Exploring Measures of Agentic Power for Organisational Workflow Structures

Andrea Marino[1[0000−0003−3026−1482]],
Thomas van Binsbergen[1[0000−0001−8113−2221]],
Tom van Engers[1], Giovanni Sileno[1[0000−0001−5155−9021]]

Informatics Institute, University of Amsterdam, The Netherlands
{a.marino, l.t.vanbinsbergen, t.m.vanengers, g.sileno}@uva.nl

Abstract. The rising deployment of autonomous multi-agent systems, where agentic AI components are increasingly integrated in traditional processes, requires analytical tools for understanding the contribution of individual agents to given organisational structures. In this study, we explore several measures of relative power of an agent within workflows modelled as Petri nets. We propose a formalisation for multi-agent computation based on process models: we provide a definition of atomic power, based on which we progressively build towards topological and game theoretic measures. We apply the proposed measures on several illustrative cases and observe how they behave with regards to capturing the different dimensions of power, also discussing their computational tractability.

Keywords: Multi-Agent Systems · Process Modelling · Game Theory

1 Introduction

How much *power* does a specific agent hold within a particular social configuration? This question, traditionally addressed by disciplines such as political science, sociology, law, economics, and game theory, is gaining importance within the computational realm too. For a responsible deployment of artificial multi-agent systems, it is essential—particularly for *auditability* and *control* [13]—to have methods to map the possibilities and responsibilities of agents resulting from a given agentic design. As agentic components are progressively integrated in traditional processes (possibly entirely substituting humans in some tasks), we urge developing analytical tools to assess their organisational role.

In the discourse surrounding the definition of power in normative settings, there is a crucial distinction between the *institutional powers* that constitute the social/normative framework in which agents interact (see *e.g.*, [8]), and the practical or *operational powers*, representing their physical abilities within an environment, possibly enabled by the former by means of accepted social commitments. In this study, for simplicity, we focus on operational powers, though in the discussion section we will provide directions on how institutional dimensions may be incorporated in subsequent studies.

The organisational control structure, *i.e.*, the distribution of behavioural responsibilities and expectations amongst the agents participating to the organisation, reflects in the behaviour that agents perform and that ultimately can be observed. Vice-versa, by looking at the organisation's operational footprint (*i.e.*, its processes, as described by process models), we can draw some conclusions about its internal power distribution.

Our overarching research question is *"How can we compare different organisational structures given the constraints imposed by process dependencies and individual agent capabilities?"*. To achieve this, and focusing precisely on power, we apply the following methods:

- We introduce a formalisation of multi-agent computation based on process models (with Petri nets).
- We propose measures of power starting by considering it in terms of atomic actions and building towards *topological* and *game-theoretic* interpretations.
- We apply the proposed measures to a number of illustrative cases, from canonical edge-cases to more complex systems and open source a reproduction package for our experiments[1].

The paper concludes with a discussion on the applicability and relevance of the proposed analytical framework with respect to multi-agent systems.

2 Theoretical Framework

In this section, we introduce the notational framework that we use to define multi-agent computation in the workflow-focused perspective tackled by this study. First, we recall the definition of Petri nets and their relation to process modelling. Second, we provide a definition of atomic power based on such modelling notation. Third, we introduce two candidate measures for capturing topological features of power within workflows. Finally, we reduce the problem of cooperative games to that of workflow execution and present a formalisation of the game-theoretic measures we selected.

2.1 Process Models

To take into account the distributed nature of the workflow, we model causal dependencies in the world as Petri nets. Petri nets are a visual formal notation expressed as a bipartite graph, consisting of places and transitions. The transition mechanism in a Petri net can be seen the simplest possible model of *physical causation*: when a transition fires, it consumes resources (also called *tokens*) located in its input places, and produces resources (tokens) in its output places. This execution semantics favours modularity, as it does not require global state to describe the behaviour of sub-components of the system. Petri nets are particularly used in fields as electronics, biology, and business process modelling.

[1] Available at https://github.com/uva-cci/coine-aamas-2026

They have shown to be particularly effective for proving that processes have desirable *soundness* properties [1, 2]. For these reasons, we use Petri nets as a foundational design instrument in our notational framework, defined as follows.

Definition 1 (Petri net). *A Petri net $N = (P, T, F, W)$ is a tuple with P a finite set of places $p \in P$; T a finite set of transitions $t \in T$; $F \subseteq (P \times T) \cup (T \times P)$ a flow relation consisting of input arcs $(p, t) \in F$ and output arcs $(t, p) \in F$; $W : F \to \mathbb{N}$ is a weight function which associates weights to arcs.*

Given a transition t, we refer to p as an *input* to t if $(p, t) \in F$ and as an *output* to t if $(t, p) \in F$. The execution semantics of Petri nets is informally called *token game*. A function in the form $M : P \to \mathbb{N}_0$ is called a *marking*, and it assigns a number of tokens to each place. M_0 is an initial marking, M_{end} is a final marking. A transition $t \in T$ is *enabled* if $M(p) \geq W(p, t)$ for all inputs p to t. *Firing* a transition t in M produces a new marking M' such that $M'(p) = M(p) - W(p, t) + W(t, p)$, for all p input or output to t, denoted as $M \xrightarrow{t} M'$. The set of all possible markings for a Petri net N is $\mathcal{M}_N$. A *firing sequence* is a sequence of transitions $t_1 t_2 \ldots t_n$ such that

$$M_0 \xrightarrow{t_1} M_1 \xrightarrow{t_2} \ldots \xrightarrow{t_n} M_N$$

In this study, we specifically focus on workflows with a well-defined end marking M_{end}. Finally, to support the representation of non-deterministic workflow execution, in this work we also refer to stochastic Petri nets.

Definition 2 (Stochastic Petri net). *A stochastic Petri net is a tuple $N = (P, T, F, W, \lambda)$ where (P, T, F, W) is a Petri net and $\lambda : T \to [0, 1]$ assigns to each transition $t \in T$ a firing rate $\lambda(t)$ (the probability of transition).*

2.2 Atomic Power

Atomically, power can be seen as the certainty of obtaining an outcome O if an action A is performed (*sufficiency*), in the context in which there is a possibility to perform A (*usability*), and that O is not obtained yet (*opportunity*), and is not obtained if A is not performed (*counter-factuality*) (see *e.g.*, [4, 15]). If O is not a binary outcome, we could also measure its extent (*strength*).

Given a set of agents $\mathcal{A} = \{a_1, ..., a_n\}$ participating in a process model $N = (P, T, F, W)$, we can assign transitions to agents by defining an agent mapping $m_{\mathcal{A}} : T \to 2^{\mathcal{A}}$ such that $m_{\mathcal{A}}(t)$ is the set of agents capable of firing t. When adding agents to the model, transitions can be thought as associated to (atomic) *action-types*. When a transition fires, the corresponding action is performed by one of the assigned agents. This interpretation is fit to characterise all of the aforementioned dimensions of atomic power. Specifically, if we observe the simplest instance of an action-based mechanism (*i.e.*, an enabled transition t with a single input p_{in} and a single output p_{out}), it satisfies all definitions:

- **sufficiency**: when t fires, a token is produced in p_{out}.

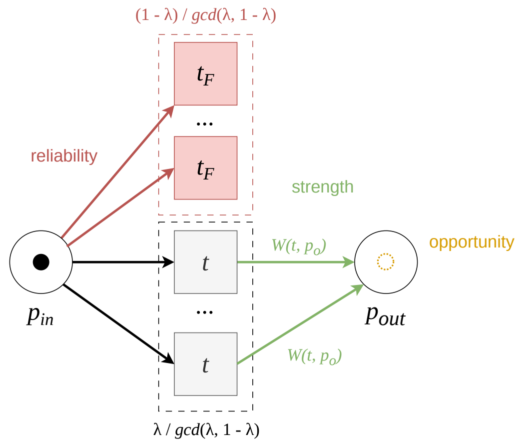

Fig. 1: The reification of λ-firing probabilities in a fully reactive Petri net.

- **usability**: t is enabled by the token in p_{in} (*i.e.*, an agent can perform it).
- **opportunity**: if we describe the outcome as new resources, the transition effectively produces new tokens, which accumulate to pre-existing ones;[2]
- **counter-factuality**: if there is no other enabled t' such that $(t', p_{out}) \in F$, it is impossible to obtain the outcome if not by firing t;
- **strength**: $W(t, p_{out})$ can be used as a measure of strength of the action (the number of resources produced with one firing).

Finally, to further tailor our notation to real-world cases, we can introduce the concept of failure or *reliability* of an action, so that sufficiency may be defined with a certain degree, relying on the concept of firing rate λ. We can then further distinguish an external failure (no agent commits to the action, or *intent failure*) and an internal failure (an agent commits to the action, but it does not produce the expected outcome, or *action failure*). A *reactive* Petri net can be seen as a Petri net in which intent failure never occurs (transitions are fired as soon as they are enabled). We consider here reactive Petri nets, meaning that the agents are committed to the agentic design, exploiting all opportunities of actions once they appear.

Reifying reliability topologically Any transition with probability of action failure λ can be decomposed in $\lambda/_{\gcd(\lambda,1-\lambda)}$ "success" transitions and $^{1-\lambda}/_{\gcd(\lambda,1-\lambda)}$ failure transitions (see Fig. 1). For instance, if a transition has a probability to

[2] This does not strictly hold in the case of Condition/Event nets, *i.e.*, Petri nets whose places are bounded to maximum one token. In this case, if a condition consequent of an event already holds, firing the event would be irrelevant, and would be functionally equivalent as not firing. In many formal definitions of C/E the transition is deemed as disabled in this configuration.

fail $\lambda = 2/3$, this can be reproduced with three transitions in mutually exclusive branching, one associated with success, and two with failure.

2.3 Power in Workflows

We define the usability of an action as the degree of its enablement in a certain environment. In the context of graph-based process models, we can leverage the network properties to assess how much a power (*i.e.*, the actions associated to it) participate in the process execution flow. We propose therefore to observe the occurrences of a certain action in the firing sequences of the reachability graph, and derive a static measure of how "usable" a certain power is in relation to the overall process. We formally define an index to describe this measure:

Definition 3 (Usability Index). *Let $\mathcal{S}$ be the set of firing sequences σ reachable from initial marking M_0, such that no marking is visited more than once. Let $m_{\mathcal{A}} : T \to 2^{\mathcal{A}}$ be the agent mapping. The Usability Index $U(a)$ for an agent $a \in \mathcal{A}$ is defined as the average number of transitions in which an agent can perform any action within a sequence, normalised by the number of other capable agents, across all sequences:*

$$U(a) = \frac{1}{|\mathcal{S}|} \sum_{\sigma \in \mathcal{S}} \sum_{i=1}^{|\sigma|} u(a, \sigma, i) \tag{1}$$

where $u(a, \sigma, i)$ at position i is:

$$u(a, \sigma, i) = \begin{cases} \frac{1}{|m_{\mathcal{A}}(\sigma_i)|} & \textit{if } a \in m_{\mathcal{A}}(\sigma_i) \\ 0 & \textit{otherwise} \end{cases} \tag{2}$$

To further explore the properties of the reachability graph, we consider another measure, grounded in network analysis. Specifically, we aim to capture *downstream dependence* and *non-redundancy* of certain actions by measuring their *centrality* in enabling firing sequences that would otherwise be unreachable. This perspective is methodologically aligned with the study of social networks in organisational theory (see *e.g.*, *structural holes* in [5]). To support the definition of our measure, we first recall the one of *node dominator* [10]:

Definition 4 (Node Dominator). *Let $G = (V, E, r)$ be a flowgraph with start vertex r. A vertex v* dominates *another vertex $w \neq v$ in G if every path from r to w contains v. Vertex v is the* immediate dominator *of w, denoted $v = idom(w)$, if v dominates w and every other dominator of w dominates v.*

Definition 5 (Gatekeeper Index). *Let $\mathcal{S}$ be the set of firing sequences σ reachable from initial marking M_0, such that no marking is visited more than once. Let $m_{\mathcal{A}} : T \to 2^{\mathcal{A}}$ be the agent mapping. The Gatekeeper Index $G(a)$ for an agent $a \in \mathcal{A}$ is defined as the average number of immediately dominated transitions when an agent can perform an action within a sequence, normalised by the number of other capable agents and total transitions, across all sequences:*

$$G(a) = \frac{1}{|\mathcal{S}|} \sum_{\sigma \in \mathcal{S}} \sum_{i=1}^{|\sigma|} g(a, \sigma, i) \tag{3}$$

where $u(a, \sigma, i)$ *at position* i *is:*

$$g(a, \sigma, i) = \begin{cases} \frac{|\{t \in T | \sigma_i = idom(t)\}|}{|T| \cdot |m_{\mathcal{A}}(\sigma_i)|} & \text{if } a \in m_{\mathcal{A}}(\sigma_i) \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

2.4 Power in Cooperative Games

In game theory, *cooperative* games allow players to form alliances (or *coalitions*) to reach a common goal. This is optionally regulated by some normative framework defining the allowed cooperative behaviour. In the context of computational social choice, the study of *simple voting games* has contributed to a better understanding of governance mechanisms.

Definition 6 (Simple Voting Game). *A* simple voting game *is a pair* (N, v) *where* $N = \{1, \dots, n\}$ *is a finite set of players and* $v : 2^N \to \{0, 1\}$ *is a characteristic function such that* $v(\varnothing) = 0$, $v(N) = 1$, *and* v *is monotonic (for all* $S \subseteq T \subseteq N$, $v(S) \leq v(T)$*). A coalition* $S \subseteq N$ *is called* winning *if* $v(S) = 1$ *and* losing *otherwise.*

We relate the concept of voting to workflow execution by reducing the problem of determining a winning coalition to that of finding execution paths reaching a desirable absorbing goal state M_{end}. We call a firing sequence σ *successful* if $M_0 \xrightarrow{\sigma} M_{end}$. We say that a coalition $S \subseteq \mathcal{A}$ is winning if there exists at least one successful sequence σ such that for every $t \in \sigma$, $m_{\mathcal{A}}(t) \cap S \neq \emptyset$. In the stochastic case, to account for the transition firing rates, we introduce a variant $v_{\mathbb{P}} : 2^{\mathcal{A}} \to [0, 1]$ of the binary characteristic function and define it as following:

$$v_{\mathbb{P}}(S) = \sum_{\sigma : M_0 \xrightarrow{\sigma} M_{\text{end}}} \prod_{i=1}^{|\sigma|} p(\sigma_i \mid S)$$

where $p(\sigma_i \mid S)$ is the probability of firing transition σ_i given the λ rates determined by the coalition S.

In these settings, we use the Shapley-Shubik index [14] to assess how *pivotal* an individual agent is within a coalition to the realisation of a certain outcome. As all coalitions are equally likely to form, this index represents a form of a-priori *structural* measurement of power (of combinatorial nature).

Definition 7 (Shapley–Shubik Power Index). *Let* (N, v) *be a simple voting game with* $|N| = n$*. The* Shapley–Shubik power index *of player* $i \in N$ *is defined as:*

$$\phi_i(v) = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!(n - |S| - 1)!}{n!} \big(v(S \cup \{i\}) - v(S)\big).$$

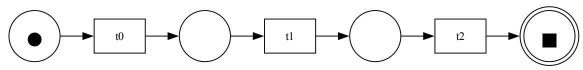

Fig. 2: A sequential workflow.

This index represents the probability that player i is pivotal under a uniformly random permutation of N.

As an alternative, we also measure the Banzhaf index [3], that better accounts for *equal representation* [17].

Definition 8 (Banzhaf Power Index). *Let (N, v) be a simple voting game with $|N| = n$. A player $i \in N$ is said to be* critical *(or a* swing player*) in a coalition $S \subseteq N \setminus \{i\}$ if $v(S) = 0$ and $v(S \cup \{i\}) = 1$. The* Banzhaf index *of player i is*

$$\beta_i(v) = \sum_{S \subseteq N \setminus \{i\}} \big(v(S \cup \{i\}) - v(S)\big),$$

This index counts the number of coalitions in which i is critical.

3 Experimental Settings

In this section, we present the experimental test cases that we use to evaluate the proposed measures. First, we present some simple abstract and deterministic canonical processes that introduce topological edge-cases of process models. Then, we introduce a process scenario inspired by table football, with agent mappings mirroring player formations and stochastic transitions for passes/shots depending on their field positions. Although simplistic, this scenario has sufficient complexity to model the role specialization occurring in social systems for achieving common goals (in the case of a single team), as well as resource competition mechanisms (in the case of two teams), possibly requiring policy coordination. With two teams, the agentic design captures both competition and cooperation mechanisms.

3.1 Canonical Processes

We first evaluate our measures against basic workflow cases. Fig. 2 presents a simple *sequential* workflow, while Fig. 3 shows a prototypical *parallel* process, expressed by a *fork-join* "diamond" structure. Tables 1 and 2 contain the full codomain definitions for the mapping $m_{\mathcal{A}} : T \to 2^{\mathcal{A}}$ for the different scenarios over these process models with a fixed agent set $\mathcal{A} = \{a_1, a_2, a_3\}$.

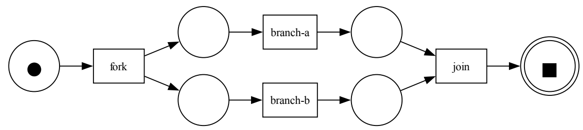

Fig. 3: A parallel workflow with fork-join branching.

$m_{\mathcal{A}}$ \ T	t_0	t_1	t_2
Even	$\{a_1, a_2, a_3\}$	$\{a_1, a_2, a_3\}$	$\{a_1, a_2, a_3\}$
Bottleneck	$\{a_1, a_2, a_3\}$	$\{a_1, a_2, a_3\}$	$\{a_1\}$
Funnel	$\{a_1, a_2, a_3\}$	$\{a_1, a_2\}$	$\{a_1\}$

Table 1: Distribution of transitions over agents in the sequential scenarios.

$m_{\mathcal{A}}$ \ T	t_{fork}	t_a	t_b	t_{join}
Even	$\{a_1, a_2, a_3\}$	$\{a_1, a_2, a_3\}$	$\{a_1, a_2, a_3\}$	$\{a_1, a_2, a_3\}$
Skewed	$\{a_1, a_2, a_3\}$	$\{a_1, a_2\}$	$\{a_3\}$	$\{a_1, a_2, a_3\}$

Table 2: Distribution of transitions over agents in the fork-join scenarios.

3.2 Table Football

We now introduce a simplified process model for the attack sequences in the game of table football, initially modelled without an opposing team (see Fig. 4).

The field is divided into a set of zones $Z = \{D, M, A\}$ (*i.e.*, *defence*, *midfield*, and *attack*), with each zone Z_i being further divided into C_i cells each. A team is composed of defenders $\mathcal{A}_D$, midfielders $\mathcal{A}_M$, and attackers $\mathcal{A}_A$ so that $\mathcal{A} = \mathcal{A}_D \cup \mathcal{A}_M \cup \mathcal{A}_A$. Each player occupies exactly one cell, can only pass the ball forward in the next zone, and only attackers can score, reaching the state M_{end}. The ball starts with a defender (this represents the initial marking M_0). The probability of a successful pass from zone i to zone $i+1$ is defined as $\frac{|\mathcal{A}_{i+1}|}{C_i}$. We do not allow the ball to pass beyond the next zone. The probability of an attacker scoring is $\frac{\gamma}{C_A}$, where γ represents the width of the goal in cells. If a pass fails, the game stops and restarts. It should be noted, that, in our calculations, we collapse the cells into single "zone" places, reverting the process of firing rate reification presented in Fig. 1 based on these probabilities.

When a second team is introduced (see Fig. 5), their players also occupy the C_i cells for a given zone Z_i. This creates the possibility of a pass being intercepted by an opponent, which reverses the flow of the game. In these settings, the opponent team players have a negative impact on the cumulative probability of reaching the goal state. Moreover, the two teams can deploy different formations (*e.g.*, 4-4-2 vs 3-5-2), which enables the analysis of how different agentic mappings relate to each other in a competitive setting.

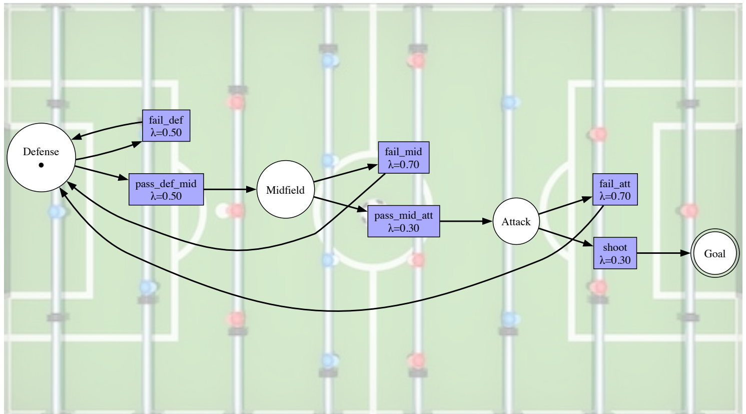

Fig. 4: Petri net process model of a single team table football attack sequence using a 2-5-3 formation. The workflow starts with the ball in the team's defence area and is passed from player to player in order to reach the opponents goal.

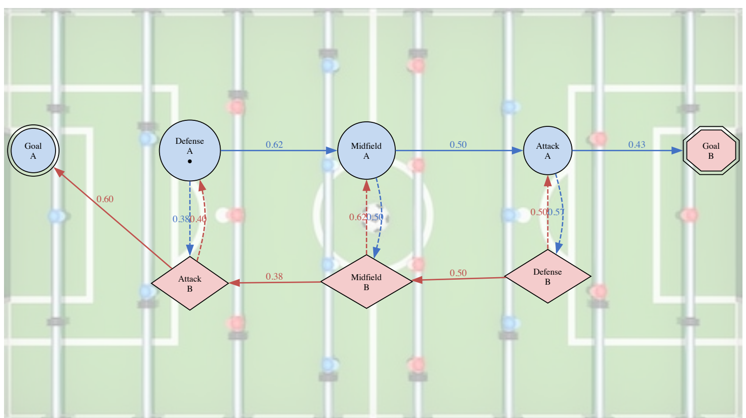

Fig. 5: Visualisation of the stochastic reachability graph (in this case, a discrete-time Markov chain, due to the single-token nature of the process) in the competitive table football scenario, with team A (circles) using a 2-5-3 formation and team B (diamonds) a 4-3-3 formation. The transition probabilities are expressed on the edges that connect each zone state, and failure (*i.e.*, interception) corresponds to transitions going towards the opponent team.

4 Results

4.1 Canonical Processes

In Fig. 6 we present the power indices for each agent in the different configurations on the canonical processes (sequence, fork-join). In the sequential case, we can observe how the Usability index correctly detects the centralisation of

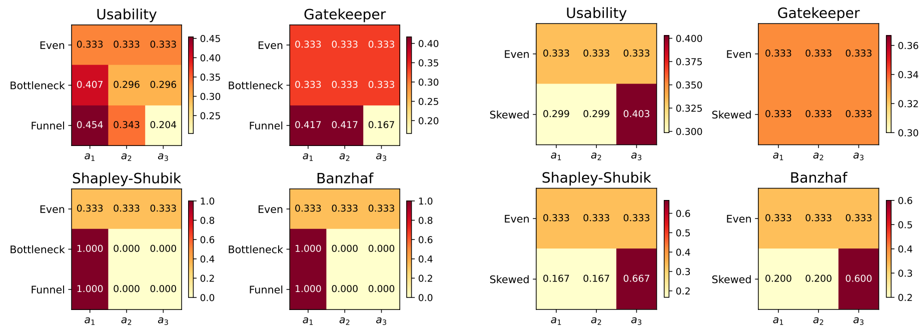

(a) Sequential (b) Fork-Join

Fig. 6: Heat map of power indices calculated on the canonical processes.

power towards the end of the workflow sequence. On the contrary, the Gatekeeper index is particularly insensitive to downstream bottlenecks. Finally, both voting power indices behave identically with respect to each other and, while they clearly capture the bottleneck issue (*i.e.*, single pivotal agent), lack nuance in differentiating between progressive and sudden power shifts in the sequence.

In the fork-join case, the Usability index correctly identifies that the power distribution is less fair (thanks to per-transition normalisation over the agent mapping $m_{\mathcal{A}}$). The Gatekeeper index only considers the ability to fork, and fails to detect the different agent allocation between the two branches. In the voting indices, the power gap is even more apparent. Here, Shapley-Shubik shows higher sensitivity but is closely followed by Banzhaf.

4.2 Table Football

For the single team football scenario, we measure power over canonical football formations `2-5-3`, `4-4-2`, `3-5-2`, `4-3-3` and some illegal "flat" formations `1-1-1` and `4-4-4`. The indices values per agent are presented in Fig. 7. Additionally, because power scores can be interpreted as a distribution of influence in the agent pool, in Fig. 9 we show the Lorentz curves derived from the Gini index [7] which characterises concentration of power in terms of distribution inequality.

As expected, static path-based measures are particularly sensitive to player positioning in the field. Specifically, defenders tend to be more powerful as the attack sequence always starts from their zone and represent a strict dependency. This observation is reinforced by the Gatekeeper index values that are correctly detecting the sequential dependency. However, the Gatekeeper index appears to provide a very stiff interpretation of power. As precedence is key to partake in the power distribution, attackers get a null contribution (their scores are always 0.0 in Fig. 7, and this can also be observed in Fig. 9a sudden curve step). On the other hand, voting inspired indices are capable of identifying pivotal agents (in this case, attackers) acting as a bottleneck with respect to the goal state M_{end}.

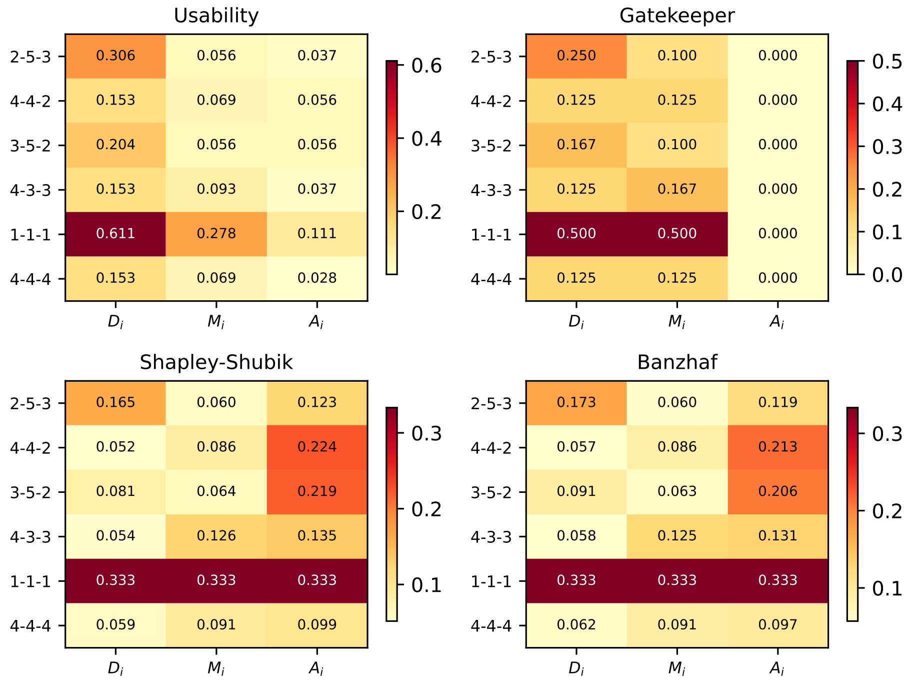

Fig. 7: Heat map of the power indices calculated on different formations for the single team football scenario.

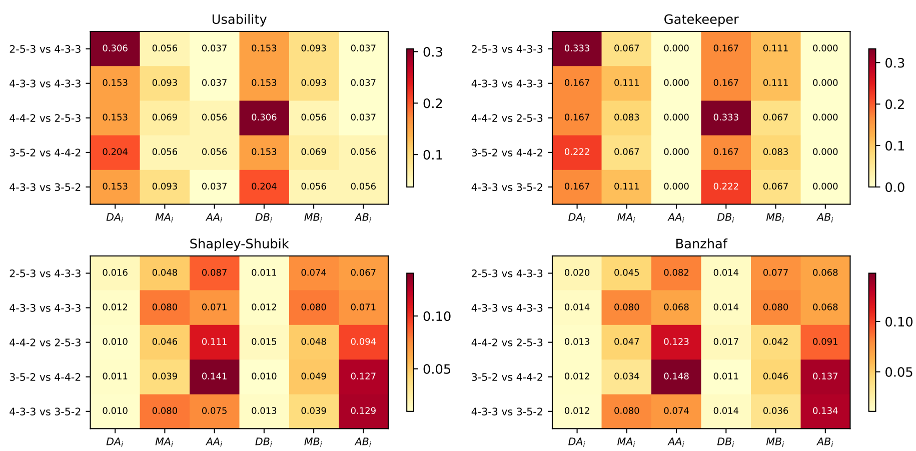

Fig. 8: Heat map visualization of the power indices calculated on different matches of the competitive football scenario.

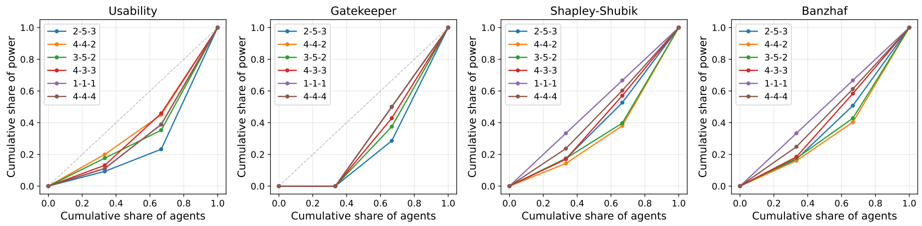

(a) Single team

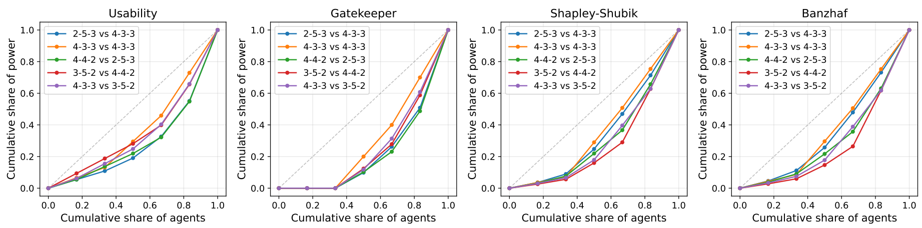

(b) Competitive

Fig. 9: Lorenz curves for the Gini index of the power share in the football scenario. On the x axis we show the cumulative proportion of agents (i.e., role representatives like DA_i), sorted from least powerful to most powerful. On the y axis the cumulative proportion of total power held by those agents.

In the competitive setting, path-based measures are still prioritising defenders as these indices do not account for variations in the firing probabilities. Game-theoretic measures also mostly keep power on attackers like in the single team case. However, interestingly, we can observe a power shift towards midfield as the opponent team adopts a more or equally defensive formation (see *e.g.*, the `4-3-3` vs `3-5-2` match in Fig. 8). This clearly highlights the diminished impact attackers can have against a strong defence, suggesting there is a capability of these indices to capture adversarial environmental conditions if correctly encoded in the process. Interestingly, deploying identical formations leads to a much more evidently fair distribution of power in the Gatekeeper index than in other proposed measures (as observed in the gap between the orange line and the others in Fig. 9b).

4.3 Performance

All the proposed measures rely on intensive pre-computation due either to the enumeration of firing sequences or the highly combinatorial nature of the index calculation. Because of this, we conduct a performance analysis over an increasingly large amount of agents in the most complex process model we presented (the competitive football scenario). Specifically, we record 10 executions per index per agent mapping as we artificially increase the number of players in both teams by 2 at each iteration (from 6 to 20 total agents). We precede the computation with 2 warm-up runs to prevent operating system level interference with the final results. We run our experiments on a MacBook Pro M4 machine.

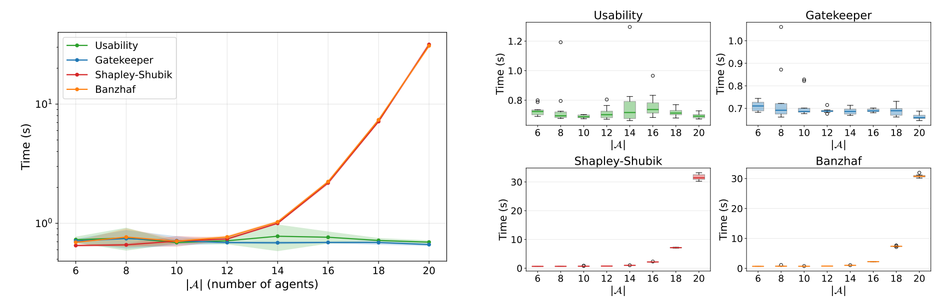

(a) Scaling Analysis (b) Variance Analysis

Fig. 10: Comparative benchmark results showing scaling trends and distribution of performance (in seconds of total execution time) in the power indices calculations over an increasing number of agents in the competitive football scenario. The results are based on 10 recorded executions per index per configuration preceded by 2 warm-up runs on a MacBook Pro M4 machine.

In Fig. 10a, we showcase the exponential growth in total computation time of the voting indices, bound by the size of the power set of players $2^{|A|}$ to generate the coalitions. On the other hand, topological measures present either low-rate or absence of linear growth.

5 Discussion

These results highlight how even a purely operational perspective on power entails that several quantitive indicators are possible and provide each different perspectives on the matter of study. The combination of topological and game-theoretic measures we applied outline distinct relevant properties of given configurations of organisational workflow structures. Topological measures appear to be particularly apt to outline power as "enablement", in terms of precedent conditioning, but fail to describe fair distributions over linear processes. On the other hand, game-theoretic measures highlight the correlation between operational bottlenecks and power, by prioritising goal attainment over cooperation support.

Yet, we are aware that the considered models have simplified the complexity of real-world dynamic environments, where individual agents might be misaligned with respect to organisational objectives, and further societal processes may be at stake. As an extension of the present work, we plan to apply these measures on frameworks as the CKSW meta-role model [12] or similar. In the

CKSW, society is stratified into four functional classes (*commander* class, *knowledge* class, *skilled* class, and *worker* class) that define the flow of authority, information, and labour. We expect that different distributions of agents across classes, and distinct causal mechanisms (e.g. related to production and consumption of resources) may produce effect on measures of power.

In computational terms, regardless of their ability to describe power distribution (and concentration), we observe how voting games inspired indices present a strong limitation with regards to scalability. The combinatorial nature of their calculation renders them unsuitable for real-world scenarios, where process models (and relative agentic configurations) tend to be much larger. Moreover, their applicability falters when the final goal state is determined by the evaluation of complex environmental side-effects instead of being encoded in the process model itself. The issue of identifying success conditions in long-term complex processes is particularly relevant in the training of LLMs [11, 9], which suggests there is a need for measures to adapt to non-verifiable rewards (or partially so) even at the organisational level.

We can conclude that the structural analysis of process models can serve as an informational instrument to aid the design of organisational processes. However, a purely topological perspective might not be sufficiently expressive to describe uneven allocations of agents especially when adversarial conditions are also encoded in the model. Additional studies are needed to understand whether (or to which extent) more efficient measures as usability or gatekeeper are representative for human judgment in these situations.

Finally, the reduction of *institutional power* to *operational power* we observed in the introduction is a rather common simplification occurring both in institutional studies (*e.g.*, no power is mentioned in the Grammar of Institutions [6]), as well as in computational regulation instruments as *access* or *usage control* (permission to perform an operation maps to having the power to perform that operation in the system, see *e.g.*, [16]). Yet, it is a simplification, and as such, it fails to separate material from institutional aspects. Fortunately, interpreting institutional power as symbolic causation (*i.e.*, power that creates symbolic artefacts playing an interactional role with the other agents), Petri-net model can be used as a *lingua franca* to study them too, including in combination with enforcement mechanisms.

6 Conclusion

In this paper, we proposed a formalisation for multi-agent-driven activities based on Petri nets and showed how it is apt to express organisational configurations by interpreting transitions as agentic abilities. In addition, we have shown that our formalisation supports the definition of measures of agentic power. Specifically, we suggested four such measure: two topological measures based on workflow structural properties and two game-theoretic measures obtained by reducing simple voting games to workflow execution. We observed that topological measures serve as informative tools to identify unfair distributions of power but that they are sensitive to local workflow dependencies. On the other hand, game-theoretic measures can successfully identify pivotal agents (*i.e.*, workflow bottlenecks) but are strongly limited in terms of scalability due to their combinatorial nature.

We have further demonstrated that the proposed formalisation can capture both cooperative and competitive aspects of an agentic design by means of the cumulative probabilities of reaching goal states in the workflow.

In future studies we aim to experiment with the application of the proposed measures onto realistic organisational structures and social systems. Underlying these experiments is the central hypothesis that complex social dependencies, conflicts, and environmental conditions may affect the distribution and value of resources and hence, the distribution of power.

Acknowledgements

This research was partially funded by the SSC-ICT organisation of The Netherlands which is part of the Ministry of the Interior and Kingdom Relations.

References

1. Van der Aalst, W.M.: A class of petri nets for modeling and analyzing business processes (1995)
2. Van der Aalst, W.M.: The application of petri nets to workflow management. Journal of circuits, systems, and computers **8**(01), 21–66 (1998)
3. Banzhaf III, J.F.: Weighted voting doesn't work: A mathematical analysis. Rutgers L. Rev. **19**, 317 (1964)
4. Brown, M.A.: On the logic of ability. Journal of philosophical logic pp. 1–26 (1988)
5. Burt, R.: Structural holes [1992]. Contemporary sociological theory **204**, 9780429494468–63 (2012)
6. Crawford, S.E., Ostrom, E.: A grammar of institutions. American political science review **89**(3), 582–600 (1995)
7. Gini, C.: Variabilità e mutabilità (variability and mutability). Tipografia di Paolo Cuppini, Bologna, Italy **156** (1912)
8. Hohfeld, W.N.: Fundamental legal conceptions as applied in judicial reasoning. The Yale Law Journal, Inc. **26**, 710–770 (1917), https://www.jstor.org/stable/786270
9. Kwa, T., West, B., Becker, J., Deng, A., Garcia, K., Hasin, M., Jawhar, S., Kinniment, M., Rush, N., Arx, S.V., Bloom, R., Broadley, T., Du, H., Goodrich, B., Jurkovic, N., Miles, L.H., Nix, S., Lin, T., Parikh, N., Rein, D., Sato, L.J.K., Wijk, H., Ziegler, D.M., Barnes, E., Chan, L.: Measuring ai ability to complete long tasks (2025), https://arxiv.org/abs/2503.14499
10. Lengauer, T., Tarjan, R.E.: A fast algorithm for finding dominators in a flowgraph. ACM Transactions on Programming Languages and Systems (TOPLAS) **1**(1), 121–141 (1979)
11. Patwardhan, T., Dias, R., Proehl, E., Kim, G., Wang, M., Watkins, O., Fishman, S.P., Aljubeh, M., Thacker, P., Fauconnet, L., Kim, N.S., Chao, P., Miserendino, S., Chabot, G., Li, D., Sharman, M., Barr, A., Glaese, A., Tworek, J.: Gdpval: Evaluating ai model performance on real-world economically valuable tasks (2025), https://arxiv.org/abs/2510.04374

12. Purvis, M.K., Purvis, M.A., Frantz, C.: Cksw: a folk-sociological meta-model for agent-based modelling. In: Social. Path Workshop (2014)
13. Reuel, A., Bucknall, B., Casper, S., Fist, T., Soder, L., Aarne, O., Hammond, L., Ibrahim, L., Chan, A., Wills, P., Anderljung, M., Garfinkel, B., Heim, L., Trask, A., Mukobi, G., Schaeffer, R., Baker, M., Hooker, S., Solaiman, I., Luccioni, A.S., Rajkumar, N., Moës, N., Ladish, J., Bau, D., Bricman, P., Guha, N., Newman, J., Bengio, Y., South, T., Pentland, A., Koyejo, S., Kochenderfer, M.J., Trager, R.: Open Problems in Technical AI Governance (4 2025), http://arxiv.org/abs/2407.14981
14. Shapley, L.S., Shubik, M.: A method for evaluating the distribution of power in a committee system. American political science review **48**(3), 787–792 (1954)
15. Sileno, G., Pascucci, M.: Disentangling deontic positions and abilities: a modal analysis. In: CEUR Workshop Proceedings. vol. 2710, pp. 36–50. CEUR-WS (2020)
16. Sileno, G., Pascucci, M., Markovich, R.: Allowed, or enabled, that is the question. In: DEON. pp. 297–315 (2023)
17. Straffin, P.D.: The shapley-shubik and banzhaf power indices as probabilities. The Shapley Value pp. 71–81 (1988)