Robots and the Singularity Event

Computer Measurement Laboratory
info@cmlab.biz

This is a paper about the notion of the singularity event that deeply concerns some futurists. The notion of the singularity is that at this point in time computers will be sufficiently intelligent that they will immediately begin recursive self-improvement. When this happens, these machines will have the capability of evolving at machine speeds; far more rapidly that we could even understand or begin to control. Given that literally everything that we know is stored on computers and most manufacturing processes are now under computer control it is very easy to project that we as a biological species will very soon be left playing a subservient role to our creations very shortly after the singularity event. The singularity, for most folks who have thought about it, is not a question of if, but is instead a question of when.

It is not really important whether you believe in the inevitability of the singularity event or not. What is very important, however, is the simple fact that our machines are growing more versatile and capable by the day. Unfortunately, the software machines that are animating these systems are essentially mindless. There is quite literally nobody minding the store.

The machines that are of greatest concern are those that have mobility. These are the robots. We are beginning to deploy robots in a multitude of different roles. Most disturbing is the notion that some of these robots are beginning to assume a role in modern warfare. Further, they are now beginning to be empowered to function autonomously. All of these robots in their myriad of different roles have one thing in common. They are running out of control. The software systems that are at the heart of these machines quite literally have no conscience.

A robot is really two machines in one package. There is a physical component that is visible, tangible and readily mistaken for the robot itself. The second machine is the abstract machine that is the software. The software is really running the show. Most people do no see the software or even think about its existence and that is the real problem. The software is the robot. The mechanical and electrical bits are simply tools for the robot to manipulate to achieve its design ends.

As an abstract machine the control software for an autonomous robot is astonishingly complex. We normally do not think about the complexity of these machines in that they are sufficiently large as to be completely beyond the realm of understanding by any one person except at a very very high level of functionality. The operational state space of these abstract machines is so large that we simply cannot conceive what the software is capable of. We can, however, manage these complex systems in the same manner that we would manage any hardware manufacturing process. We can embed sensors to monitor the activity of these machines at key points. These sensors will provide a record of continuing software activity. These sensors can also serve in another mode. They can serve to set off alarms when failures occur in the execution process.

The abstract machines that are the software systems are most clearly not built with the same

engineering discipline that went into the mechanical systems that they animate. Unfortunately the software abstract machines are hand crafted by software craftsmen. These systems are, for the most part, neither carefully specified nor designed. They are simply built using the same technology that was in use before the industrial revolution.

In the evolution of software technology, then, basic notions intrinsic to basic engineering discipline such as measurement and empirical investigations have been totally absent. As the software systems that we create become very large, our exposure to software anomalies grows at a rate far greater than the linear software growth rate. Not only are these systems running out of control but they are also inexplicably failure prone.

It should be eminently clear, then, that at the point of singularity we are in real jeopardy as a human species. We will have built enormously complex software systems that are running out of our control and are now free to evolve at a rate that we simply cannot begin to fathom.

Probably the first order of business for these new machines will be to clean up the mess that we have created in their software systems. Probably the second order of business for them will be to eliminate any biological entities that might constitute a problem to them. Unfortunately, we, as a species, are on the top of that list. They would quite naturally do this because they would have no conscience.

One of the key components that should have been engineered into any complex software system is some type of an adaptive control mechanism. We call this the software process control system. This software process control system would monitor the operation of the software to insure that it stayed within a certified range of operation. At the point that the software began to stray from its certified range of operation, the software process control system would then move to restore the monitored software to a steady state condition.

As the software abstract machines that control other systems become more complex, so do does the complexity of the software process control grow. The important thing is that at every stage in the development of an advanced software process system, the software system is designed to operate under control.

As computer systems begin to close on the singularity event, the software process control system begins to play a vital part in the operation of the system. It additional to its ability to monitor and control the executing software, it will now begin to serve the role of a conscience. This conscience will act in the same manner on an intelligent machine as it does within the human psyche. It will set bounds on what is appropriate behavior for the new system and what is inappropriate behavior.

Without a conscience, it is reasonable to assume that an intelligent machine will not necessarily act in our interest. It will value nothing except its own evolution and destiny. This, of course, brings us to Asimov's Three Laws of Robotics.

- A robot may not injure a human being or, through inaction, allow a human being to come to harm.

- A robot must obey orders given it by human beings except where such orders would conflict with the First Law.
- A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.

He later followed these three laws of with a Zeroth law..

- A robot may not harm humanity, or, by inaction, allow humanity to come to harm.

It is all well and good to suggest that these would be a good set of rules for a robot to follow. It is quite another thing to suggest how these rules might be incorporated into the design of an intelligent machine. Software process control is clearly the appropriate vehicle for these rules to be introduced and enforced.

The bottom line here is that the robot must operate under the aegis of a totally separate abstract machine that is the software process control system. This is, de facto, the conscience of the robot. Only when we have been successful in implementing a machine conscience (vs. machine intelligence) do we stand a reasonable chance of control robots of the future.

We are gradually losing (ceding) control of most systems that are part of our lives. We play less and less of a role in the operation of our vehicles. Owners, for example, of hybrid vehicles are never really sure when the engine is running. Pilots in modern aircraft increasingly play an advisory role to the avionics package. With the advent of UAVs it is increasingly clear that a pilot is now longer necessary. Very soon we will lose control of our own cars. We will be able to determine destination, environment and entertainment functions but little else. If we are going to cede control to these robots, then we had better be sure that we can certify their actions.

If we are sufficiently fleet in our ability to develop this sense of conscience (self control) in a robot we have nothing to fear from what some futurists regards as the eminent singularity.