Compositional Assurance for Robotic Coordination

Joseph Giampapa

As autonomous physical machines, Robots must operate in unpredictable, remote, and dynamic (ever-changing) environments. It should not be a surprise, then, that assurance, or the guarantees that individual robots will perform without failure and as expected, is more of an imprecise and expensive art than a science. For example, the Mars Exploratory Rovers (MERs) were designed for a mission that was intended to last 90 days, but the products of extensive and expensive assurance activities actually kept their missions going over three years on Mars. At the opposite end of the spectrum, studies of terrestrial robots have shown failure rates as high as 40% (cited from memory) for fielded robots. Failure analyses initially blamed mechanical failure as the main cause, but a re-analysis and the consideration of more data indicated a bi-modal trend that included human error as much as mechanical defect.

One of the responses to remediating robotic failure is -- rather than invest inordinate amounts of time and energy in the assurance of individual robots -- to manufacture them in a more cost-effective way -- typically with less assurance on the device -- but deploy more of them so that collectively, the probability of mission success will increase notwithstanding the relatively high probability of individual failure. But as physical entities, robots need to share resources: space, space over the course of time, communications bandwidth, and human and non-human maintenance resources. The effective sharing of resources requires coordination; otherwise a group of robots will not be able to accomplish anything.

Coordination algorithms can be grouped into three categories: biologically-inspired "swarm" coordination, such as ant foraging and bird flocking; capability-based coordination, in which coordination among peer robots emerges ad hoc and driven by the task structure of individual plans; and team-oriented coordination, in which the robots need to dedicate communications bandwidth to maintaining their cohesion as a team. Proofs of the effectiveness of coordination algorithms are not always satisfying from an assurance perspective, however. Often they consist of abstract and limiting assumptions about the "interface layers" of the systems with which they must interoperate, the operating environment, and of the quality attributes of the task that facilitate task execution by the team. Not surprisingly, roboticists are approaching the software engineering community with requests for how to design robotic systems for quality attributes such as debuggability and predictability of individual and collective coordination behaviors.

My work in progress is to propose an approach by which robotic coordination can be assured. It proposes structure by which the operating context can be appropriately segmented so as to provide a stable environment by which coordinated robotic behaviors can be studied and predicted. It also examines the interplay of communication and team coordination, which represent one of the interfaces of the physical world with the information theoretic behaviors that are generated by the coordination of plan structures. Through this approach, it will be possible to develop a body of knowledge by which it can be known what types of assurance arguments need to be made, what evidence should be collected to support or refute such arguments, and ways of qualifying and quantifying the uncertainties associated with assurance judgments. A corollary of such results will be knowledge of what does not need assurance, as well.