

Issues on building software for hardware agents

HENRY HEXMOOR

226 Bell Hall, Department of Computer Science, SUNY at Buffalo, Buffalo, NY 14260, USA
(email:hexmore@cs.buffalo.edu)

DAVID KORTENKAMP

Metrica, Inc., Robotics and Automation Group, NASA Johnson Space Center—ER4, Houston, TX 77085, USA (email:korton@mickey.jsc.nasa.gov)

1 Introduction

James Albus states that “an architecture is a description of how a system is constructed from basic components and how those components fit together to form the whole” (Albus, 1995). A software architecture for physical agents reflects the organising principles that its designers have learned from many prior experiences in building such agents. Architectures that have been proposed for physical agents have differed greatly—from subsumption (Brooks, 1986) to Soar (Laird et al., 1987). However, a surprising consensus about architectures is beginning to emerge within the small community of researchers applying artificial intelligence to robotics. The consensus is that a multi-layer, hierarchical architecture is necessary. In particular, the community is moving towards a three-layered architecture. The lowest layer is a reactive control system inspired by subsumption (Brooks, 1986). The top layer is a traditional symbolic planning and modelling system. The middle layer is the key; it serves as a “differential” between the short-range reaction and long-range reasoning.

2 Layered architectures

Layered architectures are not a new idea. Cognitive psychologists have long noticed different time scales in human behaviour. Newell (1990) describes a *biological band*, a *cognitive band* and a *rational band* within the time-scale of human action. The biological band operates on the order of 1–10 ms, the cognitive band operates on the order of 100 ms–10 s, and the rational band operates on the order of minutes to hours. Newell also added a *social band*, which operated on the order of days to months.

The robotics community arrived at a similar functional decomposition via a different path. Early mobile robot work, such as Shakey (Nilsson, 1969), created symbolic plans that would take the robot from its current world state to a desired world state. STRIPS (Fikes & Nilsson, 1971) grew out of such work. However, researchers soon found that these symbolic reasoning systems could not react quickly enough to a dynamic world. In essence, the pendulum swung the other way and reactive, situated robots proved their effectiveness (Agre & Chapman, 1987; Brooks, 1986). However, such robots could not easily demonstrate the ability to guide their reactions in order to achieve multiple, long-range goals. For some years there was a gulf between reasoning and situated reaction. Since the late 1980s, though, there has been a series of efforts to integrate deliberation and reaction (Bonasso et al., 1992; Connell, 1992; Firby, 1987; Gat, 1992; Sanborn et al., 1989; Schoppers, 1987).

While a consensus is emerging that deliberation and reactivity need to be integrated into a single robot control architecture, a recent AAAI Spring Symposium pointed to several open research issues that hinder progress. These issues are discussed in the next few sections.

3 Bridging deliberation and reactivity

The middle layer of the three-layered architecture exists to serve as a mediator between deliberation and reactivity. The most promising system for such a role is the Reactive Action Packages (RAPS) of Jim Firby (1987, 1992, 1994). RAPS takes a set of high-level tasks generated by a planner and recursively decomposes them into a set of reactive skills that can be run on the robot. RAPS then activates and deactivates these sets of skills to accomplish tasks. RAPS is also responsible for monitoring the execution of the skills to see if they are moving the robot towards the goal. If the task is not being achieved, RAPS can choose another method for achieving the same task or take other corrective action. RAPS has been used as the middle layer in several three-layer architectures (Bonasso, 1991; Elsaessar & Slack, 1994; Firby, 1994; Gat, 1992).

Much progress has been made in defining the middle layer in the three-layer architecture, especially with regards to how to take high-level task descriptions, decompose them into primitive robot operations and then monitor those operations. However, there still remain many open issues. For example, there is no principled methodology for how the middle layer can abstract indexical perceptual information from the bottom layer into symbolic representations for the top layer. The bottom layer requires representations that are direct and fast, while the top layer requires representations that are rich and computationally slow. The middle layer will need to provide an intermediary representation that can serve both purposes. Another open issue is how to keep reaction and deliberation in synch when they operate at very different time scales. In theory, the middle layer serves as a “differential” between short-term reactions and long term reasoning, but the mechanics underlying this differential are still unknown.

4 Learning within a layered architecture

As Herbert Simon puts it, learning is changing an agent’s internal structure so that the agent can improve its long-term performance. A superior architecture organises internal structures of an agent in a way that can facilitate learning. In a three-layered architecture, learning within the lowest level improves the agent’s reflexive actions; learning in the highest level improves cognitive skills; and learning in the middle layer improves the coherence between the other levels.

We contrast this with learning that takes place in between layers. By migrating planful behaviour from the highest layer to reactive behaviour in the lowest layer, a task that is highly cognitive in the beginning (i.e. cognitive expertise) can become reactive and skilled in the future (Hexmoor et al., 1993). Also, through learning, repeated patterns of interactions (i.e. motor skills) at the low level can be redefined using cognitive concepts at a later time. The mechanisms for achieving this migration of skills between layers are still unknown. However, the middle layer of the architecture will certainly play a large role in this process.

5 Real-time issues

As architectures begin to be applied to domains requiring real-time responses, they will be subject to the real-time constraints of the environment. Most architectures do not guarantee responses in hard real-time domains (i.e. those domains that have deadlines by which responses must be made or failures will probably occur), but simply attempt to execute their primitive control rules as quickly as possible and hope that they are fast enough. There are some exceptions, however. For example, CIRCA (Musliner et al., 1995) can reason about real-time behaviours and construct a real-time plan with guaranteed execution times. Another approach is to carefully structure the

underlying control system using circuit semantics that can be used for reasoning and planning (Lee & Durfee, 1994; Nilsson, 1994; Rosenschein & Kaelbling, 1988).

The issue of hard real-time guarantees within an intelligent architecture is far from solved. In fact, a new subfield of artificial intelligence is being defined by researchers in artificial intelligence and real-time control. This subfield is surveyed in an excellent article by Musliner et al. (1993).

6 Representation

Architectures can be useful means of exploring the power and tractability of representational mechanisms. For instance, one can study the mechanisms of interfacing formal knowledge representation and reasoning systems (e.g. Laird et al., 1987) to the world. KRR systems are typically used to model deliberative parts of intelligent agency, which tend to be more abstract and rather slow in interaction with the world. On the other hand, one can explore the architectural mechanisms for interpreting the real-time data from the world, i.e. perceptual reduction.

It has been argued that multiple representational mechanisms are useful in modelling a single agent (Pavio, 1986). If we allow different representational mechanisms to be used at different levels of an architecture, an open question is how to interface representations at different levels. Another issue is that for know-how that lies partly at two levels, how will the agent coordinate use of such know-how.

7 Evaluation of architectures

Evaluating and comparing the performance of architectures is a difficult (some would argue impossible) task. Erann Gat (1995) proposes a scientific methodology using simulated robots and environments as theoretical models to make predictions. These predictions then need to be verified or refuted by experimentation with real robots in real environments. Unfortunately, robotic experiments and environments are not easily reproducible. We believe that incorporation of evaluation and simulation in early stages of robotic experiments will shape the experiments and robot design.

8 Conclusion

With an emerging consensus on overall architectural issues, there is a potential for gaining a new perspective on problems in knowledge representation, learning, real-time performance and evaluation. There are, of course, researchers working to achieve the integration of deliberation and reaction without a layered architecture; Soar is a notable example (Laird et al., 1987). Moreover, even among those who agree with a layered approach, there is still not a consensus on the kinds of representations or learning algorithms that should be used by the different layers. The theoretical idea of coordinating planning and reacting is beginning to be implemented, but the implementation details are still under investigation by many researchers.

Acknowledgements

This paper is the result of presentations and discussions at the 1995 AAAI Spring Symposium on *Lessons Learned from Implemented Software Architectures for Physical Agents*. The authors were co-chairs of the symposium and wish to acknowledge the contributions of the other members of the program committee: Ron Arkin, R. Peter Bonasso and David Musliner. Information about the symposium and related discussions can be obtained at:

<http://tommy.jsc.nasa.gov/er/er6/mrl/symposium.html>.

References

- Agre, PE and Chapman, D, 1987. "Pengi—an implementation of a theory of activity" In: *Proceedings of the Fifth National Conference on Artificial Intelligence*.
- Albus, JS, 1995. "A reference model architecture for intelligent systems" In: *Working Notes: AAAI 1995 Spring Symposium on Lessons Learned from Implemented Software Architectures for Physical Agents*.
- Bonasso, RP, 1991. "Integrating reaction plans and layered competences through synchronous control" In: *Proceedings International Joint Conferences on Artificial Intelligence*.
- Bonasso, RP, Antonisse, HJ and Slack, MG, 1992. "A reactive robot system for find and fetch tasks in an outdoor environment" In: *Proceedings of the Tenth National Conference on Artificial Intelligence*.
- Brooks, RA, 1986. "A robust layered control system for a mobile robot" *IEEE Journal of Robotics and Automation* 2 (1).
- Connell, JH, 1992. "A hybrid architecture applied to robot navigation" In: *Proceedings IEEE International Conference on Robotics and Automation*.
- Elsaesser, C and Slack, MG, 1994. "Integrating deliberative planning in a robot architecture" In: *Proceedings of the AIAA/NASA Conference on Intelligent Robots in Field, Factory, Service, and Space (CIRFFSS '94)*.
- Fikes, RE and Nilsson, NJ, 1971. "Strips: A new approach to the application of theorem proving to problem solving" *Artificial Intelligence* 2.
- Firby, RJ, 1987. "An investigation into reactive planning in complex domains" In: *Proceedings of the National Conference on Artificial Intelligence (AAAI)*.
- Firby, RJ, 1992. "Building symbolic primitives with continuous control routines" In: *Proceedings International Conference on AI Planning Systems*.
- Firby, RJ, 1994. "Task networks for controlling continuous processes" In: *Proceedings of the Second International Conference on AI Planning Systems*.
- Gat, E, 1992. "Integrating planning and reacting in a heterogeneous asynchronous architecture for controlling real-world mobile robots" In: *Proceedings of the National Conference on Artificial Intelligence (AAAI)*.
- Gat, E, 1995. "On the role of simulation in the study of autonomous mobile robots" In: *Working Notes: AAAI 1995 Spring Symposium on Lessons Learned from Implemented Software Architectures for Physical Agents*.
- Hexmoor, H, Lammens, J and Shapiro, SC, 1993. "Behavior based AI, cognitive processes, and emergent behaviors in autonomous agents" In: G Rzevski, J Pastor and R Adey (eds), *Applications of AI in Engineering VIII, Vol. 2, Applications and Techniques*. CMI/Elsevier.
- Laird, JE, Newell, A and Rosenbloom, PS, 1987. "Soar: An architecture for general intelligence" *Artificial Intelligence* 33 (1).
- Lee, J and Durfee, EH, 1994. "Structured circuit semantics for reactive plan execution systems" In: *Proceedings of the Twelfth National Conference on Artificial Intelligence*, 1232–1237.
- Musliner, DJ, Hendler, JA, Agrawala, AK, Durfee, EH, Strosnider, JK and Paul, CJ, 1995. "The challenges of real-time AI" *IEEE Computer* 28 (1).
- Musliner, DJ, Durfee, E and Skin, K, 1993. "CIRCA: A cooperative, intelligent, real-time control architecture" *IEEE Transactions on Systems, Man and Cybernetics* 23 (6).
- Newell, A, 1990. *Unified Theories of cognition*. Harvard University Press.
- Nilson, NJ, 1969. "A mobile automation: An application of AI techniques" In: *Proceedings of the International Joint Conference on Artificial Intelligence*.
- Nilsson, NJ, 1994. "Teleo-reactive programs for agent control" *Journal of Artificial Intelligence Research* 1 (1).
- Pavio, A, 1986. *Mental Representations: A dual coding approach*, Oxford University Press.
- Rosenschein, SJ and Kaelbling, LP, 1988. "The synthesis of digital machines with provable epistemic properties" In: *Proceedings of the Conference on Theoretical Aspects of Reasoning About Knowledge*. 83–98.
- Sanborn, J, Bloom, B and McGrath, D, 1989. "A situated reasoning architecture for space-based repair and replace tasks" In: *Goddard Conference on Space Applications of Artificial Intelligence (NASA Publication 3033)*.
- Schoppers, M, 1987. "Universal plans for reactive robots in unpredictable environments" In: *Proceedings of the International Joint Conferences on Artificial Intelligence (IJCAI)*.
- Simon, H, 1983. "Why should machines learn?" In: J Carbonell, R Michalski and T Mitchell (eds.), *Machine Learning: An Artificial Intelligence Approach*, Tioga Press.