mixing them at too early a stage is not optimum. The move we have made to restrict enrollment to chemical engineers in the introductory class has eliminated most of the problems arising from mismatches in basic skills. This year, considerably more of our own students elected to take the first course than have in the past, and more of them are continuing in the sequence. Additionally, the teaching loads over the entire academic year have been considerably smoothed out by closing the first course to EE and CS students.

TEXT MATERIALS

Before 1977-78 there was available virtually no single source of information which could be used as a text in a course on real-time computing. In 1977 the first volumes of the CACHE Monograph Series in Real-Time Computing, edited by this author, appeared. The monograph series, an attempt to produce a definitive treatment of each major area in the field, will consist of eight volumes initially. All of them will be available in 1979. A listing of the titles in the present series:

I. An Introduction to Real-Time Computing
II. Processes, Measurements, and Signal Processing
III. Introduction to Digital Arithmetic and Hardware
IV. Real-Time Digital Systems Architecture
V. Real-Time Systems Software
VI. Real-Time Applications Software
VII. Management of Real-Time Computing Facilities
VIII. Process Analysis, Data Acquisition, and Control Algorithms

In 1978 a good book dealing with industrial applications of real-time computers was published: *Minicomputers in Industrial Control*, T. J. Harrison (Editor), Instrument Society of America, Pittsburgh (1978).

With respect to the laboratory facilities, the author earlier documented the three real-time laboratory experiments developed at Santa Barbara in reports, a few copies of which are still available:

D. A. Mellichamp and F. Kayihan, *The Tank Pressure Experiment*, UCSB Department of Chemical and Nuclear Engineering Report C-74-1, 79 pp, (August 1974).

D. A. Mellichamp and G. P. Engelberg, *The Digital Computer Controlled Model Railroad*, UCSB Department of Chemical and Nuclear Engineering Report C-74-3, 119 pp, (October 1974).

D. A. Mellichamp and T. W. Moore, *The Heated Bar Experiment*, UCSB Department of Chemical and Nuclear Engineering Report C-76-1, 96 pp, (March, 1976). □

ChE book reviews

CHEMICAL REACTOR DESIGN FOR PROCESS PLANTS; VOLUME I, PRINCIPLES AND TECHNIQUES; VOLUME II, CASE STUDIES AND DESIGN DATA

By Howard F. Rase
Wiley-Interscience, New York, 1977
Reviewed by Charles H. Ware, Jr.,
Commercialization Insights, Poughkeepsie, N.Y.

The author has written this book, as the preface states, "for the professional engineer who either daily or periodically must deal with design or operation of chemical reactors. But in addition to serving as a reference in the personal libraries of professionals, it should also be useful as a textbook for advanced design courses, including courses taught in continuing education." It will serve all of these purposes very well.

Volume I (772 pages) is divided into four parts: basic data and principles of design; general aspects of reactor design; single-phase reactors; and design of reactors for multiphase processes. Volume II (242) pages consists of 14 case studies including three oxidation reactions, two polymerizations, and two hydrogenations.

Part 1 is devoted to reaction rate theory and applications; chemical and physical aspects of catalysis and catalysts; idealized models of reaction rates and reactor performance; and experimental methods and equipment for developing design data. Two chapters devoted to catalysis and catalysts provide a good summary of them with attention to both theoretical foundations and practical considerations. Experimental methods and equipment to obtain chemical reaction data free of transport effects are emphasized.

Part 2 is concerned with selection of reactor type and mode of operation based upon yield and safety, as well as general design considerations such as mixing of reactants, flow distribution, residence-time distribution within reactors, and briefly, vessel design.

Part 3, which comprises almost half of the text in Volume I, covers the design of CFSTRs, tubular, batch, semi-batch, fixed-bed catalytic, fluid-bed catalytic, and many special reactors. In addition to the various design equations, there are numerous drawings of actual reactors and considerable attention is given to flow and heat effects, feed systems, pressure drop, scale-up, start-up and shutdown procedures.

Part 4 consists of an excellent chapter on gas-liquid reactors plus a short account of liquid-liquid reactors. In the former, stirred tanks, sparged vessels, plate and packed columns, trickle beds, and pipeline contractors are considered. Many theoretical and practical aspects are discussed: scale-up, heat transfer, power consumption, pressure drop, design models and procedures, hold-up, mass transfer, dispersion, liquid distribution, and many others. Only the agitated reactor is treated in the last chapter.

The case studies of Volume II have been selected to illustrate various types of design problems. They are indicated in each case, often accompanied by a comment on the principal weakness of the design. The data that are needed are presented, along with intermediate results, alternatives, and bases for decisions.

Continued on page 47.