

Toward Decision-Theoretic CIRCA with Application to Real-Time Computer Security Control

Vu Ha David J. Musliner

Automated Reasoning Group
Honeywell Laboratories
3660 Technology Drive
Minneapolis, MN 55418
{vha,musliner}@htc.honeywell.com

Abstract

We report our on-going work toward extending the CIRCA (Cooperative Intelligent Real-Time Control Architecture) with decision-theoretic reasoning capabilities. By explicitly modeling uncertainty using probabilities, and goals using utilities, the new CIRCA planner can now appeal to the powerful decision-theoretic paradigm of maximizing expected utility to find the best plan. We discuss representational and computational issues encountered, as well as a preliminary application in the domain of planning for computer security.

The Original CIRCA World Model

CIRCA is an architecture for real-time intelligent control. The original CIRCA world model is introduced by Musliner, Durfee, & Shin in (1993; 1995). The CIRCA planner searches for a plan that is guaranteed to be safe, by making sure that no failure state is reachable from initial states. Traditionally, CIRCA does not reason explicitly about probabilities, but uses simpler forms of uncertainty (nondeterminism and time bounds). Younes and Musliner (2002) extended the CIRCA world model by associating a probability distribution function $F(t; \tau)$ with each transition τ , giving the probability that τ will be triggered t time units after it was last enabled. With probability distribution functions available for the transitions, we can set an arbitrary threshold θ representing the highest acceptable failure probability of a plan. The problem of plan verification now becomes a hypothesis testing problem, which can be solved using the sequential testing algorithm pioneered by Wald (1945). The samples used by the algorithm are sample execution paths generated through discrete event simulation. The advantage of the sequential approach is that on average it requires fewer samples to reach a decision, compared to non-sequential approaches.

Introducing Utilities

Up to now, the current CIRCA goal language cannot express the relative benefits of achieving or maintaining different goals. In the information security domain that we are interested in, these limitations prevent the planner from constructing defense plans that can flexibly adapt to the

uncertain nature of security threats, and the changing demands of trading off information services against security level. Decision theory, which models uncertainty with probabilities and the costs/benefits of actions with utilities, provides an attractive answer to this challenge. We construct a decision-theoretic model for the CIRCA planner by defining a goal-directed utility model (Haddawy & Hanks 1998) that has the following characteristics: 1) the utility function assigns a real number to a *finite-horizon plan execution path*, where the time horizon h is domain specific and 2) the utility function is a weighted sum of sub-utility functions, each of which is scaled to have range $[0, 1]$, and belongs to one of three categories: maintenance-goal (MG), achievement goal (AG), and repeated achievement goal (RAG): $u = \sum_i w_i u_i^{MG} + \sum_j w_j u_j^{AG} + \sum_k w_k u_k^{RAG}$, where w_i, w_j, w_k are the weights of the corresponding sub-utility functions, which are scaled to sum to 1. As an example, a CIRCA planner for running a web server may have a maintenance goal “maintain high data throughput for as long as possible”, an achievement goal “complete the Perl interpreter upgrade” and a repeated achievement goal “perform crucial data backup every night”. The utility function may be defined as: $u = .4u_1^{MG} + .1u_2^{AG} + .5u_3^{RAG}$. In this equation, u_1^{MG} may be defined proportional to the average data throughput, u_2^{AG} can simply be a binary function, and u_3^{RAG} may be defined as the number of successful backups performed before the time limit. Finally, the weights .4, .1, and .5 reflect the relative importance of the three goals in the overall utility function.

Identifying Plans with Highest Expected Utilities

With the introduction of the utility model, the CIRCA planning problem now becomes the problem of searching for the plan with highest expected utility. The key issue here is to compute the EUs. Since our CIRCA model corresponds to a GSMP¹, for which there are no known analytic methods to efficiently compute the expected utility, the only feasible approach is to use Monte Carlo sampling to approximate the EU.

¹Generalized Semi-Markov Process (GSMP) is a formalism for discrete event systems introduced by Matthes (1962).

Monte Carlo Estimation of EU

Note that by definition, the utility function has range $[0, 1]$, and as a consequence, the utility of a plan is a random variable with range $[0, 1]$. Ideally, we would like our estimate \tilde{u} of $E[u]$ to be within ϵ of the actual mean with probability of at least $1 - \delta$, where ϵ and δ are small positive real numbers: $\Pr(|\tilde{u} - E[u]| > \epsilon) < \delta$. If σ^2 denotes the variance of u (which is at most $1/4$ because u is bounded in $[0, 1]$), and \bar{u}_k denotes the mean of a k -size sample of u , we have proved (by means of Chebyshev’s inequality) that this accuracy can be achieved with a sample of size $k > 1/(4\epsilon^2\delta)$. The following algorithm finds the plan with highest expected utility:

1. Set $current_best_EU = -\infty$.
2. Generate a plan A . Simulate the execution of A up to time horizon h , compute the utility of the resulting path.
3. Repeat the above step for $k = \lceil 1/(4\epsilon^2\delta) \rceil$ times, compute the average utility \bar{u}_k . If $\bar{u}_k > current_best_EU$, set $current_best_EU = \bar{u}_k$.
4. Go back to step 2. Continue until some stopping criterion is true (e.g. there are no more plans, or plan time limit is reached).

Sequential Methods for Identifying Best Plans

The above sample upper bound of $\lceil 1/(4\epsilon^2\delta) \rceil$ is applicable for *any* random variable with range $[0, 1]$, which is important for our analysis because u is a complex function and will most likely not observe known parametric forms such as uniform or Gaussian. The downside of this generality is that this upper bound is rather high: for .01 error margin and 95% confidence ($\delta = .05$), we need to have a sample of size 200,000. One possible way to cut down on the number of samples is to appeal to sequential analytic techniques (see e.g. (Younes & Musliner 2002)). Unfortunately, there is no known effective sequential method to estimate the mean of a random variable u unless parametric assumptions are made about u . Our solution is based on the observation that in Step 4 in Algorithm 1, what we really are interested in this step is whether the current plan π is inferior to the current best plan, i.e. $E[u(\pi, h)] < current_best_EU$, which is clearly a hypothesis testing problem. If we can quickly determine, via a sequential acceptance sampling procedure that π is inferior to the current best plan, there is no need to continue estimating the EU of π ; the algorithm can move on to the next plan. One scenario where this technique does not work well is when successive plans are of increasingly better quality. Nevertheless, in our experiments so far, the best plan is identified after only examining a few plans, due to the strength of our heuristic algorithm, and thus from that point on we are able to save a significant amount of time on estimating utilities. Figure 1 illustrates these savings.

In order to add the capability to reason with utility, we need to make several modifications to the existing CIRCA planner. The most significant change is the elimination of preemptive actions, which are meant to make failure states unreachable. In the decision-theoretic model, failure states




Figure 1: The number of simulations for 432 different plans for a computer security domain, with $\epsilon = .025$ and $\delta = .05$. The graph clearly shows that for a large percentage of plans, the required number of simulation runs is less than 100, significantly less than the maximum number of runs, which is 8000. Note that here we simultaneously check if the probability of plan failure exceeds a threshold

are tolerated, as long as we don’t encounter them very often. One consequence of this is that a large number of failure-bound plans now make it to the EU simulation phase. This issue raises the importance of simulation-guided backjumping. Currently, plans are generated via a depth-first exploration of the plan space. The problem with this procedure is that if a mistake is made early in the planning process, that mistake can be corrected only after all possible onward decision combinations have been explored. The obvious way to overcome this problem is to look at the simulation traces that lead a plan to failure, identify the *earliest decision* responsible for that failure, and backjump to correct that decision.

References

- Haddawy, P., and Hanks, S. 1998. Utility models for goal-directed, decision-theoretic planners. *Computational Intelligence* 14(3).
- Matthes, K. 1962. Zur Theorie der Bedienungsprozesse. In Kožešník, J., ed., *Transactions of the Third Prague Conference on Information Theory, Statistical Decision Functions, Random Processes*, 513–528. Liblice, Czechoslovakia: Publishing House of the Czechoslovak Academy of Sciences.
- Musliner, D. J.; Durfee, E. H.; and Shin, K. G. 1993. CIRCA: A cooperative intelligent real-time control architecture. *IEEE Transactions on Systems, Man, and Cybernetics* 23(6):1561–1574.
- Musliner, D. J.; Durfee, E. H.; and Shin, K. G. 1995. World modeling for the dynamic construction of real-time control plans. *Artificial Intelligence* 74(1):83–127.
- Wald, A. 1945. Sequential tests of statistical hypotheses. *Annals of Mathematical Statistics* 16(2):117–186.
- Younes, H. L. S., and Musliner, D. J. 2002. Probabilistic plan verification through acceptance sampling. In *Proceedings of the AIPS 2002 Workshop on Planning via Model Checking*. Toulouse, France: AAAI Press.