

# Integrated Plan Tracking and Prognosis for Autonomous Production Processes

Paul Maier, Martin Sachenbacher, Thomas Rühr  
Technische Universität München  
Boltzmannstraße 3, 85748 Garching, Germany  
{maierpa, sachenba, ruehr}@in.tum.de

Lukas Kuhn  
PARC  
Palo Alto, USA  
Lukas.Kuhn@parc.com

## Abstract

Today’s complex production systems allow to simultaneously build different products following individual production plans. Such plans may fail due to component faults or unforeseen behavior, resulting in flawed products. In this paper, we propose a method to integrate diagnosis with plan assessment to prevent plan failure, and to gain diagnostic information when needed. In our setting, plans are generated from a planner before being executed on the system. If the underlying system drifts due to component faults or unforeseen behavior, plans that are ready for execution or already being executed are uncertain to succeed or fail. Therefore, our approach tracks plan execution using probabilistic hierarchical constraint automata (PHCA) models of the system. This allows to explain past system behavior, such as observed discrepancies, while at the same time it can be used to predict a plan’s remaining chance of success or failure. We propose a formulation of this combined diagnosis/assessment problem as a constraint optimization problem, and present a fast solution algorithm that estimates success or failure probabilities by considering only a limited number  $k$  of system trajectories.

## 1 Introduction

As the market demands for customized and variant-rich products, the industry struggles to implement production systems that demonstrate the necessary flexibility while maintaining cost efficiency comparable to highly automated mass production. A main cost driver in automated production is the human workforce needed for setup steps, the development of processes, and quality assurance. These high labor costs can typically only be amortized by very large lot sizes. For small lot sizes as found in prototype and highly customized production, human workers are still unchallenged in flexibility and cost by automated systems. Therefore, to facilitate the emergence of mass customization, levels of flexibility similar to the flexibility of human workers must be reached at prices only highly automated systems can achieve.

The German research cluster “Cognition for Technical

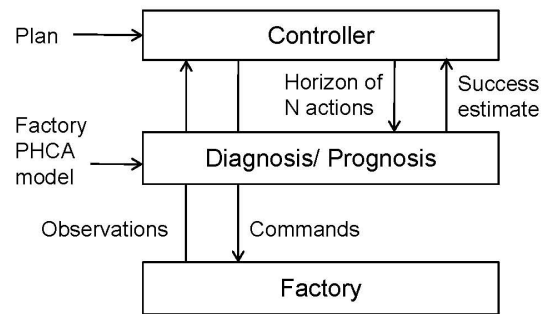



Figure 1. Model-based plan assessment.

Systems” (CoTeSys) [2] was founded to understand human cognition and make its performance accessible for technical systems. Future technical systems are expected to act robustly under high uncertainty, reliably handle unexpected events, quickly adapt to changing tasks and own capabilities. A key technology for the realization of such systems is automated planning combined with self-diagnosis and self-assessment. These capabilities can allow the system to plan its own actions, and also react to failures and adapt the behavior to changing circumstances.

From the point of view of planning, production systems are a relatively rigid environment, where the necessary steps to manufacture a product can be anticipated well ahead. However, from a diagnosis point of view, production systems are typically equipped with only few sensors, so it cannot be reliably observed whether an individual manufacturing step went indeed as planned; instead, this becomes only gradually more certain during execution of the production plan. Therefore, in the presence of faults or other unforeseen events – which become more likely in individualized production – the question arises whether plans that are ready for execution or already being executed will indeed succeed, and whether it is necessary to revise a plan or even switch to another plan.

To address this problem, we propose in this paper a model-based capability that estimates the success probability of production plans in execution (figure 1). We assume that a planner provides plans given a system model. A plan is a sequence of actions where each action is executed at its corresponding start time. Whenever the sys-



**Figure 2. Effects of cutter deterioration until breakage in machining (Image (c) Prof. Shea TUM PE).**

tem produces an observation, it is forwarded to a module that performs simultaneous plan tracking and plan prognostic using probabilistic hierarchical constraint automata (PHCA) models [11] of the system. We propose a formulation of this problem as a soft constraint optimization problem [10] over a window of  $N$  time steps that extends both into the past and the future, and present a fast but approximate solution method that enumerates only  $k$  most likely system trajectories. The resulting success or failure prognosis can then be used to autonomously react in different ways depending on the probability estimate; for instance, continue with plan execution, discard the plan, or augment the plan by adding observation-gathering actions to gain further information [5].

In the remainder of the paper, we first motivate the approach informally with an example from an automated metal machining process, and then present our algorithmic solution and experimental results.

## 2 Metal Machining and Assembly Example

As part of the CoTeSys cognitive factory test-bed, we set up a customized and extended Flexible Manufacturing System (FMS) based on the iCim3000 from Festo AG (see figure 5). The system consists of a conveyor transport and three stations: storage, machining (milling and turning), and assembly. We built a simplified model of this manufacturing system (see figure 4) which consists only of the machining and the assembly station and allows to track system behavior over time, including unlikely component faults. In particular, the machining station can transition to a “cutter blunt” composite location, where abrasions are caused during operation due to a blunt cutter. This makes it very probable that the cutter breaks, leading to flawed products (see figure 2). The assembly station model contains a composite location which models occasional abrasions. A vibration sensor at the assembly station can detect these abrasions, yielding binary signals “abrasion occurred” and “no abrasion occurred”. However, the signal is ambiguous, since the sensor cannot differentiate between the two possible causes.

Two products are produced using a single production



**Figure 3. The robotic arm product (Image (c) Prof. Shea TUM PE).**

plan  $\mathcal{P}_{prod}$ : a toy maze consisting of an alloy base plate and an acrylic glass cover, and an alloy part of a robotic arm (see figure 3).  $\mathcal{P}_{prod}$  consists of these steps: (1) cut maze into base plate (one time step), (2) assemble base plate and cover (one time step), (3,4,5,6) cut robot arm part (one to four time steps). The plan takes two to six time steps (starting at  $t = 0$ ). The plan is considered successful if both products are flawless. In our example, only a broken cutter causes the machined product to be flawed, in all other cases the production plan will succeed. Now consider the following scenario: after the second plan step (assembling the maze base plate and its cover at  $t = 2$ ) an abrasion is observed. Due to sensor ambiguity it remains unclear whether the plan is unaffected (abrasion within assembly) or whether it might fail in the future due to a broken cutter (abrasion caused by a blunt cutter), and the question for the planner is: How likely is it that the current plan will still succeed? Our new capability allows to compute this likelihood, taking into account past observations and future plan steps.

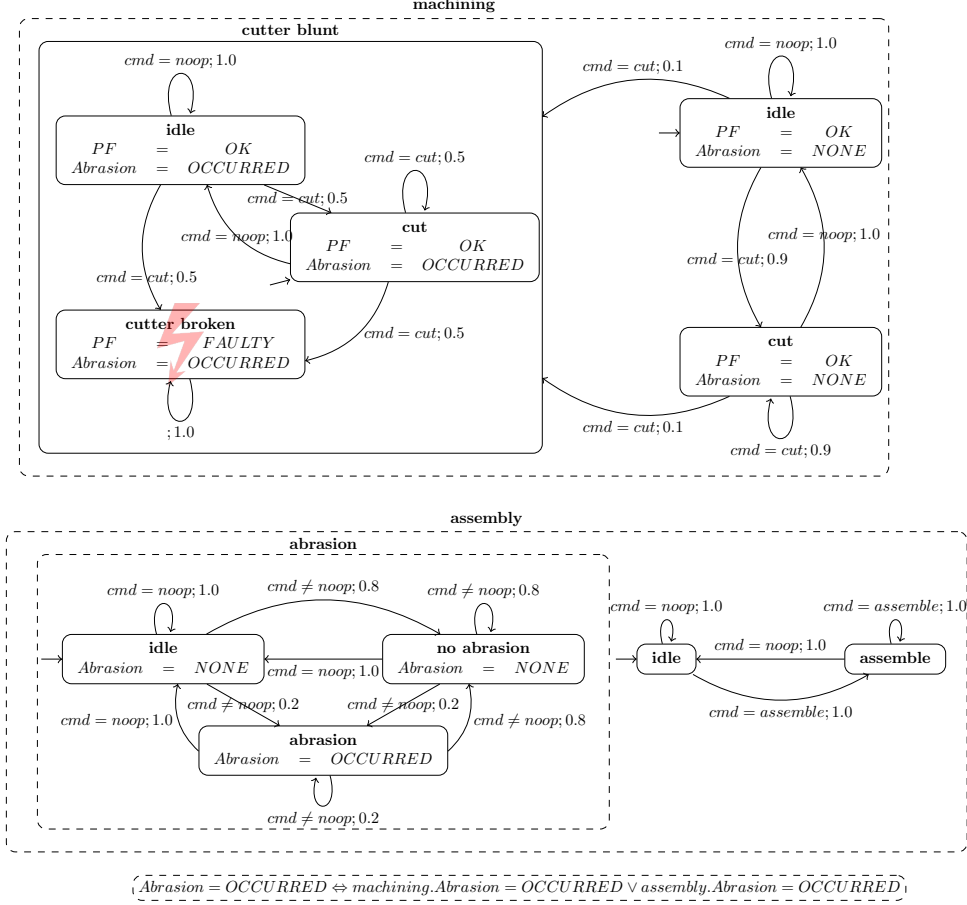
## 3 Modeling System Behavior with PHCA

Probabilistic hierarchical constraint automata (PHCA) were introduced in [11] as a compact encoding of hidden markov models (HMMs). These automata have the required expressivity to uniformly model both probabilistic hardware behavior (e.g., likelihood of component failures) and complex software behavior (such as high level control programs).

### Definition 1 (PHCA)

A PHCA is a tuple  $\langle \Sigma, P_{\Sigma}, \Pi, O, Cmd, \mathcal{C}, P_T \rangle$ , where:

- $\Sigma$  is a set of locations, partitioned into primitive locations  $\Sigma_p$  and composite locations  $\Sigma_c$ . Each com-



**Figure 4. Simplified PHCA of the manufacturing system. The machining and assembly station are modeled as parallel running composite locations (indicated by dashed borders). Variables appearing within a location are local to this location, i.e.  $machining.cmd$  refers globally to the command variable  $cmd$  within composite location  $machining$ . Note: “noop” stands for “no operation”.**

posite location denotes a hierarchical, constraint automaton. A location may be marked or unmarked. A marked location represents an active execution branch.

- $P_{\Xi}(\Xi_i)$  denotes the probability that  $\Xi_i \subseteq \Sigma$  is the set of start locations (initial state). Each composite location  $l_i \in \Sigma_c$  may have a set of start locations that are marked when  $l_i$  is marked.
- $\Pi$  is a set of variables with finite domains.  $\mathcal{C}[\Pi]$  is the set of all finite domain constraints over  $\Pi$ .
- $O \subseteq \Pi$  is the set of observable variables.
- $Cmd \subseteq \Pi$  is the set of command variables.
- $\mathcal{C} : \Sigma \rightarrow \mathcal{C}[\Pi]$  associates with each location  $l_i \in \Sigma$  a finite domain constraint  $\mathcal{C}(l_i)$ .
- $P_T(l_i)$ , for each  $l_i \in \Sigma_p$ , is a probability distribution over a set of transition functions  $T(l_i) : \Sigma_p^{(t)} \times$

$\mathcal{C}[\Pi]^{(t)} \rightarrow 2^{\Sigma^{(t+1)}}$ . Each transition function maps a marked location into a set of locations to be marked at the next time step, provided that the transition’s guard constraint is entailed.

**Definition 2 (PHCA State)** The state of a PHCA at time  $t$  is a set of marked locations called a marking  $m^{(t)} \subset \Sigma$ .

The example PHCA shown in figure 4 illustrates the PHCA definition. The main factory components  $machining$  and  $assembly$  are encoded as top level composite locations. A dashed border indicates that locations may be marked at the same time, which means they can run in parallel. There is a third top level location at the bottom of figure 4 whose behavior constraint encodes that an observed abrasion is caused by one of the two components or both. Primitive locations are for example  $machining.idle$  and  $machining.cut$ , which encode the machining station being in an idle state and working on a piece. An example for an observable variable is  $Abrasion$ , which encodes whether an abrasion has occurred or not. The dependent variables



**Figure 5. The hardware setup for experimentation, showing storage, transport, robot and machining components.**

*machining.Abrasion* and *assembly.Abrasion* encode for each component whether it caused an abrasion. A command variable is, e.g., *machining.cmd*. It occurs in the guard constraint for transition *idle*  $\rightarrow$  *cut* within composite location *machining*: *machining.cmd* = *cut*. Transition guards have the general form  $\langle \text{guard constraint} \rangle; \langle \text{transition probability} \rangle$ . The guard constraint is a logical constraint over PHCA variables, usually an assignment to command variables. The transition is non-deterministic: Given the guard is satisfied, it is taken with probability 0.9. The remaining possibility (completing the conditional probability distribution) is the transition from *idle* to the composite location *cutter blunt*, which has the same guard and is taken with probability 0.1.

#### 4 Plan Assessment as Constraint Optimization over PHCA Models

Plan assessment requires tracking of the system’s plan-induced evolution; in our case, it means tracking the evolution of PHCA markings. In previous work [9], we introduced an encoding of PHCA as soft constraints and casted the problem of tracking markings as a soft constraint optimization problem (COP) [10], whose solutions correspond to the most likely sequences of markings given the available observations. The encoding is parameterized by  $N$ , which is the number of time steps considered (for a detailed description of the encoding, see [9]).

Observations made online are encoded as hard constraints specifying assignments to observable variables at time  $t$  (e.g., *Abrasion*<sup>(2)</sup> = *OCCURRED*), and added to the constraint optimization problem.

Plans are added analogously as assignments to commandable variables at time  $t$ ; for example,  $a_{cut}^{(t)}$  and  $a_{assemble}^{(t)}$  are assignments *machining.cmd*<sup>( $t$ )</sup> = *cut*  $\wedge$  *assembly.cmd*<sup>( $t$ )</sup> = *noop* and *assembly.cmd*<sup>( $t$ )</sup> = *assemble*  $\wedge$  *machining.cmd*<sup>( $t$ )</sup> = *noop*. Note that vari-

ables appear time independent in the PHCA notation (see, e.g., figure 4), i.e. without superscript <sup>( $t$ )</sup>.

The plan’s goal  $G$  is to produce a flawless product. We encode this informal description as a logical constraint  $G \equiv \forall PF^{(t_{end})} \in RelevantFeatures(\mathcal{P}) : PF^{(t_{end})} = OK$  over product feature variables  $PF^{(t)} \in \{OK, FAULTY\}$  at the end of the execution,  $t_{end}$ . *RelevantFeatures*() is a function mapping a production plan to all product feature variables which define the product. Each system component is responsible for a product feature in the sense that if it fails, the product feature is not present ( $PF^{(t)} = FAULTY$ ). In our example, there is only a single product feature  $PF$ , which is absent if the cutter is broken. The goal constraint for the above mentioned plan (three time steps long) is accordingly  $PF^{(3)} = OK$ .

#### 4.1 Solving Constraints to Enumerate Most Likely System Trajectories

The soft constraint encoding of PHCA model, plan and observations form a COP that captures the probabilistic behavior of the system over a horizon of  $N$  time steps. The model encoding can be done offline, while the plan and the observations have to be encoded and added to the COP online. The effect of adding the plan and observations constraints is that they render certain PHCA trajectories impossible (zero probability). For example, the observation of an abrasion renders impossible the trajectory which doesn’t entail an observed abrasion. The goal constraint, however, is *not* added to the COP, since adding this constraint would render all non-goal-achieving trajectories impossible (we need these failure trajectories for normalization in computing the plan’s success probability, as shown in the next section).

For a given plan  $\mathcal{P}$  and available observations, we then enumerate the  $k$  best solutions to the COP. These correspond to the system’s most likely execution trajectories, or diagnoses, within the  $N$ -step horizon. An execution trajectory is a sequence of markings for each time step, encoded as assignment to location variables. These are the variables of interest for our COP. For example, Table 1 shows the most likely execution trajectory of the example PHCA, given production plan  $\mathcal{P}_{prod} = (a_{cut}, a_{assemble}, a_{cut})$  and observation *Abrasion*<sup>(2)</sup> = *OCCURRED*.

Technically, the  $k$ -best enumeration is done by translating the generated COP (as part of the compilation step) into the weighted CSP format as used by the soft constraint solver toolbar [3]. In the online step, we used a modified version of toolbar that implements mini-bucket elimination to generate a search heuristic for the problem. The heuristic is used by a subsequent A\* search to enumerate the  $k$ -best solutions. This approach is described in more detail in [4].

time	marking
0	$assembly.abrasion.idle_L^{(0)}, assembly.idle_L^{(0)},$ $machining.idle_L^{(0)}$
1	$assembly.abrasion.idle_L^{(1)}, assembly.idle_L^{(1)},$ $machining.cut_L^{(1)}$
2	$assembly.abrasion.abrasion_L^{(2)},$ $assembly.assemble_L^{(2)}, machining.idle_L^{(2)}$
3	$assembly.abrasion.idle_L^{(3)}, assembly.idle_L^{(3)},$ $machining.cut_L^{(3)}$

**Table 1. Most probable PHCA trajectory for production plan  $\mathcal{P}_{prod} = (a_{cut}, a_{assemble}, a_{cut})$ , given an abrasion occurred at  $t = 2$ . A shown variable  $X_L^{(t)}$  indicates a marking of location  $L$  at time  $t$ .**

## 4.2 Combining Plan Tracking and Prognosis

In the previous section, we described a method to track plan execution within an  $N$ -step time window based on a system model and observations. To assess a plan’s probability of success, we require not only to track past system behavior, but also to predict its evolution in the future. In principle, this could be accomplished in two separate steps: first, assess the system’s state given the past behavior, and then predict its future behavior given this belief state and the plan. However, this two-step approach leads to a problem. Computing a belief state (complete set of diagnoses) is intractable, thus it must be replaced by some approximation (such as considering only  $k$  most likely diagnoses [6]). But if a plan uses a certain component intensely, then this component’s failure probability is relevant for assessing this plan, even if it is very low and therefore would not appear in the approximation. In other words, the plan to be assessed determines which parts of the belief state (diagnoses) are relevant.

To address this mutual dependency, we propose a method that performs diagnosis and plan assessment *simultaneously*, by framing it as a single optimization problem. The key idea is as follows: The optimization problem formulation is independent of where the present time point is within the  $N$ -step time window. We therefore choose it such that the time window covers the remaining future plan actions as well as the past behavior. Now solutions to the COP are system trajectories which start in the past and end in the future. We then compute a plan’s success probability by summing over trajectories that achieve the goal. Again due to complexity reasons, we approximate the success probability by generating only the  $k$  most probable trajectories. But since we have only a single optimization problem now, we don’t have to prematurely cut off unlikely hypotheses and have only one source of error, compared to approximating the belief state and predicting the plan’s evolution based on this estimate.

## 4.3 Approximating the Plan Success Probability

We denote the set of all trajectories as  $\Theta$  and the set of the  $k$ -best trajectories as  $\Theta^*$ . A trajectory is considered successful if it entails the plan’s goal constraint. We define  $SUCCESS := \{\theta \in \Theta | \forall s \in \mathcal{R}_{sol}, s \downarrow_Y = \theta : F_G(s) = \text{true}\}$ , where  $\mathcal{R}_{sol}$  is the set of all solutions to the probabilistic constraint optimization problem,  $s \downarrow_Y$  their projection on marking variables, and  $F_G(s)$  is the goal constraint.  $SUCCESS^*$  is the set of successful trajectories among  $\Theta^*$ . The exact success probability is computed as

$$\begin{aligned}
P(SUCCESS|Obs, \mathcal{P}) &= \\
&\sum_{\theta \in SUCCESS} P(\theta|Obs, \mathcal{P}) = \\
&\sum_{\theta \in SUCCESS} \frac{P(\theta, Obs, \mathcal{P})}{P(Obs, \mathcal{P})} = \\
&\sum_{\theta \in SUCCESS} \frac{P(\theta, Obs, \mathcal{P})}{\sum_{\theta \in \Theta} P(\theta, Obs, \mathcal{P})} = \\
&\frac{\sum_{\theta \in SUCCESS} P(\theta, Obs, \mathcal{P})}{\sum_{\theta \in \Theta} P(\theta, Obs, \mathcal{P})}
\end{aligned}$$

The approximate success probability  $P^*(SUCCESS^*|Obs, \mathcal{P})$  is computed the same way, only  $SUCCESS$  is replaced with  $SUCCESS^*$  and  $\Theta$  with  $\Theta^*$ . We define the error of the above  $k$ -best approximation as  $E(k) := |P(SUCCESS|Obs, \mathcal{P}) - P^*(SUCCESS^*|Obs, \mathcal{P})|$ .  $E(k)$  converges to zero as  $k$  goes to infinity. Also,  $E(k) = 0$  if  $P(SUCCESS|Obs, \mathcal{P})$  is 0 or 1. However, as the example in figure 6 shows,  $E(k)$  does in general not decrease monotonically with increasing  $k$ .

## 4.4 Algorithm for Plan Evaluation

Plans are generated by the planner and then advanced until they are finished or new observations are available. In the latter case the currently executed plan is evaluated using Algorithm 1. It first computes the  $k$ -best solutions to the COP using an external solver (toolbar in our case). This results in the  $k$  most probable trajectories. Then, using these trajectories, it approximates the success probability of plan  $\mathcal{P}$  and finally compares the probability against the two thresholds  $\omega_{success}$  and  $\omega_{fail}$ . Now we have to address one of three cases: (1) The probability is above  $\omega_{success}$ , i.e. the plan will probably succeed, (2) the probability is below  $\omega_{fail}$ , i.e. the plan will probably fail or (3) the probability is in between both thresholds, which means the case cannot be decided. In the first case we simply continue execution. In the second case we have to adapt the plan to the new situation. This is done by  $REPLAN(\mathcal{P}, \Theta^*)$ , which modifies the future actions of  $\mathcal{P}$  taking into account the diagnostic information contained in  $\Theta^*$ . The third case indicates that not enough information about the system’s current state is available. As a reaction, the procedure  $REPLANPERVASIVEDIAGNOSIS(\mathcal{P}, \Theta^*)$  implements a recently developed method called *pervasive*

diagnosis [5]. It addresses this problem by augmenting a plan with information gathering actions (we do not detail the procedures REPLAN and REPLANPERVASIVEDIAGNOSIS as they are beyond this paper’s scope).

---

**Algorithm 1**

---

```

1: procedure EVALUATEPLAN( $\mathcal{R} = (X, D, C)$ ,  $Obs$ ,  $\mathcal{P}$ )
2:    $\mathcal{R}' \leftarrow$  add constraints over  $Obs$  and  $\mathcal{P}$  to  $\mathcal{R}$ 
3:    $\Theta^* \leftarrow k$ -best solutions of  $\mathcal{R}'$  for  $Y$ 
4:    $p \leftarrow P^*(SUCCESS^*|Obs, \mathcal{P})$ 
5:   if  $p > \omega_{\text{success}}$  then return
6:   else if  $p < \omega_{\text{fail}}$  then
7:     stop execution of  $\mathcal{P}$ 
8:     REPLAN( $\mathcal{P}$ ,  $\Theta^*$ )
9:   else
10:    stop execution of  $\mathcal{P}$ 
11:    REPLANPERVASIVEDIAGNOSIS( $\mathcal{P}$ ,  $\Theta^*$ )
12:  end if
13: end procedure

```

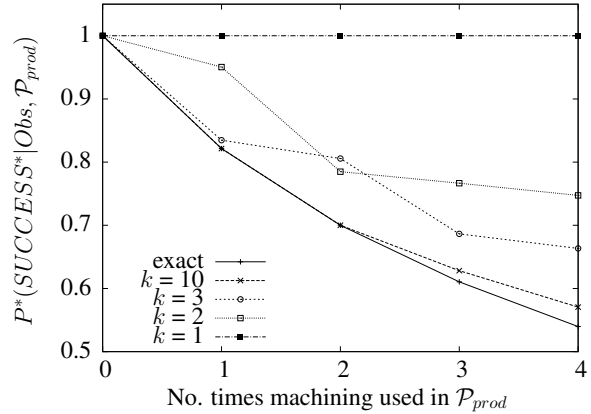
---

## 5 Experimental Results

We ran experiments for five small variations of our example scenario, where  $\mathcal{P}_{prod}$  uses the machining station zero to four times. The time window size  $N$  accordingly ranges from 2 to 6, problem sizes range from 240 to 640 variables and 240 to 670 constraints. Figure 6 shows the success probabilities for different  $\mathcal{P}_{prod}$  and  $k$ . Table 2 shows the runtime in seconds and the peak memory consumption in megabytes for computing success probabilities in the planning scenarios, additionally ranging over different values for the mini-bucket parameter  $i$ . As expected, with increasing use of the machining station,  $P^*(SUCCESS^*|Obs, \mathcal{P}_{prod})$  decreases. Also, runtime increases for larger time windows. The effect of approximation (choosing lower  $k$ ) is that  $P^*(SUCCESS^*|Obs, \mathcal{P}_{prod})$  increasingly deviates from the exact solution. In our example, the approximation tends to be optimistic. In general, however, we think that  $P^*(SUCCESS^*|Obs, \mathcal{P}_{prod})$  can be pessimistic, if success trajectories are pruned first when decreasing  $k$ . Increasing  $k$  hardly seems to affect the runtime, especially if the mini-bucket search heuristic is strong (bigger  $i$ -values). For weaker heuristics the influence increases slightly. Memory consumption is affected much stronger by  $k$ . Here also, a weaker search heuristic means stronger influence of  $k$ .

## 6 Related Work

In probabilistic verification of model-based programs [7], the problem is to determine the most likely circumstances under which a high-level control program drives the system towards a goal violating state. A plan can be understood as such a high level control program; so



**Figure 6. Approximate success probability (y-axis) of plan  $\mathcal{P}_{prod}$  against varying usage of the machining station (x-axis) after the observation of an abrasion at  $t = 2$ .**

in general, this problem is similar to the plan assessment problem. However, our problem differs in that we are interested in the *set* of all goal achieving system trajectories, from which we derive the plan’s success probability, while for the verification problem, only the single most probable goal violating trajectories are interesting. Therefore we have to go one step further, not only enumerating the trajectories, but also summing over them to compute the success probability.

McDermott [8] and Beetz’s [1] Reactive Plan Language (RPL) chooses a different approach to deal with system failures and uncertainty. It uses a hierarchical task decomposition, breaking down top level goals to a finer granularity recursively. The plan itself is not a sequence of actions but executable code. The language allows reasoning on and transformation of the plans. Heuristic routines attain the subgoals and cope with failures and unexpected events during the execution. A goal for finding a cup could e.g. look in the dishwasher after seeing that no cups are left in the cupboard. This approach is particularly promising in domains of high uncertainty, where classical planning fails. However, the RPL approach currently neglects explicit diagnosis techniques and relies on the observability of relevant environment states.

## 7 Conclusion and Future Work

We presented a model-based method that combines diagnosis of past execution steps with prognosis of future execution steps of production plans, in order to allow the production system to autonomously react to failures and other unforeseen events. The method makes use of probabilistic constraint optimization to efficiently solve this combined diagnosis/prognosis problem. Preliminary re-



$k$	$i$	No. times machining used in $\mathcal{P}_{prod}$ (window size $N$ , #Variables, #Constraints)				
		0 (2, 239, 242)	1 (3, 340, 349)	2 (4, 441, 456)	3 (5, 542, 563)	4 (6, 643, 670)
1	10	< 0.1 / 1.8	0.1 / 6.8	0.1 / 19.0	(mem)	(mem)
	15	0.1 / 1.9	0.3 / 4.2	0.5 / 7.8	0.5 / 16.6	0.8 / 32.0
	20	0.1 / 1.9	0.5 / 5.2	3.7 / 20.1	6.5 / 34.5	9.5 / 50.7
2	10	< 0.1 / 2.1	0.1 / 11.9	0.2 / 38.5	(mem)	(mem)
	15	0.1 / 2.2	0.3 / 5.4	0.5 / 9.7	0.6 / 28.0	0.8 / 52.0
	20	0.1 / 2.2	0.5 / 6.4	3.7 / 21.8	6.5 / 37.2	9.5 / 55.8
3	10	< 0.1 / 2.3 (e)	0.1 / 11.9	0.2 / 40.1	(mem)	(mem)
	15	0.1 / 2.4 (e)	0.3 / 5.4	0.5 / 11.4	0.6 / 29.9	0.9 / 55.5
	20	0.1 / 2.4 (e)	0.5 / 6.4	3.7 / 23.5	6.6 / 38.3	9.5 / 57.4
4	10	(e)	0.1 / 12.5	0.2 / 40.1	(mem)	(mem)
	15	(e)	0.3 / 5.9	0.5 / 11.4	0.6 / 30.9	0.9 / 57.2
	20	(e)	0.5 / 6.9	3.7 / 23.5	6.6 / 39.3	9.5 / 59.1
5	10	(e)	0.1 / 13.1	0.2 / 40.7	(mem)	(mem)
	15	(e)	0.3 / 6.6	0.5 / 12.0	0.6 / 33.6	0.9 / 59.5
	20	(e)	0.5 / 7.6	3.7 / 24.0	6.6 / 42.8	9.5 / 63.9
10	10	(e)	0.1 / 14.0 (e)	0.2 / 43.4 (e)	(mem)	(mem)
	15	(e)	0.3 / 6.7 (e)	0.5 / 14.7 (e)	0.6 / 36.2	0.9 / 64.8
	20	(e)	0.6 / 7.7 (e)	3.8 / 26.6 (e)	6.6 / 45.8	9.6 / 68.9

**Table 2. Runtime in seconds / peak memory consumption in megabytes. (e) indicates that the exact success probability  $P(SUCCESS|Obs, \mathcal{P}_{prod})$  could be computed with this configuration. (mem) indicates that A\* ran out of memory (artificial cutoff at > 1 GB, experiments were run on a Linux computer with a recent dual core 2.2 Ghz CPU with 2 GB RAM).**

sults for a real-world machining scenario show it can indeed be used to guide the system away from plans that rely on suspect system components. Future work will concern the integration of the method into our overall planning/execution architecture, and its extension to multiple, simultaneously executed plans. We are also interested in exploiting the plan diagnosis/prognosis results in order to update the underlying system model, for instance, to automatically adapt to parameter drifts or wear of components.

## References

- [1] M. Beetz. *Concurrent Reactive Plans: Anticipating and Forestalling Execution Failures*, volume 1772 of *Lecture Notes in Artificial Intelligence*. Springer Publishers, 2000.
- [2] M. Beetz, M. Buss, and D. Wollherr. Cognitive technical systems — what is the role of artificial intelligence? In *Proc. KI-2007*, pages 19–42, 2007.
- [3] S. Bouveret, F. Heras, S. Givry, J. Larrosa, M. Sanchez, and T. Schiex. *Toolbar: a state-of-the-art platform for wesp*. [www.inra.fr/mia/T/degivry/ToolBar.pdf](http://www.inra.fr/mia/T/degivry/ToolBar.pdf), 2004.
- [4] K. Kask and R. Dechter. Mini-bucket heuristics for improved search. In *Proc. UAI-1999*, pages 314–32, 1999.
- [5] L. Kuhn, B. Price, J. d. Kleer, M. B. Do, and R. Zhou. Pervasive diagnosis: The integration of diagnostic goals into production plans. In *Proc. AAAI-2008*, 2008.
- [6] J. Kurien and P. P. Nayak. Back to the future for consistency-based trajectory tracking. In *Proc. AAAI-2000*, pages 370–377, 2000.
- [7] T. Mahtab, G. Sullivan, and C. B. Williams. Automated Verification of Model-based Programs Under Uncertainty. In *Proceedings 4th International Conference on Intelligent Systems Design and Application*, 2004.
- [8] D. McDermott. A reactive plan language. Technical report, Yale University, Computer Science Dept., 1993.
- [9] T. Mikaelian, B. C. Williams, and M. Sachenbacher. Model-based Monitoring and Diagnosis of Systems with Software-Extended Behavior. In *Proc. AAAI-05*, 2005.
- [10] T. Schiex, H. Fargier, and G. Verfaillie. Valued constraint satisfaction problems: hard and easy problems. In *Proc. IJCAI-1995*, 1995.
- [11] B. C. Williams, S. Chung, and V. Gupta. Mode estimation of model-based programs: monitoring systems with complex behavior. In *Proc. IJCAI-2001*, pages 579–590, 2001.