

An Approach to Securely Identifying Beneficial Collaboration in Decentralized Logistics Systems

Chris Clifton

Department of Computer Science, Purdue University, West Lafayette, Indiana 47907,
clifton@cs.purdue.edu

Ananth Iyer

Krannert School of Management, Purdue University, West Lafayette, Indiana 47907,
aiyer@mgmt.purdue.edu

Richard Cho

Faculty of Business, University of New Brunswick Saint John, Saint John,
New Brunswick Canada, E2L 4L5, rcho@unbsj.ca

Wei Jiang

Department of Computer Science, Purdue University, West Lafayette, Indiana 47907,
wjiang@cs.purdue.edu

Murat Kantarcioğlu

Department of Computer Science, The University of Texas at Dallas, Richardson,
Texas 75083, muratk@utdallas.edu

Jaideep Vaidya

Management Science and Information Systems Department, Rutgers University, Newark,
New Jersey 07102, jsvaidya@rbs.rutgers.edu

The problem of sharing manufacturing, inventory, or capacity to improve performance is applicable in many decentralized operational contexts. However, the solution of such problems commonly requires an intermediary or a broker to manage information security concerns of individual participants. Our goal is to examine use of cryptographic techniques to attain the same result without the use of a broker. To illustrate this approach, we focus on a problem faced by independent trucking companies that have separate pick-up and delivery tasks and wish to identify potential efficiency-enhancing task swaps while limiting the information they must reveal to identify these swaps. We present an algorithm that finds opportunities to swap loads without revealing any information except the loads swapped, along with proofs of the security of the protocol. We also show that it is incentive compatible for each company to correctly follow the protocol as well as provide their true data. We apply this algorithm to an empirical data set from a large transportation company and present results that suggest significant opportunities to improve efficiency through Pareto improving swaps. This paper thus uses cryptographic arguments in an operations management problem context to show how an algorithm can be proven incentive compatible as well as demonstrate the potential value of its use on an empirical data set.

Key words: collaboration; routing; cryptography; space-filling curve; incentive compatible; algorithm

History: Received: January 9, 2004; accepted: December 10, 2006. Published online in *Articles in Advance* December 11, 2007.

1. Introduction

Coordination of decisions and sharing capacity across independent decision makers can provide significant cost reductions in many operations management problem contexts. In decentralized environments, there are numerous instances where individual companies would like to swap tasks or loads to gain

operational efficiencies. Keskinocak and Tayur (2001) suggest that e-marketplaces can play the role of central coordinators. They suggest, for example, a marketplace where customers post orders for paper products and manufacturers offer capacity that is shared across multiple orders to generate efficiency. They also describe an example where independent

manufacturers who have production capability for a variety of products with associated setups swap customer demands to reduce setups. Another example discussed involves intermediaries that provide consolidation and capacity trading opportunities for shippers and carriers. Similarly, Kalagnanam et al. (2000) describe a problem faced by steel companies that have to deal with surplus inventories of steel coils with specific properties. They provide an optimization model that can be solved by a coordinator who attempts to optimally allocate the surplus inventory. Keskinocak and Tayur (2001) identify that a key requirement for an intermediary is that it would have to be an independent, trustworthy entity who would keep all information confidential and who would create benefits that are equitably distributed.

The examples above suggest that one reason for the use of a broker or a marketplace is that it prevents disclosure of proprietary information across competitors. However, this information must still be revealed to the broker for coordination to occur. In addition, the broker will need to cover his costs by charging a margin, the data collected by the broker have to be secured from leakage, and the broker's incentives have to be aligned with individual company goals. In many contexts, independent companies are reluctant to share proprietary information, unless it is absolutely necessary, to protect their customer base. However, companies do realize that some sharing of data may be necessary to compensate for market fragmentation. In such contexts, it would be ideal to devise a system where all of the data are available for use but only data that are absolutely required for coordination are shared.

The cryptographic community has shown that a trusted third party is not required—it is possible to compute functions without disclosing private data to any party (Yao 1986, Goldreich et al. 1987). The result is that no party learns more than they would if a broker arranged the transactions, *and no broker is required*. Our goal is to apply state-of-the-art techniques from data encryption to logistics problems to automate the task performed by the broker. Companies learn no more than with an honest broker. However, the broker is eliminated. In fact, for the specific problem and solution given in this paper, we prove that no party learns more than the minimum needed to accomplish

the desired efficiency gains. The benefit to shipping companies and shippers is the ability to reduce collaboration costs and improve the efficiency of the overall system.

One reason for this research is that we believe that an examination of operations problems with their specific contexts can provide opportunities for the development of algorithms that exploit problem structure. In addition, by focusing on algorithms that are embedded in commercial codes, we suggest that there is potential for rapid practical adoption. Finally, there are conceptual issues regarding data leakage versus efficiency that we suggest as a rich research area, both in operations management as well as in the computer security research community.

1.1. A Specific Problem Context—The Trucking Industry

To provide a complete treatment of our approach, we examine a specific problem context involving truck routing. Truck transport is a \$462 billion industry in the United States (Wilson and Delaney 2003). However, the industry is extremely fragmented, with the largest company accounting for less than 5% of the market. The main source of inefficiency in this industry is the “deadhead” miles, or miles driven empty. The primary reason for this inefficiency is the spatial nature of this industry; i.e., for a truck to pick up a load, it has to physically be at that location. When the truck is done, it ends up at the physical drop-off point and may have to travel to a new location to be useful. Intuitively, if transport companies swap some of their loads, there is the potential for Pareto improving savings (i.e., neither company faces a higher cost and at least one company faces a lower cost).

However, attempts to collaborate and thus swap loads to get more efficient routes are often discouraged by a desire of individual companies to “share only if beneficial.” In addition, legal restrictions dealing with antitrust issues frown on information sharing and collaboration that can be potentially anticompetitive. However, antitrust considerations do permit competing firms to engage in limited information sharing and collaboration that is clearly efficiency enhancing. As reported in the *Wall Street Journal* (Bandler 2003, p. 1),

It is permissible for carriers to cooperate in certain ways. For instance, if two of them both carry chemicals for a given producer on the same route, they

may pool their capacity for the purpose of operational efficiency But cooperating to divide up markets or to affect prices would clearly fall outside these permitted arrangements.

If transport companies resort to using a broker to swap loads, the first step is for each company to independently identify loads it would like to swap. These potential loads are provided to a broker who now sees all available loads. As a result, parties will only make things available that they view as likely to be picked up in a swap. The key difference in our approach is that all available loads are provided by all companies. The algorithm is executed in a distributed manner by each company and at no time are the data entrusted to a third party. In addition, the secure protocols used in the algorithm guarantee that no information other than the swapped loads that improve efficiency are revealed to each company, and all such efficiency-improving swaps are made. This results in both lower information disclosure and higher efficiency than a traditional broker-mediated model.

Why would companies want such an approach? One motivation comes from the transport companies themselves as part of their desire to protect proprietary information while achieving maximum efficiency. However, another motivation may come from shippers who could demand such a protocol be used by their associated carriers to ensure that efficiency is enhanced while preventing any collusion by carriers regarding data that are not explicitly required to be shared and that may reduce competitiveness of the carrier market.

Thus, in this paper:

1. We provide an algorithm that ensures that sharing takes place only if each company sees its costs reduced and that the sharing scheme ensures that all potential players can engage to identify cost-reducing swaps, while ensuring no information is shared other than what can be concluded from the final swapped points.

2. We prove incentive compatibility of the associated algorithm by demonstrating that honesty on the part of the collaborators is ensured by guaranteeing that either it can be detected that one participant is cheating (and thus gets thrown out of future collaborations) or that the cheating is not incentive compatible; i.e., the cheater is worse off.

3. We show that the algorithm, implemented in a decentralized manner, affords the globally optimal split of loads for a specific setting.

4. We apply the algorithm proposed to an empirical data set from a transportation company that provided us with 11 weeks of pick-up and delivery data. Although the algorithm proposed is a heuristic in the context of vehicle routing, it uses the state-of-the-art techniques in data encryption and secure multiparty computation techniques that guarantee that the security requirements are met. The empirical data suggest the delivered value of the algorithm. Empirical results suggest that the potential to reduce costs is over 15% based on application of the algorithm.

This paper thus provides an application of cryptographic techniques as an enabler of operational efficiency in a decentralized ownership environment. We also suggest the associated methodological issues as a potentially fruitful area of research in operations management. In addition, several combinatorial problems can be solved by using the space-filling curve as a heuristic. For such problem contexts, our paper provides a template to generate secure implementations.

The next section gives a formal treatment of the algorithm, along with proofs, showing that it achieves a one-dimensionally optimal result, that nothing is disclosed, that is not obvious from the result, and that cheating is not incentive compatible; i.e., it is detrimental to the cheater. In §2.5, we show that the algorithm can be securely used among multiple parties. In §3, we provide results when this approach is applied to a set of real shipping transactions. In the appendix, we provide notation and formal proofs of the theorems in the text.

2. Problem Description

Formally, the general problem is as follows. We have N independent transport companies with company i having m_i points (corresponding to m_i unit loads) located in a two-dimensional plane that must be served by a truck. The goal is to identify a sequence of enquiries and swaps between pairs of companies that results in: (1) a set of points that when swapped between the companies guarantees that no company is worse off, (2) a situation where no information is shared other than what can be concluded from the

final swapped points, (3) a situation where the algorithm is polynomial in running time, and (4) a situation where any cheating by either party during the execution of the algorithm is either detected by the other party or results in a less efficient solution for the dishonest party, thus providing the incentive to truthfully follow the algorithm.

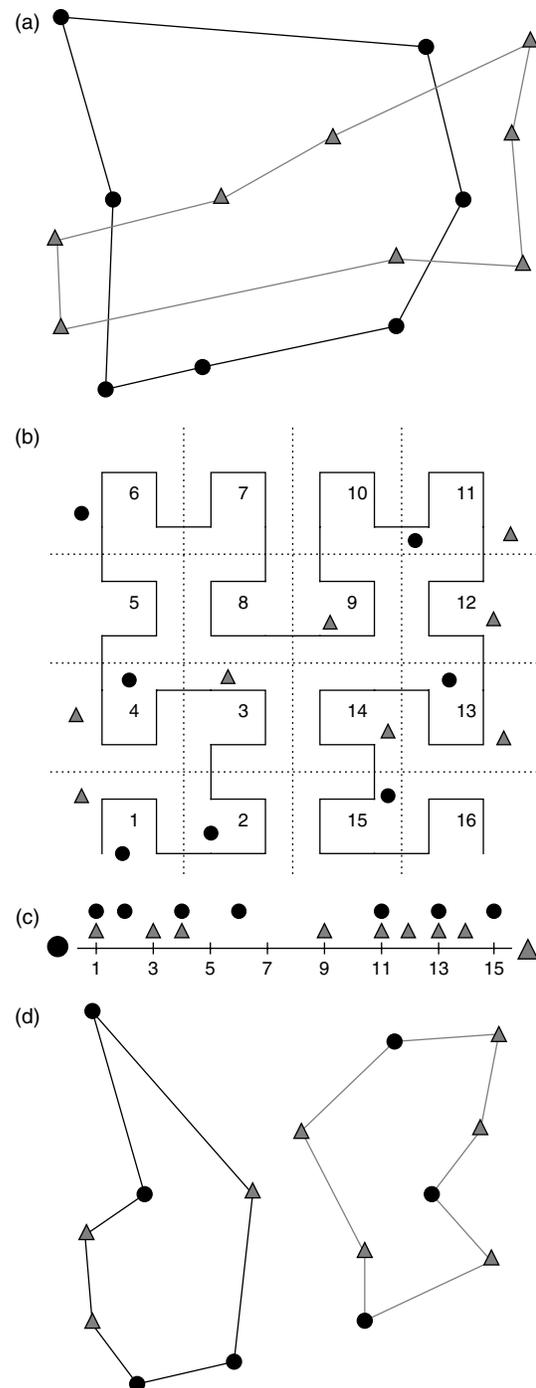
2.1. Using Space-Filling Curves

We first map the points in the initial two-dimensional problem to one dimension using a space-filling curve. Figure 1(a) shows the existing routes through the points served by each of the two transport companies. Figures 1(b) and 1(c) show a mapping of the points, via a Hilbert space-filling curve, to one dimension. The use of a space-filling curve to develop heuristics for combinatorial problems is described in Bartholdi and Platzman (1988). Bartholdi and Platzman (1988) show that this results in the following properties: (a) 25% worse than optimal worst-case performance for planar traveling salesman problems (Bartholdi and Platzman 1982), (b) solutions within one second that have a gap of less than 34% more than the best approach for large problems using two months of computing time (see Bartholdi 2006), and (c) reported implementation in logistics packages such as the ARC/Info Geographical Information System, the CAPS logistics toolkit of Baan Systems, and other commercial systems managing two-dimensional data (see Bartholdi 2006).¹ Also, because the space-filling curve is a one-to-one mapping from two dimensions to one dimension, points identified for swap in one dimension provide a unique pointer to the corresponding original point location in two dimensions.

Why map to one dimension? In general, the vehicle routing problem is an NP-hard optimization problem, even without worrying about privacy or security or both. Although proper choice of heuristics may give good solutions directly on the two-dimensional problem, choosing those heuristics requires an understanding of the characteristics of the data—and sharing this information violates goal 2. However, in one dimension the optimal solution is tractable: The search for the best solution requires a logarithmic number of steps.

¹Note that our approach will be to incorporate the encryption processes within this heuristic.

Figure 1 Execution of OROD Protocol Between Two Parties



2.2. A Swap Algorithm

The goal of the algorithm is to identify points that can be swapped between two parties, whose locations are on a line, that results in (a) both parties

decreasing their costs and (b) finding the maximum number of swaps that will optimize the total distance traveled. Note that in one dimension, given the number of points to swap, the specific points that will guarantee the lowest cost are apparent, i.e., swap the farthest points for each location. We begin with the assumption that each party has selected an end point. This choice may be influenced by their current hub locations—a natural approach would be to draw a space-filling curve with endpoints at the hubs. A list of symbols and notation is provided in Appendix A.1.

Let the k th extreme point $Extreme_POS(X, d_x, k)$ refer to the k th farthest point from the endpoint d_x in the points X . Given a number of points k , the parties compare their k th extreme points. If the locations do not cross, k is a lower bound on the number of points to swap to decrease costs and the parties try $2k$. Once the locations cross, k is an upper bound, and the search continues between the upper and lower bounds until the optimal number of points to swap is found. A formal statement of the algorithm is provided as Algorithm 1. (Note that Aggarwal et al. 2004 present a protocol for securely computing the k th-ranked element among multiple parties. The problem is different, as we do not know k in advance; the distinctions are further discussed in Appendix B.4.)

ALGORITHM 1. ONE-DIMENSIONAL RELATIVE OUTLIER DETECTION (OROD)

Require: S^1, O^1 ,
1: $lbound \leftarrow 0$
2: $ubound \leftarrow \infty$
3: $i \leftarrow 1$
4: {Lines 5 through 13 determines the maximal size of i }
5: **while** ($ubound - lbound > 1$) **do**
6: **if** $Extreme_POS(S^1, l, i) > Extreme_POS(O^1, r, i)$
 {Defined in Appendix A.2} **then**
7: $lbound \leftarrow i$
8: **else**
9: $ubound \leftarrow i$
10: **end if**
11: $i \leftarrow \min(i * 2, \lfloor (lbound + ubound) / 2 \rfloor)$
12: **end while**
13: **return** i

2.2.1. Illustration of OROD Execution. We now provide an example to illustrate the algorithm. Let

\bigcirc and Δ be two parties, with original routes shown in Figure 1(a). Figures 1(b) and 1(c) present a space transformation process via a Hilbert curve. Let \bigcirc start from the left end and Δ start from the right end. We now show the execution of the algorithm from \bigcirc 's point of view.

Initially, $lbound \leftarrow 0$, $ubound \leftarrow \infty$, and $i \leftarrow 1$. Since the difference between $lbound$ and $ubound$ is greater than 1, the execution enters the *while* loop. $Extreme_POS(\bigcirc, left, 1)$ returns 15, which is greater than the index returned from $Extreme_POS(\Delta, right, 1) = 1$, so $lbound \leftarrow 1$ and $i \leftarrow 2$. In the next iteration, $Extreme_POS(\bigcirc, left, 2) = 13$ is still greater $Extreme_POS(\Delta, right, 2) = 3$, so $lbound \leftarrow 2$ and $i \leftarrow 4$. $Extreme_POS(\bigcirc, left, 4) = 6 < Extreme_POS(\Delta, right, 4) = 9$, so $ubound \leftarrow 4$ and $i \leftarrow 3$. $Extreme_POS(\bigcirc, left, 3) = 11 > Extreme_POS(\Delta, right, 3) = 4$, so $lbound \leftarrow 3$ and $i \leftarrow 3$. Since $ubound - lbound = 1$, the execution exits with 3 as the number of points to be swapped. Thus \bigcirc swaps its (least-desired) points in locations 11, 13, and 15 for Δ 's points in 1, 3, and 4. The line is thus partitioned into two parts, each assigned to one end. Given the one-to-one mapping preserved by the space-filling curve, the corresponding three swapped points can be identified in the original two dimensions. The resulting tours are shown in Figure 1(d).

THEOREM 1. *The OROD algorithm terminates with the optimal number of points swapped; i.e., after swapping, there are no two points that could be swapped to give a lower tour length.*

PROOF. See Appendix C. \square

2.3. Secure Execution

We next focus on secure execution of the algorithm OROD. The goal in this section is to ensure that no “excess” information is revealed by either party, other than what can be deduced from the final swaps. Because Algorithm 1 only exchanges information in step 6, we want to perform the comparison of the $Extreme_POS$ functions from both parties without disclosing anything except the comparison result. As we shall show, the comparison results at each step can be deduced from the final swaps, so the result of the comparison does not disclose excess information.

The idea of the secure version of the algorithm is to provide a mechanism to check whether it is okay to

Figure 5 Example of Execution, Iteration 4

		lbound = 2, ubound = 4															
i = 3		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
		>	>	>	>	>	>	>	>	>	>	≤	≤	≤	≤	≤	≤

be the contents of boxes 0 and 1, respectively. (For simplicity, we show the contents as a single bit; B_0 and B_1 take values 0 or 1.)

The protocol is described below:

1. \bigcirc generates and publishes three numbers:

C : a random number between $1, \dots, p-1$,

p : a large prime number, and

g : the generator of p 's multiplicative group; i.e., every number between 1 and $p-1$ can be written as $g^k \bmod p$ for some k between $1, \dots, p-1$.

2. Δ picks a random number k between $1, \dots, p-1$, sets $P_\sigma = g^k \bmod p$ and $P_{1-\sigma} = (C/P_\sigma) = C * P_\sigma^{-1} \bmod p$, where P_σ^{-1} is the multiplicative inverse of P_σ (i.e. $P_\sigma^{-1} * P_\sigma = 1 \bmod p$), and sends P_0 to \bigcirc .

3. \bigcirc finds P_1 by evaluating $C/P_0 \bmod p$, creates $E_0 = (g^{r_0} \bmod p, M_0 = (H((P_0)^{r_0} \bmod p) \oplus B_0))$, $E_1 = (g^{r_1} \bmod p, M_1 = (H((P_1)^{r_1} \bmod p) \oplus B_1))$, by randomly choosing r_0, r_1 between $1, \dots, p-1$, and sends E_0, E_1 to Δ . Note that E_0, E_1 are both pairs of numbers. The hash function H is described in Appendix B.2.

4. Δ computes

$$T_\sigma = H((P_\sigma)^{r_\sigma} \bmod p) = H((g^{r_\sigma})^k \bmod p)$$

and gets B_σ by computing $T_\sigma \oplus M_\sigma$. It cannot compute the other value due to the cryptographic assumptions outlined in Appendix B.2.

In the above protocol the choice of Δ (σ) is not revealed because all \bigcirc receives is either $g^k \bmod p$ or $C/g^k \bmod p$, where k is chosen randomly. Since operations are done in $\bmod p$, both $g^k \bmod p$ or $C/g^k \bmod p$ values are uniformly distributed between $0, \dots, p-1$. Therefore, \bigcirc does not see anything more than a random number. Δ learns nothing by receiving the random C or (because of the random oracle hash function) from inspecting E_0 or E_1 .

Although Δ can decrypt E_σ to obtain the final result, by the original Diffie-Hellman assumption (Diffie and Hellman 1976; see Appendix B.2) it cannot decrypt the other box. If Δ could decrypt both E_0 and E_1 , it would mean that it would know an efficient way to find $(P_{1-\sigma})^{r_{1-\sigma}} = g^{k'r_{1-\sigma}} \bmod p$. (Note that the preceding value is needed to open the other box.)

Given $P_{1-\sigma} = g^{k'} \bmod p$ (since g is a generator we can write $P_{1-\sigma} = g^{k'} \bmod p$ for some k') and $g^{r_{1-\sigma}} \bmod p$, this violates the Diffie-Hellman assumption.

To clarify the above protocol, we give a simple example.³

EXAMPLE 1. Let the shared hash function H be $H(0) = 1, H(1) = 0, H(2) = 1, H(3) = 1, H(4) = 0$.

1. \bigcirc has $B_0 = 1, B_1 = 0$; it also generates and publishes three numbers: $C = 4, p = 5$, and $g = 2$ (note that $2^0 = 1 \bmod 5, 2^1 = 2 \bmod 5, 2^2 = 4 \bmod 5, 2^3 = 3 \bmod 5$).

2. Δ wants to learn the value of B_0 ($\sigma = 0$). Δ picks a random number $k = 3$ and sets $P_0 = 2^3 = 8 = 3 \bmod 5$ and $P_1 = (C/P_0) = 4/3 = 4 * 2 = 8 = 3 \bmod 5$ (note that $P_0^{-1} = 2$ is the multiplicative inverse of $P_0 = 3$ taken $\bmod 5$; $P_0 * P_0^{-1} = 3 * 2 = 1 \bmod 5$). Δ sends P_0 to \bigcirc .

3. \bigcirc independently calculates P_1 by evaluating $4/P_0 = 4 * 2 = 3 \bmod 5$, and randomly choosing $r_0 = 4, r_1 = 3$, calculates $g^{r_0} = 2^4 = 16 = 1 \bmod 5, g^{r_1} = 2^3 = 8 = 3 \bmod 5, (P_0)^{r_0} = 3^4 = 81 = 1 \bmod 5, (P_1)^{r_1} = 3^3 = 27 = 2 \bmod 5, M_0 = H(1) \oplus B_0 = 0 \oplus 1 = 1$, and $M_1 = H(2) \oplus B_1 = 1 \oplus 0 = 0$. \bigcirc sets $E_0 = (1, 1)$ and $E_1 = (3, 1)$, and sends E_0, E_1 to Δ .

4. Δ computes

$$\begin{aligned} T_0 &= H((P_0)^{r_0} \bmod 5) = H((3^4) \bmod 5) \\ &= H(81 \bmod 5) = H(1) = 0 \end{aligned}$$

and gets B_0 by computing $T_0 \oplus M_0 = 0 \oplus 1 = 1$.

Thus Δ gets to know the contents of the box 0 (equal to 1), whereas \bigcirc does not know which box was opened by Δ . This shows the secure 1-out-of-2 oblivious transfer. The approach above can be extended to deal with 1-out-of- n oblivious transfer. We provide details in Appendix B.3.

2.4. Formal Statements Regarding Security of the Algorithm

Formally, we divide the security into two parts. First, if each company follows the protocol correctly, we show that the algorithm is secure. This satisfies the semihonest definition of secure multiparty computation. Second, we show that each company benefits by following the protocol correctly; i.e., individual company participation by following the protocol is incentive compatible.

³We would like to stress that 1,024-bit or larger p and C values would be used in practice for secure applications.

A secure implementation requires us to prove that nothing is learned by either party during the execution except what can be deduced from the final points swapped. To prove this, we show that it is possible to build a “simulator” that uses only the final results and the data known by each individual party to identify the result of the comparison. The proof requires verification that the results at each step can indeed be deduced from the result. The logic that follows is that no new information is revealed at each step of the execution of the algorithm. This provides a formal proof of security of the algorithm if each party were to follow the protocol provided.

THEOREM 2. *Algorithm 1 is secure under the semihonest definition of secure multiparty computation.*

PROOF. See Appendix D. \square

The theorem above formally establishes that the OROD algorithm results in a secure implementation.

The secure multiparty computation literature also provides techniques secure against *malicious* parties; i.e., parties that deviate from the protocol can be detected. However, even with security against malicious parties one can “cheat” by faking one’s input and running the protocol honestly. We thus focus on the concept of *incentive compatibility*; i.e., we prove that it is in each party’s self-interest to follow the protocol and provide the correct input. In other words, trying to learn more is discouraged by the fact that the party is worse off or is detected as cheating. (Similar issues were explored in Dodis et al. 2000, who suggest that finding problems where incentive compatibility simplifies protocols is an interesting research area; this is discussed further in Appendix E.) Formally, we have to show that it is incentive compatible for each party to follow the protocol. Note that we will assume that if a party does not benefit from violating the protocol, it will follow the protocol.

THEOREM 3. *The OROD protocol is incentive compatible; i.e., any failure to correctly follow the protocol results in the dishonest party being worse off, or leaves the result unaffected and gives no additional information to the dishonest party.*

PROOF. See Appendix E. \square

2.5. Execution of OROD Among Multiple Parties

Define the optimal result from the execution of OROD among multiple parties to be the perfect partitioning among all parties’ one-dimensional data sets. In other words, any two parties’ ranges of their one-dimensional data sets are disjoint. This results in the minimal global tour (in one dimension).

We note one problem: It is not possible to draw a space-filling curve such that every party’s depot is uniquely at the opposite end of the line from every other party. The result is that a globally optimal partitioning may not be optimal (or even an improvement) for every party; we lose the incentive compatibility in Theorem 3. However, if we assign a total ordering to the parties and use this to determine which end of the line they get in any pairwise operation, repeated execution does converge to a minimum global tour.

THEOREM 4. *Execution of OROD among k parties eventually reaches optimal: k partitions (where $k \geq 2$).*

PROOF. See Appendix C. \square

However, this multiple party execution of OROD is *not* secure under the definitions of secure multiparty computation; i.e., parties may see intermediate results (swapped points) that eventually are passed to another party. To be secure, each party should see only those points that it starts with and that it ends with.

On the other hand, with one additional constraint we can utilize the two-party OROD protocol as a subroutine to build a secure multiparty protocol. The constraint is that once a party has received a swapped point, it can never swap it again.

Constraint MP: Whenever OROD is executed between any two parties, the points swapped between them should be removed from consideration for their subsequent execution with other parties. Informally, this constraint says that a subcontractor may not further subcontract the work he received as part of a swap. This constraint guarantees that contracts with the shipping company are transferred to a known (and approved) carrier. We thus assume that the MP constraint reflects the fact that only the original carrier can subcontract—this information is passed along to the shipper for payment. No one other than the original carrier or the first subcontracted carrier will

be paid by the shipper or be permitted to handle the shipment.

To ensure that parties follow the MP constraint, we introduce a mechanism that allows violation to be detected. The basic idea is to have every party provide a hash function of their input values at the beginning of the protocol. The procedure to create an appropriate hash function, and its properties are shown in the appendix. After receiving a point, its hash can be checked to see if it comes from the correct source. The formal details of the process are provided in the appendix.

Note that two malicious parties (including the final recipient) could collude to ignore the MP constraint. This corresponds to swapping points after the completion of the OROD protocol among multiple parties and could not be prevented by the protocol. However, real-world incentives protect against this, as described earlier. A shipper contracts with a particular transporter; if this transporter subcontracts to another (the result of a swap), this information will be communicated to the shipper. If a third party shows up to collect the load (the result of a malicious swap), the shipper would obviously be concerned. (Would you trust your valuable shipment to whoever tries to come collect it?)

THEOREM 5. *Executing OROD among multiple parties is secure provided that constraint MP is satisfied.*

We do not provide a formal proof of this theorem because it follows the same steps as Theorem 2. According to Theorem 2, a single execution of OROD is secure. In addition, because constraint MP must be satisfied, no swapped points are considered for the subsequent executions, and all intermediate swapped points are part of the final result. Therefore, each party's view during the execution of OROD among multiple parties can be simulated by the party's input and output, following the same steps as in the proof of Theorem 2.

Also note that securely executing OROD among multiple parties merely requires that each party executes OROD with the other parties once, provided that each execution of OROD satisfies constraint MP. The reasoning follows: Suppose party i has already executed OROD once with all other parties, and assume that i executes OROD with party j (where

$i \neq j$) for a second time. Because constraint MP is satisfied, no points would be considered during this second execution between i and j except their original points. Based on Theorem 1, the first execution of OROD between i and j creates two perfect partitions of their original points. In other words, if some of their original points were swapped after the second execution of OROD, these points would have been swapped after the first execution. Consequently, when constraint MP is guaranteed, any subsequent execution of OROD between i and j after the first execution does nothing.

2.6. Related Issues

The multiparty OROD with MP⁴ may not produce an optimal solution in the one-dimensional space. Although Theorem 5 shows that no information is disclosed except the points swapped, do the points swapped in multiparty OROD with MP reveal more information than an optimal solution would?

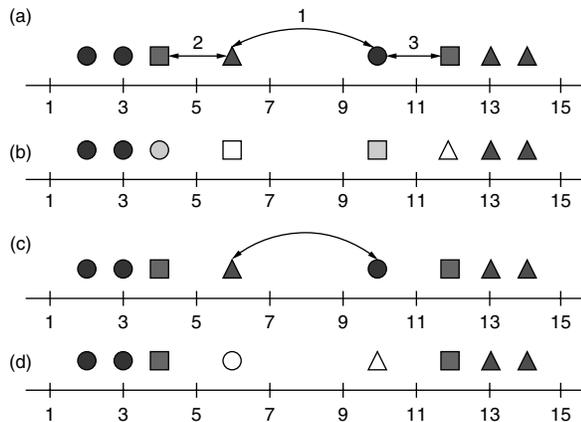
CLAIM 1. *Multiparty OROD with MP does not always leak more information than an optimal one-dimensional secure solution.*

PROOF. From Theorem 5, all that is disclosed is the points swapped.

Figure 6(a) provides three sets of shippers, whose points are identified as \circ , \triangle , and \square . As the example shows, these parties can have different numbers of points available to swap, but because swaps are one for one, their capacity constraints are satisfied after the swaps are completed. Assume that in the original OROD (with no MP constraint), we first have \circ (on the left) swapping with \triangle (on the right). The effect is to swap the triangle in location 6 for the circle in position 10. Next consider swaps between \circ (on the left) swapping with \square (on the right). The result is that the circle, now in position 6, is swapped with square in position 4. Finally, consider \square (on the left) swapping with \triangle (on the right). The result is to swap the square, in position 12, for the triangle, in position 10. We would thus have a final set of positions with the circles in positions 2, 3, and 4, the squares in positions 6 and 10, and the triangles in positions 12, 13, and 14. This solution is shown in Figure 6(b).

⁴For simplicity, the term "multiparty OROD (with MP)" means execution of OROD among multiple parties with the constraint MP.

Figure 6 Execution of OROD Protocol Among Three Parties



Note that under the MP constraint, we only have one swap that happens, i.e., the first step where \circ (on the left) swapping with \triangle (on the right). The effect is to swap the triangle in location 6 for the circle in position 10. Next consider swaps between \circ (on the left) swapping with \square (on the right). After this step, the next two steps result in no swaps, given the MP constraint. This is shown in Figures 6(c) and 6(d). Thus Figure 6 presents a situation where the number of points swapped under an optimal solution (4) > the number under multiparty OROD with MP (2).

We see that multiparty OROD may reveal less information, as part of the solution, than a broker-generated solution. In addition, knowing that a solution is optimal reveals a different kind of information. Given the points it receives, a party can learn intervals within which there are no other parties' points:

- $I_{R_i^*}$ (I_{R_i}): regions where there are no other parties' points (R_i^* and R_i are sets of points party i receives from an optimal solution and solution computed from the multiparty OROD with MP, respectively).

Refer to Figure 6(b). From this optimal solution, \circ learns $I_{R_{\circ}^*}$: between locations 2 and 4, there are no other parties' points. On the other hand, from the solution in Figure 6(d), \circ learns $I_{R_{\circ}}$: Between locations 2 and 6, there are no other \triangle points. $I_{R_{\circ}}$ cannot convey any certain information regarding \square points because \circ is not allowed to swap the previously swapped point (at location 6) with one of the \square points at location 4. Depending on the execution ordering, multiparty OROD may or may not produce an optimal solution. As a result, $I_{R_{\circ}}$ always contains less

certain information than $I_{R_{\circ}^*}$. From this point of view, we can conclude that a solution computed by multiparty OROD with MP does not always leak more information than an optimal solution. \square

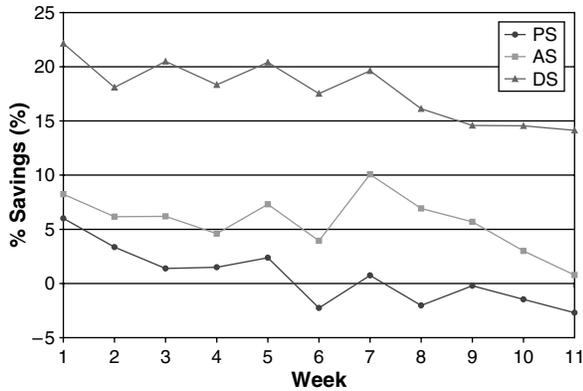
3. Experimental Analysis

In this section, we present results obtained when the algorithms described in this paper were applied to an empirical data set provided by a trucking company, representing LTL (less than truckload) loads during the first quarter of 2003. The purpose was to observe the empirical effect of running the algorithm on a data set where we have (a) individual points in two dimensions, and (b) pick-up and drop-off pairs associated with each shipment. The data set contains 81,842 data points (including both pick-up and delivery points) spanning 11 weeks. One issue in the data set is that in the actual problem, what we have to swap is *loads*, i.e., pick-up and delivery pairs, whereas the algorithm picks *points* to swap. Thus the results of the empirical tests provide an indication of the performance of the algorithm under more realistic problem conditions.

When the algorithms were executed for two companies, we randomly assigned some pick-up points to company A and the rest to company B. Based on the assigned pick-up points, each pick-up and delivery pair for each week was assigned to company A or B. Within a given weekly data set, we considered a single truck completing all pickups and then all deliveries. Actual distance being dependent on the starting point of each company, we used the leftmost pick-up point in one dimension as the starting point of company A and the rightmost point for company B. Before swapping any points, we obtained a preswap tour length (in two dimensions) of the pickup and delivery for each company by using the 2-opt algorithm (Lin and Kernighan 1973). We also used the same optimization method to calculate the tour length after swapping.

We considered three swapping methods for our analysis. One is *pair swapping* (PS), where we swap only if both pick-up and delivery points are beyond a certain point. Another is *average swapping* (AS), where companies swap data based on the average of pick-up and delivery points. As the number of delivery points is more than three times the number of pick-up points in the data set, we also consider *delivery swapping*

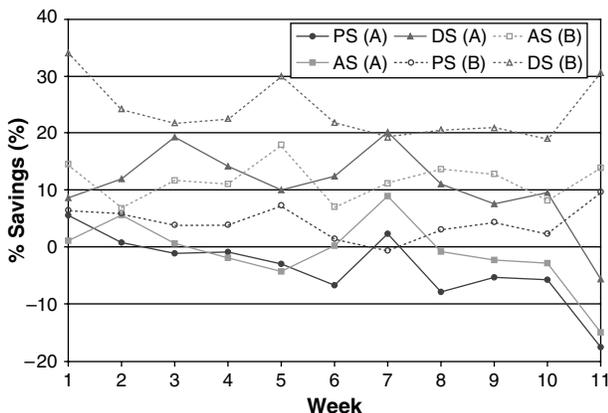
Figure 7 Savings in Total Distance



(DS) based on the delivery point. Figure 7 gives the savings in total tour length by week. Although none of our swapping method guarantees positive savings, the average savings in total distance are 0.6% for PS, 5.7% for AS, and 17.8% for DS. The results thus show that DS performs best in our empirical test. The main reason for this result is that the major distance driven in our data set is the trip to delivery points from primary pick-up locations, as there are more deliveries than pickups. This long tour has the most opportunity for improvement.

Figure 8 gives the savings by party—again it is clear that destination swapping, by optimizing the longer tour, gives the best results. It also shows an interesting real-life situation. Because *pairs* are being swapped, rather than individual points, it is possible that the tour length increases. The optimality result of Theorem 1 holds for a single tour but becomes a

Figure 8 Savings in Individual Distance



heuristic when both pickup and destination must be swapped. However, it is an effective heuristic: In these data, destination swapping always gives a global benefit and almost always benefits both parties.

3.1. Empirical Results of Multiparty Collaboration

We also evaluated the secure multiparty protocol of §2.5 on the data set. The experiments were conducted under two different settings: execution among three parties and execution among five parties. Because of the constraint MP, different execution orderings may produce different results. Therefore, under each assignments of the data (a trial), we tried multiple execution orderings. This was repeated for three trials. We present the best case, the worst case, and the average savings (in %) across all possible orderings for each trial.

Data for all experiments presented here come from the 1,470 pick-up and delivery pairs in the first week of the data set. Each party is given the same number of pick-up/delivery pairs. Based on the superior two-party performance of the DS strategy, we chose that as the heuristic.

3.1.1. Empirically Measured Savings with Three Companies.

In this section, we consider the case when the weekly data was split (randomly) across three companies for a single trial. As described earlier, there are six possible sequences of two party algorithm executions, which are represented by the possible permutations of the ordering of three companies, i.e.,

- {(C1 C2), (C1 C3), (C2 C3)},
- {(C1 C2), (C2 C3), (C1 C3)},
- {(C1 C3), (C1 C2), (C2 C3)},
- {(C1 C3), (C2 C3), (C1 C2)},
- {(C2 C3), (C1 C2), (C1 C3)}, and
- {(C2 C3), (C1 C3), (C1 C2)}.

In the description above, (C1 C2) denotes the execution of the two-party protocol between C1 and C2 with constraint MP being satisfied. The secure multiparty algorithm was executed for all six of these possible sequences. This was repeated for three trials.

Table 1 shows statistics for each trial and provides the maximum, minimum, and average savings for

Table 1 Percentage Savings by Company and Trial for Swap Among Three Companies

	Trial 1 (%)			Trial 2 (%)			Trial 3 (%)			Overall average (%)
	Min	Max	Avg	Min	Max	Avg	Min	Max	Avg	
C1	16.48	35.00	24.02	23.09	31.91	26.16	15.52	23.22	19.20	23.12
C2	9.85	27.11	16.48	-6.35	8.69	1.22	1.98	14.40	8.54	8.75
C3	24.40	29.45	24.40	27.65	31.09	27.65	31.49	35.62	31.49	27.85

each of the three companies (C1, C2, and C3) across the six possible execution orderings. For some execution orderings and initial data assignments, a company may not see benefits in the original problem context (e.g., see the minimum % change for C2 in Trial 2). But, in most cases, on average, every company observes cost reductions varying from 8.7% to 27.9%.

3.1.2. Empirically Measured Savings with Five Companies. We ran the same experiment with five companies. Thus, we divided the data randomly into five equal subsets, one for each company. Because each company tries to swap with every other company, we have

$$S = \{(C1\ C2), (C1\ C3), (C1\ C4), (C1\ C5), (C2\ C3), (C2\ C4), (C2\ C5), (C3\ C4), (C3\ C5), (C4\ C5)\},$$

i.e., 10 possible pairwise swaps. Because of the number (10!) of possible orderings of these swaps, we randomly chose 10 orderings for each of the three trials.

The following steps were executed for each trial:

- Randomly partition the weekly data set into five data subsets (one per company).
- Execute the secure multiparty protocol 10 times, each time randomly choosing the ordering to run. (a random permutation of all elements in S.)
- Record statistics of each of the 10 random execution orderings.

The above steps were repeated three times (trials). Table 2 provides, for each company, and for each trial, the minimum, maximum, and average savings across the 10 executions. The table shows that the average savings vary from 9.05% to 30.63%. As observed earlier, there are cases where companies do not see improvements for a particular sequence of orderings. Note that the algorithms we described always guarantee savings in the space-filling curve mapped dimension. The results show that the gap between the space-filling curve and its link to the pick-up and delivery travel distance in the two-dimensional problem context may result in an increase in travel distance in the original travel metric in some cases.

4. Conclusion and Future Work

We have described an algorithm that enables independent companies to identify opportunities for collaboration without sharing unnecessary data. The theoretical analysis shows that the algorithm ensures that no information can be inferred except that from the shared data, and that truth telling is incentive compatible. In addition, application of the algorithm to an empirical data set indicates a substantial potential impact. Papers such as Dodis et al. (2000) in the cryptography literature state that the creation of protocols where it is incentive compatible to correctly follow the protocol and truthfully provide the data poses interesting theoretical challenges. Our paper

Table 2 Percentage Savings by Company and Trial for Swap Among Five Companies

	Trial 1 (%)			Trial 2 (%)			Trial 3 (%)			Overall average (%)
	Min	Max	Avg	Min	Max	Avg	Min	Max	Avg	
C1	4.59	19.50	15.75	18.50	30.59	24.63	13.30	26.09	18.82	19.73
C2	6.70	26.77	16.77	0.66	20.53	12.27	-13.89	14.39	-0.13	9.64
C3	-8.87	28.78	5.58	3.08	26.59	15.23	-4.35	22.52	8.59	9.80
C4	4.35	34.60	15.41	-14.57	14.67	-4.11	0.00	36.96	15.85	9.05
C5	27.83	44.90	33.17	7.43	35.46	27.39	25.77	36.43	31.34	30.63

provides an example of an algorithm designed for an operations management problem context. We suggest that exploration of applied operations management problem contexts can provide a rich source of research problems as well as potential for practical applications.

The specific problem described in this paper is a step toward a greater integration of data encryption techniques and theory to problems in operations management. Another example is given by Atallah et al. (2003), who address capacity allocation and auctions—a supplier-consumer interaction rather than the peer-to-peer interaction presented in this paper. An interesting open research question concerns the potential trade-off between the use of a heuristic approach to problem solving that admits a tight security property versus a closer-to-optimal algorithm that permits data leakage. We leave such exploration to future research.

Another area for research is applying the concept of incentive compatibility to secure multiparty computations. Current secure multiparty computation models do not take the reward structure of the parties into account, leading to interesting results where the most secure model (malicious model) may not provide the needed security in practice. Integrating this concept into evaluation of an algorithm, as in Appendix E, leads to new opportunities for practical application of secure collaborative computation.

Acknowledgments

The authors thank Reha Uzsoy for contributions made to the formulation of this problem in initial discussions of the work. This work was supported by a grant from the Purdue University's e-Enterprise Center at Discovery Park.

Appendix A. Algorithm Details

Appendix A.1. Notations

- $S^k = \{s_1^k, s_2^k, \dots, s_m^k\}$, a set of m k -dimensional points,
- $O^k = \{o_1^k, o_2^k, \dots, o_n^k\}$, a set of n k -dimensional points,
- $SF \equiv A$ space filling curve,
- $[a, b] \equiv$ The range of SF .

Appendix A.2. Formal Definitions of Extreme Points

DEFINITION 1 (EXTREME POINTS WITH RESPECT TO X AND d_x).

- X is a set of points, d_x is a direction (i.e., left or right) related to X ;
- An extreme point is a point in X that is the farthest away from d_x , and

- i th extreme point is a point in X that is the i th furthest away from d_x .

DEFINITION 2 (*Extreme_POS*(X, d_x, i)).

- X is a set of points, d_x is a direction (i.e., left or right), and i is an integer;
- The function returns the position of the i th extreme points with respect to X and d_x , and
- If $i > |X|$, the function returns a position beyond the range limits (e.g., $-\infty, +\infty$) in the direction d_x .

Formally, *Extreme_POS*(X, d_x, i) is:

Reorder X in ascending order according to position

if $d_x = \textit{right}$ **then**

Return the position of i th item of X ($+\infty$ if $i > |X|$)

else

Return the position of the $(|X| + 1 - i)$ th item of X ($-\infty$ if $i > |X|$)

end if.

Appendix B. Secure Multiparty Computation

Substantial work has been done on secure multiparty computation. The key result is that a wide class of computations can be done securely under reasonable assumptions. Any function that can be represented by a polynomial circuit in terms of the number of input bits can be evaluated in reasonable time. We give a brief overview of this work, concentrating on material that is used in the paper. The definitions given here are from Goldreich (2004a). For simplicity, we concentrate on the two-party case. Extending the definitions to the multiparty case is straightforward.

Appendix B.1. Definition of Secure Multiparty Computation

The basic idea behind showing that a protocol is secure is that no party can learn more than in a “trusted broker” model, where all parties give their input to the broker, and the broker computes and returns the result. The problem with the trusted broker model is that the broker learns everything. A secure multiparty computation (SMC) accomplishes this *without* a broker learning anything.

The problem is that in the trusted broker model, the only information a party receives is the final result. However, in a real computation, more information must be exchanged (for example, the results of the *Extreme_POS* comparison). The goal is to prove that the information exchanged does not reveal anything other than what can be inferred from the final results.

The security of an algorithm is proved as follows: If a party, given its own input and the final result, can *simulate* the information received during an actual execution of the protocol, then it has not learned anything from the protocol. The proof of Theorem 2 demonstrates such a proof by simulation.

Note that the values chosen by the simulator may not be the same as the values seen during any *particular* execution of the protocol. What matters is that over many executions

of the protocol on the same inputs, the *distribution* of values seen (which varies based on the parties' choice of random numbers, such as encryption keys) is the same as the one generated by simulation.

Formal definitions for secure multiparty computation, and the proof that the simulation argument in fact guarantees security equivalent to the "trusted broker" model, can be found in Goldreich (2004a).

In summary, a secure multiparty protocol will not reveal more information to a particular party than the information that can be induced by looking at that party's in- and output.

Appendix B.2. Cryptographic Definitions

Security of the oblivious transfer relies on certain functions being infeasible to compute and the ability to simulate values seen during the protocol by choosing a random number. Two key assumptions used are given below.

Computational Diffie-Hellman Assumption (Diffie and Hellman 1976). Assume that p is a very large prime number and g is the generator of its multiplicative group (i.e., every number between $1, \dots, p-1$ can be written as $g^k \bmod p$ for some k between $1, \dots, p-1$). The computational Diffie-Hellman assumption states that given $g^a \bmod p$ and $g^b \bmod p$ (note that a and b are not given), there is no efficient way to compute $g^{ab} \bmod p$.

Random Oracle Assumption. In the construction of the protocol, we will use a cryptographic hash function H . We assume that this function is known to all parties (e.g., SHA (NIS 1995)), and it maps its input to what appears to be (is computationally indistinguishable from) a random output. The computationally indistinguishability of the output of commonly used cryptographic hash functions from a random output is not formally provable (yet) but is a common assumption in the cryptography literature.

The computational Diffie-Hellman assumption is the basis for the Diffie-Hellman key exchange protocol, and it is similar to the assumptions relied on by other public-key cryptosystems. The random oracle assumption is also relied on by many cryptosystems (Goldreich 2004b). If either does not hold, many commonly used encryption schemes could be broken.

Appendix B.3. 1-Out-of- n Oblivious Transfer

We now provide an approach to extend the 1-out-of-2 protocol to generate a 1-out-of- n protocol. The basic idea is that the values in the boxes are masked using an exclusive or with $\log n$ bits. Each bit is an encrypted value, with two keys for each bit. One key is used if that bit in the binary representation of the desired box number is 0, the other if it is 1. Encryption requires two inputs, the key and value, to be encrypted; for cryptographic reasons, the value being encrypted must be different for each box but known to both parties—the box number itself meets this criteria.

\bigcirc sends the value EB_i (exclusive or of encryptions and contents) for every box to Δ . To recover the key, Δ chooses a 0 or 1 as the value it will be able to recover for each bit (corresponding to the binary representation of the desired box number). Each bit of the key is transferred using 1-out-of-2 oblivious transfer; Δ will only be able to recover the key where all the bits match the actual box to be opened, and this will only correctly decrypt the desired box number. For any other box, the result of at least one bit would appear random (because of trying to get the "wrong" bit using 1-out-of-2 oblivious transfer for that bit).

We show this protocol for 16 boxes (four bits); assume that \bigcirc has B_1, \dots, B_{16} , and Δ wants to learn B_7 . We also use an encryption function E (e.g., DES (NIS 1988)) as the hash function H meeting the random oracle assumption; as will be seen below, we need to invert the hash as part of this protocol.

1. \bigcirc generates four key pairs

$$(K_1^0, K_1^1), (K_2^0, K_2^1), (K_3^0, K_3^1), (K_4^0, K_4^1),$$

where each $K_j^{\sigma[j]}$ is a randomly chosen key for E . For $1 \leq i \leq 16$, with binary representation $(\sigma[1], \sigma[2], \sigma[3], \sigma[4])$, \bigcirc creates $EB_i = B_i \oplus E_{K_1^{\sigma[1]}}(i) \oplus E_{K_2^{\sigma[2]}}(i) \oplus E_{K_3^{\sigma[3]}}(i) \oplus E_{K_4^{\sigma[4]}}(i)$.

2. \bigcirc sends all EB_i to Δ .

3. Because 7 (the box number Δ wants to open) has binary representation 0111, using four executions of OT_1^2 , Δ learns $K_1^0, K_2^1, K_3^1, K_4^1$.

4. Δ retrieves B_7 by calculating $EB_7 \oplus E_{K_1^0}(7) \oplus E_{K_2^1}(7) \oplus E_{K_3^1}(7) \oplus E_{K_4^1}(7)$.

It can easily be proven that if the encryption scheme and OT_1^2 are secure, then the above algorithm is secure. Because it only retrieves one key from each pair, Δ accurately decrypts at most one message.

For simplicity of presentation, we have shown a construction that scales linearly (because \bigcirc sends the encrypted value of every box). However, there are known logarithmic time protocols for secure comparison that can be used for step 6 of Algorithm 1; for details, see Yao (1986), Goldreich (2004a), or Ioannidis and Grama (2003).

Appendix B.4. Related Work: k th-Ranked Element

Aggarwal et al. (2004) present a protocol for securely computing the k th-ranked element among lists held by multiple parties. A special case of our problem, in which the number of points of each party is known to all, can be solved by computing the k th-ranked element. Knowing the number of points held by (in particular) the party at the "low" end of the line gives the value to use for k ; their protocol would give the location of this k th element. This location would partition the line; the parties would then know to swap points to the right (left) of that location to get the optimal solution (see Lemma 3.1).

We do not make the assumption that the number of points held by each party is public; we limit the information disclosed to *only* the swapped points. Without sharing

knowledge of the number of points belonging to at least one party, finding the k th element will not solve the problem. Although it would be possible to use the k th-ranked element protocol iteratively to find the appropriate number of points to swap, this would require revealing the location of the $k + 1$ st (nonswapped) point. Thus, their protocol is not suitable for our application.

Appendix C. Proof of Correctness

THEOREM 1. *The OROD algorithm terminates with the optimal number of points swapped; i.e., after swapping, there are no two points that could be swapped to give a lower tour length.*

PROOF. First, note that partitioning the line gives the lowest tour length: With the line partitioned, neither party needs to drive past the partition, and any swap would require that each party drives past the partition, extending the tour length (this is shown in Lemma 3.1). A simple inductive proof demonstrates that the process terminates with the number of swaps required to achieve this optimal (in one dimension) result.

First, note that it is easy to see that $lbound \leq i \leq ubound$ always holds.

We will show that the following always hold.

1. Swapping fewer than $lbound$ points does not partition the line, and
2. Swapping $ubound$ or more points does not partition the line.

First, they are trivially true initially (since we cannot swap a negative number or ∞ points).

If $Extreme_POS(S, l, i) > Extreme_POS(O, r, i)$, then swapping fewer than i points will not partition the line (since at least the i th points would still be on the wrong side of the partition). In this case, $lbound$ gets set to i , which does not violate 1. Since $ubound$ is unchanged, 2 still holds.

If $Extreme_POS(S, l, i) \leq Extreme_POS(O, r, i)$, then swapping i or more points will not partition the line, since the i th points are on the wrong side of the partition. In this case, $ubound$ gets set to i , which does not violate 2, and since $lbound$ is unchanged, 1 continues to hold.

If $ubound - lbound = 1$, then by 1 and 2, exactly $lbound$ points must be exchanged. Since at this point $i = \lfloor (lbound + lbound + 1)/2 \rfloor = lbound$, the result is exactly the number that must be exchanged to partition the line.

All that remains is to show that the algorithm does terminate. Although $ubound = \infty$, i doubles and at some point must exceed the number of points on the wrong side of the partition (since the number of points is finite). At this point, $ubound$ gets set to a finite number. As long as $ubound - lbound > 1$, $lbound < \lfloor (lbound + lbound + 1)/2 \rfloor < ubound$, and $lbound < i * 2$, forcing i to be set such that $lbound < i < ubound$. Since at the next iteration, either $lbound$ or $ubound$ is set to i , either $lbound$ increases or $ubound$ decreases. Since the values are integral, a finite number of iterations are possible before $lbound + 1 \geq ubound$, and the algorithm terminates. \square

The number of iterations can be verified to be logarithmic in the number of points swapped.

THEOREM 4. *Execution of OROD among k parties eventually reaches optimal: k partitions (where $k \geq 2$).*

PROOF. This is proven through strong induction on the number of parties.

Base case: $k = 2$. True from Theorem 1.

Inductive step: Let integer $t > 2$, and assume the claim is true for all $k < t$. We need to prove that the claim is true when $k = t$.

Let $k = t - 1$. From the induction hypothesis, the execution among the k parties creates k partitions. Then let another party (say l) join the k parties and i be any party among the k parties or k partitions. Leave i out and execute the protocol among the remaining $k - 1$ parties plus l . Then we have a new set of k partitions.

Considering i with the other parties, if the protocol is executed between i and any one of the parties whose data set range does not overlap that of i , the protocol does nothing. Let p be the number of parties whose data set ranges do overlap that of i .

If $p < k$, the execution of the protocol among these p parties plus i eventually leads to a set of disjoint partitions. These partitions, plus the previously disjoint partitions with i , form $k + 1$ disjoint partitions.

If $p = k$, during the execution between p and either the leftmost or rightmost party, the partition of the leftmost or the rightmost party will shrink. Because there are a finite number of points for each party, the number of times the partition can shrink is finite. Therefore, eventually, no party's partition overlaps with that of the leftmost or the rightmost party. The execution among the rest of the parties will lead to k partitions. Adding the partition of either the leftmost or the rightmost party, we have $k + 1$ disjoint partitions. As a result, the claim is true when $k = t$.

Because the base case and inductive step are true, the claim is true for all $k \geq 2$. \square

Appendix D. Proof of Security

THEOREM 2. *Algorithm (1) is secure under the semihonest definition of secure multiparty computation.*

PROOF. Communication occurs only at line 6. To prove the protocol is secure, we only need to show that given the information received at the end, the intermediate comparison results can be simulated. (The comparison is computed in a secure manner as described in §2.3.1, so it reveals nothing but the comparison result.)

Algorithm 2 gives the simulator for line 6 for the party at the left end of the line, S^1 .

ALGORITHM 2. Simulator for OROD Protocol

Require: S^1, n

1: $i \leftarrow 1$

2: $lbound \leftarrow 0$

```

3:  $ubound \leftarrow \infty$ 
4: while ( $i \neq n$ ) do
5:   if  $i < n$  then
6:      $Known: Extreme\_POS(S^1, l, i) >$ 
        $Extreme\_POS(O^1, r, i)$ 
7:      $lbound \leftarrow i$ 
8:      $i \leftarrow i * 2$ 
9:   else
10:     $Known: Extreme\_POS(S^1, l, i) <$ 
       $Extreme\_POS(O^1, r, i)$ 
11:     $ubound \leftarrow i$ 
12:     $i \leftarrow \lfloor (lbound + ubound) / 2 \rfloor$ 
13:   end if
14: end while
    
```

The simulator works as follows: let n be the number of points to swap and i be an integer. If i is less than n , S^1 knows i has yet to be maximal, and the protocol will double the size of i . On the other hand, if i is greater than n , S^1 knows i is too big to be maximal, and the protocol will decrease the size of i to a value that is in the middle of its upper and lower bounds. If i is the same as n , the protocol terminates.

From the simulator, it is obvious that the shared results from the protocol, plus a party's input, are sufficient to precisely simulate each execution of the protocol. Therefore, because the simulation process and execution of the protocol are computationally indistinguishable, the OROD protocol is secure. \square

Appendix E. Proof of Honesty

Cryptographic approaches and game-theoretic solutions are complementary to each other. Specifically, Dodis et al. (2000) explicitly demonstrate how cryptographic approaches can be used to solve game-theoretic problems by showing how to construct a two-player game that achieves the same payoffs as a mediated two-player game (mediated by an "honest broker"). We have shown how this approach can be used to solve a practical problem, and now prove that the payoff for this problem guarantees incentive compatibility.

There are two levels at which incentive compatibility can be shown. The first is to show that it is incentive compatible to provide the correct data for the algorithm. This is the traditional definition used in economics. The second is to show that it is optimal for each party to correctly follow the protocol; i.e., any attempt to gain more information by deviating from (compromising the security of) the protocol is self-defeating or will result in discovery by the other party. This is related to the distinction between "semihonest" and "malicious" in secure multiparty computation definitions. A protocol that does not meet the first level (giving incorrect data) is vulnerable even if it meets the standard of security against malicious parties. A protocol that meets the second form of incentive compatibility need not be secure against

malicious parties; it will be in the party's best interest to follow the protocol, and the semihonest security ensures that no excess information is disclosed. We will show that our protocol meets the first type (correct data), and with the exception of the secure comparison (for which we describe protocols secure against a malicious party), our protocol also meets the second. Thus we ensure both correct results and prevent disclosure of information by ensuring that it is in a party's best interest to follow the protocol with correct data.

We first assume that the space-filling curve heuristic is jointly agreed on by all parties because it is already built into several available mechanisms for vehicle routing. It is possible that the heuristic may cause an optimal one-dimensional result to be nonoptimal and possibly not beneficial for all parties in two dimensions. For one party to know that the heuristic is "bad" for a particular data set requires a knowledge of the other party's data; we assume that parties do not know each other's information. By the same argument, a party would not be able to manipulate the choice of heuristic to its advantage.

We now show that cheating in the protocol results in a nonoptimal one-dimensional solution and that such a solution is in neither party's best interest. We start with the following lemma.

LEMMA 3.1. *If both parties swap an equal number of points, each will achieve the greatest benefit if the points swapped partition the line.*

PROOF. Formally, the cost for a party P , $Cost(P)$ is $\{Let S$ represent the list of points owned by $P\}$.

Reorder S in ascending order according to position

$Cost = |Position(S_{|S|}) - Position(S_d)|$.

Let C be the dishonest party and H be the honest party. Let C_o be C 's points before executing the protocol, C_t be C 's points after a correct execution of the protocol, and C_c be C 's points after a false execution. H has the corresponding sets as well. Note that $|C_t| = |C_c| = |C_o|$, since an equal number of points are given up and received.

Let $C_{t|C|}$ be the farthest point for C in the true set and $C_{c|C|}$ be the farthest in the false set. Assume $Position(C_{c|C|}) < Position(C_{t|C|})$. Since $C_{t|C|} \notin C_c$, $C_{t|C|} \in H_c$. Since $|H_c| = |H_t|$, $\exists X \in H_t - H_c$, and since no points are lost, $X \in C_c$.

Since the line is partitioned in a true execution (Theorem 3), $Position(C_{t|C|}) < Position(H_{t|H|})$. But this gives

$$Position(C_{c|C|}) < Position(C_{t|C|}) < Position(H_{t|H|}) \leq X.$$

This gives a contradiction: $Position(C_{c|C|})$ cannot be the farthest point for C . \square

THEOREM 3. *The OROD protocol is incentive compatible; i.e., any failure to correctly follow the protocol results in the dishonest party being worse off, or leaves the result unaffected and gives no additional information to the dishonest party.*

PROOF. Theorem 1 shows that the algorithm produces a partition. A partition is optimal for *both* parties: If a dishonest party gets a point on the wrong side of the partition, it must travel beyond the partition, extending its tour length. We will show that any cheating leads the cost (defined as above) to increase, thus producing an incentive for honesty.

First, note that a dishonest party cannot “invent” false points to swap. If a dishonest party tries to give away a point that does not really exist, the cheating will be detected when the honest party goes to pick up the nonexistent load. Contractual obligations can enforce disincentives to such cheating. The same holds true for giving up fewer or more points than agreed on by the algorithm.

This leaves two avenues for dishonest behavior:

- Cause the algorithm to arrive at the wrong value, and
- Give away different points than those determined by the algorithm.

From Lemma 3.1, we have shown that any cheating that causes points to be swapped other than those swapped in a correct execution of the algorithm will be disadvantageous to the dishonest party.

Any cheating in the comparison at step 6 of Algorithm 1 does not reveal any information to C . At most, C can cause the comparison to give the wrong result. Looking at the proof of Theorem 1, we see that at any point, the correct number of items to swap S lies in the range $lbound \leq S < ubound$. Suppose the correct comparison result is $>$; then we are supposed to set $lbound \leftarrow i$, so we know $S \geq i$. If C causes the comparison result to be wrong (\leq), it would instead set $ubound \leftarrow i$. However, since $ubound$ can never increase, and $S \geq i$, the result of the algorithm will not be the true result. An analogous construction works with $lbound$. Therefore, any dishonesty in the execution of the algorithm either causes both parties to come to the wrong answer, or if C behaves incorrectly in another step, to come to different answers. As discussed before, if C and the honest party give up a different number of points, C will be detected to have cheated and contractual disincentives would apply. By Lemma 3.1, swapping the wrong number of points is disadvantageous.

The second possibility is that a dishonest party comes to the correct number to swap, but gives up the wrong points. Assume a point X is kept that should have been swapped. The same approach used in Lemma 3.1 can be used to show that $position(X) > position(C_{t|C})$, increasing the cost for the cheating party.

Thus, any cheating either increases the cost for the dishonest party, or is discovered and provable by the honest party. Honesty is thus incentive compatible. \square

Note that the proof is only for one dimension. It is thus possible that parties may not be better off in two dimensions. This only happens when the space-filling curve would be a poor heuristic for two-dimensional optimization, which we expect to be rare (Bartholdi and Platzman 1982). Although we have assumed no prior information

regarding the location of loads of other parties, there may be cases where prior information, about other parties load locations, may permit individual parties to infer more from the swaps. The corresponding adjustments in the algorithm and proofs of incentive compatibility are left for future work.

Appendix F. Proof of Hash Function Claim

Formally, let h be an unkeyed hash function with the following properties (Menezes et al. 1996):

1. *Ease of computation*: Given h and an input x , $h(x)$ is easy to compute.

2. *Preimage resistance*: For essentially all prespecified outputs, it is computationally infeasible to find any input, which hashes to that output, i.e., to find any preimage x' such that $h(x') = y$ when given any y for which a corresponding input is not known.

3. *Second-preimage resistance*: It is computationally infeasible to find any second input that has the same output as any specified input, i.e., given x , to find a second-preimage $x' \neq x$ such that $h(x) = h(x')$.

4. *Collision resistance*: It is computationally infeasible to find any two distinct inputs x, x' which hash to the same output, i.e., such that $h(x) = h(x')$.

Let m be the maximum number of candidate points a party may have, and k the number of parties. Let C_i be the candidate set of party i . Denote $H_i = \{h(c_i)\}$, $\forall c_i \in C_i$. (Formally, c_i is a unique identifier, such as a tracking number, for the shipment.) We also pad H_i with hashes of “fake” identifiers such that $|H_i| = k \cdot m$ (the reason will be apparent later). (By 2 above, the hashes of fake identifiers will be indistinguishable from those of real identifiers.) Before executing the OROD protocol with the MP constraint among the k parties, party i broadcast H_i to the other parties.

Suppose party i and party j execute the OROD protocol. On receiving the swapped points from party j , party i computes the hash values of the identifiers of these points and searches through H_j to see if the received values are valid (contained in H_j). If these values are not valid, party i can conclude that party j is not executing the protocol correctly and abort the protocol. Simultaneously, party j can perform the same task to see if party i has behaved maliciously. Each party then generates an $h(c)$ for each of the points it gave, pads this set with “fake” hashes from H_i (H_j) to create a set γ_i (γ_j) of m hashes, and broadcasts the set. All parties remove those values in γ_i (γ_j) from the corresponding original H_i (H_j) to ensure the points are not swapped when party i (j) executes the protocol with them. Note that the two parties will perform this check so that they can be sure multiple shippers will not attempt to pick up the same shipment, as per the MP constraint described earlier.

Collision resistance ensures that points from party i will appear only in the set H_i . (Second-preimage resistance ensures this holds even if a party j gets to see H_i before sending its own set.) This guarantees that a party can only

swap its own points; giving away a point received from another party will be detected. Preimage resistance and second-preimage resistance of the hash, combined with the fixed size $k \cdot m$ of the received H_i , ensure that the hash does not disclose information. If party j receives H_i in the protocol, it could simulate H_i by hashing the identifiers of the points actually received from i and padding to size $k \cdot m$ with values drawn randomly from the domain of the hash, allowing proof that nothing is disclosed by H_i (the same argument holds for each γ_i received after each party swaps). As stated in §2.5, with the MP constraint, each party only needs to execute OROD with the other parties once. Thus, the number of “fake” hash values in H_i is sufficient to pad each γ_i for $k - 1$ rounds if needed.

References

- Aggarwal, G., N. Mishra, B. Pinkas. 2004. Secure computation of the k th-ranked element. *Proc. IACR Eurocrypt (EUROCRYPT04)*, Interlaken, Switzerland, 40–55.
- Atallah, M. J., H. G. Elmongui, V. Deshpande, L. B. Schwarz. 2003. Secure supply-chain protocols. *IEEE Internat. Conf. E-Commerce*. Newport Beach, CA, 293–302. <http://doi.ieeecomputersociety.org/10.1109/COEC.2003.1210264>.
- Bandler, J. 2003. How seagoing chemical Haulers may have tried to divide market. *Wall Street J.* (February 20) A1.
- Bartholdi III, J. J. 2006. Some combinatorial applications of spacefilling curves. <http://www2.isye.gatech.edu/~jjb/mow/mow.html>.
- Bartholdi III, J. J., L. K. Platzman. 1982. An $O(N \log N)$ planar travelling salesman heuristic based on spacefilling curves. *Oper. Res. Lett.* **1**(4) 121–125.
- Bartholdi III, J. J., L. K. Platzman. 1988. Heuristics based on spacefilling curves for combinatorial problems in Euclidean space. *Management Sci.* **34**(3) 157–160.
- Diffie, W., M. Hellman. 1976. New directions in cryptography. *IEEE Trans. Inform. Theory* **IT-22**(6) 644–654.
- Dodis, Y., S. Halevi, T. Rabin. 2000. A cryptographic solution to a game theoretic problem. *Advances in Cryptology—CRYPTO 2000*. Springer-Verlag, Berlin, Germany.
- Goldreich, O. 2004a. General cryptographic protocols. *The Foundations of Cryptography*, Vol. 2. Cambridge University Press, New York.
- Goldreich, O. 2004b. Encryption schemes. *The Foundations of Cryptography*, Vol. 2. Cambridge University Press, New York.
- Goldreich, O., S. Micali, A. Wigderson. 1987. How to play any mental game—A completeness theorem for protocols with honest majority. *19th ACM Sympos. Theory Comput.* ACM Press, New York, 218–229.
- Ioannidis, I., A. Grama. 2003. An efficient protocol for Yao’s millionaires’ problem. *Proc. Hawaii Internat. Conf. System Sci. (HICSS-36)*, Waikoloa Village, HI, 205–210.
- Kalagnanam, J., M. Trumbo, H. S. Lee. 2000. The surplus inventory matching problem in the process industry. *Oper. Res.* **48**(4) 505–516.
- Keskinocak, P., S. Tayur. 2001. Quantitative analysis for Internet-enabled supply chains. *Interfaces* **31**(2) 70–89.
- Lin, S., B. W. Kernighan. 1973. An effective heuristic algorithm for the traveling-salesman problem. *Oper. Res.* **21**(2) 498–516.
- Menezes, A. J., P. C. van Oorschot, S. A. Vanstone. 1996. *Handbook of Applied Cryptography*. CRC Press, Boca Raton, FL.
- Naor, M., B. Pinkas. 1999. Oblivious transfer and polynomial evaluation. *Proc. Thirty-First Annual ACM Sympos. Theory Comput.* ACM Press, Atlanta, GA, 245–254.
- Naor, M., B. Pinkas. 2001. Efficient oblivious transfer protocols. *Proc. SODA 2001 (SIAM Sympos. Discrete Algorithms)*, Washington, D.C.
- NIS. 1988. Data encryption standard (DES). Technical Report FIPS PUB 46-2, National Institutes of Standards and Technology, Gaithersburg, MD. <http://www.itl.nist.gov/fipspubs/fip46-2.htm>.
- NIS. 1995. Secure hash standard. Technical Report FIPS PUB 180-1, National Institutes of Standards and Technology, Gaithersburg, MD. <http://www.itl.nist.gov/fipspubs/fip180-1.htm>.
- Wilson, R., R. V. Delaney. 2003. 14th Annual State of Logistics Report[®]: The case for reconfiguration. National Press Club, Washington, D.C. http://www.uwa.com/supply_a_005.pdf.
- Yao, A. C. 1986. How to generate and exchange secrets. *Proc. 27th IEEE Sympos. Foundations Comput. Sci.* IEEE Computer Society Press, Los Alamitos, CA, 162–167.