

Optimizations for Risk-Aware Secure Supply Chain Master Planning

Axel Schröpfer

(SAP Research Karlsruhe, Germany
axel.schroepfer@sap.com)

Florian Kerschbaum

(SAP Research Karlsruhe, Germany
florian.kerschbaum@sap.com)

Christoph Schütz

(SAP Research Karlsruhe, Germany
christoph.schuetz@sap.com)

Richard Pibernik

(Supply Chain Management Institute, EBS Wiesbaden, Germany
pibernik@supplyinstitute.org)

Abstract: Supply chain master planning strives for optimally aligned production, warehousing and transportation decisions across a multiple number of partners. Its execution in practice is limited by business partners' reluctance to share their vital business data. Secure Multi-Party Computation (SMC) can be used to make such collaborative computations privacy-preserving by applying cryptographic techniques. Thus, computation becomes acceptable in practice, but the performance of SMC remains critical for real world-sized problems. We assess the disclosure risk of the input and output data and then apply a protection level appropriate for the risk under the assumption that SMC at lower protection levels can be performed faster. This speeds up the secure computation and enables significant improvements in the supply chain.

Key Words: Linear Programming, Privacy, Security and Protection

Category: C.4, H.3

1 Introduction

Supply chain master planning (SCMP) strives for optimally aligned production, warehousing and transportation decisions across multiple partners. In practice, we can commonly observe a decentralized coordination mechanism (referred to as upstream planning) that usually only leads to local optima rather than to global supply chain optima [Dudek, Stadtler 2005]. At least in theory, optimal master plans can be generated for the whole supply chain if some planning unit has at its disposal all relevant information pertinent to the individual partners in the supply chain. It is, however, a well known fact that companies are typically not willing to share sensitive private data (e.g. cost and capacity data)

([Pibernik, Sucky 2006, Pibernik, Sucky 2007]). They perceive the risk that the central planning unit or other parties misuse data to their disadvantage in order to obtain additional benefits.

The major obstacles to centralized master planning can be removed if a mechanism for securely and privately computing the supply master plan is in place [Atallah et al. 2009]. A central planning unit, e.g. a 4th party logistics provider (4PL), could then determine globally optimal master plans and distribute these to the individual partners involved in the supply chain. To this end, Secure (Multi-Party) Computation (SMC) can be employed such that the relevant data does not need to be disclosed even to the central planning unit. This offers the ultimate level of protection, since no data sharing risk remains. In this paper we propose a framework for secure centralized supply chain master planning (SS-CMP). We introduce a basic model for centralized supply chain master planning and, from this, derive the relevant data a central planning unit requires to optimally coordinate manufacturing and transportation decisions. We then analyse this data with respect to its "criticality". Criticality refers to how sensitive certain pieces of data are and how willing the different partners will be to share this data. The criticality is determined by the perceived risks associated with data sharing and its prior public knowledge. In this context, risk can be characterized by the potential negative impact that occurs if a partner misuses the data to its own benefit and the likelihood for this to happen. We derive an overall criticality assessment for each data element that is relevant for supply chain master planning and use information about on the prior (public) knowledge of the data to determine an overall criticality score. This criticality score constitutes an input to secure computation of centralized supply chain master plans. We map criticality scores to protection levels which consist of certain technologies and parameters for SMC. Lower protection levels lead to faster SMC implementations. We propose a mixed approach to SMC combining the different protection levels in one implementation and propose modifications to Linear Programming (LP) that optimize the effort involved by selecting the pivot element based on the protection level. We experimentally verify the effectiveness of the new algorithm.

2 Related Work

Numerous works in the area of supply chain management exists on supply chain master planning as well as information sharing and collaboration in supply chains. In general, it is a well acknowledged fact that sharing relevant information and planning in a collaborative fashion can improve supply chain performance and mitigate the consequences of demand variability, especially with respect to the well-known bullwhip effect (see for example [Chen et al. 1999, Lee et al. 1997, Min et al. 2005, Yu et al. 2001]). With respect to supply chain

master planning, numerous authors have proposed multi-stage models that can be utilized to coordinate planning activities across multiple locations and firms (e.g. [Fleischmann, Meyr 2003, Jayaraman, Pirkul 2001, Shapiro 2001]). Various authors have stated that employing a centralized approach to master planning will lead to better results as compared to decentralized approaches that are most commonly employed in industry. [Simpson, Erengüç 2001], for example, analyze the disadvantages of upstream coordination in comparison with centralized coordination. They compute the average gap between centralized and upstream coordination for several test scenarios with varying cost parameters and demand patterns. Similar findings are reported in [Pibernik, Sucky 2007]. However, centralized supply chain planning has not been widely adopted in industry. [Holström et al. 2002] states: "it is difficult, or maybe even impossible, to get a large network consisting of independent companies to agree on and implement a centralised planning and control solution." Reluctance towards information sharing (a prerequisite for centralized master planning) has been identified as the main obstacle that inhibits centralized master planning ([Pibernik, Sucky 2006, Pibernik, Sucky 2007]). For this reason, alternative approaches have been developed that either build on negotiation based coordination ([Dudek, Stadtler 2005]) or hybrid forms ([Pibernik, Sucky 2006]). So far there has been no research on supply chain master planning based on mechanisms that privacy preserving data sharing and computation. To the best of our knowledge the only approach to secure multi-party computation in the area of supply chain management can be found in [Atallah et al. 2009]. The authors develop secure protocols for a Collaborative Planning, Forecasting, and Replenishment (CPFR) process. Next to the fact that we, in our paper, consider a different problem setting, a major distinction between the research presented in [Atallah et al. 2009] and our research is that they do not consider different protection levels for different risks of data to be shared. They follow the approach to provide the highest protection for all data using a specially developed protocol. Their protocols are two-party protocols, while we consider a multi-party problem. We will now review related work for SMC.

SMC allows a set of n players, $P = P_1, \dots, P_n$, to jointly compute an arbitrary function of their private inputs, $f(x_1, \dots, x_n)$. The computation is privacy preserving, i.e. nothing else is revealed to a player than what is inferable by his private input and the outcome of the function. A cryptographic protocol is then run between the players in order to carry out the computation. Even if there are adversarial players, the constraints on correctness and privacy can be proven to hold under well stated settings. These settings consider the type of adversary as well as his computing power which can be bounded or unbounded. An adversary can be passive, i.e. following the protocol correctly but trying to learn more or he can be active, by arbitrarily deviating. For the two-party case it has been proven

by Yao in [Yao 1982], that any arbitrary function is computable in privacy preserving fashion, using garbled binary circuits. This approach has been extended to the multi-party case in [Beaver et al. 1990, Goldreich et al. 1987]. Alternative approaches base on secret sharing schemes. A player's secret s is split into m shares which are then distributed to m players. Players can compute intermediate results on the shares, and in the end a reconstruction is performed in order to receive the final result. Other approaches utilize semantically secure homomorphic encryption (HE) [Damgård, Jurik 2001], a public encryption scheme, where $E(x) \cdot E(y) = E(x + y)$ and x cannot be deduced by $E(x)$.

Using the general approach leads to solutions that have high complexity and are therefore almost always not practically feasible [Li, Atallah 2006]. Thus, in order to get a practical solution, a dedicated protocol should be constructed. Atallah et al. constructed solutions for a couple of supply chain problems, e.g. planning, forecasting, replenishment, benchmarking, capacity allocation and e-auctions ([Atallah et al. 2009, Atallah et al. 2004, Atallah et al. 2003]). Their cryptographic protocols base on additive secret sharing, homomorphic encryption and garbled circuits. A contribution of Atallah et al. which is closely related to ours is that of secure linear programming [Li, Atallah 2006]. It uses the simplex method introduced by Dantzig in [Dantzig, Thapa 1997] to solve linear programs which get expressed as a matrix D . The method consists of two steps: selecting the pivot element d_{rs} and pivoting all elements d_{ij} of D over this element. The pivot step sets the new value of d_{ij} , denoted d'_{ij} , by

$$\begin{aligned} d'_{ij} &= \frac{1}{d_{rs}} && \text{for } i = r \text{ and } j = s \text{ (pivot element)} \\ d'_{ij} &= \frac{d_{ij}}{d_{rs}} && \text{for } i = r \text{ and } j \neq s \text{ (pivot row)} \\ d'_{ij} &= \frac{-d_{ij}}{d_{rs}} && \text{for } i \neq r \text{ and } j = s \text{ (pivot column)} \\ d'_{ij} &= \frac{d_{ij} - d_{is}d_{rj}}{d_{rs}} && \text{for } i \neq r \text{ and } j \neq s \text{ (all other elements)}. \end{aligned}$$

The method is repeated until the optimal solution of the LP is found (resp., it is stated that the problem is unbounded or infeasible). As input to the cryptographic protocol, matrix D gets additively split between both parties (i.e., $D = D^{(a)} + D^{(b)}$). In order to not reveal additional information (e.g. by the pivot column or row index), the matrix gets permuted at the beginning of each iteration. Details are omitted here, but can be found in [Li, Atallah 2006]. The pivot element selection and the pivot step are then carried out using cryptographic tools additive splitting, homomorphic encryption and garbled circuits.

3 Supply Chain Master Planning

In this section we first provide a basic model for centralized supply chain master planning. This model will be used to derive the relevant data that partners in

Master planning parameters (input)

D_l^n	Demand for final finished product $n \in N_I$ at customer location $l \in K_{I+1}$
$\alpha^{m,n}$	Quantity of input product m required for manufacturing one unit of output product n
β_k^n	Unit capacity requirement at location $k \in K_i$ for output of product $n \in N_i$
$cap_{i,k}$	Production capacity at location $k \in K_i$
$cp_{i,k}^n$	Unit production costs of product $n \in N_i$ at location $k \in K_i$
$cs_{i,k,l}^n$	Unit shipping costs of product $n \in N_i$ from location $k \in K_i$ to location $l \in K_{i+1}$
$ch_{i,k}^n$	Unit holding costs of product $n \in N_i$ at location $k \in K_i$

Master planning variables (output)

$x_{i,k}^n$	Production quantity of output product $n \in N_i$ manufactured at location $k \in K_i$
$y_{i,k,l}^n$	Shipping quantity of product $n \in N_i$ shipped from location $k \in K_i$ to $l \in K_{i+1}$

the supply chain need to share for centralized master planning. We then propose a simple approach to assess the criticality of the individual elements.

3.1 Model for Centralized Supply Chain Master Planning

As a basis for our subsequent analysis we utilize a simple generic supply chain master planning model presented in [Pibernik, Sucky 2007]. Although rather simple, this model is sufficient for the illustration of our concept and can easily be extended in order to account for further practical requirements and restrictions. We consider a supply chain with I stages on which different operations (e.g. manufacturing, warehousing, etc.) are performed. We use index i ($i = 1, \dots, I$) to distinguish the different stages. By $I + 1$ we denote the final customer stage. By K_i we denote the set of nodes on stage i . Every node $k \in K_i$ represents one production facility or warehouse on stage $i = 1, \dots, I$. The final customer locations are modelled through nodes $k \in K_{I+1}$ on stage $i = I + 1$. By N_i we denote the set of products produced on stage i and use $m \in N_{i-1}$ and $n \in N_i$ as indices for the input and output products of stage i . For a given supply chain, master planning determines the production and inventory quantities for every node and the material flows between the nodes for a given time period. We introduce the following additional notation to formulate a centralized master planning model:

Objective function

$$\begin{aligned} \text{Min } C = & \sum_{i=1}^I \sum_{k \in K_i} \sum_{n \in N_i} cP_{i,k}^n x_{i,k}^n + \sum_{i=1}^I \sum_{k \in K_i} \sum_{n \in N_i} ch_{i,k}^n x_{i,k}^n + \\ & \sum_{i=1}^I \sum_{k \in K_i} \sum_{l \in K_{i+1}} \sum_{n \in N_i} cS_{i,k,l}^n y_{i,k,l}^n \end{aligned} \quad (1)$$

Constraints

$$\sum_{k \in K_I} y_{I,k,l}^n = D_l^n \quad \forall n \in N_I, l \in K_{I+1} \quad (2)$$

$$x_{i,k}^n = \sum_{l \in K_{i+1}} y_{i,k,l}^n \quad \forall n \in N_i, i \in \{1, \dots, I\}, k \in K_i \quad (3)$$

$$\sum_{j \in K_{i-1}} y_{i,j,k}^m = \sum_{n \in N_i} \alpha^{m,n} x_{i,k}^n \quad \forall m \in N_{i-1}, i \in \{1, \dots, I\}, k \in K_i \quad (4)$$

$$\sum_{n \in N_i} \beta_k^n x_{i,k}^n \leq cap_{i,k} \quad \forall i \in \{1, \dots, I\}, k \in K_i \quad (5)$$

$$x_{i,k}^n, y_{i,k,l}^n \geq 0 \quad \forall n \in N_i, i \in \{1, \dots, I\}, k \in K_i \quad (6)$$

The following deterministic, linear programming model can be used to determine a supply chain master plan. The objective of the model is to minimize the total relevant costs of the SC for fulfilling final customer demand. The objective function (1) accounts for production costs, holding costs, and shipping costs for finished products. Constraints (2) ensure that the final customer demand at stage $I + 1$ is met. (3) and (4) represent finished product and intermediate product balance constraints. The capacity constraints (5) ensure that the available capacity of any location will not be exceeded. Constraints (6) ensure non-negativity of all decision variables.

The output of this model is a supply chain master plan for a single period that specifies the production quantity for the individual products in each node and the shipping quantities across the whole supply chain. From this basic model we can directly infer the relevant data that needs to be shared in order to realize centralized supply chain master planning. All parties in the supply chain need to make the above listed *input parameters* available to the central planning unit. After generating the supply chain master plan, the central planning unit has to communicate the results (i.e. the values of the *output variables*) to the corresponding partners. In typical industry settings, both the input parameters and the master planning output constitute private data that is only accessible to the planning units (firms, departments) responsible for individual nodes and arcs. The willingness to share this data will depend on the risk the individual data owners perceive. The perceived risk, however, is not identical for all of the relevant data elements. A company may, for example perceive a low risk associ-

ated with sharing forecast data, but a high risk when revealing production cost or capacity information. While SMC can overcome the risk of data sharing in theory with the highest protection level, in practice such solutions can become too slow to be useful (e.g. if the computation takes longer than what the continuous planning period is). We therefore use the result of the risk assessment, the criticality scores, to optimize the SMC, such that each data element is handled at its appropriate risk level. We achieve a significant performance improvement in our experiments.

3.2 Data Criticality and Protection Levels

In this section we illustrate a simple approach to determine protection levels for individual data elements in the context of centralized master planning. Although it is rather straightforward to see that the risk will differ across the individual data elements, it is not possible to determine general criticality levels that are valid for any supply chain setting. Whether other partners in the supply chain can use data to their benefit and to the disadvantage of the data owner depends on factors such as the distribution of power among the partners, the type of industry and product, the relative position in the supply chain, trust among partners, etc. Production costs, for example are generally considered as critical data that a data owner will not want to share. However, in many industries (e.g. for commodities) production costs are known by different partners without implying a negative impact. Because a general assessment of the criticality is not likely to be attainable, an individual assessment has to be conducted for any specific supply chain. We propose a simple scheme to support such a criticality assessment. It is based on the following questions that need to be answered for any one of the data elements identified in the previous section:

1. What disadvantage may a data owner potentially incur when sharing private data?
2. What is the probability that a partner in the SC (mis-) uses the shared data to the disadvantage of the data owner?
3. To what extent is the data prior knowledge?

With the first two questions we capture the individual components of the risk induced by sharing a certain data element. When considering the potential negative impacts (question one), we have to consider that these may vary depending on the position of the data source within the supply chain and the potential incentives other partners in the SC may have to (mis-) use the data. We differentiate between partners who are responsible for nodes on the same stage (competitors) and those who are responsible for nodes on previous or subsequent stages (supplier-buyer-relationships). For each of the aforementioned cases

Potential negative impacts

- induced by competitors: impact * probability \Rightarrow score
 - induced by suppliers: impact * probability \Rightarrow + score
 - induced by buyers: impact * probability \Rightarrow + score
- [0; 5] [0; 5] [0; 25]

Overall risk assessment (Sum of individual scores):

- Public knowledge (*pub*): score (*pub*)
 - Prior knowledge of specific stages (*spe*): + score (*spe*)
- [0; 5]

Prior knowledge (*pk*)

Prior knowledge weight: $(1 - pk/10)$

Criticality assessment

- = Overall risk assessment * prior knowledge weight:

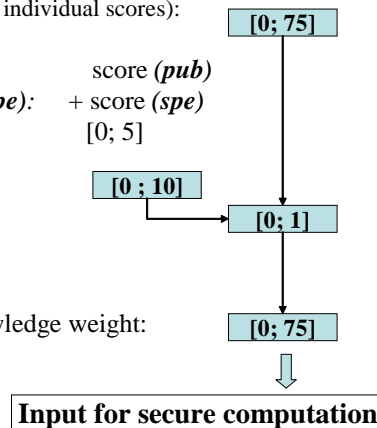


Figure 1: Determination of criticality

it is necessary to assess the likelihood of a disadvantage on the side of the data owner, i.e. the probability that another partner in the supply chain will actually make use of the knowledge of the data element (question 2). The risk cannot be considered independent of the prior knowledge about the data. It is reasonable to assume that the criticality of certain data elements is lower if the data is already accessible for some or all of the partners in the supply chain. Figure 3.2 illustrates our basic scheme for assessing the criticality of individual data elements. We propose a scoring range between zero and five to adequately assess by discrete values the potential negative impact and the expected probability of data misuse. Through multiplication of both scores, we obtain a particular risk measure for negative impacts induced by competitors, suppliers, or buyers. Their addition provides a measure for the overall risk. The overall risk for each data element is then weighted with a value that expresses the prior knowledge of data. Similarly, a scoring range from zero to five is used to measure the degree of public knowledge in general as well as specific knowledge of individual SC partners. The sum of both scores measures the level of prior knowledge. A score of zero indicates that the data is pertinent to the data owner, while higher scores indicate that the data may anyways be known prior to centralized master planning. We determine an aggregate weight for the prior knowledge as in order to derive the overall criticality level. In Figure 3.2 we provide an example of a

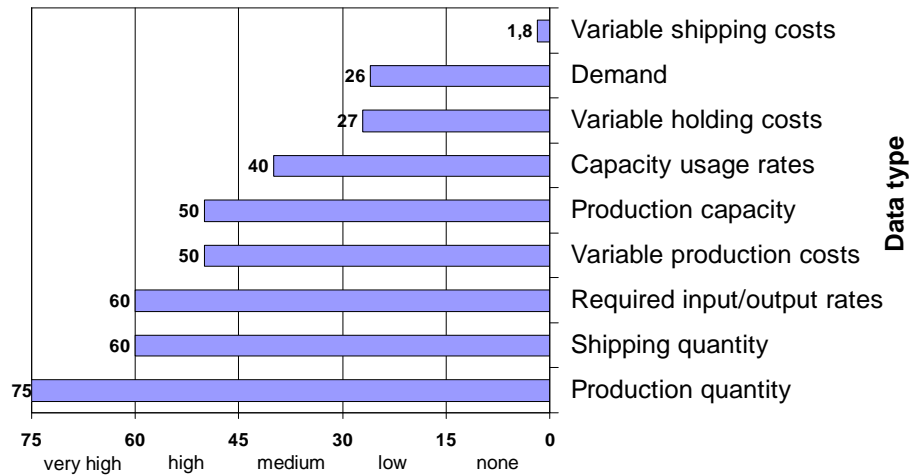


Figure 2: Criticality levels of different types of data (example)

possible outcome of our criticality assessment. With our assessment scheme, a criticality score between zero and 75 is assigned to each data element.

4 Secure Computation

4.1 Protection Levels

The data criticality analysis of section 3.2 shows that different variables in the SCMP problem have different perceived risk and in order to mitigate this risk require different protection. The data criticality scores of the variables range from zero to 75. We map a data criticality score to a *protection level*.

A protection level specifies a concrete set of SMC technologies and their parameters for protecting a variable. These technologies and parameters are: the computational *setting* (information-theoretic, cryptographic or best-effort), the cryptographic *tools* (Homomorphic Encryption, SHA-1, etc.) and the *tool parameters*. Dependencies among the different parameters of a protection level are possible, e.g. there cannot be a SMC computation that is information-theoretically secure, but uses homomorphic encryption as a tool. It is reasonable to arrange the protection levels in order of the effort for an attacker to infer the protected value and higher protection levels require higher computation and communication effort. Higher criticality scores map to higher protection levels. We currently

Table 1: Linear Mapping

Protection Level	1	2	3	4	5
Criticality Score	0-15	16-30	31-45	46-60	61-75
Number of Data	3	4	0	2	2

do not take into account costs for converting data from from one protection level to another.

For an implementation of protection levels, one may utilize different key sizes for encryption schemes, hash algorithms of different strength, different share sizes in secret sharing schemes or any other suitable security parameter.

4.2 Mapping

A monotone function maps the criticality score c to a protection level $p = f(c)$. We propose a linear mapping. Other mappings are possible and may depend on the concrete application context. We assume the ordered protection levels differ in their effort by an almost constant factor. We define a linear mapping function $f(c)$ which maps data criticality score c to m protection levels by $f(c) = 1 + \lfloor c \cdot m / (c_{max} + 1) \rfloor$, where c_{max} is 75 in Section 3.2. Applying this mapping to the criticality scores of section 3.2, we receive the values of Table 2. Applying Table 2 to the criticality scores of Figure 2, nine of eleven variables are assigned a protection level below the maximum. Thus, we expect a significant reduction in computational effort by applying risk-aware protection levels to the variables compared to the straight-forward approach of applying maximum protection to all variables.

4.3 Integration of Protection Levels

We adapt the secure linear programming protocol by Atallah et al. by introducing an additional matrix denoted P . Every element of P , p_{ij} , represents the protection level of the corresponding data element in D , d_{ij} . Recall that the LP is rewritten as a matrix D . Let

$$D = \begin{pmatrix} c^T & -z_0 \\ A & b \end{pmatrix}$$

where c^T denotes the vector of the objective function's coefficients, z_0 the outcome, A the coefficients of the constraints and b the vector of the constraint values. Atallah et al.'s secure linear programming protocol uses a slight adaptation of the Bland's Rule [Bland 1977] as pricing scheme. The rule computes the pivot column s as the left most column with a negative element.

$$s = \arg \min \{j : c_j < 0\}$$

The pivot row r is computed as the row with the minimum ratio between a row element in the pivot column and the corresponding element in b . In case of equal ratios the one with the lowest row index is selected. More formally, the pivot row is obtained by the condition

$$j_r = \min \left\{ j_i : i \in \arg \min \left\{ \frac{b_i}{a_{is}} : a_{is} > 0 \right\} \right\}.$$

All parties agree on the specifications of the protection levels and have plaintext access to the matrix P of protection levels for the elements of the data matrix D . P gets blinded and identically permuted as D in the original protocol. P may leak little information, e.g. if there is a unique occurrence of a protection level.

Whenever a pivot step is performed in order to receive a new value d'_{ij} , the new protection level value p'_{ij} is set to the highest protection level assigned to any element of D involved in the computation. According to the pivot rules of the Simplex algorithm the involved elements for element d'_{ij} with pivot element d_{rs} may be: d_{rs} , d_{rj} , d_{is} and d_{ij} . Let p_{B_r} denote the protection level of the variable assigned to row r in the current basis. The new protection level p'_{ij} for d'_{ij} then is computed as

$$\begin{aligned} p'_{ij} &= p_{ij} && \text{for } i = r \text{ and } j = s \text{ (pivot element)} \\ p'_{ij} &= \max(p_{ij}, p_{rs}, p_{B_r}) && \text{for } i = r \text{ and } j \neq s \text{ (pivot row)} \\ p'_{ij} &= \max(p_{ij}, p_{rs}) && \text{for } i \neq r \text{ and } j = s \text{ (pivot column)} \\ p'_{ij} &= \max(p_{ij}, p_{is}, p_{rj}, p_{rs}) && \text{for } i \neq r \text{ and } j \neq s \text{ (all other elements)} \end{aligned}$$

Over a number of iterations the elements of P will converge to the maximum protection value. We introduce three approaches to counter a quick convergence of P . Our goal is to minimize the effort of the cryptographic SMC operations.

4.4 An Optimized Rule for Row Selection

Existing pricing schemes for row selection only consider values of matrix D . We construct a pivot selection rule which not only uses entries of D , but also entries of P and moreover prevents P from fast convergence. Our first approach keeps the algorithm of the Bland's Rule for selecting the pivot column, but replaces the algorithm for selecting the row. We select r for $0 < r \leq m$ as

$$r = \min \left(\frac{b_r}{a_{rs}} \right) : \min \left(\sum_{i=0}^{m-1} \sum_{j=0}^{n-1} \max(p_{rs}, p_{is}, p_{rj}, p_{ij}) - p_{ij} \right).$$

This formula selects the element in the pivot column s from the set of elements with minimum ratio as described in Section 4.3 which has the lowest impact on the convergence of matrix P . Note that this rule will only deviate from Bland's rule if ties in the minimum ratio occur.

4.5 An Optimized Rule for Column Selection

Existing pricing schemes for column selection only consider values of matrix D as was the case for the row selection. Again, we change the rule to also consider values of matrix P . We base the selection of s on two conditions: one that $c_s^T < 0$ necessary for decreasing the objective function and one that c_s^T has minimal expected impact on the convergence of matrix P . We express the expected impact on the convergence of P as a weight $w(i)$ for each column i . Since the impact of the pivot selection can only be exactly determined after row selection, we employ three different weight functions and experimentally evaluate their effectiveness.

The first weight function w_{max} is the maximum protection level value in column i .

$$w_{max}(i) = \max(p_{ji}) \text{ for } 0 \leq j < m$$

The second weight function w_{sum} is the sum (average) of the protection levels in column i .

$$w_{sum}(i) = \sum p_{ji} \text{ for } 0 \leq j < m$$

The third weight function w_{freq} is the sum (average) of the squares of the protection levels in column i . The intuition is to find a compromise between average and maximum.

$$w_{freq}(i) = \sum p_{ji}^2 \text{ for } 0 \leq j < m$$

We define the rule for selecting the pivot column s to

$$s = \arg \min w(i) : c_i^T < 0 \text{ for } 0 \leq i < n.$$

This rule requires the access to the protected elements c_i^T and, as such, must be partially performed as a secure computation. We suggest to use scrambled circuits as in [Atallah et al. 2003].

4.6 Optimized Pre-sort

Instead of modifying the column selection rule and thereby modifying the SMC protocol, we can pre-sort the columns, such that Bland's rule is likely to find the column with minimal impact on the convergence of P . Recall, that Bland's rule selects the leftmost column with $c_i^T < 0$. Therefore, before engaging in the secure linear programming protocol the parties sort the columns of D in ascending order

of their initial weight $w(i)$ as determined by the risk assessment. We emphasize that the weights can change during protocol execution and it is not guaranteed that the columns remain sorted and that Bland's rule will always select the same column as our modified column selection rule. We again experimentally evaluate the effectiveness of the pre-sorting approach.

5 Experiments

We determine the most effective approach of the three described in Section 4 by experimentation based on the data gathered in Section 3. We also compare our approaches to the straight-forward approach of always applying the maximum protection level.

In our experiments we made the simplifying assumption that each product and each party uses the same protection level, since our available data of criticality scores from risk assessment is limited. In practice, one can assume a greater variety of protection levels, likely to increase the variance of the expected benefit.

5.1 Setup

We replicated the linear programming algorithm used in Atallah et al.'s protocol in local, non-secure implementation. Note that it is not necessary to run the cryptographic protocol in order to estimate the performance gain of our approaches. We estimate the cost of the protocol as the sum of the costs of the updates to matrix D . In each iteration, each element d_{ij} is updated to d'_{ij} . We estimate the cost of each update as its new protection level p'_{ij} .

We sample the problem domain by randomly generated instances of the supply chain master planning (SCMP) problem. Our random SCMP generator takes as input the number of stages, the number of producers per stage, the number of products per stage and the number of customer. As output it produces a random data matrix D and its corresponding protection level matrix P . Table 2 shows the six supply chain models used in our experiments. For each model we generated 100 random instances.

Table 3 shows the combination of approaches used in our experiments. In each run of the experiments we measured the cost of the protocol and the number of pivot steps.

5.2 Results

As a baseline we used the cost of only using Bland's rule and normalized all costs accordingly. In Table 4 we give the resulting relative effort. In brackets we denote the standard deviation. Figure 3 shows a graphical representation of the results in Table 4.

Table 2: Supply chain network characteristics (experiment)

Type	Stages	Producers (per Stage)	Products (per Stage)	Customer	Dimension
M_0	2	2	2	8	93x141
M_1	3	2	2	8	123x183
M_2	2	3	2	8	137x215
M_3	2	2	3	8	137x209
M_4	3	3	2	8	188x290
M_5	2	3	3	8	202x319

Table 3: Experiments setup details

Name	Description
Bland	Bland without any Protection Levels (i.e., P_{max})
BlandP	Bland with Protection Levels
BlandPRow	Bland with Protection Levels, optimized row selection
BlandPSortMax	Bland with Protection Levels, Pre-sort (f_{max})
BlandPSortSum	Bland with Protection Levels, Pre-sort (f_{sum})
BlandPSortFreq	Bland with Protection Levels, Pre-sort (f_{freq})
BlandPRowSortMax	Bland with Protection Levels, optimized row selection, Pre-sort (f_{max})
BlandPRowSortSum	Bland with Protection Levels, optimized row selection, Pre-sort (f_{sum})
BlandPRowSortFreq	Bland with Protection Levels, optimized row selection, Pre-sort (f_{freq})
BlandPColumnMaxRow	Bland with Protection Levels, optimized column selection (f_{max}), optimized row selection
BlandPColumnSumRow	Bland with Protection Levels, optimized column selection (f_{sum}), optimized row selection
BlandPColumnFreqRow	Bland with Protection Levels, optimized column selection (f_{freq}), optimized row selection

In Table 5 we give the number of pivot steps relative to the Bland's rule. Values in brackets again denote the standard deviation. Figure 4 shows graphical representation of the results in Table 5.

One can see that risk-aware protection levels already reduce the cost of the protocol between 30-38% compared to always applying the maximum protection level in our experiments, if only Bland's rule is used. Using our row selection approach, the costs have further been reduced by 5-12%. Using our column pre-




Figure 3: Overview of effort of all experiments

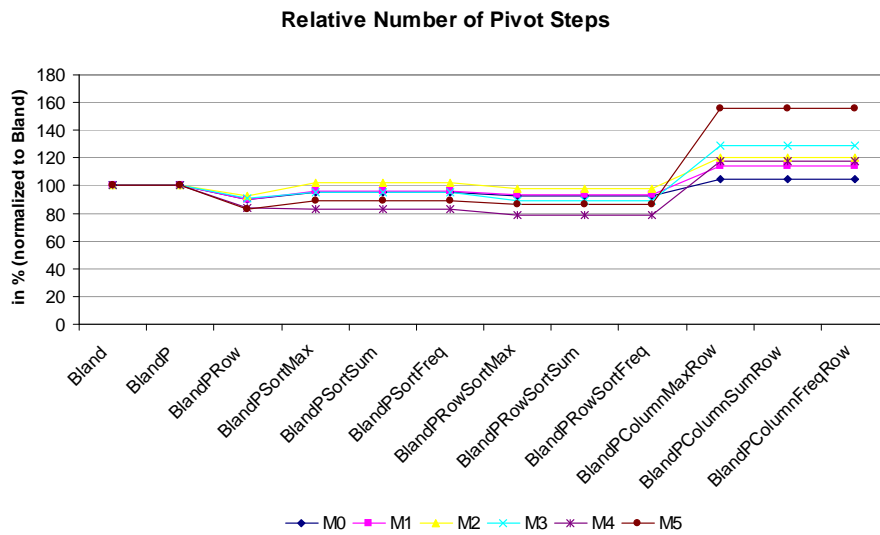


Figure 4: Overview of number of pivot steps of all experiments

Table 4: Relative effort of all experiments (in % normalized to Bland)

Setup	M_0	M_1	M_2	M_3	M_4	M_5
Bland	100	100	100	100	100	100
	(19,87)	(46,34)	(20,58)	(19,47)	(36,47)	(23,81)
BlandP	61,63	64,71	65,07	61,57	69,24	66,25
	(24,24)	(53,80)	(23,76)	(23,76)	(39,54)	(26,99)
BlandPRow	54,29	57,13	59,59	54,33	57,15	53,69
	(23,26)	(28,40)	(26,64)	(28,96)	(30,97)	(28,31)
BlandPSortMax	32,92	46,76	41,67	32,37	45,44	35,49
	(34,51)	(33,95)	(35,45)	(36,43)	(31,80)	(35,65)
BlandPSortSum	32,92	46,76	41,67	32,37	45,44	35,49
	(34,51)	(33,95)	(35,45)	(36,43)	(31,80)	(35,65)
BlandPSortFreq	32,92	46,76	41,67	32,37	45,44	35,49
	(34,51)	(33,95)	(35,45)	(36,43)	(31,80)	(35,65)
BlandPRowSortMax	32,01	45,36	41,85	31,18	43,61	38,19
	(35,70)	(34,08)	(38,81)	(39,35)	(31,72)	(35,88)
BlandPRowSortSum	32,01	45,36	41,85	31,18	43,61	38,19
	(35,70)	(34,08)	(38,81)	(39,35)	(31,72)	(35,88)
BlandPRowSortFreq	32,01	45,36	41,85	31,18	43,61	38,19
	(35,70)	(34,08)	(38,81)	(39,35)	(31,72)	(35,88)
BlandPColumnMaxRow	58,52	70,82	74,67	77,18	79,15	103,34
	(33,52)	(42,57)	(37,29)	(28,98)	(33,36)	1(35,42)
BlandPColumnSumRow	58,52	70,82	74,67	77,18	79,15	103,34
	(33,52)	(42,57)	(37,29)	(28,98)	(33,36)	1(35,42)
BlandPColumnFreqRow	58,52	70,82	74,67	77,18	79,15	103,34
	(33,52)	(42,57)	(37,29)	(28,98)	(33,36)	1(35,42)

sorting we achieved cost savings of 53-67%. We received the highest cost savings of 54-68% using column pre-sorting together with optimized row selection.

The results also show that using the optimized rule for row selection together with pre-sorting is not significantly better than using Bland's rule with pre-sorting. For the model M_5 the optimized rule with pre-sorting was even slightly worse than Bland's rule with a pre-sorting.

No improvement was achieved by our column selection rule. Using this rule for the model M_5 even resulted in an average effort greater than the use of Bland's rule with P_{max} . This seemingly paradox result can be easily explained: The application of the optimized rule for column selection led to a considerably higher average number of pivot steps. Our column selection rule needs on average 226 steps in the model M_5 , while the average number of pivot steps for the other

Table 5: Relative pivot steps of all experiments (in % normalized to Bland)

Setup	M_0	M_1	M_2	M_3	M_4	M_5
Bland	100	100	100	100	100	100
	(20,25)	(46,57)	(20,84)	(19,48)	(36,54)	(23,98)
BlandP	100	100	100	100	100	100
	(20,25)	(46,57)	(20,84)	(19,48)	(36,54)	(23,98)
BlandPRow	90,24	89,87	92,85	90,47	84,11	83,44
	(19,03)	(24,19)	(23,07)	(23,19)	(28,12)	(24,45)
BlandPSortMax	95,12	96,2	102,38	95,23	83,17	88,96
	(17,11)	(21,78)	(20,54)	(17,12)	(23,04)	(19,88)
BlandPSortSum	95,12	96,2	102,38	95,23	83,17	88,96
	(17,11)	(21,78)	(20,54)	(17,12)	(23,04)	(19,88)
BlandPSortFreq	95,12	96,2	102,38	95,23	83,17	88,96
	(17,11)	(21,78)	(20,54)	(17,12)	(23,04)	(19,88)
BlandPRowSortMax	92,68	93,67	97,61	88,88	78,5	86,89
	(17,03)	(22,22)	(22,48)	(18,59)	(23,59)	(21,22)
BlandPRowSortSum	92,68	93,67	97,61	88,88	78,5	86,89
	(17,03)	(22,22)	(22,48)	(18,59)	(23,59)	(21,22)
BlandPRowSortFreq	92,68	93,67	97,61	88,88	78,5	86,89
	(17,03)	(22,22)	(22,48)	(18,59)	(23,59)	(21,22)
BlandPColumnMaxRow	104,87	113,92	120,23	128,57	117,28	155,86
	(24,96)	(36,07)	(31,16)	(23,53)	(30,02)	(31,47)
BlandPColumnSumRow	104,87	113,92	120,23	128,57	117,28	155,86
	(24,96)	(36,07)	(31,16)	(23,53)	(30,02)	(31,47)
BlandPColumnFreqRow	104,87	113,92	120,23	128,57	117,28	155,86
	(24,96)	(36,07)	(31,16)	(23,53)	(30,02)	(31,47)

runs is 126 to 145. We attribute this to the fact that our rule ignores the values from the data matrix; only the protection level matrix is considered. Thus, a modification of the selection rule might lead, as in this case, to a greater number of pivot steps. We note that the lowest average effort does not necessarily correspond to the lowest average number of pivot steps. The application of Bland's rule together with pre-sorting led to a greater average number of pivot steps than the application of our row selection rule, while its average cost was lower. However, the cost of the protocol is commonly linear in the average number of pivot steps.

All three weight functions performed equally effective. We attribute this to the particular distribution of our initial protection levels derived from the criticality scores. Different risk assessments might result in different effectiveness of

the weight functions.

6 Conclusion

We introduced Secure Supply Chain Master Planning (SSCMP), an approach for centralized planning using secure computation. Traditional SCMP centrally computes the optimal production and transportation plan across a number of parties using Linear Programming. SSCMP can alleviate the perceived risk in SCMP due to data disclosure, since secure computation protects the confidentiality of the input values. We derived a methodology to assess these risks, the criticality score, in supply chains. Not surprisingly, the criticality scores vary for different input data.

We propose to modify the Linear Programming algorithm implemented in secure computation handling each data item at its risk level in order to increase the performance of SSCMP. We suggest three approaches modifying the Linear Programming algorithm to take advantage of the risk assessment results. We then assessed the effectiveness of our approaches in an experimental study. We conclude that a pre-sorting of the columns of the LP matrix according to their protection levels significantly and most effectively reduces the cost of the secure computation.

Future work is to apply the method to other algorithms for linear optimization, e.g. inner point methods, and to other supply chain optimization problems also adapting the risk assessment.

References

- [Atallah et al. 2009] Atallah, M., Blanton, M., Deshpande, V., Frikken, K., Li, J., Schwarz, L.: “Secure Collaborative Planning, Forecasting, and Replenishment”. Working Paper, Purdue University, 2005.
- [Atallah et al. 2004] Atallah, M., Bykova, M., Li, J., Frikken, K., Topkara, M.: “Private Collaborative Forecasting and Benchmarking”. Workshop on Privacy in the Electronic Society (WPES), 2004.
- [Atallah et al. 2003] Atallah, M., Elmongui, H., Deshpande, V., Schwarz, L.: “Secure Supply-Chain Protocols”. Proceedings of the IEEE International Conference on E-Commerce (CEC’03), 2003.
- [Beaver et al. 1990] Beaver, D., Micali, S., Rogaway, P.: “The round complexity of secure protocols”. Proceedings of 22nd STOC, 1990.
- [Bland 1977] Bland R.: “New finite pivoting rules for the simplex method”. Math. of Op. Res. 2, 1977.
- [Chen et al. 1999] Chen, F., Drezner, Z., Ryan, J., Simchi-Levi, D.: “The bullwhip-effect :managerial insights on the impact of forecasting and information on variability in a supply chain”, in: Taylor, S., Ganeshan, R., and Magazine, M. (Eds.), Quantitative Models for Supply Chain Management, Boston 1999.

- [Dantzig, Thapa 1997] Dantzig, G., Thapa, M.: "Linear Programming 1: Introduction". Springer-Verlag, 1997.
- [Damgård, Jurik 2001] Damgård, I., Jurik, M.: "A generalisation, a simplification and some applications of Paillier's probabilistic public-key system". International Workshop on Practice and Theory in Public Key Cryptography (PKC) 2001, 2001.
- [Dudek, Stadtler 2005] Dudek, G., Stadtler, H.: "Negotiation-based collaborative planning between supply chain partners", in: European Journal of Operational Research 163, 2005.
- [Fleischmann, Meyr 2003] Fleischmann, B., Meyr, H.: "Planning Hierarchy, Modeling and Advanced Planning Systems", in: De Kok, A. G., Graves, S. C. (Eds.): Supply Chain Management: Design, Coordination and Operation, Handbooks in Operations Research and Management Science, Vol. 11, Amsterdam 2003.
- [Goldreich et al. 1987] Goldreich, O., Micali, S., Wigderson, A.: "How to play any mental game". In Proceedings of the 19th annual ACM symposium on Theory of computing, 1987.
- [Holström et al. 2002] Holström, J., Främling, K., Tuomi, J., Krkkinen, M., Ala-Risku, T.: "Implementing collaboration process networks", in: The International Journal of Logistics Management 13(2), 2002.
- [Jayaraman, Pirkul 2001] Jayaraman, V., Pirkul, H.: "Planning and coordination of production and distribution facilities for multiple commodities", in: European Journal of Operational research, Vol. 133, No. 2, 2001.
- [Lee et al. 1997] Lee, H., Padmanabhan, V., Whang, S.: "Information distortion in a supply chain: the bullwhip effect", in: Management Science, Vol. 43, No. 4, 1997.
- [Li, Atallah 2006] Li, J., Atallah, M.: "Secure and Private Collaborative Linear Programming". 2nd International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom), 2006.
- [Min et al. 2005] Min, S., Roath, A., Daugherty, P., Genchev, S., Chen, H., Arndt, A., Glenn Richey, R.: "Supply chain collaboration: what's happening?", in: The International Journal of Logistics Management, Vol. 16, No. 2, 2005.
- [Pibernik, Sucky 2006] Pibernik, R., Sucky, E.: "Centralised and decentralised supply chain planning", in: International Journal of Integrated Supply Management 2(1/2), 2006.
- [Pibernik, Sucky 2007] Pibernik, R., Sucky, E.: "An approach to inter-domain master planning in supply chains", in: International Journal of Production Economics, Vol. 108, No. 1-2 2007.
- [Shamir 1979] Shamir, A.: "How to share a secret". Communications of the ACM, 1979.
- [Shapiro 2001] Shapiro, J.: "Modeling the Supply Chain". Pacific Grove, 2001.
- [Simpson, Erengüc 2001] Simpson, N., Erengüc, S.: "Modelling the order picking function in supply chain systems: formulation, experimentation, and insights", in: IIE Transactions 33(2), 2001.
- [Yao 1982] Yao, A.: "Protocols for secure computations". Proceedings of the 23rd IEEE Symposium on the Foundations of Computer Science (FOCS), IEEE, 1982.
- [Yu et al. 2001] Yu, Z., Yan, H., Cheng, T.: "Benefits of information sharing with supply chain management", in: Industrial Management and Data Systems, 2001.