

How the Estonian Tax and Customs Board Evaluated a Tax Fraud Detection System Based on Secure Multi-party Computation

Dan Bogdanov¹, Marko Jõemets¹, Sander Siim^{1,2}, and Meril Vaht¹

¹ Cybernetica, Mäealuse 2/1, Tallinn, Estonia**

{dan,markoj,sanders,meril}@cyber.ee

² Institute of Computer Science, University of Tartu, Liivi 2, Tartu, Estonia

Abstract. The Estonian Tax and Customs Board (MTA) has identified that Estonia is losing over 220 million euros a year due to avoidance of value-added tax (VAT). The parliament proposed legislation that makes companies declare their purchase and sales invoices for automated risk analysis and fraud detection. The law was vetoed by the Estonian President on the grounds of confidentiality breach and unnecessary burden to companies. In this paper, we report on our collaboration with MTA to build a tax fraud detection system prototype that uses secure multi-party computation (SMC) to remove the companies' concerns over confidentiality. We estimate that the prototype could process a month of Estonian VAT data in ten days running on 20 000 euros worth of hardware.

Keywords: tax fraud detection, risk analysis, secure multi-party computation, case study

1 Battling value-added tax fraud in a modern economy

Value-added tax (VAT) is a consumption tax on the value added to a sold product or service. To simplify, when a company sells a product, it will pay a tax on the difference of the sales price of the product and the price of materials and tools acquired to create it.

According to an estimation by the Estonian Tax and Customs Board (MTA) in 2013, Estonia has 72 000 registered taxable persons, a third of whom apply for a refund of overpaid VAT every month. Among them, there are about 9 700 enterprises with a suspicion of VAT fraud. The estimated total loss in unpaid VAT exceeds 220 million euros per year [1].

One of the main ways for avoiding VAT is to not declare sales to other companies, thus reducing the VAT liability. MTA detects such fraud by analyzing the financial records of the suspect company and its partners to determine the actual taxable sum. MTA estimated that using the existing process, it would

** This work has received funding from the Estonian Research Council through grant IUT27-1, ERDF through EXCS, and European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n 609611 (PRACTICE).

take 11 years to check a year’s worth of transaction data that has a fraud risk. The problem is amplified by the practice of forming shadow companies with no notable economic activities and the sole purpose of enabling VAT fraud.

The government reacted in 2013 by publishing updated drafts for the Value-Added Tax Act and the Accounting Act Amendment Act. These drafts describe a mandatory annex to the monthly VAT declaration form in the online tax information system deployed by MTA. According to the new law, companies must report transactions with each partner with whom the monthly sum of transactions exceeds 1000 euros. To keep automated accounting systems simpler, taxpayers can also declare all of their invoices with all partners.

For each transaction, the company has to report the registry code of the partner, date, identifier and value of the invoice. Depending on whether the transaction was a sale or purchase, it will also include the tax value, tax rate and the taxable supply for the current period of taxation. Once MTA has the VAT declarations for a month, it can match the declared sales and purchase invoices of companies with each other using the enterprise registry codes. It can then run risk analysis algorithms to find cases where a company has incorrectly declared transactions (or not declared them at all).

The government supported the acts and adjusted the budget with the prediction that the first year of activity will increase the amount of collected VAT by at least 27.5 million euros (Table 15, page 89 of the State Budget Strategy 2015—2018 [2]). Initially, the amendments would be enforced from July 1, 2014. However, opposition quickly arose from companies whose two main concerns were the administrative burden and the significant privacy risk.

The Estonian Traders Association claimed that, for large enterprises, the changes to accounting systems will require investments and time. They also conjectured that such data collection will not eliminate VAT fraud, but will force MTA employees to waste time on fixing human errors. The association was also concerned about the security of the “super database” of financial transactions. MTA has a significant employee turnover and a tax officer could copy the database to support his or her future business ambitions in the private sector.

In a controversial move, the President of Estonia blocked the legislation with the justification that “Burdening all businesses with additional costs and obligations and creating a database containing almost all of Estonia’s business secrets cannot be justified with a hypothetical, unproven conjecture that the tax hole would diminish [4].” The legislation was sent back to the parliament.

2 A solution based on secure multi-party computation

2.1 Requirements and the choice of the cryptographic platform

Examining the problem, we saw secure multi-party computation (SMC) as a solution. The companies are the input parties who have confidential data to protect. MTA is the result party, who wants to analyse the confidential inputs and learn the risk scores associated with companies. However, process-wise, MTA

may not need access to the detailed records of a company before the risk analysis has deemed that the company has a high risk score.

Therefore, we can design a system that collects VAT declaration annexes from the companies in a protected form and conducts the risk analysis while the transactions are in the encrypted domain. Only the risk scores will be published to the tax officer who can then request the detailed records for the at-risk companies. This protects the information and rights of the honest taxpayers, as their declaration annexes remain encrypted in the process.

We contacted MTA and explained our intentions of building a research prototype that would perform privacy-preserving VAT fraud detection. Fortunately, they were very supportive of our goals and ready to cooperate. We were able to work together with the developers of the new VAT declaration annex in MTA's online system. We had access to the architecture and system analysis documents and held regular meetings with the analysts and architects of the system to determine the following requirements.

1. **Privacy for companies.** To reassure the private sector in Estonia, the processing of VAT declarations must guarantee the secrecy of their contents.
2. **Data utility for investigators.** MTA must conduct automated risk analysis on VAT declarations and investigate suspicious companies in detail.
3. **Transparency.** To convince the companies in the security guarantees, they could retain some degree of control over the data and its processing.
4. **Performance.** MTA collects VAT declarations on a monthly basis and needs to complete the processing of one month of data before the next month.

2.2 Application and trust model

Figure 1 shows a proposed deployment model for the SMC-based tax fraud detection model designed based on the requirements. In the proposed system, a company would use a special tool that loads the XML file containing the invoices, applies secret sharing to each input value and uploads the shares of each value to the SMC servers. This tool can be audited by the company to ensure that good randomness is used and the correct servers are being connected to. If all shares are sent to parties with clearly non-collusive relations, the direct perception of security for data owners is greatly improved.

MTA and the Estonian Traders Association are good host candidates for the secure multi-party system, as both are motivated to keep the privacy of the data—MTA has a legal obligation and companies own the data. Both also have the capability to run IT systems—MTA will run one anyway and companies will participate if it provides them with better privacy. The latter is achieved as currently efficient SMC systems assume that all parties know the function being computed. This means that MTA will have to agree with companies on the kinds of analyses it wants to perform and all computing parties have to deploy them.

Some efficient SMC protocols also need a third party, so we need an organization that is independent from MTA as well as other companies and has the necessary motivation and resources. In Estonia, the Information Systems and




Fig. 1: Proposed deployment model for a secure tax fraud prototype

Registers Centre is a governmental agency under the Ministry of Justice tasked with maintaining security-critical registries so it is perfect for the job.

Once secret-shared inputs are stored by the computing parties, MTA can request parties to run the agreed-to risk analysis algorithms. These algorithms will run on secret-shared data and produce secret-shared results. The shares of these results will be published to the MTA. If the results show that a company has a risk, MTA needs to acquire transaction data from the secure multi-party system (in agreement with other computing parties) or the company directly.

3 Description of the prototype

3.1 Implementation platform

We chose the Sharemind secure multi-party computation system as our implementation platform [3]. While Sharemind supports protocols that are secure against an active adversary, we decided to use a passively secure protocol suite with three parties for its range of operations and performance. We solve the deficiencies of the passive model by deploying additional technical controls.

First, MTA can check the consistency of input data by comparing privately computed aggregate statistics of transactions to the public part of the VAT declaration. In our prototype, only MTA receives outputs from the computation, so actively tampering with the protocol to leak something from the outputs is not a feasible attack. This would need many queries, but MTA can not perform a multiple query attack undetected, because other parties involved in the computation can block them. The correctness of the computation can be checked with SMC auditing techniques.

The algorithms themselves are implemented in Sharemind’s programming language SecreC. We also developed a web-based interface for secret-sharing, uploading transaction files and running queries in a web browser.




Fig. 2: Example data flow within the secure risk analysis system

3.2 The privacy-preserving risk analysis process

Figure 2 shows the flow of private data during the risk analysis process. Algorithm 1 shows how sales and purchases are split into n tables that are aggregated by separate SMC processes. In the first part of the aggregation, n parallel SMC processes run Algorithm 2 on the sales and (with minor changes) purchase transactions. The results are gathered in the `sales_aggr` and `purchases_aggr` tables. The results of parallel computations are combined with Algorithm 3.

In the prototype, we implemented three risk analysis algorithms. First, to find high differences between partners' transactions, the values of a company's sales invoices per partner are added together and totals of the partners' purchase invoices connected to this company are added together. If the difference between sales and purchases is negative, the purchasing partner has the risk. Algorithm 4 shows how we calculate the risk based on aggregation results.

For the second risk, we find the proportion of the sum of all taxpayers' declared sales invoices from the sales total declared in the declaration main part (this is calculated as Pr_X in Algorithm 2). If the percentage is smaller than some estimated fixed amount (for example 30%) it is counted as a risk, implying that some invoices have been left out of the declaration.

Finally, our prototype performs one pass of risk propagation so that if a partner of a company has a risk, we also mark the company as potentially risky.

Algorithm 1: Processing a secret-shared declaration at a computing party.

Data: Secret-shared VAT declaration of company with registry code X
Result: The declaration is stored in tables $sales_i$ and $purchases_i$

- 1 Choose a joint random value $i \in 1, 2, \dots, n$ with other computing parties
- 2 **foreach** sales transaction in declaration **do**
- 3 $Y \leftarrow$ transaction partner’s registry code
- 4 Save $[X, Y, \text{transaction amount}]$ in table $sales_i$
- 5 **foreach** purchase transaction in declaration **do**
- 6 $Y \leftarrow$ transaction partner’s registry code
- 7 Save $[X, Y, \text{transaction amount}]$ in table $purchases_i$
- 8 Add X to the aggregation queue of aggregator i
- 9 **return**

Algorithm 2: Transaction data aggregation at a computing party.

Data: Aggregator index i , table $sales_i$
Result: Aggregated sales data stored in table $sales_aggr$

- 1 **foreach** non-aggregated X in the queue of aggregator i **do**
- 2 $D_X \leftarrow$ transaction data of X from $sales_i$
- 3 $Pr_X \leftarrow \frac{\text{sum}(\text{sales transactions in } D_X)}{\text{sum of sales transactions declared by company } X}$
- 4 **foreach** transaction partner Y appearing in D_X **do**
- 5 $S_{X,Y} \leftarrow \text{sum}(\text{sales transactions with } Y \text{ in } D_X)$
- 6 Save $[X, Y, S_{X,Y}, Pr_X]$ in $sales_aggr$
- 7 Mark X as aggregated and add it to the finalization queue
- 8 **return**

The analysis and implementation took 3.5 man-months of work from developers with some experience with Sharemind (but no special cryptographic training). Their main challenge was to find a suitable and efficient algorithm.

3.3 Performance of the prototype

We generated test datasets with realistic distributions to measure the performance of our prototype on three servers with 3 GHz 12-core processors connected to a 1 Gb local network. Figure 3a shows the total running times of aggregation with up to 8 parallel tasks. We see that parallel tasks improve the efficiency by a constant factor. This is explained by Figure 3b that shows the running time of different phases in the computation when using 4 aggregators. The performance of the finalization phase does not improve with parallel aggregations.

We believe we can speed up the final combination of aggregated tables significantly by using more efficient merging techniques. The bottleneck of our prototype seems to be the network channels, as the CPUs are not fully used. Thus, an increase in network bandwidth will also improve performance.

Algorithm 3: Aggregation finalization of sales at a computing party.

Data: Non-finalized aggregation tables `sales_aggr`, `purchases_aggr`

Result: Finalized aggregation table `sales_summary`

```
1  $Pr_X \leftarrow \text{sales\_aggr}$ 
2 foreach  $X$  in the finalization queue do
3   foreach transaction partner  $Y$  do
4      $S_{X,Y} \leftarrow \text{sales\_aggr}$  // sum( $X$ 's sales to  $Y$ )
5     if  $Y$  in purchases_aggr then
6        $P_{Y,X} \leftarrow \text{purchases\_aggr}$  // sum( $Y$ 's purchases from  $X$ )
7     else
8        $P_{Y,X} \leftarrow 0$ 
9     Save  $[X, Y, S_{X,Y}, P_{Y,X}, Pr_X]$  in sales_summary
10 return
```

Algorithm 4: Discrepancies between declared sales and purchases.

Data: Table `sales_summary`

Result: Registry codes of companies with risk 1 confirmed.

```
1 RiskCompanies  $\leftarrow \emptyset$ 
2 foreach pair of transaction partners  $(X, Y)$  in sales_summary do
3    $S_{X,Y}, P_{Y,X} \leftarrow \text{sales\_summary}$ 
4   if  $S_{X,Y} < P_{Y,X}$  then
5     | add  $Y$  to RiskCompanies
6 return RiskCompanies
```

4 Evaluation by the Estonian Tax and Customs Board

We gave a presentation and a technical report to MTA's management and experts from the risk analysis and IT departments. We focused on differences in processes, architecture and deployment that result from using SMC.

MTA representatives understood the security guarantees provided by SMC and accepted them as superior to what can be achieved with current technologies. However, the risk analysts were concerned with the required transparency. Today, MTA can perform risk analyses autonomously so that unauthorized parties have no knowledge of the kind of algorithms that are used. SMC would change this and MTA would have to agree on the algorithms with other hosts.

We argued that transparency will also improve the acceptance of the system. In response, the Director General admitted that his philosophy is to develop taxation so that taxpayers feel more responsibility on the grassroots level and consider paying taxes to be a social obligation. He agreed, that the SMC-based solution we described is a step in this direction, but stated that significant change would be needed in the processes of risk analysis to enable the sharing of related algorithms with other parties. Alternatively, SMC technology should become significantly more efficient at hiding the algorithm being evaluated or we should find ways to hide the class of algorithms used in risk analysis.




Fig. 3: Running times for secure tax fraud detection

Based on the calculations from MTA, 80 000 companies will upload 50 million economic transactions every month. We estimate that our prototype can process one month of Estonian economy in ten days, using about 20 000 euros worth of hardware. This was met with some concern, as today, MTA processes VAT returns in three days. The hardware cost, however, was not of much concern. With algorithmic improvements and clever hardware usage we can make the prototype an order of magnitude faster and make the running time sufficiently low for practical use.

The President of Estonia did not block an updated version of the tax legislation that granted a longer transition time to companies and solved other concerns. Thus, MTA continued to develop a non-encrypted version of the VAT declaration system for deployment in late 2014. However, MTA agreed to consider SMC as a technology for confidential data collection and analysis in future application, inspired by our prediction that the cost of deploying SMC will be further reduced in the coming years.

References

1. The transaction information system for 1000-euro purchases will be completed in November. *Ärileht*, Oct 22, 2013. <http://arileht.delfi.ee/news/uudised/mta-1000-euroste-arvete-tehinguinfo-susteem-saab-valmis-novembris.d?id=66955998> (in Estonian). Last accessed: Sept 5, 2014.
2. State Budget Strategy of Estonia 2015-2018. Available: <http://www.fin.ee/doc.php?110953>. Last accessed: Sept 5, 2014.
3. The Sharemind secure database and application server. <http://sharemind.cyber.ee/>. Last accessed: Sept 5, 2014.
4. Ilves Blocks Amendment for Sweeping Disclosures in Tax Filing. *National Public Broadcasting News*, Dec 19, 2013. Available: <http://news.err.ee/v/politics/5b358dbd-8836-43ca-992c-973d206a3ec6>. Last accessed: Sept 5, 2014.