



University of HUDDERSFIELD

University of Huddersfield Repository

Parkinson, Simon, Vallati, Mauro, Chrupa, Lukáš, Longstaff, Andrew P. and Fletcher, Simon

Planning Machine Activity Between Manufacturing Operations: Maintaining Accuracy While Reducing Energy Consumption

Original Citation

Parkinson, Simon, Vallati, Mauro, Chrupa, Lukáš, Longstaff, Andrew P. and Fletcher, Simon (2016) Planning Machine Activity Between Manufacturing Operations: Maintaining Accuracy While Reducing Energy Consumption. In: 10th Scheduling and Planning Applications woRKshop (SPARK), 14th June 2016, Kings College, London. (Unpublished)

This version is available at <http://eprints.hud.ac.uk/id/eprint/28662/>

The University Repository is a digital collection of the research output of the University, available on Open Access. Copyright and Moral Rights for the items on this site are retained by the individual author and/or other copyright owners. Users may access full items free of charge; copies of full text items generally can be reproduced, displayed or performed and given to third parties in any format or medium for personal research or study, educational or not-for-profit purposes without prior permission or charge, provided:

- The authors, title and full bibliographic details is credited in any copy;
- A hyperlink and/or URL is included for the original metadata page; and
- The content is not changed in any way.

For more information, including our policy and submission procedure, please contact the Repository Team at: E.mailbox@hud.ac.uk.

<http://eprints.hud.ac.uk/>

Planning Machine Activity Between Manufacturing Operations: Maintaining Accuracy While Reducing Energy Consumption

**Simon Parkinson, Mauro Vallati
Lukas Chrpa**
Department of Informatics
School of Computing and Engineering
University of Huddersfield, UK
s.parkinson@hud.ac.uk

Andrew Longstaff, Simon Fletcher
Centre for Precision Technologies
School of Computing and Engineering
University of Huddersfield, UK

Abstract

There has recently been an increased emphasis on reducing energy consumption in manufacturing. This is largely because of fluctuations in energy costs causing uncertainty. The increased competition between manufacturers means that even a slight change in energy consumption can have implications on their profit margin or competitiveness of quote. Furthermore, there is a drive from policy-makers to audit the environmental impact of manufactured goods from cradle-to-grave. The understanding, and potential reduction of machine tool energy consumption has therefore received significant interest as they require large amounts of energy to perform either subtractive or additive manufacturing tasks.

One area that has received relatively little interest, yet could harness great potential, is reducing energy consumption by optimally planning machine activities while the machine is not in operation. The intuitive option is to turn off all non-essential energy-consuming processes. However, manufacturing processes such as milling often release large amounts of heat into the machine's structure causing deformation, which results in deviation of the machine tool's actual cutting position from that which was commanded, a phenomenon known as thermal deformation. A rapid change in temperature can increase the deformation, which can deteriorate the machine's manufacturing capability, potentially producing scrap parts with the associated commercial and environmental repercussions. It is therefore necessary to consider the relationship between energy consumption, thermal deformation, machining accuracy and time, when planning the machine's activity when idle, or about to resume machining.

In this paper, we investigate the exploitability of automated planning techniques for planning machine activities between subtractive manufacturing operations, while being sufficiently broad to be extended to additive processes. The aim is to reduce energy consumption but maintain machine accuracy. Specifically, a novel domain model is presented where the machine's energy consumption, thermal stability, and their relationship to the overall machine's accuracy is encoded. Experimental analysis then demonstrates the effectiveness of the proposed approach using a case study which considers real-world data.

Introduction

Machine tools are complex mechanronic system used in both subtractive and additive manufacturing. Much of their

Copyright © 2016, All rights reserved.




Figure 1: Example C-Frame three-axis machine tool

performance comes from their mechanical rigidity. For example, Figure 1 illustrates the structure of a three-axis machine tool. Machine tools come in a large variety of sizes and configurations, but a common feature is their ability to position their tool in a three-dimensional space relative to the workpiece either to remove (cut, grind, etc.) or add material. Accuracy is often a primary commercial driver in the advancement of machine tools for precision, high-value manufacturing to micrometre-level tolerances. However, maintaining such high levels of accuracy requires strict control of the many factors which can cause a change in accuracy. For example, the effect of temperature change on the machine's structure can have a dramatic impact on the accuracy of the tool. Energy efficiency is also becoming an increasingly important factor in machine tool development both to reduce manufacturing costs (Draganescu *et al.* 2003; Diaz *et al.* 2011), as well as reducing environmental impact (Diaz *et al.* 2010). However, the relationship between the improvement in energy efficiency and possible reduction in machine accuracy resulting from rapid temperature change is less well explored. This is surprising considering the amount of heat generated from electrical devices and mechanical subsystems during the machining process.

The use of machine tools has been identified as the largest consumer of energy during the manufacturing of parts. It has been established that machine tools use 63% of the total energy required to manufacture a part (Hesselbach and Herrmann 2011). Additionally, energy consumption occupies over 20% of the operating costs of machine tools per year, in excess of £10,000. While it is difficult to state

how much of the 20% is consumed between manufacturing operations, it is likely that the machine will be stationary for many periods during the working-day as new parts are loaded, etc. Many researchers have investigated the potential of reducing energy consumption during the manufacturing process itself (Vijayaraghavan and Dornfeld 2010; Liu *et al.* 2014; 2015). For example, reducing energy usage during milling (Diaz *et al.* 2011). These works have largely been motivated by the fact that large forces are required to cut material, and any reduction at this stage can therefore be significant. However, one area that has received less attention is the consumption of energy between manufacturing operations, when the machine is not cutting and therefore is nominally idle. In the first instance it may appear that if the machine is idle it will be consuming no energy. However, it is often the case that many electrical components of a machine tool will continue to use energy. Furthermore, once the machine is required to operate once again, an energy-intensive warm-up cycle is often required to bring the subsystem (e.g spindle motor) into a suitable (stable) state for actual machining.

Such warm-up cycles are often required since the heat generated from the machine components during manufacturing will transfer to the machine tool's structure and cause deformation. This thermal deformation is, in the simplest case, a first-order response to the temperature step input. Heating the subsystem prior to manufacturing means that much of the deformation will take place before cutting begins, helping to reduce in-process change and increasing the accuracy of the component. A warm-up cycle is usually energy intensive, but will only be necessary should the heat-generating subsystem and surrounding structure decrease below an identified temperature. This creates an interesting possibility where keeping the machine subsystems active, at a reduced level, whilst not manufacturing can generate sufficient heat to maintain the thermal stability of the machine tool's structure and remove the need for a warm-up cycle, thus reducing the overall energy consumption. For example, Figure 2 illustrates, through Finite Element Analysis (FEA), the deformation of the machine tool's structure resulting from the release of heat generated by the spindle motor and friction in the moving mechanical elements. In the diagram, the nominal tool position and orientation are shown superimposed on the actual location; the difference between the two, caused by temperature effects, leads to a displacement at the tool tip, which is known as the thermal error. If this deformation were to take place during the manufacturing process then the resultant manufactured component would display the results of this error, leading to a requirement for rework or even scrapped parts. However, if this deformation were to occur before the manufacturing process, then the thermal error of the machine can remain stable during the process, and therefore the accuracy of the produced part is largely unaffected.

Energy consumption information for many machine tool electrical subsystems is widely available, but that from mechanical interaction (friction) is often less well defined. However, in both cases the amount of heat released into the machine's structure and its affect on machine accuracy needs




Figure 2: Deformation of the machine tool's structure due to heat generated by the spindle motor

to be established. This can be acquired by recording the temperature of the machine tool's structure while monitoring the deviation of the machine tool's cutting point. During analysis, each subsystem will typically be run at different speeds to establish the relationship between the different levels of energy consumption and heat generation, and also the relationship between heat generation at the subsystem's location and the effect on machine accuracy. Once all the data has been acquired, FEA can be used to computationally model the relationship between heat generation and deformation of the machine tool (Mian *et al.* 2011; 2013). This model can then be used to derive a series of coefficients that describe the generation of heat with increased energy consumption, and the change in machine tool accuracy from the resulting different thermal gradients.

The number of electrical subsystems, the different operational levels, the current state of the machine tool, and the required initial state of the next manufacturing operation make it challenging to consider all possible options and minimise energy consumption whilst maintaining a desired level of accuracy. This creates an interesting and novel possibility to utilise Automated Planning to automate the process, removing the requirement for expert knowledge, minimise energy consumption, and maintain the required level of accuracy. While the exploitation of planning techniques for planning machine activities between manufacturing operations has never been investigated, previous works demonstrated the potential of using automated planning for optimising different aspects of using machine tools. For example, the non-productive time (downtime) of a machine tool during calibration has been reduced through automatically constructing calibrations plans, reducing reliance on expert knowledge (Parkinson *et al.* 2012a; 2012b). Further work of encoding mechanisms to calculate measurement uncertainty (Parkinson *et al.* 2014b) created the potential to perform multi-objective optimisation (Parkinson *et al.* 2014a).

This paper is organised as follows: first, the importance of planning for activity between manufacturing operations (named *interval activity* herein) is described and motivated. Second, we provide a discussion on the importance of planning interval activity, and a domain model is provided, encoded using the Planning Domain Definition Language

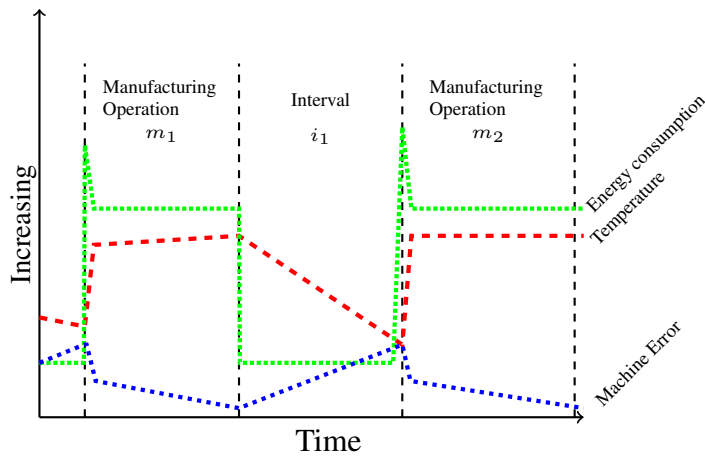


Figure 3: Illustrating how the machine tool energy consumption, structural temperature, and accuracy is changing during manufacturing and interval periods.

(PDDL) (Fox and Long 2003). Then, the effectiveness of automated planning is demonstrated using a real-world case study. Finally, conclusions are given.

Importance of Interval Activity

In this section, the importance of considering interval activity is motivated. In particular, the relationship that is of interest in this paper is that between energy consumption, generation of temperature profile, and the affect on machining accuracy. Prior knowledge of this relationship creates the potential to optimise machine tool use between manufacturing operations. For example, in some situations, it may be advantageous to keep the electronic components in use to maintain energy consumption, generate heat, and thus maintain machine accuracy.

Figure 3 provides a graphical illustration of two manufacturing operations with an interval between. The figure illustrates the relationship between increasing energy consumption (green), heat generation (red), and increasing machine error (blue) through a simplified representation. Note that although the figure is for illustration purposes, the data is a realistic, if simplified, representation of what occurs. In the figure, it can first be seen that energy consumption is at its lowest when the machine is idle, and its highest when a new manufacturing job is started. This is because a dedicated warm-up cycle is required to stabilise the machine's structure and avoid thermal change during manufacturing. It is then noticeable in the figure that as the energy consumption increases, so does the temperature of the machine's structure. The final relationship presented is that the error of the machine tool increases to a steady-state value and maintained when the temperature is stable. In practice, the number of different operations that occur during machining mean that the energy profile, and resulting temperature and error trends, will display somewhat more complex behaviour. In the remainder of this section, a more detailed analysis of each element (energy consumption, temperature and accuracy) is presented.




Figure 4: Heat generated by spindle motor during a two-hour heating and cooling cycle. Spindle bottom, top and motor indicate the normalised temperature for three surface temperature sensors mounted around the spindle. Spindle RPM reports the normalised spindle speed in RPM (0 to 9000).

The relationship between energy consumption, the generation of heat and its dissipation into the machines structure is different for every subsystem: a high-speed spindle motor uses significantly more power than a linear axis servo motor. For example, consider manufacturing an aluminium housing with the dimensions of 150mm \times 50mm \times 20mm (Heidenhain 2010). The total energy needed for the machine tool to produce the part is 20.4kW. A total of a 4.8kW for the machine tool spindle, and 0.5kW for the three axes' feed drives. Other electrical subsystems (e.g controller, coolant pump, etc.) make up the remainder. As both these components have different levels of energy consumption, they generate different amounts of heat. Each component will have different modes of operation. For example, a common spindle motor might be capable of speeds in excess of 9,000 revolutions per minutes (RPM). Figure 4 demonstrates the heat generated as the spindle speed increases on a three-axis machine tool. The figure shows the normalised spindle speed in RPM (0 to 9000), and normalised temperature for three surface temperature sensors mounted around the spindle. There are: (1) spindle bottom (21.8°C to 27.4°C), (2) spindle motor (21.6°C to 26.2°C), and (3) spindle motor (21.7°C to 33.1°C). The graph illustrates that when the spindle is used at its higher speed, the temperature of the machine tool's structure surrounding the spindle increases rapidly in temperature. Once high speed usage has finished, it can be seen that the structure of the machine tool begins to reduce in temperature.

The next relationship of interest is that of changing machine tool temperature and its effect on structural deformation of the machine tool. The heat generated by machine subsystems transfers into the machine tool's structure causing distortion. The severity of the effect of changing temperature is dependent on the material from which it is constructed. For example, steel has a high coefficient of thermal expansion ($\sim 12\mu\text{m per }^\circ\text{C}$) compared to carbon fibre ($\sim 2\mu\text{m per }^\circ\text{C}$), though much less than aluminium ($\sim 22\mu\text{m per }^\circ\text{C}$).




Figure 5: Error of a three axis machine tool generated during a two hour spindle heating and cooling cycle.

Considering the previous example of a spindle motor, Figure 5 illustrates the effect of changing temperature on the machine’s structure. In this figure the spindle of a three-axis machine tool was running at 9,000 RPM (70% utilisation) for 120 minutes, and then left to cool for a further 120 minutes. From this experiment, it is noticeable that error for each of the three axes is changing throughout the heating cycle in the first 120 minutes, and then once the spindle is disabled, the errors continue to increase as the heat is still transferred into the machine’s structure.

The examples discussed in this section demonstrate the importance of planning for machine activity between manufacturing operations. However, planning in this context is not a trivial task as any action can impact on machine accuracy and energy usage, both of which can have significant financial implications. Currently it is up to the machine operator to make the correct decision in an ad-hoc manner where they determine machine activity by knowing future manufacturing operations, as well as the energy saving policies of their manager. However, this planning for the machine operator is complicated by the large number of different machine activity actions that can be performed and their potential implications on machine accuracy and energy consumption. For example, each axis and spindle can be moved at different speeds sequentially or concurrently for different periods of time. Moving a single linear axis will transfer heat in the machine’s structure surrounding the axis and would result in thermal distortion from that location, whereas moving all three axes simultaneously would transfer heat into more of the machine’s structure and potentially result in more symmetrical expansion.

Domain Modelling

In this section, a PDDL model is developed and discussed to describe the domain of interval planning. In the presented model, the two following equations are used to determine energy consumption as well as machine accuracy. These equations require machine-specific data acquired through performing an error mapping and energy monitoring audit.

$$total\ error = total\ error + duration \times (effect\ on\ error \times energy\ consumption) \quad (1)$$

$$total\ energy = total\ energy + (duration \times energy\ consumption) \quad (2)$$

Equation 1 is used for updating the error fluent by a quantity of time in minutes, multiplied by the the effect on error in micrometres of deviation per minute of energy consumption. Here, there is a different effect on error value for each different mode of operation. Equation 2 updates the energy consumption fluent by the same duration (in minutes) multiplied by the a fluent storing the energy use of a particular component when being used in a predefined mode of operation.

The use of predefined modes has been adopted to reduce the size of the domain model, in terms of number of operators, and make it easier to be handled by state-of-the-art planning engines. Many machine subsystems, such as the spindle motor, can be run at any speed between stationary and their maximum RPM. This continuous behaviour could be encoded in PDDL+ (Fox and Long 2006); however, this would dramatically increase domain complexity as the number of heat-generating machine components increases. In addition, the number of planning systems able to handle PDDL+ is limited (see, e.g. (Coles and Coles 2014; Della Penna *et al.* 2009)), especially when compared with those capable of handling different versions of PDDL. Even more restricted is the number of solvers able to support the entire function set of PDDL+. Therefore, for the preliminary work undertaken to determine the feasibility of using automated planning in this domain, PDDL2.2 (S. Edelkamp and J. Hoffmann 2004) –an extension of PDDL2.1 (Fox and Long 2003)– is used. The International Planning Competition¹ has resulted in the existence of a significant number of planners able to solve PDDL2.2 planning problems.

Initial and Goal State

The initial state specifies energy consumption and effect on machine accuracy for each predefined level of operation through the use of numeric fluents. For example, `energy_idle ?c` and `error ?c` represent the energy consumption and the effect on accuracy for a component `?c`. In addition `time_unit` fluent is introduced to specify a predetermined duration of an action that should occur to bring about a change in accuracy and energy consumption. The `total_error` and `total_energy` fluents are used in the initial state to encode information regarding the machine’s current state after finishing manufacturing. In addition, timed initial literals are also used to encode the duration of the interval. Using timed initial literals restricts the makespan to the duration of the interval, overcoming some planner’s inability to handle concurrency in durative actions.

The goal state makes use of four optional numeric conditions to impose a tolerance window on the total error and

¹<http://www.icaps-conference.org/index.php/Main/Competitions>

```

(:durative-action normal
  :parameters (?c - component)
  :duration(= ?duration (time_unit ?c))
  :condition
  (and
    (over all (in_interval))
    (at start (not(in_use ?c)))
  )
  :effect
  (and
    (at start (in_use ?c))
    (at end (not(in_use ?c)))
    (at end (increase(total_error)
      (*(time_unit ?c)
        *(error_normal ?c)
        /(energy_normal ?c)60))))
    (at end (increase(total_energy)
      (*(time_unit ?c)
        /(energy_normal ?c)60))))
  )
)
)

```

Figure 6: Durative action representing the fact that the machine component $?c$ is planned to remain in a normal state of operation for a time unit.

energy consumption. The tolerance window for the total error creates the possibility to specify the manufacturing requirements of the next job and to ensure the interval activity correctly prepares the machine’s state. For example, using $(< (total_error) 10)$ and $(< (total_energy) 4)$ in the goal state would ensure that the error must be less than $10\mu\text{m}$ (at the end of the modelled interval, thus ready for the next job) and the overall energy usage must be less than 4Wm^{-1} .

Operators

In our domain there are machine component objects, and four operators representing different levels of operation: off, idle, normal and high. Each operator is similar apart from the equation to update both error and energy fluents. Figure 6 details the normal durative action where the machine error and energy usage are adjusted based on a normal mode of operation. The action will execute for the $time_unit$, while the $in_interval$ predicate is true. The full PDDL source is available from the authors on request.

Plan Metric

In this paper the following three different metrics are used:

1. `(:metric minimize (total_error))`
2. `(:metric minimize (total_energy))`
3. `(:metric minimize (/ (+ (total_error) (total_energy)) 2))`

The first two aim to minimise the values held in the total error and total energy fluent, whereas the third metric is

Electrical Item	Off Wm^{-1} , μm	Idle Wm^{-1} , μm	Normal Wm^{-1} , μm	High Wm^{-1} , μm	Multiplier
X Servo	0, 1	0.1, 0.1	0.7, 0.5	2.8, 2	1.2E-005
Y Servo	0, 1	0.1, 0.1	0.7, 0.5	2.8, 2	1.2E-005
Z Servo	0, 3	0.2, 0.3	1.0, 1.5	4.2, 6	2.4E-005
Spindle Motor	0, 6	1.3, 0.6	6.5, 3.1	26, 12	7.7E-006

Table 1: Case study data demonstrating the different energy consumption (in Wm^{-1}) and the effect on error (in μm per minute)

used to minimise the arithmetic mean of both. This creates the potential to perform multi-objective optimisation where both error and total energy consumption are minimised for a given weighting.

Experimental Analysis

In this section, a case study is provided where interval planning is performed for a single machine tool when considering different interval scenarios. The data presented in Table 1 details the energy consumption of the machine tool, as well as the relationship between energy consumption and machine error. These values have been extracted from a similar machine tool as presented in earlier sections of this paper. As interval duration is in minutes, the data presented in Table 1 has been converted into time units. These are Watts per minute (Wm^{-1}) and the positional error in micrometres per Wm^{-1} . This is calculated using a multiplier derived from dividing the deviation in micrometres per minute by Watts per minute ($\mu\text{mWm}^{-1}\text{m}^{-1}$). For example, to calculate the micron error resulting from 60 minute high use of the spindle motor would be $12\mu\text{m}$ by using Equation 1 where $duration = 12$, $energy\ consumption = 26,000$, and $effect\ on\ error = 7.7E - 006$.

The machine-specific data presented in Table 1 is now used in the creation of several PDDL problem files to simulate the following interval scenarios. First, problem definitions are created with a duration of 30, 60, and 120 minutes. Following this, three variations of each problem are created with three different requirements on machine error. These are: tight ($<20\mu\text{m}$), medium ($<50\mu\text{m}$), and large ($>50\mu\text{m}$). These requirements are synthetically generated; however, they do provide an adequate description of different energy and machine tool accuracy requirements in a manufacturing environment. Considering the combination of each of these scenarios results in the creation of 9 different problem instances. In addition, each problem instance will be solved using each of the three metrics stated in the domain modelling section, resulting in a total of 27 different PDDL problem definition files. LPG-td (Gerevini *et al.* 2006) is used to find the best solutions (in terms of the specified metric) to the problem definitions within a 5 minute time-frame.

Instance	Metric: Error		Metric: Energy		Metric: $\overline{Er} + \overline{En}$	
	En(Wm ⁻¹)	Er(μm)	En(Wm ⁻¹)	Er(μm)	En(Wm ⁻¹)	Er(μm)
T-30	138	10	138	10	138	10
T-60	260	16	220	19	232	20
T-120	520	19	380	20	410	20
M-30	0	21	0	21	0	21
M-60	800	26	374	44	670	36
M-120	1351	45	984	47	1263	46
L-30	-	-	-	-	-	-
L-60	2578	52	-	-	-	-
L-120	3951	58	1641	61	2584	60

Table 2: Experimental results detailing both error and error values for each of the nine scenarios when using three different metrics. Entries marked with a dash (-) were not solved within the 5 minute cut-off time. Problem instances are in the format of a character to represent the scenario and the interval duration in minutes. The characters are: T= tight, M = medium, L = large.

LPG-td has been used in “anytime” configuration; it keeps increasing the quality of plan, for a given problem instance, until the available CPU-time is over.

Table 2 provides the error (Er) and energy (En) values for each problem instance and the use of the three metrics. From the Table it is evident that the majority of the problem instances were solved within the 5 minute cut-off time. After examination, it is noticeable the problem instances requiring a large error (L) are not solved within the allowed time. It is worthy reminding that, while T and M benchmarks require that the initial error value is lower than a given value, in L benchmarks the accuracy requirement is to have a value higher than a given threshold. Therefore, plans which are suitable for T and M, are not valid for the L scenario. The fact that some instances are not solvable is not because the planning problem requires more time to identify a solution, rather there is no suitable sequence of actions capable of taking the error beyond that specified in the goal state during the allocated interval duration. However, this should be seen as a useful piece of information rather than an issue: it would not be detrimental for a manufacturer to manufacture a part on a machine with a smaller error than required to satisfy the tolerance constraints of the part.

From analysing the results presented in Table 2, it can also be seen that for both the tight and medium 30 minute interval problem instances, optimising for all three metrics results in the identification of the same plan. In addition, in some instances the energy consumption is 0. Interestingly, this is because the planner is able to identify a plan where the machine is switched-off and the slow deterioration in accuracy over the 30 minute period does not take the accuracy beyond the value set in the initial state ($<20\mu\text{m}$ and $<50\mu\text{m}$). In addition, it is also possible to identify that optimising for a single metric is often at the expense of the other. For example, in T-60 it can be seen how the error is reduced to $16\mu\text{m}$ by using 260Wm^{-1} when optimising for error, whereas when optimising for energy, the error increases to $19\mu\text{m}$ and the energy usage decreases to 220Wm^{-1} . Optimising for the arithmetic mean of both metrics also results in a plan where both metrics are at their lowest.

This experimental analysis has demonstrated that the technique is useful for stabilising the machine error whilst reducing estimated energy consumption. It has also demonstrated that planning for large machine error is unnecessary as it will result in high energy consumption over short durations to generate large amounts of heat and cause structural deformation to reduce accuracy. Conversely, when planning for tight tolerances, the plan will contain actions with low energy consumption which will result in gradual heat generation, leading to thermal stabilisation and improved accuracy. For example, in the following plan excerpt it can be seen how, over a 30 minute period, the spindle is initially turned off and then switched back on and left to idle to maintain accuracy.

```
0: (OFF SPINDLE_MOTOR) [10.0000]
10: (IDLE SPINDLE_MOTOR) [10.0000]
20: (IDLE SPINDLE_MOTOR) [10.0000]
```

Conclusion

This paper presents the exploitative use of automated planning to plan for machine tool activity between manufacturing operations. This is a novel application with potential to aid machine tool operators prepare their machine for the next manufacturing task considering overall energy consumption and the effect on manufacturing accuracy. Research has previously been undertaken in the area of reducing energy consumption and improving the error of machine tools during manufacturing. However, to the best of the authors’ knowledge, little work has considered the intervals between manufacturing.

The paper provides a discussion detailing that subsystems of a machine tool (in particular servo motors) consume a large quantity of energy and that this energy results in the generation of heat. It was then discussed how this heat transfers through the machine tool’s structure causing thermal deformation, resulting in positioning error of the machine tool’s cutting point relative to the workpiece. This error then transfers to the manufactured component and can

result in the production of out-of-tolerance parts. It was discussed how changing the machine's activity during non-manufacturing intervals can help to stabilise the machine tool's structure, resulting in a reduction in error. An experimental PDDL domain model was then developed to enable planning for manufacturing intervals using available energy and error information.

Experimental analysis was then performed using the developed domain model and nine different problem instances. Each instance relates to a different combination of accuracy requirements and interval duration. The exploratory experimental analysis demonstrates good potential as accuracy, energy, and the arithmetic mean of both are minimised below the specified accuracy limit. It has also been identified that automated planning is capable of providing a viable mechanism for aiding manufacturing, albeit with a simplified domain model. The developed domain requires further development and testing to make it more accurately represent interval planning. One limitation of the presented domain is that the relationship between temperature and accuracy is non-linear, whereas for this initial proof-of-concept the relationship has been discretized into linear rates-of-change. Future work will include investigating alternative approaches to solve problems of a larger size to gain better results. For example, the use of PDDL+ and mixed integer programming will be considered, as well as undertaking further domain modelling work with the view of experimenting with different planning algorithms.

References

- Amanda Jane Coles and Andrew Ian Coles. PDDL+ planning with events and linear processes. In *Proceedings of the Twenty-Fourth International Conference on Automated Planning and Scheduling, ICAPS, 2014*.
- Giuseppe Della Penna, Daniele Magazzeni, Fabio Mercorio, and Benedetto Intrigila. UPMurphi: a tool for universal planning on PDDL+ problems. In *Nineteenth International Conference on Automated Planning and Scheduling, 2009*.
- Nancy Diaz, Moneer Helu, Stephen Jayanathan, Yifen Chen, Arpad Horvath, and David Dornfeld. Environmental analysis of milling machine tool use in various manufacturing environments. In *Sustainable systems and technology (ISSST), 2010 IEEE international symposium on*, pages 1–6. IEEE, 2010.
- Nancy Diaz, Elena Redelsheimer, and David Dornfeld. Energy consumption characterization and reduction strategies for milling machine tool use. In *Glocalized Solutions for Sustainability in Manufacturing*, pages 263–267. Springer, 2011.
- F Draganescu, M Gheorghe, and CV Doicin. Models of machine tool efficiency and specific consumed energy. *Journal of Materials Processing Technology*, 141(1):9–15, 2003.
- Maria Fox and Derek Long. PDDL2.1: an extension to PDDL for expressing temporal planning domains. *Journal of Artificial Intelligence Research*, 20:61–124, 2003.
- Maria Fox and Derek Long. Modelling mixed discrete-continuous domains for planning. *Journal of Artificial Intelligence Research*, 27:235–297, 2006.
- Alfonso Gerevini, Alessandro Saetti, and Ivan Serina. An approach to temporal planning and scheduling in domains with predictable exogenous events. *Journal of Artificial Intelligence Research*, 25:187–231, 2006.
- Heidenhain. Aspects of energy efficiency in machine tools. Technical report, 2010.
- Jürgen Hesselbach and C Herrmann. Globalized solutions for sustainability in manufacturing. In *Proceedings of the 18th CIRP International Conference on Life Cycle Engineering*, pages 2071–1050, 2011.
- Ying Liu, Haibo Dong, Niels Lohse, Sanja Petrovic, and Nabil Gindy. An investigation into minimising total energy consumption and total weighted tardiness in job shops. *Journal of Cleaner Production*, 65:87 – 96, 2014.
- Ying Liu, Haibo Dong, Niels Lohse, and Sanja Petrovic. Reducing environmental impact of production during a rolling blackout policy—a multi-objective schedule optimisation approach. *Journal of Cleaner Production*, 102:418–427, 2015.
- Naeem S Mian, Simon Fletcher, Andrew P Longstaff, and Alan Myers. Efficient thermal error prediction in a machine tool using finite element analysis. *Measurement Science and Technology*, 22(8):085107, 2011.
- Naeem S Mian, Simon Fletcher, Andrew P Longstaff, and Alan Myers. Efficient estimation by FEA of machine tool distortion due to environmental temperature perturbations. *Precision engineering*, 37(2):372–379, 2013.
- Simon Parkinson, Andrew Longstaff, Andrew Crampton, and Peter Gregory. The application of automated planning to machine tool calibration. In *Proceedings of the twenty-second international conference on automated planning and scheduling (ICAPS)*, 2012.
- Simon Parkinson, Andrew P Longstaff, Simon Fletcher, Andrew Crampton, and Peter Gregory. Automatic planning for machine tool calibration: A case study. *Expert Systems with Applications*, 39(13):11367–11377, 2012.
- Simon Parkinson, Andrew P Longstaff, Andrew Crampton, and Peter Gregory. Automated planning for multi-objective machine tool calibration: Optimising makespan and measurement uncertainty. *Proceedings of the twenty-fourth international conference on automated planning and scheduling (ICAPS)*, pages 421–429, 2014.
- Simon Parkinson, Andrew P Longstaff, and Simon Fletcher. Automated planning to minimise uncertainty of machine tool calibration. *Engineering Applications of Artificial Intelligence*, 30:63–72, 2014.
- S. Edelkamp and J. Hoffmann. PDDL2.2: The Language for the Classical Part of the 4th International Planning Competition. Technical Report 195, Albert-Ludwigs-Universität Freiburg, Institut für Informatik, 2004.
- Athulan Vijayaraghavan and David Dornfeld. Automated energy monitoring of machine tools. *CIRP Annals-Manufacturing Technology*, 59(1):21–24, 2010.