

## "Real-Time Isolation and Composition of Virtualized Resource Sets"

Aloysius K. Mok  
Department of Computer Sciences  
University of Texas at Austin  
Austin, TX 78712  
email: mok@cs.utexas.edu  
phone: 512-471-9542

### 1. What are the most important challenges?

A critical issue in designing survivable systems on the scale of Internet-connected devices is the capability to adapt system configuration without centralized control and in the face of an attack. There are a number of reasons why centralized control is impractical in this case: inasmuch as the communication latency on the Internet would preclude real-time response to local disturbances, some degree of local autonomy is unavoidable; ownership issues are also an impediment to centralized control. Thus a necessary condition before advanced feedback control strategy can be deployed to contain attacks is the ability to reconfigure system resources under distributed and semi-autonomous control. If we look back at the history of computer science, a common theme for maintaining control of system resources is to start with the definition of a configuration space. But what is a configuration space in this case?

A configuration space can be viewed as a formal description of how system resources can be organized into independent units with well defined QoS attributes. Depending on the computational and communication requirements of the feedback control scheme, a system configuration is to be selected to meet the QoS requirements of the feedback control scheme, and this must be done without centralized control.

A proven way to enforce a system configuration is to virtualize the available resources. However, traditional virtualization techniques often abstract out system attributes such as timeliness, security and fault tolerance properties. Therefore, a critical research issue is how to achieve virtualization of all types of resources on the Internet while respecting timeliness, security and fault tolerance requirements.

### 2. What are the most important information technology research needs?

We need to revisit the idea of virtualization in a fundamental way. We need to be able to virtualize physical resources to yield not only "logically independent" resource sets; the virtual resources must be made to behave as if they are slower (shared) but can support scheduling policies that are invented for applications running on dedicated resources; virtualization should also address security requirements, so that we can ideally be able to say that a system has attained a configuration with resources that can support such and such real-time, fault-tolerance and security requirements. New research is needed that must cross traditional delineation of research areas. For example, timeliness issues are often studied by the real-time systems community but the implicit assumption there is that the underlying resources must be reliable. On the other hand, the usual

assumptions underlying most practical fault-tolerant systems include at least some implicit timeliness requirements; for example, it is impossible to distinguish a crash failure from a performance failure that is due to a transient overload without a time-out mechanism. We need a framework that considers fault tolerance and timeliness issues simultaneously and we need new design paradigms that achieve simultaneously design objectives that span multiple areas of concern.

On the experimental side, we need to build and experiment with system architectures that support the new design paradigms. Aspect programming a start but it suffers from a lack of carefully thought out a framework that integrate different design concerns; it also suffers from being at too low a level of design (programming language) that may overcommit a design because of ontological reasons. We should

3. What is a possible roadmap for the next 5 to 10 years?

For the next 5 years, we should encourage research that is cross-area. This includes the integration of real-time and fault tolerance concerns, real-time and security concerns etc. We should also select a distributed process control application implemented with current SCADA-based technology and use it as a benchmark for experimentation of new design paradigms and reengineering with new architectures. We should also form blue and red teams to explore possible attacks on SCADA-base systems and to formulate a realistic framework for attack modeling. It is important to be able to SYSTEMATICALLY identify after five years the major weaknesses of current systems and prioritize research in new paradigms and architectures that address them. This sets an intermediate goal that will help us formulate a yardstick against which to measure our progress.

In 10 years, we should be able to deploy some of our research results in fortifying our systems against exploitation of the identified weaknesses in the previous five years. We should also be in a position to start establishing standards and criteria for certification, V&V and to start automating the design paradigms that we have identified. I do not see the 5 and 10 year time frames as two distinct phases because it is more likely that we shall continually learn new lessons, revise "new" paradigms as we gain more understanding. I do see the value of setting a short/medium-term time frame, say 5 years to address the weaknesses in current systems. We need to fix things fast enough to gain the public's confidence while on our way for a scientific, comprehensive solution for the long term.

\* Biography of Aloysius K. Mok

Aloysius Mok obtained his B.S in electrical engineering, a M.S. and Ph.D. degree in computer science, all from the Massachusetts Institute of Technology. Since 1983, Dr. Mok has been on the faculty of the University of Texas at Austin, Austin, Texas where he is currently Quincy Lee Centennial Professor in Computer Sciences. Dr. Mok is

known internationally for his work in real-time systems design and was a past chair of the Technical Committee on Real-Time Systems (TC-RTS) of the Institute of Electrical and Electronics Engineers. He received the IEEE TC-RTS Award for Technical and Leadership Achievements in 2002. Dr. Mok has provided expert advice to government and industry on problems in embedded systems engineering, including the USAF F-22 Raptor, F-35 Joint Strike Fighter avionics and the SBIRS-High signal processing architecture. Dr. Mok's research interests focus on formal design methods for fault-tolerant, secure, real-time systems for embedded applications.