

Practical Secure Evaluation of Semi-private Functions

Annika Paus, Ahmad-Reza Sadeghi*, and Thomas Schneider**

Horst Görtz Institute for IT-Security, Ruhr-University Bochum, Germany
{annika.paus,ahmad.sadeghi,thomas.schneider}@trust.rub.de

Abstract. Two-party Secure Function Evaluation (SFE) is a very useful cryptographic tool which allows two parties to evaluate a function known to both parties on their private (secret) inputs. Some applications with sophisticated privacy needs require the function to be known only to one party and kept private (hidden) from the other one. However, existing solutions for SFE of private functions (PF-SFE) deploy Universal Circuits (UC) and are still very inefficient in practice.

In this paper we bridge the gap between SFE and PF-SFE with SFE of what we call *semi-private functions* (SPF-SFE), i.e., one function out of a given class of functions is evaluated without revealing which one.

We present a general framework for SPF-SFE allowing a fine-grained trade-off and tuning between SFE and PF-SFE covering both extremes. In our framework, semi-private functions can be composed from several privately programmable blocks (PPB) which can be programmed with one function out of a class of functions. The framework allows efficient and secure embedding of constants into the resulting circuit to improve performance. To show practicability of the framework we have implemented a compiler for SPF-SFE based on the Fairplay SFE framework.

SPF-SFE is sufficient for many practically relevant privacy-preserving applications, such as privacy-preserving credit checking which can be implemented with our framework and compiler as described in the paper.

Keywords: SFE of semi-private functions, Yao's protocol, topology, optimization, compiler, privacy.

1 Introduction

Two-party Secure Function Evaluation (SFE) is an important and wide area of cryptographic research (see, e.g., [18,10,13,1,11,7,12]). It allows two parties to securely evaluate a common function on their private inputs without involving a trusted third party. The function is represented as a boolean circuit and evaluated based on a garbled version of the circuit which is created by one party (constructor Bob) and evaluated by the other party (evaluator Alice). Usually SFE hides the intermediate results but - as the function is known to both parties - not the structure (topology) of the function.

* The author was supported by the European Union under FP6 project SPEED.

** The author was supported by the European Union under FP7 project CACE.

In practice, however, a variety of business models require privacy properties beyond the secrecy of parties' input data to additionally keep the evaluated function private. The underlying business motivations vary from commercial incentives (e.g., protection of intellectual property) to pure security requirements to reduce the probability of credential forgery or to make insider attacks obsolete. Typical use cases are client-server applications where a user Alice inputs her private data x (hidden to Bob), the server Bob inputs his private function f (hidden to Alice), and the protocol outputs $f(x)$ to both parties such that neither party gain any information about the other party's input. Prominent examples are privacy-preserving trust negotiation schemes [3,6,4], credit checking [5], or data classification using neural networks [16].

To allow SFE of a private function, called PF-SFE [8], a universal circuit (UC) [17,8,16] is evaluated that simulates the function, and entirely hides the structure of their circuit representation. UCs require a huge overhead of $O(k \log k)$ [17], $O(k \log^2 k)$ [8], respectively $O(k^2)$ [16] additional gates, where k is the number of gates of the simulated circuit.

Fairplay [13], a state-of-the art implementation of SFE, can evaluate functions consisting of millions of gates whereas in FairplayPF [8], a recent implementation for PF-SFE, functions are restricted to a few thousand gates only due to the huge overhead for evaluating UC. Hence, a better trade-off between maximal performance (SFE) and maximal privacy of the evaluated function (PF-SFE) is desired. For many practically relevant applications (e.g., those mentioned above) it is sufficient that functions are only partly private, what we call *semi-private functions* (SPF). Basically, these applications reflect the following scenario: A user Alice has private data x , and a service provider Bob has a semi-private function $f \in \mathcal{F}$ as input, where \mathcal{F} represents a given class of functions. At the end of the protocol, Alice obtains $f(x)$ but not which specific f was evaluated and Bob obtains no information on x . This problem, called *secure function evaluation of semi-private functions* (SPF-SFE), can be reduced to Yao's protocol where circuit's topology is revealed to the evaluator but the functionality of the gates is hidden. Evaluator sees the circuit topology but can only guess which functionality each part of the circuit might evaluate. We concentrate on *relaxed-security* model, i.e., security against malicious evaluator Alice and semi-honest (honest-but-curious) constructor Bob. This model is widely used in current cryptographic literature [14,2,9] and well-justified in many practical applications where performance is crucial and constructor Bob can be assumed to behave semi-honestly by means of legal contracts or possible loss of reputation.

While SPF-SFE based on Yao's protocol has been proposed as building block in many applications (e.g., [3,5,6,4,16]), we give the first unified theory for SPF-SFE. Extending and improving previously known techniques we present a general theoretical framework for SPF-SFE together with a language and tools to specify and automatically generate SPF-SFE protocols for practical applications.

Related Work. The idea of constructing circuits for a special class of functions and evaluating them efficiently with Yao's protocol in the relaxed-security model have been used in several sub-protocols [3,5,6,4,8,16]. Frikken et al. call the

respective building blocks oblivious gates/circuits where evaluator does not know the function that each gate/circuit computes. However, they only mention the existence of several useful topologies like binary trees, comparison circuits, or universal circuits together with their asymptotic size, but do not give explicit constructions. We extend their basic ideas into a generic framework and provide a wide class of functional blocks, each with a concrete efficient implementation (topology, programming, and exact size), that can be arbitrarily combined to represent semi-private functions in many practical applications.

Existing frameworks for secure computation based on Yao’s protocol are the Fairplay SFE system [13] with a proposed extension to the malicious model [12] and another extension to private functions with UCs (PF-SFE), called FairplayPF [8]. The Fairplay compiler includes an optimizer that optimizes on the basis of the high-level Secure Function Description Language (SFDL) using peek-hole optimization, duplicate code removal, and dead code elimination. In contrast to this, our proposed optimization algorithm for constant inputs optimizes on the lower abstraction level of circuits and can also be applied to further optimize the output of circuits generated with the Fairplay compiler.

Our Contribution and Outline. We propose a general framework together with a compiler for efficient secure function evaluation of semi-private functions (SPF-SFE) in the relaxed-security model.

In §2 we describe how common SFE can be extended with building blocks that we call *Privately Programmable Blocks* (PPB) to allow practical secure evaluation of semi-private functions (SPF-SFE). A privately programmable block (§4) consists of a fixed topology of several programmable gates (with a small number of inputs) and can be programmed to evaluate different functions out of a class of functions. The evaluator learns how the blocks are connected (topology) but *not* with which of the functions of their corresponding class of functions the blocks are programmed. Hence parts of the function are hidden from the evaluator while the topology is still revealed. In §5 we show how to design efficient constructions for PPBs that also allow to securely incorporate private constants into PPBs and give concrete constructions that are of special interest for practical applications. In particular we present efficient PPB constructions to compare two numbers and a number with a private constant. Other efficient PPB constructions for arithmetic operations (add or subtract two numbers/a number and a private constant, multiply a number with a private constant) and boolean operations are given in the full version of this paper [15]. Also switching functions, e.g., permutation and selection blocks, as well as universal circuits from [8] fit into this concept. The resulting SPF-SFE protocol is as efficient as Yao’s SFE protocol while providing function privacy at the same time.

In §8 we present an optimization algorithm that incorporates constant inputs into the circuit resulting in a circuit with less inputs and smaller size having a topology which is *independent* of the values of the constant inputs. Besides the well known propagation of constant inputs, our algorithm additionally eliminates resulting gates with one input by incorporating them into surrounding gates which results in smaller circuit size. The proposed optimization algorithm applies

no cryptographic modification of circuits and hence is of independent interest. This optimization can be used in combination with Yao’s SFE protocol in the relaxed-security scenario where constant inputs might be public values known to both parties as well as the inputs of circuit constructor Bob.

In order to allow usage of SPF-SFE in many practical applications we present a general compiler framework for secure evaluation of semi-private functions, called FairplaySPF, based on the well known Fairplay SFE system [13] as described in §6. Our Secure Programmable Block Description Language (SPBDL) allows to specify the topology of interconnected programmable blocks together with their corresponding private programming. A compiler automatically compiles SPBDL descriptions to circuits described in Fairplay’s Secure Hardware Description Language (SHDL). After incorporating Bob’s inputs into the circuit with the optimization algorithm presented in §8, the circuit can securely be evaluated with the SPF-SFE protocol while hiding the programming. Also a Universal Circuit (UC) that is evaluated in PF-SFE (cf. [8]) can be seen as a PPB that is programmed with a private circuit (specified in SHDL). By incorporating UCs as programmable blocks into SPBDL, our framework becomes a general purpose framework capable of expressing SFE, SPF-SFE, and PF-SFE as well as arbitrary combinations of them where only sensitive parts of the function’s structure are hidden as shown in the example in §7. This allows a fine-grained trade-off between performance and privacy of the evaluated function.

Our framework and compiler can be applied (combining SPF-SFE and PF-SFE) to implement and improve efficiency of several applications such as privacy-preserving credit checking [5], blinded policy evaluation [3,6,4], or secure data classification [16]. In §7 of this paper we concentrate on privacy-preserving credit checking. Usually, before getting a loan from a bank a person has to reveal a substantial amount of private information. This information has to satisfy certain criteria that are defined by the bank. We show how SPF-SFE can be used to securely evaluate the trustworthiness of a borrower while ensuring that (i) the privacy of his input is preserved and (ii) nothing is revealed about the criteria of the bank used for credit checking. Instead of using a UC for the whole function as in PF-SFE we reveal the topology of the trivial part of the function (e.g., comparing attributes with thresholds) and only hide the sensitive part in a UC, which is much more efficient. The description of the function in SPBDL can automatically be compiled into SHDL code with our compiler. This can be obliviously evaluated in a one-round protocol.

2 Yao’s Protocol and Semi-private Functions

Yao’s Protocol. In the following, we concentrate on Yao’s protocol [18] for SFE. Yao’s protocol is often called *garbled circuit* protocol as a garbled version of the (boolean) circuit representing the function is created by one party (constructor Bob) and evaluated by the other party (evaluator Alice) as described in the following. For each wire of the circuit, Bob uses two random bit strings

(garbled values) that are assigned to the corresponding values 0 and 1, respectively. Note, that the garbled values do not reveal to which value they correspond as they are chosen randomly. Bob sends *only* the garbled values corresponding to his inputs (garbled inputs) to Alice. For Alice's inputs, Bob uses 1-out-of-2 oblivious transfer (OT) to send Alice *only* the garbled values corresponding to her inputs without Bob learning which strings she gets. Additionally, for each gate G_i of the circuit, Bob creates and sends to Alice a *garbled table* T_i with the following property: given garbled values for G_i 's inputs, T_i allows to recover *only* the garbled value of the corresponding output of G_i and nothing else. Afterwards, Alice uses the received garbled values of the input wires and garbled tables T_i to evaluate the garbled circuit gate by gate. The output wires of the circuit are not garbled (or the mappings from garbled values to values 0 and 1 are published by Bob), thus Alice learns (only) the output of the circuit, but no plain values of internal wires (only garbled values). Correctness and security against semi-honest adversaries of Yao's protocol are proven in [10]. It is easy to show that Yao's protocol is even secure against malicious Alice, i.e. relaxed-secure, as the only message Alice sends to Bob is within the OT protocol where Alice is unable to cheat successfully if the OT protocol is secure against malicious Alice [4, Appendix A]. An efficient OT protocol with relaxed-security is given in [2].

Yao's protocol is the kernel of existing implementations of SFE protocols [13,12] which also extend it to be secure against malicious constructor Bob via cut-and-choose, e.g., multiple circuits are garbled, correctness of some of them is verified by revealing all garbled input values (called open), and the remaining ones are evaluated. As justified in the introduction, we concentrate on the plain Yao's protocol (secure against semi-honest Bob and potentially malicious Alice) where only *one* circuit is evaluated and no circuits are opened.

Yao's Protocol for Semi-Private Functions. Observe, in Yao's protocol the garbled tables T_i consist of symmetric encryptions of the garbled output value using the corresponding garbled input values as keys. Alice can use these garbled input values to decrypt exactly the one garbled output value corresponding to these keys. All other garbled output values, i.e., entries of the garbled function table remain hidden from Alice and hence she cannot determine the type of the gate. The only information Alice learns about the function in Yao's protocol is the *topology* of the circuit, i.e., the way the different gates are connected and how many inputs each gate has.

When Alice obtains a garbled circuit from Bob, she can guess from its topology what functionality the circuit evaluates, e.g., chains of 3-input gates might be an integer comparison circuit. This can be exploited constructively by Bob to keep parts of the function private, we call this a *semi-private function*, as follows. Bob composes his intended functionality from blocks with a fixed topology that can evaluate different functionalities each, called *privately programmable blocks* (PPBs) as explained in §4. The maximum amount of information Alice can gain from the topology of a PPB is the set of functionalities the PPB might compute but not the specific functionality chosen privately by Bob.

From combining these two arguments follows that evaluation of a circuit, composed out of several PPBs representing the semi-private function, with Yao's protocol is a secure protocol for SPF-SFE.

Additionally, (semi-honest) Bob can incorporate his input values into the circuit before garbling the circuit if they are already known at that time. In §8 we give an algorithm for efficient optimization of circuits for Bob's (constant) inputs together with an example. The optimization *only* depends on the topology of the original circuit but not on Bob's input values and hence the optimized circuit does not reveal more information on Bob's input values than the original circuit. After this optimization, Bob no longer needs to transfer the garbled values corresponding to his input values to Alice and also the size of the circuit is reduced (resulting in less communication and computation).

3 Definitions and Preliminaries

Let $x \in [0, 2^\ell)$ be an unsigned ℓ -bit integer value and $\mathbf{x} = (x_1, \dots, x_\ell)$, $x_i \in \{0, 1\}$ its corresponding representation as bit vector, i.e., $x = \sum_{i=1}^{\ell} x_i 2^{i-1}$. The *length* of \mathbf{x} is $|\mathbf{x}| = \ell$. We draw a (single) *wire* with one-bit value as \longrightarrow . As usual, *multi wire* X with ℓ -bit value x is drawn as $\xrightarrow{\ell}$ and consists of ℓ wires indexed by $X[i]$, $i = 1, \dots, \ell$ with values x_i .

A *gate* G with degree d has d *inputs* and one *output*. It is the implementation of a boolean function $g : \{0, 1\}^d \rightarrow \{0, 1\}$. As special case, a *constant gate* has no inputs ($d = 0$) and outputs a constant value. The *size* of a gate G , denoted by $|G|$, is the number of function table entries needed to implement the gate, namely $|G| = 2^d$. A gate with $e > 0$ outputs can easily be combined from e gates with one output resulting in size $e \cdot 2^d$.

We consider acyclic *circuits* consisting of connected gates with arbitrary fan-out, i.e., the output of each gate can be used as input to arbitrary many gates. The *size* of a circuit, denoted by $|C|$, is the sum of the sizes of its gates. Note, communication and computation complexity of efficient SFE protocols is linear in the size of the circuit.

A *block* B_v^u is a sub-circuit with u inputs in_1, \dots, in_u and v outputs out_1, \dots, out_v . B_v^u computes function $f_B : \{0, 1\}^u \rightarrow \{0, 1\}^v$ mapping input values to output values. Blocks consist of connected gates and other sub-blocks. *Size* of block B , denoted by $|B|$, is the sum of the sizes of its sub-elements.

A *programmable gate* (PG) is a gate with an unspecified function table. *Programming* it is done by providing a specific function table with 2^d entries (one entry for each input combination). The concept of PGs corresponds to a universal circuit for simulating a single gate in Valiant's UC construction [17]. As described in the previous section, in SPF-SFE evaluator Alice is not able to extract the corresponding function table (program) from PG. Analogously, a *programmable block* (PB) is a block consisting of programmable gates or programmable sub-blocks. It is programmed by programming each of its sub-elements. As described before, in SPF-SFE evaluator Alice is unable to extract the program from PB.

4 Privately Programmable Blocks

In this section we present the concept of *Privately Programmable Blocks* (PPB) for constructing semi-private functions. Using our efficient PPB constructions given in §5 with the SPF-SFE protocol of §2 allows to preserve the privacy of the function while the protocol remains as efficient as SFE protocol.

Definition 1. A Privately Programmable Block (*PPB*) is a programmable block which can be programmed to compute any function f of a given class of functions \mathcal{F} (e.g., $\mathcal{F} = \{ADD, SUB\}$) with a corresponding program p_f (e.g., $f = ADD$). We write PPB^f for a PPB which is programmed to compute f :

$$\forall f \in \mathcal{F}, \forall (in_1, \dots, in_u) \in \{0, 1\}^u : PPB^f(in_1, \dots, in_u) = f(in_1, \dots, in_u).$$

As explained in §2 before, in SPF-SFE the function to be evaluated is composed of several PPBs. Evaluator Alice learns how the PPBs are connected (topology), but the programming of the PPBs remains to be private information of constructor Bob (that's why PPBs are called *privately* programmable). Alice can infer from the topology of a PPB at most the class of possible functionalities \mathcal{F} but not the specific functionality f chosen by Bob. Hence, from Alice's point of view the PPB can compute any functionality from \mathcal{F} and the amount of information hidden inside the PPB is $\log_2 |\mathcal{F}|$ bits. For a semi-private function which is composed from programmable blocks PPB_1, \dots, PPB_n , the program of each PPB can be combined with any programming of the other PPBs and hence the maximum (as some combinations might not make sense depending on the application) amount of information hidden in the semi-private function is $\log_2(|\mathcal{F}_1| \cdot \dots \cdot |\mathcal{F}_n|) = \sum_{i=1}^n \log_2 |\mathcal{F}_i|$ bits. Clearly, if this is not large enough (i.e., if the number of PPBs n or number of possible functionalities of PPBs $|\mathcal{F}_i|$ is small), Alice might guess the correct function with high probability or probe the system via exhaustive search which must be prohibited by other means.

Universal Circuits (UC) indeed are special PPBs that can be programmed to compute an arbitrary function. UC_k is capable of simulating *any* function corresponding to a circuit with up to k gates with two inputs each. UCs provide *full privacy* of the evaluated function as the topology is hidden entirely. However, they cause a huge overhead by increasing the size of the evaluated circuit by $O(k \log k)$ [17], $O(k \log^2 k)$ [8], or $O(k^2)$ [16] additional gates which is often intolerable in practice. Evaluating a UC programmed with a private function known by constructor Bob with a SFE protocol is called Secure Evaluation of Private Functions (PF-SFE). By combining the PPBs presented in this paper with UCs, users can find a fine-grained trade-off between efficient PPB constructions for semi-private functions (SPF-SFE) and less efficient UC constructions for completely private functions (PF-SFE) as explained in §7.

Simple PPB Construction. A straight-forward implementation of a PPB for a class of n arbitrary functionalities $\mathcal{F} = \{f_1, f_2, \dots, f_n\}$ can directly be derived from the definition of PPBs in Definition 1 as shown in Fig. 1(a). Each functionality f_i is computed by a circuit C_i and a $n : 1$ multiplexer (*MUX*) is




Fig. 1. PPB constructions

programmed to select the intended output. The *MUX* block can be constructed from v parallel selection blocks S_1^n (as defined in [8]) for each of the v outputs that can be programmed to select any of their n inputs as outputs.

If the program p_f is known by Bob beforehand it can directly be incorporated into the circuit as described in §8. After optimization, each of the v selection blocks consists of a chain of $n - 1$ programmable 2-input gates which can be programmed to select either their left or right input as output each [8]. The size of this simple PPB construction is $|PPB^{simple}| = 4v(n - 1) + \sum_{i=1}^n |C_i|$.

Efficient PPB Constructions. *Efficient PPB constructions* can be obtained by choosing special classes of functionalities having circuits with the same (or at least a similar) topology. This allows to re-use (parts of) the same circuit C for the different functionalities f_i as shown in Fig. 1(b). For instance, the topology of an adder is the same as that of a subtractor and hence for $\mathcal{F} = \{ADD, SUB\}$ the same topology can be used. Based on the intended functionality $f \in \mathcal{F}$, the sub-elements of C are programmed differently while the topology is the same. This efficient PPB construction has size $|PPB^{efficient}| = |C| \approx |C_i| \ll |PPB^{simple}|$.

When a private constant c is incorporated into a PPB, the *value of the constant can not be extracted from PPB's topology* and hence is hidden from Alice in the SPF-SFE protocol, e.g., circuits to add/subtract an input with a s -bit constant c have the same topology. To simplify notation, we parametrize the class of possible functionalities with parameter c and write $\mathcal{F}_c = \{f_{1c}, \dots, f_{nc}\}$ for $\mathcal{F} = \{f_1|_{c=0}, \dots, f_1|_{c=2^s-1}, f_2|_{c=0}, \dots, f_2|_{c=2^s-1}, \dots, f_n|_{c=0}, \dots, f_n|_{c=2^s-1}\}$, e.g., $\mathcal{F}_c = \{ADD_c, SUB_c\}$ in the example given above. The amount of information hidden inside a PPB is

$$\log_2 |\mathcal{F}| = \log_2 |\mathcal{F}_c| + |c| = \log_2(n) + s \text{ bits.} \quad (1)$$

5 Practical Efficient PPB Constructions

In this section we show how to construct several efficient PPBs that are useful in practical applications (cf. §7). All these building blocks are already implemented in our framework for practical SPF-SFE described in §6. In the following we

present two efficient PPB constructions for *arithmetic operations*: compare two numbers and a number with a private constant. Other efficient PPB constructions for arithmetic operations (add or subtract two numbers/a number and a private constant, multiply a number with a private constant) and boolean operations are given in the full version of this paper [15]. Our SPF-SFE framework also provides PPBs for *Switching Functions* (i.e., permutation and selection blocks) and *Universal Circuits* for which we refer to the definitions, descriptions, and constructions in [8]. A list of efficient PPB constructions provided for implementation in our framework is given in the full version of this paper [15].

For each PPB we give the *Interface* specifying the functionality of the block, its number of inputs and outputs, and the different possibilities for programming \mathcal{F}_c . The *Implementation* describes the topology of the corresponding efficient PPB construction, how to program it, and its size. The inputs are called \mathbf{x} , \mathbf{y} and the potential private constant is called \mathbf{c} , where $|\mathbf{x}| = m$, $|\mathbf{y}| = n$, and $|\mathbf{c}| = s$. To simplify presentation we assume w.l.o.g. $m = n$, respectively $m = s$ in the following descriptions. The other cases can easily be derived from these by padding the shorter input with zeros and optimizing constant inputs afterwards as described in §8. Recall, evaluator Alice can neither extract the chosen function $f_c \in \mathcal{F}_c$, nor the value of the possibly embedded private constant $c \in \{0, 1\}^s$, from the topology of any PPB. The amount of information hidden inside the PPB is given by equation (1).

The main idea underlying efficient PPB constructions is to combine functionalities that have structurally equivalent recursive definitions that directly translate into programmable gates of equivalent topologies. E.g., comparison if two m -bit numbers \mathbf{x}, \mathbf{y} of bitlength m are less or equal is defined recursively as

$$(x \leq y) \Leftrightarrow \left((x_m < y_m) \vee ((x_m = y_m) \wedge ((x_{m-1}, \dots, x_1) \leq (y_{m-1}, \dots, y_1))) \right). \quad (2)$$

Whether two numbers are greater or equal is defined recursively as

$$(x \geq y) \Leftrightarrow \left((x_m > y_m) \vee ((x_m = y_m) \wedge ((x_{m-1}, \dots, x_1) \geq (y_{m-1}, \dots, y_1))) \right) \quad (3)$$

which is structurally equivalent and translates into the same topology (Fig. 2(b)).

5.1 PPB:COMP - Compare Two Numbers

Interface (Fig. 2(a)). PPB_{COMP} implements $z = f(x, y) = x \bowtie y$, where $\bowtie \in \{<, >, =, \leq, \geq, \neq\}$ and $|z| = 1$. The corresponding class of functions is $\mathcal{F} = \{L, G, E, LE, GE, NE\}$.

Implementation (Fig. 2(b)). Topology of PPB_{COMP} consists of a chain of m programmable gates PG_i (full comparers) with input bits x_i, y_i , and carry-in t_{i-1} and output carry-out t_i . The output of PPB_{COMP} is $z = t_m$ and the first carry $t_0 = 1$ can be directly incorporated into PG_1 . The carry t_i propagates whether for the i least significant bits $x_{<i} = x \bmod 2^i$ and $y_{<i} = y \bmod 2^i$ the corresponding relation is fulfilled ($t_i = 1$) or not ($t_i = 0$). In the following we describe the




Fig. 2. PPB:COMP

programming for the cases $=$, \leq , and \geq ; the corresponding cases \neq , $>$, and $<$ can be easily derived from this by negating output t_m in PG_m . In case $f = E$, PG_i is programmed to compute $t_i = (x_i = y_i) \wedge (x_{<i} = y_{<i}) = (x_i = y_i) \wedge t_{i-1}$. Analogously, in case $f = LE$, PG_i computes $t_i = (x_i < y_i) \vee [(x_i = y_i) \wedge t_{i-1}]$ and in case $f = GE$, PG_i computes $t_i = (x_i > y_i) \vee [(x_i = y_i) \wedge t_{i-1}]$. Note, these function table entries correspond exactly to the recursive definitions in equation (2) and (3). This block has size $|PPB_{COMP}| = (m - 1) \cdot 2^3 + 2^2 = 8m - 4$.

5.2 PPB:COMP_c - Compare Number with Private Constant

Interface (Fig. 3(a)). PPB_{COMP_c} implements $z = f_c(x) = x \bowtie c$, where $\bowtie \in \{<, >, =, \leq, \geq, \neq\}$, c is a private constant hidden inside PPB, and $|z| = 1$. The corresponding class of functions is $\mathcal{F}_c = \{L_c, G_c, E_c, LE_c, GE_c, NE_c\}$.

Implementation (Fig. 3(b)). Topology of PPB_{COMP_c} is exactly the same as that of PPB_{COMP} described in the previous section, however, each programmable gate PG_i has no input for y_i which is replaced by the internal constant c_i . The programming is the same as for PPB_{COMP} with constant c_i instead of input y_i . This block has size $|PPB_{COMP_c}| = (m - 1) \cdot 2^2 + 2^1 = 4m - 2$.




Fig. 3. PPB:COMP_c

6 FairplaySPF - A General Framework for SPF-SFE

We have implemented a general framework for Secure Evaluation of Semi-Private Functions (SPF-SFE) called *FairplaySPF*¹ by extending the Fairplay SFE

¹ FairplaySPF is available for download at <http://www.trust.rub.de/FairplaySPF>.

framework [13], both written in JAVA. Fairplay provides two languages: The high-level Secure Function Description Language (SFDL) allows users to specify the functionality to be computed with elements known from other high-level hardware description languages like VHDL or Verilog (e.g., variables, arrays, procedures, arithmetic- and logic expressions, control structures, etc.). Fairplay optimizes the function described in SFDL and automatically transforms it into a boolean circuit described in Fairplay’s low-level Secure Hardware Description Language (SHDL). This language consists of wires, input wires, gates, and output gates only. Using the SHDL circuit as input for both parties, Alice and Bob invoke their respective programs of the Fairplay runtime environment to execute the two-party SFE protocol. These programs evaluate the function on their respective private inputs over a TCP connection.

FairplaySPF Framework. In FairplaySPF, we extend the Fairplay framework [13] to secure evaluation of semi-private functions that are known to Bob only. In the following we describe the workflow of the FairplaySPF framework. Bob composes his semi-private function from several available privately programmable blocks (as described in §5) in our newly designed *Secure Programmable Block Description Language (SPBDL)* explained later in this section. Our FairplaySPF compiler automatically translates this SPBDL program into an SHDL circuit. Alternatively, SHDL circuits that are generated by the original Fairplay compiler from SFDL descriptions can be used. Bob’s private input data is automatically incorporated into the SHDL circuit and optimized afterwards by the FairplaySPF circuit optimizer as described in §8 resulting in a smaller SHDL circuit. This optimized SHDL circuit (containing the combination of Bob’s semi-private function and his private data) is evaluated by the FairplaySPF runtime environment (RE) which is only a slight modification of the Fairplay RE for semi-private functions: In FairplaySPF RE *only Bob* inputs the SHDL circuit but not Alice. The topology of the circuit (but without the types of the gates) is sent to Alice and afterwards the SPF-SFE protocol as described in §2 is executed between Alice and Bob over a TCP connection.

Secure Programmable Block Description Language (SPBDL). Our new SPBDL language allows to easily specify semi-private functions by combining different PPBs. SPBDL extends the basic functionality of SHDL to input wires (`input`), multi-wires (`vector`), privately programmable blocks (`block`), programmable gates (`gate`), and output wires (`output`). The formal specification of the *syntax* of SPBDL in Extended Backus-Naur Form (EBNF) is given in the full version of this paper [15]. In the following, we briefly describe the *semantics* of SPBDL. Please see Fig. 4 for an example SPBDL description of a semi-private function. As in SHDL, each line of a SPBDL program starts with a line number beginning with 0. In following lines, this number refers to the output of the element defined in this line. Line comments start with `//`.

A SPBDL program starts with the definition of inputs as `input Player [w]`, where `Player` defines from which party the input is given (`alice` or `bob`). The optional parameter `[w]` specifies that the input consists of w bits (default $w = 1$).

Afterwards, three kinds of elements can be specified - **gate**, **vector**, and **block**: A programmable gate is defined as `gate in [Wires] p [Bits]`, where `Wires` is its list of inputs and `Bits` is the programming of its function table. A list of `Wires` can be grouped into a vector with `vector [Wires]`. The single wires of a vector can be accessed via `Vector.Index`, e.g., `4.2` denotes the second wire of vector 4. A PPB is defined as `block [Btype] out Num in [Vects] p [Bprog]`, where `Btype` is the type of the PPB (e.g., `comp` for `PPB_COMP` described in §5), `Num` specifies the number of output bits, and `Vects` is the list of input vectors. The programming of the PPB specified in `Bprog` depends on the type of the PPB `Btype`. All types of PPBs `Btype` and corresponding programming parameters `Bprog` available in SPBDL are given in the full version of this paper [15]. Finally, outputs are defined as `output Player Vect`, where `Player` defines which party obtains the output (`alice` or `bob`) and `Vect` is the vector to be output.

7 Applications

Our general framework and tools for SPF-SFE presented in this paper can be used to specify and implement many privacy-preserving applications. Examples are *Blinded Policy Evaluation* [3,6,4], *Privacy-Preserving Credit Checking* [5], or provably secure evaluation of *Private Neural Networks* [16].

In the following we concentrate on privacy-preserving credit checking [5] which demonstrates how the evaluated function can be partitioned into semi-private and private parts which are both supported by our framework.

Privacy-Preserving Credit Checking. Typically, before granting a loan from a lender (Bob), the credit worthiness of the borrower (Alice) is checked to have the confidence that she will be able to pay it back later. The borrower is asked for her credit report that contains a large amount of private information including for example gender, age, income, salary, or other sensitive information like how many trade lines she owns, the number of overdrafts, or the number of late payments. However, revealing this data should be avoided as the lender may not always be a credible organization or, even worse, dishonest employees (so called insiders) could sell such private information on customers to third parties.

Additionally, the evaluation criteria of the lender are highly sensitive information that must be protected as revelation of these may cause loss of intellectually property or loss of repudiation for the lender.

As suggested by Frikken et al. [5], this scenario can be reduced to SPF-SFE, where Alice inputs her private credit report and Bob evaluates his semi-private function that checks if the credit report fulfills his criteria. To ensure that Alice inputs correct data into the SPF-SFE protocol, the authors describe how to replace the oblivious transfer step by a Credit Report Agency, i.e., a trusted third party, that checks and accredits Alice’s inputs instead.

Bob’s semi-private credit checking function can be expressed in our framework for SPF-SFE as shown in the tiny example of Fig. 4 which is due to space limitations not intended to give the complete solution but merely to show the

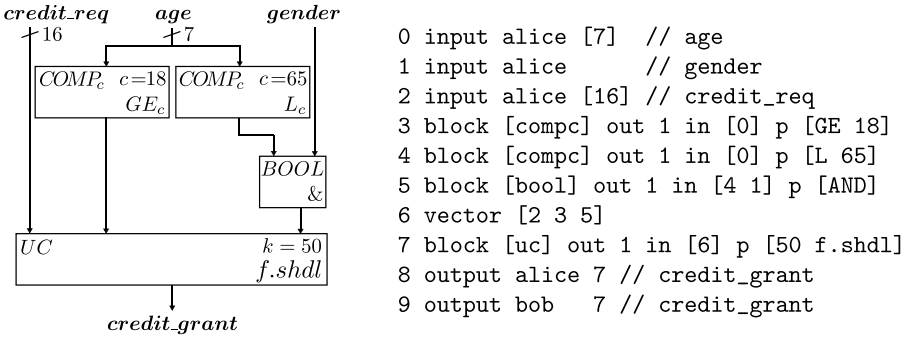


Fig. 4. Example for Privacy-Preserving Credit Checking

main concepts. The upper part of the circuit performs some obvious computation on Alice’s data, e.g., compare her *age* with a private constant, or combine this result with her *gender*. The sensitive information in this part of the function are the private constants, e.g., grant credit only to female persons (*gender* = 1) that are younger than 65 (*age* < 65), which are hidden from Alice, whereas the obvious topology can safely be revealed.

The highly sensitive part of the functionality that combines these results depending on the amount of credit requested (*credit_req*) is hidden entirely from Alice within the universal circuit *UC*. This PPB can be programmed to compute *any* functionality computable by a circuit of up to $k = 50$ gates with arbitrary topology. The specific functionality intended by Bob is the SHDL circuit described in `f.shdl`, which can automatically be generated from a high-level description in SFDL with the Fairplay compiler.

This example shows how our framework for SPF-SFE can be used to implement an application-specific, reasonable tradeoff between efficiency while revealing irrelevant information (SPF-SFE with PPBs) and complete function privacy (PF-SFE with UC).

Comparison of SPF-SFE and PF-SFE. Revealing the topology of obvious parts of the functionality while hiding the sensitive parts in a UC results in a smaller circuit as UC overhead can be substantially reduced due to less simulated gates k and less inputs into UC. This reduced size of the evaluated circuit directly translates into corresponding speedups in *any implementation* of the underlying SPF-SFE protocol as their performance must be at least linear in the size of the evaluated circuit.

As concrete example, Table 1 shows the number of gates that can be saved in the privacy-preserving credit checking example of Fig. 4 compared to hiding the functionality entirely in a UC in PF-SFE. For different maximum size k (row A) of the part of the functionality which is hidden in UC we give the achieved performance improvements when extracting the obvious part of the functionality into the upper part of the circuit ($COMP_c$ blocks and $BOOL$ block in Fig. 4). In our example, these blocks consist of 14 gates, i.e., row B contains the fraction

Table 1. Improved UC Overhead in the Example of Fig. 4

A) Gates hidden in UC, k	25	50	100
B) Gates extracted, $14/(k + 14)$	35.9%	21.9%	12.3%
C) UC overhead in PF-SFE (UC type)	1,861 (M3)	3,720 (M3)	8,264 (M3)
D) UC overhead in SPF-SFE (UC type)	850 (M1)	2,571 (M3)	6,797 (M3)
E) Improvement SPF-SFE vs. PF-SFE	1,011 (54.3%)	1,149 (30.9%)	1,467 (17.8%)

of the functionality which is revealed: $14/(k + 14)$. Row C shows how many gates are needed to hide the whole functionality of $14 + k$ gates in a UC with 24 inputs (for *credit_req*, *age*, and *gender*) using the most efficient UC construction of [16] which is denoted in parentheses. Row D shows how many gates are needed to implement the UC in our mixed approach as shown in Fig. 4, where UC has 18 inputs and simulates k gates. The resulting improvements compared to the PF-SFE solution (row E) supersede the fraction of the gates extracted (row B) as the number of inputs into UC is also reduced.

8 Optimization of Circuits with Constant Inputs

We describe a general optimization algorithm that incorporates constant inputs into a block (sub-circuit) B . The topology of the resulting optimized block B_{opt} is *independent* of the values of the constant inputs and its number of inputs and size are smaller, i.e., the number of gates respectively their degree is reduced as shown in Fig. 5. Besides the well known propagation of constant inputs (step 1), our algorithm additionally eliminates resulting gates with one input by incorporating them into surrounding gates (steps 2 and 3), which results in a smaller circuit size. The optimization algorithm is a non-cryptographic transformation of circuits and hence of independent interest. As outlined in §2, one possible

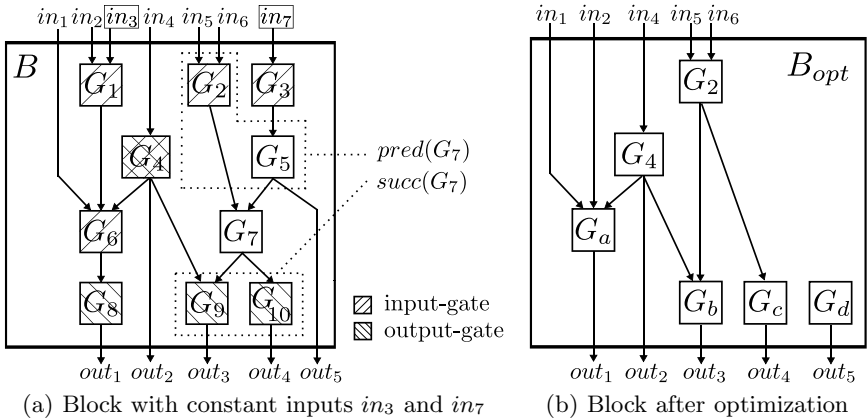


Fig. 5. Example for circuit optimization with Algorithm 1

application is to use this optimization to improve Yao’s protocol. In this application, constant inputs might be public constant values known to both parties as well as the private inputs of (semi-honest) circuit constructor Bob (if known at the time of construction of the garbled circuit).

Terminology. The following terminology is visualized in Fig. 5(a). Assume the gates $G_i, i = 1, \dots, n$ of a block B are numbered in topological order, i.e., gate G_i has no inputs that are outputs of gates with larger index $G_{j>i}$. Otherwise, this order can be obtained efficiently via topological sorting in $O(n)$.

An *output gate* is a gate whose output is also an output of B . Similarly, an *input gate* is a gate, which has at least one input that is also an input of B . For gate $G_i, pred(G_i)$ denotes the set of its predecessors, i.e., gates whose output is an input into G_i . Analogously, $succ(G_i)$ denotes the set of G_i ’s successors, i.e., gates having the output of G_i as input. The fan-out of a gate G_i is the number of its successors, i.e., $fanout(G_i) = \#succ(G_i)$.

Optimization. We refer to the running example of Fig. 5 that optimizes a block B with constant inputs in_3 and in_7 in the following description of Algorithm 1.

Step 1 - Eliminate constant inputs. The first step of Algorithm 1 eliminates all constant inputs $c_j, j = 1, \dots, c$ of block B with respective constant value $v_j \in \{0, 1\}$. For all gates G_i with degree d_i having c_j as k_i -th input, the function **eliminateConstInput**(G_i, k_i, v_j) is called that eliminates the corresponding input of G_i . Only the lines of the function table of G_i with value v_j in the k_i -th position are used while the other entries are eliminated, i.e., the modified gate $G_{i'}$ computes $g_{i'}(in_1, \dots, in_{k_i-1}, in_{k_i+1}, \dots, in_{d_i}) = g_i(in_1, \dots, in_{k_i-1}, v_j, in_{k_i+1}, \dots, in_{d_i})$.

Algorithm 1. Optimize block B with constant inputs

Input: Block B of gates G_1, \dots, G_n in topological order

Output: Optimized block B_{opt}

begin

```

1  # Eliminate constant inputs
   forall constant inputs  $c_j$  with constant value  $v_j$  that are not outputs of  $B$  do
     forall gates  $G_i$  having  $c_j$  as  $k_i$ -th input do
       eliminateConstInput( $G_i, k_i, v_j$ )
2  # Eliminate non-output gates with one input
   forall non-output gates  $G_i$  with  $d_i = 1$  do
     integrateInSucc( $G_i$ )
3  # Eliminate output gates with one input
   forall output gates  $G_i$  with  $d_i = 1$  do
     let  $\{G_p\} = pred(G_i)$ 
     if  $G_i$  is not input gate and  $fanout(G_p) = 1$  then
       integrateInPred( $G_i, G_p$ )

```

end

$|G_i|$ shrinks by a factor of two for each of its constant inputs. Let $\#c_i$ denote the number of constants of the d_i inputs of G_i , then $|G_{i'}| = 2^{d_i - \#c_i}$ after Step 1 of Algorithm 1 has eliminated all constant inputs. For an efficient implementation of Algorithm 1 it is crucial that **eliminateConstInput()** does not copy the entire function table of a gate G_i for each elimination of a constant input as this would result in runtime $O(\#c_i \cdot |G_i|)$ for each gate. Instead, the constant inputs are marked in runtime $O(\#c_i)$ and afterwards all constant inputs are eliminated simultaneously in runtime $O(|G_i|)$ by copying the corresponding elements of the function table. This results in runtime $O(|G_i|)$ per gate. Constant gates $G_{i'}$ with $d_{i'} = 0$ are propagated into their successors by recursively calling **eliminateConstInput($G_s, k_s, g_i(v_j)$)** for all $G_s \in \text{succ}(G_{i'})$ having $G_{i'}$ as k_s -th input. If constant gate $G_{i'}$ is not an output gate it is eliminated afterwards.

In the example of Fig. 5, constant input in_3 is input into gate G_1 whose size is reduced by half when eliminating the second input ($k_1 = 2$). The resulting gate G'_1 has one non-constant input in_2 and hence no further optimization is performed. The other constant input in_7 is input into G_3 which is optimized into a constant gate G'_3 by eliminating the constant input. Hence, **eliminateConstInput()** is called recursively for successor G_5 and G'_3 is eliminated. Similarly to G_3 , gate G_5 is reduced to a constant gate G'_5 and **eliminateConstInput()** is called for successor G_7 which eliminates its second input. As the output of G'_5 is also output of B it is not eliminated and remains as constant gate G_d .

After termination of Step 1 there might be gates G_i with one input left. The next two steps of Algorithm 1 try to remove these by incorporating their functionalities into their successors (Step 2) or predecessors (Step 3).

Step 2 - Eliminate non-output gates with one input. Step 2 of Algorithm 1 eliminates non-output gates with $d = 1$. The functionality of each one-input gate G_i which is not an output gate is incorporated into its successors $G_s \in \text{succ}(G_i)$ by the function **integrateInSucc(G_i)**. This function eliminates G_i by replacing it with a wire and incorporating the functionality of g_i into the function tables of all its successors $G_s \in \text{succ}(G_i)$: Let the output of G_i be the k -th input of G_s and d the degree of G_s . Then, the modified gate G'_s computes $g'_s(in_1, \dots, in_k, \dots, in_d) = g_s(in_1, \dots, g_i(in_k), \dots, in_d)$. Note that, independent of the functionality g_i , the resulting gate G'_s has the same size as G_s but additionally incorporates the functionality of g_i while not revealing any additional information on it. As in Step 1, for runtime $O(|G_i|)$ per gate the modifications of the function tables are not applied directly but first marked and then done simultaneously.

In the running example of Fig. 5, Step 2 eliminates G_1 by replacing it with a wire and modifying the function table of G_6 correspondingly. Analogously, gate G'_1 which only has one input from G_2 left after the optimizations performed in Step 1 is replaced by a wire. The function tables of its successors $G_9 \rightarrow G_b$ and $G_{10} \rightarrow G_c$ are modified correspondingly.

Step 3 - Eliminate output gates with one input. The third step of Algorithm 1 tries to eliminate output gates with $d = 1$. The functionality of each output gate G_i with one input is incorporated into its predecessor G_p . This is only possible

if G_i is the only successor of G_p , i.e., $\text{fanout}(G_p) = 1$. In this case, function **integrateInPred** (G_i, G_p) is called which eliminates gate G_i by replacing it with a wire and incorporates its functionality into gate G_p with d inputs. The modified gate G'_p computes $g'_p(\text{in}_1, \dots, \text{in}_d) = g_i(g_p(\text{in}_1, \dots, \text{in}_d))$. As in Step 2, this optimization step is independent of the functionality g_i and the resulting gate G'_p has the same size as G_p but additionally incorporates the functionality of g_i while not revealing any additional information on it.

In the running example of Fig. 5, Step 3 eliminates G_8 by replacing it with a wire and modifying the function table of $G_6 \rightarrow G_a$ correspondingly. In contrast to this, gate G_c cannot be incorporated into its predecessor G_2 as G_c is not its only successor ($\text{fanout}(G_2) = 2$). The optimized block B_{opt} produced by Algorithm 1 is shown in Fig. 5(b). It has size $|B_{\text{opt}}| = 21$ which is less than 62% of the size of the original block $|B| = 34$.

Correctness, efficiency and security of Algorithm 1 are summarized in the following theorem. Its proof is given in the full version of this paper [15].

Theorem 1. *Algorithm 1 efficiently eliminates all $c > 0$ constant inputs that are not outputs of block B in runtime $O(|B|)$. The optimized block B_{opt} has smaller size and computes the same functionality as B . The topology of B_{opt} does not reveal anything about the values of the constant inputs.*

Acknowledgements. We would like to thank Vladimir Kolesnikov and anonymous reviewers of ACNS'09 for helpful comments on the paper.

References

1. Ahn, L.v., Hopper, N.J., Langford, J.: Covert two-party computation. In: ACM Symposium on Theory of Computing (STOC 2005), pp. 513–522. ACM Press, New York (2005)
2. Aiello, W., Ishai, Y., Reingold, O.: Priced oblivious transfer: How to sell digital goods. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 119–135. Springer, Heidelberg (2001)
3. Frikken, K.B., Atallah, M.J., Li, J.: Hidden access control policies with hidden credentials. In: ACM Workshop on Privacy in the Electronic Society (WPES 2004), p. 27. ACM Press, New York (2004)
4. Frikken, K.B., Atallah, M.J., Li, J.: Attribute-based access control with hidden policies and hidden credentials. IEEE Trans. Comput. 55(10), 1259–1270 (2006)
5. Frikken, K.B., Atallah, M.J., Zhang, C.: Privacy-preserving credit checking. In: ACM conference on Electronic Commerce (EC 2005), pp. 147–154. ACM Press, New York (2005)
6. Frikken, K.B., Li, J., Atallah, M.J.: Trust negotiation with hidden credentials, hidden policies, and policy cycles. In: Network and Distributed System Security Symposium (NDSS 2006) (2006)
7. Kolesnikov, V., Schneider, T.: Improved garbled circuit: Free XOR gates and applications. In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfssdóttir, A., Walukiewicz, I. (eds.) ICALP 2008, Part II. LNCS, vol. 5126, pp. 486–498. Springer, Heidelberg (2008)

8. Kolesnikov, V., Schneider, T.: A practical universal circuit construction and secure evaluation of private functions. In: Tsudik, G. (ed.) FC 2008. LNCS, vol. 5143, pp. 83–97. Springer, Heidelberg (2008), <http://thomaschneider.de/FairplayPF>
9. Laur, S., Lipmaa, H.: A new protocol for conditional disclosure of secrets and its applications. In: Katz, J., Yung, M. (eds.) ACNS 2007. LNCS, vol. 4521, pp. 207–225. Springer, Heidelberg (2007)
10. Lindell, Y., Pinkas, B.: A proof of Yao’s protocol for secure two-party computation. ECCO Report TR04-063, Electr. Coll. on Comp. Complexity (2004)
11. Lindell, Y., Pinkas, B.: An efficient protocol for secure two-party computation in the presence of malicious adversaries. In: Naor, M. (ed.) EUROCRYPT 2007. LNCS, vol. 4515, pp. 52–78. Springer, Heidelberg (2007)
12. Lindell, Y., Pinkas, B., Smart, N.: Implementing two-party computation efficiently with security against malicious adversaries. In: Ostrovsky, R., De Prisco, R., Visconti, I. (eds.) SCN 2008. LNCS, vol. 5229, pp. 2–20. Springer, Heidelberg (2008)
13. Malkhi, D., Nisan, N., Pinkas, B., Sella, Y.: Fairplay — a secure two-party computation system. In: USENIX (2004), <http://www.cs.huji.ac.il/project/Fairplay/>
14. Naor, M., Pinkas, B.: Oblivious transfer and polynomial evaluation. In: ACM Symposium on Theory of Computing (STOC 1999), pp. 245–254. ACM Press, New York (1999)
15. Paus, A., Sadeghi, A.-R., Schneider, T.: Practical secure evaluation of semi-private functions. Cryptology ePrint Archive, Report 2009/124 (2009), <http://eprint.iacr.org/>
16. Sadeghi, A.-R., Schneider, T.: Generalized universal circuits for secure evaluation of private functions with application to data classification. In: Lee, P.J., Cheon, J.H. (eds.) ICISC 2008. LNCS, vol. 5461, pp. 336–353. Springer, Heidelberg (2008)
17. Valiant, L.G.: Universal circuits (preliminary report). In: Proc. 8th ACM Symp. on Theory of Computing (STOC 1976), pp. 196–203. ACM, New York (1976)
18. Yao, A.C.: How to generate and exchange secrets. In: Proc. 27th IEEE Symp. on Foundations of Comp. Science (FOCS 1986), Toronto, pp. 162–167. IEEE, Los Alamitos (1986)