Generalizing Multi-party Contract Signing*

Sjouke Mauw[1] and Saša Radomirović[2]

[1] CSC/SnT, University of Luxembourg
sjouke.mauw@uni.lu
[2] Institute of Information Security, Dept. of Computer Science, ETH Zurich
sasa.radomirovic@inf.ethz.ch

Abstract. Multi-party contract signing (MPCS) protocols allow a group of signers to exchange signatures on a predefined contract. Previous approaches considered either completely linear protocols or fully parallel broadcasting protocols. We introduce the new class of DAG MPCS protocols which combines parallel and linear execution and allows for parallelism even within a signer role. This generalization is useful in practical applications where the set of signers has a hierarchical structure, such as chaining of service level agreements and subcontracting.
Our novel DAG MPCS protocols are represented by directed acyclic graphs and equipped with a labeled transition system semantics. We define the notion of *abort-chaining sequences* and prove that a DAG MPCS protocol satisfies fairness if and only if it does not have an abort-chaining sequence. We exhibit several examples of optimistic fair DAG MPCS protocols. The fairness of these protocols follows from our theory and has additionally been verified with our automated tool.
We define two complexity measures for DAG MPCS protocols, related to execution time and total number of messages exchanged. We prove lower bounds for fair DAG MPCS protocols in terms of these measures.

1 Introduction

A multi-party contract signing (MPCS) protocol is a communication protocol that allows a number of parties to sign a digital contract. The need for MPCS protocols arises, for instance, in the context of service level agreements (SLAs) and in supply chain contracting. In these domains (electronic) contract negotiations and signing are still mainly bilateral. Instead of negotiating and signing one multi-party contract, in practice, multiple bilateral negotiations are conducted in parallel [21]. Because negotiations can fail, parties may end up with just a subset of the pursued bilateral contracts. If a party is missing contracts with providers or subcontractors, it faces an *overcommitment* problem. If contracts with customers are missing, it has an *overpurchasing* problem [8]. Both problems can be prevented by using fair multi-party contract signing protocols.

Existing optimistic MPCS protocols come in two flavors. *Linear* MPCS protocols require that at any point in time at most one signer has enough information to proceed in his role by sending messages to other signers. *Broadcast*

* This is the extended version of [14].

MPCS protocols specify a number of communication rounds in each of which all signers send or broadcast messages to each other. However, neither of the two kinds of protocols is suitable for SLAs or supply chain contracting. The reason is that in both domains, the set of contractors typically has a hierarchical structure, consisting of main contractors and levels of subcontractors. It is undesirable (and perhaps even infeasible) for the main contracting partners and their subcontractors to directly communicate with another partner's subcontractors. This restriction immediately excludes broadcast protocols as potential solutions and forces linear protocols to be impractically large.

In this paper we introduce MPCS protocol specifications that support arbitrary combinations of linear and parallel actions, even within a protocol role. The message flow of such protocols can be specified as a directed acyclic graph (DAG) and we therefore refer to them as *DAG* MPCS protocols.

A central requirement for MPCS protocols is *fairness*. This means that either all honest signers get all signatures on the negotiated contract or nobody gets the honest signers' signatures. It is well known that in asynchronous communication networks, a deterministic MPCS protocol requires a trusted third party (TTP) to achieve fairness [5]. In order to prevent the TTP from becoming a bottleneck, protocols have been designed in which the TTP is only involved to resolve conflicts. A conflict may occur if a message gets lost, if an external adversary interferes with the protocol, or if signers do not behave according to the protocol specification. If no conflicts occur, the TTP does not even have to be aware of the execution of the protocol. Such protocols are called *optimistic* [1]. We focus on optimistic protocols in this paper.

DAG MPCS protocols not only allow for better solutions to the subcontracting problem, but also have further advantages over linear and broadcast MPCS protocols and we design three novel MPCS protocols that demonstrate this. One such advantage concerns communication complexity. Linear protocols can reach the minimal number of messages necessary to be exchanged in fair MPCS protocols at the cost of a high number of protocol "rounds". We call this the *parallel complexity*, which is a generalization of the round complexity measure for broadcast protocols, and define it in Section 4.3. Conversely, broadcast protocols can attain the minimal number of protocol rounds necessary for fair MPCS, but at the cost of a high message complexity. We demonstrate that DAG MPCS protocols can simultaneously attain best possible order of magnitude for both complexity measures.

As discussed in our related work section, the design of fair MPCS protocols has proven to be non-trivial and error-prone. We therefore not only prove our three novel DAG MPCS protocols to be fair, but we also derive necessary and sufficient conditions for fairness of any optimistic DAG MPCS protocol. These conditions can be implemented and verified automatically, but they are still non-trivial. Therefore, for a slightly restricted class of DAG protocols, we additionally derive a fairness criterion that is easy to verify.

Contributions. Our main contributions are (i) the definition of a syntax and interleaving semantics of DAG MPCS protocols (Section 4.1); (ii) the definition

of the message complexity and parallel complexity of such protocols (Section 4.3); (iii) a method to derive a full MPCS specification from a *skeletal graph*, including the TTP logic (Section 5); (iv) necessary and sufficient conditions for fairness of DAG MPCS protocols (Section 5.3); (v) minimal complexity bounds for DAG MPCS protocols (Section 6.1); (vi) novel fair MPCS protocols (Section 6.2); (vii) a software tool that verifies whether a given MPCS protocol is fair (described in Appendix A). [3]

2 Related Work

We build on the body of work that has been published in the field of fair optimistic MPCS protocols in asynchronous networks. The first such protocols were proposed by Baum-Waidner and Waidner [2], viz. a round-based broadcast protocol and a related round-based linear protocol. They showed subsequently [3] that these protocols are round-optimal. This is a complexity measure that is related to, but less general than, parallel complexity defined in the present paper.

Garay et al. [6] introduced the notion of *abuse-free* contract signing. They developed the technique of *private contract signature* and used it to create abuse-free two-party and three-party contract signing protocols. Garay and MacKenzie [7] proposed MPCS protocols which were later shown to be unfair using the model checker Mocha and improved by Chadha et al. [4]. Mukhamedov and Ryan [17] developed the notion of *abort chaining attacks* and used such attacks to show that Chadha et al.'s improved version does not satisfy fairness in cases where there are more than five signers. They introduced a new optimistic MPCS protocol and proved fairness for their protocol by hand and used the NuSMV model checker to verify the case of five signers. Zhang et al. [22] have used Mocha to analyze the protocols of Mukhamedov and Ryan and of Mauw et al. [15].

Mauw et al. [15] used the notion of abort chaining to establish a lower bound on the message complexity of linear fair MPCS protocols. This complexity measure is generalized in the present paper to DAG MPCS protocols. Kordy and Radomirović [10] have shown an explicit construction for fair linear MPCS protocols. The construction covers in particular the protocols proposed by Mukhamedov and Ryan [17] and the linear protocol of Baum-Waidner and Waidner [2], but not the broadcast protocols. The DAG MPCS protocol model and fairness results developed in the present paper encompass both types of protocols. They allow for arbitrary combinations of linear and parallel behaviour (i.e. partial parallelism), and in addition allow for parallelism within signer roles. MPCS protocols combining linear and parallel behaviour have not been studied yet.

Apart from new theoretical insights to be gained from designing and studying DAG MPCS protocols, we anticipate interesting application domains in which multiple parties establish a number of related contracts, such as SLAs. Emerging business models like Software as a Service require a negotiation to balance a customer's requirements against a service provider's capabilities. The result of

[3] The tool and models of our protocols are available at the following website: `http://people.inf.ethz.ch/rsasa/mpcs`

such a negotiation is often complicated by the dependencies between several contracts [13] and multi-party protocols may serve to mitigate this problem. Karaenke and Kirn [8] propose a multi-tier negotiation protocol to mitigate the problems of overcommitment and overpurchasing. They formally verify that the protocol solves the two observed problems, but do not consider the fairness problem. SLAs and negotiation protocols have also been studied in the multi-agent community. An example is the work of Kraus [11] who defines a multi-party negotiation protocol in which agreement is reached if all agents accept an offer. If the offer is rejected by at least one agent, a new offer will be negotiated.

Another interesting application area concerns *supply chain contracting* [12]. A supply chain consists of a series of firms involved in the production of a product or service with potentially complex contractual relationships. Most literature in this area focuses on economic aspects, like pricing strategies. An exception is the recent work of Pavlov and Katok [9] in which fairness is studied from a game-theoretic point of view. The study of multi-party signing protocols and multi-contract protocols has only recently been identified as an interesting research topic in this application area [20].

3 Preliminaries

3.1 Multi-party contract signing

The purpose of a multi-party contract signing protocol is to allow a number of parties to sign a digital contract in a fair way. In this section we recall the basic notions pertaining to MPCS protocols. We use A to denote the set of signers involved in a protocol, $\mathfrak{C}$ to denote the contract, and T to denote the TTP.

A signer is considered *honest* (cf. Definition 5) if it faithfully executes the protocol specification. An MPCS protocol is said to be *optimistic* if its execution in absence of adversarial behaviour and failures and with all honest signers results in signed contracts for all participants without any involvement of T. Optimistic MPCS protocols consist of two subprotocols: the *main* protocol that specifies the exchange of *promises* and *signatures* by the signers, and the *resolve* protocol that describes the interaction between the agents and T in case of a failure in the main protocol. A promise made by a signer indicates the intent to sign $\mathfrak{C}$. A promise $\wp_P(m, x, Q, \mathrm{T})$ can only be generated by signer $P \in A$. The content (m, x) can be extracted from the promise and the promise can be verified by signer $Q \in A$ and by T. A signature $\mathcal{S}_P(m)$ can only be generated by P and by T, if T has a promise $\wp_P(m, x, Q, \mathrm{T})$. The content m can be extracted and the signature can be verified by anybody. Cryptographic schemes that allow for the above properties are digital signature schemes and private contract signatures [6].

MPCS protocols must satisfy at least two security requirements, namely *fairness* and *timeliness*. An optimistic MPCS protocol for contract $\mathfrak{C}$ is said to be *fair* for an honest signer P if whenever some signer $Q \neq P$ obtains a signature on $\mathfrak{C}$ from P, then P can obtain a signature on $\mathfrak{C}$ from all signers participating in the protocol. An optimistic MPCS protocol is said to satisfy *timeliness*, if each

signer has a recourse to stop endless waiting for expected messages. The fairness requirement will be the guiding principle for our investigations and timeliness will be implied by the communication model together with the behaviour of the TTP. A formal definition of fairness is given in Section 5.3.

A further desirable property for MPCS protocols is abuse-freeness which was introduced in [6]. An optimistic MPCS protocol is said to be *abuse-free*, if it is impossible for any set of signers at any point in the protocol to be able to prove to an outside party that they have the power to terminate or successfully complete the contract signing. Abuse-freeness is outside the scope of this paper.

3.2 Graphs

Let $G = (V, E)$ with $E \subseteq V \times V$ be a directed acyclic graph. Let $v, w \in V$ be vertices. We say that v *causally precedes* w, denoted $v \prec w$, if there is a directed path from v to w in the graph. We write $v \preceq w$ for $v \prec w \vee v = w$. We extend *causal precedence* to the set $V \cup E$ as follows. Given two edges $(v, w), (v', w') \in E$, we say that (v, w) *causally precedes* (v', w') and write $(v, w) \prec (v', w')$, if $w \preceq v'$. Similarly, we write $v \prec (v', w')$ if $v \preceq v'$ and $(v, w) \prec v'$ if $w \preceq v'$. Let $x, y \in V \cup E$. If x causally precedes y we also say that y *causally follows* x. We say that a set $S \subseteq V \cup E$ is *causally closed* if it contains all causally preceding vertices and edges of its elements, i.e., $\forall x \in S, y \in V \cup E : y \prec x \implies y \in S$.

By $\mathrm{in}(v) \subseteq E$ we denote the set of edges incoming to v and by $\mathrm{out}(v) \subseteq E$ the set of edges outgoing from v. Formally, we have $\mathrm{in}(v) = \{(w, v) \in E \mid w \in V\}$ and $\mathrm{out}(v) = \{(v, w) \in E \mid w \in V\}$.

3.3 Assumptions

The communication between signers is asynchronous and messages can get lost or be delayed arbitrary long. The communication channels between signers and the TTP T are assumed to be *resilient*. This means that the messages sent over these channels are guaranteed to be delivered eventually. In order to simplify our reasoning, we assume that the channels between protocol participants are confidential and authentic. We consider the problem of delivering confidential and authentic messages in a Dolev-Yao intruder model to be orthogonal to the present problem setting.

We assume that $\mathfrak{C}$ contains the contract text along with fresh values (contributed by every signer) which prevent different protocol executions from generating interchangeable protocol messages. Furthermore we assume that $\mathfrak{C}$ contains all information that T needs in order to reach a decision regarding the contract in case it is contacted by a signer. This information contains the protocol specification, an identifier for T, identifiers for the signers involved in the protocol, and the assignment of signers to protocol roles in the protocol specification.

We assume the existence of a designated resolution process per signer which coordinates the various resolution requests of the signer's parallel threads. It will ensure that T is contacted at most once by the signer. After having received the first request from one of the signer's threads, this resolution process will contact

T on behalf of the signer and store T's reply. This reply will be forwarded to all of the signer's threads whenever they request resolution.

4 DAG Protocols

Our DAG protocol model is a multi-party protocol model in an asynchronous network with a TTP and an adversary that controls a subset of parties.

4.1 Specification and Execution Model

A *DAG protocol specification* (or simply, a *protocol specification*) is a directed acyclic graph in which the vertices represent the state of a signer and the edges represent either a causal dependency between two states (an ε-edge) or the sending of a message. A vertex' outgoing edges can be executed in parallel. Edges labeled with *exit* denote that a signer contacts T.

Definition 1. *Let R be a set of roles such that $\mathrm{T} \notin R$ and M a set of messages. Let ε and exit be two symbols, such that $\varepsilon, exit \notin M$. By M_ε^{exit} and R_T we denote the sets $M_\varepsilon^{exit} = M \cup \{\varepsilon, exit\}$ and $R_\mathrm{T} = R \cup \{\mathrm{T}\}$, respectively. A* DAG protocol specification *is a labeled directed acyclic graph $\mathcal{P} = (V, E, r, \mu, \delta)$, where*

1. *(V, E) is a directed acyclic graph;*
2. *$r \colon V \to R_\mathrm{T}$ is a labeling function assigning roles to vertices;*
3. *$\mu \colon E \to M_\varepsilon^{exit}$ is an edge-labeling function that satisfies*
   (a) $\forall (v, v') \in E \colon \mu(v, v') = \varepsilon \implies r(v) = r(v')$,
   (b) $\forall (v, v') \in E \colon \mu(v, v') = exit \implies r(v') = \mathrm{T}$;
4. *$\delta \colon M^* \to M$ is a function associated with exit-labeled edges.*

A message edge (v, v') specifies that $\mu(v, v') = m$ is to be sent from role $r(v)$ to role $r(v')$. An ε-edge (v, v') represents internal progress of role $r(v) = r(v')$ and allows to specify a causal order in the role's events. An exit edge denotes that a role can contact the TTP. The TTP then uses the function δ to determine a reply to the requesting role, based on the sequence of messages that it has received. In Section 5 exit messages and the δ function are used to model the resolve protocol of the TTP.

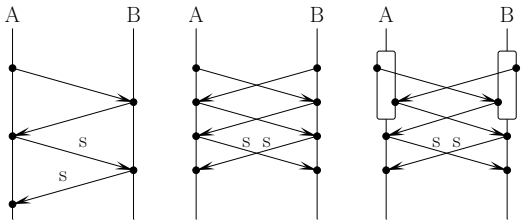

Fig. 1: Linear, broadcast, and the novel DAG MPCS protocols.

We give three examples of DAG protocols in Figure 1, represented as Message Sequence Charts (MSCs). The dots denote the vertices, which we group vertically below their corresponding role names. Vertical lines in the MSCs correspond to ε-edges and horizontal or diagonal edges represent message edges. We mark edges labeled with signing messages with an "s" and we leave out the edge labels of promise messages. We do not display exit edges, they are implied by the MPCS protocol specification. A box represents the splitting of a role into two parallel threads, which join again at the end of the box. We revert to a traditional representation of labeled DAGs if it is more convenient (see, e.g., Figure 2).

The first protocol in Figure 1 is a classical linear 2-party contract signing protocol. It consists of one round of promises followed by a round of exchanging signatures. The second protocol is the classical broadcast protocol for two signers. It consists of two rounds of promises, followed by one round of signatures. The third protocol is a novel DAG protocol, showing the power of in-role parallelism. It is derived from the broadcasting protocol by observing that its fairness does not depend on the causal order of the first two vertices of each of the roles.

Let $\mathcal{P} = (V, E, r, \mu, \delta)$ be a protocol specification. The *restriction* of $\mathcal{P}$ to role P, denoted by $\mathcal{P}_P$, is the protocol specification $(V_P, E_P, r_P, \mu_P, \delta_P)$, where

$$E_P = \{(v, v') \in E \mid r(v) = P \vee r(v') = P\}\,, \quad V_P = \{v, v' \in V \mid (v, v') \in E_P\}\,,$$
$$r_P(v) = r(v) \text{ for } v \in V_P, \quad \mu_P(e) = \mu(e) \text{ for } e \in E_P, \quad \text{and} \quad \delta_P = \delta.$$

The execution state of a protocol consists of the set of events, connected to vertices or edges, that have been executed.

Definition 2. *Let $\mathcal{P} = (V, E, r, \mu, \delta)$ be a protocol specification. A* state *of $\mathcal{P}$ is a set $s \subseteq V \cup E$. The set of states of $\mathcal{P}$ is denoted by $\mathcal{S}_\mathcal{P}$. The* initial state *of $\mathcal{P}$ is defined as $s_0 = \emptyset$.*

In order to give DAG protocols a semantics, we first define the *transition relation* between states of a protocol.

Definition 3. *Let $\mathcal{P} = (V, E, r, \mu, \delta)$ be a protocol specification, $L = \{\varepsilon, send, recv, exit\}$ the set of transition labels, and $s, s' \in \mathcal{S}_\mathcal{P}$ the states of $\mathcal{P}$. We say that $\mathcal{P}$ transitions with label α from state s into s', denoted by $s \overset{\alpha}{\rightsquigarrow} s'$, iff $s \neq s'$ and one of the following conditions holds*

1. *$\alpha = recv$ and $\exists v \in V$, such that $\mathrm{in}(v) \subseteq s$ and $s' = s \cup \{v\}$,*
2. *$\alpha = send$ and $\exists m \in M, e \in E$, such that $\mu(e) = m$, and $s' = s \cup \{e\}$,*
3. *$\alpha = \varepsilon$ and $\exists e = (v, v') \in E$, such that $\mu(e) = \varepsilon$, $v \in s$ and $s' = s \cup \{e\}$,*
4. *$\alpha = exit$ and $\exists e \in E$, such that $\mu(e) = exit$ and $s' = s \cup \{e\}$.*

In Definition 3, receive events are represented by vertices, all other events by edges. By the first condition, a receive event can only occur if all events assigned to the incoming edges have occurred. In contrast, the sending of messages (second condition) can take place at any time. The third condition states that an ε-edge can be executed if the event on which it causally depends has been executed.

Finally, like send events, an exit event can occur at any time. Every event may occur at most once, however. This is ensured by the condition $s' \neq s$.

The transitions model all possible behavior of the system. The behavior of honest agents in the system will be restricted as detailed in the following subsection. We denote sequences by $[a_0, a_1, \ldots, a_l]$ and the concatenation of two sequences σ_1, σ_2 by $\sigma_1 \cdot \sigma_2$.

Definition 4. *Let $\mathcal{P} = (V, E, r, \mu, \delta)$ be a protocol specification and $L = \{\varepsilon, send, recv, exit\}$ a set of labels. The* semantics *of $\mathcal{P}$ is the labeled transition system $(\mathcal{S}_\mathcal{P}, L, \rightsquigarrow, s_0)$, which is a graph consisting of vertices $\mathcal{S}_\mathcal{P}$ and edges $\rightsquigarrow$ with start state s_0. An* execution *of $\mathcal{P}$ is a finite sequence $\rho = [s_0, \alpha_1, s_1, \ldots, \alpha_l, s_l]$, $l \geq 0$, such that $\forall i \in \{0, \ldots, l-1\}$: $s_i \overset{\alpha_{i+1}}{\rightsquigarrow} s_{i+1}$. The set of all executions of $\mathcal{P}$ is denoted by* $\mathrm{Exe}(\mathcal{P})$.

If $\rho = [s_0, \alpha_1, s_1, \ldots, \alpha_l, s_l]$ is an execution of $\mathcal{P}$ and $\mathcal{P}_P$ is the restriction to role P, then the *restricted* execution ρ_P is obtained inductively as follows.

1. $[s]_P = [s \cap (V_P \cup E_P)]$ for a state s.
2. $([s, \alpha, s'] \cdot \sigma)_P = \begin{cases} [s]_P \cdot \sigma_P & \text{if } [s]_P = [s']_P \\ [s]_P \cdot [\alpha] \cdot ([s'] \cdot \sigma)_P & \text{else.} \end{cases}$

Commutativity of restriction and execution is asserted by the following lemma.

Lemma 1. *Let $\mathcal{P}$ be a protocol specification and $\mathcal{P}_P$ the restriction to role P. Then every restricted execution ρ_P is an execution of $\mathcal{P}_P$.*

4.2 Adversary Model

An honest agent executes the protocol specification faithfully. The following definition specifies what this entails for a DAG protocol: the agent waits for the reception of all causally preceding messages before sending causally following messages, does not execute an *exit* edge attached to a vertex v if all messages at v have been received and never executes more than one *exit* edge (which in the context of MPCS protocols corresponds to contacting the TTP at most once), and does not send any messages which causally follow a vertex from which the *exit* edge was executed.

Definition 5. *Let $\mathcal{P}$ be a DAG protocol specification. An agent P is* honest *in an execution ρ of $\mathcal{P}$, if all states s of the restricted execution ρ_P satisfy the following conditions:*

1. *s contains at most one exit edge.*
2. *If s contains no exit edge, then s is causally closed.*
3. *If s contains an exit edge $e = (v, w)$, $\mu(e) = exit$, then $v \notin s$ and $s \setminus \{e\}$ is causally closed.*

A dishonest agent is only limited by the execution model. Thus a dishonest agent can send its messages at any time and in any order, regardless of the causal precedence given in the protocol specification. A dishonest agent can execute multiple *exit* edges and may send and receive messages causally following an exit edge. Dishonest agents are controlled by a single adversary, their knowledge is shared with the adversary. The adversary can delay or block messages sent from one agent to another, but the adversary cannot prevent messages between agents and the TTP from being delivered eventually. All communication channels are authentic and confidential.

4.3 Communication Complexity

To define measures for expressing the communication complexity of DAG protocols, we introduce the notion of *closed executions*. A closed execution is a complete execution of the protocol by honest agents.

Definition 6. *Let $\mathcal{P} = (V, E, r, \mu, \delta)$ be a protocol specification and $(\mathcal{S}_\mathcal{P}, L, \rightsquigarrow, s_0)$ be the semantics for $\mathcal{P}$. Given $\rho = [s_0, \alpha_1, s_1, \ldots, \alpha_l, s_l] \in \mathrm{Exe}(\mathcal{P})$, we say that ρ is* closed *if the following three conditions are satisfied*

1. *s_i is causally closed, for every $0 \leq i \leq l$,*
2. *$\alpha_i \neq exit$, for every $1 \leq i \leq l$,*
3. *$\nexists \alpha \in L \setminus \{exit\}, s \in \mathcal{S}_\mathcal{P} \colon \rho \cdot [\alpha, s] \in \mathrm{Exe}(\mathcal{P})$.*

The set of all closed executions of $\mathcal{P}$ is denoted by $\mathrm{Exe_C}(\mathcal{P})$.

Let $\rho = [s_0, \alpha_1, s_1, \ldots, \alpha_l, s_l]$ be an execution of a protocol $\mathcal{P}$. By $|\rho|_{send}$ we denote the number of labels α_i, for $1 \leq i \leq l$, such that $\alpha_i = send$.

Lemma 2. *For any two closed executions ρ and ρ' of a protocol $\mathcal{P}$ we have $|\rho|_{send} = |\rho'|_{send}$.*

The proof is given in the appendix. The first measure expressing the complexity of a protocol $\mathcal{P}$ is called *message complexity*. It counts the overall number of messages that have been sent in a closed execution of a protocol $\mathcal{P}$.

Definition 7. *Let $\mathcal{P}$ be a protocol specification and let $\rho \in \mathrm{Exe_C}(\mathcal{P})$. The* message complexity *of $\mathcal{P}$, denoted by $MC_\mathcal{P}$, is defined as $MC_\mathcal{P} = |\rho|_{send}$.*

Lemma 2 guarantees that the message complexity of a protocol is well defined.

The second complexity measure is called *parallel complexity*. It represents the minimal time of a closed execution assuming that all events which can be executed in parallel are executed in parallel. The parallel complexity of a protocol is defined as the length of a maximal chain of causally related send edges.

Definition 8. *The* parallel complexity *of a protocol $\mathcal{P}$, denoted by $PC_\mathcal{P}$, is defined as*

$$PC_\mathcal{P} = \max_{n \in \mathbb{N}} \exists_{[e_1, e_2, \ldots, e_n] \in E^*} : \forall_{1 \leq i \leq n} : \mu(e_i) = send \wedge \forall_{1 \leq i < n} : e_i \prec e_{i+1}.$$

Example 1. The message complexity of the first protocol in Figure 1 is 4, which is known to be optimal for two signers [19]. Its parallel complexity is 4, too. The message complexity of the other two protocols is 6, but their parallel complexity is 3, which is optimal for broadcasting protocols with two signers [3].

5 DAG MPCS protocols

We now define a class of optimistic MPCS protocols in the DAG protocol model.

5.1 Main Protocol

The key requirements we want our DAG MPCS protocol specification to satisfy, stated formally in Definition 9, are as follows. The messages exchanged between signers in the protocol are of two types, promises, denoted by $\wp()$, and signatures, denoted by $\mathcal{S}()$. Every promise contains information about the vertex from which it is sent. This is done by concatenating the contract $\mathfrak{C}$ with the vertex v the promise originates from and is denoted by $(\mathfrak{C}, v)$. The signers can contact the TTP at any time. This is modeled with exit edges: Every vertex $v \in V$ such that $r(v) \in A$ (the set of all signers) is adjacent to a unique vertex $v_{\mathrm{T}} \in V$, $r(v_{\mathrm{T}}) = \mathrm{T}$. The communication with T is represented by δ. The set of vertices with outgoing signature messages is denoted by *SigSet*.

Definition 9. *Let $\mathcal{P} = (V, E, r, \mu, \delta)$ be a protocol specification, $A \subset R$ be a finite set of signers, $\mathfrak{C}$ be a contract, and $SigSet \subseteq V$. $\mathcal{P}$ is called a* DAG MPCS protocol specification for $\mathfrak{C}$, *if* [4]

1. $\exists!\ v_{\mathrm{T}} \in V : r(v_{\mathrm{T}}) = \mathrm{T} \wedge \forall v \in V \setminus \{v_{\mathrm{T}}\} : (v, v_{\mathrm{T}}) \in E$,
2. $\forall v, w \in V : v \prec w \Rightarrow (v, w) \in E \vee \exists u \in V : v \prec u \prec w \wedge r(u) \in \{r(v), r(w)\}$,
3. $\forall (v, w) \in E : \mu(v, w) =$

$$\begin{cases} \varepsilon, & \text{if } r(v) = r(w), \\ \mathit{exit}, & \text{if } w = v_{\mathrm{T}}, \\ \mathcal{S}_{r(v)}(\mathfrak{C}), & \text{if } v \in \mathit{SigSet} \wedge r(v) \neq r(w) \neq \mathrm{T}, \\ \wp_{r(v)}(\mathfrak{C}, v, r(w), \mathrm{T}), & \text{else.} \end{cases}$$

4. $\delta : M^* \to \left\{ \text{“}\mathit{abort}\text{”}, (\mathcal{S}_P(\mathfrak{C}))_{P \in A} \right\}$, *where $(\mathcal{S}_P(\mathfrak{C}))_{P \in A}$ denotes a list of signatures on $\mathfrak{C}$, one by each signer.*

We write $SigSet(\mathcal{P})$ for the largest subset of SigSet which satisfies

$$v \in \mathit{SigSet}(\mathcal{P}) \Rightarrow \exists w \in V : (v, w) \in E, \mu(v, w) \in M.$$

The set $SigSet(\mathcal{P})$ is called the signing set.

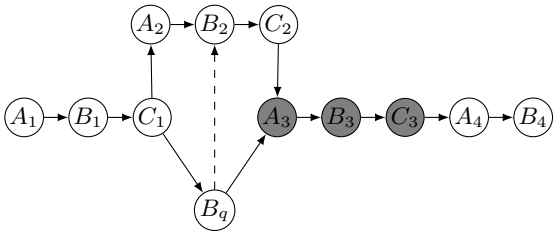

(a) Skeletal graph.

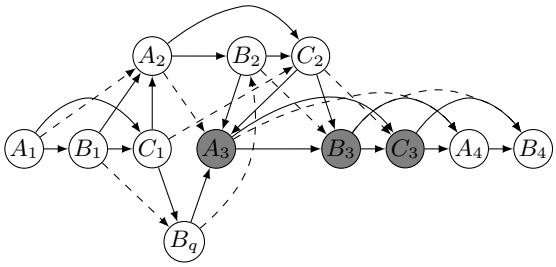

(b) Full graph.

Fig. 2: Skeletal and full representation of a DAG MPCS protocol.

We represent DAG MPCS protocols as *skeletal* graphs as shown in Figure 2a. The full graph, shown in Figure 2b, is obtained from the skeletal graph by adding all edges required by condition 2 of Definition 9 and extending μ according to condition 3. The ε edges are dashed in the graphs. The shaded vertices in the graphs indicate the vertices that are in the signing set. We define the *knowledge* $K(v)$ of a vertex v to be the set of message edges causally preceding v, and incoming to a vertex of the same role. The knowledge of a vertex represents the state right after its receive event.

$$K(v) = \{(w, v') \in E \mid \mu(w, v') \in M, v' \preceq v, r(v') = r(v)\}$$

We define the *pre-knowledge* $K_\prec(v)$ of a vertex v to be $K_\prec(v) = \{(w, v') \in K(v) \mid v' \prec v\}$. The pre-knowledge represents the state just *before* the vertex' receive event has taken place. We extend both definitions to sets $S \subseteq V$:

$$K(S) = \bigcup_{v \in S} K(v) \quad \text{and} \quad K_\prec(S) = \bigcup_{v \in S} K_\prec(v).$$

We define the *initial set* of $\mathcal{P}$, denoted $\mathit{InitSet}(\mathcal{P})$ to be the set of vertices of the protocol specification for which the pre-knowledge of the same role does not contain an incoming edge by every other role. Formally,

$$v \in \mathit{InitSet}(\mathcal{P}) \iff \{r(w) \in A \mid (w, v') \in K_\prec(v)\} \cup \{r(v)\} \neq A$$

[4] We write $\exists!$ for unique existential quantification.

The *end set* of $\mathcal{P}$, denoted $\mathit{EndSet}(\mathcal{P})$, is the set of vertices of the protocol specification at which the corresponding signer possesses all signatures.

$$v \in \mathit{EndSet}(\mathcal{P}) \iff \{r(w) \in A \mid (w, v') \in K(v), w \in \mathit{SigSet}(\mathcal{P})\} \cup \{r(v)\} = A$$

5.2 Resolve Protocol

Let $\mathcal{P} = (V, E, r, \mu, \delta)$ be a DAG MPCS protocol specification. The resolve protocol is a two-message protocol between a signer and the TTP T, initiated by the signer. The communication channels for this protocol are assumed to be resilient, confidential, and authentic. T is assumed to respond immediately to the signer. This is modeled in $\mathcal{P}$ via an exit edge from a vertex $v \in V \setminus \{v_{\mathrm{T}}\}$ to the unique vertex $v_{\mathrm{T}} \in V$. T's response is given by the δ function, $\delta : M^* \to \{\text{“abort”}, (\mathcal{S}_P(\mathfrak{C}))_{P \in A}\}$. If $m_1, \ldots, m_n$ is the sequence of messages sent by the signers to T, then $\delta(m_1, \ldots, m_n)$ is T's response for the last signer in the sequence. The function will be stated formally in Definition 10.

We denote the resolve protocol in the following by $\mathit{Res}(\mathfrak{C}, v)$. The signer initiating $\mathit{Res}(\mathfrak{C}, v)$ is $r(v)$. He sends the list of messages assigned to the vertices in his pre-knowledge $K_{\prec}(v)$, prepended by $\wp_{r(v)}(\mathfrak{C}, v, r(v), \mathrm{T})$, to T. This demonstrates that $r(v)$ has executed all receive events causally preceding v. We denote $r(v)$'s message for T by p_v:

$$p_v = \big(\wp_{r(v)}(\mathfrak{C}, v, r(v), \mathrm{T}), (\mu(w, v'))_{(w,v') \in K_{\prec}(v)}\big) \tag{1}$$

The promise $\wp_{r(v)}(\mathfrak{C}, v, r(v), \mathrm{T})$, which is the first element of p_v, is used by T to extract the contract $\mathfrak{C}$, to learn at which step in the protocol $r(v)$ claims to be, and to create a signature on behalf of $r(v)$ when necessary. All messages received from the signers are stored. T performs a deterministic decision procedure, shown in Algorithm 1, on the received message and existing stored messages and *immediately* sends back “abort” or the list of signatures $(\mathcal{S}_P(\mathfrak{C}))_{P \in A}$.

Our decision procedure is based on [10, 17]. The input to the algorithm consists of a message m received by the T from a signer and state information which is maintained by T. T extracts the contract and the DAG MPCS protocol specification from m. For each contract $\mathfrak{C}$, T maintains the following state information. A list $\mathit{Evidence}_{\mathfrak{C}}$ of all messages received from signers, a set $I_{\mathfrak{C}}$ of vertices the signers contacted T from, a set $\mathit{Dishonest}_{\mathfrak{C}}$ of signers considered to be dishonest, and the last decision made $\mathit{decision}_{\mathfrak{C}}$. If T has not been contacted by any signer regarding contract $\mathfrak{C}$, then $\mathit{decision}_{\mathfrak{C}} =$ “abort”. Else, $\mathit{decision}_{\mathfrak{C}}$ is equal to “abort” or the list $(\mathcal{S}_Q(\mathfrak{C})_{Q \in A})$ of signatures on $\mathfrak{C}$, one by each signer.

T verifies that the request is legitimate in that the received message m is valid and the requesting signer P is not already considered to be dishonest. If these preliminary checks pass, the message is appended to $\mathit{Evidence}_{\mathfrak{C}}$. This is described in Algorithm 1 in lines 1 through 9. The main part of the algorithm, starting at line 10, concerns the detection of signers who have continued the main protocol execution after executing the resolve protocol. If P has not received a promise from every other signer in the protocol (i.e. the if clause in line 10 is not

satisfied), then T sends back the last decision made (line 17). This decision is an "abort" token unless T has been contacted before and decided to send back a signed contract. If P has received a promise from every other signer, T computes the set of dishonest signers (lines 11 through 13) by adding to it every signer which has carried out the resolve protocol, but can be seen to have continued the protocol execution (line 12) based on the evidence the TTP has collected. If P is the only honest signer that has contacted T until this point in time, the decision is made to henceforth return a signed contract.

Algorithm 1: TTP decision procedure δ_0

input : $m, r, \mathit{decision}_{\mathfrak{C}}, \mathit{Evidence}_{\mathfrak{C}}, I_{\mathfrak{C}}, \mathit{Dishonest}_{\mathfrak{C}}$
output: $r, \mathit{decision}_{\mathfrak{C}}, \mathit{Evidence}_{\mathfrak{C}}, I_{\mathfrak{C}}, \mathit{Dishonest}_{\mathfrak{C}}$

1 **if** $m \neq (\wp_P(\mathfrak{C}, v, P, \mathrm{T}), \mathit{list})$ **then**
2 $\quad r =$ "abort";
3 $\quad$ **return output**;
4 **if** $P \in \mathit{Dishonest}_{\mathfrak{C}} \vee \forall w \in V : m \neq p_w \vee \exists w' \in I_{\mathfrak{C}} : P = r(w')$ **then**
5 $\quad \mathit{Dishonest}_{\mathfrak{C}} := \mathit{Dishonest}_{\mathfrak{C}} \cup \{P\}$;
6 $\quad r =$ "abort";
7 $\quad$ **return output**;
8 $I_{\mathfrak{C}} := I_{\mathfrak{C}} \cup \{v\}$;
9 $\mathit{Evidence}_{\mathfrak{C}} := (\mathit{Evidence}_{\mathfrak{C}}, m)$;
10 **if** $v \notin \mathit{InitSet}(\mathcal{P})$ **then**
11 $\quad$ **for** $w \in I_{\mathfrak{C}}$ **do**
12 $\quad\quad$ **if** $w \prec (w', x) \in K_{\prec}(I_{\mathfrak{C}}) \wedge r(w') = r(w)$ **then**
13 $\quad\quad\quad \mathit{Dishonest}_{\mathfrak{C}} := \mathit{Dishonest}_{\mathfrak{C}} \cup \{r(w)\}$;
14 $\quad$ **if** $\forall w \in I_{\mathfrak{C}} : r(w) \notin \mathit{Dishonest}_{\mathfrak{C}} \implies r(w) = P$ **then**
15 $\quad\quad \mathit{decision}_{\mathfrak{C}} := (\mathcal{S}_Q(\mathfrak{C}))_{Q \in A}$;
16 $r = \mathit{decision}_{\mathfrak{C}}$;
17 **return output**;

Definition 10. *Let $\mathcal{P} = (V, E, r, \mu, \delta)$ be a DAG MPCS protocol specification and δ_0 the TTP decision procedure from Algorithm 1. Then $\delta : M^* \to M$ is defined for $m_1, \ldots, m_n \in M$ by*

$$\delta(m_1, \ldots, m_n) = \pi_1(\delta_1(m_1, \ldots, m_n)),$$

where π_1 is the projection to the first coordinate and δ_1 is defined inductively by

$$\delta_1() = (\text{“abort”}, \text{“abort”}, \emptyset, \emptyset, \emptyset)$$
$$\delta_1(m_1, \ldots, m_n) = \delta_0(m_n, \delta_1(m_1, \ldots, m_{n-1})).$$

Thus the δ function represents the response of the TTP in the $\mathit{Res}(\mathfrak{C}, v)$ protocol for all executions of $\mathcal{P}$.

5.3 Fairness

We say that a DAG MPCS protocol execution is fair for signer P if one of the following three conditions is true: (i) No signer has received a signature of P; (ii) P has received signatures of all other signers; (iii) P has not received an "abort" token from the TTP. The last condition allows an execution to be fair as long as there is a possibility for the signer to receive signatures of all other signers.

The key problem in formalizing these conditions is to capture under which circumstances the TTP responds with an "abort" token to a request by a signer. The TTP's response is dependent on the decision procedure which in turn depends on the order in which the TTP is contacted by the signers. Since the decision procedure is deterministic, it follows that the δ function can be determined for every execution $\rho = [s_0, \alpha_1, s_1, \ldots, s_n]$ by considering the pre-knowledge of the vertices from which the *exit* transitions are taken. Abusing notation, we will write $\delta(\rho)$ instead of $\delta(m_1, \ldots, m_k)$ where m_i are the messages sent to the TTP at the i-th *exit* transition in the execution.

Definition 11. *Let* T *be the TTP. An execution* $\rho = [s_0, \alpha_1, \ldots, s_n]$ *of* $\mathcal{P}$ *is* fair *for signer* P *if one of the following conditions is satisfied:*

1. *P has not sent a signature and no signer has received signatures from* T*.*

$$\delta(\rho) = \textit{"abort"} \wedge \forall (v, w) \in s_n : r(v) = P, r(w) \neq P \Longrightarrow v \notin \mathit{SigSet}(\mathcal{P})$$

2. *P has received signatures from all other signers.*

$$\exists v \in s \cap \mathit{EndSet}(\mathcal{P}) : r(v) = P$$

3. *P has not received an "abort" token from* T*.*

$$\exists (v, w) \in s : r(v) = P \wedge r(w) = \mathrm{T} \Rightarrow \delta([s_0, \ldots, s_k, \mathit{exit}, s_k \cup \{(v, w)\}]) \neq \textit{"abort"}$$

If none of these conditions are satisfied, the execution is unfair for P.

Definition 12. *An MPCS protocol specification $\mathcal{P}$ is said to be fair, if every execution ρ of $\mathcal{P}$ is fair for all signers that are honest in ρ.*

5.4 Sufficient and necessary conditions

By the TTP decision procedure, T returns an "abort" token if contacted from a vertex $v \in \mathit{InitSet}(\mathcal{P})$. Thus a necessary condition for fairness is that an honest signer executes all steps of the initial set causally before all steps of the signing set that are not in the end set:

$$\forall v \in \mathit{InitSet}(\mathcal{P}), w \in \mathit{SigSet}(\mathcal{P}) \setminus \mathit{EndSet}(\mathcal{P}) : r(v) = r(w) \implies v \prec w \quad (2)$$

If P contacts T from a vertex $v \notin \mathit{InitSet}(\mathcal{P})$, then T responds with an "abort" token if it has already issued an "abort" token to another signer who is not in the set $\mathit{Dishonest}_{\mathfrak{C}}$. This condition can be exploited by a group of dishonest signers

in an *abort chaining attack* [16]. The following definition states the requirements for a successful abort chaining attack. For ease of reading, we define the predicate $\mathrm{hon}(v, I)$. The predicate is true if there is no evidence (pre-knowledge) at the vertices in I that the signer $r(v)$ has sent a message at or causally after v:

$$\mathrm{hon}(v, I) \equiv \neg\exists (x, y) \in K_{\prec}(I) : v \prec (x, y) \wedge r(v) = r(x)$$

This is precisely the criterion used by T to verify honesty in Algorithm 1, line 12.

Definition 13. *Let $\mathfrak{C}$ be a contract and $l \leq |A|$. A sequence $(v_1, \ldots, v_l \mid s)$ over V is called an abort-chaining sequence (AC sequence) for $\mathcal{P}$ if the following conditions hold:*

1. *Signer $r(v_1)$ receives an abort token: $v_1 \in \mathit{InitSet}(\mathcal{P})$*
2. *No signer contacts* T *more than once: $\forall_{i \neq j}\ r(v_i) \neq r(v_j)$*
3. *The present and previous signer to contact* T *are considered honest by* T*:*

$$\forall i \leq l : \mathrm{hon}(v_i, \{v_1, \ldots, v_i\}) \wedge \mathrm{hon}(v_{i-1}, \{v_1, \ldots, v_i\})$$

4. *The last signer to contact* T *has not previously received all signatures:*

$$\forall v \prec v_l : r(v) = r(v_l) \implies v \notin \mathit{EndSet}(\mathcal{P})$$

5. *The last signer to contact* T *has sent a signature before contacting* T *or in a parallel thread:*

$$s \in \mathit{SigSet}(\mathcal{P}) \setminus \mathit{EndSet}(\mathcal{P}) : r(s) = r(v_l) \wedge v_l \not\preceq s$$

The AC sequence represents the order in which signers execute the resolve protocol with T. A vertex v_i in the sequence implies an exit transition via the edge (v_i, v_{T}) in the protocol execution. An abort chaining attack must start at a step in which T has no choice but to respond with an abort token due to lack of information. Condition 1 covers this. Each signer may run the resolve protocol at most once. This is covered by Condition 2. To ensure that T continues to issue "abort" tokens, Condition 3 requires that there must always be a signer which according to T's evidence has not continued protocol execution after contacting T. To complete an abort chaining attack, there needs to be a signer which has issued a signature (Condition 5), but has not received a signature (Conditions 4 and 5) and will not receive a signed contract from T because there is an honest signer (by Condition 3) which has received an "abort" token.

It is not surprising (but nevertheless proven in the appendix) that a protocol with an AC sequence is unfair. However, the converse is true, too.

Theorem 1. *Let $\mathcal{P}$ be a DAG MPCS protocol. Then $\mathcal{P}$ is fair if and only if it has no AC sequences.*

The proof of this and the following theorems is given in the appendix.

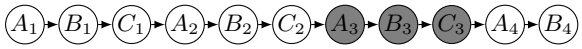

(a) A three-party MPCS protocol from a signing sequence [10].

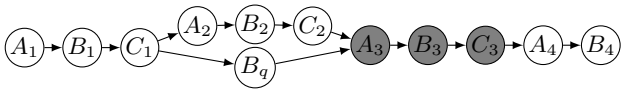

(b) Adding a vertex

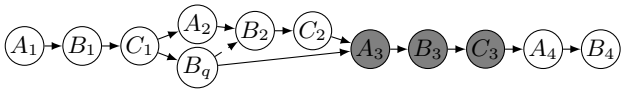

(c) Adding an ε edge.

Fig. 3: Skeletal graphs of fair protocols (a, c) and an unfair protocol (b).

5.5 Fairness criteria

Theorem 1 reduces the verification of fairness from analyzing all executions to verifying that there is no AC-sequence (Definition 13). This, however, is still difficult to verify in general. The following two results are tools to quickly assess fairness of DAG MPCS protocols. The first is an unfairness criterion and the second is a fairness criterion for a large class of DAG MPCS protocols.

The following theorem states that in a fair DAG MPCS protocol, the union of paths from the initial set to every vertex $v \in \mathit{SigSet}(\mathcal{P})$ must contain all permutations of all signers (other than $r(v)$) as subsequences. In the class of linear MPCS protocols, considered in [10], this criterion was both necessary and sufficient. We show in Example 2 below that this criterion is not sufficient for fairness of DAG MPCS protocols.

For $I \subseteq V$, $v \in V$, we denote by $\mathrm{path}(I, v) = \{(v_1, \ldots, v_k) \in V^* \mid v_1 \in I, v_k = v, \forall_{1 \le i < k} : (v_i, v_{i+1}) \in E\}$ the set of all directed paths from a vertex in I to v. If $p = (v_1, \ldots, v_k)$ is a sequence of vertices, we denote by $r(p) = (r(v_1), \ldots, r(v_k))$ the corresponding sequence of signers. The sequences of signers corresponding to the paths from I to v is denoted by $\mathrm{seq}(I, v) = \{r(p) \in A^* \mid p \in \mathrm{path}(I, v)\}$.

Theorem 2. *Let $k = |A|$. Let $\mathcal{P}$ be an optimistic fair DAG MPCS protocol,*

$$I = \{v \in \mathit{InitSet}(\mathcal{P}) \mid (v \prec w \wedge r(v) = r(w)) \Rightarrow w \notin \mathit{InitSet}(\mathcal{P})\}\,.$$

If $v \in \mathit{SigSet}(\mathcal{P})$, then for every permutation $(P_1, \ldots, P_{k-1})$ of signers in $A \setminus \{r(v)\}$, there exists a sequence in $\mathrm{seq}(I, v)$ which contains $(P_1, \ldots, P_{k-1})$ as a (not necessarily consecutive) subsequence.

The converse of the theorem is not true as the following example shows. In particular, this example demonstrates that the addition of a vertex to a fair DAG MPCS protocol does not necessarily preserve fairness.

Example 2. The protocol in Figure 3a is fair by the results of [10]. By Theorem 2, for every vertex $v \in \mathit{SigSet}(\mathcal{P})$ every permutation of signers in $A \setminus \{P\}$ occurs as a subsequence of a path in $\mathrm{seq}(I, v)$. The protocol in Figure 3b is obtained by adding the vertex B_q as a parallel thread of signer B. Thus the permutation property on the set of paths is preserved, yet the protocol is not fair: An AC sequence is $(B_q, C_3, A_4|A_3)$. The vertex B_q is in $\mathit{InitSet}(\mathcal{P})$, the evidence presented to the TTP at C_3 includes the vertices causally preceding C_2, thus B is considered to be honest. The evidence presented by signer A at A_4 are the vertices causally preceding A_3 proving that B is dishonest, but C is honest. Thus A has sent a signature at A_3 but will not receive signatures from B and C.

If a protocol has no in-role parallelism, then the converse of Theorem 2 is true. Thus we have a simple criterion for the fairness of such protocols.

Theorem 3. *Let $\mathcal{P}$ be an optimistic DAG MPCS protocol without in-role parallelism. Let*

$$I = \{v \in \mathit{InitSet}(\mathcal{P}) \mid (v \prec w \wedge r(v) = r(w)) \Rightarrow w \notin \mathit{InitSet}(\mathcal{P})\}\,.$$

If all paths from I to $v \in \mathit{SigSet}(\mathcal{P})$ contain all permutations of $A \setminus \{r(v)\}$ then $\mathcal{P}$ is fair for $r(v)$.

Example 3. By adding a causal edge between vertex B_q and vertex B_2 of the protocol in Figure 3b, as shown in Figure 3c, we obtain again a fair protocol.

6 Protocols

In this section we illustrate the theory and results obtained in the preceding sections by proving optimality results and constructing a variety of protocols.

6.1 Minimal complexity

We prove lower bounds for the two complexity measures defined in our model, viz. parallel and message complexity.

Theorem 4. *The minimal parallel complexity for an optimistic fair DAG MPCS protocol is $n + 1$, where n is the number of signers in the protocol.*

Proof. By Theorem 2, every permutation of signers in the protocol must occur as a subsequence in the set of paths from a causally last vertex in the initial set to a vertex in the signing set. Since a last vertex v in the initial set must have a non-empty knowledge $K(v)$, there must be a message edge causally preceding v. There are at least $n - 1$ edges in the path between the vertices associated with the n signers in a permutation and there is at least one message edge outgoing from a vertex in the signing set. Thus a minimal length path for such a protocol must contain $n + 1$ edges.

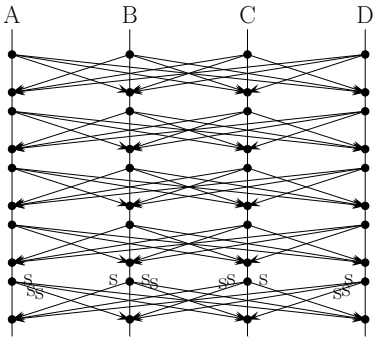

Fig. 4: A minimal 4-party fair broadcasting protocol.

The minimal parallel complexity is attained by the broadcast protocols of Baum-Waidner and Waidner [2]. An example is shown in Figure 4.

Theorem 5. *The minimal message complexity for an optimistic fair DAG MPCS protocol is* $\lambda(n) + 2n - 3$*, where* n *is the number of signers in the protocol and* $\lambda(n)$ *is the length of the shortest sequence which contains all permutations of elements of an* n*-element set as subsequences.*

The minimal message complexities for $2 < n < 8$ *are* $n^2 + 1$*. The minimal message complexities for* $n \geq 10$ *are smaller or equal to* n^2*.*

Note that while broadcasting protocols have a linear parallel complexity, they have a cubic message complexity, since in each of the $n + 1$ rounds each of the n signers sends a message to every other signer. Linear protocols on the other hand have quadratic minimal message and parallel complexities. In the following we demonstrate that there are DAG protocols which attain a linear parallel complexity while maintaining a quadratic message complexity.

6.2 Protocol constructions

Single contractor, multiple subcontractors. A motivation for fair MPCS protocols given in [10] is a scenario where a single entity, here referred to as a contractor, would like to sign k contracts with k independent companies, in the following referred to as subcontractors. The contractor has an interest in either having all contracts signed or to not be bound by any of the contracts. The subcontractors have no contractual obligations towards each other. It would therefore be advantageous if there is no need for the subcontractors to directly communicate with each other.

The solutions proposed in [10] are linear protocols. Their message and parallel complexities are thus quadratic. Linear protocols can satisfy the requirement that subcontractors do not directly communicate with each other only by greatly increasing the message and parallel complexities.

The protocol we propose here is a DAG, its message complexity is $2(n + 1)(n-1)$ and its parallel complexity is $2n+2$ for n signers. It therefore combines

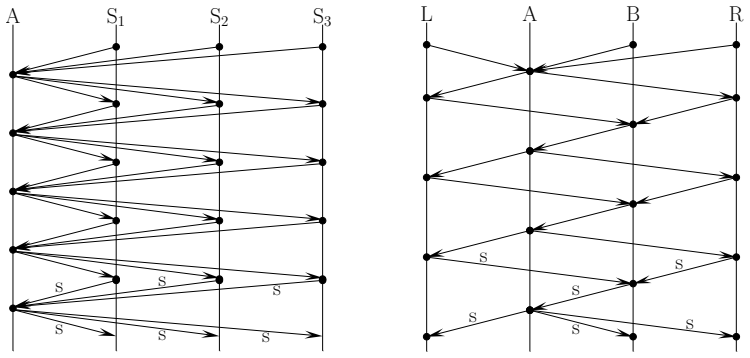

(a) A single contractor and three subcontractors.

(b) Two joint subcontractors.

Fig. 5: Two examples of novel, fair DAG MPCS protocols.

the low parallel complexity typically attained by broadcasting protocols with the low message complexity of linear protocols. Additionally, the protocol proposed does not require any direct communication between subcontractors.

Figure 5a shows a single contractor with three subcontractors. The protocol can be subdivided into five rounds, one round consisting of the subcontractors sending a message to the contractor followed by the contractor sending a message to the subcontractors. In the first four rounds promises are sent, in the final round signatures are sent. The protocol can be easily generalized to more than three subcontractors. For every subcontractor added, one extra round of promises needs to be included in the protocol specification.

The protocol is fair by Theorem 3. The MSC shown in Figure 5a resembles the skeletal graph from which it was built. The message contents can be derived by computing the full graph according to Condition 2 of Definition 9. The result is as follows. In each round of the protocol, each of the subcontractors sends to the contractor a promise for the contractor and for each of the other subcontractors. The contractor then sends to each of the subcontractors all of the promises received and his own promise. The final round is performed in the same manner, except that promises are replaced by signatures.

Two contractors with joint subcontractors. Figure 5b shows a protocol where two contractors want to sign a contract involving two subcontractors. The subcontractors are independent of each other.

After the initial step, where all signers send a promise to the first contractor A, there are three protocol rounds, one round consisting of the contractor A sending promises to the two subcontractors L and R in parallel which in turn send promises to the second contractor B. A new round is started with the second contractor sending the promises received with his own promise to contractor A.

This protocol, too, can be generalized to several independent subcontractors. For every subcontractor added, one extra protocol round needs to be included in the protocol specification and each protocol step of the subcontractors executed analogously.

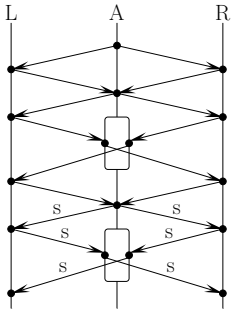

Fig. 6: In-role parallelism.

Parallelism within a role. Figure 6 shows an example of a subcontracting protocol with in-role parallelism for the contractor role. The contractor initiates the protocol. In the indicated parallel phase, the contractor may immediately forward a promise by one of the subcontractors along with his own promise to the other subcontractor without waiting for the latter subcontractor's promise. The same is true in the signing phase. The fairness property for this protocol has been verified with a tool (described in Appendix A) which tested fairness for each signer in all possible executions.

7 Conclusion

We have identified fair subcontracting as a challenging new problem in the area of multi-party contract signing. We have made first steps towards solving this problem by introducing DAG MPCS protocols and extending existing fairness results from linear protocols to DAG protocols. For three typical subcontracting configurations we propose novel DAG MPCS protocols that perform well in terms of message complexity and parallel complexity. Fairness of our protocol schemes follows directly from our theoretical results and we have verified it for concrete protocols with our automatic tool.

There are a number of open research questions related to fair subcontracting that we haven't addressed. We mention two. The first concerns the implementation of multi-contracts. In our current approach we consider a single abstract contract shared by all parties. However, in practice such a contract may consist of a number of subcontracts that are accessible to the relevant signers only. How to cryptographically construct such contracts and what information these contracts should share is not evident. Second, a step needs to be made towards putting our results into practice. Given the application domains identified in

this paper, we must identify the relevant signing scenarios and topical boundary conditions in order to develop dedicated protocols for each application area.

Acknowledgement

We thank Barbara Kordy for her many helpful comments on this paper.

References

1. N. Asokan. *Fairness in electronic commerce*. PhD thesis, Univ. of Waterloo, 1998.
2. B. Baum-Waidner and M. Waidner. Optimistic asynchronous multi-party contract signing. Research Report RZ 3078 (#93124), IBM Zurich Research Laboratory, Zurich, Switzerland, November 1998.
3. B. Baum-Waidner and M. Waidner. Round-optimal and abuse free optimistic multi-party contract signing. In *Automata, Languages and Programming — ICALP 2000*, volume 1853 of *LNCS*, pages 524–535. Springer, July 2000.
4. R. Chadha, S. Kremer, and A. Scedrov. Formal analysis of multi-party contract signing. In *CSFW'04*, page 266, Washington, DC, USA, 2004. IEEE.
5. S. Even and Y. Yacobi. Relations among public key signature systems. Technical Report 175, Computer Science Dept., Technion, Haifa, Isreal, March 1980.
6. J. Garay, M. Jakobsson, and P. MacKenzie. Abuse-free optimistic contract signing. In *CRYPTO'99*, volume 1666 of *LNCS*, pages 449–466. Springer, Aug. 1999.
7. J. A. Garay and P. D. MacKenzie. Abuse-free multi-party contract signing. In *13th Int. Symp. Distr. Computing*, volume 1693 of *LNCS*, pages 151–165. Springer, 1999.
8. P. Karaenke and S. Kirn. Towards model checking & simulation of a multi-tier negotiation protocol for service chains. In *AAMAS 2010*, pages 1559–1560. Int. Found. for Autonomous Agents and Multiagent Systems, 2010.
9. E. Katok and V. Pavlov. Fairness in supply chain contracts: a laboratory study. *J. of Operations Management*, 31:129–137, 2013.
10. B. Kordy and S. Radomirović. Constructing optimistic multi-party contract signing protocols. In *CSF 2012*, pages 215–229. IEEE Computer Society, 2012.
11. S. Kraus. Automated negotiation and decision making in multi-agent environments. In *ACM multi-agent systems and applications*, pages 150–172, 2001.
12. H. Krishnan and R. Winter. The economic foundations of supply chain contracting. *Foundations and Trends in Technology, Information and Operations Management*, 5(3-4):147–309, 2012.
13. K. Lu, R. Yahyapour, E. Yaqub, and C. Kotsokalis. Structural optimisation of reduced ordered binary decision diagrams for SLA negotiation in IaaS of cloud computing. In *ICSOC 2012*, volume 7636 of *LNCS*, pages 268–282. Springer, 2012.
14. S. Mauw and S. Radomirović. Generalizing Multi-party Contract Signing. In *POST 2015*, 2015.
15. S. Mauw, S. Radomirović, and M. T. Dashti. Minimal message complexity of asynchronous multi-party contract signing. In *CSF'09*, pages 13–25. IEEE, 2009.
16. A. Mukhamedov and M. Ryan. Improved multi-party contract signing. In *Financial Cryptography*, volume 4886 of *LNCS*, pages 179–191. Springer, 2007.
17. A. Mukhamedov and M. D. Ryan. Fair multi-party contract signing using private contract signatures. *Inf. Comput.*, 206(2-4):272–290, 2008.

18. S. Radomirović. A construction of short sequences containing all permutations of a set as subsequences. *Electronic J. of Combinatorics*, 19(4):Paper 31, 11 pp., 2012.
19. M. Schunter. *Optimistic Fair Exchange*. Phd thesis, Universität des Saarlandes, 2000.
20. R. Seifert, R. Zequiera, and S. Liao. A three-echelon supply chain with price-only contracts and sub-supply chain coordination. *Int. J. of Production Economics*, 138:345–353, 2012.
21. E. Yaqub, P. Wieder, C. Kotsokalis, V. Mazza, L. Pasquale, J. Rueda, S.G.Gómez, and A. Chimeno. A generic platform for conducting SLA negotiations. In *Service Level Agreements for Cloud Computing*, pages 187–206. Springer, 2011.
22. Y. Zhang, C. Zhang, J. Pang, and S. Mauw. Game-based verification of contract signing protocols with minimal messages. *Innovations in Systems and Software Engineering*, 8(2):111–124, 2012.

A DAG MPCS Verification Tool

We have developed a prototype tool in Python 2 that model checks a skeletal protocol graph for the fairness property (Definition 11) in the execution model defined in Section 4.1. The tool, along with specifications for the protocols presented in this paper, is available at `http://people.inf.ethz.ch/rsasa/mpcs`.

The tool's verification procedure works directly on the execution model and the TTP decision procedure (Algorithm 1). It therefore provides evidence for the correctness of the protocols shown in Section 6, independent of the fairness proofs given in this paper.

The verification is performed as follows. For each specified signer, the tool analyzes a set of executions in which the signer is honest and all other signers dishonest. The tool does not analyze all possible executions. It starts the analysis from the state where all promises of dishonest signers have been sent, but no protocol step has been performed by the honest signer. By analyzing this type of executions only, we do not miss any attacks, because the honest signers' fairness is not invalidated until he has sent a signature. In this reduced set of executions, the dishonest signers retain the possibility to contact the TTP from any of their vertices and all these possibilities are explored by the tool.

We note that the same type of verification could be achieved with an off-the-shelf model checker and we would expect better performance in such a case. However, the code complexity and room for error when encoding a given protocol and TTP decision procedure in a model checker's input language is comparable to the code complexity of this self-contained tool.

B Technical Details and Proofs

B.1 Technical Details

Parallelism within a role The MPCS protocols designed in this work allow for parallelism during the execution of the protocol. The specification language allows even for parallel threads to occur within a signer role. This allows us

to model the case where a signer role represents multiple branches of the same entity. A signature issued by any branch represents the signature of the entire entity. We expect that the signing processes across branches are not easily synchronizeable with each other. Such parallelism can be implemented in multiple ways. We discuss the various options and explain the choices made for this paper.

The first decision to be made is whether parallel threads of a signer role should be assumed to have shared knowledge. In this paper, we choose the weaker assumption: memory for a signer's parallel threads is local to the threads. This is in accordance with our expectation that parallel-threads are not easily synchronizeable and allows us, for instance, to specify and analyze protocols in which representatives of a signing entity can independently carry out parallel protocol steps without the need to communicate and synchronize their combined knowledge. Causal dependence between two actions of a signer is explicitly indicated in the protocol specification.

This design decision leads to three options for handling protocol failures.

1. All threads of a signer immediately synchronize and stop executing whenever any of the threads intends to issue a resolve request to the TTP. A designated resolution process per signer will be required to continuously schedule all threads and take care of the interaction with the TTP.
2. Threads of a signer contact the signer's designated resolution process only when they intend to issue a resolve request. The resolution process will take of contacting the TTP (only once per signer) and distributing the TTP's reply upon request of the threads.
3. Threads of a signer are considered fully independent. A signer's threads are not orchestrated. The TTP may take into account that two requests can originate from the same signer, but from different (causally not related) threads.

In this paper we adopt the second option, which keeps the middle between the fully synchronized and fully desynchronized model. This will on the one hand allow for independent parallel execution of the threads and on the other hand minimize the impact of the signer's threading on the TTP's logic. From an abstract point of view, one could even argue that the second and third option are equivalent if we consider the signer's designated resolution processes just as part of a distributed TTP. We assume that the communication between a thread and the designated resolution process is resilient.

The class of DAG MPCS protocols The class of DAG MPCS protocols defined in Section 5 is restricted by condition 2 of Definition 9. It requires that every signer P sends a message to all subsequent, causally following signers occurring before signer P's next step. While there are fair DAG MPCS protocols which do not belong to this restricted class, such protocols are not going to have a lower communication complexity. The reason for this is that each message received by a signer serves as evidence for the TTP that the sender has executed the protocol up to a certain step. Skipping such a message thus lengthens the protocol, because the evidence is available only at a later vertex.

Furthermore, the restriction simplifies the reasoning about fairness in that causal precedence $v \prec w$ between vertices v, w is enough to guarantee that there is a message sent from signer $r(v)$ to signer $r(w)$ at some point between the execution of v and the execution of w. Finally, it also permits one to design, characterize, and represent protocols using skeletal graphs rather than full graphs as displayed in Figure 2.

B.2 Proofs

The set $\text{maxset}(S) = \{v \in S \mid \forall w \in V : v \prec w \Rightarrow w \notin S\}$ is the set of vertices in S which do not have any causally following vertices in S and we will refer to it as the set of *maximal vertices* of S. Similarly, $\text{minset}(S) = \{v \in S \mid \forall w \in V : w \prec v \Rightarrow w \notin S\}$ is the set of vertices in S which do not have any causally preceding vertices in S and will be referred to as the set of *minimal vertices* of S.

Theorem. *If there exists an AC sequence for a DAG MPCS protocol, then the protocol is not fair.*

Proof. Conditions 1 through 3 imply that the TTP decision procedure leads to an "abort" token for the last signer to contact the TTP. The remaining two conditions imply that the last signer has sent a signature, but not received a signature.

To complete the proof, we need to construct an execution in which the exit transitions occur in the order indicated by the AC sequence and signer $r(v_l)$ is honest. Let $(v_1, \ldots, v_l \mid v)$ be an AC sequence. For each vertex v_i in the AC sequence, let $\overline{V_i}$ be the causal closure of $\{v, v_i\}$ in $V \cup E$. Note that the union of causally closed sets is causally closed. Let ρ_i be the sequence of transitions $\bigcup_{j<i} \overline{V_j} \overset{\alpha}{\rightsquigarrow} \ldots \overset{\alpha'}{\rightsquigarrow} \bigcup_{j \leq i} \overline{V_j}$ without exit transitions and such that all states are causally closed.

For $1 \leq i \leq l$ and $\rho_i = [s_0, \alpha_1, \ldots, s_k]$, let $\rho'_i = [s_0 \cup \{(v_1, v_\mathrm{T}), \ldots, (v_{i-1}, v_\mathrm{T})\}, \alpha_1, \ldots, s_k \cup \{(v_1, v_\mathrm{T}), \ldots, (v_{i-1}, v_\mathrm{T})\}, \mathit{exit}]$. That is, ρ'_i is equal to ρ_i, except for an additional exit transition $s \overset{\mathit{exit}}{\rightsquigarrow} s \cup \{(v_i, v_\mathrm{T})\}$ and additional exit edges in all states which stem from exit transitions added to $\rho'_1, \ldots, \rho'_{i-1}$. Finally, for $\rho'_l = [s_0, \alpha_1, \ldots, s_k, \mathit{exit}]$, let $\rho''_l = [s_0 \setminus \{v_l\}, \alpha_1, \ldots, s_k \setminus \{v_l\}, \mathit{exit}, s_k \cup \{(v_l, v_\mathrm{T})\} \setminus \{v_l\}]$.

Then $\rho = \rho'_1 \cdots \rho'_{l-1} \cdot \rho''_l$ is an execution in which signer $r(v_l)$ is honest, since the restricted execution is by construction causally closed in all states before the last state and the single exit transition occurs in the last transition.

Unfairness for $r(v_l)$ follows since $r(v_l)$ has sent a signature at v, not received all signatures from the other signers and received an "abort" from the TTP.

Proof (of Lemma 2). Let $\rho = [s_0, \alpha_1, s_1, \ldots, \alpha_l, s_l]$ be an execution of $\mathcal{P}$. It is sufficient to show that if ρ is closed, it contains all send events exactly once. According to Definition 3, we know that for every $i \in \{0, \ldots, l-1\}$ we have $s_i \overset{\alpha_{i+1}}{\rightsquigarrow} s_{i+1} \implies s_i \neq s_{i+1}$. This implies that, in any execution, each step of the protocol (in particular every send event) can be executed at most once.

Furthermore, if ρ is closed, the third condition from Definition 6 implies that, every send event has already occurred in ρ. Otherwise, there exists $e \in E$ such that $\mu(e) = \mathit{send}$ and ρ can be extended to $\rho \cdot [\mathit{send}, s_l \cup \{e\}] \in \mathrm{Exe}(\mathcal{P})$, which contradicts the closedness of ρ.

Lemma 3. *Let $\mathcal{P}$ be an optimistic fair DAG MPCS protocol. Let v, v', v'' be pairwise distinct vertices assigned to the same signer such that*

1. *$v \in \mathit{SigSet}(\mathcal{P}) \setminus \mathit{EndSet}(\mathcal{P})$,*
2. *v'' is a maximal common ancestor of v and v', i.e., $v'' \prec w \prec v, v' \Rightarrow r(w) \neq r(v)$, and*
3. *for every signer $P \neq r(v)$ there exists a vertex $w \succ v''$ with $r(w) = P$.*

Then for every permutation $(P_1, \ldots, P_{k-1})$ of signers in $A \setminus \{r(v)\}$, there exists a sequence in $\mathrm{seq}(I, v'')$ which contains $(P_1, \ldots, P_{k-1})$ as a (not necessarily consecutive) subsequence.

Proof. Suppose there exists a permutation $(P_1, \ldots, P_{k-1})$ of signers in $A \setminus \{r(v)\}$ which is not a subsequence of any sequence in $\mathrm{seq}(I, v'')$. We construct an AC sequence as follows. Let V_1 be the set of all vertices of P_1 in I. For $i > 1$, let V_i be the minset of all vertices of P_i which causally follow a vertex of V_{i-1}, i.e. $V_i = \mathrm{minset}(\{w \in V \mid r(w) = P_i \wedge \exists w' \in V_{i-1} : w' \prec w\})$. Since for every signer there exists a vertex which causally follows v'', it follows that for some j, there exists a vertex $v_j \in V_j$ with $v'' \prec v_j$. (Else we have contradiction to the assumption that $(P_1, \ldots, P_{k-1})$ is a missing permutation in $\mathrm{seq}(I, v'')$.) Thus we obtain a sequence $(v_1, \ldots, v_j, w|v)$, where $v'' \prec w \preceq v'$, $r(w) = r(v)$, which is an AC sequence.

Proof (of Theorem 2). It suffices to verify the statement for a subset of all vertices in $\mathit{SigSet}(\mathcal{P})$ by the following two facts: **Fact 1:** Let $v \in \mathit{SigSet}(\mathcal{P})$ be a causally earliest vertex of a signer from which a signature is sent, i.e. $\forall w \in \mathit{SigSet}(\mathcal{P}) : w \prec v \Rightarrow r(w) \neq r(v)$. If $\mathrm{seq}(I, v)$ contains all permutations of signers in $A \setminus \{r(v)\}$, then $\mathrm{seq}(I, w)$ contains all such permutations of signers for all $w \succ v$ with $r(w) = r(v)$. **Fact 2:** If $v \in \mathit{SigSet}(\mathcal{P})$ such that for every signer $P \in A \setminus \{r(v)\}$ there is a vertex $w \prec v$ for which $\mathrm{seq}(I, w)$ contains all permutations of signers in $A \setminus \{r(w)\}$, then $\mathrm{seq}(I, v)$ contains all permutations of signers in $A \setminus \{r(v)\}$.

Thus, we may assume that $v \in \mathit{SigSet}(\mathcal{P})$ is a causally earliest vertex of a signer from which a signature is sent (by Fact 1) and that $v \notin \mathit{SigSet}(\mathcal{P}) \setminus \mathit{EndSet}(\mathcal{P})$ (by Fact 2).

Since v is a causally earliest vertex of a signer from which a signature is sent, it follows by the fact that the protocol is optimistic that for every signer other than $r(v)$ there exists a vertex which causally follows v or that there exists another vertex v'' of signer $r(v)$ from which a signature is sent such that $v'' \not\preceq v$ and $v \not\preceq v''$. We consider these two cases separately.

1. For every signer other than $r(v)$, there exists a vertex which causally follows v.

We split this case into two separate subcases depending on whether there exists a vertex v' of signer $r(v)$ which causally follows v.

(a) $\exists v' \succ v : r(v') = r(v)$. Let $(P_1, \ldots, P_{k-1})$ be a permutation of signers in $A \backslash \{r(v)\}$ and suppose towards a contradiction that the permutation does not appear as a subsequence of any sequence in $\mathrm{seq}(I, v)$. We construct an AC sequence as follows. Let V_1 be the set of all vertices of P_1 in I. For $i > 1$, let V_i be the minset of all vertices of P_i which causally follow a vertex of V_{i-1}, i.e. $V_i = \mathrm{minset}(\{w \in V \mid r(w) = P_i \wedge \exists w' \in V_{i-1} : w' \prec w\})$.
Since for every signer there exists a vertex which causally follows v, it follows that for some j there exists a vertex $v_j \in V_j$ with $v \prec v_j$, else we have contradiction to the assumption that the permutation $(P_1, \ldots, P_{k-1})$ is not a subsequence of any sequence in $\mathrm{seq}(I, v)$.
By construction, there exists a vertex in V_{j-1} which causally precedes v_j and thus we obtain a sequence $(v_1, \ldots, v_j, v'|v)$ which is an AC sequence.

(b) $\neg\exists v' \succ v : r(v') = r(v)$.
Since the protocol is optimistic, there exists a vertex assigned to signer $r(v)$ such that $v' \in \mathit{EndSet}(\mathcal{P})$. Since $v \notin \mathit{EndSet}(\mathcal{P})$, it follows that v' is not causally related to v. By the remark preceding Lemma 3, there exists a common ancestor v'' or v and v' and v, v', v'' satisfy the hypothesis of the Lemma. Thus there exists a vertex w causally preceding v such that $\mathrm{seq}(I, w)$ contains all permutations of signers in $A \setminus \{r(v)\}$ and therefore $\mathrm{seq}(I, v)$ contains all such permutations.

2. There are causally unrelated vertices of signer $r(v)$ from which signatures are sent.
Let $v' \neq v$ be such a vertex. By Equation (2) in Section 5.4, there is a vertex w assigned to signer $r(v)$ which causally precedes all vertices of $r(v)$ which are in $\mathit{SigSet}(\mathcal{P})$. Let v'' be a maximal such vertex, i.e. for any vertex w' assigned to signer $r(v)$, there exists a vertex in $\mathit{SigSet}(\mathcal{P})$ of signer $r(v)$ which does not causally follow v''.
Since the protocol is optimistic, for every signer P in the protocol, there exists a vertex w'', $r(w'') = P$ which causally follows v''.
Then the vertices v, v', v'' satisfy the hypothesis of Lemma 3, thus there exists a vertex w causally preceding v such that $\mathrm{seq}(I, w)$ contains all permutations of signers in $A \backslash \{r(v)\}$ and therefore $\mathrm{seq}(I, v)$ contains all such permutations.

Proof (of Theorem 3). Suppose that the protocol is not fair. Consider a shortest AC sequence, $(v_1, \ldots, v_l|\_)$, $r(v_l) = r(v)$. Since the sequence is a shortest sequence, we have that $v_2 \notin \mathit{InitSet}(\mathcal{P})$, else $(v_2, \ldots, v_l|\_)$ would be a shorter AC sequence. Consider the permutation of signers $(P_1, \ldots, P_l)$ corresponding to the AC sequence, i.e. $P_i = r(v_i)$.

Let w_l be the unique vertex in

$$\mathrm{minset}(\{w \in \mathit{SigSet}(\mathcal{P}) \mid r(w) = r(v_l)\}).$$

Existence of a vertex in the set follows from the fact that the protocol is optimistic, uniqueness follows from the fact that there is no in-role parallelism, i.e.

the vertices assigned to a particular signer are totally ordered. By hypothesis, the set of paths from I to w_l contains all permutations of signers $A \setminus \{r(v)\}$. Let $u_1, \ldots, u_l$ be the vertices associated with one such permutation. Note that either $u_1 \in I$ or we can find $u_1' \in I$, $u_1' \prec u_1$ and $r(u_1') = u_1$. Thus we may assume $u_1 \in I$. We have $u_l \preceq w_l \prec v_l$. We also have $w_l \prec v_{l-1} \prec v_l$, else condition 3 for $(v_1, \ldots, v_l|\_)$ being an AC sequence (Definition 13) would be violated.

Thus we have $u_l \preceq w_l \prec v_{l-1} \prec v_l$. This forms the basis for the inductively constructed sequence $w_1, \ldots, w_l$: Given $w_{i+1}, \ldots, w_l$, satisfying $u_{i+1} \preceq w_{i+1} \prec v_i \prec v_{i+1}$, let w_i be the unique vertex in $\mathrm{maxset}(\{w \prec w_{i+1} \mid r(w) = r(v_i)\})$. Existence of a vertex in the set follows from $u_i \prec u_{i+1} \preceq w_{i+1}$ and uniqueness follows from the lack of in-role parallelism. By construction, $u_i \preceq w_i \prec v_i$. If $i > 1$, then we also have $u_i \preceq w_i \prec v_{i-1} \prec v_i$, else condition 3 for $(v_1, \ldots, v_l|\_)$ being an AC sequence (Definition 13) would be violated.

Thus, we have constructed a sequence $w_1, \ldots, w_l$ satisfying $u_1 \preceq w_1 \prec w_2 \prec v_1 \prec v_2$. This is not possible, since $r(u_1) = r(v_1)$ and $u_1 \prec v_1 \in \mathit{InitSet}(\mathcal{P})$, contradicting $u_1 \in I$.

Lemma 4. *Let $G = (V, E)$ be the DAG of a fair optimistic DAG MPCS protocol for two or more signers. Let $G' = (V', E')$, where $V' = V \setminus \{v_\mathrm{T}\}$ and $E' = E \setminus \{(v, w) \in E \mid v = v_\mathrm{T} \vee w = v_\mathrm{T}\}$, be the DAG obtained by removing the TTP vertex and corresponding edges. Then G' is a single connected component.*

Proof. Suppose there are more than one connected components in G'. Let $v \in \mathit{SigSet}(\mathcal{P})$ be a causally earliest vertex from which a signature is sent, i.e. $\forall w \in V : w \prec v \Rightarrow w \notin \mathit{SigSet}(\mathcal{P})$.

Let w be a vertex in the $\mathit{InitSet}(\mathcal{P})$ of a different connected component than v. We have two cases:

- $r(w) = r(v)$. Then $(w|v)$ is an AC sequence.
- $r(w) \neq r(v)$. Let v' be a vertex such that $r(v) = r(v')$ and $v' \not\preceq v$. Such a vertex exists, because the protocol is optimistic, thus there must be a vertex of signer $r(v)$ receiving a signature. But such a vertex cannot precede v, because v is a causally earliest vertex from which a signature is sent.
  Consider two cases:
  - $w \not\prec v'$: Then $(w, v'|v)$ is an AC sequence.
  - $w \prec v'$: Then w and v' are in the same connected component and v is in another connected component. If $v' \notin \mathit{InitSet}(\mathcal{P})$, then let $v'' \prec v'$ be such that $r(v') = r(v'')$ and $v'' \in \mathit{InitSet}(\mathcal{P})$. Else let $v'' = v'$.
    Then $(v''|v)$ is an AC sequence.

Proof (of Theorem 5). The minimal message complexity has been derived for optimistic fair linear protocols in [10, 15]. Since these protocols are a subset of DAG MPCS protocols we see that the same message complexity can be attained. We need to show that there are no optimistic DAG MPCS protocols with lower message complexity. By Theorem 2, every permutation of signers in the protocol must occur as a subsequence in the set of paths from a maximal vertex of the set of vertices of a signer in the initial set to a vertex in the signing set.

Consider any fair optimistic DAG MPCS protocol $\mathcal{P} = (V, E, r, \mu, \delta)$. Construct a linear DAG (V, E') by choosing any topologically sorted list $(v_1, \ldots, v_k)$ of the vertices in (V, E) and setting $E' = \{(v_i, v_{i+1}) | 1 \leq i \leq k\}$. Since all permutations of signers occur along the paths in the DAG (V, E) under the labelling $r : V \to A$, they also occur in the topologically sorted list $(v_1, \ldots, v_k)$ under the same labelling. Since the DAG is a single connected component by Lemma 4, the number of edges in E' is smaller or equal to the number of edges in E. Thus the message complexity of $\mathcal{P}$ is greater than or equal to the message complexity of a protocol based on the linear DAG (V, E').

The specific numbers for message complexity follow from [10, 18].