



This document is the original author manuscript of a paper submitted to an IFIP conference proceedings or other IFIP publication by Springer Nature. As such, there may be some differences in the official published version of the paper. Such differences, if any, are usually due to reformatting during preparation for publication or minor corrections made by the author(s) during final proofreading of the publication manuscript.

Protocol Scheduling

Kasper Dokter and Farhad Arbab

Centrum Wiskunde & Informatica, Amsterdam, Netherlands

Abstract. Interactions amongst different processes in concurrent software are governed by a protocol. The blocking I/O operations involved in a protocol may temporarily suspend the execution of some processes in an application. Scheduling consists of the allocation of available processors to the appropriate non-suspended processes in an application, such that some specified criteria (e.g., shortest execution time or highest throughput) are met. We use a generic, game-theoretic scheduling framework to find optimal non-preemptive schedules for an application. We then show how such schedules themselves can be encoded as protocols, which in our framework, can be composed with the original application protocol. The resulting composed protocol restricts the number of ready processes to the number of available processors, which enables standard preemptive schedulers of modern operating-systems to closely approximate the behavior and the performance of the optimal non-preemptive scheduler of the application. We evaluate our work by comparing the throughput of two versions of a cyclo-static dataflow network: one version with the usual protocol, and the other version with a restricted protocol.

1 Introduction

Context Reliable concurrent software is notoriously complex to develop. The main factor that contributes to this complexity is the vastness of the number of possible sequences of interactions amongst concurrent processes. Many such sequences of interactions lead to bad states. The protocol in a concurrent software aims to prevent occurrences of such sequences.

In most programming languages, a protocol manifests implicitly as relatively few operations performed on locks and buffers, intermixed with a manifold larger number of computational operations. Consequently, it is practically impossible to directly analyze the protocol and verify its correctness. Exogenous coordination languages (such as Reo [2–4]) define a protocol explicitly as the medium through which all application processes interact. As such, concurrency protocols expressed in these languages can be developed and analyzed in isolation.

Problem By construction, the protocol introduces dependencies amongst the different processes in an application. For example, the rate at which a producer process can fill a buffer of bounded capacity depends on the rate at which a consumer process drains it. For optimal performance of a concurrent software, the scheduler must take these dependencies into account.

However, operating systems schedulers are application-independent and remain oblivious to the dependency information inherent in a protocol, even if such information is available. At best, these schedulers detect consequential effects of a protocol, such as blocking on an I/O-operation or waiting for a lock.

Contribution In our current work, we refine the protocol of a concurrent application such that it enforces a custom schedule and derives an improved performance from the operating system scheduler. Our work demonstrates a close connection between a protocol and a schedule. Our protocol refinement uses the fact that the operating system scheduler executes only processes that are not blocked or waiting. We enforce a custom schedule by blocking all application processes, except for those that should run according to the application’s desired schedule. We block a process by prolonging existing blocking operations like I/O operations or waiting for a lock. Therefore, our approach assumes that the application’s desired custom schedule is non-preemptive.

Outline We represent (Section 2) a concurrent application as a work automaton [9], which is a finite-state machine together with a finite set of real-valued variables, called jobs, each of which measures the progress of a single process.

We synthesize non-preemptive schedules for concurrent applications using algorithms for games on graphs (Section 3). A graph game is a two-player zero-sum game played by moving a token on a directed graph, wherein each vertex is owned by one of the players. Typically, the ownership of the nodes along a path through the graph alternates between the two players. If the token is at a vertex owned by a player, this player moves the token along one of the outgoing edges to the next vertex, after which the process repeats. The winning condition determines the player that wins based on the resulting path of the token.

Following our earlier results [10], we represent the scheduling problem as a graph game between the scheduler and the application (Section 4). The scheduler selects the processes that can execute, and the application resolves possible non-determinism within the processes of the concurrent program. Game-theoretic machinery computes a strategy for the scheduler that optimizes the execution of the concurrent program (according to an objective function that embodies some desired performance measure). Here, we consider only the objective of maximizing throughput, but similar techniques can also optimize for other scheduling performance measures, like fairness, context-switches, or energy consumption.

We view the resulting non-preemptive scheduling strategy itself as a scheduling protocol, and we compose the original protocol with this scheduling protocol to obtain a composite, scheduled protocol (Section 5). To evaluate the effect of the restricted protocol on a practical situation, we implement a reference version and a scheduled version of a simple cyclo-static dataflow network. This network consists of four processes, called actors, that interact asynchronously via five buffers (Figure 1). A buffer can handle overflows by dropping items to match its capacity. We measure the throughput (i.e., the time between consecutive productions) of both versions of the program. The throughput of the reference

version varies significantly, and is on average worse than the throughput of the scheduled version, which, moreover, shows only a small variation.

Finally, we conclude and point to future work in Section 6.

Related work There is a wealth of literature on different models for concurrent software, ranging from the well-known Petri nets [15], to the lesser-known higher-dimensional automata [16, 12]. The syntax of work automata equals that of timed automata [1], but their semantics differ significantly, as jobs in work automata may progress independently at differing rates.

Our implementation of the scheduler as a composition is very similar to the definition of a scheduler defined by Goubault [13]. In his work, Goubault specifies the application as a higher dimensional automaton and views the scheduler as a subautomaton. While our scheduling framework builds on graph games, his work depends on the solution of a particular problem on huge sparse matrices.

2 Work automata

The foundation of our scheduling framework consists of work automata [9], which provide a precise, formal specification of a concurrent application. A work automaton records the state and the progress of a finite set of processes. The state is recorded as an element, q , from a finite set of states, Q . The progress is recorded as an assignment $p : J \rightarrow \mathbb{R}$ of real-number values to a finite set J of jobs. A tuple (q, p) defines the complete configuration of the concurrent program.

A work automaton restricts the progress of its jobs by means of boolean expressions, which are conjunctions and disjunctions of atomic expressions of the form $j \sim n$, with $j \in J$ a job, $n \in \mathbb{N}$ an integer, and $\sim \in \{\leq, =, \geq\}$ a relation. The precise shape of the boolean expressions is irrelevant for our work; we just assume the existence of a set \mathcal{B}_J of boolean expressions over J , together with a satisfaction relation $p \models \phi$ that relates assignments $p : J \rightarrow \mathbb{R}$ and expressions $\phi \in \mathcal{B}_J$.

Definition 1. *A work automaton is a tuple $(Q, \Sigma, J, T, I, c_0)$, where Q is a finite set of states, Σ is a finite set of signals, J is a finite set of jobs, $T \subseteq Q \times \Sigma \times \mathcal{B}_J \times 2^J \times Q$ is a (finite) transition relation, $I : Q \rightarrow \mathcal{B}_J$ an invariant, and $c_0 \in Q \times \mathbb{N}_0^J$ is an initial configuration.*

Although a work automaton consists of finitely many elements, it defines a continuously evolving system, due to the real-valued progress of jobs.

The semantics of work automata requires some auxiliary notation of job assignments. For every subset $X \subseteq J$ of jobs, the job assignment $[X] : J \rightarrow \mathbb{R}$ maps every $j \in J$ to 0 or 1, such that $[X](j) = 1$ if $j \in X$, and $[X](j) = 0$, otherwise. Furthermore, we allow composition of job assignments via component-wise addition, multiplication, and scalar multiplication: for any two job assignments $p, p' : J \rightarrow \mathbb{R}$ and any scalar $t \in \mathbb{R}$, we have $(p + p')(j) = p(j) + p'(j)$, $(p \cdot p')(j) = p(j) \cdot p'(j)$, and $(t \cdot p)(j) = t \cdot p(j)$, for all $j \in J$.




Fig. 1. Petri net representation of a small cyclo-static dataflow graph.

Definition 2. The semantics $\llbracket A \rrbracket$ of a work automaton $A = (Q, \Sigma, J, T, I, c_0)$ is a labeled graph, with vertices $C = Q \times \mathbb{R}^J$, labels $L = \Sigma \cup \mathbb{R}_+^J$, initial vertex $c_0 \in C$, and edges $\rightarrow \subseteq C \times L \times C$ defined by the following two rules:

$$\frac{d : J \rightarrow \mathbb{R}_+, \quad \forall t \in [0, 1] (p + t \cdot d \models I(q))}{(q, p) \xrightarrow{d} (q, p + d)}$$

and

$$\frac{(q, \sigma, g, R, q') \in T, \quad p \models g, \quad p \cdot [J \setminus R] \models I(q')}{(q, p) \xrightarrow{\sigma} (q', p \cdot [J \setminus R])}$$

We demonstrate the utility of work automata by formalizing a simple cyclo-static dataflow network, which serves as our running example.

Example 1. Figure 1 shows the Petri net representation [15] of a small cyclo-static dataflow (CSDF) graph from [5, p. 29]. Recall that a Petri net consists of places (depicted as circles) that contain zero or more tokens, and transitions (depicted as rectangles). A transition firing consumes a single token from each of its input places and produces a token in each of its output places.

The system in Figure 1 consists of 4 actors (A_1 to A_4), which are connected via 5 buffers (E_1 to E_5). The original example in [5] does not make any assumptions on the nature of the buffers: they can be FIFO, LIFO, lossy, or priority buffers. For our purpose, we assume that writing to a full buffer loses the written data. For simplicity, Figure 1 represents each buffer as a place in the Petri net. The losing behaviour is made precise in the work automaton in Figure 2(e).

Each actor A_i , for $1 \leq i \leq 4$, cycles through a number of phases (hence the name cyclo-static). Figure 1 represents each phase as a transition A_i^j , for some index j , in the Petri net. Each phase of an actor requires time to execute. Although the actual times may vary, we consider only the worst-case execution times (WCET). We assume that all phases of a given actor have the same WCET. The integer value next to each transition in Figure 1 specifies its WCET. \diamond

Although the functional behavior (as a Petri net) of the dataflow network in Figure 1 is very precise, its non-functional behavior (the timing of transitions) is still unclear. The following example makes this precise by providing a work automaton for each actor and buffer in Figure 1:

Example 2. Figure 2 shows the work automata that encode the informal description of the behavior of the CSDF graph in Example 1. The work automaton for A_i , with $1 \leq i \leq 4$, has a real-valued job variable j_i that measures the progress of its respective actor. The initial condition $j_1 = 0$ in Figure 2(a) shows that actor A_1 must first perform 5 units of work on j_1 before it can produce on E_1 and E_3 . In contrast, the initial condition $j_4 = 4$ in Figure 2(b) shows that actor A_4 can immediately consume tokens from E_3 and E_4 . The work automata for A_2 and A_3 in Figures 2(c) and 2(d) are very similar. Each first consumes a datum from its input buffer, then executes for a given amount of time, and finally produces a datum in its output buffer. Figure 2(e) shows the work automaton for a buffer of capacity 4 (buffers of other capacity are defined similarly). The self-loop transition on state s_4 loses the data, if the buffer is full. \diamond




Fig. 2. Work automata for the dataflow graph in Figure 1, without the idling transitions $(s_i, \emptyset, \top, \emptyset, s_i)$, for all states i . The self-loop transition in (e) loses the data, if the buffer is full.

One of the important features of work automata (which we use repeatedly in our work) is that they can be composed to obtain the specification of a larger system. The composition operator combines two work automata by synchronizing on their respective signals [9]. The original definition of work automata synchronizes signals in a specific way, but this is again not important for the purpose of our work. Therefore, we simply assume the existence of a composition operator, \bowtie , on sets of signals. For any two sets of (local) signals Σ_1 and Σ_2 , we obtain a set of (global) signals $\Sigma_1 \bowtie \Sigma_2$, and an injective function $\pi : \Sigma_1 \bowtie \Sigma_2 \rightarrow \Sigma_1 \times \Sigma_2$ that decomposes any global signal $\sigma \in \Sigma_1 \bowtie \Sigma_2$ into two local signals $\pi(\sigma) = (\sigma_1, \sigma_2) \in \Sigma_1 \times \Sigma_2$. Injectivity of π ensures that the global signal, σ , is completely determined by its local components σ_1 and σ_2 .

Definition 3. Let $A_i = (Q_i, \Sigma_i, J_i, T_i, I_i, q_{0i}, \phi_{0i})$, for $i \in \{1, 2\}$, be two work automata. Their composition is the work automaton $A_1 \bowtie A_2 = (Q_1 \times Q_2, \Sigma_1 \bowtie \Sigma_2, J_1 \cup J_2, T, I, (q_{01}, q_{02}), \phi_{01} \wedge \phi_{02})$, where \wedge is conjunction; $I(q_1, q_2) = I(q_1) \wedge I(q_2)$, for $(q_1, q_2) \in Q_1 \times Q_2$, is the invariant; and the transition relation T is defined by the following rule:

$$\frac{(q_i, \sigma_i, g_i, R_i, q'_i) \in T_i \text{ for all } i \in \{1, 2\} \quad \pi(\sigma) = (\sigma_1, \sigma_2)}{((q_1, q_2), \sigma, g_1 \wedge g_2, R_1 \cup R_2, (q'_1, q'_2)) \in T}$$

The composition of work automata is synchronous: both components in a composition produce their signals simultaneously. However, such a synchronous execution is not always desirable. To simulate asynchronous execution, we can add auxiliary idling transitions of the form $(q, \varepsilon, \top, \emptyset, q)$, where $\varepsilon \in \Sigma$ is an ‘empty signal’, and $\top \in \mathcal{B}_J$ is a tautology. For readability, we do not show these idling transitions in Figure 2.

3 Graph games

The second ingredient for our scheduling framework consists of algorithms that construct non-preemptive schedules. As shown in [10], we can use games on graphs to construct these schedules. A graph game is a two-player game of infinite duration. The possible move sequences are characterized by a safety automaton:

Definition 4. A safety automaton is a tuple $(Q, \Sigma, \delta, q_0, F)$, with Q a finite set of states, Σ a finite set of moves, $\delta : Q \times \Sigma \rightarrow Q$ a transition function, $q_0 \in Q$ an initial state, and $F \subseteq Q$ a set of accepting states, such that for every state $q \in Q$, we have $q \in F$ if and only if $\delta(q, \sigma) \in F$, for some move $\sigma \in \Sigma$.

In other words, a safety automaton is a deterministic finite automaton (DFA), wherein accepting states cannot be reached from non-accepting states, and every accepting state has an accepting successor.

As usual, we extend the transition function δ to finite move sequences by defining $\delta(q, \lambda) = q$ and $\delta(q, s\sigma) = \delta(\delta(q, s), \sigma)$, for every state $q \in Q$, finite move sequence $s \in \Sigma^*$, move $\sigma \in \Sigma$, and empty sequence $\lambda \in \Sigma^*$.

Since we are interested in infinite games only, we define the accepted language of a safety automaton as follows:

Definition 5. *The accepted language $L(A)$ of a safety automaton A is the set of infinite words $\sigma_1\sigma_2\cdots \in \Sigma^\omega$, such that $\delta(q_0, \sigma_1\cdots\sigma_n) \in F$, for all $n \geq 0$.*

Instantiating Definition 5 for $n = 0$, we see that $\delta(q_0, \sigma_1\cdots\sigma_n) = \delta(q_0, \lambda) = q_0 \in F$, which means that the accepted language is non-empty if and only if the initial state q_0 is accepting.

Consider a state $q \in Q$ of a safety automaton. While all moves are possible in q (the transition function is defined for all pairs of states and moves), not every move is enabled in the sense that it leads to an accepted word. In view of Definition 5, we consider the set $\Sigma_q = \{\sigma \mid \delta(q, \sigma) \in F\}$ of enabled moves at q . By Definition 4, Σ_q is non-empty, for accepting states $q \in F$.

Based on safety automata, we construct a graph game as follows:

Definition 6. *A graph game is a tuple (A, C, W) , with $A = (Q, \Sigma, \delta, q_0, F)$ a safety automaton, $C \subseteq Q$ a set of controlled states, and $W \subseteq \Sigma^\omega$ a winning condition.*

A graph game (A, C, W) is played by two players (say Player 1 and Player 2) who take turns to move a token from state to state. We consider Player 1 a protagonist and Player 2 an antagonist. To simplify notation, we write $C_1 = C$ for the states controlled by Player 1, and we write $C_2 = Q \setminus C$ for the states controlled by Player 2. Initially, the token is in state $q_0 \in Q$. If the token is in state $q \in C_k$, with $k \in \{1, 2\}$, then Player k selects an enabled move $\sigma \in \Sigma_q$ and moves the token to state $\delta(q, \sigma)$. As a result, the token moves along a path

$$q_0 \xrightarrow{\sigma_1} q_1 \xrightarrow{\sigma_2} q_2 \xrightarrow{\sigma_3} \cdots$$

through the safety automaton. Player 1 wins if the sequence $\sigma_1\sigma_2\sigma_3\cdots \in \Sigma^\omega$ of moves is contained in the winning condition W . Otherwise, Player 2 wins.

A joint strategy (for both players) is a function $\zeta : \Sigma^* \rightarrow \Sigma$ that selects a move, for every finite move sequence. A strategy for Player $k \in \{1, 2\}$ is a function $\zeta_k : P_k \rightarrow \Sigma$, with $P_k = \{\sigma_1\cdots\sigma_n \in \Sigma^* \mid \delta(q_0, \sigma_1\cdots\sigma_n) \in C_k\}$, that selects a move, for every finite move sequence that leads to a state controlled by Player k . If Player k follows a strategy ζ_k , then the resulting move sequence is contained in the set

$$L(\zeta_k) = \{\sigma_1\sigma_2\sigma_3\cdots \in L(A) \mid \zeta_k(\sigma_1\cdots\sigma_n) = \sigma_{n+1} \text{ whenever defined}\}$$

of outcomes of the game that are ensured by following strategy s_k . A strategy ζ_1 for Player 1 is winning if $L(\zeta_1) \subseteq W$. That is, strategy ζ_1 ensures that the resulting move sequence is contained in W , irrespective of the moves of Player 2. Winning strategies for Player 2 are defined similarly. A strategy ζ_k for Player k is optimal, if it is winning or if no winning strategy for Player k exists.

Of course, for a given game, it is impossible that both players have a winning strategy. However, for general winning conditions, it is possible that neither player has a winning strategy. In this case, the game is not determined. Nevertheless, Martin proved [14] that graph games are determined if the winning

condition is a Borel set¹. Unfortunately, the results by Martin are descriptive and do not suggest a practical algorithm to find a winning strategy.

For simpler winning conditions, such as the ratio objective, we do have algorithms that compute a winning strategy:

Definition 7. *A ratio game is a graph game with a winning condition*

$$W_{s/t \geq v} = \left\{ \sigma_1 \sigma_2 \sigma_3 \cdots \in \Sigma^\omega \mid \liminf_{n \rightarrow \infty} \frac{\sum_{i=1}^n s(\sigma_i)}{\sum_{i=1}^n t(\sigma_i)} \geq v \text{ and } t(\sigma_1) \neq 0 \right\}$$

for some functions $s : \Sigma \rightarrow \mathbb{Z}$ and $t : \Sigma \rightarrow \mathbb{N}_0$, and value $v \in \mathbb{Q}$.

To the best of our knowledge, ratio games are introduced by Bloem et al. for the synthesis of robust systems [6]. The winning condition $W_{s/t \geq v}$ stipulates that, for every $\epsilon \geq 0$ there exists some time $N \geq 0$, such that for all $n \geq N$ after this time, the fraction $(\sum_{i=1}^n s(\sigma_i)) / (\sum_{i=1}^n t(\sigma_i))$ is at most ϵ less than v .

If $t(\sigma) = 1$, for all moves $\sigma \in \Sigma$, then we obtain a mean-payoff game. Much research has been devoted to finding efficient algorithms that solve mean-payoff games. One of the best known algorithms for mean-payoff games is due to Brim et al. [7]. Their solution for mean-payoff games generalize easily to ratio games. We provide a brief explanation of this algorithm, and refer to [7] for full details.

A classical result by Ehrenfeucht and Mycielsky [11], called memoryless determinacy, states that there exists a positional optimal joint strategy for mean-payoff games (and ratio games).

Definition 8. *A joint strategy $\zeta : \Sigma^* \rightarrow \Sigma$ is positional, if $\delta(q_0, s_1) = \delta(q_0, s_2)$ implies $\zeta(s_1) = \zeta(s_2)$, for all move sequences $s_1, s_2 \in \Sigma^*$.*

A positional strategy depends only on the current state $\delta(q_0, s)$, instead of the full history $s \in \Sigma^*$. If both players follow some positional strategy, the outcome of the game is a path

$$q_0 \xrightarrow{\sigma_1} \cdots \rightarrow q_k \xrightarrow{\sigma_k} \cdots \rightarrow q_n \xrightarrow{\sigma_n} q_k \xrightarrow{\sigma_k} \cdots \quad (1)$$

in the safety automaton that ends with a cycle in the distinct states q_k, \dots, q_n . The outcome of a ratio game is winning for a value $v = a/b \in \mathbb{Q}$ if and only if $(\sum_{i=k}^n s(\sigma_i)) / (\sum_{i=k}^n t(\sigma_i)) \geq v = a/b$, which is equivalent to $\sum_{i=k}^n w(\sigma_i) \geq 0$, with $w(\sigma) = b \cdot s(\sigma) - a \cdot t(\sigma)$, for all moves $\sigma \in \Sigma$.

Brim et al. observed that positional winning strategies for Player 1 in a ratio game correspond with consistent valuations:

Definition 9. *A valuation $f : Q \rightarrow \mathbb{N}_0 \cup \{\infty\}$ is consistent in state $q \in Q$ iff*

1. $q \in C_1$ implies $f(q) + w(\sigma) \geq f(\delta(q, \sigma))$, for some move $\sigma \in \Sigma_q$, or
2. $q \in C_2$ implies $f(q) + w(\sigma) \geq f(\delta(q, \sigma))$, for every move $\sigma \in \Sigma_q$,

¹ A Borel set is a subset of Σ^ω obtained from languages of safety automata by repeated complements and countable intersections.

Algorithm 1: Synthesis problem for ratio games.

Input : A ratio game $(A, C, W_{s/t \geq v})$, with $A = (Q, \Sigma, \delta, q_0, F)$, functions $s : \Sigma \rightarrow \mathbb{Z}$ and $t : \Sigma \rightarrow \mathbb{N}_0$, and a value $v = \frac{a}{b} \in \mathbb{Q}$.
Output: A the largest quasi-strategy $\zeta : Q \rightarrow 2^\Sigma$ that is winning.

- 1 **foreach** $\sigma \in \Sigma$ **do** $w(\sigma) \leftarrow b \cdot s(\sigma) - a \cdot t(\sigma)$;
- 2 **foreach** $q \in Q$ **do** $f(q) \leftarrow 0$;
- 3 **foreach** $q \in C$ **do** $c(q) \leftarrow |\{\sigma \in \Sigma_q \mid f(q) + w(\sigma) < f(\delta(q, \sigma))\}|$;
- 4 $B \leftarrow \sum_{q \in Q} \max(\{0\} \cup \{-w(q, \sigma) \mid \sigma \in \Sigma_q\})$;
- 5 $L \leftarrow \{q \in Q \mid f \text{ inconsistent in } q\}$;
- 6 **while** $L \neq \emptyset \neq \{q \in Q \mid f(q) = 0\}$ **and** $f(q_0) < \infty$ **do**
- 7 Pick $q \in L$, with $f(q)$ minimal;
- 8 $f_q \leftarrow f(q)$;
- 9 $f(q) \leftarrow \min\{n \in \{1, \dots, B, \infty\} \mid f[q \mapsto n] \text{ consistent in } q\}$;
- 10 $L \leftarrow L \setminus \{q\}$;
- 11 **if** $q \in C$ **then** $c(q) \leftarrow |\{\sigma \in \Sigma_q \mid f(q) + w(\sigma) < f(\delta(q, \sigma))\}|$;
- 12 **foreach** (p, σ) with $\delta(p, \sigma) = q \neq p$ **and** $f(p) + w(\sigma) < f(q)$ **do**
- 13 **if** $p \in C$ **then**
- 14 **if** $f(p) + w(\sigma) \geq f_q$ **then** $c(p) \leftarrow c(p) - 1$;
- 15 **if** $c(p) = 0$ **then** $L \leftarrow L \cup \{p\}$;
- 16 **if** $p \in Q_0$ **then** $L \leftarrow L \cup \{p\}$;
- 17 **foreach** $q \in Q$ **do** $\zeta(q) \leftarrow \{\sigma \in \Sigma_q \mid f(q) + w(\sigma) \geq f(\delta(q, \sigma)), \text{ and } f(q) < \infty\}$;

A valuation is consistent if it is consistent in every state.

Suppose that there exists a consistent valuation $f : Q \rightarrow \mathbb{N}_0 \cup \{\infty\}$. Repeated application of Definition 9 to the path in Equation (1) yields $f(q_k) + \sum_{i=1}^k w(\sigma_i) \geq f(\delta(q_k, \sigma_1 \cdots \sigma_n)) = f(q_k)$. If $f(q_k) < \infty$ is finite, then the outcome is winning for Player 1.

Brim et al. suggest a value iteration method to find the smallest possible valuation (we compare valuations pointwise: $f \leq f'$ iff $f(q) \leq f'(q)$, for all states $q \in Q$). Our Algorithm 1 shows a variation of their algorithm with only a few minor, but novel, adjustments.

The first modification is on Lines 6 and 7. Let $a = \min_{q \in Q} f_*(q)$ be the smallest value of the smallest valuation f_* . If $a < \infty$, then the valuation $f' : Q \rightarrow \mathbb{N}_0$ defined as $f'(q) = f_*(q) - a$, for all $q \in Q$, is less than or equal to f_* . Minimality of f_* shows that $a = 0$, which means that there exists a state $q \in Q$ with $f_*(q) = 0$. We refer to $q \in Q$ with $f_*(q) = 0$ as a pivot state. If there are no more pivot states, we can terminate the value iteration.

The second (minor) modification is on Line 9, which becomes apparent for states $q \in Q$ with a negative self loop transition (i.e., some $\sigma \in \Sigma$ with $\delta(q, \sigma) = q$ and $w(\sigma) < 0$). While the original algorithm by Brim et al. repeatedly adds $-w(\sigma)$ to the valuation $f(q)$ at state q , Algorithm 1 immediately jumps to the smallest valuation that resolves the inconsistency at q .

Value problem For given functions s and t , we have a family $\{W_{s/t \geq v} \mid v \in \mathbb{Q}\}$ of winning conditions. Since Player 1 wishes to maximize the ratio between the




Fig. 3. Work automaton specifying cooperative behavior of a job j . Signal x_j executes job j , and signal y_j yields job j . The bound $B \in \mathbb{N}_0$ ensures job j eventually yields.

cumulatives of s and t , it is natural to look for the largest value $v \in \mathbb{Q}$ for which there exists a winning strategy. This problem is known as the value-problem. The set of values that are winning for Player 1 is a half-open interval $(-\infty, v_*]$, with v_* the optimal value. Using Algorithm 1, we can test the query $v \geq v_*$ for any value $v \in \mathbb{Q}$. Hence, the value problem can be solved by a binary search. Comin and Rizzi [8] improved this idea by reusing results from earlier queries.

4 Scheduling Game

We now apply the results in Sections 2 and 3 to define a game-theoretic framework for the synthesis of non-preemptive schedules. We assume that a concurrent program is given in the form of a work automaton $(Q, \Sigma, J, T, I, c_0)$. Of course, we assume that the work automaton accurately models the real application. It may be nontrivial to verify whether or not the work automaton in fact models the application sufficiently accurately, but this concern is beyond the scope of our work on scheduling. In the worst case, one can ensure the application's compliance with the work automaton model by means of runtime verification.

By definition, non-preemptive scheduling relies on the cooperation of the application for managing its execution. We therefore, assume that the work automaton is cooperative:

Definition 10. *A work automaton A is cooperative if and only if for every job $j \in J$ of A , we have that $A \bowtie A_j$ and A are identical up to renaming of states, where A_j is the work automaton in Figure 3.*

To obtain a scheduling game from a given cooperative work automaton A , we compose A with the auxiliary work automaton G_A in Figure 4 that allows us to determine which player is to move and what moves are allowed by each player. In the composition $A \bowtie G_A$, every state is either controlled (if G_A is in state 0) or uncontrolled (if G_A is in state 1).

In the sequel, we assume without loss of generality that the work automaton G_A is already integrated in the specification of the cooperative work automaton:

Definition 11. *A cooperative work automaton is playable iff $A \bowtie G_A$ and A are identical up to state renaming.*

The semantics of a playable cooperative work automaton A from Definition 2 cannot be used directly in a graph game, because there are infinitely many ways




Fig. 4. Work automaton G_A specifying the rules of the scheduling game for a work automaton A . The scheduler selects any signal $X \subseteq N$ that intersects $X_J = \{x_j \mid j \in J\}$. The application selects a signal $Y \subseteq N \setminus X_J$ that intersects $Y_J = \{y_j \mid j \in J\}$ or a signal $Z \subseteq N \setminus X_J$ that does not intersect Y_J .

to make progress via the d -transitions. Of course, not all progressions are equally likely to happen. In fact, if sufficiently many processors are available, we expect that all running processes make an equal amount of progress. We assume that these processes run uninterruptedly and at equal speeds. We denote the expected transition in the semantics with a double arrow:

Definition 12. *The expected semantics $\llbracket A \rrbracket_e$ of a work automaton A is the subgraph of $\llbracket A \rrbracket$ with the \Rightarrow edges defined by the rules*

$$\frac{c \xrightarrow{[S]} c' \quad \text{if } c \xrightarrow{[S']} c' \text{ and } S' \neq S \text{ then } S \not\subseteq S'}{c \xRightarrow{[S]} c'} \quad \text{and} \quad \frac{c \xrightarrow{\sigma} c'}{c \xRightarrow{\sigma} c'}$$

The expected transition relation is just a subrelation of the real transition relation of the semantics of a work automaton. Note that there is no guarantee that the real execution follows the expected one. However, we show in Section 5 that deviations of real execution from expected execution do not cause deadlocks.

Lemma 1. *The expected semantics of a cooperative work automaton is finite.*

Proof. Since A is cooperative, the progress of every job is bounded by $B \in \mathbb{N}_0$ in every state (Figure 3). By a simple induction, it follows that the configurations of the expected semantics are contained in the finite set $Q \times \{0, \dots, B\}^J$. \square

The final ingredient for a scheduling game is an objective. The only restriction that we impose on a scheduling objective is that it must be expressible as a ratio objective $W_{s/t \geq v}$, in terms of some functions $s : M \rightarrow \mathbb{Z}$ and $t : M \rightarrow \mathbb{N}_0$.

Example 3. The composition of the work automata in Figures 2 and 4 yields a game graph. We maximize the throughput, which we define as the ratio between the number of productions and the number of time steps (ticks). We can count the productions by counting how often e_5 fires. Hence, we define

$$s(a_1 \cdots a_n) = |\{i \mid e_5 \in a_i \in \Sigma\}| \quad \text{and} \quad t(a_1 \cdots a_n) = |\{i \mid a_i \in \mathbb{R}_+^J\}|.$$

Algorithm 1 finds a subgame for which every play is of optimal throughput.

As the CSDF graph is a deterministic program, all non-determinism is controlled by the scheduler. Since all options ensure optimal throughput, we can resolve the non-determinism arbitrarily. We resolve non-determinism by preferring idling. The deterministic scheduling strategy for this example can be presented as a Gantt chart, shown in Figure 5. \diamond




Fig. 5. Generated schedule. The vertical dashed line indicates the start of the period.

5 Protocol restriction

Consider a playable cooperative work automaton A from Section 2 and a non-preemptive scheduling strategy $\zeta : C \rightarrow \Sigma$ from Section 4, with C the set of configurations of the expected semantics $\llbracket A \rrbracket_e$ of A . Now we intend to implement the schedule ζ as a protocol $A(\zeta)$ that blocks precisely those jobs that are not supposed to make progress, such that the operating system scheduler has to closely follow the desired scheduling strategy.

The solution to ratio games in Algorithm 1 returns a subgraph of the original game graph. For a scheduling game, vertices are configurations of A , and the edge come from the expected transition relation. Hence, the resulting scheduling strategy can be easily transformed back into a work automaton.

Definition 13. *The scheduling work automaton $A(\zeta)$ of a scheduling strategy ζ is defined as the tuple $(C, \Sigma, \emptyset, \rightarrow, I, c_0)$ with states $C = Q \times \mathbb{N}_0^J$, trivial invariant $I(c) = \top$ for all states $c \in C$, and transition relation defined by the rule*

$$\frac{c \in C \text{ controlled} \quad c \xrightarrow{a} c' \quad \zeta(c) \text{ defined implies } a = \zeta(c)}{c \xrightarrow{a, \top, \emptyset} c'}$$

We enforce the scheduling strategy by considering the composition $A \bowtie A(\zeta)$. The composition $A \bowtie A(\zeta)$ does not introduce any new, undesired executions, but merely restricts the composition to a subset of desired executions of A . By construction, A and $A(\zeta)$ synchronize only on signals, which agrees with the fact that the schedule ζ is non-preemptive.

For the construction of the scheduling game in Section 4, we assume that scheduled jobs run as expected, i.e., they run continuously and at constant speed. However, it is actually very likely that the actual execution deviates from the expected execution. A natural question is whether such deviations may confuse the scheduler enough to introduce deadlocks.

Theorem 1. *If A is a composition of simple work automata, then $A \bowtie A(\zeta)$ is deadlock free.*

A simple work automaton is a work automaton with at most one job and no silent transitions that in every configuration can either make progress or fire a transition, but cannot do both. The work automata in Figure 2 are examples of simple work automata.




Fig. 6. Throughput of original program (a) and scheduled program (b). The horizontal axis is the time in milliseconds between successive firings of A_4^3 (grouped in bins of 10 milliseconds), and the vertical axis is the frequency of each bin.

Proof (of Theorem 1). The state-space of $A \bowtie A(\zeta)$ is defined by tuples consisting of a state $q \in Q$ and a configuration (q', p') of $A(\zeta)$. Let $c = ((q, (q', p')), p) \in (Q \times (Q \times \mathbb{N}_0^J)) \times \mathbb{R}^J$ be a reachable configuration of $A \bowtie A(\zeta)$. The absence of silent transitions ensures that the scheduler $A(\zeta)$ knows the state of A (i.e., we have $q = q'$). The progress $p : J \rightarrow \mathbb{R}$ of the jobs may still be unknown (i.e., $p \neq p'$ is possible). By construction, there exists some expected execution of $A \bowtie A(\zeta)$ that passes through the same state q (although the progress may be different from p) and enables a transition $t = (q, \sigma, g, R, q')$. Since the work automaton for each job is simple, the guard g of transition t states that the progress of a subset of jobs is maximal (while the progress of other jobs is irrelevant). Hence, from c we can make sufficient progress to enable transition t , which implies that c is not a deadlock. \square

Example 4. To evaluate the schedule in Figure 5, we implement the CSDF graph in Figure 1 in Python. We performed the experiment on a 64-bit Windows 10 Home Edition with a Intel[®] Core[™] i7-7700HQ CPU at 2.80 GHz and 16 GB RAM. We executed the source code with a 64-bit Python 3.9.0 interpreter.

The source code of the scheduled application is a manual, but mechanical, transformation of the original source code. The transformation adds barrier synchronizations between the actors and a scheduler process. This scheduler process implements the non-preemptive schedule in Figure 5.

Figure 6 shows the histogram (with a bin-size of 10 ms) of the output of both versions of the program. We measure the throughput (i.e., the time between successive productions) of each version. Both the expected value and the standard deviation of the two versions differ. The original version has an expected throughput of 441 ms with a standard deviation of 95 ms. The scheduled version has an expected throughput of 386 ms and a standard deviation of 6 ms.

The quality of the schedule alone does not explain the improvement of the expected throughput; the characteristics of the original protocol are the most important factor in this example. Recall that we use overflow buffers, which means that an actor loses its datum, if its respective buffer is full. If a datum

is lost, all effort invested in its production is also lost. The general-purpose scheduler of the operating system is unaware of these losses. \diamond

Although the game-theoretic framework in Section 3 aims to optimize the expected throughput, Example 4 shows that the most significant improvement in the scheduled protocol is in its impact on the standard deviation of the throughput. The time between successive productions is much more predictable for the scheduled version compared to the original version. Since we force the operating system scheduler to closely follow a fixed, deterministic schedule, predictability of the throughput is not surprising.

Predictable timing is a requirement for many systems. For example, a pacemaker must assist a patient’s heart to beat at a regular and predictable rate, and a self-driving car must sense and analyze its environment at a predictable rate to avoid collisions. The results of Example 4 show that the unpredictable behavior of the scheduler of the operating systems can be made tightly predictable by restricting the protocol through composition with a scheduling protocol.

6 Conclusion

Protocols contain valuable information indispensable for construction of optimal (non-preemptive) schedules for allocation of resources to execute a concurrent application. Exogenous languages like Reo express a protocol as an explicit software construct, which makes this scheduling information accessible. We express protocols together with their scheduling information in terms of the work automaton semantics of Reo. We construct a generic scheduling framework based on ratio games to find optimal non-preemptive schedules for an application defined as a work automaton. By composing such a scheduling protocol with the original protocol of an application, we obtain a composite scheduled protocol that forces generic operating system preemptive schedulers to closely follow the desired optimal schedule. The exogenous nature of Reo guarantees that the application code remains oblivious to the substitution of the composite scheduled protocol for its original protocol. An experiment shows that a scheduled version of a cyclo-static dataflow network (with the composite protocol) has higher and more predictable throughput compared to its original version.

Future work The algorithm by Brim et al. to solve ratio games requires the full state-space of an application. In view of the state-space explosion problem, we can use other schemes (like Monte-Carlo tree search) to find good schedules.

Although the work presented in our current paper focuses on maximizing throughput, the ratio objective can express many other scheduling performance measures, like fairness, context-switches, or energy consumption. We intend to express these performance measures in terms of ratio objectives.

Acknowledgements We thank Benjamin Lion for his suggestion to use the overflow semantics for the buffers in the CSDF graph and the anonymous reviewers of FSEN 2021 for their constructive comments.

References

1. Alur, R., Dill, D.L.: Automata for modeling real-time systems. In: Paterson, M. (ed.) Proc. of ICALP. LNCS, vol. 443, pp. 322–335. Springer (1990), <https://doi.org/10.1007/BFb0032042>
2. Arbab, F.: Reo: a channel-based coordination model for component composition. *Math. Struct. Comput. Sci.* 14(3), 329–366 (2004), <https://doi.org/10.1017/S0960129504004153>
3. Arbab, F.: Abstract behavior types: a foundation model for components and their composition. *Sci. Comput. Program.* 55(1-3), 3–52 (2005), <https://doi.org/10.1016/j.scico.2004.05.010>
4. Arbab, F.: Proper protocol. In: Ábrahám, E., Bonsangue, M.M., Johnsen, E.B. (eds.) *Theory and Practice of Formal Methods - Essays Dedicated to Frank de Boer on the Occasion of His 60th Birthday*. LNCS, vol. 9660, pp. 65–87. Springer (2016), https://doi.org/10.1007/978-3-319-30734-3_7
5. Bamakhrama, M.A.M.: On hard real-time scheduling of cyclo-static dataflow and its application in system-level design. Ph.D. thesis, Leiden University (2014)
6. Bloem, R., Greimel, K., Henzinger, T.A., Jobstmann, B.: Synthesizing robust systems. In: Proc. of FMCAD. pp. 85–92. IEEE (2009), <https://doi.org/10.1109/FMCAD.2009.5351139>
7. Brim, L., Chaloupka, J., Doyen, L., Gentilini, R., Raskin, J.: Faster algorithms for mean-payoff games. *Formal Methods Syst. Des.* 38(2), 97–118 (2011), <https://doi.org/10.1007/s10703-010-0105-x>
8. Comin, C., Rizzi, R.: Improved pseudo-polynomial bound for the value problem and optimal strategy synthesis in mean payoff games. *Algorithmica* 77(4), 995–1021 (2017), <https://doi.org/10.1007/s00453-016-0123-1>
9. Dokter, K., Arbab, F.: Exposing latent mutual exclusion by work automata. In: Mousavi, M.R., Sgall, J. (eds.) Proc. of TTCS. LNCS, vol. 10608, pp. 59–73. Springer (2017), https://doi.org/10.1007/978-3-319-68953-1_6
10. Dokter, K., Jongmans, S., Arbab, F.: Scheduling games for concurrent systems. In: Lluch-Lafuente, A., Proença, J. (eds.) Proc. of COORDINATION. LNCS, vol. 9686, pp. 84–100. Springer (2016), https://doi.org/10.1007/978-3-319-39519-7_6
11. Ehrenfeucht, A., Mycielski, J.: Positional strategies for mean payoff games. *Int. J. Game Theory* 8(2), 109–113 (1979)
12. van Glabbeek, R.J.: On the expressiveness of higher dimensional automata. *Theor. Comput. Sci.* 356(3), 265–290 (2006), <https://doi.org/10.1016/j.tcs.2006.02.012>
13. Goubault, E.: Schedulers as abstract interpretations of higher-dimensional automata. In: Proc. of PEPM. pp. 134–145. ACM (1995)
14. Martin, D.A.: Borel determinacy. *Annals of Mathematics* pp. 363–371 (1975)
15. Murata, T.: Petri nets: Properties, analysis and applications. *Proceedings of the IEEE* 77(4), 541–580 (1989)
16. Pratt, V.R.: Modeling concurrency with geometry. In: Wise, D.S. (ed.) Proc. of POPL. pp. 311–322. ACM Press (1991), <https://doi.org/10.1145/99583.99625>