Proceedings of the 2000 International Material Handling Research Colloquium

RETHINKING WAREHOUSE DESIGN RESEARCH

Leon F. McGinnis, Marc Goetschalckx, Gunter Sharp,
Doug Bodner, T. Govindaraj
Keck Virtual Factory Lab
School of Industrial and Systems Engineering
The Georgia Institute of Technology
Atlanta, GA 30332-0205

ABSTRACT

The Keck Virtual Factory Lab (KVFL) was created to be a platform for computationally-based research on industrial logistics systems, and currently engages nine faculty and over a dozen graduate students. Initially, the KVFL has focused on warehousing and the development of integrated computational tools to support warehouse design and optimization, and warehousing-related courses. Recent collaboration with expert warehouse designers has caused a rethinking of the warehouse design process and the research likely to have the greatest impact on practice. This presentation will describe briefly the KVFL, present a conceptual framework for warehouse design research, and describe several ongoing research projects addressing tools for warehouse design.

1.0 INTRODUCTION

Since 1990, The Material Handling Research Colloquium has provided an international forum for researchers to present their latest work in the area of designing, implementing, optimizing, and improving material handling applications. An often unstated assumption is that material handling is to support the larger industrial systems of which material handling is a part.

Without question, these are important issues. Material handling systems are critical to modern enterprise; without effective, efficient material handling, firms experience excessive operating costs, excessive investments in materials, work-in-process and finished goods inventories, product damage in handling, slow customer response, and poor customer service.

In 1998, the National Research Council (NRC) published a report—*Visionary Manufacturing Challenges for 2020* [8]—that identified several “grand challenges,” among them the achievement of “reconfigurable enterprises,” i.e., the ability for firms to reconfigure rapidly to meet changing needs, opportunities, or threats. Achieving this high degree of reconfigurability in manufacturing firms carries with it some implications for supply chains, not the least of which is a growing pressure for warehouses to be similarly responsive to changing requirements. Thus designing, redesigning, and managing warehouses will grow ever more important to the success of the overall supply chain.

Among the nine “strategic technology” areas identified in the NRC report was “enterprise wide modeling and simulation.” Nowhere is this need more apparent than in the modeling and simulation of enterprise logistics, including the material handling systems and their controls. The Material Handling Research Colloquium, for the past decade, has been a major forum for the presentation of new modeling and simulation results that respond directly to the NRC's grand challenges.

The research presented at the Colloquium often is of the highest quality, and often leads to archival journal publications. A cursory review of prior Colloquium proceedings reveals the breadth and depth of the research presented. This research community can be proud of the progress it has made in the past decade.

And yet, there is clearly room for improvement. If we look at the current state of practice, we are hard pressed to find specific examples where specific research results find widespread deployment. If we examine the software tools that have widespread application in the design and analysis of material handling systems, there is little evidence that material handling research, *per se*, has contributed in a fundamental way to that software.

This apparent "disconnect" between the research community and the practice community has been a topic of considerable discussion at previous Colloquia. A consistent conclusion has been that research is not "deployment-ready," i.e., it is not in a form that is usable by the practice community. One reason is that the extant research addressing warehouse design is fragmented, in the sense that isolated parts of the problem are studied in detail without a coherent approach that looks at the bigger picture in context of design.

Our research team in the Keck Virtual Factory decided in 1999 to directly address these issues in the context of warehouse design. Warehouse design requires a number of decisions that are classical topics in material handling research, including the design of storage systems, the specification and management of lift trucks and AGVs, the design and management of pick-lines, order-pick management systems, and the design of sortation systems. We wanted to answer the question "What must we do in warehouse design research to make a significant impact on practice, i.e., to see research results widely deployed to improve the design of warehouses?"

The process we initiated to try to answer this question is still ongoing, but we have reached some initial conclusions that we believe are worth sharing with the research community, and discussing more broadly. In the following pages we will describe the process, present the significant initial conclusions, and describe the research projects we have initiated as a result.

1.1 The Process

The question we asked was, "What must we do in warehouse design research to make a significant impact on practice?" To answer this question, we set out to establish a Warehousing Roundtable, which would bring practitioners together with the Keck Virtual Factory Lab team. The Warehouse Roundtable would meet on a fairly regular basis—approximately quarterly—to discuss problems, practices, and research results.

There have been three roundtable meetings to date, attended by practitioners from firms in design/consulting, WMS vendor, and user categories. Each roundtable meeting is characterized by generic presentation of an issue, intense brainstorming, and summarization of conclusions. The roundtable meetings typically have lasted four to six hours, with a working lunch.

Topics addressed to date in the Warehouse Roundtable have included the overall process of warehouse design, understanding the warehouse requirements, and representation of the warehouse.

1.2 Preliminary Conclusion

On the basis of the experience to date with the Warehouse Roundtable, we have reached one significant (albeit preliminary) conclusion.

> ***The relative lack of impact on practice from material handling research (in the context of warehouse design) is not necessarily because the research is not "user ready;" rather, the research, by and large, addresses problems that expert warehouse designers consider secondary, while failing to address problems they consider primary.***

This conclusion is based on our current understanding of the process followed by expert warehouse designers, which we will outline below. Based on this conclusion, we have initiated several specific research activities aimed at:

1. Formalizing our understanding of the process experts follow in designing warehouses;
2. Developing a suite of computational tools to support experts in their primary warehouse design decision-making
3. Developing an integrating framework for engineering tools to support the secondary design decision-making in a manner that is consistent with primary design decisions.

These research projects also will be briefly described, with regard to their activities and objectives.

2.0 EXPERT WAREHOUSE DESIGN

Consider the scenario in which a warehouse is to be designed to support an existing business. Suppose the client (the prospective warehouse owner) has extensive data on the items to be stored, and on the historical pattern of customers orders. The expert warehouse designer begins by carefully examining the data to identify patterns in the nature of the items stored in the warehouse and the nature of the customer orders. These identified patterns then become key factors in the primary warehouse design decisions. (It may be the case that the client has no existing data, in which case the expert designer attempts to ascertain what the patterns might have been if the client had data to be analyzed. This point is explored further below.)

How does the expert designer analyze the data, and what are the primary design decisions? These are not simple questions, and our answers are, at this point, preliminary and relatively subjective. As discussed below, we are in the process of a more formal and objective analysis of expert warehouse design.

We believe the expert designer conceptualizes the warehouse, at least implicitly, as a network of functions, as illustrated in Figure 1. In the figure, each node corresponds to a specific warehouse function, which may be receiving, a storage function, a value-adding function, a pick-line, a sortation/accumulation, a pack function, shipping, or some other relevant function. Storage functions may be distinguished in terms of the technology used, or the types of products stored. The arrows connecting the nodes in the figure represent the aggregate flow of material between the functions. We refer to this network as the ***function flow network,*** or FFN, for the warehouse. Note that the function flow network can be readily identified for an existing warehouse.

The primary warehouse design decisions are the specification of the function flow network, i.e., identification of the nodes in the network, and the arcs connecting them. In other words, the primary decisions are “architectural,” having to do with the “shape” of the function flow network, and the selection of the technologies for each function. Clearly, these architectural decisions are driven by the nature of the items that are expected to flow through the warehouse, the nature of the orders to be shipped, and the nature of the warehouse replenishment process. A host of other factors are important as well, such as the availability and cost of suitable labor, the cost and location of the facility, the level of automation or information integration that is desired, and the projected future loads on the warehouse.

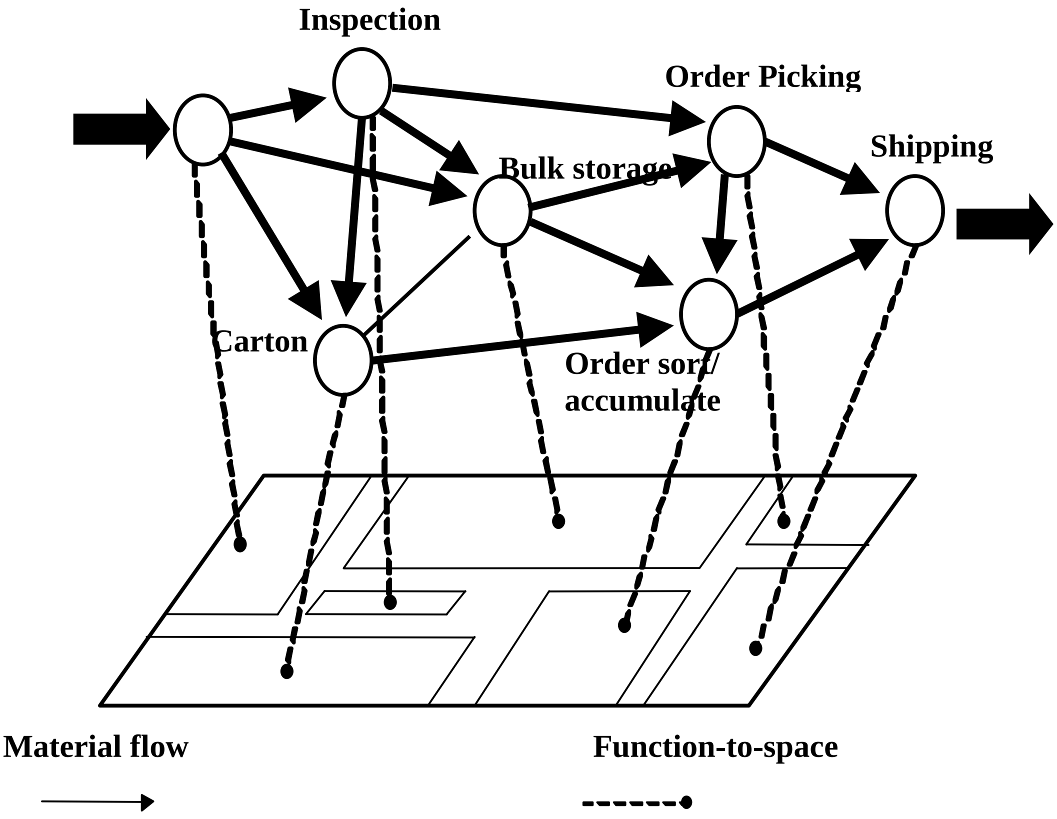

Figure 1 Illustration of Function Flow Network

Once these architectural decisions have been made, the remainder of the design involves detailed engineering. For example, the specific set of items to be assigned to the pick-line must be determined, along with the allocation of space to each item. Or the specific items to assign to a specific storage technology must be determined, and the size and configuration of the technology specified. This process is illustrated in figure 2, which emphasizes the iterative nature of the expert design process.

Figure 2 illustrates four elements of the expert warehouse design process. ***Profiling*** is the assessment of the requirements the warehouse must satisfy, and usually involves careful analysis of historical data on customer orders and the items to be stored. More is said about profiling in McGinnis and Mulaik [7]. In figure 2, ***architecture*** refers to the specification of the function flow network, i.e., the structure of the network and the technologies assigned to each node, as well as the partitioning of the items into item families and the assignment of item families to storage functions. ***System specification and optimization*** is the detailed specification of each function (and perhaps each item family). The element identified as ***functional space*** is intended to represent what is, today, largely an intuitive aspect of expert warehouse design. It appears that experts retain a great deal of information, some quantitative and some qualitative, gleaned from prior designs or other relevant experiences. This intuitive or implicit data is used along with results from profiling analysis to reach architectural decisions. Note that this view of the warehouse design process is different from, but not necessarily inconsistent with that presented by Mantell and Rouwenhorst [6].

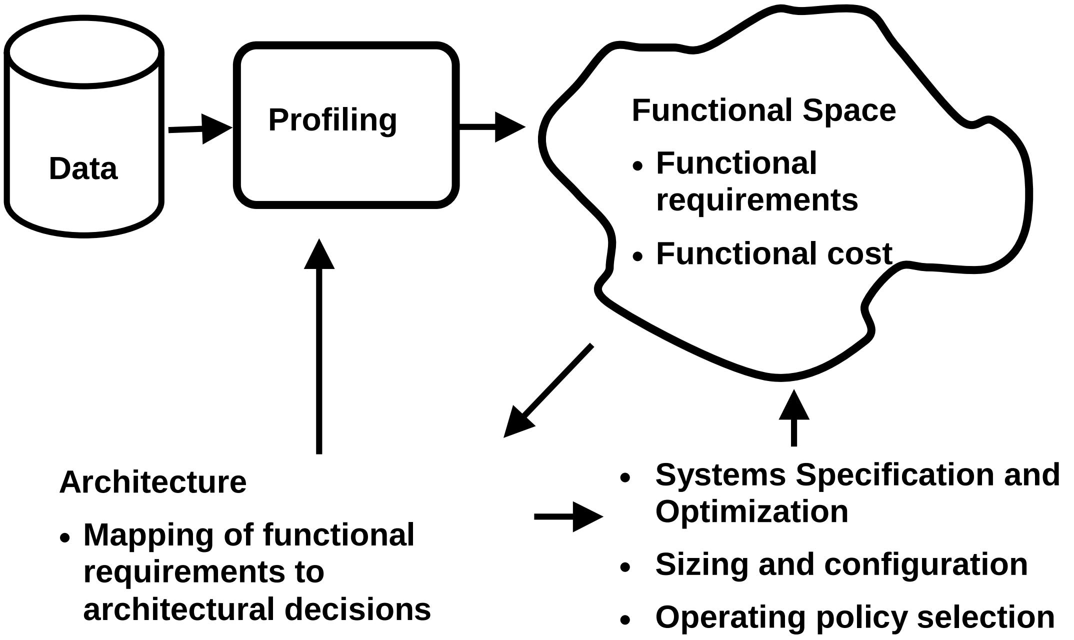

Figure 2 Expert Warehouse Design Process

We suspect that for many, if not most warehouses, the criterion function for evaluating design is relatively "flat" in a fairly large neighborhood around the optimum specification for the individual functions. In other words, marginal changes in the number of storage locations, the footprint and configuration for a storage zone, the assignment of items to the pick-line, etc., will not have a dramatic impact on the quality of the design. This, of course, is conjecture, and remains to be verified or refuted.

In contrast, decisions such as whether or not to have a pick line, or to have more than one storage technology, or the specific selection of storage technology do have a potentially significant impact on warehouse cost and performance. This is the primary focus of the expert designers with whom we have worked. They also appear to believe that the detailed design of the individual warehouse functions, while important and perhaps intricate, are relatively straightforward, once the higher level architectural decisions have been made. Of course, as the details of the design are elaborated, it may become desirable or necessary to revisit some of the architectural decisions, for example, to add an additional storage technology, or to consolidate storage from two technologies into one.

It appears also that expert designers recognize that the detailed specification of the individual warehouse functions is based on data that is, at best, subject to considerable uncertainty and error. Furthermore, it often is possible to make accommodations in the operation of the warehouse for errors in the design. For example, if the storage capacity is discovered to be inadequate, operations often can reduce the amount of inventory required by replenishing more frequently, or use storage more effectively by combining low volume items in a single storage location. All three considerations support the notion that the architectural decisions are the most critical, while the detailed function specifications are important but less critical.

3.0 CONCEPTUAL FRAMEWORK

Together, figures 1 and 2 provide a conceptual framework for warehouse design as the specification of the function flow network, and the elaboration of each function and flow, employing a process of profiling, architectural design, and detailed engineering. In this framework, ***functions*** correspond to the transformation of material, either by receiving it (and breaking it down), changing its configuration by sorting it from other material or accumulating it with other material, adding value to it, storing it, packing it, or shipping it. In this framework, a function has a corresponding area and/or equipment in the warehouse where the function is performed. Thus, specifying the function involves specifying the technology to be used, and eventually, the details of the implementation, i.e., the configuration of the technology and perhaps associated control functions.

Flows between warehouse functions correspond to warehouse ***operations***. For example, put-away is an operation that moves material from the receiving function to the storage function. Order picking is an operation that moves material from a storage function or pick-line function to a sortation or accumulation function. An operation may require equipment, for example, lift trucks, order picking trucks, conveyors, or other equipment. It also may require other resources, such as labor or computer systems. In addition, operations will require a ***method***, i.e., a system of planning, managing, and controlling the operation. For example, an order picking operation may be specified to use single order picking or batch picking with picking zones.

The process of design involves translating some description of the “to be” warehouse requirements into a specification of its function flow network, i.e., of its functions and operations. In the design process, the understanding of requirements comes from profiling. The function flow network description requires, at a minimum, identifying the functions, the technology for each function, the identification of the items (item types or item families) for each flow (or operation), and the corresponding technologies and methods associated with each flow. Detailed engineering of each function and operation will complete the design by specifying the configuration of technologies, including their footprint in the warehouse, and any essential operating policies or control rules.

Within this conceptual framework, there are at least two distinct strategies for research to support warehouse design. One strategy would be to develop a broad range of specific decision support tools for the detailed engineering design of the functions and operations. That is, to some degree, the strategy inherent in much of the existing warehouse research, and it places the burden on the user to understand the research and to find a way to bring the research results into the warehouse design process. In other words, the predominant approach in the past leaves the issue of integration to the potential user of the research.

A quite different strategy is to attempt to support design decisions making from the initial profiling of client requirements, through architectural design, to detailed design. In other words, this approach explicitly attempts to integrate concepts and tools for the potential user. This is the strategy currently being pursued by the Keck Virtual Factory Lab.

4.0 KVFL RESEARCH AGENDA

The Keck VFL team has initiated a set of related research projects, based on the conclusions outlined above. Working with our industry partners, our goal is to develop the conceptual models, the methodology, and the computational tools that will find acceptance in industry, and make a significant difference in the way warehouse design is performed. Four specific projects address that goal, and each is described below.

4.1 Ethnographic Studies

Recognizing that the informal conclusions we’ve reached about the process of warehouse design are inadequate as a foundation for future research, we are conducting a two year study of expert designers at work. The project is supported by a grant from the National Science Foundation, and will “seek first to understand industrial logistics system design as practiced by experts, and then to propose a conceptual framework for industrial logistics system design, and finally to explore specific key elements of a computational platform for industrial logistics system design.” The project addresses three principal issues:

(1) How do expert designers conceptualize the industrial logistics system design problem; what concepts or abstractions do they employ as they formulate the key design decision problems? We believe the "function space model" illustrated in Figure 2 is a valid and useful framework for understanding and organizing the warehouse system design process. In conducting the proposed research, we expect either to validate the model, or to modify and elaborate it as necessary.

(2) How do expert designers formulate alternatives for the key design decisions; what information (both explicit and implicit) and criteria are used? Not only will we catalog the answer, we will explore the development of formal data models for

creating and manipulating alternative decisions and their instantiation in generic databases.

(3) What criteria do expert designers use in selecting among design alternatives and how do they evaluate each alternative; how do they test the resulting system designs?

At the conclusion of this project, we will have two important results. First, we will have a documented methodology for studying expert warehouse design, which can be repeated for other design projects, obtaining results that can be compared and correlated. Second, we will have explicit models for the way that at least one group of expert designers conceptualize the problem, formulates alternatives, and evaluates alternatives. This knowledge forms the foundation for academic research that can truly serve the needs of practicing designers, and also the basis for teaching warehouse design to students.

4.2 Profiling

The term "profiling" is used here to represent the process of developing a description of the warehouse requirements which incorporates the relevant characteristics of both the items to be handled, and the orders to be filled. In the ideal situation, warehouse design begins with a SKU master, a database describing each item in detail, and an order master, a database with a representative sample of customer orders for, say, one year. These two databases could be sorted, queried, partitioned, and statistically analyzed to answer such questions as:

(1) Can the items be grouped according to size? Weight? Type (i.e., flammable, refrigerated, etc)? Frequency of orders? Size of orders?

(2) Can the orders be categorized according to number of lines? Size? Weight? Value?

(3) Can the customers be grouped according to frequency of orders? Value of orders?

Clearly, these are not difficult questions to answer, at least for an *ad hoc* analysis, since they require at most queries to a database and creating distributions or perhaps Pareto charts. More discussion of profiling may be found in McGinnis and Mulaik [7]. One element of the profiling project is simply to create a generic database tool that can be applied to any given data for the purpose of answering the profiling questions.

Where profiling becomes more interesting, and perhaps even more important, is when either the available data is "dirty," i.e., incomplete or containing errors, or when there is no data available, e.g., for a dot.com startup company. In the case of "dirty data," what is needed is a set of methods that can be applied to "rectify" the data, i.e., to remove any spurious information, and to fill in estimates where possible.

When there is no data, warehouse design proceeds with "synthetic results" from profiling, i.e., the designers either explicitly or implicitly answer the questions posed above in order to generate an architecture and a detailed design. This presents an interesting research question, namely, "Is it possible to develop a synthetic order generator that can be driven by a relatively small set of parameters, yet create synthetic orders that are indistinguishable from "real" orders?" If so, then understanding which parameters are needed for the synthetic order generator enables designers to ask the right questions of clients who have no data.

The profiling project will yield both a generic database tool for rectifying and analyzing data, and a methodology for generating synthetic orders.

4.3 Design Tool Integration

Making the results of warehouse design research accessible to prospective users requires solving at least two types of problems:

(1) The absence of a standard "reference model" for warehouse design, i.e., a standard way of representing a warehouse design in terms of its data description; and

(2) The absence of an existing computational infrastructure into which a range of modeling and analysis tools can be integrated.

Without a reference model for warehouse design, there is no uniform common frame of reference for prospective users. Inconsistent terminology and *ad hoc* data organization would make it literally impossible to integrate various design tools at the conceptual or methodological levels. Even if a reference model is in place, there remains the daunting challenge of creating a corresponding computational infrastructure that will enable the integration of standard data, user data, standard design/analysis tools, and *ad hoc* design/analysis tools.

The KVFL team is approaching this challenge by adopting an industry standard approach—UML or Unified Modeling Language [5]—and a commitment to developing an infrastructure that is distributed using web technologies, object-oriented, open, and extensible.

The UML methodology applied to the software integration of warehouse design tools first requires an analysis of the interactions that a user would have with the software system being developed. These "use cases" are then the basis for beginning to design the objects, or class hierarchy, for representing the warehouse design and the design functions or tools with which the user interacts. The KVFL team has developed an initial set of use cases, and initial object definitions. These are described in greater detail in Amirhosseini et al [1].

Within the KVFL, we are exploring a wide range of internet technologies that hold promise for enabling the deployment of new design/analysis tools to a broad audience, both in industry and in education and research. We are committed to a client/server model of computing, supporting the distribution of user and design data among users and a variety of server. While we recognize the need to develop some software solutions, such as the infrastructure itself and the implementation of new analysis or design algorithms, we also are committed to incorporating as much commercially available software as possible. So for example, work in the Keck Lab currently utilizes a wide range of tools, from application development tools such as WebObjects™ and Visual Café™, to databases such as Access™ and PostGreSQL™, communication protocols such as CORBA and Orbacus™, simulation, optimization, and visualization via VRML. More detail on our use of internet technologies may be found in Bodner, et al [3, 4].

4.4 System Design Decision Optimization

In warehouse design, there are many trade-offs that should be considered. More automation may reduce the labor cost, but increase the investment cost. More storage zones may improve the space utilization, but require more control. Zone picking may improve the productivity of the stock selectors, but requires more planning and control of picking. Trade-offs like these abound.

Ideally, it would be possible to formulate an optimization model that would consider the trade-offs between all the available options and select the decisions leading to the "best" solution, according to the specified criterion. This optimization model would enable us to deal explicitly with customer service levels, either as a constraint or as part of the criterion, which means it would require a formal definition and computable metrics customer service. Such an ideal formulation may be possible some day, but for the time being, we need more pragmatic approaches.

Consider, for example, the design of three key functions in the warehouse—the reserve storage, pick-line, and sortation/accumulation functions. Under what assumptions can we formulate each functional design problem using consistent definitions? If we can create "uniform" formulations, can we then identify a relatively small set of parameters, decisions, or performance values that can be used to coordinate the solutions of the three functional design problems?

These are the questions that motivate our current efforts on warehouse design decision optimization. We are seeking to develop a heuristic procedure that will allow us to coordinate the (sub)optimization of a set of functions in the warehouse.

5.0 SUMMARY AND CONCLUSIONS

Warehouse design and operation has been a rich source of research questions and problems for the past twenty years, and warehousing research is a major component of the archival literature in the IE/OR disciplines. As the NRC report [8] makes clear, the importance of this research to practice will only increase in the future. Building upon that vast resource, and exploiting recent breakthroughs in computing technology, the Keck Virtual Factory Lab is defining a research agenda that will lead to widespread deployment of research results in the field. The agenda is based on rethinking of warehouse design research, considering a user's perspective, to identify the intellectual and engineering issues that must be resolved to enable deployment. Key elements of the agenda include developing reference models, creating integrative computational infrastructure, developing specific components of the integrated solution, and, of course, continuing to refine our understanding of the real problem of warehouse design through interaction with industry partners.

6.0 ACKNOWLEDGEMENTS

The authors wish to acknowledge the W. M. Keck Foundation for its funding of the facilities and activities of the Keck Virtual Factory Lab at Georgia Tech. In addition, this work has been partially funded by UPS Worldwide Logistics, the Ford Foundation, the Georgia Tech Foundation, and by the National Science Foundation under Grant No. DUE 9950301.

7.0 REFERENCES

1. Amirhosseini, M., D. Bodner, M. Goetschalckx, and T. Govindaraj, “Object Models and Design Databases for Warehouse Applications,” IIE Solutions 2000, May 23, 2000, Cleveland, OH.
2. Bodner, D. A., T. Govindaraj, L. McGinnis, “Industrial Logistics Systems Design: Conceptual and Computational Foundation,” proposal to National Science Foundation, October 1, 1999.
3. Bodner, D. A., T. Govindaraj, E. Blanco, M. Goetschalckx, L. McGinnis, G. Sharp, “Web-Based Tools In Warehouse Design,” IIE Solutions 2000, May 23, 2000, Cleveland, OH.
4. Bodner, D.A., T. Govindaraj, E. Kemahlioglu, E. Blanco, M. Goetschalckx, L. McGinnis, and G. Sharp, “Using Internet Technology for Design of Facilities and Material Handling Systems,” Material Handling Research Colloquium 2000, York, PA, June 12-14, 2000.
5. Fowler, M. and K. Scott. *UML Distilled, Second Edition*, Reading, MA: Addison-Wesley, 2000.
6. Mantel, R., and B. Rouwenhorst, “A Warehouse Design Model,” *Progress in Material Handling Research: 1998*, Graves, et al (eds), Braun-Brumfield, Inc., Ann Arbor, MI, 1998.
7. McGinnis, L., and S. Mulaik, “Your Data and How to Analyze It,” IIE Solutions 2000, May 23, 2000, Cleveland, OH.
8. National Research Council, *Visionary Manufacturing Challenges for 2020*, Washington, DC: National Academy Press, 1998.