

REAL-TIME CONTROL OF FLEXIBLY AUTOMATED PRODUCTION SYSTEMS

Spyros Reveliotis
Assistant Professor
School of Industrial & Systems Eng.
Georgia Institute of Technology

ABSTRACT

Effective deployment and operation of flexibly automated production systems is of ever-increasing importance in contemporary competitive manufacturing. Yet, currently we lack an adequate methodology for effective deployment, reconfiguration and control of these systems, partly due to the fact that past research has not systematically modeled and analyzed the logical/structural aspects characterizing their behavior. The importance and criticality of these problems becomes even clearer once it is realized that they are directly related to the operationalization and effective management of the flexibility inherent in these systems. This paper seeks to formally introduce the control problem related to the establishment of logically correct and robust behavior for flexibly automated production systems as a separate class of problems in the hierarchical decomposition framework typically used for the real-time control of these environments. It also outlines an analytical framework able to support the systematic study of these behavior/structure-oriented problems, and surveys some key results developed in the area. The last part of the paper considers the integration of structure-oriented control with the performance-oriented control - which is the more "traditional" concern in production planning and control -, demonstrates some interesting effects arising from the combined study of these two issues, and suggests directions for future research.

INTRODUCTION

Effective and efficient deployment and operation of flexibly automated production systems is of ever-increasing importance in today's competitive manufacturing. Even though the current globalization of the economy has led to significant reductions of the human labor costs, automation of the production processes is still desirable on the basis of the predictability and the controllability that it brings - at

least, in principle - to the system behavior. Furthermore, in certain industries, production automation is a necessity imposed by strong technological constraints. For instance, it is deemed [1] that in the next generation of the semiconductor fabs, the adopted 300mm wafer diameter will lead to lots too large and heavy to be handled efficiently by human operators. Similarly, in the emerging industry of bioengineering, processed items are too small to be handled and processed by humans. In other cases, contamination considerations regarding either the product or the human operators make the process automation a competitive option.

Flexibility as an attribute of contemporary manufacturing systems is necessitated by the time-based competition underlying the current manufacturing strategy. Typically, contemporary production environments must present significant economies of scope, i.e., they must be able to accommodate the concurrent production of a large variety of parts, quite often with extensively customized features. Furthermore, production shop floors must be easily reconfigured in order to accommodate large demand fluctuations and operational contingencies, by effectively reallocating production capacities.

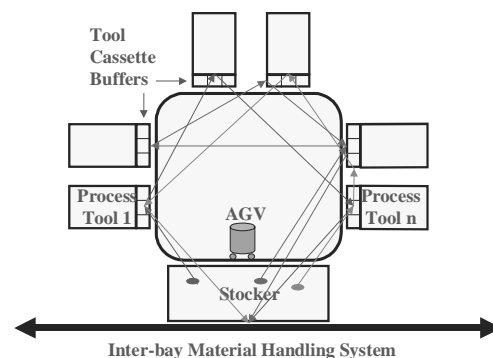


Figure 1 – Intra-Bay Material Flow Control in 300mm-wafer fabs

As a more concrete example of a flexibly automated production system and its underlying operational

characteristics, consider the intra-bay material flow of the emerging 300mm wafer fab. The layout of such a bay area is depicted in Figure 1, and it consists of: (i) a *stocker*, where cassettes of wafers requiring the bay processes are kept stored until their loading and removal from the bay, (ii) a number of *process tools*, that constitute the bay's main processing units, and (iii) the interconnecting material handling system – typically an *Automated Guided Vehicle (AGV)* or *Track Guided Vehicle (TGV)* system – transferring wafer cassettes among the process tools and the stocker. The bay process tools are characterized by different processing capabilities, and each wafer cassette is routed through a distinct sequence of process tools, determined by its accompanying recipe. Furthermore, in order to reduce the system cycle times and contamination risks, it is desirable that a cassette wafer, once released in the bay area, completes its entire processing before it returns to the stocker. Since, however, the operations taking place at the system process tools are quite lengthy, a number of wafer cassettes with different product recipes – and therefore, routing schemes – will be concurrently processed in the system. The wafer cassette flow through the system is constrained by the limited buffering capacity of the bay process tools, which is typically restricted to no more than four or five wafer cassettes. On the other hand, the system presents considerable routing flexibility, since many of the bay process tools have overlapping processing capabilities, established by the availability of appropriate auxiliary equipment (e.g., process masks). Finally, quite often, the system experiences operational contingencies, like decalibration of certain pieces of equipment, or the arrival of expedient (“hot”) jobs; effective reaction to these contingencies requires the drastic reconfiguration of the system logical and/or physical structure.

Beyond providing a concrete example of a flexibly automated production system, the above description also demonstrates the scope and complexity of the problems involved in the effective real-time management of such an environment. Using an analogy from the computational system models, we claim that the effective deployment and operation of flexibly automated production systems essentially requires the development of an “*operating system*” for these environments, that will ensure:

- *consistency* of the system operation with respect to physical, technological, and user-imposed constraints,
- *robustness* and *graceful degradation* with respect to various operational contingencies

through appropriate exploitation of the system operational flexibility, and

- *efficient utilization* of the system resources through pertinent exploitation of the system production capacity.

However, currently we lack an adequate methodology for effective tactical planning, reconfiguration and control of these environments. Quoting from a recently written paper: “*currently, most control implementations of flexible manufacturing cells have been developed specially to a particular facility, and no generic format or tools exist for systematic creation and planning of control software. In general, these systems are developed as “turnkey” systems by personnel other than the manufacturing engineers responsible for their operation.*” [2] As a result, the deployment and operation of flexibly automated production systems is currently characterized by (i) high implementation times and costs, (ii) lack of portability of the resulting control software, and eventually (iii) limited operational flexibility. In a similar vein, a recently conducted survey regarding the support of manufacturing flexibility in printed circuit board assemblies, indicated that “*production technology turned out to be significantly related to mix and new-product flexibility, but with a pattern opposite to what [the authors] expected, given the capabilities of the technology... This touches a very important point: The fact that automated and programmable equipment in the [authors’] sample tends to be used to run the largest production batches, instead of being used in a more flexible way.*” [3]

Developing an effective integrated solution to the production planning and control of flexibly automated production systems is admittedly an extremely complex task. All the same, we believe that a the key reason for the rather limited success of past efforts towards the development of such a control framework, has been their complete focus to performance-related issues, and their inadequate consideration of the logical/structural aspects of the system behavior. It is interesting to notice that the *hierarchical decomposition* framework [4,5], the most widely adopted “analytical” framework for the production planning and control problem, discerns:

- *strategic* decisions, concerning the initial design and subsequent expansions of the production infrastructure,
- *tactical* decisions, concerning the system reconfiguration (resetting) in order to effectively track the externally imposed demand patterns,

- *operational* decisions, concerning real-time issues, like the job induction and dispatching to the system workstations,

all of which address performance objectives, and presume the logically consistent and robust behavior of the production shop floor. In an extensively automated environment, however, the establishment of a behavior that is logically correct and robust to the various operational contingencies, is definitely a responsibility of the underlying control logic. In fact, this logical analysis and control of the system behavior gives rise to a new suite of control problems in the production planning and control framework, which will be collectively referred to as the *structural control (SC)* of flexibly automated production systems. The incorporation and the role of structural control in the hierarchical decomposition framework is depicted in Figure 2. A verbal characterization of the relationship implied in this figure is that *the imposed structural control logic/policy defines the admissible operational space over which the more “traditional” performance-oriented control will seek to optimize the system behavior.*

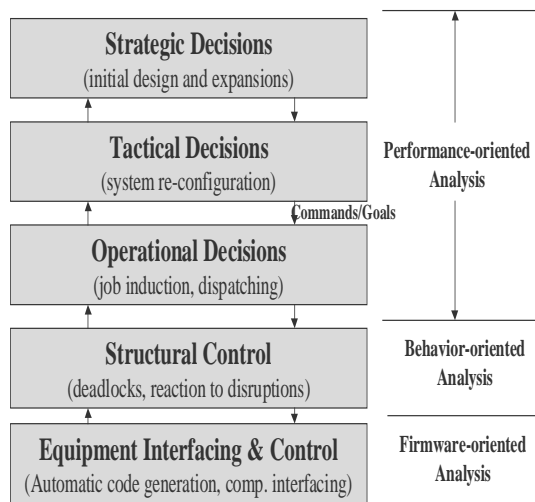


Figure 2 – An Extended Hierarchical Decomposition Framework for Production Planning and Control in Flexibly Automated Manufacturing Systems

Since structural control is the control component immediately responsible for the structuring of the system behavior, it is also a key component for implementing and managing the system operational flexibility. It turns out that the complexity of the design of effective structural control policies (SCPs) is directly related to the amount and types of flexibility that are to be supported by the system operation. Hence, one of the main challenges for the current research in the SC area, is to provide correct and robust system behavior, while allowing for

effective exploitation of the system inherent flexibility.

In the next sections of the paper, we provide an overview of the problem of structural control of flexibly automated production systems, and survey the current research results in the area. We also provide a framework for the integration of structural and performance-oriented control, identify some interesting effects resulting from the combined considerations of these two control components, and outline future research needs.

STRUCTURAL CONTROL OF FLEXIBLY AUTOMATED PRODUCTION SYSTEMS

As it was explained in the previous section, the main objective of structural control is to ensure the logical correctness of the system behavior, and its consistency to any externally imposed operational requirements, by constraining appropriately the system operational space. Currently, the most prevailing problems in the area are:

- the establishment of *deadlock-free* behavior, and
- the development of reaction strategies for the *accommodation of operational contingencies*, like machine breakdowns and the arrival of expedient jobs (cf., the 300mm wafer fab example in the introductory section), through effective exploitation of the system operational flexibility.

DEADLOCK AVOIDANCE IN FLEXIBLY AUTOMATED PRODUCTION SYSTEMS




Figure 3 – A manufacturing system deadlock

To demonstrate the problem of deadlock avoidance arising in flexibly automated production systems, consider the very simple system presented in Figure 3. It consists of two workstations and an interconnecting single-vehicle AGV system. In its current configuration, the system supports the concurrent production of two job types, J_1 and J_2 , with

the corresponding job routes annotated in the figure. Assuming that each of the system workstations and the AGV can buffer only one part at a time, it is easy to see that under the loading scheme presented in Figure 3, the system is permanently stuck: each of the two depicted parts waits upon the other to release its currently allocated resource, in order to proceed to its next stage. This circular waiting situation is characterized as a (manufacturing system) *deadlock*, and it essentially results from the interaction of (i) the *arbitrary* routing of jobs through the system with (ii) the *finite* buffering capacity of the system equipment.




Figure 4 – Developing the STD underlying the operation of the manufacturing system of Figure 3

The formal characterization and analysis of the manufacturing deadlock problem requires the abstraction of the underlying operational environment to a *resource allocation system (RAS)*. Specifically, each of the system workstations and material handling components defines a particular resource type R_i , with finite buffering capacity C_i . The job types processed through the system are characterized by the sequences of the system resources required for their completion. In the special - but quite general - case where each job processing stage is uniquely corresponded to a resource set, and the only resource in that set is a single unit of buffering capacity on the supporting piece of equipment (workstation or MHS), the resulting RAS is characterized as *Single-Unit (SU)*. SU-RAS is the simplest class of RAS modeling the behavior of flexibly automated production systems, and they will be used for the further exposition for the manufacturing deadlock-related concepts.




Figure 5 – Characterization of the reachable state space and the optimal deadlock avoidance policy for the SU-RAS corresponding to the manufacturing system of Figure 3

The SU-RAS *state*, s , is defined by the processing stages of all currently executed jobs. Notice that this information is sufficient for determining the current allocation of the system resources to the various job instances, but also, the potential evolution of the system resource allocation status. Specifically, the system state changes by:

- *loading* a new job in the system,
- *advancing* a job to its next processing stage, or
- *unloading* a finished job.

The evolution of the system state through the occurrence of the three event types defined above can be formally described by the *state transition diagram (STD)* of a *finite state automaton (FSA)* [6]. The systematic development of the STD corresponding to the SU-RAS of Figure 3 is depicted in Figure 4, while its complete deployment is given in Figure 5.

As it is indicated in Figure 4, an SU-RAS *deadlock* can be formally defined as a system state in which there exists a set of resources, R^a , filled to capacity with jobs requesting the allocation of a unit of another resource in R^d . Notice, however, that in Figure 4, there is also another RAS state, which will unavoidably lead to deadlock, even though this state itself is not characterized as deadlock, according to the previous definition. We characterize this broader class of states from which deadlock is unavoidable as *unsafe*, and as it is suggested by Figure 5, the objective of an effective *deadlock avoidance policy (DAP)* is to identify and prevent the system transition to its unsafe operational space. In particular, the policy that disables only transitions to the system unsafe region, while allows full accessibility of the entire reachable and safe subspace is characterized as *optimal*, since it supports maximal operational flexibility of the production system.

In practice, the optimal DAP for any given SU-RAS should be implemented by controlling the admissibility of any tentative resource allocation, based on the safety of the resulting RAS state. It can be shown, however, that in the SU-RAS class, state

safety is an NP-complete problem [7]. Therefore, the computational cost of the real-time implementation of the optimal DAP, according to the one-step look-ahead scheme suggested above, will be prohibitive for practically sized systems. In the light of this remark, current research on the problem of deadlock avoidance in SU-RAS has focused on:

- the identification of special RAS structure able to admit polynomially computable optimal deadlock avoidance, and
- for the remaining cases, the development of suboptimal, yet efficient deadlock avoidance policies, able to be implemented in real-time with polynomial computational cost.

The next two sections provide a brief outline of the key results obtained through these two lines of research.

Special SU-RAS structure admitting polynomially computable optimal DAP

The key mechanism for the identification of these special structures has been the observation that even though the state safety problem is NP-complete in SU-RAS, the *deadlock detection* problem is of polynomial complexity [8]. Hence, if there are any special cases of SU-RAS in which the class of unsafe and deadlock states coincide - i.e., *there are no unsafe states which are not already deadlocks* - then, the optimal deadlock avoidance is polynomially implementable through one-step look-ahead for transitions to deadlock states. Indeed, past research has identified three SU-RAS subclasses with the aforesaid property:

- SU-RAS in which every resource has capacity strictly greater than 1. [8]
- SU-RAS in which every resource has a single predecessor and/or successor resource w.r.t. the supported job routings. [9]
- SU-RAS which correspond to re-entrant lines s.t. (i) every “forward” step is from resource R_i to R_{i+1} , and (ii) every “backward” step is immediately followed by at least one “forward” step. [10]

Notice that most of these conditions are of a very practical nature, and therefore, they can be easily implemented in the design of many flexibly automated production systems.

Suboptimal, polynomially computable DAPs for more general SU-RAS

The basic reasoning behind the development of these policies is the observation that, since the identification of the safe subspace in these broader SU-RAS classes is computationally intractable, we should strive to limit the system operation in a region of the safe space which is polynomially identifiable. However, such a control policy will be logically consistent if and only if (i) it contains the initial idle and empty RAS state, s_0 , and (ii) the policy-admissible state space which is reachable from the initial state s_0 , is strongly connected. A suboptimal DAP satisfying these two conditions is said to be *correct*. Furthermore, due to the basic logic driving their development, this class of policies is also characterized as *Polynomial-Kernel (PK) DAPs*. A correct suboptimal DAP for the SU-RAS of Figure 3, is depicted in Figure 6.




Figure 6 – A correct suboptimal Polynomial-Kernel DAP for the SU-RAS of Figure 3

Given the current emphasis on operational flexibility, an additional requirement for the development of PK-DAPs is that they are efficient, i.e., they do not give up too much of the system safe subspace. The detailed development of a series of PK-DAPs and the evaluation of their efficiency can be found in [11, 12, 13, 14].

FURTHER DEVELOPMENTS IN THE STRUCTURAL CONTROL OF FLEXIBLY AUTOMATED PRODUCTION SYSTEMS

In this section we briefly outline some additional results pertaining to the problem of structural control of flexibly automated production systems.

Extending the basic RAS modeling assumptions

Remember that the basic assumptions underlying the SU-RAS model used in the developments of the

previous section are, that: (i) every processing stage can be executed in only one way (i.e., it can be supported by only one resource), and (ii) the only resource actually required is one unit of buffering capacity in the supporting piece of equipment. However, as we saw in the introductory example, many contemporary production systems are characterized by extensive *routing flexibility*, i.e., each stage can be executed in more than one workstation. Furthermore, in some cases, the successful execution of a certain processing step might require, in addition to the supporting buffering capacity, the exclusive allocation of some *auxiliary equipment*, like fixtures, cutting tools and process masks. The allocation of this auxiliary equipment, when it is centrally managed and in scarce quantities, can also give rise to circular waiting patterns, i.e., lead to deadlocks.

The accommodation of the two additional operational features outlined above in the development of effective DAPs, has led to the extension of the basic SU-RAS model to the *Disjunctive* and the *Conjunctive* RAS, respectively. Deadlock avoidance in these new RAS classes is characterized by higher complexity than the complexity of the corresponding problem in the SU-RAS. Specifically, in the disjunctive RAS case, one has also to deal with the increased space complexity resulting from the potentially exponential number of possible routes for each job type. Results on deadlock avoidance in Disjunctive RAS can be found in [15, 16]. In the Conjunctive RAS case, the increased problem complexity of deadlock avoidance essentially arises from the concurrent job requests for more than one resource type. Deadlock avoidance in the Conjunctive RAS context is the least understood among the three RAS classes, but a formal problem characterization and some correct Polynomial Kernel DAPs for this class of systems can be found in [8, 17].

AGV RAS and Hierarchical RAS structures

Except from disturbing the smooth part flow among the system workstations, deadlocks can also be a problem arising in the internal operation of specific components of these environments. In particular, deadlock is a prominent problem in the operation of *zone-controlled* AGV systems, typically characterized as the problem of establishing *conflict-free routing*.

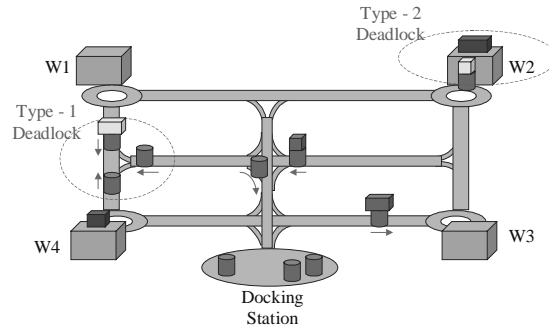


Figure 7 – The AGV RAS

The situation is depicted as the *type-1* deadlock of Figure 7, where three vehicles, having traversed their currently allocated zones, are entangled in a permanent blocking condition, at the zone-intersecting node. The modeling and analysis of the type-1 AGV deadlock along the general RAS theory developed in the previous section, and the development of an effective and efficient DAP for this environment, is undertaken in [18]. It is interesting to notice that in the developed AGV RAS model, processed jobs correspond to the vehicle trips among the source, destination, and potentially the docking station, while the allocated resources are the zones of the system guidepath network.

Another interesting structural control problem is the resolution of the *type-2* deadlock depicted in Figure 7. This situation arises whenever a loaded vehicle arrives to its destination station only to find it fully allocated. If the vehicle persists in its effort to unload its part to the receiving station, it essentially blocks the arrival of any other vehicle that could have retrieved a finished part from that station, and thus, alleviate the problem. Hence, we are faced with another situation of circular, and therefore, indefinite waiting. The resolution of the type-2 deadlock requires the interaction of the controller managing the flow among the system workstations, and the controller managing the vehicle dispatching and the internal traffic of the AGV system. The problem is formally modeled and analyzed in [19], through the development of a hierarchical, distributed structural control framework, depicted in Figure 8. The main challenge underlying the approach proposed in [19] is the development of a communication and command protocol among the subsystem controllers that (i) will guarantee logically correct and robust behavior across the entire system, (ii) without taxing excessively the operational flexibility that would have been (theoretically) supported by a more cumbersome, monolithic, centralized design.

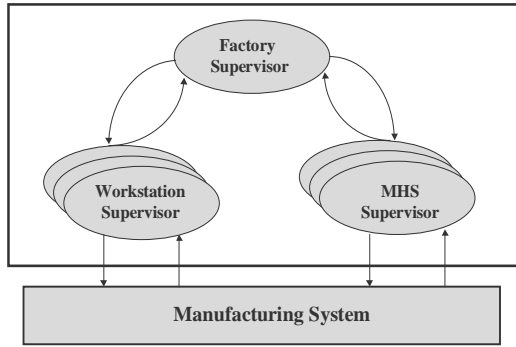


Figure 8 – A hierarchical structural control architecture

Accommodating operational contingencies through system reconfiguration and routing flexibility

The logical characterization and control of the behavior of a flexibly automated production system through RAS modeling, provides also a structured way for organizing the system reaction to various operational contingencies, like the breakdown of a piece of equipment, or the arrival of an expedient job. From the SC perspective, the resulting problem is defined by the requirement to meet any additional contingency imposed restrictions with the minimal disruption, while maintaining consistency to the applied SCP. Analytically, the incurred disruption is measured by the number – or more generally, the “value” – of the jobs that might have to be unloaded temporarily, in order to maintain compliance to the applied SCP, under the new operational conditions. Additional mechanisms that can be employed in the accommodation of the occurring contingencies through the minimum possible disruption are: (i) the establishment of new routes/process plans for the remaining stages of the currently processed jobs (assuming that the system configuration presents some routing flexibility), and/or (ii) the reconfiguration of the applied SCP itself, by appropriate restructuring of the policy imposed constraints, to fit better the new system structure (this is more applicable in the case of equipment breakdowns, since they are essentially changing the underlying RAS structure). A detailed analytical treatment of the problem can be found in [20].

INTEGRATING STRUCTURAL WITH PERFORMANCE-ORIENTED CONTROL

Remember that the key “mission” of structural control is to define the admissible operational space over which the performance-oriented control will seek to optimize the system behavior. A natural implementation of this statement in an event driven, closed-loop real-time control architecture is provided in Figure 9.




Figure 9 – The proposed real-time control architecture

As it is indicated in this figure, the current resource allocation status, combined with the system physical and/or logical operational configuration, defines the set of *feasible* actions. This set is subsequently filtered through the applied SCP, in order to determine the set of *admissible* actions. Finally, it is the responsibility of the performance-oriented control policy to pick the particular admissible action that will be eventually dispatched for execution. At that point, the controller updates also its internal representation of the system state, and the whole decision making cycle is repeated over the new state, upon the reception of the next triggering event.

From an implementational standpoint, the control framework proposed in Figure 9 can easily accommodate the control logic of (i) most commonly used dispatching rules [21] – e.g., Last Buffer First Serve, Shortest Remaining Processing Time, Earliest Due Date, Least Slack, etc. – or (ii) the more recently proposed Non-Idling-Non-Exceeding tracking policies [22]. However, some recently obtained results [23] render questionable the efficiency of these policies in the considered operational context. Hence, the detailed implementation of the performance-oriented control logic that will lead to the most efficient operation of the system, and to the most effective exploitation of its operational flexibility, is an open research issue.

CONCLUSIONS AND FUTURE RESEARCH

The starting point for the research program outlined in this paper has been the observation that currently we lack an adequate methodology for the systematic deployment and operation of flexibly automated production systems, to a large extent, due to the failure of past research to address the logical modeling, analysis and control of the system behavior. To this effect, an entire new area in the real-time control of these systems has been developed, collectively known as their structural control. Currently, the most prevailing issues in the SC of flexibly automated production systems has been the effective resolution of the manufacturing system deadlock, and the development of reaction strategies for the accommodation of different operational contingencies. The paper surveyed the key research results w.r.t. these two issues, and it also proposed a real-time control architecture for the integration of structural and the more “traditional”, performance-oriented control.

The ultimate objective of this research program is the development of a complete production planning and control framework for these environments, based on the typical hierarchical decomposition control paradigm. However, in order to guarantee the feasibility and relevance of the developed control scheme, our design has adopted a bottom-up approach, which differs from the top-down approaches typically taken by past research in the area (e.g., [4,5]). Specifically, our approach seeks to (i) initially develop high fidelity representations for the controlled system, (ii) elicit logical and robust behavior from it through the implementation of effective structural control, (iii) evaluate and materialize the system production capacity through the application of pertinent performance-oriented control, and finally, at a higher level, (iv) use the established capability to measure and control the system behavior and performance, in order to track the externally imposed demand, in the most efficient manner. A visual representation of the pursued approach is provided in Figure 10.




Figure 10 - The proposed bottom-up approach to production planning and control in flexibly automated production systems

REFERENCES

1. M. Weiss, “*Semiconductor Factory Automation*”, **Solid State Technology**, pgs 89-96, 1996.
2. S.B. Joshi, E.G. Mettala, J.S. Smith and R.A. Wysk, “*Formal models for control of flexible manufacturing cells: Physical and System models*”, **IEEE Trans. On Robotics & Automation**, vol. 11, pgs 558-570, 1995.
3. F.F. Suarez, M.A. Cusumano and C.H. Fine, “*An empirical study of manufacturing flexibility in printed circuit board assembly*”, **Operations Research**, vol. 44, pgs 223-240, 1997.
4. K.E. Stecke, “*Design, planning, scheduling and control problems of flexible manufacturing systems*”, **Annals of Operations Research**, vol. 3, 1985.
5. S.B. Gershwin, “*Hierarchical flow control: A framework for scheduling and planning discrete events in manufacturing systems*”, **Proceedings of the IEEE**, vol. 77, pgs 195-209, 1989.
6. J.E. Hopcroft and J.D. Ullman, “*Introduction to Automata Theory, Languages and Computation*”, **Addison-Wesley**, Reading, MA, 1979.
7. M.R. Garey and D.S. Johnson, “*Computers and Intractability: A guide to the theory of NP-completeness*”, **W.H. Freeman & Co.**, NY, NY, 1979.

8. S.A. Reveliotis, M.A. Lawley and P.M. Ferreira, “*Polynomial Complexity Deadlock Avoidance Policies for Sequential Resource Allocation Systems*”, **IEEE Trans. on Automatic Control**, vol. 42, pgs 1344-1375, 1997.
9. M.P. Fanti, B. Maione, S. Mascolo and B. Turchiano, “Event-based feedback control for deadlock avoidance in flexible production systems”, **IEEE Trans. on Robotics & Automation**, vol. 13, pgs 347-363, 1997.
10. M.A. Lawley and S.A. Reveliotis, “*Deadlock Avoidance in Sequential Resource Allocation Systems: Hard and easy cases for a new subclass*”, **working paper**, 1999.
11. Z.A. Banaszak and B.H. Krogh, “Deadlock avoidance in flexible manufacturing systems with concurrently competing process flows”, **IEEE Trans. on Robotics & Automation**, vol. 6, pgs 724-734, 1990.
12. M.A. Lawley, S.A. Reveliotis and P.M. Ferreira, “*FMS structural control and the Neighborhood policy, Part 1: correctness and scalability*”, **IIE Trans.**, vol. 29, pgs 877-887, 1997.
13. M.A. Lawley, S.A. Reveliotis and P.M. Ferreira, “*A correct and scalable deadlock avoidance policy for flexible manufacturing systems*”, **IEEE Trans. on Robotics & Automation**, vol. 14, pgs 796-809, 1998.
14. M.A. Lawley, S.A. Reveliotis and P.M. Ferreira, “*The application and evaluation of Banker’s algorithm for deadlock-free buffer space allocation in flexible manufacturing systems*”, **Intl. Jrnl of FMS**, vol. 10, pgs 73-100, 1998.
15. M.A. Lawley, “*Deadlock avoidance for production systems with flexible routing*”, **IEEE Trans. on Robotics & Automation** (to appear).
16. J. Ezpeleta, J.M. Colom and J. Martinez, “*A Petri-net based deadlock prevention policy for flexible manufacturing systems*”, **IEEE Trans. on Robotics & Automation**, vol. 11, pgs 173-184, 1995.
17. S.A. Reveliotis and M.A. Lawley, “*Efficient implementations of Banker’s algorithm for deadlock avoidance in flexible manufacturing systems*”, **Proceedings of ETFA’97**, pgs 214-220, IEEE, 1997.
18. S.A. Reveliotis, “*Conflict resolution in AGV systems*”, **Tech. Report**, School of Industrial & Systems Eng. Georgia Institute of Technology (submitted to IIE Trans.), 1998.
19. J. Park, S.A. Reveliotis, D.A. Bonder and L.F. McGinnis, “*Event-based control of flexibly automated manufacturing: Integrating logical and performance control*”, **working paper**, 1999.
20. S.A. Reveliotis, “*Accommodating FMS operational contingencies through routing flexibility*”, **IEEE Trans. on Robotics & Automation**, vol. 15, pgs 3-19, 1999.
21. P.R. Kumar, “*Scheduling semiconductor manufacturing plants*”, **IEEE Control Systems Magazine**, vol. 14 (6), pgs 33-40, 1994.
22. A. Sharifnia, “*Stability and Performance of Distributed production control methods based on continuous-flow models*”, **IEEE Trans. on Automatic Control**, vol. 39, pgs 725-737, 1994.
23. S.A. Reveliotis, “*The destabilizing effect of blocking due to finite buffering capacity in multi-class queueing networks*”, **IEEE Trans. on Automatic Control** (to appear).

BIOGRAPHY

SPYROS REVELIOTIS, Assistant Professor, has been with the School of Industrial & Systems Engineering, of the Georgia Institute of Technology, since the fall of 1996. Dr. Reveliotis holds a Ph.D. in Industrial Engineering from the University of Illinois, and an M.S. in Computer Systems Engineering from Northeastern University. His research interests are in the area of Discrete Event Systems theory and applications, with special emphasis on the control of flexible automation. He is a member of IEEE, IIE and INFORMS. He is also the recipient of a number of distinctions and awards, the more recent being the 1998 IEEE Intl. Conference on Robotics & Automation Kayamori Best Paper Award, for his

paper “*Accommodating FMS Operational Contingencies through Routing Flexibility*”.