

50 & 25 YEARS AGO



EDITOR ERICH NEUHOLD
University of Vienna
erich.neuhold@univie.ac.at



JUNE 1970

In the early years, *Computer* was only published bimonthly. Therefore, we will have to skip our interesting and/or informative extractions for June. The next one will appear in the an upcoming issue of *Computer*, and we hope you will eagerly wait for our next publication of this column.

JUNE 1995

www.computer.org/csdl/mags/co/1995/05/index.html

Binary Critic: Where the Big Money Is (p. 6) “The really big bucks are in consumer software products (about \$90 billion per year), not information processing (\$30–40 billion), ... Everyone knows about Bill Gates and the Dream Team, but the little guys in San Jose, Mountain View, Sunnyvale, and Palo Alto are also standing in line to autograph fat contracts with the “content kings.” (p. 7) “The trend is obvious: Mass media is sinking into oblivion, while products based on narrowcasting (for example, VCR tapes) are surging onto the market. ... The size of the VCR tape rental market far exceeds the box office theater market, and just about everything else. This “content” will quickly find its way onto long-playing CD-ROMs for your home computer.” [Editor’s note: *The move toward narrowcasting (personalization) has taken place, but TV stations, cable TV, PCs, and newer tablets and smartphones came and are still here. In that respect, the prediction toward PCs only has not been correct.*]

Computer Telephony Integration (p. 7) “Even though computer telephony integration (CTI) has been around since 1970, high hardware and software prices have limited it to specialized applications. However, as these prices fall dramatically, CTI is poised to explode into the mainstream. ... There are several key CTI applications that will start to generate interest in this trend, Glassman said. The first is screen-based telephony, which basically lets someone use a keyboard or click a mouse on an icon to have a computer dial a telephone number. The

second is call-based data selection, where an organization uses a phone company’s caller ID system to identify a caller. The phone company’s computer accesses a database and shows all relevant customer information before the call is answered.” [Editor’s note: *Despite some truth in this prediction, what was missed was the advent of the Internet and the rise of smartphones that brought multimedia communication to the end user, bypassing the CTI protocols envisioned in 1995.*]

Implications of Classical Scheduling Results for Real-Time Systems (p. 16) “The scheduling theory literature is so vast that we can’t pretend to be comprehensive, but this article does present a minimum set of results and their implications. The set includes Jackson’s rule, Smith’s rule, McNaughton’s theorem, Liu and Layland’s rate-monotonic rule, Mok’s theorems, and Richard’s anomalies.” (p. 24) “Most multiprocessor scheduling problems are NP, but for deterministic scheduling this is not a major problem. We can use a polynomial algorithm and develop an optimal schedule if the specific problem is not NP-complete, or we can use off-line heuristic search techniques based on classical theory implications. These off-line techniques usually need to find only feasible schedules, not optimal ones. Many heuristics perform well in the average case and only deteriorate to exponential complexity in the worst (rare) case. Good design tools would allow users to provide feedback and redesign the task set to avoid the rare case.” [Editor’s note: *This article provides a very useful analysis of scheduling processes and specifies around 20 known theorems. Thus, it should help a designer to select the best-fit methods to solve a specific scheduling task. Some of these insights are certainly still useful today.*]

A Specification-Driven Architectural Design Environment (p. 26) “In this article, we introduce an environment we’ve developed that helps designers represent, model, and explore trade-offs at the architectural abstraction level and synthesize designs at the behavioral level. The Design Analysis and Synthesis Environment (DASE) accomplishes this by supporting design capture, design space exploration, and

validation for the final design synthesis.” (p. 28) “DSL, the internal specification language in the Design Analysis and Synthesis Environment (DASE), is a Prolog metalanguage that is interpreted through a processor. Modular object-oriented design entities called modules are its primitive building blocks.” (p. 33) “Abstracting hierarchy can facilitate simulation and modeling in varying detail. For example, while a systems analyst might not focus on a model’s lower-level details, a hardware designer would be concerned with the simulation’s timing details. By defining an observation level, the DSL simulator enables dynamic alteration of the abstraction level viewed by the user. ... Module behavior can be translated to a netlist based on petri nets, which can help analyze hardware or software systems. The petri net formalism simplifies definition of asynchronous concurrent communications.” [Editor’s note: The article represents the methodology in detail and also claims, more briefly, that it has been used in the development of various forms of asynchronous transfer mode switches.]

Interrupt Processing in Concurrent Processors (p. 36) “To help designers systematically explore options for handling interrupts and help researchers compare interrupt processing strategies, we offer a taxonomy (or classification) of implementation choices. The approach we’ve developed broadly classifies interrupt-processing techniques and implementations into six phases.” (p. 38) “Interrupt-processing systems can be implemented in many ways. Some techniques, like history buffers, are general enough that the designer could make all processor interrupts precise, depending on the desired end result. With other strategies, such as adding special-purpose registers to a processor, the designer might make only one type of interrupt precise—for example, I/O interrupts.” (p. 42) “Interrupt Processing Phases: • Detect the interrupt. • Finish pending instructions. • Undo process state changes. • Save the process state. • Run the interrupt handler. • Resume the interrupted process.” (p. 45) “We can therefore conclude that, to be completely accurate, each possible processor interrupt (or type of interrupt) should be classified on an individual basis. As a whole, processors tend to resist being conveniently classified because of the kinds of special cases just described.” [Editor’s note: The article essentially analyzes the different kinds of interrupts and how to treat them. The six phases mentioned above are analyzed in detail with different alternatives of implementation explored. Unfortunately, the article concludes with a remark that the classification of existing implementations is not reasonable, as they contain too many different choices.]

Fault Injection (p. 47) “Fault injection is an effective solution to the problem of validating highly reliable computer systems. Tools such as React are facilitating its application. ... In critical applications, such as aircraft flight control, nuclear reactor monitoring, medical life support, business transaction processing, and telecommunications switching, computing resource failures can cost lives and/or money.”


(p. 48) “Taxonomy of Experiments: Fault injection experiments can be classified according to three general attributes: system abstraction, fault model and injection method, and dependability measure.” (p. 50) “Not until the mid-1980s did academia begin actively using fault injection to conduct experimental research. Initial work concentrated on understanding error propagation and analyzing the efficiency of new fault-detection mechanisms. Research has since expanded to include characterization of dependability at the system level and its relationship to work load.” (p. 53) “The studies discussed so far focused on validating existing systems, but fault injection can also evaluate the dependability of proposed designs. We have used simulated fault injection to analyze the reliability of several alternative TMR (Triple Modular Redundant) architectures.” [Editor’s note: The article analyzes a range of possible fault-injection mechanisms, investigates applications of those mechanisms, and then discusses various tools that have been used in those and other applications. Again, this is a survey article that has stayed interesting over the 25 years since its publication.]

International Survey: Virtual-Environment Research (p. 57) “A key development goal is to involve many—eventually all of our senses in acquiring information, although initially the focus has been on visual information display. Technologies for presenting information to other sensory modalities are being developed, such as acoustic displays that present 3D sound environments to a person’s sense of hearing and haptic displays that provide tactile and force feedback to a person’s sense of touch. Technologies for presenting information to other senses, such as smell and taste, are occasionally discussed but have rarely been incorporated into a VE.” (p. 65) “VIRTUAL ENVIRONMENT RESEARCH CLEARLY INVOLVES a broad range of technologies and potential applications with the intriguing potential to affect us in previously undreamed-of ways. Internationally, many researchers are independently conducting ongoing projects to develop these technologies and to incorporate them into effective VE systems. The recent emergence of commercial VE applications and the enterprises needed to support them indicates a maturing VE industry that is ready for expansion.” [Editor’s note: Interestingly, the techniques and applications discussed in this survey are not much different from what exists today. Of course, price and performance stood in the way of widespread utilization. As in other cases, the gaming industry finally led to mass-marketed availability of many of the discussed methodologies.]

Resources for Networks in Less-Industrialized Nations (p. 66) “True global connectivity is years away. Countries must first have an internal network, no matter how small, before linking electronically to the rest of the world’s networks. ... The good news is the Internet’s phenomenal growth. The bad news is that the growth is concentrated in the Northern Hemisphere: North America, Western Europe, and parts of Asia...

COMPUTING THROUGH TIME

ERGUN AKLEMAN



DURING MEDIEVAL TIMES, PLAGUE DOCTORS WORE A BEAK-LIKE MASK DURING A PANDEMIC. THESE MASKS, FILLED WITH AROMATIC ITEMS SUCH AS HERBS, WERE SUPPOSED TO PROTECT THEM FROM PUTRID AIR, BASED ON MIASMATIC THEORY OF DISEASE. THIS THEORY HELD THAT THE ORIGIN OF EPIDEMICS WAS DUE TO A MIASMA (A NOXIOUS FORM OF BAD AIR), EMANATING FROM ROTTING ORGANIC MATTER.

Digital Object Identifier 10.1109/MC.2020.2984375
Date of current version: 4 June 2020


LINs are sparsely connected. Furthermore, connectivity is unevenly dispersed in industrial nations. For example, children in some Latin American schools have access to the Internet while those in US inner cities seldom do. ... Consider the evolution of the Relcom (Reliable Communications) network in the ex-Soviet Union." (p. 67) "While it has been successful, Relcom is unlike a more industrialized nation's national network. Where a typical US network uses large computers and workstations connected with dedicated high-speed communication links, Relcom uses mostly PCs and dialup telephone connections." (p. 70) "Eventually, outside experts must be supplanted by local expertise and self-sufficiency. When a network has been established, the need for people resources shifts from technical to managerial and financial skills. ... Pre-university education is an area of vast discrepancies. Ivan Illich and others have pointed out that education is often used to maintain the social status quo. Nations such as Cuba, Malaysia, and Costa Rica have ambitious computer literacy programs to reach all people - rich, poor, rural, and urban." [Editor's note: This interesting article shows us again what the situation was in 1995 and what tremendous progress has happened the world over. However, despite the fact that Wi-Fi was around


then, the article totally ignores the effect that mobile telephony has had in the development of the Internet for all.]

On-Line Monitoring: A Tutorial (p. 72) "On-line monitoring can complement formal techniques to increase application dependability. This tutorial outlines the concepts and identifies the activities that comprise event-based monitoring, describing several representative monitoring systems. ... Monitoring gathers information about a computational process as it executes and can be classified by its functionality. Dependability-Performance enhancement-Correctness checking-Security-Control-Debugging and testing-Performance evaluation. ...I focus on (first) four of the seven functional areas." (p. 77) "ON-LINE MONITORING IS INCREASINGLY SEEN AS A VIABLE means of increasing application dependability." [Editor's note: The tutorial discusses a number of monitoring systems available in 1995. They are used to increase system dependability. However, the same techniques are now used to monitor a vast number of applications, among them manufacturing, logistics, traffic control, people movement, crowd control, and so forth. Unfortunately, some of these can be used maliciously for the invasion of individuals as well as corporate privacy.]

Open Channel: The Virtual and the Paperless Office

(p. 120) “Among the most promising developments of the telecommunications age is the virtual office. The potential benefits are remarkable, we hear: less traveling, better use of time, higher productivity, fewer dry cleaning bills. ... There is still one major factor, however, that could spoil the whole virtual office idea: the quality of those remote interpersonal exchanges. ... Team productivity today hinges on those inef-
fable informal exchanges (sometimes known as chit-chat)

fostered by physical proximity. ... Will cyberspace ever replace such contact? ... There is another precedent here of a technology that just isn’t good enough to be exploited in full. Think about WYSI-WYG, it’s been around for years, yet how many of us still printout hard-copy drafts of our work, the better to review?” [Editor’s note: *The progress over the last 25 years has not eliminated physical office space and personal contacts (look at the amount of business travel). Also, the use of paper in office work is far from being eliminated. Just ask the paper industry.*] 



BPCnet




RESOURCE PORTAL



The BPCnet Resource Portal is a clearinghouse for the community to learn about and engage with ongoing projects to diversify computing. BPCnet serves two key roles:

- Increasing the visibility of ongoing broadening participation efforts and adoption of effective efforts.
- Assisting NSF CISE PIs in planning meaningful broadening participation in computing (BPC) plans in their proposals.

<http://bpcnet.org/>

An initiative by the Computing Research Association (CRA) with support from the National Science Foundation's (NSF) Directorate for Computer and Information Science and Engineering (CISE). CRA and CISE are committed to broadening participation in computing (BPC).

Resources on BPCnet are designed to help NSF PIs and departments learn about, and engage in, BPC activities and develop impactful BPC plans that address issues in their context.

Check out BPCnet.org for upcoming workshops on developing BPC plans!