

Throughput optimization for micro-factories subject to failures

Anne Benoit¹, Alexandru Dobrila^{2,*}, Jean-Marc Nicod² and Laurent Philippe²

¹ ENS Lyon, Université de Lyon, LIP laboratory (ENS, CNRS, INRIA, UCBL)

² Laboratoire d'Informatique de Franche-Comté, Université de Franche-Comté

* Contact author: dobrila@lifc.univ-fcomte.fr

LIP Research Report RR-2009-02

Abstract

In this paper, we study the problem of optimizing the throughput for micro-factories subject to failures. The challenge consists in mapping several tasks onto a set of machines. One originality of the approach is the failure model for such applications, in which tasks are subject to failures rather than machines. If there is exactly one task per machine in the mapping, then we prove that the optimal solution can be computed in polynomial time. However, the problem becomes NP-hard if several tasks can be assigned to the same machine. Several polynomial time heuristics are presented for the most realistic specialized setting, in which tasks of a same type can be mapped onto a same machine. Experimental results show that the best heuristics obtain a good throughput, far from the throughput obtained with a random mapping. Moreover, we obtain a throughput close to the optimal solution in the particular cases on which the optimal throughput can be computed.

1 Introduction

It is usual to say that distributed systems are a good support for fault-tolerance but that their correct management needs to take faults into account. Standard distributed systems mainly focus on processor dependent faults. It means that the common assumption used to model fault tolerance in distributed systems is to assume that faults are generated by the execution platform and thus that the fault model used must be linked to the processor or to the resource that performs a task. In this case, a stochastic fault model defining the fault probability is usually attached to the processor. This model fits distributed computing environments such as parallel

platforms where failures come from the processors or the nodes of the platform.

If we look however at a more general definition of a distributed system, we can state that this model does not always fit. In some distributed platforms the fault model may be attached to the task rather than to the processor or node. For example, in production systems, a task may be complex to perform due to some hard manipulation and that impacts its success ratio. If the same robot is able to perform different tasks, it will generate less faults on simple tasks than on difficult ones.

In this paper we are interested in studying the impact of a fault model linked to task. The application context is more a production system than a distributed computing system. Our specific use case is a micro-factory but the results presented in this paper are more generally applicable to distributed production systems or distributed systems where the fault probability is attached to tasks rather than resources.

Micro-factories are production units designed to produce pieces composed of micro-metric elements [8]. Today's micro-factories are composed of micro-robots able to carry out basic operations through elementary actuators as piezo-electric beams (e.g. for gripping), stick-slip systems, etc. As these robots are usually teleoperated by a human operator, only simple tasks can be done. To perform more complex operations and to improve their efficiency, micro-factories need to be automated and robots grouped in cells. Then cells will be put together and they will cooperate to produce complex assembled pieces, as we do it for macroscopic productions. Due to the pieces, actuators and cells size, it is however impossible for human operators to directly interfere with the physical system. So it needs a highly automated command and the complexity of this command makes it mandatory to develop a distributed system to support this control. So, the cell group results in a distributed system that is very similar to a distributed computing platform. Note that

the physical constraints that act on the components and on the environment are different compared to the macro world. As these constraints are not easily controlled, it is difficult to guarantee the command, leading to a poor reliability and to the need to take faults into account as in computing systems.

The main issue for fault tolerant systems [6] is to overcome the failure of a node, a machine or a processor. To deal with those faulty machines the most common method in distributed systems is to replicate [1] the data. Those models assume that the failure is attached to a machine. Indeed when a replicated task is executed on several machines, the probability to get one product as a result is highly increased. Once all the replicated jobs are done, a vote algorithm [7] is often used to decide which result is the right one. In real time systems, another model called Window-Constrained [9] model can be used. In this model one considers that, for y messages, only x ($x \leq y$) of them will reach their destination. The y value is called the Window. The losses are not considered as a failure but as a guarantee: for a given network a Window-Constrained Scheduling [10] can guarantee that no more than x messages will be lost for every y sent messages. In this paper, the Window-Constrained based failure model is adapted to a distributed system, the micro-factory. So the issue is to guarantee the output of a given number of products. With failures attached to tasks, we can compute the number of products needed as input of the system and guarantee the output for the desired number of products.

The paper is organized as follows. Section 2 gives the characteristics and a more formal presentation of the context of micro-factories and the failure model. Section 3 presents the optimization problems tackled in the paper. The complexity study and results are given in Section 4. Heuristics to solve the problem are proposed in Section 5 and simulation results for these heuristics are given and commented in Section 6. Finally, we conclude in Section 7.

2 Framework

We outline in this section the characteristics of the applicative framework and target platform. Finally, we describe and motivate the failure model that we use in this work.

2.1 Applicative framework

We consider a set \mathcal{N} of n tasks: $\mathcal{N} = \{T_1, T_2, \dots, T_n\}$. Each task T_i ($1 \leq i \leq n$) is applied

successively on a set of products, numbered from 1 to B . We wish to produce B_o products as an output, and the total number of products B being processed may depend on the allocation ($B \geq B_o$, losses due to failure as explained later in Section 2.3). Note that all products are identical. When the context is not ambiguous, we may also design task T_i by i for clarity, as for instance in the figures.

A type is associated to each task as the same operation may be applied several times to the same product. Thus, we have a set \mathcal{T} of p task types with $n \geq p$ and a function $t : [1..n] \rightarrow \mathcal{T}$ which returns the type of a task: $t(i)$ is the type of task T_i , for $1 \leq i \leq n$.

The application is a directed acyclic graph (DAG) in which the vertices are tasks, and edges represent dependencies between tasks. An example of application with $n = 5$ tasks is represented on Figure 1. In the top branch of the DAG, we need to finish the processing of task T_1 on one product before proceeding to task T_2 . The join to task T_4 corresponds to the merge of two products, which produces a new unit of product. Typically one product from each predecessor in the graph is required to process with the joining task. We do not consider forks, and rather restrict to in-trees. Indeed, micro-factories assemble objects (joins), but a task cannot split a micro-component: these are physical objects that cannot be replicated to be processed by more than one task.




Figure 1: Example of application.

2.2 Target platform

The platform consists in a set \mathcal{M} of m machines: $\mathcal{M} = \{M_1, M_2, \dots, M_m\}$. All machines can be interconnected: the platform graph is a complete graph. A machine handles some of the tasks at a given speed: machine M_u can perform the task T_i onto one product in a time $w_{i,u}$. We also consider that tasks of the same type have the same execution time on a given machine, since they correspond to the same action to be performed on the products. Thus, we have:

$$\forall i, i' \in [1, n], \forall u \in [1, m], t(i) = t(i') \Rightarrow w_{i,u} = w_{i',u}$$

We neglect the communication time required to transfer a product from one machine to another. If a communication may not be negligible, we can always model it

as a particular task with a dedicated machine (the machine responsible of the transfer of the product).

We are interested in producing the desired number of products rather than producing a particular instance of a product. So we consider that products are not identified: two products, on which the same sequence of tasks has been done, are exactly similar and we can use one or the other indifferently for further operations.

2.3 Failure model

An additional characteristic of our framework is that tasks are subject to failure. It may happen that a product is lost or damaged while a task is being executed on this product. For instance electrostatic strength may be accumulated on the actuator, and thus the piece will be pushed away rather than caught. Indeed, we work at a scale such that these electrostatic strengths are stronger than gravity.

Due to our application setting, we deal only with transient failures, as defined in [4]. The tasks are failing for some of the products, but we do not consider a permanent failure of the machine responsible of the task, as this would lead to a failure for all the remaining products to be processed and the inability to finish them.

One classical technique used to deal with failures is replication [1]. However, while replication is very useful for hardware failures of machines, we cannot use it in our framework since the products are not a data such as a numerical image that we need to process, but it is a physical object. Thus, the only solution consists in processing more products than needed, so that at the end, the required number of finished products are output.

The failure rate of task T_i is characterized by the percentage of failure for this task. More precisely, the failure is denoted $f_i = \frac{a_i}{b_i}$, where a_i is the number of products that fail each time b_i products have been processed. $r_i = b_i - a_i$ is the number of successful products, and b_i is also called *period* of task T_i .

This failure model is very close to the Window-Constrained model [9], in which an application might tolerate a losses every b products.

We enforce that two tasks of similar type are likely to fail at the same rate with the following equation:

$$\forall i, i' \in [1, n] \quad t(i) = t(i') \Rightarrow f_i = f_{i'}.$$

Since we advocate the computation of more products than needed, we explain in the following how to compute the number B of products that should be processed in order to get B_o products as an output, and we illustrate it on the example of Figure 2. For instance task




Figure 2: Example of a linear chain application with failure.

T_2 has one failure every 6 products that are being processed by this task. Given these failure rates, the number of products that should be given as an input to task T_i in order to have B_o products out of the system is denoted \bar{x}_i . Thus, $B = \max_{1 \leq i \leq n} \bar{x}_i$.

The computation of \bar{x}_i is done backward: if we know the number of products that should be output by task T_i and its failure rate, we can compute the number of input products that should be given to this task to guarantee this output. As we work only with linear chain and in-trees, each task has an unique outgoing edge, thus the number of products to be output by T_i is \bar{x}_{i+1} (or B_o for T_4 in the example).

To determine \bar{x}_i , we need to sum both the number of products which will be successfully processed (i.e., \bar{x}_{i+1}), and the number of products which fail during the processing phase. Thus we must compute the number of periods of the task T_i , which is an integer number being greater than the number of output products divided by the number of successful products computed each period: $\left\lceil \frac{\bar{x}_{i+1}}{r_i} \right\rceil$. In the worst case, failures occur for the first a_i products of the period, thus the number of products to be computed and which will fail is $a_i \times \left\lceil \frac{\bar{x}_{i+1}}{r_i} \right\rceil$.

Finally we can deduce the total number of products that should be computed:

$$\bar{x}_i = \bar{x}_{i+1} + a_i \times \left\lceil \frac{\bar{x}_{i+1}}{r_i} \right\rceil \quad (1)$$

As an extension of this formula, we can also deduce the completion time $L_{\bar{x}_{i+1}, i}$ needed to exit \bar{x}_{i+1} products out of task T_i , once it has been assigned to a machine, M_u . The completion time $L_{\bar{x}_{i+1}, i}$ is the maximum between the time needed to compute \bar{x}_i products on task T_i (i.e., $\bar{x}_i \times w_{i,u}$) and the sum of the time to output \bar{x}_i product out of the task T_{i-1} (i.e., $L_{\bar{x}_i, i-1}$) and the completion time of the last product on task T_i (i.e., $w_{i,u}$).

$$L_{\bar{x}_{i+1}, i} = \max(L_{\bar{x}_i, i-1} + w_{i,u}, \bar{x}_i \times w_{i,u}) \quad (2)$$

3 Optimization problems

Now that the framework has been clarified, we formalize in this section the various optimization problems that we wish to solve. Our goal is to assign tasks to machines so as to optimize some key performance criteria. The solution to a problem is thus an allocation function $a : [1..n] \rightarrow [1..m]$ which returns for each task the machine on which it is executed. Thus, if $a(i) = u$, task T_i is executed on machine M_u , and the processing of one product for this task takes a time $w_{i,u}$.

We first discuss the objective criteria that we want to optimize. Then we introduce the different rules of the game that can be used in the definition of the allocation function a . Finally, Section 3.3 gives a summary of all problem variants, combining framework characteristics and rules of the game. The complexity of these various problems is discussed in Section 4.

3.1 Objective function

In our framework, several objective functions could be optimized. For instance, one may want to produce a mapping of the tasks on the machines as reliable as possible, i.e., minimize the total number of products to input in the system, B . Rather, we consider that products are cheap, and we focus on a performance criteria, the throughput. The goal is to maximize the number of products processed per time unit, making abstraction of the initialization and clean-up phases. This objective is important when a large number of products must be produced.

Rather than maximizing the throughput of the application, we rather deal with the *period*, which is the inverse of the throughput. First we need to introduce the fractional number \tilde{x}_i , which is the average number of products required to output one product out of the system for task T_i . Similarly to the computation of the \bar{x}_i performed in Section 2.3, we can compute the \tilde{x}_i recursively for any application DAG, setting the number of final products $B_o = 1$. Indeed, if task T_i needs to output B_i products, then $\tilde{x}_i = \frac{b_i}{r_i} \times B_i$ (the fraction represents the number of products needed per successful product). Starting from the nodes with no successor (and thus $B_i = B_o = 1$), we can then compute \tilde{x}_i for all i . Note that $\tilde{x}_i \leq \bar{x}_i$ for $B_o = 1$ since \bar{x}_i is an upper integer part which accounts for the worst case failures.

The computation of \tilde{x}_i and \bar{x}_i for the example of Figure 2 is illustrated in Table 1. For instance, $\tilde{x}_4 = 9/(9-2) = 9/7 \simeq 1.3$.

We are now ready to define the *period* of a machine: it is the time needed by a machine to execute all the tasks

Table 1: Values of \bar{x}_i and \tilde{x}_i for the example of Figure 2, with $B_o = 1$.

Task number	1	2	3	4
\bar{x}_i	7	5	4	3
\tilde{x}_i	$\simeq 2.2$	$\simeq 1.8$	$\simeq 1.5$	$\simeq 1.3$

allocated onto this machine in order to produce one final product out of the system. Formally, we have

$$period(M_u) = \sum_{a(i)=u} \tilde{x}_i w_{i,u} \quad (3)$$

The period of machine M_u is the sum, for each task allocated to that machine, of the average number of products (\tilde{x}_i) needed to output one product, multiplied by the speed ($w_{i,u}$) of that task onto that machine. The slowest machine will slow down the whole application, thus we aim at minimizing the largest machine period. The machines realizing this maximum are called *critical machines*. If M_c is a critical machine, then $period = period(M_c) = \max_{M_u \in \mathcal{M}} period(M_u)$. Note that minimizing the period is similar to maximizing the throughput.

3.2 Rules of the game

In this section, we classify several variants of the optimization problem that has been introduced. For *one-to-one* mappings, we enforce that a single task must be mapped onto each machine. Then we consider the case of specialized machines: several tasks of the same type can be mapped onto the same machine; such mappings are called *specialized* mappings. Finally, *general* mappings have no constraints: any task (no matter the type) can be mapped on any machine.

3.2.1 One-to-one mappings

In this first class of problems, a single task is mapped on each machine. This rule of the game is enforced with the following constraint, meaning that a machine cannot compute two different tasks:

$$\forall 1 \leq i, i' \leq n \quad i \neq i' \Rightarrow a(i) \neq a(i')$$

On Figure 3, we have an application graph (a) that must be mapped on a platform graph (b). The result is shown in (c) where we can see that one machine can handle only one task. Thus this mapping is quiet restrictive because we must have at least as many machines as tasks.




Figure 3: One-to-one mapping.

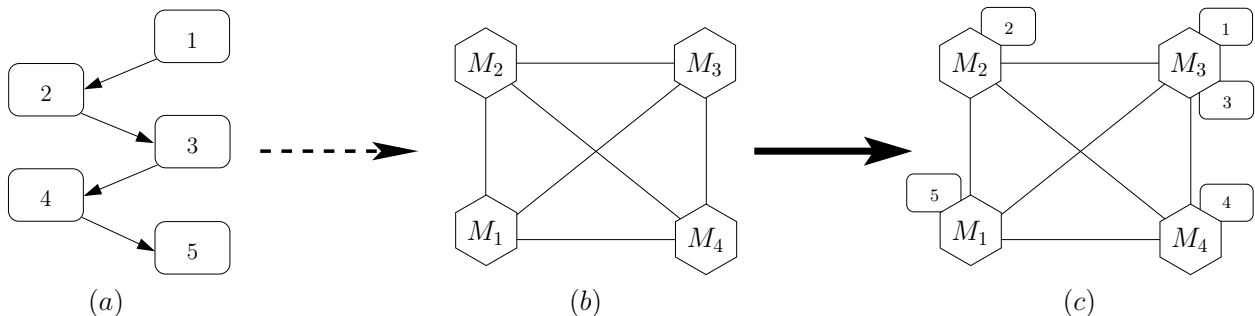


Figure 4: Specialized mapping, with $t(1)=t(3)=t(5)=1$ and $t(2)=t(4)=2$.

3.2.2 Specialized mappings

We have dedicated machines that can realize only one type of tasks. But task types are not dedicated to machines, so two machines may compute different tasks of the same type.

For instance, let us consider five tasks T_1, T_2, T_3, T_4, T_5 with the following types: $t(1) = t(3) = t(5) = 1$ and $t(2) = t(4) = 2$ (see Figure 4). If the machine M_3 computes task T_1 , it could also execute T_3 and T_5 but not T_2 and T_4 . As types are not dedicated to machines, T_5 could also be assigned to another machine, for instance M_1 .

The following constraint expresses the fact that a machine cannot compute two tasks of different types:

$$\forall 1 \leq i, i' \leq n \quad t(i) \neq t(i') \Rightarrow a(i) \neq a(i')$$

3.2.3 General mappings

A machine can compute any task regardless of its type, thus there are no constraints. Such a mapping is illustrated on the example of Figure 5. The problem with this kind of mappings is the reconfiguration cost, since a single machine must perform several types of operations (tasks).

3.3 Problem classification

We summarize in this section the optimization problems which arises from our application. The two important parameters of a problem are

- the rules of the game (*one-to-one* or *specialized* or *general* mapping);
- and the degree of heterogeneity of machines and tasks: the time to compute one product of task T_i on machine M_u may be identical for each task/machine (w), depend only on the task (w_i) or the machine (w_u), or be fully general (w_{iu}).

4 Complexity results

Complexity results are classified depending on the mapping rules. We start with one-to-one mappings, then we focus on specialized and general ones. Finally, we compare one-to-one mappings with general and specialized ones.

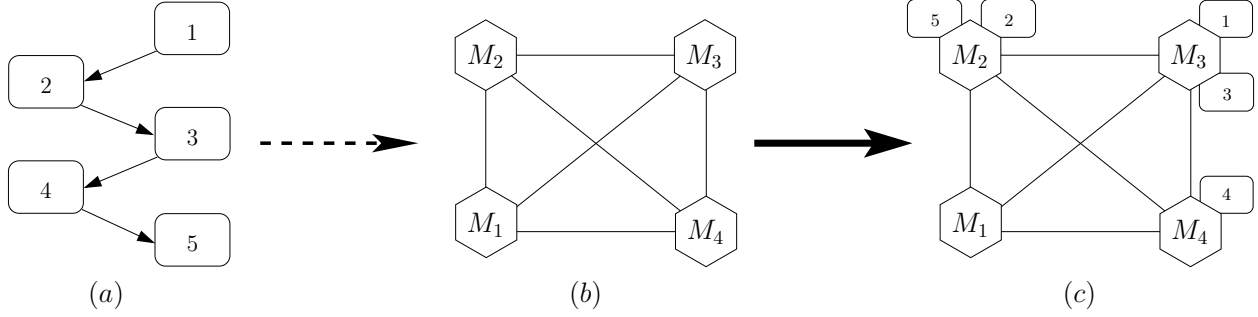


Figure 5: General mapping, with $t(1)=t(3)=1$, $t(2)=t(4)=2$ and $t(5) = 3$.

4.1 Complexity of one-to-one mappings

Theorem 1. *Given an application and a set of machines, finding the one-to-one mapping which maximizes the throughput can be done in polynomial time.*

Proof. We can compute the average number of products \tilde{x}_i needed to output one product out of task T_i , as explained in Section 3.1. Since the mapping is required to be one-to-one, we create a bipartite graph with one node per task on one side, one node per machine on the other side. The cost of an edge from task T_i to machine M_u is then set to $\tilde{x}_i w_{i,u}$, which corresponds to the period of machine M_u if task T_i is assigned to this machine. Since the period of the mapping is the maximum of the periods of each machine, the problem is equivalent to a maximum weight matching in bipartite graphs, which can be found in polynomial time, for instance using the Hungarian method [2, 5]. \square

4.2 Complexity of specialized and general mappings

Theorem 2. *Finding the optimal specialized or general mapping is NP-hard, even with constant processing costs w .*

Proof. We consider the following decision problems: given a period K , is there a general/specialized mapping whose period does not exceed K ? The problem is obviously in NP: given a period and a mapping, it is easy to check in polynomial time whether it is valid or not. The NP-completeness is obtained by reduction from 2-PARTITION [3]. Let \mathcal{I}_1 be an instance of 2-PARTITION: given a set $\{a_1, \dots, a_n\}$ of n integers, does it exist a subset I such that $\sum_{i \in I} a_i = \frac{1}{2} \sum_{1 \leq j \leq n} a_j$? We construct the instance \mathcal{I}_2 with n tasks ordered as a linear chain, 2 machines, and $w = 1$. All tasks are of the

same type, thus there is no difference between general and specialized mappings, and both problems are tackled simultaneously. We assume that $a_1 \geq a_2 \geq \dots \geq a_n$ (the sort can be done in polynomial time), and then we fix:

- $f_n = \frac{a_n - 1}{a_n}$; $\forall 1 \leq i \leq n - 1, f_i = \frac{a_i - a_{i+1}}{a_i}$;
- $K = \frac{1}{2} \sum_{1 \leq j \leq n} a_j$;

First we prove by induction that $\tilde{x}_i = a_i$ for $1 \leq i \leq n$. For $i = n$, we have $\tilde{x}_n = 1 \times b_n / r_n = a_n$. For $1 \leq i \leq n - 1$, if $\tilde{x}_j = a_j$ for $j > i$, then $\tilde{x}_i = \tilde{x}_{i+1} \times b_i / r_i = a_{i+1} \times a_i / a_{i+1} = a_i$.

The size of \mathcal{I}_2 is polynomial in the size of \mathcal{I}_1 . Suppose that \mathcal{I}_1 has a solution I . We construct the allocation function a such that: $\forall i, a(i) = 1 \iff i \in I$. Since $w = 1$ and $\tilde{x}_i = a_i$ for all i , the period of the mapping is thus $P = \max\{\sum_{i \in I} a_i, \sum_{i \notin I} a_i\}$, that means $P = K$ and \mathcal{I}_2 has a solution.

Suppose now that \mathcal{I}_2 has a solution. Let $I = \{a_i | a(i) = 1\}$. By hypothesis, we have $\sum_{i \in I} a_i \leq K$ and $\sum_{i \notin I} a_i = 2K - \sum_{i \in I} a_i \leq K$. We can conclude that $\sum_{i \in I} a_i = \frac{1}{2} \sum_{1 \leq j \leq n} a_j$. Then, \mathcal{I}_1 has a solution. This concludes the proof. \square

4.3 Comparison of mapping rules

In this section, we compare the three mapping strategies, namely one-to-one, specialized and general mappings. The first thing that we want to point out is that one-to-one mappings are a particular case of specialized mappings, which are themselves a particular case of general mapping. Thus, an optimal one-to-one mapping cannot be better than an optimal specialized mapping, which itself cannot be better than a general mapping.

Why not restrict to general mappings? The problem of these general mappings is that they are not realistic, because if a machine is processing tasks of different types, one needs to reconfigure the machine between

operations, and this cost is unaffordable in most micro-factories. Thus, in the following, an emphasis is given to one-to-one and specialized mappings.

Since the optimal one-to-one mapping can be found in polynomial time (see Theorem 1), why not restrict to such mappings? The problem arises when $m \leq n$, i.e., there are many tasks and not so many machines. In such cases, it is mandatory to execute several tasks on the same machine. When there are enough machines ($m \geq n$), one-to-one allocations are a good way to tackle the problem (see the following theorem), but they can be arbitrarily worse than a specialized allocation in the general case.

Theorem 3. *If $m \geq n$, and for problems with w_i ($w_{i,u} = w_i$ for $1 \leq u \leq m$), there is an optimal specialized or general mapping which performs a one-to-one allocation of tasks onto machines. In other words, one-to-one mappings are dominant in this case.*

Proof. The proof is simply done by an exchange argument. Suppose that there is an optimal mapping which is not a one-to-one mapping. For instance, tasks T_i and T_j are mapped onto the same machine, M_u . Since $m \geq n$, there is at least one free machine, say M_v , and the period can be decreased from $\tilde{x}_i w_i + \tilde{x}_j w_j$ to $\max(\tilde{x}_i w_i, \tilde{x}_j w_j)$ if task T_j is assigned to M_v instead of M_u . This concludes the proof. \square

Note that this is not true if the completion time also depends on the processor. For instance, consider a problem with w_u ($w_{i,u} = w_u$ for $1 \leq i \leq n$) in which there is one machine with $w_1 = 1$ and a second one with $w_2 = K$, where K is arbitrary large. If the application consists in two tasks of same type with no failures, then the optimal throughput can be obtained by mapping both tasks onto machine 1, resulting in a period of $1 + 1 = 2$, while a one-to-one mapping must use machine 2 and thus its period cannot be better than K , which can be arbitrarily greater than 2.

5 Heuristics

As explained in Section 4.3, general mappings are not realistic in the context of micro-factories, because of the unaffordable reconfiguration costs. When the number m of machines is greater than the number p of task types, it is always possible to find a specialized mapping, since each machine is able to process all the tasks of a same type. The key point is thus to find m (or less) groups of tasks of the same type to be assigned to the m machines of the platform. The best solution may be a one-to-one

mapping (cases in which such mappings are optimal, see Theorem 3).

As shown before, finding the optimal specialized mapping is NP-hard (see Theorem 2). Thus, we present in the following five heuristics that returns a mapping, by grouping tasks of same type onto machines.

H1: Random heuristic — The m groups are made by using a random assignment function. We randomly choose p tasks, such that $t(i) \neq t(j)$ for all chosen tasks T_i and T_j , and we randomly assign them to p machines of the platform. Then we can randomly assign the rest of the tasks T_i either on a machine which is free or already specialized to the same task type $t(i)$.

H2: Task group heuristic — p groups are made by assigning all the tasks of the same type to the same group. While the number of groups is less than m , the number of machines, the group which consists in the larger number of tasks is divided into two groups to balance the workload. Then, an assignment of groups to machines is performed using the *one-to-one* mapping algorithm.

H3: Binary search heuristic 1 — This heuristic performs a binary search on the period, and aims at finding a correct mapping with the smallest period as possible. The minimum value used in the binary search is 0, and the maximum value is the worst time required to perform sequentially all the tasks on a machine. Then we look for a greedy assignment that optimizes the use of the potential of every resource and respects the target period. We begin by checking the type that asks the larger number of processing tasks. The heuristic stops when a given precision has been reached.

H4: Binary search heuristic 2 — This heuristic is based on the same idea as the previous one. Now we sort the machines with their $w_{i,u}$ in ascending order and once again, we check the task types in the same order as before.

H5: Binary search heuristic 3 — This heuristic is the same as H4 except the fact that the machine are sorted by their heterogeneity level in descending order when considering the assignment phase.

The next section presents experimental results that compare these heuristics.

6 Experiments

In this section, we compare together the 5 heuristics that give sub-optimal solutions to the specialized mapping problem with $w_{i,u}$. The performance of each heuristic is measured by its period in *ms* (see Section 3.1).

Recall that m is the number of machines, p the number of types, and n the number of tasks. Each point in the figures is an average value of 50 simulations where the $w_{i,u}$ are randomly chosen between 100 and 1000 *ms*, for $1 \leq i \leq n$ and $1 \leq u \leq m$. Similarly, failure rates f_i ($1 \leq i \leq n$) are randomly chosen between 0.5 and 2 % (i.e., 1/200 and 1/50). To analyze the impact of the platform heterogeneity ratio, the same experiments have been run with a smaller duration interval (100 to 200 *ms*) in order to simulate less heterogeneous platforms; however results are very similar, except the scale (see Figures 9, 10, 11, 21, 22 and 23).

In the first set of experiments, m and p are fixed, and we plot the period for each heuristic as a function of the number of tasks n . Figures 12, 13 and 14 show that the random heuristic H1 returns very large periods, compared to the 4 other heuristics. This remains true for all experiments: H1 shows very poor performance. Thus H1 is removed from the curves for readability.

Figures 7, 8, 16, 17, 20 and 28 show that the performance of H2 is very similar to that of H4 and H5 when the difference between the number of machines and the number of types is small. Indeed, H2 tries to use all the machines and thus it splits the groups until it has as many groups as machines. For instance, in the experiment of Figure 17, the way the groups are split does not influence the performance so much because only 2 extra groups will be created (20 machines for 18 types). H3 is clearly the best heuristic in such a case. On the contrary, when the number of machines m is much greater than the number of types p , the performance of H2 decreases, as we can see in Figures 6, 15, 18, 19, 24, 25, 26 and 27. In these experiments, H3 is always returning the best period; H4 is sometimes having a result close to that of H3, and H5 is always slightly less good.

In the last experiment (Figure 29), we fix $m = n = 100$, and we plot the period as a function of the number of types p . Moreover, we randomly chose values $w_{i,u}$ such that the duration of a task is machine-independent ($w_{i,u} = w_{i,u'}$ for $1 \leq u, u' \leq m$). In this case, we know that there is an optimal one-to-one mapping (see Theorem 3) and we are able to compute it (see Theorem 1). Thus we are able to assess the absolute performance of the heuristics by computing the optimal period, obtained

with a one-to-one mapping (Hungarian algorithm). The results show that H3 and H4 return a mapping whose period is very close to the optimal, which is a very good result. Indeed, we expect this behaviour to be similar in a more heterogeneous context, thus assessing the performance of our heuristics. H5 is always returning greater period, thus showing that faster machines must be considered first to find a good mapping (recall that H5 sorts machines in decreasing order of $w_{i,u}$).

7 Conclusion

In this paper, we have investigated a throughput optimization problem in the context of micro-factories subject to failures. The problem consists in assigning tasks to machines, either performing a one-to-one mapping (one task per machine), or a specialized mapping (several tasks of the same type per machine), or a general mapping. On the theoretical side, we proved that the optimal one-to-one mapping can be found in polynomial time, while the problem becomes NP-hard for specialized and general mappings. Since general mappings are not usable in practice because of reconfiguration costs, we focused on specialized mappings and proposed several polynomial heuristics to solve the problem. Experimental results suggest that some heuristics return mappings with a throughput close to the optimal, and the sophisticated heuristics return results much better than a random mapping.

As future work, we plan to investigate other mapping rules, as for instance the mapping of one task onto several machines. In such a case, different instances of the task would be handled by different machines. This would allow to obtain a better throughput when a task is time consuming (bottleneck). Also, it would be interesting to consider a failure model in which the failure rate is also machine-dependent (rates $f_{i,u}$ depending on both the task T_i and the machine M_u on which the task is mapped). Finally, other objective functions could be considered, as for instance the total time required to obtain a given number of products, or the average time needed to output one product.

References

- [1] W. Cirne, F. Brasileiro, D. Paranhos, L. F. W. Góes, and W. Voorsluys. On the efficacy, efficiency and emergent behavior of task replication in large distributed systems. *Parallel Computing*, 33(3):213–234, 2007.
- [2] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. The MIT Press, 2001.
- [3] M. R. Garey and D. S. Johnson. *Computers and Intractability, a Guide to the Theory of NP-Completeness*. W.H. Freeman and Company, 1979.
- [4] P. Jalote. *Fault Tolerance in Distributed Systems*. Prentice Hall, 1994.
- [5] H. W. Kuhn. The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2:83–97, 1955.
- [6] V. P. Nelson. Fault-tolerant computing: Fundamental concepts. *Computer*, 23(7):19–25, 1990.
- [7] B. Parhami. Voting algorithms. *Reliability, IEEE Transactions on*, 43(4):617–629, Dec 1994.
- [8] M. Tanaka. Development of desktop machining microfactory. *Journal RIKEN Rev*, 34:46–49, April 2001. ISSN:0919-3405.
- [9] R. West and C. Poellabauer. Analysis of a window-constrained scheduler for real-time and best-effort packet streams. In *Proc. of the 21st IEEE Real-Time Systems Symp.*, pages 239–248. IEEE, 2000.
- [10] R. West, Y. Zhang, K. Schwan, and C. Poellabauer. Dynamic window-constrained scheduling of real-time streams in media servers, 2004.




Figure 6: $m = 10, p = 2$.




Figure 9: $m = 10, p = 2, 100 < w_{i,u} < 200$.




Figure 7: $m = 10, p = 5$.

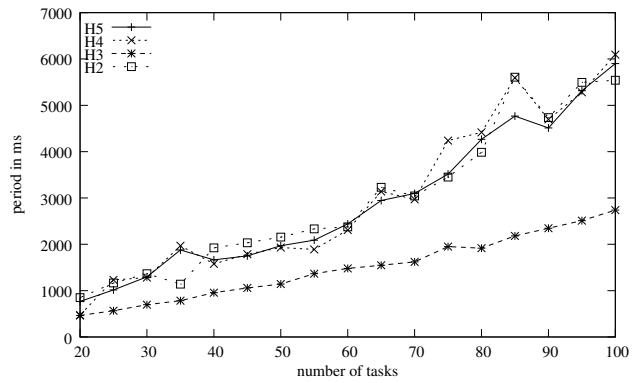


Figure 10: $m = 10, p = 5, 100 < w_{i,u} < 200$.




Figure 8: $m = 10, p = 8$.




Figure 11: $m = 10, p = 8, 100 < w_{i,u} < 200$.




Figure 12: $m = 10, p = 5$.




Figure 15: $m = 20, p = 5$.




Figure 13: $m = 50, p = 25$.

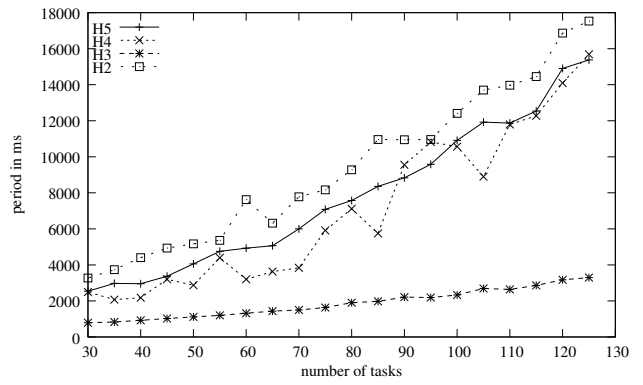


Figure 16: $m = 20, p = 10$.




Figure 14: $m = 50, p = 45$.




Figure 17: $m = 20, p = 18$.




Figure 18: $m = 50, p = 5$.




Figure 21: $m = 50, p = 5, 100 < w_{i,u} < 200$.




Figure 19: $m = 50, p = 25$.




Figure 22: $m = 50, p = 25, 100 < w_{i,u} < 200$.




Figure 20: $m = 50, p = 45$.




Figure 23: $m = 50, p = 45, 100 < w_{i,u} < 200$.




Figure 24: $m = 100, p = 5$.




Figure 27: $m = 100, p = 75$.




Figure 25: $m = 100, p = 25$.




Figure 28: $m = 100, p = 90$.




Figure 26: $m = 100, p = 50$.




Figure 29: $m = n = 100$, with $w_{i,u} = w_{i,u'}$.