PRODUCTION AND OPERATIONS MANAGEMENT

POMS

Vol. 20, No. 2, March–April 2011, pp. 165–180
ISSN 1059-1478 | EISSN 1937-5956 | 11 | 2002 | 0165

DOI 10.3401/poms.1080.01188

Outsourcing Manufacturing: Secure Price-Masking Mechanisms for Purchasing Component Parts

Vinayak Deshpande, Leroy B. Schwarz
Krannert School of Management, Purdue University, West Lafayette, Indiana 47907-2056, USA
vinayak@purdue.edu, lschwarz@purdue.edu

Mikhail J. Atallah
Department of Computer Science, Purdue University, West Lafayette, Indiana 47907-2066, USA, mja@cs.purdue.edu

Marina Blanton
Department of Computer Science and Engineering, University of Notre Dame, Notre Dame, Indiana 46556, USA,
mblanton@cse.nd.edu

Keith B. Frikken
Department of Computer Science and Software Engineering, Miami University of Ohio, Oxford, Ohio 45056, USA,
frikkekb@muohio.edu

This paper develops and tests a privacy-preserving business process that supports the selection of a contract manufacturer by an original equipment manufacturer (OEM), and the determination of whether the OEM or the chosen contract manufacturer will procure each of the components to be used in the manufacture of the OEM's branded product. Our "secure price-masking (SPM)" technology contributes to procurement theory and practice in four significant ways: First, it preserves the privacy of every party's individual component prices. Second, SPM assures that the contract manufacturers will bid their own private purchase cost (i.e., not add a margin to their cost). Third, SPM is not invertible; i.e., none of the participants can "solve" for the private inputs of any other participant based on its own inputs and the outputs provided to it by SPM. Fourth, the posterior distribution of any other participant's private inputs is practically indistinguishable from its prior distribution. We also describe the results of a proof-of-concept implementation.

Key words: supply chain; outsourcing; procurement; price masking
History: Received: January 2008; Accepted: August 2008, after 2 revisions.

1. Introduction

Electronic product contract manufacturing has become a US $100 billion business worldwide. Indeed, although most consumers are ignorant about their existence, it is contract manufacturers—now, more generally known as "Electronic Manufacturing Service (EMS) providers"—with names like Flextronics and Jabil that manufacture most of the electronics branded as Dell, Motorola, IBM, Hewlett-Packard, etc.

This paper uses techniques from secure multiparty computation (SMC) and game theory to develop and test a business process for negotiating the procurement of component parts to be used by an EMS provider in the manufacture/assembly of the branded products of another company; i.e., an original equipment manufacturer (OEM). More specifically, this business process determines which company will procure which components and what prices will be paid for them.

As described by Amaral et al. (2006), some OEMs delegate the procurement of component parts entirely to the EMS under a "turnkey" arrangement. Amaral et al. (2006) point out several hazards associated with this practice. More sophisticated OEMs procure some or all of the component parts themselves and provide them to the EMS under a "price-masking" program (described below). Amaral et al. (2006) point out that price masking partially mitigates the hazards of "turnkey." However, *none* of these programs: (1) assure the privacy of both the OEM and the EMS's component prices; or (2) assure that the EMSs will bid their own purchase cost (e.g., not add a margin to their cost) in the negotiation process. These characteristics are provided by our "secure price-masking (SPM) mechanisms."

Outsourcing, of course, is not new: Manufacturing companies have historically delegated the fabrication of components and sub-assemblies to other manufacturers. However, as the term is used today, when an OEM like Xerox or Nokia outsources to an EMS, it delegates the *entire* manufacturing/assembly process. Moreover, as their name implies, EMS providers no longer limit themselves to contract manufacturing:

most also offer design, packaging, repair, maintenance, and logistics services.

Compared to in-house manufacture, EMSs offer many advantages to OEMs. Small OEMs typically do not own, and cannot afford, the equipment and expertise to manufacture their products. Large OEMs, which can afford to acquire their own equipment and expertise, often use EMSs to avoid increasing their investment in fixed assets and, possibly, reducing their return on investment. Indeed, in some cases, OEMs have sold their manufacturing assets to EMSs in order to reduce their asset base. See Amaral et al. (2006) for discussion of these advantages.

Whatever advantages EMSs offer to OEMs, they also pose hazards. See Amaral et al. (2006) for a general discussion. One major challenge is the loss of competitive advantage in purchasing components (e.g., semiconductors, circuit boards) from component suppliers (CSs). The challenge arises as follows: Every CS has a catalog price for each of the components that it sells. However, a CS's large customers—OEMs and EMSs—will usually have negotiated contracts to purchase their components at a discount; and, typically, these discounts are different for different customers. Now, consider the bill of material (BOM) for a cell phone that, say, Nokia, would like to have assembled by, say, Flextronics. Given the fact that each has negotiated its own discounts with the CSs, the component cost for the cell phone is likely to be different, depending on whether the OEM or the EMS procures them. Naturally, Nokia would like to minimize component cost. Hence, in the absence of price masking, Nokia and Flextronics share component cost information with one another. If Flextronics can purchase some components for less than Nokia can, then it does so. If not, then Nokia provides them to Flextronics. Under this practice, Nokia benefits from being able to purchase every component in its cell phone either at its own price or Flextronics' price, whichever is lower.

Despite its obvious short-term benefit to the OEM, note that because price information is shared, the EMS has learned which components its OEM customers are able to buy at lower prices than it can, and it has learned what those prices are. Henceforth, the EMS will use that information to negotiate the same or similar prices from those CSs; and then, offer those lower prices to *other* OEMs. As a consequence, Nokia has lost any component cost advantage it may enjoy over, say, Motorola and Ericsson. Furthermore, to the extent that Flextronics is thereby able to increase its volume of business with the CSs, Flextronics may be able to secure lower component prices than *any* of the OEMs. Finally, the CSs lose revenue because they are obliged to offer lower prices to more customers.

As a consequence, companies such as Motorola, IBM, and Sony have adopted "price-masking" programs:

- "Under the new global buy/sell program, variants of which can be found at top-tier OEMs like IBM Corp. and Hewlett-Packard Co., Motorola will negotiate prices for direct materials with its vendors, while 'masking' the terms from its contract manufacturers. Motorola will then resell the components to its EMS providers at what Metty, an ex-IBM executive and former head of supply-chain management for Motorola's Personal Communication Systems group, described last week as the 'street price.' The policy will affect about 80% of the company's $11 billion annual direct procurement spend, moving to 100% at some point in the future, Metty said. … 'Companies are giving away their competitive advantage associated with the bill of materials,' she said. 'That BOM, depending on your business model, is between 50% and 70% of a company's cost.' (Sullivan 2003)
- "Hewlett-Packard masks the price it pays for the parts from its EMS providers because it believes it gets world-class prices and wants to keep the prices confidential." Purchasing.com, June 6, 2004.

According to a 2004 iSuppli survey (Pick 2004), six out of 15 OEMs, i.e., 40%, responded that they are using price-masking programs with their EMS providers.

OEMs mask prices in basically three different ways: "buy/sell," "rebate," and "consignment." Under a buy/sell program, the OEM buys components from the CS and then sells them to the EMS at their "street" prices. Under a rebate program, the CS sells components to the EMS for its higher price and then rebates the difference to the OEM. A consignment program works like a buy/sell program except that the EMS maintains ownership of the component throughout the manufacturing/assembly process (see Figure 1). Amaral et al. (2006) describe each program in more detail and describe their advantages and disadvantages. Also see Lee and Tang (1996).

One fundamental aspect of *any* price-masking program is to determine whether the OEM or the EMS has the lower price for each component. In some cases, the EMS is willing to quote prices on each component on the OEM's BOM. If so, the OEM then "cherry picks" the items for which it has the lower price and provides them to the EMS under, say, a buy/sell program. Under this scheme, the EMS can infer which components the OEM is able to purchase at lower prices, but not what those lower prices are. More often, the EMS is not willing to disclose its prices. Hence, the OEM must guess which components to cherry-pick.

We next provide an overview of our approach for the SPM process.

Figure 1 Conventional Price-Masking Techniques

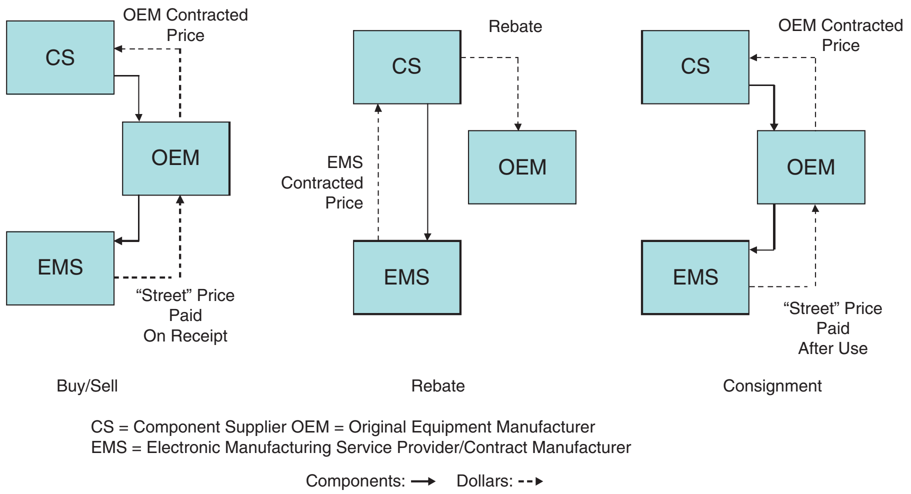
CS
OEM Contracted Price
OEM
EMS
"Street" Price Paid On Receipt
Buy/Sell
CS
Rebate
EMS Contracted Price
OEM
EMS
Rebate
CS
OEM Contracted Price
OEM
EMS
"Street" Price Paid After Use
Consignment
CS = Component Supplier OEM = Original Equipment Manufacturer
EMS = Electronic Manufacturing Service Provider/Contract Manufacturer
Components: →
Dollars: - - ▸

1.1. Overview/Summary

This paper uses techniques from SMC and game theory to develop and test a business process that selects which of N competing EMSs will assemble a given product for a given OEM and determine which components the OEM and the EMS will procure. The product's BOM is known to all the participants. The business process is structured as an auction in which each EMS privately submits "encrypted" bids for each component part. For each EMS, secure protocols partition the BOM into two sets: those for which the OEM has the lower price and those for which the EMS has the lower price. This process yields the minimum total component cost to the OEM if that EMS is eventually chosen to manufacture/assemble the product. A secure mechanism is then used to identify which EMS, in combination with the OEM, will provide the lowest total cost to the OEM.

Our model involves a single OEM, N EMSs, and a single product with $M-1$ unique components. (Multiple copies of a given component are modeled as a single component "kit.") The assembly of these $M-1$ components into a finished product is represented as component M. And, since the OEM does want to outsource assembly, the OEM's cost for component M is set to an arbitrarily large number.

Each of the $N+1$ participants (the OEM and N EMSs) has private information about its cost for each of the M components. If the EMSs were willing to share this information with the OEM, then, for each EMS, the OEM would "cherry pick" between the EMS's costs and its own for each component in order to minimize its total component cost, and, by comparing the corresponding total component costs among all the EMSs, (trivially) determine the EMS who should manufacture the product. However, none of the participants wants to share its private prices with any of the other participants.

The goal of the OEM is to select an EMS to manufacture its product and to partition the sourcing of its components in order to minimize total component cost; and to do so without disclosing its private cost information. The goal of each EMS is to be selected by the OEM, provided that this yields a non-negative margin; and, if selected, to maximize that margin (i.e., EMS price minus its cost) on the set of components it provides; and to do so without disclosing its private cost information. Although it is expected that all $N+1$ participants will attempt to learn the private information of the others, based on its own private information and whatever information is disclosed during the business process, we assume that none of the participants gains any utility from sabotaging the process itself.

To summarize, the contributions of this paper are as follows:

- We develop and demonstrate a business process for price masking that preserves the private information of all the participants (see section 5).
- We address incentive issues in price masking. In particular, we (1) demonstrate that simple adaptations of a secure Vickrey auction are not incentive compatible; and (2) construct an incentive-compatible auction (see section 4).
- We demonstrate that the SPM process is not invertible; i.e., that it is impossible for any participant to determine the private prices (i.e., costs) of any other participant using its own private information and the outputs by the SPM process. This is an important practical consideration; i.e., why use SMC techniques if participants can infer their partners' private inputs from the output of the SPM process? See section 7.
- We provide "information-leakage" analysis of the SPM process. Even though a business process is

not invertible, it may, nonetheless, give other participants some probabilistic information about its partners' private inputs. We provide a measure of ''information leakage,'' and establish conditions under which the SPM process is practically leak-proof (see section 7).

- We recommend an architecture for implementing the SPM process and describe the results of a proof-of-concept implementation (see sections 5 and 6).

The rest of this paper is organized as follows. In the next section, we summarize related work in procurement, auctions, and SMC. Section 3 introduces our model. Section 4 describes mechanisms that induce the EMSs to submit bids equal to their costs: first, for a business scenario in which a single EMS will be chosen to manufacture all of the units to be produced; and second, for a business scenario in which the OEM wants to allocate production among two or more EMS providers. Section 5 provides an overview of how these mechanisms preserve the privacy of all the participants. Section 6 describes the results of a proof-of-concept implementation. Section 7 demonstrates that the SPM process is not invertible and examines the issue of information leakage. Section 8 summarizes our contributions and suggests additional work.

2. Related Work

Our work draws on four streams of literature: procurement mechanisms, auction theory, secure multiparty computation, and supply chain management.

2.1. Procurement Mechanisms

Given its strategic importance, there is a vast trade literature in industrial procurement (Cavinato and Kauffman 1999, Nelson et al. 2005), including outsourcing (Jenster 2005). Procurement decisions also play an important role in supply chain management practice (Simchi-Levi 2004) and research (Simchi-Levi et al. 2004).

Elmaghraby (2000) provides an overview of the operations research and economics literature on sourcing policies and contract competition. Within Elmagrahby's framework, our work involves a fixed contract, a single selection (i.e., time) period, and the selection of either a single source (section 4.1) or the selection of a fixed number of multiple sources (section 4.2) given a fixed partition of volume among the sources (e.g., that the lowest bidder will get *X*% of the volume, the second-lowest bidder, *Y*%, etc.). Given the EMS chosen by the mechanism we propose, the partition of components between the chosen EMS and the OEM can be viewed as an outsourcing/insourcing decision, although, in practice, both the EMS and the OEM outsource component production to the CS specified in the BOM. The selection criterion is the minimization of total parts cost. See Bichler and Steinberg (2007) and Rothkopf and Whinston (2007) for recent work on e-auctions for procurement operations.

The business process we propose contributes to procurement theory and practice in two significant ways: first, by preserving the privacy of every party's individual parts prices. This is significant in removing the issue of trust in supplier selection (see Kramer and Tyler 1996 for a review of the issue of trust in organizations). It also removes the threat posed by industrial spies and communications hackers. Second, the incentive-compatible mechanisms we propose motivate the EMSs to bid their own private purchase cost (i.e., not add a margin to their cost).

2.2. Auction Theory

Adverse-selection models, due to information asymmetry between principal and agents, have been well studied in economics (see Fudenberg and Tirole 2000). The classic references on adverse-selection models and signaling include Akerlof (1970), Rothschild and Stiglitz (1976), and Spence (1974). Our model also draws on classical auction theory as described in the seminal papers by Vickrey (1961), Myerson (1981), Riley and Samuelson (1981), and Milgrom and Weber (1982). See Klemperer (1999) for a more recent review on the theory of auctions. Elmaghraby (2007) presents a survey of current industry practices in designing and running auctions as part of e-sourcing events.

Most of the above papers focus on analyzing auctions for either a single indivisible unit or multiple units of a single object. Our model can be viewed as an auction of heterogeneous objects (components of a finished product). Several papers have analyzed multi-object auctions. Palfrey (1983) analyzes the sellers' preferences for bundling heterogeneous objects. Other papers that analyze multi-object auctions include Armstrong (1996, 2000) and Avery and Hendershott (2000). Our model has the feature that the BOM can be split between the winning bidder and the OEM to reduce total cost. There is also a large stream of literature on combinatorial auctions (see de Vries and Vohra 2003 for a review). Our model is not a combinatorial auction since the total package bid is the sum of prices of individual parts in the package.

The Vickrey–Clarke–Groves (VCG) mechanism has been extensively studied in the literature due to its properties such as efficiency. Rothkopf (2007) pointed out several limitations of the VCG mechanism, which include information revelation of private prices and concern about cheating by the bid-taker. Our SPM process solves both of these problems of the VCG process by constructing a verifiable process that preserves the privacy of individual part prices.

2.3. Introduction to SMC

Cryptographic techniques have revolutionized the way consumers and businesses interact, particularly on the Internet. Examples include the encryption of credit card information and the use of digital certificates and signatures. These techniques are routinely used in a manner almost entirely transparent to the user, and are already an integral part of the user's daily computing experience.

The sub-area of cryptography that is most relevant to our work is SMC, a form of cooperative-distributed computing. SMC *protocols* are step-by-step procedures, which, if faithfully followed by all the participants, can be used to evaluate any function computable by a single party who has all of the inputs (i.e., apply any computable decision rule) while preserving the privacy of all of the participants' inputs. In other words, at the end of the protocol, the only value/s the participants learn is the output of the function (i.e., the decision(s)). The security that SMC seeks to achieve is the same *as if* the protocol had taken place through a trusted third party, to whom all participants submit their private inputs, and trust not to divulge their private information to any other participant. Note that SMC seeks to achieve this *without* the use of a trusted third party. Thus an SMC protocol is *secure* if, at the end of the protocol, the protocol reveals to the participants only the value of the function that was computed (i.e., the decision made), but not its inputs. This notion of security is illustrated by the following simple example.

A group of N professors is sitting around a table, commiserating about being underpaid by their universities. They share an interest in computing their average salary. However, each participant wants to keep her/his own salary private. The following SMC protocol is well known: Professor 1 selects a very large random number, denoted Z_0. She adds her own salary to this number, writes the sum, Z_1, on a slip of paper, and passes this slip of paper to Professor 2. Note that Professor 2 learns nothing about Professor 1's salary from Z_1: to him Z_1 looks like a random number. Professor 2 adds his salary to Z_1, and passes the sum, Z_2, written on a different slip of paper, to Professor 3, etc. Eventually, Professor 1 receives Z_N from Professor N. Note that Z_N is equal to Professor 1's chosen random number, Z_0, plus the sum of the N professors' salaries. Professor 1 determines the average salary by subtracting Z_0 from Z_N and dividing the remainder by N. Typically, SMC protocols are considerably more complex, and are performed by a network of computers, not by paper and pencil.

Note that the effectiveness of the average salary protocol requires that all participants follow the protocol faithfully. This is known as the "honest-but-curious" framework in the SMC literature. (Alternatively, one or more participants might want to stop or disrupt the process. For example, one of the professors might interfere with the passing of the paper slips. Such protocol "attacks" are studied in the SMC literature, but are outside the boundaries of this work.) In addition, the accuracy of the protocol's computed result requires that all participants provide their *true* salaries as inputs.

2.3.1. Our SPM Model. Our SPM business process assumes that the goal of the OEM is to select an EMS to manufacture its product and to partition the sourcing of its components in order to minimize its total cost; and to do so without disclosing its private cost information. Correspondingly, the goal of each EMS is to be selected by the EMS, provided that this yields a non-negative margin; and if selected, to maximize that margin (i.e., EMS price minus its cost) on the set of components it provides; and to do so without disclosing its private cost information. In particular, we assume that none of the participants gains any utility from stopping or disrupting the prescribed protocols. Hence, none of the participants will attack the protocols. We assure that the participants will be honest through the modified Vickrey auction described in section 4. Hence, the participants will not only faithfully follow the prescribed steps of the protocol, but will also provide their *true* costs to the protocols.

Equally important, in order to be of practical value, our SPM business process is "non-invertible" and practically "leak-proof." In order to describe these concepts, we return to the average salary protocol, but with only $N = 2$ professors. As described, neither professor will attack the protocol and, given an appropriate incentive mechanism, both professors will honestly report their salaries to the protocol. And, as demonstrated, this process is cryptographically secure to the extent that neither learns anything more than if a trusted third party had been used.

Yet the process is "invertible"; i.e., either professor could (easily) determine her/his colleague's private input (i.e., her/his salary) based on the knowledge of her/his own private input (i.e., salary) and the securely computed output of the protocol (i.e., the average salary of both). Adding a third professor makes the process "non-invertible." In other words, any attempt to determine the private inputs of any of the other two professors from the securely computed average salary and information about one's own salary will yield an infinity of possibilities.

However, even with $N = 3$ participants, the protocol has "leaked" some information. That is, any of the three professors can determine the average salary of her/his two colleagues and the range in which each of their salaries must fall. Adding more professors to the process leaks less information. Indeed, as the number of professors increases, the protocol becomes "leak-proof." In other words, every participant's posterior distribution of any of her/his colleagues' salaries is practically indistinguishable from the actual distribution. Hence, for a sufficiently large number of professors and/or variance in the professors' salaries, the average salary protocol is practically leak-proof.

In section 7 we will demonstrate that the SPM business process is non-invertible for all of the participants. Specifically, that neither the OEM nor any EMS can determine the private costs of its partners based on its own inputs and the outputs of the protocols. Furthermore, we will describe conditions under which SPM is practically leak-proof.

2.3.2. Selected SMC Literature. The history of the SMC problem is extensive since it was introduced by Yao (1982) and extended by Yao (1986), Goldreich et al. (1987), and many others. Broadly speaking, it has been established that there exists a secure protocol to evaluate any well-defined function, no matter how complex. Recent results (Damgård and Ishai 2005) have shown promise in evaluating many functions in an efficient manner. Furthermore, there is a substantial volume of work on the application of SMC techniques to auctions. For example, Franklin and Reiter (1996) describe SMC protocols that ensure that an auctioneer will be able to extract the winning bid without learning anything about the losers' bids until after the bidding period. Others (Brandt and Sandholm 2005, Decker et al. 2001, Elkind and Lipmaa 2004, Jakobsson and Juels 2000, Naor et al. 1999) have extended the research on secure auctions. A more detailed description of work related to secure auctions can be found in supporting information Appendix S1.

We utilize existing SMC techniques to implement SPM. We do not need a special-purpose protocol for computing our specific auction mechanism, because the general results in SMC are reasonably efficient for the calculations required. This efficiency is partly due to the manner in which we have defined our auction mechanism.

2.4. SMC Applications to Operations/Supply Chain Management

To date, there are very few applications of SMC techniques in operations or supply chain management. To the best of our knowledge, Atallah et al. (2003) were the first to apply SMC to an operations management problem. They develop secure protocols for allocating the fixed capacity of a supplier among N retailers. Their allocation protocols are both incentive compatible and privacy preserving with respect to the supplier's capacity and the retailers' demand drivers. Clifton et al. (2008) examined a problem faced by independent trucking companies that have separate pickup and delivery tasks. They describe a secure protocol that finds opportunities to swap loads without revealing any information except the loads to be swapped.

More recently, Deshpande et al. (2009) used SMC techniques to develop secure protocols for the collaborative planning, forecasting, and replenishment business process. In their model, N retailers and their supplier engage in secure protocols that result in: (1) customer-demand forecasts that use each of the retailers' and the supplier's privately observed demand signals; and (2) order/shipment quantities based on system-wide costs and inventory levels (and on the joint forecasts) that minimize supply chain expected cost/period. Our business scenario is distinctly different from those above. In addition, both the incentive mechanisms developed (transfer payment versus auction) and recommended secure multiparty techniques (custom protocols versus circuit simulation) are different.

3. Model

In section 3.1, we formally state our SPM model using the language and rhetoric of game theory. This lays the foundation for Theorem 1; i.e., that the mechanisms proposed in section 4 are incentive compatible. This formalism is not used once Theorem 1 has been established. Readers already familiar with this formalism may wish to skip to section 4. Readers not familiar with this formalism and some of its seemingly unnecessary notation (e.g., probability distributions representing the OEM's uncertainty about EMS prices for components) may also want to proceed directly to section 4.1.

3.1. Stylized Model of Outsourcing Scenario

Let $i = 0$ denote the OEM, while $i = 1, \ldots, N$ denote the EMSs. Given the game-theoretic analysis to follow, we will refer to participants as "players." Let $j = 1, \ldots, M$ denote the components. We use v_{ij} to denote player i's (true) cost for component j and b_{ij} to denote player i's bid for component j. Note that these values correspond to the costs (resp., bids) for the production of some given number of units of the OEM's product.

The OEM does not know the cost of each EMS i for each component j. We assume that the OEM's uncertainty about the EMS's cost can be captured by a continuous probability distribution with density $f_{ij}(\cdot)$. We also assume that this is common knowledge, i.e., other EMS's share the same beliefs. Let $\mathbf{T} = (\mathbf{t}_1, \ldots, \mathbf{t_N})$ denote the matrix of component prices, where $\mathbf{t_i}$ is the vector of component prices for EMS i. The joint density function for the matrix $\mathbf{T}$ is given by

$$f(\mathbf{T}) = \prod_{\forall i} \prod_{\forall j} f_{ij}(t_{ij}).$$

Also, let $\mathbf{T_{-k}}$ represent the cost vector of all EMSs other than EMS k. Then the joint density of costs of all EMSs other than k is given by

$$f_{-k}(\mathbf{T_{-k}}) = \prod_{\forall i;\, k \neq i} \prod_{\forall j} f_{ij}(t_{ij}).$$

EMS *i*'s procurement cost for component *j*, conditional on the matrix of component prices **T**, is

$$v_{ij}(\mathbf{T}) = t_{ij}.$$

Thus EMS *i*'s total procurement cost, conditional on the matrix of component prices **T**, is given by

$$v_i(\mathbf{T}) = \sum_{\forall j} t_{ij}.$$

Our goal is to design *feasible* mechanisms that have the following properties: *incentive compatibility, security, non-invertibility,* and *non-negative profits*. The OEM desires that the mechanism be *incentive compatible*; i.e., that it induces the EMSs to bid truthfully. Also, the EMSs will participate in the mechanism only if they are guaranteed non-negative profits (i.e., *participation constraints*). The OEM and the EMS desire the mechanism to be *secure*; i.e., the process of computing the winner, and the transfer payments should reveal no more information than what would be learned if the process had been carried out by a trusted third party. Finally, the OEM and all the EMSs desire *non-invertibility,* i.e., the output of the mechanism should not reveal the individual component prices of one party to any other party.

We restrict our attention to *direct-revelation mechanisms*. In these mechanisms, the EMSs simultaneously and confidentially bid their component prices. The OEM then determines who wins the contract based on the bids received, and how much each EMS will get paid, as functions of the announced matrix **T**. Thus, a direct-revelation mechanism is of the form $(\mathbf{Y}, \mathbf{x}, \mathbf{P})$, where y_{ij} is the fraction of component *j* volume allocated to player *i*, x_i is the fraction of finished product volume allocated to EMS *i*, and P_i is the payment received by EMS *i*. Note that the mechanism $(\mathbf{Y}, \mathbf{x}, \mathbf{P})$ is a function of **t**. The following constraints define the set of feasible mechanisms.

$$\sum_{i=0}^{N} y_{ij} = 1, \quad \forall j, \tag{1}$$

$$y_{ij} \geq 0 \quad \forall i,\ \forall j, \tag{2}$$

$$y_{ij} \leq x_i, \quad i = 1, \ldots, N, \quad j = 1, \ldots, M, \tag{3}$$

$$\sum_{i=1}^{N} x_i = 1. \tag{4}$$

Constraint (1) requires that the total volume for each component must be allocated among the players; Constraint (2), that the allocation for each player be non-negative. Constraint (3) makes sure that the fraction of a component volume allocated to a player does not exceed its fraction of the finished product volume. Finally, (4) requires that the sum of the finished product volumes allocated across all players should add up to one.

Initially, we consider the problem where there is only one winner among the EMSs, i.e., only one EMS will be chosen. In this case

$$x_i = 0, 1 \quad i = 1, \ldots, N. \tag{5}$$

Thus, given EMS *i*'s cost vector $\mathbf{t_i}$, its expected profit from the auction mechanism is

$$\Pi_i(\mathbf{Y}, \mathbf{x}, \mathbf{P}, \mathbf{t_i}) = E_{-i}\left\{ P_i(\mathbf{T}) - \sum_j v_{ij}(\mathbf{T}) y_{ij}(\mathbf{T}) \right\}. \tag{6}$$

Here the expectation is over the (unknown) costs of all players other than EMS *i*.

3.1.1. Participation.

The participation constraint, (7), guarantees that each EMS *i* will make a non-negative expected profit

$$\Pi_i(\mathbf{Y}, \mathbf{x}, \mathbf{P}, \mathbf{t_i}) \geq 0 \quad i = 1, \ldots, N. \tag{7}$$

3.1.2. Incentive compatibility.

The revelation mechanism can be implemented only if no EMS can gain by lying about its component prices. Hence, truth-telling must form a Bayesian Nash equilibrium strategy. In what follows, let $\mathbf{b_i}$ be a vector of component bids of EMS *i*. Incentive compatibility is then captured by Constraint (8):

$$\begin{aligned} \Pi_i(\mathbf{Y}, \mathbf{x}, \mathbf{P}, \mathbf{t_i}) \geq E_{-i}\Bigg\{ & P_i(\mathbf{T_{-i}}, \mathbf{b_i}). \\ & - \sum_j v_{ij}(\mathbf{T_{-i}}, \mathbf{b_i}) y_{ij}(\mathbf{T_{-i}}, \mathbf{t_i}) \Bigg\} \forall \mathbf{b_i}. \end{aligned} \tag{8}$$

3.1.3. Non-Invertibility.

The non-invertibility constraint states that no participant *i* should be able to infer the true component prices, $t_{i'j}$, $i' \neq i$, of any other participant from the output of the mechanism $(\mathbf{Y}, \mathbf{x}, \mathbf{P})$. This can be written mathematically as

$$Prob_i(\mathbf{t_{i'j}} = t_{i'j} | (\mathbf{Y}, \mathbf{x}, \mathbf{P})) < 1 \quad \forall i' \neq i, \quad i = 0, \ldots, N. \tag{9}$$

3.1.4. Security.

The mechanism must be implemented with a protocol that does not reveal the EMS or OEM component prices other than what can be deduced from the outcome of the mechanism. More specifically, the protocol should not use a trusted third party but should be equivalent from a security standpoint to a protocol where the OEM and the EMSs reveal their prices to a trusted third party who then reveals the outcome to the participants. For a formal definition of this property, see Goldreich (2004).

Equations (1)–(9) define the set of *feasible* mechanisms.

3.2. Properties of Potential and Existing Mechanisms

Before describing the mechanisms we propose to solve (1)–(9), it is appropriate to consider the "natural" Vickrey auction, an auction when there is a single component (i.e., $M=1$). However, generalizing the single-item case to multiple items is non-trivial. To demonstrate this, we will describe two generalizations that fail.

One generalization is to compute the total bid of each participant (i.e., the sum of all part bids) and use a Vickrey auction on these totals. Hence, if the OEM has the lowest cost, it procures everything; otherwise, the lowest bidder procures everything and is paid the second total lowest bid. It is easy to show that this mechanism is incentive compatible. However, it forces the OEM to pay more than necessary. To demonstrate this, consider the example bids in Table 1, with $N=2$ and $M=2$.

The total bids of EMS-1, EMS-2, and the OEM are, respectively, US$18, US$19, and US$24, and thus, EMS-1 is the winner. EMS-1 will procure both parts and be paid US$19. It is clear that the OEM would be better off obtaining Part 1 itself and having EMS-1 obtain Part 2.

Another generalization is to use the total bid mechanism to find the winner, but then use a "cherry-picking" algorithm to determine who—the winning EMS or the OEM—procures each part. The EMS is still paid the second-lowest bid, but the OEM reduces this payment for items that the OEM procures by the amount that the EMS bid for these items. Returning to our example, EMS-1 wins the auction and its base payment is US$19. However, it is only paid US$19 – US$10 = US$9 and will procure only Part 2. Thus the OEM's cost is US$4+US$9 = US$13. This mechanism appears to overcome the weakness of the first mechanism. Unfortunately, this mechanism is not incentive compatible (at least not in terms of a dominant strategy). Consider the example when EMS-2 lies about its bid on Part 1, claiming that it is US$5 instead of US$9. In this case, EMS-2 wins the auction (its bid is US$15), and its base payment is US$18. It will procure Part 2 only and will be paid (US$18 – US$5) = US$13 to obtain Part 2. Thus, its profit (from information rent) is US$13 – US$10 = US$3. On the other hand, if EMS-2 is honest, then it obtains US$0 profit. Clearly, EMS-2 is better off by lying.

Table 1 Example Bids for a Generalization of the Vickrey Auction to Multiple Items

Party	Part 1 (US$)	Part 2 (US$)	Total (US$)
EMS-1	10	8	18
EMS-2	9	10	19
OEM	4	20	24

EMS, electronic manufacturing service; OEM, original equipment manufacturer.

Hence, the challenge is to allow some form of "cherry-picking" by the OEM, while retaining EMS incentive compatibility.

4. Proposed Mechanisms and Incentive Compatibility

In this section, we describe our proposed mechanisms, which satisfy the desired properties stated in section 3, for two cases: the case of a single winner, and the case where there are multiple winners.

4.1. The Case of a Single Winner

Auction Mechanism: For each EMS i compute:

$$c_i = \sum_{j=1}^{M} \min\{b_{ij}, b_{0j}\}.$$

Note that for EMS i, c_i is the OEM's total parts cost for the finished product assuming that the OEM procures part j if it has the lower price. Let

$$i^* = \arg\min_{i \in N} c_i,$$

$$i^{*(2)} = \arg\min_{i \in N, i \neq i^*} c_i.$$

Hence, i^* is the index of the EMS with the lowest c_i, and $i^{*(2)}$ is the index of the EMS with the second lowest c_i. Also

$$y_{ij} = \begin{cases} 1 & \text{if } b_{ij} \leq b_{0j}, \\ 0 & \text{otherwise.} \end{cases}$$

Player i^* wins the auction, procures items $J^* = \{j | y_{i^*j} = 1\}$, and gets paid

$$P_{i^*} = c_{i^{*(2)}} - \sum_{j=1}^{m} (1 - y_{i^*j}) b_{0j}.$$

A property of our proposed mechanism is that it results in a (weakly) dominant strategy equilibrium for each EMS; i.e., each EMS's bidding strategy is independent of the bidding strategy used by other EMSs or the OEM. As a result, the knowledge of the probability distribution of part prices, $f_{ij}(\cdot)$, is not needed by any EMS in deciding its bidding strategy. Our proposed mechanism can be considered as a modified, but not identical, version of the VCG mechanism. A VCG mechanism results in an *efficient* allocation, when all participants bid their true valuations. However, since our mechanism is not OEM incentive compatible, it may not result in an efficient allocation. Note that while the proposed mechanism may not result in an efficient allocation, it has the potential to decrease the OEM's expected cost if the OEM were to bid properly.

An astute reader will also note that although this mechanism is incentive compatible among the EMS's that participate, a given EMS may refuse to participate, preferring, instead, a Vickrey auction for the entire BOM. Such a refusal would leave the OEM worse off if it increased the second price among the remaining participants. And, of course, by refusing to participate, an EMS gives up the opportunity to become the chosen partner in the proposed "cherry-picking" mechanism, and, hence, earn the difference between its cherry-picked potentially lowest bid and that of the second-lowest bidder. Finally, in comparing the proposed "cherry-picking" and Vickrey mechanisms: (1) the OEM will be worse off with the Vickrey mechanism if it has the lowest price for any of the components in the BOM (as noted above); and (2) a given EMS may be better off with the Vickrey mechanism (e.g., if there is some subset of components for which its prices are significantly lower than all other participants) or worse off (e.g., if its prices for the set of OEM-supplied components are significantly lower than those of the other EMSs). In summary, without knowledge of every participant's individual component prices, it is impossible to determine which participant, if any, would be better or worse off with a Vickrey mechanism for the entire BOM.

4.2. The Case of Multiple Winners

In some business situations, it is desirable to avoid having a single EMS produce all T units, so as to (i) avoid EMS acquiring large market power over time, and (ii) keep other bidders in business for other strategic reasons (such as diversification—maintaining a plurality of viable future suppliers, or perhaps to honor special long-term strategic partnership agreements that require some minimal order flow).

We, therefore, consider the case of ℓ non-zero values among the x_i's, where $1 \leq \ell < N$ and is set by the OEM ahead of time. Thus, there are ℓ "winners" who each obtain a non-zero fraction of the finished product's volume T. The ℓ winners are, of course, the ℓ lowest bidders, so the main design issue is how much volume each winner should get, and how much it gets paid, so as to solve constraints (1)–(9) for ℓ non-zero x_i's.

Let $r(k)$ to denote the EMS having the kth smallest c_i value (we assume there are no ties, for convenience and to avoid unnecessarily cluttering the exposition). Hence, $c_{r(1)} < c_{r(2)} < \ldots c_{r(N)}$. The auction specification is then as follows:

- The OEM determines fractions of the total volume $\lambda_2, \ldots, \lambda_{\ell+1}$ that add up to 1. For $k = 2, \ldots, \ell+1$, each λ_k specifies the fraction of the total volume T for which the OEM pays according to the kth lowest price.
- The λ_k fraction is split evenly among the $k-1$ EMSs who have the $k-1$ lowest c_i values; i.e., among the EMSs $r(1), r(2), \ldots, r(k-1)$. Each EMS $r(t)$ $(1 \leq t \leq k-1)$ then gets its $\lambda_k/(k-1)$ fraction and is paid for that fraction an amount equal to:

$$\left(c_{r(k)} - \sum_{j=1}^{M} (1 - y_{r(t),j}) b_{0,j}) (\lambda_k/(k-1) \right).$$

Table 2 A Sample Specification of the Proposed Mechanism for Auctions with Multiple Winners

	EMS $r(1)$	EMS $r(2)$	EMS $r(3)$	All EMSs
Fraction at second lowest price	0.55	—	—	$\lambda_2 = 0.55$
Fraction at third lowest price	0.15	0.15	—	$\lambda_3 = 0.30$
Fraction at fourth lowest price	0.05	0.05	0.05	$\lambda_4 = 0.15$
Total x_i	$x_{r(1)} = 0.75$	$x_{r(2)} = 0.20$	$x_{r(3)} = 0.05$	1

EMS, electronic manufacturing service.

The above implies that for $t = 1, \ldots, \ell$, the EMS $r(t)$ obtains a fraction of the total volume equal to

$$x_{r(t)} = \sum_{k=t}^{\ell} \lambda_{k+1}/k.$$

As an alternative (and equivalent) way of specifying the auction, the OEM might determine a sequence of values $x_{r(1)}, \ldots, x_{r(\ell)}$ (subject to $x_{r(t)} \geq x_{r(t+1)}$ and $\sum_{t=1}^{\ell} x_{r(t)} = 1$), with $x_{r(t)}$ being the fraction of the total volume T that the EMS with the tth lowest bid procures. Each λ_k can then be computed from x_i's as $\lambda_k = (x_{r(t-1)} - x_{r(t)})(k-1)$.

To illustrate this, consider $\ell = 3$, $\lambda_2 = 0.55$, $\lambda_3 = 0.30$, and $\lambda_4 = 0.15$. Then Table 2 shows how the total volume is distributed among EMSs $r(1)$, $r(2)$, and $r(3)$.

EMS $r(t)$ gets paid an amount equal to

$$\sum_{k=t}^{\ell} \left(c_{r(k+1)} - \sum_{j=1}^{M} (1 - y_{r(t),j}) b_{0,j} \right) (\lambda_{k+1}/k).$$

The profit of EMS $r(t)$ is

$$\hat{P}_{r(t)} = \sum_{k=t}^{\ell} \left(c_{r(k+1)} - \sum_{j=1}^{M} ((1 - y_{r(t),j}) b_{0,j}) + y_{r(t),j} v_{r(t),j} \right) (\lambda_{k+1}/k).$$

It is easily observed that this profit is non-negative (as every term of the outer summation is non-negative).

Theorem 1. *The above auction mechanism is EMS incentive compatible for the multiple winner case, i.e., for each EMS it is a (weakly) dominant strategy to bid $b_{i,j} = v_{i,j}$ $\forall j$.*

Proof. Provided in supporting information Appendix S1.

COROLLARY 1. *The auction mechanism described for the single winner case is EMS incentive compatible, i.e., for each EMS it is a (weakly) dominant strategy to bid* $b_{ij} = v_{ij}\ \forall j$.

PROOF. This is a special case of Theorem 1. □

5. Secure Implementation of Price Masking

For a given OEM, a given product, and its M-component BOM, the goal of SPM is twofold. First, for any given EMS, to determine a partition of the components in the BOM that will yield the lowest total component cost for the OEM. Second, for a set of competing EMSs, given each EMS-specific partition of components, to determine which EMS, in partnership with the OEM, will provide the lowest total component cost to the OEM.

Figure 2 demonstrates the overall SPM process. The public input for the protocol is the OEM's BOM, which is known to all participants. For each component in the BOM, the OEM and each EMS enter their private bids (i.e., $b_{i,j}$). At the end of the protocol, the winning EMS is informed that it has been chosen and the others will be informed that they have not been chosen. The winning EMS will be given a list of the components, J^*, that it is to provide and the single, total price, P_{i^*}, that the OEM will pay the EMS for all of the parts in J^*. The OEM will be informed which EMS has been chosen along with P_{i^*} and J^*. The non-winning EMS(s) will not learn any of these values.

In this paper, we recommend[1] building the SPM process via a semi-trusted third party (STTP); a similar idea has been used in the secure auction literature (Naor et al. 1999). This third party is only semi-trusted in that it *does not learn* the bidders' information nor the results of the auction. The only trust assumption for this third party is that it will not collude with other participants. It is reasonable to believe that such STTP can be found, because there are many parties that receive no benefit from this SPM process and, hence, are not easily coerced into collusion. To provide a participation incentive to such a third party, the OEM (or EMSs) may have to pay a fee.

Figure 2 The Secure Price-Masking Process

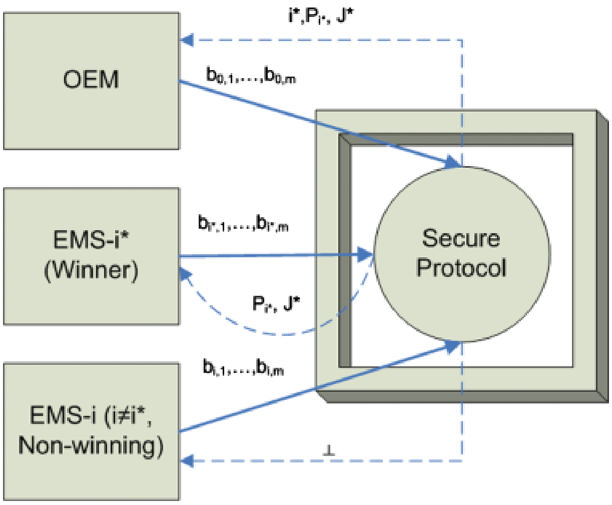

Figure 3 First Step of Protocol for Secure Price-Masking

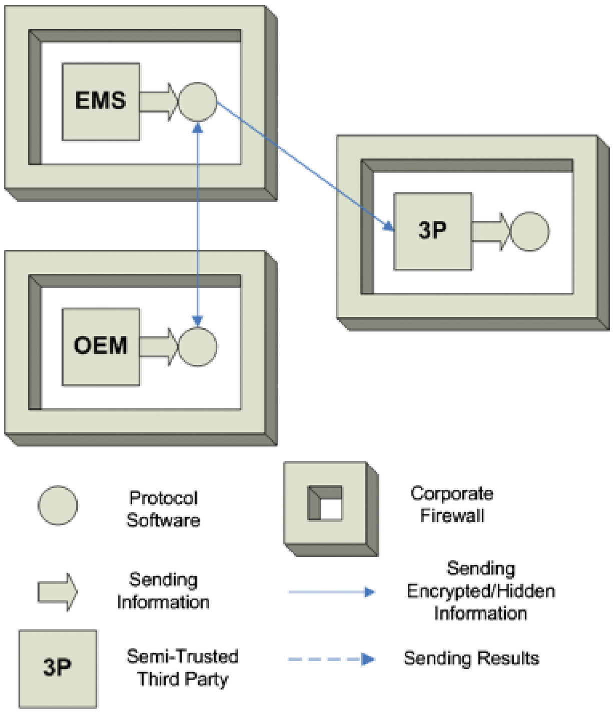

We now describe how the protocol works in more detail. Our protocol operates in two phases. In the first phase, each EMS and the OEM engage in a two-party protocol to compute: (i) the bid value for the EMS, (ii) the items that the EMS will procure, and (iii) the cost that is deducted from the bid price for the items procured by the OEM. These values are then hidden from the participants by "splitting them" between (see below) the OEM and the untrusted third party. This phase is depicted in Figure 3. In the second phase, the OEM and the third party engage in a protocol to compute the winning EMS, the procurement set, and the payment to the EMS and the OEM. This is depicted in Figure 4.

To help clarify this protocol further, we describe the process in more detail:

Phase 1: Each EMS (EMS-i) and the OEM engage in a two-party protocol to calculate c_i along with $y_{i,1}, \ldots, y_{i,m}$. The result is split between EMS-i and the OEM so that neither party knows the values. EMS-i then sends it shares of the results from the previous

Figure 4 Second Step of Protocol for Secure Price-Masking

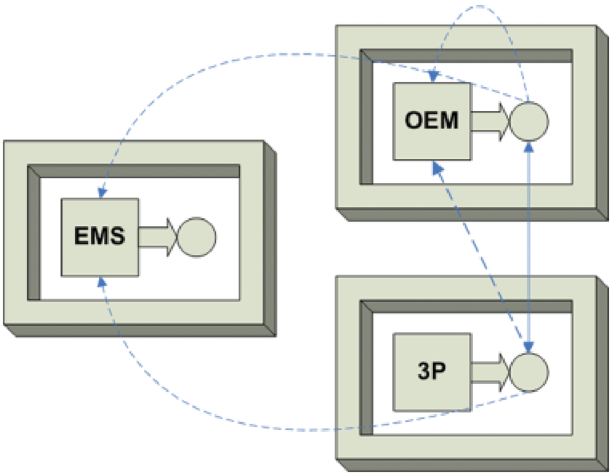

step to the third party (so that it is now split between the third party and the OEM).

Phase 2: The third party and the OEM engage in a protocol to calculate i^*, J^*, and P_{i^*}. Note that like the previous step the results are split. These values are then opened to the winning EMS and the OEM.

Suppose that we have a two-part BOM and that the OEM's prices are 8 and 8, that EMS-1 has prices 6 and 10, and that EMS-2 has prices 12 and 5. After Phase 1, EMS-1 and the OEM have the following values split between them, $c_1 = 14$, $y_{11} = 1$, and $y_{12} = 0$. Assuming that all values are being stored additively split modulo 16, one way of doing this would be by having EMS-1's values be $c'_1 = 6$, $y'_{11} = 12$, and $y'_{12} = 14$ and the OEM's values be $c''_1 = 8$, $y''_{11} = 5$, and $y''_{12} = 2$ (note that we are just using the ''and'' notation to denote the respective shares of the two participants). EMS-2's values would be split in a similar fashion. After phase 1, EMS-1 and EMS-2 would send their shares to the STTP. Now the STTP and the OEM engage in a protocol to find i^*, J^*, and P_{i^*} in a split fashion just as EMS-1 and the OEM split the values above. These values can then be revealed to the OEM (i.e., the STTP reveals its shares to the OEM) who then can reveal the appropriate results to the appropriate EMSs.

Theorem 2. *Assuming that the STTP and the OEM do not collude, the protocol described above is secure.*

Proof. In the first step of the protocol, the EMSs split their bids between the OEM and the untrusted third party. There are well-known techniques for splitting a value so that no individual party has any information about the value. The second part of the protocol uses a scrambled circuit evaluation protocol (such as Yao 1986) to evaluate the mechanism, and such techniques have been proven to be secure (Lindell and Pinkas 2004). □

6. Results from a Proof-of-Concept Implementation

In this section we provide experimental results that show how our proposed business process performs. Our results are based on Fairplay (Malkhi et al. 2004): a tool that takes as input a specification of the desired functionality and creates the corresponding computer program. Such a circuit can be built for evaluating any computable function, but the approach becomes impractical for complex functions or for dealing with large volumes of data. There are alternatives to this approach, but the message that we would like to convey with this proof-of-concept implementation is that the simplicity of the computations used in our protocol allows us to invoke general mechanisms from SMC literature with reasonable performance.

Our implementation considers the single winner case (section 4.1). The programs that specify the computation steps and that were input into Fairplay are provided in supporting information Appendix S1. The first program corresponds to the first phase of the protocol executed by the OEM and an EMS i (for $i = 1, \ldots, N$), which computes y_{ij}'s and c_i and stores the result in a hidden form, split between the OEM and the EMS. The second program is executed by the OEM and a third party and finishes the computation by determining the winner and the second lowest price. In this second phase, the third party is assumed to have the (hidden) shares of c_i's that it obtains from the EMSs.

The goal of our experiments is to measure the complexity and overhead associated with this solution. The timing results we report exclude delays due to communication. Since communication capabilities of parties can vary significantly, we assume that the setup will be such that the parties will be able to be connected to each other through fast links with appropriate bandwidth.

Figures 5 and 6 show the experimental results of phases 1 and 2, respectively. Each figure contains two plots.

The plot on the left-hand side shows the size of the Boolean circuit for secure function evaluation and the plot on the right-hand side shows the runtime of the protocol. In case of phase 1 (Figure 5), the plots correspond to the protocol between a single EMS and the OEM with a variable number of components; in case of phase 2 (Figure 6), they correspond to the protocol between the OEM and a third party with a variable number of EMSs. Note that the first phase will be executed between the OEM and each EMS and the second phase is independent of the number of components. As it can be seen from the plots, the

Figure 5 Experimental Results for Phase 1 of the Protocol with a Varying Number of Components

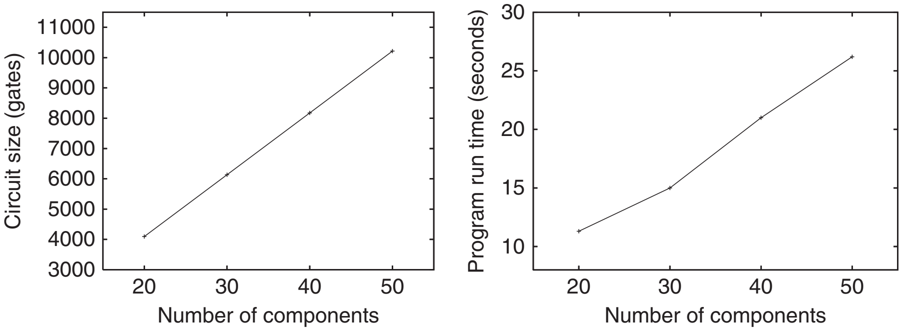

runtime and the circuit size in the first phase increase linearly with the number of components; the computational load of the OEM also increases linearly with the number of EMSs (since the OEM executes this protocol with each EMS). In the second phase, the runtime and the circuit size also increase linearly with the number of EMSs, but are independent of the number of components.

The protocols were executed on a dual Mac G5 1.8 Hz computer with 2 GB of RAM. This should be considered to be rather modest computing resources for the setup we target. Nevertheless, the protocol performed well. Unfortunately, our computing resources and the use of this specific tool for generating circuits did not allow us to conduct experiments with a larger number of components (due to massive memory resources that Fairplay requires). However, if these mechanisms are implemented and used ''in practice,'' we expect that a secure circuit would be custom built and optimized getting around Fairplay's explosive memory requirements. For products consisting of very large numbers of component parts, such circuits can be built in stages and then combined to achieve desired functionality. This, however, is specific to each problem instance and is beyond the scope of the paper.

Example: To better illustrate how the above performance results can be used, consider a scenario with 5 EMSs and a BOM with 100 parts. Before protocol execution, the OEM and EMSs agree on the third party, such that the OEM and this third party are unlikely to collude. The OEM then compiles the circuit for phase 1 and executes it with each of the EMSs. After the execution, each EMS sends its (hidden) output to the third party. The third party then executes phase 2 with the OEM and the OEM announces the winner (and obtains the list of parts the winning EMS is to procure).

To estimate the runtime of the protocols, assume that each party has a machine twice as fast as the desktop computer in our experiments. Then phase 1 is estimated to take 26 seconds on 100 parts (with each EMS). Runtime of phase 2, however, is not affected by the number of parts and will be under 2 seconds. With the continuing trend of doubling computing power every 2 years, it will be possible to run this solution much faster in just a few years.

7. Non-Invertibility and Information Leakage

In this section we demonstrate that the SPM business process is not invertible; i.e., that it is impossible for any participant to *determine* the private prices (i.e., costs) of any other participant given its own private information and the outputs of the SPM process. We also address the issue of information ''leakage''; i.e., to what extent does information revealed by the SPM

Figure 6 Experimental Results for Phase 2 of the Protocol with a Varying Number of Electronic Manufacturing Services

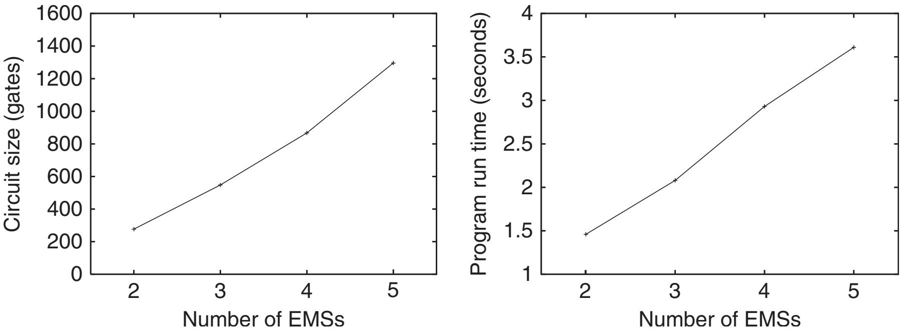

process permit any participant to *more accurately estimate* the private inputs of any other participant after the SPM process than they could before. These questions are managerially significant, since, regardless of how secure the SPM process may be from a computer science viewpoint, the SPM process is only of value if it is non-invertible and leak-proof.

The invertibility and information leakage discussion can be carried out from two points of view: What is learned by the OEM, or what is learned by an EMS. Since what the OEM learns about an EMS j is a superset of what any other EMS learns about EMS j, we carry out the analysis from the point of view of the OEM. Moreover, since the OEM can learn more about EMSs i^* and $i^{*(2)}$ (the EMSs with the lowest and second-lowest bids) than about EMSs other than i^* and $i^{*(2)}$, we focus on what the OEM learns about EMSs i^* and $i^{*(2)}$.

7.1. Non-Invertibility

As described above, the OEM is informed by the protocol of (i) which EMS has been chosen; (ii) the list J^* of the components that the chosen EMS is to provide; and (iii) the total price P_{i^*} that the OEM will pay the EMS for the components in J^*.

We start with (i): The fact that the EMS i^* was chosen unavoidably reveals that i^*'s bid c_{i^*} is smaller than any of the other c_j (which is also unknown to the OEM). Note that c_{i^*} itself is not revealed to the OEM, since the P_{i^*} depends on $c_{i^{*(2)}}$ rather than on c_{i^*}.

We now turn to (ii): From the list J^* the OEM learns that, for every $j \in J^*$ (resp., $j \notin J^*$), its own b_{0j} is an upper (lower) bound on the b_{i^*j} of EMS i^*. This does *not* reveal to the OEM the magnitude of the difference between the b_{i^*j} and the OEM's own b_{0j}.

Finally, (iii) $P_{i^*} = c_{i^{*(2)}} - \sum_{j \notin J^*} b_{0j}$. Since the OEM knows $\sum_{j \notin J^*} b_{0j}$, the OEM learns $c_{i^{*(2)}}$. If there are only 2 EMSs, then the OEM further learns which EMS this $c_{i^{*(2)}}$ comes from, so we need to examine carefully what the OEM learns about $i^{*(2)}$'s bids. As noted, the OEM learns

$$c_{i^{*(2)}} = \sum_{j=1}^{m} \min\{b_{i^{*(2)}j}, b_{0j}\}.$$

If we denote by J the set $\{j : b_{i^{*(2)}j} < b_{0j}\}$ and by $\bar{J}$ the complement of J $(=\{1,\ldots m\} - J)$, then $c_{i^{*(2)}}$ can be written as

$$c_{i^{*(2)}} = \sum_{j \in J} b_{i^{*(2)}j} + \sum_{j \notin J} b_{0j}.$$

The OEM, however, knows neither J nor any of the $b_{i^{*(2)}j}$'s. All it knows is that there is some vector $(b_{i^{*(2)}1}, \ldots, b_{i^{*(2)}m})$ (none of whose components it knows) and some subset J of the indices $\{1,\ldots,m\}$ (a subset it does not know either) for which it knows the value of the summation $\sum_{j \in J} b_{i^{*(2)}j} + \sum_{j \notin J} b_{0j}$. The worst case for this is when $|J| = 1$, in which case the OEM would learn the bid of EMS $i^{*(2)}$ for the single ("assembly") component. This can happen even if m is large—the EMS $i^{*(2)}$ could be very competitive for the "assembly" component and have higher cost for all the other components. The OEM has no way of detecting when $|J| = 1$, and it is quite unlikely that an EMS will have higher cost for all the components listed in a BOM, but an EMS who is notoriously uncompetitive for everything in a BOM and super-competitive for assembly should know that it may disclose its cost of assembly to the OEM in case it happens to be the EMS $i^{*(2)}$ (if this is unacceptable, then that EMS should not participate).

If $|J| > 1$, then the above-mentioned invertibility by the OEM no longer holds: The OEM cannot invert even if it can guess which components are in J, because it has one equation in more than one unknown—it only knows the sum $b_{i^{*(2)}1} + b_{i^{*(2)}2} + \cdots + b_{i^{*(2)}|J|}$ (we assumed WLOG that J consists of components $1, 2, \ldots, |J|$).

Non-invertibility is, however, not enough: We need to quantify, even when we know there is non-invertibility, the information leaked to the OEM about a specific element in the sum (say, about $b_{i^{*(2)}1}$) by the OEM's knowledge of the sum $b_{i^{*(2)}1} + b_{i^{*(2)}2} + \cdots + b_{i^{*(2)}|J|}$. This quantification is done next.

7.2. Information Leakage

Our protocol's information leakage to the OEM is similar to that of the average salary example given earlier: The OEM learns the sum of private values, and the issue is quantifying how much is revealed about the individual private values from the knowledge of their sum. Hence, we need to quantify the information leaked to the OEM about a private value $X = b_{i^{*(2)}1}$, when the OEM learns the sum of this private value and another private value $Y = b_{i^{*(2)}2} + \cdots + b_{i^{*(2)}|J|}$.

In information theory, entropy is used to quantify information—in fact it is well known that any information measure that satisfies a basic set of axioms that one would expect an information measure to satisfy has to be mathematically expressed as the entropy (to within a multiplicative constant factor). We refer the reader to Cover and Thomas, 1991 for the basic definition of information based on entropy.

Let X and Y be independent random variables, and let $Z = X+Y$. The information about X that is leaked by revealing $X+Y$ is given by the mutual information between X and Z (Cover and Thomas 1991), expressed as

$$I(X;Z) = H(X) + H(Z) - H(X,Z),$$

where H denotes entropy. Recall (Cover and Thomas 1991) that the entropy function $H(X)$ measures the amount of uncertainty about X, and that

$$H(X) = -\sum_{i} P(X = i)\log P(X = i),$$

where P denotes probability (a similar expression holds for continuous variables, with an integral replacing the discrete summation). Also recall that $H(X, Z)$, the joint entropy of the two discrete random variables X and Z, is the entropy of their pairing (X, Z):

$$H(X, Z) = -\sum_{i,j} P(X = i, Z = j) \log P(X = i, Z = j).$$

The above $I(X, Z)$ equation has a simple and intuitive meaning: $I(X, Z)$ is the amount of information shared by X and Z, i.e., the reduction in uncertainty about either one of $\{X, Z\}$ from the knowledge of the other. Note, that if $I(X; Z)$ is zero, then knowing the value of Z does not reveal any additional information about X. Moreover, the above specific functional form for entropy rigorously follows from a set of basic axioms that one expects any information measure to have.

Because $Z = X + Y$, and X and Y are independent, we have $H(X, Z) = H(X) + H(Y)$. Using this in the above equation for $I(X, Z)$ gives

$$I(X; Z) = H(Z) - H(Y).$$

If X (resp., Y) is normal with mean μ_X (μ_Y) and variance σ_X^2 (σ_Y^2), then Z is normal with mean $\mu_Z = \mu_{X+}\mu_Y$ and variance $\sigma_Z^2 = \sigma_X^2 + \sigma_Y^2$. As the entropy of a normal distribution of variance σ is $2^{-1}(1 + \ln(2\pi\sigma^2))$ (Lazo and Rathie 1978), we get the following:

$$I(X; Z) = 2^{-1} \ln\left(1 + \frac{\sigma_X^2}{\sigma_Y^2}\right).$$

Using in the above $X = b_{k1}$ and $Y = b_{k2} + \cdots + b_{k|J|}$, where $k = i^{*(2)}$, gives

$$I(b_{k1}; b_{k1} + b_{k2} + \cdots + b_{k|J|}) = 2^{-1} \ln\left(1 + \frac{\sigma_{k1}^2}{\sigma_{k2}^2 + \cdots + \sigma_{k|J|}^2}\right).$$

Note that typical BOMs contain hundreds of components and hence $|J|$ is typically also in the hundreds, and in such cases the $I(b_{k1}; b_{k1} + b_{k2} + \cdots + b_{k|J|})$ is small, i.e., little information leakage takes place. Similarly, if the individual variances σ_{kj} are large then even for a relatively small $|J|$, the information leaked is also small.

This analysis can be useful in guiding each participant in deciding whether to participate in the protocol. Each EMS who knows in which components it is particularly competitive can roughly "guess" its J and estimate the information that would be leaked by the process if that EMS were to be $i^{*(2)}$: If $I(b_{k1}; b_{k1} + b_{k2} + \cdots + b_{k|J|})$ is large and the EMS wants to jealously guard its competitive edge in component 1, then such an EMS may decide that too much information leakage may take place and may decline to participate in the process.

8. Conclusions

This paper has described a business process for negotiating the procurement of component parts to be used by an EMS in the manufacture/assembly of the branded products of an OEM.

Our secure business process has four distinctive characteristics. First, it assures the privacy of both the OEM's and the EMSs' individual component prices; i.e., none of the parties in the negotiation learn the individual parts prices of any other party. More specifically, each party's prices remain inside their respective firewalls. The protocol works only on encrypted information, which is of no value to anyone (e.g., hackers, other parties) who might gain access to it. Second, the business process motivates the EMSs to bid their cost of each part, i.e., each EMS is incentivized to bid its own cost to purchase the part. Third, our SPM business process is non-invertible; i.e., none of the participants can "solve" for the private inputs of any other participant based on its own inputs and the outputs provided to it by the business process. Fourth, the posterior distribution of any other participant's private inputs is practically indistinguishable from its prior distribution.

Several extensions are planned. First, the development of mechanisms and secure protocols for multiple time periods. In practice, individual products are contracted for over a single time period. However, a multiple time-period model permits the explicit modeling of some of the motives behind awarding individual contracts to multiple EMSs (section 4.2); e.g., to maintain a plurality of bidders on future contracts or to honor long-term partnership agreements. Multiple time periods also pose the potential for parties to infer individual parts prices, despite the fact that their privacy is preserved during every single negotiation.

A second extension involves the development of mechanisms to handle volume-dependent prices. The existence of different levels of discounts for different levels of volume is, after all, one of the basic motives underlying the practice of price masking. In most cases, the incremental volumes required for incremental discounts are larger than the volume for each component involved in a single contract. Hence, our assumption of constant evaluations (i.e., prices) for each party is appropriate in most cases. However, for some components in some scenarios, valuations and, hence, the corresponding bids depend on volume. A third extension is the correlation of valuations across components. For example, a party's valuation for component A may depend on whether or not that party is chosen to provide component B.

Acknowledgments

Portions of this work were supported by Grants IIS-0325345 and CNS-0627488 from the National Science Foundation, and by sponsors of the Center for Education and Research in Information Assurance and Security. We also thank Chris Tang and two referees for their feedback on an earlier version of the paper.

Note

[1] See Appendix S1 for a discussion of trade-offs between various architectures.

References

Akerlof, G. 1970. The market for lemons: Quality uncertainty and the market mechanism. *Q. J. Econ.* **89**: 488–500.

Amaral, J., C. Billington, A. Tsay. 2006. Safeguarding the promise of production outsourcing. *Interfaces* **36**(3): 220–233.

Armstrong, M. 1996. Multiproduct nonlinear pricing. *Econometrica* **64**(1): 51–75.

Armstrong, M. 2000. Optimal multi-object auctions. *Rev. Econ. Stud.* **67**(3): 455–481.

Atallah, M., V. Deshpande, H. Elmongui, L. Schwarz. 2003. Secure supply-chain protocols. International Conference on E-Commerce, June 24–27, Newport Beach, CA, pp. 293–302.

Avery, C., T. Hendershott. 2000. Bundling and optimal auctions of multiple products. *Rev. Econ. Stud.* **67**(3): 483–497.

Bichler, M., R. Steinberg, eds. 2007. Special issue on e-auctions and procurement operations. *Prod. Oper. Manag.* **16**(4): 401–403.

Brandt, F., T. Sandholm. 2005. Efficient privacy-preserving protocols for multi-unit auctions. International Conference on Financial Cryptography and Data Security (FC), February 28–March 3, Roseau, The Commonwealth of Dominica. LNCS, Vol. 3570, pp. 298–312.

Cavinato, J., R. Kauffman. 1999. *The Purchasing Handbook*. New York, McGraw-Hill.

Clifton, C., A. Iyer, R. Cho, J. Vaidya, J. Wei, M. Kantivoglu. 2008. An approach to identifying beneficial collaboration securely in decentralized logistics systems. *Manuf. Serv. Oper. Manage.* **10**: 108–125.

Cover, T., J. Thomas. 1991. *Elements of Information Theory.* New York, Wiley-Interscience.

Damgård, I., Y. Ishai. 2005. Constant-round multiparty computation using a black-box pseudorandom generator. Advances in Cryptology—CRYPTO 2005, August 14–18, Santa Barbara, CA. LNCS, Vol. 3621, pp. 378–394.

Decker, B. De, G. Neven, F. Piessens, E. Van Hoeymissen. 2001. Second price auctions, a case study of secure distributed computing. IFIP International Working Conference on Distributed Applications and Interoperable Systems, September 17–19, Krakow, Poland, pp. 217–228.

Deshpande, V., L. Schwarz, M. Atallah, M. Blanton, K. Frikken, J. Li. 2009. Secure-computations for collaborative planning, forecasting and replenishment (scpfr). Purdue University, Working paper.

de Vries, S., R. Vohra. 2003. Combinatorial auctions: A survey. *INFORMS J. Comput.* **15**(3): 284–309.

Elkind, E., H. Lipmaa. 2004. Interleaving cryptography and mechanism design: The case of online auctions. Conference on Financial Cryptography (FC), February 9–12, Key West, FL, pp. 117–131.

Elmaghraby, W. 2000. Supply contract competition and sourcing policies. *Manuf. Serv. Oper. Manage.* **2**(4): 350–371.

Elmaghraby, W. 2007. Auctions within e-sourcing events. *Prod. Oper. Manag.* **16**(4): 409–422.

Franklin, M., M. Reiter. 1996. The design and implementation of a secure auction service. *IEEE Trans. Softw. Eng.* **22**(5): 302–312.

Fudenberg, D., J. Tirole. 2000. *Game Theory.* MIT Press, Cambridge, MA.

Goldreich, O. 2004. *Foundations of Cryptography. Volume 2: Basic Applications*. New York, Cambridge University Press.

Goldreich, O., S. Micali, A. Wigderson. 1987. How to play any mental game or a completeness theorem for protocols with honest majority. ACM Symposium on Theory of Computing (STOC), May 25–27, New York, pp. 218–229.

Jakobsson, M., A. Juels. 2000. Mix and match: Secure function evaluation via ciphertexts, ASIACRYPT 2000, December 3–7, Kyoto, Japan, pp. 162–177.

Jenster, P. 2005. *Outsourcing-Insourcing: Can Vendors Make Money from the New Relationship Opportunities?* J. Wiley, New York, NY.

Klemperer, P. 1999. Auction theory: A guide to the literature. *J. Econ. Rev.* **3**: 227–260.

Kramer, R., T. Tyler. 1996. *Trust in Organizations: Frontiers of Theory and Research*. Sage, Thousand Oaks, CA.

Lazo, A., P. Rathie. 1978. On the entropy of continuous probability distributions (corresp.). *IEEE Trans. Inf. Theory* **24**(1): 120–122.

Lee, H. L., C. S. Tang. 1996. Managing supply chains with contract manufacturing. *Asian J. Bus. Inf. Syst.* **1**(1): 11–22.

Lindell, Y., B. Pinkas. 2004. A proof of yao's protocol for secure two-party computation. Cryptology ePrint Archive, Report 2004/175. Available at http://eprint.iacr.org/ (accessed date May 6, 2010).

Malkhi, D., N. Nisan, B. Pinkas, Y. Sella. 2004. Fairplay—A secure two-party computation system. Usenix Security Conference, August 9–13, San Diego, CA, pp. 287–302.

Milgrom, P., R. Weber. 1982. A theory of auctions and competitive bidding. *Econometrica* **50**(5): 1089–1122.

Myerson, R. 1981. Optimal auction design. *Math. Oper. Res.* **6**(1): 58–73.

Naor, M., B. Pinkas, R. Sumner. 1999. Privacy preserving auctions and mechanism design. ACM Conference on Electronic Commerce, November 3–5, Denver, CO, pp. 129–139.

Nelson, D., P. Moody, J. Stegner. 2005. *The Incredible Payback: Innovative Sourcing Solutions that Delivery Extraordinary Results*. AMACOM, New York, NY.

Palfrey, T. 1983. Bundling decisions by a multiproduct monopolist with incomplete information. *Econometrica* **51**(2): 463–483.

Pick, A. 2004. Survival of the fittest? EMSnow.com. Available at http://www.emsnow.com/newsarchives/archivedetails.cfm?ID=6633 (accessed date May 6, 2010).

Riley, J., W. Samuelson. 1981. Optimal auctions. *Am. Econ. Rev.* **71**(3): 381–392.

Rothkopf, M. H. 2007. Thirteen reasons why the Vickrey-Clarke-Groves process is not practical. *Oper. Res.* **55**(2): 191–197.

Rothkopf, M. H., A. B. Whinston. 2007. On e-auctions for procurement operations. *Prod. Oper. Manag.* **16**(4): 404–408.

Rothschild, M., J. Stiglitz. 1976. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *Q. J. Econ.* **80**: 629–649.

Simchi-Levi, D. 2004. *Managing the Supply Chain: The Definitive Guide for the Business Professional*. New York, McGraw-Hill.

Simchi-Levi, D., S. Wu, Z. Shen. 2004. *Handbook of Quantitative Supply Chain Analysis*. New York, Springer Science-Business Media, Inc.

Spence, A. 1974. *Market Signaling*. Harvard University Press, Cambridge, MA.

Sullivan, L. 2003. Motorola rewrites rules for EMS: New policy removes contractors from pricing negotiations. *Electronics Supply and Manufacturing*. Available at: www.eetimessupplynetwork.com/16000660 (accessed date October 5, 2010).

Vickrey, W. 1961. Counterspeculation, auctions and competitive sealed tenders. *J. Finance* **16**(1): 8–37.

Yao, A. 1982. Protocols for secure computations. ACM Symposium on Theory of Computing (STOC), May 5–7, San Francisco, CA, pp. 128–136.

Yao, A. 1986. How to generate and exchange secrets. IEEE Symposium on Foundations of Computer Science. October 27–29, Toronto, Canada, pp. 162–167.

Supporting Information

Additional supporting information may be found in the online version of this article:

Appendix S1: Technical Appendix.

Please note: Wiley-Blackwell is not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.