

NISTIR 8091

Predicting Robot System Performance by Measuring and Composing Subsystem Attributes

Michael Shneier

This publication is available free of charge from:
<http://dx.doi.org/10.6028/NIST.IR.8091>

NISTIR 8091

Predicting Robot System Performance by Measuring and Composing Subsystem Attributes

Michael Shneier
*Intelligent Systems Division
Engineering Laboratory*

This publication is available free of charge from:
<http://dx.doi.org/10.6028/NIST.IR.8091>

November 2015



U.S. Department of Commerce
Penny Pritzker, Secretary

National Institute of Standards and Technology
Willie May, Under Secretary of Commerce for Standards and Technology and Director

Abstract

A robotic system typically consists of a number of subsystems, for example for sensing, manipulation, and motion. Performance measurements are frequently obtained independently for each of the subsystems but, because of conflicting interactions, it may be difficult to predict the performance of the whole system from these measures. It is often useful to separate the independent subsystem measures from those that interact and to trade off performance between interacting subsystems to achieve global performance goals. This paper discusses the representation and composition of performance measures that aim at achieving such a global understanding of a robot system's performance.

Keywords: Manufacturing System; Performance Evaluation; Performance Measures; Robotics; Taxonomy.

Disclaimer: Commercial equipment and materials are identified in order to adequately specify certain procedures. In no case does such identification imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the materials or equipment identified are necessarily the best available for the purpose.

1. Introduction

Robotic workcells are complex systems made up of components that may themselves be complex. Setting up a workcell typically requires the expertise of a robot integrator who selects the components and designs their interactions to achieve the desired functionality and throughput. This process requires a lot of experience and there are few guidelines for how to select components. Usually there are many competing criteria that must be taken into consideration, including cost, throughput, reliability, ease of setup and use, and quality of the products produced. Ideally, the system would be made up of subsystems, each of which is the best available within the budget, resulting in a total system that is, similarly, as good as can be obtained for the price. Unfortunately, the definition of “best” for a subsystem is usually dependent on the role the subsystem plays in the overall workcell and there are often conflicts between the requirements of the subsystems that may make it impossible to use them in the same system. As a result, it is very difficult to develop a principled approach to workcell design and most implementations involve tradeoffs simply to get the system to work.

While many design problems are indeed intractable, it would be very useful to know which ones are not and to understand what makes a given problem unsolvable. This would enable a more structured approach to workcell design and let tradeoffs be made with greater understanding of what the real issues are. Given that high-level system performance goals are frequently much less constrained than lower-level goals, this might make it easier to achieve an approximation of the ideal solution. This paper proposes an approach to design that starts with performance measures for the subsystems and overall system performance goals. In principle, it should be possible to decompose the problem into one of three classes, from simplest to most difficult:

- There is no interaction between the subsystems or any interactions do not place constraints on the impacted subsystems.
- The subsystems interact but the constraints induced by the interactions are satisfiable.
- The subsystems interact and the constraints are not satisfiable (note that establishing this may not be decidable).

In the first case, it is easy to build a system by selecting subsystems that meet the overall goals such as performance, cost, and reliability. In the second case, it is more difficult because selecting a subsystem might place constraints on how the other subsystems may be chosen or may bound the parameters required for another subsystem (e.g., by requiring a maximum time to complete a step of the task or a minimum quality of fit to prevent later jamming). This leads to a combinatorial problem, but one that can be solved either analytically or numerically. The last case is not solvable but it may be possible to determine what is preventing a solution and develop an approximation, either by modifying the overall system goals or by limiting the impact to a defined class of parts (e.g., if the system was designed to make a range of parts, perhaps a subset of the range can still be produced).

The need for establishing global performance measures through composition of local measures is not restricted to manufacturing robotics. For example, a recent report from the United States Department of Defense Research and Engineering Community of Interest on Automation [1] describes the need in the test and evaluation of autonomous systems for composable measures of performance, although the report does not suggest any specific measures.

2. Related Work

In general, the performance of a robotic system as a whole is measured with respect to the task being accomplished or the goals of the company making the products, while the performance of each subsystem is measured against more limited and specific criteria. There may not be a direct relationship between the overall goals and the individual subsystem goals, but a path must be established that links them if it is to be possible to predict the overall performance from that of the components. While this is frequently an *ad hoc* process, a number of multi-objective optimization techniques have been applied to try to bridge the gap. Other approaches have included using simulations or methods based on fuzzy reasoning or genetic algorithms. Much of the literature is focused on the problem of job shop scheduling, which involves assigning tasks to a set of machines so as to optimize production. While this is not exactly the same problem discussed here, the solution approaches are applicable.

Several researchers have developed simulation-based systems for manufacturing planning. In a European Esprit project, Valckenaers, et al. [2] developed a benchmarking system that made use of a web-based simulation. It consists of three parts: a testbench assistant that interactively helps set up the problem, the emulator that actually carries out the simulation and collects statistics, and the performance evaluation component that computes a range of measures that can be used to evaluate how well different control schemes work. Xing, et al. [3] define a simulation model to solve the multi-objective flexible job shop scheduling problem. They developed their simulation using Matlab and developed a model consisting of six components: the input component that defines the problem and its associated data; the operation assignment component that optimizes the assignment of operations to machines; the operation sequencing component that provides the flow of operations through machines; the objective evaluation component that creates a schedule of operations; the control component that is responsible for flow control and ensures that the scheduled operation sequence is feasible; and the output component that presents the optimization results to the user. In a different application, that of Input/output subsystems in a computer operating system, Ganger and Patt [4] show how a simulation enables understanding of subtle issues that arise when performance measures of subsystems are rolled up into performance measures for a whole system.

A system that is designed with goals very much in line with what is proposed here is Shin and Park [5]. Their goal is to address engineering problems in which multiple subsystems are involved, each of which uses different performance measures. The relationships between subsystems are such that inputs for one subsystem are outputs from another. In their approach, the system is decomposed into subsystems corresponding to the parts for which performance values have been defined. An optimization is carried out for each

subsystem using its own design variables, objective function, and constraints. In these optimizations, the coupling terms with other subsystems are treated as constants. The coupled physics is solved separately and transferred to the individual subsystems and the process is iterated until the convergence criteria are satisfied.

Genetic algorithms are another commonly-used approach to the challenges of combining multiple components into a working system that addresses the constraints imposed by subsystems. Sivakumar, et al. [6] apply a genetic approach to the problem of tolerance allocation across a number of manufacturing subsystems. Here, the goal is to maintain the tolerance requirements of parts as they are processed during manufacture. This affects the choice of machines to use and processing methods to apply. They use two different genetic algorithms to search for optimal solutions and show numerically the success of their proposed methods.

Another genetic algorithm approach is presented by Taboada, et al. [7], who explore two problems in flexible manufacturing systems. In the first, a range of products will be made and the goal is to select the machines to be included (and the vendors) and to determine the best overall configuration of the system. The second problem is that of redundancy allocation in which multiple systems can perform the same task but with different characteristics and costs. The goal is to maximize the overall system reliability or availability while satisfying a constraint for some other system characteristics such as system cost. A multi-objective optimization framework is defined to address these problems based on availability analysis. Availability is used to select a particular configuration that maximizes the probability of meeting the demand for each task or the expected productivity for each task. Three examples are presented to demonstrate the effectiveness of the technique.

Fuzzy logic has also been used for multi-object optimization problems. Lakhdari and Sculfort [8] define a hierarchical model for optimizing a manufacturing system based on subjective and objective attributes. Their overall goal is to measure the effectiveness of a system in terms of lead time, quality, and cost. These attributes are broken into the sub-attributes that contribute to them. For example, lead time is considered to be dependent on human resources, material resources, and flexibility. These second-tier attributes are further dependent on basic, measurable items such as machine reliability, availability of raw materials, scrap rates, etc. By combining the basic measures through a structured hierarchy, different scenarios for manufacturing can be compared.

The choice of the function to optimize has a very large effect on the way production is optimized. Cochran, et al. [9] show how traditional performance measures at the factory level have led to manufacturing system designs oriented towards large-scale production. By aligning the performance measures towards specific corporate objectives, they show how more flexible production can be achieved by explicitly incorporating functional requirements such as quality or delivery time. The functional requirements impose a set of design parameters whose achievement becomes the goal of the performance measurement system.

The most relevant and practical engineering approach to designing manufacturing systems is provided by Whitney [10]. He defines an assembly as “a chain of coordinate frames on parts designed to achieve certain dimensional relationships, called key characteristics, between some of the parts or between features on those parts.” He describes a number of ways in which the constraints between the coordinate frames can be used to guide the assembly design. The last few chapters of the book address manufacturing systems and design for manufacturing and assembly. Some of the ideas from this book are discussed below and incorporated into the proposed approach.

3. Approach

Given a set of subsystems with performance measures, the goal is to construct a system from those components to execute some class of assembly operations. There may be many such systems that can be constructed from the subsystems and the one selected should be optimal in some sense. Without imposing some structure on the problem, it is hard to provide any guidance about how to proceed. The approach suggested here is to make use of Whitney’s methodology augmented with a taxonomy of assembly operations with associated performance metrics. Whitney describes the basic factors in system design as:

1. Analyze the product and determine alternate assembly methods and sequences, including for subassemblies. Determine the assembly process requirements, flexibility requirements, and problematic assembly steps.
2. Select an assembly sequence.
3. Determine the required production capacity of the system.
4. List the feasible assembly techniques for each operation and estimate cost and time for each.
5. Select a set of equipment or people that can make the product at the required rate for a reasonable cost.
6. Conduct an economic analysis and proceed to a detailed system design.

Whitney then provides guidance for all the above steps, while acknowledging that there are no approaches that will work in all cases. The underlying assumption is that it is indeed possible to manufacture the assembly for a reasonable cost but a significant amount of analysis, simulation, and approximation will be needed to determine the best approach. It is also expected that multiple iterations will be necessary, not just to converge on an initial system, but to improve the system over time after production begins.

The role of a taxonomy of assembly operations and of the equipment that carries out those operations is to provide alternative techniques that can be combined to make the desired product. Such a taxonomy is the subject of an earlier paper, Shneier, et al. [11]. The taxonomy provides a hierarchical structure that can be traversed to link different components of a manufacturing system together to achieve the required capabilities. It does not include a way of combining performance measures, nor does it resolve the problem of different measures being important at different levels of the hierarchy, but it does constrain which items can be considered at each stage and gives some idea of the

costs incurred and capabilities added when a selection is made. This taxonomy needs to be augmented with assembly-relevant constraints to allow the performance of the combined system to be estimated.

The performance criteria for an assembly system at the highest level are economic and include cost, product quality, flexibility, throughput, reliability, and ease of setup and use (and re-use) of the equipment. The criteria for individual subsystems include the same economic measures and, more importantly, characteristics that determine the suitability of the subsystem to contribute to making the product. Performance measures for a manipulator or a sensor must enable a decision to be made about whether or not it will contribute to a step in the assembly and, if so, what constraints it will impose on the solution and how adding it will affect the total system's cost and performance. An incomplete list of the kinds of measures that are needed includes:

- Tolerances (e.g., variability by which a manipulator reaches a commanded position, uncertainty in the reported output of a sensor)
- Time needed for execution (e.g., a manipulator picking up a part, a sensor locating a part)
- Variability (e.g., in time, position)
- Weight (e.g., load capacity of a manipulator, constraints on capacity in different parts of its reachable volume)
- Size (e.g., the footprint of a machine or the field of view of a sensor, the largest-sized object a robot can pick up)
- Reachability of equipment (e.g., the working volume of a manipulator)
- Accuracy of equipment (e.g., Is a robot capable of mating parts well enough to achieve the assembly? Can a sensor measure to the required accuracy to acquire a part?)
- Speed of equipment (e.g., can the subsystem execute within the time allotted?)
- Part feeding (e.g., does the equipment need precise feeding? Does a sensor need to be added to locate the parts?)
- Required system capacity (e.g., can the system as configured achieve the necessary throughput?)
- Uptime of equipment (is the system reliable enough to ensure that the required capacity will be achieved? If a subsystem fails, what is the impact on the system as a whole?)
- Tool changes required (if tool changes are needed, can the order of operations be modified to minimize them?)
- Testing and inspection (can testing and inspection be incorporated in other steps or is a separate subsystem needed and, if so, how many?)
- Constraints—does adding this subsystem limit how other subsystems can be used (e.g., a manipulator requires a sensor to be able to locate a part to within a specific tolerance)

The taxonomy of subsystems must represent the information necessary for these measures to be computed for the given situation. This is the case not only for equipment such as manipulators, sensors, or transportation devices, but also for the software that controls those devices. Then, when step 5 of Whitney's procedure is reached, it should be possible to compare the different options to be added for each step of the task and choose the best feasible alternative.

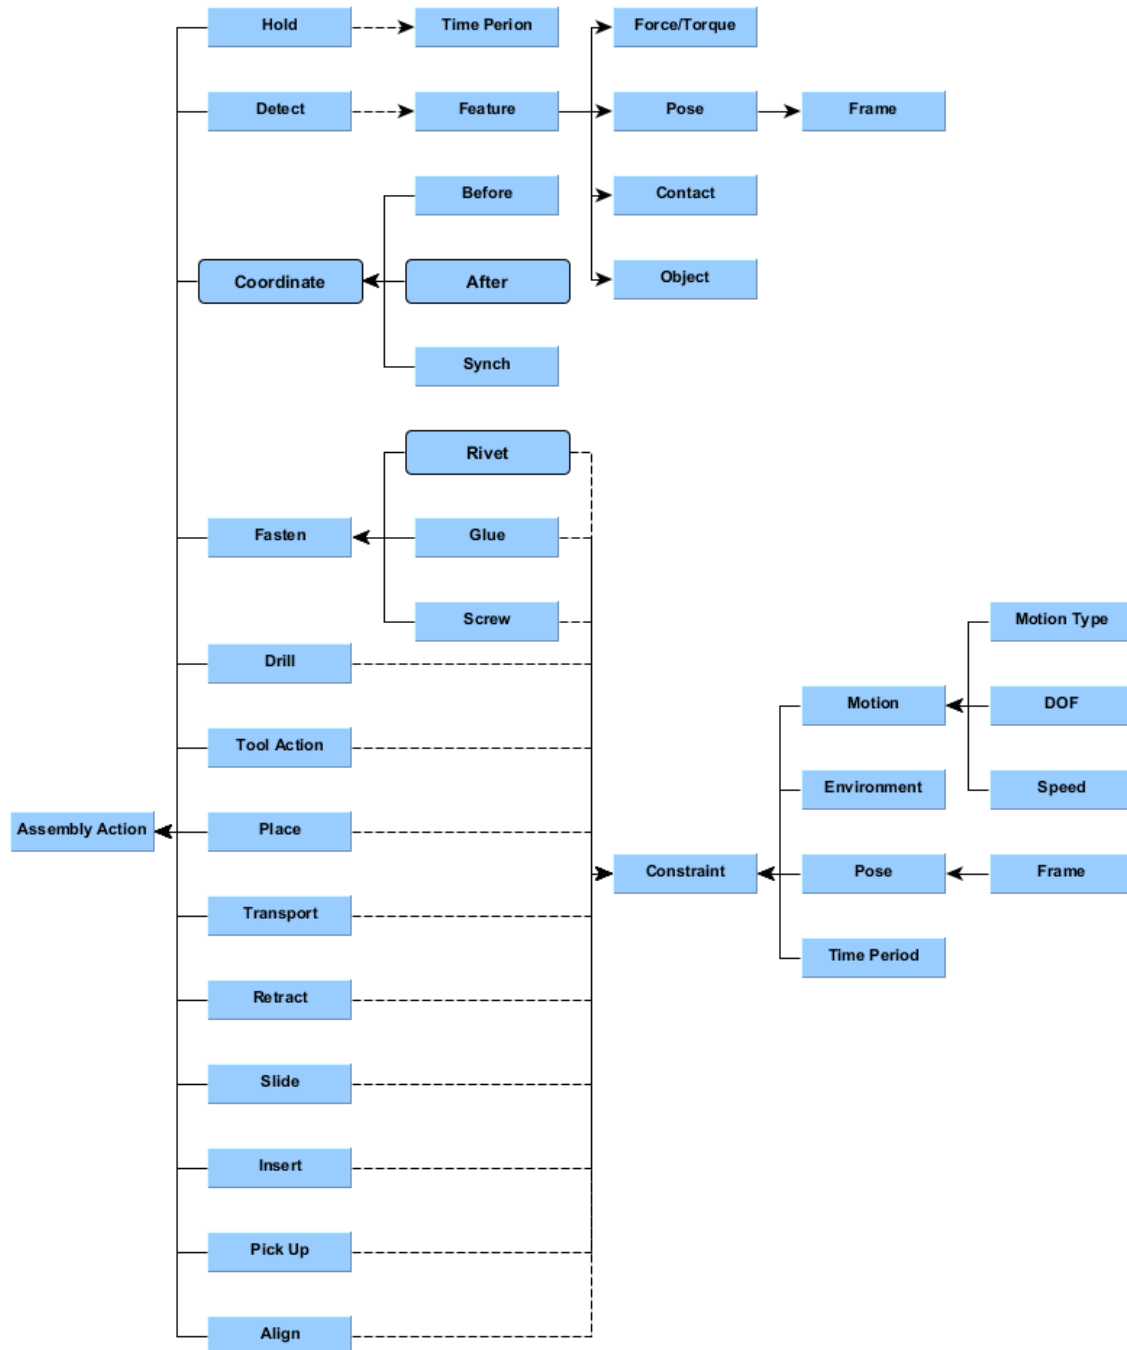


Figure 1. The assembly taxonomy from Shneier, et al. [11].

Figure 1 shows the assembly taxonomy. For each action in the hierarchy, there may be several available pieces of equipment that can achieve the action. Each of them should have relevant parameters from the list above that characterizes its performance. The actions also have constraints that contain further parameters limiting the applicability of a given piece of equipment.

3.1. Example

To illustrate the approach, we present a simple example. Suppose the goal is to build a gearbox out of the set of parts shown in Figure 2. Assume that there are two robots available for this task together with two cameras and a mobile delivery vehicle (Automated Guided Vehicle, or AGV). One of the robots has force sensing, while the other does not. A person is also available to work on the assembly. The goal is to build the gearbox from the components, which will be delivered in a kit. A kit is a collection of the necessary parts, in this case placed in a flat tray but in unknown positions within the tray. The tray may be delivered to the work area by the vehicle and from the kit tray to the assembly area by one of the robots. Either or both robots may be selected to build the gearbox, but they have different capabilities. Similarly, any of the sensors may be used to locate the parts and inspect the results, but they have different fields of view and resolutions.




Figure 2. A set of parts to be assembled into a gearbox.




Figure 3. A view of the gearbox with all parts except the top cover assembled.

1. Analyze the product and determine alternate assembly methods and sequences, including subassemblies. Determine the assembly process requirements, flexibility requirements, and problematic assembly steps.

The assembly is challenging for several reasons. There are a number of small parts that may be difficult for a robot to pick up and manipulate; the gears fit tightly and require a strategy for meshing; they are also similar in size and color. Visual sensing may have trouble with the parts because internal features may be

difficult to detect and there may not be enough contrast with the background. The base of the gearbox (the part on the right of **Error! Reference source not found.**) can move as other parts are attached to it unless it is held or fixtured. A tool is required to screw the top onto the base. The small parts must be removed from the bag before use.

A good starting heuristic is to look for steps that can only be accomplished in one way. Several of the challenges must be addressed by having a person assist with the assembly. There seems no alternative to having the person remove the small parts from the bag and insert the pins into the assembly. The person will also have to insert the screws because the robots have no tools to accomplish this step. There also seems not to be a choice about using the robot with force sensing to insert the gears because of the tight fit and need for meshing. The vision system must be able to recognize the parts and locate them accurately enough to direct the manipulator to grasp them. It must also be able to confirm that each step of the assembly has been completed correctly before the system moves on to the next step. This may not be possible if the field of view of the sensor cannot cover both the location where the kit is delivered and the place where the assembly is carried out. Thus, two vision systems may be required. Given that the kit of parts is delivered by an AGV, two vision systems will be used.

There are a few alternatives with respect to how to proceed with the assembly. One of the robots can fetch the parts while the other inserts them, or one of the robots can act as a fixture and hold the base while the other robot does all the other activities, or one robot can fetch a part and then hold the base while it is added. The person can insert the pins that act as axes for the two smaller gears before the robot inserts the large gear or afterwards or can insert each pin just before the robot adds the corresponding gear.

2. Select an assembly sequence.

The taxonomy in Shneier, et al. [11] defines a set of possible actions to be used in the assembly. Each of these actions is associated in the taxonomy with a set of performance measures that are used to help select suitable equipment to perform the action. The italicized categories in parentheses below are actions from the taxonomy. For all of these actions, a cost can be associated with the taxonomy, and a weighted combination can be used to combine the costs into the overall task cost.

In this example, the assembly sequence starts with the delivery of the kit (*Transport*). Then, the base of the gearbox must be picked up (*Pick Up*) and placed (*Place*) in the work area of one of the robots. This is done with visual input from a camera (*Detect Pose*). Given that inserting the gears will require force sensing, the only robot that can be used is the one that has a force sensor. The

other robot is thus available to move the parts (*Transport*) from the kit to the work area, which may save a little time if it has nothing else to do.

The gears can be inserted in two different orders. The big gear can be installed before the medium-sized gear, or vice versa. It is easier for an automated system to install the medium gear first because the alternative requires a sideways motion with very tight tolerance and, in any case, cannot be completed if the big gear is fully seated. Thus, the sequence is to install the middle gear, followed either by the big gear or the small gear.

There are two options for installing the gears. Either the shafts can be inserted into the gears before they are installed, or the shafts can be installed in the base and then the gears inserted. In either case, the shaft insertions will be done by a person because of the small size of the parts and the need for significant dexterity. If the shafts are inserted into the base before the gears are added, there is a potential problem because the shafts may slide in their holes and it could be difficult for a manipulator to ensure that they stay in place as the force-based search is carried out to seat the gear. The alternative of first inserting the shaft into the gear and then installing the resulting subassembly is also problematic because it is difficult to ensure that the shaft will not slide as it is being installed, making it difficult to seat it in the base. A gripper with a third finger that could block the shaft from moving is a possible solution. It turns out that the shafts fit very tightly into the base so the first option, inserting the shaft in the base is unlikely to cause a problem. The hole in the base into which the shaft fits is very short, so locating the shaft over it is difficult, which is another reason for a person to carry out this operation. Thus, the next step is for the person to pick up the bag of small parts, open it, and insert the shafts that will hold the medium gear and the small gear (it is more efficient to do them both in the same operation) (*Detect Pose, Pick Up, Transport, Detect Pose, Align, Insert, Detect Force*). While this is being done, the second robot can fetch the medium gear (*Detect Pose, Pick Up, Transport, Place*).

The medium gear is installed next. This requires the robot with the force sensor to pick up the part, align it with the shaft, insert it, and push it down until it is fully seated (using force sensing). Since this is the first gear to be inserted, no meshing is needed but the base must be held in place using the second robot (*Detect Pose, Pick Up, Detect Pose, Align, Insert, Detect Force, Retract, Hold*).

Inserting the large gear requires both a force-mediated fit over the shaft on the base and meshing with the medium gear. The robot with the force sensor thus needs to pick up (*Detect Pose, Pick Up*) the big gear from where the second robot has put it down and carry out an insertion operation (*Detect Pose, Align, Insert, Detect Force*) first to fit the hole in the gear over the shaft and then, while pushing down, rotate the gear until the teeth mesh (*Detect Pose, Insert, Detect*

Force, Retract). The second robot must hold (*Hold*) the base to keep it in place during these operations.

Once the other two gears are installed, the small gear can be installed onto the base. This requires picking up (*Detect Pose, Pick Up*) the gear and moving it (*Transport*) to the assembly area, which is done by the second robot. When putting down the part (*Place*), this robot should turn it over so that it is in the right orientation for insertion. The robot with the force sensor then picks up (*Detect Pose, Pick Up*) the component and uses force sensing to assemble it (*Detect Pose, Align, Insert, Detect Force*), while the second robot holds the base (*Hold*). This requires meshing with the teeth of the medium gear while locating the small gear over its shaft. This completes the gear subassembly.

The next step is to install the top cover. This is a difficult operation because all the components have to fit into their proper places. It involves a press fit that will be completed by the robot and a set of screwing operations that will be done by the person. As before, the second robot picks up the part at the kitting area (*Detect Pose, Pick Up*) and moves it to the assembly location (*Transport*). There, the assembly robot picks it up (*Detect Pose, Pick Up*), moves it to align with the assembly (*Transport, Detect Pose, Align*), and uses force sensing to position it on the base (*Detect Pose, Place, Detect Force*). The positioning has to ensure that the top cover is directly over the base, that the shafts of the medium and small gears fit into their sockets on the top cover, and that the large gear protrudes from the cover. The second robot holds (*Hold*) the base during this operation. The person then inserts the screws while the robot continues to hold the top in place (*Detect Pose, Insert, Fasten, Detect Force*). Then the person and assembly robot move away (*Retract*). The final, completed gearbox is then removed from the assembly area and moved to the AGV (*Transport*). Figure 3 shows the assembled gearbox with only the top cover left to attach.

In practice, it is unlikely that two robots would be used for this task. With only one robot, all moves from the kitting area to the assembly area would be carried out by that robot as well as the assembly tasks. The biggest change to the sequence above would be that a fixture would need to be installed in the assembly area into which the base would be placed in the first step. The fixture would hold the base still while the rest of the parts were assembled onto it.

For high-speed, large volume production, a better approach would likely be to feed the gears using vibratory mechanisms and to deliver the base and top cover on trays or pallets. It would also be helpful to incorporate the gear shafts into the extrusion process for the base, which would reduce the need for a person in the assembly. Placing the top cover at the end would be done with fixed automation, and the screws could also be inserted automatically if the top was properly

fixtured, or a different means, such as a snap fit, would be used to attach the top cover.

3. Determine the required production capacity of the system

The capacity of the system is not important for this example, but a real system would need to be able to assemble the gearboxes in one minute or less to be cost-effective (a person can do it that fast).

4. List the feasible assembly techniques for each operation and estimate cost and time for each.

The feasible techniques include manual assembly, automated assembly, and mixed assembly. As described above, either manual or mixed assembly is feasible for the task. It would be very expensive to implement a completely automated assembly station although it would be technically feasible. A fully manual assembly would likely be fastest, but a reasonable throughput should be possible with mixed assembly.

5. Select a set of equipment or people that can make the product at the required rate for a reasonable cost.

This is where the taxonomy is most useful. Given the requirements of the task and the available resources (or, if new equipment is to be purchased, the specifications of candidate machines), a subset is chosen with the capability to make the product. Then, the performance of each piece of equipment is examined to determine if it can do so within the required time and cost. To do this, the performance of each individual machine is checked using the measures in the taxonomy associated with its tasks to be sure that it meets minimal requirements. If there are constraints such as limited conditions under which the performance is acceptable, a study is required to be sure that the limitations will not adversely affect the whole system. Once a piece of equipment is verified to be able to carry out parts of the assembly, its performance is checked against any other equipment that can perform the same parts and a choice is made based on global performance measures such as cost and time. If no machine can perform part of the task effectively, adding a person should be considered.

For the example assembly, there are multiple robots, multiple cameras, an AGV, and a person available for the task. As indicated in 2 above, two robots will be used. Both have the necessary reach and weight capacity and both are repeatable enough and fast enough for their roles. One of the robots has force sensing so it is the only one able to perform the gear insertions and meshing. The other robot performs pick and place operations and acts as a fixture when parts are being added to the base. The actions described above are associated with performance measures for each step in the taxonomy. Shneier, et al. [11] contains a number of tables that provide metrics for each of the actions as well as properties and parameters that are important for establishing how well a given piece of

equipment can carry it out. Applying the information in these tables should allow equipment to be selected and a sequence of actions to be implemented that meets the overall system objectives. There are no guarantees, however, that this will always be possible, nor that the resulting system will be optimal for the task. The steps in the task require visual sensing to locate the parts and ensure they are placed correctly. Two visual sensors are needed, one at the location where the parts are picked up from the kit and the other at the assembly station. The sensor at the kit must have a large enough field of view to see the whole kit plus a region around it in case something is dropped when a part is picked up. It must have enough resolution to locate parts to within the grasp uncertainty of the robot and its gripper. For picking up parts, the resolution can be lower than for assembly because the part will be re-acquired by the other robot. Note that the resolution is a function of the camera itself (the size of the pixels on the chip), of the lens, of the height of the camera above the work surface, and of the algorithm used to locate the parts (typically, algorithms can locate parts to better than 0.1 pixel). Thus, there is a lot of flexibility in the selection of a camera.

The camera used for assembly has similar requirements except that the location accuracy must be significantly higher. It is also harder to position the camera to be able to see the state of the assembly and still be able to monitor the handover of parts from one robot to the other.

Force sensing is critical to this assembly, being necessary both for the peg-in-hole insertion of the gears and for gear meshing. The forces are not expected to be very large, however, and commercial force sensors should be able to handle the requirements.

6. Conduct an economic analysis and proceed to a detailed system design.
This step is beyond the scope of the paper but Whitney [10] contains a good discussion of how to conduct such an analysis.

4. Discussion

Each product presents unique factors to be addressed when determining how best to manufacture it. Each manufacturer has slightly different criteria for system performance and changing markets further complicate the picture, especially when a long-term investment is being considered. Without a careful analysis of the options and their costs and benefits, the risk of failure is significantly increased. Lacking measures that can be used to compare different options, such an analysis is largely subjective. Thus, the availability of performance measures, together with a principled way of applying them, is of great value, especially to smaller manufacturers who are unlikely to have the expertise and experience to address this complex multi-objective optimization problem.

The approach described here provides a framework for the analysis although it doesn't supply a mathematical procedure for combining the performance measures of subsystems to generate an overall system measure. There are too many unique attributes to each problem for one method to apply to all. The general framework at least gives an approach

that helps address all the aspects and shows where individual performance measures come into play and which aspects of those measures are important. As the approach is used, it is likely that gaps will become more obvious both in the method itself and in the taxonomy and performance metrics that it uses. This will help improve the approach and may also lead to better performance measures for subsystems. Some of the methods described in the related work section may be relevant, especially simulation, which can be used to explore the candidate solutions and see how well they work in the actual factory environment.

5. Conclusions

As robots are applied more frequently to assembly in manufacturing, and as the need for quick changes in products and styles increases, there is greater need for ways of designing and implementing automated manufacturing workcells that are capable and cost effective. Workcells are complex and there are many vendors of components, each of which represents a product's capabilities in a way that emphasizes its best aspects. In building an assembly workcell, it is important to be able to compare candidate equipment objectively and to have enough information to select the appropriate components. This paper has presented an approach based on the work of Whitney and on a taxonomy of assembly actions with associated performance measures. It is aimed at enabling small and medium-sized manufacturers who lack the resources and in-house expertise to conduct a complete analysis each time they need to build a new workcell. Over time, it is expected that the performance measures will improve and that the approach will be expanded to include more computational methods and guidelines based on experience.

References

- [1] Office of the Assistant Secretary of Defense For Research & Engineering, Test and Evaluation, and Verification and Validation (TEVV) Working Group, "Technology Investment Strategy 2015-2018," 2015.
- [2] P. Valckenaers, S. Cavalieri, B. Germain, P. Verstraete, Hadeli, R. Bandinelli, *et al.*, "A benchmarking service for the manufacturing control research community," *Journal of Intelligent Manufacturing*, vol. 17, pp. 667-679, 2006.
- [3] L.-N. Xing, Y.-W. Chen, and K.-W. Yang, "Multi-Objective Flexible Job Shop Schedule: Design and Evaluation by Simulation Modeling," *Applied Soft Computing*, vol. 9, pp. 362-376, 2009.
- [4] G. R. Ganger and Y. N. Patt, "Using System-Level Models to Evaluate I/O Subsystem Designs," *IEEE Transactions on Computers*, vol. 47, pp. 667-678, 1998.
- [5] M.-K. Shin and G.-J. Park, "Multidisciplinary design optimization based on independent subspaces," *International Journal for Numerical Methods in Engineering*, vol. 64, pp. 599-617, 2005.
- [6] K. Sivakumar, C. Balamurugan, and S. Ramabalan, "Concurrent multi-objective tolerance allocation of mechanical assemblies considering alternative manufacturing process selection," *International Journal of Advanced Manufacturing Technology*, vol. 53, pp. 711-732, 2011.
- [7] H. Taboada, J. Espiritu, D. Coit, and G. Levitin, "A multi-objective evolutionary algorithm for determining optimal configurations of multi-task production systems," Rutgers University Industrial Engineering Working Paper 07-014, 2007.

- [8] K. Lakhdari and J. L. Sculfort, "Performance Evaluation Model for Manufacturing System Based on Approximate Reasoning," presented at the IEEE International Conference on Service Operations and Logistics, and Informatics. IEEE/SOLI, 2008.
- [9] D. S. Cochran, Y.-S. Kim, and J. Kim, "The Impact of Performance Measurement on Manufacturing System Design," presented at the First International Conference on Axiomatic Design, Cambridge, MA, 2000.
- [10] D. E. Whitney, *Mechanical Assemblies: Their Design, Manufacture, and Role in Product Development*: Oxford University Press, 2004.
- [11] M. Shneier, E. Messina, C. Schlenoff, F. Proctor, T. Kramer, and J. Falco, "Measuring and Representing the Performance of Manufacturing Assembly Robots," National Institute of Standards and Technology (NIST), Internal Report, 2015.