The Application of Automated Planning to Machine Tool Calibration

Simon Parkinson and **Andrew P. Longstaff**
Centre for Precision Technologies
School of Computing and Engineering
University of Huddersfield, UK
s.parkinson@hud.ac.uk
a.p.longstaff@hud.ac.uk

Andrew Crampton and **Peter Gregory**
Department of Informatics
School of Computing and Engineering
University of Huddersfield, UK
a.crampton@hud.ac.uk
p.gregory@hud.ac.uk

Abstract

Engineering companies working with machine tools will often be required to calibrate those machines to international standards. The calibration process requires various errors in the machine to be measured by a skilled expert. In addition to conducting the tests, the engineer must also plan the order in which the tests should take place, and also which instruments should be used to perform each test. It is critical to find as short a calibration plan as possible so that the machine is not out of service for too long.

In this work, automated planning is applied to the problem of generating machine tool calibration plans. Given a description of a machine, and its various axes, we produce a calibration plan that minimises the time taken to measure all of the errors of a machine. We also consider the case when there is not enough time to test all errors of the machine, and the calibration plan must maximise the importance of the set of errors tested in the limited time available.

Introduction

Generating process plans automatically is a challenging task. However, the economic advantages are seen as significant to engineers (Gupta, Nau, and Regli 1998). This is because the ability to create both an efficient and complete process plan can result in minimising the risk of problems occurring that could ultimately result in excessive expenditure. This is true for the process of machine tool calibration planning (Bringmann, Besuchet, and Rohr 2008).

The requirement to manufacture more accurate parts and minimise manufacturing waste is resulting in the continuing requirement for machine tools which are more accurate. Therefore, machine tool calibration is required regularly to gain an understanding of a machine's capability. When planning a machine tool calibration, an engineer will derive a calibration plan based on many influencing factors. Reducing the time taken to perform a calibration is fundamental to many engineering companies. For example, the machine downtime cost for a large company operating a production line could be in excess of a £1000 per hour.

In this paper, a solution to the problem of machine tool calibration planning is presented by modelling the problem using a Hierarchical Task Network (HTN) (which is implemented using the SHOP2 (Nau et al. 2003) architecture) and also by using PDDL.

Calibration planning can naturally be thought of as being a sequence of tasks, each one requiring the set-up of an instrument, measurement of an error, and finally the tear-down of the instrument. In this way, machine tool calibration seems well suited to being represented by an HTN. SHOP2, however, has no explicit model of time, and so true temporal action concurrency is not possible. We also encode the problem using PDDL to exploit the potential concurrency in the problem and to solve the problem using state-of-the-art domain independent planners. The PDDL model is interesting as it combines several different PDDL requirements: quantification, numbers, time and timed initial literals. And since it models a real-world problem very closely, it provides an interesting challenge to existing and new planners. The model and sample problem instances will be made available to the planning community as a new benchmark.

In the next sections, we provide details of the machine tool calibration problem, and our HTN and PDDL models. We provide details of the constraints on this new planning benchmark, with some initial results. Exploiting the concurrency in the problem improves performance considerably. We also show results maximising the importance of the errors tested when there is not enough time to test all of the errors in the machine. The most complex version of our problem is not solved in the cutoff time, suggesting that the domain is a useful benchmark.

Related Work

There has been previous interest (Parkinson et al. 2011) in encoding the process of machine tool calibration in first-order logic. However, that work uses this information for querying purposes in a decision support tool for an engineer constructing a calibration plan. In this paper, we propose going further than this and use an automated planner to construct the calibration plan for the engineer.

Machine tool calibration is just one of the important problems studied in metrology. Automated planning is a good fit for solving many of the other complex planning problems in dimensional metrology. Precision measurement of the complex geometries of various artifacts is important to many different organisations and National Measurement Institutions.

Many authors have developed models for planning the measurement of artifacts. Muelaner et al. (Muelaner, Cai, and Maropoulos 2010) developed a model in a procedural programming language to aid with instrumentation selection for dimensional measurement of large volume artifacts. However, this model, like others, is static in its approach.

The IMACS planning system (Gupta, Nau, and Regli 1998) and the CAPP planning system (Deak et al. 2001) are examples of the use of planning within a manufacturing environment. Automated planning has also been applied to the area of business process management (Hoffman, Weber, and Kraft 2010). Other related work includes recent successful deployments of Automated Planning techniques in real-world applications. The use of planning in modular printing (Ruml et al. 2011) for example, and the use of planning to control multi-cell batteries (Fox, Long, and Magazzeni 2011).

Machine Tool Calibration

As previously identified in the introduction, machine tool calibration is based on many influencing factors. For the context of this paper, the following section contains enough information regarding the influencing factors of machine tool calibration to allow the reader to understand the planning problem in sufficient detail.

Machine Configuration

A machine tool can be designed and constructed in many different ways to perform its task. Figure 1(a) shows a cross-table machine tool with three perpendicular linear axes, while Figure 1(b) shows a gantry machine tool with three perpendicular linear axes and two rotary axes. Different configurations are required for manufacturing components of different size and material. For example, it is less efficient to use a machine configuration where the workpiece is situated on a cross-table for the manufacturing of large, heavy parts. This is because the machine would be required to move the workpiece during manufacturing. Conversely, a gantry machine would move around the workpiece, reducing the amount of energy required to machine the item as well as reducing the structural strain on the machine from the workpiece, but has a larger overall footprint.

In addition to the number of linear and rotary axes, the configuration (stacking) of these axes can cause errors to propagate differently throughout the machine. The configuration of a machine tool will determine how many error components it has. While there are few common machine configurations, there are a lot of different configurations which require in-depth consideration to identify all of their error components. The configuration is heavily dependent on the item that it will be manufactured, meaning that the variety of machine tool configurations in somewhat proportional to the variety of items that they manufacture. It is also possible that a machine tool might have additional auxiliary axes to load an unload the workpiece from the machine, and these will also have errors which require calibration.

Each error component will have an effect on the overall geometric accuracy of machine tool. However, depending on the machine configuration, each error component has a different significance. For example, taking the three-axis machine as seen in Figure 1(a), the roll of the vertical Z-axis can be regarded as having a lower importance when compared to the roll error of the Y- and X-axis. This is because any Z about Z (roll) movement along the Z-axis will only affect the rotation of machine's tool position (which during milling rotates anyway), whereas any roll in the X- and Y-axis would result in the rotation of the workpiece. However, the significance of the Z about Z error component for a five-axis gantry machine, as seen in Figure 1(b), is greater. This machine has a kinematic chain where the C-axis is mounted on the Z-axis and the A-axis is mounted on the C-axis. Any error of the Z-axis roll will be propagated down the kinematic chain resulting in the incorrect positioning of the A-axis, which would be directly evident on any machining procedures that involve the rotation of the A-axis.

The configuration of the machine's constituent parts determines the potential geometric errors that a machine might have. The geometric errors associated with linear and rotary axes are well known (Bohez et al. 2007). Each linear axis will have the following quasi-static errors (not including spindle errors). For example purposes we are referring to the X-axis of a Cartesian machine as illustrated in Figure 2.

EXX	Linear positioning
EXZ	Vertical straightness
EXY	Horizontal straightness
EBX	Yaw (X about Y)
EAY	Pitch (X about Z)
ECX	Roll (X about X)

In addition to the six-degrees-of-freedom errors, there will the cross-axis errors of:

EXY, EXZ, EYZ Squareness with each perpendicular axis

If we consider the addition of rotary axes to the machine's kinematic chain, the following error components, as seen in Figure 3, are introduced for the C-axis:

ECC	Angular positioning error
EYC, EXC	Radial error motion in X and Y
EAC, EBC	Tilt error around both the X and Y axis
EZC	Axial error

In addition to error components, there are also location errors in respect to the machine's coordinate system. For the same C-axis, these are:

XOC	X position of C
YOC	Y position of C
AOC	Squareness of C to Y
BOC	Squareness of C to X

From this it is possible to deduce that a three axis machine tool as seen in Figure 1(a) will have in total 21 geometric

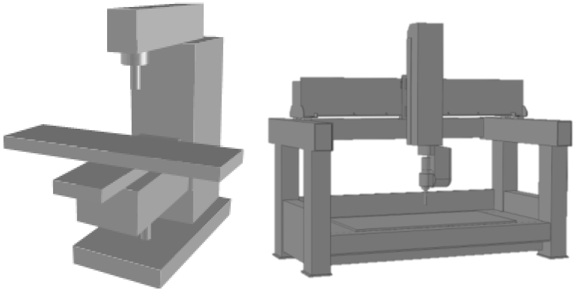

(a) Three axis (b) Five axis

Figure 1: Two different machine configurations

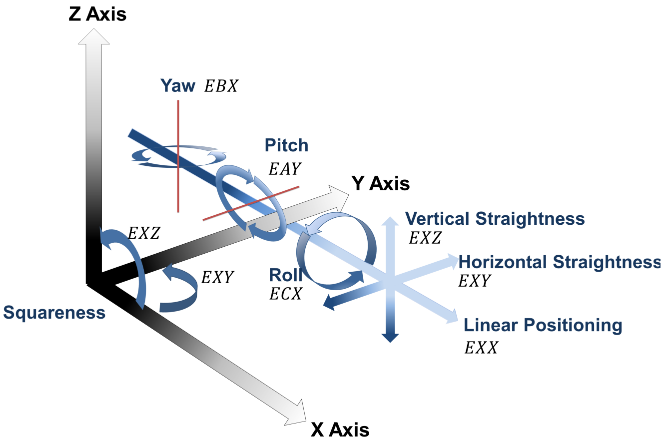

Figure 2: Six-degrees-of-freedom and squareness errors for the X-axis of a machine tool with three perpendicular linear axes

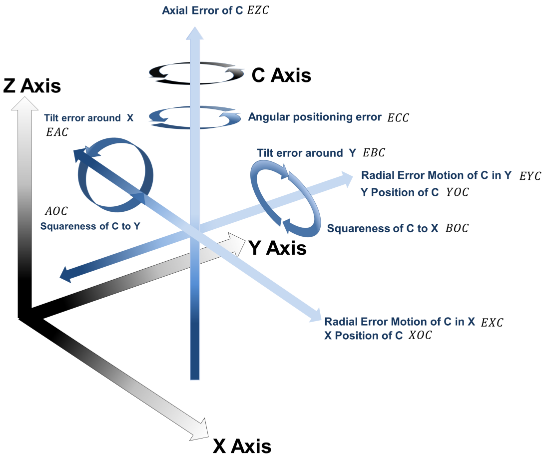

Figure 3: Motion and location errors of an axis of rotation (C-axis)

errors (Ramesh, Mannan, and Poo 2000), and that a five-axis machine tool as seen in Figure 1(b) has the error count of 41. A machine tool will, however, actually experience more error sources such as thermal, dynamic and non-rigid (Mekid 2009). For the scope of this paper, only the calibration planning problem for geometric errors in machine tools is considered.

Error Measurement

The measurement of an error component will involve setting-up the equipment and the movement of the associated axes during measurement. However, when performing the measurement it is essential that axes not being tested are kept stationary. This is because any movement of the machine could result in non-rigid and dynamic errors affecting those geometric errors that are being tested for. It is an important fundamental in metrology to maintain a high level of measurement repeatability, and systematically ensuring that only the relevant parts of the machine are moving can help achieve this. In addition to maintaining a quality measurement process, other factors such as instrument interference require the ordering of measurements to be sequential. However, there are conditions which would allow for multiple error components to be measured concurrently, while ensuring that a high level or repeatability is maintained. For example, it could be possible that the linear positioning error component, using laser interferometry, and the roll error component, using an precision level, could be tested for concurrently. This is providing that the machine's axis has enough physical room to locate both the instruments, and that the test parameters agree.

When testing for quasi-static errors of each axis, tests will involve moving the machine in specified increments before taking a measurement. For example, when testing the linear positioning, the difference between the machine's expected and actual position is examined at multiple locations throughout the travel range of the axis. If the machine has enough physical space to allow for multiple instruments to be installed at any one time, then there is the possibility of concurrent measurement. This would require that the following test parameters agree:

1. The speed at which the machine moves between the target positions (feedrate).
2. The position of the first and last target on the axis.
3. The number of target positions throughout the axis travel between the first and last position.
4. The distance between the target positions.
5. The dwell time at each target, which includes the time required for the machine to stabilise from any dynamic motion errors and the time to take the measurement.

To extend the example of testing for roll and linear positioning concurrently, the following equation taken from (ISO 230-2:2006) describes the selection of target positions

$$Pi = (i-1)p + r$$

where i is the number of current target positions; p the nominal interval based on a uniform spacing of target positions over the measurement travel; r a random number within ± the amplitude of possible periodic errors.

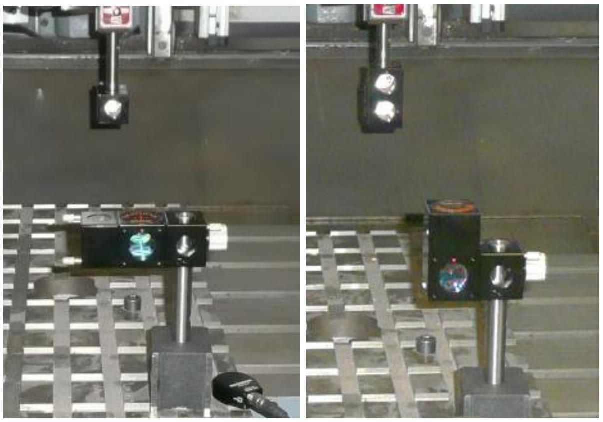

(a) Linear Position Optics (b) Pitch Optics

Figure 4: Two different tests for errors.

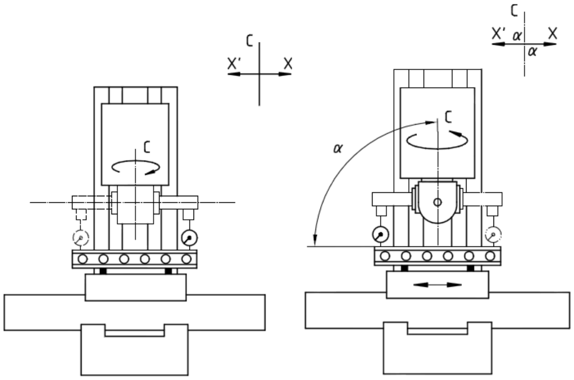

(a) XY plane Squareness (b) Squareness between C-axis and X-axis motion

Figure 5: Displacement sensor setup for measuring multiple errors (ISO 10791-2:2001)

If these parameters for multiple measurements do agree, and the machine has no physical restrictions, then the measurements could take place concurrently.

Instrumentation

The extensive variety of instrumentation available for performing a machine tool calibration adds complexity to deciding the optimum solution when measuring each error component. There are many different reasons why a specific instrument might be selected. The following list supplies some criteria which would influence the instrumentation selection.

1. The time to install and align the equipment may be lower, which could help to minimise machine downtime.
2. The resolution and accuracy of the instrument might be greater, making it a better choice to improve the quality of the calibration.
3. The instrument can measure multiple degrees-of-freedom concurrently, allowing for the simultaneous measurement of multiple error components.
4. The instrument might be better suited to the machine's physical characteristics which could place restrictions on the available space.

It is important to ensure that the instrumentation chosen can help to reduce the overall calibration time, because for most organisations, machine downtime can significantly reduce or even halt manufacturing. Therefore, identifying where the use of certain instrument can reduce the set-up and measurement time, is regarded as a beneficial quality. For example, measuring the Y-axis linear positioning error using the Renishaw XL-80 laser interferometer would require the configuration of the optics as seen in Figure 4(a). Next, measuring the Y-axis pitch error would require the use of the optics aligned as seen in Figure 4(b). However, because the optics and the laser are already aligned, it is possible to carefully exchange the optics and the laser will still be aligned parallel with the axis under examination.

In addition to the consideration taken towards the setup time, the instrument's ability to measure multiple error components (degree-of-freedom) concurrently is considered. For example, the test setup involving two displacement sensors in Figure 5(a) can be used for the test procedure as seen in Figure 5(b). Even though the second procedure will require the movement of the X-axis, meaning that it can not be carried out simultaneous to the first, the instrumentation will not require any repositioning or adjustment.

So far we have described the requirements on instrumentation selection based on the improving the calibration process and reducing the overall time. However, there are logistical reasons as to why the distribution of instrumentation is important. A machine tool calibration company might have many concurrent calibration jobs for many different clients. In this situation, the optimisation of instrumentation selection should take consideration to the overall distribution of instrumentation as well as for each calibration job. It could be that selecting the most efficient instrumentation for one job could lead to a greater decrease in efficiency for another calibration job, so company-wide instrumentation allocation should be considered.

HTN Model

As identified in the introduction, the planning problem of machine tool calibration is well suited to being represented as an HTN. The following section shows how machine tool calibration was broken down into smaller tasks to create an HTN.

Task Decomposition

Task decomposition is the process of breaking tasks into smaller tasks until primitive actions are reached. Figure 6 shows the abstract task decomposition for calibrating a machine tool, which takes into consideration what has been regarded as the main calibration tasks. A description for each primitive subtask can be found in the following list:

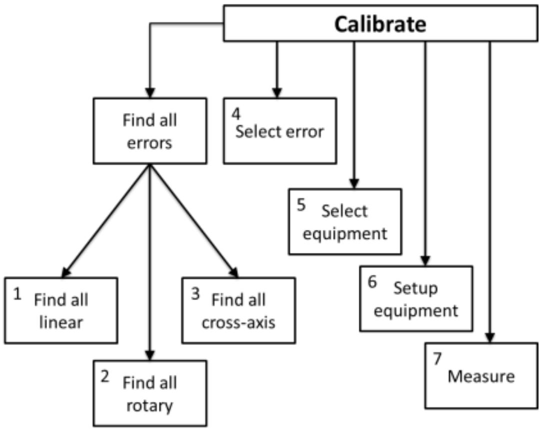

Figure 6: Task decomposition tree

1. Find all linear errors based on the machine's configuration.
2. Find all rotary errors based on the machine's configuration.
3. Find all cross-axis errors based on the configuration of the linear and rotary axes.
4. Select an error component for measuring.
5. Select the suitable equipment for measuring the error component.
6. Set-up the equipment in a suitable way to measure the error component.
7. Measure the error component using the instrumentation and the current setup

The process of performing this manual task decomposition to convert the nonprimitive task of machine tool calibration into the primitive tasks will serve as the basis for creating an HTN network. We omit the low-level description of the HTN.

Plan Metric

A SHOP2 operator also expresses a cost for performing the primitive task. The operators used in the machine tool calibration HTN have a cost assigned which is originally acquired from the initial state facts. The motivation behind an operators cost is explained below:

Error selection This is the importance of an error component. An error component that is regarded as having a high significance, or that should be measured first, is assigned a lower cost value.

Equipment set-up cost The cost in minutes that is required for setting up the instrumentation out of the box.

Equipment adjustment cost This is the cost in minutes for adjusting the equipment if it is already set up on the axis. For example, realigning the optics of a laser interferometer.

Performing the measurement cost This is the cost in minutes for measuring the error component using the selected equipment.

PDDL Model

We have also constructed a PDDL model of the problem. In the HTN encoding, the durations of the actions are encoded as part of the cost function, as it is important to minimise the time taken to calibrate the machine. In the PDDL model we encode time using durative actions. This provides the benefit of allowing concurrent activities. It also necessitates encoding various temporal constraints which limit the potential concurrency. The temporal constraints on the problem are as follows:

T1 All instruments that are currently set-up must be set up on the same axis. This constraint is necessary as running a test on one axis means that instrumentation set-up on other axes will need to be set-up again.

T2 Each instrument has a maximum number of tests which it can perform simultaneously. Most instruments can only perform one test at a time. However, some instrumentation can perform several tests simultaneously. As an example, consider the displacement sensors on a rotary axis (discussed previously).

T3 On each axis, it is not possible to use certain equipment simultaneously. This may be because of space restrictions on a small axis working area. It could also be that certain instruments interfere with other instruments (for example, by blocking a laser).

T4 Testing can happen over a number of days. However, all tests conducted on each day must be contained within an eight hour period (modelling a working day).

There are also spatial constraints on the problem. These are due to how each machine is configured. For example, there may be no way to place optics on a particular axis, as it could be blocked by another part of the machine. The spatial constraints are defined as follows:

S1 Certain instruments are blocked for each axis. In this case, that instrument cannot take any measurements on the errors on this axis.

S2 The operating range of an instrument must be greater than the axis travel length. The travel length of the longest axis in a machine tool operating on an aircraft wing, for example, is over 30 meters. Many optical tests cannot be used over this distance.

It would have also been possible to encode specific action for each instrument. However, we have designed the planning model to be flexible to new instrumentation. New instrumentation could be developed that tests a combination of errors currently not tested together by current instrumentation. Our model is general enough to allow this hypothetical instrument to be included.

PDDL Operators

In our model, there are three object types: axis, instrument and error. There are four operators: set-up, adjust, measure

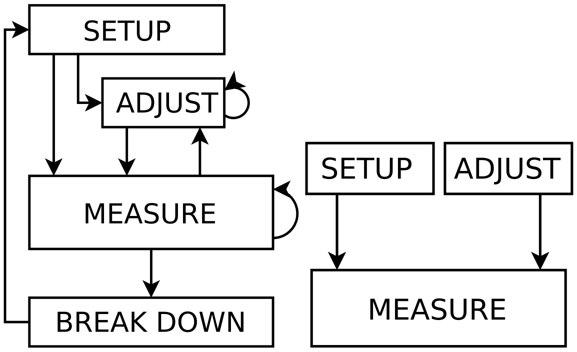

(a) Instrument Object (b) Error Object

Figure 7: Diagrammatic illustration of the timeline of instrument and error objects in the PDDL model.

and teardown. Figures 7(a) and 7(b) show at a high level how the operators manipulate objects of the instrument and error types. Taking the instrument type first (Figure 7(a)), the instrument must first be set-up on an axis. Once set-up on an axis (and given that it can measure multiple errors simultaneously, by constraint **T2**) it may be adjusted so that multiple measurements can be taken. Measurements can then be taken using the measure action. Once finished, the instrument can be adjusted to take more measurements. The number of 'measure' actions will balance the sum of the 'set-up' and 'adjust' actions. After all measurements are taken, the instrument is removed from the axis, using the 'tear down' action. When the instrument is used again, it must be set-up again, either on the same or a different axis. As can be seen in Figure 7(b), each error object is either set-up, or adjusted, to measure. Once measured, the error cannot be measured again.

We now describe the operators in more detail. We have labelled the preconditions and effects relating to the temporal and spatial constraints with the constraint labels defined above. We omit the PDDL source: however, this is available from the authors on request.

The set-up action sets an instrument up on an axis:

SET-UP	
parameters	?i - instrument ?a - axis ?e - error
duration	set-up time of ?i on ?a
preconditions	**S1**: ?i is not blocked on ?a **S2**: operating range of ?i is sufficient **T1**: for all set-up instruments ?j, ?j is set-up on ?a **T3**: for all set-up instruments ?j, ?j is compatible with ?i **T4**: set-up occurs during the working day
effects	?i is set-up to test ?e on ?a **T2**: increment the number of tests being performed by ?i

The adjust action adjusts an instrument to measure an error when the instrument is already set-up on an axis:

ADJUST	
parameters	?i - instrument ?a - axis ?e - error
duration	adjustment time of ?i on ?a
preconditions	?i is set-up on ?a **T2**: the number of tests ?i is currently set-up for is fewer than the maximum allowed for ?i **T4**: set-up occurs during the working day
effects	?i is set-up to test ?e on ?a **T2**: increment the number of tests being performed by ?i

The measure action measures an error given that the instrument is set-up to measure that error on an axis.

MEASURE	
parameters	?i - instrument ?a - axis ?e - error
duration	measurement time of ?i for ?e on ?a
preconditions	?i is set-up on ?a for ?e ?e has not been measured on ?a **T4**: set-up occurs during the working day
effects	?e is measured on ?a **T2**: decrement the number of tests being performed by ?i increase the global importance by the importance of ?e on ?a

The tear down action removes an instrument from an axis, provided that the instrument is not currently set-up to measure any of the errors on that axis.

TEAR-DOWN	
parameters	?i - instrument ?a - axis
preconditions	?i is set-up on ?a **T2**: the number of tests being performed by ?i is 0
effects	?i is not set-up on ?a

In order to encode constraint **T4**, we introduce a zero-arity predicate *working-day*. In the initial state, this predicate is true. Then, by using timed initial literals, we alternate the value of the predicate from true to false. At time 480 (as there are 480 minutes in an eight hour working day), we set the predicate to false. For the sake of convenience, we set it to true at 1000 minutes - this constitutes the second working day. We follow this pattern for as many working days as we choose.

Plan Metric

Each of the different errors on a machine has a different importance value. Depending on the configuration of a machine, different axes will also hold more importance. This is typically due to which axis holds the other axes. When there is insufficient time to fully calibrate a machine, it is still desirable to test the most important errors in the time available. Therefore, we maintain a 'global importance' fluent that sums the importance of the errors measured in the plan. The importance of an error for an axis is taken as the product of the importance of the axis and the importance of the error independent of a particular configuration.

Initial and Goal States

The goal state for the task of measuring all errors is that all errors on each axis are measured. For the optimisation version of the task, where we wish to maximise the importance of the errors measured in a limited time, we pose an empty goal as it is not known in advance which errors should be measured to maximise the importance value.

The initial state in both cases consists of predicates and fluents that encode the setup and adjustment times for each instrument; which instruments measure which errors; the errors which need to be measured on each axis and the importance of each error and axis.

Experimental Analysis

The benchmarks that we provide are artificial to some degree. However, we are working with calibration engineers, who will test our system. Calibration engineers have compared the plans we produce to actual calibration plans and report they are similar and sensible. Also, the benchmarks are close to reality; the configurations of the machines are common configurations and the timings are derived from similar real machines. We do have expert-constructed calibration plans to compare to for real machines, however this is commercially sensitive data.

Our instances are based on four different machine configurations. Firstly, we test on two machine configurations with three linear axes. Each linear axis will have six geometric plus one squareness error component. There are a total of eight different instruments available, and each error component can be measured by using at least two of the available instruments. Secondly, we test on two five axis machine configuration with three linear and two rotary axes. Each linear axis will have six geometric plus one squareness error components, and each rotary axis will have ten error components. There will also be a total of eight different instruments available, and each error component can be measured by using at least two of the available instruments. For each machine, we have three different instances (denoted A, B and C in the tables) which correspond to models with different timings for setting up and adjusting the instruments. Even for the same machine, depending on the experience of the engineer, setting up and adjusting instrumentation will take a variable amount of time.

We have implemented the HTN model using the SHOP2 (Nau et al. 2003) planner. When searching, we use the built-in branch and bound search, optimising the makespan. We use the LPG-TD planner for the PDDL model (Gerevini, Saetti, and Serina 2006) as it handles all of our requirements (timed initial literals (Edelkamp and Hoffman 2004), durative actions, plan metrics (Fox and Long 2003) and quantified preconditions). We compare the quality of our plans to the SHOP2 plans. All experiments are run on a desktop computer with an Intel Core i5 CPU (2.80 GHz.) and 8GB of system RAM. To run SHOP2 we use Steelbank Common Lisp 1.0.54. In all of the experiments, we use a ten minute cutoff.

Instance	SHOP2	LPG_S	LPG_C
3AX-01A	1817	1984	760
3AX-01B	1603	1628	762
3AX-01C	1635	1748	716
5AX-01A	3275	3235	1809
5AX-01B	2759	2979	1793
5AX-01C	2759	3105	1706
3AX-02A	1774	1828	794
3AX-02B	1561	1661	680
3AX-02C	1180	1092	494
5AX-02A	3040	2793	1516
5AX-02B	2820	2216	1506
5AX-02C	2220	2272	1249

Table 1: Comparison Between SHOP2 and LPG-TD on 12 Machine Tool Calibration Instances. Six of the instances are from 3-axis machines, six from 5-axis machines. The results show the length of the plans in minutes. LPG_S and LPG_C are LPG finding sequential and concurrent plans, respectively.

We conduct two sets of experiments: the first to compare the calibration plans produced by the HTN planner and the calibration plans produced by LPG. We do not include the **T4** constraints, as these could not be encoded in the HTN model. We provide results with concurrent actions both allowed and disallowed for LPG, showing whether or not any benefit is gained from this approach.

Table 1 shows the results of running these tests. This work is not intended to show the relative merits of planning using HTN and PDDL encodings. Our results show, that in the sequential case, the HTN typically provides plans with shorter makespans. However, the differences are typically quite small, and it is clearly possible to find good solutions with either planning technique. Once concurrency is allowed, the PDDL model allows us to find much shorter plans, typically halving the plan length.

Our second set of experiments show the effect of introducing the working day constraints (**T4**) in the model. We show the results of minimising the makespan of the plan when the number of days available exceeds the minimum days required to calibrate the machine. We also show the result of maximising the importance of the tests carried out in the case when there is a limited time to carry out the calibration. Table 2 shows the result of introducing these constraints. We firstly allow the planner sufficient time in order to satisfy all of the goals. The makespan of these plans is shown in the first column ('Time'). Note that not all of the problems are solved: the 3-axis problems are solved, but the 5-axis ones are currently too challenging when the working

		1 Day		2 Days	
Instance	Time	Quality	Tests	Quality	Tests
3AX01A	2,155	5654	13	6689	19
3AX01B	2,123	5206	11	5902	15
3AX01C	1,350	6358	14	7416	21
5AX01A	–	4374	10	6013	18
5AX01B	–	5222	12	5286	15
5AX01C	–	4838	12	5517	18
3AX02A	2,270	5877	15	6285	17
3AX02B	1,477	5585	13	6492	21
3AX02C	1,116	6837	17	7416	21
5AX02A	–	4372	15	4594	20
5AX02B	–	4548	10	4905	15
5AX02C	–	5097	17	3568	15

Table 2: The results of solving the test instances with the working day constraints enabled. The first 'Time' column shows the makespan (as *days*,*minutes*) of the best quality found global solution. Instances marked with a dash were not solved within the cutoff. The remainder of the table shows the overall quality (in terms of sum importance of errors measured) of the plans found within a restricted makespan, and also the number of errors that were measured in those plans.

day constraints are taken into consideration. Because some of the problems cannot be solved, we believe that they form an interesting benchmark set and can help motivate planner development.

When solving the problems with limited makespan, we pose no goals, but we maximize the global importance. As can be seen in the results, LPG does solve these problems whilst taking into account the metric function. When given extra time, it solves the problem with a higher metric value. In some cases, for the 3-axis machines, there is sufficient time to satisfy all of the goals, these are the cases when 21 errors are measured.

Discussion and Future Work

In our model, we have provided plans that measure the geometric errors in a machine tool. We used this as a starting point because geometric errors are considered to be the most important. We will extend our model to take all other errors (thermal, rigid and dynamic) into account, and do not foresee the need to greatly change our model beyond adding a greater number of errors and instruments.

Currently, in industry, calibration plans are constructed for a single machine at a single site. An engineering company specializing in machine tool calibration could have multiple engineers working at various sites. Creating a calibration plan for multiple machine tools across different sites can help minimise the overall costs for an engineering company by optimizing instrumentation and labour use across those sites. If instrumentation is limited, it could be shared between different sites on different days, for example.

Efforts are being spent to develop an interface for engineers to use to generate calibration plans. This will enable us to test our approach in the field.

Conclusions

In our work we have determined that planning is a suitable technology for solving the problem of machine tool calibration, and also found challenging examples that could not be solved at all within our time-limit. However, the problem is an optimisation problem (either optimising makespan, or optimising utility in the limited-time version). Our benchmarks contribute to the planning community by providing challenging real-world instances in the area of optimal planning in temporal domains with numerics.

Machine tool calibration requires careful planning in order to minimize costs. Automated planning is an effective way to aid engineers engaged in this activity. There are two main economic beneficiaries when using an automated planning approach to solve this problem. Firstly, to the owner of the machine tool, who minimises the down-time of the machine. Secondly, to the engineer carrying out the calibration, who saves the time which would have been taken constructing a calibration plan manually and also the time saved in performing the calibration. When there is insufficient time to test all of the errors in a machine tool, the owner of the machine tool also benefits because automated planning tools can maximise the importance of the errors tested. Finally, use of automated planning enables an engineer to be more responsive to change if any instrumentation fails, or if the machine configuration is different to what was expected.

Calibrating machine tools is a time-consuming and expensive process. Skilled engineers are required to both perform the calibrations and also to construct the calibration plans. We have proposed using automated planning techniques to solve the problem of constructing calibration plans. We have modelled the problem in both an HTN language and PDDL. We have demonstrated that exploiting the possible concurrency in the problem can lead to much shorter plans than solving sequential versions of the same problem. We have demonstrated that, with all constraints enabled, calibration problems are at the boundary of what can currently be solved. We provide this new domain, with complex requirements, as a challenging new benchmark for automated planners.

References

Bohez, E. L. J.; Ariyajunya, B.; Sinlappcheewa, C.; Shein, T. M. M.; Lap, D. T.; and Belforte, G. 2007. Systematic geometric rigid body error identification of 5-axis milling machines. *Computer-Aided Design* 39(4):229–244.

Bringmann, B.; Besuchet, J.; and Rohr, L. 2008. Systematic evaluation of calibration methods. *CIRP Annals - Manufacturing Technology* 57(1):529–532.

Deak, F.; Kovacs, A.; Vancza, J.; and Dobrowiecki, T. 2001. Hierarchical Knowledge-Based Process Planning in Manufacturing. In *PROLAMAT*, 428–439.

Edelkamp, S., and Hoffman, J. 2004. PDDL2.2: The Language for the Classical Part of the 4th International Planning

Competition. Technical report, Albert-Ludwigs-Universität Freiburg.

Fox, M., and Long, D. 2003. PDDL2.1: An Extension to PDDL for Expressing Temporal Planning Domains. *J. Artif. Intell. Res. (JAIR)* 20:61–124.

Fox, M.; Long, D.; and Magazzeni, D. 2011. Automatic Construction of Efficient Multiple Battery Usage Policies. In *ICAPS*, 74–81.

Gerevini, A.; Saetti, A.; and Serina, I. 2006. An Approach to Temporal Planning and Scheduling in Domains with Predictable Exogenous Events. *JAIR* 25:187–231.

Gupta, S. K.; Nau, D. S.; and Regli, W. S. 1998. IMACS: A Case Study in Real-World Planning. *IEEE Intelligent Systems* 13(3):49–60.

Hoffman, J.; Weber, I.; and Kraft, F. M. 2010. SAP Speaks PDDL. In *AAAI*, 1096–1101.

ISO 10791-2:2001. *Test conditions for machining centres Part 2: Geometric tests for machines with vertical spindle or universal heads with vertical primary rotary axis Z-axis*. ISO, Geneva, Switzerland.

ISO 230-2:2006. *Test code for machine tools Part 2: Determination of accuracy and repeatability of positioning numerically controlled axes*. ISO, Geneva, Switzerland.

Mekid, S., ed. 2009. *Introduction to Precision Machine Design and Error Assessment*. CRC Press.

Muelaner, J.; Cai, B.; and Maropoulos, P. 2010. Large volume metrology instrument selection and measurability analysis. In Huang, G.; Mak, K.; and Maropoulos, P., eds., *Proceedings of the 6th CIRP-Sponsored International Conference on Digital Enterprise Technology*, volume 66 of *Advances in Intelligent and Soft Computing*. Springer Berlin / Heidelberg. 1027–1041.

Nau, D.; Au, T.-C.; Ilghami, O.; Kuter, U.; Murdock, J. W.; Wu, D.; and Yaman, F. 2003. SHOP2 : An HTN Planning System. *JAIR* 20:379–404.

Parkinson, S.; Longstaff, A. P.; Crampton, A.; Andrew, P.; Fletcher, S.; Allen, G.; and Myers, A. 2011. Representing the Process of Machine Tool Calibration in First-order Logic. In *Automation and Computing*.

Ramesh, R.; Mannan, M.; and Poo, A. 2000. Error compensation in machine tools a review. *International Journal of Machine Tools and Manufacture* 40(9):1235–1256.

Ruml, W.; Do, M. B.; Zhou, R.; and Fromhertz, M. P. J. 2011. On-line Planning and Scheduling : An Application to Controlling Modular Printers. *JAIR* 40:415–468.