Issues in Agent-based Coordination of Public Logistics Networks

Michael G. Kay and Ashish Jain

Department of Industrial Engineering
North Carolina State University
Raleigh, NC 27695-7906

October 21, 2002

Abstract—A public logistics network is proposed as means to provide fast, flexible and low-cost parcel transport. Similar to the dynamic pricing used to sell airline seats, a price for each available space on a truck and storage space at a distribution center (DC) could be negotiated in real time for each individual item. A third party can search the network for any type of item in transit, negotiate the purchase of the item, and redirected it to a new destination. As compared to private logistics networks, scale economies could be realized in performing each logistics function since each element of the network has access to potentially all of the network's demand. It would be possible for small cities and towns to have access to the same low-cost parcel delivery services that are currently only available in larger cities, and the need for on-site inventory storage would be reduced since finished-goods and safety stocks would be mobile and could be re-routed only when needed. An important research issue is to determine what is the best combination of market and pricing mechanisms for the transport services in a public logistics network. This report describes some of the issues involved in research that is just starting to design architecture and protocols that can be used to coordinate the operation of this type of network in order to facilitate adaptive routing and in-transit trade.

1. Introduction

A public logistics network is proposed as a means to extend many of the features associated with public warehouses to the entire supply chain. In addition to providing traditional warehousing and storage functions for hire, a public logistics network would make it possible to negotiate with multiple firms on a load-by-load basis in order to determine the most efficient means of providing the resources needed to complete each stage of a load's transit through the network. Items could continuously negotiate with the logistics resources of the network using simultaneous auctions in order to determine the best route, cost, and schedule. Similar to the dynamic pricing used to sell airline seats, a price for each available space on a truck and storage space at a distribution center (DC) could be negotiated in real time for each individual item. A unique capability of such a network is that a third party can search the network for any type of item in transit. Once located, negotiations can take place and the item might be resold to the third party and redirected to a new destination. The potential utility of this search and negotiate capability depends on the characteristics of items being transported: it is not likely to be needed to locate low-cost, ubiquitous items like toothbrushes because they can be expected to be available at every local store; nor is it needed to locate custom-made, one-of-a-kind products because there are so few of the items available, of uncertain quality, that the use of traditional private logistics networks is likely to still be efficient. A public logistics network is likely to be most suitable for managing the multitude of commodity-like items (replacement parts, etc.) that fall in the middle ground between ubiquitousness and uniqueness. See Kay and Parlikad [1] for a more detailed discussion of public logistics networks, and Parlikad [2] and Gandlur [3] for some initial research in this area.

Figure 1 shows an example of how a public logistics network might be established. Focusing on the southeastern United States, a total of thirty-six interstate DCs could cover the region. Each DC would be located next to an interstate highway interchange in order to enable direct access to and from the DCs adjacent to the DC. Each of the interstate DCs in Figure 1 would serve as a hub in a sub-network of local DCs (not shown) covering the region surrounding the interstate DC.

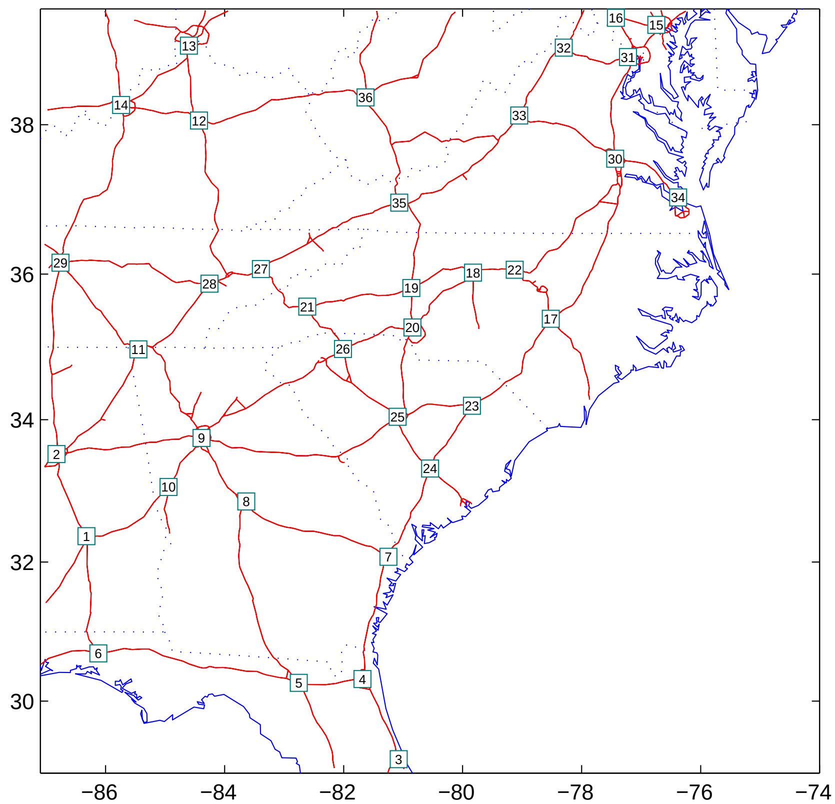

Figure 1. Hypothetical public logistics network showing 36 public DCs covering the southeastern portion of the USA and connected via interstate highways.

2. Packages vs. Packets

The goal of the research is "to do for transportation what the Internet has done for communication." Using the analogy between the packages transported in a public logistics network and the packets transmitted in the Internet, the research is focused on implementing a packet-switched-like, connectionless network for production and distribution. The motivation for this work comes from the observation that progress in the development of new types of communication and transportation networks has followed similar paths. Prior to around 1840, the most common type of communication network was the mail system, a network that could be characterized as providing connectionless communications and having distributed control. At this same time, the most widespread type of transportation network used in production and distribution included small-scale manufacturers at one end selling goods to all-purpose general merchants who passed the goods to individual peddlers at the other end. Similarly to the mail system, the transportation network of this time could be characterized as providing connectionless production and distribution and having distributed control. With the development of the telegraph and the railroad around 1840 and the telephone a few years later, the scale of the services they provided by communication and transportation networks significantly increased while their costs decreased. These networks could be characterized as connection-oriented and having centralized control.

With the development of the minicomputer around 1965, it became possible to again create a network that provided connectionless communications, this time using packet switching. In contrast to a circuit-switched network like the telephone network, where communication link resources are reserved for each call and control is centralized, packet switching has lowered the cost of communication by enabling efficient sharing of high-speed, high-capacity communication links and by allowing distributed control. There has not yet been a corresponding development of packet-switched-like transportation networks. Although companies like FedEx and UPS utilize the Internet and have very sophisticated proprietary tracking and control infrastructures, these firms can still be characterized as providing connection-oriented distribution services because a single firm handles a package throughout its transport and the control of the network is highly centralized. Even for large loads, it is still most common for a single logistics firm to handle a load throughout its transport. The most notable features of these connection-oriented networks are that a single firm controls the network and that much of the technology used to coordinate the operation of the network is proprietary. The principal competitive advantage that these companies have is that a very large scale of operation (national or international) is required in order to be able to underwrite the development of propriety technologies. Nevertheless, a single firm, unless it becomes a monopoly, is ultimately limited in the scale of its operation.

A question then arises: what would be the impact if much of the coordination of production and distribution networks could be implemented as a connectionless network. In particular, what would be the impact of making packet-switched-like transportation networks an alternative means for coordinating production and distribution? The most salient impact is likely to be it would make it possible to separate the different functions of the network so that a single firm is not required for coordination. This would enable scale economies to be realized in performing each logistics function since each element of the network has access to potentially all of the network's demand. In terms of connectionless, packet-switched communications via the Internet, it makes it economic to develop a fiber-optic backbone that would not be justifiable if only a single firm's traffic were using the link; in terms of connectionless, packet-switched transportation networks, the increase in scale might make it economic to ship in full truckloads throughout the network as opposed to more costly less-than-truckload shipments. This could be possible because a single truck could be used to transport all of the demand associated with a lane (or link) in the network. Links in the network could be served by trucks owned and operated by different firms, and each transshipment point (or node) in the network could be an independently operated facility. Due to the increase in scale, it would be economic to have many more transshipment points. Distribution centers and public warehouses could be established in small cities and towns that would never have such facilities if they were served as part of a proprietary production and distribution networks.

3. Research Issues

This research will focus on two specific issues that are critical to the coordination of public logistics networks:

1. **Adaptive Routing:** Suppose the package, on its way to its destination, finds out that there is another route, from its current location to the destination, which is cheaper than the current route. In that case, the package can be re-routed to the new path. It can also be re-routed in the case of natural calamity or any other disruption on its original path. So the objective is that the package reaches its destination with the minimum cost possible and in the case of any disruption on its route, the system can be re-adjusted.
2. **In-transit trade:** In-transit trade is possible if there is an immediate demand of a material (package) at a DC and it is not available near the DC. The DC will then look out for the required material (package), which can be on its way to its current destination. Once the package is located, if the current owner of the package is ready to negotiate for trading and if the prices after negotiations are acceptable to the parties involved, the DC will become the new owner and the package will be re-routed to the location of the DC. The previous owner can also earn some profit from this trading. Therefore, it can be beneficial for all the parties involved. In the case of multiple DCs requiring the same package, auction/bidding can also take place.

The principal issue that will be addressed as part of the development of the proposed adaptive routing protocol is how prices are determined for each logistics resource. Initially, it will be assumed that spot prices for transport along each arc are posted at each DC. As these spot prices change, the package routes will adaptively change. To this point, the research will be similar to much of what has been developed for communications networks. What is more interesting are the following alternatives to spot pricing that will be investigated as part of the major focus of the research:

1. **Spot vs. forward reserve pricing.** What is the impact of being able to reserve a posted forward price as opposed to accepting a spot price whenever a package reaches a DC. If a price is reserved, what penalty costs for cancellation are necessary? Would the creation of some type of insurance mechanism reduce the risk associated with spot pricing?
2. **Gaming the system.** The potential for self-interested agents (in case of agent-based architecture) "gaming" the logistics network is significant. One example would be an independently owned and operated truck. Once it reaches at DC and starts to load packages, it is in the interest of the truck to try to remain at the DC and continue to wait for more packages as long as it is the only truck currently at the DC that is traveling to a particular destination. Although a possible coordination mechanism to eliminate this would be a "rule of the house" that all trucks have to depart after some maximum waiting time or when it reaches a specified capacity, a better approach might be to price the transport of each package so that its price decreases over time; thus, it will be in the trucks best interest, if it has high priority items onboard, to sometimes depart a DC when it is only partially filled. This is just one example of a possible gaming situation.
3. **Options for transport futures.** Is the trading of options [4] for transport futures a feasible hedging mechanism for reducing risk? What is impact of price speculation? Can derivative instruments be created that include an entire route as a single, composite service that is priced using some type of combinatorial auction?
4. **Posted pricing vs. negotiations.** If posted pricing at each DC is not feasible, what other pricing mechanism (e.g., auctions, bidding, etc.) is better, and how the negotiations between different entities like package, trucks, brokers etc will take place?

5. **Off-peak pricing.** Since items can be stored at DCs while they are in transit, it becomes feasible for non-critical items to travel only during off-peak periods. What is the potential impact of this on reducing congestion in the logistics network? How should storage space in the DC is priced?

4. Logistics vs. Communications Routing

The adaptive routing protocol that will be developed as part of the research is similar to "distance-vector" protocols like RIP used in packet-switched communications networks because each node tells its neighbors its via cost (mentioned in the section "Agent-based Coordination – Possible Solution") as opposed to "link-state" protocols like OSPF, where each node tells every other node the cost to its neighbors [5]. The principal differences are the following:

1. In a communications network, the items being routed (packets) are the same items used to control the network; thus, failure of a link or node has to be accounted for in the protocol. The protocol for the logistics network will run on top of the communications network, and communication failures can be ignored.
2. There is no single point of control in a communications network; in the proposed logistics network, the Home Agent provide a single point of control.
3. The cost to only a single destination (the via cost) need be stored at each node as compared to the costs to all other nodes.
4. In the proposed logistics network, items can be temporally stored at the DCs. This capability adds to the flexibility of the resulting routing possibilities.
5. The long time periods between actually having to make a routing decision while the package in on the truck serve to dampen any oscillations that might occur in adaptive routing [6].

5. Future Research

Based on the issues described in this report, the following is a tentative list of some of the future research activities:

1. **National public logistics network.** Expand the 36-DC network developed in [1] to cover the entire U.S. This expanded network will be the base network structure that will be used in all subsequent work.
2. **Spot-price routing.** Develop distributed protocols for adaptive routing using spot pricing. Since a key issue is scalability, the communications impact of using centralized routing versus decentralized routing will be explored.
3. **Pricing/market alternatives.** The issues raised in Section 3 will be investigated as part of the attempt to determine what is the best combination of market and pricing mechanisms for the transport services in a public logistics network. This activity will constitute the major portion of the future research. A Matlab-based simulation using the different mechanisms will be used for the initial evaluation. Determining how the different mechanisms should be compared is not known and will be an important contribution of the research.
4. **Architecture.** Alternatives architectures will be investigated as the means to implement adaptive routing and in-transit trade, including a pure agent-based architecture, a mix of web services and agents, and traditional centralized optimization techniques. Although agents

would likely be the best means to implement the complex negotiations involved in in-transit trade, other means might be more efficient for obtaining simple price quotes.

5. **Implementation.** A computational laboratory containing 25 networked PCs will be used to test a prototype implementation of the adaptive routing and in-transit trade protocols developed for this research, using the network developed in Activity 1, above, to generate demand data.

6. Agent-based Coordination—Possible Solution

This section describes the software agents and the probable approach for the solution of the problem if the agents are used for different coordination activities in the Public Logistics Network.

Intelligent software agents are computer processes "… situated in some environment, and that are capable of autonomous action in this environment in order to meet its design objectives." [7, p. 15]. A variety of different models for the coordination of intelligent agents have been developed recent years, along with a number of associated technologies and applications [8]. Agents provide an ideal means to implement the software infrastructure needed to coordinate the operation of a public logistics network because they can move via the Internet to host server computers located at each DC. As envisioned, the servers at the DC will implement an agent-mediated marketplace for access to the logistics resources located in the immediate vicinity of the DC. By being located at the DC, a package's agent is available to immediately respond to any disruption without concern for communication network latencies—as in programmed stock trading, milliseconds are likely to be critical for effective price negotiations.

Intelligent agents representing each package will negotiate with agents representing each manufacturer, customer, truck, and distribution center in the logistics network in order to minimize its individual transport cost. This makes it possible to focus on the pure transport-related arbitrage opportunities that a public logistics network can provide. In particular, it will be determined whether, in equilibrium, the logistics network does operate in a least cost manner and, most importantly, whether the network can re-optimize through self organization after being subject to a variety of disturbances, ranging from the simple breakdown of a truck to the logistical challenges associated with a major natural disaster (e.g., a hurricane).

The agent-based architecture will address the two specific issues of this research in the following manner:

- **Adaptive routing:** When a package is to be transported through the network, intelligent software agents would be spawned and sent ahead to each DC along the route to the package's destination. Based on the local price of transport and storage at each DC, the "best" route for the package is determined. All of the package agents that are at DCs not along the intended route would remain active until the package reaches its destination in order to be available to determine alternative routes for the package in case of any disruption (cost increase) along its intended route or a significant cost decrease along an alternative route.

- **In-transit trade:** All of a package's agents that are located at the DCs along both its intended and alternative routes are available as trading agents for the possible resale of the package and its subsequent re-routing to a new destination. Triggered updates will provide each trading agent with the minimum cost required to transport the package from its current location to the DC where the trading agent is located; as a result, the exact, accurate transport cost can be added to the FOB (free-on-board) price of the package from its origin, making it possible to instantaneously negotiate using price quotes that represent the delivered price at the DC where the trading agent is located.

6.1. Related Research

Swaminathan et al. [9] have created a generic multi-agent supply chain-modeling framework, with software components for manufacturing and transport agents, control elements, and interaction protocols; unfortunately, the framework is too generic to be directly utilized in the research. Swarm intelligence [10] and ant system optimization [11] have been proposed as an effect means of implementing distributed adaptive routing. It is not a feasible approach for the research because it relies on having agents make repeated visits through a network in order to generate, over time, optimal routes; in a public logistics network, the route for each individual package is likely to be different; thus there is no time to build up a route via ant optimization.

The proposed agent-mediated market for logistics resources is most similar in spirit to the work of Kephart et al. [12] on dynamic pricing by software agents. They investigated various mixtures of automated pricing agents and the interplay between optimization, market dynamics, and optimization in information economies with billions of agents, and have shown how machine-learning can reduce the harmful effects of cyclical price-war dynamics for economies with two sellers. They are currently working on extending their analysis of learning to economies with more than two sellers. Their approach of using the "tools of modeling, analysis, and simulation to study and redesign agent strategies, protocols, and market mechanisms in the laboratory before releasing agents and agent infrastructures into the world's economy" is the same approach that will be utilized in the research.

While Kephart et al. [12] have focused on single-item market mechanisms, the research will focus on the more complex issue of determining the price of a package's route, an example of a "composite good" made up of the combined price of each segment of the route. McCabe et al. [13] have studied the problem of buying natural gas that must be transported though a pipeline network. The transport of the gas is an example of a composite good, similar to the transport of a package in the research. They found that a composite goods market was more efficient than a sequence of bilateral markets. The entire issue of what is the best market mechanism—posted prices, auctions, bidding, and various hybrids—will be one of the most important issues explored in the research (see [14] for a general discussion of market mechanisms, [15] for a computationally efficient procedure for the "combinational auction" that can be used to purchase compound goods, and [16] for a discussion of the links between combinatorial auctions and Lagrangean relaxation for the job shop scheduling problem).

Unlike most of the agent-mediated markets that have been envisioned for the competitive allocation of computer resources (e.g., [17], [18], [19]) and the cooperative allocation of manufacturing resources within a single firm (e.g., [20]), the costs and currency used in a public logistics network would correspond to the real, actual costs associated with providing the service; in most other applications, some type of transfer price and virtual currency is used because there is no directly available measure of actual costs. In the research, real prices would be negotiated and accounts settled via a continuous series of micro-payments of legal tender currency. Bredin et al. [21] have developed a model for a market that does use actual micro-payments in order to induce an open environment where computational systems are willing to serve as hosts for mobile agents.

6.2. Example

Figure 2 shows an example of adaptive routing and in-transit trade for a single package being transported from DC 4 (Jacksonville, FL) to DC 30 (Richmond, VA). The current "Home Agent" for the package is located at DC 24. The Home Agent coordinates the interactions between all of a package's other agents, routes the package, and, in the case of the possible resale of the package, forwards final offers to the current owner of the package (possibility an agent at a factory in Jacksonville, FL) and, if accepted, arranges for the re-routing of the package to its new destination.

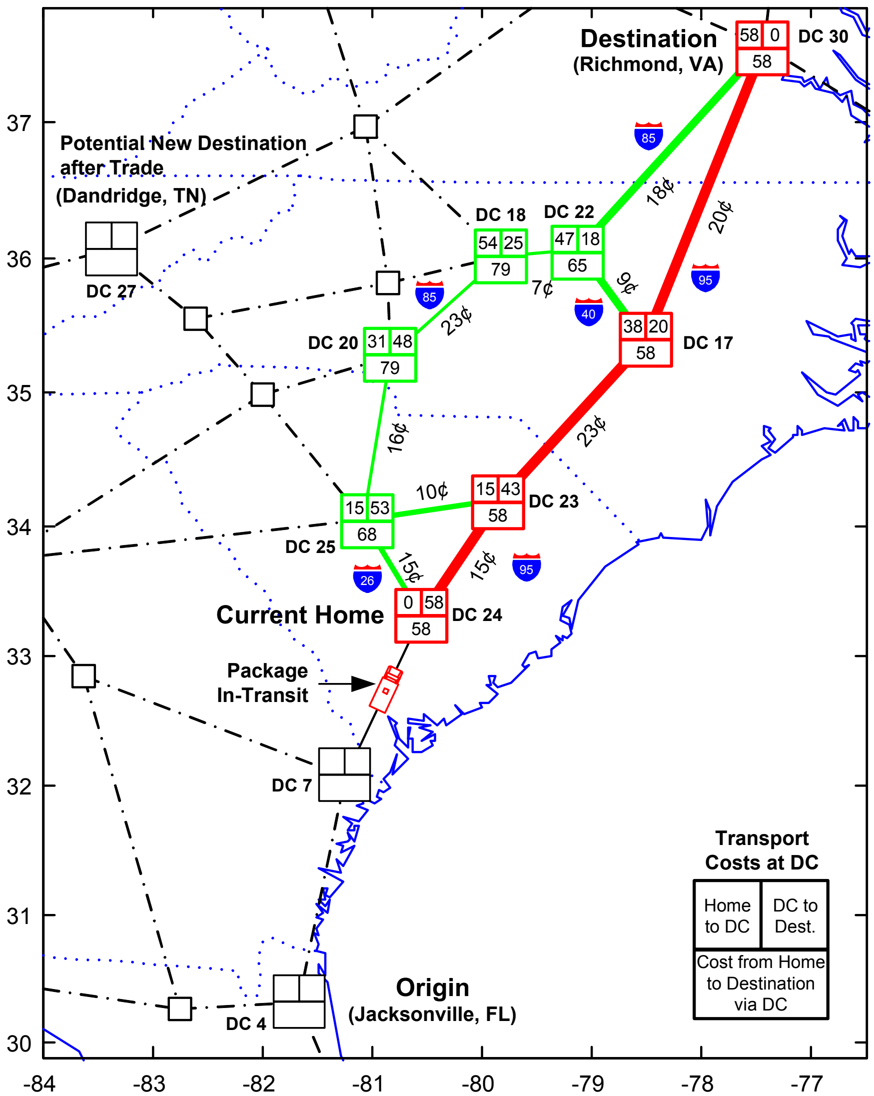

Figure 2. Example of adaptive routing.

In Figure 2, the package itself is onboard a truck traveling north along I-95 heading to DC 24. The package's transport cost from DC 4 to DC 24 has already been paid via micro-payments and is a sunk cost. The relevant cost for routing purposes is the future cost of transporting the package from DC 24 to its destination; this cost is currently \$0.58, assuming the package will travel via DCs 24, 23, 17, and 30. Once the package leaves DC 24, its Home Agent moves (as a mobile agent) to the next hop, DC 23.

The package agents at DCs 25, 20, 18, and 22 are available to determine alternative routes if necessary. The cost reaching the destination (DC 30) via these agents is greater than along the intended route. Each agent at each DC has a "via cost," the cost reaching the destination via the DC. Table 1 shows the effects of four possible disruption scenarios: Scenarios 1–3 correspond to decreases in the transport costs from DC 20 (Charlotte, NC) to DC 18 (Greensboro, NC), possibly due to an oversupply of trucks headed to Charlotte; decreases of $0.05, $0.08, and $0.16 are shown in the table. Scenario 4 corresponds to an increase of $1.00 in the transport cost from DC 23 to DC 17, possibly due to an accident along I-95 that will require trucks to be re-routed via non-interstate roads.

Scenarios 3 and 4 result in changes that are significant enough to trigger the re-routing of the package; Scenarios 1 and 2 terminate without re-routing. What is most important from an implementation point of view is that the disruptions are only propagated locally. The DC at the source of the disruption propagates its new via cost to all of its neighbors. Each neighbor DC needs to continue to propagate the change to its neighbors only if it would result in a change to its predecessor or successor. Any change along the intended route will reach the Home Agent, at which point the agent can be re-routed if necessary. As can be seen in the table, only a significant change will result in significant communications between the DCs; most minor changes will be terminated locally. Since hundreds of millions of agents could be active at any time coordinating tens of millions of packages, reducing communications requirements is an important feature for any routing protocol.

With respect to in-transit trade, any of the DC agents in Figure 2 are available as trading agents. For example, if a customer at DC 27 (Dandridge, TN) is interested in purchasing the package, it can spawn search agents. These agents would first reach the package's agents at DCs 25, 20, and 18. The cost of re-routing the package from its current location at DC 24 to these DCs is immediately known to be $0.15, $0.31, and $0.54, and can be used as part of the delivered price in the resale negotiations.

Table 1. Adaptive Routing Scenarios

Scenario	DC	24	23	17	25	20	18	22	30
0: Current	Via Cost (¢)	58	58	58	68	79	79	65	58
	Pred. DC	–	24	23	24	25	20	17	17
	Succ. DC	23	17	30	23	18	22	30	–
1: 23¢ → 18¢ = ↓5¢	Via Cost (¢)	–	–	–	–	74	74	–	–
along arc (20,18)	Pred. DC	–	–	–	–	–	–	–	–
⇒ No change	Succ. DC	–	–	–	–	–	–	–	–
2: 23¢ → 11¢ = ↓12¢	Via Cost (¢)	–	–	–	67	67	67	–	–
along arc (20,18)	Pred. DC	–	–	–	–	–	–	–	–
⇒ No change	Succ. DC	–	–	–	20	–	–	–	–
3: 23¢ → 1¢ = ↓22¢	Via Cost (¢)	57	–	–	57	57	57	57	57
along arc (20,18)	Pred. DC	–	–	–	–	–	–	18	22
⇒ Re-route	Succ. DC	25	–	–	20	–	–	–	–
4: 23¢ → $1.23 = ↑$1.00	Via Cost (¢)	79	89	90	79	79	79	79	79
along arc (23,17)	Pred. DC	–	–	22	–	–	–	18	22
⇒ Re-route	Succ. DC	25	25	–	20	–	–	–	–

7. References

1. Kay, M.G. and Parlikad, A.N., "Material flow analysis of public logistics networks," 7th Int. Material Handling Res. Colloq., Portland, ME, June 1–5, 2002 (`http://www.ie.ncsu.edu/kay/pln/IMHRC02.pdf`, accessed Sept. 2002).

2. Parlikad, A.N., "Performance Analysis of Intelligent Supply Chain Networks," Master's Thesis, Dept. of Industrial Eng., North Carolina State Univ., Raleigh, NC, 2002, (accessed Sept. 2002, `http://www.lib.ncsu.edu/theses/available/etd-05162002-151641/`).
3. Gandlur, K.S., "Implementation of Adaptive Routing in Public Logistics Networks," Master's Thesis, Dept. of Industrial Eng., North Carolina State Univ., Raleigh, NC, 2002 (accessed Sept. 2002, `http://www.lib.ncsu.edu/theses/available/etd-08152002-140219/`).
4. Reilly, F.K. and Brown, K.C., *Investment Analysis and Portfolio Management*, 5th ed., Dryden, Phil., 1997.
5. Keshav, S., *An Engineering Approach to Computer Networking: ATM Networks, the Internet, and the Telephone Network*, Addison Wesley, Reading, MA, 1997.
6. Bertsekas, D. and Gallager, R., *Data Networks*, Prentice-Hall, Englewood Cliffs, NJ, 1987.
7. Wooldridge, M., *An Introduction to Multiagent Systems*, Wiley, Chichester, England, 2002.
8. Omicini, A., Zambonelli, F., Klusch, M., and Tolksdorf, R., Eds., *Coordination of Internet Agents: Models, Technologies, and Applications*, Springer, Berlin, 2001.
9. Swaminathan, J.M., Smith, S.F., and Sadeh, N.M., "Modeling supply chain dynamics: A multiagent approach," *Decision Sci.*, 29(3): 607–632, 1998.
10. Bonabeau, E. and Theraulaz, G., "Swarm smarts," *Sci. Amer.*, March 2000, pp. 72–79.
11. Dorigo, M., Maniezzo, V., and Colorni, A., "Ant system: Optimization by a colony of cooperating agents," *IEEE Trans. Syst., Man, and Cyber.–Part B: Cyber.*, 26(1): 29–41, 1996.
12. Kephart, J.O., Hanson, J.E., and Greenwald, A.R., "Dynamic pricing by software agents," *Computer Networks*, 32: 731–752, 2000.
13. McCabe, K., Rassenti, S., and Smith, V., "Designing auction institutions for exchange," *IIE Trans.*, 31: 803–811, 1999.
14. Milgrom, P., "Auctions and bidding: A primer," *J. Eco. Perspectives*, 3(3): 3–22, 1989.
15. Rothkopf, M.H., Pekec, A., and Harstad, R.M., "Computationally manageable combinational auctions," *Management Sci.*, 44(8): 1131–1147, 1998.
16. Kutanoglu, E. and Wu, S.D., "On combinatorial auction and Lagrangean relaxation for distributed resource scheduling," *IIE Trans.*, 31: 813–825, 1999.
17. Sutherland, I.E., "A futures market in computer time," *Comm. ACM*, 11(6): 449–451, 1968.
18. Miller, M.S. and Drexler, K.E., "Markets and computation: Agoric open systems," in *The Ecology of Computation*, Huberman, B.A., Ed., North Holland, Amsterdam, 1988, pp. 133–176 (`http://www.agorics.com/Library/agoricpapers.html`).
19. Huberman, B.A. and Hogg, T., "Distributed computation as an economic system," *J. Eco. Perspectives*, 9(1): 141–152, 1995.
20. Lu, T.P. and Yih, Y., "An agent-based production control framework for multiple-line collaborative manufacturing," *Int. J. Prod. Res.*, 39(10): 2155–2174, 2001.
21. Bredin, J. et al., "A market-based model for resource allocation in agent systems," in *Coordination of Internet Agents: Models, Technologies, and Applications*, Omicini, A., Zambonelli, F., Klusch, M., and Tolksdorf, R., Eds., Springer, Berlin, 2001, pp. 426–442.

Acknowledgement

This research is supported, in part, by the National Science Foundation under Grant CMS-0229720.