Real-Time Computing: A Critical Enabling Technology

by

Jack Stankovic, IEEE Fellow – Univ of Massachusetts
Kang Shin, IEEE Fellow – Univ of Michigan
Herman Kopetz, IEEE Fellow – Tech. Univ of Vienna
Krithi Ramamritham – Univ of Massachusetts
Lui Sha – SEI-CMU
Doug Locke – IBM
Jane Liu – Univ of Illinois
Al Mok – Univ of Texas, Austin
Susan Davidson – Univ of Pennsylvania
Insup Lee – Univ of Pennsylvania
Jay Strosnider – CMU

Executive Statement

Real-time computing is a generic enabling technology for many important application areas including: flexible manufacturing, multimedia, intelligent vehicle highway systems, telecommunications (the information super highway), robotics, process control, avionics, aircraft flight programs, air traffic control, real-time simulation, virtual reality, space applications, nuclear power plants, intensive care monitoring, and defense applications such as command, control and communications. In particular, almost all safety critical systems are real-time systems. It is **astounding** that such a critical and central technology is not the **focus** of any current research agenda of any major funding agency. The effect is that for each of these application areas real-time research is funded in a piecemeal and unfocused manner, thereby sometimes *re-inventing* the wheel, slowing down progress, introducing higher than necessary cost, and possibly producing inferior (or even dangerous) solutions. A strong research and transfer of technology program in real-time computing would serve to produce generic results that can be used in all these important applications, as well as in new ones that develop in the future. The benefit of such a research program is likely to be immense since the results are potentially usable by so many industries critical to the economy, public safety, and the nation's defense. A focused research program should be funded and put into place as soon as possible to capitalize on the current research infrastructure and momentum formed in part by ONR's recent, but expired research initiative on real-time computing.

1 Real-Time Computing as an Enabling Technology

Real-time computing is an enabling technology for many important applications. Several key examples serve to illustrate the safety critical and important financial impact of real-time computing.

Real-time computing is playing a key role in many industrial sectors such as the automobile industry. Auto manufacturers can only remain *competitive* if they incorporate state of the art real-time computing systems into their cars. In the future distributed real-time control systems will replace and enhance many of the conventional control systems of the car, making the cars more efficient and adding to public safety. Before distributed real-time control systems can be used there are a number of significant research challenges that must be addressed: precise real-time response to the microsecond, fault tolerance under strict timing requirements, maintainability, and testability under competitive pricing pressures of the automotive industry. A strong research and transfer of technology program can help provide the new technical solutions needed more quickly than if left to the automotive industry itself (resulting in quicker and greater economic and safety benefits), and can facilitate the exploitation of current real-time technology in this application domain.

Another important real-time system on which the economy is *highly* dependent is Air Traffic Control. This application area must continuously manage massive amounts of data. Unlike some large data management systems, such as airline reservations, the air traffic control data is constantly changing, and has extremely high value (related to public safety) for very short amounts of time (response requirements varying from a few milliseconds for radar data to several seconds for flight control information). At completion, the air traffic control system itself is estimated to cost over five billion dollars. This system is a good example of where a combination of recent real-time and fault-tolerance generic research results have been transferred to significantly contribute to the FAA's Advanced Automation System. However, the system is so large and complex (e.g., the system will have between 1 and 2 million lines of code and thousands of consoles) that new real-time research is needed to improve the safety even further, lower the cost of the system and its maintenance, and provide for its continual evolution as the system grows in size and complexity.

Maintaining technological leadership in manufacturing is vital to industrial competitiveness. Two areas of importance are flexible and intelligent manufacturing systems. It has been estimated to cost as much as 5 billion dollars for research and development of these technologies. While many areas need to be developed to support flexible and intelligent manufacturing, one key ingredient often overlooked is support for real-time computing. Poor support for real-time computing lowers efficiency of the plant, sometimes causes catastrophic line breakdowns and lowers profits. Good real-time support can link CAD front-ends with the plant via real-time database techniques, improve just-in-time delivery of subassemblies to various workcells, guarantee completed products in time, keep workcells highly utilized, and increase profit margins. Many research results that can improve automobile control and air traffic control are also beneficial to manufacturing. For example, techniques for timing analysis can reduce costs significantly and increase the safety and quality of the system, regardless of whether it is an automobile, air traffic control, or flexible manufacturing.

Major programs in areas such a flexible manufacturing, intelligent vehicle highway systems (IVHS), dependable computing, etc. have real-time computing as a *prerequisite* to success. Many aspects of these systems deal with software issues. Capers Jones, President of the Software Productivity Research Inc., Burlington, Mass., estimates that one-half of software programming is in information systems applications, one-fourth in the military, and one-fourth in real-time control applications. Since much of the military software deals with real-time and so does an increasing portion of the information systems applications, it is evident that over 50% of all software deals with real-time issues. This translates into a financial impact in the billions of dollars. New techniques are needed to produce real-time software with both timing and logical correctness. For example, consider the following statement by Dr. Larry Druffel, Director, Software Engineering Institute, CMU,

> "I am convinced that the application of engineering discipline to software systems with real-time constraints is the most important and the most difficult technical issue facing the software community. It is vital to defense and commercial competitiveness, and is worthy of the best minds."

While much research must still be performed in real-time computing, current research results should also be utilized to a greater extent. To facilitate this transfer of technology there should be a strong transfer component in *new* funding efforts. As one example of some recent real-time results that has had major impact on real world systems, consider the advances to rate monotonic scheduling. The rate monotonic scheduling algorithm has been used as a key technology in the Space Station Freedom Project, and the FAA Advanced Automation System (AAS), and has influenced the specification of the IEEE Futurebus+. The DoD's 1991 Software Technology Strategy refers to Rate Monotonic Scheduling (RMS) as a "major payoff" and states that "system designers can use this theory to predict whether task deadlines will be met long before the costly implementation phase of a project begins. It also eases the process of making modifications to application software." The Acting Deputy Administrator of NASA, Aaron Cohen, stated in a 1992 speech titled, Charting The Future: Through the development of Rate Monotonic Scheduling,

> "we now have a system that will allow (Space Station) Freedom's computers to budget their time, to choose between a variety of tasks, and decide not only which one to do first but how much time to spend in the process."

Other examples of successful utilization of research results are the use of new real-time systems formal methods in Navy applications and the emergence of new specification, development and analysis tools based on real-time scientific theory. This technology was developed under the recently expired ONR focused research program and has subsequently benefited many different areas. There is a pressing economic and public safety need for a new research program to continue to produce such successes.

2 Real-Time Activities and Need for Immediate Action

To demonstrate that real-time computing is an extremely active field, we discuss a few recent activities and why a new funding initiative should be put in place as soon as possible.

First, as mentioned above, the Office of Naval Research has recently completed a five year research initiative supporting real-time computing. While this initiative was relatively small in total funds and has now expired, it was extremely successful in developing novel concepts, algorithms, and technologies for important problems in real time computing. As a result of this research program many new researchers entered the field creating a relatively large set of new results. The need for new funding is immediate so that the momentum is not lost on this work. Funding is needed to integrate existing solutions, study their applicability to various application areas, transfer results to industry, identify deficiencies in real-time technology, and based on the deficiencies identified, create new results for current and new application areas.

Another important development is that there are major application-driven research initiatives in flexible manufacturing, multimedia computing, and intelligent vehicle highway systems, each with substantial funds due to their obvious potential economic impact, but without any specific and focused effort on real-time computing. This is an oversight giving rise to a pressing need for such a program because a specific research focus on real-time could pay dividends in many ways including:

- safety critical systems can be made more robust, considerably improving public safety
- products can be manufactured more cheaply and with better quality, improving competitiveness
- services can be delivered on-time (time is money)
- command, control and communication for defense applications can be improved to meet new mission requirements of the 21st century.

Japan and Europe have been extremely active in real-time computing activities. For example, a Japanese consortium has designed a set of interfaces based on a real-time kernel called TRON. Using this kernel they are developing a autotraffic information system, an intelligent real-time controlled house, office building, and city. The economic benefits of such systems, if they are successful, would be extremely large. In Europe, many companies, government labs, and universities have research groups focussed on real-time computing; the Esprit project funds real-time computing with some of the projects having an aggressive agenda at focusing on real-time computing and leveraging results into applications and products.

Real-time computing is now becoming a recognized discipline in academic departments as new courses are being defined. It is expanding via a growing number of professional conferences. The following is a quote from the announcement for the EUROMICRO Workshop on Real-Time Computing for 1994, "The sixth Euromicro Workshop is a forum that covers the state-of-the-art research and development in real-time computing, which is becoming an essential discipline in the field of computer science and engineering."

3 What is Required

Since real-time computing is a prerequisite to so many extremely important and safety critical systems and so prevalent in the software field, basic research in real-time computing should have a very high priority. In the past, practitioners were satisfied with seat-of-the-pants solutions that twiddled code to reduce overhead and speed up the execution. Complexity has now overtaken this approach. It is also well recognized that fast computing is not real-time computing, so these applications cannot depend on efforts such as high performance computing (HPC) to solve all their problems. More scientific results are needed and these results must be transferred to products. The funding is required now because there are so many applications for which real-time computing is a prerequisite and the research field is primed to make important contributions to these real world applications. If there is a delay the current momentum may be lost.

A proposed research and transfer of technology program could focus on the following:

- Re-engineering existing real-time systems with the state-of-art research results. This serves to move the more mature results into practice.
- Concept demonstration for some important research results by implementing and evaluating them for typical applications/scenarios. This would take generic results and adjust them for benefit to various industries. New research and novel benefits are likely to emerge.
- Experimental research such as testbeds and prototype implementations. This enables development of theoretical and applied results under constraining (environmental, hardware, software, and cost) conditions.
- Understanding the explicit real-time computing issues in key applications such as flexible manufacturing, IVHS, process control, and safety critical systems. Research programs in each of these areas should devote significant funding to real-time computing and should interact with the more generic and basic research effort being called for in this paper. In this way the funding available is expanded and the research itself is better served by strong ties to the particular application domains.
- Development and dissemination of sound engineering methodologies and supporting tools for real-time design, coding, implementation, and evaluation.
- Use of formal methods in real-time systems.
- Development of requirements and internal documentation which can state real-time requirements precisely and can be used for maintenance and inspection of safety critical systems.
- Integration of timeliness, fault-tolerance, and security.
- Execution time scheduling and analysis in concurrent processing environments built with contemporary powerful/complex hardware and software.

- Combination of real-time technology, artificial intelligence, and database and multimedia systems.
- Developing the scientific principles underlying real-time computing.

4 Summary

Real-time computing is an *enabling technology* for many current and future applications which affect public safety, competitiveness, the economy, and life style. Many results have been developed, but difficult research and transfer of technology issues remain. A major funding initiative in real-time computing at this time would capitalize on current results, establish generic technology for the future, and pay large dividends for the economy.