

# FINAL REPORT

## 1 General Information

DFG reference number: ME 3704/5-1 and MU 4467/1-1

Project number: 419340256

Project title: Privacy-preserving Kidney Donor Exchange

Names of the applicants: Prof. Dr.-Ing. Ulrike Meyer  
Dr. Anja Mühlfeld

Official addresses: Mies-van-der-Rohe Strasse 15, 52074 Aachen  
Pauwelsstraße 30, 52074 Aachen

Name(s) of the co-applicants:

Name of the cooperation partners: Prof. Dr.-Ing. Susanne Wetzel

Reporting period (entire funding period): 08.2020 - 11.2024

## 2 Summary

Kidney exchange enables patients with medically incompatible living donors to still receive a compatible kidney transplant, by finding matches between multiple patient-donor pairs such that these can exchange their donors among each other. The existing kidney exchange systems today face significant security challenges, as they neither prevent manipulation of the exchange computation nor adequately protect the sensitive data of patients and donors.

In this project, we devised a new model of a privacy-preserving kidney exchange system that protects against these security issues. Our model follows a decentralized approach in that the computation of exchanges is distributed among a set of computing peers. At discrete points in time the computing peers then execute a secure multi-party computation (SMPC) protocol among each other in order to compute a set of exchanges among the patients' and donors' that are registered with the system. This setup guarantees that a computing peer is neither able to learn any information on the input data of the patients and donors nor to manipulate the computation of exchanges. We developed, implemented, and evaluated five different SMPC protocols for kidney exchange, using different algorithmic approaches. All of them are able to cover all desirable functional requirements discussed with medical transplant experts w.r.t. the exchange structures, the matching criteria, and the prioritization criteria supported by them.

To evaluate the impact of the run time overhead induced by SMPC and the influence of the different algorithmic approaches on the number of transplants that can be achieved over time, we developed a simulation framework that accounts for the many different parameters that influence the performance of a kidney exchange system (e.g., the interval at which new patients are registered). We used a real-world data set, which we obtained from the United Network for Organ Sharing (UNOS) in the USA, to simulate both the performance of our model of a privacy-preserving kidney exchange system as well as the performance of the non-privacy-preserving reference model that mimics the existing centralized kidney exchange systems. Based on these simulations, we were able to show that our approach only induces a small and sometimes even negligible impact on the number of found transplants over time for most parameter combinations that are found in practice. At the same time it provides for significantly stronger security guarantees compared to the existing centralized systems.

This project was carried out between RWTH Aachen University, the University Hospital of RWTH Aachen University, and Stevens Institute of Technology. It was independently funded by the DFG (project number 419340256) and the NSF (grant CCF-1646999).

## 2.1 German

In diesem Projekt haben wir einen privatsphärewahrenden Ansatz für Nierenspendertausch entwickelt. Die Grundidee von Nierenspendertausch ist es Patienten, welche nur einen medizinisch inkompatiblen Lebenspender gefunden haben, trotzdem zu ermöglichen, eine kompatible Nierenspende zu erhalten. Dazu wird zwischen mehreren Paaren von Patient und inkompatiblen Spender nach möglichen Konstellationen gesucht, in denen die Patienten ihre Spender untereinander tauschen können. Zentralisierte Systeme für Nierenspendertausch sind bereits in vielen Ländern etabliert. Sie weisen jedoch schwerwiegende Sicherheitsprobleme auf, da sie weder die Daten der involvierten Patienten und Spender ausreichend schützen, noch resistent gegen die Manipulation der Berechnung von Tauschkonstellationen sind.

In diesem Projekt haben wir ein Modell eines privatsphärewahrenden Systems entwickelt, welches diese Sicherheitsprobleme behebt. Unser Modell folgt einem dezentralen Ansatz, in dem die Berechnung der Tauschkonstellationen auf mehrere Berechnungsparteien verteilt wird. Dieser garantiert, dass eine Berechnungspartei weder die Eingabedaten der Patienten und Spender erfährt, noch die Berechnung manipulieren kann. Das Kernelement unseres Modells ist ein Protokoll zur sicheren Mehrparteienberechnung (SMPC), welches zwischen den Berechnungsparteien ausgeführt wird, um Tauschkonstellationen zwischen den registrierten Patient-Spender Paaren zu ermitteln. Die fünf von uns entwickelten SMPC Protokolle, verfolgen jeweils einen anderen algorithmischen Ansatz.

Um den Einfluss der erhöhten Laufzeit sowie der unterschiedlichen algorithmischen Ansätze unserer Protokolle auf die Anzahl an Transplantationen über Zeit evaluieren zu können, haben wir zudem ein Simulationsframework entwickelt. Dieses simuliert die Anzahl der über die Zeit gefundenen möglichen Transplantation in Abhängigkeit verschiedener Parameter basierend auf echten medizinischen Daten (bereitgestellt vom United Network for Organ Sharing (UNOS) in den USA). Diese Simulation erlaubt den Vergleich die Performanz unseres Modells mit der der zentralisierter Systeme. Dabei konnten wir zeigen, dass unser Modell für die meisten Parameterkombinationen, die in der Praxis auftreten, nur einen geringen (manchmal sogar vernachlässigbaren) Einfluss auf die Anzahl der Transplantationen über Zeit hat. Gleichermaßen bietet unser Modell jedoch erheblich stärkere Sicherheitsgarantien.

Das Projekt wurde in Kooperation zwischen der RWTH Aachen, dem Uniklinikum der RWTH Aachen und dem Stevens Institute of Technology durchgeführt, und von der DFG (Projektnummer 419340256) und der NSF (Grant CCF-1646999) finanziert.

### 3 Progress Report

**Background and Main Objectives:** If a patient is diagnosed with end-stage kidney disease, the only possible options of treatment are dialysis or a kidney transplant. While dialysis is costly and heavily impacts the patients' quality of life, the waiting lists for kidney organs are very long in most countries. For example, at the end of 2023 there were 6,196 patients on the deceased donor waiting list for a kidney transplant in Germany [? ]. An alternative to a post-mortem kidney donation is a living donation, typically from a friend or relative of the patient. However, this is sometimes impeded by the patient and the potential donor being medically incompatible.

Kidney exchange provides a solution for this problem. The idea is to consider multiple pairs of patient and incompatible living donor (so-called *patient-donor pairs*) and identify constellations in which the pairs can exchange their kidney donors among each other. The simplest form of this is a crossover exchange, where only two pairs exchange their donors among each other. Many existing kidney exchange programs today, however, support exchange cycles involving up to three pairs as well, with each donor donating to the patient of the succeeding pair in the cycle. More recently, also so-called exchange chains have become common. These are initiated by an altruistic donor who donates a kidney without expecting anything in return. This donor donates to the patient of the first pair in the chain. The chain then continues with each donor donating to the patient of the succeeding pair, until the last donor either contributes to the waiting list or becomes a bridge donor, i.e., a future altruistic donor. The challenge of finding an optimal set of exchange cycles and chains, based on a pre-defined set of criteria, is known as the kidney exchange problem (KEP) [? ].

Many countries already have nationwide kidney exchange systems in place that allow patients to find exchange partners [? ? ]. These systems use a centralized approach, where transplant centers register their patient-donor pairs and altruistic donors with a central platform, which then periodically solves the KEP among all registered patient-donor pairs and altruistic donors and afterwards informs the transplant centers of the computed exchange partners. However, these centralized kidney exchange systems face severe security risks, such as susceptibility to manipulation and data breaches. To address these issues, our project proposed a distributed, privacy-preserving approach to kidney exchange with six main objectives: (1) The system was to use a distributed infrastructure, removing the need for a central authority or database to store patient and donor data. (2) Patient and donor privacy was to be ensured by keeping all data encrypted during the entire exchange computation. (3) The system was to pre-

vent manipulation of the exchange computation as well as manipulation of the data of any donor or patient to the advantage of a specific patient. (4) Fairness in the exchange computation was to be guaranteed, with the selection strictly adhering to medical, ethical, and legal prioritization criteria. (5) The system's fairness, correctness, and security properties was to be formally provable to allow for checking them prior to deployment. (6) Cryptographic mechanisms for auditing were to be implemented based on criteria defined in collaboration with medical experts.

This decentralized approach was to be implemented based on Secure Multi-Party Computation (SMPC), which is a cryptographic technique that allows a set of parties to jointly compute a functionality in a distributed fashion such that each party only learns its private input and output and what can be deduced from both. The security of such an SMPC protocol in the presence of an adversary, who is able to corrupt a subset of the parties, is then typically either proven in the semi-honest or in the malicious model. The former assumes the corrupted parties to strictly follow the protocol specification and gather all information they learn during the protocol execution. Based on this information, the adversary then tries to infer any information on the other parties' private inputs and outputs. In the significantly stronger malicious security model, the corrupted parties may additionally arbitrarily deviate from the protocol specification. Originally, the plan was to devise the protocols with security in the semi-honest model mainly due to the performance impact induced by considering the malicious model.

The work on the project was carried out jointly between the Research Group IT-Security at RWTH and Anja Mühlfeld at the University Hospital of the RWTH on the German side, and the group of Susanne Wetzel from Stevens Institute of Technology in Hoboken New Jersey. The work plan of the joint proposal was therefore structured in seven main work packages as illustrated in Figure 1.

**Deviations from the Original Concept:** By and large, we followed the anticipated work plan with a few exceptions which mainly resulted from our requirements analysis with medical experts on the one hand and novel advances in the context of SMPC on the other hand.

We fully covered WP1, WP2, WP3, WP4, WP6, and WP7. However, we kept our study of protocol specifically designed to support cross-over exchanges alone short.

This is due to the fact that while at the time the proposal was written and in the first part of the project runtime, we did not anticipate the German Transplant Law to be changed in the near future, leaving Germany with cross-over transplant as only viable donor exchange. However, much to our surprise the representative of the German Ministry for Health we invited to the

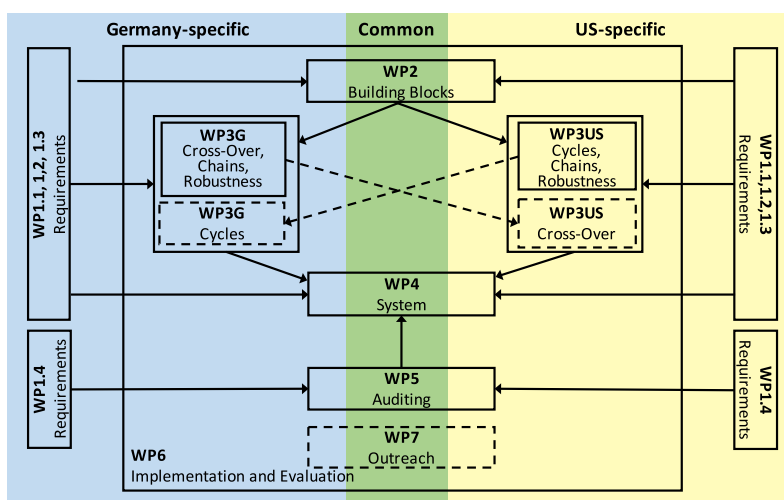


Figure 1: Project Overview - Cooperation and Distribution of Work.

first Workshop we organized during the project informed us of a change in the German government's view towards changing the German Transplant Law (TPG) to support kidney donor exchange. We therefore decided to no longer focus on the special case of crossover exchange for Germany but instead to immediately start working on larger exchange cycles and chains in a combined effort on the German and the US side. This resulted in work packages WP3G and WP3US merging into a single work package WP3.

The second line of change stems from recent advances in SMPC, which made protocols with security in the malicious model much more efficient [? ], allowing us to design the entire system with malicious security rather than semi-honest security as originally planned. This shift in design removed the need for additional auditing of the computation and thus made the cryptographic mechanisms we originally planned to develop in WP5 obsolete. The same applies to the zero-knowledge proofs that we planned to develop in Research Direction 2.3. At the same time this change allowed us to explore more than one algorithmic approach in detail and study not only protocols that exactly solve the kidney exchange problem but also ones that find an approximate solution to it.

**Project-Specific Results and Findings:** In WP1 we conducted an initial requirements analysis with Dr. Mühlfeld on potential matching and prioritization criteria as well as exchange structures to be supported. We further discussed these requirements in the two workshops we organized as part of WP7 (see below). Based on these discussions we developed two major privacy-preserving building blocks as part of WP2. First, a privacy-preserving protocol that effi-

ciently computes the medical compatibility between a patient and a donor based on their blood types, as well as the donor's Human Leukocyte Antigens (HLA), and the patient's antibodies against these. Second, a generic privacy-preserving building block for computing a prioritization score between a patient and a donor. As our requirements analysis revealed that the prioritization criteria used by different systems vary widely in practice. We therefore developed this building block such that it can be configured to the needs of any particular kidney exchange system. We published these two building blocks in [C, H].

In WP3, we then devised multiple SMPC protocols for solving the KEP. Following the original research plan, we started with the development of a protocol for the special case that only crossover exchanges are allowed [F]. However, as both medical experts and the German Ministry of Health were in favour of aiming for a change of the German transplant law that would allow for longer cyclic exchanges, we did not further explore this special case and instead focused on the development of efficient protocols for any restriction on the maximum cycle or chain length. Since the KEP is an NP-complete problem [? ], this is a highly non-trivial task already in the non-privacy-preserving case. In addition, the best performing protocols in the conventional setting does not necessarily lead to a well-performing privacy-preserving protocol, as the latter requires changes to the algorithmic approach that makes the protocol's runtime independent of the input data and thus often destroys the algorithmic advantage of a well-performing conventional algorithm. Therefore, we devised multiple SMPC protocols for solving the KEP, each following a different algorithmic approach such that all categories of algorithmic approaches that are used in existing non-privacy-preserving kidney exchange systems are covered and comparatively analyzed them.

Overall, we devised four different SMPC protocols for solving the KEP with cycles and chains. All protocols support prioritization, exchange cycles and chains, and are secure in the malicious model. As such these protocols support all exchange structures that are also supported in any of the existing kidney exchange systems. The choice of any specific protocol to be used in our model of a privacy-preserving kidney exchange system then always forms a trade-off between run time performance, quality of the solution, and data obliviousness (i.e., whether the flow of execution is independent from the input or not). We published the specifications of these protocols in [A, B, C, E, H] and won the **Best Paper Award** for [B] at the A-ranked Asia Conference on Computer and Communications Security in 2024.

In order to assess the performance of our protocols in practice, we implemented them and

extensively evaluated their run time performance. As we developed our SMPC protocols based on secret-sharing with security in the malicious model, it did not make sense to follow our original plan to use our existing Java library for the implementation. Instead, we switched to the state-of-the-art SMPC benchmarking framework MP-SPDZ [10]. This made our implementation independent from the specific underlying SMPC primitive. Thus, we were able to extensively evaluate the performance trade off between different numbers of computing peers, different security models, and different cryptographic primitives. Our main result w.r.t. the run time evaluation was then that the most efficient run times are achieved by our approximation protocol from [A, B]. While this protocol only computes an approximate solution for the KEP, it achieves run times below 4 hours for up to 200 patient-donor pairs and altruistic donors, even if security in the malicious model and the WAN setting with a latency of 20ms and a bandwidth restriction of 1Gbps are considered. This shows that our protocol even scales for the numbers of patients and donors that are found in large kidney exchange systems such as the one operated by the United Network for Organ Sharing (UNOS) in the USA. Note that run times of four hours are still feasible for kidney exchange in practice since the existing kidney exchange systems execute match runs at most at a daily basis [11].

Aside from the protocol run time, we also empirically evaluated the approximation quality obtained by our approximation protocol. To this end, we used a large real-world data set that we acquired from UNOS in order to determine the approximation quality for a real-world scenario. We observed that on average the approximation quality is always at least around 80%.

Based on our SMPC protocols, in WP4 we then devised our model of a privacy-preserving kidney exchange system. To this end, we thoroughly analyzed the model implemented by existing non-privacy-preserving kidney exchange systems. We identified two main security issues stemming from their centralized approach. First, the central platform operator alone is responsible for the exchange computation, allowing the operator or any attacker who corrupts the operator to manipulate the exchange computation, e.g., to favor a specific patient. Second, all data of the patients and donors is stored in plaintext at the central platform, leading to severe consequences of any data breach, be it forced or accidental.

Considering these security issues, we developed the first model of a privacy-preserving kidney exchange system. Our model follows a decentralized approach based on SMPC in the client-server model, where a set of input peers (clients) use secret-sharing to distribute their private input among a set of computing peers (servers), who execute an SMPC protocol for

computing the desired functionality on the private inputs among each other. Afterwards, they provide each input peer with its output. During the whole process, it is guaranteed that the computing peers do not learn anything about the private input and output of the input peers. In our model of a privacy-preserving kidney exchange system, the input peers correspond to the transplant centers which the patient-donor pairs and altruistic donors are associated with, whereas the computing peers can be hosted by any set of independent institutions such as universities, large transplant centers, or governmental institutions. In order to register their patient-donor pairs and altruistic donors, the transplant centers provide the sensitive input data of their associated patient-donor pairs and altruistic donors to the computing peers using secret-sharing. Then, similar to the existing kidney exchange systems, at discrete points in time, the computing peers execute so-called match runs, where they use one of our SMPC protocols for solving the KEP to determine a set of potential exchange partners for all registered patient-donor pairs and altruistic donors. Since our SMPC protocols are secure in the malicious model, our model of a kidney exchange system achieves all of the six main objectives that we aimed for in this project, even without the need for additional auditing of the exchange computation.

In order to determine the performance impact of the increased run time overhead as well as the approximative nature of our SMPC protocols when used as part of our model of a privacy-preserving kidney exchange system, we developed a novel simulation framework in Java based on the discrete event simulation framework DESMO-J. The framework supports both our model of a privacy-preserving kidney exchange system and the model implemented by existing centralized kidney exchange systems. It allows for the simulation of these systems over time based on our real-world data set from UNOS and it supports the specification of all the different parameters that influence the performance of a kidney exchange system over time. This allows us to simulate the performance of kidney exchange systems for the different design choices, requirements, and population characteristics in different countries.

The goal of our simulations was then to determine in how far our model of a privacy-preserving kidney exchange system impacts the number of transplants that can be achieved over time compared to the non-privacy-preserving model that is implemented by existing kidney exchange systems. As the existing kidney exchange systems vary widely w.r.t. the choice of different parameter values, such as the time interval in between match runs, we executed simulations for a wide variety of parameter values. Thereby, we were able to determine the performance of our model of a privacy-preserving kidney exchange system for all the different

combinations of parameter values that are found in existing kidney exchange systems. Our main result was that the best performance in the privacy-preserving case is achieved when using our approximation protocol. Despite its approximative nature, we were able to show that for most parameter combinations that are found in existing kidney exchange systems, the impact of our privacy-preserving approach on the number of transplants over time is very small, in some cases even negligible. This shows that our model of a privacy-preserving kidney exchange system provides for all the necessary means to establish a privacy-preserving kidney exchange system in practice at a very small cost in terms of number of transplants over time while providing for significantly stronger security guarantees than existing kidney exchange systems. We published these simulation results in [A, B].

As part of the in parallel running outreach work package (WP7), we presented our privacy-preserving approach to several actors from transplant medicine and government. In particular, we organized two workshops (one in Aachen and in Köln), where we invited nephrologists, surgeons, patients, computer scientists, and representatives from the German Ministry of Health to discuss the requirements on a kidney exchange system in Germany, as well as the potential establishment of a kidney exchange system based on our privacy-preserving model. As of today, there is a law under review towards changing the TPG [? ]. We contributed to this process not only by the discussions in the workshops that we organized but also by reviewing a preliminary version of the law and providing feedback to the German Ministry of Health.

## References

### 4 Published Project Results

#### 4.1 Publications with scientific quality assurance

- [A] M. Breuer, U. Meyer, and S. Wetzel. Efficient Integration of Exchange Chains in Privacy-Preserving Kidney Exchange. In *International Conference on Privacy, Security, and Trust (PST)*. IEEE, 2024
- [B] M. Breuer, U. Meyer, and S. Wetzel. Efficient Privacy-Preserving Approximation of the Kidney Exchange Problem. In *ASIA Conference on Computer and Communications Security (ASIA CCS)*. ACM, 2024. **Best Paper Award** <https://doi.org/10.1145/3634737.3645015>
- [C] M. Breuer, P. Hein, L. Pompe, U. Meyer, and S. Wetzel. Prioritization and Exchange Chains in Privacy-Preserving Kidney Exchange. *Journal of Computer Security*, 32(4), 2024. <https://doi.org/10.3233/JCS-230012>

- [D] A. Brüggemann, M. Breuer, A. Klinger, Thomas S. and U. Meyer. Secure Maximum Weight Matching Approximation on General Graphs. In *Workshop on Privacy in the Electronic Society (WPES)*. ACM, 2022. <https://doi.org/10.1145/3559613.3563209>
- [E] M. Breuer, P. Hein, L. Pompe, B. Temme, U. Meyer, and S. Wetzel. Solving the Kidney Exchange Problem Using Privacy-Preserving Integer Programming. In *International Conference on Privacy, Security and Trust (PST)*. IEEE, 2022. <https://doi.org/10.1109/PST55820.2022.9851968>
- [F] M. Breuer, U. Meyer, and S. Wetzel. Privacy-Preserving Maximum Matching on General Graphs and its Application to Enable Privacy-Preserving Kidney Exchange. In *Conference on Data and Application Security and Privacy (CODASPY)*. ACM, 2022. <https://doi.org/10.1145/3508398.3511509>
- [G] M. Breuer, U. Meyer, and S. Wetzel. Introducing a Framework to Enable Anonymous Secure Multi-Party Computation in Practice. In *International Conference on Privacy, Security and Trust (PST)*. IEEE, 2021. <https://doi.org/10.1109/PST52912.2021.9647793>
- [H] M. Breuer, U. Meyer, S. Wetzel, and A. Mühlfeld. A Privacy-Preserving Protocol for the Kidney Exchange Problem. In *Workshop on Privacy in the Electronic Society (WPES)*. ACM, 2020. <https://doi.org/10.1145/3411497.3420213>

## 4.2 Other publications and published results

- [I] M. Breuer. *Privacy-Preserving Kidney Exchange*. PhD thesis, RWTH Aachen University, 2024
- [J] M. Breuer, U. Meyer, S. Wetzel, J. Flöge, and A. Mühlfeld. Wenn der eigene Spender nicht passt ... Nierentransplantation als Cross-over-Lebendniere spende. *Nieren und Hochdruckkrankheiten*, 51(9), 2022. <https://doi.org/10.5414/NHX02252>